

THE UNIVERSITY OF CHICAGO

TESTING THE KINEMATICS AND DYNAMICS OF INFLATION USING μ -TYPE
SPECTRAL DISTORTIONS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF ASTRONOMY AND ASTROPHYSICS

BY
DAVID ZEGEYE

CHICAGO, ILLINOIS

AUGUST 2025

Copyright © 2025 by David Zegeye

To that child who longed for a home, you can rest now.

“Humankind cannot gain anything without first giving something in return. To obtain, something of equal value must be lost. That is alchemy’s first law of Equivalent Exchange. In those days, we really believed that to be the world’s one and only truth.”

— Alphonse Elric, Fullmetal Alchemist (2003)

“Maybe life has no equal trade. Maybe you can give up all you got, and get nothing back. But still, even if I can’t prove it’s true, I have tried, for your sake, Al.”

— Edward Elric, Fullmetal Alchemist (2003)

CONTENTS

LIST OF FIGURES	vii
LIST OF TABLES	x
ACKNOWLEDGMENTS	xi
ABSTRACT	xii
1 INTRODUCTION	1
2 SPECTRAL DISTORTION ANISOTROPY FROM INFLATION FOR PRIMORDIAL BLACK HOLES	7
2.1 Introduction	7
2.2 USR Consistency Relation	9
2.3 Local Expansion	13
2.4 Modulated Acoustic Power	15
2.5 Modulated Thermalization	18
2.6 Spectral Distortion Anisotropy	19
2.7 Signal-to-Noise	24
2.8 Discussion	26
3 CMB-S4: FORECASTS FOR CONSTRAINTS ON f_{NL} THROUGH μ -DISTORTION ANISOTROPY	28
3.1 Introduction	28
3.2 Background	31
3.2.1 Power Spectrum and Bispectrum	31
3.2.2 μ cross-correlations	32
3.3 Methods	37
3.3.1 Fisher Matrix	38
3.3.2 Component separation	40
3.3.3 Instrument noise	43
3.3.4 Atmosphere	44
3.3.5 Foregrounds	45
3.3.6 Calibration	50
3.4 Survey configuration	51
3.5 Results	56
3.6 Discussion	60
4 MEASURING μ DISTORTIONS FROM THE THERMAL SUNYAEV ZELDOVICH EFFECT	64
4.1 Introduction	64
4.2 CMB Spectral Distortions	66
4.2.1 μ and y Distortions	66

4.2.2	Interfrequency Calibration	68
4.2.3	Cluster model	71
4.3	Forecast	73
4.3.1	Survey Specifications	73
4.3.2	Cluster Catalog	74
4.3.3	Forecasting Method	76
4.4	Results	78
4.4.1	Baseline Noise	79
4.4.2	Relativistic corrections	84
4.4.3	CMB anisotropies	85
4.4.4	Cluster Kinetic Sunyaev-Zeldovich Effect	86
4.4.5	Extragalactic foregrounds	87
4.4.6	Galactic foregrounds	89
4.4.7	Atmosphere	92
4.4.8	Order of Operations	93
4.4.9	Interfrequency Calibration Requirement	94
4.4.10	Cluster Profiles	96
4.5	Conclusion	98
5	SQUARE KILOMETER ARRAY AS A COSMIC MICROWAVE BACKGROUND EXPERIMENT	100
5.1	Introduction	100
5.2	Theory Background	101
5.3	Methods	103
5.3.1	Fisher Matrix	103
5.3.2	Foregrounds	104
5.3.3	Moment Expansion	105
5.3.4	Survey specifications	106
5.3.5	Calibration	107
5.4	Results and Discussion	109
A	SUPPLEMENTAL MATERIAL	113
A.1	Dilation Consistency Relation at Second Order	113
A.2	Damping	119
A.3	Flat Sky Harmonics	121
A.4	Relativistic Corrections	123
	REFERENCES	125

LIST OF FIGURES

1.1	Scales of the primordial power spectrum probed by the Cosmic Microwave Background and large-scale structure of galaxies compared to the scales probed by μ -distortions. Taken from [1].	6
2.1	The PBH curvature power spectrum $\Delta_{\zeta}^2(k)$ for a USR enhancement starting at $k_{\text{dip}} = 681 \text{ Mpc}^{-1}$ which provides the largest spectral distortion still allowed (see Fig. 2.2). Vertical lines show $k_L = 0.1 \text{ Mpc}^{-1}$ (of order the largest k mode probed by the CMB), $(aH)_{\text{th}}$ the comoving size of the horizon at the end of thermalization, k_{dip} , and the dissipation wavenumber at the end of thermalization $k_D(z_{\text{th}})$. The slow roll power law extrapolation Δ_{SR}^2 is also shown for reference.	11
2.2	The average spectral distortion $\bar{\mu}$ for various PBH enhancement scales k_{dip} (see Eq. (2.4)). Values of $k_{\text{dip}} < k_{\text{min}}$ are excluded by the FIRAS bound of $\bar{\mu} \leq 9 \times 10^{-5}$ whereas allowed values that lead to enhancement over the slow roll prediction of $\bar{\mu} \approx 2 \times 10^{-8}$ follow $\bar{\mu} \propto k_{\text{dip}}^{-4}$	20
2.3	Maximum μ white noise level $C_{\ell}^{\mu\mu} = A_{\text{noise}}$ that allows for a $S/N = 1$ measurement of the maximal $\bar{\mu} = 9 \times 10^{-5}$ and fiducial $b_{\text{th}} = -22/5$ PBH signal as a function of largest multipole measured ℓ_{max} . A_{noise} saturates at 10^{-23} around $\ell_{\text{max}} \sim 10^3$ and other models can be scaled as $A_{\text{noise}} \propto b_{\text{th}}^2 \bar{\mu}^2$	23
3.1	Angular cross-power spectrum of $\mu \times T$ and $\mu \times E$ for $\langle \mu \rangle = 2 \times 10^{-8}$ and $f_{\text{NL}} = 1$. Solid lines correspond to positive values while dashed lines correspond to negative values.	36
3.2	Spectral shape of CMB anisotropy, μ distortions, and various Galactic and extragalactic foregrounds, all divided by the CMB anisotropy SED and scaled by an arbitrary amplitude. For frequencies $\nu < 100 \text{ GHz}$, we see that Galactic synchrotron sharply rises, even relative to μ distortions. At frequencies $\nu > 100 \text{ GHz}$, Galactic dust and the thermal SZ effect start growing as the μ -distortion spectrum starts leveling off.	46
3.3	Weights used to construct the T -free μ -distortion map (i.e., the μ component of \mathbf{w}) for the CMB-S4 ultra-deep patch from different frequency channels, including foregrounds and assuming 99% correlated atmosphere.	52
3.4	$C_{\ell}^{\mu\mu, N}$ for the CMB-S4 ultra-deep survey and various assumptions for the degree of atmospheric correlation between bands. Full correlation among bands closely matches the nominal case of no atmosphere. However, even minimal decorrelation among detectors significantly worsens $C_{\ell}^{\mu\mu, N}$ at low multipoles.	56
3.5	Signal-to-noise (S/N) per multipole for various assumed values of atmospheric correlation between bands. Dashed lines correspond to S/N when only including instrumental and atmospheric noise, while solid lines also include the effects of foregrounds. In all cases, the y SED has been projected out (see text for details).	57

3.6	1- σ constraints on f_{NL} as a function of the maximum multipole used in analysis, ℓ_{max} . Dashed lines correspond to $\sigma(f_{\text{NL}})$ when only including instrumental and atmospheric noise, while solid lines also include the effects of foregrounds. In all cases, the y SED has been projected out (see text for details).	58
4.1	The fractional response of the tSZ temperature spectrum $\partial \ln \Delta T / \partial \mu$ to a monopole μ distortion with anisotropy calibration as we assume in this work (solid blue line $\partial \ln g / \partial \mu$) vs. absolute calibration (dashed blue line $\partial \ln(g/C) / \partial \mu$). The dotted gray line denotes an undistorted spectrum for reference and the spike in the curves occurs at the tSZ null where the <i>fractional</i> response diverges corresponding to a finite change in the location of the null.	71
4.2	The mass detection limit M_{lim} as a function of redshift from Ref. [2], linearly interpolated between their $\Delta z = 0.1$ bins.	75
4.3	The cumulative squared signal to noise, $(S/N)_{\text{cumul}}^2$, above a given cluster mass and redshift threshold normalized to the total $(S/N)_{\text{surv}}^2$ for the three surveys from their catalog of clusters with masses and redshifts above the red line. . . .	79
4.4	Top: Frequency scaling relation for $\sqrt{C_w}$. Blue points correspond to the S4-Wide survey's specifications. Bottom: $\sigma^2(\mu)$ for different frequency pairs for a single cluster with $y_c = 10^{-4}$ and $\theta_c = 1$ arcminute. We see that the constraints on μ improve when the channels are separated from each other, rather than being closely spaced. In general, we see that as long as there is a separation, lower frequency channels provide more sensitivity to μ in this baseline case of white detector noise only, in accordance with Fig. 4.1.	83
4.5	The fractional response of the tSZ spectrum, with 1st-order and 4th-order relativistic corrections, to the temperature of the cluster. Here we take $\mu = 0$ for a cluster with $k_B T_e = 5$ keV and another with $k_B T_e = 12$ keV.	85
5.1	C_ℓ^N for SKA1-MID and LiteBIRD, neglecting foregrounds. Solid lines correspond to $C_\ell^{\mu\mu, N}$, while dashed lines are $C_\ell^{TT, N}$. Despite SKA1-MID's temperature noise being almost two orders of magnitude worse than LiteBIRD's, SKA1-MID has substantially lower noise for μ owing to its low-frequency coverage.	108
5.2	Forecasted upper bounds on squeezed-limit f_{NL} at different scales from various observables. These scales are approximations and purely for illustration. CMB bispectrum constraints are from CMB-S4 with scales derived from $\ell_{\text{min}} = 2$ and $\ell_{\text{max}} = 1000$. Galaxy bispectrum constraints are from SPHEREx [3] with $k_{\text{min}} \approx 0.0015 \text{ Mpc}^{-1}$ and $k_{\text{max}} = 0.3 \text{ Mpc}^{-1}$ [4]. For μ -anisotropies from SKA1+LiteBIRD, "optimistic" is the second column of Table 5.1, while "pessimistic" is the third column. For reference, the gray dashed line is $f_{\text{NL}} = 1$. Figure inspired by [5, 6].	111

A.1 Change in the angular power spectrum C_ℓ depending on the damping factor. For illustrative purposes, we choose an undamped spectrum that is nearly scale invariant, with a slight spectral tilt. The spectrum with a damping scale $q_{\mu,D} = 0.11 \text{ Mpc}^{-1}$ corresponds to our treatment of the damping of μ -anisotropies. The spectrum using the damping factor from Ref. [7] has more severe damping at $\ell = 200$, but has a slightly less steep slope compared to the μ -anisotropy damping scale. The black line corresponding to $\ell = 200$ is a reference scale for the reader. 120

LIST OF TABLES

3.1	Survey configuration for the large-aperture telescope (LAT) in the CMB-S4 ultra-deep survey.	51
3.2	Survey configuration for the small-aperture telescopes (SATs) in the CMB-S4 ultra-deep survey.	53
3.3	Angular resolution and noise levels assumed for <i>Planck</i> (reproduced from Table 4 of Ref. [8]). Note that we assume <i>Planck</i> noise is white down to the ℓ_{\min} of our forecasting.	53
3.4	Constraints on f_{NL} from the CMB-S4 ultra-deep survey and <i>Planck</i> , considering the effects of atmosphere and foregrounds. In the case where neither are present, $\sigma(f_{\text{NL}}) = 48$. Including atmosphere and foregrounds worsens this constraint by at least an order of magnitude.	60
3.5	Constraints on f_{NL} when projecting out the y SED, from the CMB-S4 ultra-deep survey and <i>Planck</i> . We see in the case of no foregrounds that constraints on f_{NL} are worse than the ones given in Table 3.4. Including foregrounds, we see that constraints are very similar to those obtained if we did not project out the y SED.	61
4.1	Specifications for the CMB-S4 Wide and Deep surveys and a CMB-HD-like survey, taken from [2].	74
4.2	Forecasted constraint on μ for the baseline white detector noise of each experimental configuration and its cumulative degradation from additional effects.	78
4.3	Atmosphere parameters for the CMB-S4 Wide and Deep surveys and a CMB-HD-like survey, taken from [2].	93
4.4	Forecasted fractional improvement to $\sigma(\mu)$ (see last line of Tab. 4.2 for the values for each experimental configuration) when excluding certain individual effects.	93
5.1	Constraints on f_{NL} from different survey configurations of SKA1-MID and Lite-BIRD, with $\ell_{\max} = 1000$. We list constraints considering no foregrounds, isotropic foregrounds, and foregrounds with moment expansion terms from III.C. Constraints not in parentheses assume $\ell_{\min} = 2$, while the ones in parentheses are given for ($\ell_{\min} = 50/\ell_{\min} = 100$)	109
A.1	Percent degradation relative to our baseline constraints on f_{NL} (shown in Table 3.5) when using the damping scale from Ref. [7] instead of the fiducial value of 0.11 Mpc^{-1} . The degradation in the constraints is less than 15% in all cases.	121

ACKNOWLEDGMENTS

I have obtained and lost on my way here. Such is the dialectic between oneself and the world. I recognize though that here marks a post for rest and reflection in my journey of understanding the world.

I would like to first thank my co-advisors, Thomas Crawford and Wayne Hu. Thank you Tom for responding to my email in 2017, you have provided immense scientific and personal support since then. Thank you Wayne for your patience and encouraging me to question my assumptions, you helped me push through my insecurities in solving advanced problems. To my two other these committee members: Thank you Austin Joyce for taking time out of your day to chat about cosmological correlators and helping me understand why non-Gaussianities are exciting. Thank you Dan Hooper for being my faculty mentor and letting me keep a draft of your textbook, I am still grateful you interviewed me for admissions.

To my godparents, Hajite and Jemal. Thank you for taking care of me since I was a child. I am happy you are able to welcome Kamil into this world.

To my mother, Abaynesh Aimro. You fostered my passion of studying the cosmos. For that, I am eternally grateful. I hope one day we may be able to understand each other.

To my sisters, Hana Zegeye and Kitty. If I could sacrifice my leg, my arm, my whole self to bring you back, I would. So I will cherish your gift of a heart made Fullmetal.

ABSTRACT

Diffusion damping of the cosmic microwave background (CMB) power spectrum results from imperfect photon-baryon coupling in the pre-recombination plasma. Energy release at redshifts $5 \times 10^4 < z < 2 \times 10^6$ can create μ -type spectral distortions of the CMB. These μ -distortions trace the underlying photon density fluctuations, probing the primordial power spectrum in short-wavelength modes k_S over the range $50 \text{ Mpc}^{-1} \lesssim k \lesssim 10^4 \text{ Mpc}^{-1}$. In this thesis, I will discuss my work on using μ -distortions to probe inflation. First, I will discuss how models that seek to greatly enhance small scale power in order to form primordial black holes (PBHs) will induce a local curvature in the plasma and spatially modulate the beginning of the μ era, thus producing anisotropic μ -distortions. Another method of generating μ -anisotropies is if small-scale power modulated by long-wavelength modes originates from squeezed-limit non-Gaussianities, parameterized by f_{NL} . My second work will forecast how well the Cosmic Microwave Background - Stage IV (CMB-S4) experiment can constrain cross-correlations between μ -anisotropies and temperature anisotropies from f_{NL} . Then, I will discuss how μ -distortions will modify the thermal Sunyaev-Zel'dovich (tSZ) effect, a spectral distortion of the CMB resulting from inverse Compton scattering of CMB photons with electrons in the medium of galaxy clusters, and the ability of CMB experiments to measure μ -distortions from this effect. Lastly, I will discuss a novel idea of using the Square Kilometer Array (SKA) to measure μ -anisotropies at radio wavelengths in order to improve constraints on f_{NL} compared to traditional CMB experiments.

CHAPTER 1

INTRODUCTION

The cosmic microwave background (CMB) [9] is the faint afterglow emitted 380,000 years after the universe was in an initial dense and thermal state, known as the hot Big Bang. The CMB is characterized as following a perfect blackbody spectrum, with a temperature of $T = 2.725$ K, as measured by the Far Infrared Absolute Spectrophotometer on the *Cosmic Background Explorer* (COBE/FIRAS) [10]. Digging closer, one finds that the CMB temperature has anisotropies of order $\sim \mathcal{O}(10^{-5})$ K. Surprisingly, these disconnected hot and cold spots are found to be correlated and nearly scale-invariant across the sky. The correlation and near scale-invariance of what should be causally disconnected regions provides evidence that these anisotropies originate from quantum-scale fluctuations stretched to macroscopic scales, driven by a process known as inflation [11]. With the simplest model being that of a single scalar field rolling down a nearly-flat potential, cosmic inflation is the most compelling model for the earliest moments of our Universe and provides a mechanism for seeding the initial density perturbations that have evolved to form all the structure we see in the Universe today, from temperature anisotropies in the CMB to the large-scale structure (LSS) of galaxy surveys [12].

Cosmologists can test the kinematics of inflation by measuring the power spectrum of temperature anisotropies along momentum scales k . Modeled as a single scalar field known as the inflaton, observations can probe physics up to $\sim 10^{14}$ GeV [9], far beyond the reach of any terrestrial experiment. Cosmologists can also test the dynamics of inflation by measuring higher-order statistics, such as the bispectrum, of temperature anisotropies. These measurements, corresponding to effective scales of $10^{-4} \text{ Mpc}^{-1} \leq k \leq 10^{-1} \text{ Mpc}^{-1}$, are close to being saturated and only cover a small window of the inflationary epoch. My research has focused on extending our abilities to probe the scales and particle content of inflation by measuring μ -distortions of the CMB, which is a distortion of the blackbody spectrum of the CMB.

In the early Universe, energy injected into the plasma will efficiently thermalize through double Compton scattering and bremsstrahlung, producing a blackbody distribution with a new temperature. These two processes are highly efficient until a redshift $z_i \approx 2 \times 10^6$, at which point the Universe has expanded enough that the density necessary for double Compton and bremsstrahlung to frequently occur becomes too low [13].¹ This leaves elastic Compton scattering, which conserves photon number, as the primary method for the photon-baryon bath to reach thermal equilibrium. The inefficient thermalization introduces a chemical potential, where the mixing of different blackbody spectra at different temperatures produces a Bose-Einstein rather than Planckian distribution [15]. At small chemical potentials, the Bose-Einstein distribution can be approximated as a distorted blackbody spectrum in which

$$n(\nu) = \left[e^{h\nu/(k_B T)} - 1 \right]^{-1} \rightarrow \left[e^{h\nu/(k_B T) + \mu} - 1 \right]^{-1}, \quad (1.1)$$

where h is Planck's constant and k_B is Boltzmann's constant. This distortion of the CMB spectrum from a typical blackbody is known as a μ -type distortion. It occurs until a redshift $z_f \approx 5 \times 10^4$, where even thermalization through single Compton becomes inefficient and distortions of the y -type start being produced [16].

μ distortions can be generated from a variety of mechanisms. The primary contribution we will consider in this work is from diffusion damping of small-scale power. The μ distortions produced by diffusion damping can be related to the primordial curvature perturbations via

$$\mu(x) = \int_{\vec{k}_1} \int_{\vec{k}_2} \zeta(\vec{k}_1) \zeta(\vec{k}_2) W(\vec{k}_1, \vec{k}_2) e^{i(\vec{k}_1 + \vec{k}_2) \cdot \vec{x}}, \quad (1.2)$$

where $W(\vec{k}_1, \vec{k}_2)$, the window function, captures the weighted amount of dissipated modes

1. In reality, double Compton and bremsstrahlung processes will always be efficient at sufficiently low frequencies, restoring a full blackbody spectrum in equilibrium with the electrons deep into the Rayleigh-Jeans tail. However, at $z_i \lesssim 2 \times 10^6$, the low-frequency photons produced are neither efficiently transported towards higher frequencies (where the main distortion signals is present) nor do they give the required total photon injection to restore equilibrium [14].

with wavenumber k [17, 18, 19]. To relate the amount of μ distortions to the primordial power spectrum, we take the ensemble average of $\mu(x)$ and $k_1 = k_2$

$$\langle \mu(x) \rangle \approx \frac{1}{2\pi^2} \int dk P_\zeta(k) k^2 W(k). \quad (1.3)$$

Assuming that μ distortions are generated from sub-horizon modes at the time of dissipation, the window function is

$$W(k) \approx -4.54k^2 \int_0^\infty dz \frac{dk_{\text{D}}^{-2}}{dz} \mathcal{J}_\mu(z) e^{-\frac{2k^2}{k_{\text{D}}^2(z)}}, \quad (1.4)$$

where k_{D} is the damping scale for energy injection from diffusion damping,

$$k_{\text{D}}(z) \approx 4.1 \times 10^{-6} (1+z)^{3/2} \text{Mpc}^{-1}, \quad (1.5)$$

and \mathcal{J}_μ is the time window function for μ distortions [17, 18, 19]. The time window function is well approximated by an analytical Green's function given in Ref. [20]

$$\mathcal{J}_\mu(z) \approx \left[1 - \exp \left(- \left[\frac{1+z}{5.8 \times 10^4} \right]^{1.88} \right) \right] e^{-(z/z_i)^{5/2}}, \quad (1.6)$$

where $z_i \approx 2 \times 10^6$ is defined above as the beginning of the μ -distortion era. For single-field, slow-roll inflation, the damping of the small-scale power spectrum results in the amount of μ distortions to be roughly $\langle \mu(x) \rangle \approx 2 \times 10^{-8}$ [21, 22, 16]. We find that μ distortions are tracers of primordial perturbations in the range $50 \text{ Mpc}^{-1} \lesssim k_\mu \lesssim 1 \times 10^4 \text{ Mpc}^{-1}$ [21, 22, 16].

Measuring the average, or monopole, frequency spectrum of the CMB is extremely challenging, because an experiment is required to be extremely stable and measure the absolute power received from the sky, not just the difference in power between different sky locations. For these reasons, cosmologists often assumed that such measurements require building a new, space-based experiment. My research focused on methods to measure the μ -distortion

monopole and anisotropy by using experiments that perform differential measurements of the sky.

A method for generating μ -anisotropies is if a long-wavelength mode ζ_L modulates the amplitude of the local small-scale power, resulting in a spatially-varying power spectrum. The corresponding bispectrum is parameterized by a coupling strength $f_{\text{NL}}^{\text{local}}$ and induces a correlation between μ -distortions, which are at k_S , and CMB modes at k_L [1]. These constraints on $f_{\text{NL}}^{\text{local}}$ also correspond to $50 \text{ Mpc}^{-1} \lesssim k_\mu \lesssim 1 \times 10^4 \text{ Mpc}^{-1}$, highly complementary to the constraints on $f_{\text{NL}}^{\text{local}}$ from the CMB temperature (T) anisotropies at $10^{-4} \text{ Mpc}^{-1} \leq k \leq 10^{-1} \text{ Mpc}^{-1}$. There is a minimal amount of non-Gaussianity expected due to the modulation of small-scale power being sourced by gravitational interactions with the long-wavelength mode. This is both a target for future experiments, as well as a bias to probing additional fields that could induce couplings between long and short wavelengths. Fortunately, μ -anisotropies, correlated with T -anisotropies, are uniquely immune to this bias [18]. As such, μ - T correlations are a clean probe for testing single-field inflation.

Chapter 2 focuses on how primordial black hole (PBH) models can generate μ -anisotropies. PBH models generically have a rising slope for the power spectrum, preceded by a dipped region, at small scales. Previous work by Ref. [23] argues the dip can generate non-Gaussianities larger than standard inflationary predictions, suggesting upcoming CMB experiments can rule out various PBH models. While the PBH model may change the associated minimal value of non-Gaussianity, this doesn't change the fact that the signal is still unmeasurable using μ - T . We prove this to second-order in perturbation theory, which highlights μ - T remaining a clean probe for constraining new particles and fields present during inflation [19]. Once correcting for this error, we point out that μ - T can still test PBH models given they can enhance the μ -monopole, which will have a μ -anisotropy counterpart that won't bias μ - T signals from additional particles.

Next, Chapter 3 discusses how to measure μ -anisotropies from the ground. μ -anisotropy maps have been produced from *Planck* [24] and COBE/FIRAS [25], both of which are space-

based missions. While *Planck* currently provides the best constraints on μ -anisotropies, the signal is strongly obscured by foregrounds. The upcoming Cosmic Microwave Background - Stage IV (CMB-S4) experiment will improve on *Planck* by measuring the CMB with lower noise, and at higher resolution. One challenge that ground-based experiments face is that the atmosphere will severely suppress sensitivities to wide-area measurements of the sky, which is where much of the signal-to-noise for μ -anisotropies is located. I led the CMB-S4's second collaboration paper to study using the experiment for measuring μ -anisotropies [26]. CMB-S4's exquisite sensitivity, larger frequency coverage, and ability to measure deep patches of the sky makes up for the experiment's difficulty in using wide-area measurements of μ . We found that CMB-S4 will obtain better constraints on $f_{\text{NL}}^{\text{local}}$ from μ -anisotropies compared to *Planck* and future space-based missions, such as LiteBIRD [27, 28].

In Chapter 4, I discuss constraining the μ -monopole by measuring their distortion of the Sunyaev-Zel'dovich (SZ) effect [15], a late-time spectral distortion due to inverse Compton scattering of CMB photons with electrons in the medium of galaxy clusters. The SZ effect is measured by differencing the photons scattered by the cluster from the CMB. Cosmologists typically assume the incoming CMB photon follows a blackbody spectrum. However, we know the CMB's spectrum is expected to contain μ -distortions. This should introduce additional contributions to the SZ effect, which could be used to constrain the average μ [29]. While this effect is small for an individual cluster's spectrum, upcoming CMB surveys are expected to observe tens of thousands of clusters with the SZ effect and we found that we can place meaningful constraints on the average μ relative to *COBE/FIRAS*[10].

Lastly, while working on these previous projects, I noticed that the μ -distortion signal is the brightest in the radio regime (≤ 20 GHz). This suggests a radio survey, which can have 1000s of frequency channels and observe at frequency ranges where the effect of the atmosphere is negligible, can search for μ -distortions of the CMB. In Chapter 5, I propose using the Square Kilometer Array (SKA), combined with a traditional CMB experiment such as LiteBIRD, in order to constrain $f_{\text{NL}}^{\text{local}}$ from μ -anisotropies [30].




Figure 1.1: Scales of the primordial power spectrum probed by the Cosmic Microwave Background and large-scale structure of galaxies compared to the scales probed by μ -distortions. Taken from [1].

CHAPTER 2

SPECTRAL DISTORTION ANISOTROPY FROM INFLATION FOR PRIMORDIAL BLACK HOLES

This chapter is reproduced from Ref. [19]

2.1 Introduction

There has been much recent interest in primordial black holes (PBHs) from an extremely large enhancement of small scale fluctuations during inflation (e.g. [31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49]) given their potential to explain dark matter and the binary black hole mergers detected by LIGO-Virgo [50, 51, 52, 53, 54, 55]. Moreover primordial fluctuations on scales much smaller than those probed by the cosmic microwave background (CMB) and large-scale structure are currently relatively poorly constrained [56, 57, 58, 59].

Spectral distortions in the CMB are one way to constrain primordial power on small scales. After the thermalization epoch where photon number changing processes in the plasma drop out of equilibrium, the energy dissipated in small scale acoustic waves leaves observable distortions in the frequency spectrum [60, 61, 62, 63, 20]. The amplitude of these acoustic waves is itself enhanced if the PBH scale is not much smaller than the dissipation scale at the end of thermalization. The current bounds for chemical potential or μ distortions from COBE-FIRAS already place strong bounds on such models and rule out PBHs as a significant fraction of the dark matter between $10^4 - 10^{13} M_\odot$ [64, 65].

While current spectral distortion limits can be greatly improved with future space-based spectrometers [66, 67], an absolute measurement is limited by contamination from foregrounds and systematics in addition to instrument sensitivity. Anisotropy in spectral distortions, if they are large and correlated with CMB anisotropy in temperature and polarization, provides a promising complementary approach that uses cross correlation and differential

measurement to mitigate these issues [1, 68, 69, 17, 18, 70, 28].

These correlations can arise if the amplitude of small scale power is modulated by long wavelength fluctuations due to the squeezed bispectrum. However in single-field slow roll inflation, the squeezed bispectrum obeys the Maldacena consistency relation [71], where long wavelength curvature fluctuations modulate the scale and not the amplitude of the small scale power spectrum. Moreover this scale modulation is unobservable locally since a coordinate system established using clocks and rulers locally cannot reference the global coordinate system in which the long-wavelength mode is embedded [72]. Since spectral distortions depend only on the amount of power dissipated and not the globally referenced comoving scale from which it originates, single-field slow roll inflation does not produce spectral distortion anisotropy at leading order [1, 18].

Recently it has been suggested that single field PBH models may violate the conditions that make the spectral distortion anisotropy vanish at leading order in the squeezed bispectrum [73, 23]. All such PBH models must violate the slow roll assumption in order to enhance small scale power sufficiently rapidly [39]. Proposed PBH models typically have a period of non-attractor behavior which causes a violation of the Maldacena consistency relation for long wavelength modes that exit the horizon sufficiently close to or in the non-attractor phase [74, 75, 76, 43].

In this work we show that for scales relevant for correlation with CMB anisotropy, which are much, much larger than the horizon at the onset of the non-attractor phase, the Maldacena consistency relation holds in this limit and spectral distortion anisotropy is suppressed by the square of the ratio between the long wavelength scale and a characteristic short wavelength scale. Nevertheless, these suppressed effects can still be much larger than they are in slow roll inflation [18] and appear mainly due to the modulation of the dissipation scale at the end of thermalization by the long wavelength mode.

The outline of this paper is as follows. In §2.2, we review non-attractor “ultra-slow roll” mechanism for enhancing small scale power during inflation and discuss the modulation of

short-wavelength modes by long-wavelength modes in the form of the consistency relation. For the relevant long-wavelength scales we show that the usual Maldacena consistency relation holds and, in Appendix A.1, we explicitly verify that once combined with its impact on short-wavelength acoustic evolution, the zeroth order effect of the long-wavelength mode is a dilation of scales that is unobservable locally. In §2.4 and §2.5, we calculate the leading order effect of the long wavelength density fluctuation on the local amplitude of acoustic oscillations and their dissipation at the end of thermalization respectively. We discuss the implications of this modulation on spectral distortion anisotropy in §2.5. In §2.7 we assess the prospects for detecting this signal and discuss these results in §3.6.

2.2 USR Consistency Relation

A common, albeit tuned, aspect of many inflationary models that enhance the small scale curvature power spectrum by the orders of magnitude that would be required to later form PBHs in the radiation dominated epoch is a transient period of so-called “ultra slow-roll” (USR) [77]. During USR, the inflaton ϕ rolls faster than can be sustained by the slope of its potential $V(\phi)$,

$$\left| \frac{d\phi}{dN} \right| \gg \left| \frac{V'(\phi)}{H^2} \right|, \quad (2.1)$$

where $H = d \ln a / dt = dN / dt$ is the Hubble parameter and N is the efold. This excess kinetic energy can arise for example from a very flat potential around an inflection point (e.g. [32]) or from a sudden increase in kinetic energy due to a downward feature in the potential [48]. In either case, the excess kinetic energy of the field then redshifts away as $\epsilon M_{\text{Pl}}^2 \equiv (d\phi/dN)^2/2 \propto a^{-6}$ and the curvature fluctuation in unitary gauge grows as $\zeta = -\delta\phi/(d\phi/dN) \propto a^3$ once it crosses the horizon. Here $M_{\text{Pl}} = (8\pi G)^{-1/2}$ is the reduced Planck mass.

Modes that were already outside the horizon at the beginning of the USR phase also experience growth to the extent that they have not completely frozen out. Following the

treatment in Ref. [43], the formal solution for the evolution of ζ

$$\frac{d\zeta}{dN} = -\frac{1}{a^3\epsilon H} \left[\int dN a^3 \left(\frac{k}{aH} \right)^2 (\epsilon H)\zeta + \text{const} \right] \quad (2.2)$$

shows that during the preceding slow roll (SR) phase where ϵ is slowly varying superhorizon curvature fluctuations evolve to a constant as $d\ln\zeta/dN \approx -(k/aH)^2$. Once USR begins the term in brackets becomes constant and $d\zeta/dN \propto a^3$ and so the value of $(k/aH)_{\text{USR}}^2$, where the subscript denotes the beginning of the USR epoch, determines the growth of these modes. Given the opposite sign of $d\ln\zeta/dN$ during SR, there is a mode for which this growth is just sufficient to overcome the initially constant SR piece and cause a near zero crossing in the curvature power spectrum at $k = k_{\text{dip}}$ where

$$\left(\frac{k_{\text{dip}}}{aH} \right)_{\text{USR}}^2 \sim e^{-3\Delta N_{\text{USR}}}, \quad (2.3)$$

and ΔN_{USR} is the number of efolds in the USR phase. For example for the $\sim 10^7$ enhancement of power typical to PBH models $\Delta N_{\text{USR}} \approx \ln(10^7)/6$ and so $k_{\text{dip}} \sim 10^{-2}(aH)_{\text{USR}}$. Note that these scalings also imply that the enhancement in power for $k_{\text{dip}} < k \lesssim (aH)_{\text{USR}}$ scales as k^4 [42]. Thus the curvature power spectrum for $k \lesssim (aH)_{\text{USR}}$ can be parameterized as

$$\Delta_{\zeta}^2(k) \equiv \frac{k^3 P_{\zeta}(k)}{2\pi^2} = \Delta_{\text{SR}}^2(k) \left[1 - \left(\frac{k}{k_{\text{dip}}} \right)^2 \right]^2, \quad (2.4)$$

where Δ_{SR}^2 is the SR power spectrum which we take for simplicity to be a pure power law

$$\Delta_{\text{SR}}^2(k) = A_s \left(\frac{k}{0.05 \text{ Mpc}^{-1}} \right)^{n_s-1}, \quad (2.5)$$

with $A_s = 2.1 \times 10^{-9}$ and $n_s = 0.965$ consistent with CMB measurements [78]. Note that this form can be derived more rigorously in specific models and it holds to leading order in the downward step PBH model (see [49] Eq. (39)). For $k \gtrsim (aH)_{\text{USR}}$ the result becomes




Figure 2.1: The PBH curvature power spectrum $\Delta_\zeta^2(k)$ for a USR enhancement starting at $k_{\text{dip}} = 681 \text{ Mpc}^{-1}$ which provides the largest spectral distortion still allowed (see Fig. 2.2). Vertical lines show $k_L = 0.1 \text{ Mpc}^{-1}$ (of order the largest k mode probed by the CMB), $(aH)_{\text{th}}$ the comoving size of the horizon at the end of thermalization, k_{dip} , and the dissipation wavenumber at the end of thermalization $k_D(z_{\text{th}})$. The slow roll power law extrapolation Δ_{SR}^2 is also shown for reference.

dependent on the specific PBH model but for the modes relevant for spectral distortions that satisfy current observational constraints Eq. (2.4) suffices.

In Fig. 2.1, we show the power spectrum (2.4) for $k_{\text{dip}} = 681 \text{ Mpc}^{-1}$, which we shall see below in Fig. 2.2 is the largest scale allowed by current constraints on spectral distortions which arise from the power near the dissipation scale at end of thermalization $k_D(z_{\text{th}}) \approx 10^4 \text{ Mpc}^{-1}$. Notice the large hierarchy of scales between these values and the smallest scale accessible to measurements of the primary CMB anisotropy $k_L \sim 0.1 \text{ Mpc}^{-1}$ due to the dissipation scale at recombination.

The wavenumber k_{dip} separates two very different regimes for the impact of long wavelength fluctuations on much shorter wavelength power or equivalently the squeezed bispectrum. For a long wavelength $k_L \ll k_{\text{dip}}$ the squeezed bispectrum obeys the usual Maldacena consistency relation [71]

$$\lim_{k_L/k_S \rightarrow 0} B_\zeta(k_L, k_S, k_S) = - \left[\frac{d \ln \Delta_\zeta^2(k_S)}{d \ln k_S} + \mathcal{O} \left(\frac{k_L}{k_{\text{dip}}} \right)^2 \right] \times P_\zeta(k_L) P_\zeta(k_S), \quad (2.6)$$

which can be demonstrated by explicit calculations using the in-in or δN formalisms [43]. The interpretation of this relation is that the *scale* of the short wavelength modes are dilated by the nearly constant long wavelength curvature perturbation, which acts as a spatial fluctuation in the local scale factor

$$\begin{aligned} \Delta_\zeta^2(k_S, \mathbf{y}) &\approx \Delta_\zeta^2(k_S(1 - \zeta_L(\mathbf{y}))) \\ &\approx \Delta_\zeta^2(k_S) \left(1 - \frac{d \ln \Delta_\zeta^2(k_S)}{d \ln k_S} \zeta_L(\mathbf{y}) \right), \end{aligned} \quad (2.7)$$

where local spatial variations are denoted by the comoving coordinate \mathbf{y} . Notice that the amplitude of the power spectrum does not actually change locally, just the comoving scale that the power is associated with. Physically, the long wavelength mode just changes the efold at which the short wavelength mode enters into USR, not any of the dynamics due to the USR growth. As shown in Ref. [18] this change in scale cannot modulate spectral distortions, which depend only on the amplitude of the power dissipated, not the scale. In the Appendix we explicitly verify that a constant ζ_L generates a locally unobservable dilation for both the primordial non-Gaussianity and the subsequent dynamics of short wavelength modes to second order in perturbations.

For $k_L \gtrsim k_{\text{dip}}$, the leading effect of the long wavelength mode is not simply a dilation of scales and hence the consistency relation (2.6) no longer holds [74, 75]. Physically, the

inflaton field fluctuation of the long wavelength mode changes the number of e-folds that the short wavelength mode experiences USR since there is no longer an attractor solution that makes it equivalent to a shift along the background phase space trajectory [43]. For example, a constant backwards fluctuation of the field means that locally the short wavelength modes see more e-folds of USR growth and so the actual amplitude of the small scale power spectrum is modulated by the long wavelength mode. This then would produce a spatial modulation in spectral distortions.

We shall see that for modes relevant for correlation with CMB anisotropy and models that satisfy spectral distortion constraints, $k_L \ll k_{\text{dip}}$. Even in this regime, there is always some dynamical effect of short wavelength growth in the long wavelength field to the extent that $\zeta_L \neq \text{const}$. As we have seen, here the evolving part of ζ_L during USR is suppressed by k_L^2/k_{dip}^2 . This expectation is consistent with an explicit calculation of the leading order correction the Maldacena consistency relation given in Eq. (2.6) in an inflection point PBH model [43]. However since again the relevant $k_L \ll k_{\text{dip}}$ this correction produces a very small μT correlation (see Fig. 2.1 and §4.2.1).

2.3 Local Expansion

After inflation, there are $(k_L/aH)^2$ suppressed modulations of short wavelength physics by the long wavelength curvature fluctuations ζ_L , analogous to the corrections to the primordial bispectrum from $k_{\text{dip}} \propto (aH)_{\text{USR}}$ in the previous section. To the local observer, the long wavelength mode appears as a change to the background cosmology induced by the density perturbation that it carries. Since synchronous observers are freely falling test particles, we can absorb these synchronous gauge adiabatic density fluctuations δ_L [79, 80],

$$\delta_L = \frac{1}{3} \left(\frac{k}{aH} \right)^2 \zeta_L \propto a^2 \tag{2.8}$$

during radiation domination into a new background or “separate universe”.

Following the construction of the separate universe for the late time growth of structure [81, 82], we can define a local background density $\bar{\rho}_L$ as

$$\bar{\rho}(1 + \delta_L) = \bar{\rho}_L, \quad (2.9)$$

which implies that the local observer sees a scale factor a_L , such that $\bar{\rho}_L \propto a_L^{-4}$, that is related to the global scale factor a at equal times as

$$a_L \approx a(1 - \delta_L/4), \quad (2.10)$$

where we have chosen the normalization such that $a_L \rightarrow a$ at early times and made use of the fact that synchronous gauge observers in the global universe measure time in the same way as local observers. Notice that this normalization removes the dilation effect of ζ_L in Eq. (2.7) by measuring scales locally so that they coincide when $\delta_L \ll 1$ at the end of inflation.

By virtue of this normalization and conservation of particles, at the same numerical value of the scale factors (or efolds) from the end of inflation in the local and global universe, all particle number and energy densities are the same. However the scale factors do not coincide at the same time in the local and global universe.

To extract the cosmological parameters of the separate universe, we can express the expansion rate as a function of the local scale factor. With the definition $H_L \equiv d \ln a_L / dt$ and the radiation dominated growth of $\delta_L \propto a^2$ we have at equal times

$$H_L^2 \approx H^2(1 - \delta_L). \quad (2.11)$$

In the global universe let us define a reference epoch a_r in the radiation dominated limit where $H^2(a) = H_r^2(a_r/a)^4$ so that we can express the local expansion rate in terms of the

local scale factor as

$$H_L^2 = H_r^2 \left(\frac{a_r}{a}\right)^4 (1 - \delta_L) \approx H_r^2 \left(\frac{a_r}{a_L}\right)^4 (1 - 2\delta_L). \quad (2.12)$$

In the local universe this takes the form of the Friedmann equation with $\delta_L \propto a^2 \approx a_L^2$ playing the role of spatial curvature to linear order in δ_L , specifically

$$H_0^2 \Omega_{KL} = -2\delta_L(a_r) H_r^2 a_r^2 \quad (2.13)$$

so that $\delta_L(a_r)$ at the reference epoch defines the comoving curvature scale in units of the comoving Hubble scale at that epoch. Finally it is useful to express the local conformal time as

$$\eta_L = \int \frac{d \ln a_L}{a_L H_L} \approx \frac{a_L}{a_r} \frac{1}{a_r H_r} \left(1 + \frac{\delta_L}{3}\right). \quad (2.14)$$

These relations now determine the modulation of all short wavelength observables by a long wavelength curvature fluctuation ζ_L which to leading order scales as $(k_L/aH)^2 \zeta_L$. The specific size of the modulation will then depend on the epoch at which it influences the short wavelength observable the most. We shall see that for μ distortions and PBH models, this is the end of the thermalization epoch.

2.4 Modulated Acoustic Power

The first step in understanding the local modulation of spectral distortions is to determine the impact of the long wavelength curvature perturbation ζ_L on the amplitude of short wavelength acoustic oscillations in the CMB during radiation domination. These oscillations then dissipate via diffusion damping leaving a spectral distortion after the thermalization epoch. We shall see that since acoustic oscillations are generated at horizon crossing of the short wavelength mode k_S , the impact of the long wavelength mode occurs at horizon crossing of the short wavelength mode $aH = k_S$ and therefore scales as $(k_L/k_S)^2 \zeta_L$. The

impact of CMB scale wavenumbers k_L is therefore highly suppressed for the modes that contribute to spectral distortions (see Fig. 2.1).

We can analytically understand this scaling in the simple case where the photons dominate the radiation density, i.e. neglecting the effect of neutrinos which only change the numerical factors and not the overall scaling. In this case we can solve the perturbation equations in terms of the continuity and Euler equations for the photon fluid under self gravity in the photon-dominated local universe as (see [83] Eq. (10))

$$\begin{aligned}\Delta'_\gamma - \frac{y'}{y}\Delta_\gamma &= -\frac{4}{\sqrt{3}}\frac{1 - y''/y + 2(y'/y)^2}{1 + 6/(f_K y^2)}V_\gamma, \\ V'_\gamma + \frac{y'}{y}V_\gamma &= \left(1 - \frac{6}{f_K y^2}\right)\frac{\sqrt{3}}{4}\Delta_\gamma,\end{aligned}\tag{2.15}$$

where Δ_γ is the comoving gauge photon density perturbation, $y = (\Omega_\gamma H_0^2)^{-1/2} a_L k_S$, $' = d/dx$ with $x = k_S \eta_L / \sqrt{3}$, and $f_K = 1 + 3\Omega_{KL} H_0^2 / k_S^2$. Here we have again made use of the fact that at the same value for the scale factor, the physical density of the photons is the same in the local and global universe. The synchronous and comoving gauge differ in their density perturbations outside the horizon due to radiation pressure so $\lim_{x \rightarrow 0} \Delta_\gamma \neq \delta_L$ but they do approach each other for $x \gg 1$ in the regime relevant for spectral distortions.

Without the curvature perturbation induced by the long wavelength mode, $x = y/\sqrt{3}$ during photon domination and the solution is analytic

$$\lim_{y \rightarrow \sqrt{3}x} \Delta_\gamma = 4 \left(\frac{\sin x}{x} - \cos x \right) \zeta_S,\tag{2.16}$$

from which we can read off the usual transfer function¹ for acoustic oscillations $-4 \cos x$ at $x \gg 1$

We can now solve for the leading order correction from the small curvature induced by

1. This transfer function is reduced to $-4/(1 + 4R_\nu/15) \sim -3.61$ when neutrinos with $R_\nu = \rho_\nu/(\rho_\nu + \rho_\gamma)$ are included [84]

the long wavelength mode. Since we neglect neutrinos $H_r^2 = \Omega_\gamma H_0^2 a_r^{-4}$ and we have

$$y \approx \sqrt{3}x \left(1 - \frac{\delta_L}{3}\right) \approx \sqrt{3}x \left(1 - \frac{x^2}{3}\alpha\right), \quad (2.17)$$

so that Eq. (2.15) becomes to leading order:

$$\Delta_\gamma'' + \left(1 - \frac{2}{x^2} - \frac{4}{3}\alpha\right) \Delta_\gamma = 0 \quad (2.18)$$

with

$$\alpha \equiv \left(\frac{k_L}{k_S}\right)^2 \zeta_L \quad (2.19)$$

and $f_K = 1 - 2\alpha$ constant in time.

Notice that at $x \gg 1$ this takes the form of an oscillator equation with a perturbed constant frequency

$$\tilde{x} = \left(1 - \frac{2}{3}\alpha\right) x. \quad (2.20)$$

We can now iterate to solve Eq. (2.18) to first order in α using the zeroth order solution (2.16) to determine $\alpha\Delta_\gamma$ as an external source,

$$\Delta_\gamma \approx 4\zeta_S \left(1 + \frac{4}{3}\alpha\right) \left(\frac{\sin \tilde{x}}{\tilde{x}} - \cos \tilde{x}\right). \quad (2.21)$$

We can explicitly verify that this form solves (2.18) and satisfies $\Delta_\gamma \approx 4\zeta_S x^2/3$ in $x \ll 1$ to linear order in α .

Therefore there is an $\alpha = (k_L/k_S)^2 \zeta_L$ change in the amplitude and frequency of the acoustic wave. Since $k_L/k_S \ll 1$, this is a large suppression factor and produces a negligible change in the local spectral distortion once the acoustic modes have dissipated. We can therefore hereafter assume that the power in acoustic modes at k_S is effectively the same in the global and local universe at the same value of the scale factors.

2.5 Modulated Thermalization

We can now compute the local dissipation of energy from the acoustic waves into μ spectral distortions in the separate universe. Changes in thermalization due to the local background induce larger local variations in μ for PBH models, than the primordial effect from k_{dip} or the acoustic growth, since we shall see they scale with the comoving horizon size at the end of thermalization $(k_L/aH)_{\text{th}}^2 \zeta_L$.

Let us now see how thermalization is altered by the long-wavelength modulation in the local universe. The thermalization rate for the joint action of double Compton scattering $e^- + \gamma \leftrightarrow e^- + 2\gamma$, which changes photon number, and Compton scattering $e^- + \gamma \leftrightarrow e^- + \gamma$, which redistributes energy, to establish a blackbody scales as (see e.g. [85, 13])

$$\Gamma_{\text{th}} \propto T^{3/2} n_e \propto a_L^{-9/2}, \quad (2.22)$$

where n_e is the free electron density and T is the plasma temperature and recall that at the same numerical value of the scale factors all particle densities are the same in the local and global universe.

Let us define the thermalization time by the condition

$$\int dt \Gamma_{\text{th}} = \int_{\ln a_{\text{th}}}^0 d \ln a_L \frac{\Gamma_{\text{th}}}{H_L} = 1. \quad (2.23)$$

Using Eq. (2.12), the change in the efold of thermalization is given by

$$\Delta \ln a_{\text{th}} = 2\delta_L(a_{\text{th}}) \quad (2.24)$$

such that in an overdensity, thermalization continues to a later efold due to a slower expansion rate.

The energy in acoustic waves is dissipated when the photon diffusion length crosses the wavelength. If this occurs after the thermalization epoch but before the Compton scattering

becomes inefficient at redistributing energy, then this energy is transferred into a μ spectral distortion (see Eq. (2.31) below). The diffusion wavenumber in the radiation epoch is given by

$$k_D^{-2} \propto \int \frac{d \ln a}{aH} \frac{1}{n_e \sigma_T a}, \quad (2.25)$$

where σ_T is the Thomson cross section (see Eq. (2.29) below for the general expression). Again let us make use of the fact that the particle densities are the same in the local and global universe at the same efold. We then get the change in the diffusion wavenumber as

$$\begin{aligned} \Delta \ln k_D(a_{\text{th}}) &= -\frac{3}{2} \Delta \ln a_{\text{th}} - \frac{3}{10} \delta_L(a_{\text{th}}) \\ &= -\frac{33}{10} \delta_L(a_{\text{th}}) = -\frac{11}{10} \left(\frac{k_L}{aH} \right)_{\text{th}}^2 \zeta_L, \end{aligned} \quad (2.26)$$

where recall that in the global universe $k_D \propto a^{-3/2}$. Therefore in an overdensity the diffusion wavenumber decreases. The net result is that the maximal wavenumber that dissipates into μ distortions given the same small scale power spectrum in local coordinates as global coordinates is modulated by $(k_L/aH)_{\text{th}}^2 \zeta_L$. Since the power spectrum is strongly blue tilted, it is this modulation that changes the local value of μ . This is in contrast to single field slow roll inflation where the larger horizon size at the end of the μ epoch makes the modulation at that time the dominant effect [18].

2.6 Spectral Distortion Anisotropy

With the local change in the thermalization and dissipation scales due to the long wavelength curvature perturbation as calculated in the previous section, we can now determine how it modulates the μ spectral distortion. Because long wavelength curvature perturbations also generate CMB anisotropy, this leads to a potentially observable μT correlation in PBH models.

First we calculate the average spectral distortion in the global universe with the PBH

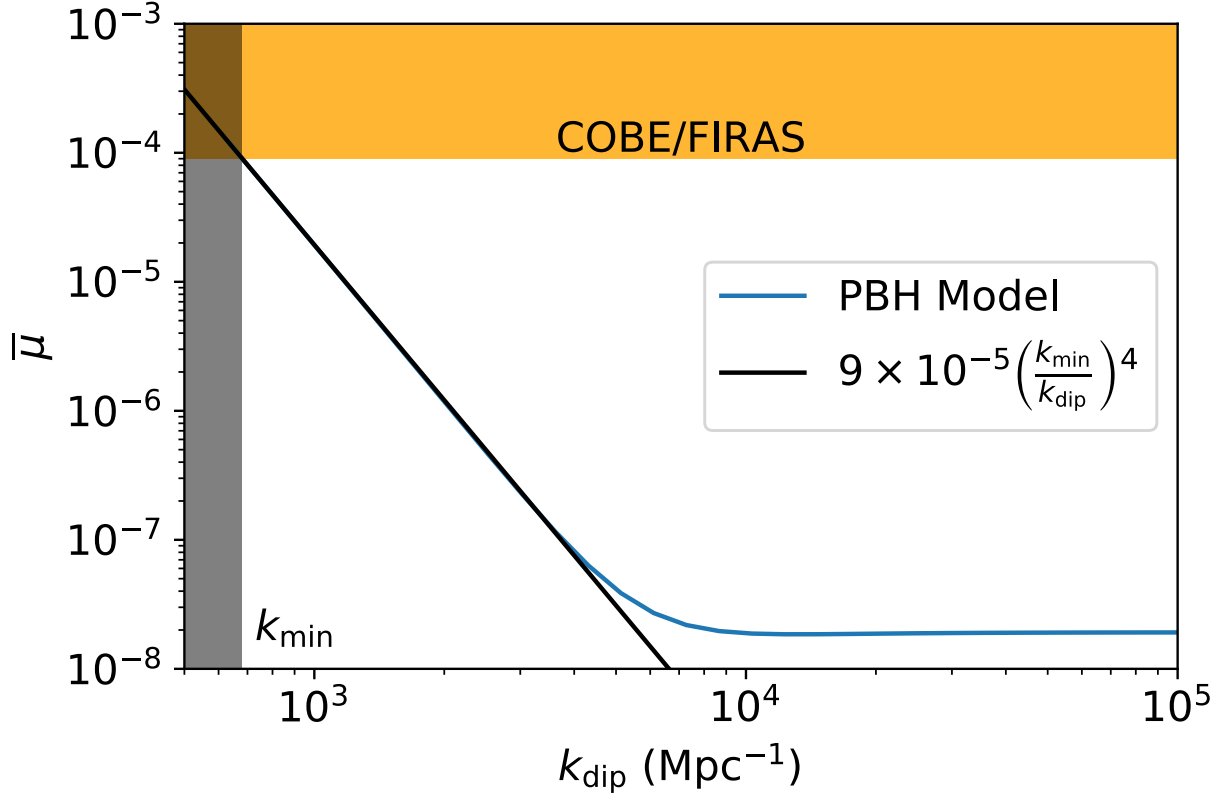


Figure 2.2: The average spectral distortion $\bar{\mu}$ for various PBH enhancement scales k_{dip} (see Eq. (2.4)). Values of $k_{\text{dip}} < k_{\text{min}}$ are excluded by the FIRAS bound of $\bar{\mu} \leq 9 \times 10^{-5}$ whereas allowed values that lead to enhancement over the slow roll prediction of $\bar{\mu} \approx 2 \times 10^{-8}$ follow $\bar{\mu} \propto k_{\text{dip}}^{-4}$.

power spectrum of Eq. (2.4). Following [18, 17],

$$\bar{\mu} = \int d \ln k \Delta_{\zeta}^2(k) W(k), \quad (2.27)$$

where to good approximation

$$W(k) \approx -4.54 k^2 \int_0^{\infty} dz \frac{dk_D^{-2}}{dz} \mathcal{J}_{\mu}(z) e^{-2k^2/k_D^2(z)}. \quad (2.28)$$

Here the diffusion wavenumber k_D is given by

$$k_D^{-2}(z) = \frac{1}{6} \int_z^\infty \frac{dz}{H} \frac{1}{n_e \sigma_{Ta}} \frac{R^2 + 16(1+R)/15}{(1+R)^2}, \quad (2.29)$$

where $R = 3\rho_b/4\rho_\gamma$. In the radiation dominated epoch and with the best fit Λ CDM parameters

$$k_D \approx 4.05 \times 10^{-6} (1+z)^{3/2} \text{ Mpc}^{-1}. \quad (2.30)$$

The Green function \mathcal{J}_μ for μ distortions is well approximated by [20]

$$\mathcal{J}_\mu(z) \approx \left[1 - \exp\left(-\left[\frac{1+z}{5.8 \times 10^4}\right]^{1.88}\right) \right] \Theta(z - z_{\text{rec}}) \times e^{-(z/z_{\text{th}})^{5/2}}, \quad (2.31)$$

where we have included a Θ step function at recombination z_{rec} since below this redshift there are no acoustic waves to dissipate. Here $z_{\text{th}} \approx 2 \times 10^6$ is the thermalization redshift in the global universe and the quantity in brackets of Eq. (2.31) determines the gradual transition from the μ epoch to the Compton y epoch. Notice that $\bar{\mu}$ receives contributions from the initial inflationary power spectrum at wavenumber k mainly when it crosses k_D but these contributions are sharply cut off by the thermalization process above $k_D(z_{\text{th}})$.

In Fig. 2.2 we show $\bar{\mu}$ as a function of k_{dip} in Eq. (2.4). The COBE-FIRAS constraint $\bar{\mu} < 9 \times 10^{-5}$ (95% CL) [10] places a limit of

$$k_{\text{dip}} > k_{\text{min}} \approx 681 \text{ Mpc}^{-1} \quad (2.32)$$

and

$$\bar{\mu} \approx 9 \times 10^{-5} \left(\frac{k_{\text{min}}}{k_{\text{dip}}} \right)^4 \quad (2.33)$$

for models with $k_{\text{min}} < k_{\text{dip}} \lesssim 5000 \text{ Mpc}^{-1}$ whereas for even larger k_{dip} , $\bar{\mu}$ asymptotes to its slow-roll value of $\bar{\mu} \approx 2 \times 10^{-8}$ because enhanced scales have already dissipated by the

thermalization epoch. Notice that even this smallest allowed value of k_{dip} is much greater than the comoving horizon at the end of thermalization $(aH)_{\text{th}} \approx 4.3 \text{ Mpc}^{-1}$.

The modulation of $k_D(z_{\text{th}})$ due to the long wavelength mode therefore modulates the local value of μ from its background value $\bar{\mu}$. Given that the PBH power spectrum rises as approximately $\Delta_{\zeta}^2 \propto k^{4-(1-n_s)}$ in the region which can enhance μ ,² we get the fractional change in the local value of μ as

$$\delta \ln \mu = [4 - (1 - n_s)] \Delta \ln k_D \equiv b_{\text{th}} \left(\frac{k_L}{aH} \right)_{\text{th}}^2 \zeta_L, \quad (2.34)$$

where from Eq. (2.26) $b_{\text{th}} \approx -22/5 + 11(1 - n_s)/10$. Note that this approximation can in the future be improved by numerically recalibrating the Green function \mathcal{J}_μ in the separate universe so we parameterize the result as a “bias” factor b_{th} , similar to the slow-roll calculation of Ref. [18] but with respect to the horizon scale at the end of thermalization rather than the end of the μ epoch. For example if we assume that the functional form of \mathcal{J}_μ remains the same and only z_{th} changes according to Eq. (2.24), then $b_{\text{th}} \approx -4.1$ for $k_{\text{dip}} = k_{\text{min}}$.

Given the level of precision in these estimates, we simply take

$$b_{\text{th}} = -22/5 \quad (2.35)$$

as our fiducial bias.

Following Ref. [18], we can characterize the long wavelength correlation between CMB temperature anisotropy and μ anisotropy with the angular cross power spectrum

$$C_\ell^{\mu T} \equiv \bar{\mu} b_{\text{th}} C_\ell^{\mu T, b_{\text{th}}}, \quad (2.36)$$

2. For consistency with Eq. (2.4) we have retained $1 - n_s$ here but note that any actual $\mathcal{O}(1 - n_s)$ correction to the local slope will depend on the details of the model.




Figure 2.3: Maximum μ white noise level $C_\ell^{\mu\mu} = A_{\text{noise}}$ that allows for a $S/N = 1$ measurement of the maximal $\bar{\mu} = 9 \times 10^{-5}$ and fiducial $b_{\text{th}} = -22/5$ PBH signal as a function of largest multipole measured ℓ_{max} . A_{noise} saturates at 10^{-23} around $\ell_{\text{max}} \sim 10^3$ and other models can be scaled as $A_{\text{noise}} \propto b_{\text{th}}^2 \bar{\mu}^2$.

where

$$C_\ell^{\mu T, b_{\text{th}}} = \frac{4\pi}{(aH)_{\text{th}}^2} \int_0^\infty dk k \Delta_\zeta^2(k) \Delta_\ell^\mu(k) \Delta_\ell^T(k). \quad (2.37)$$

Here Δ_ℓ^T is the CMB temperature transfer function which we take from CAMB³ and

$$\Delta_\ell^\mu(k) = e^{-k^2/q_{\mu,D}^2} j_\ell(\eta_0 - \eta_{\text{rec}}), \quad (2.38)$$

where $q_{\mu,D} \approx 0.084 \text{ Mpc}^{-1}$ is the dissipation scale of μ inhomogeneities at recombination

3. <https://camb.info>

[69].⁴ Notice that this damping factor is comparable to that of the temperature transfer function at $k_D(z_{\text{rec}})$ and in combination they limit the integral in Eq. (2.36) to the long wavelength k_L values of the CMB. It is straightforward to generalize this result for the cross correlation with CMB polarization with the polarization transfer function which we leave that to a future work.

Finally notice that the primordial deviation from the dilation or consistency relation and the separate universe growth of acoustic oscillations takes the same form as Eq. (2.36) but are suppressed by factors of $(aH/k_{\text{dip}})_{\text{th}}^2 \lesssim 4 \times 10^{-5}$ and $(aH/k_D)_{\text{th}}^2 \approx 10^{-7}$ respectively and hence provide a negligible correction to $C_\ell^{\mu T}$ as calculated in Eq. (2.36) from the thermalization bias.

2.7 Signal-to-Noise

The μT cross power spectrum as calculated in the previous section is enhanced in PBH models by the k^4 rise in the small scale curvature power spectrum as long as $k_{\text{dip}} \ll k_D(z_{\text{th}})$ but also suppressed by the smallness of the density perturbation associated with ζ_L at the thermalization epoch $(k_L/aH)_{\text{th}}^2 \zeta_L$. Consequently unlike μT correlations from nearly scale invariant perturbations in multifield inflation [1, 68, 69, 17, 18, 28], the signal-to-noise is dominated by the smallest angular scales at which the correlation can be detected, namely the damping scale of primary CMB anisotropy $k_L \sim 0.1 \text{ Mpc}^{-1}$.

To estimate the signal-to-noise in μT , we take the Gaussian approximation for μ anisotropy

$$\left(\frac{S}{N}\right)^2 = \sum_{\ell=2}^{\ell_{\text{max}}} (2\ell + 1) \frac{(C_\ell^{\mu T})^2}{(C_\ell^{\mu T})^2 + C_\ell^{\mu\mu} C_\ell^{TT}}, \quad (2.39)$$

where the $C_\ell^{\mu\mu}$ and C_ℓ^{TT} terms in the denominator include both the sample variance of the signal and any instrumental noise from their measurement.

4. This approximation can also be refined in the future with the Green function for spectral spatial anisotropy (Chluba, private communication).

Given the smallness of the μ anisotropy, $C_\ell^{\mu\mu}$ will be noise dominated for the foreseeable future.⁵ To assess the maximal white noise level $C_\ell^{\mu\mu} = A_{\text{noise}}$ at which the signal is barely detectable at $S/N = 1$, we can drop the $C_\ell^{\mu T}$ sample variance term in the denominator of Eq. (2.39) and solve for white noise term

$$\begin{aligned} A_{\text{noise}} &= \sum_{\ell=2}^{\ell_{\text{max}}} (2\ell + 1) \frac{(C_\ell^{\mu T})^2}{C_\ell^{TT}} \\ &= b_{\text{th}}^2 \bar{\mu}^2 \sum_{\ell=2}^{\ell_{\text{max}}} (2\ell + 1) \frac{(C_\ell^{\mu T, b_{\text{th}}})^2}{C_\ell^{TT}}. \end{aligned} \quad (2.40)$$

To estimate the maximal noise level for detection, we assume that the TT measurement is cosmic variance limited to ℓ_{max} . In Fig. 3.5, we show this maximal noise amplitude as a function of ℓ_{max} for the maximal signal $\bar{\mu} = 9 \times 10^{-5}$ and $b_{\text{th}} = -22/5$. Notice that the result saturates at around $\ell_{\text{max}} \approx 10^3$ at a level of $\sim 10^{-23}$ since both T and μ anisotropies are damped by diffusion. To constrain other PBH models, this result can be scaled as $A_{\text{noise}} \propto b_{\text{th}}^2 \langle \mu \rangle^2$ using Eq. (2.33).

As this represents the maximal signal, detecting this effect will be challenging experimentally. First, to optimize detection, an experiment would need at least several arcminute scale resolution since

$$C_\ell^{\mu\mu} = A_{\text{noise}} e^{\frac{\ell^2 \theta_b^2}{8 \ln 2}}, \quad (2.41)$$

where θ_b is the full width half max of the beam in radians.

Second, the required noise level $A_{\text{noise}} < 10^{-23}$ is quite stringent even ignoring foregrounds and systematics. For reference, a PIXIE-like mission [86, 87] which aims to measure $\bar{\mu}$ in single field slow roll with $\theta_b \approx 1.6^\circ$, will at best achieve $A_{\text{noise}} \sim 10^{-15}$ which would suffice to constrain a squeezed bispectrum signal from multifield inflation to $|f_{\text{NL}}| \lesssim 3000$

5. The ultimate limit comes from the residual fluctuations from the averaging across CMB scales of the patches that dissipate [1, 69, 18] but note that in the PBH context the patches that contribute most to $\bar{\mu}$ are much smaller in size than the SR contributions at the end of the μ epoch and do not appreciably enhance the noise.

at 68% CL [1, 68]. To constrain the maximal PBH signal, one would require the equivalent sensitivity to $|f_{\text{NL}}| < 0.3$ but with a much higher angular resolution. LiteBIRD can potentially achieve a detector sensitivity of $A_{\text{noise}} \approx 10^{-18}$ but still with only $\sim 0.5^\circ$ resolution and subject to foreground contamination [28].

2.8 Discussion

We have shown that in spite of the large squeezed bispectrum due to the enhancement of small scale power in single field inflationary PBH models, the spectral distortion anisotropy is highly suppressed since for scales relevant to CMB cross correlation, it represents an unobservable modulation of global scales rather than of the local amplitude of the short wavelength modes. Nonetheless the μ anisotropy can be larger than in single field slow roll inflation because of the large enhancement of small scale power itself. The largest effect comes from the local modulation of the expansion rate at a given locally measured efold from the end of inflation and hence the end of the thermalization epoch. This modulation provides a spatial variation in the amount of power in acoustic waves dissipated near thermalization that causes a μ distortion that is correlated with CMB anisotropy itself.

This leading order correlation is enhanced by the k^4 rise in small scale power in PBH models but suppressed by the square of the ratio of the comoving horizon at end of thermalization to observable CMB scales, $(k_L/aH)_{\text{th}}^2 \lesssim 0.0005$. On the other hand, the enhancement of the average μ itself can be up to ~ 5000 and still satisfy current COBE-FIRAS bounds. These compensating factors leave the signal potentially observable but still well beyond the capability of proposed space-based instruments like PIXIE and LiteBIRD. Moreover to detect the correlation at the smallest observable CMB scales, where the signal peaks, would require a telescope with several arcminute scale resolution or better.

These properties of the PBH μT correlation suggest that in the future, ground based instruments may provide a competitive path forward, given the rapid advance in the scale of detectors deployed on large telescopes into the CMB-S4 [88] era and beyond. Unlike

the absolute measurement of $\bar{\mu}$, systematics and foregrounds can also be mitigated by differential measurements and cross correlation [68]. Furthermore the μT correlation can be supplemented by polarization cross correlation [89, 90, 28]. We leave these studies to a future work.

CHAPTER 3

CMB-S4: FORECASTS FOR CONSTRAINTS ON f_{NL} THROUGH μ -DISTORTION ANISOTROPY

This is a slightly truncated version of Ref. [91] in order to reduce redundancy.

3.1 Introduction

The paradigm of cosmic inflation presents arguably the most compelling and plausible scenario for the earliest moments of the existence of our Universe and provides a mechanism for creating the density perturbations that have grown under the influence of gravity to form all the structure we see in the Universe today. (For an overview of the theoretical background and observational evidence for inflation, see, e.g., the review by Ellis & Wands in Chapter 23 in Ref. [92].) If a period of inflation did occur, observations of the cosmic microwave background (CMB) and tracers of the density field can give us clues about the nature of the fields involved in inflation and their dynamics. By extension, these observations can probe physics at the energy scale of the inflationary potential, far beyond the energy reach of any terrestrial experiment. The CMB has most famously been used to constrain inflation models through measurements of the power spectrum of density fluctuations and constraints on the imprint of inflationary gravitational-wave background on the polarization of the CMB (e.g., [93]), but CMB observations can also shed light on inflationary dynamics through searches for signatures of primordial non-Gaussianity (PNG, e.g., [94, 95, 96, 97, 98, 99, 100, 101, 102, 103, 104, 105, 106]).

The most well-studied form of PNG is the so-called squeezed-limit bispectrum (or Fourier-domain three-point function), which describes a configuration of triangles in Fourier space in which two of the k modes have a much larger value than the third. This corresponds to a real-space configuration in which a long-wavelength mode modulates the amplitude of small-scale fluctuation power. Mathematically, this is usually expressed in terms of the

relationship of the bispectrum of curvature perturbations to the curvature power spectrum (e.g., [107]):

$$B_{\zeta}(k_1, k_2, k_3)_{k_3 \ll k_1 \approx k_2} = \frac{12}{5} f_{\text{NL}} P_{\zeta}(k_1) P_{\zeta}(k_3), \quad (3.1)$$

where f_{NL} parameterizes the amplitude of local non-Gaussianity and is defined explicitly through the first-order expression for real-space local PNG:

$$\zeta(x) = \zeta_g + \frac{3}{5} f_{\text{NL}} (\zeta_g^2 - \langle \zeta_g^2 \rangle), \quad (3.2)$$

where ζ_g is a random Gaussian field. One of the reasons the squeezed-limit configuration is so well-studied is that models of inflation with a single scalar field whose kinetic energy is always much less than its potential energy (“single-field, slow-roll inflation”) produce vanishingly small amounts of PNG of the squeezed-limit type [108]. A detection of this type of PNG at levels of $f_{\text{NL}} \gtrsim 0.01$ would thus rule out large classes of inflation models, including many of the most popular and viable models (e.g., [88]).

The best current limits on the value of f_{NL} in the squeezed limit come from analyses of the primary CMB from the *Planck* satellite [109]: $f_{\text{NL}} = -0.9 \pm 5.1$ (68% CL). Even cosmic-variance-limited maps of the CMB temperature and polarization out to $\ell = 4000$ would only improve these limits by roughly a factor of five [110]. Interest is thus high in other methods of determining f_{NL} . From a pure mode-counting perspective, there is much more information in the distribution of matter in the local and moderate-redshift Universe, but the non-Gaussianity caused by non-linear growth complicates bispectrum measurements of the galaxy distribution significantly. Previous work has shown that squeezed-limit PNG causes a unique scale-dependent bias in the galaxy distribution [111], and this signature is a target of upcoming galaxy surveys (e.g., [112]), as is a direct measurement of the matter bispectrum supplanted with careful modeling (e.g., [113]).

Another promising avenue for detecting or constraining local PNG involves distortions to the blackbody spectrum of the CMB [114, 115, 13, 63]. Injection of energy into the primordial

plasma will distort the spectrum of the CMB, and a known source of energy release in the early Universe is the dissipation of small-scale acoustic waves [116, 61, 62, 21] from photon diffusion or Silk damping [117]. The authors of Refs. [1, 68] pointed out that a modulation of small-scale power by long-wavelength modes that underlie squeezed-limit PNG would result in *anisotropy* in the μ distortion and that a cosmic-variance-limited measurement of the μ -distortion – temperature cross-spectrum ($\mu \times T$) could principally limit local-type PNG to $f_{\text{NL}} \lesssim 10^{-3}$. Furthermore, because the μ -distortion signal is created by perturbations with wavenumbers $k \approx 50 - 10^4 \text{ Mpc}^{-1}$, scales that are inaccessible by other cosmological probes [21, 118, 119], constraints on f_{NL} from $\mu \times T$ are highly complementary to constraints on large-scale PNG from primary CMB or even large-scale structure.

In Refs. [70, 28], the authors performed realistic forecasting for the limits on f_{NL} using μ -distortion anisotropy measurements from proposed space missions, including the effects of foreground contamination. Similar methods have also been used to derive limits on PNG from *Planck* [120, 24] and COBE/FIRAS [25], with the most stringent limit, $f_{\text{NL}} \lesssim 6800$ (95% C.L.), coming from Ref. [24].

In this work we consider the potential of measuring μ -distortion anisotropy from the ground. Ground-based measurements have not been considered in previous works, mainly because the calibration and stability requirements for measuring the global μ signal are so stringent that it has been assumed that only in the exquisitely stable environment of space, with no intervening atmosphere, would such a measurement be possible. Following Ref. [68], we point out in this work that measuring the μ -distortion *anisotropy* is a fundamentally different task. Since anisotropy measurements can be made with a differencing radiometer, most of the stability and calibration requirements are converted to requirements on instantaneous sensitivity.

In this work we focus on the potential of the upcoming CMB-S4 experiment [9] to measure μ -distortion anisotropy—specifically, the correlation between μ and CMB temperature and polarization anisotropies. We calculate the constraints on f_{NL} from CMB-S4 $\mu \times T$ and $\mu \times E$

measurements, first considering raw sensitivity alone then adding the effects of atmospheric contamination and foregrounds.

The paper is organized as follows. We discuss the theoretical background in Section 5.2; we describe the methods we use for forecasting in Section 5.3; we present the survey configuration for CMB-S4 in Section 3.4; we present our results in Section 5.4; and we conclude in Section 3.6.

3.2 Background

In this section, we review some of the key theoretical concepts in calculating and forecasting the $\mu \times T$ and $\mu \times E$ correlation arising from squeezed-limit non-Gaussianity.

3.2.1 Power Spectrum and Bispectrum

In the limit of purely Gaussian-distributed primordial curvature perturbations $\zeta(\vec{k})$, the auto-correlation of these perturbations is given by

$$\langle \zeta_{\vec{k}_1} \zeta_{\vec{k}_2} \rangle = (2\pi)^3 \delta(\vec{k}_1 + \vec{k}_2) P_\zeta(k_1). \quad (3.3)$$

The power spectrum $P_\zeta(k)$ of these perturbations from single-field inflation models is given by

$$P_\zeta(k) = \frac{2\pi^2}{k^3} \Delta_\zeta^2(k), \quad \Delta_\zeta^2(k) \equiv A_s \left(\frac{k}{k_p} \right)^{n_s-1}. \quad (3.4)$$

Here we use best-fit numbers from the *Planck* 2018 data release [121]: $A_s = 2.1 \times 10^{-9}$; and $n_s = 0.965$, reported for a pivot scale $k_p = 0.05 \text{ Mpc}^{-1}$.

In the squeezed limit, where the global power spectrum of two modes on short scales k_S are modulated by a long-wavelength mode k_L , the long-wavelength modes generate a local non-Gaussianity and the power spectrum becomes position-dependent

$$P_\zeta(k_S, x) = P_\zeta(k_S) \left[1 + \frac{d \ln P_\zeta(k_S)}{d \zeta_L} \zeta_L(x) \right]. \quad (3.5)$$

One common measure of the level of non-Gaussianity in the curvature distribution is the Fourier-domain statistic known as the curvature bispectrum $B_\zeta(k_1, k_2, k_3)$. Analogous to the power spectrum and the auto-correlation, the bispectrum is defined through its relationship to the three-point correlation function:

$$\langle \zeta_{\vec{k}_1} \zeta_{\vec{k}_2} \zeta_{\vec{k}_3} \rangle = (2\pi)^3 \delta(\vec{k}_1 + \vec{k}_2 + \vec{k}_3) B_\zeta(k_1, k_2, k_3). \quad (3.6)$$

In the squeezed limit, $k_1/k_3 \rightarrow 0$, the bispectrum can be expressed as in Eq. (3.1), with f_{NL} parameterizing the amount of local non-Gaussianity coupling the power spectrum at long and short wavelengths. In the squeezed limit, f_{NL} obeys a consistency relationship with the primordial power spectrum [71]:

$$\lim_{k_L/k_S \rightarrow 0} \frac{12}{5} f_{\text{NL}}(k_L, k_S, k_S) = -\frac{d \ln \Delta_\zeta^2(k_S)}{d \ln k_S}. \quad (3.7)$$

For slow-roll single-field inflation, the consistency relationship becomes

$$\lim_{k_L/k_S \rightarrow 0} \frac{12}{5} f_{\text{NL}}(k_L, k_S, k_S) = 1 - n_s. \quad (3.8)$$

For the specific observable $\mu \times T$, however, the type of PNG produced in single-field inflation results in a vanishingly small signal, far below what is predicted by the consistency relation [1, 18]. This means any measurable signal of $\mu \times T$ would rule out single-field inflation models. We review the reasoning behind this result in Section 3.2.2.

3.2.2 μ cross-correlations

Anisotropies of the CMB probe the primordial curvature perturbations of inflation. CMB anisotropies are typically decomposed into spherical harmonics. For example, the real-space temperature anisotropy field $\Theta(\hat{n}) = \delta T(\hat{n})/T$, can be decomposed into spherical harmonics $\Theta(\hat{n}) = \sum_{\ell m} a_{\ell m}^T Y_{\ell m}(\hat{n})$. The spherical harmonic coefficients of the various CMB anisotropy

fields are related to (Fourier-space) primordial curvature perturbations through

$$a_{\ell m}^X = 4\pi i^{-\ell} \int_{\vec{k}} e^{i\vec{k}\cdot\vec{x}} \zeta(\vec{k}) \Delta_{\ell}^X(k) Y_{\ell m}^*(\vec{k}), \quad (3.9)$$

where X denotes the type of CMB anisotropy field (we are considering only T and E here), and $\Delta_{\ell}^X(k)$ is the transfer function connecting primordial perturbations to CMB anisotropies. The coefficients can then be correlated with each other:

$$\left\langle \left(a_{\ell m}^X \right)^* a_{\ell' m'}^Y \right\rangle = \delta_{\ell\ell'} \delta_{mm'} C_{\ell}^{XY}, \quad (3.10)$$

where C_{ℓ}^{XY} is the angular cross-power spectrum of CMB anisotropy fields X and Y . Observations of μ -distortion anisotropies can also be decomposed into spherical harmonics:

$$a_{\ell m}^{\mu} = 4\pi i^{-\ell} \int_{\vec{k}} e^{i\vec{k}\cdot\vec{x}} \mu(\vec{k}) \Delta_{\ell}^{\mu}(k) Y_{\ell m}^*(\vec{k}), \quad (3.11)$$

where $\Delta_{\ell}^{\mu}(k)$ is the transfer function capturing the leading order evolution of μ distortions anisotropies [69, 17, 89, 18]:

$$\Delta_{\ell}^{\mu}(k) \approx e^{-k^2 / (q_{\mu, D}^2(z_{\text{rec}}))} j_{\ell}(k\Delta\eta), \quad (3.12)$$

where $q_{\mu, D}(z_{\text{rec}}) \approx 0.11 \text{ Mpc}^{-1}$ [17] is the diffusion damping scale of the μ anisotropies, and $\Delta\eta \equiv \eta_0 - \eta_{\text{rec}}$, where η_0 and η_{rec} are the conformal times corresponding to redshifts $z = 0$ and recombination, respectively.

A few words about the μ -distortion transfer function approximation in Eq. (3.12) are in order. First, so far no self-consistent treatment of the distortion sourcing or evolution in the context of PNG has been carried out. The framework for formulating this problem is now in principle available [122, 123, 124], but a detailed solution and discussion is beyond the scope of the present work. Second, from previous works [69, 17] and also recently [123],

it is clear that the distortion transfer functions behave differently than temperature modes. The former undergo spectral and spatial evolution (akin to a rotation in spectral parameters under each scattering event), while the latter (at first order in perturbation theory) evolve purely spatially. As such, the propagation of distortion modes and the related distortion damping process differs. The quoted damping scale $q_{\mu,D}(z_{\text{rec}}) \approx 0.11 \text{ Mpc}^{-1}$ is based on μ -distortion evolution in the tight-coupling limit where no Doppler terms appear in the distortion dipole [69, 17]. As such, no standing waves form, and distortion damping occurs directly in the dipole. This is in stark contrast to CMB temperature perturbations, which damp primarily in the quadrupole [7], implying $q_{T,D}(z_{\text{rec}}) \approx 0.14 \text{ Mpc}^{-1}$. However, even if a self-consistent treatment of the problem should certainly be carried out in the future using the formalism of Refs. [122, 123, 124], overall the precise value of the damping scale does not affect our conclusions at a significant level. We demonstrate this explicitly in Appendix A.2. Furthermore, we anticipate other effects related to the sourcing of the distortion anisotropies, window-function approximations and other simplifications to be more significant [17, 124].

While CAMB includes damping of CMB anisotropies from the finite width of the surface of last-scattering (SLS), we also do not include this effect for μ anisotropies. Properly accounting for the effect would again require performing a full spectro-spatial calculation of the μ thermalization history [122, 123, 124], which is beyond the scope of this paper. We estimate how the additional damping, along with other possible modifications to μ damping, would change our results in Appendix A.2. We find that our constraints on f_{NL} are within 15% of our fiducial results in all cases, suggesting the effect of the additional damping will be minor for μ anisotropies.

The angular correlation of μ distortions with CMB anisotropies is given by

$$\begin{aligned}
\langle a_{\ell m}^{\mu} (a_{\ell' m'}^X)^* \rangle &= \delta_{\ell \ell'} \delta_{m m'} C_{\ell}^{\mu X} = \\
(4\pi)^2 i^{-\ell + \ell'} \int_{\vec{k}_S} \int_{\vec{k}_L} e^{i(\vec{k}_S - \vec{k}_L) \cdot \vec{x}} \times \\
\langle \mu(\vec{k}_S) \zeta(-\vec{k}_L) \rangle \Delta_{\ell}^{\mu}(\vec{k}_S) \Delta_{\ell'}^X(\vec{k}_L) Y_{\ell m}^*(\hat{k}_S) Y_{\ell' m'}(\hat{k}_L),
\end{aligned} \tag{3.13}$$

where we have explicitly indicated here that, because of the very different transfer functions, this correlation probes the connection of the curvature power on very small scales through $\mu(\vec{k}_S)$ with the large-scale curvature $\zeta(\vec{k}_L)$. In other words, this measurement is sensitive to the correlation of a large-scale mode with two extremely small-scale modes, i.e., the bispectrum in the ultra-squeezed limit. The non-Gaussianities being probed are at scales of $50 \text{ Mpc}^{-1} \lesssim k_\mu \lesssim 1 \times 10^4 \text{ Mpc}^{-1}$, much smaller than scales probed by *Planck* measurements of the primary CMB [109].

The ensemble average of $\langle \mu(\vec{k}_S) \zeta(-\vec{k}_L) \rangle$ is

$$\langle \mu(\vec{k}_S) \zeta(-\vec{k}_L) \rangle = \langle \mu \rangle P_\zeta(k_L) \frac{12}{5} f_{\text{NL}}. \quad (3.14)$$

We use this to re-express $C_\ell^{\mu X}$ as

$$C_\ell^{\mu X} = \frac{24 \langle \mu \rangle}{5\pi} f_{\text{NL}} \int_0^{+\infty} dk P_\zeta(k) k^2 \Delta_\ell^\mu(k) \Delta_\ell^X(k). \quad (3.15)$$

The angular power spectra $C_\ell^{\mu T}$ and $C_\ell^{\mu E}$ for $\langle \mu \rangle = 2 \times 10^{-8}$ and $f_{\text{NL}} = 1$ are plotted in Figure 3.1.¹

$\mu \times T$ in single-field inflation

As discussed in Section 3.2.1, even in single-field inflation, squeezed-limit non-Gaussianity can be generated at the level of $f_{\text{NL}} \sim 1 - n_s$, where n_s is the spectral index of the curvature power spectrum [71]. The mechanism that produces this squeezed-limit non-Gaussianity, however, produces vanishingly small amounts (far smaller than the level of $f_{\text{NL}} = 1 - n_s$) of $\mu \times T$ correlations. For this reason, $\mu \times T$ measurements have the potential to rule out single-field inflation more stringently than other probes. This was recognized by the authors

1. This value for $\langle \mu \rangle$ includes the small negative μ -distortion, $\mu_{\text{cool}} \approx -0.3 \times 10^{-8}$ caused by the adiabatic cooling of baryons [63]. Since this contribution does not partake in the PNG evolution [17], our quoted results are somewhat conservative.




Figure 3.1: Angular cross-power spectrum of $\mu \times T$ and $\mu \times E$ for $\langle \mu \rangle = 2 \times 10^{-8}$ and $f_{\text{NL}} = 1$. Solid lines correspond to positive values while dashed lines correspond to negative values.

of Ref. [18]; we briefly summarize their argument here.

In single-field inflation, squeezed-limit non-Gaussianity is produced by long-wavelength modes that are outside of the Hubble radius modulating the global small scale power spectrum by acting as a modulation of the global scale, or a coordinate transformation, as seen in Eq. (3.5). For μ distortions, this results in a shift in the observed value of μ in direction \hat{n} of the form:

$$\mu(\hat{n}) = \mu(z_f, \vec{x}_{\text{rec}}) + \zeta_L(z_f, \vec{x}_{\text{rec}}) \hat{n} \cdot \nabla_{\hat{n}} \mu(z_f, \vec{x}_{\text{rec}}), \quad (3.16)$$

where $\vec{x}_{\text{rec}} = \hat{n}(\eta_0 - \eta_{\text{rec}})$. However, the value of μ distortion at a given position is a local phenomenon and only traces the amount of power dissipated, which in turn depends on the value of the local small-scale power spectrum, which is invariant in the case of no other

form of modulation. Thus, $\mu(\vec{x}) = \langle \mu \rangle$, and, relating the long wavelength mode ζ_L to the large-scale temperature anisotropy via $\zeta_L = -\frac{9}{2}\Theta_L$, in the Sachs-Wolfe approximation [18], we can write

$$\begin{aligned}\mu(\hat{n}) &\sim \langle \mu \rangle + \Theta_L \hat{n} \cdot \nabla_{\hat{n}} \langle \mu \rangle, \\ &= \langle \mu \rangle + 0,\end{aligned}\tag{3.17}$$

This makes it clear that μ distortions from diffusion damping cannot correlate with a long wavelength temperature mode via this mechanism:

$$\langle \mu T \rangle \sim \langle \Theta_L \rangle \langle \mu \rangle + \langle \Theta_L^2 \rangle \nabla_{\hat{n}} \langle \mu \rangle = \langle \Theta_L \rangle \langle \mu \rangle = 0.\tag{3.18}$$

This means that $\mu \times T$ is vanishingly small from non-Gaussianity in single-field inflation, and the only terms that survive are ones with $(k_L/k_S)^2$ suppression [18]. Therefore, any measurable signal of $\mu \times T$ would rule out single-field models. This similarly applies to $\mu \times E$.

3.3 Methods

Our aim is to forecast constraints on f_{NL} using joint analysis of $\mu \times T$ and $\mu \times E$ correlations from CMB-S4. In this section, we will describe the key techniques we use to forecast sensitivity. First we describe the Fisher-matrix formalism used to convert μ , T , and E power spectra to constraints on f_{NL} . We then describe the component-separation technique that allows us to predict μ and T or E power spectra given the total signal + noise + foreground covariance matrix for a particular experimental configuration. Finally, we describe how we model the noise and foreground contributions to the band-band covariance matrix.

3.3.1 Fisher Matrix

If we would like to know how accurately we can measure a given parameter p_i in a data set, we can assume the likelihood \mathcal{L} of measuring the parameter follows a Gaussian distribution:

$$\mathcal{L} \propto \exp \left[-\frac{1}{2} (p_i - \hat{p}_i) F_{ij} (p_j - \hat{p}_j) \right], \quad (3.19)$$

where p_i is a fiducial value of the parameter and \hat{p}_i is the measured value. F_{ij} is the Fisher matrix, which captures the covariance of measured parameters. We discuss the validity of the assumption of Gaussian likelihood in Section 3.3.2.

We make the approximation that the amplitude of the $\mu \times T$ and $\mu \times E$ spectra are controlled by a single free parameter f_{NL} and write

$$C_\ell^{\mu X} = f_{\text{NL}} C_\ell^{\mu X} \Big|_{f_{\text{NL}}=1}. \quad (3.20)$$

The Fisher “matrix” in this case is a scalar:

$$\begin{aligned} F_{ij} &= -\frac{\partial^2 \ln \mathcal{L}}{\partial p_i \partial p_j} \\ &= -\frac{\partial^2 \ln \mathcal{L}}{\partial f_{\text{NL}}^2} \delta_{ij} \equiv F, \end{aligned} \quad (3.21)$$

and the 1σ uncertainty on f_{NL} is

$$\sigma(f_{\text{NL}}) = \frac{1}{\sqrt{F}}. \quad (3.22)$$

The expected noise at a given multipole ℓ is given by

$$\begin{aligned} \sigma_\ell^2 &= \left\langle \left(C_\ell^{\mu X} \right)^2 \right\rangle - \left\langle C_\ell^{\mu X} \right\rangle^2 \\ &= \frac{C_\ell^{\mu\mu} C_\ell^{XX} + \left(C_\ell^{\mu X} \right)^2}{(2\ell + 1) f_{\text{sky}}}, \end{aligned} \quad (3.23)$$

where f_{sky} is the fraction of the full sky observed in the survey. Our fiducial model for

forecasting is $C_\ell^{\mu X} = 0$, and we neglect the contribution of the noise part of $C_\ell^{\mu X}$ to the variance, as it will always be much smaller than the product of $C_\ell^{\mu\mu}$ and C_ℓ^{XX} . The likelihood is thus

$$-2 \ln \mathcal{L} = \sum_{\ell=\ell_{\min}}^{\ell_{\max}} (2\ell + 1) f_{\text{sky}} \frac{\left(f_{\text{NL}} C_\ell^{\mu X} \Big|_{f_{\text{NL}}=1} \right)^2}{C_\ell^{\mu\mu} C_\ell^{XX}}. \quad (3.24)$$

For the $C_\ell^{\mu X}$ cross-spectrum (the numerator or signal part of the Fisher matrix calculation), we use the formulation in Eq. (3.15). For our fiducial forecasts, we adopt $\langle \mu \rangle = 2 \times 10^{-8}$. For the denominator or noise part of the Fisher calculation, we note that the auto-power-spectra C_ℓ^{XX} and $C_\ell^{\mu\mu}$ can be separated into signal and noise terms

$$C_\ell = C_\ell^S + C_\ell^N, \quad (3.25)$$

where S and N denote signal and noise, respectively. At CMB-S4 noise levels, measurements of both C_ℓ^{TT} and C_ℓ^{EE} are signal-dominated for $\ell \leq 2000$ (at which point our constraints on f_{NL} are well saturated, see Figure 3.6). Therefore, we neglect the effects of noise and foregrounds on our temperature and E -mode anisotropy maps, and can make the approximation $C_\ell^{XX} \approx C_\ell^{XX,S}$. On the other hand, measurements of $C_\ell^{\mu\mu}$ will be noise-dominated for the foreseeable future, such that $C_\ell^{\mu\mu} \approx C_\ell^{\mu\mu,N}$. Calculations for $C_\ell^{XX,S}$ are taken from CAMB² [125], while $C_\ell^{\mu\mu,N}$ is dependent on instrument and observation parameters, including instrumental and atmospheric noise levels. We describe how we obtain $C_\ell^{\mu\mu,N}$ from noise and foreground models in the following sections.

The constraining power of $C_\ell^{\mu T}$ and $C_\ell^{\mu E}$ for measuring f_{NL} are comparable to each other. Rather than having independent constraints on f_{NL} from $C_\ell^{\mu T}$ or $C_\ell^{\mu E}$, we can use the fact that they both probe the same underlying correlator $\langle \mu \zeta \rangle$ and perform a joint analysis. The authors of Ref. [89] and Ref. [28] demonstrated that the differing behavior of $C_\ell^{\mu T}$ and $C_\ell^{\mu E}$ with ℓ provides a better constraint on f_{NL} than an independent analysis. We follow their

2. <http://camb.info>

method and modify the likelihood function to include the correlations between T and E in both the signal and the covariance. The final likelihood is then

$$\begin{aligned}
-2 \ln \mathcal{L} \approx & \sum_{\ell=\ell_{\min}}^{\ell_{\max}} \frac{(2\ell+1)f_{\text{sky}}}{C_{\ell}^{\mu\mu,N} \left[C_{\ell}^{TT} C_{\ell}^{EE} - (C_{\ell}^{TE})^2 \right]} \times \\
& \left[C_{\ell}^{TT} \left(f_{\text{NL}} C_{\ell}^{\mu E} |_{f_{\text{NL}}=1} \right)^2 + C_{\ell}^{EE} \left(f_{\text{NL}} C_{\ell}^{\mu T} |_{f_{\text{NL}}=1} \right)^2 - \right. \\
& \left. 2C_{\ell}^{TE} f_{\text{NL}}^2 C_{\ell}^{\mu T} |_{f_{\text{NL}}=1} C_{\ell}^{\mu E} |_{f_{\text{NL}}=1} \right].
\end{aligned} \tag{3.26}$$

3.3.2 Component separation

For the purposes of our forecasting, observed maps of the total intensity of the sky in direction \hat{n} and frequency ν_i , $I_i(\hat{n})$, can be described as a linear combination of the CMB μ and T signals and a noise contribution:

$$I_i(\hat{n}) = a_{\mu,i} s_{\mu}(\hat{n}) + a_{T,i} s_{\text{CMB}}(\hat{n}) + n_i(\hat{n}), \tag{3.27}$$

where $a_{\mu,i}$ and $a_{T,i}$ are the μ and temperature spectral energy distributions (SEDs) at different frequency bands i , and $s_{\mu}(\hat{n})$ and $s_{\text{CMB}}(\hat{n})$ are the true underlying μ and temperature anisotropy maps. We treat all astrophysical signals that are not CMB μ or T as noise and include them in n . We can choose to work in spherical harmonic space, defining $I_{\ell m}$ such that $x(\hat{n}) = \sum_{\ell} \sum_m I_{\ell m} Y_{\ell m}(\hat{n})$. We can then rewrite Eq. (3.27) as

$$I_{\ell m,i} = a_{\mu,i} s_{\ell m,\mu} + a_{T,i} s_{\ell m,\text{CMB}} + n_{\ell m,i}. \tag{3.28}$$

Traditionally, data from CMB experiments are calibrated such that maps in all frequency bands have the same response to primary CMB temperature anisotropy—i.e., the maps are in units of CMB fluctuation temperature ΔT or fractional CMB fluctuation $\Delta T/T$. In the

latter case, the CMB temperature SED $a_{T,i}$ is given by

$$a_{T,i} = T_{\text{CMB}} \quad (3.29)$$

for all bands. We follow Ref. [70] and approximate the μ -distortion SED at frequencies of 20 GHz and above as

$$a_{\mu,i} = T_{\text{CMB}} \left(\frac{1}{2.19} - \frac{1}{x_i} \right), \quad (3.30)$$

where

$$x \equiv \frac{h\nu}{k_{\text{B}}T_{\text{CMB}}}. \quad (3.31)$$

We can obtain a temperature-free μ map or its spherical harmonic transform through component separation using a constrained internal linear combination method (CILC) [126]. This method takes advantage of the known SEDs of temperature and μ -distortion anisotropies and calculates weights \mathbf{w} that, when applied to observed frequency maps, result in a T -free μ map and a μ -free T map. For example, if we assign the μ weights to the $i = 0$ component of \mathbf{w}_{ij} , then the T -free μ map is given by

$$\begin{aligned} \hat{\mu}_{\ell m}^{T\text{-free}} &= \sum_i \mathbf{w}_{0i} I_{\ell m, i} \\ &= 1 * s_{\mu} + 0 * s_{\text{CMB}} + \sum_i \mathbf{w}_{0i} n_{\ell m, i} \end{aligned} \quad (3.32)$$

It is shown in Ref. [70] that the weights that enforce unit response to μ distortions and zero response to temperature anisotropy, and minimize total variance, are given by

$$\mathbf{w}^T = \mathbf{e}^t \left(\mathbf{A}^T \mathbf{C}^{-1} \mathbf{A} \right)^{-1} \mathbf{A}^T \mathbf{C}^{-1} \quad (3.33)$$

where

$$\mathbf{A} = \begin{pmatrix} \mathbf{a}_\mu & \mathbf{a}_{\text{CMB}} \end{pmatrix} \quad (3.34)$$

$$\mathbf{e}^T = \begin{pmatrix} 1 & 0 \end{pmatrix}, \quad (3.35)$$

$\mathbf{C} = \mathbf{C}_\ell^{ij}$ is the frequency-frequency ℓ -space covariance matrix, and T denotes transpose. The above weights can be generalized to null additional components \mathbf{b}_m by generalizing $\mathbf{A} \rightarrow \begin{pmatrix} \mathbf{a}_\mu & \mathbf{a}_{\text{CMB}} & \mathbf{b}_1 & \dots & \mathbf{b}_m \end{pmatrix}$ and $\mathbf{e}^T \rightarrow \begin{pmatrix} 1 & 0 & 0 & \dots & 0 \end{pmatrix}$. We note that the weights depend on multipole number ℓ , i.e., $\mathbf{w} \rightarrow \mathbf{w}_\ell$. For simplicity of notation, we leave the ℓ dependence implicit and continue to use \mathbf{w} .

The reduction of the full $\mathbf{C}_{\ell m \ell' m'}^{ij} = \langle I_{\ell m}^i I_{\ell' m'}^j \rangle$ to \mathbf{C}_ℓ^{ij} rests on the assumption that all sources of variance in the maps are isotropic, stationary, and Gaussian. With detector noise only and assuming no correlations between detector noise at different frequencies, the covariance matrix is diagonal ($\mathbf{C}_\ell^{ij} = \mathbf{C}_\ell^{ii} \delta_{ij}$). Including foregrounds and atmosphere introduces correlated fluctuations between frequency bands and requires the full frequency-frequency matrix.

The assumption of Gaussianity and statistical isotropy is reasonably satisfied by detector noise, atmospheric emission, and extragalactic foregrounds, but not especially well by Galactic foregrounds. The CMB-S4 “ultra-deep” survey, which is the main survey we present forecasts for in this work (see Section 3.4) is located in an area of the sky with very low Galactic foreground emission. Furthermore, non-Gaussianity in the true foreground emission not accounted for in the covariance will not bias the component separation, it will only make it slightly suboptimal. We note that the Gaussianity of the likelihood (Eq. 3.27) depends on the power spectrum of the noise sources being Gaussian, not the pixel values or modes of the noise sources themselves, which is a lighter burden owing to the Central Limit Theorem.

We can then obtain $C_\ell^{\mu\mu}$ by applying the μ -distortion weights to the frequency-frequency

covariance matrix:

$$C_\ell^{\mu\mu} = \sum_{ij} \mathbf{w}_{0i} \mathbf{C}_\ell^{\text{ij}} \mathbf{w}_{0j}. \quad (3.36)$$

We discuss the various contributions to the frequency-frequency covariance matrix $\mathbf{C}_\ell^{\text{ij}}$ in the following sections.

3.3.3 Instrument noise

Assuming detector noise that is white (uncorrelated between time samples), uniform sky coverage, and a Gaussian instrument beam or point-spread function, the statistics of the map noise for frequency band i in spherical harmonic or ℓ space are given by

$$n_\ell^i = N^i e^{\ell^2 \theta_i^2 / (16 \ln(2))}, \quad (3.37)$$

where N_i and θ_i is are the white noise level in the map and the angular resolution (beam FWHM) for frequency band i . The contribution to the frequency-frequency covariance matrix from this source will be

$$\begin{aligned} \mathbf{C}_\ell^{\text{ij},\text{N}} &= \langle n^i n^j \rangle \\ &= N^i N^j e^{\ell^2 \theta_i \theta_j / (8 \ln(2))} \end{aligned} \quad (3.38)$$

and, for detector noise that is uncorrelated between bands,

$$\mathbf{C}_\ell^{\text{ij},\text{N}} = (N^i)^2 e^{\ell^2 \theta_i^2 / (8 \ln(2))} \delta_{ij}, \quad (3.39)$$

where δ_{ij} is the Kronecker delta. For white detector noise only, the μ -distortion power spectrum reduces to

$$C_\ell^{\mu\mu} = \sum_i \mathbf{w}_{0i}^2 (N^i)^2 e^{\ell^2 \theta_i^2 / (8 \ln(2))}. \quad (3.40)$$

3.3.4 Atmosphere

A major source of noise that must be considered in ground-based observations of the CMB is the emission from blobs of poorly mixed water vapor in the Earth's atmosphere (i.e., clouds). For detailed discussions of this effect and measurements of the impact at various sites, see, e.g., Refs. [127], [128], and [129]. The spectrum of cloud sizes is such that the noise power from this source is much larger at large angular scales, and it is often modeled as a power law in ℓ (e.g., Ref. [130]). The total detector + atmosphere noise power in frequency band i can then be parameterized with three numbers, namely the white noise level N_i , the multipole value at which the detector and atmosphere noise levels are equal $\ell_{\text{knee},i}$, and the power-law index of the atmosphere noise $\alpha^{\text{atmo},i}$.

$$\mathbf{C}_\ell^{\text{ij},\text{N}} = \left[1 + \left(\frac{\ell_{\text{knee},i}}{\ell} \right)^{\alpha^{\text{atmo},i}} \right] (N^i)^2 e^{\ell^2 \theta_i^2 / (8 \ln(2))} \delta_{ij}. \quad (3.41)$$

Implicit in the formulation of Eq. (4.44) is the assumption that the atmospheric noise is uncorrelated between bands. In fact, nearly the opposite is the case, at least for detectors for which the beam patterns mostly overlap at the height of the atmospheric emission. For example, internal CMB-S4 analysis of data from the SPT-3G receiver on the South Pole Telescope (SPT, [131, 132]) found that for detectors in different frequency bands but co-located in a focal-plane pixel, the long-timescale fluctuations in the time-ordered data were over 99% correlated. To model this, we can introduce an atmospheric correlation parameter η and rewrite the noise + atmosphere contribution to the frequency-frequency covariance

matrix as

$$\begin{aligned}
\mathbf{C}_\ell^{\text{ij,N}} = & \left[1 + \left(\frac{\ell_{\text{knee},i}}{\ell} \right)^{\alpha^{\text{atmo},i}} \right] \times \\
& (N^i)^2 e^{\ell^2 \theta_i^2 / (8 \ln(2))} \delta_{ij} + \\
& + \left(\frac{\ell_{\text{knee},i}}{\ell} \right)^{\alpha^{\text{atmo},i}/2} \left(\frac{\ell_{\text{knee},j}}{\ell} \right)^{\alpha^{\text{atmo},j}/2} \times \\
& N^i N^j e^{\ell^2 \theta_i \theta_j / (8 \ln(2))} \pi_{ij},
\end{aligned} \tag{3.42}$$

where $\pi_{ij} = \eta(1 - \delta_{ij})$. In Sec. 5.4, we will show forecasts using values of η ranging from 0 to 1. As the correlation η increases, the component separation algorithm defined in Sec. 3.3.2 is more effective in reducing the atmospheric contribution to the final $C_\ell^{\mu\mu}$ covariance.

3.3.5 Foregrounds

In our forecasting pipeline, we also consider the effects of foregrounds. Foregrounds contribute to the frequency-frequency covariance matrix as linear additions (because they are not correlated with the other sources of covariance) such that

$$\mathbf{C}_\ell^{\text{ij}} = \mathbf{C}_\ell^{\text{ij,N}} + \mathbf{C}_\ell^{\text{ij,fore}} \tag{3.43}$$

For this work, we approximate all foregrounds as 100% correlated across frequency bands (though we approximate each foreground source as uncorrelated with the others), such that for a given foreground type (call it “type X”)

$$\mathbf{C}_\ell^{\text{ij,X}} = \sqrt{C_\ell^{ii,\text{X}} C_\ell^{jj,\text{X}}}. \tag{3.44}$$

Foreground sources can be separated into Galactic and extragalactic sources, and we treat each of these in turn below. Often in the literature, foreground amplitudes and ℓ -space

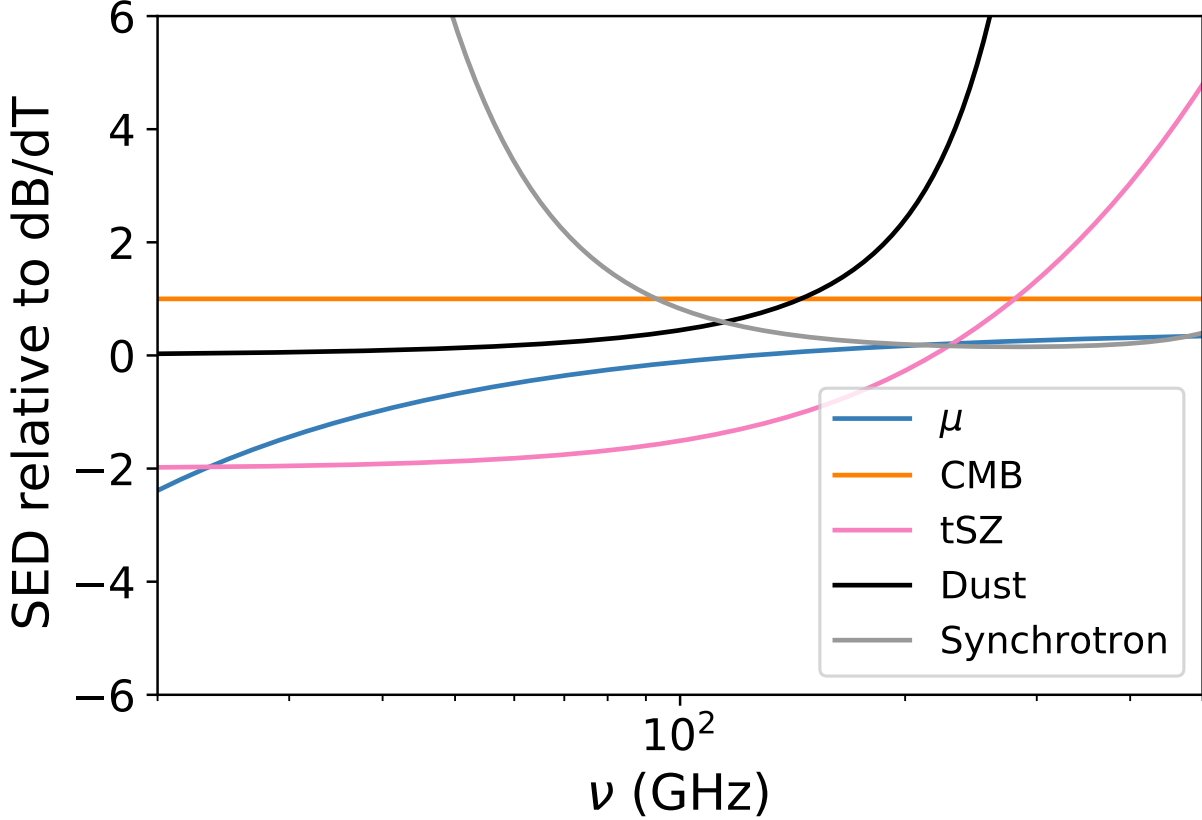


Figure 3.2: Spectral shape of CMB anisotropy, μ distortions, and various Galactic and extragalactic foregrounds, all divided by the CMB anisotropy SED and scaled by an arbitrary amplitude. For frequencies $\nu < 100$ GHz, we see that Galactic synchrotron sharply rises, even relative to μ distortions. At frequencies $\nu > 100$ GHz, Galactic dust and the thermal SZ effect start growing as the μ -distortion spectrum starts leveling off.

behavior are quoted in $D_\ell = \frac{\ell(\ell+1)}{2\pi} C_\ell$. When we adopt such parameterizations, we keep the description in D_ℓ but convert to C_ℓ when actually implementing the model. We show the SEDs of the primary foreground sources, along with the T and μ SEDs, in Figure 3.2.

Galactic Sources

The primary sources of Galactic contamination at CMB observing frequencies are thermal dust emission and synchrotron emission. Interstellar dust heated by starlight emits as a quasi-thermal modified blackbody. We follow Ref. [130] and parameterize the frequency

behavior and ℓ -space shape of thermal dust emission as

$$D_{\ell,\nu_1,\nu_2}^{\text{d}} = D_{80,\nu_0}^{\text{d}} \epsilon_{\nu_1,\nu_2} \left(\frac{\nu_1 \nu_2}{\nu_0^2} \right)^{\alpha^{\text{d}}} \times \frac{B_{\nu_1}(T_{\text{d}}) B_{\nu_2}(T_{\text{d}})}{B_{\nu_0}^2(T_{\text{d}})} \left(\frac{\ell}{80} \right)^{-0.4}, \quad (3.45)$$

where B is Planck’s law, D_{80,ν_0}^{d} is the value of D at $\ell = 80$ and the reference frequency ν_0 , T_{d} is the dust temperature, and α^{d} is the dust spectral index. Following Ref. [133], we define $D_{80}^{\text{d}} = 3.253 \mu\text{K}^2$, $\alpha^{\text{d}} = 1.6$, $T_{\text{d}} = 19.6$, and $\nu_0 = 145$ GHz. Finally, ϵ_{ν_1,ν_2} relates the brightness of CMB fluctuations at ν_1 , ν_2 , and ν_0 :

$$\epsilon_{\nu_1,\nu_2} \equiv \frac{\left[\frac{dB(\nu_0, T_{\text{CMB}})}{dT} \right]^2}{\frac{dB(\nu_1, T_{\text{CMB}})}{dT} \frac{dB(\nu_2, T_{\text{CMB}})}{dT}}. \quad (3.46)$$

Because we have defined our μ and T SEDs assuming that the input maps are calibrated to CMB fluctuation amplitudes, we must also account for this in the foreground modeling. The ratio $1/(dB/dT|_{\nu, T_{\text{CMB}}})$ converts the source radiance of a given foreground at frequency ν to an equivalent CMB temperature anisotropy.

We use a similar parameterization for Galactic synchrotron, again following Ref. [130]

$$D_{\ell,\nu_1,\nu_2}^{\text{sync}} = D_{80,\nu_0}^{\text{sync}} \epsilon_{\nu_1,\nu_2} \left(\frac{\nu_1 \nu_2}{\nu_0^2} \right)^{\alpha^{\text{sync}}} \left(\frac{\ell}{80} \right)^{-0.4}, \quad (3.47)$$

and we again adopt values from Ref. [133]: $D_{80}^{\text{sync}} = 0.005 \mu\text{K}^2$, $\alpha^{\text{sync}} = -1.10$, and $\nu_0 = 93$ GHz. We note that both the synchrotron and dust amplitudes used here are estimated for the CMB-S4 “ultra-deep” $f_{\text{sky}} = 0.03$ survey (see next section for details).

Because of the shape of the μ -distortion SED, lower-frequency channels are particularly important for recovering the signal, and it is possible that other Galactic foregrounds such as free-free emission and “anomalous microwave emission” (AME) could be important con-

taminants. We investigate the behavior of these additional foregrounds in the CMB-S4 3% sky region using PySM [134], which is based on the *Planck* Sky Model [135]. We find that the AME SED has a double-peaked shape, which we parameterize as

$$f_{\text{AME}}^2(\nu) = \frac{e^{-[\ln(\nu)-\ln(\nu_{a1})]^2/2\sigma_{a1}^2} + Ae^{-[\ln(\nu)-\ln(\nu_{a2})]^2/2\sigma_{a2}^2}}{1 + Ae^{-[\ln(\nu_{a1})-\ln(\nu_{a2})]^2/2\sigma_{a1}^2}}, \quad (3.48)$$

with $\nu_{a1} = 10$ GHz, $\sigma_{a1} = 0.43$ GHz, $\nu_{a2} = 22$ GHz, $\sigma_{a2} = 0.35$ GHz, and $A = 0.0065$. We assume similar ℓ -space behavior as the thermal dust and write

$$D_{\ell,\nu_1,\nu_2}^{\text{AME}} = D_{80,\nu_0}^{\text{AME}} f_{\text{AME}}(\nu_1) f_{\text{AME}}(\nu_2) \left(\frac{\ell}{80}\right)^{-0.4}, \quad (3.49)$$

with $D_{80,\nu_0}^{\text{AME}} = 1.0 \times 10^4 \mu\text{K}^2$ at $\nu_0 = 10$ GHz. The free-free emission in the CMB-S4 3% sky region appears to be dominated by point-like sources (either residual contributions from extragalactic radio sources or emission from compact HII regions). Because CMB-S4 will have the sensitivity and resolution to find and mask such sources, we neglect the contribution of free-free emission in this analysis.

Extragalactic foregrounds

In addition to emission from our Galaxy, we also need to consider foreground emission from extragalactic sources. We treat four independent types of extragalactic foregrounds: the thermal Sunyaev-Zel'dovich (tSZ) effect, the clustered cosmic infrared background (CIB), the unclustered (“shot-noise” or “Poisson”) component of the CIB, and synchrotron-emitting active galactic nuclei (or “radio sources”), the clustering of which is assumed to be negligible. We neglect the contribution from the kinetic Sunyaev-Zel'dovich (kSZ) effect, which has the same frequency spectrum as the CMB and will be nulled by the CILC.

The tSZ effect, a type of y distortion, arises from CMB photons scattering off electrons

in the intra-cluster medium of galaxy clusters and other collapsed structures. This shifts the re-emitted photons to higher frequencies compared to the CMB spectrum. We parameterize the tSZ as

$$D_{\ell, \nu_1, \nu_2}^{\text{tSZ}} = D_{3000, \nu_0}^{\text{tSZ}} T(\ell) \frac{f(\nu_1) f(\nu_2)}{f^2(\nu_0)}, \quad (3.50)$$

where

$$f(\nu) = x \frac{e^x + 1}{e^x - 1} - 4, \quad (3.51)$$

$$x = \frac{h\nu}{k_{\text{B}}T}, \quad (3.52)$$

and $T(\ell)$ is the tSZ template used in Ref. [136]. Following Ref. [137], we adopt $D_{3000, \nu_0}^{\text{tSZ}} = 3.4 \mu\text{K}^2$ at $\nu_0 = 150$ GHz.

Again following Ref. [137], we parameterize the clustered CIB as

$$D_{\ell, \nu_1, \nu_2}^{\text{C}} = D_{3000}^{\text{C}}(\nu_0) \epsilon_{\nu_1, \nu_2} \left(\frac{\nu_1 \nu_2}{\nu_0^2} \right)^{\alpha^{\text{P}}} \left(\frac{\ell}{3000} \right)^{0.8}, \quad (3.53)$$

where $D_{3000}^{\text{C}}(\nu_0) = 3.46 \mu\text{K}^2$, $\alpha^{\text{P}} = 4.27$, and $\nu_0 = 150$ GHz. And we parameterize the Poisson component of the CIB as

$$D_{\ell, \nu_1, \nu_2}^{\text{P}} = D_{3000}^{\text{P}}(\nu_0) \epsilon_{\nu_1, \nu_2} \left(\frac{\nu_1 \nu_2}{\nu_0^2} \right)^{\alpha^{\text{P}}} \left(\frac{\ell}{3000} \right)^2, \quad (3.54)$$

with $D_{3000}^{\text{P}} = 9.16 \mu\text{K}^2$, $\alpha^{\text{P}} = 3.27$, $\nu_0 = 150$ GHz. For this initial work, we do not consider spatial correlation between the tSZ and CIB. This does not cause any direct bias to our final result, as we explicitly null any signal with a tSZ spectrum in our final constraints (see Section 5.4 for details). For more discussion of the effects of tSZ-CIB correlation on $\mu \times T$

measurements, see Ref. [24].

Finally, we parameterize radio source power as

$$D_{\ell, \nu_1, \nu_2}^r = D_{3000}^r(\nu_0) \epsilon_{\nu_1, \nu_2} \left(\frac{\nu_1 \nu_2}{\nu_0^2} \right)^{\alpha^p} \left(\frac{\ell}{3000} \right)^2, \quad (3.55)$$

with $D_{3000}^r(\nu_0) = 0.02 \mu\text{K}^2$, $\alpha^p = -0.7$, $\nu_0 = 150 \text{ GHz}$. This is significantly lower than amplitudes quoted in, e.g., Ref. [137]. This is because the radio Poisson power is dominated by the brightest individual sources in the map, and masking and removing the brightest radio sources will reduce the Poisson term. For an experiment similar to CMB-S4, which will achieve roughly $1 \mu\text{K-arcmin}$ map noise, radio sources can be cleaned down to roughly 0.2 mJy , which reduces their Poisson amplitude to $D_{3000}^r(\nu_0) = 0.02 \mu\text{K}^2$.

3.3.6 Calibration

Historically, measurements of absolute brightness or temperature at microwave/millimeter-wave frequencies have been successfully carried out only by space- or balloon-borne telescopes, because of the stringent requirements on calibration accuracy and stability (e.g., Refs. [138, 139]). Proposed measurements of the absolute μ -distortion amplitude, such as with the PIXIE satellite [140], are designed with similar constraints in mind. It was pointed out, however, in Ref. [68] that measuring the *anisotropy* of μ distortions, and particularly the correlated anisotropy of μ and temperature, effectively converts the absolute calibration requirement to a relative calibration requirement, and the bias on the absolute measurement to an uncertainty on the measurement of anisotropy.

A calibration error in CMB-S4-like data will result primarily in leakage of the much larger temperature signal (or foregrounds) into the component-separated μ map, resulting in a component of $T \times T$ in the $\mu \times T$ cross-spectrum. The $T \times T$ spectrum does not have the same shape as $\mu \times T$ (which crosses zero many times over the ℓ range we consider), so there will be on average no bias from this leakage, just excess variance. We have investigated the

level to which the relative calibration must be known for this variance not to dominate the error on f_{NL} , but a simple scaling argument tells us that, because the relative calibration will come from enforcing equal response to the CMB temperature in every frequency band, the calibration precision in any one band will be equal to the S/N on the CMB temperature anisotropy in that band. Thus, the contribution to uncertainty on f_{NL} from calibration errors will be on the order of the contribution from noise divided by the square root of the number of bands. For this reason, we ignore calibration uncertainty in our main results.

A related concern is the knowledge of the instrumental bandpasses. Even in the limit of perfect relative calibration off the CMB, imperfect knowledge of the instrument bandpasses could lead to a different level of foreground contamination in the μ map than the component-separation algorithm predicts. As with CMB T leakage, this will not generally have the same shape as the $\mu \times T$ spectrum and will thus not cause bias on average. Furthermore, as we are not explicitly projecting out foregrounds in the component separation, the extra leakage of foregrounds into the μ map from bandpass uncertainty is likely to be small. Nevertheless, we will update the forecasting machinery to include this effect in a future paper.

3.4 Survey configuration

Frequency (GHz)	20	27	39	93	145	225	278
Angular resolution (arcmin)	11	8.4	5.8	2.5	1.6	1.1	1.0
White noise level in temperature (μK -arcmin)	9.31	4.6	2.94	0.45	0.41	1.29	3.07
ℓ_{knee} for TT	400	400	400	1200	1900	2100	2100
α for TT	4.2	4.2	4.2	4.2	4.1	4.1	3.9

Table 3.1: Survey configuration for the large-aperture telescope (LAT) in the CMB-S4 ultra-deep survey.

In this work, we are forecasting $\sigma(f_{\text{NL}})$ using the parameters of the upcoming CMB-S4 experiment [9]. There are two major surveys planned for CMB-S4, an “ultra-deep” survey

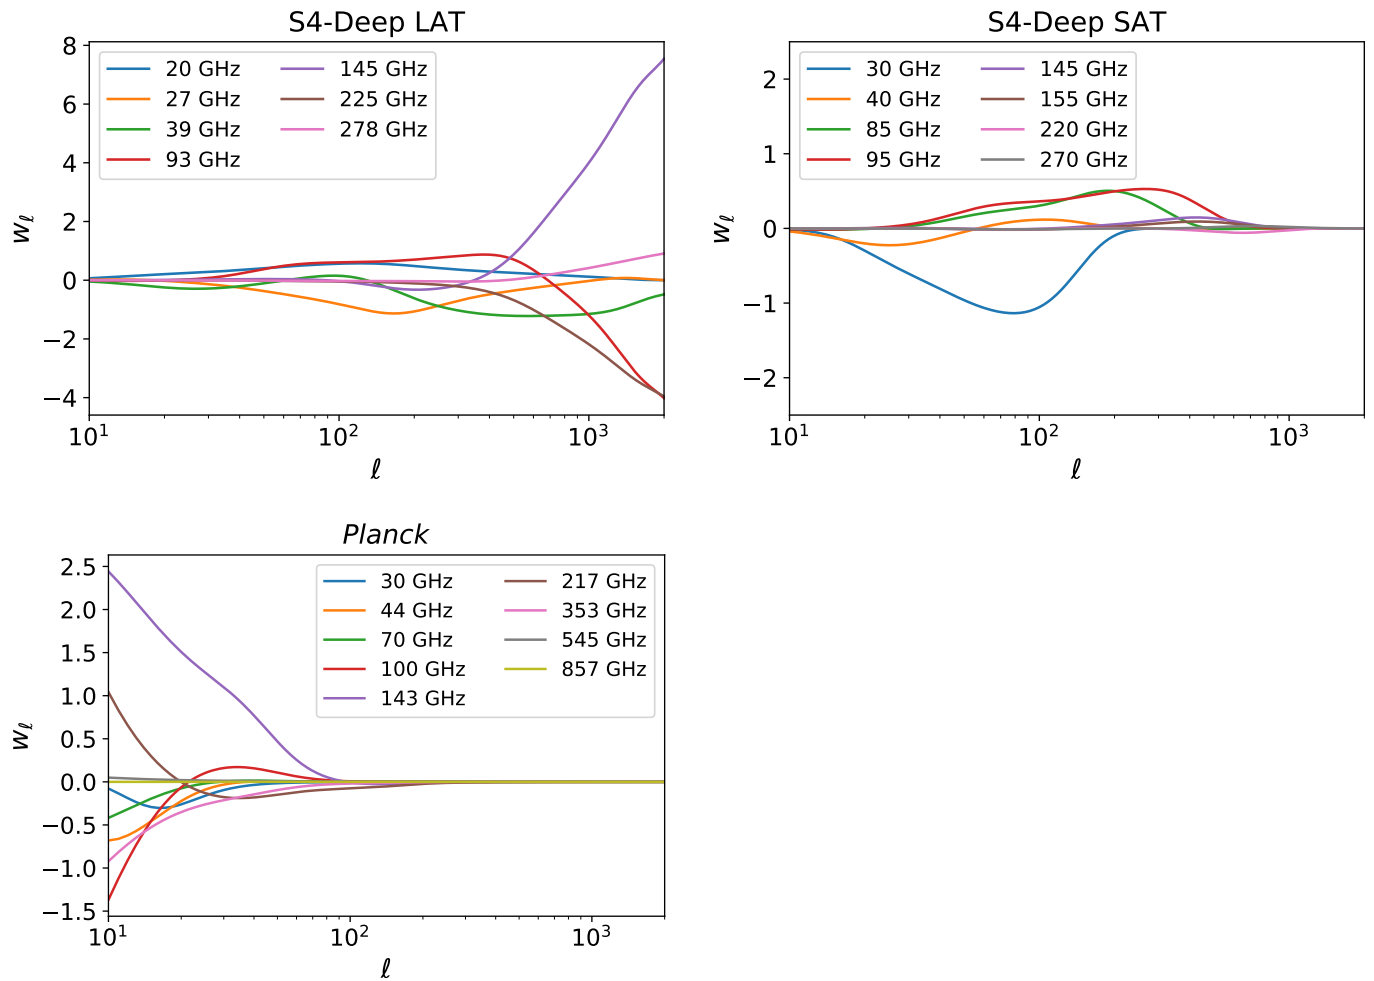


Figure 3.3: Weights used to construct the T -free μ -distortion map (i.e., the μ component of \boldsymbol{w}) for the CMB-S4 ultra-deep patch from different frequency channels, including foregrounds and assuming 99% correlated atmosphere.

Frequency (GHz)	30	40	85	95	145	155	220	270
Angular resolution (arcmin)	72.8	72.8	25.5	22.7	25.5	22.7	13.0	13.0
Noise level in polarization (μK -arcmin)	3.53	4.46	0.88	0.78	1.23	1.34	3.48	5.97
Noise level in temperature (μK -arcmin)	2.50	3.15	0.62	0.55	0.87	0.95	2.46	4.22
ℓ_{knee} for TT	400	400	1200	1200	1900	1900	2100	2100
α for TT	4.2	4.2	4.2	4.2	4.1	4.1	4.1	3.9

Table 3.2: Survey configuration for the small-aperture telescopes (SATs) in the CMB-S4 ultra-deep survey.

Frequency (GHz)	30	44	70	100	143	217	353	545	857
Angular resolution (arcmin)	32.4	27.1	13.3	9.69	7.30	5.02	4.94	4.83	4.64
Noise level for TT (μK -arcmin)	150	162	210	77.4	33.0	46.8	154	815	2.98×10^4

Table 3.3: Angular resolution and noise levels assumed for *Planck* (reproduced from Table 4 of Ref. [8]). Note that we assume *Planck* noise is white down to the ℓ_{min} of our forecasting.

of roughly 3% of the Southern sky, conducted from the South Pole with both small-aperture telescopes (SATs) and a large-aperture telescope (LAT), and a “deep and wide” survey of roughly 60% of the sky, conducted from Chile with LATs only.

Our fiducial forecasting will be for the $f_{\text{sky}} = 0.03$ survey. The $f_{\text{sky}} = 0.6$ survey involves a similar investment of total detector number and observing time, but spread across a sky patch that is 20 times larger. In the limit of detector noise only, this introduces a factor f_{sky} into the covariance matrix, leading to

$$\mathbf{C}_{\ell}^{\text{ij},\text{N}}(f_{\text{sky}}) = \mathbf{C}_{\ell}^{\text{ij},\text{N}}(f_{\text{sky}} = 0.03) \frac{f_{\text{sky}}}{0.03}. \quad (3.56)$$

Because incomplete sky surveys also observe fewer total sky modes and hence have fewer samples of the various power spectra, the Fisher matrix calculations also include an f_{sky} term, as in Eq. (3.23).

As discussed in the previous section, we can make the approximation that $C_{\ell}^{\mu\mu}$ is noise-

dominated, while C_ℓ^{TT} is signal-dominated for $\ell \leq 2000$. This means that $C_\ell^{\mu\mu}$ is the only term in the denominator of the $\mu \times T$ Fisher calculation that has an f_{sky} dependence:

$$-2 \ln \mathcal{L} = \sum_{\ell=2}^{\ell_{\text{max}}} 0.03 (2\ell + 1) \frac{f_{\text{sky}}}{0.03} \times \frac{\left(f_{\text{NL}} C_\ell^{\mu T} \Big|_{f_{\text{NL}}} \right)^2}{(f_{\text{sky}}/0.03) C_\ell^{\mu\mu, \text{N}} \Big|_{f_{\text{sky}}=0.03} C_\ell^{TT}}. \quad (3.57)$$

In this limit, f_{sky} cancels out, leaving our Fisher matrix independent of f_{sky} in the detector-noise-only case. We note that this will be true for any measurement that depends on correlating a very small signal (that is below the detection threshold for a given experiment) with a much larger one (that is measured at high S/N).³ When we add atmosphere and foregrounds, the situation becomes more complicated, because the statistics of the atmospheric noise are different at the two CMB-S4 sites, and because it is more difficult to avoid bright parts of our Galaxy when more sky is observed.

The parameters of the high-resolution (LAT) part of the ultra-deep ($f_{\text{sky}} = 0.03$) survey, including band centers, angular resolutions, and noise levels, are given in Table 3.1. The corresponding values for the degree-scale (SAT) ultra-deep survey are shown in Table 3.2.

The $f_{\text{sky}} = 0.03$ ultra-deep survey for CMB-S4 will be conducted from the South Pole. In terms of atmospheric emission at CMB frequencies, the South Pole is the best large, developed site on Earth [141]. In Tables 3.1 and 3.2, we show the expected values of ℓ_{knee} and α (see Eq. 4.44) for the CMB-S4 bands at the South Pole, derived from CMB-S4 internal analysis of SPT-3G data. We note that the Deep Survey LAT value for ℓ_{knee} for bands below 40 GHz in official CMB-S4 documents is 1200, while we use 400 here. The value of 1200 is a conservative choice made assuming no improvement in atmospheric noise from the lowest SPT band of 95 GHz. The ℓ_{knee} values for the Wide Survey LAT low-frequency bands, which

3. One caveat to this is that the survey needs to be big enough to resolve the largest scale that is important for measuring the signal.

were derived by scaling the atmospheric noise power with levels of precipitable water vapor, are close to 400, and thus we adopt that value for the Deep Survey here. We discuss the impact of this choice in Section 5.4.

In addition to reducing raw noise levels, the authors of Ref. [70] suggest expanding detector frequency coverage, in order to lower $C^{\mu\mu,N}$. To increase the upper frequency range of CMB-S4 measurements, we include *Planck* 2018 data in the forecast, which extends up to 857 GHz. This should provide valuable complementary information that can help mitigate foregrounds at frequencies that CMB-S4 is unable to observe. We implement *Planck* bands in our forecasting pipeline in the same way as CMB-S4 (as independent frequency channels, over the same 3% of the sky as CMB-S4 Deep), using the parameters shown in Table 3.3. In Figure 4.4, we show the components of \mathbf{w} for the T -free μ map contributed from each individual CMB-S4 or *Planck* band, assuming 99% atmospheric correlation and the foregrounds described in the previous section.

In Figure 3.4, we show the expected noise power in the μ map from CMB-S4 ($C_\ell^{\mu\mu,N}$) for various values of atmospheric correlation among frequency bands. In the scenario where frequency bands are fully correlated and the atmosphere can be fully subtracted out, the associated noise curve closely matches the scenario with no atmosphere. However, even minimal decorrelation among detectors significantly worsens $C_\ell^{\mu\mu,N}$ at low multipoles. Residual atmosphere acts as an effective ℓ_{\min} when summing over multiple observation bands to constrain f_{NL} .

We note that currently fielded experiments such as Advanced ACTPol [142] and SPT-3G, as well as the upcoming Simons Observatory (SO) experiment [143] can be approximated in this forecasting framework as versions of either the CMB-S4 Deep or Wide surveys, but with higher noise and (in the case of currently fielded experiments) reduced frequency coverage. In particular, the SO LAT “goal” survey is similar to the CMB-S4 Wide survey with 2-3 times higher noise and slightly less sky area, thus we would thus expect to forecast roughly 2-3 times worse constraints on f_{NL} for SO compared to CMB-S4.




Figure 3.4: $C_{\ell}^{\mu\mu, N}$ for the CMB-S4 ultra-deep survey and various assumptions for the degree of atmospheric correlation between bands. Full correlation among bands closely matches the nominal case of no atmosphere. However, even minimal decorrelation among detectors significantly worsens $C_{\ell}^{\mu\mu, N}$ at low multipoles.

3.5 Results

Using the noise, atmosphere, and foreground parameterizations described in the previous sections, we forecast $\sigma(f_{\text{NL}})$ from the combination of $\mu \times T$ and $\mu \times E$ for the CMB-S4 experiment, combined with data from the *Planck* satellite [8]. For our results, we assume a fiducial value of $f_{\text{NL}} = 1$. We will generally show the cumulative constraint on f_{NL} as a function of ℓ_{max} , the highest multipole value considered in the calculation. In all cases, we assume $\ell_{\text{min}} = 10$, though our final results are not sensitive to this exact choice.

We first show, in Figure 3.5, the signal-to-noise (S/N) as a function of multipole ℓ . The solid lines include detector and atmospheric noise and foreground residual, with color




Figure 3.5: Signal-to-noise (S/N) per multipole for various assumed values of atmospheric correlation between bands. Dashed lines correspond to S/N when only including instrumental and atmospheric noise, while solid lines also include the effects of foregrounds. In all cases, the y SED has been projected out (see text for details).

differentiating the amount of assumed atmospheric correlation between frequency bands. The dashed lines correspond to detector and atmospheric noise only, with no foreground residuals.

For the scenario where atmosphere is perfectly correlated between frequency bands we see that most of our signal is at the lowest ℓ modes. The addition of foregrounds reduces our overall signal by about an order of magnitude, agreeing with previous results [70] that foregrounds are a primary obstacle for experiments constraining μ distortions.

The main foregrounds that impact our constraints are from Galactic sources, such as dust, synchrotron radiation, and AME. Galactic foregrounds have a larger effect on reducing sensitivities to μ , compared to extragalactic sources [28]. Even though several Galactic




Figure 3.6: $1\text{-}\sigma$ constraints on f_{NL} as a function of the maximum multipole used in analysis, ℓ_{max} . Dashed lines correspond to $\sigma(f_{\text{NL}})$ when only including instrumental and atmospheric noise, while solid lines also include the effects of foregrounds. In all cases, the y SED has been projected out (see text for details).

and extragalactic foregrounds are brighter at lower frequencies, the shape of the μ SED becomes more distinct for these frequency bands. This suggests for optimal detection of μ distortions, future experiments should concentrate on improving observations with lower frequency bands rather than higher frequency ones. In spite of foreground residuals, in the limit of 100% atmospheric correlation between bands, we find that the lowest ℓ modes provide the best leverage for constraining f_{NL} .

However, when we introduce any atmospheric decorrelation, residual atmosphere has a more significant impact on our μ maps than residual foregrounds. While foregrounds reduce the S/N across the entire ℓ range, atmosphere suppresses the S/N primarily at low ℓ , where

the raw S/N is highest. This means that we lose significant leverage in constraining f_{NL} from atmospheric contamination.

Another potential contamination to measurements of $\mu \times T$ correlations are y -type spectral distortions. In our default component-separation algorithm, y distortions are treated as noise and not explicitly projected out of the μ or T map. However, the authors of Ref. [89] showed that there are $y \times T$ correlations induced from late Integrated Sachs-Wolfe effects. This means that $y \times T$ correlations will leak into our $\mu \times T$ correlations and potentially bias $\sigma(f_{\text{NL}})$, unless we deproject them explicitly. This is easily accomplished in the algorithm by adding the y -distortion SED as one of the components in the matrix A in Eq. (3.33), but we pay some noise penalty for this.

Figure 3.6 shows our constraints on f_{NL} (including projecting out a y component) after summing over all multipoles up to ℓ_{max} . We consider different values of atmospheric correlation and the presence of foregrounds. We see that at $\ell_{\text{max}} = 2000$, we have effectively saturated our constraints on f_{NL} for all configurations. Therefore, in Tables 3.4 and 3.5, we report $\sigma(f_{\text{NL}})$ for $\ell_{\text{max}} = 2000$.

In the ideal case in which we can ignore atmosphere, foregrounds, and y distortions, we find $\sigma(f_{\text{NL}}) = 48$. The addition of foregrounds worsens our constraints to $\sigma(f_{\text{NL}}) = 461$, or 462 if we project out y . This is better than forecasts for *LiteBIRD*, which is forecasted to achieve $\sigma(f_{\text{NL}}) = 825$ [28]. Adding atmosphere and assuming some level of decorrelation across frequency bands, we find our constraints on f_{NL} noticeably degrade even in the absence of foregrounds. For 1% decorrelation ($\eta = 99\%$) we find $\sigma(f_{\text{NL}}) = 165$, or 231 with projecting out y . The addition of foregrounds further reduces constraints, with $\sigma(f_{\text{NL}}) = 738$ or 750 in the case of 1% atmospheric decorrelation, comparable to the forecast for *LiteBIRD*. We note that, because of the effect of atmosphere on CMB-S4 constraints and the low angular resolution of *LiteBIRD*, it is likely that the f_{NL} constraints from the two experiments will come primarily from independent regions of ℓ space, in which case we will be able to improve the individual constraints by nearly a factor of $\sqrt{2}$ by combining them.

$\sigma(f_{\text{NL}}),$ $\ell_{\text{max}} = 2000$	S4-Deep + <i>Planck</i> (no foregrounds)	S4-Deep + <i>Planck</i>
No atmosphere	48	461
100% correlated	54	463
99% correlated	165	738
90% correlated	236	883
70% correlated	292	1002

Table 3.4: Constraints on f_{NL} from the CMB-S4 ultra-deep survey and *Planck*, considering the effects of atmosphere and foregrounds. In the case where neither are present, $\sigma(f_{\text{NL}}) = 48$. Including atmosphere and foregrounds worsens this constraint by at least an order of magnitude.

Foreground residuals can in principle be reduced through expanded frequency coverage. If the foreground SEDs are smooth and require minimal degrees of freedom to model, with enough observations across unique frequency bands, one can constrain the foreground and CMB SEDs. Pairing CMB-S4 with additional surveys that aim to accurately model CMB foregrounds can greatly improve our constraints on f_{NL} . However, even with reduced foreground residuals, CMB-S4 will be limited by atmospheric noise to $\sigma(f_{\text{NL}}) > 100$, unless the correlation between bands is $> 99\%$.

3.6 Discussion

In this work, we have presented forecasted constraints on f_{NL} , at effective scales of $k \approx 740 \text{ Mpc}^{-1}$, with CMB-S4 using correlations of μ -distortion anisotropies with CMB temperature and E -mode polarization. We find that with this ground-based experiment we are able to achieve comparable results to forecasts for the *LiteBIRD* satellite [28], depending on the amount of atmospheric correlation among frequency bands. CMB-S4 will not have prior knowledge of the amount of atmospheric correlation, which is expected to depend on the optical design and observing strategy. Co-located detectors will likely be more correlated, since they will observe along the same incident angle and will observe similar parts of the atmo-

$\sigma(f_{\text{NL}})$, “null- y ” $\ell_{\text{max}} = 2000$	S4-Deep + <i>Planck</i> (no foregrounds)	S4-Deep + <i>Planck</i>
No atmosphere	68	462
100% correlated	72	465
99% correlated	231	750
90% correlated	332	912
70% correlated	411	1050

Table 3.5: Constraints on f_{NL} when projecting out the y SED, from the CMB-S4 ultra-deep survey and *Planck*. We see in the case of no foregrounds that constraints on f_{NL} are worse than the ones given in Table 3.4. Including foregrounds, we see that constraints are very similar to those obtained if we did not project out the y SED.

sphere. Conversely, detectors located further from each other will observe through slightly different columns of the atmosphere and are expected to be less correlated. Understanding the correlation properties of atmospheric emission between frequency bands is therefore vital for producing realistic forecasts of $\sigma(f_{\text{NL}})$ from μ -distortion anisotropies.

Independent of the details of atmospheric correlation, we have demonstrated that CMB-S4 has the potential to constrain small-scale Gaussianities down to $\sigma(f_{\text{NL}}) \lesssim 1000$ with $\mu \times T$ and $\mu \times E$. We note that correlations between μ -distortion and primary CMB anisotropies represent one of our only probes for understanding non-Gaussianities on extremely small scales or, equivalently, the ultra-squeezed-limit of the bispectrum. This presents an opportunity to constrain inflationary models that predict specific behavior at very small scales. One such area that can be constrained with small-scale non-Gaussianities are inflationary models that contain additional fields. Certain curvaton models can naturally generate $f_{\text{NL}} \gg 1$ on small scales while preserving current upper limits on $f_{\text{NL}}(0.05 \text{ Mpc}^{-1})$ from *Planck* [144]. If massive fields have direct couplings to the inflaton, they can also induce oscillatory features in the squeezed-limit bispectrum at small scales [106]. Modified (non-Bunch-Davies) initial vacuum states can also produce enhanced small-scale non-Gaussianity [68] and can be constrained with $\mu \times T$ and $\mu \times E$ correlations. Finally, $\mu \times T$ and $\mu \times E$ correlations can place

stringent constraints on the running of the spectral index n_s and general scale-dependent non-Gaussianities [145], such as have been proposed to explain recent discoveries from JWST [146].

It is also worth noting that the constraint on f_{NL} from $\mu \times T$ and $\mu \times E$ is linearly proportional to the mean value of μ (cf. Eqs. 3.22 and 3.15), so any non-standard model that boosts $\langle \mu \rangle$ will correspondingly improve the forecasted constraints on f_{NL} [17]. One example of such a model is given in Ref. [23], in which the authors used $\mu \times T$ to constrain primordial black hole (PBH) models. PBH models generically have a rising slope for the power spectrum at small scales, preceded by a region with lower power. For scales close to the ones μ distortions probe, the rising power can enhance $\langle \mu \rangle$, while the dip can generate non-Gaussianities larger than slow-roll predictions. These combined effects mean that upcoming CMB experiments can potentially rule out some PBH models. Work in Ref. [19] has shown that the main signal PBH models imprint on $\mu \times T$ is from local modulation of acoustic dissipation by long wavelength modes, which induce a bias in μ distortions. We will examine this potential signal in future work.

Our analysis probes non-Gaussianities from scalar perturbations $\langle \zeta \zeta \zeta \rangle$. Distortions of the μ type should also be generated by primordial gravitational waves (GW) injecting energy into the photon-baryon fluid. This allows μ to probe the primordial tensor power spectrum. Since gravitational waves are free streaming throughout the radiation-dominated era, tensor $\mu^{(t)}$ -distortions probe scales in the range $1 \text{ Mpc}^{-1} \lesssim k \lesssim 10^6 \text{ Mpc}^{-1}$, complementing the gap between CMB scales and the scales probed by upcoming GW interferometers [147]. Recent work in Ref. [90] has demonstrated that $\mu \times B$ correlations can probe tensor and mixed tensor-scalar non-Gaussianities. Using FIRAS data, Ref. [25] attempted the first measurement of $\mu \times B$, which they found was consistent with zero at existing noise levels.

The forecasts presented here are not the last word on f_{NL} from $\mu \times T$ and $\mu \times E$. As shown in Ref. [18], a cosmic-variance limited experiment can in principle constrain $\sigma(f_{\text{NL}})$ to $\lesssim 10^{-3}$ with $\mu \times T$, although this will also require a significantly improved limit on the

average μ -distortion amplitude to break the degeneracy between f_{NL} and $\langle\mu\rangle$ [17] and ensure that no other source of distortion anisotropies is present [123]. Putting this in context, the best constraints a cosmic-variance limited CMB experiment could place on squeezed-limit non-Gaussianity with $\langle TTT \rangle$ is $\sigma(f_{\text{NL}}) \sim 1$. One promising avenue for improvement on the constraints presented here is suggested in Ref. [70], namely expanding coverage to lower frequency ranges where the relative amplitude of the μ -distortion SED is higher. In particular, the addition of bands at and below 10 GHz would significantly improve detection. One possibility is to combine CMB-S4 with upcoming low-frequency radio surveys such as the Square Kilometer Array (SKA) [148] to improve constraints on f_{NL} , an aspect we will explore in future work.

CHAPTER 4

MEASURING μ DISTORTIONS FROM THE THERMAL SUNYAEV ZELDOVICH EFFECT

This chapter is reproduced from [26]

4.1 Introduction

The cosmic microwave background (CMB) provides us with vital information about the origin and evolution of our observable universe, and of the underlying physical laws that govern it. We have greatly improved measurements of CMB temperature and polarization anisotropy over the last 20 years with experiments such as *WMAP* [149] and *Planck* [8].

On the other hand, our constraints on the frequency spectrum of the CMB have not improved since the measurements of the Far Infrared Absolute Spectrophotometer on the *Cosmic Background Explorer* (*COBE/FIRAS* hereafter) [10]. Although the measured CMB spectrum closely matches a blackbody, the CMB in fact is expected to have some small distortion away from a blackbody spectrum.

Energy injections in the form of diffusion damping of small-scale anisotropies, resulting from imperfect photon-baryon coupling in the pre-recombination plasma, during periods of inefficient thermalization ($z \lesssim 2 \times 10^6$) will slightly distort the spectrum. For $2 \times 10^6 \gtrsim z \gtrsim 5 \times 10^4$, distortions of the μ -type are generated by this process, while for $5 \times 10^4 \gtrsim z \gtrsim 1100$ distortions of the y -type are generated. Using an internal blackbody as a calibrator, *COBE/FIRAS* was able to confirm the CMB spectrum closely follows a blackbody distribution and place upper limits of $|y| < 1.5 \times 10^{-5}$ and $|\mu| < 9 \times 10^{-5}$ (95% CL) (see also follow up analyses [150, 25]).

While diffusion damping of fluctuations from slow-roll inflation is one small ($\mu \gtrsim 10^{-8}$) but guaranteed method for generating spectral distortions before recombination [60, 61, 62, 21], other possibilities include annihilating particles [151], diffusion damping in inflationary

models that generate primordial black holes [19], primordial black hole evaporation [152], and primordial gravitational waves [153].

The CMB radiation can also be distorted by post-recombination sources, for example through inverse Compton scattering off of the hot electron gas in galaxy clusters, resulting in cluster-scale distortions of the CMB spectrum, a phenomenon referred to as the thermal Sunyaev-Zel'dovich (tSZ) effect [60]. In general spectral distortions allow one to probe any process associated with energy injection into the CMB after the thermalization epoch.

Measuring the mean, or monopole, frequency spectrum of the CMB is extremely challenging, because it requires an experiment to retain information about the absolute power received from the sky, not just the difference in power between different sky locations. Absolute measurements require exquisite stability over long timescales and tight control over any spatially varying sources of emission. Note that these stability requirements remain even for an experiment that does not require an overall absolute gain calibration (e.g., [154, 155]). For these reasons, it is often assumed that such measurements can only be made from space.

One way around these requirements is to measure the mean distortion of an anisotropic signal that can be measured differentially, such as the CMB dipole [156] or primary CMB anisotropy. The issue with using CMB temperature anisotropy is that most differential CMB experiments use the temperature anisotropy (either the dipole—or, more precisely, the annual modulation of the dipole—or the degree-scale and smaller anisotropy) as a calibration source, with the underlying assumption that the photon distribution follows a perfect blackbody. This effectively destroys any sensitivity to spectral distortions from the dipole or primary anisotropy, because the calibrated spectrum of the anisotropy will be forced to look like the derivative of a blackbody. Put another way, experiments designed to measure spectral distortions in the dipole or primary CMB anisotropy must find a different way of calibrating the relative response between observing frequencies.

In this work, we investigate the prospect for using the tSZ effect to measure monopole spectral distortions. This method, first proposed by [157], was used recently by [158] to

forecast constraints on the (primordial) y -distortion of the CMB from distortions of the (local-universe) tSZ effect. As discussed in [29], this technique can in principle be applied to y - or μ -type distortions, and was also proposed in [159] to test the validity of early measurements indicating large spectral distortions near the blackbody peak, later demonstrated by *COBE*/FIRAS to be spurious. Similar works have explored constraining the primordial recombination radiation [160] and the redshift evolution of the CMB temperature from the distortion of the tSZ effect. This paper focuses on the potential constraints on the mean value of μ -type distortions from measurements of the tSZ effect in the direction of massive clusters of galaxies using calibration from primary CMB temperature anisotropy under the blackbody assumption. We will forecast constraints on this quantity from the upcoming CMB-S4 experiment [9] as well as one based on the proposed CMB-HD experiment [161].

4.2 CMB Spectral Distortions

4.2.1 μ and y Distortions

At early epochs, any changes in the photon phase space distribution f are efficiently thermalized to a blackbody distribution through the joint action of the photon-number-changing processes double Compton scattering and Bremsstrahlung, and the energy-exchanging process (single) Compton scattering. Number-changing processes fall out of equilibrium at a redshift $z_i \sim 2 \times 10^6$ after which the photon distribution evolves mainly under the Kompaneets equation [162] (see Appendix A.4 for relativistic corrections)

$$\frac{\partial f}{\partial \tau} = \frac{k_B T_e}{m_e c^2} \frac{1}{x_e^2} \frac{\partial}{\partial x_e} \left[x_e^4 \left(\frac{\partial f}{\partial x_e} + f(1+f) \right) \right], \quad (4.1)$$

where τ is the Thomson optical depth, and $x_e = h\nu/k_B T_e$ for a thermal distribution of electrons at temperature T_e . The equilibrium distribution under the Kompaneets equation is a Bose-Einstein distribution. Any changes to the number or energy density of the photons

thereafter lead to a μ -type distortion

$$f = (e^{x+\mu} - 1)^{-1}, \quad (4.2)$$

where $x = h\nu/k_B T$ with the temperature of the photons $T = T_e$. For example a fractional energy injection of $\Delta\rho/\rho$ to the photons leads to $\mu \sim 1.4\Delta\rho/\rho$. Energy exchange via Compton scattering falls out of equilibrium at around $z_f \sim 5 \times 10^4$. After this epoch, we can solve the Kompaneets equation by plugging in the unperturbed spectrum (4.2) into the right hand side of Eq. (4.1) and integrating [29]

$$\Delta f(x, \mu, y) = \int d\tau \frac{\partial f}{\partial \tau} \approx y x e^{x+\mu} f^2 g(x, \mu), \quad (4.3)$$

with

$$g(x, \mu) = x \coth\left(\frac{x+\mu}{2}\right) - 4, \quad (4.4)$$

where the Comptonization parameter,

$$y = \int d\tau \frac{k_B(T_e - T)}{m_e c^2}, \quad (4.5)$$

is assumed to be $|y| \ll 1$. This generalizes the standard expression for the y -type distortion to the case where $\mu \neq 0$, i.e. the photons possess an initial μ -type distortion. Notice that the spectrum only changes when $T_e \neq T$, e.g. when the electrons are heated after z_f . In particular we are interested in the case where the hot electrons exist in galaxy clusters and produce the late-time y -type distortions known as the tSZ effect. Our generalization implies that in principle the initial μ value can be determined from a precise measurement of the tSZ spectrum.

4.2.2 Interfrequency Calibration

As discussed in §4.1, most differential CMB experiments derive their interfrequency calibration from CMB anisotropy, either the CMB dipole or the primary temperature anisotropy, under the assumption that the background photon distribution is a pure blackbody. Experiments that have access to very large angular scales, such as the *Planck* and *WMAP* satellites, calibrate off of the annual modulation of the dipole from the Earth’s motion around the Sun. When compared to predictions using our precise knowledge of the current CMB temperature T_0 and the Earth’s orbital velocity, and assuming a blackbody background, this provides both an interfrequency calibration and a calibration of the overall intensity scale. Experiments that use the primary anisotropy for interfrequency calibration (as is the case for most ground-based CMB experiments) need a separate reference for the absolute intensity scale, but since the inference for μ depends on the relative frequency dependence for a given amplitude y , an accurate relative calibration of channels is more important than the overall absolute calibration. For the specific measurement envisioned in this work, the absolute scale is effectively marginalized over, and we neglect it hereafter.

In practice, for the case of calibration off of the annual modulation of the dipole, the signal in each frequency band is scaled to agree with predictions assuming a pure blackbody background. The situation is similar for calibration off of the primary anisotropy: maps at every observing frequency ν are compared to each other in a region of the sky and a range of angular scales in which the primary CMB anisotropy is the dominant signal, and the maps are calibrated so that the signal follows the expected spectrum of temperature fluctuations in a background blackbody with mean temperature T_0 . In both cases, the true spectrum of the calibration source is that of temperature fluctuations in the true background, and the result of calibrating assuming a blackbody background is that the measured, calibrated dipole and/or primary CMB anisotropy is forced to follow the spectrum of temperature fluctuations in a blackbody.

Let us examine the case of calibrating off of the observed dipole in the presence of a

monopole μ distortion in the background spectrum, while assuming the background spectrum is a blackbody. (The results in the case of calibrating off of the primary CMB anisotropy are identical.) In the case of dipole calibration, the Lorentz invariance of f implies that the specific intensity in the boosted frame $I_\nu^d \propto \nu^3 f$ obeys

$$I_\nu^d \propto \frac{\nu^3}{e^{h\nu_{\text{rest}}/k_B T + \mu} - 1} \quad (4.6)$$

where

$$\nu_{\text{rest}} = \left(\frac{1 - \beta \cos \theta}{\sqrt{1 - \beta^2}} \right) \nu, \quad (4.7)$$

and θ is the angle between the line of sight and the velocity. Notice that we can absorb the Doppler shift into a temperature anisotropy as usual and to first order in β , $T(\theta) \approx T(1 + \beta \cos \theta)$. The change in the specific intensity becomes

$$\Delta I_\nu^d \approx (\beta \cos \theta) T \frac{\partial I_\nu}{\partial T}. \quad (4.8)$$

The frequency dependence involves the derivative of I_ν , and this result holds for calibration involving any type of temperature anisotropy by suitably generalizing the anisotropy source, not just a dipole due to a boost. Note that we are ignoring higher-order terms in the expansion of the blackbody fluctuation spectrum, which are negligible at least for the order 10^{-5} anisotropy in the CMB.

If a blackbody background distribution is assumed in the calibration process, then the anisotropy-calibrated specific intensity I_ν^c differs from the true specific intensity I_ν by

$$I_\nu^c = C(x, \mu) I_\nu, \quad (4.9)$$

where the miscalibration from the true spectrum is characterized as

$$C(x, \mu) = \frac{\partial B_\nu / \partial T}{\partial I_\nu / \partial T}. \quad (4.10)$$

Notice that this anisotropy calibration factor involves the spectral shape of the derivative of the specific intensity not the specific intensity itself. Thus, while this particular calibration procedure removes any information about spectral distortions from the primary anisotropy signal, distortions of signals that do not have the spectrum of the temperature derivative of the CMB monopole spectrum can still be measured.

Counterintuitively, this observability includes the μ -distortion of the CMB monopole itself:

$$\frac{I_\nu^c(\mu)}{B_\nu} = \frac{I_\nu}{B_\nu} \frac{\partial B_\nu / \partial T}{\partial I_\nu / \partial T} = e^{-\mu} \frac{e^{x+\mu} - 1}{e^x - 1}, \quad (4.11)$$

and the correction for $|\mu| \ll x \ll 1$ goes as μ/x . In practice, as discussed in §4.1, since this measurement requires a non-differential measurement on the sky, it remains challenging from the ground.

Now let us apply this sort of calibration to the tSZ distortion of a μ -distorted background in the direction of a galaxy cluster, a signal which can be measured differentially. In terms of the calibrated apparent CMB temperature fluctuation at frequency ν , ΔT , we obtain

$$\begin{aligned} \Delta T(x, \mu) &\equiv \frac{\Delta I_\nu^c}{\partial B_\nu / \partial T} = \frac{\Delta I_\nu}{\partial I_\nu / \partial T} = \frac{\Delta f}{\partial f / \partial T} \\ &= yT_0 g(x, \mu), \end{aligned} \quad (4.12)$$

where we have used Eq. (4.3) for Δf . Notice that the anisotropy calibrated ΔT differs from the absolutely calibrated temperature fluctuation, and Eq. (4.4) for $g(x, \mu)$ carries the measurable frequency dependence under anisotropy calibration. This difference is illustrated in Fig. 4.1. Notice also that in both cases the response to μ increases at low frequency but with the opposite sign.




Figure 4.1: The fractional response of the tSZ temperature spectrum $\partial \ln \Delta T / \partial \mu$ to a monopole μ distortion with anisotropy calibration as we assume in this work (solid blue line $\partial \ln g / \partial \mu$) vs. absolute calibration (dashed blue line $\partial \ln(g/C) / \partial \mu$). The dotted gray line denotes an undistorted spectrum for reference and the spike in the curves occurs at the tSZ null where the *fractional* response diverges corresponding to a finite change in the location of the null.

4.2.3 Cluster model

Using Eq. (4.5), our model for the value of the Compton y parameter in the direction of an isothermal cluster ($T_e = \text{const.} \gg T$) is

$$y(\theta) = \frac{k_B T_e}{m_e c^2} \tau(\theta), \quad (4.13)$$

where θ is the angular distance from the center of the cluster. For the optical depth profile $\tau(\theta)$ we follow e.g., [163], and adopt a spherically symmetric β model or King profile with

$\beta = 1$ and express Eq. (4.13) as

$$y(\theta) = y_c \left[1 + \left(\frac{\theta}{\theta_c} \right)^2 \right]^{-1}. \quad (4.14)$$

Here, the angular size of the cluster's core is given by $\theta_c = r_c/D_A$, with D_A being the angular diameter distance and r_c the core radius of the cluster, all in comoving coordinates for later convenience. We follow [164, 165] and we adopt the relation $r_c \sim 0.2R_{500c}$, where R_{500c} is the radius at which the enclosed spherically averaged density is 500 times the critical density $\rho_c(z) \equiv 3H^2(z)/8\pi G$.

For y_c we adopt the self-similar scaling relation

$$y_c = A\tilde{E}^2(z) \left(\frac{M_{500c}}{10^{14}M_\odot} \right) \quad (4.15)$$

where

$$\tilde{E}(z) \equiv \frac{H(z)}{70\text{km/s/Mpc}}, \quad (4.16)$$

and the normalization A from X-ray cluster observations of luminosity and temperature at low z [166] to calibrate the universal pressure profile (Eq. 6 in Ref. [167])

$$A = 0.97 \times 10^{-5} h^{-3/2}. \quad (4.17)$$

Note using this normalization in the context of Eq. (4.14) is approximate given differences with the universal pressure profile [168]. We also adopt the temperature-mass relation [169]

$$k_B T_e = 2.28 \left(\frac{M_{500c}}{10^{14}M_\odot} \tilde{E}(z) \right)^{0.585} \text{keV}. \quad (4.18)$$

Because the noise in our forecasted surveys is expected to be diagonal in spherical harmonic (ℓ, m) space, we choose to work in that basis. To transform Eq. (4.14) into ℓ, m space, we note that since even the most massive and low-redshift clusters only subtend a small

angle on the sky, we can use the flat-sky approximation. As detailed in Appendix A.3, in coordinates centered on the cluster at $\theta = 0$, the spherical harmonic-space cluster profile is given by

$$y_{\ell m} = \sqrt{\frac{2\ell + 1}{4\pi}} \delta_{m,0} y(\ell), \quad (4.19)$$

where

$$y(\ell) = y_c 2\pi\theta_c^2 K_0(\ell\theta_c), \quad (4.20)$$

and $K_n(x)$ is the modified Bessel function of the second kind.

4.3 Forecast

4.3.1 Survey Specifications

We forecast our constraints on μ from tSZ cluster measurements using instrument configurations based on the upcoming CMB-S4 experiment and the proposed CMB-HD experiment. CMB-S4 will conduct two surveys: the Wide Survey conducted from Chile will cover 67% of the sky, while the Deep Survey will concentrate a similar amount of total observing weight on 3% of the sky from the South Pole. From here on, we will refer to these two CMB-S4 surveys as “S4-Wide” and “S4-Deep,” respectively. Both surveys will have similar beam sizes and differ mainly in the noise in the sky maps. The CMB-HD-like survey we forecast for here covers 50% of the sky. For all three surveys, we use the instrument configuration parameters from Tab. 1 of [2], which we reproduce in Tab. 4.1. We note that for both S4-Wide and CMB-HD, the galactic plane will significantly contaminate our maps of tSZ clusters and reduce our ability to accurately measure the cluster spectrum. Therefore, for these surveys we assume $f_{\text{sky}} = 0.5$.

Channels (GHz)		30	40	90	150	220	270
Survey	f_{sky}	$\theta_{\text{FWHM}} \ \& \ \sqrt{C_w} \ (\mu\text{K-arcmin})$					
S4-Wide	50%	7.3'	5.5'	2.3'	1.5'	1.0'	0.8'
		21.8	12.4	2.0	2.0	6.9	16.7
S4-Deep	3%	8.4'	5.8'	2.5'	1.6'	1.1'	1.0'
		4.6	2.94	0.45	0.41	1.29	3.07
CMB-HD	50%	1.4'	1.05'	0.45'	0.25'	0.2'	0.15'
		6.5	3.4	0.73	0.79	2.0	2.7

Table 4.1: Specifications for the CMB-S4 Wide and Deep surveys and a CMB-HD-like survey, taken from [2].

4.3.2 Cluster Catalog

In addition to specifications on map noise, angular resolution, and sky fraction, to forecast constraints on μ from the distortion of the tSZ spectrum we also need to define a sample of galaxy clusters. For each of the three surveys considered here, we use the expected cluster catalog for that survey, based on work from [2, 170].

Underlying the expected number of clusters detected by a given CMB experiment is the halo mass function $dn/d\ln M$, the number density of host dark matter halos at a given redshift z over a logarithmic mass interval $d\ln M$. We adopt for this quantity the Tinker mass function [171] as implemented in the publicly available code `Colossus`¹ [172]. Our cosmological parameters are taken from *Planck* 2018 [8], where $\Omega_m = 1 - \Omega_\Lambda = 0.31$, $\Omega_b = 0.049$, $H_0 = 67.7\text{km/s/Mpc}$, $\sigma_8 = 0.81$, $\tau = 0.054$, and $n_s = 0.965$.

A given experiment will have a selection function in mass and redshift which we approximate here as a simple mass limit as a function of redshift $M_{\text{lim}}(z)$. For each of the three surveys we forecast, we use the values of $M_{\text{lim}}(z)$ calculated in [2]. These limits are reproduced in Fig. 4.2. The jaggedness of the curves reflects the $\Delta z = 0.1$ binning in Ref. [2] as does our effective $z_{\text{min}} = 0.05$, but we will show in later sections that this effective redshift

1. <http://www.benediktdiemer.com/code/colossus/>




Figure 4.2: The mass detection limit M_{lim} as a function of redshift from Ref. [2], linearly interpolated between their $\Delta z = 0.1$ bins.

limit does not affect our results significantly. The general trend of the $M_{\text{lim}}(z)$ curves—which is the opposite of mass-limit curves from, e.g., X-ray-selected cluster samples—is discussed in Section 3.2.1 of [170].

We model the expected number of total detected clusters for each survey as

$$N_{\text{tot}} = 4\pi f_{\text{sky}} \int_{z_{\text{min}}}^{\infty} dz \frac{D_A^2(z)}{H(z)} \int_{M_{\text{lim}}(z)}^{\infty} \frac{dM}{M} \frac{dn}{d \ln M}, \quad (4.21)$$

where f_{sky} is the fraction of sky measured by the experiment. We find that for our fiducial cosmology: $N_{\text{tot}} = 1.04 \times 10^5$ for S4-Wide; 1.10×10^4 for S4-Deep, and 4.63×10^5 for CMB-HD. Our number of clusters agrees with [2] to within $\sim 3\%$ for S4-Wide, $\sim 7\%$ for S4-Deep, and $\sim 10\%$ for CMB-HD.

4.3.3 Forecasting Method

We forecast constraints on μ from the distorted tSZ spectrum in the direction of massive clusters using a Fisher matrix technique. First, we define the likelihood per cluster in the catalog. Given the expression for the measured, calibrated tSZ spectrum from Eq. (4.12), we model the cluster likelihood \mathcal{L} as

$$\begin{aligned}
 -2 \ln \mathcal{L} = & \sum_{ij, \ell m, \ell' m'} [\Delta T_{i, \ell m} - y_{\ell m} T_0 g(x_i, \mu)] \\
 & \mathbf{C}_{ij, \ell m \ell' m'}^{-1} [\Delta T_{j, \ell' m'} - y_{\ell' m'} T_0 g(x_j, \mu)], \tag{4.22}
 \end{aligned}$$

where i and j run over frequency bands, $\Delta T_{i, \ell m}$ is the measured, calibrated (spherical harmonic-space) temperature fluctuation in band i in the direction of the cluster, $y_{\ell m}$ is the spherical harmonic-space cluster profile, and we have approximated the sources of noise as Gaussian by characterizing the likelihood with the covariance matrix \mathbf{C} . Using Eq. (4.20) for the cluster profile and assuming statistical isotropy there is no azimuthal dependence in the model or the covariance, and the covariance will be diagonal in ℓ , in which case we can write

$$\begin{aligned}
 -2 \ln \mathcal{L} = & \sum_{ij, \ell} \frac{2\ell + 1}{4\pi} [\Delta T_{i, \ell} - y(\ell) T_0 g(x_i, \mu)] \\
 & (\mathbf{C}_\ell)_{ij}^{-1} [\Delta T_{j, \ell} - y(\ell) T_0 g(x_j, \mu)]. \tag{4.23}
 \end{aligned}$$

For the noise covariance matrix, we begin with a baseline of just uncorrelated white noise and write

$$(\mathbf{C}_\ell)_{ij} \rightarrow (\mathbf{C}_\ell)_{ij}^w = \delta_{ij} \frac{C_{w,i}}{B_{\ell,i}^2}, \tag{4.24}$$

where $C_{w,i}$ is the map noise variance in band i , and the Gaussian beam profile is

$$B_{\ell,i}^2 \approx \exp \left[-\frac{\ell(\ell+1)}{8 \ln 2} \theta_{\text{FWHM},i}^2 \right]. \quad (4.25)$$

In this case, the likelihood reduces to

$$-2 \ln \mathcal{L} = \sum_{i,\ell} \frac{2\ell+1}{4\pi} \frac{B_{\ell,i}^2}{C_{w,i}} [\Delta T_{i,\ell} - y(\ell) T_0 g(x_i, \mu)]^2. \quad (4.26)$$

More generally we can include other noise terms, indexed by X , as additional contributions to the covariance matrix

$$(\mathbf{C}_\ell)_{ij} = (\mathbf{C}_\ell)_{ij}^w + \sum_X (\mathbf{C}_\ell)_{ij}^X \quad (4.27)$$

and in particular for various foreground noise contributions that are fully correlated in frequency space, we take

$$(\mathbf{C}_\ell)_{ij}^X = \sqrt{C_X(\ell, \nu_i) C_X(\ell, \nu_j)}, \quad (4.28)$$

where $C_X(\ell, \nu_i)$ is the angular power spectrum of component X at frequency ν_i . We often characterize such contributions using their logarithmic power spectrum

$$D_X(\ell, \nu_i) \equiv \frac{\ell(\ell+1)}{2\pi} C_X(\ell, \nu_i). \quad (4.29)$$

To forecast measurement errors on μ we employ the Fisher matrix

$$\mathbf{F}_{\alpha\beta} = - \left\langle \frac{\partial^2 \ln \mathcal{L}}{\partial p_\alpha \partial p_\beta} \right\rangle, \quad (4.30)$$

where in our baseline study we take the parameters as $p_\mu \in y_c, \mu$ and evaluate the parameter derivatives around a fiducial model with $\mu = 0$ and the expected $y_c(M, z)$. We include T_e as a parameter when considering relativistic corrections in §4.4.2. In general, the forecasted

$\sigma(\mu)$ assuming:	S4-Wide	S4-Deep	CMB-HD
Baseline noise only	1.6×10^{-4}	1.4×10^{-4}	2.8×10^{-5}
+ 1st order rSZ	2.1×10^{-4}	1.9×10^{-4}	3.6×10^{-5}
+ CMB & background kSZ	2.5×10^{-4}	2.5×10^{-4}	4.4×10^{-5}
+ cluster kSZ	2.8×10^{-4}	2.6×10^{-4}	4.6×10^{-5}
+ extragalactic foregrounds	3.5×10^{-4}	7.0×10^{-4}	1.2×10^{-4}
+ galactic foregrounds	9.2×10^{-4}	9.1×10^{-4}	1.6×10^{-4}
+ atmosphere	1.3×10^{-3}	9.9×10^{-4}	1.9×10^{-4}

Table 4.2: Forecasted constraint on μ for the baseline white detector noise of each experimental configuration and its cumulative degradation from additional effects.

error on μ then comes from the $\mu\mu$ element of the matrix inverse of \mathbf{F} ,

$$\sigma_k(\mu) = \sqrt{(\mathbf{F}^{-1})_{\mu\mu}}, \quad (4.31)$$

where k indexes the cluster so that the combined result of the independent clusters in the catalogue is given by

$$\sigma^{-2}(\mu) = \sum_k \sigma_k^{-2}(\mu). \quad (4.32)$$

Since the sum over identical clusters involves the same σ_k , in practice we sum over mass and redshift bins that are narrow enough so as to provide results that is sufficiently close to the full sum once weighted by the expected number of clusters per bin.

4.4 Results

In this section, we present our main forecasting results. We begin by providing the forecasted constraints on μ for each of the three experimental configurations in the idealized or “baseline” case of white detector noise only. We then introduce real-world complexities that an experiment will have to address, including relativistic corrections to the tSZ effect, CMB background anisotropy, cluster-associated kSZ signal, foreground sources, and atmospheric

contamination. We report the degradation of constraints from each of these cumulatively. Since we do not analyze each effect separately, the ordering of the cumulative contributions can matter in the interpretation of which is seemingly the most significant. We choose this approach to instead emphasize which complexities, in descending order, are fundamental to the measurement and which ones are contaminants to specific experiments.

4.4.1 Baseline Noise

Constraints on μ for each survey configuration for the baseline case of white detector noise only are shown in the first row of Tab. 4.2. These represent the most optimistic possible projections from each survey, and the baseline against which we compare the degraded constraints from successive real-world effects in the rest of the table and section.

We notice a few interesting results with regards to our baseline constraints on μ . For the two CMB-S4 surveys in this ideal forecast, $\sigma(\mu)$ is comparable to the bounds from *COBE*/FIRAS, which constrain $|\mu| < 9 \times 10^{-5}$ (95% CL, [10]). With a CMB-HD-like configuration, we start to see improved constraints on μ relative to *COBE*/FIRAS, indicating that, from a raw sensitivity standpoint, this method of constraining μ has some promise.




Figure 4.3: The cumulative squared signal to noise, $(S/N)_{\text{cumul}}^2$, above a given cluster mass and redshift threshold normalized to the total $(S/N)_{\text{surv}}^2$ for the three surveys from their catalog of clusters with masses and redshifts above the red line.

In addition, we note that S4-Deep provides a slightly better constraint on μ than S4-

Wide, despite the fact that the constraint comes from an order of magnitude fewer clusters. This can be understood from the fact that, for a fixed set of frequency bands, the per-cluster $\mu\mu$ Fisher matrix element (Eq. 4.31) will scale as the square of the total signal-to-noise (S/N) on the tSZ signal from the cluster. In the ideal white-noise-only case, for the i th frequency channel and a cluster of a given mass and redshift, the squared, per-cluster tSZ S/N is given by

$$\begin{aligned} \left(\frac{S}{N}\right)_i^2 &= \sum_{\ell} \frac{2\ell + 1}{4\pi} \frac{[y_{\ell} T_0 g(x_i, 0) B_{\ell, i}]^2}{C_{w, i}} \\ &= \frac{[y_c T_0 g(x_i, 0)]^2}{C_{w, i}} \sum_{\ell} \frac{2\ell + 1}{4\pi} [2\pi\theta_c^2 K_0(\ell\theta_c) B_{\ell, i}]^2. \end{aligned} \quad (4.33)$$

As expected, this scales as $y_c^2/C_{w, i}$.

The total μ constraint for a given survey will scale with this quantity summed over all the clusters in the catalog and frequency, $(S/N)_{\text{surv}}^2$. For S4-Wide $(S/N)_{\text{surv}}^2$ is 1.5×10^9 , for S4-Deep it is 1.9×10^9 , and for CMB-HD it is 2.3×10^{10} . The CMB-S4 Wide survey covers 17 times more sky than the Deep survey, so for any mass and redshift bin above the detection limit of both surveys, the Wide survey will have 17 times more clusters in the catalog. But the square of the ratio of map noise in the main CMB bands in the two surveys—and, by extension the squared S/N per cluster—is over 20, so it is not surprising that the Deep survey attains slightly better μ constraints.

This line of reasoning ignores the fact that the CMB-S4 Deep survey also has a lower mass limit and a higher cluster number density, which in principle could lead to an even larger difference between the μ constraints from the Deep and Wide surveys. All of the clusters that will be in the S4-Deep catalog but not the S4-Wide catalog are low-mass systems with $z > z_{\text{min}} = 0.05$, but as we shall see next, these clusters do not significantly improve the constraint.

To understand which clusters are providing most of the constraining power, we calculate the cumulative $(S/N)^2$ above a given mass and redshift and plot that quantity in Figure

4.3. Specifically we calculate the per-cluster $(S/N)^2$, calculated for frequency channel i using Eq. (4.33), sum over frequency channels and clusters above a given mass M and redshift z in the catalogue, and plot this cumulative $(S/N)_{\text{cumul}}^2(M, z)$.

Note that Figure 4.3 extends the $(S/N)_{\text{cumul}}^2$ to below our fiducial values for $M_{\text{lim}}(z)$ and z_{min} (red lines) so that the ratio with the given survey $(S/N)_{\text{surv}}^2$ can exceed unity. Nonetheless, in each case half of $(S/N)^2$ at z_{min} comes from cluster masses well above $M_{\text{lim}}(z_{\text{min}})$ and at $M_{\text{lim}}(z)$ from cluster redshifts below $z < 0.5$. This implies that the clusters around $M_{\text{lim}}(z)$ for each survey are not providing much constraining power on μ if $z_{\text{min}} = 0.05$. It is only for $z < 0.05$ and masses substantially below $M_{\text{lim}}(z_{\text{min}})$ that the cumulative $(S/N)^2$ changes noticeably, but even then only by 20 – 30%.²

In this work we produce forecasts for fixed instrument configurations, but it is possible that small modifications to one or more of the configurations could improve the μ constraints. In particular, it is not obvious from just the total S/N which frequency bands are contributing most to the constraint, and where more bands could potentially help. We note that, when the frequency band allocation is not fixed, the total μ constraint depends not just on the total S/N but also includes the sensitivity of bands to the μ distortion. We can write the $\mu\mu$ Fisher matrix element as

$$\mathbf{F}_{\mu\mu} = \sum_i \left(\frac{S}{N}\right)_i^2 \left(\frac{\partial \ln g(x_i, \mu)}{\partial \mu}\right)^2 \quad (4.34)$$

where recall $\partial \ln g / \partial \mu$ is shown in Fig. 4.1. Of course, the final constraint on μ depends on the other contributions to the signal that must be marginalized over.

In the simple case where only y_c is marginalized over, we can build intuition for which frequencies contribute most to the μ constraint by considering the scenario with only two channels, in which case the squared uncertainty on μ (or, equivalently, the $\mu\mu$ element of the

2. At least part of this signal could be recovered by augmenting the internal cluster catalogs with external detections in the optical and X-ray bands. For example, the cluster mass limit for the all-sky survey of the currently operating *eROSITA* mission is $\lesssim 2 \times 10^{13} M_{\odot}$ at $z < 0.1$ (see, e.g., Figure 5.1.1 in [173]).

inverse Fisher matrix) is given analytically by

$$\sigma^2(\mu) = \frac{(S/N)_1^2 + (S/N)_2^2}{(S/N)_1^2 (S/N)_2^2} \left[\frac{\partial}{\partial \mu} \ln \frac{g(x_1, \mu)}{g(x_2, \mu)} \right]^{-2}. \quad (4.35)$$

From Eq. (4.35), it is clear that both sensitivity and frequency lever arm are important for constraining μ , as $\sigma^2(\mu)$ blows up when the S/N in either of the two bands gets too small or when $g(x, \mu)$ is similar enough between the bands that the spectral signature becomes indistinguishable from that of y_c . For estimation purposes, we find that the expression

$$\left(\frac{S}{N} \right)_i^2 \approx \frac{(y_c T_0)^2}{C_{w,i}} \frac{\pi \theta_c^2 g^2(x_i, 0)}{1 + (\theta_{\text{FWHM},i}/4\theta_c)^{1.6}} \quad (4.36)$$

approximates Eq. (4.33) to within a few percent for all clusters and instrument configurations discussed here. To further illuminate scaling results we can also roughly scale $\sqrt{C_w}$ and θ_{FWHM} with frequency from 150GHz to mimic CMB-S4 Wide survey specifications:

$$\frac{\sqrt{C_w}(\nu)}{2 \mu\text{K-arcmin}} = 1 + 11.5 \left(\frac{\nu}{150 \text{ GHz}} - 1 \right)^2, \quad (4.37)$$

$$\frac{\theta_{\text{FWHM}}(\nu)}{1.5 \text{ arcmin}} = \left(\frac{\nu}{150 \text{ GHz}} \right)^{-1}. \quad (4.38)$$

We plot Eq. (4.37) in Fig. 4.4 (upper panel, curve) and compare it against the actual CMB-S4 Wide channel noise (points). Using this noise curve and Eq. (4.38), in Fig. 4.4 (bottom panel), we show $\sigma^2(\mu)$ for a cluster with $y_c = 1 \times 10^{-4}$ and $\theta_c = 1$ arcminute in this two-channel case, as a function of the frequency of the second channel ν with the first fixed at either $\nu_1 = 30, 90,$ or 150 GHz.

Notice that $\sigma^2(\mu)$ diverges whenever the two frequencies have the same value of $\partial \ln(\Delta T)/\partial \mu$ (see Fig. 4.1), causing μ to become degenerate with y_c in the fit. This occurs by definition when the two frequencies are coincident, and, for a lower frequency ν_1 below the null, it occurs again for a specific upper frequency ν_2 . In the limit where the lower frequency goes to zero and the μ response diverges, this second degeneracy between μ and y_c occurs when




Figure 4.4: Top: Frequency scaling relation for $\sqrt{C_w}$. Blue points correspond to the S4-Wide survey's specifications. Bottom: $\sigma^2(\mu)$ for different frequency pairs for a single cluster with $y_c = 10^{-4}$ and $\theta_c = 1$ arcminute. We see that the constraints on μ improve when the channels are separated from each other, rather than being closely spaced. In general, we see that as long as there is a separation, lower frequency channels provide more sensitivity to μ in this baseline case of white detector noise only, in accordance with Fig. 4.1.

the upper frequency approaches the null. For a lower frequency around 150 GHz, the degeneracy disappears since the response in Fig. 4.1 is near the local minimum where it is single-valued in frequency. The degeneracy is “accidental” in the sense that it only exists for pairs of frequency channels and is resolved once there are three or more channels. As we shall see next, the more complexity we add onto this baseline case the more multiple frequency channels are required to distinguish the μ signal.

4.4.2 Relativistic corrections

So far when forecasting μ -distortions from the distorted tSZ spectrum, we have used the non-relativistic limit of the tSZ frequency spectrum. In reality, the hottest clusters, from which most of our constraining power is derived, are going to have non-negligible relativistic corrections, sometimes called the relativistic Sunyaev-Zeldovich (rSZ) effect, especially compared to the small level of distortion that μ introduces. We show in Appendix A.4 that the μ -distorted rSZ spectrum modifies Eq. (4.12) for the anisotropy calibrated temperature fluctuation to

$$\Delta T = yT_0g(x, \mu, \theta_e), \quad (4.39)$$

where $\theta_e = k_B T_e / m_e c^2$. This generalizes from the form $g(x, \mu, 0) = g(x, \mu)$ given in Eq. (4.4).

We then marginalize over T_e per cluster bin, around the central values given by Eq. (4.18), along with y_c using the first order in θ_e expression for g from Eq. (A.40). In Fig. 4.5, we show the corresponding fractional change in the y distortion as a function of frequency for a range of cluster temperatures for comparison with Fig. 4.1 for the μ distortion. Marginalizing over T_e per cluster bin degrades our constraint on μ by $\sim 30\%$ (see Tab. 4.2). While this is not a large effect by itself, marginalizing over T_e has the effect of using up another linear combination of frequency bands to help break the degeneracy with μ , as was the case with y_c .

As noted by Ref. [174], the convergence of relativistic corrections as a Taylor expansion

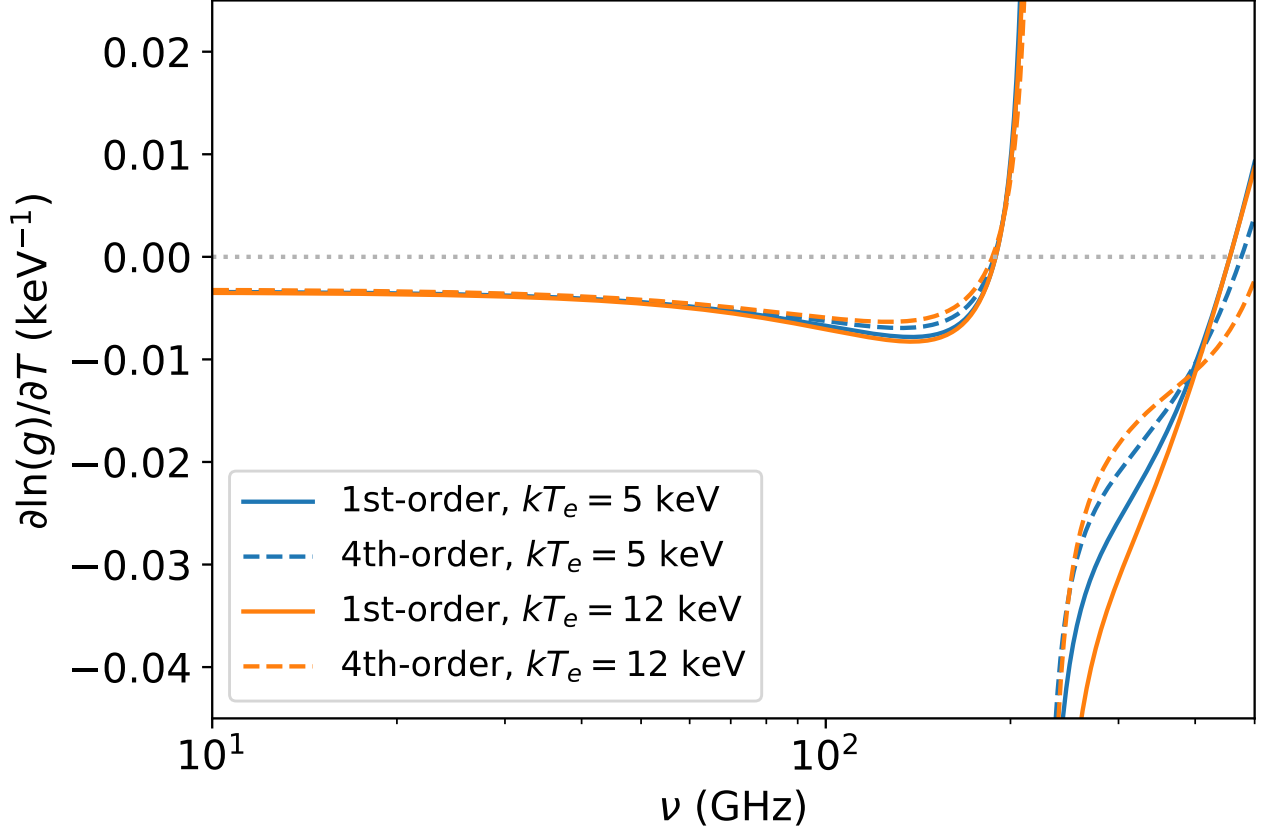


Figure 4.5: The fractional response of the tSZ spectrum, with 1st-order and 4th-order relativistic corrections, to the temperature of the cluster. Here we take $\mu = 0$ for a cluster with $k_B T_e = 5$ keV and another with $k_B T_e = 12$ keV.

in powers of θ_e is slow at frequencies around the null and above. In Fig. 4.5, we also show the spectral shape of the relativistic correction at 4th order. The small change in the shape associated with the central frequencies of the surveys, which are below the null, implies a correspondingly small change in the μ constraints. We find that going to 4th order makes a 5% change for S4-Wide, a 6% change for S4-Deep, and a 3% change for CMB-HD.

4.4.3 CMB anisotropies

Another source of variance in the measurement of the tSZ effect from clusters is primary CMB anisotropy. As can be inferred from Fig. 4.3, the constraint on μ is dominated by high-mass, low-redshift clusters. These clusters are sufficiently extended on the sky that

primary anisotropy is a potential concern.

There are also secondary CMB anisotropies arising from the kinetic Sunyaev-Zeldovich (kSZ) effect due to scattering off of gas through its bulk rather than thermal motion after recombination. The kSZ effect is intrinsically a Doppler shift and has the same spectrum as primary anisotropies (see §4.2.2). In this section, we only consider the background kSZ signal rather than the contribution specific to the clusters in our catalog. The isotropic kSZ signal is subdominant relative to the primary CMB anisotropies until $\ell \gtrsim 4000$.

Since these sources are statistically isotropic, we model their effects on our constraints by including the CMB temperature power spectrum and kSZ power spectrum in our noise covariance matrix. As described in §4.2.1, the spectrum for CMB anisotropies will look like the derivative of a blackbody given that our experiments will calibrate off the anisotropies. Such sources have a constant temperature across frequencies by definition. This means that the temperature power spectra will act as frequency independent, fully-correlated noise across frequency channels in Eqs. (5.5), (4.29). We use CAMB to generate the primary CMB logarithmic power spectrum $D_{TT}(\ell)$ for the fiducial cosmology. For the kSZ power spectrum, we consider a scale invariant spectrum with a constant $D_{\text{kSZ}}(\ell)$. We use the amplitude measured from the South Pole Telescope (SPT) at $\ell = 3000$ [137], $D_{\text{kSZ}}(3000) = 2.9 \mu\text{K}^2$.

As shown in Tab. 4.2, $\sigma(\mu)$ slightly degrades when including the CMB and kSZ background, but not by a large amount. For both CMB-S4 surveys, the individual contribution of the CMB and kSZ effect are similar to one another. The smaller beams and wider ℓ coverage of CMB-HD make the kSZ effect relatively more important than the primary CMB.

4.4.4 Cluster Kinetic Sunyaev-Zeldovich Effect

In the previous section, we only considered the kSZ signal corresponding to a statistically isotropic background. Since each cluster also has a specific kSZ profile associated with its gas profile and peculiar motion, we cannot treat it as statistically isotropic noise as we can with foregrounds that are not associated with the cluster.

We could model the kSZ signature from each cluster as part of the cluster signal and marginalize over the peculiar velocity of the cluster as we did for y_c and T_e . On the other hand, the spectrum of the kSZ is perfectly known, and we can effectively marginalize over a signal with a known spectrum by adding a component with that spectrum and artificially high amplitude to the covariance matrix [175]. Since kSZ has the same spectrum as the primary CMB anisotropy, which already is in the covariance matrix with an amplitude much larger than the instrumental noise, we expect this procedure to have minimal impact on our μ constraints. In practice, we marginalize over any signal with the spectrum of primary CMB anisotropy or kSZ by multiplying the $D_{TT} + D_{\text{kSZ}}$ spectra by a sufficiently large constant that the resulting $\sigma(\mu)$ saturates to its asymptotic value. Any part of the μ -signal that comes from the combination of frequency channels with a blackbody spectrum will have effectively infinite noise and not contribute to the constraint.

The fourth row of Table 4.2 shows $\sigma(\mu)$ when we implement this procedure. We see that our constraints hardly change for all three CMB surveys. This suggests that, as expected, any contribution to the μ constraint from the combination of frequency band information corresponding to the spectrum of primary CMB anisotropy is already made negligible by the inclusion of the fiducial $D_{TT} + D_{\text{kSZ}}$ in the covariance.

4.4.5 *Extragalactic foregrounds*

We treat three independent types of extragalactic foregrounds $X \in \text{c, p, r}$: “c”, the clustered cosmic infrared background (CIB); “p”, the spatially unclustered or “Poisson” component of the CIB; and “r”, the radio sources or synchrotron-emitting active galactic nuclei, the clustering of which is assumed to be negligible. For extragalactic foregrounds and, in the next section, galactic foregrounds, we parameterize the logarithmic power spectrum of foreground X at multipole ℓ and frequency ν as

$$D_X(\ell, \nu) \propto \left[\frac{f_X(\nu)}{(\partial B / \partial T)_\nu} \right]^2 \ell^{\beta_X}, \quad (4.40)$$

and provide the normalization $D_X(\ell_X, \nu_X)$ at a fiducial multipole value ℓ_X and frequency ν_X . Here β_X is the index of the assumed power-law multipole dependence, $f_X(\nu)$ encodes the frequency dependence in specific intensity units, and $(\partial B/\partial T)_\nu$ converts specific intensity to CMB temperature units at the relevant frequency ν assuming a blackbody CMB spectrum.³ We assume 100% correlation between observing bands for all individual extragalactic and galactic foreground components, such that we only need to specify the behavior at a single frequency, with the cross-frequency components of the covariance given by Eq. (5.5).

For all three extragalactic foregrounds considered here, we assume a power-law frequency dependence

$$f_X(\nu) = \nu^{\alpha_X}. \quad (4.41)$$

Following [137], for the clustered CIB we adopt $D_c(3000, 150 \text{ GHz}) = 3.5 \mu\text{K}^2$ and $\alpha_c = 4.3$, while for the Poisson component of the CIB we adopt $D_p(3000, 150 \text{ GHz}) = 9.2 \mu\text{K}^2$ and $\alpha_p = 3.3$, and for the radio Poisson component, we adopt $D_r(3000, 150 \text{ GHz}) = 2.0 \times 10^{-2} \mu\text{K}^2$ and $\alpha_r = -0.70$. For the ℓ dependence, we adopt $\beta_c = 0.80$ (again following [137]), while for the Poisson terms $\beta_p = \beta_r = 2$ by definition.

We note that the assumed radio amplitude we adopt is significantly lower than the best-fit radio source power quoted in [176] (the follow-up paper to [137], which did not constrain the radio amplitude), owing to the assumption of a much lower source cut threshold in the experiments treated here. The 5σ point source threshold in the 150 GHz channel of the CMB-S4 Deep survey should be roughly $50\times$ lower than that used in [176], and the slope of the number counts of radio sources is such that the Poisson power should scale roughly linearly with flux cut. As such, we adopt a radio amplitude value $50\times$ smaller than the $1.0 \mu\text{K}^2$ value from [176]. This number will be slightly optimistic for CMB-S4 Wide and slightly pessimistic for CMB-HD.

We see in Tab. 4.2 that extragalactic foregrounds have a larger impact on our constraints

3. Note that we neglect the μ -dependent anisotropy calibration C -factors from Eq. (4.10) here, since we are considering the foregrounds as noise rather than signal. In principle measuring anisotropy-calibrated foregrounds with known absolute spectra could themselves be used to measure μ .

of μ than the effects of the previous sections. Moreover CMB-HD no longer provides improvements on μ compared to the constraints from *COBE*/FIRAS. More specifically, in all three of our surveys, radio point sources are the dominant extragalactic foreground contaminating our constraints on μ . Both the clustered and Poisson contributions of the CIB have a negligible effect on $\sigma(\mu)$. Thus, efforts to reduce the effects of extragalactic foregrounds should prioritize mitigating the effects of radio sources. One clear path forward is to exploit the available lower-frequency data, both in the CMB surveys themselves and in planned contemporaneous radio surveys such as the Square Kilometer Array (SKA, [148]), the source detection threshold of which will be such that masking of SKA-detected sources in CMB-S4 or CMB-HD data will be limited by the number of independent pixels or resolution elements in the map. Using current source models (e.g., [177]), the source density at flux cut levels a factor of several lower than those assumed here still only reaches hundreds per square degree, still feasible for masking in CMB-HD data.

Note that we are implicitly treating extragalactic foregrounds as a statistically isotropic background to the cluster signal. In fact, galaxy clusters will likely be overdensities of “foreground” contamination as well as the desired tSZ signal. A potential method to account for this would be to parameterize the covariance matrix with amplitude parameters for each extragalactic component and marginalize over these parameters per cluster. With sufficiently informative priors from observations in other surveys and at other wavelengths, this could be achieved with a minimal degradation of the eventual μ constraint.

4.4.6 *Galactic foregrounds*

The sources of galactic contamination that have traditionally been considered most important at CMB observing frequencies are thermal dust emission and synchrotron emission, but we also include a component of “anomalous microwave emission” (AME) because of its importance at low frequencies (e.g., [178]).

Because galactic foregrounds are not statistically isotropic, we adopt separate sets of

values for foreground amplitudes for the $f_{\text{sky}} = 0.03$ S4-Deep survey and the $f_{\text{sky}} = 0.5$ region targeted by the S4-Wide survey (which we also adopt for the CMB-HD survey).⁴

Interstellar dust heated by starlight emits as a quasi-thermal modified blackbody. We follow Ref. [130] and parameterize the frequency behavior of thermal dust emission as

$$f_{\text{d}}(\nu) = \nu^{\alpha_{\text{d}}} B_{\nu}(T_{\text{d}}), \quad (4.42)$$

where $T_{\text{d}} = 19.6\text{K}$ is the dust temperature, and $\alpha_{\text{d}} = 1.6$. Also following that work, we set $\beta_{\text{d}} = -0.4$. Following Ref. [133], we adopt $D_{\text{d}}(80, 145 \text{ GHz}) = 3.3 \mu\text{K}^2$ for S4-Deep and $1.2 \times 10^3 \mu\text{K}^2$ for S4-Wide and CMB-HD.⁵ This very large increase from Deep to Wide is attributed at least partly to the requirement adopted in the CMB-S4 Wide survey to restrict observing elevation to $\geq 40^\circ$. If we impose no elevation restriction and instead choose the 50% of the sky at highest galactic latitude (using the publicly available PySM simulations [134] as in [133]), we find $D_{\text{d}}(80, 145 \text{ GHz}) = 63 \mu\text{K}^2$. This would similarly reduce the impact of AME on the wide surveys (see below for details).

Again following Ref. [130], we parameterize synchrotron as a pure power law in frequency (as in Eq. 4.41) and adopt $\alpha_{\text{s}} = -1.1$ and $\beta_{\text{s}} = -0.4$. Likewise following Ref. [133], we adopt $D_{\text{s}}(80, 93 \text{ GHz}) = 5.0 \times 10^{-3} \mu\text{K}^2$ for S4-Deep and $5.5 \times 10^{-2} \mu\text{K}^2$ for S4-Wide and CMB-HD. We note that the synchrotron amplitude does not vary as strongly across the sky in PySM as the dust amplitude: The ratio of power in the Wide and Deep areas is only ~ 10 for synchrotron, compared to over 300 for dust. Similarly, if we use $|b| > 30^\circ$ instead of the official CMB-S4 Wide region, we find that the synchrotron amplitude decreases by less than a factor of two (compared to ~ 20 for dust).

Because of the potential importance of low-frequency information in our μ constraint, we

4. Because of the statistically anisotropic nature of the galactic foregrounds, in a real data set, the covariance for clusters in different parts of the sky would be potentially quite different, and using the mean covariance across the sky for all clusters is not strictly correct.

5. This is technically for $f_{\text{sky}} = 0.58$, but if we recalculate for $f_{\text{sky}} = 0.50$, the value only decreases by $\sim 30\%$.

also consider the impact of AME. We investigate the behavior of AME in the CMB-S4 3% sky region using PySM. We find that the AME SED has a double-peaked shape, which we parameterize as

$$f_{\text{a}}^2(\nu) = e^{-[\ln(\nu/\nu_1)]^2/2\sigma_1^2} + Ae^{-[\ln(\nu/\nu_2)]^2/2\sigma_2^2}, \quad (4.43)$$

with $\nu_1 = 10$ GHz, $\sigma_1 = 0.43$ GHz, $\nu_2 = 22$ GHz, $\sigma_2 = 0.35$ GHz, and $A = 6.5 \times 10^{-3}$. We assume $\beta_{\text{a}} = -0.4$ (as would be expected if AME were from spinning dust grains and traced the thermal dust emission). From PySM, we estimate $D_{\text{a}}(80, 10 \text{ GHz}) = 1.0 \times 10^4 \mu\text{K}^2$ for the S4-Deep survey and, assuming the same scaling between deep and wide found for the thermal dust, $D_{\text{a}}(80, 10 \text{ GHz}) = 3.6 \times 10^6 \mu\text{K}^2$ for S4-Wide and CMB-HD.

The inclusion of galactic foregrounds has a larger impact on S4-Wide compared to S4-Deep and CMB-HD, which are more impacted by extragalactic than galactic foregrounds due to their μ constraint being weighted toward higher multipoles (see Tab. 4.2). More specifically, S4-Wide and S4-Deep now have comparable constraints on μ , despite vastly different galactic foreground amplitudes. Most notably, AME is responsible for most of the degradation in $\sigma(\mu)$ for all three surveys. Synchrotron provides some contribution, while being sub-dominant to AME, and dust has a negligible effect on $\sigma(\mu)$. As discussed above, relaxing restrictions on observing elevation in the wide surveys can help mitigate the impact of AME. But this result also motivates a more careful investigation into the spectral and spatial behavior of AME, beyond the simple ansatz made in this work.

We also note that, because of the large variation in Galactic foreground amplitudes across the wide survey region, using a single covariance matrix for the wide-area surveys is almost certainly sub-optimal. The true foreground covariance across the full region will have off-diagonal elements in ℓ space, owing to the statistical anisotropy across the region, but even including these off-diagonal elements, the use of a single covariance matrix implicitly assumes statistical isotropy across the field. Because of this, the constraints presented here for the wide surveys are likely somewhat pessimistic. We could potentially recover some of the μ constraints lost to Galactic foregrounds (and make the use of covariance matrices that are

diagonal in ℓ more appropriate) by dividing the wide surveys into sub-patches based on foreground amplitude and analyzing each sub-patch individually, with different covariance matrices. This would effectively weight different parts of the wide survey region inversely by foreground amplitude.

4.4.7 Atmosphere

We saw in our S/N contour plot that most of the signal is from high-mass, low redshift clusters, and they can subtend a large angle in the sky. In addition to being potentially confused with primary CMB fluctuations (see § 4.4.3), signals from objects this large on the sky are also impacted (in ground-based measurements) by emission from poorly mixed water vapor in the atmosphere. The amplitude of water-vapor fluctuations in the atmosphere is higher at large spatial scales than small spatial scales, and the emission thus behaves as “red noise” in CMB maps, often modeled as a power law in ℓ . The total detector + atmosphere noise power in frequency band i can then be parameterized with three numbers, namely the white noise level C_w , the multipole value at which the detector and atmosphere noise levels are equal ℓ_{knee} , and the power-law index of the atmosphere noise α_{atmo} :

$$C_{w,i} \rightarrow C_{w,i} \left[1 + \left(\frac{\ell_{\text{knee},i}}{\ell} \right)^{\alpha_{\text{atmo},i}} \right]. \quad (4.44)$$

Our values of ℓ_{knee} and α_{atmo} for the three surveys are taken from [2] and given in Tab 4.3.

While Eq. (4.44) describes atmospheric emission as uncorrelated between frequency bands, physical intuition and empirical evidence (e.g., [179]) argue that it should in fact be strongly correlated between bands, at least for instruments in which the beam patterns for detectors at different frequencies overlap in the atmosphere. The effects of atmosphere could in principle be reduced using the correlation between frequency bands to project out much of the atmospheric contamination.

Channels (GHz)		30	40	90	150	220	270
Survey	f_{sky}	$\ell_{\text{knee}} \& \alpha_{\text{atmo}}$					
S4-Wide	50%	400	400	1900	3900	6700	6800
		3.5	3.5	3.5	3.5	3.5	3.5
S4-Deep	3%	400	400	1200	1900	2100	2100
		4.2	4.2	4.2	4.1	4.1	3.9
CMB-HD	50%	400	400	1900	3900	6700	6800
		3.5	3.5	3.5	3.5	3.5	3.5

Table 4.3: Atmosphere parameters for the CMB-S4 Wide and Deep surveys and a CMB-HD-like survey, taken from [2].

4.4.8 Order of Operations

Our chosen ordering of cumulative effects may give the impression that certain effects are negligible because they are when implemented early in the ordering. However, these effects could prove significant when implemented last, after the survey’s constraining power is used to fix other effects. To help gauge the impact each effect has on the end result, we calculate $\sigma(\mu)$ when excluding individual effects from the end result. These results are shown in Tab. 4.4.

$\sigma(\mu)$ when excluding relativistic contributions, compared to including 1st-order corrections, improves by $\sim 40\%$ for both CMB-S4 surveys, whereas the results improve by $\sim 30\%$

Excluded	S4-Wide	S4-Deep	CMB-HD
1st order rSZ	$\times 0.60$	$\times 0.58$	$\times 0.68$
CMB & all kSZ	$\times 0.90$	$\times 0.94$	$\times 0.97$
cluster kSZ	$\times 0.97$	$\times 0.99$	$\times 0.98$
extragal. fore.	$\times 0.99$	$\times 0.79$	$\times 0.92$
gal. fore.	$\times 0.41$	$\times 0.83$	$\times 0.71$

Table 4.4: Forecasted fractional improvement to $\sigma(\mu)$ (see last line of Tab. 4.2 for the values for each experimental configuration) when excluding certain individual effects.

for CMB-HD. If we include the effect up to 4th order, the results are almost identical to 1st-order results.

The CMB and the kSZ effect, both the isotropic and cluster component, have a negligible impact on $\sigma(\mu)$ when excluded at the end. Interestingly, extragalactic foregrounds also have a negligible impact on μ for S4-Wide and CMB-HD. S4-Deep’s constraint improves when we exclude extragalactic foregrounds, but only by about 25%.

The exclusion of galactic foregrounds improves S4-Wide and CMB-HD’s constraints on μ , but only modestly improves S4-Deep’s constraints. This is likely because contamination from galactic foregrounds is much worse for S4-Wide and CMB-HD, which include observations near the galactic plane.

Our results suggest that the largest way to improve constraints on μ for all surveys is to address relativistic corrections. For S4-Wide and CMB-HD, galactic foregrounds are a major challenge to improving constraints. For S4-Deep, galactic and extragalactic foregrounds present similar levels of degradation. Generally, addressing these challenges requires additional frequency channels in order to help isolate the μ signal. Finally, we note that priors on T_e can in principle be obtained from external data such as X-ray observations.

4.4.9 *Interfrequency Calibration Requirement*

While we can effectively perfectly account for the mis-calibration induced from assuming the background photon distribution is a blackbody when it is in fact a Bose-Einstein distribution, in a real instrument there will also be mis-calibration from the fact that the observation of the calibration source is not noise-free. If we parameterize this calibration error as

$$C_{\text{meas}} = C_{\text{true}}(1 + \delta_{\text{cal}}) \equiv 1 + \delta_{\text{cal}}, \quad (4.45)$$

then in a real experiment, the measured, (mis-)calibrated signal from a single cluster will be

$$\Delta T(x, y_c, \mu, \delta_{\text{cal}}) = yT_0 g(x, \mu)(1 + \delta_{\text{cal}}). \quad (4.46)$$

The requirement for interfrequency calibration is most obvious in the Rayleigh-Jeans limit (and in the limit $\mu \ll x$), in which $g(x, \mu) = -2(1 + \mu/x)$. The basic information used to constrain μ is the ratio of the cluster signal in two bands. If we assume perfect calibration in one band and a mis-calibration in the other, we find

$$\begin{aligned} R(x_1, x_2, y_c, \mu, \delta_{\text{cal}}) &= \frac{\Delta T(x_2, y_c, \mu, \delta_{\text{cal}})}{\Delta T(x_1, y_c, \mu)} \\ &= \frac{-2y_c T_0 (1 + \mu/x_2)(1 + \delta_{\text{cal}})}{-2y_c T_0 (1 + \mu/x_1)} \\ &= \frac{1 + \delta_{\text{cal}} + \mu/x_2 + \delta_{\text{cal}}\mu/x_2}{1 + \mu/x_1} \\ &\simeq \left(1 + \delta_{\text{cal}} + \frac{\mu}{x_2}\right) \left(1 - \frac{\mu}{x_1}\right) \\ &\simeq 1 + \mu \left(\frac{1}{x_2} - \frac{1}{x_1}\right) + \delta_{\text{cal}}. \end{aligned} \quad (4.47)$$

It is clear from this formulation that to constrain μ to some level $\sigma(\mu)$, we need calibration uncertainty smaller than $\sigma(\mu) \left| (x_2^{-1} - x_1^{-1}) \right|$. For the experimental configurations considered in this work, that means we need calibration better than $10^{-4} - 10^{-5}$ in the bands around the peak of the CMB blackbody. A full-sky experiment with noise levels of $1 \mu\text{K-arcmin}$ has S/N per band on the primary CMB temperature anisotropy approaching 10^6 , so in principle this level of interfrequency calibration is achievable.

Additionally, this level of calibration precision must be maintained over the full survey area. For surveys that cover a large fraction of the sky, different parts of the survey are in general surveyed at widely separated times and possibly under different atmospheric conditions. This places an effective requirement on calibration stability; alternatively, different parts of the survey can be calibrated independently, in which case the S/N requirement on

the CMB is per independently calibrated patch.

Finally, we note that this calibration requirement imposes a practical minimum size of the survey area. Fig. 4.3, taken at face value, implies that an efficient strategy for constraining μ with tSZ could be to make incredibly deep measurements on a single very massive and low-redshift cluster (or a handful of such clusters). If, however, the calibration for such a survey is to come from CMB anisotropy, the survey must contain enough sky in which the signal is dominated by primary CMB to achieve the required calibration precision. This disfavors strategies along the lines of pointing a powerful interferometer (such as ALMA) at a small number of massive clusters.

4.4.10 Cluster Profiles

The model for the angular distribution of the Compton- y signal adopted in §4.3 is an oversimplification in at least three ways:

1. Real clusters of the same mass and redshift will not all have the same overall normalization to their Compton- y signal.
2. Real clusters of the same mass and redshift will not all have the same profile shape.
3. The Compton- y signal in real clusters will not be described perfectly by an isothermal β -model profile.
4. Real clusters are not spherically symmetric.

One important point to remember here is that, in our forecasting, the cluster model serves effectively as a weighting of angular modes, and getting this weighting wrong will result in a lower S/N and thus a worse constraint on μ but will not bias the result.

The impact of the first two effects, scatter in the overall normalization and profile shape at fixed mass and redshift, can be estimated using the expression in Eq. (4.33). In a given mass and redshift bin, if there is scatter in overall normalization or scale radius, but they are

treated as uniform—for example, if the analysis is performed on a stacked cluster profile—the total achieved $(S/N)^2$ in a bin will be equal to $N_{\text{cl}}(S/N)_{\text{mean}}^2$, where N_{cl} is the number of clusters in the bin, and $(S/N)_{\text{mean}}^2$ is the result of Eq. (4.33) using the mean normalization or scale radius in the bin. The maximum achievable $(S/N)^2$, if each cluster was treated individually and its properties fully known, would be $\sum_i (S/N)_i^2$, and the degradation in $(S/N)^2$ from stacking is the ratio of these, $N_{\text{cl}}(S/N)_{\text{mean}}^2 / \sum_i (S/N)_i^2$. Because the constraint on μ scales as the inverse square root of this quantity (i.e., $\sigma^2(\mu) = F_{\mu\mu}^{-1} \propto (S/N)^2$, the degradation on $\sigma(\mu)$ will be 1/2 the degradation on $(S/N)^2$ (for small degradation factors).

We estimate the degradation factor for 10% scatter in overall normalization (roughly the upper limit to the intrinsic scatter in the $Y - M$ relation found in [180]) and for 30% scatter in scale radius (see Tab. A2 in [166]). The effect of scatter in the normalization can be estimated analytically and should result in a degradation of order $\sigma_{y_0}^2/y_0^2$ in total $(S/N)^2$, or $0.5\sigma_{y_0}^2/y_0^2$ in σ_μ . This relation predicts a 1% degradation in total $(S/N)^2$ or 0.5% in $\sigma(\mu)$ for 10% scatter, and we find results consistent with this using Eq. (4.33) for a Gaussian distribution with $\sigma_{y_0}/y_0 = 0.1$. For a Gaussian distribution of scale radii with $\sigma_{\theta_c}/\theta_c = 0.3$, we find a degradation of roughly 10% in total $(S/N)^2$ or 5% in $\sigma(\mu)$.

Assuming the wrong cluster model or assuming a spherically symmetric cluster when the true cluster is elliptical or triaxial will result in a degradation in S/N even with no scatter in cluster shape within a mass and redshift bin. For a fiducial cluster with mass $M_{500} = 5 \times 10^{14} M_\odot$ and redshift $z = 0.5$, if we assume the true underlying profile is the universal pressure profile from [166], but we use a β model in the fitting procedure, we incur a penalty in $(S/N)^2$ of roughly 3%. If we assume the true underlying model is a β model with typical ellipticity (defined as ratio of minor to major axis) of 0.8 (cf. Tab. 2 in [181]), we find a degradation in $(S/N)^2$ of roughly 1%. All of these effects are small compared to the impact of the configuration choices described earlier and shown in Tab. 4.2.

4.5 Conclusion

In this study, we have demonstrated that the spectrum of the tSZ effect in the direction of massive clusters of galaxies can be used to constrain the μ -distortion monopole. We have shown that this can in principle be achieved without measuring the mean intensity across the sky and instead using a differential experiment that calibrates off of the CMB anisotropies, even when assuming the underlying CMB is an undistorted blackbody. We forecasted constraints on μ using the tSZ spectrum for the upcoming CMB-S4 experiment, using both the Wide and Deep surveys, as well as the proposed CMB-HD experiment. We found that the most massive clusters at the lowest redshifts provide the strongest constraints on the μ -distortion monopole. In terms of raw sensitivity, we found that all three surveys closely match or outperform *COBE*/FIRAS in constraining the μ -distortion monopole. Extragalactic and galactic foregrounds significantly degrade these constraints to the point where CMB-S4 performs worse than *COBE*/FIRAS, and CMB-HD delivers roughly equivalent constraints to *COBE*/FIRAS. Specifically, we found that radio point sources heavily impact low-noise surveys such as S4-Deep and CMB-HD, whereas the inclusion of AME significantly degrades S4-Wide constraints on μ .

To improve on these constraints, foreground removal is a priority. Improving foreground removal, in general, requires additional frequency channels to help distinguish signal from foregrounds. In regards to specific foregrounds, improved masking of radio point sources using higher-resolution surveys should reduce their impact, while to reduce the effects of galactic foregrounds such as AME, the most straightforward strategy is to perform deep sky observations that avoid the galactic plane. We see in Fig. 4.1 and Fig. 4.4 that the distortion of the tSZ spectrum increases at lower frequencies. This suggests additional coverage at low-frequencies should also improve constraints on μ . Finally, the inclusion of external data (particularly X-ray data) could be useful both in filling in the low-redshift gaps in the CMB experiments' cluster selection and in providing external priors on the temperature of individual clusters, helping to break degeneracies between μ , central tSZ decrement, and

cluster temperature.

Based on the above discussion, an experiment that would improve on these current constraints should have many frequency channels to remove foregrounds, with some channels dedicated to frequencies below 30 GHz if possible. The experiment should have beams comparable to the targeted cluster sample with white noise levels comparable to or better than CMB-S4. This implies observations with radio instruments combined with a CMB experiment have the potential to improve measurements of μ . While our results suggest a deep observation of individual low-redshift clusters would be ideal for obtaining better constraints on μ , we caution that one would need to also measure in the same observation the CMB at a high enough SNR for all frequencies to calibrate off of CMB anisotropy.

Certain assumptions we make in our forecasts may turn out to be overly optimistic. For example, the assumption of 100% correlation between the foreground power across all frequency bands must break down at some level. Although the level of decorrelation in galactic dust at these frequencies has been limited to be very small [182], even a low level of decorrelation could degrade precision constraints significantly. On the other hand, analyzing the foregrounds in sub-patches should improve the statistical modeling and more optimally weight the data. We also show that the degradation in $\sigma(\mu)$ due to scatter in the overall normalization and cluster shape is small compared to other effects we model. Although it is small, this degradation could be recovered in future work by fitting the true profile to each cluster in order to recover the maximum achieved $(S/N)^2$.

The low-frequency enhancement of μ -distortions of the tSZ effect suggests that a synergistic combination of CMB and radio telescope data could further improve constraints on the μ monopole using this technique. To realize this promise with specific radio surveys, future studies can use the forecasting framework presented here to address the calibration and foreground-mitigation requirements of the combined data set.

CHAPTER 5

SQUARE KILOMETER ARRAY AS A COSMIC MICROWAVE BACKGROUND EXPERIMENT

This chapter is reproduced from [30]

5.1 Introduction

The Square Kilometre Array (SKA, [183])¹ is a planned array of radio telescopes, aiming to probe the Universe with unprecedented resolution and sensitivity, primarily through observations of the neutral hydrogen 21-cm hyperfine transition [183]. SKA 21-cm data have the potential to make groundbreaking contributions to our understanding of cosmology and astrophysics. More fundamentally, SKA will make maps of the sky at many frequencies (from 50 MHz to 15 GHz) with high spectral resolution. Contemporary experiments designed to measure the cosmic microwave background (CMB) typically observe only down to ~ 30 GHz, limited primarily by the challenge of adapting standard CMB detector technology to long wavelengths [184]. While this is sufficient for many science goals, certain applications of CMB data suffer from the lack of lower-frequency information, in particular the search for μ -type spectral distortions [185, 91] and the characterization of the Galactic synchrotron signal [186, 187].

We propose using SKA data not to trace neutral hydrogen but as a low-frequency CMB experiment. SKA is well-matched and complementary to current and planned CMB experiments in several ways:

- *Raw sensitivity:* The expected noise in SKA wide survey maps, in CMB units, is below $10 \mu\text{K-arcmin}$ for SKA-1 (and even lower for SKA-2) for sufficiently wide frequency bands, which approaches the projected map depths for, e.g., CMB-S4 [9] and LiteBIRD [27].

1. <https://www.skao.int>

- *Range of angular scales:* Cosmologically relevant information in the CMB is concentrated at angular scales of roughly tens of degrees down to several arcminutes (multipoles $10 \lesssim \ell \lesssim 3000$). The SKA will operate in two modes: Interferometric mode, which will access the high- ℓ part of this range, and single-dish mode, which will recover the low- ℓ part of this range.
- *Spectral resolution:* The ability to channelize the SKA data will provide flexibility in treating foregrounds. If foregrounds are more complex than anticipated, SKA data can be analyzed at high spectral resolution, with some noise penalty; if foregrounds turn out to be relatively simple, SKA data can be combined into wider bands to gain sensitivity.

While there have been previous suggestions of using SKA as a CMB experiment [188], to our knowledge, no detailed forecasts exist.

In the worked example of constraining primordial non-Gaussianity at effective scales of $k \approx 740 \text{ Mpc}^{-1}$ from correlations between μ distortion and CMB temperature and E -mode polarization [1, 68], we find that the combination of SKA with planned CMB missions can significantly improve constraints, even in the case of spatially varying foreground properties. This example can likely be generalized to other CMB science cases, such as detecting B -modes generated from primordial gravitational waves during inflation.

5.2 Theory Background

μ -distortions are generated when energy injection into the early-universe photon-baryon plasma cannot be efficiently thermalized although Comptonization is still efficient. This leads to a Bose-Einstein distribution $n(\nu) = \left[e^{h\nu/(k_B T) + \mu(\nu)} - 1 \right]^{-1}$ [189, 115, 13, 63]. The frequency-dependent chemical potential, $\mu(\nu)$, is approximately constant above 500 MHz, but exponentially decays at lower frequencies due to double Compton scattering and Bremsstrahlung [15, 14]. A μ -type distortion is formed until redshift $z \simeq 5 \times 10^4$, after which a y -type dis-

tortion is produced [15, 16].

In standard cosmology, the primary energy injection that generates μ -distortions is diffusion damping of the primordial power spectrum $P_\zeta(k)$ at small scales, where the diffusion of photons effectively extracts energy from hotter regions and injects the energy to colder regions [189, 61, 62, 21]. The average μ -distortion from diffusion damping is related to the power spectrum as:

$$\langle \mu \rangle \propto \int_0^\infty dk \int_0^\infty dz \frac{P_\zeta(k) k^4}{2\pi^2} \frac{dk_{\text{D}}^{-2}}{dz} e^{-\frac{2k^2}{k_{\text{D}}^2}} \mathcal{J}_\mu(z), \quad (5.1)$$

where k_{D} is the damping scale, and \mathcal{J}_μ is the time window function for μ distortions [119, 17].

If inflation is driven by a single field initially in a Bunch–Davies vacuum, with no additional interactions, the resulting primordial curvature perturbations $\zeta(\vec{k})$ are purely Gaussian, and the power spectrum is statistically isotropic. The presence of additional fields and interactions during inflation introduces higher-order correlations between curvature perturbations, introducing statistical anisotropy, or non-Gaussianity. In the scenario where a long-wavelength curvature mode k_L is correlated with much smaller modes k_S , the long mode modulates the power spectrum at small scales. In Fourier space, this correlation, represented by a triangle configuration of momentum with two sides given by k_S and one by k_L , is known as the *squeezed-limit bispectrum*. In real space, this results in a spatially varying small-scale power spectrum, thus inducing *anisotropies* in the μ -distortion, which can be correlated with CMB modes [1, 68, 190, 191, 17, 144, 90, 192].

Following [18, 91], the angular cross-power spectrum of μ anisotropies and CMB modes is

$$C_\ell^{\mu X} = \frac{24\langle \mu \rangle}{5\pi} f_{\text{NL}} \int_0^\infty dk P_\zeta(k) k^2 \Delta_\ell^\mu(k) \Delta_\ell^X(k), \quad (5.2)$$

where f_{NL} parameterizes the amplitude of the correlation, $X \in \{T, E\}$ and Δ_ℓ^X is the corresponding transfer function, calculated using CAMB [125]. For the μ -anisotropy transfer function, we use $\Delta_\ell^\mu \approx e^{-k^2/(q_{\mu, \text{D}}^2(z_{\text{rec}}))} j_\ell(k\eta_0 - k\eta_{\text{rec}})$ [69, 17, 89, 18], with damping scale

$q_{\mu,D}(z_{\text{rec}}) \approx 0.11 \text{ Mpc}^{-1}$ [17].

5.3 Methods

5.3.1 Fisher Matrix

Following [70, 28, 91], we forecast constraints on f_{NL} from measurements of the $\mu \times T$ and $\mu \times E$ cross-spectra $C_\ell^{\mu T}$ and $C_\ell^{\mu E}$ using a Fisher-matrix approach. For our single-parameter model (given by Eq. 5.2 with fixed value for $\langle \mu \rangle$ and free parameter f_{NL}), the Fisher “matrix” is a scalar. Here, the shape of the spectrum is fixed, and f_{NL} controls the overall amplitude, yielding an analytical expression for the projected 1σ uncertainty on f_{NL} :

$$\sigma(f_{\text{NL}}) = \left(\sum_{\ell=\ell_{\text{min}}}^{\ell_{\text{max}}} \frac{(2\ell+1)f_{\text{sky}}}{C_\ell^{\mu\mu} [C_\ell^{TT} C_\ell^{EE} - (C_\ell^{TE})^2]} \times \right. \quad (5.3)$$

$$\left. \left[C_\ell^{TT} (C_\ell^{\mu E}|_{f_{\text{NL}}=1})^2 + C_\ell^{EE} (C_\ell^{\mu T}|_{f_{\text{NL}}=1})^2 - 2C_\ell^{TE} C_\ell^{\mu T}|_{f_{\text{NL}}=1} C_\ell^{\mu E}|_{f_{\text{NL}}=1} \right]^{-1/2} \right).$$

As in [91], we assume measurements of the primary CMB power spectra C_ℓ^{TT} , C_ℓ^{TE} , and C_ℓ^{EE} are signal-dominated at the scales of interest to this work, while $C_\ell^{\mu\mu}$ will be dominated by noise and foregrounds. Similarly to [91, 26], we model the contribution to the band-band ℓ -space covariance matrix $\mathbf{C}_\ell^{\text{ij}}$ from noise as

$$\mathbf{C}_\ell^{\text{ij},\text{N}} = (N^i)^2 e^{\ell^2 \theta_i^2 / (8 \ln(2))} \delta_{ij}, \quad (5.4)$$

where N^i is the white noise level in the map from frequency band i , θ_i is the beam FWHM in frequency band i , and from foreground X as

$$\mathbf{C}_\ell^{\text{ij},\text{X}} = \sqrt{C_{\text{X}}(\ell, \nu_i) C_{\text{X}}(\ell, \nu_j)}, \quad (5.5)$$

where $C_X(\ell, \nu_i)$ is the angular power spectrum for a given foreground at frequency ν_i . The $\mu \times \mu$ covariance matrix is then given by

$$C_\ell^{\mu\mu} = \sum_{ij} \mathbf{w}_{0i} \mathbf{C}_\ell^{ij} \mathbf{w}_{0j}, \quad (5.6)$$

where \mathbf{w} are the weights used to produce a T -free μ map from the individual-band intensity maps, using the component-separation formalism described in, e.g., Refs. [70, 91]. μ -distortions do not contain a polarized component, so we do not also construct an E -free μ -map.

5.3.2 Foregrounds

We closely follow the Galactic and extragalactic foreground treatment of [91, 26]. For Galactic foregrounds, we consider synchrotron, dust, and anomalous microwave emission. For extragalactic foregrounds, we consider the thermal Sunyaev-Zeldovich (tSZ) effect, the cosmic infrared background (CIB), and synchrotron-emitting active galactic nuclei. We model each foreground component as a power law in both frequency and ℓ space, adopting the values for power-law indices and “Wide survey” amplitude values in [26]. As indicated by Eq. 5.5, we assume each foreground component is 100% correlated across frequency bands.

Two low-frequency foregrounds that we neglect here are Galactic free-free emission and a potential source related to the low-frequency excess emission reported by the ARCADE team [193]. For free-free emission, we note that the free-free templates in the *Planck* Sky Model [135] and PySM [134] are dominated by point-like sources, either real HII regions (which can be masked) or contamination from extragalactic radio sources (which are already in our foreground model). If the ARCADE excess is a true astrophysical background that clusters at some level [194, 195, for discussion], it will contaminate the searches discussed here, but also make a new scientific target that SKA will be well-positioned to constrain.

5.3.3 Moment Expansion

Our foreground treatment implicitly assumes constant amplitude and spectral indices across the sky. While this roughly holds for extragalactic foregrounds, galactic foregrounds have significant spatial variations across the sky. When attempting to characterize extremely faint signals such as primordial B modes or μ distortions, the isotropic approximation can lead to significant biases on the parameters of interest [196, 197].

At the Fisher level, this can be treated using a moment expansion formalism [198], which we apply to account for spatially varying spectral indices for dust and synchrotron. We do not consider variations in their amplitude, given that the isotropic signal is so bright the covariance matrix effectively marginalizes out the corresponding spectral energy distribution (SED) [175]. We also ignore spatial variations in the dust temperature T_d since that is nearly degenerate with changes in amplitude in the frequency bands of interest to μ -distortions.

In this work, we consider the two largest contributions to the covariance from auto- or cross-power spectra of moment terms: 1×1 and 0×2 , in addition to 0×0 , the contribution from the component with the mean spectral behavior. We follow [199] and calculate these as:

$$\begin{aligned}
C_{\ell, \nu_1, \nu_2}^{X, 1 \times 1} &= \log\left(\frac{\nu_1}{\nu_0}\right) \log\left(\frac{\nu_2}{\nu_0}\right) \times \\
&\sum_{\ell_1 \ell_2} \frac{(2\ell_1 + 1)(2\ell_2 + 1)}{4\pi} \begin{pmatrix} \ell & \ell_1 & \ell_2 \\ 0 & 0 & 0 \end{pmatrix}^2 C_{\ell_1, \nu_1, \nu_2}^{X, 0 \times 0} C_{\ell_2}^{\beta_X}, \\
C_{\ell, \nu_1, \nu_2}^{X, 0 \times 2} &= \frac{\sigma_{\beta_X}^2}{2} \left[\log\left(\frac{\nu_1}{\nu_0}\right)^2 + \log\left(\frac{\nu_2}{\nu_0}\right)^2 \right] C_{\ell, \nu_1, \nu_2}^{X, 0 \times 0}, \\
\sigma_{\beta_X}^2 &\equiv \sum_{\ell} \frac{2\ell + 1}{4\pi} C_{\ell}^{\beta_X}, \quad C_{\ell}^{\beta_X} = B_X \left(\frac{\ell}{\ell_0}\right)^{\gamma_X}.
\end{aligned} \tag{5.7}$$

for pairs of frequency channels ν_1 and ν_2 . Here, the foreground components $X \in \{\text{dust, synchrotron}\}$ are normalized at the pivot frequency $\nu_0 = \{145, 93\}$ GHz, $\sigma_{\beta_X}^2$ is the variance in spectral index across the sky, and $C_{\ell}^{\beta_X}$ is the angular power spectrum of the spectral index variations, assumed to be a power law in ℓ with amplitude B_X and index γ_X .

5.3.4 Survey specifications

We forecast our constraints on f_{NL} from $\mu \times T$ and $\mu \times E$ based on Phase 1 of the planned SKA Observatory, augmented with data from existing or upcoming “traditional” CMB experiments. Phase 1 of SKA is under construction and encompasses two arrays: a “low-frequency” array (SKA1-LOW) in Australia that will observe from 50-350 MHz, and a “mid-frequency” array (SKA1-MID) in South Africa observing at 350 MHz to 15.4 GHz, with the goal of extending to 24 GHz.² Since the μ -distortion SED peaks below 1 GHz, both arrays have the potential to isolate and identify μ -distortions from foregrounds. Given the challenges of modeling the foreground behavior at low frequencies, we only consider the SKA1-MID array.

One of two proposed cosmology surveys for SKA1-MID is a “wide survey” of the southern sky [200]. In the default configuration for the wide survey, each telescope operates in single-dish mode, acting as an individual detector. The highest signal-to-noise ratio for $\mu \times T$ and $\mu \times E$ is at low ℓ , so single-dish mode is the preferred configuration for our purposes.

The angular resolution of each dish at center frequency ν_i is $\theta_i = 1.22c(D_{\text{dish}}\nu_i)^{-1}$, where $D_{\text{dish}} \approx 14$ meters is the dish diameter, and c is the speed of light. The noise for SKA1-MID in single-dish mode at frequency channel i is

$$N^i = \sqrt{\frac{T_{\text{sys},i}^2 S}{N_{\text{d}} N_{\text{b}} N_{\text{p}} t \Delta\nu}}, \quad (5.8)$$

where $S = 20,000 \text{ deg}^2$ is the area of the wide survey; $t = 10^4$ hours is the total on-sky time; $\Delta\nu$ is the channel bandwidth; N_{d} is the number of dishes, which for SKA1-MID is 197 [200]; N_{b} is the number of simultaneous observing beams, which is 1 for SKA1-MID; and N_{p} is the number of independent Stokes I measurements, which is 2. T_{sys} is the system temperature, with contributions from receiver noise T_{rcvr} , spill-over $T_{\text{spl}} \approx 3\text{K}$, the CMB (assuming negligible contribution from spectral distortions) $T_{\text{CMB}} = 2.73\text{K}$, and Galactic emission $T_{\text{gal}} \approx 25\text{K}(408 \text{ MHz}/\nu_i)^{2.75}$ [200]. Above 1 GHz, we assume $T_{\text{rcvr}} = 7.5\text{K}$

2. skao.int/en/science-users/118/ska-telescope-specifications

and below we set $T_{\text{RCVR}} = 15\text{K} + 30\text{K} \left(\frac{\nu_i}{1\text{GHz}} - 0.75 \right)^2$ [200]. We neglect the atmospheric contribution, noting that at Karoo, South Africa it will be subdominant in the total sky contribution compared to T_{CMB} and T_{gal} .³

The SKA1-MID receiver bands will be sub-divided into 65,000 frequency channels, useful for identifying and removing RFI and other systematics, and can then be combined into wider bands. For frequencies below 2.5 GHz, we set $\Delta\nu = 100$ MHz to better isolate the peak of the μ spectrum, while for frequencies above 2.5 GHz we set $\Delta\nu = 1$ GHz to obtain lower noise for improved calibration off of the CMB.

We propose to pair SKA1-MID with a wide survey at traditional CMB frequencies ($\nu \gtrsim 20$ GHz), from which we can obtain signal-dominated maps of T and E anisotropies, and to further enhance foreground subtraction in the μ map. Given that we are targeting low- ℓ ($\ell \lesssim 100$) μ -anisotropies, sensitivity at large angular scales is more important than angular resolution. For foreground removal, more individual frequency channels are preferred. One survey that satisfies these criteria is the all-sky survey planned for the space-based telescope LiteBIRD [27], the primary science goal of which is constraining low- ℓ CMB B-mode polarization. The angular resolution θ_i and noise level N^i of each frequency band for LiteBIRD is given in Table 1 of [28]. SKA1-MID’s wide survey corresponds to a $f_{\text{sky}} = 0.48$, which we will limit our forecast of LiteBIRD to. Fig. 5.1 shows SKA1-MID and LiteBIRD fill highly complementary roles: LiteBIRD is more sensitive to signals with a blackbody SED, while the low-frequency coverage of SKA1-MID results in much better sensitivity to μ . This also means that signals at much higher- ℓ will provide diminishing returns on constraining $\sigma(f_{\text{NL}})$. As such, we take $\ell_{\text{max}} = 1000$.

5.3.5 Calibration

Precise inter-frequency calibration is needed to extract the very small μ -distortion signal from the much larger T and foreground signals. A small gain mismatch between bands will

3. skao.int/sites/default/files/documents/Anticipated %20Performance%20of%20the%20SKA.pdf




Figure 5.1: C_ℓ^N for SKA1-MID and LiteBIRD, neglecting foregrounds. Solid lines correspond to $C_\ell^{\mu\mu,N}$, while dashed lines are $C_\ell^{TT,N}$. Despite SKA1-MID’s temperature noise being almost two orders of magnitude worse than LiteBIRD’s, SKA1-MID has substantially lower noise for μ owing to its low-frequency coverage.

cause $T \rightarrow \mu$ leakage in the component separation, resulting in a $T \times T$ component to the $\mu \times T$ spectrum. Even if this can be modeled, it will cause excess variance in the $\mu \times T$ spectrum and degrade the constraints on f_{NL} . This implies that the leaked T should be kept below the level of the noise in the μ map, which in turn implies that the precision of the relative calibration between bands should be comparable to the map noise divided by the T signal or better. This is satisfied by maps in which the dominant signal is the CMB temperature anisotropy, which is true for the workhorse bands of LiteBIRD, but is not for SKA. For SKA bands, which are expected to be dominated by synchrotron emission at the multipole ranges of interest, a potential strategy is to first model the synchrotron and

then calibrate on the CMB using cross-spectra with the CMB-S4 maps. The synchrotron model could be varied, with the calibration effectively marginalized over the uncertainty in the synchrotron model parameters. In principle, this procedure can also provide a check on the model of the instrument angular response function (or beam) and its dependence on SKA observing frequency. This overall calibration strategy will be tested in upcoming work; for the purposes of this forecast we assume perfect calibration and knowledge of instrument beams.

5.4 Results and Discussion

$\sigma(f_{\text{NL}}), \ell_{\text{max}} = 1000$	No foregrounds	Foregrounds	+Moment expansion
LiteBIRD	91	826	923
SKA1-MID	8 (19/33)	31 (56/130)	234 (505/805)
LiteBIRD & SKA1-MID	6 (13/19)	23 (61/124)	92 (141/181)
+ 15-24 GHz	6 (13/18)	22 (56/100)	68 (100/126)

Table 5.1: Constraints on f_{NL} from different survey configurations of SKA1-MID and LiteBIRD, with $\ell_{\text{max}} = 1000$. We list constraints considering no foregrounds, isotropic foregrounds, and foregrounds with moment expansion terms from III.C. Constraints not in parentheses assume $\ell_{\text{min}} = 2$, while the ones in parentheses are given for ($\ell_{\text{min}} = 50/\ell_{\text{min}} = 100$)

Table 5.1 summarizes the forecasted constraints on f_{NL} from SKA1-MID, LiteBIRD, and their combination. In all cases, adding SKA1-MID data improves the f_{NL} constraint over LiteBIRD alone by a factor of at least 10. In the most ideal cases, SKA1-MID alone is nearly as powerful as the combination, but in the most realistic foreground case, the combination is a factor of 3-4 more powerful than SKA1-MID, highlighting the synergy between SKA and traditional CMB experiments.

In terms of the absolute level of the f_{NL} constraints, constraints approach $\sigma(f_{\text{NL}}) = 6$ in the most ideal case, comparable to what a cosmic variance-limited measurement of the CMB bispectrum can achieve for large-scale non-Gaussianity, but at much smaller scales ($k \approx 740 \text{ Mpc}^{-1}$, see Fig. 3.6). In the most realistic case, combined constraints degrade to $\sigma(f_{\text{NL}}) = 92$, which would still be the strongest constraint on non-Gaussianity at such

small scales, improving current constraints from $\mu \times T$ [201] by a factor of $\simeq 30$. The largest degradation in combined constraints comes from introducing spatially varying foregrounds (specifically synchrotron—removing dust entirely from the covariance has almost no effect on $\sigma(f_{\text{NL}})$).

We have made certain assumptions in this analysis that may be optimistic when compared to real data. We have ignored the effects of instrumental “ $1/f$ ” or “red” noise (see, e.g., [202]). In single-dish mode, instrumental $1/f$ noise will corrupt angular modes with scales larger than $\sim v_{\text{scan}}/f_{\text{knee}}$ in the scan direction, where v_{scan} is the telescope scanning velocity, and f_{knee} is the frequency at which the $1/f$ noise equals the white noise. The SKA precursor experiment MeerKAT has achieved a raw $f_{\text{knee}} \simeq 0.1 \text{ Hz}$ [203], which would contaminate modes with $\ell \lesssim 50$, for a scan velocity of 1 degree/s. To account for $1/f$ noise with a plausible range of f_{knee} and scan speed values, we include in parentheses in Table 5.1 results for $\sigma(f_{\text{NL}})$ with $\ell_{\text{min}} = 50$ and 100 to indicate the degradation. In addition, we have ignored free-free emission and the ARCADE excess as foregrounds, choices motivated in §5.3.2. Finally, we also ignored calibration uncertainties, a more robust estimate of which will inform future forecasting for SKA.

There are also aspects of our analysis that might be overly pessimistic. The assumed bands are much wider than SKA1-MID’s planned capability of dividing bands into $\Delta\nu \sim 10$ kHz channels. The use of very narrow bands can improve our ability to isolate the μ -anisotropy signal. In addition, we are assuming a constant synchrotron amplitude across the Southern sky, using the mean value over the $f_{\text{sky}} = 0.48$ region treated in [26]. A more careful choice of observing region, and dividing the region into multiple patches, could reduce the impact of synchrotron significantly.

This single worked example of constraining primordial non-Gaussianity through correlations between μ -distortion anisotropy and CMB temperature and polarization anisotropies demonstrates the impressive potential of SKA as a CMB experiment, particularly when combined with a traditional, higher-frequency CMB experiment such as LiteBIRD. Our forecasts

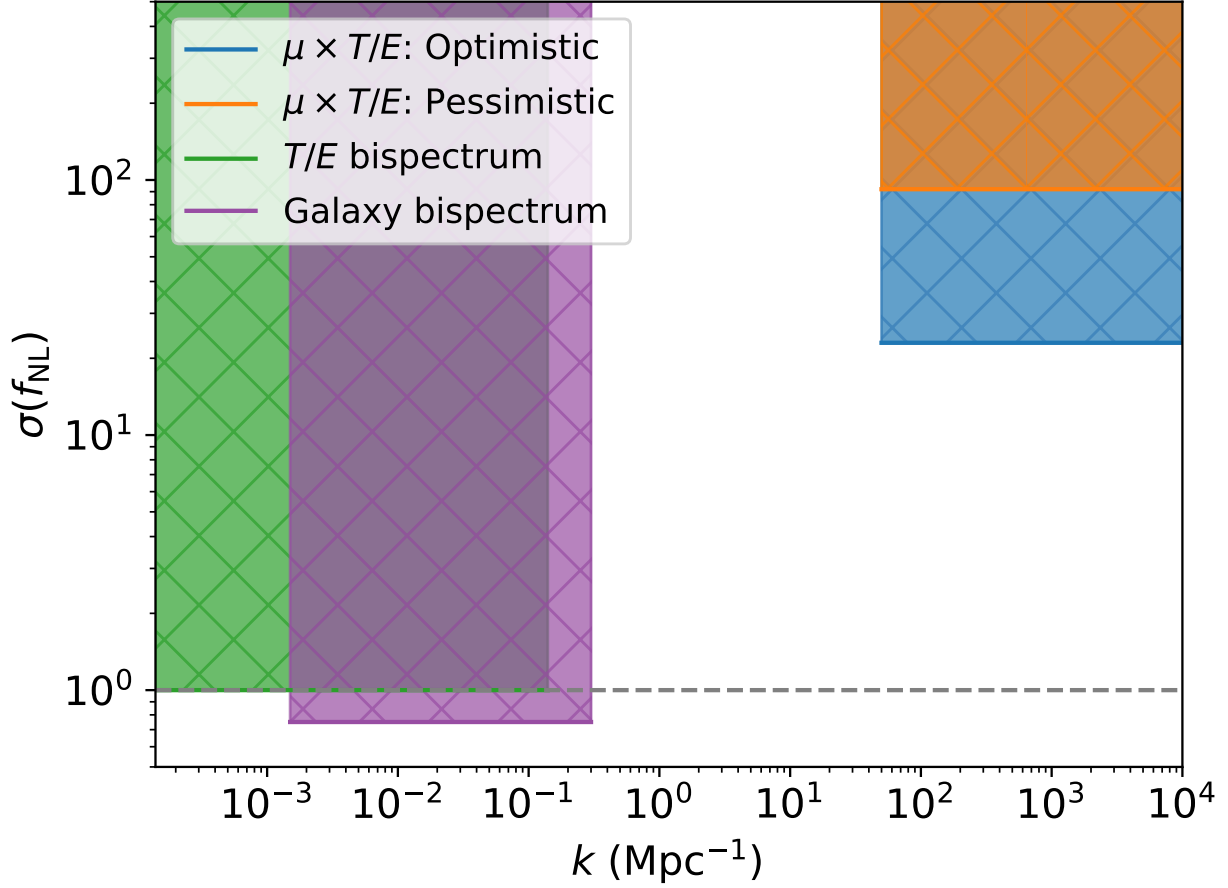


Figure 5.2: Forecasted upper bounds on squeezed-limit f_{NL} at different scales from various observables. These scales are approximations and purely for illustration. CMB bispectrum constraints are from CMB-S4 with scales derived from $\ell_{\text{min}} = 2$ and $\ell_{\text{max}} = 1000$. Galaxy bispectrum constraints are from SPHEREx [3] with $k_{\text{min}} \approx 0.0015 \text{ Mpc}^{-1}$ and $k_{\text{max}} = 0.3 \text{ Mpc}^{-1}$ [4]. For μ -anisotropies from SKA1+LiteBIRD, “optimistic” is the second column of Table 5.1, while “pessimistic” is the third column. For reference, the gray dashed line is $f_{\text{NL}} = 1$. Figure inspired by [5, 6].

can also be used to constrain other sources of μ -anisotropies, such as modulated thermalization in the μ -era by a long-wavelength curvature mode [18, 19] or due to energy injections [122, 123, 124]. As seen in Fig. 5.2, our constraints on f_{NL} place bounds on primordial non-Gaussianity at scales inaccessible to other cosmological observables. If squeezed-limit non-Gaussianity grows at smaller scales (i.e has a blue-tilted spectrum), it results in a proportional improvement on our constraints on f_{NL} .

Future forecasts will expand on this single example to other CMB science cases, including B -modes from primordial gravitational waves (PGW). A simplistic forecast for PGW shows that SKA1-MID alone has the raw statistical power to constrain the tensor-to-scalar ratio r to $\sigma(r) \leq 10^{-3}$, which is comparable to the expected constraints from LiteBIRD [27] and CMB-S4 [130]. Small-scale information from SKA1-MID in interferometer mode can be used to reconstruct the CMB gravitational lensing potential out to very small scales, both for use in cleaning the single-dish data from B modes induced by lensing [204] and to probe dark matter models [205]. Future forecasts will also include the even more impressive raw sensitivity of the planned SKA2 upgrade³, which will be at least an order of magnitude more sensitive than SKA1 and should improve the f_{NL} and PGW constraints by roughly the same factor.

APPENDIX A

SUPPLEMENTAL MATERIAL

A.1 Dilation Consistency Relation at Second Order

In this Appendix, we show that at second order in Newtonian gauge, the response of the short wavelength density perturbation to the long wavelength ζ_L is a pure dilation and when combined with the Maldacena consistency relation there is no locally observable effect of a constant ζ_L . Since we focus on the μ -distortion, mainly produced during the radiation dominated epoch, we concretely calculate the second-order perturbations in that era. Note that the calculation and notation here, which differs from the main text, is based on Ref. [206].

First, let us summarize our notation in this appendix. In Newtonian gauge, we can write the scalar parts of the metric perturbations as

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu = a^2 \left\{ -(1 + 2\Phi^{(1)} + \Phi^{(2)}) d\eta^2 + \left[(1 - 2\Psi^{(1)} - \Psi^{(2)}) \delta_{ij} \right] dy^i dy^j \right\}, \quad (\text{A.1})$$

where the superscript denotes the order in perturbations. Here, we assume a perfect fluid for simplicity, which leads to $\Phi^{(1)} = \Psi^{(1)}$, and thus ignore the correction from neutrinos. Then, we can express the energy-momentum tensor as

$$T^\mu_\nu = (\rho + P) u^\mu u_\nu + P \delta^\mu_\nu, \quad (\text{A.2})$$

where ρ is the energy density, P is the pressure, and u_μ is the 4-velocity. We take the following notation for their perturbations up to the second order:

$$\rho = \rho^{(0)} + \delta\rho^{(1)} + \frac{1}{2}\delta\rho^{(2)}, \quad (\text{A.3})$$

$$P = P^{(0)} + \delta P^{(1)} + \frac{1}{2}\delta P^{(2)}, \quad (\text{A.4})$$

$$u^i = \frac{1}{a} \left(\delta v^{(1),i} + \frac{1}{2}\delta v^{(2),i} + \frac{1}{2}\delta v_V^i{}^{(2)} \right), \quad (\text{A.5})$$

where δv is the velocity potential and δv_V^i is the vector part of the velocity perturbation. We define the density perturbation, $\delta^{(a)}$, as

$$\delta^{(a)} \equiv \frac{\delta \rho^{(a)}}{\rho^{(0)}}. \quad (\text{A.6})$$

At the first order in perturbations, this is related to the curvature perturbation as

$$\delta_{\mathbf{k}}^{(1)}(\eta) = -\frac{2}{3}\zeta_{\mathbf{k}}^{(1)}T_{\delta}(x), \quad (\text{A.7})$$

where the subscript \mathbf{k} represents the Fourier mode whereas perturbations are in real space otherwise in this Appendix, $x \equiv k\eta$ with $k \equiv |\mathbf{k}|$, and $\zeta^{(1)} = -(3/2)\Phi^{(1)}$ in the superhorizon limit. The transfer function is defined as

$$T_{\delta}(x) \equiv \frac{6x(-6+x^2)\cos(\frac{x}{\sqrt{3}}) - 12\sqrt{3}(-3+x^2)\sin(\frac{x}{\sqrt{3}})}{x^3}. \quad (\text{A.8})$$

Next, let us calculate the second-order density perturbations. The density perturbation is related to the other perturbations as [206]

$$\begin{aligned} \delta^{(2)} = & -2\Psi^{(2)} + 2N_j^i B_j^{i(2)} - \frac{2}{\mathcal{H}}\Psi^{(2)'} + \frac{2}{3\mathcal{H}^2}\Psi^{(2),i} + \frac{2}{\mathcal{H}^2}\left(\Phi^{(1)'}\right)^2 + \frac{16}{3\mathcal{H}^2}\Phi^{(1)}\Phi^{(1),i} \\ & + \frac{1}{\mathcal{H}^2}\left(2 - \frac{8}{9(1+w)}\right)\Phi^{(1),i}\Phi^{(1)},i - \frac{8}{9(1+w)\mathcal{H}^3}(\Phi^{(1),i}\Phi^{(1)},i)' - \frac{8}{9(1+w)\mathcal{H}^4}\Phi^{(1),i'}\Phi^{(1),i'}, \end{aligned} \quad (\text{A.9})$$

where the prime here denotes the derivative with respect to η , $\mathcal{H}(\equiv a'/a)$ is the conformal Hubble parameter, and w is the equation of state parameter. B_j^i and N_j^i are defined as

$$B_j^{i(2)} \equiv \left[\frac{4(5+3w)}{3(1+w)}\Phi^{(1),i}\Phi^{(1)},j + \frac{8}{3(1+w)\mathcal{H}}\left(\Phi^{(1),i}\Phi^{(1)},j\right)' + \frac{8}{3(1+w)\mathcal{H}^2}\Phi^{(1),i'}\Phi^{(1)},j' \right], \quad (\text{A.10})$$

$$\begin{aligned}
N_i^j A_j^i(\mathbf{x}) &\equiv \frac{3}{2} \nabla^{-2} \left(\frac{\partial^j \partial_i}{\nabla^2} - \frac{1}{3} \delta_i^j \right) A_j^i(\mathbf{x}) \\
&= \int \frac{d^3 k}{(2\pi)^3} \left(-\frac{3}{2k^2} \right) \left(\frac{k^j k_i}{k^2} - \frac{1}{3} \delta_i^j \right) A_{\mathbf{k}}^i{}^j,
\end{aligned} \tag{A.11}$$

where A_j^i is an arbitrary tensor. Since $\delta^{(2)}$ depends on $\Psi^{(2)}$, we need to calculate $\Psi^{(2)}$ first. The equation of motion for $\Psi^{(2)}$ is [206]

$$\Psi^{(2)''} + 3(1 + c_s^2) \mathcal{H} \Psi^{(2)'} + \left[2\mathcal{H}' + (3c_s^2 + 1) \mathcal{H}^2 \right] \Psi^{(2)} - c_s^2 \Psi^{(2),i}{}_{,i} = S^{(2)}, \tag{A.12}$$

$$\begin{aligned}
S^{(2)} &\equiv \left(3c_s^2 - \frac{1}{3} \right) \Phi^{(1),i} \Phi^{(1)}{}_{,i} + 8c_s^2 \Phi^{(1)} \Phi^{(1),i}{}_{,i} + (3c_s^2 + 1) \left(\Phi^{(1)'} \right)^2 + \\
&\quad \left[(3c_s^2 + 1) \mathcal{H}^2 + 2\mathcal{H}' \right] N_i^j B_j^i{}^{(2)} + \mathcal{H} N_i^j \left(B_j^i{}^{(2)} \right)' + \frac{1}{3} N_i^j \left(B_j^i{}^{(2)} \right)'{}_{,k} \\
&\quad + \left(\frac{1}{3} - c_s^2 \right) \frac{4}{3(1+w) \mathcal{H}^2} \left(\mathcal{H} \Phi^{(1),i} + \Phi^{(1),i'} \right) \left(\mathcal{H} \Phi^{(1)}{}_{,i} + \Phi^{(1)'}{}_{,i} \right),
\end{aligned} \tag{A.13}$$

where c_s is the sound speed.

Here, we focus on the second-order perturbations induced by a superhorizon mode “ L ” on the much smaller mode $k_L/k \ll 1$. This case corresponds to the μ -distortion production in the presence of the superhorizon perturbations. In this case, we can approximate Eq. (A.12) as

$$\Psi^{(2)''} + 4\mathcal{H} \Psi^{(2)'} - \frac{1}{3} \Psi^{(2),i}{}_{,i} \simeq \frac{8}{3} \Phi_L^{(1)} \Phi^{(1),i}{}_{,i}, \tag{A.14}$$

where we have substituted $c_s^2 = w = 1/3$. The other contributions in $S^{(2)}$ are sub-leading because they include spatial derivative on the superhorizon mode “ L ”. In the Fourier space, the equation becomes

$$\Psi_{\mathbf{k}}^{(2)''} + \frac{4}{\eta} \Psi_{\mathbf{k}}^{(2)'} + \frac{k^2}{3} \Psi_{\mathbf{k}}^{(2)} \simeq -\frac{8}{3} \Phi_L^{(1)} k^2 \Phi_{\mathbf{k}}^{(1)}, \tag{A.15}$$

where we have assumed that the superhorizon perturbation as the constant quantity, Φ_L , because it does not evolve in the Newtonian gauge. Note that Φ_L here and below is still a function of spatial comoving coordinates but variations are small locally compared to the short wavelength mode so we have ignored them in the Fourier transform. Then, we can solve $\Psi_{\mathbf{k}}^{(2)}$ using the Green function method:

$$\Psi_{\mathbf{k}}^{(2)}(\eta) \simeq \Psi_{\mathbf{k}}^{(2)}(0)T(x) + \int_0^\eta d\bar{\eta} \left(\frac{a(\bar{\eta})}{a(\eta)} \right)^2 G(k, \eta; \bar{\eta}) \left(-\frac{8}{3} \Phi_L k^2 \Phi_{\mathbf{k}}^{(1)}(\bar{\eta}) \right), \quad (\text{A.16})$$

where the concrete expression of the Green function is given by

$$kG(k, \eta; \bar{\eta}) = -\Theta(\eta - \bar{\eta}) \frac{x\bar{x}}{\sqrt{3}} \left[j_1(x/\sqrt{3})y_1(\bar{x}/\sqrt{3}) - j_1(\bar{x}/\sqrt{3})y_1(x/\sqrt{3}) \right]. \quad (\text{A.17})$$

The initial condition of $\Psi^{(2)}$ is given by

$$\Psi_{\mathbf{k}}^{(2)}(0) = -\frac{16}{9} \zeta_L^{(1)} \zeta_{\mathbf{k}}^{(1)}, \quad (\text{A.18})$$

where we have here assumed a Gaussian distribution of ζ , which means that the contribution from primordial non-Gaussianity from the Maldacena consistency relation is not included in $\Psi^{(2)}$ (and $\delta^{(2)}$) in this expression. We will independently take into account that contribution later. Then, we obtain $\Psi^{(2)}$ as

$$\Psi_{\mathbf{k}}^{(2)}(\eta) \simeq \left[\frac{16\sqrt{3}}{3} \frac{\sin(x/\sqrt{3})}{x} + \frac{64 \cos(x/\sqrt{3})}{x^2} \right] \zeta_L^{(1)} \zeta_{\mathbf{k}}^{(1)}, \quad (\text{A.19})$$

where we have neglected the contribution of $\mathcal{O}(x^{-3})$ because we want to know the evolution of the second-order density perturbations on subhorizon scales, which is related to the μ -distortion production. In the large x limit, Eq. (A.9) can be approximated in the Fourier

space as

$$\begin{aligned}\delta_{\mathbf{k}}^{(2)} &\simeq -\frac{2x^2}{3}\Psi_{\mathbf{k}}^{(2)} - 2x\frac{d\Psi_{\mathbf{k}}^{(2)}}{dx} - \frac{16x^2}{3}\Phi_L^{(1)}\Phi_{\mathbf{k}}^{(1)} \\ &\simeq \left(-\frac{32\sqrt{3}}{9}x\sin(x/\sqrt{3}) - 32\cos(x/\sqrt{3})\right)\zeta_L^{(1)}\zeta_{\mathbf{k}}^{(1)}.\end{aligned}\quad (\text{A.20})$$

Then, we finally get the following expression for the total energy density perturbation:

$$\begin{aligned}\delta_{\mathbf{k}}(\eta) &= \delta_{\mathbf{k}}^{(1)}(\eta) + \frac{1}{2}\delta_{\mathbf{k}}^{(2)}(\eta) \\ &\simeq -\frac{2}{3}\left(1 + \frac{4}{3}\zeta_L^{(1)}\right)T_\delta\left(x\left(1 - \frac{4}{3}\zeta_L^{(1)}\right)\right)\zeta_{\mathbf{k}}^{(1)},\end{aligned}\quad (\text{A.21})$$

where the approximate equality is valid in $x \gg 1$.

In the following, we relate this result to the dilation transformation. The dilation consistency relation in the Newtonian gauge is given by [207]

$$\lim_{\mathbf{q} \rightarrow 0} \frac{\langle \pi_{\mathbf{q}} \delta_{\mathbf{k}_1} \cdots \delta_{\mathbf{k}_N} \rangle^c}{P_\pi(q)} \epsilon(\eta) = \sum_{a=1}^N \left[-\epsilon(\eta_a) \left(\left. \frac{\bar{\rho}'}{\bar{\rho}} \right|_{\eta_a} + \partial_{\eta_a} \right) + \lambda (3 + \mathbf{k}_a \cdot \partial_{\mathbf{k}_a}) \right] \langle \delta_{\mathbf{k}_1} \cdots \delta_{\mathbf{k}_N} \rangle^c, \quad (\text{A.22})$$

where $\langle \pi_{\mathbf{q}} \pi_{\mathbf{p}} \rangle = (2\pi)^3 \delta_D(\mathbf{q} + \mathbf{p}) P_\pi(q)$, π is the velocity potential ($\pi = \delta v$), and the superscript c of the bracket means the connected part. The ϵ and λ are the coordinate transformation parameters to go from a gauge where the effect is a pure dilation and are given by $y^\mu \rightarrow y'^\mu = y^\mu + \xi^\mu$ with $\xi^0 = \epsilon$ and $\xi^i = \lambda y^i$. During a radiation dominated epoch, we can derive the following relations: $\lambda = -3\epsilon/\eta$, $\bar{\rho}'/\bar{\rho} = -4/\eta$, and $\pi \simeq \eta\zeta/3$ in the superhorizon limit. Here, we take $\lambda = \zeta_L$ and $N = 2$ to determine the modulation of short wavelength power by ζ_L , such that the transformation is from comoving gauge to Newtonian gauge and

rewrite the consistency relation as

$$\lim_{\mathbf{q} \rightarrow 0} \frac{\langle \zeta_{\mathbf{q}} \delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \rangle^c}{P_{\zeta}(q)} \zeta_L^{(1)} = \sum_{a=1}^2 \left[\left(\frac{4}{3} \zeta_L^{(1)} - \frac{1}{3} \zeta_L^{(1)} \eta \partial_{\eta} \right) - \zeta_L^{(1)} (3 + \mathbf{k}_a \cdot \partial_{\mathbf{k}_a}) \right] \langle \delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \rangle^c. \quad (\text{A.23})$$

Note again that ζ_L is in real space as $\zeta_L(\mathbf{y})$, but its perturbation scale is much larger than the scales of \mathbf{k}_1 and \mathbf{k}_2 so that we can approximately consider it constant except when considering the correlation to the large scale mode \mathbf{q} . Using Eq. (A.21), we can rewrite the left-hand side as

$$\begin{aligned} \lim_{\mathbf{q} \rightarrow 0} \frac{\langle \zeta_{\mathbf{q}} \delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \rangle^c}{P_{\zeta}(q)} \zeta_L^{(1)} &= \frac{4}{9} \left[\left(1 + \frac{8}{3} \zeta_L^{(1)} \right) T_{\delta}^2 \left(x_1 \left(1 - \frac{4}{3} \zeta_L^{(1)} \right) \right) - T_{\delta}^2(x_1) \right] (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2) P_{\zeta}(k_1) \\ &\quad - \zeta_L^{(1)} \frac{4}{9} T_{\delta}^2(x_1) \frac{d \ln \Delta_{\zeta}^2}{d \ln k} \Big|_{k=k_1} (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2) P_{\zeta}(k_1), \end{aligned} \quad (\text{A.24})$$

where we have used

$$\begin{aligned} \langle \zeta_{\mathbf{q}}^{(1)} \zeta_L^{(1)}(\mathbf{y}) \rangle^c &= \int \frac{d^3 k}{(2\pi)^3} e^{i\mathbf{k} \cdot \mathbf{y}} \langle \zeta_{\mathbf{q}}^{(1)} \zeta_{\mathbf{k}}^{(1)} \rangle^c \\ &= e^{-i\mathbf{q} \cdot \mathbf{y}} P_{\zeta}(q). \end{aligned} \quad (\text{A.25})$$

The first line in Eq. (A.24) corresponds to the contribution from $\delta^{(2)}$. On the other hand, the second line corresponds to the Maldacena consistency relation [71], that is, it just comes from the contribution proportional to $\langle \zeta_{\mathbf{q}} \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} \rangle$, not related to the evolution of the second-order scalar perturbations. On the other hand, once we substitute

$$\langle \delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \rangle^c = \frac{4}{9} T_{\delta}^2(x_1) (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2) P_{\zeta}(k_1) \quad (\text{A.26})$$

into the right-hand side of Eq. (A.23), we can see that the right-hand side is the same as Eq. (A.24) at least at the lowest order in the perturbations, $\mathcal{O}(\zeta_L P_{\zeta})$. From this, we can see that the dilation consistency relation is satisfied, once both the second-order evolution and

primordial non-Gaussianity of the Maldacena relation are included, which indicates that the constant ζ_L does not give locally observable effects. Although our calculation is based on the perfect fluid assumption, we can expect that the dilation consistency relation would be satisfied even in imperfect fluid.

A.2 Damping

To understand how different effects of damping impact the angular power spectrum C_ℓ , we see that the spectrum scales as

$$C_\ell \propto \int_0^{+\infty} dk P_\zeta(k) k^2 [\Delta_\ell(k)]^2, \quad (\text{A.27})$$

where the transfer function $\Delta_\ell(k)$ is

$$\Delta_\ell(k) = D_\ell \times j_\ell(k\Delta\eta), \quad (\text{A.28})$$

and D_ℓ is the damping factor of anisotropies. For temperature anisotropies, the formulation that accounts for both diffusion damping and the finite width of the SLS is given by Ref. [7], which can be approximated by their Eq. (16).

For μ anisotropies, since we can only reliably calculate the effects of diffusion damping, we approximate the damping factor as $D_\ell = e^{-k^2/(q_{\mu,D}^2)}$, where $q_{\mu,D}$ is the damping scale of μ -distortion anisotropies. For the diffusion damping scale of μ -anisotropies, we use the value from Ref. [17], $q_{\mu,D} = 0.11 \text{ Mpc}^{-1}$. This can be compared to the diffusion-only damping scale for temperature anisotropies, $q_{T,D} = 0.14 \text{ Mpc}^{-1}$.

We show the effect of the two damping functions (the diffusion-only μ -anisotropy damping function and the temperature anisotropy damping function from Ref. [7]) on an undamped spectrum in Figure A.1. At $\ell = 200$ we see that the damping corresponding to temperature anisotropies is more severe than the damping corresponding to μ anisotropies. However, damping due to the finite width of the surface of last-scattering makes the slope of the

damping less steep [7]. This results in the μ anisotropy spectrum being more severely damped than the temperature anisotropy spectrum at $\ell \approx 1000$.

Table A.1 shows our constraints on f_{NL} when using the damping factor from Ref. [7]. We find a degradation of our constraints on f_{NL} of 15% or less compared to our nominal results in Table 3.5. This change is much smaller than the effects of foregrounds and the atmosphere.




Figure A.1: Change in the angular power spectrum C_ℓ depending on the damping factor. For illustrative purposes, we choose an undamped spectrum that is nearly scale invariant, with a slight spectral tilt. The spectrum with a damping scale $q_{\mu,D} = 0.11 \text{ Mpc}^{-1}$ corresponds to our treatment of the damping of μ -anisotropies. The spectrum using the damping factor from Ref. [7] has more severe damping at $\ell = 200$, but has a slightly less steep slope compared to the μ -anisotropy damping scale. The black line corresponding to $\ell = 200$ is a reference scale for the reader.

$\sigma(f_{\text{NL}})$, “null-y” $\ell_{\text{max}} = 2000$	S4-Deep + <i>Planck</i> (no foregrounds)	S4-Deep + <i>Planck</i>
No atmosphere	3%	6%
100% correlated	3%	6%
99% correlated	12%	12%
90% correlated	14%	13%
70% correlated	15%	14%

Table A.1: Percent degradation relative to our baseline constraints on f_{NL} (shown in Table 3.5) when using the damping scale from Ref. [7] instead of the fiducial value of 0.11 Mpc^{-1} . The degradation in the constraints is less than 15% in all cases.

A.3 Flat Sky Harmonics

For a function on the sky $y(\vec{n} = \{\theta, \phi\})$ with support only on a small area $\theta \ll 1$ around the pole (in the main text, center of the cluster), we can directly relate the spherical harmonic $y_{\ell m}$ and flat sky y_ℓ harmonic coefficients

$$\begin{aligned}
y(\vec{n}) &= \sum_{\ell m} y_{\ell m} Y_{\ell m}(\vec{n}) \\
&\approx \int \frac{d^2\ell}{(2\pi)^2} y(\vec{\ell}) e^{i\vec{n}\cdot\vec{\ell}},
\end{aligned}
\tag{A.29}$$

using an approximation for $Y_{\ell m}$ itself in an elaboration of the derivation in Ref. [208]. This approximation follows from the relation ([209], 8.722.2)

$$\ell^m P_\ell^{-m}(\cos\theta) \approx J_m(\ell\theta)
\tag{A.30}$$

for $m \geq 0$ and $\ell \gg 1$. We can use the fact that

$$J_{-m}(x) = (-1)^m J_m(x),
\tag{A.31}$$

$$P_\ell^{-m} = (-1)^m \frac{(\ell - m)!}{(\ell + m)!} P_\ell^m,
\tag{A.32}$$

and

$$Y_{\ell m} = \sqrt{\frac{2\ell+1}{4\pi} \frac{(\ell-m)!}{(\ell+m)!}} P_{\ell}^m(\cos\theta) e^{im\phi} \quad (\text{A.33})$$

to obtain for all m

$$Y_{\ell m} \approx \ell^{-|m|} \sqrt{\frac{(\ell+|m|)!}{(\ell-|m|)!}} (-1)^m \sqrt{\frac{2\ell+1}{4\pi}} J_m(\ell\theta) e^{im\phi}. \quad (\text{A.34})$$

When transforming functions with support only near the pole only $|m| \ll \ell$ modes contribute substantially due to the rapid variation of higher modes with ϕ , so it is a good approximation to cancel the factorials with $\ell^{-|m|}$ and use

$$Y_{\ell m} \approx (-1)^m \sqrt{\frac{2\ell+1}{4\pi}} J_m(\ell\theta) e^{im\phi}, \quad (|m| \ll \ell). \quad (\text{A.35})$$

Note that we can always orient the pole of the spherical coordinate system to align with the region of support. We can now obtain the desired relation between the two coefficients in Eq. (A.29) [208]

$$\begin{aligned} y(\vec{\ell}) &\approx \sqrt{\frac{4\pi}{2\ell+1}} \sum_m i^{-m} y_{\ell m} e^{im\varphi_{\ell}}, \\ y_{\ell m} &\approx \sqrt{\frac{2\ell+1}{4\pi}} \int \frac{d\varphi_{\ell}}{2\pi} e^{-im\varphi_{\ell}} y(\vec{\ell}), \end{aligned} \quad (\text{A.36})$$

where φ_{ℓ} is the azimuthal angle ϕ that $\vec{\ell}$ points at the pole. In particular if the function is azimuthally symmetric around the pole only $m = 0$ coefficients contribute and

$$y(\vec{\ell}) = 2\pi \int \theta d\theta J_0(\ell\theta) y(\theta) \approx \sqrt{\frac{4\pi}{2\ell+1}} y_{\ell 0}. \quad (\text{A.37})$$

It is common in the literature to slightly improve on the accuracy of the underlying approximation (A.30) at low ℓ by taking the argument of the Bessel function as $\ell\theta \rightarrow (\ell + 1/2)\theta$ and correspondingly $\ell^2 \rightarrow \ell(\ell + 1)$, e.g. in the Gaussian beam profile formula (4.25).

A.4 Relativistic Corrections

Following Refs. [210, 174], we can derive the relativistic corrections to the y distortion of an initial μ distortion using the generalized Kompaneets equation which is the expansion of the Compton collision term to the Boltzmann equation in the small energy transfer due to scattering. To first order in $\theta_e \equiv k_B T_e / m_e c^2$, Eq. (4.1) is generalized to

$$\frac{\partial f}{\partial \tau} = \theta_e \sum_{n=1}^4 x_e^n I_n \left[(1+f) \left(\frac{\partial}{\partial x_e} + 1 \right)^n - \frac{\partial^n f}{\partial x_e^n} \right] f \quad (\text{A.38})$$

with

$$\begin{aligned} I_1 &= 4 - x_e + \left(10 - \frac{47}{2}x_e + \frac{21}{5}x_e^2\right)\theta_e, \\ I_2 &= 1 + \left(\frac{47}{2} - \frac{63}{5}x_e + \frac{7}{10}x_e^2\right)\theta_e, \\ I_3 &= \left(\frac{42}{5} - \frac{7}{5}x_e\right)\theta_e, \\ I_4 &= \frac{7}{10}\theta_e. \end{aligned} \quad (\text{A.39})$$

We can again find the change Δf in the $|y| \ll 1$ regime by plugging in an initial Bose-Einstein distribution to the right hand side of Eq. (A.38) to obtain $\Delta f = y x e^{x+\mu} f^2 g$, where

$$\begin{aligned} g &= X - 4 + \theta_e \left[-10 + \frac{47}{2}X - \frac{42}{5}X^2 + \frac{7}{10}X^3 \right. \\ &\quad \left. + S^2 \left(-\frac{21}{5} + \frac{7}{5}X \right) + \frac{7x^2}{10} (6 - X) \frac{T}{T_e} \right] \end{aligned} \quad (\text{A.40})$$

and

$$X = x \coth[(x + \mu)/2], \quad S = x \operatorname{csch}[(x + \mu)/2]. \quad (\text{A.41})$$

Recall that the Comptonization parameter y was defined in Eq. (4.5) to vanish for $T = T_e$.

For the tSZ effect in clusters where $T_e \gg T$, the expression becomes even simpler, reproducing and generalizing the $\mu = 0$ result found in Ref. [174]. This same rule for

generalizing g in the presence of μ through the modification to X and S in Eq. (A.41) applies to the higher order in θ_e terms of Ref. [174] for $T_e \gg T$ as we have explicitly checked to 4th order.

REFERENCES

- [1] Enrico Pajer and Matias Zaldarriaga. A New Window on Primordial non-Gaussianity. *Phys. Rev. Lett.*, 109:021302, 2012.
- [2] Srinivasan Raghunathan. Assessing the Importance of Noise from Thermal Sunyaev–Zel’dovich Signals for CMB Cluster Surveys and Cluster Cosmology. *Astrophys. J.*, 928(1):16, 2022.
- [3] Chen Heinrich, Olivier Dore, and Elisabeth Krause. Measuring f_{NL} with the SPHEREx Multi-tracer Redshift Space Bispectrum. 11 2023.
- [4] Olivier Doré et al. Cosmology with the SPHEREX All-Sky Spectral Survey. 12 2014.
- [5] Marilena LoVerde, Amber Miller, Sarah Shandera, and Licia Verde. Effects of Scale-Dependent Non-Gaussianity on Cosmological Structures. *JCAP*, 04:014, 2008.
- [6] Nashwan Sabti, Julian B. Muñoz, and Diego Blas. First Constraints on Small-Scale Non-Gaussianity from UV Galaxy Luminosity Functions. *JCAP*, 01:010, 2021.
- [7] Wayne Hu and Martin J. White. The Damping tail of CMB anisotropies. *Astrophys. J.*, 479:568, 1997.
- [8] N. Aghanim et al. Planck 2018 results. I. Overview and the cosmological legacy of Planck. *Astron. Astrophys.*, 641:A1, 2020.
- [9] Kevork Abazajian et al. CMB-S4 Science Case, Reference Design, and Project Plan. 7 2019.
- [10] D. J. Fixsen, E. S. Cheng, J. M. Gales, John C. Mather, R. A. Shafer, and E. L. Wright. The Cosmic Microwave Background spectrum from the full COBE FIRAS data set. *Astrophys. J.*, 473:576, 1996.
- [11] Alan H. Guth. The Inflationary Universe: A Possible Solution to the Horizon and Flatness Problems. *Phys.Rev.*, D23:347–356, 1981.
- [12] Simone Ferraro, Noah Sailer, Anze Slosar, and Martin White. Snowmass2021 Cosmic Frontier White Paper: Cosmology and Fundamental Physics from the three-dimensional Large Scale Structure. 3 2022.
- [13] Wayne Hu and Joseph Silk. Thermalization and spectral distortions of the cosmic background radiation. *Phys. Rev. D*, 48:485–502, 1993.
- [14] Jens Chluba. Refined approximations for the distortion visibility function and μ -type spectral distortions. *Mon. Not. Roy. Astron. Soc.*, 440(3):2544–2563, 2014.
- [15] Ya. B. Zeldovich and R. A. Sunyaev. The Interaction of Matter and Radiation in a Hot-Model Universe. *Astrophys. Space Sci.*, 4:301–316, 1969.

- [16] Jens Chluba. Which spectral distortions does Λ CDM actually predict? *Mon. Not. Roy. Astron. Soc.*, 460(1):227–239, 2016.
- [17] Jens Chluba, Emanuela Dimastrogiovanni, Mustafa A. Amin, and Marc Kamionkowski. Evolution of CMB spectral distortion anisotropies and tests of primordial non-Gaussianity. *Mon. Not. Roy. Astron. Soc.*, 466(2):2390–2401, 2017.
- [18] Giovanni Cabass, Enrico Pajer, and Drian van der Woude. Spectral distortion anisotropies from single-field inflation. *JCAP*, 08:050, 2018.
- [19] David Zegeye, Keisuke Inomata, and Wayne Hu. Spectral distortion anisotropy from inflation for primordial black holes. *Phys. Rev. D*, 105(10):103535, 2022.
- [20] Jens Chluba. Distinguishing different scenarios of early energy release with spectral distortions of the cosmic microwave background. *Mon. Not. Roy. Astron. Soc.*, 436:2232–2243, 2013.
- [21] Jens Chluba, Rishi Khatri, and Rashid A. Sunyaev. CMB at 2x2 order: The dissipation of primordial acoustic waves and the observable part of the associated energy release. *Mon. Not. Roy. Astron. Soc.*, 425:1129–1169, 2012.
- [22] Giovanni Cabass, Alessandro Melchiorri, and Enrico Pajer. μ distortions or running: A guaranteed discovery from CMB spectrometry. *Phys. Rev. D*, 93(8):083515, 2016.
- [23] Ogan Özsoy and Gianmassimo Tasinato. CMB μ T cross correlations as a probe of primordial black hole scenarios. *Phys. Rev. D*, 104(4):043526, 2021.
- [24] Aditya Rotti, Andrea Ravenni, and Jens Chluba. Non-Gaussianity constraints with anisotropic μ distortion measurements from Planck. *Mon. Not. Roy. Astron. Soc.*, 515(4):5847–5868, 2022.
- [25] Federico Bianchini and Giulio Fabbian. CMB spectral distortions revisited: A new take on μ distortions and primordial non-Gaussianities from FIRAS data. *Phys. Rev. D*, 106(6):063527, 2022.
- [26] David Zegeye, Thomas Crawford, and Wayne Hu. Measuring μ -Distortions from the Thermal Sunyaev-Zeldovich effect. 5 2023.
- [27] E. Allys et al. Probing Cosmic Inflation with the LiteBIRD Cosmic Microwave Background Polarization Survey. *PTEP*, 2023(4):042F01, 2023.
- [28] Mathieu Remazeilles, Andrea Ravenni, and Jens Chluba. Leverage on small-scale primordial non-Gaussianity through cross-correlations between CMB E -mode and μ -distortion anisotropies. 10 2021.
- [29] G. De Zotti, M. Negrello, G. Castex, A. Lapi, and M. Bonato. Another look at distortions of the Cosmic Microwave Background spectrum. *JCAP*, 03:047, 2016.

- [30] David Zegeye, Thomas Crawford, Jens Chluba, Mathieu Remazeilles, and Keith Grainge. Square Kilometer Array as a cosmic microwave background experiment. *Phys. Rev. D*, 111(6):063517, 2025.
- [31] Masahiro Kawasaki, Alexander Kusenko, Yuichiro Tada, and Tsutomu T. Yanagida. Primordial black holes as dark matter in supergravity inflation models. *Phys. Rev.*, D94(8):083523, 2016.
- [32] Juan Garcia-Bellido and Ester Ruiz Morales. Primordial black holes from single field models of inflation. *Phys. Dark Univ.*, 18:47–54, 2017.
- [33] Jose Maria Ezquiaga, Juan Garcia-Bellido, and Ester Ruiz Morales. Primordial Black Hole production in Critical Higgs Inflation. *Phys. Lett. B*, 776:345–349, 2018.
- [34] Kristjan Kannike, Luca Marzola, Martti Raidal, and Hardi Veermäe. Single Field Double Inflation and Primordial Black Holes. *JCAP*, 1709(09):020, 2017.
- [35] Cristiano Germani and Tomislav Prokopec. On primordial black holes from an inflection point. *Phys. Dark Univ.*, 18:6–10, 2017.
- [36] Guillermo Ballesteros and Marco Taoso. Primordial black hole dark matter from single field inflation. *Phys. Rev.*, D97(2):023501, 2018.
- [37] Keisuke Inomata, Masahiro Kawasaki, Kyohei Mukaida, and Tsutomu T. Yanagida. Double inflation as a single origin of primordial black holes for all dark matter and LIGO observations. *Phys. Rev. D*, 97(4):043514, 2018.
- [38] Mark P. Hertzberg and Masaki Yamada. Primordial Black Holes from Polynomial Potentials in Single Field Inflation. *Phys. Rev. D*, 97(8):083509, 2018.
- [39] Hayato Motohashi and Wayne Hu. Primordial Black Holes and Slow-Roll Violation. *Phys. Rev. D*, 96(6):063503, 2017.
- [40] Shi Pi, Ying-li Zhang, Qing-Guo Huang, and Misao Sasaki. Scalon from R^2 -gravity as a heavy field. *JCAP*, 05:042, 2018.
- [41] Shu-Lin Cheng, Wolung Lee, and Kin-Wang Ng. Superhorizon curvature perturbation in ultraslow-roll inflation. *Phys. Rev. D*, 99(6):063524, 2019.
- [42] Christian T. Byrnes, Philippa S. Cole, and Subodh P. Patil. Steepest growth of the power spectrum and primordial black holes. *JCAP*, 06:028, 2019.
- [43] Samuel Passaglia, Wayne Hu, and Hayato Motohashi. Primordial black holes and local non-Gaussianity in canonical inflation. *Phys. Rev. D*, 99(4):043536, 2019.
- [44] Manuel Drees and Yong Xu. Overshooting, Critical Higgs Inflation and Second Order Gravitational Wave Signatures. *Eur. Phys. J. C*, 81(2):182, 2021.
- [45] Pedro Carrilho, Karim A. Malik, and David J. Mulryne. Dissecting the growth of the power spectrum for primordial black holes. *Phys. Rev. D*, 100(10):103529, 2019.

- [46] Samuel Passaglia, Wayne Hu, and Hayato Motohashi. Primordial black holes as dark matter through Higgs field criticality. *Phys. Rev. D*, 101(12):123523, 2020.
- [47] Kin-Wang Ng and Yi-Peng Wu. Constant-rate inflation: primordial black holes from conformal weight transitions. 2 2021.
- [48] Keisuke Inomata, Evan McDonough, and Wayne Hu. Primordial Black Holes Arise When The Inflaton Falls. 4 2021.
- [49] Keisuke Inomata, Evan McDonough, and Wayne Hu. Amplification of Primordial Perturbations from the Rise or Fall of the Inflaton. 10 2021.
- [50] Simeon Bird, Ilias Cholis, Julian B. Muñoz, Yacine Ali-Haïmoud, Marc Kamionkowski, Ely D. Kovetz, Alvise Raccanelli, and Adam G. Riess. Did LIGO detect dark matter? *Phys. Rev. Lett.*, 116(20):201301, 2016.
- [51] Sebastien Clesse and Juan García-Bellido. The clustering of massive Primordial Black Holes as Dark Matter: measuring their mass distribution with Advanced LIGO. *Phys. Dark Univ.*, 15:142–147, 2017.
- [52] Misao Sasaki, Teruaki Suyama, Takahiro Tanaka, and Shuichiro Yokoyama. Primordial Black Hole Scenario for the Gravitational-Wave Event GW150914. *Phys. Rev. Lett.*, 117(6):061101, 2016.
- [53] A. Kashlinsky. LIGO gravitational wave detection, primordial black holes and the near-IR cosmic infrared background anisotropies. *Astrophys. J. Lett.*, 823(2):L25, 2016.
- [54] A. Kashlinsky, R. G. Arendt, F. Atrio-Barandela, N. Cappelluti, A. Ferrara, and G. Hasinger. Looking at cosmic near-infrared background radiation anisotropies. *Rev. Mod. Phys.*, 90(2):025006, 2018.
- [55] Juan García-Bellido, José Francisco Nuño Siles, and Ester Ruiz Morales. Bayesian analysis of the spin distribution of LIGO/Virgo black holes. *Phys. Dark Univ.*, 31:100791, 2021.
- [56] Amandeep S. Josan, Anne M. Green, and Karim A. Malik. Generalised constraints on the curvature perturbation from primordial black holes. *Phys. Rev.*, D79:103520, 2009.
- [57] Hooshyar Assadullahi and David Wands. Constraints on primordial density perturbations from induced gravitational waves. *Phys. Rev.*, D81:023527, 2010.
- [58] Keisuke Inomata and Tomohiro Nakama. Gravitational waves induced by scalar perturbations as probes of the small-scale primordial spectrum. *Phys. Rev.*, D99(4):043511, 2019.
- [59] Masahiro Kawasaki, Hiromasa Nakatsuka, and Kazunori Nakayama. Constraints on small-scale primordial density fluctuation from cosmic microwave background through dark matter annihilation. 10 2021.

- [60] R. A. Sunyaev and Ya. B. Zeldovich. Small-Scale Fluctuations of Relic Radiation. *Astrophys. and Space Sci.*, 7(1):3–19, April 1970.
- [61] R. A. Daly. Spectral distortions of the microwave background radiation resulting from the damping of pressure waves. *Astrophys. J.*, 371, 1991.
- [62] Wayne Hu, Douglas Scott, and Joseph Silk. Power spectrum constraints from spectral distortions in the cosmic microwave background. *Astrophys. J. Lett.*, 430:L5–L8, 1994.
- [63] J. Chluba and R. A. Sunyaev. The evolution of CMB spectral distortions in the early Universe. *Mon. Not. Roy. Astron. Soc.*, 419:1294–1314, 2012.
- [64] Tomohiro Nakama, Bernard Carr, and Joseph Silk. Limits on primordial black holes from μ distortions in cosmic microwave background. 2017.
- [65] Bernard Carr, Kazunori Kohri, Yuuiti Sendouda, and Jun’ichi Yokoyama. Constraints on Primordial Black Holes. 2020.
- [66] J. Chluba et al. New horizons in cosmology with spectral distortions of the cosmic microwave background. *Exper. Astron.*, 51(3):1515–1554, 2021.
- [67] M. Hazumi et al. LiteBIRD: JAXA’s new strategic L-class mission for all-sky surveys of cosmic microwave background polarization. *Proc. SPIE Int. Soc. Opt. Eng.*, 11443:114432F, 2020.
- [68] Jonathan Ganc and Eiichiro Komatsu. Scale-dependent bias of galaxies and μ -type distortion of the cosmic microwave background spectrum from single-field inflation with a modified initial state. *Phys. Rev. D*, 86:023518, 2012.
- [69] Enrico Pajer and Matias Zaldarriaga. A hydrodynamical approach to CMB μ -distortion from primordial perturbations. *JCAP*, 02:036, 2013.
- [70] M. Remazeilles and J. Chluba. Extracting foreground-obscured μ -distortion anisotropies to constrain primordial non-Gaussianity. *Mon. Not. Roy. Astron. Soc.*, 478(1):807–824, 2018.
- [71] Juan Martin Maldacena. Non-Gaussian features of primordial fluctuations in single field inflationary models. *JHEP*, 05:013, 2003.
- [72] Enrico Pajer, Fabian Schmidt, and Matias Zaldarriaga. The Observed Squeezed Limit of Cosmological Three-Point Functions. *Phys. Rev. D*, 88(8):083502, 2013.
- [73] Ogan Özsoy and Gianmassimo Tasinato. Consistency conditions and primordial black holes in single field inflation. 11 2021.
- [74] Mohammad Hossein Namjoo, Hassan Firouzjahi, and Misao Sasaki. Violation of non-Gaussianity consistency relation in a single field inflationary model. *EPL*, 101(3):39001, 2013.

- [75] Jerome Martin, Hayato Motohashi, and Teruaki Suyama. Ultra Slow-Roll Inflation and the non-Gaussianity Consistency Relation. *Phys. Rev. D*, 87(2):023514, 2013.
- [76] Yi-Fu Cai, Xingang Chen, Mohammad Hossein Namjoo, Misao Sasaki, Dong-Gang Wang, and Ziwei Wang. Revisiting non-Gaussianity from non-attractor inflation models. *JCAP*, 05:012, 2018.
- [77] William H. Kinney. Horizon crossing and inflation with large eta. *Phys. Rev. D*, 72:023515, 2005.
- [78] N. Aghanim et al. Planck 2018 results. VI. Cosmological parameters. *Astron. Astrophys.*, 641:A6, 2020.
- [79] Wayne Hu, Chi-Ting Chiang, Yin Li, and Marilena LoVerde. Separating the Universe into real and fake energy densities. *Phys. Rev. D*, 94(2):023002, 2016.
- [80] Wayne Hu and Austin Joyce. Separate Universes beyond General Relativity. *Phys. Rev. D*, 95(4):043529, 2017.
- [81] Edwin Sirko. Initial conditions to cosmological N-body simulations, or how to run an ensemble of simulations. *Astrophys. J.*, 634:728–743, 2005.
- [82] Yin Li, Wayne Hu, and Masahiro Takada. Super-Sample Covariance in Simulations. *Phys. Rev. D*, 89(8):083519, 2014.
- [83] Wayne Hu and Martin J. White. Acoustic signatures in the cosmic microwave background. *Astrophys. J.*, 471:30–51, 1996.
- [84] Wayne Hu and Naoshi Sugiyama. Small scale cosmological perturbations: An Analytic approach. *Astrophys. J.*, 471:542–570, 1996.
- [85] L. Danese and G. de Zotti. Double Compton process and the spectrum of the microwave background. *Astronomy and Astrophysics*, 107(1):39–42, March 1982.
- [86] A. Kogut et al. The Primordial Inflation Explorer (PIXIE): A Nulling Polarimeter for Cosmic Microwave Background Observations. *JCAP*, 07:025, 2011.
- [87] A. Kogut and D. J. Fixsen. Calibration Method and Uncertainty for the Primordial Inflation Explorer (PIXIE). *JCAP*, 05:041, 2020.
- [88] Kevork N. Abazajian et al. CMB-S4 Science Book, First Edition. 10 2016.
- [89] Andrea Ravenni, Michele Liguori, Nicola Bartolo, and Maresuke Shiraishi. Primordial non-Gaussianity with μ -type and γ -type spectral distortions: exploiting Cosmic Microwave Background polarization and dealing with secondary sources. *JCAP*, 09:042, 2017.
- [90] Giorgio Orlando, P. Daniel Meerburg, and Subodh P. Patil. Primordial tensor bispectra in μ -CMB cross-correlations. 9 2021.

- [91] David Zegeye et al. CMB-S4: Forecasting Constraints on f_{NL} Through μ -distortion Anisotropy. *Phys. Rev. D*, 108(10):103536, 2023.
- [92] R. L. Workman et al. Review of Particle Physics. *PTEP*, 2022:083C01, 2022.
- [93] P. A. R. Ade et al. Improved Constraints on Primordial Gravitational Waves using Planck, WMAP, and BICEP/Keck Observations through the 2018 Observing Season. *Phys. Rev. Lett.*, 127(15):151301, 2021.
- [94] Eva Silverstein and David Tong. Scalar speed limits and cosmology: Acceleration from D-cceleration. *Phys. Rev. D*, 70:103505, 2004.
- [95] Nima Arkani-Hamed, Paolo Creminelli, Shinji Mukohyama, and Matias Zaldarriaga. Ghost inflation. *JCAP*, 04:001, 2004.
- [96] Mohsen Alishahiha, Eva Silverstein, and David Tong. DBI in the sky. *Phys. Rev. D*, 70:123505, 2004.
- [97] Xingang Chen, Min-xin Huang, Shamit Kachru, and Gary Shiu. Observational signatures and non-Gaussianities of general single field inflation. *JCAP*, 01:002, 2007.
- [98] Clifford Cheung, Paolo Creminelli, A. Liam Fitzpatrick, Jared Kaplan, and Leonardo Senatore. The Effective Field Theory of Inflation. *JHEP*, 03:014, 2008.
- [99] Leonardo Senatore, Kendrick M. Smith, and Matias Zaldarriaga. Non-Gaussianities in Single Field Inflation and their Optimal Limits from the WMAP 5-year Data. *JCAP*, 01:028, 2010.
- [100] Xingang Chen and Yi Wang. Quasi-Single Field Inflation and Non-Gaussianities. *JCAP*, 04:027, 2010.
- [101] Andrew J. Tolley and Mark Wyman. The Gelaton Scenario: Equilateral non-Gaussianity from multi-field dynamics. *Phys. Rev. D*, 81:043502, 2010.
- [102] Sera Cremonini, Zygmunt Lalak, and Krzysztof Turzyski. Strongly Coupled Perturbations in Two-Field Inflationary Models. *JCAP*, 03:016, 2011.
- [103] Ana Achucarro, Jinn-Ouk Gong, Sjoerd Hardeman, Gonzalo A. Palma, and Subodh P. Patil. Features of heavy physics in the CMB power spectrum. *JCAP*, 01:030, 2011.
- [104] Daniel Baumann and Daniel Green. Signatures of Supersymmetry from the Early Universe. *Phys. Rev. D*, 85:103520, 2012.
- [105] Neil Barnaby and Sarah Shandera. Feeding your Inflaton: Non-Gaussian Signatures of Interaction Structure. *JCAP*, 01:034, 2012.
- [106] Nima Arkani-Hamed and Juan Maldacena. Cosmological Collider Physics. 3 2015.
- [107] Kendrick M. Smith, Marilena LoVerde, and Matias Zaldarriaga. A universal bound on N-point correlations from inflation. *Phys. Rev. Lett.*, 107:191301, 2011.

- [108] Paolo Creminelli and Matias Zaldarriaga. Single-Field Consistency Relation for the 3-Point Function. *JCAP*, 0410:006, 2004.
- [109] Y. Akrami et al. Planck 2018 results. IX. Constraints on primordial non-Gaussianity. *Astron. Astrophys.*, 641:A9, 2020.
- [110] Alba Kalaja, P. Daniel Meerburg, Guilherme L. Pimentel, and William R. Coulton. Fundamental limits on constraining primordial non-Gaussianity. *JCAP*, 04:050, 2021.
- [111] Neal Dalal, Olivier Dore, Dragan Huterer, and Alexander Shirokov. The imprints of primordial non-gaussianities on large-scale structure: scale dependent bias and abundance of virialized objects. *Phys. Rev. D*, 77:123514, 2008.
- [112] James Annis, Jeffrey A. Newman, and Anže Slosar. Snowmass2021 Cosmic Frontier: Report of the CF04 Topical Group on Dark Energy and Cosmic Acceleration in the Modern Universe. 9 2022.
- [113] Olivier Doré et al. Cosmology with the SPHEREX All-Sky Spectral Survey. 12 2014.
- [114] R. A. Sunyaev and Ya. B. Zeldovich. The Interaction of matter and radiation in the hot model of the universe. *Astrophys. Space Sci.*, 7:20–30, 1970.
- [115] Carlo Burigana, L. Danese, and G. De Zotti. Formation and evolution of early distortions of the microwave background spectrum - A numerical study. *Astron. Astrophys.*, 246(1):49–58, 1991.
- [116] R. A. Sunyaev and Ya. B. Zeldovich. Small scale entropy and adiabatic density perturbations ? Antimatter in the Universe. *Astrophys. Space Sci.*, 9(3):368–382, 1970.
- [117] Joseph Silk. Cosmic black body radiation and galaxy formation. *Astrophys. J.*, 151:459–471, 1968.
- [118] Rishi Khatri, Rashid A. Sunyaev, and Jens Chluba. Mixing of blackbodies: entropy production and dissipation of sound waves in the early Universe. *Astron. Astrophys.*, 543:A136, 2012.
- [119] Jens Chluba, Adrienne L. Erickcek, and Ido Ben-Dayan. Probing the inflaton: Small-scale power spectrum constraints from measurements of the CMB energy spectrum. *Astrophys. J.*, 758:76, 2012.
- [120] Rishi Khatri and Rashid Sunyaev. Constraints on μ -distortion fluctuations and primordial non-Gaussianity from Planck data. *JCAP*, 09:026, 2015.
- [121] N. Aghanim et al. Planck 2018 results. VI. Cosmological parameters. *Astron. Astrophys.*, 641:A6, 2020. [Erratum: *Astron. Astrophys.* 652, C4 (2021)].
- [122] Jens Chluba, Thomas Kite, and Andrea Ravenni. Spectro-spatial evolution of the CMB. Part I. Discretisation of the thermalisation Green’s function. *JCAP*, 11:026, 2023.

- [123] Jens Chluba, Andrea Ravenni, and Thomas Kite. Spectro-spatial evolution of the CMB. Part II. Generalised Boltzmann hierarchy. *JCAP*, 11:027, 2023.
- [124] Thomas Kite, Andrea Ravenni, and Jens Chluba. Spectro-spatial evolution of the CMB. Part III. Transfer functions, power spectra and Fisher forecasts. *JCAP*, 11:028, 2023.
- [125] Antony Lewis, Anthony Challinor, and Anthony Lasenby. Efficient computation of CMB anisotropies in closed FRW models. *Astrophys. J.*, 538:473–476, 2000.
- [126] Mathieu Remazeilles, Jacques Delabrouille, and Jean-Francois Cardoso. CMB and SZ effect separation with Constrained Internal Linear Combinations. *Mon. Not. Roy. Astron. Soc.*, 410:2481, 2011.
- [127] Oliver P. Lay and Nils W. Halverson. The impact of atmospheric fluctuations on degree-scale imaging of the cosmic microwave background. *Astrophys. J.*, 543:787–798, 2000.
- [128] Robert S. Bussmann, W. L. Holzapfel, and C. L. Kuo. Millimeter wavelength brightness fluctuations of the atmosphere above the South Pole. *Astrophys. J.*, 622:1343–1355, 2005.
- [129] Thomas W. Morris et al. The Atacama Cosmology Telescope: Modeling bulk atmospheric motion. *Phys. Rev. D*, 105(4):042004, 2022.
- [130] Kevork Abazajian et al. CMB-S4: Forecasting Constraints on Primordial Gravitational Waves. 8 2020.
- [131] J. E. Carlstrom et al. The 10 Meter South Pole Telescope. *Publ. Astron. Soc. Pac.*, 123:568–581, 2011.
- [132] J. A. Sobrin et al. The Design and Integrated Performance of SPT-3G. *Astrophys. J. Supp.*, 258(2):42, 2022.
- [133] Karia R. Dibert et al. Forecasting ground-based sensitivity to the Rayleigh scattering of the CMB in the presence of astrophysical foregrounds. *Phys. Rev. D*, 106(6):063502, 2022.
- [134] Ben Thorne, Jo Dunkley, David Alonso, and Sigurd Naess. The Python Sky Model: software for simulating the Galactic microwave sky. *Mon. Not. Roy. Astron. Soc.*, 469(3):2821–2833, 2017.
- [135] J. Delabrouille et al. The pre-launch Planck Sky Model: a model of sky emission at submillimetre to centimetre wavelengths. *Astron. Astrophys.*, 553:A96, 2013.
- [136] K. T. Story et al. A Measurement of the Cosmic Microwave Background Damping Tail from the 2500-square-degree SPT-SZ survey. *Astrophys. J.*, 779:86, 2013.

- [137] E. M. George et al. A measurement of secondary cosmic microwave background anisotropies from the 2500-square-degree SPT-SZ survey. *Astrophys. J.*, 799(2):177, 2015.
- [138] John C. Mather et al. Measurement of the Cosmic Microwave Background spectrum by the COBE FIRAS instrument. *Astrophys. J.*, 420:439–444, 1994.
- [139] Alan Kogut et al. ARCADE: Absolute Radiometer for Cosmology, Astrophysics, and Diffuse Emission. *New Astron. Rev.*, 50:925–931, 2006.
- [140] Alan J. Kogut et al. The Primordial Inflation Explorer (PIXIE) Mission. *Proc. SPIE Int. Soc. Opt. Eng.*, 7731:77311S, 2010.
- [141] Chao-Lin Kuo. Assessments of Ali, Dome A, and Summit Camp for Mm-wave Observations Using MERRA-2 Reanalysis. *Astrophys. J.*, 848(1):64, 2017.
- [142] S. W. Henderson et al. Advanced ACTPol Cryogenic Detector Arrays and Readout. *J. Low Temp. Phys.*, 184(3-4):772–779, 2016.
- [143] Peter Ade et al. The Simons Observatory: Science goals and forecasts. *JCAP*, 02:056, 2019.
- [144] Emanuela Dimastrogiovanni and Razieh Emami. Correlating CMB Spectral Distortions with Temperature: what do we learn on Inflation? *JCAP*, 12:015, 2016.
- [145] Razieh Emami. Probing the Running of Primordial Bispectrum and Trispectrum using CMB Spectral Distortions. *Phys. Rev. D*, 100(8):083021, 2019.
- [146] M. Biagetti, G. Franciolini, and A. Riotto. The JWST High Redshift Observations and Primordial Non-Gaussianity. 10 2022.
- [147] Thomas Kite, Andrea Ravenni, Subodh P. Patil, and Jens Chluba. Bridging the gap: spectral distortions meet gravitational waves. *Mon. Not. Roy. Astron. Soc.*, 505(3):4396, 2021.
- [148] A. Weltman et al. Fundamental physics with the Square Kilometre Array. *Publ. Astron. Soc. Austral.*, 37:e002, 2020.
- [149] C. L. Bennett et al. Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. *Astrophys. J. Suppl.*, 208:20, 2013.
- [150] M. Gervasi, M. Zannoni, A. Tartari, G. Boella, and G. Sironi. TRIS II: search for CMB spectral distortions at 0.60, 0.82 and 2.5 GHz. *Astrophys. J.*, 688:24, 2008.
- [151] Boris Bolliet, Jens Chluba, and Richard Battye. Spectral distortion constraints on photon injection from low-mass decaying particles. *Mon. Not. Roy. Astron. Soc.*, 507(3):3148–3178, 2021.

- [152] Sandeep Kumar Acharya and Rishi Khatri. CMB spectral distortions constraints on primordial black holes, cosmic strings and long lived unstable particles revisited. *JCAP*, 2002(02):010, 2020.
- [153] Atsuhisa Ota, Tomo Takahashi, Hiroyuki Tashiro, and Masahide Yamaguchi. CMB μ distortion from primordial gravitational waves. *JCAP*, 10:029, 2014.
- [154] Suvodip Mukherjee, Joseph Silk, and Benjamin D. Wandelt. FSD: Frequency Space Differential measurement of CMB spectral distortions. *Mon. Not. Roy. Astron. Soc.*, 477(4):4473–4482, 2018.
- [155] Suvodip Mukherjee, Joseph Silk, and Benjamin D. Wandelt. How to measure CMB spectral distortions with an imaging telescope. *Phys. Rev. D*, 100(10):103508, 2019.
- [156] S. A. Balashev, E. E. Kholupenko, J. Chluba, A. V. Ivanchik, and D. A. Varshalovich. Spectral distortions of the CMB dipole. *Astrophys. J.*, 810(2):131, 2015.
- [157] Y. Rephaeli. On the determination of the degree of cosmological Compton distortions and the temperature of the cosmic blackbody radiation. *Astrophys. J.*, 241:858–863, November 1980.
- [158] Gemma Luzzi, Emanuele D’Angelo, Hervé Bourdin, Federico De Luca, Pasquale Mazzotta, Filippo Oppizzi, and Gianluca Polenta. Cosmology with the SZ spectrum: Measuring the Universe’s temperature with galaxy clusters. *EPJ Web Conf.*, 257:00028, 2022.
- [159] E. L. Wright. Was the Big Bang hot? In G. O. Abell and G. Chincarini, editors, *Early Evolution of the Universe and its Present Structure*, volume 104, pages 113–116, January 1983.
- [160] E. E. Kholupenko, S. A. Balashev, A. V. Ivanchik, and D. A. Varshalovich. The thermal Sunyaev–Zeldovich effect of primordial recombination radiation. *Mon. Not. Roy. Astron. Soc.*, 446(4):3593–3607, 2015.
- [161] Simone Aiola et al. Snowmass2021 CMB-HD White Paper. 3 2022.
- [162] A. S. Kompaneets. The Establishment of Thermal Equilibrium between Quanta and Electrons. *Soviet Journal of Experimental and Theoretical Physics*, 4(5):730–737, May 1957.
- [163] L. E. Bleem et al. Galaxy Clusters Discovered via the Sunyaev-Zel’dovich Effect in the 2500-square-degree SPT-SZ survey. *Astrophys. J. Suppl.*, 216(2):27, 2015.
- [164] T. Plagge et al. Sunyaev-Zel’dovich Cluster Profiles Measured with the South Pole Telescope. *Astrophys. J.*, 716:1118–1135, 2010.
- [165] J. Liu et al. Analysis of Sunyaev–Zel’dovich effect mass–observable relations using South Pole Telescope observations of an X-ray selected sample of low-mass galaxy clusters and groups. *Mon. Not. Roy. Astron. Soc.*, 448(3):2085–2099, 2015.

- [166] M. Arnaud, G. W. Pratt, R. Piffaretti, H. Boehringer, J. H. Croston, and E. Pointecouteau. The universal galaxy cluster pressure profile from a representative sample of nearby systems (REXCESS) and the Y_{SZ}-M₅₀₀ relation. *Astron. Astrophys.*, 517:A92, 2010.
- [167] Matthew Hasselfield et al. The Atacama Cosmology Telescope: Sunyaev-Zel’dovich selected galaxy clusters at 148 GHz from three seasons of data. *JCAP*, 07:008, 2013.
- [168] Daisuke Nagai, Andrey V. Kravtsov, and Alexey Vikhlinin. Effects of Galaxy Formation on Thermodynamics of the Intracluster Medium. *Astrophys. J.*, 668:1–14, 2007.
- [169] Monique Arnaud, E. Pointecouteau, and G. W. Pratt. The Structural and scaling properties of nearby galaxy clusters. 2. The M-T relation. *Astron. Astrophys.*, 441:893–903, 2005.
- [170] Srinivasan Raghunathan, Nathan Whitehorn, Marcelo A. Alvarez, Han Aung, Nicholas Battaglia, Gilbert P. Holder, Daisuke Nagai, Elena Pierpaoli, Christian L. Reichardt, and Joaquin D. Vieira. Constraining Cluster Virialization Mechanism and Cosmology Using Thermal-SZ-selected Clusters from Future CMB Surveys. *Astrophys. J.*, 926(2):172, 2022.
- [171] Jeremy L. Tinker, Andrey V. Kravtsov, Anatoly Klypin, Kevork Abazajian, Michael S. Warren, Gustavo Yepes, Stefan Gottlober, and Daniel E. Holz. Toward a halo mass function for precision cosmology: The Limits of universality. *Astrophys. J.*, 688:709–728, 2008.
- [172] Benedikt Diemer. COLOSSUS: A python toolkit for cosmology, large-scale structure, and dark matter halos. *Astrophys. J. Suppl.*, 239(2):35, 2018.
- [173] A. Merloni et al. eROSITA Science Book: Mapping the Structure of the Energetic Universe. 9 2012.
- [174] Naoki Itoh, Yasuharu Kohyama, and Satoshi Nozawa. Relativistic Corrections to the Sunyaev-Zel’dovich Effect for Clusters of Galaxies. *Astrophys. J.*, 502:7–15, 1998.
- [175] J. R. Bond, Andrew H. Jaffe, and L. Knox. Estimating the power spectrum of the cosmic microwave background. *Phys. Rev. D*, 57:2117–2137, 1998.
- [176] C. L. Reichardt et al. An Improved Measurement of the Secondary Cosmic Microwave Background Anisotropies from the SPT-SZ + SPTpol Surveys. *Astrophys. J.*, 908(2):199, 2021.
- [177] Gianfranco De Zotti, Roberto Ricci, Dino Mesa, Laura Silva, Pasquale Mazzotta, Luigi Toffolatti, and Joaquin Gonzalez-Nuevo. Predictions for high-frequency radio surveys of extragalactic sources. *Astron. Astrophys.*, 431:893, 2005.
- [178] Clive Dickinson et al. The State-of-Play of Anomalous Microwave Emission (AME) Research. *New Astron. Rev.*, 80:1–28, 2018.

- [179] W. L. Holzapfel, T. M. Wilbanks, P. A. R. Ade, S. E. Church, M. L. Fischer, P. D. Mauskopf, D. E. Osgood, and Andrew E. Lange. The Sunyaev-Zel'dovich infrared experiment: A Millimeter wave receiver for cluster cosmology. *Astrophys. J.*, 479:17, 1997.
- [180] Patrick M. Motl, Eric J. Hallman, Jack O. Burns, and Michael L. Norman. The Integrated Sunyaev-Zeldovich effect as the superior method for measuring the mass of clusters of galaxies. *Astrophys. J. Lett.*, 623:L63–L66, 2005.
- [181] Rocco Piffaretti, Philippe Jetzer, and Sabine Schindler. Aspherical galaxy clusters: Effects on cluster masses and gas mass fractions. *Astron. Astrophys.*, 398:41–48, 2003.
- [182] Christopher Sheehy and Anže Slosar. No evidence for dust B-mode decorrelation in Planck data. *Phys. Rev. D*, 97(4):043522, 2018.
- [183] Robert Braun, Tyler Bourke, James A Green, Evan Keane, and Jeff Wagg. Advancing Astrophysics with the Square Kilometre Array. *PoS, AASKA14:174*, 2015.
- [184] D. Benford, H. Moseley, and J. Zmuidzinas. Direct detectors for the Einstein inflation probe. *J. Phys. Conf. Ser.*, 155:012001, 2009.
- [185] Maximilian H. Abitbol, Jens Chluba, J. Colin Hill, and Bradley R. Johnson. Prospects for Measuring Cosmic Microwave Background Spectral Distortions in the Presence of Foregrounds. *Mon. Not. Roy. Astron. Soc.*, 471(1):1126–1140, 2017.
- [186] N. Krachmalnicoff, C. Baccigalupi, J. Aumont, M. Bersanelli, and A. Mennella. Characterization of foreground emission on degree angular scales for CMB B-mode observations - Thermal dust and synchrotron signal from Planck and WMAP data. *Astron. Astrophys.*, 588:A65, 2016.
- [187] Aritra Basu, Dominik J. Schwarz, Hans-Rainer Klöckner, Sebastian von Hausegger, Michael Kramer, Gundolf Wieching, and Blakesley Burkhart. CMB foreground measurements through broad-band radio spectro-polarimetry: prospects of the SKA-MPG telescope. *Mon. Not. Roy. Astron. Soc.*, 488(2):1618–1634, 2019.
- [188] Ravi Subrahmanyam and R. D. Ekers. Cmb observations using the ska. 9 2002.
- [189] R. A. Sunyaev and Ya. B. Zeldovich. Small scale fluctuations of relic radiation. *Astrophys. Space Sci.*, 7:3–19, 1970.
- [190] Razieh Emami, Emanuela Dimastrogiovanni, Jens Chluba, and Marc Kamionkowski. Probing the scale dependence of non-Gaussianity with spectral distortions of the cosmic microwave background. , 91(12):123531, June 2015.
- [191] Rishi Khatri and Rashid Sunyaev. Constraints on μ -distortion fluctuations and primordial non-Gaussianity from Planck data. *JCAP*, 09:026, 2015.
- [192] Ciaran McCulloch, Enrico Pajer, and Xi Tong. A cosmological tachyon collider: enhancing the long-short scale coupling. *JHEP*, 05:262, 2024.

- [193] D. J. Fixsen et al. ARCADE 2 Measurement of the Extra-Galactic Sky Temperature at 3-90 GHz. *Astrophys. J.*, 734:5, 2011.
- [194] J. Singal et al. The Radio Synchrotron Background: Conference Summary and Report. *Publ. Astron. Soc. Pac.*, 130(985):036001, 2018.
- [195] J. Singal et al. The Second Radio Synchrotron Background Workshop: Conference Summary and Report. *Publ. Astron. Soc. Pac.*, 135(1045):036001, 2023.
- [196] M. Remazeilles, C. Dickinson, H. K. K. Eriksen, and I. K. Wehus. Sensitivity and foreground modelling for large-scale cosmic microwave background B-mode polarization satellite missions. *Mon. Not. Roy. Astron. Soc.*, 458(2):2032–2050, 2016.
- [197] Mathieu Remazeilles, Aditya Rotti, and Jens Chluba. Peeling off foregrounds with the constrained moment ILC method to unveil primordial CMB B -modes. *Mon. Not. Roy. Astron. Soc.*, 503(2):2478–2498, 2021.
- [198] Jens Chluba, J. Colin Hill, and Maximilian H. Abitbol. Rethinking CMB foregrounds: systematic extension of foreground parametrizations. *Mon. Not. Roy. Astron. Soc.*, 472(1):1195–1213, 2017.
- [199] S. Azzoni, M. H. Abitbol, D. Alonso, A. Gough, N. Katayama, and T. Matsumura. A minimal power-spectrum-based moment expansion for CMB B-mode searches. *JCAP*, 05:047, 2021.
- [200] David J. Bacon et al. Cosmology with Phase 1 of the Square Kilometre Array: Red Book 2018: Technical specifications and performance forecasts. *Publ. Astron. Soc. Austral.*, 37:e007, 2020.
- [201] Aditya Rotti, Andrea Ravenni, and Jens Chluba. Non-Gaussianity constraints with anisotropic μ distortion measurements from Planck. *Mon. Not. Roy. Astron. Soc.*, 515(4):5847–5868, 2022.
- [202] Stuart Harper, Clive Dickinson, Richard Battye, Sambit Roychowdhury, Ian Browne, Yin-Zhe Ma, Lucas Olivari, and Tianyue Chen. Impact of Simulated $1/f$ Noise for HI Intensity Mapping Experiments. *Mon. Not. Roy. Astron. Soc.*, 478(2):2416–2437, 2018.
- [203] Yichao Li, Mario G. Santos, Keith Grainge, Stuart Harper, and Jingying Wang. HI intensity mapping with MeerKAT: $1/f$ noise analysis. *Mon. Not. Roy. Astron. Soc.*, 501(3):4344–4358, 2021.
- [204] Toshiya Namikawa, Daisuke Yamauchi, Blake Sherwin, and Ryo Nagata. Delensing Cosmic Microwave Background B-modes with the Square Kilometre Array Radio Continuum Survey. *Phys. Rev. D*, 93(4):043527, 2016.
- [205] Amanda MacInnis and Neelima Sehgal. CMB-HD as a Probe of Dark Matter on Sub-Galactic Scales. 5 2024.

- [206] Keisuke Inomata. Analytic solutions of scalar perturbations induced by scalar perturbations. *JCAP*, 03:013, 2021.
- [207] Bart Horn, Lam Hui, and Xiao Xiao. Soft-Pion Theorems for Large Scale Structure. *JCAP*, 09:044, 2014.
- [208] Wayne Hu. Weak lensing of the CMB: A harmonic approach. *Phys. Rev. D*, 62:043007, 2000.
- [209] I. S. Gradshteyn and I. M. Ryzhik. *Table of integrals, series and products*. 1994.
- [210] A. D. Challinor and A. N. Lasenby. Relativistic corrections to the sunyaev-zel'dovich effect. *Astrophys. J.*, 499:1, 1998.