

Challenges in probing inflation with primordial gravitational waves



Ben Thorne

St John's College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2019

For my family,

Declaration of Authorship

I declare that no part of this thesis has been accepted, or is currently being submitted, for any degree or diploma or certificate or any other qualification in this University or elsewhere. Except where explicit reference is made to the work of others, the work contained in this thesis is my own, and is not the outcome of work done by others.

Chapter 2 is based on work published in Ref. [1]. This work is wholly my own, except for the likelihood analysis presented in Section 2.5, Figures 2.12 and 2.13, and minor modifications reflecting comments and edits from collaborators and referees.

Chapter 3 is based on work published in Ref. [2]. This work is wholly my own, except for Figures A.1a, A.1b, and A.2, and minor modifications reflecting comments and edits from collaborators and referees.

Chapter 5 is based on work that has been submitted to Physical Review D for publication, and is currently under review. This work is an independent extension of work done by the Simons Observatory collaboration, which was published in Ref [3]. The work in Chapter 5 is wholly my own, excepting some modifications reflecting comments and edits from collaborators within the Simons Observatory.

Acknowledgements

Above all, I am privileged to have been supervised by Jo Dunkley. I can not imagine a better role model as I begin my scientific career.

During my DPhil I have been fortunate to travel the world, and I would like to start by thanking the staff in the Denys Wilkinson Building, the Kavli IPMU, and Peyton Hall, for making this possible. Above all, I would like to thank Ashling Morris and Polly Strauss.

From an early age, I knew what I wanted to do, but it was made possible by educators like Jon Brown, Catherine Carroll, and Joe Conlon, all of whom have contributed to any scientific success I might obtain. I would also like to thank Adriana de Lorenzo-Cáceres for her supervision of my first research project, and my collaborator, David Alonso, for his patience and insights.

Thank you to New College, and my friends: Aneesh, Matt, Jacob, Sam and Adam. Sam, five years was too short a time to spend with such an admirable friend, good luck on your antipodean jaunt. Thanks to Catherine, Naomi, and all the other members of my cohort in Oxford who made grappling with the first year of graduate school so entertaining.

I would like to thank everyone at the IPMU for their hospitality during my stay in Japan, especially: Nobuhiko Katayama, Tomotake Matsumura, and Maresuke Shiraishi; as well as my friends Robin, Li, Tom, Hiroko, and Jennifer, with whom it was a pleasure to explore Japan.

Of all the places I could have landed to finish my DPhil, I am glad that it was Princeton. I am grateful for every member of the graduate student community in Peyton Hall, as well as my erstwhile housemates: Deepen and Celia.

As a frequent visitor, and now transplant, to California, I would like to thank Julian Borrill, Reijo Keskitalo, Ted Kisner, and Andrea Zonca for always welcoming me at Berkeley Lab, and having a desk available.

After meeting on the first day of graduate school, Zeenia and I have found our paths intertwined over several years, and as many continents. Without her, this thesis would still have been written, but the grammar would have been even worse. Thank you to Bomi and Sonya for your support during my stays in Menlo Park, where this thesis was constructed.

In all my endeavours, I have enjoyed the unwavering support of my family. Anything I achieve would not have been possible without the opportunities afforded to me by them, and my grandparents: Mavis & Richard, and Mollie & Roy. I have dedicated this thesis to my parents, my grandparents, and to my sister, Charlie.

Abstract

It is still not yet understood how structure in the Universe was seeded. The theory of inflation provides a framework to understand the origins of structure as quantum fluctuations of the vacuum, whilst maintaining an attractive simplicity, and giving a natural explanation of how the Universe is so flat, homogeneous, and isotropic. If inflation occurred, it would have produced a gravitational wave background, indirectly observable today through its effect on the polarization of the cosmic microwave background.

We present a new model for the polarized microwave sky, combining data from the *Planck* and *WMAP* satellites, and legacy radio surveys, with theoretical models for Galactic foreground SEDs. We present an example use-case of this model by forecasting the effect of Galactic foregrounds on a future satellite mission.

We then move on to consider an extension to the simplest inflationary scenario in which an additional set of $SU(2)$ gauge fields are present during inflation, and may source a chiral gravitational wave background. We show that future CMB missions will not be able to distinguish such a scenario from the simplest single field slow-roll model using standard two-point statistical techniques, and that space-based laser interferometers will provide a complementary probe by constraining the short wavelength behavior of the blue-tilted gravitational wave spectrum.

Finally, we demonstrate the efficacy of a new method to account for complex Galactic foregrounds in the analysis of high-sensitivity polarized CMB data from the upcoming Simons Observatory. We first show that ignoring the spatial variation could lead to a false detection of a gravitational wave background, and that this can be corrected for by introducing additional parameters describing the spatial variation of the foreground spectral energy dependence.

Contents

1	Introduction	1
1.1	Concordance Cosmology	1
1.1.1	FLRW metric	2
1.1.2	Energy content of the Universe	4
1.1.3	Perturbations to the FLRW metric	5
1.1.3.1	Scalar Perturbations	8
1.1.3.2	Tensor Perturbations	8
1.1.4	Problems with Big Bang Cosmology	9
1.2	Inflation	10
1.2.1	Single Field Slow Roll Inflation	11
1.2.1.1	Background evolution	11
1.2.1.2	Perturbations	12
1.2.2	Observational status of inflation	14
1.2.3	Alternative Models of Inflation	15
1.3	Cosmic Microwave Background	16
1.3.1	The anisotropic CMB	16
1.3.2	CMB anisotropies from scalar perturbations	17
1.3.3	Polarization	19
1.3.3.1	From scalar perturbations	21
1.3.3.2	From tensor perturbations	22
1.3.3.3	From gravitational lensing	22
1.3.4	CMB observations	24
1.3.4.1	Past and current observations	24
1.3.4.2	Future observations	27
1.4	CMB Foregrounds	27
1.4.1	Thermal dust emission	29
1.4.2	Synchrotron emission	30

1.4.3	Anomalous microwave emission	30
1.4.4	Free-free	31
1.5	Statistical Methods	31
1.5.1	Bayesian statistics	31
1.5.1.1	Parameter inference	31
1.5.1.2	Markov Chain Monte Carlo	33
1.5.2	Fisher forecasts	34
1.5.3	Power spectrum estimation	36
1.5.3.1	Quadratic estimator and minimum-variance	36
1.5.3.2	Pseudo-spectrum techniques	38

2 Python Sky Model: software for simulating the Galactic microwave sky **39**

2.1	Introduction	40
2.2	Large-scale simulations	41
2.2.1	Synchrotron	41
2.2.1.1	Model 1: Nominal index	42
2.2.1.2	Model 2: Spatially steepening index	44
2.2.1.3	Model 3: Curvature of index	44
2.2.2	Thermal dust emission	45
2.2.2.1	Model 1: Nominal index	46
2.2.2.2	Models 2 and 3: Spatially variable index	46
2.2.2.3	Model 4: Two dust temperatures	49
2.2.3	Anomalous microwave emission	50
2.2.3.1	Model 1: Nominal unpolarised AME	50
2.2.3.2	Model 2: Polarised AME	50
2.2.4	Free-free	51
2.2.5	CMB	52
2.2.6	Instrument	52
2.3	Small-scale simulations	52
2.3.1	Polarisation	55
2.3.2	Intensity: Synchrotron	57
2.3.3	Intensity: Free-free	58
2.3.4	Intensity: Thermal dust and AME	60
2.4	Frequency decorrelation	60
2.5	Discussion	63

2.5.1	Forecast example	63
2.5.2	Conclusion	65
3	Finding the chiral gravitational wave background of an axion-SU(2) inflationary model using CMB observations and laser interferometers	66
3.1	Introduction	67
3.2	Theory	70
3.3	CMB	73
3.3.1	Cosmic-variance limited case	74
3.3.2	Including instrument noise and foreground contamination . . .	77
3.3.2.1	BB Signal-to-Noise	77
3.3.2.2	TB+EB Signal-to-Noise	77
3.3.3	Simultaneous detection and self-calibration	78
3.3.4	CMB Results	82
3.4	Laser Interferometers	83
3.4.1	Interferometer notation	85
3.4.2	Interferometer response	86
3.4.2.1	One constellation	89
3.4.2.2	Two-constellations	90
3.4.3	Interferometer signal-to-noise	91
3.4.4	Interferometer results	93
3.5	Discussion	94
4	ACT foregrounds	96
4.1	Introduction	97
4.2	Data	98
4.3	Method	98
4.3.1	Theory	98
4.3.2	Power spectrum estimation	101
4.3.3	Calculation of uncertainties	103
4.3.4	Likelihood	103
4.4	Results	105
4.4.1	Temperature	105
4.4.2	Polarization	109
4.5	Conclusions	109

5	Removal of Galactic foregrounds for upcoming primordial gravitational wave searches	111
5.1	Introduction	112
5.2	Simulations	113
5.2.1	Galactic simulations	114
5.2.2	CMB	115
5.2.3	Sky area and expected noise levels	115
5.3	Component Separation and parameter estimation	117
5.3.1	Component separation	117
5.3.2	B-mode power spectrum estimation	119
5.3.3	Cosmological parameter likelihood	120
5.4	Results	121
5.4.1	Fiducial sky simulations	121
5.4.2	Establishing source of bias	122
5.4.2.1	Masking the Galactic plane	122
5.4.2.2	Cross-correlating with foreground templates	123
5.4.3	Spatially varying spectral indices	124
5.4.3.1	Fitting spectral indices on a HEALPix grid	125
5.4.3.2	With prior information on spatial variation of indices	126
5.5	Conclusions	128
6	Conclusions	130
A	Derivation of the template for GW power spectrum	133
B	Calculation of the covariance matrix	137
C	CMB noise spectrum	138
D	Frequency dependence of CMB sensitivity	141
	Bibliography	143

List of Figures

1.1	Figure taken from Ref. [4] showing the stages of evolution of a perturbation: generation, exiting the horizon during inflation, superhorizon evolution as a comoving curvature perturbation, and reentry of the horizon as the Universe expands.	7
1.2	Marginalized 68 % and 95 % constraints on the n_s - r plane from a combination of 2018 <i>Planck</i> BICEP / Keck, and baryon acoustic oscillation data, presented in Ref. [5].	15
1.3	Figure taken from Ref. [6] showing various contributions to the temperature power spectrum. The rise at multipoles $\ell < 10$ is due to the late time integrated Sachs-Wolfe effect due to the cosmological constant. The Sachs-Wolfe Plateau is an almost scale-invariant interval of perturbations corresponding to the initial curvature perturbations on scales larger than the η_* . The acoustic peaks correspond to the interval $\eta_* < \lambda < \eta_{\text{Silk}}$ in the baryon-photon fluid behaves as a harmonic oscillator supported by photon pressure, under gravitational forcing. At scales $\ell \gtrsim 1000$ Silk damping due to photon diffusion results in exponential suppression of the acoustic peaks.	20
1.4	Figure taken from Ref. [7] comparing constraints from different experiments on: TT , EE , and BB in the top panel; TE in the middle panel; and $\phi\phi$ in the bottom panel. from various experiments, many of which are described in Section 1.3.4.1. The dashed line is the best-fit model from the <i>Planck</i> results [8].	25

1.5	Figure comparing the theoretical B -mode power spectra for lensing and primordial contributions, with constraints from different observatories, and observed foreground levels from the Planck experiment (the range corresponds to observations covering 24% to 71% of the sky). The primordial B mode spectrum generated at the surface of last scattering, for several different values of r is shown as the dot-dash line. The lensing contribution is shown as the dashed black line, and the sum of lensing and $r = 0.01$ is shown in solid black. Constraints from BICEP2+Keck (circle), POLARBEAR (diamond), and SPTpol (square) are overplotted.	28
1.6	Figure taken from Ref. [9], showing the RMS variation in the polarized foregrounds as a function of frequency. Synchrotron (green) dominates at low frequencies, and falls off as $\sim \nu^{-3}$, dust (red) dominates at high frequencies, and within the CMB frequencies of interest has a power law dependence of $\sim \nu^{1.5}$. The CMB emission (teal) is roughly constant below $\nu \sim 100$ GHz, and falls off above this limit. The width of each of the lines represents the range of foreground signal for f_{sky} between 27% and 83%. The gray bands show the bands in which the Planck experiment made observations [9]. Note that there is not sky fraction, and no frequency at which the CMB dominates over foreground emission.	32
2.1	The frequency scaling laws for the individual components of PySM; we show the nominal and alternative models as solid and dashed lines respectively. We show only the alternative models which have a significant impact on the shape of the spectrum. These spectra are calculated by producing masked maps of each component at each frequency, smoothing to FWHM 1° in intensity and $40'$ in polarisation, and then computing the RMS. The mask used in intensity is the WMAP 9 year KQ85 mask, and the polarisation mask is the <i>Planck</i> polarisation confidence mask CPM83.	42
2.2	Template maps used in the PySM models. All emission templates are in units of μK_{RJ} and all dust temperature templates are in K. Intensity templates are plotted on a log scale, and the polarisation templates on a linear scale.	43

2.3	Normalised histograms of the dust spectral index, β_d , calculated for noisy simulations of different PySM models with varying intrinsic index dispersion. We see that the resulting dispersions are very similar, indicating noise-dominated data.	48
2.4	Mask used in calculation of β_d in section 2.2.2.2. This mask is an approximation to the one used in the <i>Planck</i> analysis (figure 1 of Ref. [10]).	48
2.5	<i>Left</i> : synchrotron polarisation spectra in a square region centred on RA, DEC = [0, -55] of size 1600 deg ² . The errors shown are cosmic variance only. The best-fit power-law signal plus noise model from Eqn. 2.15 is shown. The BB model minimum is used to estimate the scale l_* to smooth the maps. <i>Right</i> : synchrotron polarisation spectra computed with the <i>WMAP</i> polarisation analysis mask, and best-fit model. The dashed lines are the extrapolated power laws used in the small-scale simulation.	53
2.6	<i>Left</i> : dust polarisation spectra as in Figure 2.5, but for a smaller patch of 800 square degrees. <i>Right</i> : dust polarisation spectra as in Figure 2.5, but using the <i>Planck</i> Gal 80 Galactic plane mask.	53
2.7	<i>Left</i> : synchrotron <i>BB</i> spectra using the 1600 square degree region centred on RA, DEC = [0, -55]. We show the original template, the smoothed template, the small scale realisation, and the final map with small scales added. The dashed red line shows the shape of the power law of the small scale realisation to guide the eye. <i>Right</i> : synchrotron <i>BB</i> spectra over 75% of the sky using the <i>WMAP</i> polarisation analysis mask.	54
2.8	<i>Left</i> : dust <i>BB</i> spectra in the 825 square degree region centred on RA, DEC = [0, -55], as in Figure 2.7. <i>Right</i> : dust <i>BB</i> spectra in the Gal 80 region, as in Figure 2.7.	54
2.9	Normalisation map, $N(\hat{\mathbf{n}})$, for the dust <i>Q</i> map.	55
2.10	Gnomic projection of dust <i>Q</i> maps in a patch centred at RA, DEC = (0, -55); 40 degrees to a side. The left panel is the original map, the middle panel has been smoothed (M_0), and the right hand panel has had small scales added ($M_0 + M_{ss}$). The maps are plotted in histogram-equalisation in units of μK_{RJ}	56

2.11	Synchrotron (top), free-free (middle), and AME (bottom) simulated intensity maps in a patch of side 40° centred at RA, DEC as indicated. Left: original template; right: simulation including small scales. These have been plotted in histogram-equalisation to increase the dynamic range.	59
2.12	Examples of the frequency dependence of a given component for different correlation lengths ξ , normalised by the uncorrelated case, given a template at $\nu = 300$ GHz. The sharp kinks at low frequencies for small values of ξ are caused by the frequency correlation length being smaller than the sampling rate.	61
2.13	Forecasted posterior distribution for r estimated for one of the LiteBIRD-like simulations described in Section 2.5. The black line shows the results assuming perfect knowledge of the foreground spectral indices, while the red line corresponds to the case of spatially-varying spectral indices on pixels of $N_{\text{side}} = 16$, which more than doubles the uncertainty on r	63
3.1	The predicted tensor spectrum, $\mathcal{P}_h^{\text{L,Sourced}}$, for three sets of parameters: (blue: $\sigma = 2$, $r_* = 0.07$, $k_p = 0.005 \text{ Mpc}^{-1}$), (orange: $\sigma = 2$, $r_* = 0.07$, $k_p = 0.0005 \text{ Mpc}^{-1}$), (green: $\sigma = 2$, $r_* = 0.07$, $k_p = 7 \times 10^{-5} \text{ Mpc}^{-1}$).	72
3.2	LEFT PANEL: C_ℓ^{BB} for the same three sets of parameters used in Figure 3.1: (blue: $\sigma = 2$, $r_* = 0.07$, $k_p = 0.005 \text{ Mpc}^{-1}$), (orange: $\sigma = 2$, $r_* = 0.07$, $k_p = 0.0005 \text{ Mpc}^{-1}$), (green: $\sigma = 2$, $r_* = 0.07$, $k_p = 7 \times 10^{-5} \text{ Mpc}^{-1}$) compared to the LiteBIRD noise spectrum with 2% foregrounds (solid black), the lensing BB spectrum (dashed black), and the standard vacuum fluctuation $C_\ell^{BB}(r = 0.07)$ consistent with the BKP $r < 0.07$ (95% C.L.) (dash-dot black). The axion-SU(2) spectra contain a contribution from vacuum fluctuations with $r = 10^{-5}$, as is used in the text. RIGHT PANEL: $ C_\ell^{TB} $ (solid colour) and $ C_\ell^{EB} $ (dashed colour) spectra for the same three sets of parameters. Shown in black is an example of the spurious TB signal induced by polarimeter miscalibration for an angle of one arcminute, as discussed in Section 3.3.3.	75
3.3	Signal-to-noise of TB + EB spectra assuming the perfect case of $f_{\text{sky}} = 1$, with no foreground contamination, noiseless observations, and no delensing. The black dashed line indicates the bounds placed by $r_* < 0.07$	76

3.4	Signal-to-noise of BB spectrum assuming no delensing and 2% foreground contamination and LiteBIRD instrumental noise added using method described in Appendix C. The dashed line refers to the observational constraint of $r_* = 0.07$. The signal-to-noise achieved in BB is much larger than that in TB+EB as the cosmic variances of BB and TB are proportional to $(C_\ell^{BB})^2$, and $C_\ell^{BB}C_\ell^{TT}$, respectively. The factor of C_ℓ^{TT} means cosmic variance in the TB spectrum is much more significant than in BB.	78
3.5	Signal-to-noise of TB + EB spectra assuming no delensing and 2% foreground contamination and LiteBIRD instrumental noise added using method described in Appendix C. The dashed line refers to the observational constraint of $r_* = 0.07$	79
3.6	Correlation coefficient α (left) and 1- σ uncertainty on the polarimeter calibration (right) for LiteBIRD, calculated over the allowed parameter space of the model assuming 2% foreground residual and no delensing. The dashed line shows the observational constraint of $r_* = 0.07$	82
3.7	Comparison of the sensitivity curves for LiteBIRD, Planck, LISA, and BBO corresponding to a signal-to-noise of one at a given frequency in intensity (I) or polarization (V). At the top horizontal axis we also show the corresponding wavenumber computed via $\frac{k}{\text{Mpc}^{-1}} = 6.5 \times 10^{14} \frac{f}{\text{Hz}}$. Also plotted are the primordial spectra for the parameters: $k_p = 10^{13} \text{ Mpc}^{-1}$, $\sigma = 9$, $r_* = 835$ (black dotted) and $k_p = 10^{11} \text{ Mpc}^{-1}$, $\sigma = 8$, $r_* = 0.15$ (black dash-dotted). Note that below $f = 10^{-17} \text{ Hz}$ transfer function of the fractional gravitational wave background energy density changes due to the transition between matter and radiation dominated eras. We see that even for the large values of σ required by the large k_p values of the axion-SU(2) model LISA and BBO can make a detection that would still be inaccessible at CMB scales. This motivates the evaluation of signal-to-noise for the interferometers. Note that the CMB sensitivity curves have been smoothed with a Gaussian kernel due to the sharp oscillations introduced by the transfer function (Equation D.2). . . .	84

3.8	Possible designs for future space-based laser interferometers. The blue arrows show the laser links used in the Michelson interferometer signals defined in Equation 3.19. LEFT PANEL: One constellation design of a space-based interferometer, this corresponds to the baseline LISA design. The points \mathbf{x}_i show the i^{th} satellite. RIGHT PANEL: An advanced stage design of LISA or BBO with two constellations. The points \mathbf{x}_{ij} show the j^{th} satellite on the i^{th} constellation.	87
3.9	Signal-to-noise contours obtained using Equation 3.24 for a LISA-like experiment described in §3.4.3. The primordial spectrum has $k_p = 1 \times 10^{11} \text{ Mpc}^{-1}$	93
3.10	Signal-to-noise contours obtained using Equation 3.24 for a BBO-like experiment described in §3.4.3. The primordial spectrum has $k_p = 1 \times 10^{13} \text{ Mpc}^{-1}$	94
4.1	Map from Ref. [11], showing the various ACTPol fields overplotted on a map of dust emission. In this chapter we consider the D56, D8 and BOSS-N regions, and not D1, D2, or D9. D56 overlaps with the ACT-Equatorial region from previous observation seasons (e.g. Ref. [12]), and the D8 region overlaps with the ACT-Southern region.	99
4.2	Figure from [12] showing the dust levels in the ACT Equatorial and ACT Southern regions, at 148 GHz and 218 GHz. Shown in the teal line are the results for ACT-E at 148 GHz, which is most similar to our analysis of D56, and in purple the results for ACT-S at 148 GHz, which is most similar to our analysis of D8.	104
4.3	Figure taken from Ref. [13], showing a computation of the difference between the 353 GHz Planck TT spectrum, and the 149 GHz ACT TT spectrum, which is attributable to dust + CIB. Overplotted in the shaded region is the Das 2014 dust model, with dust amplitude $A_d = 0.8 \mu\text{K}^2$	105
4.4	Figure showing the TT data (black points), and the fitted dust + CIB model (black), for the D56 patch. The best-fit dust model, marginalized over CIB amplitudes, is shown in dashed black. These are compared to the results of Ref. [12] for the ACT equatorial region (blue), which overlaps with D56. Shown in orange is the CIB prior model from [14].	106

4.5	Figure showing the TT data (black points), and the fitted dust + CIB model (black) for the D8 patch. The best-fit dust model, marginalized over CIB amplitudes, is shown in dashed black. These are compared to the results of Ref. [12] for the ACT southern region (blue), which borders D8. Shown in orange is the CIB prior model from [14].	107
4.6	Figure showing the TT data (black points), and the fitted dust + CIB model (black) for the BOSS-N patch. The best-fit dust model, marginalized over CIB amplitudes, is shown in dashed black. These are compared to the results of Ref. [12] for the ACT equatorial region (blue). Shown in orange is the CIB prior model from [14].	108
4.7	Figure showing the calculated dust EE, BB, and TE, spectra for BOSS-N. BOSS-N is the largest sky region we consider, and contains more dust in intensity than D56 or D8, as established in Section 4.4.1.	110
5.1	The analytic model SEDs of the CMB, thermal dust, and synchrotron described in Section 5.2.1. Amplitudes have been rescaled to compare just the shapes of the curves. The vertical lines indicate frequencies at which SO is due to make observations; two channels characterize the low frequency synchrotron, two channels characterize the high frequency dust, and two channels observe the CMB around 100 GHz.	114
5.2	Expected noise curves for SAT polarized observations, from Ref. [3], for baseline sensitivity and optimistic ℓ_{knee} parameter. The lensing power spectrum is also shown for comparison.	117
5.3	Maps in Galactic coordinates showing simulations of the Q Stokes parameter for SIMSET1 with the baseline sensitivity and optimistic ℓ_{knee} configuration of SO. Note that the colorscale is a combination of a linear scale between $(-0.1, 0.1) \mu\text{K}$ and a log scale outside this range, to display structure over a large dynamic range. In the 27 GHz channel, the large scale structure of synchrotron emission is visible, with the North Galactic Spur clearly visible above the Galactic plane. At the CMB frequencies 93 GHz and 145 GHz, the CMB E-mode polarization becomes visible at high Galactic latitudes. At the highest frequency channels the morphology of dust emission becomes dominant.	118

5.4	Forecasted constraints on r , for a model with $r = 0$, from simulated observations of SIMSET1 for the four Simons Observatory noise levels. The results from this work, assuming no spatial variation of the spectral indices (black), are compared to results from the <code>xForecast</code> method presented in SO19 (grey).	121
5.5	Forecasted constraints on r as a function of sky area, from simulated observations of SIMSET1 (diamonds) and SIMSET2 (squares), when foregrounds are removed assuming no spatial variation of spectral indices. Each point represents a different level of Galactic masking when computing the power spectrum of the cleaned CMB maps. The bias is reduced when the brightest sky regions are masked.	123
5.6	The cross-correlation of cleaned CMB maps with Galactic templates can be used to test for residual foregrounds. This shows the cross-spectrum when cleaning with a spatially constant dust SED and synchrotron spectral index (diamonds) or accounting for spatial variation (circles). We see a non-zero signal at the largest scales when using the spatially constant models, which correlates with a bias in r . (Left) This uses SIMSET1, and correlates with a template of Galactic dust. (Right) This uses SIMSET2 and correlates with a synchrotron template.	124
5.7	(Left and middle) The fitting regions used to fit different dust SEDs, as defined by the algorithm described in 5.4.3.2, for $N_{\text{spec}} = 2, 6$. (Right) The fitting regions using a Healpix $N_{\text{side}} = 2$ grid.	125
5.8	Forecasted constraints on r when fitting spectral indices defined by HEALPix grids at increasing resolution. Since SO observes a fraction of the sky, the number of coarse pixels is given by $12f_{\text{sky}}N_{\text{side}}^2$. As N_{side} increases we find the projected uncertainty on r increases significantly: this is not a good choice for dividing up the sky area.	126
5.9	(Left) Forecasted constraints on r when fitting models with increasing independent regions defined from the β_{d1}^{dust} template, for the optimistic ℓ_{knee} goal and baseline sensitivities. The bias is removed with a modest increase in parameters, if these regions are known a-priori. Right: using SIMSET2 and fitting models with regions defined from the β_{s2}^{sync} template, a similar effect is seen	127

A.1	The comparison between the full numerical result of $\mathcal{P}_h^{\text{Sourced}}(k)$ (blue solid line) and the template eq. (3.3) with eq. (A.4) and $\sigma^2 = 0.15\Delta N^2$ (red dashed line). In the left (right) panel, $\Delta N = 5(10)$, $m_* = 4$, $\epsilon_{B*} \approx 9 \times 10^{-4}$ and the peak amplitude reaches the tensor-to-scalar ratio, $r_* = 0.05$. The Hubble parameter is set as $H_{\text{inf}} = 8 \times 10^{11}$ GeV which corresponds to $r = 10^{-5}$ without the sourced GW. In the case of $\Delta N = 5$, the derived formula slightly underestimate the peak amplitude and the width, while the fit is excellent for $\Delta N \gtrsim 10$	135
A.2	Peak tensor-to-scalar ratio r_* as a function of ϵ_{B*} and m_* for $k_p = 0.005 \text{ Mpc}^{-1}$	136

Chapter 1

Introduction

Cosmology has, over the past few decades, experienced an influx of data that has revolutionized the field. Using many different physical probes, from the earliest light emitted, to the gravitational waves created by colliding black holes, we have rigorously tested the best theories of our Universe's 13.7 billion year history. The 'concordance' model, Λ CDM, describes a Universe originating from a hot, dense, state, in which it was almost perfectly homogeneous and isotropic. Small inhomogeneities were imprinted upon the density of the Universe by some primordial process, and over time these fluctuations grew into the galaxies, and clusters of galaxies, we see today. The process that imprinted the initial density fluctuations is not yet understood. A popular theory, *inflation*, suggests that a period of rapid expansion amplified quantum fluctuations in the density, which were then 'frozen-in' on cosmological scales [15, 16, 17]. However, recent works have suggested a collapsing previous Universe may be responsible [18]. During inflation, fluctuations in the fabric of spacetime (gravitational waves) would have been amplified too, creating a gravitational wave background (GWB) [19, 20]. If present, the GWB would leave a characteristic fingerprint on the linear polarization of the cosmic microwave background (CMB). Measuring this curling pattern is a major driving force behind many current and proposed CMB observatories.

1.1 Concordance Cosmology

The best model that we currently have to explain our observations of the Universe is the Λ CDM model. This model describes a flat Universe, composed of dark matter,

baryons, radiation in the form of photons and neutrinos, and dark energy in the form of a cosmological constant, Λ .

In this section we introduce the metric of a homogeneous and isotropic Universe, and derive the equations that govern the evolution of matter and the metric itself with time. We will go on to the framework to understand how small perturbations to the uniform background metric and matter distribution are formed. Section 1.2 will go into detail about how perturbations to this background may be sourced, and how they evolve into the cosmological observables we use today.

1.1.1 FLRW metric

An important principle within cosmology, the *cosmological principle*, states that the laws of physics apply similarly throughout the Universe, or equivalently, that that Universe is expected to be *isotropic* and *homogeneous*. This principle started out as an assertion, but over many years, observations of the structure in our Universe have found that the Universe is indeed homogeneous and isotropic on scales larger than ~ 100 Mpc. The most generic metric that satisfies these constraints, and solves Einstein's equations, is the Friedmann-Lemaître-Robertson-Walker (FLRW) metric [21, 22, 23, 24]:

$$ds^2 = c^2 dt^2 - a^2(t) \left[\frac{dr^2}{1 - kr^2} + r^2(d\theta^2 + \sin^2 \theta d\phi^2) \right], \quad (1.1)$$

where t is proper time, and $a(t)$ is a time-dependent scale factor normalized to be unity today, and k describes the spatial curvature of 3-surfaces at constant time. The constraints of homogeneity and isotropy require that curvature be constant, and so we have only three possibilities for the spatial hypersurfaces: *flat* ($k = 0$), *open* ($k = -1$), or *closed* ($k = 1$). It will be useful to define the spatial part of the metric separately:

$$\gamma_{ij} dx^i dx^j = \left[\frac{dr^2}{1 - kr^2} + r^2(d\theta^2 + \sin^2 \theta d\phi^2) \right]. \quad (1.2)$$

From this point on we will work in natural units, in which $c = 1$. We also define *conformal time*, η , as: $\eta = \int_0^t \frac{dt}{a}$. This can be interpreted as the maximum distance traveled by light since $t = 0$, the *horizon*. Conformal time allows us to write the FLRW metric as:

$$ds^2 = a(\eta) [d\eta^2 - \gamma_{ij} dx^i dx^j], \quad (1.3)$$

The Einstein equations relate the metric to the matter content of space [25]:

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = 8\pi GT_{\mu\nu} + g_{\mu\nu}\Lambda, \quad (1.4)$$

where $R_{\mu\nu}$ is the Ricci tensor (which is a function purely of the metric), R is the fully-contracted Ricci tensor, $T_{\mu\nu}$ is the stress-energy tensor, G is the gravitational constant, and Λ is a cosmological constant. Under the constraints on homogeneity and isotropy, we can describe the energy content of the Universe as a perfect fluid at equilibrium and at rest, with pressure and energy density that is a function only of time:

$$T^{\mu}_{\nu} = \begin{pmatrix} \rho & 0 & 0 & 0 \\ 0 & -P & 0 & 0 \\ 0 & 0 & -P & 0 \\ 0 & 0 & 0 & -P \end{pmatrix}, \quad (1.5)$$

where P is pressure, and ρ is energy density. In addition, the energy-momentum tensor is subject to the continuity equation, $\partial_{\mu}T^{\mu}_{\nu} = 0$, which simplifies to:

$$\dot{\rho} + 3\frac{\dot{a}}{a}(P + \rho) = 0. \quad (1.6)$$

Substituting the FLRW metric in Equation 1.1 and the stress-energy tensor in Equation 1.5 into the Einstein equations (1.4) results in two independent equations for the evolution of the scale factor, the *Friedmann equations*:

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho + \frac{\Lambda}{3} - \frac{k^2}{a^2(t)}, \quad (1.7)$$

$$\frac{\ddot{a}}{a} + \frac{1}{2}\left(\frac{\dot{a}}{a}\right)^2 = -4\pi Gp + \frac{\Lambda}{3}, \quad (1.8)$$

Equations (1.6) to (1.8) may be solved for the background evolution of a homogeneous, isotropic Universe with a given energy content.

1.1.2 Energy content of the Universe

As described in the introduction, the Λ CDM model describes a Universe made up of radiation, cold matter, and a cosmological constant. The bulk behavior of each of these energy contents will be different, but all can be parameterized by a simple *barotropic* equation of state:

$$\frac{P}{\rho} = w, \quad (1.9)$$

where the constant w takes a different value depending on the energy content:

$$w = \begin{cases} 1/3 & \text{radiation,} \\ 0 & \text{matter,} \\ -1 & \text{cosmological constant.} \end{cases} \quad (1.10)$$

It should be noted that w is not restricted to these values, and it is possible to imagine other matter contents which take on different values. Using the barotropic parameterization, we can apply Equation (1.6) to determine how the energy density of a single-component Universe depends on the scale factor: $\rho \propto a^{\frac{1}{1+w}}$. Curvature does not enter Equation (1.6), however it is clear from Equation (1.7) that it plays an important role dynamically. It is conventional to express curvature as an effective ‘energy density’ by absorbing it into the ‘ ρ ’ of Equation (1.7). Doing so, it is clear that $\rho_k \propto a^{-2}$. However, the interpretation of a pervading energy density due to the intrinsic curvature of space should be avoided.

Setting $\Lambda = 0$ in Equation 1.7, we see that the curvature will always be zero for the *critical density*: $\rho_c = 3H^2/(8\pi G)$, which today is roughly 10^{-29} g cm⁻³. Using this as a reference density, we then define the dimensionless quantities $\Omega_i = \frac{\rho_i}{\rho_c}$ where $i = m, r, \Lambda, k$, labels matter, radiation, dark energy, or curvature, respectively. Equation 1.7 may now be rewritten as:

$$H^2 = H_0^2 [\Omega_{r,0}a^{-4} + \Omega_{m,0}a^{-3} + \Omega_{k,0}a^{-2} + \Omega_{\Lambda,0}], \quad (1.11)$$

where the subscript ‘zero’ indicates a quantity measured today (recalling that $a(t_0) \equiv 1$).

For mixtures of components, this equation generally has to be solved numerically for the evolution of $a(t)$. However, in simple cases where one form of energy density

dominates, informative solutions may be found analytically:

$$a \propto \begin{cases} t^{\frac{2}{3(1+w)}} & w \neq -1. \\ e^{Ht} & w = -1. \end{cases} \quad (1.12)$$

There are several interesting qualitative conclusions that we can draw from Equation (1.12). The first derivative of the scale factor is positive for all types of energy density we have considered, indicating that the Universe is always expanding. The deceleration parameter, $q \equiv -\frac{\ddot{a}}{a^2}a$, can be calculated from a combination of Equations (1.7) and (1.8), to find $q = \frac{2}{3}(\frac{1}{3} + w)$. So, for matter or radiation-dominated Universes, the expansion is decelerating, whereas for $w < -1/3$, which includes cosmological constant-dominated Universes, the expansion is accelerating.

The present-day values have been measured by the *Planck* satellite to be¹ [8]:

$$\begin{aligned} \Omega_{\Lambda,0} &= 0.68 \pm 0.01, \\ \Omega_{m,0} &= 0.32 \pm 0.01, \\ \Omega_{k,0} &= 0.0007 \pm 0.0019 \\ \Omega_{r,0} &< 10^{-4} \\ H_0 &= (67.27 \pm 0.60) \text{ km s}^{-1} \text{ Mpc}^{-1} \end{aligned} \quad (1.13)$$

Measuring the values of these parameters today allows us to then work backwards in time, piecing together the various epochs dominated by different energy contents, to form a history of our Universe.

1.1.3 Perturbations to the FLRW metric

Structure in the Universe began as small deviations from homogeneity and isotropy, which grew over time by gravitational collapse. In this section we will outline the process of deriving predictions for the evolution of perturbations to the metric, and to

¹It should be noted that there is an appreciable $\sim 3.6 \sigma$ tension between high-redshift and low-redshift probes of the Hubble parameter. Refs. [26, 27, 28] use observations of galactic and extragalactic Cepheid variable stars to build a local distance ladder, from which they infer a Hubble constant of $H_0 = (73.48 \pm 1.66) \text{ km s}^{-1} \text{ Mpc}^{-1}$.

the stress-energy tensor.

Given a smooth background metric, $\tilde{g}_{\mu\nu}$, we may write down a generic perturbation to Equation 1.3 for a flat Universe:

$$ds^2 = a^2(\eta) [(1 + 2A)d\eta^2 - 2B_i dx^i d\eta - (\delta_{ij} + h_{ij})dx^i dx^j], \quad (1.14)$$

where A has one degree of freedom, B_i has three, and due to symmetry requirements the tensor h_{ij} has 6. These degrees of freedom can be largely reduced by taking advantage of *gauge freedom*. Under a gauge transformation, the coordinates are transformed as $X^\mu \rightarrow X^\mu + \xi^\mu(\eta, \mathbf{x})$, and one can derive how the metric transforms by considering the invariant line element, $ds^2 = g_{\mu\nu} dx^\mu dx^\nu$. The gauge transformation properties of the perturbations in Equation 1.14 are more easily leveraged when the perturbations are decomposed using scalar-vector-tensor decomposition [29]. This decomposition splits three-vectors into the sum of a divergenceless vector, and the gradient of a scalar: $B_i = \partial_i B + \hat{B}_i$. Similarly, for tensor perturbations: $h_{ij} = 2C\delta_{ij} + 2(\partial_i \partial_j - \frac{1}{3}\delta_{ij}\nabla^2)E + \partial_i \hat{E}_j + \partial_j \hat{E}_i + \hat{E}_{ij}$. The result of this decomposition is: four scalar degrees of freedom, A, B, C, E ; four vector degrees of freedom, the components of \hat{B}_i and \hat{E}_i ; and two tensor degrees of freedom, the components of \hat{E}_{ij} .

Whilst vector perturbations may be produced by some physical processes, they will generally decay quickly in an expanding Universe, and so we will not consider them any further here.

A similar perturbation expansion may be followed for the stress-energy tensor. The perturbation $\delta T_{\mu\nu}$ is comprised of the two scalar quantities $\delta\rho$ and δP , the three-vector v^i , and the anisotropic stress tensor Π^{ij} .

The scalar degrees of freedom can now be reduced by a judicious choice of gauge. We may either form gauge-invariant combinations of perturbations, or fix the gauge such that some perturbations are zero. Some examples of gauge choices are:

Newtonian gauge: here we set $B = E = 0$, such that we have two remaining scalar degrees of freedom. The perturbed FLRW metric then becomes:

$$ds^2 = a^2(\eta) [(1 + 2\Phi)d\eta^2 - (1 + 2\Psi)\delta_{ij}dx^i dx^j]. \quad (1.15)$$

It transpires that in this choice of gauge, we can identify A and C with *Bardeen* potentials Φ and Ψ , which are gauge invariant quantities, and therefore provide an interface to results derived in other gauges.




Figure 1.1: Figure taken from Ref. [4] showing the stages of evolution of a perturbation: generation, exiting the horizon during inflation, superhorizon evolution as a comoving curvature perturbation, and reentry of the horizon as the Universe expands.

Spatially-flat gauge: here we set $C = E = 0$. This gauge is especially useful for deriving the fluctuations in the inflaton field, ϕ , that will set the perturbations Φ , and will be used in Section 1.2.

Comoving-gauge: This gauge is defined such that the three-velocity of the matter perturbations is zero: $v^i = 0$. This can be especially useful when deriving the perturbation's initial conditions.

Whilst some examples of useful gauge-invariant quantities are:

Comoving curvature perturbation: It will transpire in the next section that forming a gauge-invariant combination of the matter and metric perturbations will allow us to track the evolution of perturbative modes as they exit and enter the horizon. The comoving curvature perturbation is popular in the literature for this reason:

$$\mathcal{R} = C - \frac{1}{3}\nabla^2 E + \mathcal{H}(B + v) \quad (1.16)$$

With a perturbed metric in hand, the task is then to understand how the perturbations evolve dynamically. To answer this question, we must once again solve the Einstein equations, Equation 1.4, for the perturbed metric in Equation 1.15, and the perturbed stress energy tensor.

1.1.3.1 Scalar Perturbations

Einstein's equations may be applied to scalar and tensor perturbations separately, allowing us to consider them one-at-a-time. The key will be to consider perturbations to the potential Φ and the matter content together in a gauge-invariant combination that does not evolve at super-horizon scales. This will allow us to track perturbations as they are generated at early times, leave the horizon, and then reenter at some point in the future as the horizon increases. The *comoving curvature perturbation*, \mathcal{R} in the Newtonian gauge is given by:

$$\mathcal{R} = -\Phi + \mathcal{H}v, \quad (1.17)$$

which can be combined with the perturbed Einstein equations to show that for adiabatic matter content $d\ln(\mathcal{R})/d\ln(a) \propto (k/H)^2$, to wit., for modes outside of the horizon, \mathcal{R} is constant.

Theoretical predictions of \mathcal{R} will be specified by in terms of the power spectrum:

$$\langle \mathcal{R}_{\mathbf{k}} \mathcal{R}_{\mathbf{k}'} \rangle = (2\pi^3) \delta(\mathbf{k} - \mathbf{k}') P_{\mathcal{R}}(k), \quad (1.18)$$

where angle brackets indicate an ensemble average. Finally we define the primordial power spectrum as:

$$\mathcal{P}(k) \equiv \Delta_{\mathcal{R}}^2(k) = \frac{k^2}{2\pi^2} P_{\mathcal{R}}(k) \quad (1.19)$$

1.1.3.2 Tensor Perturbations

We now consider the tensor perturbations to the metric, \hat{E}_{ij} . Via the SVT decomposition, we are able to treat tensor perturbations independently from the scalar sector. The matter content will therefore only affect the background on which tensor perturbations propagate, and the metric perturbations can be treated as small perturbations on this background. We will expand $E_{ij}(t, \mathbf{x})$ using a transverse traceless set of basis tensors, with normalization $e_{ij}^A e_{ij}^{A'} = 2\delta_{AA'}$, where $A = +, \times$:

$$\hat{E}_{ij} = \sum_A h_A(\mathbf{x}, t) e_{ij}^A \quad (1.20)$$

where $A = +, \times$ labels the linear polarization state of the tensor. It transpires that the amplitudes h_A are gauge-invariant, and evolve independently of one another,

according to:

$$\ddot{h}_A(k) + 3H\dot{h}_A(k) - \frac{\eta k^2}{a^2}h_A(k) = 0. \quad (1.21)$$

The initial conditions for tensor perturbations will later be specified in terms of the power spectrum of the expansion coefficients h_A :

$$\langle h_k^A h_{k'}^{A'} \rangle = (2\pi)^3 \delta_{AA'} \delta^3(\mathbf{k} + \mathbf{k}') P_h^A(k) \quad (1.22)$$

The primordial tensor spectrum is comprised of two contributions, one from each of the two linear polarizations:

$$\mathcal{P}_t(k) \equiv \Delta_h^2(k) = \sum_A \frac{k^3}{2\pi^2} \mathcal{P}_h^A(k) \quad (1.23)$$

1.1.4 Problems with Big Bang Cosmology

Horizon problem: When we observe two patches of the Universe, that are separated by a distance greater than the horizon, we see that they have roughly the same average properties. For example, our observations of the CMB show it to be well described by a blackbody spectrum, solely defined by its temperature. That temperature has been well measured to be $T_{\text{CMB}} = (2.72548 \pm 0.00057)$ K [30], with tiny spatial variations of order 10^{-5} . The implication of this uniformity is that all regions of the CMB must have been in causal contact prior to the surface of last scattering, in order to equilibrate to the same temperature. However, assuming only FLRW expansion, the horizon size when the CMB was formed can be calculated to find $\theta_{\text{horizon}} \sim 1$ degree, indicating that in an FLRW universe we would expect uniformity only on scales smaller than 1 degree.

Flatness problem: As was discussed in Section 1.1.2, the contribution of curvature to the energy density of the Universe can be expressed as

$$\Omega_k(t) = \frac{\Omega_{k,0}}{a^2} \frac{1}{\Omega_{\text{tot}}(t)}. \quad (1.24)$$

We noted that the *Planck* satellite measured a value of $\Omega_{k,0}$ that is consistent with zero (quoted in Equation (1.13)). Applying Equation (1.24), this value can be scaled back to its initial value at extremely early times. Such a calculation reveals that $\Omega_k(t_{\text{GUT}}) = 10^{-57}$, where GUT refers to the epoch of Grand Unified

Theory. Since there is no reason to prefer values of Ω_k close to zero, this requires a large degree of fine-tuning.

1.2 Inflation

Inflation proposes a period of exponential expansion of the Universe very early on in its history, before $\sim 10^{-34}$ s, as a solution to the problems outlined in Section 1.1.4 [15, 16, 17]. Inflation was first suggested by Ref. [15], and subsequently improved upon by Refs. [16, 17] to solve some technical problems, resulting in the flavor of inflation referred to as *slow roll inflation*. For a more pedagogical introduction consider the review of Ref. [31], and the lectures of Ref. [4]. For a review of the various inflationary models currently considered, see the review [32], and for the current constraints from the *Planck* satellite, Ref. [5].

The important quantity when considering which regions of the Universe may be in causal contact is the comoving *Hubble radius*, $1/aH$, which measures the distance over which signals can pass in one Hubble time. The particle horizon is the integral of the Hubble radius, and measures the distance traveled by light from some initial time, t_0 , which we take to be $a(t_0) = 0$, and so defines the size of regions in the Universe that may have equilibrated by time t :

$$\eta(a) = \int_{\ln a(t_0)}^{\ln a(t)} d \ln a' \frac{1}{a'H}. \quad (1.25)$$

Applying the Friedmann equations, it can be shown that the integrand is an increasing function of the scale factor:

$$\eta(a) = \int_0^a \frac{da'}{H_0} (\Omega_{m,0} a'^{-3} + \Omega_{r,0} a'^{-4} + \Omega_\Lambda)^{-1/2}, \quad (1.26)$$

and so has dominant contributions from late times. This is problematic, as it indicates that regions of the Universe causally disconnected at some time t , will have always been causally disconnected in the past. The solution offered by inflation is to amend Equation 1.25 with a preceding era, inflation, in which the Hubble radius shrinks, and large contributions can be made to the particle horizon:

$$\eta(a) = \int_{\ln a(t_0)}^{\ln a(t_e)} d \ln a' \frac{1}{a'H} + \int_{\ln a(t_e)}^{\ln a(t)} d \ln a' \frac{1}{a'H}, \quad (1.27)$$

where t_e denotes the time at the the end of the inflationary period. In order to receive a sufficiently large contribution to the particle horizon during the inflationary period, the Hubble radius must shrink:

$$\frac{d}{dt}(aH)^{-1} < 0 \rightarrow q < 0.$$

So, inflation may equivalently be thought of as an accelerating expansion. As per the discussion of Section 1.1.2, $q < 0$ can only be satisfied for $w < -1/3$. As has previously been discussed, a Universe consisting of only a cosmological constant satisfies these conditions. Indeed, a de Sitter Universe is equivalent to inflation occurring forever. We know, however, that inflation must have stopped, and by observing the value of Ω_Λ today, we know that the cosmological constant was subdominant during the inflationary epoch. Therefore, an additional type of matter, so far unobserved, with $w < -1/3$, is required to produce the inflationary conditions.

1.2.1 Single Field Slow Roll Inflation

The simplest field that satisfies the cosmological principle is a single scalar field. We will now consider the dynamics of a Universe in which the energy content is dominated by such a field, and show that it may produce the conditions necessary for inflation to occur.

1.2.1.1 Background evolution

The action for a scalar field in a background metric $g_{\mu\nu}$ is given by the Einstein-Hilbert action:

$$\mathcal{S} = \int d^4x \sqrt{-|g|} \left[\frac{R}{16\pi G} + \mathcal{L}_\phi \right] \quad (1.28)$$

where we assume the metric of a flat FLRW universe, $|g|$ is the determinant of said metric, and the Lagrangian for the scalar field is:

$$\mathcal{L}_\phi = \frac{1}{2} g^{\mu\nu} \partial_\mu \partial_\nu \phi - V(\phi). \quad (1.29)$$

Varying the action with respect to the metric recovers the Einstein Equation (1.4) with density and pressure given by the components of the stress-energy tensor for the

field ϕ . Applying Noether's theorem allows the stress-energy tensor to be derived [33]:

$$T^{\mu\nu} = \partial^\mu \phi \partial^\nu \phi - g_{\mu\nu} \mathcal{L}_\phi. \quad (1.30)$$

Considering the time-time and spatial parts of the Einstein equations give expressions for the pressure and energy density of the scalar field:

$$\rho_\phi = \frac{1}{2} \dot{\phi}^2 + V(\phi), \quad p_\phi = \frac{1}{2} \dot{\phi}^2 - V(\phi) \quad (1.31)$$

where we have dropped spatial derivatives due to the homogeneity constraints of the FLRW solution. Varying the action with respect to the scalar field, and again dropping spatial derivatives, we obtain the evolution equation of the scalar field:

$$\ddot{\phi} + 3H\dot{\phi} + \frac{\partial V}{\partial \phi} = 0 \quad (1.32)$$

This is the Klein-Gordon equation, and can also be derived from the Friedmann and continuity equations. Taking the ratio of the density and pressure gives us the equation of state of the field ϕ :

$$\frac{P_\phi}{\rho_\phi} = \frac{\frac{1}{2} \dot{\phi}^2 + V(\phi)}{\frac{1}{2} \dot{\phi}^2 - V(\phi)}. \quad (1.33)$$

Under the condition $\frac{1}{2} \dot{\phi}^2 \ll V(\phi)$, $P_\phi/\rho_\phi \rightarrow -1$, and inflation may occur. As the kinetic energy term is subdominant to the potential, this is known as the 'slow-roll' approximation. In the limit that the potential is completely flat the de Sitter solution is recovered. For small deviations from flatness, the de Sitter solution can be expanded in a set of small parameters. Defining the *slow-roll parameters* [34, 35, 4, 36]:

$$\epsilon \equiv -\frac{\dot{H}}{H^2}, \quad \eta \equiv \frac{d \ln \epsilon}{dN}, \quad (1.34)$$

one can expand the de Sitter solution for small deviations from a flat potential.

1.2.1.2 Perturbations

Inflation is able to solve the horizon and flatness problems described in Section 1.1.4. In addition to that, it is able to provide the initial seeds of inhomogeneity that grew into the structure we see today. In brief, the quantum behavior of the inflaton field

leads to some variation in the time at which inflation ended in different parts of the Universe, and consequentially the local over density. In this section we review the key results when perturbing the inflaton field, the derivation is too detailed to present and can be found in greater detail in Refs. [37, 38, 39, 4].

The background evolution we considered in Section 1.2.1.1 gives the homogeneous evolution of the field, $\bar{\phi}$, as a function of time. Expanding the background solution:

$$\phi(\mathbf{x}, t) = \bar{\phi}(t) + \delta\phi(\mathbf{x}, t). \quad (1.35)$$

In order to compute the evolution of the perturbations to ϕ , one can substitute Equation 1.35 into the action 1.28, and collect orders of $\delta\phi$ to recover the *Mukhanov-Sasaki equation*:

$$f_k'' + \left(k^2 - \frac{2}{\eta^2}\right) f_k = 0, \quad (1.36)$$

where $f(\mathbf{x}, t) = a \delta\phi(\mathbf{x}, t)$. The Mukhanov-Sasaki equation is generally solved numerically for a given background solution $\bar{\phi}(t)$, $a(t)$. However, it has informative limiting behavior for subhorizon and superhorizon modes. For perturbations of wavelength $\lambda \gg 1/(aH)$, the dominant solution gives a constant curvature perturbation, $\mathcal{R} \propto const.$, indicating that perturbations that exit the horizon are ‘frozen-in’ during the subsequent evolution, and before re-entering the horizon. On the other hand, for $\lambda \ll 1/(aH)$, Equation 1.36 reduces to that of a harmonic oscillator for each of the modes of f_k . The zero-point oscillations of this harmonic oscillator lead to the perturbations in the curvature and metric at the end of inflation.

So far we have applied classical perturbation theory. In order to derive the amplitude of the zero-point fluctuations, f_k must be quantized. The details are beyond the scope of this thesis, and may be found in Ref. [35]. The result of second quantization is a calculation of the variance of the mode f_k due to zero-point perturbations. This variance is evaluated as the mode f_k crosses the horizon, and freezes-in the classical perturbation \mathcal{R} :

$$\Delta_{\mathcal{R}}^2(k) = \frac{1}{8\pi^2} \frac{1}{\epsilon} \frac{H^2}{M_{\text{pl}}^2} \Big|_{k=aH}. \quad (1.37)$$

Since $\Delta_{\mathcal{R}}^2(k)$ is a function of k only through the Hubble parameter and the slow-roll parameter, ϵ , it is a weak function of wavenumber. Therefore, the scalar power

spectrum is often approximated as an almost scale-invariant spectrum:

$$\Delta_{\mathcal{R}}^2(k) = A_s \left(\frac{k}{k_*} \right)^{n_s - 1}, \quad (1.38)$$

where we have chosen some arbitrary pivot scale, k_* , and the scalar index ($n_s \approx 1$) measures the deviation from perfect scale-invariance.

The computation for the tensor perturbations proceeds exactly the same way, resulting in an equivalent equation of motion to Equation 1.36. The power spectrum for each of the two tensor polarizations is:

$$\Delta_h^2(k) = \frac{1}{\pi^2} \frac{H^2}{M_{\text{pl}}^2} \Big|_{k=aH}, \quad (1.39)$$

which is also parameterized as an almost scale-invariant spectrum:

$$\Delta_t^2(k) = A_t \left(\frac{k}{k_*} \right)^{n_t}. \quad (1.40)$$

It is common in the literature to quote constraints on the tensor power spectrum in terms of the ratio between the tensor and scalar amplitudes:

$$r = \frac{A_t}{A_s}, \quad (1.41)$$

which depends on the pivot scale k_* .

1.2.2 Observational status of inflation

The most stringent test of inflation so far comes from the 2018 *Planck* data release. Ref. [5] find that there is no evidence for deviation from a single field slow roll model of inflation with Einstein gravity. The *Planck* 2018 data is consistent with a purely power law power spectrum in the range of scales probed by the CMB, and a scalar spectral index of $n_s = 0.9649 \pm 0.0042$ (68% C.L.). When the *Planck* data is combined with *WMAP* [40] and BICEP / Keck [41] observations, a limit of $r < 0.06$ (95% C.L.) is achieved. Inflationary theories predict that the scalar spectral index and tensor-to-scalar ratio are both determined by the equation of state during inflation, and




Figure 1.2: Marginalized 68 % and 95 % constraints on the n_s - r plane from a combination of 2018 *Planck* BICEP / Keck, and baryon acoustic oscillation data, presented in Ref. [5].

therefore of the inflationary potential, and so are not independent. Constraints on specific inflationary models can then be plotted in the $n_s - r$ plane, as shown in Figure 1.2.

1.2.3 Alternative Models of Inflation

So far we have concentrated on the simplest case of inflation, in which a single scalar field drives the inflationary expansion. However, similar models with more complicated matter content remain possible. For example, in Chapter 3 we will study a model in which an additional axion field and set of SU(2) fields are present during inflation [42].

Whilst the overall expansion will be similar in both cases, the predicted observable effects can be quite different, for example by producing a chiral gravitational wave background [43, 44], or sourcing tensor perturbations from matter fluctuations. Making correct inferences about the inflationary epoch from a detection of non-zero inflationary B -modes will therefore require a careful accounting of these models.

1.3 Cosmic Microwave Background

By the time the Universe was roughly 300,000 years old, it had cooled to roughly 3000 K, and free electrons and protons formed neutral hydrogen. After this transition, the mean free path for photons became larger than the Hubble radius, and photons were able to free stream. As the Universe has evolved over the past billions of years, these photons have been redshifted, and now form an almost uniform background, with a blackbody spectral energy dependence characterized by a temperature of $T_{\text{CMB}} = (2.72548 \pm 0.00057) \text{ K}$ [30]. This background was first detected in 1963 as excess noise in the 20-foot horn-reflector antenna at the Crawford Hill Laboratory in New Jersey by Arno Penzias and Robert Wilson [45]. The *COBE* satellite established the blackbody spectrum of the CMB, and detected deviations from the mean temperature of order 10^{-5} [46]. Since *COBE*, precision cosmology has been largely based on ever-more-sensitive measurements of these tiny anisotropies, and their polarization, by satellite missions such as *WMAP* [47] and *Planck* [8], and ground-based telescopes such as the POLARBEAR [48], SPTPol [49], ACTPol [50], and BICEP2 / Keck Array [51].

In this section we introduce CMB phenomenology, and establish the link between the perturbations of inflation and CMB anisotropies.

1.3.1 The anisotropic CMB

The fundamental observable of the CMB is its temperature, T , in direction $\hat{n} = (\theta, \phi)$, where θ , ϕ are the polar and azimuthal angles in a spherical coordinate system. We have established already that the CMB is incredibly uniform, with spatial variations of only 10^{-5} . Therefore, we will refer to the deviations from the average temperature,

$$\Theta(\hat{n}) \equiv \frac{T(\hat{n}) - \bar{T}}{\bar{T}},$$

as the primary observable. The CMB is a two-dimensional image on the sphere, and its natural basis is the spherical harmonics:

$$\Theta(\hat{n}) = \sum_{\ell m} Y_{\ell m}(\hat{n}) \Theta_{\ell m}, \quad \Theta_{\ell m} = \int d\Omega Y_{\ell m}^*(\hat{n}) \Theta(\hat{n}), \quad (1.42)$$

where the spherical harmonics, $Y_{\ell m}(\hat{n})$ are given by:

$$Y_{\ell m}^*(\hat{n}) \equiv \sqrt{\frac{2\ell + 1}{4\pi} \frac{(\ell - m)!}{(\ell + m)!}} P_{\ell m}(\cos \theta) e^{im\phi},$$

and $P_{\ell m}(\cos \theta)$ are the Legendre polynomials, for which $m \in \{-\ell.. \ell\}$. The multipole, ℓ , acts as the equivalent to the wavenumber, k , for a spherical harmonic decomposition. It is conjugate to the angular size, θ , on the sky: $\ell \sim 1/\theta$. We expect the CMB anisotropies to be Gaussian-distributed, and so all the physical information will be contained in the two-point statistics of the fluctuations. In harmonic space, this is the power spectrum:

$$C_{\ell}^{\Theta\Theta} = \langle \Theta_{\ell m} \Theta_{\ell' m'}^* \rangle \delta_{\ell\ell'} \delta_{mm'}, \quad (1.43)$$

where the angle brackets indicate an ensemble average across realizations of the CMB. For a given sky realization, this can be estimated by:

$$\hat{C}_{\ell}^{\Theta\Theta} = \sum_{m=-\ell}^{\ell} \frac{|\Theta_{\ell m}|^2}{2\ell + 1}. \quad (1.44)$$

1.3.2 CMB anisotropies from scalar perturbations

Before recombination occurred, the Universe was composed of a photon-baryon fluid, tightly-coupled by Compton scattering, and decoupled dark matter. Perturbations in the dark matter density are free to evolve independently from the radiation, and begin to grow, forming gravitational potential wells. The dynamics of the photon-baryon fluid is determined by a balance between these potential wells, and support from photon pressure. Overdensities in the photon-baryon fluid begin to collapse, driven by the gravitational potential. As the fluid is compressed, the photon pressure acts as a restoring force, establishing the motion of a *driven harmonic oscillator* [52]. The overall size of the fluctuations is set by the amplitude of the initial density perturbations.

For modes well within the horizon, and in the tight coupling limit, we can write the temperature perturbation in the conformal Newtonian gauge as [53, 54, 55]:

$$\Theta \propto \underbrace{\frac{\Delta_{\gamma}}{4} + \Psi}_{\text{Sachs-Wolfe}} + \underbrace{\hat{n} \cdot (\mathbf{v}_{\text{obs}} - \mathbf{v}_b)}_{\text{Doppler}} + \underbrace{\int_{\eta}^{\eta_0} d\eta (\Psi' + \Phi')}_{\text{ISW}} \quad (1.45)$$

where we have used the Bardeen potentials introduced in the last Section, \mathbf{v}_{obs} is the velocity of the observer relative to the CMB, and \mathbf{v}_b is the bulk baryon velocity.

Since the perturbations are small, there is no coupling between modes, and different wavelength perturbations can evolve separately on the same background. The phenomenology of different perturbation modes can be separated based on their wavelength. Modes outside of the horizon at last scattering, $\lambda \gg \eta_*$, will not have had time to collapse, and will be unaffected by the physics governing the photon-baryon fluid. There are, however, purely general relativistic effects that process that largest CMB modes. As the photons propagate over time-varying potentials, they will be red or blue shifted [56]. This so-called *Sachs-Wolfe* effect accumulates over the photon's path from the surface of last scattering until now, and receives a few major contributions. Firstly, the first two terms of Equation 1.45 cancel in the radiation-dominated era to give $\frac{\Delta\gamma}{4} + \Psi = \frac{\Phi}{3}$. This represents the shift required to escape the potential of the curvature perturbation at last scattering. Secondly, the *integrated Sachs-Wolfe* effect is represented by the integral in Equation 1.45 over the photon path between last scattering and observation. This integral receives two major contributions: i) at early times during the tail-end of the radiation -dominated era, ii) after the transition to cosmological constant domination [57].

Modes within the particle horizon, $\lambda \ll \eta_*$, at the surface of last scattering will have had time to undergo oscillations, and will be seen at different stages of their oscillations, depending on their wavelength. Modes at maximum compression or rarefaction correspond to maxima in the power spectrum of the CMB, whilst modes at velocity maxima (the Doppler term in Equation 1.45) correspond to troughs.

For perturbation wavelengths smaller than the photon mean free path at last scattering, the tight coupling approximation breaks down, and *Silk damping* of the power spectrum as a result of photon diffusion must be taken into account [58]. Scales similar to, and smaller than, the Silk radius, η_{Silk} will be affected. At recombination the electrons are non-relativistic, and have a mass much larger than the average photon energy (the ionization energy of hydrogen is ~ 13 eV c.f. the electron rest mass of $m_e = 511$ MeV), therefore Thomson scattering is the appropriate description of the electron-photon interactions. Assuming a random walk for an individual photon, and given the electron number density, n_e , and cross-section to Thomson scattering, σ_T , one can estimate the Silk damping length

$$\eta_{\text{Silk}} \approx 8 \text{ Mpc},$$

which corresponds to $\ell \gtrsim 1000$ on the CMB.

In order to solve for the full evolution of linear perturbations on an FLRW background to better than the 1% level [59] one must solve the equations numerically [37, 6]. In practice this is done by expanding the relevant temperature and metric perturbations in their Legendre coefficients, as an infinite set of coupled Boltzmann equations [60, 61, 62, 63]. This can be done with one of several public and well-tested codes, such as **CMBFAST** [64], **CAMB** [65], and **CLASS** [66]. **CAMB** and **CLASS** currently achieve an accuracy of 0.01% [67].

1.3.3 Polarization

Thomson scattering is an intrinsically polarized process, the cross section of scattering is a function of both incoming and outgoing polarization states. If the incoming radiation is isotropic, the scattered radiation will be unpolarized. However, in the presence of a quadrupolar radiation field, the scattered radiation will be linearly polarized [68]. Before recombination, Thomson scattering was very efficient and dissipated quadrupole anisotropies [37]. During recombination, the tight-coupling of the photon-baryon fluid breaks down, and scattering electrons begin to see a quadrupole moment in the radiation field, driven by the velocity gradients in the surrounding baryon fluid [69]. The primary CMB is expected to be polarized at the level of about 10% [68, 70, 71], and was first detected by DASI at the South Pole in 2002 [72], and has since been measured by the *WMAP* and *Planck* satellites, as well as a host of ground-based observatories.

The linearly polarized CMB forms a spin-2 field on the sphere, requiring two numbers to describe for a given (θ, ϕ) . There are a number of useful bases providing descriptions of linear polarization. The Stokes parameters $(Q(\hat{n}), U(\hat{n}))$ are often used as they are close to the quantities measured in observations, however, they depend on the instrument used in the observation. A more natural approach is to decompose the vector field into a curl-free $(E(\hat{n}))$ and divergence-free $(B(\hat{n}))$ contribution [19, 73, 20]:

$$(Q \pm iU)(\hat{n}) = \sum_{\ell m} (a_{\ell m}^E \pm ia_{\ell m}^B)_{\pm 2} Y_{\ell m}(\hat{n}), \quad (1.46)$$

where $_{\pm 2} Y_{\ell m}(\hat{\mathbf{n}})$ are the spin-weighted spherical harmonics, and $a_{\ell m}^B, a_{\ell m}^E$ are the




Figure 1.3: Figure taken from Ref. [6] showing various contributions to the temperature power spectrum. The rise at multipoles $\ell < 10$ is due to the late time integrated Sachs-Wolfe effect due to the cosmological constant. The Sachs-Wolfe Plateau is an almost scale-invariant interval of perturbations corresponding to the initial curvature perturbations on scales larger than the η_* . The acoustic peaks correspond to the interval $\eta_* < \lambda < \eta_{\text{Silk}}$ in the baryon-photon fluid behaves as a harmonic oscillator supported by photon pressure, under gravitational forcing. At scales $\ell \gtrsim 1000$ Silk damping due to photon diffusion results in exponential suppression of the acoustic peaks.

spherical harmonic coefficients of the (pseudo)scalar field (B)E:

$$E(\hat{n}) = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^E Y_{\ell m}(\hat{n}) \quad B(\hat{n}) = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^B Y_{\ell m}(\hat{n}) \quad (1.47)$$

It is often simpler to interpret observations in the (E, B) basis, because observations by different instruments can be compared, and certain physical processes produce characteristic signatures in the E and B fields. As for the temperature case, we expect the primordial polarized anisotropies to be fully specified by their two-point statistics, but in this case between all three fields, T, E, B :

$$\langle a_{\ell m}^X a_{\ell' m'}^{X'} \rangle = \delta_{\ell\ell'} \delta_{mm'} C_{\ell}^{XX'} \quad (1.48)$$

where $X, X' = T, E, B$.

1.3.3.1 From scalar perturbations

The conditions for polarization existed only for a short time during recombination, and only for scales below the photon mean free path [74]. Modes much larger than the mean free path remain unpolarized (*there is no Sachs-Wolfe effect for polarization*), and scales much smaller than the mean free path are exponentially suppressed by Silk damping. A full calculation of the scale dependence requires a detailed computation [75], but the result is that the polarization signal peaks around the scale of the mean free path, $100 < \ell < 1000$. Polarized anisotropies are generated by the same acoustic waves as the temperature anisotropies, however, since they trace the velocity of the baryon fluid, which is exactly $\pi/2$ out of phase with the density perturbations, the peaks in the polarized anisotropies occur out of phase with temperature.

Around redshift $z \approx 7-8$, the first stars formed, and began emitting UV light that *reionized* the Universe, populating it with free electrons once again. CMB photons were scattered by the free electrons, imposing an optical depth to the surface of last scattering. The free electrons at reionization were sensitive to the CMB quadrupole at $z \sim 7-8$, and therefore the scattered light was polarized [74].

1.3.3.2 From tensor perturbations

The great advantage of the Helmholtz decomposition is that B -mode polarization is only produced by tensor perturbations to the metric [19, 20]. The B mode power spectrum in the presence of tensor perturbations is given by [19]:

$$C_\ell^{BB} \propto \int d \ln k \mathcal{P}_h(k) \left[\Delta_\ell^{(t)}(k) \right]^2, \quad (1.49)$$

where $\Delta_\ell^{(t)}(k)$ is the transfer function describing the evolution of the k^{th} mode of the tensor perturbation, and its projection onto the multipole ℓ . Tensor perturbations make similar contributions to the TT and EE tensor power spectra [74]

We can provide a qualitative description of the features of the tensor BB power spectrum, which is shown in Figure 1.5. As the Universe expands gravitational radiation redshifts in a similar way to electromagnetic radiation, and the amplitude of tensor modes within the horizon decays with the inverse of the scale factor [76, 77]. The peak at $\ell \sim 100$ corresponds to modes that enter the horizon at recombination, whilst perturbations larger than the horizon at recombination do not contribute and are suppressed. On small scales ($\ell > 100$) modes that have already reentered the horizon by recombination, and so have been redshifted by expansion [77, 78].

The *reionization* bump at $\ell \lesssim 10$ corresponds to polarization generated during reionization when free electrons scatter the background CMB at $z \sim 8$ [74].

The overall amplitude of the BB spectrum is directly proportional to the tensor power spectrum, and therefore to the tensor-to-scalar ratio, r . A detection of B -mode polarization of the CMB would therefore directly probe the inflationary potential (Equation 1.39), providing direct constraints on GUT scale physics.

1.3.3.3 From gravitational lensing

In the linear perturbation theory discussed so far, we do not expect scalar metric perturbations to contribute to B -mode polarization. The most significant higher-order effect that must be considered is that of gravitational lensing [60, 79, 80]. As photons travel from the surface of last scattering to an observer today their trajectory may be perturbed by gravitational interactions with the intervening large scale structure. The effect on the observed CMB temperature is a remapping: $T(\hat{n}) \rightarrow T(\hat{n} + \delta\hat{n})$, and the effect on the polarization is a similar remapping whilst also parallel transporting the polarization vector along the geodesic connecting the points \hat{n} and $\hat{n} + \delta\hat{n}$.

In the so called *Born approximation*, lensing effects can be summarized by the gravitational potential evaluated along the null geodesic of the unperturbed photon [81, 82]. In the absence of anisotropic stress and curvature, the lensing potential, ϕ , is given by the integral [80]:

$$\phi(\hat{n}) = -2 \int_0^{\chi_*} d\chi \frac{\chi_* - \chi}{\chi_* \chi} \Psi(\eta\hat{n}; \eta_0 - \chi) \quad (1.50)$$

where χ_* is the conformal distance to the source, and $\eta_0 - \chi$ is the conformal time when the photon was at position $\chi\hat{n}$, and Ψ is the familiar scalar perturbation in conformal Newtonian gauge. The lensing deflection field, $\delta\hat{n}$, is then given by the covariant derivative of the lensing potential, $\delta\hat{n} = \nabla\phi$.

The contribution of lensing to the BB power spectrum can then be evaluated by expanding the remapped polarization map, $(Q \pm iU)(\hat{n} + \delta\hat{n})$, to linear order in the deflection field [83]:

$$C_\ell^{BB,\text{Lens}} = \frac{1}{2} \sum_{\ell_1, \ell_2} \frac{({}_2F_{\ell\ell_1\ell_2})^2}{2\ell + 1} C_{\ell_1}^{\phi\phi} [C_{\ell_2}^{EE} + C_{\ell_2}^{BB} - (-1)^L (C_{\ell_2}^{EE} - C_{\ell_2}^{BB})] \quad (1.51)$$

where $L = \ell + \ell_1 + \ell_2$, and the function ${}_2F_{\ell\ell_1\ell_2}$ can be found in Ref. [83]. So the effect of lensing is to convolve a combination of EE and BB spectra with the $\phi\phi$ lensing spectrum, and add this to the unlensed BB spectrum. In Λ CDM the EE spectrum is expected to dominate over BB by an order of magnitude, and the lensing contribution to BB dominates over the primordial contribution for values of r allowed by current observations [84, 41]. In this regime lensing forms a noise floor of $\sim 5 \mu\text{K a}_{\text{min}}$, due to the increased cosmic variance in the observed C_ℓ^{BB} . A theoretical calculation of the lensing spectrum is shown in Figure 1.5, compared to theoretical primordial spectra for a range of r values.

The instrumental noise in modern CMB experiments is low enough (of order $5 \mu\text{K a}_{\text{min}}$ or less) that lensing is a dominant uncertainty in the constraints on r [85, 84] (see Figure 1.5). In order to reduce lensing uncertainty, two broad classes of methods have been developed to *delens* B -mode observations:

External delensing methods estimate the lensing potential in Equation 1.50 using a tracer of the underlying matter field, such as galaxy surveys [86, 87], and the CIB [88, 89, 90]. Recently, the CIB has been shown to be correlated with CMB lensing [91, 92, 93], and has been used to delens B mode observations.

Multi-tracer probes of the CIB, using *Planck* and WISE data, have removed up to 40 to 50 % of the lensing power [90, 94] .

Internal delensing methods estimate $C_\ell^{\phi\phi}$ from within the CMB data set itself using higher-order correlators of both the temperature and polarization fields [95, 96, 97, 98]. This has been performed on the *Planck* data [99], achieving a 7% reduction in lensing power.

1.3.4 CMB observations

In this section we provide an overview of the results from past and ongoing CMB observatories, and outline the design and goals of proposed future observatories in various stages of development.

1.3.4.1 Past and current observations

Here we list some key CMB missions, and describe the science goals they accomplished. A comparison of the power spectrum measurements made by many of these experiments is made in Figure 1.4.

COBE: The Cosmic Background Explorer (*COBE*) made observations from 1989 to 1993. As was discussed in the introduction to this Section, *COBE* established that the CMB behaves as an almost-perfect blackbody, with tiny deviations in its temperature of order 10^{-5} [46]. The study of the structure of these deviations has formed the basis of modern precision cosmology.

WMAP: The Wilkinson Microwave Anisotropy Probe (*WMAP*) made observations in five bands from 2001 to 2010. Its results were key in establishing the success of the Λ CDM model of a flat Universe with a cosmological constant [100], and made the first detection of a cosmic neutrino background [101].

Planck: The *Planck* satellite made observations in nine bands from 2009 to 2013, with greater resolution and sensitivity than *WMAP*. The wide frequency coverage allowed *Planck* to better remove various astrophysical foregrounds [9], and its greater sensitivity and resolution allowed the measurement of smaller scale anisotropies, and consequently even tighter constraints on the key Λ CDM parameters [5]. *Planck* reported the first detection of the correlation between gravitational lensing and the cosmic infrared background [102].




Figure 1.4: Figure taken from Ref. [7] comparing constraints from different experiments on: TT , EE , and BB in the top panel; TE in the middle panel; and $\phi\phi$ in the bottom panel. from various experiments, many of which are described in Section 1.3.4.1. The dashed line is the best-fit model from the *Planck* results [8].

Atacama Cosmology Telescope: The Atacama Cosmology Telescope (*ACT*) is a six-meter telescope on Cerro Toco, in the Atacama Desert, Chile. *ACT* made its first observations in 2007, and has undergone several upgrades to its detectors in the intervening years: *ACTPol*, and *AdvACT*. *ACT* made the first detection of the power spectrum of gravitational lensing [103, 104], and when combined with *WMAP* data, provided the first detection of dark energy from the CMB alone [105], as well as the first detection of the correlation between CMB lensing and galaxy lensing [106], detection of lensing of the CMB by galaxy clusters [107], and the first detection of the kinematic Sunyaev-Zel'dovich effect [108].

BICEP and Keck Array: The BICEP and Keck Array (*BK*) are a series of experiments (BICEP1, BICEP2, BICEP3, Keck Array, and BICEP Array) based at the South Pole. In its various forms, *BK* has been observing in three bands (95 GHz, 100 GHz, and 150 GHz) since 2006, with the latest incarnation, BICEP Array, installed in 2017/2018. *BK* has focused on making high-sensitivity measurements of the B-mode polarization of the CMB, in a small patch of sky. Combined with the *Planck* and *WMAP* data, *BK* provides the best constraints on the tensor-to-scalar ratio [41]: $r < 0.06$ (95% C.L.).

South Pole Telescope: The South Pole Telescope (*SPT*) is a 10-meter telescope situated at the South Pole [109, 110]. *SPT* has consisted of three stages: *SPT-SZ*, *SPTpol*, and the current *SPT-3G*. *SPT-SZ* focused on a 2500 square degree survey to detect clusters of galaxies at various redshifts using the Sunyaev-Zel'dovich effect [111, 112, 113], but also allowed the detection of CMB lensing by galaxy clusters [114]. *SPTpol* upgraded the receivers with polarization-sensitive detectors, and focused on high-resolution measurements of the B-mode polarization of the CMB in a 500 square degree region of sky. This survey resulted in the first detection of CMB B-modes induced by gravitational lensing, and demonstrated that lensing B-modes could be predicted by combining the primary E-modes with an external tracer of the projected mass density [91, 115]. The *SPT-3G* survey will observe a larger 1500 square degree region, and combine its analysis with the overlapping BICEP Array observations. The combined frequency coverage of BICEP Array, and the high-resolution of *SPT-3G* will allow a joint foreground removal and delensing of primordial B-mode observations.

1.3.4.2 Future observations

Simons Observatory: The Simons Observatory (*SO*) will combine the efforts of the existing Chilean telescopes in the Atacama desert, to build a set of three small aperture telescopes, similar in size to the BICEP3 telescope, and a large aperture telescope, similar in size to *ACT*. *SO* seeks to make almost cosmic-variance limited observations of polarized anisotropies down to the arcminute scale. There will also be a large-aperture telescope, making high-resolution observations of the CMB with the aim of constraining dark energy and neutrino properties, and delensing capabilities for the large-scale survey [116, 3]. Construction for *SO* is planned to complete in 2020.

LiteBIRD: *LiteBIRD* is a proposed satellite mission lead by the Japanese space agency, JAXA [117, 118, 119]. If selected, it would launch in the 2020s to observe the large-scale polarized CMB at fifteen frequencies between 34 GHz and 448 GHz. It aims to achieve as broad frequency coverage as possible in order to remove Galactic foregrounds, which will be the leading barrier to its sensitivity to inflationary B-modes.

1.4 CMB Foregrounds

The information contained in the CMB may not be accessible if it is contaminated by radiation from sources within our own Milky Way. A variety of processes in the interstellar medium produce microwave radiation in the *foreground* of CMB observations. The confusion of such sources with genuine CMB anisotropies will lead to incorrect cosmological parameters inferences.

Observations have established that the CMB temperature anisotropies dominate over Galactic foregrounds at high Galactic latitudes, and so are quite robust to the effects of Galactic foregrounds [120]. However, foregrounds are significantly polarized by interactions with the Galactic magnetic field. Current observations indicate that the CMB *B* mode signal is subdominant to Galactic foregrounds at all frequencies and in all parts of the sky (for example see Figure 1.5) [9, 121, 122, 123, 41]. Understanding both the spectral behaviour, as well as the morphology of Galactic foregrounds, will be critical in designing upcoming experiments such as Simons Observatory [3], CMB-S4 [124], and LiteBIRD [125], as well as in the analysis of current data [41, 121].

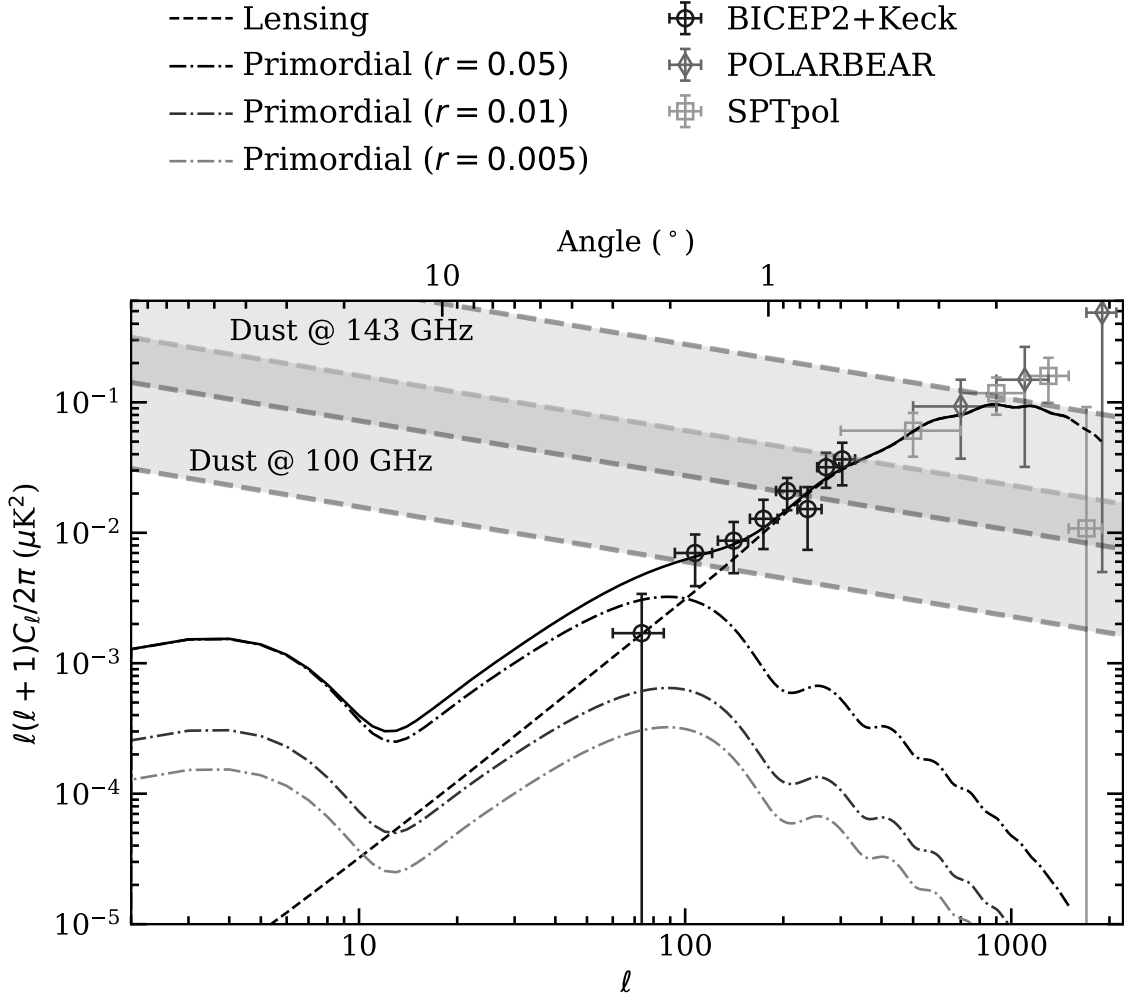


Figure 1.5: Figure comparing the theoretical B -mode power spectra for lensing and primordial contributions, with constraints from different observatories, and observed foreground levels from the Planck experiment (the range corresponds to observations covering 24% to 71% of the sky). The primordial B mode spectrum generated at the surface of last scattering, for several different values of r is shown as the dot-dash line. The lensing contribution is shown as the dashed black line, and the sum of lensing and $r = 0.01$ is shown in solid black. Constraints from BICEP2+Keck (circle), POLARBEAR (diamond), and SPTpol (square) are overplotted.

In this section we will review the physical basis of each of the important foreground sources, the constraints from current observations, and the foreground modeling currently available.

1.4.1 Thermal dust emission

At frequencies above ~ 70 GHz the dominant foreground is thermal emission from warm *dust* grains, composed of graphites, silicates, and polycyclic aromatic hydrocarbons (PAHs) [126]. Grains, absorb light from the interstellar radiation field, and emit it in the infrared and microwave [127]. The process of emission is very nearly black body, and assuming the dust clouds to be optically thin, the spectral radiance of a dust cloud along a given line of sight may be written as a function of its temperature and emissivity:

$$I_\nu \propto (1 - \exp(-\tau_\nu))B_\nu(T_d) \approx \nu^{\beta_d} B_\nu(T_d) \quad (1.52)$$

where T_d is the temperature of the dust population, expected to be $T_d \sim 21$ K, and β_d is the power law index describing the frequency-dependence of the dust emissivity, and is expected to be in the range $\beta_d \sim 1 - 2$ [128]. Populations of dust exist in many different astrophysical environments, and with varying chemical compositions [129]. Therefore, the temperature and emissivity is expected to vary spatially. Observations by the *Planck* satellite have confirmed that both the temperature and emissivity of dust within the Milky Way are spatially variable at the level of $\sim 10\%$ [9].

Polarization of starlight due to preferential extinction by aligned interstellar grains was discovered in the mid half of the previous century [130, 131], and was quickly attributed to interactions between aspherical dust grains and the ambient Galactic magnetic field [132, 133]. There have been many competing theories for the mechanism of this alignment, see for example the review of Ref. [134], and the current favored mechanism is that of radiative alignment torque, based on the early work of Ref. [135] and later developed by Refs. [136, 137].

Dust grains are aspherical, with a spin-axis perpendicular to their longest axis. As the spin is, on average, aligned with the Galactic magnetic field, the thermal emission of dust grains is linearly polarized perpendicular to the Galactic magnetic field projected on the sky [138]. The polarized thermal emission of dust grains is a significant contaminant of observations of the polarized CMB [139, 41, 129].

Observations of the Galaxy are sensitive only to the emission integrated along the line of sight through the Galaxy, and averaged over the experimental beam. Multiple

populations of dust may exist along a given line of sight, or below the resolution limit of the instrument, leading to blending of distinct sources. Even if each source acts as a power law, their sum will not, and Equation 1.52 will no longer describe the frequency dependence of dust emission [140, 141, 129]. If the *decorrelation* effect is not considered, it can bias constraints on the tensor -to-scalar ratio [140, 129]. A detection of decorrelation in the *Planck* data was claimed in Ref. [142], however a careful reanalysis of the data, including additional systematic effects, found that the data was consistent with no decorrelation given the noise level [143]. Modeling decorrelation [144, 145], and developing techniques to mitigate its effects [146, 147], is an ongoing area of research.

1.4.2 Synchrotron emission

Synchrotron radiation is the dominant radiation mechanism in polarisation at frequencies $\lesssim 50$ GHz [148]. Fast moving electrons in the Galactic magnetic field experience an acceleration perpendicular to their velocity, producing a trajectory spiraling in a plane perpendicular to the magnetic field. The emitted flux depends both on the square of the ambient magnetic field, and on the energy distribution of the injected cosmic rays, $N(E) \propto E^{-p}$, and can be approximated as:

$$I_\nu \propto B^{\frac{p+1}{2}} \nu^{\beta+C \log(\nu)}, \quad (1.53)$$

where $\beta = -\frac{(p+3)}{2}$ [149], and the logarithmic term accounts for the steepening of the spectral index at higher frequencies due to aging sources [148]. The spectral index, β , depends on the spectrum of the energy spectrum of the underlying electrons through p , and so will have some spatial variability.

The spiral motion of synchrotron electrons is confined to a two-dimensional plane, and is intrinsically polarized, achieving polarization fractions of up to 40 % [139], however, the observed polarization fraction will depend on frequency due to the depolarizing effects of Faraday rotation.

1.4.3 Anomalous microwave emission

In the late nineties, observations made by the DIRBE instrument on the *COBE* satellite, probing Galactic dust in the infrared, were combined with low-frequency observations made by the Saskatoon experiment [150], Owens Valley Radio Observatory [151], and

other radio surveys [152], suggested the existence of a Galactic component correlated with far-infrared dust emission, peaking at frequencies in the range 10 GHz - 60 GHz, and uncorrelated with synchrotron or free-free emission. A natural explanation for the so-called *anomalous microwave emission* (AME) is rotational emission by ultra small, rapidly-spinning, dust grains [153, 154]. However, there may also be a contribution due to thermal fluctuations in the magnetization of dust grains [155].

AME is not expected to be strongly polarized, and observations of compact sources have established upper limits on the polarization fraction of $\sim 1\%$ [156, 157].

1.4.4 Free-free

Scattering interactions between electrons and charged particles in the interstellar medium produce free-free emission, also known as thermal bremsstrahlung [149]. This radiation is characterized by a power-law with spectral index very close to 2, and is expected to be unpolarized due to the random orientations of the scattering process [158, 159].

1.5 Statistical Methods

In this Section we introduce a few of the statistical methods upon which we will call in future Sections. We will review: *Bayesian statistics* and *sampling* in Section 1.5.1; power spectrum estimation and *purification* in Section 1.5.3; and the Fisher forecast methodology in Section 1.5.2.

1.5.1 Bayesian statistics

Bayesian statistics is an interpretation of probability theory, usually stated in contrast to the *frequentist* interpretation, that uses prior knowledge of a system to predict its future statistical behavior. Bayes' theorem allows one to compute the probability of future events, combining any measured data with prior information, and to update this estimate with new information as it becomes available.

1.5.1.1 Parameter inference

A common problem in physics and astrophysics is the determination of model parameters from a data set. Given the data \mathbf{d} , the probability that the underlying model




Figure 1.6: Figure taken from Ref. [9], showing the RMS variation in the polarized foregrounds as a function of frequency. Synchrotron (green) dominates at low frequencies, and falls off as $\sim \nu^{-3}$, dust (red) dominates at high frequencies, and within the CMB frequencies of interest has a power law dependence of $\sim \nu^{1.5}$. The CMB emission (teal) is roughly constant below $\nu \sim 100$ GHz, and falls off above this limit. The width of each of the lines represents the range of foreground signal for f_{sky} between 27% and 83%. The gray bands show the bands in which the Planck experiment made observations [9]. Note that there is not sky fraction, and no frequency at which the CMB dominates over foreground emission.

parameters, represented by the n -dimensional vector θ , is given by Bayes' Theorem [160]:

$$P(\boldsymbol{\theta}|\mathbf{d}) = P(\mathbf{d}|\boldsymbol{\theta})\frac{P(\boldsymbol{\theta})}{P(\mathbf{d})} \quad (1.54)$$

where $P(\boldsymbol{\theta}|\mathbf{d})$ is the *posterior* probability for the model parameters $\boldsymbol{\theta}$, $P(\mathbf{d}|\boldsymbol{\theta})$ is the *likelihood* of measuring the data \mathbf{d} given a set of underlying model parameters $\boldsymbol{\theta}$, and we will later work with the log-likelihood, $\mathcal{L}(\boldsymbol{\theta}) \equiv -\ln P(\boldsymbol{\theta}|\mathbf{d})$. $P(\boldsymbol{\theta})$ is the *prior* knowledge one has about the model parameters, which could be due to a previous measurement made by a different experiment, or due to physical intuition (e.g. imposing the requirement that the tensor-to-scalar ratio be positive: $r > 0$). $P(\mathbf{d})$ is the prior probability of obtaining the data \mathbf{d} .

The posterior probability for a single parameter, θ_i can be obtained from the joint posterior by *marginalizing* over the other variables:

$$P(\theta_i|\mathbf{d}) = \left(\prod_{j \neq i} \int d\theta_j \right) P(\boldsymbol{\theta}|\mathbf{d})P(\mathbf{d}). \quad (1.55)$$

‘Fitting’ the model parameters requires that we:

- Determine the most probable value of $\boldsymbol{\theta}$: $\boldsymbol{\theta}_{\text{ML}} = \arg \max [P(\boldsymbol{\theta}|\mathbf{d})]$, where the subscript ML indicates ‘maximum likelihood’.
- Characterize the uncertainty in $\boldsymbol{\theta}_{\text{ML}}$, by exploring the parameter space close to $\boldsymbol{\theta}_{\text{ML}}$.

In special cases this can be done analytically, however, it will usually require a numerical computation of the likelihood in Equation 1.54 over a suitable subset of the parameter space. The computational requirement of evaluating Equation 1.54 for many values of $\boldsymbol{\theta}$ is in large part what made Bayes’ theorem an intractable approach until the revolution in computational resources over the last few decades.

1.5.1.2 Markov Chain Monte Carlo

A naive strategy to fitting $\boldsymbol{\theta}$ could be to compute the posterior on an n -dimensional grid, and infer $\boldsymbol{\theta}_{\text{ML}}$ and its uncertainty from these samples. However, this approach scales very poorly as n is increased. Different strategies have been developed to efficiently sample the posterior, with the most successful group Markov Chain Monte Carlo (MCMC) techniques. The idea behind MCMC is to form a Markov Chain

from samples from $P(\boldsymbol{\theta}|\mathbf{d})$, and demonstrate that its equilibrium state is given by the underlying probability distribution. The details of the MCMC approach are beyond the scope of this thesis, but can be found in Ref. [161]. A few important MCMC algorithms are:

- *Metropolis-Hastings (MH)* is one of the most widely-used algorithms. It uses a proposal density and rejection criterion to determine $\boldsymbol{\theta}_{i+1}$ from the ratio of the likelihoods at the points $\boldsymbol{\theta}_i$ and $\boldsymbol{\theta}_{i+1}$ [162, 163, 164, 165].
- *Gibbs sampling* is a special case of Metropolis-Hastings in which samples can be drawn directly from the conditional distributions of some parameters, allowing a much more efficient exploration of the parameter space [166].
- *Affine-invariant sampling* [167] is an improvement over the standard MH algorithm as it does away with the long ‘burn-in phase’ [168] MH algorithms tend to require. It significantly outperforms naive MH implementations, and requires the setting of only a couple of hyperparameters. In the rest of this thesis, unless otherwise stated, we perform sampling of posterior distributions using the `emcee`² package [168], a Python implementation of this technique.

1.5.2 Fisher forecasts

A common task within cosmology and astronomy is determining the design of an instrument, and its survey parameters. For example, one may have to determine the balance between depth and coverage of a survey, and how to distribute detector effort across different frequency bands. Fisher forecasts are a fast way to narrow down a large experiment parameter space. The Fisher methodology has been used widely within physics to optimize surveys [169, 170, 171, 172, 173, 174, 124], or to assess the synergy of combined probes [169, 175]. Fisher forecasts are a useful tool to narrow down large parameter spaces that would be otherwise computationally cumbersome. However, they neglect many real-world considerations, making the results usually indicative and not wholly correct [176].

For a data set \mathbf{d} , and parameters of a predictive model, $\boldsymbol{\theta}$, the *Fisher information* matrix is defined as [177]:

$$\mathcal{F}_{ij} \equiv \left\langle \frac{\partial^2 \mathcal{L}(\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} \right\rangle \bigg|_{\boldsymbol{\theta}_{\text{ML}}}, \quad (1.56)$$

²<https://github.com/dfm/emcee>

where the angle brackets indicates an ensemble average of the data, evaluated at the maximum-likelihood point. It can be shown that [178, 179]:

- For any unbiased estimator, the Cramér-Rao bound states that $\sigma(\theta_i) \geq \sqrt{(F^{-1})_{ii}}$.
- If there is a best unbiased estimator, then it is $\boldsymbol{\theta}_{\text{ML}}$, or a function of it.
- $\boldsymbol{\theta}_{\text{ML}}$ is asymptotically the best unbiased estimator.

The close relationship between the maximum-likelihood approach and the Fisher information can be elucidated by expanding the log likelihood around $\boldsymbol{\theta}_{\text{ML}}$. By definition, the first derivative is zero, and so the likelihood is:

$$P(\mathbf{d}|\boldsymbol{\theta}) = \exp\left(\mathcal{L}(\boldsymbol{\theta}_{\text{ML}}) + (\boldsymbol{\theta} - \boldsymbol{\theta}_{\text{ML}})^T \mathcal{F}(\boldsymbol{\theta}_{\text{ML}})(\boldsymbol{\theta} - \boldsymbol{\theta}_{\text{ML}}) + \dots\right). \quad (1.57)$$

Written in this way, we see that the posterior in Equation 1.56 is approximately Gaussian in the model parameters at the maximum likelihood point, and the parameter covariance is given by the Hessian evaluated at $C_{ij} = (\mathcal{F}^{-1}(\boldsymbol{\theta}_{\text{ML}}))_{ij}$. The approximation remains good for all $\boldsymbol{\theta}$ if the likelihood drops off rapidly away from $\boldsymbol{\theta}_{\text{ML}}$.

In the case of the CMB, the data usually constitute a set of temperature and polarization power spectra. The Fisher matrix can be written as [172, 179]:

$$\mathcal{F}_{ij} = \sum_{\ell} \frac{2\ell+1}{2} f_{\text{sky}} \text{Tr} \left(\mathbf{C}_{\ell}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}_{\ell}}{\partial \theta_i} \mathbf{C}_{\ell}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{C}_{\ell}}{\partial \theta_j} \right), \quad (1.58)$$

where:

$$\mathbf{C}_{\ell} \equiv \begin{pmatrix} C_{\ell}^{TT} + N_{\ell}^{TT} & C_{\ell}^{TE} & C_{\ell}^{T\phi} & 0 \\ C_{\ell}^{TE} & C_{\ell}^{EE} + N_{\ell}^{EE} & 0 & 0 \\ C_{\ell}^{T\phi} & 0 & C_{\ell}^{\phi\phi} + N_{\ell}^{\phi\phi} & 0 \\ 0 & 0 & 0 & C_{\ell}^{BB} + N_{\ell}^{BB} \end{pmatrix} \quad (1.59)$$

In Section 3 we will apply this formulation to study the constraints future CMB satellite missions will place on inflationary physics.

1.5.3 Power spectrum estimation

The workhorse of cosmological inference is the two-point function, in the form of two-point correlations or power spectra, from which most cosmological information is extracted. In this thesis we will frequently estimate power spectra from observed or simulated maps, and in this section we briefly review the pseudo- C_ℓ method used.

Whilst it is possible to write down optimal estimators for the power spectrum, based on maximum likelihood techniques [180, 181], or on quadratic estimators [182, 183], these methods scale as ℓ_{\max}^6 and become computationally prohibitive at the resolution of modern CMB experiments. Therefore, approximate pseudo- C_ℓ methods were developed to scale as ℓ_{\max}^3 [184, 181, 185, 186, 187]. The pseudo- C_ℓ framework can be applied to both temperature and polarization data, and may be corrected for the effects of incomplete sky coverage [188, 189]. There are several publicly available implementations of the pseudo- C_ℓ approach, such as `PolSpice` [187], and `NaMaster` [189], both of which are used in later chapters.

1.5.3.1 Quadratic estimator and minimum-variance

Given a fields, $\mathbf{f}(\hat{n})$, we wish to estimate the various cross spectra of the temperature and polarization fields T , E and B . In vector form, the spherical harmonic decomposition of $\mathbf{f}(\hat{n})$ is:

$$\mathbf{f}_{\ell m} = \begin{pmatrix} f_{\ell m}^T \\ f_{\ell m}^E \\ f_{\ell m}^B \end{pmatrix} = \int d\Omega \mathbf{Y}_{\ell m}^\dagger(\hat{n}) \mathbf{f}(\hat{n}) \quad (1.60)$$

where $\mathbf{Y}_{\ell m}(\hat{n})$ is a matrix with elements composed of combinations of spherical harmonics and spin-weighted spherical harmonics to perform the decomposition of $\mathbf{d}(\hat{n})$ into the E - B basis [189]. As in Equation 1.48, the power spectra are defined by:

$$\langle \mathbf{f}_{\ell m} \mathbf{f}_{\ell' m'}^\dagger \rangle \equiv \mathbf{C}_\ell \delta_{\ell\ell'} \delta_{mm'}, \quad (1.61)$$

where \mathbf{C}_ℓ is a matrix containing the cross spectra between different fields. Assuming the data are composed of a signal and noise part:

$$\mathbf{f}_i = \mathbf{s}_i + \mathbf{n}_i \quad (1.62)$$

where \mathbf{s} is the signal, and \mathbf{n} is the noise, both of which are assumed to be Gaussian distributed with pixel-pixel covariances:

$$\mathbf{N} = \langle \mathbf{nn}^\dagger \rangle \quad (1.63)$$

$$\mathbf{S} = \langle \mathbf{ss}^\dagger \rangle = \sum_i \mathbf{P}^i \mathbf{C}_i \quad (1.64)$$

$$\mathbf{C} = \langle \mathbf{dd}^\dagger \rangle = \mathbf{N} + \sum_i \mathbf{P}^i \mathbf{C}_i \quad (1.65)$$

where \mathbf{P}_i is a bandpower response function, and \mathbf{C}_i are the bandpowers to be determined. One can write down an estimator, \mathbf{E}_i , for a given bandpower that is quadratic in the data [183, 180]:

$$\hat{\mathbf{C}}_i = \mathbf{f}^\dagger \mathbf{E}_i \mathbf{f} - \mathbf{b}_i, \quad (1.66)$$

where \mathbf{b}_i is chosen to remove the noise bias. By considering the expectation value of Equation 1.66, one finds that:

$$\langle \hat{\mathbf{C}}_i \rangle = \sum_{i'} \text{Tr} [\mathbf{E}_i \mathbf{P}_{i'}] \mathbf{C}_{i'}, \quad (1.67)$$

and that the variance is:

$$\langle \mathbf{C}_i \mathbf{C}_{i'} \rangle = \frac{1}{2} \text{Tr} [\mathbf{C} \mathbf{P}^i \mathbf{C} \mathbf{P}^{i'}]. \quad (1.68)$$

Minimizing the variance of Equation 1.68 yields a specific form for the estimator, up to an overall normalization:

$$\mathbf{E}_i = \frac{1}{2} \mathbf{C}^{-1} \mathbf{P}^i \mathbf{C}^{-1}. \quad (1.69)$$

Plugging this back into Equation 1.67 and assuming that the data is indeed Gaussian distributed, we find that the expectation value of the estimator in Equation 1.66 is related to the true bandpower through the Fisher matrix, \mathbf{F} :

$$\langle \hat{\mathbf{C}}_i \rangle = \mathbf{F}_{ij} \mathbf{C}_j. \quad (1.70)$$

As we discussed in the previous section, no unbiased estimator may have smaller error bars than those given by the Fisher information matrix. In this sense the quadratic estimator in Equation 1.69 is optimal. In summary, one can extract optimal

minimum-variance estimates of some bandpowers C_i using:

$$\mathbf{F}_{ij}C_i = \frac{1}{2}\mathbf{f}^\dagger\mathbf{C}^{-1}\mathbf{P}_i\mathbf{C}^{-1}\mathbf{f} - \frac{1}{2}\text{Tr}[\mathbf{C}^{-1}\mathbf{P}_i\mathbf{C}^{-1}\mathbf{N}]. \quad (1.71)$$

1.5.3.2 Pseudo-spectrum techniques

In order to improve on the speed of power spectrum estimation over the quadratic techniques of the previous section, the pseudo-spectrum approach assumes that the pixel-pixel covariance is diagonal [184, 190, 189]. The result is that the matrix C^{-1} can now be represented by a single weight map of length n_{pix} . The assumption of independent pixels will be closest to the optimal solution when the maps are dominated by white Gaussian noise, or when there is an almost-white underlying signal. In practical applications, the deviation from optimality is only expected to be of order 10-20% [191, 189].

The assumption of independent pixel noise now allows us to work purely with the field \mathbf{f} , and the associated weight-map, \mathbf{w}_f . Denoting the masked field $\mathbf{f}^w(\hat{n})$, we can write its harmonic coefficients as $\mathbf{f}_{\ell m}^w = \sum_{\ell_1 m_1 \ell_2 m_2} D_{\ell m \ell_1 m_1 \ell_2 m_2} \mathbf{f}_{\ell_1 m_1} \mathbf{w}_{\ell_2 m_2}$. The pseudo-spectrum, \tilde{C}_ℓ , is the power spectrum of these masked harmonics:

$$\tilde{C}_\ell = \frac{1}{2\ell + 1} \sum_m \mathbf{f}_{\ell m}^w \mathbf{f}_{\ell m}^{w \dagger}. \quad (1.72)$$

\tilde{C}_ℓ is a biased estimate of the underlying power spectrum C_ℓ due to the mode-coupling induced by the mask \mathbf{w}_f [184, 189]:

$$\tilde{C}_\ell = \sum_{\ell'} \mathbf{M}_{\ell\ell'} C_{\ell'}. \quad (1.73)$$

The mode-coupling matrix, $\mathbf{M}_{\ell\ell'}$, is purely a function of the weight map, and may be calculated exactly and inverted. This is the main computational step involved in the pseudo-spectrum technique, but only has to be done once for each mask [192, 193]. Expressions for the mode-coupling matrix, and details of the numerical implementation of its calculation may be found in Refs. [189, 193].

Chapter 2

Python Sky Model: software for
simulating the Galactic microwave
sky

2.1 Introduction

The CMB temperature anisotropy dominates over foreground emission from the Galaxy in a broad range of frequencies. In contrast, the polarised CMB signal is weaker than the strongly polarised Galactic thermal dust and synchrotron radiation. In particular, the divergence-free B -mode polarisation signal sourced by primordial gravitational waves at recombination is predicted to be at least several orders of magnitude weaker than the polarised foregrounds, averaged over the sky, and is a subdominant signal even in the cleanest sky regions [10].

To optimise our ability to extract the CMB polarisation signal from upcoming and future experiments we rely on realistic models of the Galactic emission to simulate observations of these components at a range of frequencies. Several sky simulation tools are already publicly available, including the Planck Sky Model [194] and the Global Sky Model [195, 196]. While our work was in preparation a similar modelling and software effort was presented in Ref. [197] for polarised Galactic emission.

With the new code presented here we build on existing efforts, providing a flexible and easily used tool for simulating Galactic emission that includes recent public data from the *Planck* satellite. Unique to this code, we define a set of viable alternative models that are consistent with current data, in combination with a new prescription to simulate smaller-scale realisations of the components. We do not attempt to physically model the emission in three dimensions, via for example, integrating a dust or electron density over a Galactic magnetic field [198, 199, 200, 201, 202, 203]. Instead we adopt empirical models that describe the frequency scaling of each component with simple forms consistent with current data, using high signal-to-noise maps of each component as templates at frequencies far from the foreground minimum. These simulations will not therefore capture all the complexity present in the true emission.

The structure of this chapter is as follows: in §2.2 we describe the structure of the code together with the models and alternatives used for each component. In §2.3 we describe a procedure to add small-scale anisotropy to the simulated maps. In §2.4 we describe a method used to introduce a decorrelation of the signal in each component as a function of frequency; and in §2.5 we demonstrate the usefulness of these simulations by forecasting the impact of foreground complexity on the uncertainty of the recovered tensor-to-scalar ratio for the LiteBIRD experiment as an example.

2.2 Large-scale simulations

We simulate Galactic diffuse emission in intensity and polarisation from four Galactic components: thermal dust, synchrotron, free-free, and anomalous microwave emission (AME). We also include a gravitationally lensed CMB realisation and white instrument noise. Maps can be integrated over a top-hat bandpass describing the response of each experimental channel, and smoothed with a Gaussian beam.

The user specifies a set of observation frequencies, beam widths, bandpass widths, noise and chosen output components and units. The code simulates each component at each frequency using a phenomenological model. One or more emission template maps are defined at pivot frequencies, and then the extrapolation in frequency is performed using scaling laws and maps of spectral parameters. A lensed CMB realisation can be included by calling the `Taylens` software [204] directly, or using a pre-calculated realisation.

In this section we describe the suite of models available in `PySM`. These models are designed to cover a range of new complexities found to be consistent with the most recent *WMAP* and *Planck* data. For example we provide a new model of synchrotron spectral index steepening consistent with *WMAP* nine year data, and a model of dust decorrelation with frequency, an effect detected in the most recent *Planck* analysis [205] and compatible with spatially-varying spectral indices. Careful consideration of such effects is vital in the design of the next generation of CMB experiments seeking to constrain the tensor-to-scalar ratio in order to avoid bias introduced by model mismatch. Furthermore, we provide for the first time simulations of small scale power modulated by the signal-dominated large scale power, making these maps suitable for both full sky analysis and the analysis of small patches in high Galactic latitude regions.

The intensity and polarised emission as a function of frequency for the models we consider is summarised in Figure 2.1, together with the template maps in Figure 2.2.

2.2.1 Synchrotron

Synchrotron radiation is the dominant radiation mechanism in polarisation at frequencies $\lesssim 50$ GHz [148]. It is produced by cosmic rays spiralling around Galactic magnetic fields and radiating. The power and spectral energy distribution depends on both the strength of the local magnetic field, and the energy distribution of the injected cosmic rays. The polarisation of the radiation depends on the orientation




Figure 2.1: The frequency scaling laws for the individual components of PySM; we show the nominal and alternative models as solid and dashed lines respectively. We show only the alternative models which have a significant impact on the shape of the spectrum. These spectra are calculated by producing masked maps of each component at each frequency, smoothing to FWHM 1° in intensity and $40'$ in polarisation, and then computing the RMS. The mask used in intensity is the *WMAP* 9 year KQ85 mask, and the polarisation mask is the *Planck* polarisation confidence mask CPM83.

of the intervening magnetic field. The predicted dependence of the spectrum on the magnetic field for a population of cosmic rays with energy distribution $N(E) \propto E^{-p}$ is, in antenna temperature units:

$$I_\nu \propto B^{\frac{p+1}{2}} \nu^\beta, \quad (2.1)$$

where $\beta = -\frac{(p+3)}{2}$ [149]. The spectral index, β , is expected to have some spatial variability and to vary with frequency. As synchrotron sources age their spectral energy distribution (SED) steepens, since high frequency radiation corresponds to higher energy particles which radiate energy away most rapidly. Along a line of sight there will likely also be multiple synchrotron components, and the stacking of their spectra can lead to flattening of the SED. The spectrum can also be flattened through effects of synchrotron self-absorption, which tends to be more significant towards the Galactic centre.

2.2.1.1 Model 1: Nominal index

The nominal PySM model assumes that the synchrotron intensity is a scaling of the degree-scale-smoothed 408 MHz Haslam map [206, 207], reprocessed by Ref. [208]. It models the polarisation as a scaling of the *WMAP* 9-year 23 GHz Q and U maps

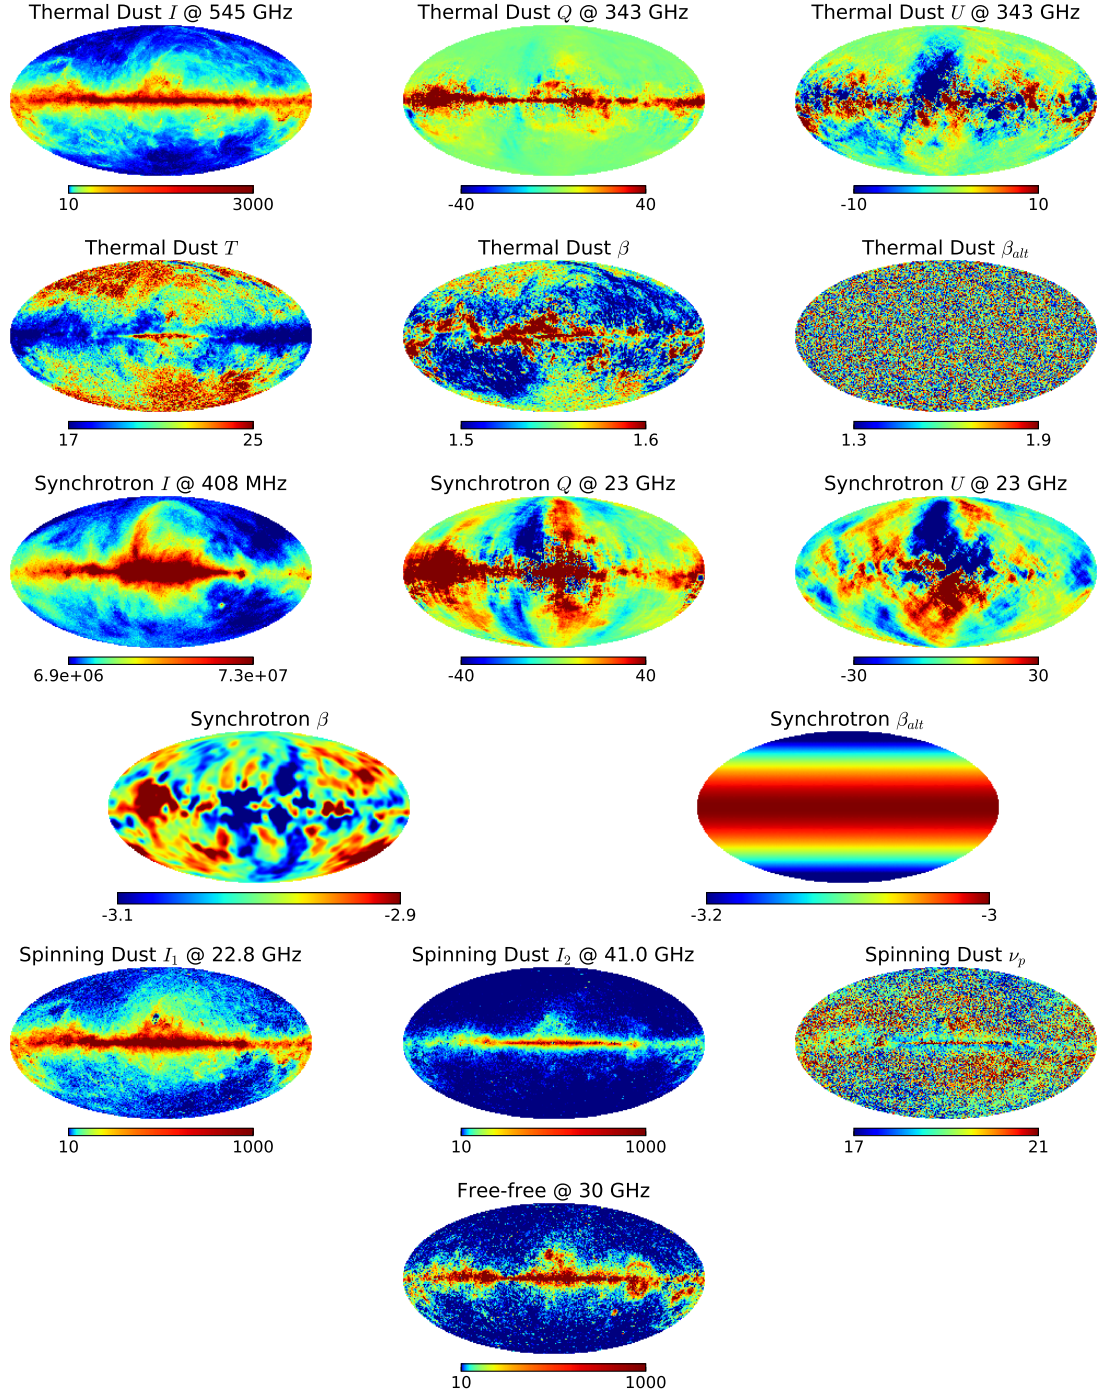


Figure 2.2: Template maps used in the PySM models. All emission templates are in units of μK_{RJ} and all dust temperature templates are in K. Intensity templates are plotted on a log scale, and the polarisation templates on a linear scale.

[100], smoothed to three degrees. Both of these maps have small scales added using the prescription described in §2.3.

In the nominal model we simulate the spectral index as being a power-law in every direction, such that

$$I_{\nu}^{\text{Synch}}(\hat{\mathbf{n}}) = A_{\nu_0}(\hat{\mathbf{n}}) \left(\frac{\nu}{\nu_0} \right)^{\beta_s(\hat{\mathbf{n}})}. \quad (2.2)$$

As in the nominal Planck Sky Model v1.7.8 simulations, we use the spectral index map from ‘Model 4’ of Ref. [209], calculated from a combination of Haslam and *WMAP* 23 GHz polarisation data using a model of the Galactic magnetic field. We assume that the index is the same in temperature and polarisation, although the true sky will most likely be more complicated than this. The template maps and index map are shown in Fig 2.2.

2.2.1.2 Model 2: Spatially steepening index

The cosmic rays responsible for synchrotron radiation are thought to be energised by processes such as supernovae, which are more common in the Galactic plane. Synchrotron emission observed at higher latitudes will therefore likely be produced by older cosmic rays which have diffused out of the Galactic plane, and therefore lost more energy. This is expected to result in the steepening of the synchrotron spectral index away from the plane [210, 126]. Evidence for steepening in the polarisation emission has been seen in Refs. [148, 211, 212] using *WMAP* and QUIET data.

We parameterise the steepening with a smoothly varying index described by a gradient δ_β that scales with Galactic latitude, b , such that $\beta_s = \beta_{s,b=0} + \delta_\beta \sin |b|$. In Model 2 we use $\delta_\beta = -0.3$, consistent with *WMAP* polarisation data [148, 213]. The simulated index varies from $\beta_s = -3.0$ at the equator to $\beta_s = -3.3$ at the poles in both intensity and polarisation, as shown in Figure 2.2.

2.2.1.3 Model 3: Curvature of index

The synchrotron emission may be better modelled by a curved spectrum that either flattens or steepens with frequency. Model 3 simulates the steepening or flattening of the spectral index above a frequency, ν_c as:

$$I_{\nu}^{\text{Synch}}(\hat{\mathbf{n}}) = A_{\nu_0}(\hat{\mathbf{n}}) \left(\frac{\nu}{\nu_0} \right)^{\beta_s(\hat{\mathbf{n}}) + C \ln(\frac{\nu}{\nu_c})}, \quad (2.3)$$

where positive C corresponds to flattening and negative C to steepening.

Ref. [210] fits this model to a small patch of sky with ten overlapping radio frequency sky surveys and *WMAP* 23 GHz data, finding best-fit values of $\beta = -2.64 \pm 0.03$, $C = -0.052 \pm 0.005$ at 0.31 GHz. This corresponds to a steepening of about 0.57 between 408 MHz and 94 GHz. Evaluating the spectral index at 23 GHz, Ref. [210] finds $\beta_{23} = -3.09 \pm 0.05$. This is consistent with the index map in model 1 which has a mean and standard deviation of -3.00 ± 0.06 . Therefore, for simplicity we use the same map as model 1 from Ref. [209] for $\beta(\hat{\mathbf{n}})$, and a baseline curvature value of $C = -0.052$ at $\nu_c = 23$ GHz.

2.2.2 Thermal dust emission

At frequencies greater than ≈ 70 GHz the polarised foreground spectrum is dominated by thermal dust emission. The dust grains are thought to be a combination of carbonaceous and silicate grains, and polycyclic aromatic hydrocarbons (PAHs). The total emission results from the interaction of these species with the interstellar radiation field: the grains are heated by absorption in the optical and cool by emitting in the far infrared [127]. The thermal dust is polarised since aspherical dust grains preferentially emit along their longest axis, which tend to align perpendicular to magnetic fields.

In the frequency range of interest for CMB experiments, the spectrum is well approximated by a modified blackbody with a power-law emissivity, such that

$$I = A\nu^{\beta_d} B_\nu(T_d), \quad (2.4)$$

for spectral index β_d and temperature T_d , where B_ν is the Planck function. A single component at $T = 15.9$ K fits the *Planck* data well [214], with different indices preferred by the intensity ($\beta = 1.51 \pm 0.01$) and polarisation (1.59 ± 0.02) data. This difference indicates the presence of multiple components with different polarisation properties.

In intensity the two component model of Ref. [215], with a hot and cold component at 9.4 K and 16 K, is marginally preferred [216]. They use this model to extrapolate 100 μm emission and 100/240 μm flux ratio maps to microwave frequencies. The exact physical model is not well constrained by current observations, including the number of components, spatial variability of spectral index, and spatial variation of the dust temperature.

2.2.2.1 Model 1: Nominal index

Our nominal model uses template maps at 545 GHz in intensity and 353 GHz in polarisation. We use the templates estimated from the *Planck* data using the ‘Commander’ code [214]. In polarisation these maps closely match the 353 GHz *Planck* data which is dominated by thermal dust. We use the N_{side} 2048 dust intensity map degraded to N_{side} 512, and the polarisation product smoothed to two degrees FWHM in polarisation with small scale variations added by the procedure described in section §2.3.

In the nominal simulations, we model the frequency scaling as a single component, using the best-fit emissivity estimated by the **Commander** fit. The emission model is given by

$$\begin{aligned} I_\nu^{\text{d}}(\hat{\mathbf{n}}) &= A_{\text{I},\nu_I}(\hat{\mathbf{n}})(\nu/\nu_I)^{\beta_d(\hat{\mathbf{n}})} B_\nu(T_d(\hat{\mathbf{n}})) \\ \{Q_\nu^{\text{d}}(\hat{\mathbf{n}}), U_\nu^{\text{d}}(\hat{\mathbf{n}})\} &= \{A_{\text{Q},\nu_P}(\hat{\mathbf{n}}), A_{\text{U},\nu_P}(\hat{\mathbf{n}})\} \times \\ &\quad (\nu/\nu_P)^{\beta_d(\hat{\mathbf{n}})} B_\nu(T_d(\hat{\mathbf{n}})) \end{aligned} \quad (2.5)$$

Here $\nu_I = 545$ GHz and $\nu_P = 353$ GHz. We assume that the intensity and polarisation share the same index, as was assumed in the **Commander** fitting process. Both β_d and T_d vary spatially; the maps are shown in Fig 2.2.

This model will not capture all the of the physical complexity as it is likely that silicate and carbonaceous grains have distinct emissivities. They also likely have different degrees of polarisation, since the efficiency of the grain alignment varies with the size and shape of grain. This would result in the polarisation fraction in dust being a function of frequency, with some evidence for this shown in Ref. [10].

2.2.2.2 Models 2 and 3: Spatially variable index

The dust index is expected to vary spatially, in particular in polarisation, but current data cannot strongly constrain this possible variation. We perform a test to assess how well a varying index can be detected by the *Planck* data given the current noise levels.

We simulate a spectral index map with degree-scale variation drawn from a Gaussian of mean 1.59 and dispersion σ . We then simulate polarised dust emission in Stokes Q and U at 217 and 353 GHz at $N_{\text{side}} = 128$ for σ in the range 0.05 to 0.7. We produce noise maps at 217 GHz and 353 GHz using the *Planck* half-mission

Model	Mean	Std. Dev.
Nominal	1.53	0.22
$\sigma(\beta) = 0.2$	1.58	0.24
$\sigma(\beta) = 0.3$	1.58	0.23
Uniform	1.58	0.23

Table 2.1: Statistics of dust polarisation index calculated from different simulations of dust polarisation at 217 GHz and 353 GHz containing instrumental noise compatible with the corresponding *Planck* channels.

full-sky maps at 217 and 353 GHz. We first degrade these to $N_{\text{side}} = 128$ and then at each frequency take the difference of the two half-mission maps and divide by a factor of 2. Finally we smooth each noise map with a Gaussian kernel of one degree FWHM. We then estimate the index from these maps in circles of radius ten degrees centred on HEALPix $N_{\text{side}} = 8$ pixels, using

$$\beta_d(\hat{\mathbf{n}}) = \frac{\ln\left(\frac{[Q,U]_1(\hat{\mathbf{n}}) B(\nu_2, T(\hat{\mathbf{n}}))}{[Q,U]_2(\hat{\mathbf{n}}) B(\nu_1, T(\hat{\mathbf{n}}))}\right)}{\ln\left(\frac{\nu_2}{\nu_1}\right)} + 2. \quad (2.6)$$

This follows a similar method used in the comparable *Planck* analysis in Ref. [10], except we do not use the 143 GHz channel and do not add CMB and synchrotron, nor fit for them. We use a similar region as the *Planck* analysis, shown in Fig 2.4. The dispersion of the indices for a uniform input index of 1.59, and for an input index map with degree-scale variation of standard deviation of 0.2 is shown in Figure 2.3, and can be compared to Figure 9 in [10]. The statistics of the recovered index distributions for these two models, the nominal model, and a model with a larger standard deviation of 0.3, are shown in Table 2.1.

The distributions for β_d are similar since the data are noise-dominated. The dispersion due to noise is ~ 0.22 compared to the value of 0.17 found in Ref. [10], and the value of 0.22 found in a comparable calculation by Ref. [141]. This indicates that models of the dust spectral index with significant spatial variation on degree scales are still consistent with the data. Furthermore, a recent analysis of decorrelation of the *Planck* half-mission and detector set maps found an intrinsic variation of 0.07 in the dust index [205]. Models 2 and 3 therefore modify the nominal dust model with a different spectral index map. The spectral index of model 2 (3) is a Gaussian random field with mean of 1.59 and $\sigma = 0.2(0.3)$ varying on degree scales for both intensity and polarisation.




Figure 2.3: Normalised histograms of the dust spectral index, β_d , calculated for noisy simulations of different PySM models with varying intrinsic index dispersion. We see that the resulting dispersions are very similar, indicating noise-dominated data.




Figure 2.4: Mask used in calculation of β_d in section 2.2.2.2. This mask is an approximation to the one used in the *Planck* analysis (figure 1 of Ref. [10]).

2.2.2.3 Model 4: Two dust temperatures

One can also consider a number N_d of dust components with their own temperatures and spectral indices:

$$I(\hat{\mathbf{n}}, \nu) = \sum_{a=1}^{N_d} I_a(\hat{\mathbf{n}}) \left(\frac{\nu}{\nu_*} \right)^{\beta_a} \frac{B_\nu(T_a(\hat{\mathbf{n}}))}{B_{\nu_*}(T_a(\hat{\mathbf{n}}))}, \quad (2.7)$$

and similarly for polarisation. For our fourth dust model we use $N_d = 2$, using the best-fit model templates estimated by Ref. [216] from the *Planck* data, using the model from Ref. [215].

The model as proposed in these references can be written as

$$I(\hat{\mathbf{n}}, \nu) = I_{\nu_0}(\hat{\mathbf{n}}) \frac{\sum_{a=1}^2 f_a q_a \left(\frac{\nu}{\nu_0} \right)^{\beta_a} B_\nu(T_a(\hat{\mathbf{n}}))}{\sum_{b=1}^2 f_b q_b B_{\nu_0}(T_b(\hat{\mathbf{n}}))}, \quad (2.8)$$

where I_{ν_0} is the intensity template at 100 μm ($\nu_0 = 3000$ GHz), β_k are constant spectral indices, T_k are spatially varying dust temperatures, q_k is the IR/optical ratio for each species, f_k is the fraction of power absorbed from the interstellar radiation field and emitted in the FIR by each component, and we have omitted the colour correction factors. In order to adapt this to the model in Eq. 2.7 we generate the separate amplitude templates $I_a(\hat{\mathbf{n}})$, at $\nu_* = 545$ GHz in terms of I_{ν_0} and $T_k(\hat{\mathbf{n}})$ as

$$I_a(\hat{\mathbf{n}}) = I_{\nu_0}(\hat{\mathbf{n}}) \frac{\left(\frac{\nu_*}{\nu_0} \right)^{\beta_a} f_a q_a B_{\nu_*}(T_a(\hat{\mathbf{n}}))}{\sum_{b=1}^2 f_b q_b B_{\nu_0}(T_b(\hat{\mathbf{n}}))}. \quad (2.9)$$

In polarisation, we construct the polarisation simulations using the polarisation angles and fractional polarisation from the 353 GHz template maps in Model 1, such that

$$\begin{aligned} Q(\nu, \hat{\mathbf{n}}) &= f_d(\hat{\mathbf{n}}) I(\nu, \hat{\mathbf{n}}) \cos(2\gamma(\hat{\mathbf{n}})) \\ U(\nu, \hat{\mathbf{n}}) &= f_d(\hat{\mathbf{n}}) I(\nu, \hat{\mathbf{n}}) \sin(2\gamma(\hat{\mathbf{n}})). \end{aligned} \quad (2.10)$$

where $f_d = \sqrt{Q^2 + U^2}/I$ at 353 GHz in Model 1.

2.2.3 Anomalous microwave emission

Anomalous microwave emission refers to emission with a spectral distribution not well approximated by known foreground models. It has been detected in compact objects, and in the diffuse sky, with early measurements by Refs. [150, 151]. It is spatially correlated with dust, and primarily important in the 20 - 40 GHz range, with variable peak frequency [217].

A likely model for the emission is rapidly spinning dust grains. Ref. [154] explain the emission by a population of grains of size $< 3 \times 10^{-7}$ cm, with modest electric dipole moments. A candidate for these grains is polycyclic aromatic hydrocarbons (PAHs) that are detected in vibrational emission in the range $3 - 12 \mu\text{m}$. The theoretical SED for such spinning PAH grains have been successfully fit to AME observations [218], but recent analysis of the *Planck* data has cast some doubt on their nature [219]. A second candidate for AME is magnetic dipole radiation due to thermal fluctuations of magnetisation in small silicate dust grains [155].

2.2.3.1 Model 1: Nominal unpolarised AME

We model the AME using the *Planck* templates derived from the **Commander** parametric fit to the *Planck* data [214], using the **Commander** model:

$$\begin{aligned}
 I_{\nu}^{\text{AME}}(\hat{\mathbf{n}}) &= A_{\nu_{0,1}}(\hat{\mathbf{n}})\epsilon(\nu, \nu_{0,1}, \nu_{p,1}(\hat{\mathbf{n}}), \nu_{p0}) \\
 &+ A_{\nu_{0,2}}(\hat{\mathbf{n}})\epsilon(\nu, \nu_{0,2}, \nu_{p,2}, \nu_{p0}).
 \end{aligned}
 \tag{2.11}$$

Here the first component has a spatially varying emissivity, and the second component a spatially constant emissivity. Both these emissivity functions are calculated using **SpDust2** [220, 221], evaluated for a cold neutral medium and shifted in $\log(\nu) - \log(I)$ space. The two template maps are shown in Figure 2.2. This nominal AME model is unpolarised.

2.2.3.2 Model 2: Polarised AME

AME is not thought to be strongly polarised, and the polarisation fraction has been constrained to be below 1 - 3% in the range 23 - 41 GHz by observations of the Perseus molecular complex using *WMAP* 7-year data [159]. More recent observations of AME

emission from the molecular complex W43 by the QUIJOTE experiment have placed a 0.39% upper limit on its polarisation fraction, which falls to 0.22% when combined with *WMAP* data [222]. Ref. [223] found that neglecting a 1% level of polarised AME can bias the derived value of the tensor-to-scalar ratio by non-negligible amounts for satellite missions.

To construct a template we use the dust polarisation angles, γ_d , calculated from the *Planck* Commander 2015 thermal dust Q and U maps at 353 GHz. The AME polarisation is then

$$Q_a = f_a I_\nu \cos(2\gamma_{353}), \quad U_a = f_a I_\nu \sin(2\gamma_{353}). \quad (2.12)$$

In this model we assign a global polarisation fraction of 2%; the fraction can also be easily changed by varying the f_a parameter.

2.2.4 Free-free

Free-free emission is caused by electrons scattering off ions in the interstellar medium [149]. The frequency scaling is well approximated by a function of the electron temperature and emission measure [127]. This is very close to a power law of -2.14 at frequencies greater than 1 GHz, and flattens abruptly at lower frequencies [214].

Free-free has been measured in *WMAP* and *Planck* intensity data, and it should be unpolarised since the scattering is independent of direction. However, there are small effects at the edges of dense ionised clouds due to the non-zero quadrupole moment in the electron temperature, which can cause up to 10% polarisation [224]. The net polarisation over the sky is estimated to be below 1% [225].

The PySM nominal model for free-free emission assumes it is unpolarised, and uses the degree-scale smoothed emission measure and effective electron temperature *Commander* templates [214]. We apply the analytic law presented in Ref. [127] to produce an intensity map at 30 GHz, which we then scale with a spatially constant power law index. We choose this index to be -2.14 consistent with *WMAP* and *Planck* measurements for electrons at ~ 8000 K [100, 214]. This gives

$$I_\nu^{\text{ff}}(\hat{\mathbf{n}}) = A_{\nu_0}^{\text{ff}}(\hat{\mathbf{n}}) \left(\frac{\nu}{\nu_0} \right)^{-2.14}. \quad (2.13)$$

Different behaviour will be expected below ~ 0.01 GHz, where the *Commander* model flattens [214].

2.2.5 CMB

We use the `Taylens` code [204] in `PySM` to generate a lensed CMB realisation. The input to `Taylens` is a set of C_l 's ($C_{TT}, C_{EE}, C_{BB}, C_{TE}, C_{\phi\phi}, C_{T\phi}, C_{E\phi}$) which have been calculated using the CAMB numerical code [65]. The nominal model uses Λ CDM cosmological parameters that best fit the *Planck* 2015 data. We incorporate the functions of `Taylens` into the `PySM` code for portability, so some functionality is removed¹. We scale the CMB emission between frequencies using the blackbody function.

The user can opt to either run `Taylens` during the simulation, or use a pre-computed temperature and polarisation map supplied with the code or generated by the user. If using `Taylens`, the CMB map can also be artificially delensed, with the expected lensing signal suppressed by a chosen factor.

2.2.6 Instrument

We describe the instrument response with a simple top-hat bandpass, Gaussian white noise, and Gaussian beam profile. The user specifies a central frequency, ν , and a width per band, $\Delta\nu$. The output signal is calculated using

$$I_{\nu, \Delta\nu}(\hat{\mathbf{n}}) = \int_{\nu - \frac{\Delta\nu}{2}}^{\nu + \frac{\Delta\nu}{2}} \frac{I_{\nu'}(\hat{\mathbf{n}})}{\Delta\nu} d\nu'. \quad (2.14)$$

The white noise level is set per band for both intensity and polarisation. The beam is characterised by a FWHM per channel. This instrument model will not capture realistic noise realisations or realistic bandpasses; the code is designed to be easily modifiable to incorporate such features.

2.3 Small-scale simulations

Ground-based CMB experiments often observe only small patches of sky, and current data limit how well we can predict the small-scale behaviour of the foregrounds in high latitude regions at the $\ell \sim 100$ scales of interest. Here we describe our method for simulating sky maps at a higher resolution than the available data. Our approach is to extrapolate the angular power spectrum of the available data to smaller scales,

¹The original code is available at <https://github.com/amaurea/taylens>




Figure 2.5: *Left*: synchrotron polarisation spectra in a square region centred on RA, DEC = [0, -55] of size 1600 deg². The errors shown are cosmic variance only. The best-fit power-law signal plus noise model from Eqn. 2.15 is shown. The BB model minimum is used to estimate the scale l_* to smooth the maps. *Right*: synchrotron polarisation spectra computed with the *WMAP* polarisation analysis mask, and best-fit model. The dashed lines are the extrapolated power laws used in the small-scale simulation.




Figure 2.6: *Left*: dust polarisation spectra as in Figure 2.5, but for a smaller patch of 800 square degrees. *Right*: dust polarisation spectra as in Figure 2.5, but using the *Planck* Gal 80 Galactic plane mask.




Figure 2.7: *Left*: synchrotron BB spectra using the 1600 square degree region centred on RA, DEC = [0, -55]. We show the original template, the smoothed template, the small scale realisation, and the final map with small scales added. The dashed red line shows the shape of the power law of the small scale realisation to guide the eye. *Right*: synchrotron BB spectra over 75% of the sky using the *WMAP* polarisation analysis mask.




Figure 2.8: *Left*: dust BB spectra in the 825 square degree region centred on RA, DEC = [0, -55], as in Figure 2.7. *Right*: dust BB spectra in the Gal 80 region, as in Figure 2.7.

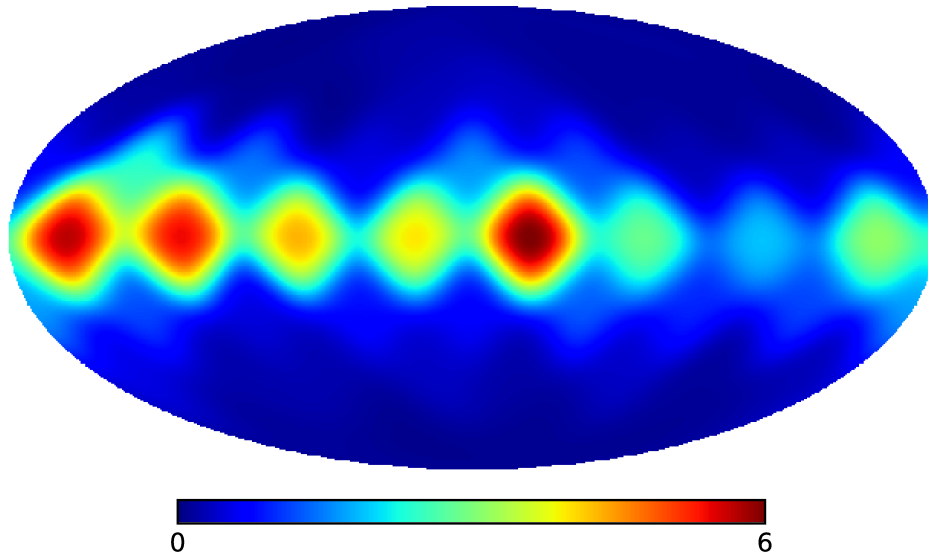


Figure 2.9: Normalisation map, $N(\hat{\mathbf{n}})$, for the dust Q map.

drawing a Gaussian realisation from this spectrum. Other similar methods have been implemented in Refs. [209, 194, 208, 197].

We simulate intensity and polarisation maps using $M = M_0 + M_{\text{ss}}$ where M_0 is the original smoothed data and M_{ss} is our small-scale simulation. We implement different methods in polarisation and intensity for generating M_{ss} . Although the real sky will be non-Gaussian, we limit these small-scale simulations to Gaussian or lognormal realisations.

2.3.1 Polarisation

The *WMAP* and *Planck* polarisation templates used in *PySM* are all noise dominated at degree scales at high Galactic latitudes. To add power to the Q and U maps at small scales we determine the multipole, ℓ_* , to which the original template is limited in resolution, smooth the maps to this scale, and add a realisation of a model power law spectrum to the smoothed templates. We compute angular power spectra on masked skies using the *PolSpice* code² [187].

The scale ℓ_* varies spatially, but here we adopt a single global ℓ_* , which we determine by computing the polarisation power spectra in a region centred on RA, DEC = [0, -55], chosen as the location of the BICEP2/Keck patch. We choose a square region of side 40 degrees for synchrotron, and 30 degrees for dust, with a larger region

²The *PolSpice* code is available at <http://www2.iap.fr/users/hivon/software/PolSpice/>




Figure 2.10: Gnomonic projection of dust Q maps in a patch centred at RA, DEC = (0, -55); 40 degrees to a side. The left panel is the original map, the middle panel has been smoothed (M_0), and the right hand panel has had small scales added ($M_0 + M_{ss}$). The maps are plotted in histogram-equalisation in units of μK_{RJ} .

for synchrotron as the maps are noisier. We fit the spectra with a signal-plus-noise model,

$$\frac{\ell(\ell+1)}{2\pi}C_\ell^{\text{BB}} = A\ell^\gamma + N\frac{\ell(\ell+1)}{2\pi}, \quad (2.15)$$

approximating the uncertainties on the spectrum as due only to cosmic variance. The power law model has previously been found to be an excellent fit to many observations of dust and synchrotron angular power spectra on large scales [226, 10, 205]. We choose to model the small scale behaviour as an extension of this power law. We fit for three free parameters A , γ and N , and estimate ℓ_* as the scale at which this model is minimal in BB or EE . The masked synchrotron and dust EE and BB spectra are shown in Figures 2.5 and 2.6. We find $\ell_*^{\text{synch}} = 36$ and $\ell_*^{\text{dust}} = 69$.

We generate the large-scale templates M_0 by smoothing the original maps with a Gaussian kernel of FWHM $\theta_{\text{fwhm}} = 180/\ell_*$ deg. We then construct M_{ss} by assuming that the small-scales follow a power law behaviour with $\frac{\ell(\ell+1)}{2\pi}C_\ell^{\text{XX}} = A^{\text{XX}}\ell^{\gamma^{\text{XX}}}$. We find A^{XX} and γ^{XX} by fitting this model to the EE and BB spectra calculated on the original template with a Galactic mask. We use the *WMAP* polarisation analysis mask for synchrotron [227], and the 80% mask provided in the second *Planck* data, which we refer to here as Gal80. We find $\gamma^{\text{synch,EE}} = -0.66$, $\gamma^{\text{synch,BB}} = -0.62$, $\gamma^{\text{dust,EE}} = -0.31$, $\gamma^{\text{dust,BB}} = -0.15$.

We multiply these power law spectra by the window function $1 - W_\ell(\ell_*)$, where $W_\ell = \exp(-\sigma^2(\ell_*)\ell^2)$ with $\sigma = \theta_{\text{fwhm}}/\sqrt{8\ln(2)}$, such that it be added to the large-scale map that has been smoothed by the window function W_ℓ . We then draw a pair of Q

and U Gaussian random fields, δ_G , from this spectrum using the HEALPix³ routine *synfast*.

We expect the true small-scale power to be modulated by the large-scale power, so we multiply the Gaussian random field by a spatially varying normalisation such that

$$M_{\text{SS}} = N(\hat{\mathbf{n}})\delta_G(\hat{\mathbf{n}}). \quad (2.16)$$

We choose $N(\hat{\mathbf{n}})$ by dividing the sky into HEALPix $N_{\text{side}} = 2$ pixels and computing the angular power spectrum in each patch, $C_\ell(\hat{\mathbf{n}})$, and smoothing this map with FWHM 10° to avoid sharp pixel boundaries. We define

$$N(\hat{\mathbf{n}}) = \sqrt{\frac{C_{\ell_*}(\hat{\mathbf{n}})}{A\ell_*^\gamma}} \quad (2.17)$$

so that the small-scale realisation is normalised by the large-scale power in each patch. The $N(\hat{\mathbf{n}})$ for the dust Q template is shown in Figure 2.9.

A patch of the resulting Q map for dust is shown in Figure 2.10, illustrating the large-scale and additional small scale components. We also show the power spectra of the maps in Figures 2.7 and 2.8, both for the masked all-sky maps and the smaller regions centred at $[0, -55]$. In both regions the power law behaviour is continuous at $\ell = \ell_*$.

We note that a limitation of this method is that it does not capture spatial variations in the modulation of the small-scale signal on scales smaller than $N_{\text{side}} = 2$ pixels, so the normalisation will not be accurate in these small regions.

2.3.2 Intensity: Synchrotron

We use a similar procedure for simulating the intensity at small-scales, but we use a lognormal rather than Gaussian distribution because it guarantees that the final map will be positive. It is also possible to generate a lognormal distribution from a Gaussian random field, and maintain the shape of the Gaussian field's angular power spectrum to a good approximation. In these simulations we do not impose a correlation between the intensity and polarisation at small scales.

For synchrotron, the Haslam template is provided at 57 arcminute resolution, which defines M_0 . As for the polarisation we fit a power law to the signal, finding

³The HEALPix code is available at <http://healpix.sourceforge.net>

$\gamma = -0.55$. We draw a Gaussian realisation δ_G with variance σ_G^2 , but here we generate M_{ss} using a lognormal distribution with

$$M_{\text{ss}} = M_0^{\text{min}}[\exp(R(\hat{\mathbf{n}})\delta_G(\hat{\mathbf{n}}) - \sigma_G^2/2) - 1], \quad (2.18)$$

motivated by the requirement that intensity be positive, a form commonly used in large scale structure simulations [228]. Here $R(\hat{\mathbf{n}})$ normalises the small scales. Instead of using the local power spectrum, we normalise the small-scale intensity map by the large-scale intensity smoothed to 4° and raised to a power

$$R(\hat{\mathbf{n}}) = \left[\frac{M_0(\hat{\mathbf{n}})}{\langle M_0 \rangle} \right]^\alpha.$$

We find the best-fit $\alpha = 0.6$ that results in a total power spectrum of $\ell(\ell+1)C_\ell \propto \ell^\gamma$, fit in the multipole range $200 < \ell < 1000$. An example of the synchrotron maps are shown in Figure 2.11.

2.3.3 Intensity: Free-free

The free-free template is smoothed at degree scales, which defines M_0 . We found the lognormal procedure to be unsuitable for generating small scales for the free-free maps, as the comparatively larger dynamic range in small patches caused the exponential term to yield unrealistically large variation on small scales. We also found the free-free angular spectrum to be flatter than the synchrotron, so a direct extrapolation of the power law to smaller scales produced excess power at small scales that is likely not physical. We therefore fixed the gradient of the free-free power spectrum to be $\gamma = -0.5$, and used this to generate a δ_G realisation with variance σ_G^2 . We then take the small scale map to be:

$$M_{\text{ss}}(\hat{\mathbf{n}}) = R(\hat{\mathbf{n}})\delta_G(\hat{\mathbf{n}}) \quad (2.19)$$

where $R(\hat{\mathbf{n}}) = \langle M_0 \rangle (M_0(\hat{\mathbf{n}}) / \langle M_0 \rangle)^\alpha / 4\sigma_G$. We find that $\alpha = 1.15$ is the best-fit value to recover the correct power law behaviour of the power spectrum in the range $200 < \ell < 1000$. We redrew δ_G for any negative pixels from additional full-sky realisations until we have positive values everywhere. This was necessary for $< 0.5\%$ of pixels. An example is shown in Figure 2.11.




Figure 2.11: Synchrotron (top), free-free (middle), and AME (bottom) simulated intensity maps in a patch of side 40° centred at RA, DEC as indicated. Left: original template; right: simulation including small scales. These have been plotted in histogram-equalisation to increase the dynamic range.

2.3.4 Intensity: Thermal dust and AME

The *Planck* thermal dust map has a power spectrum in the low-foreground [RA, DEC = 0, -55] region that falls off approximately as a power law. This indicates that the thermal dust intensity map is signal dominated in this high Galactic latitude region at small scales, and we do not add additional components.

The AME templates are limited to degree resolution, so we use the high resolution thermal dust product as a proxy for the AME small scales. We produce the final AME map by multiplying the two intensity templates by the ratio of the high resolution thermal dust template and the dust template smoothed to one degree FWHM. An example is shown in Figure 2.11. The resulting AME templates therefore have the same small-scale morphology as the thermal dust template. Since the AME polarisation templates are produced from the thermal dust polarisation products we do not simulate AME polarisation separately.

2.4 Frequency decorrelation

The spatial variation of spectral parameters discussed in Section 2.2 reflects the stochastic distribution of dust clouds and of the turbulent component of the magnetic field in the Galaxy. This spatial variability will realistically occur both across the sky and along the line of sight, and it will produce a frequency decorrelation that manifests as a departure from simple emission laws [229, 230]. We implement this possible frequency decorrelation in PySM using two different models.

Multiple components. For each component type described in Section 2.2, PySM supports the inclusion of an arbitrary number of components with different amplitude and spectral index maps. This mimics the superposition of emission laws caused by the variation of spectral parameters along the line of sight.

Stochastic decorrelation. Let $m_\nu(\hat{\mathbf{n}}) = v_\nu(\hat{\mathbf{n}})f_\nu(\beta)$ be the sky emission of a given component as a function of frequency, where f_ν is any of the emission laws described in Section 2.2, governed by a set of spectral parameters β . The component v_ν is therefore the sky emission normalised by the emission law, so $v_\nu = v_{\nu'}$ in the absence of frequency decorrelation. In what follows, we write the harmonic coefficients of v_ν as a vector \mathbf{v} of N_ν components, where N_ν is the total number of frequencies, for each multipole (ℓ, m) .




Figure 2.12: Examples of the frequency dependence of a given component for different correlation lengths ξ , normalised by the uncorrelated case, given a template at $\nu = 300$ GHz. The sharp kinks at low frequencies for small values of ξ are caused by the frequency correlation length being smaller than the sampling rate.

We assume that the covariance between different frequencies takes the form

$$R_{\nu\nu'} \equiv \langle v_\nu v_{\nu'}^* \rangle = \langle |v_\nu|^2 \rangle \exp \left[-\frac{1}{2} \left(\frac{\log(\nu/\nu')}{\xi} \right)^2 \right], \quad (2.20)$$

following Ref. [144], where ξ is the correlation length in frequency space. Let us now decompose \mathbf{v} as $\mathbf{v} = (c_1, c_2, \dots, c_{N_c}, u_1, u_2, \dots, u_{N_u})$, where c_i are sky maps corresponding to N_c observed templates at different frequencies, and u_i correspond to the sky emission in $N_u = N_\nu - N_c$ unobserved bands. The goal is to generate the unconstrained maps \mathbf{u} with the correct correlation properties, as in Eq. 2.20, subject to the N_c constraints of the observed sky maps \mathbf{c} .

We simplify the process by making two assumptions:

- \mathbf{u} is a Gaussian random variable.
- $\langle |v_\nu|^2 \rangle$ does not depend on ν (although it may depend on ℓ, m).

The solution is then to draw a sample map from a multivariate Gaussian distribution given by:

$$\mathbf{u} \leftarrow \mathcal{N}(\bar{\mathbf{u}}, \mathbf{C}_{uu}), \quad (2.21)$$

where the mean $\bar{\mathbf{u}}$ and covariance \mathbf{C}_{uu} are given by:

$$\mathbf{C}_{uu} = ((\mathbf{R}^{-1})_{uu})^{-1}, \quad \bar{\mathbf{u}} = \mathbf{C}_{uu} \cdot (\mathbf{R}^{-1})_{uc} \cdot \mathbf{c}. \quad (2.22)$$

Here \mathbf{R} is the frequency covariance matrix in Eq. 2.20 and the subindices uu and uc select the $\mathbf{u} - \mathbf{u}$ and $\mathbf{u} - \mathbf{c}$ elements of the corresponding matrix.

In the current version of PySM we simplify this process further by considering only one constrained template at a single frequency for each component ($N_c = 1$), and by assuming that the angular distribution of all components is fully deterministic such that all maps in \mathbf{u} are proportional to c_1 . The solution for the unconstrained maps then takes the simplified form

$$\mathbf{u}(\hat{\mathbf{n}}) = \left[((\mathbf{r}^{-1})_{uu})^{-1/2} \mathbf{x} + ((\mathbf{r}^{-1})_{uu})^{-1} (\mathbf{r}^{-1})_{uc} \right] c_1(\hat{\mathbf{n}}), \quad (2.23)$$

where \mathbf{x} is an array of N_u Gaussian random numbers with unit variance and

$$r_{\nu\nu'} = \exp \left[-\frac{1}{2} \left(\frac{\log(\nu/\nu')}{\xi} \right)^2 \right]. \quad (2.24)$$

It is worth noting that, for large correlation lengths, inverting \mathbf{r} above is a numerically unstable operation. PySM solves this by adding a small uncorrelated contribution to the diagonal of \mathbf{r} only large enough to ensure that all of its eigenvalues are strictly positive. This ensures a stable matrix inverse while preserving the large-scale correlation structure. Figure 2.12 shows examples of the corresponding emission laws for different values of ξ , normalised by the fully-correlated case, for a template at $\nu = 300$ GHz. For comparison, ξ can be roughly related to the scatter in the spectral indices as $\xi \simeq \sigma_\beta^2$, and therefore a scatter of $\sigma_\beta \sim 0.2$ (see Section 2.2.2.2) would correspond to $\xi_{\text{dust}} \sim 5$. Finally, the Gaussian model used to generate this stochastic decorrelation may cause the simulated fluctuations to take negative values for small correlation lengths and far




Figure 2.13: Forecasted posterior distribution for r estimated for one of the LiteBIRD-like simulations described in Section 2.5. The black line shows the results assuming perfect knowledge of the foreground spectral indices, while the red line corresponds to the case of spatially-varying spectral indices on pixels of $N_{\text{side}} = 16$, which more than doubles the uncertainty on r .

away from the constrained frequencies. This is an undesirable feature, particularly in intensity (where positivity is a physical requirement), although it should not affect simulations with well-motivated correlation lengths (e.g. see above), assuming the constrained frequencies are sufficiently close to the frequency range under study.

2.5 Discussion

2.5.1 Forecast example

To illustrate the application of PySM, we use it as a basis for forecasting the detectability of primordial tensor perturbations by a LiteBIRD-like satellite experiment [232]. To

Frequency [GHz]	Noise rms (P) [$\mu\text{K arcmin}$]	FWHM [arcmin]
40, 50, 60	53.4, 32.3, 25.1	70
68, 78, 89	19.6, 15.3, 12.4	
100, 119, 140	15.6, 12.6, 8.3	
166, 195, 235	8.7, 6.7, 8.6	
280, 337, 402	19.0, 21.9, 52.3	

Table 2.2: Specifications for LiteBIRD (taken from [231]). Note that noise levels are provided for polarisation. All maps were smoothed to the lowest resolution of 70 arcmin FWHM, which does not correspond to the resolution achievable by LiteBIRD.

do this we generate simulated maps containing Galactic synchrotron (according to Model 1), thermal dust (according to Model 1) and CMB (with tensor-to-scalar ratio $r = 0$), and use the instrument specifications described in Table 2.2 to generate 100 different noise realisations.

We remove the foregrounds from each simulation using the Bayesian component separation algorithm described in [233], fitting for spatially-varying spectral parameters (β_s, β_d, T_d) on HEALPix pixels of size $N_{\text{side}} = 16$ which corresponds to angular scales of ~ 4 deg. Following [233], the foreground-cleaned maps are then used to generate a map of the expected variance after foreground removal by averaging over all simulations. Using the variance map, we generate CMB simulations with the appropriate noise levels, and then estimate the posterior distribution for r for each simulation using a simplified pixel-based likelihood [234], marginalising over the amplitude of lensing B -modes. Specifically, we fit for two parameters, r and A_{lens} , that define the B -mode power spectrum as:

$$C_\ell = rC_\ell^{\text{prim}}(r = 1) + A_{\text{lens}}C_\ell^{\text{B,lens}}, \quad (2.25)$$

where C_ℓ^{prim} and $C_\ell^{\text{B,lens}}$ are templates for the primordial and lensing contributions to the B -mode power spectrum. All other parameters are kept fixed to the best-fit values of [235], except for the optical depth to reionisation, which we fix to $\tau = 0.06$ in agreement with [236]

We compare the resulting posterior distribution for r with the one expected in the absence of foregrounds. This has a noise level of $4.5\mu\text{K-arcmin}$ in polarisation coadded over frequency channels. The result for one of these simulations is shown in Fig. 2.13. We find that, in the ideal case of perfect knowledge of the foreground emission laws, a LiteBIRD-like experiment would be able to achieve a 1σ uncertainty of $\sigma(r) = 1.6 \times 10^{-4}$. We find that this gets degraded after accounting for spatially-varying spectral indices to $\sigma(r) = 3.7 \times 10^{-4}$, for the simulations considered in this

example. This non-negligible inflation of errors demonstrates the importance of including realistic foreground levels in forecasts for future experiments.⁴

2.5.2 Conclusion

We have presented new software to simulate the Galactic microwave sky in polarisation and intensity. The nominal models reflect the current understanding of Galactic foregrounds, and we have included a set of simple alternative models that capture physical extensions to these models and are still consistent with current data. There are many more possible alternatives that are not included, but we provide the public code in a way that makes adding further astrophysical complications straightforward. The code is also fast, portable, and easy to install and begin using.

We have developed methods for the addition of simulated small scale variation in polarisation and intensity, recovering power law behaviour of the polarised components in sky patches of low and high signal, with minimal noise biasing. These simulations may aid in forecasting for ground-based observations limited to partial sky coverage. These small-scale simulations have certain limitations. Different simulated components are not correlated, and the small-scale procedure loses information in high signal-to-noise regions by smoothing at a single scale. Incorporating the spatially varying signal-to-noise into the definition of this smoothing scale would provide more accurate simulations. The small-scales will also be non-Gaussian in practice, which we do not account for.

There are other approaches to foreground modelling. PySM uses 2D sky maps and parametric models to extrapolate single frequency maps to different frequencies, including the possibility of frequency decorrelation. This will be limited in its ability to replicate the polarised nature of Galactic foregrounds. Due to the combination of the complex three-dimensional structure of the Galaxy's magnetic field and the stacking of different sources along any given line of sight we may expect the polarisation fraction of any component to be a function of frequency. Even on a microphysical level there is good evidence that the polarisation spectrum of thermal dust is frequency dependent [10], as carbonaceous and silicate grains may align with the Galactic magnetic field with different efficiencies. More realistic simulations could be derived from three dimensional realisations of the Galaxy's magnetic field and source distributions.

⁴These values were found as the standard deviation of the maximum-likelihood values of r found in the different simulations.

Chapter 3

Finding the chiral gravitational
wave background of an axion-SU(2)
inflationary model using CMB
observations and laser
interferometers

3.1 Introduction

It is the aim of the next generation of CMB experiments to better measure the polarization of the CMB in order to detect its primordial B-mode polarization, parametrized by r , the ratio between tensor and scalar perturbations, which would provide strong evidence for the presence of a primordial gravitational wave background (GWB) (see e.g. [4, 237, 238] for review). Normally, the GWB is produced only by quantum fluctuations of the vacuum during inflation, and is consequently simply related to the energy density of inflation : $\rho_{\text{inf}}^{1/4} \sim \left(\frac{r}{0.01}\right)^{1/4} 1.04 \times 10^{16}$ GeV. A measurement of the power spectrum of tensor perturbations to the metric would therefore be an extremely powerful probe of physics at GUT scales $\sim 10^{16}$ GeV.

The next round of CMB experiments, such as the LiteBIRD satellite [239], and the ground-based Stage-4 [124] effort, seek to push constraints on r to $\sim 10^{-3}$. Interestingly, this search for B-modes may also be sensitive to the dynamics of subdominant fields other than the inflaton, considering the possibility of alternative gravitational wave generation scenarios. Some particular matter fields present during inflation can produce primordial tensor perturbations similar to those sourced by vacuum fluctuations. Therefore, in the event of a detection of r , we must first understand its source.

Recent efforts to provide alternative models for the generation of gravitational waves, which are also consistent with existing observations, have introduced the coupled system of the axion and gauge fields as the spectator sector in addition to the inflaton sector [240, 42, 241, 242, 243, 244]. Such a setup is quite natural from the point of view of particle physics, since many high energy theories contain axion fields and its coupling to some gauge fields, namely the Chern-Simons term: $(\chi/f)F^{\mu\nu}\tilde{F}_{\mu\nu}$. In particular, string theory typically predicts the existence of numerous axion fields. From the view point of low energy effective field theory, at the same time, such dimension five interaction term is expected to exist, because it respects the shift symmetry of the axion field, $\chi \rightarrow \chi + \text{constant}$. Therefore it is strongly motivated to investigate the observational consequence of their dynamics during inflation in light of the role of inflation as a unique probe of high energy physics.

Interestingly enough, the GWB produced by the additional axion-gauge sector has several characteristic features, including non-Gaussianity, scale-dependence, and chirality. A model involving a U(1) gauge field was studied first, and it was confirmed that the resulting GWB is amplified to the same level as the scalar perturbation [240, 245] and hence visible in CMB B-mode observations [246] and interferometer experiments [247]. Recently, a more intriguing model due to a SU(2) gauge field was

also examined, achieving a surpassing GWB production against the scalar sector [42]. This yields more rich phenomenology, and thus motivates us toward the assessment of its detectability.

Gravitational waves may be decomposed into modes with left (L) and right (R) handed polarization. A GWB produced by conventional vacuum fluctuations would have equal amplitudes of L and R, but the effect of the Chern-Simons term in the theory is to allow their amplitudes to differ [248, 240, 42]. Such a chiral GWB would have signatures observable both in CMB polarization and by laser interferometers. CMB polarization may be decomposed into modes of opposing parity: E and B [20, 19]. A detection of a correlation between E and B modes (EB), or between temperature and B modes (TB), would therefore be strong evidence of a parity-violating GWB [248, 43, 44, 249, 246]. To-date observational constraints using the CMB are consistent with no parity-violation and are dominated by systematic uncertainty [249, 250, 251, 252]. An alternative to using the CMB is to directly probe the circular polarization of the GWB, denoted with the circular polarization Stokes parameter $V(f)$, using gravitational interferometers. Interferometers are sensitive to the strain induced in their arms by passing gravitational waves, and for certain detector geometries are sensitive to the polarization of the passing wave [253, 254, 255, 256].

In this chapter we seek to provide a realistic forecast of the ability of LiteBIRD to distinguish this SU(2) model proposed in Ref. [42] from the conventional GWB generation by vacuum fluctuations. LiteBIRD is a proposed CMB satellite mission with the primary science goal of detecting the GWB with $r < 10^{-3}$ [239, 118, 257]. Therefore its sensitivity will be focused in the lowest two hundred multipoles where the B-mode signal is both strong and relatively uncontaminated by gravitational lensing. We exclude Stage 4 from the analysis as we found that the chirality signal is contained in the multipole range $2 \lesssim \ell \lesssim 30$.

Since Stage 4 experiments will have B-mode surveys over the range $\ell \gtrsim 30$ [124], they will be ill-suited to constrain chirality, and we do not consider them further. Ref. [258] consider a simple model for detecting primordial chirality using the CMB, and conclude that ground-based small-scale experiments are not well-suited for pursuing this signal. We also considered a COrE-type experiment ¹, the results of which we do not include in our analysis, as they are similar to LiteBIRD due to the dominant impact of large scale foreground residuals for both instruments. In our analysis we include four contributions to the uncertainty in a measurement of the chiral

¹COrE was a proposed European mission at the time this work was carried out; it is now no longer being considered for funding.

GWB: instrumental noise, foreground residuals from the imperfect cleaning of multi-channel data, gravitational lensing, and the joint self-calibration of the instrument’s polarimeter. This provides a robust assessment of LiteBIRD’s capability to detect primordial chirality.

On the other hand, laser interferometer gravitational wave observatories are sensitive to the GWB today, and provide probes of much smaller scales: $k_{\text{CMB}} \sim 10^{-4} \text{ Mpc}^{-1}$ vs. $k_{\text{interf}} \sim 10^{13} \text{ Mpc}^{-1}$ [247].

In the case of single-field slow-roll inflation, the tensor spectrum is expected to be slightly red-tilted ($n_T = -r/8$, where n_T is the tilt of the tensor spectrum $P_h \sim k^{n_T}$), in which case modern interferometers would not be sensitive enough to make a detection. However, given the scale-dependence of the model of Ref. [42] for part of the parameter space the small scale tensor spectrum is comparatively large. For symmetry reasons the nominal designs of space-based gravitational interferometers are insensitive to the circular polarization of gravitational waves. Since we are interested in constraining chirality we therefore consider ‘advanced’ stages of the nominal design of LISA [259, 260], following the proposed designs of Ref. [256] which provide equal sensitivity to both intensity and polarization of the GWB. In this chapter we show that interferometers and CMB observations provide complementary probes at different scales of the axion-SU(2) ’s primordial tensor spectrum. We then consider the sensitivities of two designs of an advanced stage LISA mission, and compare to constraints achieved using the CMB.

In Section 3.2 we review the model proposed by Ref. [42] and its prediction for the GWB. In Section 3.3 we forecast the ability of a LiteBIRD-like CMB satellite mission to detect the TB and EB correlations expected due to the chiral tensor spectrum, in the presence of foreground contamination, gravitational lensing, instrument noise, and simultaneous self-calibration of the telescope’s polarimeter. In Section 3.4 we analyse the sensitivity of space-based gravitational interferometers to the chiral gravitational background expected by this model. Finally, in Section 3.5 we summarize our findings and discuss our conclusions.

3.2 Theory

In this section we will briefly review the axion-SU(2) model proposed in Ref. [42]. The model is described by the following Lagrangian:

$$\mathcal{L} = \mathcal{L}_{\text{inflaton}} + \frac{1}{2}(\partial_\mu \chi)^2 - \mu^4 \left[1 + \cos\left(\frac{\chi}{f}\right) \right] - \frac{1}{4}F_{\mu\nu}^a F^{a\mu\nu} + \frac{\lambda}{4f}\chi F_{\mu\nu} \tilde{F}^{a\mu\nu}, \quad (3.1)$$

where $\mathcal{L}_{\text{inflaton}}$ denotes the unspecified inflaton sector which realizes inflation and the generation of the curvature perturbation compatible with the CMB observation, χ is a pseudo-scalar field (axion) with a cosine type potential, μ and f are dimensionful parameters and λ is a dimensionless coupling constant between the axion and the gauge field. $F_{\mu\nu}^a \equiv \partial_\mu A_\nu^a - \partial_\nu A_\mu^a - g\epsilon^{abc}A_\mu^b A_\nu^c$ is the field strength of SU(2) gauge field and $\tilde{F}^{a\mu\nu} \equiv \epsilon^{\mu\nu\rho\sigma}F_{\rho\sigma}^a/(2\sqrt{-g})$ is its dual. Here, g is the self-coupling constant of the gauge field and ϵ^{abc} and $\epsilon^{\mu\nu\rho\sigma}$ are the completely asymmetric tensors, $\epsilon^{123} = \epsilon^{0123} = 1$.

In the axion-SU(2) model in the FLRW universe:

$$g_{\mu\nu} = \text{diag}(1, -a^2(t), -a^2(t), -a^2(t)). \quad (3.2)$$

The SU(2) gauge fields naturally take an isotropic background configuration, $A_0^a = 0$, $A_i^a = a(t)Q(t)\delta_i^a$ by virtue of the coupling to the axion χ , and the transverse and traceless part of its perturbation, $t_{ij} = \delta A_j^i$ (not to be confused with the time variable, t), sources gravitational waves at the linear order. Interestingly, either of the two circular polarization modes of t_{ij} , namely t_R or t_L , undergo a transient instability around the horizon crossing and gets substantially amplified. Subsequently, only the corresponding polarization mode of the gravitational wave, h_R or h_L , is significantly sourced by t_{ij} and fully chiral gravitational waves are generated. Note that the parity ($R \leftrightarrow L$) symmetry is spontaneously broken by the background evolution of the axion (i.e. the sign of $\partial_t \chi(t)$). In this chapter we assume the left hand modes are produced for definiteness. In Appendix A, we derive the following expression for the sourced GW power spectrum:

$$\begin{aligned} \mathcal{P}_h^{\text{L, Sourced}}(k) &= r_* \mathcal{P}_\zeta \exp \left[-\frac{1}{2\sigma^2} \ln^2 \left(\frac{k}{k_p} \right) \right] \\ \mathcal{P}_h^{\text{R, Sourced}}(k) &\simeq 0, \end{aligned} \quad (3.3)$$

where the amplitude is parameterized by the tensor-to-scalar ratio r_* at the peak

scale $k = k_p$, σ is the width of the Gaussian-shaped spectrum, and \mathcal{P}_ζ is the power spectrum of curvature perturbations. We treat r_* and σ as free parameters in our analysis, while they can be rewritten in terms of more fundamental parameters m_* , ϵ_{B*} and ΔN , as discussed in Appendix A. Note that, there is no theoretical bound on r_* , while the possible values of σ are restricted by k_p as Eq. (A.11). Figure A.2 gives an example of how the amplitude r_* is degenerate in m_* and ϵ_{B*} , and we show an example plot of $P_h^{\text{Sourced}}(k)$ for three sets of these parameters in Figure 3.1.

Here we define the power spectrum of primordial tensor perturbations to be:

$$\langle h_{\mathbf{k}}^A h_{\mathbf{k}'}^{A'} \rangle = (2\pi)^3 \frac{2\pi^2}{k^3} \mathcal{P}_h^A(k) \delta^{(3)}(\mathbf{k} + \mathbf{k}') \delta_{AA'}, \quad (3.4)$$

where A refers to the circular polarization of the gravitational wave with the momentum vector \mathbf{k} : $A = L, R$. For the rest of this chapter we model the primordial tensor spectrum as being the sum of two contributions: a completely polarized sourced contribution to the tensor spectrum $\mathcal{P}_h^{\text{Sourced}}$: and a contribution from the vacuum fluctuations, which we take to be unpolarized and which we do not vary:

$$\begin{aligned} \mathcal{P}_h^{\text{vac}} &= A_T \left(\frac{k}{k_*} \right)^{n_T} \\ \mathcal{P}_\zeta^{\text{vac}} &= A_S \left(\frac{k}{k_*} \right)^{n_S - 1}, \end{aligned} \quad (3.5)$$

where $A_T = r_{\text{vac}} A_S$, $A_S = 2.2 \times 10^{-9}$, $n_S = 0.96$, $k_* = 0.05 \text{ Mpc}^{-1}$ are taken from the best-fit Planck cosmology [261]. We fix $r_{\text{vac}} = 10^{-5}$ which corresponds to the inflationary Hubble scale $H_{\text{inf}} = 8 \times 10^{11} \text{ GeV}$ and the tensor tilt is given by the consistency relation $n_T = -r_{\text{vac}}/8$. Note that r_{vac} is not required to be so small compared to the sourced contribution; for larger values of r_{vac} the chiral contribution would be more difficult to detect on the CMB due to the vacuum contribution to the BB spectrum. Therefore, we make the simplifying assumption of a small r_{vac} . In summary:

$$\begin{aligned} \mathcal{P}_\zeta(k) &= \mathcal{P}_\zeta^{\text{vac}} \\ \mathcal{P}_h(k, k_p, r_*, \sigma) &= \mathcal{P}_h^{\text{vac}}(k) + \mathcal{P}_h^{\text{Sourced}}(k, k_p, r_*, \sigma) \\ \mathcal{P}_h^L(k) - \mathcal{P}_h^R(k) &= \mathcal{P}_h^{\text{Sourced}}(k, k_p, r_*, \sigma). \end{aligned}$$

It is found that contrary to the tensor perturbation, the scalar perturbations in the axion-SU(2) sector do not have any instability for $m_Q \geq \sqrt{2}$ and they are even




Figure 3.1: The predicted tensor spectrum, $\mathcal{P}_h^{\text{L,Sourced}}$, for three sets of parameters: (blue: $\sigma = 2$, $r_* = 0.07$, $k_p = 0.005 \text{ Mpc}^{-1}$), (orange: $\sigma = 2$, $r_* = 0.07$, $k_p = 0.0005 \text{ Mpc}^{-1}$), (green: $\sigma = 2$, $r_* = 0.07$, $k_p = 7 \times 10^{-5} \text{ Mpc}^{-1}$).

suppressed compared to the vacuum fluctuation of a massless scalar field due to their kinetic and mass mixing [42, 262, 263]. Since the axion- $SU(2)$ sector is decoupled from the inflaton and its energy density is subdominant, its contribution to the curvature perturbation is negligible. It is possible that the energy fraction of the axion $\Omega_\chi \equiv \rho_\chi/\rho_{\text{total}}$ grows after inflation and χ becomes a curvaton if σ is very large and the decay of the axion is suppressed more than that of the inflaton [264, 265, 266]. In that case, the contribution from the scalar perturbations in the axion- $SU(2)$ sector to the curvature perturbation may be significant and hence it would be interesting to investigate such cases. However, it is beyond the scope of this chapter. Therefore, we can simply consider that the curvature perturbation produced by the inflaton is not affected by the axion and the $SU(2)$ gauge fields in this model. We may then take the TT, EE, and TE spectra to be given by constrained cosmological parameters (which we take to be: $h = 0.675$, $\Omega_b = 0.022$, $\Omega_c = 0.12$, $n_s = 0.96$, $\tau = 0.09$, $A_s = 2.2 \times 10^{-9}$), and investigate only the B-mode spectra: TB, EB, and BB.

3.3 CMB

In this section, we study the CMB phenomenology of the model introduced previously in §3.2. The interesting CMB features of this are the non-zero TB and EB spectra produced by the chiral tensor spectrum. We will calculate the expected TB and EB spectra and make forecasts of their detectability by the LiteBIRD satellite in the presence of cosmic-variance, residual foregrounds, instrumental noise, gravitational lensing, and polarimeter self-calibration.

The anisotropies on the CMB are calculated by the integration of the primordial perturbation spectra over the transfer functions describing the evolution of perturbations with time. The tensor contribution to the angular power spectra of the anisotropies are [248, 43, 44, 249]

$$\begin{aligned} C_\ell^{t,X_1X_2} &= 4\pi \int d(\ln k) [\mathcal{P}_h^L(k) + \mathcal{P}_h^R(k)] \Delta_{X_1,\ell}^t(k) \Delta_{X_2,\ell}^t(k), \\ C_\ell^{t,Y_1Y_2} &= 4\pi \int d(\ln k) [\mathcal{P}_h^L(k) - \mathcal{P}_h^R(k)] \Delta_{Y_1,\ell}^t(k) \Delta_{Y_2,\ell}^t(k). \end{aligned} \quad (3.6)$$

where $X_1X_2 = \{TT, TE, EE, BB\}$ and $Y_1Y_2 = \{TB, EB\}$, and $\Delta_{X,\ell}^t(k)$ indicates the tensor transfer function [267]. To calculate these spectra we use the CLASS code [268], making the necessary modifications for it to calculate TB and EB spectra. In Figure 3.2 we plot examples of the BB and TB spectra calculated in this way for a few different combinations of the model parameters, and compare them to the noise contributions from lensing, instrument noise and foreground residuals that we will consider later.

In this chapter we assess the detectability of the chirality of the primordial GWB over the parameter space spanned by (r_*, k_p, σ) . Therefore, we calculate the expected signal-to-noise of the combined TB and EB spectra [246]:

$$\left(\frac{S}{N}\right)_{\text{TB+EB}}^2 = \sum_{\ell=2}^{\ell_{\max}} \sum_{\substack{X_1X_2 \\ X_3X_4}} C_\ell^{X_1X_2} [\xi^{-1}]_\ell^{X_1X_2, X_3X_4} C_\ell^{X_3X_4}, \quad (3.7)$$

where $X_1X_2, X_3X_4 = \{TT, EE, BB, TE, TB, EB\}$, and ξ is the covariance of our estimate of the power spectra given a certain theoretical and experimental setup: $\xi^{X_1X_2X_3X_4} = \langle (\hat{C}_\ell^{X_1X_2} - C_\ell^{X_1X_2})(\hat{C}_\ell^{X_3X_4} - C_\ell^{X_3X_4}) \rangle = \frac{1}{(2\ell+1)f_{\text{sky}}} (\tilde{C}_\ell^{X_1X_3} \tilde{C}_\ell^{X_2X_4} + \tilde{C}_\ell^{X_1X_4} \tilde{C}_\ell^{X_2X_3})$, where tildes indicate the observed spectrum: $\tilde{C}_\ell^{XX'} = C_\ell^{XX'} + N_\ell^{XX'}$, with $N_\ell^{XX'}$ de-

noting the noise spectrum, and the calculation of ξ is detailed in Appendix B. ℓ_{\max} denotes the highest multipole we consider, which in this case is 500.

Similarly, we can calculate the detectability of the primordial GWB, as opposed to its chirality, by calculating the signal-to-noise of its contribution to the BB spectrum. In the case of no lensing, this is simply:

$$\left(\frac{S}{N}\right)_{\text{BB}}^2 = f_{\text{sky}} \sum_{\ell=2}^{\ell_{\max}} \frac{(2\ell+1)}{2} \left[\frac{C_{\ell}^{BB}}{\tilde{C}_{\ell}^{BB}} \right]^2, \quad (3.8)$$

However, one of the major sources of uncertainty in a measurement of the BB spectrum is due to gravitational lensing. As the CMB propagates to us from the surface of last scattering it is gravitationally lensed by the intervening matter density, converting primary E-mode anisotropies to secondary B-mode anisotropies, which then need to be accounted for in measurements of BB [79].

We can separate the contributions to BB into $C_{\ell}^{BB} = C_{\ell}^{BB, \text{Prim}} + C_{\ell}^{BB, \text{Lens}}$, where ‘Prim’, ‘Lens’ refer to the primordial and lensed contributions respectively. We are interested in measuring $C_{\ell}^{BB, \text{Prim}}$, and in effect $C_{\ell}^{BB, \text{Lens}}$ acts as an extra source of noise, with an unknown amplitude. The modification required to calculate the signal-to-noise of the primordial BB signal is to consider the 2×2 matrix:

$$\mathcal{F}_{ij} = \sum_{\ell=2}^{\ell_{\max}} \frac{(2\ell+1)f_{\text{sky}}}{2} \frac{C_{\ell}^{BB, i} C_{\ell}^{BB, j}}{(\tilde{C}_{\ell}^{BB})^2}, \quad (3.9)$$

such that:

$$\left(\frac{S}{N}\right)_{\text{BB}, i}^2 = \frac{1}{(\mathcal{F}^{-1})_{ii}} \quad (3.10)$$

where the indices i, j run pver ‘Prim’, ‘Lens’. Note that we will assume that the temperature spectrum is perfectly known over the range of scales we are interested in, and that the sourced contribution to the scalar spectrum is negligible [42]: $\tilde{C}_{\ell}^{\text{TT}} = C_{\ell}^{\text{TT}}$.

3.3.1 Cosmic-variance limited case

Here, we discuss the signal-to-noise of the TB, EB, and BB spectra in the case of cosmic variance-limited observations: $\tilde{C}_{\ell}^{XX'} = C_{\ell}^{XX'}$. In this scenario, in the absence of lensing, Equation 3.8 has the simple analytic form $\left(\frac{S}{N}\right)_{\text{BB}}^2 = f_{\text{sky}}(\ell_{\max} + 3)(\ell_{\max} - 1)/2$. The signal-to-noise of the TB and EB spectra calculated using Equation 3.7 are




Figure 3.2: LEFT PANEL: C_ℓ^{BB} for the same three sets of parameters used in Figure 3.1: (blue: $\sigma = 2$, $r_* = 0.07$, $k_p = 0.005 \text{ Mpc}^{-1}$), (orange: $\sigma = 2$, $r_* = 0.07$, $k_p = 0.0005 \text{ Mpc}^{-1}$), (green: $\sigma = 2$, $r_* = 0.07$, $k_p = 7 \times 10^{-5} \text{ Mpc}^{-1}$) compared to the LiteBIRD noise spectrum with 2% foregrounds (solid black), the lensing BB spectrum (dashed black), and the standard vacuum fluctuation $C_\ell^{BB}(r = 0.07)$ consistent with the BKP $r < 0.07$ (95% C.L.) (dash-dot black). The axion-SU(2) spectra contain a contribution from vacuum fluctuations with $r = 10^{-5}$, as is used in the text. RIGHT PANEL: $|C_\ell^{TB}|$ (solid colour) and $|C_\ell^{EB}|$ (dashed colour) spectra for the same three sets of parameters. Shown in black is an example of the spurious TB signal induced by polarimeter miscalibration for an angle of one arcminute, as discussed in Section 3.3.3.

shown in Figure 3.3 for the parameter space of the model, assuming a lensed BB spectrum with $f_{\text{sky}} = 1$. We consider only $r_* < 0.07$, in line with current observational constraints on the scale-invariant tensor-to-scalar ratio $r_{0.05} < 0.07$ (95% C.L.), where the subscript indicates the pivot scale in Mpc^{-1} [51]. Figure 3.2 demonstrates that the shape of C_ℓ^{BB} is strongly dependent on the position of the peak in the GW spectrum, k_p , and also on the width of the peak, σ . Therefore, the BKP bound on r does not simply imply the same bound on r_* ; a small value of k_p and σ could allow a large value of r_* without exceeding the BKP limit, due to the small scale damping of C_ℓ^{BB} . However, excepting $\mathcal{O}(1)$ underestimation for small k_p , the BKP bound provides a useful guide as to what is allowed by current observations.

The values of $k_p = 7 \times 10^{-5} \text{ Mpc}^{-1}$ and $k_p = 5 \times 10^{-3} \text{ Mpc}^{-1}$ were chosen as they probe different scales to which the CMB is sensitive. $\mathcal{P}_h^{\text{L, Sourced}}(k)$ is more sharply peaked for smaller σ and so for a given r_* the signal-to-noise decreases with decreasing σ . As σ increases the tensor spectrum becomes almost scale-invariant over the range of scales accessible with the CMB and so the signal-to-noise does not depend on σ for large values of σ . Figure 3.3b shows that the maximum achievable signal-to-noise is ~ 3 and that the chirality is undetectable with $\frac{S}{N} \lesssim 1$ for $r_* \lesssim 0.01$.

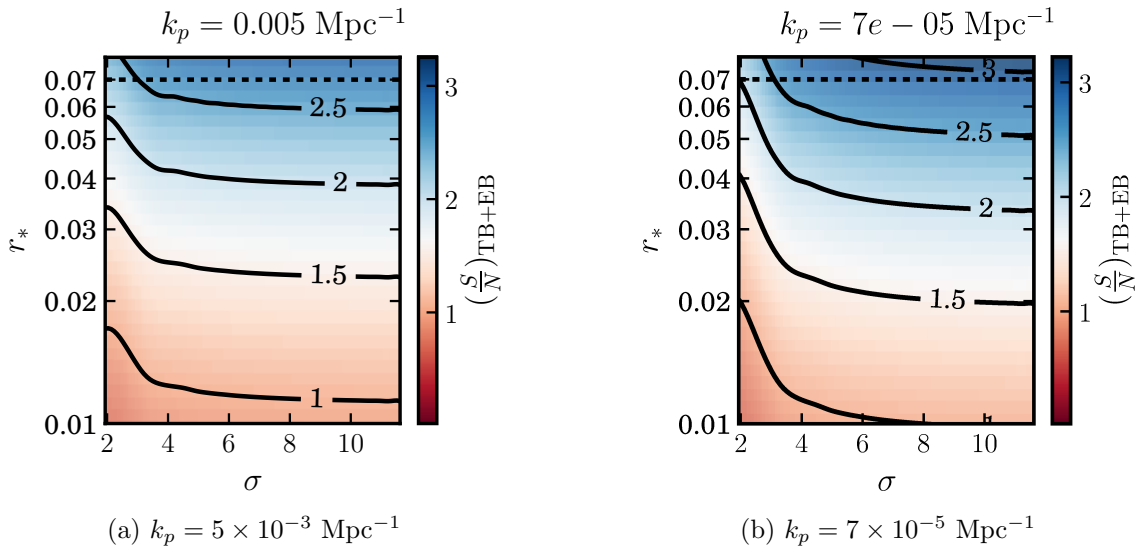


Figure 3.3: Signal-to-noise of TB + EB spectra assuming the perfect case of $f_{\text{sky}} = 1$, with no foreground contamination, noiseless observations, and no delensing. The black dashed line indicates the bounds placed by $r_* < 0.07$.

3.3.2 Including instrument noise and foreground contamination

We now consider instrument noise, contamination of the spectrum due to imperfect foreground separation, and assume that we are unable to perform any ‘delensing’.

The model we use for the noise spectrum includes the instrument noise in the CMB channels, the residual foregrounds in the final CMB map (assumed to be at a level of 2 %, following Refs. [118, 269, 246, 270]) and the instrumental noise from channels used for foreground cleaning that is introduced into the CMB channels by the cleaning process. The details of how we combine these factors to produce a final noise contribution to the measured CMB spectrum, as well as the instrument specifications for LiteBIRD can be found in Appendix C. In the left panel of Figure 3.2 we show the contributions to the BB noise spectrum, N_ℓ^{BB} , from lensing, LiteBIRD instrumental noise, and foreground residuals compared to the primordial C_ℓ^{BB} .

3.3.2.1 BB Signal-to-Noise

We calculate Equation 3.10 over the available parameter space and show the result in Figure 3.4. In a similar way to the TB and EB signal-to-noise we see that there is some dependence on σ , especially in the case of smaller k_p . This is expected since $k_p = 7 \times 10^{-5} \text{ Mpc}^{-1}$ is slightly smaller than those scales to which we expect the CMB to be sensitive [247]. Therefore, we expect that reducing σ for this value of k_p will eventually exclude the tensor perturbations from contributing to CMB scales, explaining the sharp decrease in S/N for low σ and a given r_* . From Figure 3.4 it is clear that we can detect the primordial contribution to BB for $r_* > 10^{-3}$, which is consistent with the aim of LiteBIRD to achieve an uncertainty on the null case of $r = 0$ of less than 10^{-3} .

3.3.2.2 TB+EB Signal-to-Noise

Lensing affects the TB and EB signal-to-noise only through $\tilde{C}_\ell^{\text{BB}}$, since the direct lensing contributions to TB and EB are negligible [43, 246]. We calculate Equation 3.7 over the available parameter space, now including instrument noise for a LiteBIRD-type experiment (with parameters shown in Table C.2), foreground residuals, and gravitational lensing, and show the result in Figure 3.5. Over the allowed parameter space, the maximum achievable signal-to-noise is $\frac{S}{N} \sim 2$. Whilst for $r_* \lesssim 0.03$

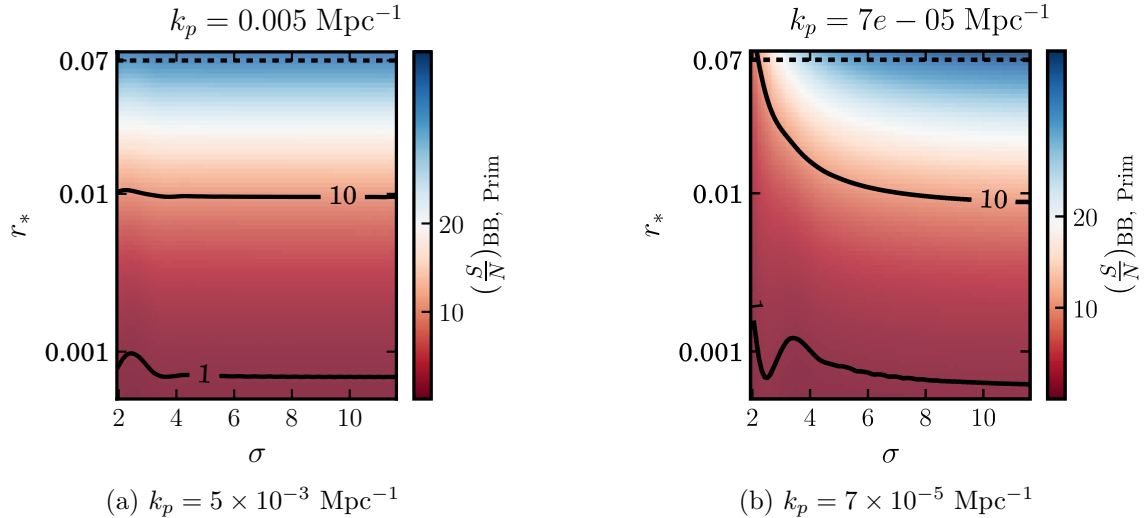


Figure 3.4: Signal-to-noise of BB spectrum assuming no delensing and 2% foreground contamination and LiteBIRD instrumental noise added using method described in Appendix C. The dashed line refers to the observational constraint of $r_* = 0.07$. The signal-to-noise achieved in BB is much larger than that in TB+EB as the cosmic variances of BB and TB are proportional to $(C_\ell^{BB})^2$, and $C_\ell^{BB}C_\ell^{TT}$, respectively. The factor of C_ℓ^{TT} means cosmic variance in the TB spectrum is much more significant than in BB.

LiteBIRD can not detect chirality in this model, compared to $r_* \lesssim 0.01$ in the CV-limited case. The right panel of Figure 3.2 demonstrates that the TB and EB signal peaks at $\ell \lesssim 10$, making the large scale foreground residual contribution to the noise, shown in the left panel of Figure 3.2, the dominant factor causing this reduction in sensitivity.

Improvements in foreground cleaning algorithms could reduce the level of foreground contamination, and perhaps allow a larger sky fraction to be used in the analysis. However, even with perfect control of these factors, the cosmic-variance limit of Figure 3.3 can not be beaten. We conclude from this study that the most important factor limiting the sensitivity of CMB observations to the chirality of the GWB is the large cosmic variance of the TB and EB spectra due to large scalar T and E signals, respectively.

3.3.3 Simultaneous detection and self-calibration

In order to achieve its baseline performance target, LiteBIRD will need to achieve an uncertainty on the polarimeter calibration angle of less than one arcminute [271, 272].




Figure 3.5: Signal-to-noise of TB + EB spectra assuming no delensing and 2% foreground contamination and LiteBIRD instrumental noise added using method described in Appendix C. The dashed line refers to the observational constraint of $r_* = 0.07$.

There are several methods that have been used in the past to calibrate polarimeters such as polarized astrophysical sources like the Crab Nebula (Tau A), or man-made sources such as a polarization selective mesh. There are many factors preventing such methods achieving calibrations better than one degree. For example, Tau A is the best candidate for a point-like polarized source, but this provides a calibration uncertainty of ~ 0.5 degrees [273], and with these it is hard to achieve a calibration uncertainty better than one degree [274]. The polarization of Tau A also has a poorly understood frequency dependence, and is ultimately an extended source, making it poorly suited to a characterization of the polarized beam [275]. Man-made sources on the other hand must often be placed in the near field and are unstable over long time frames. However, a recent proposal of a balloon-borne artificial polarization source in the far field of ground-based experiments may ameliorate this problem for ground-based telescopes [275].

LiteBIRD plans to self-calibrate its polarimeter using the EB spectrum, which is assumed to have zero contribution from primordial perturbations [276]. Unfortunately this makes assumptions about cosmology, and uses part of the constraining power to calibrate the instrument, instead of for science. Furthermore, residual foreground contributions to TB and EB may result in a biasing of the calibration angle. Ref. [277]

shows that a miscalibration angle of 0.5 degrees can result in a bias in the recovered value of r of 2×10^{-3} , which is significant for LiteBIRD's aim to push constraints on r to $r \sim 10^{-3}$. However, Ref. [277] also finds that TB and EB are consistent with zero in a study of the low-foreground BICEP2 region. Furthermore, in a study of the Planck data Ref. [205] finds that TB and EB are both consistent with zero for sky fractions up to $f_{\text{sky}} = 0.3$, and that TB increases to significant levels only for larger sky fractions, whilst EB is only marginally non-zero for $f_{\text{sky}} = 0.7$. Therefore whilst foregrounds must be considered, they do not necessarily limit the use of this approach to calibration.

We want to study the detectability of primordial TB and EB correlations when taking self-calibration into account. The self-calibration process is carried out by zeroing the miscalibration $\Delta\psi$ by measuring its contribution to the TB and EB spectra. In this analysis we will assume that residual foreground contributions to TB and EB are negligible.

If the angle of the polarimeter is miscalibrated by some angle $\Delta\psi$ the measured Q , U will be rotated. We work with the spin-2 quantities $(Q \pm iU)(\hat{n})$ which have the transformation properties under rotation:

$$(\tilde{Q} \pm i\tilde{U})(\hat{n}) = e^{\pm i2\Delta\psi}(Q \pm iU)(\hat{n}).$$

E and B can be computed to find:

$$\begin{pmatrix} \tilde{a}_{\ell m}^T \\ \tilde{a}_{\ell m}^E \\ \tilde{a}_{\ell m}^B \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(2\Delta\psi) & -\sin(2\Delta\psi) \\ 0 & \sin(2\Delta\psi) & \cos(2\Delta\psi) \end{pmatrix} \begin{pmatrix} a_{\ell m}^T \\ a_{\ell m}^E \\ a_{\ell m}^B \end{pmatrix}$$

which give the resulting rotations of the angular power spectra:

$$\begin{aligned} \begin{pmatrix} C_{\ell}^{TE} \\ C_{\ell}^{TB} \end{pmatrix}_{\text{rot}} &= \begin{pmatrix} \cos(2\Delta\psi) & -\sin(2\Delta\psi) \\ \sin(2\Delta\psi) & \cos(2\Delta\psi) \end{pmatrix} \begin{pmatrix} C_{\ell}^{TE} \\ C_{\ell}^{TB} \end{pmatrix} \\ \begin{pmatrix} C_{\ell}^{EE} \\ C_{\ell}^{BB} \\ C_{\ell}^{EB} \end{pmatrix}_{\text{rot}} &= \begin{pmatrix} \cos^2(2\Delta\psi) & \sin^2(2\Delta\psi) & -\sin(4\Delta\psi) \\ \sin^2(2\Delta\psi) & \cos^2(2\Delta\psi) & \sin(4\Delta\psi) \\ \frac{\sin(4\Delta\psi)}{2} & -\frac{\sin(4\Delta\psi)}{2} & \cos(4\Delta\psi) \end{pmatrix} \begin{pmatrix} C_{\ell}^{EE} \\ C_{\ell}^{BB} \\ C_{\ell}^{EB} \end{pmatrix}. \end{aligned} \quad (3.11)$$

We then replace the primordial spectra in our expression for \tilde{C}_ℓ with the rotated spectra:

$$\tilde{C}_\ell^{XX'} = C_{\text{rot}\ell}^{XX'} + N_\ell^{XX'}.$$

We jointly estimate the uncertainty on the miscalibration angle and the recovered amplitude of the TB and EB spectra parametrized by r_* using the Fisher information:

$$\mathcal{F}_{ij} = \sum_{X_1 X_2, X_3 X_4} \sum_{\ell=2}^{\ell_{\max}} \frac{\partial C_\ell^{X_1 X_2}}{\partial a_i} [\xi^{-1}]^{X_1 X_2 X_3 X_4} \frac{\partial C_\ell^{X_3 X_4}}{\partial a_j} \quad (3.12)$$

where $a_i, a_j = \Delta\psi, r_*$. The uncertainty on the miscalibration angle is then given by $(\sigma_{\Delta\psi})^2 = (\mathcal{F}^{-1})_{\Delta\psi\Delta\psi}$. We can easily calculate the derivatives with respect to $\Delta\psi$ in Equation 3.12 using Equation 3.11. In order to calculate the derivatives with respect to r_* we write: $C_\ell^{TB/EB} = r_* C_\ell^{TB/EB}(r_* = 1)$. In order to study the interaction of the miscalibration angle and primordial chirality we calculate the correlation coefficient:

$$\alpha \equiv \frac{\mathcal{F}_{\Delta\psi r_*}}{\sqrt{\mathcal{F}_{\Delta\psi\Delta\psi} \mathcal{F}_{r_* r_*}}}.$$

We now calculate the 1-sigma uncertainty in a measurement of the miscalibration angle $\Delta\psi$ and α over the allowed parameter space of the model and show the resulting contour plots in Figure 3.6.

We find that for LiteBIRD $\sigma_{\Delta\psi} < 1$ arcmin for all of the allowed space, making the simultaneous calibration of the polarimeter and detection of the parity-violation possible. The correlation coefficient is less than 0.03 for the allowed parameter space, indicating that the effects of primordial parity-violation and miscalibration are easily separable. This can be understood from the right panel of Figure 3.2 where it is clear that the primordial signal is a large scale effect, with maximum signal at $\ell \sim 10$, whereas the contribution to TB from miscalibration is a small scale effect which dominates at $\ell > 100$. This is supported by the σ dependence of α in the left panel of Figure 3.6. The two effects become more correlated for larger values of σ which correspond to flatter spectra, and hence more power at small scales. Varying k_p has little effect on the result that the effects are separable, but does introduce some interesting dependence on σ . This indicates that a sufficiently high ℓ_{\max} is necessary for the separation of these effects. For smaller values of k_p , $\sigma_{\Delta\psi}$ becomes more dependent on σ . For example, with $k_p = 7 \times 10^{-5} \text{ Mpc}^{-1}$, for a given r_* , $\sigma_{\Delta\psi}$ increases with σ since the flatter spectra of large σ become more important when k_p is further away

from the small scales at which the miscalibration effect occurs. On the other hand when $k_p = 5 \times 10^{-3} \text{ Mpc}^{-1}$ the dependence on σ is reversed. This is because the miscalibration effect peaks at $\ell \sim 100$, which corresponds to contributions from modes around $k \sim 100/\eta_0 = 7 \times 10^{-3} \text{ Mpc}^{-1}$, where η_0 is the comoving distance to the surface of last scattering. Therefore an increase of σ for $k_p \sim 7 \times 10^{-3} \text{ Mpc}^{-1}$ will make the signals less correlated as the flatter spectra will introduce more power at larger scales.

In conclusion, any reduction in sensitivity to TB and EB due to the calibration requirements is negligible, and is ignored in the results we quote for LiteBIRD. Our results are in agreement with Refs [258, 44], which also find that the primordial and miscalibration contributions are readily separable.




Figure 3.6: Correlation coefficient α (left) and 1- σ uncertainty on the polarimeter calibration (right) for LiteBIRD, calculated over the allowed parameter space of the model assuming 2% foreground residual and no delensing. The dashed line shows the observational constraint of $r_* = 0.07$.

3.3.4 CMB Results

Here, we summarize the findings of the CMB section and provide a prognosis of the usefulness of CMB observations in detecting gravitational wave chirality.

In the case of cosmic variance-limited ultimate observations we found that over the parameter space of the model the maximum signal-to-noise achievable was ~ 3 for the largest values of r_* , and that the chirality is undetectable for $r_* \lesssim 0.01$, in agreement

with previous studies of simpler models of chiral GWBs with nearly scale-invariant spectra [44, 249, 43]. Moving on to the realistic case of a LiteBIRD-like experiment with no delensing capability, a 2% level of foreground residuals, and a simultaneous self-calibration, we find that for the largest allowable values of r_* it may achieve a signal-to-noise of 2.0, making the chirality detectable. The chirality is undetectable by LiteBIRD for $r_* \lesssim 0.03$.

Though a detection with a two sigma significance may be of interest, it is only achievable for a small part of the parameter space, $0.03 \lesssim r_* \lesssim 0.07$, and in any event we have demonstrated that we may not exceed a $\frac{S}{N}$ of 3 using CMB two-point statistics. We also investigated a CoRE+ design with the same level of foreground residuals as LiteBIRD and found that is performed very similarly to LiteBIRD since both instruments would be limited by foreground residuals on the large scales we are interested in. As stated in §3.1 we will not gain anything extra from Stage 4 observations, as they are limited to $\ell \gtrsim 30$. Therefore, in order to make stronger statistical detections of this model using the CMB, higher order statistical techniques taking advantage of the model’s non-Gaussianity may have more success as shown for the axion-U(1) model [246].

Alternatively, we can investigate different physical probes altogether. In the next section, we consider complementary constraints on the axion-SU(2) model from space-based laser interferometer gravitational wave observatories.

3.4 Laser Interferometers

Due to the strong scale-dependence of the tensor spectrum, it may be possible to study the case of large k_p using laser interferometer gravitational wave observatories. Previous studies have indicated that the scale-invariant spectrum of single-field slow-roll inflation would be too weak at interferometer scales to be detected by current generation interferometers such as LIGO [278], VIRGO [279], and LISA [260]. However, the model we consider has a large feature at k_p , therefore for $k_p \sim 10^{11} - 10^{13} \text{ Mpc}^{-1}$, current generation interferometers may be sensitive to the GWB of the axion-SU(2) model.

It should be noted that it is difficult to have a sourced gravitational wave spectrum with a sharp peak on interferometer scales. This is because of the attractor behaviour of the background axion field coupled to the SU(2) gauge fields (see Appendix A). As a result, we consider the rather flat spectra seen in Figure 3.7. For such flat




Figure 3.7: Comparison of the sensitivity curves for LiteBIRD, Planck, LISA, and BBO corresponding to a signal-to-noise of one at a given frequency in intensity (I) or polarization (V). At the top horizontal axis we also show the corresponding wavenumber computed via $\frac{k}{\text{Mpc}^{-1}} = 6.5 \times 10^{14} \frac{f}{\text{Hz}}$. Also plotted are the primordial spectra for the parameters: $k_p = 10^{13} \text{ Mpc}^{-1}$, $\sigma = 9$, $r_* = 835$ (black dotted) and $k_p = 10^{11} \text{ Mpc}^{-1}$, $\sigma = 8$, $r_* = 0.15$ (black dash-dotted). Note that below $f = 10^{-17} \text{ Hz}$ transfer function of the fractional gravitational wave background energy density changes due to the transition between matter and radiation dominated eras. We see that even for the large values of σ required by the large k_p values of the axion-SU(2) model LISA and BBO can make a detection that would still be inaccessible at CMB scales. This motivates the evaluation of signal-to-noise for the interferometers. Note that the CMB sensitivity curves have been smoothed with a Gaussian kernel due to the sharp oscillations introduced by the transfer function (Equation D.2).

spectra one may expect any signal detectable with interferometers would also be detectable on CMB scales, making the use of interferometers redundant. We therefore first demonstrate the complementarity of our CMB and interferometer studies. We compare their sensitivities as a function of the frequency f of the gravitational wave background. The quantity we use to compare sensitivities is the minimum detectable fractional energy density in primordial gravitational waves today:

$$\Omega_{\text{GW}}(f) \equiv \frac{1}{\rho_c} \frac{\partial \rho_{\text{GW}}}{\partial \ln(f)} \quad (3.13)$$

where ρ_c is the critical density to close the Universe evaluated today, and $\rho_{\text{GW}} = \frac{c^2}{32\pi^2 G} \langle \dot{h}_{ij} \dot{h}_{ij} \rangle$, where $h_{ij} \equiv \delta g_{ij}^{TT}/a^2$. The calculation for the CMB is detailed in Appendix D, and for interferometers in the remainder of this section. Figure 3.7 displays the minimum detectable fractional energy density using the CMB and interferometers for Planck, LiteBIRD, an advanced LISA [260] and BBO [280]. We see that LiteBIRD has a much improved sensitivity to chirality, compared to Planck, which is due to its much lower instrumental noise. The two plotted theoretical spectra are clearly detectable by LISA or BBO, without being detectable at CMB scales, making interferometers an independent, complementary probe of the primordial spectrum of the axion-SU(2) model.

3.4.1 Interferometer notation

Laser interferometers consist of a set of test masses placed at nodes and linked by laser beams. Interferometry is used to measure the change in the optical path length between test masses. A passing gravitational wave induces a time-dependent oscillation in the optical path length, which can be isolated from noise by taking cross-correlations between detectors.

The metric perturbation at point \mathbf{x} at time t , $h_{ij}(t, \mathbf{x})$, can be decomposed into a superposition of plane waves [281]:

$$h_{ij}(t, \mathbf{x}) = \sum_P \int d^3\mathbf{k} C_P(\mathbf{k}) \sin(ckt - \mathbf{k} \cdot \mathbf{x} + \Phi(k)) e_{ij}^P(\hat{\Omega}),$$

where we use the transverse traceless basis tensors with normalization $e_{ij}^P(\hat{\Omega}) e_{ij}^{P'}(\hat{\Omega}) = 2\delta_{PP'}$, and $P = +, \times$. It is more convenient to deal with complex values, and so we

rewrite this as:

$$h_{ij}(t, \mathbf{x}) = \sum_P \int_{-\infty}^{\infty} df \int d^2\hat{\Omega} h_P(f, \hat{\Omega}) \exp(2\pi i f(t - \frac{\mathbf{x} \cdot \hat{\Omega}}{c})) e_{ij}^P(\hat{\Omega}),$$

where $ck = 2\pi f$, $\mathbf{k} \cdot \mathbf{x} = 2\pi f \frac{\hat{\Omega} \cdot \mathbf{x}}{c}$, and $\hat{\Omega}$ is a unit vector in the direction of propagation of the gravitational wave. Since the coefficients satisfy $h_P(f, \hat{\Omega}) = h_P^*(-f, \hat{\Omega})$, $h_{ij}(t, \mathbf{x})$ is explicitly real. The theory we are dealing with produces a highly non-Gaussian GWB [282]. We can summarize the two-point statistics using the following expectation values of the Fourier coefficients, but this will not capture all the available information:

$$\begin{aligned} \langle h_P(f, \hat{\Omega}) h_{P'}^*(f', \hat{\Omega}') \rangle &= \frac{1}{2} \delta(f - f') \frac{\delta^{(2)}(\hat{\Omega} - \hat{\Omega}')}{4\pi} S_h^{PP'}(f) \\ \begin{pmatrix} \langle h_+(f, \hat{\Omega}) h_+^*(f', \hat{\Omega}') \rangle & \langle h_+(f, \hat{\Omega}) h_\times^*(f', \hat{\Omega}') \rangle \\ \langle h_\times(f, \hat{\Omega}) h_+^*(f', \hat{\Omega}') \rangle & \langle h_\times(f, \hat{\Omega}) h_\times^*(f', \hat{\Omega}') \rangle \end{pmatrix} &= \frac{1}{2} \delta(f - f') \frac{\delta^{(2)}(\hat{\Omega} - \hat{\Omega}')}{4\pi} \\ &\times \begin{pmatrix} I(f) & iV(f) \\ -iV(f) & I(f) \end{pmatrix}, \end{aligned} \quad (3.14)$$

where $I(f)$ and $V(f)$ are the Stokes parameters for intensity and circular polarization respectively. As shown below, $V(f)$ quantifies the difference between the amplitudes of two circular polarization states and hence is a clean observable for the chiral GWB [253, 254, 255].

3.4.2 Interferometer response

In this section, we present the design of the interferometers for which we will forecast the sensitivity to a polarized gravitational wave background. This analysis uses the designs proposed by Ref. [256]. We summarise some of the results of Ref. [256] here, however for further details we refer readers to Ref. [256].

Let us consider a set of masses placed at positions \mathbf{x}_i , and the phase change, $\Delta\phi_{ij}$, of light as it travels from mass i at time t_i to mass j arriving at time t [283]:

$$\Delta\phi_{ij}(t) = \int_{-\infty}^{\infty} df \int d^2\hat{n} \sum_P \tilde{h}_P(f, \hat{n}) e_{ab}^P(\hat{n}) e^{i2\pi f(t_i - \hat{n} \cdot \mathbf{x}_i)} D^{ab}(\hat{u}_{ij} \cdot \hat{n}, f), \quad (3.15)$$

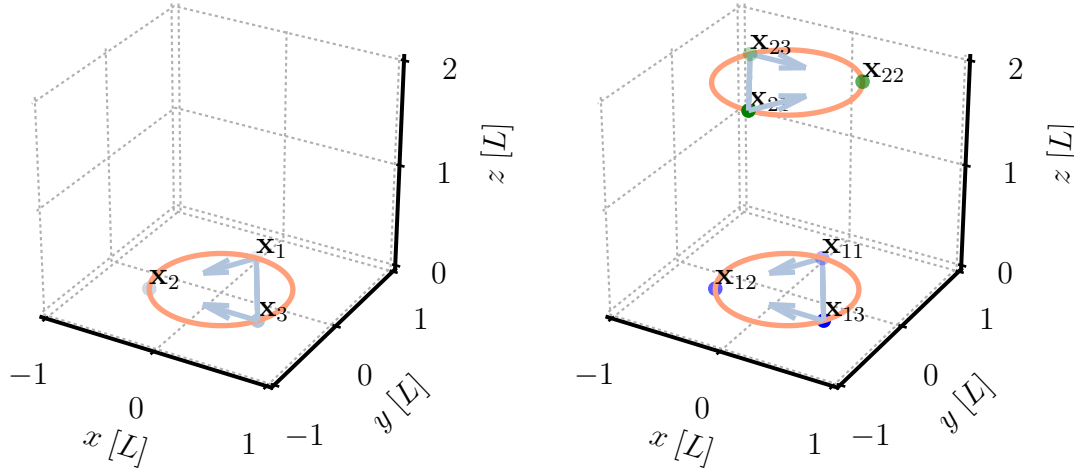


Figure 3.8: Possible designs for future space-based laser interferometers. The blue arrows show the laser links used in the Michelson interferometer signals defined in Equation 3.19. LEFT PANEL: One constellation design of a space-based interferometer, this corresponds to the baseline LISA design. The points \mathbf{x}_i show the i^{th} satellite. RIGHT PANEL: An advanced stage design of LISA or BBO with two constellations. The points \mathbf{x}_{ij} show the j^{th} satellite on the i^{th} constellation.

where D^{ab} is the single-arm transfer function which contains all the geometric information about the instrument and must be derived individually for each interferometer set-up [284], and \hat{u}_{ij} is a unit vector pointing from detector i to detector j . We now define the Fourier transform of a signal $g(t)$ observed for a time T : $g(f) = \int_{-T/2}^{T/2} g(t) \exp(-2\pi i f t) dt$. The Fourier transform of the phase change $\Delta\phi$ is then:

$$\begin{aligned}
\Delta\phi_{ij}(f) &= \int_{-T/2}^{T/2} dt \int_{-\infty}^{\infty} df' \int d^2\hat{\Omega} \sum_P h_P(f', \hat{\Omega}) \exp\left(i2\pi f'(t - \frac{\mathbf{x} \cdot \hat{\Omega}}{c}) - 2\pi i f t\right) \\
&\quad \times D^{ab}(\hat{u}_{ij} \cdot \hat{n}, f') \\
&= \int_{-\infty}^{\infty} df' \delta_T(f - f') \int d^2\hat{\Omega} \sum_P h_P(f', \hat{\Omega}) \exp\left(-i2\pi f' \frac{\mathbf{x} \cdot \hat{\Omega}}{c}\right) \\
&\quad \times D^{ab}(\hat{u}_{ij} \cdot \hat{n}, f'),
\end{aligned} \tag{3.16}$$

where δ_T is a finite-time approximation to the delta function defined as $\delta_T(f - f') \equiv T \text{sinc}(\pi T(f - f'))$, with the properties: $\delta_T(0) = T$, $\lim_{T \rightarrow \infty} \delta_T(f) \rightarrow \delta(f)$. We may

form a signal by constructing a linear combination of phase changes along different paths around the instrument, and then cross-correlating these signals. The signal we seek to measure is stochastic and so to distinguish it from noise we must cross-correlate the detector output with the output from a detector with independent noise properties. The expectation of the cross correlated signal will be composed of terms like:

$$\begin{aligned}
\langle \Delta\phi_{ij}(f_1)\Delta\phi_{kl}(f_2) \rangle &= \int_{-\infty}^{\infty} df' \int_{-\infty}^{\infty} df'' \int d^2\hat{\Omega}_1 \int d^2\hat{\Omega}_2 \sum_{P_1 P_2} \delta_T(f_1 - f')\delta_T(f_2 - f'') \\
&\times \langle h_{P_1}(f', \hat{\Omega}_1)h_{P_2}(f'', \hat{\Omega}_2) \rangle \\
&\times \exp\left(-2\pi i f' t \hat{\Omega}_1 \cdot \mathbf{x}_i\right) \exp\left(-2\pi i f'' t \hat{\Omega}_2 \cdot \mathbf{x}_k\right) \\
&\times D^{ab}(\hat{u}_{ij} \cdot \hat{\Omega}_1, f') D^{cd}(\hat{u}_{kl} \cdot \hat{\Omega}_2, f'') e_{ab}^{P_1}(\hat{\Omega}_1) e_{cd}^{P_2}(\hat{\Omega}_2).
\end{aligned} \tag{3.17}$$

Using $\langle h_{P_1}(f', \hat{\Omega}_1)h_{P_2}(f'', \hat{\Omega}_2) \rangle = \langle h_{P_1}(f', \hat{\Omega}_1)h_{P_2}^*(-f'', \hat{\Omega}_2) \rangle$, and $D^{ab}(\hat{u}_{ij} \cdot \hat{\Omega}, -f) = D^{ab*}(\hat{u}_{ij} \cdot \hat{\Omega}, f)$ we can write this as:

$$\begin{aligned}
\langle \Delta\phi_{ij}(f_1)\Delta\phi_{kl}(f_2) \rangle &= \frac{1}{2} \int_{-\infty}^{\infty} df' \delta_T(f_1 - f')\delta_T(f_2 - f') S_h^{P_1 P_2}(f') \mathcal{R}_{P_1 P_2}^{ijkl}(f'), \\
\mathcal{R}_{P_1 P_2}^{ijkl}(f) &= \frac{1}{4\pi} \int d^2\hat{\Omega} \exp\left(-2\pi i f \hat{\Omega} \cdot (\mathbf{x}_i - \mathbf{x}_k)\right) \\
&\times D^{ab}(\hat{u}_{ij} \cdot \hat{\Omega}, f) D^{cd}(\hat{u}_{kl} \cdot \hat{\Omega}, f) e_{ab}^{P_1}(\hat{\Omega}_1) e_{cd}^{P_2}(\hat{\Omega}).
\end{aligned} \tag{3.18}$$

$\mathcal{R}_{P_1 P_2}^{ijkl}$ is referred to as the response function of the detector. \mathcal{R} depends on the relative position and orientation of the arms $i \rightarrow j$ and $k \rightarrow l$, as well as the transfer functions of the two arms.

In the remainder of this section we consider two interferometer designs. In §3.4.2.1 we consider the baseline design for near-future space-based interferometers such as the European Space Agency-led Laser Interferometer Space Antenna (LISA) [259], and in §3.4.2.2 we consider two futuristic ‘advanced stage’ LISA-like missions similar to the proposed Big Bang Observatory (BBO) [280].

3.4.2.1 One constellation

In this section, we consider the design shown in the left panel of Figure 3.8. This is the baseline design of the LISA mission, and consists of three satellites placed at the vertices \mathbf{x}_i of an equilateral triangle of side L , and a total of six laser links between the satellites, allowing for measurement of the phase change $\Delta\phi_{ij}$ where $i, j = \{1, 2, 3\}$, $i \neq j$. We define the following three signals:

$$\begin{aligned} s^\alpha(t) &= \frac{1}{2} (\Delta\phi_{12}(t-2L) + \Delta\phi_{21}(t-L) - \Delta\phi_{13}(t-2L) - \Delta\phi_{31}(t-L)) + n^\alpha(t), \\ s^\gamma(t) &= \frac{1}{2} (\Delta\phi_{31}(t-2L) + \Delta\phi_{13}(t-L) - \Delta\phi_{32}(t-2L) - \Delta\phi_{23}(t-L)) + n^\gamma(t), \\ s^\beta(t) &= s^\alpha(t) + 2s^\gamma(t). \end{aligned} \tag{3.19}$$

The equilateral design means that the laser phase noise, which is the dominant contribution to the noise terms $n(t)$, cancels [284]. Furthermore Ref. [256] shows that signals α and β have independent noise properties. We therefore consider their cross-correlations:

$$\langle s^{X_1}(f) s^{X_2}(f') \rangle = \frac{1}{2} \delta(f - f') [\mathcal{R}_I^{X_1 X_2}(f) I(f) + \mathcal{R}_V^{X_1 X_2}(f) V(f)], \tag{3.20}$$

where $X_1, X_2 = \{\alpha, \beta\}$, and:

$$\begin{aligned} \mathcal{R}_I^{X_1 X_2}(f) &= \frac{1}{4\pi} \int d^2\hat{\Omega} \left[F_{X_1}^+(f, \hat{u} \cdot \hat{\Omega}) F_{X_2}^{+*}(f, \hat{u} \cdot \hat{\Omega}) + F_{X_1}^\times(f, \hat{u} \cdot \hat{\Omega}) F_{X_2}^{\times*}(f, \hat{u} \cdot \hat{\Omega}) \right], \\ \mathcal{R}_V^{X_1 X_2}(f) &= \frac{1}{4\pi} \int d^2\hat{\Omega} \left[F_{X_1}^+(f, \hat{u} \cdot \hat{\Omega}) F_{X_2}^{\times*}(f, \hat{u} \cdot \hat{\Omega}) - F_{X_1}^\times(f, \hat{u} \cdot \hat{\Omega}) F_{X_2}^{+*}(f, \hat{u} \cdot \hat{\Omega}) \right], \end{aligned} \tag{3.21}$$

and [284, 256, 285]:

$$\begin{aligned}
F_X^P(f, \hat{\Omega} \cdot \hat{u}) &= D_{ij}^X(f, \hat{u} \cdot \hat{\Omega}) e_{ij}^P(\hat{\Omega}), \\
D^\alpha(f, \hat{u} \cdot \hat{\Omega}) &= \frac{1}{2} \exp(-2\pi i f \hat{\Omega} \cdot \mathbf{x}_1) \left[\hat{u} \otimes \hat{u} T_{\text{MI}}(f, \hat{n} \cdot \hat{u}) - \hat{v} \otimes \hat{v} T_{\text{MI}}(f, \hat{\Omega} \cdot \hat{v}) \right], \\
D^\beta(f, \hat{u} \cdot \hat{\Omega}) &= D^\alpha(f, \hat{u} \cdot \hat{\Omega}) + \exp(-2\pi i f \hat{n} \cdot \mathbf{x}_3) \\
&\quad \times \left[\hat{u} \otimes \hat{u} T_{\text{MI}}(f, \hat{\Omega} \cdot -\hat{u}) - \hat{w} \otimes \hat{w} T_{\text{MI}}(f, \hat{\Omega} \cdot -\hat{w}) \right], \\
T_{\text{MI}}(f, \hat{\Omega} \cdot \hat{u}) &= \frac{1}{2} \text{sinc} \left(\frac{f(1 - \hat{u} \cdot \hat{\Omega})}{2f_*} \right) \exp \left(-i \frac{f}{2f_*} (3 + \hat{u} \cdot \hat{\Omega}) \right) \\
&\quad + \frac{1}{2} \text{sinc} \left(\frac{f(1 + \hat{u} \cdot \hat{\Omega})}{2f_*} \right) \exp \left(-i \frac{f}{2f_*} (1 + \hat{u} \cdot \hat{\Omega}) \right).
\end{aligned}$$

Consider the instrument's response to a gravitational wave travelling in the direction $\hat{\Omega} = (\theta, \phi)$, and another travelling in a direction with $\theta \rightarrow \pi - \theta$, i.e. reflected in the $x - y$ plane. Since the vectors $\hat{u}, \hat{v}, \hat{w}, \mathbf{x}_i$ are all in the $x - y$ plane it is easy to see that the products $\mathbf{x} \cdot \hat{\Omega}, \hat{u} \cdot \hat{\Omega}$ etc. are invariant. Under this transformation only the z part of the basis tensor $e_{ab}^+(\hat{\Omega})$ is altered. Since $D^{ab}(f, \hat{\Omega} \cdot \hat{u})$ is non-zero only in the $x - y$ part, then the product $D^{ab}(f, \hat{\Omega} \cdot \hat{u}) e_{ab}^+(\hat{\Omega})$ is invariant. On the other hand the $x - y$ part of the $e_{ab}^\times(\hat{\Omega})$ tensor changes sign, meaning that $D^{ab}(f, \hat{\Omega} \cdot \hat{u}) e_{ab}^\times(\hat{\Omega})$ changes sign. Therefore, when performing the angular integral in Equation 3.21 the terms with a single power of $F_X^\times(f, \hat{\Omega} \cdot \hat{u})$ go to zero, giving $\mathcal{R}_V^{X_1 X_2}(f) = 0$. The conclusion is that co-planar detectors are not sensitive to the circular polarization of the gravitational wave background. This is true of other types of detectors with planar geometries such as pulsar timing arrays and individual ground-based detectors such as LIGO [278].

To gain sensitivity to circular polarization we need to introduce non-co-planar detector arms. Others [254] have considered using cross-correlations between ground-based detectors like LIGO, VIRGO [279], and KAGRA [286], which have a suitable geometry. In the next subsection we consider an extension to LISA in which we add a second constellation of three satellites to break the co-planar geometry.

3.4.2.2 Two-constellations

The extended LISA set-up is shown in the right panel of Figure 3.8. It consists of two constellations of three equal-arm detectors. The two constellations are separated by a rotation of π radians and a translation of $DL\hat{z}$. The j^{th} detector on the i^{th}

constellation is at position \mathbf{x}_{ij} , and the unit vectors joining them are given by: $u_i = (\mathbf{x}_{i2} - \mathbf{x}_{i1})/L$, $v_i = (\mathbf{x}_{i3} - \mathbf{x}_{i1})/L$, $w_i = (\mathbf{x}_{i3} - \mathbf{x}_{i2})/L$. We base this analysis on the designs proposed by Ref. [256] which optimize the parameters L and D to achieve equal sensitivity to intensity and polarization of the gravitational wave background. Similar designs have also been considered by [284, 280, 287].

We use the signals defined in Equation 3.19, but α, β are now written α_i, β_i where i refers to the constellation on which we are measuring the signal. The detector transfer functions are the same as the single-constellation, but with extra indices referring to the constellation we are considering [284, 256]:

$$\begin{aligned}
D^{\alpha_i}(f, \hat{u}_i \cdot \hat{n}) &= \frac{1}{2} \exp(-2\pi i f \hat{n} \cdot \mathbf{x}_{i1}) [\hat{u}_i \otimes \hat{u}_i T_{\text{MI}}(f, \hat{n} \cdot \hat{u}_i) - \hat{v}_i \otimes \hat{v}_i T_{\text{MI}}(f, \hat{n} \cdot \hat{v}_i)], \\
D^{\beta_i}(f, \hat{u}_i \cdot \hat{n}) &= D^{\alpha}(f, \hat{u}_i \cdot \hat{n}) + \exp(-2\pi i f \hat{n} \cdot \mathbf{x}_{i3}) \hat{u}_i \otimes \hat{u}_i T_{\text{MI}}(f, \hat{n} \cdot -\hat{u}_i) \\
&\quad - \exp(-2\pi i f \hat{n} \cdot \mathbf{x}_{i3}) \hat{w}_i \otimes \hat{w}_i T_{\text{MI}}(f, \hat{n} \cdot -\hat{w}_i).
\end{aligned} \tag{3.22}$$

Following [256] we then combine Equations 3.19 to form estimators sensitive to just intensity or circular polarization:

$$\begin{aligned}
\frac{1}{2} \delta_T(f - f') \mathcal{R}_I(f) I(f) &\equiv \langle [s^{\alpha_1}(f) + s^{\beta_1}(f)] [s^{\alpha_2^*}(f') + s^{\beta_2^*}(f')] \rangle, \\
\frac{1}{2} \delta_T(f - f') \mathcal{R}_V(f) V(f) &\equiv \langle s^{\alpha_1}(f) s^{\beta_2^*}(f') - s^{\beta_1}(f) s^{\alpha_2^*}(f') \rangle.
\end{aligned} \tag{3.23}$$

We will consider two experimental configurations of the two-constellation, introduced in Ref. [256]: ‘LISA’ with $L = 1 \times 10^9$ m, $D = 7$, $T = 10$ years, and ‘BBO’ with $L = 5 \times 10^7$ m, $D = 2$, $T = 10$ years. These designs are optimized to achieve roughly equal sensitivity to I and V .

3.4.3 Interferometer signal-to-noise

Under the assumption that the signals we are cross-correlating have independent noise properties and are Gaussian-distributed, and that the noise spectrum dominates over the signal, then the signal-to-noise in the interferometer is given by [256, 284, 285]:

$$\left(\frac{S}{N}\right)_{I,V}^2 = 2T \int_0^\infty df \left(\frac{3H_0^2}{4\pi^2}\right)^2 \frac{|\mathcal{R}_{I,V}(f) \Omega_{\text{GW}}^{I,V}(f)|^2}{f^6 S_n^{I,V}(f)^2}, \tag{3.24}$$

where $S_n^{I,V}(f)$ is the power spectrum of the noise in the I, V signals, and $\Omega_{\text{GW}}^{I,V}$ is the fractional energy density of gravitational waves in intensity and circular polarization today, defined in Equation 3.13. To find the background fractional energy density today we multiply the primordial spectrum by the appropriate transfer function [288, 247, 289]: $\Omega_{\text{GW}}^{I,V}(f) = \frac{\Omega_{R,0}}{24}(\mathcal{P}^L(f) \pm \mathcal{P}^R(f))$, where $\Omega_{R,0}$ is the fractional energy density in radiation today.

Up to this point we have not discussed the noise, since it vanishes in the cross-correlations we consider. However it still contributes to the variance of the estimators in Equations 3.23. There are three major sources of noise in measurements of a particular optical path through an interferometer: shot noise $S_{n,s}(f)$, accelerometer noise $S_{n,a}(f)$, and the dominant laser phase noise, $S_{n,\phi}(f)$. As pointed out in §3.4.2 the major motivation for using equal-arm Michelson interferometers, as given in the first two lines of Equations 3.19, is the cancellation of the laser phase noise. The shot and acceleration noises can be approximated by taking the fiducial LISA [259] and BBO [280] values and scaling them to an instrument with arm length L observing at frequency f [284]. The final expressions for $S_n^{I,V}(f)$ are derived by Ref. [256]:

$$\begin{aligned} S_n^I(f) &= \frac{121}{4} \left[S_{n,s}(f) + 2S_{n,a}(f) \left(1 + \cos^2 \left(\frac{f}{f_*} \right) \right) \right]^2, \\ S_n^V(f) &= \frac{96}{121} S_n^I(f), \end{aligned} \quad (3.25)$$

where the values for $S_{n,a}(f)$ and $S_{n,s}(f)$ for LISA and BBO are given in Ref. [256]. As is the case for the CMB, our Galaxy contains sources of gravitational waves that may act as a confusion noise to a measurement of the GWB [290, 291]. It is expected that compact binary systems in our Galaxy will form a gravitational wave foreground with an amplitude in intensity of $\Omega_{\text{WD}} \sim 10^{-12}$ in the mHz regime. The shape of this spectrum is quite complicated because different periods of a binary system's evolution dominate at different frequencies and have different frequency dependences [290]. For the design of LISA we consider we expect the impact of such a foreground to be small compared to the acceleration noise [292]. The BBO design we consider peaks at $\gtrsim 0.3$ Hz, which is expected to be relatively free of such sources of noise [293, 294]. However, we are mainly interested in detecting chirality of the GWB, and this is more easily distinguished from astrophysical foregrounds, and accordingly previous studies have not considered polarised foregrounds [256, 295, 247]. Therefore, we do not consider a contribution to the noise from astrophysical foregrounds in intensity or in polarisation,




Figure 3.9: Signal-to-noise contours obtained using Equation 3.24 for a LISA-like experiment described in §3.4.3. The primordial spectrum has $k_p = 1 \times 10^{11} \text{ Mpc}^{-1}$.

but it should be noted that we expect a small degradation in the achievable intensity sensitivity of the fiducial LISA design compared to our result, due to the confusion noise of astrophysical sources.

3.4.4 Interferometer results

In Figure 3.9 we plot signal-to-noise contours for LISA assuming $k_p = 1 \times 10^{11} \text{ Mpc}^{-1}$ and in Figure 3.10 we plot the corresponding contours for BBO assuming $k_p = 1 \times 10^{13} \text{ Mpc}^{-1}$. We see that both the LISA and BBO configurations may detect a polarized gravitational wave background with signal-to-noise greater than one in a regime unavailable to the CMB. In the case of LISA the signal-to-noise for $k_p = 1 \times 10^{13} \text{ Mpc}^{-1}$ is of order one. However, we see that a BBO-like design far exceeds the sensitivity of LISA, probing a much larger range of r for the large k_p values, inaccessible to CMB experiments. A single constellation design described in 3.4.2.1 would achieve equivalent sensitivity in I to LISA and BBO, but with no V sensitivity. Therefore, the fiducial LISA design would be sensitive to the inflationary model we consider here, since a positive detection of I at these scales with no corresponding detection on CMB scales would require a strong scale dependence of the gravitational wave spectrum.




Figure 3.10: Signal-to-noise contours obtained using Equation 3.24 for a BBO-like experiment described in §3.4.3. The primordial spectrum has $k_p = 1 \times 10^{13} \text{ Mpc}^{-1}$.

3.5 Discussion

In this chapter we have considered for the first time the detectability of a new model for the production of gravitational waves proposed in Ref. [42]. Given the increasing effort to measure the B-mode spectrum of the CMB, this is an important step in establishing the origins of any detected primordial tensor perturbations. This model has a unique tensor spectrum characterized by its scale-dependence and chirality, both of which we use in order to find observational markers that allow it to be distinguished from the conventional primordial gravitational waves produced by vacuum fluctuations. If a detection of primordial gravitational waves is made, and the markers we find to be detectable are absent, we may then rule out such a model. In §3.3 we provided robust forecasts of the ability of the LiteBIRD satellite mission to detect the TB and EB correlations that result from the chiral tensor spectrum. We found that LiteBIRD would be able to detect the chirality for $r_* \gtrsim 0.03$, whilst $r_* < 0.07$ is required by current observations. The addition of Stage 4 observations has little effect as such a survey would be limited to $\ell > 30$, but the primordial chiral signal is contained almost entirely within $2 < \ell < 30$. Further, we found that for cosmic-variance limited observations the maximum achievable signal-to-noise for $r_* < 0.07$ would be ~ 3 . From these studies we conclude that the ability of CMB two-point statistics to determine the presence of a chiral GWB is fairly limited.

However, in this study we have not fully leveraged the scale-dependence of the axion-SU(2) model. Single-field slow-roll expects the tensor spectrum to have a tilt given by the self-consistency relation $n_T = -r/8$, and it would be possible to test departures from this using a combination of both CMB and interferometer constraints to provide a lever-arm [296, 297]. Such a study would be aided by future groundbased observations such as Simons Observatory or S4. In this study we found that for a peak wavenumber in the range $k_p \sim 7 \times 10^{-5} - 5 \times 10^{-3} \text{ Mpc}^{-1}$ the primordial BB spectrum is detectable by LiteBIRD with $(S/N)_{BB} \gtrsim 1$ for $r_* \gtrsim 10^{-3}$. However, the projected sensitivity on n_T for LiteBIRD alone is ~ 0.04 , which is not sufficient to test deviations from the self-consistency relation, without external constraints.

Another characteristic of the axion-SU(2) model of Ref. [42] is its intrinsic non-Gaussianity. Some studies have recently shown that higher order statistics of B-modes, such as the BBB bispectrum, may yield a $> 2\sigma$ significance for the axion-U(1) model [246, 240]. An analysis of the CMB non-Gaussianity for the axion-SU(2) model is therefore in order [282].

In §3.4 we showed that interferometers may provide a complementary probe to the CMB at much smaller scales $\sim 10^{12} \text{ Mpc}^{-1}$, even for the relatively flat spectra required by the attractor behaviour of the background axion field coupled to the SU(2) gauge field. This takes advantage of the scale-dependence of the axion-SU(2) model, which allows the spectrum to have a large excursion at some scale k_p , e.g. as shown in Figure 3.7, making the cosmological GWB of the axion-SU(2) model a viable target for interferometers with current sensitivities. We went on to consider two designs of an advanced stage LISA-like mission proposed by Ref. [256] which are sensitive to both the intensity and circular polarization of the GWB. Whilst interferometers are not in general sensitive to the same parameter space of the model as CMB probes, we found that for spectra with a very large values of k_p and σ , that would be undetectable on CMB scales, such experiments could make significant detections, and therefore complement CMB constraints.

Chapter 4

ACT foregrounds

4.1 Introduction

The Atacama Cosmology Telescope (introduced first Section 1.3.4) is a six meter telescope in the Atacama Desert, Chile. *ACT* was first commissioned in 2008, and made high-resolution Bolometer measurements of the temperature anisotropies of the CMB with the Millimeter Bolometer Array Camera (MBAC) [298]. *ACT* aimed to measure the thermal and kinematic Sunyaev-Zel'dovich effects, as well as lensing of the CMB by large scale structure. In order to extract independent cosmological constraints from the polarized anisotropies of the CMB, *ACT* was upgraded with a polarization-sensitive instrument, *ACTPol* [299, 300], which started observing four non-contiguous patches of sky in 2013, and in 2014 focused on three wider patches: D56, D8, and BOSS-N, the parameters of which are summarized in Table 4.1. Power spectra from observations in the 2013 season have been presented in Ref. [301], and have been updated to include observations from 2014 in Ref. [13].

In the previous chapters we have discussed how thermal radiation from Galactic dust grains can contaminate observations of the CMB, especially if those observations are made close to the Galactic plane, or are polarization sensitive. In order to forecast the importance of dust modeling when analyzing future *ACT* data, we calculate the expected dust power spectrum in regions observed by *ACTPol*. Previous analyses, such as Refs. [12, 13], have shown that the dust level is negligible in the *ACT* equatorial, and larger D56 regions. These analyses therefore neglected map-space modeling of the foregrounds, and simply marginalized over a dust power spectrum amplitude in the cosmological likelihood analysis. In this chapter, we use *Planck* data to assess whether this will remain a good approach when analyzing polarization data, and high-sensitivity temperature data for the D8, D56, and BOSS-N regions.

Since D8, D56, and BOSS-N are non-overlapping (see Figure 4.1, note that we do not consider D1, D2, or D9 in this analysis), and vary in size, it is possible that the amplitude of the dust power spectrum could vary significantly between them. Therefore, we consider each patch independently, and estimate the dust power spectrum at 150 GHz from high-frequency *Planck* observations.

In Section 4.2 we review the data we use to measure the dust power spectrum. In Section 4.3 we introduce the methodology applied to calculate the dust power spectrum. In Section 4.4 we present the resulting constraints, and finally in Section 4.5 we briefly summarize our conclusions.

4.2 Data

In this analysis we use the publicly available *Planck* PR2 353 GHz and 217 GHz half-mission temperature and polarization maps [302], downloaded from the *Planck* Legacy Archive ¹. In order to assess the uncertainties we also make use of the full focal plane noise simulations (referred to as FFP8.1) [303], which are available on NERSC.

To define observations regions, we create binary masks for the three ACT regions using the survey parameters summarized in Table 4.1. The binary maps are then apodized to reduce ringing in the power spectrum estimation, and to satisfy the boundary conditions required when applying a pure- B estimator [192, 193]. In a realistic analysis, the highest foreground intensity regions would be masked out in the estimation of the power spectrum, and a point source mask would be applied to remove the contribution from point sources. Therefore, in each patch we combine with the *Planck* PR2 143 GHz likelihood mask, which accounts for both of these effects.

Patch	RA (deg)	Dec (deg)	Size (sq. deg)
D8	(-4, 13)	(-50, -35)	190
D56	(-8, 41)	(-7.2, 4)	700
BOSS-N	(142, 228)	(-4, 20)	2000

Table 4.1: Specifications of the patches of sky used in this analysis, also illustrated in Figure 4.1.

4.3 Method

4.3.1 Theory

We model each of the *Planck* half-mission frequency maps, $\mathbf{g}_\nu^i(\hat{n})$, as a combination of CMB, dust, CIB, and noise:

$$\mathbf{g}_\nu^i(\hat{n}) \equiv \mathbf{g}_\nu^{\text{CIB}}(\hat{n}) + \mathbf{g}_\nu^{\text{CMB}}(\hat{n}) + \mathbf{g}_\nu^{\text{Dust}}(\hat{n}) + \mathbf{n}_\nu^i(\hat{n}), \quad (4.1)$$

¹<https://www.cosmos.esa.int/web/planck/pla>




Figure 4.1: Map from Ref. [11], showing the various ACTPol fields overplotted on a map of dust emission. In this chapter we consider the D56, D8 and BOSS-N regions, and not D1, D2, or D9. D56 overlaps with the ACT-Equatorial region from previous observation seasons (e.g. Ref. [12]), and the D8 region overlaps with the ACT-Southern region.

where $i = 1, 2$ labels the mission half. The dust contribution is assumed to be a modified black body with a single temperature and spectral index:

$$\mathbf{g}_\nu^{\text{Dust}}(\hat{n}) \equiv \alpha(\nu, \nu_0) \left(\frac{\nu}{\nu_0} \right)^\beta \frac{B(\nu, T)}{B(\nu_0, T)} \mathbf{g}_{\nu_0}^{\text{Dust}}(\hat{n}), \quad (4.2)$$

where $\alpha(\nu, \nu_0)$ is a factor converting from flux units to thermodynamic units. We may take differences of these maps to eliminate the CMB contribution:

$$\begin{aligned} \Delta \mathbf{g}^i(\hat{n}) &\equiv \mathbf{g}_{\nu_1}^i(\hat{n}) - \mathbf{g}_{\nu_2}^i(\hat{n}) \\ &= \mathbf{g}_{\nu_1}^{\text{Dust}}(\hat{n}) - \mathbf{g}_{\nu_2}^{\text{Dust}}(\hat{n}) + \mathbf{g}_{\nu_1}^{\text{CIB}}(\hat{n}) - \mathbf{g}_{\nu_2}^{\text{CIB}}(\hat{n}) + \mathbf{n}_{\nu_1}^i(\hat{n}) - \mathbf{n}_{\nu_2}^i(\hat{n}). \end{aligned} \quad (4.3)$$

To avoid noise bias in the estimation of the power spectrum, we cross correlate difference maps constructed from different half-mission splits:

$$\langle \Delta \mathbf{g}_{\ell m}^1 \Delta \mathbf{g}_{\ell m}^2 \rangle = C_\ell^{\text{Dust+CIB}}(\nu_1, \nu_1) + C_\ell^{\text{Dust+CIB}}(\nu_2, \nu_2) - 2C_\ell^{\text{Dust+CIB}}(\nu_1, \nu_2). \quad (4.4)$$

Assuming that the frequency scaling of the dust is well-described by a single modified black body component with a single temperature and spectral index in each patch ², i.e. that there is no decorrelation, we may calculate the power spectrum of dust at some frequency, $\tilde{\nu}$:

$$\begin{aligned} C_\ell^{\Delta \mathbf{g} \Delta \mathbf{g}} &\equiv \langle \Delta \mathbf{g}_{\ell m}^1 \Delta \mathbf{g}_{\ell m}^2 \rangle \\ &= f^{\text{Dust}}(\tilde{\nu}, \nu_1, \nu_2; \beta_d, T_d) C_\ell^{\text{Dust}} \\ &\quad + C_\ell^{\text{CIB}}(\nu_1, \nu_1) + C_\ell^{\text{CIB}}(\nu_2, \nu_2) - 2C_\ell^{\text{CIB}}(\nu_1, \nu_2), \end{aligned} \quad (4.5)$$

²We do not know the true values of the dust temperature or spectral index. We calculated the variance of the estimated power spectra when drawing these parameters from the distributions: $\beta_d \sim \mathcal{N}(1.6, 0.017)$, and $T_d \sim \mathcal{N}(19.6, 2)$, and found it to be negligible compared to the noise at $\ell \gtrsim 400$. The distributions chosen are representative of the scatter in the dust spectral parameters over the whole sky, as found by the Planck COMMANDER foreground separation analysis [9], and so represent the most pessimistic case for variation within the compact region. We therefore proceed by fixing $T_d = 19.6$, and $\beta_d = 1.6$.

with:

$$f^{\text{Dust}}(\nu_0, \nu_1, \nu_2; \beta_d, T_d) = \left(\frac{\alpha(\nu_1, \nu_0) \nu_1^{\beta_d} B(\nu_1, T_d) - \alpha(\nu_2, \nu_0) \nu_2^{\beta_d} B(\nu_2, T_d)}{\alpha(\tilde{\nu}, \nu_0) \nu_0^{\beta_d} B(\nu_0, T_d)} \right)^2. \quad (4.6)$$

We would like to find the dust power at frequencies observed by *ACT* so $\tilde{\nu} = 150$ GHz, $\nu_1 = 217$ GHz, and $\nu_2 = 353$ GHz.

For consistency with Ref. [12], we model the dust power spectrum as a power law, with index -2.7, and pivot multipole $\ell_0 = 3000$:

$$C_\ell^{\text{Dust}} = A_D (\ell/\ell_0)^{-2.7}. \quad (4.7)$$

Our results are not very sensitive to the choice of power law index, and it is anyway found to be a good fit by recent analyses of the *Planck* data. The CIB power spectrum is made up of a clustered and a Poisson contribution [14]:

$$\begin{aligned} \frac{\ell(\ell+1)}{2\pi} C_\ell^{\text{CIB-P}}(\nu_1, \nu_2) &= A_P \mu(\nu_1, \nu_2) (\ell/\ell_0)^2, \\ \frac{\ell(\ell+1)}{2\pi} C_\ell^{\text{CIB-C}}(\nu_1, \nu_2) &= A_C \mu(\nu_1, \nu_2) (\ell/\ell_0)^{2-n}, \end{aligned} \quad (4.8)$$

where the frequency dependence, $\mu(\nu_1, \nu_2)$, is based on a modified black body spectrum for dust galaxies, and can be found in Ref. [14]. The CIB is not polarized, and therefore $A_P = A_C = 0$ in the analysis of the polarization spectra.

With these definitions, Equation 4.5 now has three free parameters: the amplitude of dust, A_D , and the amplitudes of the Poisson and clustered CIB components, A_P and A_C . We group these together in the parameter vector, $\boldsymbol{\theta} = (A_D, A_P, A_C)$, and by inserting Equations 4.7 and 4.8 into Equation 4.5 can write an equation for the power spectrum in Equation 4.5, as a function of $\boldsymbol{\theta}$:

$$C_\ell^{\Delta\mathbf{g}\Delta\mathbf{g}} = C_\ell^{\Delta\mathbf{g}\Delta\mathbf{g}}(\boldsymbol{\theta}) \quad (4.9)$$

4.3.2 Power spectrum estimation

Equation 4.5 manifestly excludes any contribution from the CMB, and so we do not first calculate different maps, but instead proceed to calculate TT , EE , BB , and TE

power spectra for the frequency combinations: 217x217, 217x353, 353x353. We avoid noise-bias by taking cross spectra between half-mission splits of the data.

We calculate power spectra using the pseudo- C_ℓ package `NaMaster` [189]. This approach first calculates a biased *pseudo*-spectrum, $\tilde{\mathbf{C}}_\ell$, of the two fields $\mathbf{f}^w(\hat{n})$ and $\mathbf{g}^v(\hat{n})$, of spin s_f and s_g respectively, which have been weighted by the inverse variance maps $\mathbf{w}(\hat{n}), \mathbf{v}(\hat{n})$, respectively:

$$\tilde{\mathbf{C}}_\ell \equiv \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \mathbf{f}_{\ell m}^w \mathbf{g}_{\ell m}^{v\dagger}, \quad (4.10)$$

and then computes an analytic estimate of the mode-coupling matrix, $\mathbf{M}_{\ell\ell'}^{s_f s_g}$, that relates the biased estimate to the true underlying spectrum, \mathbf{C}_ℓ :

$$\langle \tilde{\mathbf{C}}_\ell \rangle = \sum_{\ell'} \mathbf{M}_{\ell\ell'}^{s_f s_g} \mathbf{C}_{\ell'} \quad (4.11)$$

where the angle brackets indicate equality over the ensemble average, and we have used bold font to indicate that \mathbf{C}_ℓ is a vector containing the cross spectra between different fields: $\mathbf{C}_\ell = (C_\ell^{TT}, C_\ell^{TE}, \dots)$.

The finite sky coverage of the *ACT* regions will impact the power spectrum estimation by limiting the resolution in ℓ space, and introducing coupling between the *E* and *B* modes. The limited resolution in ℓ space means that there are not sufficient modes to invert Equation 4.11. We may then proceed by binning the spectra and mode coupling matrix into bandpowers, defined by some set of multipoles $\boldsymbol{\ell}_b = (\ell_b^1, \dots, \ell_b^{N_b})$, and associated weights $\mathbf{w}_b = (w_b^1, \dots, w_b^{N_b})$, subject to the normalization condition $\sum_{\ell \in \boldsymbol{\ell}_b} w_b^\ell = 1$ [189, 192, 193]. Under this binning operation, we can estimate only the bandpowers, $\tilde{\mathbf{C}}_b$:

$$\tilde{\mathbf{C}}_b = \sum_{\ell \in \boldsymbol{\ell}_b} w_b^\ell \tilde{\mathbf{C}}_\ell, \quad (4.12)$$

and the binned mode-coupling matrix:

$$\mathbf{M}_{bb'} = \sum_{\ell \in \boldsymbol{\ell}_b} \sum_{\ell' \in \boldsymbol{\ell}_{b'}} w_b^\ell w_{b'}^{\ell'} \mathbf{M}_{\ell\ell'}^{s_f s_g}. \quad (4.13)$$

We bin the power spectra in uniform bins of width $\Delta\ell_b = 50$ for BOSS-N, and $\Delta\ell_b = 100$ for D56 and D8.

We correct for the E to B leakage by computing only *pure* B modes [193, 190]. This is implemented as a feature in `NaMaster`, and requires only to ensure that the weight masks $\mathbf{v}(\hat{n})$ and $\mathbf{w}(\hat{n})$ satisfy Neumann and Dirichlet boundary conditions [189]. We check the masks described in Section 4.2 by explicitly calculating their second derivative, and we check for biases on 200 mock Gaussian signal only simulations.

The final estimator for the bandpower b is then:

$$\hat{\mathbf{C}}_b = \sum_{b'} (\mathbf{M}_{bb'}^{s_f s_g})^{-1} \tilde{\mathbf{C}}_{b'}^{\text{Pure}} \quad (4.14)$$

4.3.3 Calculation of uncertainties

Since we are dealing with combinations of maps and power spectra that exclude all contributions from the CMB, the variance of Equation 4.5 is due only to noise. To quantify the uncertainty, we calculate the covariance of the bandpowers in Equation 4.14, $\mathbf{C}_{bb'}$, using 100 Monte Carlo realizations of the FFP8.1 noise simulations [303]:

$$\mathbf{C}_{bb'} \approx \langle (\hat{\mathbf{C}}_b - \langle \hat{\mathbf{C}}_b \rangle_{\text{MC}}) (\hat{\mathbf{C}}_{b'} - \langle \hat{\mathbf{C}}_{b'} \rangle_{\text{MC}}) \rangle_{\text{MC}} \quad (4.15)$$

4.3.4 Likelihood

In order to constrain the dust amplitude, we fit Equation 4.5 to the measured power spectrum using the model in Equation 4.9, and assuming a Gaussian distribution of the estimated bandpowers, which is suitable at the high multipoles constrained by *ACT* ($\ell \gtrsim 400$) [304]:

$$\ln [\mathcal{L}(\boldsymbol{\theta} | \hat{\mathbf{C}})] \propto (\mathbf{C}(\boldsymbol{\theta}) - \hat{\mathbf{C}})^T \mathbf{C}^{-1} (\mathbf{C}(\boldsymbol{\theta}) - \hat{\mathbf{C}}). \quad (4.16)$$

We use prior measurements of the CIB amplitude in similar patches of sky to impose priors on the amplitude of the CIB terms, A_P and A_C , from [14]: $A_P \sim \mathcal{N}(6.9, 0.4)$ and $A_C \sim \mathcal{N}(4.9, 0.9)$, and compute the posterior for $\boldsymbol{\theta}$ using the `emcee` [168] package to sample Equation 4.16.




Figure 4.2: Figure from [12] showing the dust levels in the ACT Equatorial and ACT Southern regions, at 148 GHz and 218 GHz. Shown in the teal line are the results for ACT-E at 148 GHz, which is most similar to our analysis of D56, and in purple the results for ACT-S at 148 GHz, which is most similar to our analysis of D8.

q'




Figure 4.3: Figure taken from Ref. [13], showing a computation of the difference between the 353 GHz Planck TT spectrum, and the 149 GHz ACT TT spectrum, which is attributable to dust + CIB. Overplotted in the shaded region is the Das 2014 dust model, with dust amplitude $A_d = 0.8 \mu K^2$

4.4 Results

We conduct this analysis separately on the BOSS-N, D56, and D8 regions, in both temperature and polarization. Previous *ACT* analyses have focused on the equatorial and southern regions ACT-E, and ACT-S [12], and D56 [13]. Refs. [12, 13] conducted a similar analysis to that presented here, and found consistent results for the overlapping ACT-E and D56 regions, with a dust amplitude of $A_D^{D56} = (0.8 \pm 0.2) \mu K^2$. Ref. [12] find that the ACT-S region has an amplitude $A_D^{ACT-S} = (0.4 \pm 0.2) \mu K^2$. Figure 4.2 shows the ACT-E and ACT-S spectra compared to Λ CDM, and Figure 4.3 shows the D56 results from Ref. [13].

4.4.1 Temperature

The results for the dust temperature-only spectra are shown in Figures 4.4, 4.6, and 4.5. The fitted dust amplitudes are summarized in Table 4.2. We find dust amplitudes consistent with the findings of [12, 13] in D56 and D8. We find that BOSS-N has a

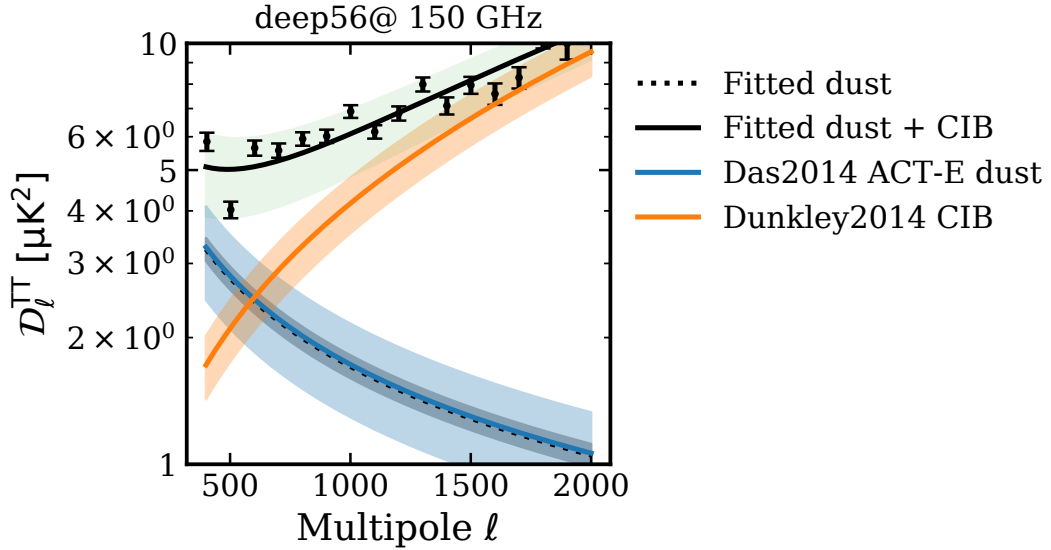


Figure 4.4: Figure showing the TT data (black points), and the fitted dust + CIB model (black), for the D56 patch. The best-fit dust model, marginalized over CIB amplitudes, is shown in dashed black. These are compared to the results of Ref. [12] for the ACT equatorial region (blue), which overlaps with D56. Shown in orange is the CIB prior model from [14].

slightly higher level, $A_D^{\text{BOSS-N}} = (1.3 \pm 0.1) \mu\text{K}^2$, as would be expected for a larger region, which overlaps more with the Galaxy. Our results suggest that dust emission will contaminate observations of the temperature anisotropies of the CMB at a similar level to the previous *ACT* studies, and does not require more advanced modeling than previous studies.

Patch	$A_D^{\text{TT}} (\mu\text{K}^2)$	$A_D^{\text{EE}} (\mu\text{K}^2)$	$A_D^{\text{BB}} (\mu\text{K}^2)$	$A_D^{\text{TE}} (\mu\text{K}^2)$
D8	0.38 ± 0.10	0.034	0.037	-0.0034 ± 0.023
D56	0.79 ± 0.05	0.060	0.034	-0.0021 ± 0.026
BOSS-N	1.3 ± 0.11	0.015 ± 0.0090	0.021 ± 0.012	0.030 ± 0.0026

Table 4.2: Summary of dust temperature power spectrum amplitude fits. Results are quoted for a pivot multipole of $\ell_0 = 3000$, for the dust model in Equation 4.7, $C_\ell^{\text{Dust}} = A_D(\ell/\ell_0)^{-2.7}$. Cells with a single number give the one-sigma upper bound on the amplitude.




Figure 4.5: Figure showing the TT data (black points), and the fitted dust + CIB model (black) for the D8 patch. The best-fit dust model, marginalized over CIB amplitudes, is shown in dashed black. These are compared to the results of Ref. [12] for the ACT southern region (blue), which borders D8. Shown in orange is the CIB prior model from [14].




Figure 4.6: Figure showing the TT data (black points), and the fitted dust + CIB model (black) for the BOSS-N patch. The best-fit dust model, marginalized over CIB amplitudes, is shown in dashed black. These are compared to the results of Ref. [12] for the ACT equatorial region (blue). Shown in orange is the CIB prior model from [14].

4.4.2 Polarization

We conduct the same exercise for polarization, the difference here being the lack of a CIB component. Without CIB contamination, the rescaled data should consist of only dust. However, the analysis is made harder by a few factors: i) the larger noise in the polarization maps, ii) the lower dust signal in polarization, iii) E to B leakage due to incomplete sky coverage. As such, at scales $\ell \gtrsim 400$ the signal to noise for these data on such small sky regions is very poor, and it has not proved possible to detect a non-zero dust amplitude in EE or BB for D56 or D8. In Figure 4.7 we show the polarization spectra for BOSS-N, for which a non-zero amplitude of EE, BB, and TE is marginally preferred (Table 4.2).

4.5 Conclusions

In this chapter we have estimated the temperature and polarization dust power spectra at 150 GHz in three regions of the sky observed by *ACT*. We validated our estimates by comparing to previous results from the literature. Our estimates of the dust temperature power spectra in the well-studied D8 and D56 regions are consistent with independent estimates made using FIRAS and DIRBE data in Refs. [12, 13]. We find that the dust temperature power spectrum in the BOSS-N region is 60% higher than in the D56 region, as might be expected, as it is roughly three times larger. We therefore do not suggest any modifications to the way in which dust is modeled in the analysis of temperature data taken as part of the *ACTPol* surveys, and instead provide our constraints on the dust amplitudes as priors to the likelihood analysis.

The polarization data provided fewer insights, due to the much poorer signal-to-noise. With only a slight preference for non-zero power in the BOSS-N region, and results consistent with no dust in D8 and D56, we are not able to make strong conclusions about the nature of dust contamination in the regions analyzed by *ACTPol*.




Figure 4.7: Figure showing the calculated dust EE, BB, and TE, spectra for BOSS-N. BOSS-N is the largest sky region we consider, and contains more dust in intensity than D56 or D8, as established in Section 4.4.1.

Chapter 5

Removal of Galactic foregrounds for upcoming primordial gravitational wave searches

5.1 Introduction

Many experiments, including ACT [305], SPT-3G [306], BICEP3 / Keck array [307], the Simons Array [308] and CLASS [268] are measuring the polarized anisotropies of the CMB with improved precision, over a range of scales. The Simons Observatory [3] is a new experiment for the 2020s that will consist of multiple telescopes in the Chilean Atacama desert. It will have a 6 m Large Aperture Telescope (LAT) with arcminute scale resolution, and an array of three 0.42 m refracting Small Aperture Telescopes (SAT). It is these that are targeted at measuring degree scale B modes. SO has a variety of science goals, and the design of the LAT and SAT, as well as their scan strategy, is optimized to achieve these goals [309]. The B mode measurements at $\ell > 500$ from the LAT will be sensitive to the growth of structure in the Universe, and so will place constraints on the dark energy equation of state [310], the sum of neutrino masses [311], and the number of relativistic species, whilst the large-scale observations of the SAT will target inflationary gravitational waves [3].

Other processes on the sky produce B modes, which may contaminate observations of CMB primary B modes, leading to a bias in the estimation of r , or an increase in the uncertainty, σ_r . Polarized radiation from sources within our own Galaxy is well known to contaminate large scale observations in all directions and at all frequencies [122]. On sub-degree scales, gravitational lensing of the CMB by structure between the surface of last scattering and today mixes primary E and B modes [80], The lensing signal contains a wealth of physical information, as it probes different epochs in the Universe's history. It may be the cleanest probe of many phenomena, such as dark energy and the sum of neutrino masses, especially when combined with external tracers of structure, such as galaxy surveys [310]. However, for observations of the primordial B -mode spectrum, the lensing contribution is a nuisance, acting as an additional source of confusion noise for primordial B -modes.

As the sensitivity of B mode observations improves, our ability to make inferences about cosmology will be limited by our modeling of both gravitational lensing and polarized Galactic foregrounds. To date, an array of methods for modeling Galactic foregrounds have been suggested (e.g. Refs. [312, 233, 313]).

In this chapter we present a foreground-removal pipeline similar to that used in [233], which fits the foregrounds parametrically. We use it to demonstrate the ability of Simons Observatory to place improved constraints on the tensor-to-scalar ratio in the presence of large-scale Galactic foregrounds. It is a new implementation of one of the algorithms already described in Ref. [3], hereafter SO19. We find that the nominal

design of SO, with simple foregrounds and in the absence of additional systematic uncertainties, should achieve a constraint of $\sigma_r \leq 0.003$. In this paper we extend the forecasts presented in SO19, exploring the effect of masking and fitting the foregrounds with spatially varying spectral parameters. We also show how cross-correlations of the cleaned maps with Galactic tracers can be used to detect residual foregrounds.

This chapter is structured as follows: in Section §5.2 we describe the synthetic observations, in §5.3 the component separation technique used, the estimation of power spectra from cleaned CMB maps, and the inference of the tensor-to-scalar ratio. In §5.4 we present the results of applying our pipeline to the SO design, and in §5.5 discuss the results and their implications.

5.2 Simulations

In this section we describe the simulations on which these forecasts are based. As in SO19 we use the PySM¹ software [1] to produce Q and U maps of Galactic dust and synchrotron emission at the n_{freqs} frequencies observed by SO. We represent these sky simulations as the vector $\mathbf{s}(\hat{n})$, containing $n_{\text{freq}} \times n_{\text{pol}} \times n_{\text{pix}}$ elements for n_{pol} maps with n_{pix} pixels. PySM models may be summarized by:

$$\mathbf{s}(\hat{n}) = f(\beta(\hat{n})) \cdot \mathbf{T}(\hat{n}), \quad (5.1)$$

where \mathbf{T} is a $n_{\text{comp}} \times n_{\text{pol}} \times n_{\text{pix}}$ vector containing templates of the emission of each of the n_{comp} components at a frequency ν_0 , f is a $(n_{\text{freq}} \times n_{\text{pol}} \times n_{\text{pix}}) \times (n_{\text{comp}} \times n_{\text{pol}} \times n_{\text{pix}})$ matrix, containing the component SEDs that scale each component from its reference frequency to the observed frequency ν , and β represents the parameters of the assumed model SED, which may be spatially varying.

We follow SO19 by modeling the effects of the instrument as a symmetric Gaussian beam, $B(\hat{n})$, parameterized by a full width at half maximum, θ_ν^{FWHM} , and non-uniform correlated noise, $\mathbf{n}(\hat{n})$:

$$\mathbf{d}(\hat{n}) = \mathbf{s}(\hat{n}) \otimes B(\hat{n}) + \mathbf{n}(\hat{n}). \quad (5.2)$$

We use the synthetic observations $\mathbf{d}(\hat{n})$ as inputs to the component separation algorithm. In the rest of this section we elaborate on the model choices made in each of the simulation steps.

¹https://github.com/bthorne93/PySM_public




Figure 5.1: The analytic model SEDs of the CMB, thermal dust, and synchrotron described in Section 5.2.1. Amplitudes have been rescaled to compare just the shapes of the curves. The vertical lines indicate frequencies at which SO is due to make observations; two channels characterize the low frequency synchrotron, two channels characterize the high frequency dust, and two channels observe the CMB around 100 GHz.

5.2.1 Galactic simulations

We consider two foreground models in this study, both of which have polarized synchrotron and dust components. The synchrotron emission is produced by fast-moving electrons interacting with the Galactic magnetic field. Infrared emission from dust grains in the interstellar medium comes from their absorption of light from the interstellar radiation field. Both components are polarized due to the Galactic magnetic field. The two PySM models we consider are:

SIMSET1 This corresponds to the ‘a1d1f1s1’ model of PySM, and the ‘standard’ model of SO19. It has power-law synchrotron emission with a spatially varying synchrotron index estimated from *WMAP* data, The dust emission is described by a modified blackbody with $I_\nu \propto \nu^\beta B_\nu(T)$, where T is the temperature of the dust, and β is the opacity index. It has a spatially varying dust emissivity and temperature estimated from *Planck* data. This model has no decorrelation of the components between

Table 1: Simons Observatory expected instrument properties, from Ref. [3]. The noise levels σ_I are the intensity white noise levels in $\mu K \text{ amin}$ for a sky area of $f_{\text{sky}} = 0.1$, with polarization noise $\sqrt{2}$ higher. The parameters ℓ_{knee} and α_{knee} quantify the $1/f$ model in Eq. 5.3. Here θ_{FWHM} is the full-width at half-maximum in arcminutes, but in this chapter we use $\theta_{\text{FWHM}} = 30'$ for all channels.

Frequency (GHz)	ℓ_{knee}		α_{knee}	σ_I ($\mu K \text{ amin}$)		θ_{FWHM} (')
	optimistic	pessimistic		goal	baseline	
27	30	15	-2.4	25	35	91
39	30	15	-2.4	17	21	63
93	50	25	-2.6	1.9	2.6	30
145	50	25	-3.0	2.1	3.3	17
225	70	35	-3.0	4.2	6.3	11
280	100	40	-3.0	10	16	9

frequencies, and no polarized anomalous dust. The emission of each component as a function of frequency is shown in Fig 5.1.

SIMSET2 This model modifies the synchrotron spectral index map of SIMSET1 by adding power at small scales using a Gaussian realization of a power law power spectrum $\propto \ell^{-2.6}$ [122]. This model is referred to as the ‘high-res β_s ’ model used in SO19.

5.2.2 CMB

The PySM code simulates the primary CMB by creating Gaussian realizations of a given set of theoretical power spectra, calculated for a given cosmology. We use a theoretical power spectrum for the fiducial Planck 2018 cosmological parameters, with no tensor-to-scalar ratio, $r = 0$. PySM uses the `healpy` implementation of SYNFAST to generate primary CMB realizations of temperature and polarization. It then uses the TAYLENS² [204] software to apply the displacements of the primary CMB caused by gravitational lensing.

5.2.3 Sky area and expected noise levels

As described in SO19, the Simons Observatory plans to conduct two surveys: a large survey covering about 40% of the sky, conducted by the LAT, and a smaller survey of

²<https://github.com/amaurea/taylens>

the cleanest $\sim 10 - 20\%$ of the sky using the SATs [309].

The noise model used in this study is described in detail in SO19, and we summarize it here. Two noise levels are considered: the ‘baseline’ design assumes a raw sensitivity based on the achieved performance of previous ground based experiments such as POLARBEAR, BICEP, and QUIET, and a ‘goal’ design which will require more ambitious detector development. These raw sensitivities are multiplied by an observing efficiency of 20%, accounting for all data cuts, observing downtime and instrument calibration, based on the efficiency achieved during observations at the same site by the ACT experiment.

The $1/f$ noise induced in the instrument by atmospheric loading and instrument systematics is parameterized as an additional term in the noise power spectrum that increases at large scales:

$$N_\ell = N_{\text{white}} \left[1 + \left(\frac{\ell}{\ell_{\text{knee}}} \right)^{\gamma_{\text{knee}}} \right], \quad (5.3)$$

where γ_{knee} and ℓ_{knee} are the knee index and multipole, respectively. The range of parameters we consider here, as in SO19, are summarized in Table 1. In Figure 5.2 we show the individual frequency noise curves, from SO19, for the optimistic knee multipole, and goal sensitivity, compared to the lensing B -mode spectrum.

Due to the large field of view of the SAT, the survey design has non-uniform depth. Therefore, in modeling the noise properties it is important to weight noise realizations by the relative hit density. We follow the same procedure as in SO18 to generate noise realizations with this non-uniform hit density, and non-white noise.

In SO19, the effect of the resolution of the lowest frequency channels was conducted. They found that the impact of the resolution of this channel did not affect the achieved sensitivity dramatically. In order to simplify the later analysis, we therefore make the assumption that all frequencies have the same 30' FWHM resolution.

We generate 200 Monte Carlo realizations of the noise and CMB, and use a common foreground realization for this suite of simulations.

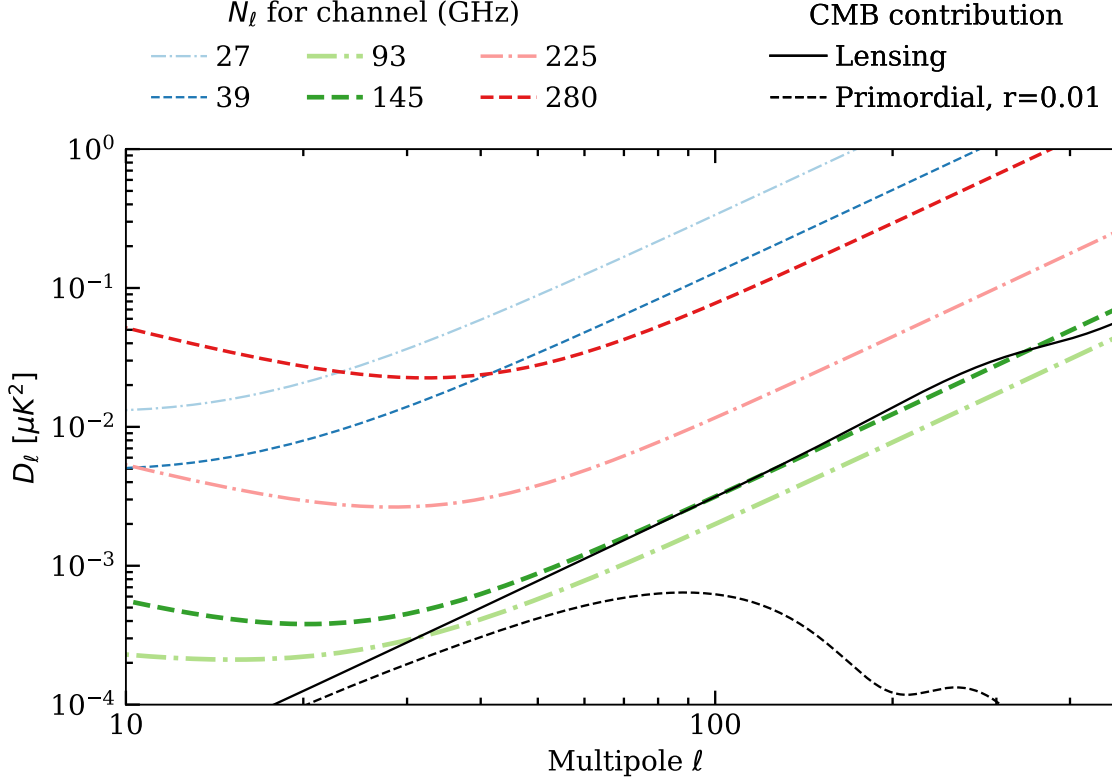


Figure 5.2: Expected noise curves for SAT polarized observations, from Ref. [3], for baseline sensitivity and optimistic ℓ_{knee} parameter. The lensing power spectrum is also shown for comparison.

5.3 Component Separation and parameter estimation

In this section we describe the map-space component separation algorithm, and the subsequent estimation of power spectra and parameters from the cleaned maps.

5.3.1 Component separation

Our component separation method follows the BFoRe method described in Ref. [233]. We take the simulated sky maps, \mathbf{d} , and model them as a linear combination of components with spatially varying SEDs, and noise:

$$\mathbf{d}(\hat{n}) = F(\theta(\hat{n})) \cdot \mathbf{T}(\hat{n}) + \mathbf{n}(\hat{n}) \quad (5.4)$$

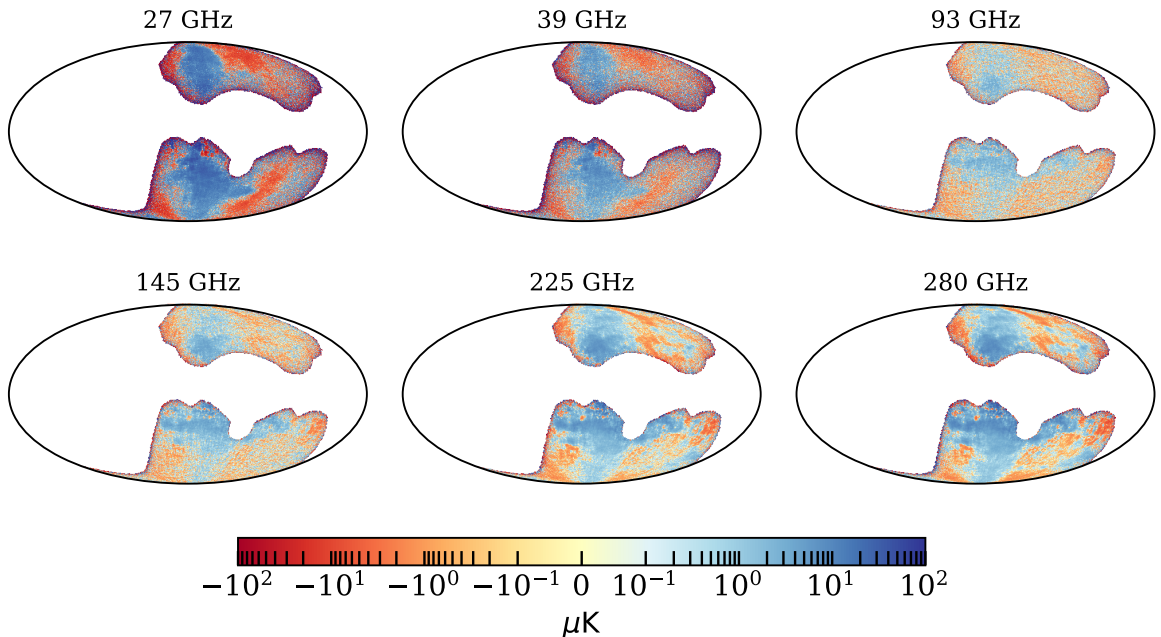


Figure 5.3: Maps in Galactic coordinates showing simulations of the Q Stokes parameter for SIMSET1 with the baseline sensitivity and optimistic ℓ_{knee} configuration of SO. Note that the colorscale is a combination of a linear scale between $(-0.1, 0.1) \mu\text{K}$ and a log scale outside this range, to display structure over a large dynamic range. In the 27 GHz channel, the large scale structure of synchrotron emission is visible, with the North Galactic Spur clearly visible above the Galactic plane. At the CMB frequencies 93 GHz and 145 GHz, the CMB E-mode polarization becomes visible at high Galactic latitudes. At the highest frequency channels the morphology of dust emission becomes dominant.

where F is the mixing matrix containing different component SEDs, with parameters $\theta = \{\beta_d(\hat{n}), T_d(\hat{n}), \beta_s(\hat{n})\}$, \mathbf{T} is a vector of component templates at a specific reference frequency, and \mathbf{n} is a noise term. Comparing to 5.2 we see that this is essentially the ‘correct’ model, modulo the permitted degree of spatial variation of the spectral parameters $\theta(\hat{n})$.

Under the assumption of Gaussian noise we can write down the likelihood for the data:

$$-2 \ln[\mathcal{L}(\mathbf{T}, \theta | \mathbf{d})] \propto (\mathbf{d} - F \cdot \mathbf{T})^T N^{-1} (\mathbf{d} - F \cdot \mathbf{T}) \quad (5.5)$$

where N^{-1} is the data covariance. We simplify the analysis by assuming that the noise is uncorrelated between pixels, allowing us to write the covariance as a diagonal matrix:

$$(N^{-1})_{(i,j,\nu),(i',j',\nu')} = \sigma_P^2 \delta_{(i,j,\nu),(i',j',\nu')}. \quad (5.6)$$

Under this assumption, the likelihood can be separated into a product over N_{spec} large

pixels in which the spectral parameters are allowed to vary. In principle these large pixels are not tied to any particular pixelization scheme, but can be implemented as any arbitrary shape intended to follow the true spatial variation of spectral parameters. The complexity of the foreground model can be increased by allowing more parameters to vary in the fit, or by increasing the number of independent patches.

Unless otherwise stated, during the rest of this chapter we maximize the likelihood in equation 5.5 by varying only the dust and synchrotron spectral indices, β_d and β_s , keeping the temperature fixed at 20K. We carry out this maximization in each large pixel, and for each Monte Carlo realization of CMB and noise. This is equivalent to marginalizing over the spectral indices for a single realization.

5.3.2 B-mode power spectrum estimation

In order to constrain the tensor-to-scalar ratio we first calculate the power spectrum from the cleaned CMB Q and U maps. The transformation from $\mathbf{P} = (Q, U)$ to $\mathbf{P}'_{\ell m} = (E_{\ell m}, B_{\ell m})$ is inherently non local as it requires the calculation of spherical harmonic transforms (SHT):

$$\mathbf{P}'_{\ell m} = \int_{4\pi} d\Omega \mathcal{Y} \cdot \mathbf{P}, \quad (5.7)$$

where \mathcal{Y} is a 2×2 matrix, with each element being a specific combination of spin-weighted spherical harmonics [193].

Ground based observations observe only part of the sky, and so can only access the true sky, multiplied by some window function, $W(\hat{n})$: $\tilde{\mathbf{P}}(\hat{n}) \equiv (W(\hat{n})Q(\hat{n}), W(\hat{n})U(\hat{n}))$. A naive calculation using the standard pseudo-power spectrum technique [185] will mix E and B modes, and if not accounted for increases the variance of the estimated B -modes, limiting the achievable constraints on the tensor-to-scalar ratio [314, 190].

To correct for this effect we use a ‘pure’ estimator of the power spectrum [190, 315, 193], which is equivalent to first calculating the naive pseudo-spectrum of the maps over some mask, $W(\hat{n})$, and then calculating and removing the leaked E modes [193]. This method assumes that the applied mask satisfies Dirichlet and Neumann boundary conditions. Therefore, we apply an additional tapering to the inverse variance map that would usually be used in the calculation of the power spectrum. In this work we use the publicly available `NaMaster`³ code. For details of the implementation see the

³<https://github.com/damonge/namaster>

NaMaster documentation and Ref. [189], and Ref. [190, 315, 193].

Due to the limited sky coverage, there is insufficient information to invert the mode coupling matrix at all multipoles. Instead we coarse-grain the matrix by defining some binning scheme. In this work we use the binning operator:

$$W_{\ell_b \ell} = \frac{1}{\Delta \ell} \Theta(\ell - \ell_b) \Theta(\ell - \ell_b + \Delta \ell), \quad (5.8)$$

where Θ is the Heaviside function, ℓ_b denotes the bandpower, and $\Delta \ell$ is the width of each bin. Then estimates of the binned power spectrum are:

$$C_{\ell_b}^{XX} = \sum_{\ell} W_{\ell_b \ell} \sum_m \frac{|\tilde{a}_{\ell m}^{XX}|^2}{2\ell + 1} \quad (5.9)$$

From the cleaned Monte Carlo simulations we calculate the expected power spectrum, and its covariance.

5.3.3 Cosmological parameter likelihood

The full posterior for the individual bandpowers is non-Gaussian. However, for high enough multipoles the central limit theorem justifies a Gaussian approximation [304]. The combination of the compact observing region and atmosphere-induced systematics, limit constraints on large scale modes. Therefore, we consider only $\ell > 30$, for which the Gaussian approximation is valid:

$$\begin{aligned} -2 \ln[\mathcal{L}(r, A_L)] &= (\hat{C}_{\ell_b}^{BB} - C_{\ell_b}^{BB}(r, A_L))^T \text{Cov}(\hat{C}_{\ell_b}^{BB}, C_{\ell_b}^{BB})^{-1} (C_{\ell_b}^{BB} - C_{\ell_b}^{BB}(r, A_L)), \\ C_{\ell_b}^{BB} &= r C_{\ell_b}^{BB, \text{prim}}(r = 1) + A_L C_{\ell_b}^{BB, \text{Lens}}(A_L = 1), \end{aligned} \quad (5.10)$$

where the measured power spectrum is indicated by the hat, and the model, $C_{\ell_b}^{BB}$, is the sum of a primordial term and a lensing term. $C_{\ell_b}^{BB, \text{Lens}}(A_L = 1)$ is a template for the lensing contribution, and $C_{\ell_b}^{BB, \text{prim}}(r = 1)$ is a template for the primordial contribution, both of which are calculated using the CLASS code with Planck 2018 cosmological parameters and $r = 1$. Equations 5.10 are given at bandpowers ℓ_b after accounting for the effects of bandpower averaging and inversion of the mode coupling matrix. We sample this likelihood for these two parameters using the emcee package,




Figure 5.4: Forecasted constraints on r , for a model with $r = 0$, from simulated observations of SIMSET1 for the four Simons Observatory noise levels. The results from this work, assuming no spatial variation of the spectral indices (black), are compared to results from the `xForecast` method presented in SO19 (grey).

and summarize the posterior on r by marginalizing over the lensing amplitude and calculating the median and standard deviation of the posterior distribution.

5.4 Results

In this section we present the results of applying the foreground separation algorithm described in Section 5.3 to the simulations described in Section 5.2.

We first present the results for SIMSET1, which are the fiducial sky simulations analyzed in SO19. As noted, the foreground-cleaning algorithm presented in this chapter is similar to the `BForE` method used in [3] and [233], but uses an independent pipeline. Therefore, we first validate our analysis, and then provide an extension of the cases studied in previous works.

5.4.1 Fiducial sky simulations

We show the resulting bias and uncertainty on r for the four Simons Observatory noise configurations in Figure 5.4, compared to the results of the `xForecast` [125] framework presented in SO19. We compare to `xForecast`, an alternative parametric

method that uses the same simulations, because the reported BFoRe forecasts in SO19 include an additional marginalization over residual foregrounds while estimating the tensor-to-scalar ratio. This comparison is therefore done for the specific `xForecast` case where no additional foreground power is marginalized over. Figure 5.4 shows that our method produces consistent forecasts with those in SO19. The nominal Simons Observatory design is biased by 1σ , and the most sensitive design is biased by 2σ . The errors from our method and `xForecast` are not identical; we have checked (private communication) that our forecasted errors agree with the BFoRe code in the case of no-marginalization.

In the rest of this section we focus on the case of optimistic ℓ_{knee} and baseline sensitivity. We demonstrate that the source of the bias in r is due to foreground contamination, primarily from residual dust. In SO19 the bias was largely removed at the power spectrum level by marginalizing over a template power spectrum for the foreground contamination. Relying on this method for an unbiased detection of primordial gravitational waves could be problematic as the shape of the residual foreground power spectrum is not known *a priori*.

In the rest of this section we establish that neglecting spatial variation of the indices is the root cause of the bias, and that it may be removed with additional masking of the higher-foreground regions, or by introducing only a few more parameters describing the spatial variation of the dust spectral parameters.

5.4.2 Establishing source of bias

5.4.2.1 Masking the Galactic plane

To establish the primordial origin of a non-zero value of r , the result must be robust against different splits of the data. In this section we establish that increasing the masking of the Galactic plane when estimating the power spectrum from cleaned CMB maps reduces the spurious detection of r , and we show that sufficient masking can remove all of the bias in r , at the expense of a small increase in the uncertainty.

We use a set of Galactic masks based on thresholds of the *Planck* intensity maps at 353 GHz, downloaded from the Planck Legacy Archive ⁴. The masks range from leaving 20 % to 80 % of the sky unmasked. We combine each *Planck* mask with the SO hits map from SO19, resulting in four maps with effective sky area from $\sim 6\%$ to

⁴HFI.Mask.GalPlane-apo0_2048_R2.00.fits




Figure 5.5: Forecasted constraints on r as a function of sky area, from simulated observations of SIMSET1 (diamonds) and SIMSET2 (squares), when foregrounds are removed assuming no spatial variation of spectral indices. Each point represents a different level of Galactic masking when computing the power spectrum of the cleaned CMB maps. The bias is reduced when the brightest sky regions are masked.

$\sim 16\%$, as calculated by:

$$f_{\text{sky}} = \frac{\langle w^2 \rangle^2}{\langle w^4 \rangle} \quad (5.11)$$

For each mask we repeat the power spectrum estimation and cosmological parameter fitting to estimate r , and present the results in Figure 5.5. As the masks become more aggressive, the bias in r is reduced substantially for both SIMSET1 and SIMSET2. Since we expect foreground residuals to be localized to the Galactic plane, a decreasing bias with increased plane masking indicates the Galactic nature of the bias.

5.4.2.2 Cross-correlating with foreground templates

If the cleaned CMB maps are contaminated by foreground residuals, we may expect there to be a significant correlation with templates of the individual foreground components. In order to perform this test using only the observed data, without relying on external observations, we form templates of high and low frequency foregrounds in the following way.

A synchrotron template is constructed by differencing the 27 GHz and 40 GHz maps, and a dust template is constructed by differencing the 270 GHz and 220 GHz




Figure 5.6: The cross-correlation of cleaned CMB maps with Galactic templates can be used to test for residual foregrounds. This shows the cross-spectrum when cleaning with a spatially constant dust SED and synchrotron spectral index (diamonds) or accounting for spatial variation (circles). We see a non-zero signal at the largest scales when using the spatially constant models, which correlates with a bias in r . (Left) This uses SIMSET1, and correlates with a template of Galactic dust. (Right) This uses SIMSET2 and correlates with a synchrotron template.

maps. In thermodynamic units the CMB contribution will cancel, and we will be left with a map proportional to the synchrotron and dust populations, respectively. The foreground templates are then correlated with the cleaned CMB maps.

In Figure 5.6 we show the cross spectra of `simset1` and `simset2` with dust and synchrotron templates, respectively, using the largest 16% sky region and cleaning with a single synchrotron and dust index. We find that the residual foregrounds, which showed up as a bias in r in Figure 5.4, can be detected in the cross-spectra.

From these exercises we find that: i) a single spectral parameter is insufficient to describe the spectral energy dependence of foregrounds in the region observed by the Simons Observatory, and results in at least a bias of 1.5×10^{-3} for the most optimistic case of foreground complexity, as was found in SO19, ii) depending on the complexity of the true foregrounds, the bias can be primarily due to dust or to synchrotron contamination, iii) the bias can be mostly removed by using only $\sim 6\%$ of the sky.

5.4.3 Spatially varying spectral indices

In this sub-section we demonstrate that the residual foreground contamination occurs because the fitting process does not account for the spatial variation of the spectral




Figure 5.7: (Left and middle) The fitting regions used to fit different dust SEDs, as defined by the algorithm described in 5.4.3.2, for $N_{\text{spec}} = 2, 6$. (Right) The fitting regions using a Healpix $N_{\text{side}} = 2$ grid.

indices, and we demonstrate that such variation may be sufficiently described by only a few additional parameters if they could be chosen to accurately capture the spatial variation.

We adapt the map-space cleaning algorithm to account for spatial variation of the foregrounds by defining a set of N_{spec} patches in which to fit spectral indices. In theory $1 \leq N_{\text{spec}} \leq N_{\text{pix}}$, where $N_{\text{spec}} = 1$ corresponds to fitting a single spectral index over the whole sky, and $N_{\text{spec}} = N_{\text{pix}}$ fits a spectral index in every pixel at the resolution of the input maps. The limit of large N_{spec} would account for as much spatial variation in the foreground SEDs as possible, however requires increasing the number of fitted parameters, and therefore post-separation noise. It would therefore be ideal to keep N_{spec} as low as possible, whilst still achieving an unbiased estimation of r . These simulations were generated at N_{pix} resolution, therefore, there would be no advantage to choosing $N_{\text{spec}} > N_{\text{pix}}$. Note that, for observations of the real sky, the creation of pixelized maps requires spatial averaging both along the line of sight, and transverse to it. Such averaging could be better fitted by taking $N_{\text{spec}} > N_{\text{pix}}$, as the extra parameters could absorb some of the resulting SED curvature [146, 122].

How to distribute patches on the sky is an important consideration, as it involves a choice of how to model the spatial dependence of foreground parameters. This is unknown *a priori*, so we explore an initial exploration of possible approaches.

5.4.3.1 Fitting spectral indices on a HEALPix grid

Some previous studies have used patches corresponding to a coarse HEALPix grid for convenience [233], allowing an independent spectral parameter to be fit in each coarse pixel, as shown in the right panel of Figure 5.7. We define a set of fitting models with spectral parameters varying on grids at $n_{\text{side}} = 1, 2, 4$ (with $n_{\text{pix}} = 48, 192, 3072$), and perform the foreground cleaning. We show our resulting constraints on r in Figure 5.8. We find that the errors are significantly inflated as n_{side} increases, but determine that

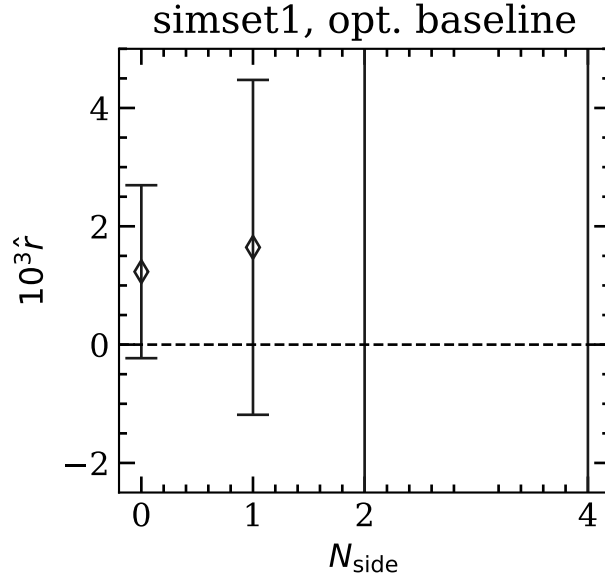


Figure 5.8: Forecasted constraints on r when fitting spectral indices defined by HEALPix grids at increasing resolution. Since SO observes a fraction of the sky, the number of coarse pixels is given by $12f_{\text{sky}}N_{\text{side}}^2$. As N_{side} increases we find the projected uncertainty on r increases significantly: this is not a good choice for dividing up the sky area.

this is partly due to many of the larger Healpix pixels having too few observed pixels, and also due to having so many additional parameters.

5.4.3.2 With prior information on spatial variation of indices

With prior information about the spatial variation of the spectral parameters we investigate how many additional parameters would be needed to mitigate the observed biases in r due to foreground residuals. Here we use the dust spectral index map, β_{d1}^{dust} , used in PySM model ‘d1’, and apply Algorithm 1, to create a template that defines the regions in which we fit spectral parameters.

We repeat Algorithm 1 for $N_{\text{spec}} \in 1, \dots, 12$, to produce twelve different template maps, with increasing spatial resolution. Figure 5.7 shows the regions for $N_{\text{spec}} = 2, 6$.

We run our foreground cleaning algorithm on the SIMSET1 sky simulations for these spectral index pixelizations described above. The resulting constraints we find on r are shown in Figure 5.9 for $N_{\text{spec}} = 1, 4, 12$. We find that increasing N_{spec} results in a reduced bias on r , as we are able to account for spatial variation of β_d , with negligible inflation of the uncertainty. No improvement of the bias is achieved past a value of $N_{\text{spec}} \approx 4$. Since the fitting code is currently designed to fit all parameters in the same

Algorithm 1 This algorithm is used to produce a set of fitting regions for an input spectral index map, β , and integer N_{spec} .

- 1: Resample β at $n_{\text{side}} = 256$ to obtain the map $\beta_{256}(i)$, where i is the pixel index .
 - 2: Restrict to the subset of pixels observed by SO, β_{256}^{SO} .
 - 3: Create histogram of β_{256}^{SO} with N_{spec} equally sized bins, with edges β_j , where $j = 0, \dots, N_{\text{spec}}$.
 - 4: Create an empty $n_{\text{side}} = 256$ map, $b(i)$.
 - 5: For each pixel in b assign $b(i) = j$, where j is the number of the bin into which $\beta_{256}^{\text{SO}}(i)$ falls.
 - 6: Return $b(i)$.
-

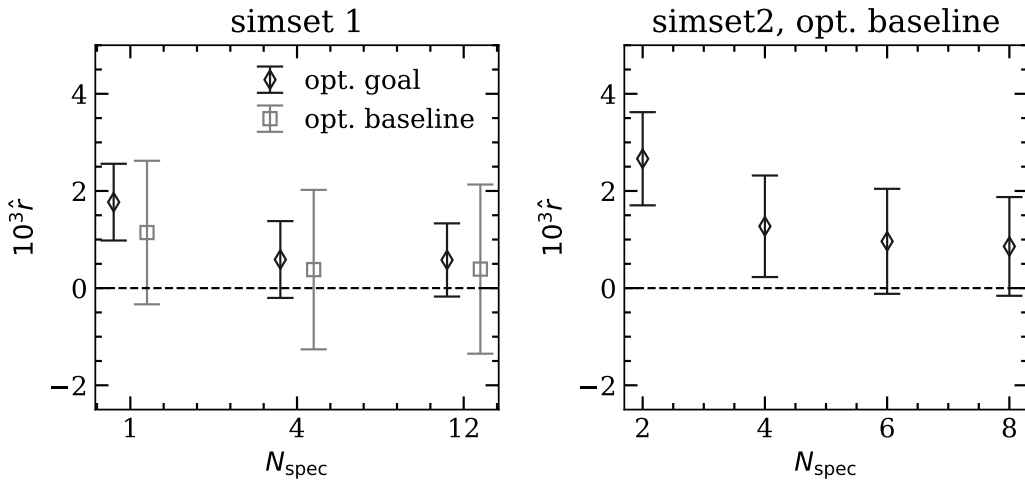


Figure 5.9: (Left) Forecasted constraints on r when fitting models with increasing independent regions defined from the β_{d1}^{dust} template, for the optimistic ℓ_{knee} goal and baseline sensitivities. The bias is removed with a modest increase in parameters, if these regions are known a-priori. Right: using SIMSET2 and fitting models with regions defined from the β_{s2}^{sync} template, a similar effect is seen

patches, this remaining bias is due to the mismatch between the dust-focused fitting regions and synchrotron spectral index variation. Modifying the fitting code to allow different physical parameters to be fitted with different constraints is a non-trivial extension, and is left to future work.

In Figure 5.6 we show the cross-correlation of the cleaned CMB maps with the dust foreground templates, as a function of number of fitting pixels N_{spec} . Increasing the number of fitting regions results in a better cleaning of the foregrounds, and consequently we find no signature of residual contamination by $N_{\text{spec}} \sim 4$.

We then repeat this study for SIMSET2, which has a synchrotron spectral index with a greater spatial variance, β_{s2}^{sync} . Ignoring spatial variation of the foreground SED in this case results in an even larger bias on r than in the case of SIMSET1. Here

we use β_{s2}^{sync} to define the sky regions used for fitting the indices, and $N_{\text{spec}} = 2, \dots, 8$ as the input for Algorithm 1. Our resulting constraints on r are shown in Figure 5.9. Again, there is a reduction in the bias on r as we increase N_{spec} , and the improvement saturates at $N_{\text{spec}} \approx 6$. For $N_{\text{spec}} \gtrsim 6$ we see no improvement as the bias is now dominated by residuals due to dust mis-modeling. Indeed the remaining bias for $N_{\text{spec}} \gtrsim 6$ is at the $\sim 1\sigma$ level found when fitting SIMSET1 with a single spectral index.

In Figure 5.9 we show the correlation of the cleaned CMB maps with a synchrotron template formed from the two lowest frequency channels. Similar to SIMSET1 we find that increasing N_{spec} removes the hints of foreground bias at low multipoles seen when ignoring spatial variation of spectral indices.

It is promising that with only 4-6 additional parameters describing the spatial variation of the synchrotron and dust indices, that a 1-3 σ bias in r might be mitigated without significant masking. In practice, though, we will not have the perfect information about their spatial variation. There may also be additional variation not captured in the existing simulations. Instead, we would need to derive the fitting regions directly from the data.

Towards this, we determined that the χ^2 of the model compared to the data, with the intent to add spatial resolution to the fitting model in areas of poor χ^2 , was not sufficiently sensitive to discriminate between areas of good and poor fit. A judicious choice of these fitting regions, based on existing and upcoming observations, will be a natural direction for future study.

5.5 Conclusions

By accounting for Galactic foreground models based on current data, we have assessed how large-scale B -mode observations of the Simons Observatory might be used to constrain the tensor-to-scalar ratio, following on from the forecast study in SO19. We showed that the nominal SO design may result in a biased estimation of r due to foreground contamination, which can either be mitigated by marginalizing over a foreground residual after cleaning the maps, or by restricting the sky area to discard the most contaminated region. This masking removes or reduces this bias to one standard deviation or less. In a further alternative approach, we defined a scheme to allow for spatial variation in the parameterization of the fitted foreground model. We found that using Healpix pixels was unsatisfactory, as the coarse-grained pixels that overlapped the observed regions were not fitted reliably, and could lead to large

residuals. We also defined a new spatial parameterization that split the sky into an arbitrary number of regions following the true morphology of the spectral behavior of the dominant foreground contaminant. With this approach, with perfect knowledge of the spectral index variation, we found that fitting 4-6 independent regions was sufficient to remove the dominant r -bias. This points to a direction for further exploration in choosing the regions from the real data. We found that the cross-correlation of the cleaned maps with the tracers of synchrotron and dust provide a useful way to check for residual foregrounds that bias the estimate for r . In practice, a comprehensive analysis of the real data from SO will likely implement a set of these bias-mitigation approaches, coupled with alternative non-parametric methods for foreground cleaning, internal null tests to check for consistency, and tests for residual non-Gaussianity of the maps.

Chapter 6

Conclusions

In Chapter 2 we introduced a new code for the simulation of Galactic foregrounds, `PySM`. The models we presented spanned a range of complexities, for use in forecasting the performance of future instruments, and validating foreground analysis methods. We then provided an example use case by forecasting the constraints on the tensor-to-scalar ratio achievable in the presence of foregrounds for a future satellite observatory. `PySM` simulations have been used to: forecast the success of the large-scale B -mode observations of the Simons Observatory [3]; study the impact of foregrounds in the CMB-Stage 4 science book [124] and Concept Definition Task Force Report¹; forecast B -mode constraints for the LiteBIRD experiment [125]; in the validation of the most recent BICEP / Keck analysis pipeline that provides the most competitive constraints yet on primordial gravitational waves [41], as well as others. The `PySM` code continues to be developed, and is currently being updated to include comprehensive extragalactic simulations, as well as new models for Galactic emission based on magnetohydrodynamic simulations of the interstellar medium. These MHD simulations allow both dust and synchrotron to be simulated with the same Galactic magnetic field, as well as providing a physically-motivated picture for the smallest scales of polarized foreground morphology, which are currently poorly constrained by observations.

Later, in Chapter 5, we considered future observations made by the Simons Observatory (SO). We simulated various models of the Galactic sky using `PySM`, and added realistic inhomogeneous non-white noise to the simulated sky signal. We then forecasted achievable constraints on the tensor-to-scalar ratio, r . We demonstrated that ignoring the spatial variation of the foreground SEDs would result in a spurious detection of r by a standard deviation or more. We showed that this bias could be removed by increased masking of the Galaxy, or marginalization over a residual template. We went on to consider a scheme designed to split up the sky into regions of similar spectral index values, based on a perfect knowledge of the true underlying variation. We showed that four to six such regions were sufficient to remove the bias in r . This success indicates useful future work could be done to define these regions from real data, without prior knowledge of the true variation foreground spectral dependencies.

In Chapter 4, we considered observations made by the Atacama Cosmology Telescope (*ACT*). We used publicly available high-frequency observations from the *Planck* satellite to calculate the expected level of dust contamination at frequencies, and in regions of the sky, surveyed by the *ACT* experiment, both in temperature and polarization. Our results for the temperature power spectrum of dust were consistent

¹CMB-S4 collaboration, 2017

with existing analyses in Refs. [12, 13]. On the other hand, the results for the polarized dust power spectra were largely inconclusive, due to the low signal-to-noise ratio of the polarized *Planck* maps, especially at the high Galactic latitudes targeted by the *ACT* instrument. We concluded that the treatment of Galactic dust in the analysis of new data from the *ACTPol* instrument will not be a limiting factor, and our constraints on the dust amplitudes will be used as priors on the foreground model used in the *ACT* cosmological likelihood.

There are many possible models of inflation compatible with current measurements. In Chapter 3 we considered a model in which an additional set of $SU(2)$ gauge fields is present during inflation. These gauge fields source only a single parity of tensor perturbations, giving rise to a chiral gravitational wave background, and breaking the simple relationship between the energy scale of inflation and the amplitude of the gravitational wave background. We considered unique observables of the $SU(2)$ model, and forecasted the ability of experiments such as S4, LiteBIRD, and LISA to distinguish it from the simplest case of single field slow-roll inflation. We found that a LiteBIRD-like experiment would achieve a signal-to-noise of ~ 2 in parity-violating signatures on the CMB (combining information from TB and EB power spectra), whilst S4 would be virtually unable to constrain the chirality of this model, as most of the information in TB and EB is contained in large-scale anisotropies. Finally, we considered an advanced stage of the Laser Interferometer Space Antenna, and a more futuristic space-based interferometer with a longer baseline, the Big Bang Observatory. These experiments probe a range of much smaller wavenumbers than the CMB, and so provide useful, complementary, constraints on the blue-tilted tensor power spectra allowed by the $SU(2)$ model.

We concluded that, in the event of a positive detection of r , it will be vital to demonstrate the true origin of tensor fluctuations by excluding such well-motivated alternative models. This will be hard to do using standard two-point statistics of the CMB, and instead will need to leverage the non-Gaussianity of such models, or the strong scale dependence of the BB power spectrum.

Appendix A

Derivation of the template for GW power spectrum

In Ref.[42], it has been shown that the power spectrum of the sourced GW is given by

$$\mathcal{P}_h^{\text{L,Sourced}}(k) = \frac{\epsilon_B H_{\text{inf}}^2}{\pi^2 M_{\text{Pl}}^2} \mathcal{F}^2(m_Q), \quad (\text{A.1})$$

where H_{inf} is the inflationary Hubble scale, $\epsilon_B \equiv g^2 Q^4 / (M_{\text{Pl}}^2 H_{\text{inf}}^2)$ roughly indicates the energy fraction of the SU(2) gauge field. $\mathcal{F}(m_Q)$ is a monotonically increasing function for $3 \leq m_Q \leq 7$ which is well approximated by

$$\mathcal{F}(m_Q) \simeq \exp [2.4308m_Q - 0.0218m_Q^2 - 0.0064m_Q^3 - 0.86], \quad (3 \leq m_Q \leq 7), \quad (\text{A.2})$$

where the value of a dynamical parameter $m_Q(t) \equiv gQ(t)/H_{\text{inf}}$ around the horizon crossing $k \sim aH_{\text{inf}}$ is substituted. Solving the background equations of motion for $\chi(t)$ and $Q(t)$ with the slow-roll approximation, one can show

$$m_Q(t) = m_* \sin^{1/3} [\chi(t)/f]. \quad (\text{A.3})$$

where $m_* \equiv (g^2 \mu^4 / 3\lambda H_{\text{inf}}^4)^{1/3}$ is the maximum value of $m_Q(t)$. From the definition of m_Q and ϵ_B , the value of ϵ_B at $m_Q = m_*$ is $\epsilon_{B*} \equiv H^2 m_*^4 / g^2 M_{\text{Pl}}^2$. Therefore the

tensor-to-scalar ratio r on the peak scale k_p of the sourced GW power spectrum is

$$r_* = \frac{\mathcal{P}_h^{\text{L,Sourced}}}{\mathcal{P}_\zeta}(k_p) = \frac{\epsilon_{B*} H_{\text{inf}}^2}{\pi^2 M_{\text{Pl}}^2 \mathcal{P}_\zeta} \mathcal{F}^2(m_*). \quad (\text{A.4})$$

Next, we consider the width of the GW spectrum. Around the peak of $m_Q(t)$ at $t = t_*$, or $\chi(t = t_*) = \pi f/2$, $\chi(t)$ is expanded as

$$\chi(t) \simeq \frac{\pi}{2} f + \dot{\chi}_*(t - t_*) \simeq f \left[\frac{\pi}{2} + \frac{2\xi_*}{\lambda} H_{\text{inf}}(t - t_*) \right], \quad (\text{A.5})$$

where $\dot{\chi}_* \equiv \dot{\chi}(t = t_*)$, $\xi_* \equiv \lambda \dot{\chi}_*/(2f H_{\text{inf}})$ and one can show $\xi_* \simeq m_* + m_*^{-1}$ in the slow-roll regime. Then we obtain the approximated equation for $m_Q(t)$ which is valid around the peak value ,

$$m_Q(t) \simeq m_* \left[1 - \frac{1}{6} \left(\frac{H_{\text{inf}}(t - t_*)}{\Delta N} \right)^2 \right], \quad (t \sim t_*), \quad (\text{A.6})$$

where we define $\Delta N \equiv \lambda/2\xi_*$. Substituting it into eq. (A.1) and using $H_{\text{inf}}(t - t_*) = \ln(k/k_p)$, we obtain the leading order result as

$$\mathcal{P}_h^{\text{L,Sourced}}(k) \simeq \frac{\epsilon_{B*} H_{\text{inf}}^2}{\pi^2 M_{\text{Pl}}^2} \mathcal{F}^2(m_*) \times \exp \left[-\mathcal{G}(m_*) \frac{\ln^2(k/k_p)}{\Delta N^2} \right], \quad (\text{A.7})$$

with $\mathcal{G}(m_*) \approx 0.666 + 0.81m_* - 0.0145m_*^2 - 0.0064m_*^3$. Note that the contribution from $\epsilon_B(t) \propto m_Q^4(t)$ in the prefactor should not be missed. Comparing it with the template eq. (3.3), one finds

$$\sigma^2 = \frac{\Delta N^2}{2\mathcal{G}(m_*)}. \quad (\text{A.8})$$

The validity of the derived expression for $\mathcal{P}_h^{\text{L,Sourced}}(k)$ is checked by the comparison with the full numerical result. Once r_* , m_* , ϵ_{B*} and ΔN are fixed, all the model parameters g , λ , μ and f are determined. Then we can numerically solve the background equation of $\chi(t)$ and $Q(t)$ as well as the equations for the perturbations $t_L(k, t)$ and $h_L(k, t)$ to obtain the power spectrum of the sourced GW. In Figure. A.1, we compare the derived expression with the full numerical result. It should be noted that eq. (A.1) and our derivation rely on the slow-roll approximation. The approximation is less accurate for a small ΔN , because ΔN characterizes the time scale of $\chi(t)$ rolling down




Figure A.1: The comparison between the full numerical result of $\mathcal{P}_h^{\text{Sourced}}(k)$ (blue solid line) and the template eq. (3.3) with eq. (A.4) and $\sigma^2 = 0.15\Delta N^2$ (red dashed line). In the left (right) panel, $\Delta N = 5(10)$, $m_* = 4$, $\epsilon_{B*} \approx 9 \times 10^{-4}$ and the peak amplitude reaches the tensor-to-scalar ratio, $r_* = 0.05$. The Hubble parameter is set as $H_{\text{inf}} = 8 \times 10^{11}$ GeV which corresponds to $r = 10^{-5}$ without the sourced GW. In the case of $\Delta N = 5$, the derived formula slightly underestimate the peak amplitude and the width, while the fit is excellent for $\Delta N \gtrsim 10$.

its potential. In Figure. A.1, one can find a small deviation in the case of $\Delta N = 5$, while the excellent agreement is seen for $\Delta N = 10$.

Finally, we discuss how long it takes χ to get to χ_* , given that the initial value of χ is negligibly small compared to $f\pi/2$. Assuming $\chi(t \approx 0) \ll f$ and using eq. (A.5), one finds

$$Ht_* \sim \frac{\pi}{2}\Delta N. \quad (\text{A.9})$$

However, it is definitely underestimated, because $\dot{\chi}(t \ll t_*)$ must be smaller than $\dot{\chi}(t = t_*)$ which is the maximum value. In fact, a full numerical calculation shows that the coefficient is somewhat larger,

$$Ht_* \approx 1.8\Delta N. \quad (\text{A.10})$$

One may wonder if $\chi(t)$ can stay on the top of its potential hill for a longer time if its initial value is small enough. However, since χ is coupled to the SU(2) gauge fields and the system quickly goes to the attractor behavior, the time scale of the motion of χ is almost solely determined by ΔN . It indicates that the peak scale k_p should be smaller than $k_i \exp[1.8\Delta N]$. Here k_i is the wave number of the mode exiting the horizon at the initial time, and it is smaller or roughly equals to the largest CMB




Figure A.2: Peak tensor-to-scalar ratio r_* as a function of ϵ_{B*} and m_* for $k_p = 0.005 \text{ Mpc}^{-1}$.

scale. Therefore we obtain the following constraint on ΔN ,

$$\Delta N \gtrsim \frac{1}{1.8} \ln \left(\frac{k_p}{k_{\text{CMB}}} \right). \quad (\text{A.11})$$

Appendix B

Calculation of the covariance matrix

For a given beam, b_ℓ , and a white noise level, $w_{X_1 X_2}^{-1}$, the expected variance of the multipoles of an observed sky is given by:

$$\langle (a_{\ell m}^{X_1})^* a_{\ell' m'}^{X_2} \rangle = (|b_\ell|^2 C_\ell^{X_1 X_2} + w_{X_1 X_2}^{-1}) \delta_{\ell \ell'} \delta_{m m'}. \quad (\text{B.1})$$

An unbiased estimator of the angular power spectrum is then:

$$\hat{C}_\ell^{X_1 X_2} = |b_\ell|^2 \left(\sum_{m=-\ell}^{\ell} \frac{(a_{\ell m}^{X_1})^* a_{\ell m}^{X_2}}{2\ell + 1} - w_{X_1 X_2}^{-1} \right) \quad (\text{B.2})$$

By considering the expectation $\langle (\hat{C}_\ell^{X_1 X_2} - C_\ell^{X_1 X_2})(\hat{C}_\ell^{X_3 X_4} - C_\ell^{X_3 X_4}) \rangle$ it can then be shown that the covariance is given by [44]:

$$\xi^{X_1 X_2 X_3 X_4} = \frac{1}{(2\ell + 1) f_{\text{sky}}} (\tilde{C}_\ell^{X_1 X_3} \tilde{C}_\ell^{X_2 X_4} + \tilde{C}_\ell^{X_1 X_4} \tilde{C}_\ell^{X_2 X_3}). \quad (\text{B.3})$$

where $\tilde{C}_\ell^{X_1 X_2} = C_\ell^{X_1 X_2} + |b_\ell|^{-2} w_{X_1 X_2}^{-1}$.

Appendix C

CMB noise spectrum

For a given set of experimental parameters such as channel frequencies, FWHM and sensitivity in polarization and temperature per channel we want to find the aggregate noise in the CMB spectra. We follow the treatment of Ref. [246], which itself closely follows Ref. [270].

There are multiple sources of noise in the final spectrum: instrumental noise in the CMB channels, residual foreground noise from incomplete cleaning, and additional systematic noise introduced from the templates used in cleaning the CMB channels.

The noise in the final CMB spectrum is:

$$N_\ell^{\text{BB}} = \left[\sum_i \frac{1}{n_\ell(\nu_i) + [C_\ell^{\text{S}}(\nu_i) + C_\ell^{\text{D}}(\nu_i)] \sigma_{\text{RF}} + n_\ell^{\text{RF}}(\nu_i)} \right]^{-1} \quad (\text{C.1})$$

where the index i runs over channels used in CMB analysis, RF refers to residual foregrounds, $n_\ell(\nu)$ is the noise spectrum in the channels used for CMB analysis, $[C_\ell^{\text{S}}(\nu_i) + C_\ell^{\text{D}}(\nu_i)] \sigma_{\text{RF}}$ is the residual foreground level in dust and synchrotron rescaled to the frequencies used in CMB analysis, and $n_\ell^{\text{RF}}(\nu_i)$ is the instrumental uncertainty in the process of foreground removal.

The simplest of the above terms is the noise in the CMB channels:

$$n_\ell(\nu) = \sigma_P^2(\nu) \exp \left[\frac{\ell(\ell+1) \left(\frac{\pi}{10800} \theta_{\text{FWHM}}(\nu) \right)^2}{8 \ln(2)} \right],$$

where $\theta_{\text{FWHM}}(\nu)$ is the FWHM of the channel ν in arcminutes. The instrumental

Parameter	Value
A_S	$4.7 \times 10^{-5} \mu\text{K}^2$
α_S	-3
β_S	-2.6
$\nu_{S,0}$	30 GHz
$\ell_{S,0}$	350
A_D	$1 \mu\text{K}^2$
α_D	2.2
β_D	-2.5
$\nu_{D,0}$	94 GHz
$\ell_{D,0}$	10
T	18 K
p	0.15

Table C.1: Spectral parameters used in noise model taken from Ref. [246].

uncertainties in the process of foreground removal are given by Ref. [270]:

$$n_\ell^{\text{RF}} = \frac{4}{N_{\text{chan}}(N_{\text{chan}} - 1)} \left[\sum_j \frac{1}{n_\ell(\nu_j)} \right]^{-1} \left[\left(\frac{\nu}{\nu_{S,ref}} \right)^{2\alpha_S} + \left(\frac{\nu}{\nu_{D,ref}} \right)^{2\alpha_D} \right],$$

where N_{chan} is the number of channels used in foreground cleaning (in this case $N_{\text{chan}} = 10$), and $\nu_{S,ref}, \nu_{D,ref}$ are the highest and lowest frequency channel used in the removal (in this case $\nu_{S,ref} = 30$ GHz, $\nu_{D,ref} = 94$ GHz). The foreground spectra are:

$$C_\ell^S(\nu) = A_S \left(\frac{\nu}{\nu_{S,0}} \right)^{2\alpha_S} \left(\frac{\ell}{\ell_{S,0}} \right)^{\beta_S}$$

$$C_\ell^D(\nu) = p^2 A_D \left(\frac{\nu}{\nu_{D,0}} \right)^{2\alpha_D} \left(\frac{\ell}{\ell_{D,0}} \right)^{\beta_D} \left[\frac{e^{\frac{h\nu_{D,0}}{k_B T}} - 1}{e^{\frac{h\nu}{k_B T}} - 1} \right].$$

These are converted into a Gaussian addition to the noise by the factor σ^{RF} such that a 2% residual level corresponds to $\sigma^{\text{RF}} = 4 \times 10^{-4}$.

The spectral parameters of the foregrounds are summarized in Table C.1. They are taken from Ref. [270], and are consistent with the 2015 Planck data.

Channel (GHz)	θ_{FWHM} (amin)	$\sigma_{\text{P}}(\nu)$ [μKamin]
40.0	69.0	36.8
50.0	56.0	23.6
60.0	48.0	19.5
68.0	43.0	15.9
78.0	39.0	13.3
89.0	35.0	11.5
100.0	29.0	9.0
119.0	25.0	7.5
140.0	23.0	5.8
166.0	21.0	6.3
195.0	20.0	5.7
235.0	19.0	7.5
280.0	24.0	13.0
337.0	20.0	19.1
402.0	17.0	36.9

Table C.2: Summary of the LiteBIRD specifications ($f_{\text{sky}} = 0.5$).

Appendix D

Frequency dependence of CMB sensitivity

When we calculate the CMB angular power spectrum we are decomposing the signal into multipoles corresponding to certain angular distance on the sky. Each multipole has contributions from all frequencies of the GWB, determined by an integral of transfer functions:

$$C_\ell^{YY'} = 4\pi \int \frac{dk}{k} [\mathcal{P}_h^L(k) - \mathcal{P}_h^R(k)] \Delta_{Y,\ell}^h(k) \Delta_{Y',\ell}^h(k).$$

This makes a direct link between multipole and frequency ambiguous. Since the transfer functions are sharply peaked at $k_\ell = \ell/\eta_0$ with η_0 denoting the comoving distance to the last scattering surface. We make the approximation:

$$\begin{aligned} C_\ell^{YY'}(k_\ell) &= \left[4\pi \int \frac{dk}{k} [\mathcal{P}_h^L(k, r_* = 1) - \mathcal{P}_h^R(k, r_* = 1)] \Delta_{Y,\ell}^h(k) \Delta_{Y',\ell}^h(k) \right] \\ &\quad \times (\mathcal{P}_h^L(k_\ell, r_*) - \mathcal{P}_h^R(k_\ell, r_*)) \\ &= C_\ell^{YY'}(r_* = 1) (\mathcal{P}_h^L(k_\ell, r_*) - \mathcal{P}_h^R(k_\ell, r_*)) \end{aligned} \tag{D.1}$$

To calculate the sensitivity to a circular background we calculate the signal-to-noise of the TB spectrum, ignoring the small contribution from EB for simplicity. The

signal-to-noise is therefore:

$$(S/N)_{\text{TB}, \ell}^2 = (2\ell + 1) f_{\text{sky}} \frac{(C_\ell^{TB})^2}{\hat{C}_\ell^{TT} \hat{C}_\ell^{BB}},$$

where over-hat indicates the observed spectrum, including foreground residuals, instrument noise, and lensing. Our assumption that the transfer function is strongly peaked at k_ℓ now allows us to write this as a function of k_ℓ instead of just ℓ :

$$(S/N)_{\text{TB}}^2(k_\ell) = (2\ell + 1) f_{\text{sky}} \frac{(C_\ell^{TB}(k_\ell))^2}{\hat{C}_\ell^{TT} \hat{C}_\ell^{BB}}.$$

Note that we still calculate the observed spectrum fully. We then ask the question: what is the required $P_h^L(k_\ell)$ (take $P_h^R = 0$) to achieve a signal-to-noise of one in the channel k_ℓ ? This will be the minimum GWB detectable with a signal-to-noise of one. So:

$$(\mathcal{P}_h^L(k_\ell, r_*) - \mathcal{P}_h^R(k_\ell, r_*))_{\text{min}} = \sqrt{\frac{\hat{C}_\ell^{TT} \hat{C}_\ell^{BB}}{(2\ell + 1) f_{\text{sky}}}} [C_\ell^{TB}(r_* = 1)]^{-1}.$$

This quantity tells us about the tensor spectrum at recombination, however in order to compare with interferometers which are sensitive to the current GWB, we have to evolve this forward in time. The tensor spectrum transfer function for CMB scales is [256, 316]:

$$\Omega_V^{\text{min}} h^2 = 1875 (\mathcal{P}_h^L(k_\ell) - \mathcal{P}_h^R(k_\ell))_{\text{min}} \left(\frac{3j_2(k_\ell \eta_0)}{k_\ell \eta_0} \frac{k_\ell}{k_*} \right)^2 \quad (\text{D.2})$$

Bibliography

- [1] Thorne, B., Dunkley, J., Alonso, D., et al. *MNRAS*, **469**(3):2821 (2017).
- [2] Thorne, B., Fujita, T., Hazumi, M., et al. *Phys. Rev. D*, **97**(4):043506 (2018).
- [3] Ade, P., Aguirre, J., Ahmed, Z., et al. *JCAP*, **2019**:056 (2019).
- [4] Baumann, D. *arXiv e-prints*, arXiv:0907.5424 (2009).
- [5] Planck Collaboration, Y., Akrami, Arroja, F., Ashdown, M., et al. *arXiv e-prints*, arXiv:1807.06211 (2018).
- [6] Scott, D. and Smoot, G. *arXiv e-prints*, astro-ph/0601307 (2006).
- [7] Planck Collaboration, Y., Akrami, Arroja, F., Ashdown, M., et al. *arXiv e-prints*, arXiv:1807.06205 (2018).
- [8] Planck Collaboration, Aghanim, N., Akrami, Y., et al. *arXiv e-prints*, arXiv:1807.06209 (2018).
- [9] Planck Collaboration, Akrami, Y., Ashdown, M., et al. *arXiv e-prints*, arXiv:1807.06208 (2018).
- [10] Planck Collaboration, Ade, P. A. R., Alves, M. I. R., et al. *A&A*, **576**:A107 (2015).
- [11] De Bernardis, F., Stevens, J. R., Hasselfield, M., et al. In ‘Observatory Operations: Strategies, Processes, and Systems VI,’ volume 9910 of *PROCSPIE*, page 991014 (2016).
- [12] Das, S., Louis, T., Nolta, M. R., et al. *JCAP*, **4**:014 (2014).
- [13] Louis, T., Grace, E., Hasselfield, M., et al. *JCAP*, **6**:031 (2017).

- [14] Dunkley, J., Calabrese, E., Sievers, J., et al. *JCAP*, **7**:025 (2013).
- [15] Guth, A. H. *Phys. Rev. D*, **23**:347 (1981).
- [16] Linde, A. D. *Physics Letters B*, **108**:389 (1982).
- [17] Albrecht, A. and Steinhardt, P. J. *Phys. Rev. Lett.*, **48**:1220 (1982).
- [18] Ijjas, A. and Steinhardt, P. J. *Classical and Quantum Gravity*, **35**(13):135004 (2018).
- [19] Zaldarriaga, M. and Seljak, U. *Phys. Rev. D*, **55**:1830 (1997).
- [20] Kamionkowski, M., Kosowsky, A., and Stebbins, A. *Phys. Rev. D*, **55**:7368 (1997).
- [21] Friedmann, A. *Zeitschrift für Physik*, **21**:326 (1924).
- [22] Lemaître, G. *Annales de la Societe Scientifique de Bruxelles*, **A53**:51 (1933).
- [23] Robertson, H. P. *ApJ*, **82**:284 (1935).
- [24] Walker, A. G. *Proceedings of the London Mathematical Society*, **s2-42**(1):90 (1937).
- [25] Einstein, A. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin)*, Seite 1030-1085. (1914).
- [26] Riess, A. G., Macri, L. M., Hoffmann, S. L., et al. *ApJ*, **826**:56 (2016).
- [27] Riess, A. G., Casertano, S., Yuan, W., et al. *ApJ*, **855**:136 (2018).
- [28] Riess, A. G., Casertano, S., Yuan, W., et al. *ApJ*, **861**:126 (2018).
- [29] Lifshitz, E. M. *Zhurnal Eksperimentalnoi i Teoreticheskoi Fiziki*, **16**:587 (1946).
- [30] Fixsen, D. J. *ApJ*, **707**:916 (2009).
- [31] Martin, J. *arXiv e-prints*, arXiv:1807.11075 (2018).
- [32] Martin, J., Ringeval, C., and Vennin, V. *Physics of the Dark Universe*, **5**:75 (2014).
- [33] Ryder, L. H. *Quantum Field Theory* (1996).

- [34] Leach, S. M., Liddle, A. R., Martin, J., et al. *Phys. Rev. D*, **66**:023515 (2002).
- [35] Baumann, D., Jackson, M. G., Adshead, P., et al. volume 1141 of *American Institute of Physics Conference Series*, pages 10–120 (2009).
- [36] Schwarz, D. J., Terrero-Escalante, C. A., and García, A. A. *Physics Letters B*, **517**:243 (2001).
- [37] Dodelson, S. *Modern cosmology*. Academic Press, San Diego, CA (2003).
- [38] Mukhanov, V. *Physical Foundations of Cosmology*. Cambridge Univ. Press, Cambridge (2005).
- [39] Steven, W. *Cosmology*. Oxford University Press (2008).
- [40] Hinshaw, G., Larson, D., Komatsu, E., et al. *ApJS*, **208**(2):19 (2013).
- [41] BICEP2 Collaboration, Keck Array Collaboration, Ade, P. A. R., et al. *Phys. Rev. Lett.*, **121**:221301 (2018).
- [42] Dimastrogiovanni, E., Fasiello, M., and Fujita, T. *JCAP*, **1**:019 (2017).
- [43] Saito, S., Ichiki, K., and Taruya, A. *JCAP*, **9**:002 (2007).
- [44] Gluscevic, V. and Kamionkowski, M. *Phys. Rev. D*, **81**(12):123529 (2010).
- [45] Penzias, A. A. and Wilson, R. W. *ApJ*, **142**:419 (1965).
- [46] Smoot, G. F., Bennett, C. L., Kogut, A., et al. *ApJL*, **396**:L1 (1992).
- [47] Hinshaw, G., Larson, D., Komatsu, E., et al. *Ap J Supp*, **208**:19 (2013).
- [48] The Polarbear Collaboration: P. A. R. Ade, Akiba, Y., Anthony, A. E., et al. *ApJ*, **794**:171 (2014).
- [49] Keisler, R., Hoover, S., Harrington, N., et al. *ApJ*, **807**:151 (2015).
- [50] Naess, S., Hasselfield, M., McMahon, J., et al. *JCAP*, **10**:007 (2014).
- [51] BICEP2 Collaboration, Keck Array Collaboration, Ade, P. A. R., et al. *Phys. Rev. Lett.*, **116**(3):031302 (2016).
- [52] Hu, W. and Sugiyama, N. *ApJ*, **444**:489 (1995).
- [53] Ma, C.-P. and Bertschinger, E. *ApJ*, **455**:7 (1995).

- [54] Seljak, U. *ApJL*, **435**:L87 (1994).
- [55] Kodama, H. and Sasaki, M. *Progress of Theoretical Physics Supplement*, **78**:1 (1984).
- [56] Sachs, R. K. and Wolfe, A. M. *ApJ*, **147**:73 (1967).
- [57] Baumann, D. and Peiris, H. V. *arXiv e-prints*, arXiv:0810.3022 (2008).
- [58] Silk, J. *ApJ*, **151**:459 (1968).
- [59] Seljak, U., Sugiyama, N., White, M., et al. *Phys. Rev. D*, **68**:083507 (2003).
- [60] Kamionkowski, M. and Kovetz, E. D. *Annual Review of Astronomy and Astrophysics*, **54**:227 (2016).
- [61] Crittenden, R., Davis, R. L., and Steinhardt, P. J. *ApJ*, **417**:L13 (1993).
- [62] Kosowsky, A. *Annals of Physics*, **246**(1):49 (1996).
- [63] Zaldarriaga, M. *Phys. Rev. D*, **55**:1822 (1997).
- [64] Seljak, U. and Zaldarriaga, M. ‘CMBFAST: A microwave anisotropy code.’ Astrophysics Source Code Library (1999).
- [65] Lewis, Antony, Challinor, Anthony and Lasenby, Anthony. *ApJ*, **538**:473 (2000).
- [66] Blas, D., Lesgourgues, J., and Tram, T. *JCAP*, **7**:034 (2011).
- [67] Lesgourgues, J. *arXiv e-prints*, arXiv:1104.2934 (2011).
- [68] Rees, M. J. *ApJL*, **153**:L1 (1968).
- [69] Zaldarriaga, M. In W. L. Freedman, editor, ‘Measuring and Modeling the Universe,’ page 309 (2004).
- [70] Kaiser, N. *MNRAS*, **202**:1169 (1983).
- [71] Hu, W. and White, M. *Phys. Rev. D*, **56**:596 (1997).
- [72] Kovac, J. M., Leitch, E. M., Pryke, C., et al. *Nature*, **420**:772 (2002).
- [73] Zaldarriaga, M. *Phys. Rev. D*, **64**(10):103001 (2001).
- [74] Zaldarriaga, M. *Phys. Rev. D*, **55**(4):1830 (1997).

- [75] Zaldarriaga, M. and Harari, D. D. *Phys. Rev. D*, **52**:3276 (1995).
- [76] Pritchard, J. R. and Kamionkowski, M. *Annals of Physics*, **318**:2 (2005).
- [77] Turner, M. S., White, M., and Lidsey, J. E. *Phys. Rev. D*, **48**:4613 (1993).
- [78] Wang, Y. *Phys. Rev. D*, **53**:639 (1996).
- [79] Zaldarriaga, M. and Seljak, U. *Phys. Rev. D*, **58**(2):023003 (1998).
- [80] Lewis, A. and Challinor, A. *Phys. Rept.*, **429**:1 (2006).
- [81] Kaiser, N. *ApJ*, **498**:26 (1998).
- [82] Bartelmann, M. and Schneider, P. *Phys. Rept.*, **340**:291 (2001).
- [83] Hu, W. *Phys. Rev. D*, **62**:043007 (2000).
- [84] Lewis, A. and Challinor, A. *Physics Reports*, **429**(1):1 (2006).
- [85] Knox, L. and Song, Y.-S. *Phys. Rev. Lett.*, **89**:011303 (2002).
- [86] Smith, K. M., Hanson, D., LoVerde, M., et al. *JCAP*, **2012**:014 (2012).
- [87] Manzotti, A. *Phys. Rev. D*, **97**:043527 (2018).
- [88] Sherwin, B. D. and Schmittfull, M. *Phys. Rev. D*, **92**:043005 (2015).
- [89] Song, Y.-S., Cooray, A., Knox, L., et al. *ApJ*, **590**:664 (2003).
- [90] Yu, B., Hill, J. C., and Sherwin, B. D. *Phys. Rev. D*, **96**:123511 (2017).
- [91] Hanson, D., Hoover, S., Crites, A., et al. *Phys. Rev. Lett.*, **111**(14):141301 (2013).
- [92] van Engelen, A., Sherwin, B. D., Sehgal, N., et al. *ApJ*, **808**:7 (2015).
- [93] Holder, G. P., Viero, M. P., Zahn, O., et al. *ApJ*, **771**:L16 (2013).
- [94] Planck Collaboration, Aghanim, N., Akrami, Y., et al. *arXiv e-prints*, arXiv:1807.06210 (2018).
- [95] Seljak, U. b. u. and Zaldarriaga, M. *Phys. Rev. Lett.*, **82**:2636 (1999).
- [96] Hu, W. *Phys. Rev. D*, **64**:083005 (2001).

- [97] Hu, W. and Okamoto, T. *ApJ*, **574**:566 (2002).
- [98] Kesden, M., Cooray, A., and Kamionkowski, M. *Phys. Rev. Lett.*, **89**:011304 (2002).
- [99] Carron, J., Lewis, A., and Challinor, A. *JCAP*, **2017**:035 (2017).
- [100] Bennett, C. L., Larson, D., Weiland, J. L., et al. *ApJS*, **208**:20 (2013).
- [101] Hinshaw, G., Weiland, J. L., Hill, R. S., et al. *ApJS*, **180**:225 (2009).
- [102] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. *A&A*, **571**:A18 (2014).
- [103] Das, S., Sherwin, B. D., Aguirre, P., et al. *Phys. Rev. Lett.*, **107**:021301 (2011).
- [104] Sherwin, B. D., van Engelen, A., Sehgal, N., et al. *Phys. Rev. D*, **95**:123529 (2017).
- [105] Sherwin, B. D., Dunkley, J., Das, S., et al. *Phys. Rev. Lett.*, **107**(2):021302 (2011).
- [106] Hand, N., Leauthaud, A., Das, S., et al. *Phys. Rev. D*, **91**:062001 (2015).
- [107] Madhavacheril, M., Sehgal, N., Allison, R., et al. *Phys. Rev. Lett.*, **114**:151302 (2015).
- [108] Hand, N., Addison, G. E., Aubourg, E., et al. *Phys. Rev. Lett.*, **109**:041101 (2012).
- [109] Carlstrom, J. E., Ade, P. A. R., Aird, K. A., et al. *Publications of the ASP*, **123**:568 (2011).
- [110] Ruhl, J., Ade, P. A. R., Carlstrom, J. E., et al. In C. M. Bradford, P. A. R. Ade, J. E. Aguirre, J. J. Bock, M. Dragovan, L. Duband, L. Earle, J. Glenn, H. Matsuhara, B. J. Naylor, H. T. Nguyen, M. Yun, and J. Zmuidzinas, editors, ‘Z-Spec: a broadband millimeter-wave grating spectrometer: design, construction, and first cryogenic measurements,’ volume 5498 of *Proceedings of the SPIE*, pages 11–29 (2004).
- [111] Staniszewski, Z., Ade, P. A. R., Aird, K. A., et al. *ApJ*, **701**:32 (2009).
- [112] High, F. W., Stalder, B., Song, J., et al. *ApJ*, **723**:1736 (2010).
- [113] Vanderlinde, K., Crawford, T. M., de Haan, T., et al. *ApJ*, **722**:1180 (2010).

- [114] Baxter, E. J., Keisler, R., Dodelson, S., et al. *ApJ*, **806**:247 (2015).
- [115] van Engelen, A., Keisler, R., Zahn, O., et al. *ApJ*, **756**:142 (2012).
- [116] Zhu, N., Orłowski-Scherer, J. L., Xu, Z., et al. In ‘Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX,’ volume 10708 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 1070829 (2018).
- [117] Matsumura, T., Akiba, Y., Borrill, J., et al. *Journal of Low Temperature Physics*, **176**:733 (2014).
- [118] Matsumura, T., Akiba, Y., Borrill, J., et al. *Journal of Low Temperature Physics*, **176**(5-6):733 (2014).
- [119] Matsumura, T., Akiba, Y., Arnold, K., et al. *Journal of Low Temperature Physics*, pages 1–8 (2016).
- [120] Bennett, C. L., Hill, R. S., Hinshaw, G., et al. *The Astrophysical Journal Supplement Series*, **148**:97 (2003).
- [121] Planck Collaboration, Akrami, Y., Ashdown, M., et al. *arXiv e-prints*, arXiv:1801.04945 (2018).
- [122] Krachmalnicoff, N., Carretti, E., Baccigalupi, C., et al. *A&A*, **618**:A166 (2018).
- [123] Choi, S. K. and Page, L. A. *JCAP*, **2015**:020 (2015).
- [124] Abazajian, K. N., Adshead, P., Ahmed, Z., et al. *ArXiv e-prints* (2016).
- [125] Errard, J. and Stompor, R. *Phys. Rev. D*, **99**:043529 (2019).
- [126] Ichiki, K. *Progress of Theoretical and Experimental Physics*, **2014**(6):1 (2014).
- [127] Draine, B. T. *Physics of the Interstellar and Intergalactic Medium*. Princeton University Press (2011).
- [128] Calzetti, D., Armus, L., Bohlin, R. C., et al. *ApJ*, **533**:682 (2000).
- [129] Hensley, B. S. and Bull, P. *ApJ*, **853**:127 (2018).
- [130] Hall, J. S. *Science*, **109**:166 (1949).
- [131] Hiltner, W. A. *Science*, **109**:165 (1949).

- [132] Davis, L. and Greenstein, J. L. *Physical Review*, **75**:1605 (1949).
- [133] Davis, L., Jr. and Greenstein, J. L. *ApJ*, **114**:206 (1951).
- [134] Andersson, B.-G., Lazarian, A., and Vaillancourt, J. E. *Annual Review of Astronomy and Astrophysics*, **53**(1):501 (2015).
- [135] Dolginov, A. Z. and Mytrophanov, I. G. *APSS*, **43**:257 (1976).
- [136] Draine, B. T. and Weingartner, J. C. *ApJ*, **470**:551 (1996).
- [137] Hoang, T. and Lazarian, A. *ApJ*, **831**:159 (2016).
- [138] Hildebrand, R. H. *Quarterly Journal of the Royal Astronomical Society*, **29**:327 (1988).
- [139] Dunkley, J., Amblard, A., Baccigalupi, C., et al. In S. Dodelson, D. Baumann, A. Cooray, J. Dunkley, A. Fraisse, M. G. Jackson, A. Kogut, L. Krauss, M. Zaldarriaga, and K. Smith, editors, ‘American Institute of Physics Conference Series,’ volume 1141 of *American Institute of Physics Conference Series*, pages 222–264 (2009).
- [140] Tassis, K. and Pavlidou, V. *MNRAS*, **451**:L90 (2015).
- [141] Poh, J. and Dodelson, S. *Phys. Rev. D*, **95**:103511 (2017).
- [142] Planck Collaboration, Aghanim, N., Ashdown, M., et al. *A&A*, **599**:A51 (2017).
- [143] Sheehy, C. and Slosar, A. *Phys. Rev. D*, **97**:043522 (2018).
- [144] Vansyngel, F., Boulanger, F., Ghosh, T., et al. *A&A*, **603**:A62 (2017).
- [145] Clark, S. E. *ApJ*, **857**:L10 (2018).
- [146] Chluba, J., Hill, J. C., and Abitbol, M. H. *MNRAS*, **472**:1195 (2017).
- [147] Ichiki, K., Kanai, H., Katayama, N., et al. *Progress of Theoretical and Experimental Physics*, **2019**(3):033E01 (2019).
- [148] Kogut, A., Dunkley, J., Bennett, C. L., et al. *ApJ*, **665**:355 (2007).
- [149] B. Rybicki, G. and P. Lightman, A. *Radiative Processes in Astrophysics*. Wiley (1979).

- [150] de Oliveira-Costa, A., Kogut, A., Devlin, M. J., et al. *ApJL*, **482**:L17 (1997).
- [151] Leitch, E. M., Readhead, A. C. S., Pearson, T. J., et al. *ApJ*, **486**(1):L23 (1997).
- [152] Kogut, A., Banday, A. J., Bennett, C. L., et al. *ApJ*, **460**:1 (1996).
- [153] Erickson, W. C. *ApJ*, **126**:480 (1957).
- [154] Draine, B. T. and Lazarian, A. *ApJ*, **508**:157 (1998).
- [155] Draine, B. T. and Lazarian, A. *ApJ*, **512**:740 (1999).
- [156] Génova-Santos, R., Rubiño-Martín, J. A., Rebolo, R., et al. *MNRAS*, **452**:4169 (2015).
- [157] Poidevin, F., Rubiño-Martín, J. A., Dickinson, C., et al. *MNRAS*, page 3277 (2018).
- [158] Pierpaoli, E., Macellari, N., Dickinson, C., et al. *MNRAS*, **418**(2):888 (2011).
- [159] Dickinson, C., Peel, M., and Vidal, M. *MNRAS*, **418**:L35 (2011).
- [160] Bayes, T. *Phil. Trans. of the Royal Soc. of London*, **53**:370 (1763).
- [161] MacKay, D. J. C. *Information Theory, Inference & Learning Algorithms*. Cambridge University Press, New York, NY, USA (2002).
- [162] Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., et al. *The Journal of Chemical Physics*, **21**(6):1087 (1953).
- [163] Hastings, W. K. *Biometrika*, **57**(1):97 (1970).
- [164] Peskun, P. H. *Biometrika*, **60**(3):607 (1973).
- [165] Peskun, P. *Journal of Computational Physics*, **40**(2):327 (1981).
- [166] Geman, S. and Geman, D. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **PAMI-6**(6):721 (1984).
- [167] Goodman, J. and Weare, J. *Communications in Applied Mathematics and Computational Science*, **5**:65 (2010).
- [168] Foreman-Mackey, D., Hogg, D. W., Lang, D., et al. *PASP*, **125**:306 (2013).
- [169] Eisenstein, D. J., Hu, W., and Tegmark, M. *ApJ*, **518**:2 (1999).

- [170] Perotto, L., Lesgourgues, J., Hannestad, S., et al. *JCAP*, **2006**:013 (2006).
- [171] Lesgourgues, J., Pastor, S., and Perotto, L. *Phys. Rev. D*, **70**:045016 (2004).
- [172] Wu, W. L. K., Errard, J., Dvorkin, C., et al. *ApJ*, **788**:138 (2014).
- [173] Errard, J., Feeney, S. M., Peiris, H. V., et al. *JCAP*, **2016**:052 (2016).
- [174] Heavens, A. F., Seikel, M., Nord, B. D., et al. *MNRAS*, **445**:1687 (2014).
- [175] Alonso, D. and Ferreira, P. G. *Phys. Rev. D*, **92**:063525 (2015).
- [176] Wolz, L., Kilbinger, M., Weller, J., et al. *JCAP*, **2012**:009 (2012).
- [177] Fisher, R. A. *Journal of the Royal Statistical Society*, **98**(1):39 (1935).
- [178] Kendall, M. G. M. G., Stuart, ., Alan, and Ord, ., J. K. (John Keith). *The advanced theory of statistics*. London : C. Griffin, 4th ed edition (1977). Includes index.
- [179] Tegmark, M., Taylor, A. N., and Heavens, A. F. *ApJ*, **480**:22 (1997).
- [180] Bond, J. R., Jaffe, A. H., and Knox, L. *Phys. Rev. D*, **57**:2117 (1998).
- [181] Wandelt, B. D. and Hansen, F. K. *Phys. Rev. D*, **67**:023001 (2003).
- [182] Hamilton, A. J. S., Tegmark, M., and Padmanabhan, N. *MNRAS*, **317**(1):L23 (2000).
- [183] Tegmark, M. *Phys. Rev. D*, **55**:5895 (1997).
- [184] Hivon, E., Górski, K. M., Netterfield, C. B., et al. *ApJ*, **567**:2 (2002).
- [185] Wandelt, B. D., Hivon, E., and Górski, K. M. *Phys. Rev. D*, **64**:083003 (2001).
- [186] Hansen, F. K., Górski, K. M., and Hivon, E. *MNRAS*, **336**(4):1304 (2002).
- [187] Chon, G., Challinor, A., Prunet, S., et al. *MNRAS*, **350**:914 (2004).
- [188] Kogut, A., Spergel, D. N., Barnes, C., et al. *The Astrophysical Journal Supplement Series*, **148**(1):161 (2003).
- [189] Alonso, D., Sanchez, J., and Slosar, A. *MNRAS*, **484**:4127 (2019).
- [190] Smith, K. M. *Phys. Rev. D*, **74**:083002 (2006).

- [191] Benoit-Lévy, A., Peiris, H. V., Leistedt, B., et al. *MNRAS*, **435**(3):1857 (2013).
- [192] Bunn, E. F., Zaldarriaga, M., Tegmark, M., et al. *Phys. Rev. D*, **67**:023501 (2003).
- [193] Grain, J., Tristram, M., and Stompor, R. *Phys. Rev. D*, **79**:123515 (2009).
- [194] Delabrouille, J., Betoule, M., Melin, J.-B., et al. *arXiv*, **astro-ph.C**:3675 (2012).
- [195] de Oliveira-Costa, A., Tegmark, M., Gaensler, B. M., et al. *MNRAS*, **388**(1):247 (2008).
- [196] Zheng, H., Tegmark, M., Dillon, J. S., et al. *MNRAS*, **464**:3486 (2017).
- [197] Hervías-Caimapo, C., Bonaldi, A., and Brown, M. L. *MNRAS*, **462**(2):2063 (2016).
- [198] Waelkens, A., Jaffe, T., Reinecke, M., et al. *A&A*, **495**:697 (2009).
- [199] Fauvet, L., Macías-Pérez, J. F., Aumont, J., et al. *A&A*, **526**:A145 (2011).
- [200] Fauvet, L., Macías-Pérez, J. F., and Désert, F. X. *Astroparticle Physics*, **36**:57 (2012).
- [201] Jansson, R. and Farrar, G. R. *ApJ*, **757**:14 (2012).
- [202] Orlando, E. and Strong, A. *MNRAS*, **436**:2127 (2013).
- [203] Beck, M. C., Beck, A. M., Beck, R., et al. *Journal of Cosmology and Astro-Particle Physics*, **2016**(5):056 (2016).
- [204] Næss, S. K. and Louis, T. *JCAP*, **9**:001 (2013).
- [205] Planck Collaboration, Adam, R., Ade, P. A. R., et al. *A&A*, **586**:A133 (2016).
- [206] Haslam, C. G. T., Klein, U., Salter, C. J., et al. *A&AS*, **100**:209 (1981).
- [207] Haslam, C. G. T., Stoffel, H., Salter, C. J., et al. *A&AS*, **47**:1 (1982).
- [208] Remazeilles, M., Dickinson, C., Banday, A. J., et al. *eprint 1411.3628*, **18**(June):1 (2014).
- [209] Miville-Deschenes, M. A., Ysard, N., Lavabre, A., et al. *A&A*, **490**(3):1093 (2008).

- [210] Kogut, A. *ApJ*, **753**:110 (2012).
- [211] Fuskeland, U., Wehus, I. K., Eriksen, H. K., et al. *ApJ*, **790**:104 (2014).
- [212] Ruud, T. M., Fuskeland, U., Wehus, I. K., et al. *ApJ*, **811**:89 (2015).
- [213] Hinshaw, G., Larson, D., Komatsu, E., et al. *ApJS*, **208**:19 (2013).
- [214] Planck Collaboration, Adam, R., Ade, P. A. R., et al. *A&A*, **594**:A10 (2016).
- [215] Finkbeiner, D. P., Davis, M., and Schlegel, D. J. *ApJ*, **524**(2):867 (1999).
- [216] Meisner, A. M. and Finkbeiner, D. P. *eprint 1410.7523* (2014).
- [217] Stevenson, M. A. *ApJ*, **781**(2):113 (2014).
- [218] Hoang, T., Lazarian, A., and Draine, B. T. *ApJ*, **741**(2):87 (2011).
- [219] Hensley, B. S., Draine, B. T., and Meisner, A. M. *ApJ*, **827**(1):45 (2016).
- [220] Ali-Haimoud, Y., Hirata, C. M., and Dickinson, C. *MNRAS*, **395**(2):1055 (2009).
- [221] Silsbee, K., Ali-Haimoud, Y., and Hirata, C. M. *MNRAS*, **411**:2750 (2011).
- [222] Génova-Santos, R., Rubiño-Martín, J. A., Peláez-Santos, A., et al. *eprint 1605.04741* (2016).
- [223] Remazeilles, M., Dickinson, C., Eriksen, H. K. K., et al. *MNRAS*, **458**:2032 (2016).
- [224] Fraisse, A. A., Brown, J.-A. C., Dobler, G., et al. In S. Dodelson, D. Baumann, A. Cooray, J. Dunkley, A. Fraisse, M. G. Jackson, A. Kogut, L. Krauss, M. Zaldarriaga, and K. Smith, editors, ‘American Institute of Physics Conference Series,’ volume 1141 of *American Institute of Physics Conference Series*, pages 265–310 (2009).
- [225] Macellari, N., Pierpaoli, E., Dickinson, C., et al. *MNRAS*, **418**:888 (2011).
- [226] Giardino, G., Banday, A. J., Fosalba, P., et al. *A&A*, **371**:708 (2001).
- [227] Gold, B., Odegard, N., Weiland, J. L., et al. *ApJS*, **192**:15 (2011).
- [228] Kayo, I., Taruya, A., and Suto, Y. *ApJ*, **561**:22 (2001).
- [229] Tegmark, M. *ApJ*, **502**:1 (1998).

- [230] Santos, M. G., Cooray, A., and Knox, L. *ApJ*, **625**:575 (2005).
- [231] Matsumura, T., Akiba, Y., Arnold, K., et al. *Journal of Low Temperature Physics*, **184**:824 (2016).
- [232] Matsumura, T. et al. (2013). [J. Low. Temp. Phys.176,733(2014)].
- [233] Alonso, D., Dunkley, J., Thorne, B., et al. *Phys. Rev. D*, **95**(4):043504 (2017).
- [234] Hinshaw, G., Nolta, M. R., and et al., B. *ApJS*, **170**:288 (2007).
- [235] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. *A&A*, **594**:A13 (2016).
- [236] Planck Collaboration, Aghanim, N., Ashdown, M., et al. *A&A*, **596**:A107 (2016).
- [237] Kamionkowski, M. and Kovetz, E. D. *Annual Review of Astronomy and Astrophysics*, **54**:227 (2016).
- [238] Guzzetti, C., M., Bartolo, N., Liguori, M., et al. *Riv. Nuovo Cim.*, **39**(9):399 (2016).
- [239] Hazumi, M. et al. *Proc. SPIE Int. Soc. Opt. Eng.*, **8442**:844219 (2012).
- [240] Namba, R., Peloso, M., Shiraishi, M., et al. *JCAP*, **1**:041 (2016).
- [241] Obata, I., Soda, J., and CLEO Collaboration. *Phys. Rev. D*, **93**(12):123502 (2016).
- [242] Ferreira, R. Z., Ganc, J., Noreña, J., et al. *JCAP*, **4**:039 (2016).
- [243] Caprini, C. and Sorbo, L. *JCAP*, **10**:056 (2014).
- [244] Mukohyama, S., Namba, R., Peloso, M., et al. *JCAP*, **8**:036 (2014).
- [245] Peloso, M., Sorbo, L., and Unal, C. *JCAP*, **9**:001 (2016).
- [246] Shiraishi, M., Hikage, C., Namba, R., et al. *Phys. Rev. D*, **94**(4):043506 (2016).
- [247] García-Bellido, J., Peloso, M., and Unal, C. *Journal of Cosmology and Astro-Particle Physics*, **2016**(12):031 (2016).
- [248] Lue, A., Wang, L.-M., and Kamionkowski, M. *Phys. Rev. Lett.*, **83**:1506 (1999).
- [249] Gerbino, M., Gruppuso, A., Natoli, P., et al. *JCAP*, **7**:044 (2016).

- [250] Planck Collaboration, Aghanim, N., Ashdown, M., et al. *A&A*, **596**:A110 (2016).
- [251] Gruppuso, A., Gerbino, M., Natoli, P., et al. *JCAP*, **6**:001 (2016).
- [252] Molinari, D., Gruppuso, A., and Natoli, P. *Physics of the Dark Universe*, **14**:65 (2016).
- [253] Seto, N. *Phys. Rev.*, **D75**:061302 (2007).
- [254] Seto, N. and Taruya, A. *Phys. Rev. Lett.*, **99**(12):121101 (2007).
- [255] Seto, N. and Taruya, A. *Phys. Rev. D*, **77**(10):103001 (2008).
- [256] Smith, T. L. and Caldwell, R. *Phys. Rev. D*, **95**:044036 (2017).
- [257] Matsumura, T. et al. *Journal of Low Temperature Physics* (2016).
- [258] Ferté, A. and Grain, J. *Phys. Rev. D*, **89**(10):103516 (2014).
- [259] Amaro-Seoane, P., Aoudia, S., Babak, S., et al. *GW Notes, Vol. 6, p. 4-110*, **6**:4 (2013).
- [260] Bartolo, N. et al. *JCAP*, **1612**(12):026 (2016).
- [261] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. *A&A*, **594**:A13 (2016).
- [262] Adshead, P., Martinec, E., and Wyman, M. *Phys. Rev. D*, **88**(2):021302 (2013).
- [263] Dimastrogiovanni, E. and Peloso, M. *Phys. Rev. D*, **87**(10):103501 (2013).
- [264] Moroi, T. and Takahashi, T. *Phys. Rev. D*, **66**(6):063501 (2002).
- [265] Lyth, D. H. and Wands, D. *Physics Letters B*, **524**:5 (2002).
- [266] Enqvist, K. and Sloth, M. S. *Nuclear Physics B*, **626**:395 (2002).
- [267] Pritchard, J. R. and Kamionkowski, M. *Annals Phys.*, **318**:2 (2005).
- [268] Blas, D., Lesgourgues, J., and Tram, T. *JCAP*, **7**:034 (2011).
- [269] Katayama, N. and Komatsu, E. *ApJ*, **737**:78 (2011).
- [270] Oyama, Y., Kohri, K., and Hazumi, M. *JCAP*, **2**:008 (2016).
- [271] O’Dea, D., Challinor, A., and Johnson, B. R. *MNRAS*, **376**:1767 (2007).

- [272] Shimon, M., Keating, B., Ponthieu, N., et al. *Phys. Rev. D*, **77**(8):083003 (2008).
- [273] Kaufman, J. P., Miller, N. J., Shimon, M., et al. *Phys. Rev. D*, **89**(6):062006 (2014).
- [274] Planck Collaboration, Adam, R., Ade, P. A. R., et al. *A&A*, **594**:A8 (2016).
- [275] Nati, F., Devlin, M. J., Gerbino, M., et al. *ArXiv e-prints* (2017).
- [276] Keating, B. G., Shimon, M., and Yadav, A. P. S. *ApJL*, **762**:L23 (2013).
- [277] Abitbol, M. H., Hill, J. C., and Johnson, B. R. *MNRAS*, **457**(2):1796 (2016).
- [278] Abbott, B. P., Abbott, R., Adhikari, R., et al. *Reports on Progress in Physics*, **72**(7):076901 (2009).
- [279] Accadia, T., Acernese, F., Alshourbagy, M., et al. *Journal of Instrumentation*, **7**(03):P03012 (2012).
- [280] Crowder, J. and Cornish, N. J. *Phys. Rev. D*, **72**(8):083005 (2005).
- [281] Allen, B. In J.-A. Marck and J.-P. Lasota, editors, ‘Relativistic Gravitation and Gravitational Radiation,’ page 373 (1997).
- [282] Agrawal, A., Fujita, T., and Komatsu, E. *Journal of Cosmology and Astro-Particle Physics*, **2018**(6):027 (2018).
- [283] Finn, L. S. *Phys. Rev. D*, **79**(2):022002 (2009).
- [284] Cornish, N. J. *Phys. Rev. D*, **65**(2):022004 (2001).
- [285] Romano, J. D. and Cornish, N. J. *Living Reviews in Relativity*, **20**(1):2 (2017).
- [286] Somiya, K. *Classical and Quantum Gravity*, **29**(12):124007 (2012).
- [287] Cornish, N. J. and Larson, S. L. *Classical and Quantum Gravity*, **18**:3473 (2001).
- [288] Barnaby, N., Moxon, J., Namba, R., et al. *Phys. Rev. D*, **86**(10):103508 (2012).
- [289] Boyle, L. A. and Steinhardt, P. J. *Phys. Rev. D*, **77**(6):063504 (2008).
- [290] Farmer, A. J. and Phinney, E. S. *MNRAS*, **346**:1197 (2003).
- [291] Sathyaprakash, B. S. and Schutz, B. F. *Living Reviews in Relativity*, **12**:2 (2009).

- [292] Klein, A., Barausse, E., Sesana, A., et al. *Phys. Rev. D*, **93**(2):024003 (2016).
- [293] Ferrari, V., Matarrese, S., and Schneider, R. *MNRAS*, **303**:247 (1999).
- [294] Ungarelli, C. and Vecchio, A. *Phys. Rev. D*, **63**:064030 (2001).
- [295] Seto, N. *Phys. Rev. D*, **75**(6):061302 (2007).
- [296] Meerburg, P. D., Hložek, R., Hadzhiyska, B., et al. *Phys. Rev. D*, **91**(10):103505 (2015).
- [297] Lasky, P. D., Mingarelli, C. M. F., Smith, T. L., et al. *Phys. Rev. X*, **6**(1):011035 (2016).
- [298] Swetz, D. S., Ade, P. A. R., Amiri, M., et al. *The Astrophysical Journal Supplement Series*, **194**(2):41 (2011).
- [299] Niemack, M. D., Ade, P. A. R., Aguirre, J., et al. In ‘Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V,’ volume 7741 of *Proceedings of the SPIE*, page 77411S (2010).
- [300] Thornton, R. J., Ade, P. A. R., Aiola, S., et al. *ApJS*, **227**(2):21 (2016).
- [301] Naess, S., Hasselfield, M., McMahon, J., et al. *JCAP*, **10**:007 (2014).
- [302] Planck Collaboration, Adam, R., Ade, P. A. R., et al. *A&A*, **594**:A1 (2016).
- [303] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. *A&A*, **594**:A12 (2016).
- [304] Hamimeche, S. and Lewis, A. *Phys. Rev. D*, **77**(10):103013 (2008).
- [305] Thornton, R. J., Ade, P. A. R., Aiola, S., et al. *ApJS*, **227**:21 (2016).
- [306] Benson, B. A., Ade, P. A. R., Ahmed, Z., et al. In ‘Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII,’ volume 9153 of *Proceedings of the SPIE*, page 91531P (2014).
- [307] Kang, J. H., Ade, P. A. R., Ahmed, Z., et al. In ‘Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX,’ volume 10708 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 107082N (2018).
- [308] Suzuki, A., Ade, P., Akiba, Y., et al. *Journal of Low Temperature Physics*, **184**:805 (2016).

- [309] Stevens, J. R., Goeckner-Wald, N., Keskitalo, R., et al. In ‘Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX,’ volume 10708 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 1070841 (2018).
- [310] Smith, K. M., Hu, W., and Kaplinghat, M. *Phys. Rev. D*, **74**(12):123002 (2006).
- [311] Lesgourgues, J. and Pastor, S. *Phys. Rept.*, **429**:307 (2006).
- [312] Remazeilles, M., Banday, A. J., Baccigalupi, C., et al. *JCAP*, **4**:023 (2018).
- [313] Eriksen, H. K., Jewell, J. B., Dickinson, C., et al. *ApJ*, **676**:10-32 (2008).
- [314] Challinor, A. and Chon, G. *MNRAS*, **360**:509 (2005).
- [315] Smith, K. M. and Zaldarriaga, M. *Phys. Rev. D*, **76**(4):043001 (2007).
- [316] Watanabe, Y. and Komatsu, E. *Phys. Rev. D*, **73**(12):123515 (2006).