

COMPRESSED LIKELIHOODS AND EARLY
UNIVERSE CONSTRAINTS FOR COSMIC
MICROWAVE BACKGROUND EXPERIMENTS

HEATHER PRINCE

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
ASTROPHYSICAL SCIENCES
ADVISER: PROFESSOR JO DUNKLEY

SEPTEMBER 2022

© Copyright by Heather Prince, 2022.

All rights reserved.

Abstract

In this dissertation we present cosmic microwave background likelihood tools for the *Planck* and BICEP3 data, and constrain extensions to the Big Bang Λ CDM cosmological model using *Planck* in combination with current and upcoming ground-based experiments. We begin with an overview of the current Standard Model of cosmology and the CMB. We present CosMOPED, a compressed likelihood code for *Planck* data at $\ell \geq 30$, which uses the Massively Optimized Parameter Estimation and Data compression technique (MOPED) method to reduce the dimensions of the data space to one number per cosmological parameter of interest. We then construct a binned likelihood for the *Planck* low- ℓ temperature and *E*-mode polarization, called Planck-low-py. We fit the bandpowers in two temperature bins and seven polarization bins with shifted log-normal distributions, and use these bins in a differentiable Python likelihood to facilitate ease of use of the *Planck* 2018 large-scale data. The Λ CDM parameters recovered with CosMOPED and Planck-low-py are consistent with the uncompressed *Planck* likelihoods, and a 7-parameter extended model is similarly well-constrained. We also examine some of the foreground modeling choices made in the BICEP/Keck primordial gravitational wave analysis (BK18) and estimate foreground-marginalized CMB B-mode bandpower amplitudes in the BICEP3 sky region. We use these bandpowers to construct a marginalized likelihood with no nuisance parameters. The tensor-to-scalar ratio inferred from this BK18-lite likelihood matches that using the public BK18 likelihood. Additionally, we use data from the Atacama Cosmology Telescope, the South Pole Telescope, and the *Planck* satellite to constrain different types of initial conditions. First we explore an adiabatic model with a broken power law instead of a single power law, and then we look at a model with both adiabatic and isocurvature fluctuations, allowing the isocurvature to vary independently in five bins. Finally, we forecast the upcoming Simons Observatory's ability to improve on our binned isocurvature constraints at small scales.

Acknowledgements

I would like to begin by thanking Jo Dunkley, my advisor, for her insight and encouragement over the past six years, and especially for her enthusiasm and support as I became a parent, and her patience through the challenges of a global pandemic and various mental health ups and downs. Thank you for hearing out my morning sickness woes and potty-training tales before we got down to discussing research, and for being excited about my plots! Thank you also to Fara, Harriet, and Kate for hosting Thanksgiving for the international crowd, and for bird bingo.

Thank you to all the members of the astrophysics department who have made Peyton Hall a welcoming place to do research, and to my fellow basement-dwellers for all the hikes, ski trips, sushi dinners and good company over the years. Thank you to my original grad cohort: Avery Bailey, Brianna Lacy, Christina Kreisch, Hayk Hakobyan, Luke Bouma, Mackenzie Moody, and Zack Li for the excellent dinners (even if we got food poisoning from those scallops) and fun tourist days, complete with recreating the CMB power spectrum on Liberty Island and finding space-themed clothing at a consignment store. Thank you to Amy Secunda and Roohi Dalal (as well as Brianna, Mackenzie and Zack) for being excellent office buddies and for packing Edward's lunches when Andy was away. Thank you to Goni Halevi for the coffee shop visits, co-working sessions, and truck parade company. Thank you to Aida Behmard, Alwin Mao, Ben Thorne, Cole Holcomb, Emilie Storer, Eric Moseley, Erin Flowers, Erin Kado-Fong, Felipe Ardilla, Kris Pardo, Lachlan Lancaster, Manu Schaan, Melinda Soares-Furtado, Miles Cranmer, Naomi Robertson, Oliver Philcox, Patryk Pjanka, Qiana Hunt, Rodrigo Córdova Rosado, Rui Xu, Sam Yee, Scott Carlsten, Shaquann Seadrow, Vassilis Tsiolis, Wenrui Xu and Zach Atkins for your support and camaraderie. Thank you to Adri Duivenvoorden, Alice Pisani, Andrina Nicola, Brandon Hensley, Jahmour Givans, Jia Liu, Mat Madhavacheril, Simone Aiola and Vera Gluscevic for your support during your time as post-docs in Princeton (and

thank you Alice for all the baby clothes and equipment). Thank you to the incredible staff in the department for making everything run smoothly, and especially to Polly for taking such wonderful care of the graduate students and for the spectacular birthday cakes!

Thank you to past and current members of my thesis committee, including Michael Strauss, Peter Melchior, Lyman Page and David Spergel for your excellent questions and suggestions at my committee meetings. Thank you to members of the ACT and SO collaborations for helpful information and questions. Thank you to my former advisor Kavilan Moodley for his ongoing support during my time at UKZN and beyond. Thank you to Erminia Calabrese for the good suggestions and feedback on drafts. Thank you to Renée Hlozek for passing the torch of Peyton's South African representative to me; as far as I know that position remains open for next year.

Thank you to my family: to Andy for everything, to Edward for the vicarious pure toddler joy and enforced work-life balance (and for my ever-growing knowledge of truck taxonomy), to my sisters Erica and Jo for their endless support and excellent memes, and to my father Colin for showing me cool science experiments when I was little and always encouraging me. Thank you to Shir for keeping me company during my first lonely Christmas in a new country and for making me go outside, and to Zeke for being the best (and also neediest) inadvertent emotional support dog.

Thank you to everyone at Trinity Episcopal Church in Cranford for being there for us when we had a baby during a global pandemic with our families halfway across the world. Thank you to Anthony Rafaniello and to Carol, John and Sophie Beier for the good meals and conversations, and to Arlene Fricke for the handmade beanies and sweaters and for many, many trucks and books. Thank you to Piper, Patrick, Maeve and Esme for the beach days that made the pandemic more bearable and for sharing snacks in church, and to Sam, Brandon, Phoebe and Iris for the playground adventures and backyard water castles.

To Andy and Edward

Relation to Published Work

Chapter 2 is based on a publication written in collaboration with Jo Dunkley, *Data compression in cosmology: A compressed likelihood for Planck data*. in Physical Review D, Volume 100 Issue 8 (2019), [arXiv:1909.05869](https://arxiv.org/abs/1909.05869).

Chapter 3 is based on a publication written in collaboration with Jo Dunkley, *Compressed Python likelihood for large scale temperature and polarization from Planck*. in Physical Review D, Volume 105 Issue 2 (2022), [arXiv:2104.05715](https://arxiv.org/abs/2104.05715).

Contents

Abstract	iii
Acknowledgements	iv
Relation to Published Work	vii
List of Tables	xii
List of Figures	xiii
1 Introduction	1
1.1 Cosmology	1
1.1.1 Λ CDM model	1
1.1.2 Perturbations to the smooth background	5
1.1.3 The early universe	13
1.1.4 Cosmic Microwave Background	15
1.2 CMB analysis	21
1.2.1 Effect of cosmological parameters on the CMB power spectrum	21
1.2.2 Foregrounds	24
1.2.3 Bayes theorem and data analysis	25
1.3 CMB experiments	28
1.4 Dissertation Overview	33
2 Data compression in cosmology: A MOPED compressed likelihood for <i>Planck</i> data	34

2.1	<i>Planck</i> likelihood and low- ℓ binning	36
2.1.1	Low- ℓ temperature bins	37
2.1.2	Low- ℓ polarization (τ prior)	39
2.1.3	Parameter constraints	39
2.2	MOPED compression vectors	41
2.3	Likelihood and parameters	45
2.3.1	Format of the likelihood	45
2.3.2	Parameter constraints	46
2.3.3	Effect of fiducial model	46
2.3.4	Non- Λ CDM cases	47
2.4	Discussion	48
3	A binned Python likelihood for large scale temperature and polarization from <i>Planck</i>	49
3.1	The <i>Planck</i> likelihood and low- ℓ binning	51
3.1.1	Low- ℓ temperature bins	53
3.1.2	Low- ℓ <i>E</i> -mode polarization	57
3.1.3	Software products	61
3.2	Parameter constraints	62
3.3	Discussion	63
4	Testing the BICEP/Keck foreground model, and a foreground-marginalized ‘BK18-lite’ likelihood	64
4.1	The BICEP/Keck data and likelihood	66
4.1.1	Data and Window Functions	66
4.1.2	Foreground model	69
4.1.3	Likelihood	70
4.2	Further exploration of the BK18 likelihood	73

4.2.1	Consistency of r from 95 GHz and 150 GHz	74
4.2.2	Consistency of r cleaned with 220 GHz or 353 GHz	75
4.2.3	Impact of assuming spatially power-law dust	76
4.2.4	Effect of fixing the dust temperature	77
4.2.5	Relaxing the prior on the synchrotron index	78
4.3	‘BK18-lite’ - a compressed likelihood	79
4.3.1	Bandpower estimation	82
4.3.2	The BK18-lite likelihood	85
4.4	Conclusion	88
5	Primordial adiabatic and isocurvature power spectrum constraints with the Atacama Cosmology Telescope DR4: beyond the single power law	89
5.1	Data and likelihoods	91
5.2	Primordial Adiabatic Power Spectrum	93
5.2.1	Power law primordial power spectrum	94
5.2.2	A broken power law spectrum	97
5.3	Primordial isocurvature perturbations	100
5.3.1	Isocurvature mode phenomenology	101
5.3.2	Parameterizing isocurvature perturbations	104
5.3.3	Constraints on CDM isocurvature	105
5.3.4	Constraints on neutrino isocurvature	108
5.4	Discussion	114
6	Forecasts for constraining primordial isocurvature with the Simons Observatory	115
6.1	Methodology	117
6.1.1	Simulated Simons Observatory data	117

6.1.2	Simulated Planck data	120
6.1.3	Isocurvature model	120
6.2	Results	121
6.3	Discussion	122
7	Conclusion	124
	Bibliography	128

List of Tables

3.1	Parameters for the log-normal EE bins, and estimated D^{EE} power. . .	61
5.1	The difference in χ^2 (or $-2 \ln L$) for the best-fitting adiabatic model compared to the the best-fitting model with non-zero isocurvature. The overall improvement in χ^2 is only 5.8 for 5 extra degrees of freedom for the NDI model, so the adiabatic model is still statistically preferred. The <i>Planck</i> data drives the improvement in χ^2	113

List of Figures

1.1 The present-day matter power spectrum $P(k)$. The power spectrum turns over at scales corresponding to the horizon size at matter-radiation equality. During the radiation dominated era modes that entered the horizon would decay as the universe expanded, which is why the linear power spectrum, shown as the solid black curve, falls off at large k (small scales). Modes that were on superhorizon scales until after matter-radiation equality are preserved, in this case as a power law at small k . The wiggles in the power spectrum just past the peak correspond to baryon acoustic oscillations. On small scales non-linear contributions begin to be significant as $\delta\rho/\bar{\rho}$ becomes larger, and a more complicated non-linear treatment is needed. The nonlinear power spectrum is shown as the dotted blue curve. These curves were computed using CLASS, the Cosmic Linear Anisotropy Solving System (Blas et al., 2011a), with the halofit approximation used to calculate the non-linear power spectrum (Smith et al., 2003; Bird et al., 2012)	12
---	----

1.2 The cosmic microwave background temperature (black) and E- and B-mode polarization (blue and green) angular power spectra for the <i>Planck</i> best-fit Λ CDM cosmological model (Planck Collaboration et al., 2020c). The solid curves depict the gravitationally lensed spectra and the dashed curves show the unlensed spectra. These spectra were computed using CLASS (Blas et al., 2011a).	19
1.3 The angular power spectrum of dominant foregrounds in temperature (left) and polarization (right). The left plot shows the CMB temperature power spectrum (black solid curve) and the temperature power spectrum at 143 GHz of residual undetected Poisson-distributed point sources that have not been masked in the <i>Planck</i> data (blue dashed curve). The right plot shows the B-mode polarization power spectrum for $r=0.01$ (black solid curve) and $r=0$ (black dot-dashed curve), together with the level of B-mode polarization from galactic dust in the BICEP region of the sky (blue dashed curve).	26
1.4 The CMB temperature power spectrum from the first year of data from The Wilkinson Microwave Anisotropy Probe (<i>WMAP</i>) satellite from Spergel et al. (2003). The points with errorbars are the binned power spectrum. The grey dots show the unbinned points. The solid curve is the best fit theory spectrum.	29

1.5 CMB power spectrum constraints from Planck’s 2018 legacy data release (orange) and BICEP2/Keck (green), together with forecasted errors for the upcoming Simons Observatory ground-based experiment (SO, blue) and the LiteBIRD satellite (red). Simons Observatory will measure temperature and polarization fluctuations for a wide range of angular scales, including small scales (large ℓ). LiteBIRD will measure the large scale (low- ℓ) polarization, including the reionization bump at $\ell < 10$ which cannot be measured from the ground. Figure source: Lee et al. (2019)	32
2.1 The <i>Planck</i> 2015 temperature power spectrum from (Planck Collaboration et al., 2016a). The non-Gaussian $\ell < 30$ bandpowers are shown with their asymmetrical errors in gray. The two black points at $\ell < 30$ are obtained by estimating the binned spectrum in two angular bins, and have approximately Gaussian distributions. We use these two low- ℓ bins for the likelihoods described in this chapter. The $\ell \geq 30$ binned data are the foreground-marginalized temperature bandpowers from the Plik_lite likelihood for <i>Planck</i> 2015. The theoretical power spectrum for the <i>Planck</i> 2015 TT+lowTEB best fit parameters (Planck Collaboration et al., 2016b) computed using CLASS (Blas et al., 2011b) is shown in blue.	36
2.2 The probability distribution for D_ℓ for two low- ℓ temperature bins using the <i>Planck</i> 2015 low- ℓ Commander likelihood. The best-fit Gaussian distribution is shown in blue.	38

2.3	Posteriors on the Λ CDM parameters from Planck-lite-py, a Python implementation of the <i>Planck</i> Plik_lite likelihood with two Gaussian low- ℓ bins (black solid curve) compared to the combined Plik_lite TT + low- ℓ Commander <i>Planck</i> 2015 temperature likelihoods (blue-dashed curve). The parameter constraints agree to within 0.1σ , showing that our two binned low- ℓ data points capture enough information to give equivalent constraints to the full low- ℓ temperature likelihood. A Gaussian prior of $\tau = 0.067 \pm 0.023$ is included in the black and blue curves. We also show results from the public chains from sampling the TT+lowTEB <i>Planck</i> 2015 likelihood (gray solid curve) for reference. The small difference is the effect of imposing a prior on the optical depth.	40
2.4	Compression vectors for the Λ CDM parameters. The sharp features are caused by binning of the power spectrum. These compression vectors can be applied to the CMB temperature power spectrum to give six numbers that contain as much information about the cosmological parameters as the binned temperature power spectrum. The compression vectors depend on the order in which they are computed, as they are intentionally orthogonal to one another.	43
2.5	The diagonal of the binned temperature inverse covariance matrix for the 217 temperature bins (2 for low- ℓ and 215 for high- ℓ). The diagonal elements are small for $\ell < 1000$, then rise steeply for $1500 < \ell < 2000$ before dropping again when the experimental noise approaches the signal.	44

2.6 Posteriors on the Λ CDM parameters from the compressed CosMOPED likelihood (red-dashed curves) and the Planck-lite-py Python implementation of the <code>Plik_lite</code> likelihood with two low- ℓ bins (black solid curves). The parameter constraints agree to within 0.1σ , validating the CosMOPED compression scheme. A Gaussian prior of $\tau = 0.067 \pm 0.023$ is included for both curves.	45
2.7 The constraints on running of the scalar spectral index from the compressed CosMOPED likelihood (red dashed) agree to within 0.1σ with the uncompressed Planck-lite-py (black). A Gaussian prior of $\tau = 0.067 \pm 0.023$ is used for CosMOPED and Planck-lite-py. The distribution using the public <i>Planck</i> chains (TT+lowTEB, gray) is slightly broader due to the more accurate treatment of low- ℓ polarization.	47
3.1 The <i>Planck</i> 2018 temperature and polarization power spectra and best-fit theory curves, from Planck Collaboration et al. (2020b). The log-normal binned $\ell < 30$ <i>TT</i> and <i>EE</i> spectra estimated in this paper, described in section 3.1, are shown on the left of the upper panel. They compress the per- ℓ likelihoods (grey markers) described by the <i>Planck Commander</i> and <i>SimA11</i> likelihoods respectively. The Gaussian, foreground-marginalized $\ell \geq 30$ data (black markers on the right) are used in the <code>Plik_lite</code> cosmological likelihood and are shown for reference. The low- ℓ <i>TE</i> cross spectrum is not used in the <i>Planck</i> cosmological analysis.	52

3.2 The low- ℓ <i>Planck</i> temperature power spectrum $D_\ell^{TT} = \ell(\ell+1)C_\ell^{TT}/2\pi$. The binned power and errors are shown in black. The <i>Planck</i> unbinned power spectrum and errors are in grey. The errorbars are asymmetric because on these scales the power spectrum is non-Gaussian. The value of the power spectrum is taken from the peak of the probability distribution and the errorbars come from the half-maximum values. The theory curve for the best fit Λ CDM model is shown in red. The power spectrum is close to flat on these large angular scales due to the Sachs-Wolfe effect.	54
3.3 The peak normalized posterior distribution for $D_\ell^{TT} = \ell(\ell+1)C_\ell^{TT}/2\pi$ for two low- ℓ temperature bins ($\ell = 2\text{-}15$ and $16\text{-}29$) using the <i>Planck</i> 2018 Commander low- ℓ temperature likelihood. The best-fitting log-normal distributions are shown in red.	54
3.4 The low- ℓ E-mode polarization power spectrum, with the D_ℓ^{EE} estimated in seven bins shown in black ($(\ell = 2\text{-}3, 4, 5, 6\text{-}7, 8\text{-}11$ and $10\text{-}29)$. The <i>Planck</i> per- ℓ power spectrum is shown in grey. The values and errors for these non-Gaussian data points come from the peak and half-maximum of the conditional posterior distribution for each D_ℓ . The theory curve for the best fit model (red) and for high and low values of the optical depth to reionization τ are shown. The reionization bump in the power spectrum comes from polarization caused by Thompson scattering off free electrons since reionization. This signal is higher for greater τ	57

3.5	The peak normalized posterior distribution for $D_\ell^{EE} = \ell(\ell+1)C_\ell^{EE}/2\pi$ for seven low- ℓ E-mode polarization bins, estimated using the <i>Planck</i> 2018 <code>SimAll</code> low- ℓ polarization likelihood. The best-fit offset log-normal distribution is shown in red for each bin. The ℓ ranges of the bins are chosen to capture the low- ℓ polarization features imprinted from varying the optical depth to reionization, τ	58
3.6	Posterior distributions of the Λ CDM parameters estimated using the low- ℓ <i>Planck</i> 2018 temperature (<code>Commander</code>) and EE (<code>SimAll</code>) likelihoods (black), compared to those using <code>Planck-low-py</code> (our log-normal compressed low-ell likelihood) (red dashed). In both cases <code>Plik_lite</code> is used at $\ell > 30$. The parameters shown are the optical depth to reionization τ , the amplitude of the primordial power spectrum A_s , the scalar spectral index n_s , the Hubble constant H_0 in units of $\text{km s}^{-1} \text{ Mpc}^{-1}$, the baryon density $\Omega_b h^2$ and the dark matter density $\Omega_c h^2$. The parameter constraints agree to within less than 0.1σ	60
3.7	Posterior distributions for extensions to Λ CDM. Left: constraint on the running of the spectral index. Right: constraint on the amplitude of the neutrino density isocurvature power spectrum at $k = 0.1 \text{ Mpc}^{-1}$ assuming that the isocurvature power spectrum is a power law and is uncorrelated with the adiabatic power spectrum (using the parameterization from Planck Collaboration et al. (2020e) with $P_{\mathcal{R}\mathcal{I}} = 0$). The black curves show the constraints using the full <i>Planck</i> 2018 low- ℓ likelihoods and the red dashed curves show the constraints using the binned likelihoods described in this paper.	60

4.1 BK18 bandpower window functions for the 9 bins. They differ for spectra estimated in the larger BICEP3 region (black) versus the smaller BICEP2/Keck region (red). Cross-spectra between the two regions have a third typical window function shape (orange). The *Planck* and *WMAP* maps are cut to the BICEP3 region of the sky and have the same window functions as BICEP3. The BICEP2 and Keck Array data for 95, 150 and 220 GHz have almost identical window function shapes, but different normalizations. The suite of cross-spectra window functions are also similar in shape but not identical, and have different normalizations.

4.2 Auto-spectra for the 95 and 220 GHz Keck array data (grey dots and purple squares respectively), 150 GHz BICEP2/Keck data (blue triangles), and 95 GHz BICEP3 data (black circles). Also shown is the B-mode power spectrum for a Λ CDM cosmology with $r=0$ (solid black curve, just lensing B-modes) and $r=0.03$ (dashed black curve). The 220 GHz channel is above the theory curves because it is dominated by dust, and the 150 GHz and 95 GHz channels have significant dust contamination on larger scales ($\ell \lesssim 200$). Dust can cause an excess of power on degree scales, which can be confused with primordial B-modes unless the foregrounds are constrained using a range of frequency channels. The BICEP3 95 GHz spectrum provides the most constraining power for the CMB B-mode power spectrum because it has both high sensitivity (small error bars) and relatively low dust contamination. . .

4.3	Reproduction of the BK18 r and foreground parameters (dust and synchrotron amplitude, spatial spectral index and frequency spectral index, and dust-synchrotron correlation parameter) for fixed cosmological parameters. Only the dust frequency index, $\beta_{B,dust}$, is now correlated with the tensor-to-scalar ratio, r . The blue curve in the top left panel shows the constraint on r when holding the foreground parameters fixed at their best-fit values.	74
4.4	Distribution for r estimated from 95 GHz or 150 GHz, compared to the nominal combination. The foreground frequencies are included in all cases and the standard seven foreground parameters are estimated. The current limit is dominated by the 95 GHz BICEP3 data.	75
4.5	Parameter estimates are consistent whether 220+217 GHz or 353 GHz data are used to clean the dust. The best-fit value of r is non-zero in the no-220 GHz case, but $r = 0$ is still an acceptable fit. Scatter is expected since the dust-tracing data are independent in the two cases.	76
4.6	Top left: there is minimal impact on r when we expand the spatial power-law dust model to one with an ell-by-ell amplitude, and negligible impact when we relax the prior on the power-law exponent. Top right: the power-law estimated with an expanded prior can still be constrained by the data. Bottom: estimated dust amplitude in bins, defined at 353 GHz, together with the best-fitting power-law-model	78
4.7	Varying the dust temperature, T_d , does not impact the constraint on r . The temperature cannot be bounded from above by this dataset, but is weakly anti-correlated with the dust index, β_{dust}	78

4.8 Left: Removing the Gaussian prior on the synchrotron index has little effect on r . Center/right: no synchrotron emission is detected in the BK18 data, but with a uniform prior on the index, the larger volume of models at more negative values of β are preferred.	79
4.9 Constraints on the tensor-to-scalar ratio r and the seven foreground parameters from the full BICEP likelihood, using data from both the BICEP3 and BICEP2/Keck sky regions (black) and just the BICEP3 sky region maps (red). Dropping the BICEP2/Keck 150 and 220 GHz channels broadens the foreground constraints, particularly for the dust spectral index, but the r constraint is not significantly affected. . . .	81
4.10 Estimates of the nine B-mode bandpowers (here just for the BICEP3 region) marginalized over foregrounds (red) or estimated for the best-fit foregrounds (blue).	83
4.11 Distributions of the foreground parameters are consistent whether estimating the bandpowers or the tensor-to-scalar ratio r directly. . . .	84
4.12 Probability distributions for the nine CMB bandpowers in the BICEP3 sky region, each fit with a three-parameter offset log-normal distribution. 85	

4.13 Foreground-marginalized CMB bandpowers (blue circles) plotted with the theoretical B-mode power spectrum for $r = 0$ (solid black curve) and $r = 0.03$ (dashed black curve). The BICEP3 95 GHz bandpowers are shown for comparison (light blue dots). The values of the foreground-marginalized bandpowers are taken from the mode of the best-fit log-normal distribution, and the errorbars for the foreground-marginalized bandpowers are calculated using the 16th and 84th percentiles of the distribution. These errors are asymmetric, with the upper errors larger than the lower error, although this is difficult to see with the log scale on the y-axis. The errorbars plotted for the BICEP3 95GHz channel, taken from the publicly available BICEP/Keck data, are symmetric and reflect the width and not the shape of the bandpower distribution.	86
4.14 The constraints on r constraints from our new foreground-marginalized ‘BK18-lite’ likelihood (black) agree with the nominal BK18 likelihood, restricted to just the maps that use the BICEP3 window functions (red). Neglecting the covariance between neighboring bins in the BK18-lite likelihood gives the blue curve.	87
5.1 1σ and 2σ probability distribution contours for the spectral index, n_s , and the baryon density, $\Omega_b h^2$, using ACT DR4 (light blue), <i>Planck</i> (purple) and <i>WMAP</i> (yellow) data reproduced from Figure 14 of Aiola et al. (2020). <i>Planck</i> and <i>WMAP</i> prefer lower values of n_s than ACT, which probes smaller angular scales.	95

5.2 Parameter constraints using ACT DR4 (blue) and ACT DR4 combined with <i>Planck</i> 2018 (black). ACT prefers a higher value of n_s than <i>Planck</i> , as seen in the middle panel of the top row, but there is no significant tension between the two datasets. The agreement of our parameter estimates with distributions from the official ACT DR4 data release Aiola et al. (2020) are shown for comparison (dashed purple and grey curves).	96
5.3 Parameter constraints for a broken power-law primordial power spectrum that is scale invariant on small scales, with a slope equal to n_s at scales $k < k_{\text{break}}$ (given in units of Mpc^{-1}). <i>Planck</i> prefers a single power law and therefore pushes the transition scale, k_{break} , to higher values (smaller scales) that are not well-constrained by <i>Planck</i> . Adding ACT or ACT+SPT data selects a preferred transition scale at $k \sim 0.1 \text{Mpc}^{-1}$, although a range of larger k -values are still allowed.	98
5.4 Estimated primordial power spectra for a broken power law (blue) and a single power law (green), with k in units of Mpc^{-1} . The maximum likelihood curve for ACT+ <i>Planck</i> is shown as the solid curve, together with a sample of 500 other points in the MCMC chains whose density is proportional to the model's posterior probability.	98
5.5 Parameter constraints for a broken power law primordial power spectrum where both spectral indices are allowed to vary, with a break at k_{break} , for <i>Planck</i> +ACT, and <i>Planck</i> +ACT+SPT. A broken power law that is scale invariant on small scales ($n_s = 1$) has the highest probability, although models where the two spectral indices are equal, i.e., the primordial power spectrum has a single power law, are allowed by the data.	99

5.6 The TT, TE and EE angular power spectra of the CMB for adiabatic, neutrino density isocurvature (NDI) and cold dark matter density isocurvature (CDI) initial conditions. In each case the primordial power spectrum $P_{\mathcal{R}\mathcal{R}}(k)$ is scale invariant ($n_s = 1$) with scalar amplitude $A_s = 2.3 \times 10^{-9}$. The isocurvature C_ℓ 's have been multiplied by a factor of 10 relative to the adiabatic C_ℓ 's so that their shapes are easier to compare.	102
5.7 The 95% upper limits on the cold dark matter density isocurvature (CDI) primordial power spectrum as a function of wavenumber, k , with five bins smoothed by convolving the binned power spectrum with a Gaussian. The best-fitting adiabatic primordial power spectrum is shown for comparison. The ACT DR4 data improve the constraints from <i>Planck</i> at the smallest scales.	105
5.8 CDI constraints: posterior distributions for the amplitudes of the power-law adiabatic power spectrum at two reference scales P_{RR}^1 and P_{RR}^2 , Hubble constant H_0 , baryon density parameter $\Omega_b h^2$, cold dark matter density parameter $\Omega_c h^2$ and the optical depth to reionization τ , as well as the binned cold dark matter density isocurvature amplitudes in five bins, $P_{\mathcal{I}\mathcal{I}}^1 \dots P_{\mathcal{I}\mathcal{I}}^5$	107

5.9 The 95% upper limits on the contribution of isocurvature to the linear matter power spectrum at $z = 0$. The solid curve is the total adiabatic and isocurvature power spectrum, the red dotted curve shows the upper limits on the cold dark matter density isocurvature contribution, which is below 1% in the first few bins but jumps up to 25% at $k \sim 0.2 \text{ h/Mpc}$ and the blue dashed curve shows the upper limits on the neutrino density isocurvature contribution, which remains below 1% at all scales. This plot shows the linear matter power spectrum; the contributions to the total (including nonlinear) matter power spectrum are of the same order of magnitude but have a more complex as a function of k	109
5.10 The 95% upper limits on the neutrino density isocurvature (NDI) primordial power spectrum as a function of wavenumber, k (in Mpc^{-1}), for <i>Planck</i> compared to ACT+ <i>Planck</i> . The best-fitting adiabatic power is shown for comparison.	110
5.11 Neutrino density isocurvature (NDI) constraints: posterior distributions for the amplitudes of the power-law adiabatic power spectrum at two reference scales $P_{\mathcal{R}\mathcal{R}}^1$ and $P_{\mathcal{R}\mathcal{R}}^2$, Hubble constant H_0 , baryon density parameter $\Omega_b h^2$, cold dark matter density parameter $\Omega_c h^2$ and the optical depth to reionization τ , as well as the binned neutrino density isocurvature amplitudes in five bins, $P_{\mathcal{I}\mathcal{I}}^1 \dots P_{\mathcal{I}\mathcal{I}}^5$	111
5.12 The contribution of cold dark matter and neutrino density isocurvature (scaled up by a factor of 100) relative to the adiabatic temperature and polarization power spectra, for the model with 95% upper limits on the amount of isocurvature.	112

5.13 Top: CMB temperature power spectrum for the best-fitting adiabatic model, together with the ACT DR4 data points and compared to models with the 95% upper limit of NDI and CDI perturbations. The models with isocurvature have a modest excess in the $\ell \sim 2000$ region. Bottom: difference between these models, and residuals of the ACT DR4 data with each model.	113
6.1 Noise levels in temperature (top) and E-mode polarization (bottom) for the <i>Planck</i> and Simons Observatory experiments. The blue curve shows the noise curve for the Simons Observatory Large Aperture Telescope (LAT), while the orange curve shows the noise for the Small Aperture Telescopes (SATs). The <i>Planck</i> white noise curve is shown for comparison in grey. The theory curve for the CMB power spectrum is shown in black. The LAT has strong signal-to-noise on smaller angular scales (larger multipoles ℓ) than <i>Planck</i> , while the SAT has lower noise on similar angular scales. These noise levels can be combined with the cosmic variance term to obtain the uncertainty on the CMB power spectrum.	118
6.2 Forecast upper limits on the isocurvature power spectrum from <i>Planck</i> (grey) compared to <i>Planck</i> + SO LAT (red) for CDM isocurvature (left) and neutrino density isocurvature (right), for a set of five wavenumber bins.	121

6.3 Forecast posterior distributions for the Λ CDM cosmological parameters, showing how SO LAT data (red) are expected to improve constraints compared *Planck* data alone (black) in the case that the perturbations are not assumed to be perfectly adiabatic. The left panel shows the constraints when binned cold dark matter density isocurvature is allowed to vary, and the right panel show the constraints with neutrino density isocurvature. The left and right panels are very similar to one another because small amounts of isocurvature are preferred.

..... 122

Chapter 1

Introduction

At least I know I'm bewildered about the really fundamental and important facts of the universe.

Terry Pratchett *Equal Rites*

Cosmology is “the science or theory of the universe as an ordered whole” ([Simpson & Weiner, 1989](#)). In this chapter I will describe the current widely accepted model of the universe, the Λ CDM model, discuss some of the ways that we can study cosmology, including with the cosmic microwave background (CMB), and give an update on the state of CMB observations.

1.1 Cosmology

1.1.1 Λ CDM model

Copernicus proposed in the sixteenth century that Earth orbits the Sun, suggesting that the universe does not revolve around humans and our planet. The Copernican principle as applied to cosmology states that there are no special observers: the universe as observed from Earth is a fair sample of the universe as a whole, the same scientific laws apply everywhere, and the conclusions reached by studying the uni-

verse from Earth would be consistent if repeated at any location in the universe. In other words the universe is isotropic and homogeneous on large scales. The observed near-isotropy of the cosmic microwave background (Smoot et al., 1992) and the homogeneity of galaxy surveys on large scales (Scaramella et al., 1991; Davis, 1997; Martinez et al., 1998) provide observational support for the cosmological principle.

Another key foundation to our understanding of the universe is Einstein's general theory of relativity, which relates the energy and momentum of matter and radiation to the geometry of spacetime (Einstein, 1915). This relationship is described by the Einstein field equations:

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}. \quad (1.1)$$

The geometry of spacetime is encapsulated by the Einstein tensor $G_{\mu\nu}$, which depends on the spacetime metric $g_{\mu\nu}$. Λ is a cosmological constant which was introduced by Einstein to produce a static universe (Einstein, 1917), but is now understood as a vacuum energy that causes the accelerated expansion of the universe. $T_{\mu\nu}$ is the stress-energy tensor, which is taken to be that of a perfect isotropic fluid with energy density ρ and pressure p .

Measurements in general relativity use the metric tensor $g_{\mu\nu}$ to transform from time and space coordinates to the spacetime interval $ds^2 = g_{\mu\nu} dx^\mu dx^\nu$, where $dx^0 = cdt$ is the time coordinate and $dx^i, i \in \{1, 2, 3\}$ are the three spatial coordinates. The Friedmann-Lemaître-Robertson-Walker (FLRW) metric for a homogeneous and isotropic universe is given by (Friedmann, 1922; Lemaître, 1931; Robertson, 1935; Walker, 1937)

$$ds^2 = -c^2 dt^2 + a^2(t) [d\chi^2 + f_K^2(\chi)(d\theta^2 + \sin^2 \theta d\phi^2)], \quad (1.2)$$

where χ are spatial comoving coordinates (which are fixed for a comoving observer who sees the universe as isotropic and homogeneous at any given time t), $a(t)$ is the

scale factor relating the comoving separation between two observers to the physical distance between them, c is the speed of light, and $f_K(\chi)$ depends on the curvature of the universe,

$$f_K(\chi) = \begin{cases} K^{-\frac{1}{2}} \sin(K^{\frac{1}{2}}\chi) & K > 0 \text{ (closed universe)} \\ \chi & K = 0 \text{ (flat universe)} \\ (-K)^{-\frac{1}{2}} \sinh((-K)^{\frac{1}{2}}\chi) & K < 0 \text{ (open universe)} \end{cases} \quad (1.3)$$

Applying Einstein's field equations to the FLRW metric allows us to study the evolution of the scale factor $a(t)$ and therefore the expansion history of the universe. Each component x of the universe is described by an equation of state relating the pressure of that component to its energy density: $p = p(\rho) = w\rho c^2$. Radiation r and relativistic matter have $w = \frac{1}{3}$, non-relativistic matter m has $w = 0$, and the cosmological constant Λ has $w = -1$. The energy density of each component can be converted to a fraction of the critical density ρ_c required for the universe to be spatially flat: $\Omega_x(t) = \frac{\rho_x(t)}{\rho_c}$. The Hubble parameter $H(t) \equiv \frac{\dot{a}}{a}$, then evolves as

$$\frac{H^2(t)}{H_0^2} = \Omega_r(t) + \Omega_m(t) + \Omega_K(t) + \Omega_\Lambda, \quad (1.4)$$

where the energy density of each component evolves according to its equation of state and the Einstein field equations as: $\Omega_r(t) = \frac{\Omega_{r,0}}{a^4(t)}$ for relativistic matter, $\Omega_m(t) = \frac{\Omega_{m,0}}{a^3(t)}$ for dust, $\Omega_K = \frac{\Omega_{K,0}}{a^2(t)}$ for curvature and the energy density associated with the cosmological constant Ω_Λ is, as the name suggests, constant. The radiation energy density decays the fastest because the wavelength of light is stretched with the expansion of the universe in addition to the number density reducing as the volume of the universe increases with expansion as $n \propto a^{-3}$. Because the energy density of the different components evolves differently with time, the universe was radiation dominated at early times, then matter dominated, and now dark energy is the dominant component.

The curvature of the universe can be constrained by studying the observed angular scale of baryon acoustic oscillations, which imprint a physical scale on the cosmic microwave background power spectrum and galaxy clustering. CMB data and baryon acoustic oscillation measurements both prefer a $K=0$ spatially flat universe (Spergel et al., 2003; Planck Collaboration et al., 2020c; Alam et al., 2021).

Observations of distant galaxies show that they are receding from us, with more distant galaxies retreating faster, i.e., the universe is expanding. This was predicted by Lemaître (1927) and observed by Hubble (1929) using Cepheid variable stars to calculate distances. In a matter or radiation dominated universe this expansion would be decelerating with time. However in a universe dominated by the cosmological constant the expansion is exponentially increasing, with $a(t) \propto e^{Ht}$. Type Ia supernovae can be used as ‘standard candles’ to compare their observed redshifts to the distances inferred by their standardized luminosities. Riess et al. (1998) and Perlmutter et al. (1999) showed that the expansion of the universe is accelerating, in a way that is consistent with a cosmological constant dominating the current expansion of the universe. This cosmological constant is interpreted as being dark energy, i.e. an energy associated with the vacuum whose density remains constant even as the universe expands.

Zwicky (1933) studied velocity dispersion of galaxies in the Coma cluster and found that there was more mass present than could be accounted for by normal ‘baryonic’ matter such as stars, gas and dust. The presence of this dark matter component explains the observed shapes of galaxies (Ostriker & Peebles, 1973), rotation curves of galaxies (Rubin & Ford, 1970), gravitational lensing measurements (Clowe et al., 2006), and constraints on the total and baryonic matter density from the CMB (Spergel et al., 2003). This mysterious component of the universe interacts gravitationally, but does not seem to emit or absorb radiation, leading to the term ‘dark matter’. Most of the mass in the universe is therefore in the form of dark matter.

The components of the Λ CDM model of the universe include dark energy in the form of a cosmological constant (Λ), cold dark matter (CDM) which is ‘cold’ because it is non-relativistic, and baryons. Radiation and three light neutrino species dominated the total energy density early on but the radiation component is very small now as it falls off quickly with expansion. The expansion rate of the universe in the Λ CDM model, shown in Equation 1.4, is determined by general relativity in a homogeneous and isotropic universe.

1.1.2 Perturbations to the smooth background

Although the universe is homogeneous on very large scales, galaxy redshift surveys show that it is irregular and clumpy on the scales of galaxy clusters and below (Davis et al., 1982). The structures that we see today developed from very small density fluctuations at early times, which grew through gravitational collapse (Peebles, 1982; Peebles & Yu, 1970). Overdense regions of dark matter gravitationally attracted more dark matter while the universe was so young and hot that baryons were still coupled to photons. After recombination the baryons were free to fall into the dark matter overdensities, and eventually they formed stars and galaxies. With time this grew into the cosmic web of dark matter halos and filaments populated by galaxies that we see today.

Suppose we have some perturbations around a smooth background (the next section will explore theories about how these perturbations came about). The metric tensor and the stress-energy tensor can be broken down into an unperturbed background value (denoted by a bar) and a small perturbation (denoted by δ):

$$g_{\mu\nu} = \bar{g}_{\mu\nu} + \delta g_{\mu\nu} \quad (1.5)$$

$$T_{\nu}^{\mu} = \bar{T}_{\nu}^{\mu} + \delta T_{\nu}^{\mu}. \quad (1.6)$$

The metric perturbations propagate into perturbations to the Einstein tensor

$$G_\nu^\mu = \bar{G}_\nu^\mu + \delta G_\nu^\mu. \quad (1.7)$$

At early times the perturbations and their derivatives are small so we can neglect any nonlinear terms. To study the evolution of these perturbations, we use a combination of the Einstein, Boltzmann and fluid equations. The Einstein field equations capture how the stress-energy perturbations and the geometric spacetime perturbations affect one another, and the fluid and Boltzmann equations describe the evolution of the non-relativistic (baryons and cold dark matter) and relativistic (photons and neutrinos) components respectively in phase space.

The perturbations can be divided into scalar, vector and tensor modes, based on how they transform under rotations (Ma & Bertschinger, 1995). These scalar, vector and tensor modes evolve independently in linear perturbation theory, so we can examine them separately. Scalar perturbations are important for structure formation, and tensor perturbations come from gravitational waves and are relevant for constraining early universe models. Vector perturbations decay with time so we ignore them.

The stress-energy tensor perturbations include density, pressure and velocity perturbations that retain the form of a perfect fluid for the stress-energy tensor, and anisotropic stress that stop it from being a perfect fluid. The metric and the stress-energy tensor each have four scalar, four vector, and two tensor degrees of freedom. Some of these degrees of freedom correspond to a choice of coordinates, or gauge. Bardeen (1980) tackled this by constructing quantities that are invariant under gauge transformations called the Bardeen potentials Φ and Ψ . The source of gauge freedom is that for a given coordinate system in the unperturbed background spacetime, there are multiple possible coordinate systems in the perturbed spacetime. For example the synchronous gauge is a choice of coordinates in which freely falling observers have

their coordinates remain constant. However because there are different ways to assign these coordinates some gauge freedom remains which can lead to terms in the density perturbations that change depending on this choice (Ma & Bertschinger, 1995). The conformal Newtonian gauge (Mukhanov et al., 1992) is a useful choice for studying scalar perturbations, although it can be generalized to incorporate vector and tensor perturbations. In the conformal Newtonian gauge there are two scalar potentials ϕ and ψ that describe the perturbations, appearing in the line element for a flat universe as follows

$$ds^2 = a^2(\eta)[-(1+2\psi)d\eta^2 + (1-2\phi)\delta_{ij}dx^i dx^j], \quad (1.8)$$

where η denotes conformal time. The metric $g_{\mu\nu}$ is diagonal in the conformal Newtonian gauge, and the potentials ϕ and ψ correspond to the gauge-invariant Bardeen potentials so there are no spurious gauge dependent modes as in the synchronous gauge. These potentials are also interpretable: in the Newtonian limit ψ is the gravitational potential, and ϕ is the perturbation to the spatial curvature (Dodelson, 2003). If there is no anisotropic stress in the stress-energy perturbation (i.e., the perturbed components are still a perfect fluid) then the two potentials are equal, $\psi = \phi$, and the Newtonian potential is the only scalar perturbation.

Evolution of scalar perturbations and structure formation

Consider a universe containing baryons, cold dark matter, photons, massless neutrinos and dark energy. The total energy density is

$$\rho = \rho_b + \rho_c + \rho_\gamma + \rho_\nu + \rho_{DE}. \quad (1.9)$$

The pressures are given by $p_b = 0 = p_c$, $p_\gamma = \frac{\rho_\gamma}{3}$, $p_\nu = \frac{\rho_\nu}{3}$ and $p_{DE} = -\rho_{DE}$. We assume no perturbations in the dark energy. Different Fourier modes \vec{k} of linear perturbations evolve independently, so we can treat each mode separately and work with

the density contrast $\delta_{\vec{k}} = \frac{\delta\rho_{\vec{k}}}{\bar{\rho}}$ and velocity $v_{\vec{k}}$ of a specific wavenumber k , dropping the subscript \vec{k} for convenience.

The cold dark matter perturbations only interact through gravity, and their evolution can be described by the conservation equations of a pressureless perfect fluid

$$\delta'_c = -kv_c + 3\phi' \quad (1.10)$$

$$v'_c = -\mathcal{H}v_c + k\phi, \quad (1.11)$$

where $\mathcal{H} \equiv \frac{a'}{a}$ is the conformal Hubble parameter and $x' \equiv \frac{dx}{d\eta}$ denotes a derivative with respect to conformal time, while $\dot{x} \equiv \frac{dx}{dt}$ denotes a derivative with respect to coordinate time. The conformal Hubble parameter is related to the Hubble parameter $H \equiv \frac{\dot{a}}{a}$ by $\mathcal{H} = aH$. The conformal time η is related to the coordinate time by $d\eta = \frac{dt}{a}$.

The photons and baryons are coupled to one another by Thomson scattering:

$$\delta'_b = -kv_b + 3\Phi' + \text{ (collision term)} \quad (1.12)$$

$$v'_b = -\mathcal{H}v_b + k\Phi + \text{ collision term} \quad (1.13)$$

$$\delta'_\gamma = -\frac{4}{3}kv_\gamma + 4\Phi' + \text{ (collision term)} \quad (1.14)$$

$$v'_\gamma = \frac{1}{4}k\delta_\gamma + k\Phi + \text{ collision term.} \quad (1.15)$$

The collision terms in brackets can also be neglected, as the only collision effect that turns out to be important is momentum transfer between photons and baryons, captured in the collision terms in the velocity equations (Ma & Bertschinger, 1995). We can neglect interactions between the neutrinos and other components, such that

$$\delta'_\nu = -\frac{4}{3}kv_\gamma + 4\Phi' \quad (1.16)$$

$$v'_\nu = \frac{1}{4}k\delta_\gamma + k\Phi. \quad (1.17)$$

The conservation equations of fluid dynamics are sufficient to describe the evolution of dark matter (collisionless, Equations (1.10-1.11)) and baryons (collisional, Equations (1.12-1.13)). The Boltzmann equations are needed to capture the phase-space behaviour of the relativistic photons and neutrinos (Equations (1.14-1.17)). In addition to these equations for the evolution of the density contrast and velocity of each component, we also use the Einstein Field Equations.

For cold dark matter, the density contrast δ_c remains constant on superhorizon scales $k \ll \mathcal{H}$. As the comoving Hubble scale \mathcal{H}^{-1} increases, modes enter the horizon and begin to grow, at first slowly in the radiation dominated era as $\delta_c^N \propto \ln(a)$ and then more rapidly when matter begins to dominate, as $\delta_c^N \propto a$. The superscript N specifies the conformal Newtonian gauge.

Matter power spectrum

The amplitude of perturbations on different spatial scales can be studied in harmonic space. The Fourier transform of the matter density contrast $\delta_m = \frac{\delta\rho_m}{\bar{\rho}_m}$ gives the amplitude as a function of wavevector \vec{k} :

$$\delta_m(\vec{k}) = \int d^3\vec{x} e^{-i\vec{k}\cdot\vec{x}} \delta_m(\vec{x}). \quad (1.18)$$

The power at this wavenumber is obtained by squaring the amplitude and taking an expectation value:

$$\langle \delta_m(\vec{k}) \delta_m^*(\vec{k}') \rangle \equiv (2\pi)^3 \delta_D^3(\vec{k} - \vec{k}') P_m(k), \quad (1.19)$$

The matter power spectrum $P_m(k)$ depends only on the amplitude, not the direction of wavevector \vec{k} . δ_D^3 is the three dimensional Dirac delta function. We can also take

the power spectrum of the gravitational potential

$$\langle \phi(\vec{k})\phi^*(\vec{k}') \rangle \equiv (2\pi)^3 \delta_D^3(\vec{k} - \vec{k}') P_\phi(k). \quad (1.20)$$

δ_m and ϕ are related by the Poisson equation

$$k^2 \phi(k) = 4\pi G \bar{\rho}_m a^2 \delta_m(k). \quad (1.21)$$

The primordial power spectrum is assumed to be a power law (and this is later shown to be a good fit to the data).

$$P_m^{initial}(k) = A_m k^{n_s}, \quad (1.22)$$

where $n_s = 1$ for a scale-invariant Harrison-Zel'dovich-Peebles power spectrum with equal power in equal log bins in k (Harrison, 1970; Sunyaev & Zeldovich, 1970; Peebles & Yu, 1970). From the Poisson equation,

$$k^3 P_\phi^{initial}(k) \propto k^{n_s - 1}, \quad (1.23)$$

so $k^3 P_\phi(k)$ is constant for a scale-invariant potential power spectrum (Dodelson, 2003).

In this section we focus on adiabatic initial conditions, which are the dominant fluctuations in the early universe. Adiabatic perturbations have cold dark matter, baryon, radiation, and neutrino overdensities proportional to one another, so that the ratios between the different energy densities is the same everywhere in space. With adiabatic initial conditions, an overdensity in photons has a corresponding overdensity in dark matter and the other components. Isocurvature perturbations, in which the

fluctuations in two different components cancel one another out instead of tracking one another, are discussed in more detail in later chapters.

We can study the evolution of these initial perturbations by considering how the amplitude of each mode of the potential $\phi(k)$ is affected by when their wavelength entered the horizon (the transfer function), and how the overdensities grow with time as matter falls into the potential wells (the growth function). On superhorizon scales (where the wavelength of the Fourier mode is larger than the Hubble scale), different regions are not in causal contact and the potential remains constant, both during radiation and matter domination, although it drops by 10% during the transition through matter-radiation equality. The comoving Hubble scale grows with time so modes that are initially outside the horizon enter the horizon. On subhorizon scales the potential oscillates from the interplay between radiation pressure and gravity, and decays because the universe is expanding. The transfer function therefore drops off for small-scale modes (large k) that enter the horizon earlier and scales approximately as $T(k) \propto \frac{1}{k^2}$ for modes that enter during radiation domination ($k \gg k_{eq}$), neglecting a $\log k$ factor in the transfer function to show approximate scaling for illustrative purposes. The transfer function is constant for large-scale modes that enter the horizon well after matter-radiation equality ($k \ll k_{eq}$) as these modes do not decay. Applying this transfer function to a scale-invariant primordial power spectrum $P^m \propto k$ gives $P^m \propto k$ for $k \ll k_{eq}$ where the transfer function is constant, and $P^m \propto \frac{1}{k^3}$ for $k \gg k_{eq}$ where the transfer function drops off at small scales. This scaling of the matter power spectrum on different scales can be seen in Figure 1.1, with the turnover at $k \approx k_{eq}$.

After the shape of the power spectrum is determined using the transfer function, the power spectrum is evolved through time using the scale-independent growth function D . A change to the growth function would shift the amplitude of the power spectrum in Figure 1.1 up or down but keeps its shape the same. During matter dom-

ination, the matter fluctuations grow as $\delta_m(a) \propto a$, and so in a matter-dominated universe $D(a) \propto a$ and the power spectrum grows with time as $P(k, t) \propto a^2(t)$.




Figure 1.1: The present-day matter power spectrum $P(k)$. The power spectrum turns over at scales corresponding to the horizon size at matter-radiation equality. During the radiation dominated era modes that entered the horizon would decay as the universe expanded, which is why the linear power spectrum, shown as the solid black curve, falls off at large k (small scales). Modes that were on superhorizon scales until after matter-radiation equality are preserved, in this case as a power law at small k . The wiggles in the power spectrum just past the peak correspond to baryon acoustic oscillations. On small scales non-linear contributions begin to be significant as $\delta\rho/\bar{\rho}$ becomes larger, and a more complicated non-linear treatment is needed. The nonlinear power spectrum is shown as the dotted blue curve. These curves were computed using CLASS, the Cosmic Linear Anisotropy Solving System (Blas et al., 2011a), with the halo fit approximation used to calculate the non-linear power spectrum (Smith et al., 2003; Bird et al., 2012)

Evolving the perturbations from an initially scale invariant power spectrum using the transfer function to capture the scale-dependence and the growth function to determine the amplitude gives the linear matter power spectrum shown in Figure 1.1. A nonlinear treatment must be used to obtain accurate results on small scales at late times, when gravitational collapse has grown the fluctuations so that δ is no longer small enough to be treated linearly.

Tensor perturbations

Tensor perturbations evolve independently of scalar perturbations and can therefore be treated separately. Tensor perturbations affect only the spatial part of the metric

$$\delta g_\nu^\mu = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & a^2 h_+ & a^2 h_\times & 0 \\ 0 & a^2 h_\times & -a^2 h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (1.24)$$

for a choice of axes in which the perturbations occur in the x and y directions and the wavevector points in the z direction. The components of the tensor perturbation h_+ and h_\times obey a wave equation with a damping term that depends on the expansion of the universe (Dodelson, 2003). The solutions to this equation are gravitational waves that begin to decay when they enter the horizon. Because of this damping, only the large-scale tensor perturbations persist for long enough to leave an imprint in the cosmic microwave background. Early universe models that predict a significant primordial gravitational wave amplitude can be tested observationally.

1.1.3 The early universe

If we turn back the clock on an expanding universe with growing scale factor $a(t)$, then at early times a was very small, and if a keeps shrinking as we go back further in time then we reach a singularity, $a = 0$. This initial singularity is known as the Big Bang.

Let us take the Λ CDM universe described earlier, evolving it forward from $t = 0$ through a period of radiation domination, through matter domination, to recombination and the free-streaming of CMB photons at $t = 380000$. The particle horizon at that time, or the furthest distance light could have traveled since the $t = 0$, would be

small enough so that regions of the CMB that we observe to be very close to isotropic to one part in 10^5 (Smoot et al., 1992) would not have been in causal contact at last scattering. Regions that had no way of communicating with one another would not have had time to come to the same temperature. This is known as the **horizon problem** (Misner, 1969).

Another mystery that we have not addressed yet is the source of the **primordial perturbations** that seed the growth of large scale structure. A model of the early universe must generate these perturbations, and they must look like the perturbations we observe: dominated by scalar, adiabatic Gaussian fluctuations with an approximately scale-invariant power law.

Thirdly, we have the **flatness problem**: why is the universe observed to be so close to flat today? Any deviation from flatness would grow with time, and the current constraints mean that the universe had to be much closer to the exact critical density in the past to still be so close to flat now. This presents a fine tuning problem: the universe has to be made very close to flat at early times with no apparent mechanism for causing this flatness.

Inflation

An early period of rapid expansion called inflation was proposed to solve these problems (Guth, 1981; Guth & Weinberg, 1983). Inflation solves the horizon problem because at the beginning of inflation the entire observable universe was much smaller and causally connected. Inflation also predicts quantum fluctuations that can source the perturbations observed in the CMB. Some of these fluctuations are stretched to super-horizon scales during inflation, so both sub-horizon and super-horizon perturbations are generated. As the universe is stretched, it is also made more flat.

There are a range of theories of inflation, but the simplest one involves a single scalar field, the inflaton, that rolls slowly down a shallow potential (with kinetic energy

smaller than potential energy, i.e. the inflaton’s potential energy dominates the energy density of the universe at early times) (Linde, 1982; Albrecht & Steinhardt, 1982). Inflation generates both scalar and tensor perturbations from quantum fluctuations. The scalar perturbations seed the large scale structure we observe today. Slow-roll single-field inflation predicts adiabatic, Gaussian, close to scale invariant primordial scalar perturbations. The tensor perturbations to the metric are gravitational waves which produce a specific signature in the cosmic microwave background polarization.

Other early universe models

Questions of fine-tuning problems within inflation, as well as concerns about the lack of predictive ability of a multiverse scenario, prompted alternate models of the early universe to solve the horizon, flatness and inhomogeneity problems, such as cyclic models with periods of contraction and expansion (Steinhardt & Turok, 2002). One such example is a bouncing cosmology (Ijjas & Steinhardt, 2018), with a period of contraction followed by a classical bounce (the universe does not get small enough to require quantum gravity), leading to expansion. This model does not predict primordial gravitational waves, and may therefore be observationally distinguished from inflationary models.

1.1.4 Cosmic Microwave Background

At early times the universe was hot and dense. Photons were energetic enough to ionize hydrogen and helium, and the photons and baryons interacted with one another via Thomson scattering between photons and electrons, and Coulomb scattering between electrons and protons, creating a coupled photon-baryon fluid. Photons had a short mean free path between interactions.

As the universe expanded and cooled, protons and electrons could form neutral atoms and there were fewer free electrons to scatter the photons. The mean free

path of photons became long and the ‘primordial fireball radiation’ was released to free-stream through the universe (Dicke et al., 1965). This radiation cooled further with the expansion of the universe and is observed at the present time as the cosmic microwave background (Penzias & Wilson, 1965). Today the CMB has a blackbody spectrum with temperature $T = 2.725$ K (Mather et al., 1992).

The amplitude of the CMB temperature anisotropy in different directions on the sky can be decomposed into the basis of spherical harmonics

$$\frac{\delta T(\theta, \phi)}{\bar{T}} = \frac{T(\theta, \phi) - \bar{T}}{\bar{T}} = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\theta, \phi). \quad (1.25)$$

The distribution of the $a_{\ell m}$ ’s for each $\ell > 0$ has zero mean and variance C_{ℓ} :

$$\langle a_{\ell m} \rangle = 0 \quad (\ell > 0) \quad (1.26)$$

$$\langle a_{\ell m} a_{\ell' m'}^* \rangle = \delta_{\ell \ell'} \delta_{m m'} C_{\ell}. \quad (1.27)$$

To make a perfect measurement of the variance of each mode we could imagine a scenario where we could study an ensemble of many universes with the same cosmology as ours but different realizations of the initial fluctuations (different draws from the same distribution). Averaging over this ensemble of universes would allow us to measure the variance of each $a_{\ell m}$ as precisely as we wanted to.

In practice, we only have one universe to study but we can use the fact that for each ℓ we can measure $2\ell + 1$ spherical harmonic amplitudes $a_{\ell m}$ (for $m = -\ell \dots 0 \dots \ell$), each of which is drawn from a distribution with the same variance. The variance C_{ℓ} of these modes gives us an estimate of the power of each ℓ mode. Thus instead of taking an ensemble average over different realizations of the universe, we average over

different values of m for each ℓ :

$$C_\ell = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} a_{\ell m} a_{\ell m}^* \quad (1.28)$$

For large values of ℓ we have many values of m , and therefore good statistical errors on our measurement of the variance C_ℓ . However for small values of ℓ we are limited by the number of modes that we can measure. The uncertainty caused by averaging over a finite number of modes is called cosmic variance. It is largest for small values of ℓ and is the limiting factor for CMB experiments on large angular scales. The uncertainty due to cosmic variance scales inversely with the square root of the number of samples for each ℓ :

$$\Delta C_\ell = \sqrt{\frac{2}{2\ell + 1}} C_\ell. \quad (1.29)$$

The variance C_ℓ as a function of angular wavenumber ℓ is called the angular power spectrum. The two dimensional angular power spectrum can be obtained from the three dimensional power spectrum using spherical Bessel functions to project onto the plane of the sky

$$C_\ell = \frac{2}{\pi} \int_0^\infty dk \, k^2 P(k) \left| \frac{\theta_\ell(k)}{\delta(k)} \right|^2 \quad (1.30)$$

where the transfer function converting from $P(k)$ to $C(\ell)$ depends on an integral involving the source function S and the spherical Bessel function j_ℓ :

$$\theta_\ell(k, \eta_0) = \int_0^{\eta_0} d\eta S(k, \eta) j_\ell(k(\eta_0 - \eta_*)). \quad (1.31)$$

The power at angular wavenumbers ℓ depends on the peak of the spherical Bessel function $j_\ell(kd_A^*)$, where d_A^* is the comoving angular diameter distance to the surface of last scattering. This function peaks at ℓ close to but slightly less than kd_A^* (Dodelson, 2003).

The power spectra of the CMB temperature and polarization are shown in Figure 1.2 for a Λ CDM model. For visualization purposes we usually plot $\ell(\ell + 1)C_\ell/2\pi$, so that a scale-invariant initial power spectrum would translate on large scales (where there isn't any interesting recombination physics happening) to a flat line in the angular power spectrum.

Temperature power spectrum features

The CMB temperature power spectrum is shown in the black curve in Figure 1.2 (the top curve). The plotted power spectrum begins at the quadrupole $\ell = 2$. The amplitude of the monopole $\ell = 0$ gives the average temperature of the CMB, which is $T = 2.725 \pm 0.001$ K (Mather et al., 1992). The dipole $\ell = 1$ is a few milliKelvin and is caused by a Doppler shift due to our motion with respect to the CMB, which is usually subtracted.

The Sachs-Wolfe plateau at $\ell < 100$ is caused by the gravitational redshift of CMB photons climbing out of potential wells for perturbations at angular scales larger than the horizon at recombination (Sachs & Wolfe, 1967a). Approximately scale-invariant primordial perturbations make this section of the power spectrum flat in $\ell(\ell + 1)C_\ell$.

The power spectrum then has a series of acoustic peaks, caused by oscillations in the photon-baryon fluid before recombination. Gravity pulls the photons and baryons into potential wells, and then pressure resists this infalling and pushes the fluid outwards. This happens at all scales, with different timescales of oscillation and the modes that are at their extrema at recombination show up with more power as the acoustic peaks. The first peak at $\ell \sim 200$ (corresponding to angular separations of about a degree on the sky) comes from the modes that entered the horizon and then had time to collapse to maximum compression at recombination. The second peak corresponds to the mode that went through compression and was then caught at maximum rarefaction at recombination. Baryons weight these oscillations with-




Figure 1.2: The cosmic microwave background temperature (black) and E- and B-mode polarization (blue and green) angular power spectra for the *Planck* best-fit Λ CDM cosmological model (Planck Collaboration et al., 2020c). The solid curves depict the gravitationally lensed spectra and the dashed curves show the unlensed spectra. These spectra were computed using CLASS (Blas et al., 2011a).

out contributing to the pressure, making the compressions more significant than the rarefactions. This enhances the amplitude of the odd numbered peaks over the even peaks, hence the second peak is lower than the first and the third peak.

During the process of recombination, photons executed a random walk as they scattered off the remaining free electrons. Therefore on small scales, hot and cold spots were mixed together, washing out the CMB’s acoustic peaks. This Silk damping causes the power spectrum to fall off at large ℓ (Silk, 1968).

Some secondary effects are imprinted on the CMB as the photons travel from the last scattering surface to us. During reionization, when there are once more free electrons present, some photons scatter out of our line of sight, suppressing the power spectrum by a factor that depends on the optical depth to reionization τ . The CMB photons are also gravitationally lensed by the matter that they pass. The effect of

lensing is to slightly smooth the peaks of the CMB power spectrum and to shift power to small scales in the damping tail (Lewis & Challinor, 2006). The unlensed power spectra are shown as the dotted curves in Figure 1.2. Additionally, the thermal and kinetic Sunyaev-Zeldovich (SZ) effects cause distortions to the CMB temperature in the directions of galaxy clusters, arising from photons inverse Compton scattering off hot electrons and moving electrons along the line of sight (Sunyaev & Zeldovich, 1972).

Polarization power spectrum features

Thomson scattering in the presence of a quadrupole anisotropy at last-scattering creates an E-mode (curl-free) polarization pattern in the CMB. The E-mode power spectrum is shown as the blue curve in Figure 1.2. The T and E fluctuations are correlated because they are both caused by the scalar perturbations at last scattering, so the TE power spectrum (not shown) is also useful in cosmology. The E-mode peaks come from modes that are at their velocity extrema at recombination, so they are out of phase with the temperature peaks which come from displacement extrema.

Gravitational lensing affects the E-mode power spectrum in a similar way to the temperature power spectrum: spreading out the peaks and transferring power to small scales. The unlensed E-mode power spectrum is shown in dotted blue in Figure 1.2. Lensing also converts some E-mode polarization to B-mode polarization. This is the dominant contribution to the B-mode power spectrum at small scales, peaking at $\ell \sim 1000$. The lensed B mode power spectrum is the solid green curve in Figure 1.2.

Tensor perturbations in the form of primordial gravitational waves give rise to a primordial B-mode polarization signal; scalar perturbations do not generate primordial B-modes. This tensor signal has not been detected yet, although the lensed B-modes have been measured, so we only have an upper limit on the magnitude of the primordial tensor to scalar ratio r (Ade et al., 2021). During reionization, an

additional large-scale polarization signal is generated, giving the reionization bump at $\ell < 10$ in the E-mode and B-mode power spectrum. The green dotted curve in Figure 1.2 shows the unlensed B-mode power spectrum, with the reionization bump at $\ell < 10$ and the gravitational wave signal for a tensor to scalar ratio of $r = 0.01$. The gravitational wave signal peaks at around $\ell = 100$ and then falls off.

1.2 CMB analysis

1.2.1 Effect of cosmological parameters on the CMB power spectrum

The physics of the CMB power spectrum depends on many parameters, which can then be constrained using CMB observations.

Λ CDM model parameters

The following parameters are measured by CMB experiments in the standard Λ CDM model of the universe, which requires just 6 parameters to describe it.

Scalar amplitude A_s : Changing the amplitude of the primordial fluctuations scales the entire power spectrum without changing its shape or shifting it to different scales. The amplitude of the temperature power spectrum is degenerate with the optical depth to reionization τ . This degeneracy is broken in the polarization spectra. *Planck* finds a best fit scalar amplitude of $A_s = 2.1 \times 10^{-9}$ ([Planck Collaboration et al., 2020c](#)).

Scalar spectral index n_s : The tilt of the power spectrum affects the tilt of the CMB power spectrum and the relative power on small and large scales. *Planck* finds the best-fit scalar spectral index to be close to scale-invariant (which corresponds to $n_s = 1$) but with a slight red-tilt: $n_s = 0.965$ ([Planck Collaboration et al., 2020c](#))

Optical depth τ : Reionization lowers all of the acoustic peaks uniformly, and in this way is degenerate with the amplitude A_s . However reionization also produces polarization on large scales, so the low- ℓ E-mode power spectrum provides a powerful constraint on the optical depth which in turn improves our ability to measure the amplitude. *Planck* measures a best fit value of $\tau = 0.054$ ([Planck Collaboration et al., 2020c](#)).

Baryon density $\Omega_b h^2$: An increase in the baryon density weights the oscillations of the photon baryon fluid before recombination, making compressions of the fluid more pronounced than rarefactions, and thus increasing the height of the odd-numbered acoustic peaks relative to the even numbered peaks. *Planck* finds that baryons make up about five percent of the total energy content of the universe ([Planck Collaboration et al., 2020c](#)).

Dark matter density $\Omega_c h^2$: Increasing the dark matter content while keeping the radiation and baryon energy density fixed makes the effect of the potentials more important in the acoustic oscillations relative to the radiation pressure, reducing the amplitude of the oscillations and lowering the height of the acoustic peaks. The baryon weighting also becomes relatively more important, enhancing the odd peaks relative to the even peaks. *Planck* finds that cold dark matter makes up about 26% of the total energy content of the universe, around five times more than baryons ([Planck Collaboration et al., 2020c](#)).

Angular acoustic scale θ_* : The physical scale of the first acoustic peak corresponds to the sound horizon at recombination, so changing the sound speed before recombination would shift the peaks to lower or higher values of ℓ . The observed angular scale of this peak depends on the geometry of the universe and the distance to the last scattering surface, so the peak's position can be shifted by changing the curvature or expansion history of the universe. *Planck* finds that this angular scale is

$\theta_* = 0.59650^\circ$ to 0.03% precision, corresponding to a physical scale of $r_* = 144.4$ Mpc (Planck Collaboration et al., 2020c).

Derived Λ CDM parameters

These are parameters that can be calculated from the six basic parameters in the Λ CDM model, and include the following:

Hubble constant H_0 : Changing the Hubble constant changes the expansion history of the universe and therefore the distance to the last scattering surface and the angular scale of the acoustic peaks. Thus a change in the Hubble constant would shift the peaks to lower or higher ℓ . *Planck* measures the Hubble constant to be $H_0 = 67.4 \pm 0.5$ km s $^{-1}$ Mpc $^{-1}$, which is lower than the value preferred by Type 1a Supernovae observations of $H_0 = 73.04 \pm 1.04$ km s $^{-1}$ Mpc $^{-1}$ (Planck Collaboration et al., 2020c; Riess et al., 2021).

Dark energy density Ω_Λ : The dark energy content also affects the expansion history and therefore the scale of the acoustic peaks. In flat Λ CDM the dark energy fraction can be calculated from the cold dark matter and baryon densities, as the total density is then the critical density for a flat universe. *Planck* finds that 68.5% of the energy content in the universe is in dark energy, with the remaining 31.5% made up of matter in the form of cold dark matter and baryons (Planck Collaboration et al., 2020c).

Additional parameters

These parameters are extensions to the basic Λ CDM parameterization above. We give a few examples below. The curvature and effective number of neutrino species are fixed to their standard values when constraining Λ CDM parameters, but may be allowed to vary in extensions to the standard model. The tensor to scalar ratio r has not been detected by current experiments, so only upper limits exist.

Curvature Ω_K : the curvature of the universe affects how a fixed physical scale is translated into an observed angular scale, so changing the curvature will change the position of the first peak (corresponding to the physical sound horizon at recombination) and shift it to a different angular wavenumber ℓ . The curvature is very well constrained by the CMB, in combination with galaxy surveys, to be close to 0, so the universe is taken to be flat in the standard model.

Effective number of neutrino species N_{eff} : Increasing the effective number of neutrino species reduces small scale power of the CMB oscillations because neutrinos do not provide pressure to contribute to acoustic oscillations like photons do. Changing N_{eff} also changes the phases of the acoustic oscillations and therefore shifts the positions of the acoustic peaks ([Abazajian et al., 2015](#)).

Tensor to scalar ratio r : The tensor to scalar ratio measures the power in gravitational waves in the early universe, and therefore increases in r will increase the primordial amplitude of the CMB B-mode polarization. $r = 0$ corresponds to no primordial B-modes.

1.2.2 Foregrounds

The observed temperature and polarization power spectra depend on the cosmological parameters described above, but also on radiation from other galaxies, the imprints of hot electrons in galaxy clusters, and the dust in our own Galaxy. This means that CMB experiments can probe a range of interesting Galactic and extragalactic physics. It also means that we need to account for and separate out these effects when studying cosmological parameters. Many of these foregrounds have different frequency spectra to the CMB, so having a range of frequency channels helps to constrain models of the foregrounds in addition to the CMB.

The level of point sources (AGN and dusty galaxies) in the *Planck* temperature foreground power at 143 GHz is shown in the left panel of Figure 1.3, with foreground

amplitudes from [Planck Collaboration et al. \(2020a\)](#). This level will be lower for higher resolution experiments which can detect and mask out more point sources. The best-fit B-mode polarization of Galactic dust at 150 GHz in the BICEP region is shown in the right panel of Figure 1.3. This is a relatively clean patch of the sky in terms of Galactic dust. The dust is modeled as a modified blackbody in frequency with spectral index β_d and temperature $T_d = 19.6K$ ([Ade et al., 2021](#))

$$I_\nu \propto \nu^{\beta_d} B(\nu, T_d), \quad (1.32)$$

where B is the blackbody function. The dust is assumed to follow a power law as a function of angular wavenumber ℓ with spectral index α_d

$$\ell(\ell + 1)C_\ell/2\pi \propto \ell^{\alpha_d}. \quad (1.33)$$

These and other foregrounds are included as part of the model of the observed CMB, and their parameters are marginalized over when studying cosmological constraints.

1.2.3 Bayes theorem and data analysis

When analyzing experimental data, we want to find the model that best describes the data (model selection) and the parameter values for that model that give the best fit to the data (parameter estimation). Bayes theorem for the probability of the parameters θ of a model given the data \mathcal{D} is

$$p(\theta|\mathcal{D}) = \frac{p(\mathcal{D}|\theta)p(\theta)}{p(\mathcal{D})}. \quad (1.34)$$

This states that the posterior probability $p(\theta|\mathcal{D})$ of parameters θ given the data \mathcal{D} is proportional to the product of the likelihood of the data given the parameters $p(\mathcal{D}|\theta)$ and the prior probabilities of the parameters $p(\theta)$. The normalization factor is




Figure 1.3: The angular power spectrum of dominant foregrounds in temperature (left) and polarization (right). The left plot shows the CMB temperature power spectrum (black solid curve) and the temperature power spectrum at 143 GHz of residual undetected Poisson-distributed point sources that have not been masked in the *Planck* data (blue dashed curve). The right plot shows the B-mode polarization power spectrum for $r=0.01$ (black solid curve) and $r=0$ (black dot-dashed curve), together with the level of B-mode polarization from galactic dust in the BICEP region of the sky (blue dashed curve).

the evidence $p(\mathcal{D})$ which can be computed from the likelihood and the prior $p(\mathcal{D}) = \int p(\mathcal{D}|\theta)p(\theta)d\theta$ (MacKay, 2003).

The likelihood $p(\mathcal{D}|\theta)$ is a measure of how well the observed data match the model prediction with parameters θ . For CMB experiments we use Boltzmann codes to generate the theoretical CMB power spectrum for different values of the parameters θ , add modeled foregrounds, bin the theory power spectrum in the same way as the data and then compare the theory power spectrum to the data for different values of the parameters. At large ℓ the angular power spectrum likelihood can be approximated as Gaussian by the central limit theorem, giving

$$\mathcal{L} \propto \exp \left(-\frac{1}{2} (C_\ell^{theory} - C_\ell^{data}) Q^{-1} (C_\ell^{theory} - C_\ell^{data}) \right), \quad (1.35)$$

where the covariance matrix Q includes instrumental noise, cosmic variance and the effects of mode coupling. On large scales (low- ℓ) the likelihood for each C_ℓ is not

Gaussian and a different likelihood is needed; this is discussed further in later chapters of this dissertation.

The prior on each parameter captures what we know about the parameter value before performing the experiment. If our data are very constraining then the posterior distribution should not depend strongly on the prior. Uninformative priors are often used for parameters we expect to constrain well, so that our constraints are independent of previous experiments and can be compared. However for parameters that are difficult to constrain with the data, priors from previous experiments can be very helpful, especially in the case of parameter degeneracies. For example the optical depth to reionization τ is constrained well by the *Planck* satellite's measurements of large-scale E-mode polarization, which is difficult to obtain from the ground. Ground based CMB experiments can use the *Planck* constraints in their τ priors to prevent the degeneracy between A_s and τ from broadening their constraints.

Once we have defined the likelihood and the priors we can study the distribution of the posterior probability of the parameters given the data, to find the best-fit parameters and their errors. The evidence $p(\mathcal{D})$ can be ignored at the parameter estimation stage as it changes the overall normalization and not the shape of the distribution. One way to obtain samples from the posterior distribution is using Markov Chain Monte Carlo methods such as the Metropolis-Hastings algorithm (MacKay, 2003), which executes a random walk around the parameter space, deciding whether to move to the next proposed location based on the relative probability density of the proposed and current locations. After an initial burn-in period, the Metropolis-Hastings algorithm produces a chain of samples from the posterior distribution. These samples can be used to study the probability distributions of the parameters and find their best-fit values and errors.

Choosing which out of a selection of possible models best describes your data involves a trade-off between how well the model matches your data and how many

parameters the model has. A model with too many parameters can overfit the data, so for each additional parameter you need a corresponding significant increase in goodness-of-fit to justify including the extra parameter. Models can be compared using the model evidence in the denominator of Bayes theorem $p(\mathcal{D})$ which could be written more precisely as $p(\mathcal{D}|\mathcal{M})$, or the probability of the observed data given the model \mathcal{M} . Because the integral for the evidence is computationally challenging to compute, quantities such as the Bayesian Information Criterion (BIC) or Akaike Information Criterion (AIC) can also be compared for competing models. The BIC is given by

$$\text{BIC} = k \ln n - 2 \ln \mathcal{L}_{max}, \quad (1.36)$$

where k is the number of parameters in the model, n is the number of data points, and \mathcal{L}_{max} is the value of the likelihood for the best-fit model parameters. Both the BIC and AIC weigh the best-fit likelihood of a model against how many parameters it has.

1.3 CMB experiments

Relic radiation from the primeval fireball was predicted by [Dicke et al. \(1965\)](#) and first measured by [Penzias & Wilson \(1965\)](#) as a mysterious 3 Kelvin source of noise in their radio telescope. Scientists predicted anisotropies in the CMB corresponding to fluctuations that seeded the large-scale structure we see today. These anisotropies were observed by the Cosmic Background Explorer (COBE) satellite as faint fluctuations at the level of one part in 100 000 in the CMB temperature maps ([Smoot et al., 1992](#)). COBE also measured the spectrum of the CMB at a range of frequencies, finding it consistent with a blackbody with temperature $T = 2.725\text{K}$ ([Mather et al., 1992](#)).




Figure 1.4: The CMB temperature power spectrum from the first year of data from The Wilkinson Microwave Anisotropy Probe (*WMAP*) satellite from [Spergel et al. \(2003\)](#). The points with errorbars are the binned power spectrum. The grey dots show the unbinned points. The solid curve is the best fit theory spectrum.

The Wilkinson Microwave Anisotropy Probe (*WMAP*) satellite measured the first few peaks of the CMB temperature power spectrum, shown in Figure 1.4, and used this together with the temperature-polarization TE cross spectrum to constrain the six-parameter Λ CDM model of the universe ([Spergel et al., 2003](#)).

The *Planck* era

The *Planck* satellite was the third-generation of space-based CMB experiments. *Planck* scanned the sky in nine frequency bands (ranging from 30 GHz to 857 GHz) from August 2009 to October 2013. The different frequency bands were used to characterize and mitigate the foregrounds. *Planck* measures the first seven acoustic peaks in the CMB temperature, and also measures the E-mode polarization power spectrum and the TE cross spectrum ([Planck Collaboration et al., 2020a](#)). The *Planck* temperature power spectrum is cosmic variance limited up to $\ell \lesssim 1600$. The

Planck 2018 legacy data release is consistent with a 6 parameter Λ CDM model of the universe, and provides tight constraints on those parameters (Planck Collaboration et al., 2020c). *Planck* also strongly constrains extensions to the Λ CDM model, preferring the basic six-parameter model to extensions with more parameters.

Planck was the first experiment to detect $n_s < 1$ at high significance, measuring $n = 0.965 \pm 0.004$, consistent with the close to scale-invariant but slightly red spectrum predicted by inflation (Planck Collaboration et al., 2020a).

Because of the challenges of measuring very large angular scales from the ground, the low- ℓ *Planck* data is important for combining with current and future ground-based experiments.

Ground-based and balloon experiments

The next phase of CMB experiments tend to have one of two broad aims: either to measure the small scale anisotropies in temperature and E-mode polarization with high resolution and sensitivity over a sufficiently large area of the sky to obtain tight cosmological constraints, or to measure the large scale B-mode polarization signal with very low noise in a small patch of the sky (to enable deep observations) to go after primordial gravitational waves.

The Atacama Cosmology Telescope (ACT) is a six meter telescope located in the Atacama Desert in Chile. It began observing in 2007, and was upgraded to be polarization sensitive in 2013. Two advantages of ground-based experiments are that the dishes can be large, giving high resolution maps, and that detectors can be upgraded as the technology improves, increasing the sensitivity of the instrument over time. The latest ACT data release (ACT DR4) used temperature and polarization measurements over 6000 deg^2 on the sky (Aiola et al., 2020).

The South Pole Telescope is a 10 meter telescope located in Antarctica. The SPT-3G survey covers 1500 deg^2 on the sky. SPT and ACT observations supplement

the *Planck* data at high ℓ and also serve as a consistency check in the ℓ range covered by both *Planck* and ground based experiments. The CMB spectra and cosmology constraints from ACT and SPT are consistent with those from *Planck*.

BICEP (Background Imaging of Cosmic Extragalactic Polarization) and the Keck Array are located at the Amundsen–Scott South Pole Station in Antarctica like SPT, but focus on measuring primordial B-modes on large scales. They have placed the tightest upper limit so far on the tensor-to-scalar ratio of $r < 0.036$ but have not yet made a detection of primordial gravitational waves ([Ade et al., 2021](#)).

Balloon-borne experiments such as SPIDER are an intermediate option between ground-based experiments and satellites, with the advantage of minimizing the effects of the atmosphere and benefiting from the latest detector technology without being as costly as satellite experiments. They fly at high altitudes in the stratosphere for a few weeks to gather data which can then be used to map the CMB and galactic dust polarization and put constraints on the tensor to scalar ratio r ([SPIDER Collaboration et al., 2021](#)).

Neither ground nor balloon experiments can reach low enough in ℓ to measure the E-mode or B-mode reionization bump down to the quadrupole $\ell = 2$, so until another CMB satellite is launched the *Planck* data remains our best measurement on those scales.

Upcoming CMB experiments

The Simons Observatory (SO) is an upcoming CMB experiment with two components: the LAT (large aperture telescope) to measure a large fraction of the sky with high resolution and place tight constraints on cosmological parameters, and the SATs (small aperture telescopes) to measure or place an upper limit on primordial gravitational waves ([Lee et al., 2019](#)). SO is currently being constructed in the Atacama Desert. Figure 1.5 from [Lee et al. \(2019\)](#) shows the *Planck* and BICEP2/Keck data




Figure 1.5: CMB power spectrum constraints from Planck’s 2018 legacy data release (orange) and BICEP2/Keck (green), together with forecasted errors for the upcoming Simons Observatory ground-based experiment (SO, blue) and the LiteBIRD satellite (red). Simons Observatory will measure temperature and polarization fluctuations for a wide range of angular scales, including small scales (large ℓ). LiteBIRD will measure the large scale (low- ℓ) polarization, including the reionization bump at $\ell < 10$ which cannot be measured from the ground. Figure source: [Lee et al. \(2019\)](#)

and the forecasted SO errorbars. The BICEP/Keck collaboration installed the first of four receivers that will constitute the six-frequency BICEP Array in 2020 ([Moncelsi et al., 2020](#)), and will continue to search for a primordial gravitational wave signal in the CMB B-modes alongside the SO SATs.

The next generation CMB ground based experiment after SO will be CMB S4, a collaboration involving many of the current ground based CMB teams ([Abazajian et al., 2016](#)). The next planned satellite experiment is LiteBIRD ([Hazumi et al., 2020](#)), which is expected to launch in the late 2020s.

1.4 Dissertation Overview

The rest of this dissertation is structured as follows. In Chapter 2 we present Cos-MOPED, a *Planck* likelihood that has been compressed using the MOPED technique to require only one number per parameter of interest, thereby reducing the need to handle large covariance matrices. In Chapter 3 we fit binned bandpowers to the low- ℓ *Planck* temperature and polarization likelihoods, showing that they closely approximate a shifted log-normal distribution, and construct differentiable Python likelihoods for the *Planck* low- ℓ data. In Chapter 4 we explore the effects of different foreground modeling choices on BICEP’s tensor-to-scalar ratio constraint, and we present a foreground-marginalized likelihood for BICEP that relies on just 9 CMB bandpowers and has no nuisance parameters. Chapter 5 shows constraints on the primordial power spectrum using ACT, SPT and *Planck* data for a broken power law adiabatic model, as well as a model with both adiabatic and isocurvature fluctuations in which the isocurvature fluctuations vary independently on different scales. Chapter 6 shows the improved limits that the upcoming Simons Observatory will be able to place on binned isocurvature models, especially on small scales. We conclude in Chapter 7.

Chapter 2

Data compression in cosmology: A MOPED compressed likelihood for *Planck* data

When analyzing a large dataset with many different data points, but only a few relevant parameters, it can be useful to compress the data, retaining all the information that is relevant to the parameters of interest. These compressed data can then be used for parameter inference. [Tegmark et al. \(1997\)](#) showed that if there is one parameter of interest, a dataset can be compressed to one number without losing any information about that parameter. The Massively Optimized Parameter Estimation and Data compression technique (MOPED) described in [Heavens et al. \(2000\)](#) extends this to multiple parameters, developing optimal linear compression for Gaussian data provided the covariance is independent of the parameters. [Alsing & Wandelt \(2018\)](#) show that compression to the score function (the gradient of the log likelihood) preserves the information content even for non-Gaussian data and data for which the covariance depends on the parameters. Thus the full dataset of N data points can

be compressed to n numbers, where n is the number of parameters of interest, while preserving the Fisher information about those parameters.

This compression step is useful both as a way to generate a simplified likelihood function, and as a step towards likelihood-free inference when the form of the likelihood is not known precisely (e.g., [Alsing & Wandelt, 2018](#)). Compressions of this form have been used to study the star formation history of galaxies ([Reichardt et al., 2001](#); [Heavens et al., 2004](#); [Panter et al., 2007](#)) and considered for exoplanet transit detections ([Protopapas et al., 2005](#)), gravitational wave studies with *LISA* ([Graff et al., 2011](#)), and covariance matrix estimation ([Heavens et al., 2017](#)). [Gupta & Heavens \(2002\)](#) proposed applying MOPED compression to cosmic microwave background (CMB) data, and [Zablocki & Dodelson \(2016\)](#) applied a similar compression scheme to the temperature power spectrum of *WMAP*.

In this chapter we apply the MOPED compression to the *Planck* CMB power spectrum [Planck Collaboration et al. \(2016a\)](#), and show that the standard Λ CDM cosmological parameter constraints can be derived from a compressed likelihood of just six Gaussian-distributed data points. We go beyond earlier analyses of CMB data by precompressing the non-Gaussian large-scale temperature power spectrum into two approximately Gaussian data points. We make the software for the MOPED-compressed likelihood publicly available, as well as for the likelihood computed directly from the binned power spectrum with the inclusion of new large angular scale bins.¹

The outline of the chapter is as follows. In §2.1 we describe the public *Planck* likelihood and the precompression we implement to better approximate it as Gaussian. In §2.2 we describe the MOPED compression scheme we apply, and in §2.3 show

¹At <https://github.com/heatherprince/cosmoped> and <https://github.com/heatherprince/planck-lite-py>. Both codes have been updated with the 2018 temperature and polarization Plik-lite likelihood

parameter constraints for the Λ CDM model and an example extended model with running of the primordial spectral index. We conclude in §2.4.

2.1 *Planck* likelihood and low- ℓ binning

The current state-of-the-art CMB data come from the *Planck* satellite. The latest cosmological analysis is reported in [Planck Collaboration et al. \(2020c\)](#), with public data from the earlier 2015 analysis described in [Planck Collaboration et al. \(2016a\)](#). The *Planck* temperature power spectrum is shown in Fig. 2.1. The likelihood function, describing the probability of the data given some model, is separated into two main parts for the *Planck* power spectrum analysis, with different approaches for large scales and smaller scales. We summarize these components briefly here.




Figure 2.1: The *Planck* 2015 temperature power spectrum from ([Planck Collaboration et al., 2016a](#)). The non-Gaussian $\ell < 30$ bandpowers are shown with their asymmetrical errors in gray. The two black points at $\ell < 30$ are obtained by estimating the binned spectrum in two angular bins, and have approximately Gaussian distributions. We use these two low- ℓ bins for the likelihoods described in this chapter. The $\ell \geq 30$ binned data are the foreground-marginalized temperature bandpowers from the Plik_lite likelihood for *Planck* 2015. The theoretical power spectrum for the *Planck* 2015 TT+lowTEB best fit parameters ([Planck Collaboration et al., 2016b](#)) computed using CLASS ([Blas et al., 2011b](#)) is shown in blue.

At $\ell \geq 30$ (corresponding to scales smaller than several degrees on the sky) the likelihood \mathcal{L} for the temperature and E -mode polarization power spectra (TT, TE, and EE) is modeled as a Gaussian distribution, with

$$-2 \ln \mathcal{L} = (C_b^{\text{th}} - C_b^{\text{data}})^T Q^{-1} (C_b^{\text{th}} - C_b^{\text{data}}) \quad (2.1)$$

to within an overall additive constant, with binned data C_b^{data} , binned theory C_b^{th} , and binned covariance matrix Q . For the Plik_lite likelihood (Planck Collaboration et al., 2016a), these data spectra represent an estimate of the CMB bandpowers, with foregrounds already marginalized over using the approach of Dunkley et al. (2013).

At $\ell < 30$ the distribution of the angular power spectrum is non-Gaussian. For the 2015 data release, the *Planck* team released a joint pixel based likelihood for temperature and polarization for $\ell \leq 29$ ('lowTEB'). There is also a standalone temperature low- ℓ likelihood based on the foreground-cleaned Commander temperature map, which we use in this chapter. These likelihoods are computed in map space since the distribution of the power spectrum on these scales is non-Gaussian. The 2018 likelihood uses a similar low- ℓ temperature likelihood, and a separate low- ℓ polarization likelihood built from simulations (Planck Collaboration et al., 2020b).

2.1.1 Low- ℓ temperature bins

The compression approach we adopt, which we describe in the next section, is optimal for Gaussian distributions. Since we are interested in a lightweight compression to estimate simple cosmological models, we first compress the $\ell < 30$ *Planck* TT data into two bins with approximately Gaussian distributions. We do this by conditionally sampling the posterior distribution for the power in each bin, estimating

$$p(\theta|d) \propto p(d|\theta)p(\theta). \quad (2.2)$$




Figure 2.2: The probability distribution for D_ℓ for two low- ℓ temperature bins using the *Planck* 2015 low- ℓ Commander likelihood. The best-fit Gaussian distribution is shown in blue.

Here the parameters θ are the binned values $D_{2 \leq \ell \leq 15}$ and $D_{16 \leq \ell \leq 29}$, where $D_\ell = \ell(\ell + 1)C_\ell/2\pi$, assuming a constant value for D in each bin. The likelihood $p(d|\theta)$ is the *Planck* $\ell < 30$ temperature likelihood function. We assume uniform priors on θ . The binning is performed on D_ℓ rather than C_ℓ because D_ℓ is approximately constant for the low- ℓ temperature power spectrum. When binned values of C are required we convert from the binned D values by dividing by the mean of $\ell(\ell + 1)/2\pi$ in the bin.

The distributions of the two low- ℓ power spectrum bins are shown in Fig. 2.2, together with the best-fitting Gaussian distributions. We find

$$\begin{aligned} D_{2 \leq \ell \leq 15} &= 827 \pm 74 \text{ } \mu\text{K}^2 \\ D_{16 \leq \ell \leq 29} &= 854 \pm 49 \text{ } \mu\text{K}^2. \end{aligned} \tag{2.3}$$

Their distributions are close to Gaussian, unlike the distributions for the individual multipoles. These bandpowers are also indicated in Fig. 2.1 (the first two black points), together with the unbinned low- ℓ power spectrum in gray. We chose this binning scheme before sampling parameters; other choices that produce approximately Gaussian distributions would be expected to give similar results.

2.1.2 Low- ℓ polarization (τ prior)

The amplitude of the large-scale polarization signal depends primarily on the optical depth to reionization, as well as the primordial amplitude. To include the low- ℓ polarization data we compress the 2015 polarization information into a single Gaussian prior on the optical depth to reionization, adopting $\tau = 0.067 \pm 0.023$ derived from the *Planck* low- ℓ likelihood using the Low Frequency Instrument (Planck Collaboration et al., 2016a). This is an approximation since τ is correlated with other cosmological parameters, in particular the primordial amplitude A_s . Improved measurements of the optical depth have since been made from the *Planck* High Frequency Instrument Planck Collaboration et al. (2016c, 2020c). However, the purpose of this study is to compress the public 2015 likelihood, and we defer a future refinement of our compression code to include the 2018 polarization information that was recently made public.

2.1.3 Parameter constraints

The effect of describing the low- ℓ temperature data using two Gaussian bins and using a prior on τ in place of the low- ℓ polarization likelihood is shown in Fig. 2.3, which shows the posterior probabilities for the six Λ CDM parameters (the Hubble constant, baryon density, cold dark matter density, amplitude and spectral index of primordial fluctuations, and optical depth to reionization) obtained by sampling three different likelihood combinations.

Replacing the low- ℓ temperature likelihood with a Gaussian likelihood based on two low- ℓ temperature bins results in parameter constraints that agree well, to within 0.1σ (black versus blue-dashed in Fig. 2.3). Here, the black solid curve is derived using our Python implementation of Plik_lite with the additional two Gaussian low- ℓ bins included (Planck-lite-py). The blue-dashed curve shows the posteriors obtained by sampling the *Planck* high- ℓ temperature Plik_lite and low- ℓ temperature-only Comman-




Figure 2.3: Posteriors on the Λ CDM parameters from Planck-lite-py, a Python implementation of the *Planck* Plik_lite likelihood with two Gaussian low- ℓ bins (black solid curve) compared to the combined Plik_lite TT + low- ℓ Commander *Planck* 2015 temperature likelihoods (blue-dashed curve). The parameter constraints agree to within 0.1σ , showing that our two binned low- ℓ data points capture enough information to give equivalent constraints to the full low- ℓ temperature likelihood. A Gaussian prior of $\tau = 0.067 \pm 0.023$ is included in the black and blue curves. We also show results from the public chains from sampling the TT+lowTEB *Planck* 2015 likelihood (gray solid curve) for reference. The small difference is the effect of imposing a prior on the optical depth.

der likelihoods using the CosmoSIS cosmological parameter estimation code (Zuntz et al., 2015). In both cases we compute the theoretical CMB power spectrum using the Cosmic Linear Anisotropy Solving System (CLASS) (Blas et al., 2011b), sample the likelihood using the emcee (Foreman-Mackey et al., 2013) Python implementation of Markov chain Monte Carlo (MCMC) affine-invariant ensemble sampling (Goodman & Weare, 2010), and impose a Gaussian prior on the optical depth $\tau = 0.067 \pm 0.023$ from the *Planck* low- ℓ likelihood (Planck Collaboration et al., 2016a). The low- ℓ data provide an important anchor for constraints on parameters such as spectral index n_s and Hubble parameter h . The Planck-lite-py parameter constraints agree with the Plik_lite+lowT constraints to within 0.1σ , showing that for the standard cosmological

model the full low- ℓ temperature likelihood can be replaced by a Gaussian likelihood with two bins and a diagonal covariance matrix without a significant effect on the parameter constraints.

The gray curve in Fig. 2.3 shows the public *Planck* 2015 chains from the TT+lowTEB likelihood. When comparing parameters constraints from the full TT+lowTEB likelihood with the TT+lowT+ τ prior likelihood some of the parameters shift by up to 0.2σ because we are using a tau prior which is only an approximation to the full low- ℓ polarization likelihood.

We make available Planck-lite-py, a Python implementation of the *Planck* Plik_lite likelihood that includes these low- ℓ binned temperature data.²

2.2 MOPED compression vectors

We use the Massively Optimized Parameter Estimation and Data compression technique (MOPED) described in [Heavens et al. \(2000\)](#) to compress the *Planck* power spectrum. This linear compression is optimal in the sense that the Fisher information content is preserved for Gaussian data when the covariance matrix is independent of the parameters of interest.

The *Planck* data vector used in this analysis is the foreground-marginalized binned temperature angular power spectrum. It would be straightforward to include the TE and EE angular power spectra in the data vector for a combined temperature and polarization compressed likelihood. The binning is performed using a constant weighting in D_ℓ ([Planck Collaboration et al., 2016a](#)), which corresponds to

$$C_b = \sum_{\ell=\ell_b^{min}}^{\ell_b^{max}} w_b^\ell C_\ell, \quad (2.4)$$

²<https://github.com/heatherprince/planck-lite-py>.

where

$$w_b^\ell = \frac{\ell(\ell+1)}{\sum_{\ell=\ell_b^{min}}^{\ell_b^{max}} \ell(\ell+1)}. \quad (2.5)$$

The binned angular power spectrum is the sum of a signal component that depends on the cosmological parameters $\boldsymbol{\mu} = C_b^{th}$ (the noise-free theoretical binned angular power spectrum), as well as a noise component \boldsymbol{n} . The total data vector is thus

$$\boldsymbol{x} = \boldsymbol{\mu} + \boldsymbol{n}. \quad (2.6)$$

The data vector can be compressed into a single number while preserving the information about the first cosmological parameter of interest θ_1 [Tegmark et al. \(1997\)](#)

$$y_1 = \boldsymbol{b}_1^t \boldsymbol{x} \quad (2.7)$$

with

$$\boldsymbol{b}_1 = \frac{\boldsymbol{Q}^{-1} \boldsymbol{\mu}_{,1}}{\sqrt{\boldsymbol{\mu}_{,1}^t \boldsymbol{Q}^{-1} \boldsymbol{\mu}_{,1}^t}}, \quad (2.8)$$

where $\boldsymbol{\mu}_{,1}$ is the derivative of the signal (the theoretical binned temperature angular power spectrum) with respect to the first cosmological parameter, \boldsymbol{Q} is the covariance matrix, and the normalization of the compression vector has been chosen such that $\boldsymbol{b}_1^t \boldsymbol{Q} \boldsymbol{b}_1 = 1$.

Additional compression vectors can be found that produce linear combinations which capture information about the other cosmological parameters, while being orthogonal to the other compression vectors so that each linear combination y_m is




Figure 2.4: Compression vectors for the Λ CDM parameters. The sharp features are caused by binning of the power spectrum. These compression vectors can be applied to the CMB temperature power spectrum to give six numbers that contain as much information about the cosmological parameters as the binned temperature power spectrum. The compression vectors depend on the order in which they are computed, as they are intentionally orthogonal to one another.

uncorrelated with the others [Heavens et al. \(2000\)](#), giving

$$\mathbf{b}_m = \frac{\mathbf{Q}^{-1} \boldsymbol{\mu}_{,m} - \sum_{q=1}^{m-1} (\boldsymbol{\mu}_{,m}^t \mathbf{b}_q) \mathbf{b}_q}{\sqrt{\boldsymbol{\mu}_{,m} \mathbf{Q}^{-1} \boldsymbol{\mu}_{,m} - \sum_{q=1}^{m-1} (\boldsymbol{\mu}_{,m}^t \mathbf{b}_q)^2}}. \quad (2.9)$$

The compression vectors for the six standard Λ CDM parameters are shown in Fig. 2.4. The oscillatory behavior comes from the effect of the acoustic peaks on the derivatives of the theoretical CMB power spectrum. Most of the signal comes from $1000 < \ell < 2500$ because this is where the diagonal of the binned inverse covariance matrix is large (Fig. 2.5). At low ℓ cosmic variance dominates the noise while at high ℓ experimental noise takes over. The noise in each bin is also dependent on the bin width, which varies for different multipoles ([Planck Collaboration et al., 2016a](#)).




Figure 2.5: The diagonal of the binned temperature inverse covariance matrix for the 217 temperature bins (2 for low- ℓ and 215 for high- ℓ). The diagonal elements are small for $\ell < 1000$, then rise steeply for $1500 < \ell < 2000$ before dropping again when the experimental noise approaches the signal.

Applying these compression vectors to the data vector

$$y_m = \mathbf{b}_m^t \mathbf{x} \quad (2.10)$$

gives a set of M numbers $y_m, m = 1, \dots, M$ which contain as much information about the cosmological parameters of interest θ_m as the full angular power spectrum with N bins. For the binned temperature power spectrum the data vector has length $N = 217$, and for the standard Λ CDM cosmology the number of parameters of interest is $M = 6$.

The *Planck* power spectrum and covariance matrix that are used to compute the compression vectors are already binned, so the compressed statistics for the data come from applying binned compression vectors to the power spectrum. To compress the theoretical CMB power spectrum we use a version of the compression vectors that includes the binning, weighting each multipole appropriately as per Eq. (2.4), so that the binning and compression are achieved in the same step.




Figure 2.6: Posteriors on the Λ CDM parameters from the compressed CosMOPED likelihood (red-dashed curves) and the Planck-lite-py Python implementation of the Plik_lite likelihood with two low- ℓ bins (black solid curves). The parameter constraints agree to within 0.1σ , validating the CosMOPED compression scheme. A Gaussian prior of $\tau = 0.067 \pm 0.023$ is included for both curves.

2.3 Likelihood and parameters

We now describe the compressed likelihood and compare it to our Planck-lite-py implementation of the *Planck* 2015 Plik_lite likelihood.

2.3.1 Format of the likelihood

Each compressed statistic y_m is Gaussian distributed with unit variance. The y_m 's are uncorrelated with each other by design, so the total likelihood is the product of the likelihoods from each statistic. The likelihood of the parameters given the (compressed) data thus takes a simple form

$$-2 \ln \mathcal{L} = \sum_{m=1}^M \frac{(y_m - \langle y_m \rangle)^2}{2} + \text{constant}, \quad (2.11)$$

where y_m is the compressed statistic from the data and $\langle y_m \rangle = \mathbf{b}_m^t \boldsymbol{\mu}$ is the corresponding compressed statistic from the model (the theoretical temperature power spectrum).

2.3.2 Parameter constraints

The parameter constraints for the six-parameter Λ CDM model are shown in Fig. 2.6. The compressed likelihood (red-dashed curve) and the Python Plik_lite implementation (black solid curve) agree to within 0.1σ for each parameter. Both likelihoods were sampled with emcee, with a Gaussian prior on the optical depth $\tau = 0.067 \pm 0.023$ and the low- ℓ bins described in Sec. 2.1. The MOPED compression that we have applied thus results in a likelihood that is equivalent to the uncompressed case. As we showed earlier, the Python Plik_lite implementation is in good agreement with the full *Planck* temperature likelihood.

2.3.3 Effect of fiducial model

The theory vector $\boldsymbol{\mu}$ used to compute the compression vectors depends on the fiducial model parameters used. If the fiducial model is wrong then the compression is no longer optimal and the M compressed statistics are not exactly independent. However, in practice using a different fiducial model does not have a significant effect on the compressed likelihood, in agreement with the findings of [Zablocki & Dodelson \(2016\)](#). A shift of order 3σ in the fiducial parameters has an insignificant effect on the conditional probability slices obtained from the compressed likelihood.




Figure 2.7: The constraints on running of the scalar spectral index from the compressed CosMOPED likelihood (red dashed) agree to within 0.1σ with the uncompressed Planck-lite-py (black). A Gaussian prior of $\tau = 0.067 \pm 0.023$ is used for CosMOPED and Planck-lite-py. The distribution using the public *Planck* chains (TT+lowTEB, gray) is slightly broader due to the more accurate treatment of low- ℓ polarization.

2.3.4 Non- Λ CDM cases

We demonstrate the application of this compression technique to a one parameter extension to the Λ CDM case by sampling running of the scalar spectral index $d \ln n_s / d \ln k$ in addition to the six Λ CDM parameters shown above.

The results are shown in Fig. 2.7, which compares the CosMOPED constraints to the Planck-lite-py constraints (with the same τ prior as above), as well as the *Planck* TT+lowTEB chains from the 2015 data release. The likelihoods show excellent agreement. The CosMOPED and Planck-lite-py posteriors are slightly narrower than for the *Planck* TT+lowTEB chains; this is because the Gaussian prior on τ which we use in the CosMOPED likelihood comes from the low- ℓ likelihood assuming the Λ CDM model. For Λ CDM extension models with parameters that correlate with τ , this prior is slightly too narrow.

2.4 Discussion

We have demonstrated that the low- ℓ *Planck* temperature data is well represented by two Gaussian bins for simple cosmological models. We have also shown that applying MOPED linear compression to the *Planck* 2015 binned temperature power spectrum allows us to construct a simple likelihood that depends on just one compression vector and one compressed statistic per parameter of interest, and which is equivalent to the Plik-lite temperature likelihood with two low- ℓ bins included. Because we do not directly incorporate the low- ℓ polarization likelihood, we recommend including a prior on the optical depth to reionization τ when using either of these likelihoods.

We provide two public codes. The first is Planck-lite-py, an implementation of *Planck*'s Plik_lite likelihood in Python, with the option to include the low- ℓ temperature data as two Gaussian bins. The second is CosMOPED, which calculates the MOPED compression vectors for the CMB temperature power spectrum and computes the compressed *Planck* likelihood.

This method can easily be extended to incorporate the high- ℓ TE and EE data which are also Gaussian distributed. It can also be used for the *Planck* 2018 likelihood, and the publicly available code will be updated accordingly.

The MOPED data compression scheme provides a lightweight likelihood that can easily be combined with other datasets. In addition, the compressed data can be incorporated into a likelihood free inference framework which allows parameters to be inferred based on forward simulations, without knowledge of the form of the likelihood. In likelihood free inference it is useful to have informative compressed statistics, because this makes the comparison of simulations and data much less computationally intensive.

Chapter 3

A binned Python likelihood for large scale temperature and polarization from *Planck*

The current best fit model of the universe is Λ CDM , which provides a remarkable fit to a variety of cosmological data with only six cosmological parameters. The cosmic microwave background (CMB) radiation from the early universe provides strong constraints on these parameters, and the CMB temperature and polarization have been measured over the full sky by the *Planck* satellite, with its legacy data released in 2018 ([Planck Collaboration et al., 2020b,c](#)). Various current and planned CMB experiments are focused on measuring the CMB from the ground, improving on *Planck* in resolution and instrumental noise. However, because of atmospheric noise and limited sky coverage, it is challenging to measure large angular scale CMB fluctuations from the ground. Both galaxy surveys and ground-based CMB experiments will thus continue to benefit in the coming decade from including the *Planck* data in constraints, especially on large scales.

The *Planck* legacy release includes likelihood functions describing the temperature and E -mode polarization data on large scales, where the probability distribution of the angular power spectrum is non-Gaussian (Planck Collaboration et al., 2020b). While the evaluation time of these likelihoods for the *Planck* low- ℓ data is not a limiting factor compared to the time needed, for example, to run Boltzmann codes, the existing package includes code in both C and Fortran 90. To facilitate the ease of use of the data, for both analysis and forecasting applications, we thus present an alternative option that compresses the data to a set of independent log-normal bins. This form of the *Planck* likelihood can be automatically differentiated in Python, so it could be used in combination with future differentiable Boltzmann codes to sample parameters using efficient Hamiltonian Monte Carlo methods, such as the No U-Turn Sampler (NUTS, Hoffman & Gelman, 2011), requiring fewer likelihood evaluations and therefore fewer iterations of the theory code than traditional Markov Chain Monte Carlo sampling methods.

This log-normal compression extends work presented in Prince & Dunkley (2019) which included a compression of the temperature spectrum. The low- ℓ (large scale) E -mode polarization primarily constrains the optical depth to reionization, τ , and the low- ℓ temperature provides a lever arm to constrain the tilt of the primordial power spectrum, n_s . For most models of interest, one therefore only needs a few numbers to describe the *Planck* low- ℓ data (e.g., Heavens et al., 2000; Alsing & Wandelt, 2018). Here we use two bins to describe the low- ℓ temperature power spectrum and seven bins for E -mode polarization, leading to a simple likelihood code written in Python that is portable and easy to combine with other cosmological data. The method could also be applied to the alternative *Planck* data processing from the NPIPE maps described in Planck Collaboration (2020). Our code, `Planck-low-py`, is publicly available on GitHub¹. It can be used in combination with `Planck-lite-py`²,

¹<https://github.com/heatherprince/planck-low-py>

²<https://github.com/heatherprince/planck-lite-py>

our Python implementation of the *Planck* team’s `Plik_lite` foreground-marginalized likelihood for the high- ℓ data. Use of either code should reference the [Planck Collaboration et al. \(2020b\)](#) data. Other Python implementations of *Planck* likelihoods, including `Plik_lite`, are also already available with the public cosmological sampling code `cobaya`³.

We describe the *Planck* data and our low- ℓ binned likelihoods in section 3.1. In section 3.2 we compare Λ CDM parameter constraints from `Planck-low-py` with those from the legacy *Planck* likelihoods; extended models were discussed in [Prince & Dunkley \(2019\)](#). We conclude in section 3.3.

3.1 The *Planck* likelihood and low- ℓ binning

The likelihood function used for *Planck*’s cosmological analysis, describing the probability of the data given some model, is separated into two regimes, with different approaches for large and small angular scales ([Planck Collaboration et al., 2020b](#)).

At $\ell \geq 30$ (corresponding to scales smaller than several degrees on the sky) the likelihood \mathcal{L} for the temperature and E -mode polarization power spectra and cross spectrum (TT , EE and TE) is modeled as a Gaussian distribution, with

$$-2 \ln \mathcal{L} = (C_b^{\text{th}} - C_b^{\text{data}})^T Q^{-1} (C_b^{\text{th}} - C_b^{\text{data}}), \quad (3.1)$$

to within an overall additive constant, with binned data C_b^{data} , binned theory C_b^{th} , and binned covariance matrix Q . For the `Plik_lite` likelihood ([Planck Collaboration et al., 2020b](#)), these data spectra represent an estimate of the CMB bandpowers, with foregrounds already marginalized over. The bandpowers for these binned high- ℓ data are shown in Fig. 3.1, as the black markers in the right panels. The multi-frequency

³https://cobaya.readthedocs.io/en/latest/likelihood_planck.html




Figure 3.1: The *Planck* 2018 temperature and polarization power spectra and best-fit theory curves, from [Planck Collaboration et al. \(2020b\)](#). The log-normal binned $\ell < 30$ *TT* and *EE* spectra estimated in this paper, described in section 3.1, are shown on the left of the upper panel. They compress the per- ℓ likelihoods (grey markers) described by the *Planck Commander* and *SimAll* likelihoods respectively. The Gaussian, foreground-marginalized $\ell \geq 30$ data (black markers on the right) are used in the *Plik_lite* cosmological likelihood and are shown for reference. The low- ℓ *TE* cross spectrum is not used in the *Planck* cosmological analysis.

high- ℓ *Plik* likelihood does not pre-marginalize over foregrounds, and gives results in agreement with *Plik_lite*.

At $\ell < 30$ (large angular scales on the sky) the distribution of the angular power spectrum is non-Gaussian (see, e.g., [Bond et al., 2000](#)). The *Planck* low- ℓ temperature likelihood is derived using the *Commander* framework ([Eriksen et al., 2008](#)), which uses Gibbs sampling to explore the joint distribution of the CMB temperature map, CMB temperature angular power spectrum, and foreground parameters, and then uses a Gaussianized Blackwell-Rau estimator to describe the likelihood of the modeled temperature angular power spectrum given the data.

The *Planck* low- ℓ polarization uses a separate `SimAll` likelihood built from simulations (Planck Collaboration et al., 2020b). This likelihood uses the E -mode polarization angular cross spectrum between the 100 GHz and 143 GHz channels, computed using the quadratic maximum likelihood (QML) approach (Tegmark & de Oliveira-Costa, 2001; Efstathiou, 2006), with templates for synchrotron and dust contamination used to remove the foregrounds. The likelihood is then constructed from this data power spectrum using a suite of simulations.

The curl-like polarization B -modes are not used for the baseline *Planck* cosmological analysis. The large scale temperature-polarization correlation (captured in the low- ℓ TE cross spectrum) is also excluded due to performing poorly in null tests, indicating that some systematic or foreground effects remain unaccounted for. In this paper we thus focus just on compressing the $\ell < 30$ temperature and E -mode polarization likelihoods.

3.1.1 Low- ℓ temperature bins

The low- ℓ temperature power spectrum, shown in Fig. 3.2, is approximately flat due to the Sachs-Wolfe effect (Sachs & Wolfe, 1967b). In the Λ CDM model, which is an excellent fit to *Planck* data, the low- ℓ temperature power provides a large-scale anchor for constraining the spectral index of the primordial fluctuation power spectrum, n_s , which tilts the resulting temperature power spectrum. Within simple models, including Λ CDM, the lack of complicated structure on these large scales suggests that we could bin the low- ℓ data without losing information about the cosmological parameters. In Prince & Dunkley (2019) we demonstrated that compressing to two bins was sufficient to reproduce parameters for the *Planck* 2015 temperature data. There the compression to two Gaussian bins was useful for the data compression technique applied to the full ℓ range of the *Planck* power spectrum.




Figure 3.2: The low- ℓ *Planck* temperature power spectrum $D_\ell^{TT} = \ell(\ell+1)C_\ell^{TT}/2\pi$. The binned power and errors are shown in black. The *Planck* unbinned power spectrum and errors are in grey. The errorbars are asymmetric because on these scales the power spectrum is non-Gaussian. The value of the power spectrum is taken from the peak of the probability distribution and the errorbars come from the half-maximum values. The theory curve for the best fit Λ CDM model is shown in red. The power spectrum is close to flat on these large angular scales due to the Sachs-Wolfe effect.




Figure 3.3: The peak normalized posterior distribution for $D_\ell^{TT} = \ell(\ell+1)C_\ell^{TT}/2\pi$ for two low- ℓ temperature bins ($\ell = 2-15$ and $16-29$) using the *Planck* 2018 Commander low- ℓ temperature likelihood. The best-fitting log-normal distributions are shown in red.

Here we refine this compression using a log-normal distribution to approximate the likelihood, motivated by e.g., [Bond et al. \(2000\)](#). We find that the log-normal bins give a slightly better fit to the shape of the posterior of the power in each bin, and is also consistent with how we treat the low- ℓ polarization data in Section 3.1.2.

As in [Prince & Dunkley \(2019\)](#) we form two temperature bins, one for $2 \leq \ell \leq 15$ and one for $16 \leq \ell \leq 29$. We find the probability distribution for the power spectrum $D_\ell = \ell(\ell+1)C_\ell/2\pi$ by conditionally sampling the posterior distribution for the power in each bin using the *Planck 2018 Commander* likelihood using the *Cobaya* cosmological sampling code ([Torrado & Lewis, 2020](#)), estimating

$$p(\theta|d) \propto p(d|\theta)p(\theta). \quad (3.2)$$

Here the parameters θ are the binned values $D_{2 \leq \ell \leq 15}$ and $D_{16 \leq \ell \leq 29}$, assuming a constant value for D in each bin. We assume uniform priors on θ . In practice we vary one bin at a time (giving every ℓ value in that bin the same value and then varying that value) while holding the other bin constant, and use the full Planck low- ℓ likelihood to find the shape of that bin, so we use the likelihood to sample D_{bin} instead the parameters.

The probability distributions for the two low- ℓ power spectrum bins are shown in Fig. 3.3, together with the best-fit log-normal probability distributions. A log-normal distribution for the likelihood of D_{bin} means that $\ln(D_{\text{bin}})$ is close to normally distributed, or

$$\mathcal{L}(x) = p(x) = \frac{1}{x\sigma\sqrt{2\pi}} e^{-(\ln x - \mu)^2/(2\sigma^2)}, \quad (3.3)$$

for $x = D_{\text{bin}}$. Although there is some motivation for using a combination of log-normal and Gaussian distributions to describe the probability distribution of the unbinned angular power spectra ([Verde et al., 2003](#)), we find that the log-normal fit to the bins gives an acceptable fit and accurately reproduces parameter constraints. The

best-fitting parameters for the two bins are

$$\begin{aligned}\mu_1 &= 6.717, & \sigma_1 &= 0.09247, \\ \mu_2 &= 6.734, & \sigma_2 &= 0.06038.\end{aligned}\tag{3.4}$$

The covariance between bins is small, so we treat the likelihood of each bin independently. To compute the likelihood for a given theory power spectrum, we convert the theory C_ℓ^{th} to D_ℓ^{th} , bin D_ℓ^{th} into two bins, D_1^{TT} and D_2^{TT} , compute the log-normal likelihood for each bin, and then multiply them together (corresponding to adding the log-likelihoods), giving

$$\begin{aligned}\ln \mathcal{L}^{TT} &= \ln \mathcal{L}_1^{TT} + \ln \mathcal{L}_2^{TT} \\ &= \ln \left(\frac{1}{D_1^{TT} \sigma_1 \sqrt{2\pi}} e^{-(\ln D_1^{TT} - \mu_1)^2 / (2\sigma_1^2)} \right) + \\ &\quad \ln \left(\frac{1}{D_2^{TT} \sigma_2 \sqrt{2\pi}} e^{-(\ln D_2^{TT} - \mu_2)^2 / (2\sigma_2^2)} \right).\end{aligned}\tag{3.5}$$

This compression approximates the cosmic variance contribution to the uncertainty on the power spectrum as being independent of the theory, as in the $\ell > 30$ *Planck* likelihood.

The bandpowers for these two log-normal bins are indicated in Fig. 3.2 (the large black points), together with the unbinned low- ℓ power spectrum. For plotting purposes the D_{bin} at the peak of the distribution (the mode) is shown, which for a log-normal distribution is at $D_{\text{bin}} = e^{\mu - \sigma^2}$. We plot errors using the D values at which the probability drops to 0.61 of its maximum, which would be at 1σ for a Gaussian distribution, showing

$$\begin{aligned}D_{2 \leq \ell \leq 15}^{TT} &= 819_{-72}^{+79} \mu\text{K}^2, \\ D_{16 \leq \ell \leq 29}^{TT} &= 837_{-49}^{+52} \mu\text{K}^2.\end{aligned}\tag{3.6}$$




Figure 3.4: The low- ℓ E-mode polarization power spectrum, with the D^{EE} estimated in seven bins shown in black ($(\ell = 2-3, 4, 5, 6-7, 8-11$ and $10-29$). The *Planck* per- ℓ power spectrum is shown in grey. The values and errors for these non-Gaussian data points come from the peak and half-maximum of the conditional posterior distribution for each D_ℓ . The theory curve for the best fit model (red) and for high and low values of the optical depth to reionization τ are shown. The reionization bump in the power spectrum comes from polarization caused by Thomson scattering off free electrons since reionization. This signal is higher for greater τ .

3.1.2 Low- ℓ E-mode polarization

The CMB E -mode polarization signal is created by Thomson scattering of CMB photons off electrons, both at the last scattering surface and during reionization when electrons are once again free (Bond & Efstathiou, 1984; Zaldarriaga, 1997). Thomson scattering during reionization creates large-scale E -mode polarization, the amplitude of which depends primarily on the optical depth to reionization, τ . This ‘reionization bump’ can be seen in Fig. 3.4, which shows the Λ CDM theoretical EE power spectrum for a few different values of τ , including the best-fit $\tau = 0.054$ from Planck Collaboration et al. (2020c). As the optical depth to reionization increases, so does the low- ℓ EE power.




Figure 3.5: The peak normalized posterior distribution for $D_\ell^{EE} = \ell(\ell+1)C_\ell^{EE}/2\pi$ for seven low- ℓ E-mode polarization bins, estimated using the *Planck 2018 SimAll* low- ℓ polarization likelihood. The best-fit offset log-normal distribution is shown in red for each bin. The ℓ ranges of the bins are chosen to capture the low- ℓ polarization features imprinted from varying the optical depth to reionization, τ .

The *Planck* 2018 low- ℓ polarization likelihood was built using simulations to determine the probability distributions for the EE and BB spectra.

Here we compress the low- ℓ EE -mode polarization data into seven bins, three that primarily constrains the height of the reionization bump using $2 \leq \ell \leq 5$, two that constrain the width of the bump using $6 \leq \ell \leq 11$, and one wider bin for $11 \leq \ell \leq 29$ where the EE power spectrum has less structure. We do not compress the B -mode polarization data, as it is not used in the main *Planck* likelihood combination for cosmology constraints.

The probability distribution for the power in each bin is shown in Fig. 3.5. We find that the distributions are sufficiently well described by an offset log-normal distribution, which modifies equation (3.3) for the probability distribution of $x = D_{\text{bin}}^{EE}$ to

$$\mathcal{L}(x) = p(x) = \frac{1}{(x - x_0)\sigma\sqrt{2\pi}} e^{-(\ln(x-x_0)-\mu)^2/(2\sigma^2)}. \quad (3.7)$$

We experiment with other distributions, including a combination of a log-normal and a Gaussian probability distribution [Verde et al. \(2003\)](#)

$$\ln \mathcal{L} = \frac{2}{3} \ln \mathcal{L}_{LN} + \frac{1}{3} \ln \mathcal{L}_G, \quad (3.8)$$

as well as a spline, to capture the exact shape of the likelihood, but find that the log-normal distribution is sufficient to capture the sensitivity to cosmological parameters. The best-fit parameters for the log-normal bins are shown in table 3.1.

We neglect the covariance between bins just as we do for the temperature, so the log-normal likelihoods from the seven EE bins are combined independently to give the compressed low- ℓ E -mode polarization likelihood:

$$\ln \mathcal{L}^{EE} = \sum_{i=1}^7 \ln \mathcal{L}_i^{EE}. \quad (3.9)$$




Figure 3.6: Posterior distributions of the Λ CDM parameters estimated using the low- ℓ *Planck* 2018 temperature (Commander) and EE (SimA11) likelihoods (black), compared to those using *Planck-low-py* (our log-normal compressed low-ell likelihood) (red dashed). In both cases *Plik_lite* is used at $\ell > 30$. The parameters shown are the optical depth to reionization τ , the amplitude of the primordial power spectrum A_s , the scalar spectral index n_s , the Hubble constant H_0 in units of $\text{km s}^{-1} \text{Mpc}^{-1}$, the baryon density $\Omega_b h^2$ and the dark matter density $\Omega_c h^2$. The parameter constraints agree to within less than 0.1σ .




Figure 3.7: Posterior distributions for extensions to Λ CDM. Left: constraint on the running of the spectral index. Right: constraint on the amplitude of the neutrino density isocurvature power spectrum at $k = 0.1 \text{ Mpc}^{-1}$ assuming that the isocurvature power spectrum is a power law and is uncorrelated with the adiabatic power spectrum (using the parameterization from [Planck Collaboration et al. \(2020e\)](#) with $P_{\mathcal{R}\mathcal{I}} = 0$). The black curves show the constraints using the full *Planck* 2018 low- ℓ likelihoods and the red dashed curves show the constraints using the binned likelihoods described in this paper.

ℓ range	μ^{EE}	σ^{EE}	$x_0^{EE} [\mu K^2]$	$D^{EE} [\mu K^2]$
2-3	-2.190	0.5615	-0.02244	$0.059^{+0.062}_{-0.035}$
4	-3.445	0.8534	-0.00789	$0.008^{+0.021}_{-0.009}$
5	-2.781	0.7335	0.00121	$0.037^{+0.039}_{-0.019}$
6-7	-4.076	0.5327	-0.01355	≤ 0.008
8-11	-4.393	0.4075	-0.01330	≤ 0.003
12-15	-5.875	0.9196	-0.00203	≤ 0.001
16-29	-3.094	0.1161	-0.03623	$0.008^{+0.005}_{-0.005}$

Table 3.1: Parameters for the log-normal EE bins, and estimated D^{EE} power.

We also neglect any covariance between the low- ℓ TT and EE likelihoods, combining them independently. These best-fitting parameters are used to plot the mode and errors and the upper limits in Fig. 3.4, also reported in Table 3.1.

The *Planck* calibration parameter is included in the low- ℓ likelihoods in the same way as in `Plik_lite`, as an overall normalizing factor applied to the theory power spectrum.

3.1.3 Software products

We release a public Python likelihood code, `Planck-low-py`,⁴ which uses these independent log-normal likelihoods to describe the *Planck* low- ℓ temperature and polarization bins. The temperature and polarization likelihoods can be used separately or together. This likelihood code provides an alternative to the full *Planck* low- ℓ likelihood functions, for models without unusually complex large-scale behavior. Any use of it should reference the [Planck Collaboration et al. \(2020b\)](#) data. The code can be used in combination with any $\ell > 30$ *Planck* likelihoods, including our Python implementation of `Plik_lite`, `planck-lite-py`⁵, or the Python implementations of *Planck*'s likelihoods in `Cobaya`.

⁴<https://github.com/heatherprince/planck-low-py>

⁵<https://github.com/heatherprince/planck-lite-py>

3.2 Parameter constraints

We test the compressed `Planck-low-py` by comparing the constraints on Λ CDM parameters with those from the full *Planck* 2018 low- ℓ *TT* and *EE* likelihoods, using the same high- ℓ `Plik_lite` likelihood in both cases. The results are plotted in Fig. 3.6, which shows the posterior probabilities for the six Λ CDM parameters (the Hubble constant, baryon density, cold dark matter density, amplitude and spectral index of primordial fluctuations, and optical depth to reionization). These parameter constraints were obtained using `Cobaya` (Torrado & Lewis, 2020)⁶ with the `CAMB` cosmological Boltzmann code (Lewis et al., 2000; Howlett et al., 2012) and the MCMC sampler developed for `CosmoMC` (Lewis & Bridle, 2002; Lewis, 2013).

The reference constraints (black) are obtained from the high- ℓ temperature and polarization foreground-marginalized `Plik_lite` likelihood, the low- ℓ *EE* likelihood (`SimAll`), and the low- ℓ temperature-only likelihood (`Commander`) of *Planck*'s 2018 data release (Planck Collaboration et al., 2020b). The red dashed lines show our constraints from combining `Plik_lite` with our `Planck-low-py` compressed likelihood. We find that all of the Λ CDM cosmological parameters are consistent to within 0.1σ . We also tested two extensions to the standard Λ CDM model. Constraints on the running of the spectral index and on the amplitude of a neutrino density isocurvature component are shown in Fig. 3.7; parameters estimated from the full low- ℓ likelihoods and our likelihood agree.

We find that using just one or two lognormal bins to describe the low- ℓ *EE* likelihood is insufficient to reproduce the parameter constraints. Using three bins is better able to capture the structure of the reionization bump (see Fig. 3.4), and reproduces constraints with accurate values but slightly inflated errors on the optical depth, as well as on A_s through its correlation with τ . Using seven bins accurately reproduces both the means and widths of the probability distributions of the parameters. The

⁶<https://ascl.net/1910.019>

`Planck-low-py` E -mode likelihood can also be used as an alternative to a τ prior, with the advantage of directly using the amplitude of the EE power spectrum. For the equivalent of imposing a broader prior on τ , as done in [Aiola et al. \(2020\)](#) for example, one could inflate the error bars of the EE bins.

3.3 Discussion

We have demonstrated that the *Planck* 2018 low- ℓ temperature and E -mode polarization data can be effectively compressed to two and seven log-normal bins respectively while retaining accurate constraints on Λ CDM cosmological parameters. This compression would also be appropriate for non- Λ CDM models which do not have unusual large-scale features. In our explorations of the minimal bin number required to capture the information, we found that as few as three bins in polarization was sufficient to reproduce the means of the parameters to 0.1σ , but with slightly broader error bars.

We provide a public Python likelihood code, `Planck-low-py` to describe the *Planck* low- ℓ data which can be used as a light-weight alternative to the `Commander` and `SimAll` likelihoods.

The *Planck* data will provide the community with the tightest constraint on the optical depth to reionization τ , and the leading large-scale temperature measurements, for some time. Light-weight versions of the *Planck* low- ℓ likelihoods can thus be useful for future explorations of combined datasets, as well as for forecasting purposes in the design of upcoming experiments.

Chapter 4

Testing the BICEP/Keck foreground model, and a foreground-marginalized ‘BK18-lite’ likelihood

The theory of inflation predicts that tensor modes would be present in the early universe in addition to the scalar modes that form the seeds of large scale structure and source the primary CMB temperature and E-mode polarization. These primordial tensor modes propagate as primordial gravitational waves that would produce a signal in both E and B polarization ([Seljak & Zaldarriaga, 1997](#); [Kamionkowski & Kovetz, 2016](#)). Because scalar modes predict only primary E-mode polarization, B-modes are a cleaner way to search for primordial tensor perturbations. Any primordial signal would be expected to come from gravitational waves, with B-mode power peaking at $\ell \sim 100$ (degree scales) and a very large-scale bump from scattering of the primordial signal at reionization at $\ell \sim 5$. The B-mode signal also includes gravitationally lensed E-modes, peaking at $\ell \sim 1000$ (a few arcminutes in separation on the sky). In

addition to these cosmological B-mode sources, polarized radiation from our Galaxy also produces B-mode radiation at millimeter wavelengths (e.g., [Planck Collaboration et al., 2020f](#)). Some of this foreground radiation comes from thermal emission from dust grains, which is polarized because the dust grains tend to be aligned with their long axes perpendicular to the Galactic magnetic field. Another source of B-mode polarization at these wavelengths is synchrotron emission from charged particles that are accelerated by magnetic fields in the Galaxy.

The BICEP (Background Imaging of Cosmic Extragalactic Polarization) and Keck telescopes are being used to search for primordial B-modes from the South Pole Station in Antarctica. Data from the BICEP2 experiment showed an excess of B-modes at 150 GHz over those expected from lensing ([BICEP2 Collaboration et al., 2014b](#)). A joint analysis with *Planck* multifrequency data subsequently showed that the observed excess B-modes were from Galactic dust, and placed an upper limit on the tensor-to-scalar ratio r ([BICEP2/Keck Collaboration et al., 2015](#)). The tightest constraint on primordial gravitational waves to date comes from BICEP3 observations, together with data from BICEP2, the Keck Array, *Planck* and *WMAP*, giving a 95% confidence upper limit on the tensor-to-scalar ratio r at 0.05 Mpc^{-1} of $r < 0.036$ ([Ade et al. \(2021\)](#), hereafter BK18). Independent of BK18, there are also 95% upper limits of $r < 0.16$ from *Planck* B-modes, and $r < 0.056$ from *Planck* T, E and B-modes, derived from the NPIPE maps ([Tristram et al., 2021](#)). The first flight of the SPIDER balloon experiment gave $r < 0.11$ ([SPIDER Collaboration et al., 2021](#)).

These current constraints on r are informed by the treatment of foregrounds, especially polarized dust emission. For the BK18 analysis, different foreground models were tested by the BICEP/Keck collaboration, demonstrating the stability of the estimated for r . In this chapter we explore further some of the assumptions made in the BK18 foreground model, finding little impact on r . To constrain the foregrounds, eleven polarization maps at a range of frequencies from BK, *WMAP* and

Planck are combined to give a total of 11 auto-spectra and 55 cross-spectra, which are jointly analysed in the nominal BK18 likelihood to simultaneously constrain the seven parameter foreground model and the tensor-to-scalar ratio r . In this chapter we construct a ‘BK18-lite’ likelihood based on foreground-marginalized bandpowers for the B-mode angular power spectrum, following a similar approach adopted for ACT and *Planck* (Dunkley et al., 2013; Calabrese et al., 2013; Planck Collaboration et al., 2016a), which appears sufficient to reproduce the tensor-to-scalar ratio constraints.

This chapter is laid out as follows. We describe the BICEP/Keck data and likelihood in §4.1. §4.2 explores some of the modeling choices and how the foreground model affects the gravitational wave constraints. In §4.3 we estimate foreground-marginalized CMB B-mode bandpower amplitudes and construct a ‘BK-lite’ likelihood. We conclude in §4.4.

4.1 The BICEP/Keck data and likelihood

In this section we draw on Barkats et al. (2014); BICEP2 Collaboration et al. (2014a); Ade et al. (2021) to review the details of the BK18 likelihood analysis, including the data products, foreground model, and form of the likelihood.

4.1.1 Data and Window Functions

BICEP2/Keck: BICEP2 mapped the ~ 400 square degree BICEP2/Keck region of the sky first at 150 GHz from 2010–2012 (BICEP2 Collaboration et al., 2014a). The Keck Array mapped the same region at 150 GHz and then at 95 and 220 GHz from 2012–2019 (BICEP2 and Keck Array Collaborations et al., 2015; BICEP2 Collaboration et al., 2018). For the previous joint analysis of BICEP2, Keck Array, and *Planck* (BK15), BICEP2 Collaboration et al. (2018)) the apodization mask on the sky was the geometric mean of the BICEP/Keck inverse noise variance maps, and the same

mask was applied to *WMAP* and *Planck* data. The most recent BK18 analysis keeps the BICEP2/Keck apodization masks for the BICEP2 and Keck maps, but uses the BICEP3 mask for the BICEP3, *WMAP* and *Planck* data (Ade et al., 2021). The BICEP2/Keck observations contribute three polarization maps to the BK18 analysis, one each at 95, 150 and 220 GHz. At 150 and 220 GHz the B-mode power spectra of the maps are dominated by polarized Galactic dust emission, and are used to constrain and marginalize over dust foreground parameters.

BICEP3: BICEP3 measured the CMB sky at 95 GHz, and its 3-year dataset of observations from 2016–2018 were used to make the deepest CMB polarization map at 95 GHz so far, with a depth of $2.8 \mu\text{K}\text{-arcmin}$ (Ade et al., 2022). The BICEP3 scan region consists of ~ 600 square degrees and includes the BICEP2/Keck region (Ade et al., 2021). The 95 GHz channel is the least contaminated by foregrounds out of the BICEP/Keck 95, 150 and 220 GHz channels.

WMAP: The likelihood includes spectra computed from 23 and 33 GHz maps from the *WMAP* satellite, masked using the BICEP3 inverse noise variance. These maps are used to constrain synchrotron radiation at low frequencies.

Planck: Spectra are used from the 30, 44, 143, 217 and 353 GHz polarization maps from the *Planck* satellite, masked using the BICEP3 inverse noise variance. These maps are used to constrain synchrotron radiation at low frequencies (30 and 44 GHz) and polarized Galactic dust emission at high frequencies (143–343 GHz).

In total there are 11 polarization maps included in the analysis: three 400 deg^2 maps at different frequencies from BICEP2/Keck; one 600 deg^2 BICEP3 map at 95 GHz; *WMAP* maps at two frequencies and *Planck* maps at five frequencies, each covering the 600 deg^2 BICEP3 footprint. This results in 11 auto-spectra and 55 cross-spectra for BB. Each spectrum is composed of 9 bandpowers, resulting in $9 \times 66 = 594$ bandpowers, which are fit simultaneously in the likelihood analysis.




Figure 4.1: BK18 bandpower window functions for the 9 bins. They differ for spectra estimated in the larger BICEP3 region (black) versus the smaller BICEP2/Keck region (red). Cross-spectra between the two regions have a third typical window function shape (orange). The *Planck* and *WMAP* maps are cut to the BICEP3 region of the sky and have the same window functions as BICEP3. The BICEP2 and Keck Array data for 95, 150 and 220 GHz have almost identical window function shapes, but different normalizations. The suite of cross-spectra window functions are also similar in shape but not identical, and have different normalizations.

Window functions: The window functions that are used to bin a theory curve to compare to the data power spectra, with $D_b = \sum_\ell w_{b\ell} D_\ell$, include the effect of transforming the apodization mask into Fourier space, as the mask affects how different modes are coupled (Knox, 1999; Barkats et al., 2014). Thus there are effectively three distinct sets of window function shapes: one for the larger BICEP3 field which applies to the BICEP3, *WMAP* and *Planck* auto-spectra, one for the smaller BICEP2/Keck field which applies to the BICEP2/Keck auto-spectra, and one for cross spectra between the two different fields. These window functions are shown in Figure 4.1. The BICEP2/Keck window functions are shifted to slightly higher ℓ than the BICEP3 window functions, and have somewhat different shapes. The BICEP2/Keck window

functions also each have different normalizations, although their shapes are almost the same, at the three different frequencies. The cross spectra between BICEP3 regions and BICEP2/Keck regions also have different normalizations depending on which BICEP2/Keck frequency is used¹.

4.1.2 Foreground model

The foreground model assumes two contaminating components in polarization: thermal dust and synchrotron emission from the Galaxy, such that $D_\ell^{ij} = D_\ell^{\text{CMB}} + D_\ell^{\text{FG},ij}$, and

$$D_\ell^{\text{FG},ij} = D_\ell^{\text{dust},ij} + D_\ell^{\text{sync},ij} + D_\ell^{\text{dust-sync},ij} \quad (4.1)$$

for cross-spectrum between maps i and j . The model assumes that there is no polarization signal from anomalous microwave or free-free emission. Extragalactic contamination is expected to be negligible at these angular scales. In brightness temperature units, the model terms are given by the following; these are converted to thermodynamic units by integrating across the passbands.

Dust: Polarized galactic dust is assumed to follow a modified blackbody spectrum with temperature $T_d = 19.6\text{K}$. The dust has three parameters: the amplitude $A_{d,353}$ at 353 GHz at $\ell = 80$, the spectral index in ℓ -space α_d , and the spectral index of the modified blackbody function in frequency space β_d . In this model these spectral indices are assumed to be constant over the sky region. The cross-power is given by

$$D_\ell^{\text{dust},ij} = A_d \left(\frac{\ell}{\ell_0} \right)^{\alpha_d} \left[\frac{\nu_i^{\beta_d} B_{\nu_i}(T_d)}{\nu_0^{\beta_d} B_{\nu_0}(T_d)} \right] \left[\frac{\nu_j^{\beta_d} B_{\nu_j}(T_d)}{\nu_0^{\beta_d} B_{\nu_0}(T_d)} \right] \quad (4.2)$$

with emissivity index β . The function $B_\nu(T_d)$ is the Planck function at frequency ν for effective dust temperature T_d .

¹In BK15 the window functions were normalized to unity; the BK18 window functions are not normalized.

Synchrotron: Polarized synchrotron emission is assumed to follow a power law in frequency and multipole space. Synchrotron emission is parameterized by its amplitude $A_{\text{sync},23}$ at 23 GHz at $\ell = 80$, its spectral index in ℓ -space α_s and its spectral index in frequency space β_s

$$D_\ell^{\text{sync},ij} = A_s \left(\frac{\ell}{\ell_0} \right)^{\alpha_s} \left(\frac{\nu_i \nu_j}{\nu_0^2} \right)^{\beta_s}. \quad (4.3)$$

Dust/synchrotron correlation: an additional foreground parameter ϵ captures any correlation between the synchrotron and dust emission, with the cross-spectrum given by

$$D_\ell^{\text{dust-sync},ij} = \epsilon \sqrt{A_d A_s} \left(\frac{\ell}{\ell_0} \right)^{\sqrt{\alpha_s \alpha_d}} \left[\frac{\nu_i^{\beta_d} B_{\nu i}(T_d) \nu_j^{\beta_s}}{\nu_0^{\beta_d} B_{\nu 0}(T_d) \nu_0^{\beta_s}} + \frac{\nu_j^{\beta_d} B_{\nu j}(T_d) \nu_i^{\beta_s}}{\nu_0^{\beta_d} B_{\nu 0}(T_d) \nu_0^{\beta_s}} \right]. \quad (4.4)$$

4.1.3 Likelihood

The BK18 analysis computes auto- and cross-spectra for all the frequency pairs from the *WMAP*, *Planck*, BICEP and Keck Array data ([Hamimeche & Lewis, 2008](#)), and uses the Hamimeche-Lewis approximation to model the likelihood. Each BB spectrum consist of 9 bins ranging from multipoles $\ell = 20$ to 330. The Hamimeche-Lewis likelihood uses a pre-computed covariance matrix calculated from simulations using a fiducial model, although the parameter constraints do not depend strongly on small variations in the choice of fiducial model.

The likelihood for the model bandpowers, \mathcal{D}_b , given the data bandpowers, $\hat{\mathcal{D}}_b$, in the Hamimeche-Lewis approximation is given by

$$-2 \log \mathcal{L}(\mathcal{D}_b | \hat{\mathcal{D}}_b) = X_c \mathcal{M}_{cc'} X_{c'}, \quad (4.5)$$

where the indices c and c' run over all bins in ℓ for all temperature and polarization combinations (i.e.m TT, EE, BB, TE, TB, EB) ([Barkats et al., 2014](#)). The indices

b run over the 9 bins used and the different frequency auto- and cross-combinations. For the baseline BK18 analysis only the BB spectra are used, so the indices c are equivalent to b .

X_b is a vector of transformed bandpowers:

$$X_b^{BB} = (\mathcal{D}_b^f)^{1/2} U_b g(D_b) U_b^\dagger (\mathcal{D}_b^f)^{1/2}, \quad (4.6)$$

where \mathcal{D}_b^f are fiducial bandpowers calculated from simulations of the signal and noise using the fiducial Λ CDM parameter values. U_b is a matrix of the eigenvectors of $\mathcal{D}_b^{-1/2} \hat{\mathcal{D}}_b \mathcal{D}_b^{-1/2}$, and D_b is a diagonal matrix of the eigenvalues of $\mathcal{D}_b^{-1/2} \hat{\mathcal{D}}_b \mathcal{D}_b^{-1/2}$. The function g , given by

$$g(x) = \text{sign}(x - 1) \sqrt{2(x - \ln x - 1)}, \quad (4.7)$$

is applied to each element of the diagonal matrix D_b to give $g(D_b)$. The bandpower covariance $M_{cc'}$ is calculated from the same simulations as \mathcal{D}_b^f , and then terms are added to account for the gain (G_c) and beam width (S_c) calibration uncertainties:

$$\mathcal{M}_{cc'} = M_{cc'} + G_c G_{c'} \hat{\mathcal{D}}_c \hat{\mathcal{D}}_{c'} + S_c S_{c'} \hat{\mathcal{D}}_c \hat{\mathcal{D}}_{c'}. \quad (4.8)$$

The BICEP/Keck team use 499 signal and noise simulations to calculate the fiducial covariance matrix and bandpowers (Ade et al., 2021). The transformed bandpowers X_c and the modified covariance matrix $\mathcal{M}_{cc'}$ are then used to calculate what is essentially a χ^2 statistic for the log-likelihood in Equation (4.5).

In the likelihood analysis the data (9 bins for each of 66 auto- and cross-spectra for BB) are compared to a model BB spectrum containing primordial B-modes from tensor perturbations with a varying tensor-to-scalar ratio r , lensed B-modes from gravitational lensing of the primordial Λ CDM E-modes, dust, and synchrotron. The lensed

B-modes are calculated using the *Planck* best-fit Λ CDM parameters (Planck Collaboration et al., 2020c). The tensor-to-scalar ratio is measured at $k = 0.05 \text{ Mpc}^{-1}$, and the tensor spectral index is set to $n_t = 0$, which corresponds to a scale free spectrum (the tensor and scalar spectral indices have different definitions; $n_s = 1$ for scale-invariant scalars).




Figure 4.2: Auto-spectra for the 95 and 220 GHz Keck array data (grey dots and purple squares respectively), 150 GHz BICEP2/Keck data (blue triangles), and 95 GHz BICEP3 data (black circles). Also shown is the B-mode power spectrum for a Λ CDM cosmology with $r=0$ (solid black curve, just lensing B-modes) and $r=0.03$ (dashed black curve). The 220 GHz channel is above the theory curves because it is dominated by dust, and the 150 GHz and 95 GHz channels have significant dust contamination on larger scales ($\ell \lesssim 200$). Dust can cause an excess of power on degree scales, which can be confused with primordial B-modes unless the foregrounds are constrained using a range of frequency channels. The BICEP3 95 GHz spectrum provides the most constraining power for the CMB B-mode power spectrum because it has both high sensitivity (small error bars) and relatively low dust contamination.

The BICEP3 95 GHz and the BICEP2/Keck 95, 150 and 220 GHz binned power spectra are shown in Figure 4.2, together with theory curves for the B-mode power spectrum from lensing in the absence of primordial gravitational waves ($r = 0$) and

the B-mode power spectrum for an example model with $r = 0.03$, with a combination of primordial B-mode power, a bump at low ℓ from reionization, and the lensing signal rising to peak on smaller scales ($\ell \sim 1000$). The 220 GHz data are dominated by dust, and lie above the CMB-lensing theory curve. The 150 GHz data are dominated by dust on degree scales $\ell \lesssim 200$ before the lensing B-modes rise above the dust signal at $\ell \gtrsim 200$. The 95 GHz channel is the cleanest for probing the CMB, although it still has dust contamination on degree scales. Spectral information is needed to distinguish between a gravitational wave signal from non-zero tensor-to-scalar ratio r and Galactic dust contamination.

We initially use the BK18 likelihood to reproduce the r constraints from Ade et al. (2021). Figure 4.3 shows the parameter constraints on r and the seven foreground parameters, and also compares the distribution of r to one where the foreground model parameters are fixed to their best-fit values; the dust and CMB are now almost uncorrelated with only a slight correlation with the dust index, and the error on r only increases by 10% when marginalizing over foregrounds. This is consistent with the BICEP/Keck collaboration’s results (Ade et al., 2021). We also check that we recover the same distribution for r as BK18 when adding the *Planck* PR3 likelihood and marginalizing over the Λ CDM parameters.

4.2 Further exploration of the BK18 likelihood

A suite of tests of the likelihood were reported in BK18, including internal consistency of splits of the data, and consistency of the signal measured in the deeper versus wider sky region. No significant impact on r was found, compared to the nominal analysis choices, from imposing a dust index prior, including dust decorrelation, freeing the lensing amplitude, including EE data, or dropping parts of the *WMAP* or *Planck* data. BK18 find that the data are well fit by a 3-parameter dust model, with no




Figure 4.3: Reproduction of the BK18 r and foreground parameters (dust and synchrotron amplitude, spatial spectral index and frequency spectral index, and dust-synchrotron correlation parameter) for fixed cosmological parameters. Only the dust frequency index, $\beta_{B,dust}$, is now correlated with the tensor-to-scalar ratio, r . The blue curve in the top left panel shows the constraint on r when holding the foreground parameters fixed at their best-fit values.

evidence yet for synchrotron emission, or for a primordial B -mode signal. In this section we explore a few further assumptions that can be tested at the likelihood level with the public data products².

4.2.1 Consistency of r from 95 GHz and 150 GHz

An important consistency test for any future detection of r will be to show agreement of the measurement from different frequencies. We show in Figure 4.4 the distribu-

²Maps of the sky would be needed for various further tests, for example analyzing different regions separately, or testing for non-Gaussian statistics.

tion for r estimated from just the BK18 95 GHz data, or just the 150 GHz data. In each case we include the other frequency data (23–44 and 217–353 GHz) to clean foregrounds. As noted in BK18, the latest BICEP3 95 GHz data dominate the constraint, with the overall $r < 0.036$ upper limit driven by the 95 GHz data (which give $r < 0.038$ without the 150 GHz data). The 150 GHz data, in combination with the foreground tracers, give $r < 0.072$; this larger uncertainty currently limits a frequency null test.




Figure 4.4: Distribution for r estimated from 95 GHz or 150 GHz, compared to the nominal combination. The foreground frequencies are included in all cases and the standard seven foreground parameters are estimated. The current limit is dominated by the 95 GHz BICEP3 data.

4.2.2 Consistency of r cleaned with 220 GHz or 353 GHz

Since dust is the dominant foreground, we would also hope to see consistency of r if different frequency maps are used independently as the dust tracer. In Figure 4.5 we show r and the three dust parameters for the case where the 217/220 GHz is discarded, and only 353 GHz is used to trace the dust ('no-220') versus discarding 353 GHz and using only 217/220 GHz as the dust tracer ('no-353'). In both cases the usual synchrotron parameters are also sampled. The estimated model parameters appear statistically consistent, with parameter means differing by only $\sim 1\sigma$. The

best-fit value of r is non-zero in the no-220 GHz case, but $r = 0$ is still an acceptable fit. For both cases all three dust parameters can be constrained by the data to a similar level of precision; the no-220 case prefers a steeper spectral index and shallower ℓ -dependence for the dust.




Figure 4.5: Parameter estimates are consistent whether 220+217 GHz or 353 GHz data are used to clean the dust. The best-fit value of r is non-zero in the no-220 GHz case, but $r = 0$ is still an acceptable fit. Scatter is expected since the dust-tracing data are independent in the two cases.

4.2.3 Impact of assuming spatially power-law dust

From observations by *Planck*, the dust is found to follow a power-law spatially in the angular range $40 < \ell < 600$ (Planck Collaboration et al., 2016d, 2020f), scaling as ℓ_d^α . However, the *Planck* data do show some departure from power law BB at larger angular scales, and also show statistically significant variations in the power

law over different sky regions. As data improve, we might expect to detect a deviation from a pure power-law in the BICEP/Keck region, which may impact the estimated r . To explore this, we extend the BK18 likelihood to include a different amplitude of dust in each of the nine bins. We replace the two parameters A_d and α_d , where $D_\ell = A_d(\ell/80)^{\alpha_d}$ at 353 GHz, with nine amplitudes that are sampled simultaneously with r and the other foreground parameters (making 15 parameters altogether instead of eight). We find negligible impact on r , as shown in Figure 4.6. The dust amplitudes are also shown in Figure 4.6, and are consistent with estimates in BK18 where the dust and non-dust parameters are estimated separately in each bin³. We find no significant departure from power law.

Within the power-law model, with $A_d(\ell/80)^{\alpha_d}$ at 353 GHz, the BK18 analysis imposes a hard prior $-1 < \alpha_d < 0$, motivated by *Planck* data. Given the improved data quality in BK18, we expand the prior on the dust slope, α_d , so that it is not limited by the hard boundary at $\alpha_d = -1$. The BK18 data can now constrain this slope from the data, as shown in Figure 4.6. This choice also has a negligible impact on r .

4.2.4 Effect of fixing the dust temperature

The fiducial model assumes a dust temperature of 19.6 K, which is the mean temperature estimated by *Planck* over the sky. This temperature is expected to have some spatial variation, and may depart from the mean value in the particular BK18 region. We test the impact of varying the dust temperature in Figure 4.7, over a broad range $0 < T_d < 50$ K. It has a negligible effect on the estimated r , and simply shows up as a correlation with the dust emissivity index, β_d .

³Here we are simultaneously estimating r and a single dust index, and common synchrotron parameters, together with the dust amplitudes.




Figure 4.6: Top left: there is minimal impact on r when we expand the spatial power-law dust model to one with an ell-by-ell amplitude, and negligible impact when we relax the prior on the power-law exponent. Top right: the power-law estimated with an expanded prior can still be constrained by the data. Bottom: estimated dust amplitude in bins, defined at 353 GHz, together with the best-fitting power-law model




Figure 4.7: Varying the dust temperature, T_d , does not impact the constraint on r . The temperature cannot be bounded from above by this dataset, but is weakly anti-correlated with the dust index, β_{dust} .

4.2.5 Relaxing the prior on the synchrotron index

In the nominal model there is a Gaussian prior on the synchrotron index, with $\beta_s = -3.1 \pm 0.3$ estimated from *WMAP* data. The upper limit on the synchrotron

amplitude is $A_s < 1.5 \mu\text{K}^2$ at 95% CL, defined as the power at 23 GHz and $\ell = 80$; as noted in BK18 there is not yet a detection of polarized synchrotron in BB . The data can still be used to put an upper limit on the synchrotron index, assuming a flat prior, since although a flat index $\beta_s = 0$ should be able to fit the data as A_s tends to zero, there is a far smaller volume of models at this high- β limit. Figure 4.8 shows that there is minimal impact on r if we allow the synchrotron index to be constrained by the data within some prior range $4 < \beta < 0$. There is an upper limit on β_s , but no lower limit, since an arbitrarily steep index can be accommodated by having a negligible synchrotron amplitude at 95 GHz.




Figure 4.8: Left: Removing the Gaussian prior on the synchrotron index has little effect on r . Center/right: no synchrotron emission is detected in the BK18 data, but with a uniform prior on the index, the larger volume of models at more negative values of β are preferred.

4.3 ‘BK18-lite’ - a compressed likelihood

In this section we construct a foreground-marginalized likelihood that depends on the BK18 data’s best estimate of the CMB power in each of the nine ℓ -bins. This has been done in the past for the ACT likelihoods and for the Plik-lite *Planck* likelihood (Dunkley et al., 2013; Calabrese et al., 2013; Planck Collaboration et al., 2016a; Prince & Dunkley, 2019). There are a few different factors that must be accounted for in this process. Firstly, the probability distribution for each binned BICEP bandpower is non-Gaussian, so it is key that we accurately capture the shape of the probability dis-

tribution of the bandpower $p(C_{bin})$ for each bin. For ACT and *Planck* the likelihoods were Gaussian. Secondly, we need to evaluate the covariance between bandpowers to see if the likelihoods for each bandpower can be combined independently, or whether we need to account for covariances between bins. Thirdly, we need to account for the three different window function shapes, and the different window function normalizations, based on the different regions of the sky observed by BICEP2, Keck Array, and BICEP3. A sample of these window functions is shown in Figure 4.1. The window functions for the smaller BICEP2/Keck patch of sky are slightly shifted in ℓ and have different shapes to the window function for the BICEP3 region. Although these differences are fairly minor, they do result in differences in the CMB bandpowers for each window function that are significant enough to affect the likelihood for the tensor-to-scalar ratio r .

We sidestep the issue of different window functions for now by using only the BICEP3 region maps (including the BICEP3 95 GHz map, as well as *WMAP* and *Planck* maps at different frequencies), which have a common window function. The constraints on r from these eight maps are very similar to the constraints from the full 11 maps that use the BICEP2 and Keck array 150 GHz and 220 GHz data, even though the foreground constraints are broader, as shown in Figure 4.9 (see also Appendix E2 and Figure 20 of Ade et al. (2021)). An alternative approach would be to estimate 27 foreground-marginalized bandpowers, nine for each window function shape, and use these to construct the likelihood. This would require accounting for covariances between highly overlapping but not identical ℓ ranges for different maps, and including the different normalizations of the window functions for the BICEP2/Keck frequencies. Another possible approach could use the BICEP2/Keck 150 GHz and 220 GHz data to put a prior on the foreground parameters before estimating the bandpowers from the BICEP3 region data. Other methods for describing the likelihood could also be explored, such as the Hamimeche-Lewis formalism used in




Figure 4.9: Constraints on the tensor-to-scalar ratio r and the seven foreground parameters from the full BICEP likelihood, using data from both the BICEP3 and BICEP2/Keck sky regions (black) and just the BICEP3 sky region maps (red). Dropping the BICEP2/Keck 150 and 220 GHz channels broadens the foreground constraints, particularly for the dust spectral index, but the r constraint is not significantly affected.

the full likelihood (Hamimeche & Lewis, 2008). We leave such extensions for future work, and concentrate only on the BICEP3 region for the remainder of this chapter.

4.3.1 Bandpower estimation

We use the BK18 likelihood described in [Ade et al. \(2021\)](#) to estimate the CMB bandpowers in the nine ℓ -bins simultaneously, marginalizing over the foreground parameters. In practice this involves modifying Cobaya to estimate 18 parameters via the usual Metropolis-Hastings algorithm: 9 bandpowers and 7 foreground parameters. For the data vector, we use only the maps that have the wider BICEP3 sky area, which includes the BICEP3 95 GHz map and the *Planck* and *WMAP* maps at low and high frequencies to constrain synchrotron and dust foregrounds. These maps all have the same bandpower window functions, and the BICEP3 95GHz provides the most constraining power on the tensor-to-scalar ratio r .

The estimated bandpowers are shown in Figure [4.10](#), both for the foreground-marginalized case in which the seven foreground parameters are allowed to vary in addition to the nine bin amplitudes, as well as for a fixed-foreground case in which the foreground parameters are set to their best-fit values. The bandpower distributions are very similar in both cases. There are mild correlation between neighbouring bins in ℓ , which is to be expected because the foregrounds are correlated between bins. In Figure [4.11](#) we show that the foreground parameters are consistent when estimating the bandpowers or when directly constraining r ; this checks that the method is not driving the foreground parameters to different values.

The one-dimensional distributions for the bandpowers are shown in Figure [4.12](#). We find that they are well fit by an offset log-normal distribution with the form:

$$p(D_b) = \frac{1}{(D_b - D_0)\sigma\sqrt{2\pi}} e^{-(\ln(D_b - D_0) - \mu)^2/(2\sigma^2)}, \quad (4.9)$$

where D_0 is the offset that makes $\ln(D_b - D_0)$ normally distributed for each bin, and μ and σ are the mean and standard deviation of $\ln(D_b - D_0)$ respectively. These best-fitting log-normal distributions are also shown in Figure [4.12](#). We also estimate the




Figure 4.10: Estimates of the nine B-mode bandpowers (here just for the BICEP3 region) marginalized over foregrounds (red) or estimated for the best-fit foregrounds (blue).

covariance between the bandpowers, finding the neighboring bins to have correlation coefficients of order 15-20%, but non-neighboring bins below 10%.

The foreground-marginalized bandpowers are plotted together with theory curves for different values of the tensor-to-scalar ratio in Figure 4.13. Our method of simultaneously estimating the nine bandpowers while marginalizing over the foregrounds is different from the approach the BICEP team take to decompose each bin independently into CMB, dust, and an upper limit on synchrotron for plotting purposes (see e.g., Figure 16 of Ade et al. (2021)). We obtain smaller errors because we use data from all the bins to constrain the foregrounds simultaneously, instead of handling




Figure 4.11: Distributions of the foreground parameters are consistent whether estimating the bandpowers or the tensor-to-scalar ratio r directly.

each bin separately. The CMB bandpowers in Figure 4.13 can be compared to the frequency channel bandpowers in Figure 4.2, which include the dust foregrounds. For convenience, the BICEP3 95 GHz channel is reproduced in Figure 4.13 as light blue dots. On scales larger than about 1 degree ($\ell < 180$) the CMB bandpowers are lower than the total 95 GHz channel because the dust contribution has been removed.




Figure 4.12: Probability distributions for the nine CMB bandpowers in the BICEP3 sky region, each fit with a three-parameter offset log-normal distribution.

4.3.2 The BK18-lite likelihood

Independently combining the log-likelihoods for each lognormal bin would give

$$\ln \mathcal{L} \propto \sum_{i=1}^9 \left(-\frac{(\ln(D_{th,i} - D_{0,i}) - \mu_i)^2}{2\sigma_i^2} - \ln(D_{th,i} - D_{0,i}) \right) + \text{constant}, \quad (4.10)$$

where $D_{th,i}$ is the theory B-mode power spectrum for a given model ($D^{BB}(\ell) = \ell(\ell + 1)C^{BB}(\ell)/2\pi$) that has been binned using the BICEP3 window functions, i.e. $D_{th,i} = \sum_{\ell} W_i(\ell)D_{th}(\ell)$ for $i = 1 \dots 9$. The r -constraints using this likelihood are shown as the blue curve in Figure 4.14. They are close to the r distribution found using the full likelihood, but not identical.

We can modify Equation 4.11 to include the covariances between bins:

$$\ln \mathcal{L} \propto -\frac{1}{2} (\ln(\mathbf{D}_{th} - \mathbf{D}_0) - \boldsymbol{\mu})^T \mathbf{Q}^{-1} (\ln(\mathbf{D}_{th} - \mathbf{D}_0) - \boldsymbol{\mu}) - \sum_{i=1}^9 \ln(D_{th,i} - D_{0,i}) + \text{constant}, \quad (4.11)$$




Figure 4.13: Foreground-marginalized CMB bandpowers (blue circles) plotted with the theoretical B-mode power spectrum for $r = 0$ (solid black curve) and $r = 0.03$ (dashed black curve). The BICEP3 95 GHz bandpowers are shown for comparison (light blue dots). The values of the foreground-marginalized bandpowers are taken from the mode of the best-fit log-normal distribution, and the errorbars for the foreground-marginalized bandpowers are calculated using the 16th and 84th percentiles of the distribution. These errors are asymmetric, with the upper errors larger than the lower error, although this is difficult to see with the log scale on the y-axis. The errorbars plotted for the BICEP3 95GHz channel, taken from the publicly available BICEP/Keck data, are symmetric and reflect the width and not the shape of the bandpower distribution.

where \mathbf{D}_{th} is a vector of the binned theory B-mode power spectrum for a given model and \mathbf{D}_0 is a vector of the best-fit offsets and means for the CMB data bandpowers. This is equivalent to assuming that using the best-fit one-dimensional offset $D_{0,i}$ for each binned data bandpower D_i , the quantity $\ln(D_i - D_{0,i})$ for $i = 1..9$ has a multivariate normal distribution with covariance matrix \mathbf{Q} and mean vector $\boldsymbol{\mu}$. We find that a log-normal likelihood that includes the covariance between neighboring bins (\mathbf{Q} is a block-diagonal matrix) is sufficient to reproduce the BICEP3 constraint on the tensor-to-scalar ratio r .




Figure 4.14: The constraints on r constraints from our new foreground-marginalized ‘BK18-lite’ likelihood (black) agree with the nominal BK18 likelihood, restricted to just the maps that use the BICEP3 window functions (red). Neglecting the covariance between neighboring bins in the BK18-lite likelihood gives the blue curve.

The constraints on the tensor-to-scalar ratio r using our BK18-lite likelihood and the full BICEP likelihood for the BICEP3 region are shown in Figure 4.14, as the black and red curves respectively. They give almost identical constraints on r , validating our compression into nine offset log-normal foreground-marginalized bins. This likelihood is useful because it does not require sampling the foreground parameters for each r constraint, the CMB-only bandpowers can be visually compared to the CMB theory power spectrum, and it can be easily used with automatically differentiable theory codes when using sampling methods like Hamiltonian Monte Carlo and the No-U-Turn Sampler.

4.4 Conclusion

In this chapter we have studied the likelihood analysis for the BICEP/Keck collaboration tensor-to-scalar ratio (Ade et al., 2021), and explored the effects of different foreground modeling choices on the results. We compare primordial B-mode constraints using the 95 GHz and 150 GHz science channels respectively, and find that they are consistent, with the 95 GHz channel having significantly more constraining power. We find that dust parameters measured primarily using complementary data at 220 and 353 GHz are consistent. Relaxing the power law assumption for the spatial variation of dust or allowing the dust temperature to vary instead of fixing it to the mean temperature estimated by *Planck* has no significant effect on the tensor-to-scalar ratio constraint. Relaxing the prior on the synchrotron spectral index also has little effect on the r constraint. Public availability of the BICEP and Keck Array maps in addition to the spectra would enable reproduction and tests of the Gaussianity and spatial variation of the dust in the region.

We then estimate the joint probability distribution of nine foreground-marginalized CMB bandpowers using BICEP3, *Planck* and *WMAP* maps in the BICEP3 region of the sky. We fit these bandpowers with offset log-normal distributions and then construct a likelihood that independently combines the nine bins before adding a term to account for the covariances between neighbouring bins. We show that this new ‘BK18-lite’ likelihood accurately reproduces the r constraint obtained with the BICEP/Keck team’s likelihood. We plan to incorporate the BICEP2/Keck data at 95, 150 and 220 GHz into our BK18-lite likelihood and release it publicly for use as a nuisance-parameter-free alternative to the full BICEP likelihood for constraining r . We also plan to further test the BK18-lite likelihood in combination with other datasets for constraining a variety of cosmological models.

Chapter 5

Primordial adiabatic and isocurvature power spectrum constraints with the Atacama Cosmology Telescope DR4: beyond the single power law

The Λ CDM model of the universe, which is a good fit to a wide range of cosmological data, assumes that the primordial scalar fluctuations in the universe are adiabatic, Gaussian, and take the form of a power law. In standard inflation models of the early universe this power law is predicted to be close to scale invariant, i.e. close to having a spectral index $n_s = 1$ (Peebles & Yu, 1970). The *Planck* 2018 legacy data release provides an estimate of the spectral index with $n_s = 0.965 \pm 0.004$, which is 8σ away from scale invariance (Planck Collaboration et al., 2020c). A slight red tilt (more power on larger scales) is preferred, consistent with predictions from inflationary models (Planck Collaboration et al., 2020e).

Primordial isocurvature fluctuations can also be present in combination with adiabatic fluctuations in the early universe (e.g., [Bucher et al., 2000](#)). In the case of adiabatic fluctuations, which are the only fluctuations that can be generated by single field inflation, the overdensities of the different components are proportional to one another, so a region that has a radiation overdensity has a corresponding overdensity in other particle species such as dark matter, baryons and neutrinos. For isocurvature fluctuations, however, an overdensity of one component, such as radiation, corresponds to an underdensity of another component, say cold dark matter, such that the sum of these components remains constant. The possibility of the primordial fluctuations being purely isocurvature modes ([Efstathiou & Bond, 1986, 1987](#)) was ruled out when the location of first acoustic peak of the CMB was measured, as it is consistent with adiabatic and not isocurvature predictions ([Enqvist et al., 2002](#)). We now know from *Planck* data that most of the primordial fluctuation power comes from purely adiabatic fluctuations, with a 95% upper limit of 1.3% and 1.7% of the total primordial power being sourced from cold dark matter density isocurvature and neutrino density isocurvature respectively in the range of scales probed by *Planck*, assuming a power law for the isocurvature fluctuations ([Planck Collaboration et al., 2020e](#)).

More recent measurements from the Atacama Cosmology Telescope (ACT) and the South Pole Telescope (SPT) ground-based experiments constrain the CMB at smaller angular scales than *Planck*. Interestingly, they are consistent with the adiabatic scalar fluctuations having $n_s = 1$ ([Aiola et al., 2020](#); [Dutcher et al., 2021](#)). These new data also have the potential to constrain isocurvature fluctuations at smaller scales than accessible to *Planck*. In this chapter we use data from the ACT DR4 data release, and from the SPT-3G 2020 release, to explore three possible extensions to the Λ CDM model. The first is purely adiabatic, but has a broken power law for its primordial power spectrum, with scale invariant behavior on small scales and a red

tilt on larger scales. The second two extensions keep the adiabatic component as a pure power law, but include either CDM or neutrino isocurvature power in bins as a function of angular scale.

This chapter is laid out as follows. We describe the CMB data and likelihoods we use in Section 5.1. Section 5.2 gives constraints on the broken power law model for the primordial adiabatic power spectrum. In Section 5.3 we present constraints on the binned CDM density and neutrino density isocurvature power. We discuss our results in Section 5.4.

5.1 Data and likelihoods

In this section we describe the data that we use from satellite (*Planck*) and ground-based (ACT and SPT) CMB experiments. These data are complementary because satellites are able to survey a large fraction of the sky and are not limited by atmospheric noise on large scales, thus giving excellent measurements of the CMB power on large angular scales. Ground-based experiments are able to use larger telescope mirrors than satellites and replace their detector arrays with ever-improving technology to obtain low-noise measurements at small angular scales. Therefore combining data from space- and ground-based telescopes gives good signal to noise over a large range of angular scales. Additionally, ACT and SPT observe different regions of the sky from their locations in the Atacama desert and at the South Pole, so their data do not have a large overlap in sky area.

ACT: We use temperature and polarization data from the fourth data release of the Atacama Cosmology Telescope, ACT DR4 (Aiola et al., 2020; Choi et al., 2020). Pyactlike, the foreground-marginalized likelihood for TT, TE and EE, is publicly available on GitHub¹. This likelihood has been obtained by marginalizing over the

¹<https://github.com/ACTCollaboration/pyactlike>

foregrounds in the multifrequency TT, TE and EE power spectra in wide and deep fields from ACT DR4 (Choi et al., 2020) and data from ACT MBAC (Das et al., 2014), resulting in binned CMB-only angular power spectra for TT, TE and EE with a covariance matrix that includes the various sources of uncertainty including noise, foregrounds, the beam, and calibration uncertainties. Pyactlike has one nuisance parameter that must be sampled in addition to the cosmological parameters of interest: the polarization efficiency y_p . The full multipole range includes bin centers at $600 \leq \ell \leq 4125$ for TT and $350 \leq \ell \leq 4125$ for TE and EE. However when we combine ACT DR4 with the full ℓ range of *Planck* we drop the ACT DR4 TT data at $\ell < 1800$ to avoid double counting scales that are observed by both experiments.

SPT: The third generation camera on the South Pole Telescope (SPT-3G) includes in its 2020 data release (Dutcher et al., 2021) the TE and EE power spectra over a multipole range of $300 \leq \ell < 3000$. We use the Python implementation of the SPT-3G likelihood² presented in La Posta et al. (2021). The likelihood combines power spectra from the three frequency channels at 95, 150, and 220 GHz with foreground modeling of Galactic dust for TE and EE and polarized point sources for EE.

Planck: We use *Planck* data and likelihoods from the 2018 legacy data release (Planck Collaboration et al., 2020b). We use the publicly available foreground-marginalized CMB-only likelihood Plik_lite, which bins the power spectrum in the range $30 \leq \ell \leq 2508$ for TT and $30 \leq \ell \leq 1996$ for TE and EE. For low- ℓ we use the Commander TT likelihood and the Sim-all EE likelihood for $2 \leq \ell < 30$. The low- ℓ EE likelihood constrains the optical depth to reionization τ . For some data combinations we also use a modified version of Planck-lite-py, our Python

²https://github.com/xgarrido/spt_likelihoods

implementation of the `Plik_lite` likelihood, that includes *Planck* TT data only from $30 \leq \ell \leq 645$. We call this likelihood `PlanckTT645` and use it in combination with the full ℓ range of ACT DR4 temperature data to provide constraints that are independent of *Planck* on small scales ($\ell > 650$).

Planck CMB lensing: We use the CMB lensing power spectrum likelihood from the *Planck* 2018 legacy release (Planck Collaboration et al., 2020d) in our isocurvature constraints in Section 5.3.

5.2 Primordial Adiabatic Power Spectrum

In this section we investigate whether a modification to the power-law model for the primordial spectrum is able to better fit the *Planck* and ground-based data. In this first extension to Λ CDM we assume the perturbations are purely adiabatic, where the photon entropy n_i/n_γ for the various components i of the universe is constant. The photon entropy perturbation (which is zero in the adiabatic case) is given by

$$S_i = \frac{\delta(n_i/n_\gamma)}{n_i/n_\gamma} = \frac{\delta_i}{1+w_i} - \frac{\delta_\gamma}{1+w_\gamma}, \quad (5.1)$$

where n_i is the number density of component i (which could be cold dark matter, neutrinos, baryons, etc), n_γ is the number density of photons, $\delta_i = \delta\rho/\bar{\rho}$ is the energy density perturbation for component i , and $w_i = \rho_i/p_i$ is the equation of state of component i . Using the appropriate values for the equation of state, we see that adiabatic perturbations correspond to $\delta_c = \delta_b = \frac{3}{4}\delta_\nu = \frac{3}{4}\delta_\gamma$, i.e. the energy densities of the components of the universe are proportional to one another.

5.2.1 Power law primordial power spectrum

The primordial adiabatic power spectrum is defined as

$$P_\phi(k) \equiv \frac{1}{(2\pi)^3} \langle \phi(\mathbf{k}) \phi(-\mathbf{k}) \rangle, \quad (5.2)$$

where ϕ is the Newtonian potential in the perturbed FRW metric. The power spectrum has units of (length)³ so to obtain a dimensionless quantity, we can multiply it by k^3 , as the wavenumber k has units of inverse length.

$$\Delta^2(k) \equiv \frac{k^3 P(k)}{2\pi^2} \quad (5.3)$$

The *Planck* 2018 data are well-fit by a Λ CDM model with a single power law power spectrum of the form

$$\Delta^2(k) = A_s \left(\frac{k}{k_*} \right)^{n_s - 1}, \quad (5.4)$$

with a best fit scalar spectral index $n_s = 0.965 \pm 0.004$ ([Planck Collaboration et al., 2020c,e](#)). The pivot scale is set to be $k_* = 0.05 \text{Mpc}^{-1}$ in the *Planck* analysis. This best-fit spectral index is $> 8\sigma$ away from the $n_s = 1$ scale-invariant Harrison-Zel'dovich-Peebles power spectrum ([Harrison, 1970](#); [Zeldovich, 1972](#); [Peebles & Yu, 1970](#)). The *Planck* collaboration finds no evidence for n_s varying with scale through a running of the spectral index $d(\ln n_s / d \ln k)$, or for a running of running of the spectral index $(d^2 \ln n_s / d(\ln k)^2)$ ([Planck Collaboration et al., 2020e](#)). These constraints were derived using the temperature power spectrum up to maximum multipole of $\ell = 2509$ and E-mode polarization and TE power spectra with a maximum multipole of $\ell = 1997$.

Ground based experiments with sensitive observations of the CMB at small scales have shorter lever-arms over which to measure the spectral index, and find best-fitting values of the spectral index that are consistent with a scale invariant power spectrum

with $n_s = 1$. The ACT DR4 data are used to estimate $n_s = 1.008 \pm 0.015$ from ACT alone using a multipole range of $350 \leq \ell \leq 4000$, and $n_s = 0.973 \pm 0.006$ with ACT and *WMAP* data (Aiola et al., 2020). The SPT-3G data provide an estimate of $n_s = 0.995 \pm 0.024$ using binned TE and EE power spectra over the multipole range $300 \leq \ell < 3000$. The values obtained by ACT and *Planck*, shown in Figure 5.1 from (Aiola et al., 2020), do not differ by a statistically significant amount, and the different results could simply be due to statistical fluctuations. However, this $2-3\sigma$ level difference could also be suggestive of a primordial power spectrum that is not purely a power law, and we explore that possibility in this section. In the ACT analysis in (Aiola et al., 2020) a running of the spectral index was investigated, and was found to be preferred at the 2.4σ level with ACT alone, but no running was preferred once the *WMAP* or *Planck* data were included.




Figure 5.1: 1σ and 2σ probability distribution contours for the spectral index, n_s , and the baryon density, $\Omega_b h^2$, using ACT DR4 (light blue), *Planck* (purple) and *WMAP* (yellow) data reproduced from Figure 14 of Aiola et al. (2020). *Planck* and *WMAP* prefer lower values of n_s than ACT, which probes smaller angular scales.

We begin our investigation by reproducing the standard power-law results reported in Aiola et al. (2020). Figure 5.2 shows the posterior distributions of the six Λ CDM

parameters that we estimate from ACT DR4 data alone and from ACT combined with *Planck* (for this combination we follow [Aiola et al. \(2020\)](#) in dropping ACT TT data below $\ell = 1800$ to avoid double counting the modes measured by *Planck*). For both of these likelihood combinations we use a Gaussian prior of $\tau = 0.065 \pm 0.015$ instead of the *Planck* low- ℓ EE data, for consistency with [Aiola et al. \(2020\)](#). We find that the parameters agree to within the expected chain-convergence level with those estimated in the ACT DR4 analysis ([Aiola et al., 2020](#)). Figure 5.2, in addition to Figure 5.1, illustrates how the ACT data alone prefer $n_s \sim 1$, while *Planck* prefers a lower value, although they are not statistically inconsistent. Combining ACT DR4 with *Planck* data brings the spectral index down to below $n_s = 1$.




Figure 5.2: Parameter constraints using ACT DR4 (blue) and ACT DR4 combined with *Planck* 2018 (black). ACT prefers a higher value of n_s than *Planck*, as seen in the middle panel of the top row, but there is no significant tension between the two datasets. The agreement of our parameter estimates with distributions from the official ACT DR4 data release [Aiola et al. \(2020\)](#) are shown for comparison (dashed purple and grey curves).

5.2.2 A broken power law spectrum

Given the slight preference for different power-law indices at large and small scales, we consider an alternative to the running of the spectral index by exploring a broken power law model. This type of parameterization has been explored in the past for *WMAP* and *Planck* data (e.g., [Bridle et al., 2003](#)). While most inflationary models predict a slowly and monotonically varying power spectrum, more complicated and fine-tuned inflationary models could introduce features at the length-scales probed by current data. This is thus an empirical model that could be connected to theories of the early universe if preferred by the data.

The simplest version we test has a scale invariant spectral index ($n_s = 1$) on small scales and a smaller spectral index on large to intermediate scales. The equation we adopt for a smooth broken power law is

$$P(k) = A_b \left(\frac{k}{k_b} \right)^{(n_{s,1}-1)} \left[\frac{1}{2} \left(1 + \left(\frac{k}{k_b} \right)^{1/\Delta} \right) \right]^{\Delta(n_{s,2}-n_{s,1})}, \quad (5.5)$$

for a power law that breaks at $k = k_b$ with spectral index $n_s = n_{s,1}$ for $k < k_b$ and $n_s = n_{s,2}$ for $k > k_b$, with Δ controlling the width of the transition between spectral indices. We assume that the power spectrum is scale invariant at small scales then $n_{s,2} = 1$. We then fit for the amplitude, power law index on large scales, and break wavenumber k_b of the broken power law, i.e., with one additional power spectrum parameter compared to the standard Λ CDM model.

In Figure 5.3 we show the posterior distributions for the break wavenumber, k_b , and large-scale spectral index, n_s (with $n_s = 1$ for $k > k_b$), as well as the other cosmological parameters. Using just *Planck* data we find a preferred value of $k_b \geq 0.15 \text{ Mpc}^{-1}$, which corresponds to $\ell \sim 2000$ which is at the upper end of *Planck*'s ℓ range. This is effectively equivalent to the *Planck* data having a single power law




Figure 5.3: Parameter constraints for a broken power-law primordial power spectrum that is scale invariant on small scales, with a slope equal to n_s at scales $k < k_{\text{break}}$ (given in units of Mpc^{-1}). *Planck* prefers a single power law and therefore pushes the transition scale, k_{break} , to higher values (smaller scales) that are not well-constrained by *Planck*. Adding ACT or ACT+SPT data selects a preferred transition scale at $k \sim 0.1 \text{Mpc}^{-1}$, although a range of larger k -values are still allowed.




Figure 5.4: Estimated primordial power spectra for a broken power law (blue) and a single power law (green), with k in units of Mpc^{-1} . The maximum likelihood curve for ACT+*Planck* is shown as the solid curve, together with a sample of 500 other points in the MCMC chains whose density is proportional to the model's posterior probability.




Figure 5.5: Parameter constraints for a broken power law primordial power spectrum where both spectral indices are allowed to vary, with a break at k_{break} , for *Planck*+ACT, and *Planck*+ACT+SPT. A broken power law that is scale invariant on small scales ($n_s = 1$) has the highest probability, although models where the two spectral indices are equal, i.e., the primordial power spectrum has a single power law, are allowed by the data.

with a slight red tilt ($n_s < 1$), and is consistent with there being no preference in the *Planck* data alone for a deviation from power-law.

When we combine the *Planck* data with ACT DR4 data, or with ACT DR4 and SPT-3G data, we find that a preferred value of $k_b = 0.1 \text{ Mpc}^{-1}$ emerges, corresponding to $\ell \sim 1400$. This best-fitting primordial power spectrum is shown in Figure 5.4, compared to the single power law models. The spectral index for $k < k_b$ in the broken power law model is slightly steeper than in the single power law model, before flattening out to scale invariance at $k > k_b$. We indicate the uncertainties in this figure by drawing a sample of 500 points from the sampled chain, and plotting the primordial power for each sample, such that the density of models is proportional to the underlying posterior probability.

Despite the preference for this broken power-law, the distribution for k_b in Figure 5.3 also shows that a high value of $k_b \approx 0.3 \text{ Mpc}^{-1}$ is not excluded by the data. This transition scale translates to $\ell \sim 4200$, which is beyond the multipole range probed by the ACT data, and thus is effectively equivalent to a single power law. There is an improvement in χ^2 of $\Delta\chi^2 = 3.8$ for the broken power-law, compared to $\Delta\chi^2 = 1.4$ for a running spectral index. However, this is not a large enough im-

provement to statistically prefer the broken power-law overall; the single power-law model is still statistically preferred. When limiting the data to ACT DR4 + *Planck* TT $\ell < 645$, the improvement in χ^2 is 10.3 for the broken power-law model, with the broken power-law mildly preferred at the $2-3\sigma$ significance level. It will be interesting to test this model with upcoming ACT DR6 data.

To check whether the smallest scales prefer a different power-law to $n_s = 1$ we also allow both slopes to vary, with Figure 5.5 showing the posterior distributions when the spectral index at small scales is allowed to vary independently. We find that a scale invariant value of $n_{s,2} = 1$ is preferred on small scales, but a range of other options are allowed by the data, including the range where $n_{s,1} = n_{s,2}$.

5.3 Primordial isocurvature perturbations

Spatial overdensities and underdensities of different particle species in the early universe can be divided into two orthogonal modes: adiabatic and isocurvature perturbations. Taking photons as a reference species, another type of particle i (for example, cold dark matter or neutrinos) whose fluctuations are purely adiabatic would have a number density n_i that is proportional to the photon number density n_γ everywhere in space, i.e., n_i/n_γ is spatially constant. Adiabatic perturbations are also known as curvature perturbations because there is a spatial variation of density which corresponds to a spatial variation of curvature. In contrast, isocurvature modes for species i occur when spatial fluctuations of the densities of this species n_i cancel out the spatial fluctuations of the photons n_γ . Multi-field inflation (or alternative early universe models) can give rise to isocurvature modes in addition to adiabatic modes (Langlois, 1999).

In this section we consider isocurvature perturbations with non-zero perturbations in the photon entropy, n_i/n_γ for two different components i , either cold dark matter

or neutrinos. The primordial isocurvature power spectrum is defined as

$$P_{\mathcal{I}\mathcal{I},i}(k) = \left(\frac{k}{2\pi}\right)^3 \langle S_i(\mathbf{k})S_i(-\mathbf{k})\rangle, \quad (5.6)$$

where S_i is the entropy perturbation defined in Eqn 5.1, and i corresponds to baryons, cold dark matter, or neutrinos. If the isocurvature and adiabatic perturbations are correlated, the primordial cross-power is given by

$$P_{\mathcal{R}\mathcal{I},i}(k) = \left(\frac{k}{2\pi}\right)^3 \langle \phi(\mathbf{k})S_i(-\mathbf{k})\rangle, \quad (5.7)$$

where ϕ is the Newtonian potential that defines the adiabatic perturbation.

Here we consider isocurvature fluctuations between the photon density and either the cold dark matter density (\mathcal{I}_{CDI}) or neutrino density (\mathcal{I}_{NDI}). Isocurvature between the photon and baryon density has an almost identical effect on the CMB as cold dark matter density isocurvature so we do not quote separate constraints on these two modes. If there was simultaneously both baryon and cold dark matter density isocurvature then the CDI mode would include them both, with $\mathcal{I}_{\text{CDI}}^{\text{effective}} = \mathcal{I}_{\text{CDI}} + \frac{\Omega_b}{\Omega_C} \mathcal{I}_{\text{BDI}}$ (Gordon & Lewis, 2003). We do not include isocurvature fluctuations between the photon and neutrino velocity, identified in Bucher et al. (2000), because there is no known way of exciting this mode.

5.3.1 Isocurvature mode phenomenology

As described in (e.g., Peebles & Yu, 1970; Bucher et al., 2000), isocurvature initial conditions produce a CMB angular power spectrum that is quite distinct to one with adiabatic initial conditions. We illustrate this in Figure 5.6, showing the shape of the adiabatic power spectrum generated using the CLASS numerical code, compared to the power spectrum from isocurvature components. All the spectra are shown for a scale-invariant primordial power spectrum ($n_s = 1$) with scalar power $A_s = 2.3 \times 10^{-9}$




Figure 5.6: The TT, TE and EE angular power spectra of the CMB for adiabatic, neutrino density isocurvature (NDI) and cold dark matter density isocurvature (CDI) initial conditions. In each case the primordial power spectrum $P_{\mathcal{R}\mathcal{R}}(k)$ is scale invariant ($n_s = 1$) with scalar amplitude $A_s = 2.3 \times 10^{-9}$. The isocurvature C_ℓ 's have been multiplied by a factor of 10 relative to the adiabatic C_ℓ 's so that their shapes are easier to compare.

defined at a pivot scale of $k = 0.05 \text{ Mpc}^{-1}$, but we multiply the isocurvature spectra by 10 so that their overall shapes and phases relative to the acoustic peaks of the adiabatic power spectrum can be seen.

\mathcal{I}_{CDI} produces a CMB power spectrum that is almost exactly out of phase with the adiabatic power spectrum, due to the phase shift in the initial conditions: its troughs are at multipoles close to the adiabatic peaks. Because of this, a small amount of CDI, in addition to the adiabatic power, has the effect of smoothing the acoustic peaks in a similar way to gravitational lensing of the CMB. In temperature, the cold dark matter isocurvature spectrum is large on large scales (low ℓ) and falls off quickly with ℓ in both temperature and polarization, dropping much more steeply than the adiabatic or neutrino density isocurvature modes. For a given amplitude A_s , cold dark matter isocurvature produces less power in the CMB, resulting in suppressed power spectra for TT, TE, and EE compared to adiabatic and neutrino density isocurvature except at very small scales. These differing CMB amplitudes arise due to the different transfer functions for each component.

The \mathcal{I}_{NDI} power spectrum has peaks approximately in between the adiabatic and CDI peaks. It is about a factor of 10 lower than the adiabatic CMB power for the same initial primordial amplitude in $P(k)$, but has a similar behavior with ℓ , with the angular power falling off at around the same rate at high- ℓ , including at $\ell > 2000$ which is beyond the range included in Figure 5.6.

The adiabatic and isocurvature primordial perturbations are small and so add linearly, such that the total primordial CMB power spectrum is obtained by combining the adiabatic, isocurvature, and cross-correlation terms,

$$C_\ell = C_\ell^{\text{Ad}} + C_\ell^{\text{Iso}} + C_\ell^{\text{Ad} \times \text{Iso}}. \quad (5.8)$$

Some early universe models predict spatial correlations between adiabatic and isocurvature modes (e.g., [Lyth & Wands, 2003](#); [Kitajima et al., 2017](#)), which can be either negative or positive, and either partially or fully correlated depending on the physical mechanism sourcing the perturbations.

5.3.2 Parameterizing isocurvature perturbations

A single, generally correlated isocurvature mode was modeled for the *Planck* analysis in [Planck Collaboration et al. \(2020e\)](#) by specifying the amplitude of the isocurvature power and the adiabatic power at two reference scales and assuming that they take the form of a power law. The scales are fixed to wavenumbers $k_1 = 0.002 \text{ Mpc}^{-1}$ and $k_2 = 0.1 \text{ Mpc}^{-1}$ and then the relevant sampled parameters are $P_{\mathcal{R}\mathcal{R}}^{(1)}$ and $P_{\mathcal{R}\mathcal{R}}^{(2)}$ for the adiabatic power, $P_{\mathcal{I}\mathcal{I}}^{(1)}$ and $P_{\mathcal{I}\mathcal{I}}^{(2)}$ for the isocurvature power, and $P_{\mathcal{R}\mathcal{I}}^{(1)}$ for the correlation power between the adiabatic and isocurvature modes. The spectral indices $n_{\mathcal{R}\mathcal{R}}$ and $n_{\mathcal{I}\mathcal{I}}$ can be calculated from these amplitudes. The correlation power is assumed to also be a power law, with a spectral index given by the average of the adiabatic and isocurvature indices $n_{\mathcal{R}\mathcal{I}} = (n_{\mathcal{R}\mathcal{R}} + n_{\mathcal{I}\mathcal{I}})/2$.

Given that we do not a priori know the scale dependence of any possible isocurvature component, we extend this model to constrain five bins equally spaced in the logarithm of the wavenumber, k , and allow the isocurvature power to vary independently in these bins. We smooth the bin transitions by convolving with a Gaussian. We assume that the adiabatic and isocurvature modes are uncorrelated, consistent with the $P_{\mathcal{R}\mathcal{I}}$ constraints from *Planck* ([Planck Collaboration et al., 2020e](#)). As in the power law case, we allow one isocurvature mode to be present at a time, so we constrain the CDI and NDI modes separately. One advantage of using independent bins to constrain the isocurvature power is that it is possible to explore which scales the constraints are driven by, and where systematic errors may be arising. Additionally, because the small scale CMB power is well constrained by ACT but not by *Planck*, we can see whether there is some power there from a model that was not able to be constrained by *Planck*. We choose bins in the range $0.001 < k < 0.5 \text{ Mpc}^{-1}$, which maps onto $10 < \ell < 7000$.




Figure 5.7: The 95% upper limits on the cold dark matter density isocurvature (CDI) primordial power spectrum as a function of wavenumber, k , with five bins smoothed by convolving the binned power spectrum with a Gaussian. The best-fitting adiabatic primordial power spectrum is shown for comparison. The ACT DR4 data improve the constraints from *Planck* at the smallest scales.

5.3.3 Constraints on CDM isocurvature

We first test our parameter estimation framework by checking that we reproduce the same parameter distributions as the *Planck* collaboration, using the power-law assumption for the isocurvature power, and allowing the correlation power $P_{\mathcal{R}\mathcal{I}}$ to vary. We find the same results as reported in the red contours of Figure 38 of [Planck Collaboration et al. \(2020e\)](#), with $P_{\mathcal{R}\mathcal{I}}^1 < 5.7 \times 10^{-11}$ at $k_1 = 0.002 \text{ Mpc}^{-1}$ and $P_{\mathcal{R}\mathcal{I}}^2 < 3.3 \times 10^{-9}$ at $k_1 = 0.1 \text{ Mpc}^{-1}$ for CDI, and $P_{\mathcal{R}\mathcal{I}}^1 < 1.5 \times 10^{-10}$ and $P_{\mathcal{R}\mathcal{I}}^2 < 6.6 \times 10^{-10}$ for NDI. Within this restricted power-law model, we also find – as expected – that the ACT DR4 data do not add noticeable constraining power to *Planck*.

For the five-bin model, with uncorrelated CDM isocurvature, we find 95% upper limits on the binned power for CDI as shown in Figure 5.7, together with the *Planck* best-fitting adiabatic primordial power spectrum for comparison. We find the

following upper limits as a function of scale for ACT+*Planck*:

$$\begin{aligned}
P_1(k) &< 1.1 \times 10^{-10} & (0.001 - 0.0035 \text{ Mpc}^{-1}) \\
P_2(k) &< 1.7 \times 10^{-10} & (0.0035 - 0.012 \text{ Mpc}^{-1}) \\
P_3(k) &< 2.9 \times 10^{-10} & (0.012 - 0.04 \text{ Mpc}^{-1}) \\
P_4(k) &< 1.5 \times 10^{-9} & (0.04 - 0.14 \text{ Mpc}^{-1}) \\
P_5(k) &< 6.9 \times 10^{-8} & (0.14 - 0.5 \text{ Mpc}^{-1})
\end{aligned} \tag{5.9}$$

The ACT data adds negligible information to the larger-scale constraints, but improves the limits on the two bins at $k > 0.04 \text{ Mpc}^{-1}$ by about 35%.

The posterior distributions for all the sampled parameters are shown in Figure 5.8. The top two rows show the usual cosmological parameters where the power law adiabatic power spectrum is parameterized by its amplitudes $P_{\mathcal{R}\mathcal{R}}$ at two different scales k_1 and k_2 instead of by an amplitude and spectral index. All power spectrum amplitudes are renormalized by a factor of 10^{-10} so that the numerical values for most of the isocurvature bin errors are around a few in these units, while the primordial adiabatic power is $\sim 2 \times 10^{-9} = 20 \times 10^{-10}$.

We find that the adiabatic amplitudes are shifted slightly lower when isocurvature is included, with *Planck* + ACT preferring lower values of the adiabatic amplitudes $P_{\mathcal{R}\mathcal{R}}$ when isocurvature is included. The isocurvature power in the first three k -bins, corresponding to $\ell \sim 10$ to 600, is constrained to be less than a few $\times 10^{10}$. On smaller scales the upper bound on the isocurvature modes is less constrained. In k -bin 4, corresponding to $\ell \sim 600$ to 2000, the cold dark matter density isocurvature upper bound is around 2×10^{-9} , which is the same as the best-fitting adiabatic power. The fifth bin, which corresponds to $\ell > 2000$, is not well constrained in the CDI case, because CDI power contributes little to the CMB power spectrum.




Figure 5.8: CDI constraints: posterior distributions for the amplitudes of the power-law adiabatic power spectrum at two reference scales P_{RR}^1 and P_{RR}^2 , Hubble constant H_0 , baryon density parameter $\Omega_b h^2$, cold dark matter density parameter $\Omega_c h^2$ and the optical depth to reionization τ , as well as the binned cold dark matter density isocurvature amplitudes in five bins, $P_{II}^1 \dots P_{II}^5$.

Although a model with zero isocurvature is still an excellent fit to the data, we do find a mild preference for a non-zero isocurvature amplitude in the fifth bin, but at low significance ($< 2\sigma$). In adding isocurvature, we also find that the best-fitting matter density is reduced compared to the model with purely adiabatic perturbations,

and the Hubble constant is increased, but the parameter shifts are only a fraction of a standard deviation.

We follow [Planck Collaboration et al. \(2020e\)](#) in quantifying the amount of isocurvature power relative to the total power in the CMB. The non-adiabatic fraction, α , of the observed CMB variance is defined by

$$\alpha = \frac{\sum_{\ell=2}^{4000} (2\ell + 1) |C_{\text{iso},\ell}^{XY}|}{\sum_{\ell=2}^{4000} (2\ell + 1) |C_{\text{tot},\ell}^{XY}|}, \quad (5.10)$$

where $XY=TT$, TE or EE ; here we extend the multipole range from $\ell_{\text{max}} = 2500$ used by *Planck* up to $\ell_{\text{max}} = 4000$. We find 95% upper limits on the contributions of CDI to the TT, TE and EE power spectra up to $\ell = 4000$ to be 1.9%, 1.7%, and 0.6% respectively.

Finally we also check that the upper limits would not be inconsistent with current large-scale structure measurements, by plotting the matter power spectrum at $z = 0$ for the model with the 95% upper limit of CDM isocurvature. Shown in Figure 5.9, we find that the isocurvature component is still well below the adiabatic level, due to the suppressed matter transfer function for CDI compared to adiabatic perturbations.

5.3.4 Constraints on neutrino isocurvature

We repeat the same analysis for the NDI mode, and show the 95% upper limits on the primordial isocurvature power in Figure 5.10. We find the following upper limits as a function of scale for ACT+*Planck*:




Figure 5.9: The 95% upper limits on the contribution of isocurvature to the linear matter power spectrum at $z = 0$. The solid curve is the total adiabatic and isocurvature power spectrum, the red dotted curve shows the upper limits on the cold dark matter density isocurvature contribution, which is below 1% in the first few bins but jumps up to 25% at $k \sim 0.2$ h/Mpc and the blue dashed curve shows the upper limits on the neutrino density isocurvature contribution, which remains below 1% at all scales. This plot shows the linear matter power spectrum; the contributions to the total (including nonlinear) matter power spectrum are of the same order of magnitude but have a more complex as a function of k .

$$\begin{aligned}
 P_1(k) &< 3.4 \times 10^{-10} & (0.001 - 0.0035 \text{ Mpc}^{-1}) \\
 P_2(k) &< 2.2 \times 10^{-10} & (0.0035 - 0.012 \text{ Mpc}^{-1}) \\
 P_3(k) &< 3.4 \times 10^{-10} & (0.012 - 0.04 \text{ Mpc}^{-1}) \\
 P_4(k) &< 3.4 \times 10^{-9} & (0.04 - 0.14 \text{ Mpc}^{-1}) \\
 P_5(k) &< 1.5 \times 10^{-9} & (0.14 - 0.5 \text{ Mpc}^{-1})
 \end{aligned} \tag{5.11}$$

Like for the CDI mode, the ACT DR4 data has little impact at large scales, but improves the upper limits in the final two bins by 20%. At all scales the primordial power is below the adiabatic best-fit, but the levels are comparable in the smallest-scale bin.




Figure 5.10: The 95% upper limits on the neutrino density isocurvature (NDI) primordial power spectrum as a function of wavenumber, k (in Mpc^{-1}), for *Planck* compared to *ACT+Planck*. The best-fitting adiabatic power is shown for comparison.

In Figure 5.11 we show the posterior distributions for all the sampled parameters. The best-fitting adiabatic power amplitudes are reduced when isocurvature is included, but there is almost no change in the other cosmological parameters, compared to the adiabatic-only case. Without the ACT data, the Hubble constant and baryon density are slightly increased due to their correlation with the isocurvature amplitudes.

Similar to the CDI model, there is some preference for non-zero isocurvature in the final bin, P_{II}^5 , and the error bar is roughly halved with the inclusion of the ACT data. There is also a mild preference for non-zero isocurvature in the third bin, P_{II}^3 , driven by the *Planck* data; this was also noted in [Planck Collaboration et al. \(2020e\)](#). We explore the phenomenological behavior of this non-zero isocurvature model in more detail by examining the power spectra of the best-fitting model which has the 95% upper level of isocurvature power. The adiabatic and isocurvature components of this model is plotted in Figure 5.12, and the residuals compared to the data are




Figure 5.11: Neutrino density isocurvature (NDI) constraints: posterior distributions for the amplitudes of the power-law adiabatic power spectrum at two reference scales $P_{\mathcal{R}\mathcal{R}}^1$ and $P_{\mathcal{R}\mathcal{R}}^2$, Hubble constant H_0 , baryon density parameter $\Omega_b h^2$, cold dark matter density parameter $\Omega_c h^2$ and the optical depth to reionization τ , as well as the binned neutrino density isocurvature amplitudes in five bins, $P_{\mathcal{I}\mathcal{I}}^1 \dots P_{\mathcal{I}\mathcal{I}}^5$.

shown in Figure 5.13 for the TT spectrum (along with the CDI model). The non-zero isocurvature model has more TT power at $\ell \sim 2000$ than the pure-adiabatic best-fit; in TE and EE the models are very similar.




Figure 5.12: The contribution of cold dark matter and neutrino density isocurvature (scaled up by a factor of 100) relative to the adiabatic temperature and polarization power spectra, for the model with 95% upper limits on the amount of isocurvature.

We test how much better this model fits the data, by looking at the goodness of fit compared to the best-fitting adiabatic model. The χ^2 for both models is given in table 5.1, for each element of the data. The dominant impact comes from the *Planck* data (Plik-lite), which has an improvement of $\Delta\chi^2 = 6$ for the NDI model. This is not statistically significant, however, since the model has 5 additional degrees of freedom compared to the adiabatic-only case. As noted in [Planck Collaboration et al. \(2020e\)](#), this model also has some phenomenological similarity to an enhancement in gravitational lensing in the spectrum, which was seen with the *Planck* data at 2–3 σ significance.




Figure 5.13: Top: CMB temperature power spectrum for the best-fitting adiabatic model, together with the ACT DR4 data points and compared to models with the 95% upper limit of NDI and CDI perturbations. The models with isocurvature have a modest excess in the $\ell \sim 2000$ region. Bottom: difference between these models, and residuals of the ACT DR4 data with each model.

Likelihood	Adiabatic	NDI 5 bins	Difference
Planck 2018 lowTT	22.52	23.13	0.62
Planck 2018 lowEE	395.82	395.76	-0.06
Planck 2018 highTTTEEE (Plik-lite)	587.69	581.75	-5.94
Planck 2018 lensing	9.06	9.43	0.37
ACT DR4 TTTEEE($\ell > 1800$ for TT)	238.96	238.17	-0.79
Total χ^2	1254.05	1248.24	-5.81

Table 5.1: The difference in χ^2 (or $-2 \ln L$) for the best-fitting adiabatic model compared to the the best-fitting model with non-zero isocurvature. The overall improvement in χ^2 is only 5.8 for 5 extra degrees of freedom for the NDI model, so the adiabatic model is still statistically preferred. The *Planck* data drives the improvement in χ^2 .

We find 95% upper limits on the contributions of NDI perturbations to the TT, TE and EE power spectra to be $\alpha < 1.7\%$, 1.8% , and 0.7% respectively. Like for the CDI mode, the effect of NDI isocurvature at the upper level allowed by ACT and *Planck* on the present day matter power spectrum, $P(k)$, is shown in Figure 5.9. The contribution from isocurvature is well below the total power spectrum.

5.4 Discussion

The assumption made about primordial scalar perturbations in the standard Λ CDM model is that they are Gaussian, adiabatic, and that their power spectrum can be described by a power-law. This simple form of perturbations provides an excellent fit to the *Planck* data, and so any deviation must be small. In this paper we have explored whether there is any evidence from current data for such deviations, in the form of departures from a power-law primordial power spectrum, or departures from purely adiabatic perturbations over a range of scales. With the ACT DR4 data we were able to extend the range of scales probed by CMB data, out to $k \sim 0.3 - 0.4 \text{ Mpc}^{-1}$. We find no evidence for a break in the power-law, or a running index, or for non-zero isocurvature perturbations in the range $0.001 < k < 0.5 \text{ Mpc}^{-1}$. Within the broken power-law model we find a mild preference for a break at $k = 0.1 \text{ Mpc}^{-1}$, for a transition between $n_s(k < 0.1) = 0.965$ to an $n_s(k > 0.1) = 1$, but overall the power-law model is still statistically preferred. Similarly, within the non-zero isocurvature model there is a mild preference for non-zero neutrino density isocurvature at $k > 0.002 \text{ Mpc}^{-1}$, but the model is not preferred compared to a pure-adiabatic model. The addition of ACT DR4 data tightens the upper limits compared to *Planck* by $\sim 30\%$ at the smaller scales. New data from ACT, and from future Simons Observatory and CMB-S4 data, will be of great interest in better scrutinizing these extended models, both with temperature and polarization CMB anisotropy data.

Chapter 6

Forecasts for constraining primordial isocurvature with the Simons Observatory

The CMB gives us a window into the early universe at the time when photons were released at last scattering, as well as containing information about the photons' interactions during their journeys to our telescopes. The six-parameter Λ CDM model of the universe has been constrained with increasing precision by the *WMAP* (Hinshaw et al., 2013) and *Planck* (Planck Collaboration et al., 2020c,b) satellites, and by ground-based experiments including the Atacama Cosmology Telescope (Aiola et al., 2020; Choi et al., 2020) and the South Pole Telescope (Dutcher et al., 2021).

The primordial perturbations that seed the large scale structure we see in the universe today can theoretically be sourced by contributions from both adiabatic and isocurvature fluctuations. The standard Λ CDM model with single field inflation predicts solely adiabatic fluctuations in the early universe: the overdensities of the different components are proportional to one another and a region that has a radiation overdensity has a corresponding overdensity in dark matter, baryons, neutrinos

etc. In contrast, for isocurvature fluctuations an overdensity of one component, such as radiation, corresponds to an underdensity of another component, say cold dark matter, such that the sum of these components remains constant.

The *Planck* satellite has constrained the contribution of non-adiabatic fluctuations to being less than 1.3% and 1.7% of the CMB temperature power spectrum over a multipole range of $2 \leq \ell \leq 2500$ at 95% confidence level for cold dark matter density and neutrino density isocurvature respectively (Planck Collaboration et al., 2020e), assuming a generally correlated power-law model for the isocurvature primordial power spectrum. In Chapter 5 we showed how these constraints were improved with data from ACT DR4, if we relaxed the assumption of power-law fluctuations. The Simons Observatory (SO) is a new ground-based CMB experiment that is under construction in the Atacama Desert in Chile (Galitzki et al., 2018; Ade et al., 2019), with observations planned to begin in 2023–24, with a wide range of science goals, from precision cosmology measurements to extragalactic as well as Milky Way science (Ade et al., 2019; Hensley et al., 2022). SO consists of a large aperture telescope (LAT), as well as three small aperture telescopes (SATs). With the LAT it will observe the CMB intensity and polarization with ten times lower noise and at five times smaller angular scales than *Planck*, enabling us to better constrain any contributions from isocurvature from the early universe at small angular scales. The CMB E-mode polarization will also be well-measured by the SO SATs at large angular scales, leading to additional constraining power on the isocurvature modes.

In this chapter we focus on the constraining power that the SO LAT will have for primordial cold dark matter and neutrino density isocurvature (CDI and NDI) perturbations. Like in Chapter 5, we use a binned model for the isocurvature perturbations, instead of the power law model used for the *Planck* analysis, to assess the possible isocurvature amplitude at different scales. This chapter starts with a description of SO and *Planck* mock data and likelihoods in §6.1, together with a description of the

isocurvature model. We present the results of our forecasts in §6.2 and conclude in §6.3.

6.1 Methodology

6.1.1 Simulated Simons Observatory data

The SO LAT will have a six meter dish and will observe at least 40% of the sky with arcminute angular resolution. It will have six frequency bands from 27 to 280 GHz to enable robust removal of foregrounds, with resolution from 7.4' to 0.9'. We use the publicly available SO multi-frequency likelihood code `MFLike`¹, as well as power spectra simulated for the 93, 145 and 225 GHz auto and cross spectra for TT, TE, and EE². We summarize the main features of these simulated data here, and refer to Ade et al. (2019) for further details.

The noise angular power spectrum assumed in `MFLike` is modeled to include contributions from atmospheric or instrumental $1/f$ noise as well as instrumental white noise, described in Ade et al. (2019) by

$$N_\ell = N_{\text{red}} \left(\frac{\ell}{\ell_{\text{knee}}} \right)^{\alpha_{\text{knee}}} + N_{\text{white}}. \quad (6.1)$$

The first term on the right hand side describes the red noise and the second term is the detector white noise. This includes the impact at small scales from the instrument beam (Knox, 1997; Tegmark, 1997), such that

$$N_{\text{white}} = w^{-1} e^{\theta_b^2 \ell(\ell+1)/2}. \quad (6.2)$$

¹https://github.com/simonsobs/LAT_MFLike

²https://portal.nersc.gov/cfs/sobs/users/MFLike_data/




Figure 6.1: Noise levels in temperature (top) and E-mode polarization (bottom) for the *Planck* and Simons Observatory experiments. The blue curve shows the noise curve for the Simons Observatory Large Aperture Telescope (LAT), while the orange curve shows the noise for the Small Aperture Telescopes (SATs). The *Planck* white noise curve is shown for comparison in grey. The theory curve for the CMB power spectrum is shown in black. The LAT has strong signal-to-noise on smaller angular scales (larger multipoles ℓ) than *Planck*, while the SAT has lower noise on similar angular scales. These noise levels can be combined with the cosmic variance term to obtain the uncertainty on the CMB power spectrum.

Here $w^{-1} = (4\pi s^2)(t_{\text{obs}})$, where s is the detector sensitivity in $\mu K s^{1/2}$, and t_{obs} is the observation time. The full width at half maximum of the Gaussian beam is given by $\theta_{\text{FWHM}} = \sqrt{8 \ln 2} \theta_b$. These noise curves are shown in Figure 6.1 for the coadded 93 and 145 GHz science channels, together with the *Planck* 143 GHz noise curves, for temperature and E-mode polarization. The expected effect of the atmosphere can be seen in the SO noise at low ℓ (large scales), and the effects of the instrument beam due to limited angular resolution at high ℓ (small scales). The *Planck* noise, with a larger beam but no atmospheric contamination, becomes larger than the temperature signal at $\ell \gtrsim 2000$ for temperature and $\ell \gtrsim 1000$ for E-mode polarization. In contrast,

the SO LAT noise is larger than the TT power spectrum for $\ell \gtrsim 3000$ and the EE power spectrum for $\ell \gtrsim 2000$, enabling CMB measurement on smaller scales than *Planck*. Given these properties, the optimal cosmological constraints are expected to be derived from a combination of *Planck* and SO data. In our joint forecasts for *Planck* and SO we remove the SO TT power at $\ell < 1800$ to avoid double counting modes that are well-measured by both *Planck* and SO.

To create the inputs for **MFLike**, lensed CMB and foreground signal maps are simulated by drawing Gaussian realizations from a best-fitting Λ CDM theory spectrum, plus a model for the multi-frequency spectra that includes foregrounds, following [Dunkley et al. \(2013\)](#). The temperature power spectrum includes contributions from the kinetic and thermal Sunyaev-Zel'dovich (SZ) effects, radio galaxies as Poisson-distributed point sources, clustered dusty star forming point source galaxies as part of the cosmic infrared background (CIB), the cross-correlation between the tSZ and CIB, and emission from Galactic dust. The E-mode polarization power spectrum and the TE cross spectrum include polarized emission from Galactic dust and from radio galaxies. The templates for these components have 14 foreground parameters. A Gaussian noise realization at each frequency is added to the simulated signal maps, and spectra and analytic covariance matrices are estimated using pseudo- $C\ell$ methods as in [Li et al. \(2021\)](#). These are then used to form a Gaussian likelihood for the multi-frequency power spectra. Here, the error on the angular power spectrum at each frequency approximately scales as

$$\sigma(C_\ell) = \sqrt{\frac{2}{(2\ell+1)f_{\text{sky}}}} (C_\ell + N_\ell), \quad (6.3)$$

where the first term in Eqn. 6.3 captures the contribution of cosmic variance, and the second term includes the effect of detector noise and the beam of the instrument.

This expression is modified in practice to account for the mode coupling matrix that is introduced when estimating spectra on the cut sky (Alonso et al., 2019).

In this chapter we do not include forecasts for the SATs (Ali et al., 2020), leaving this to future work. The three SATs have diameters of 0.5 m and will observe larger scales in polarization (with angular resolution of 10-91') on approximately 10% of the sky. Figure 6.1 includes the SAT noise curves for comparison; they have low noise in polarization on large scales, but the beam cuts off at lower ℓ than for *Planck*.

6.1.2 Simulated Planck data

We use modified versions of our Planck-lite-py³ and Planck-low-py⁴ likelihood codes, described in Chapters 2 and 3, where we replace the CMB power spectrum with the Planck best-fitting Λ CDM power spectrum in each bin, and use the actual errors and covariances for *Planck*. This likelihood gives constraints that are consistent with the full *Planck* likelihood for the Λ CDM parameters but has best-fitting values of zero for the binned isocurvature power.

6.1.3 Isocurvature model

Following Chapter 5, we use a model with 5 bins equally spaced in $\log(k)$ and allow the isocurvature power to vary independently in these bins, as $P_{\mathcal{I}\mathcal{I}}^{(i)}$, for $i = 1, \dots, 5$. We set the correlation between the isocurvature and adiabatic modes to zero. We allow one isocurvature mode to be present at a time, either cold dark matter density isocurvature (CDI) or neutrino density isocurvature (NDI). This isocurvature model is described in more detail in the previous chapter. The primordial adiabatic power spectrum is still modeled as a power law, but is parameterized by its amplitude at two reference wavenumbers $P_{RR}^{(1)}$ at $k = 0.002 \text{ Mpc}^{-1}$ and $P_{RR}^{(2)}$ at $k = 0.1 \text{ Mpc}^{-1}$. These

³<https://github.com/heatherprince/planck-lite-py>

⁴<https://github.com/heatherprince/planck-low-py>

can be converted to the usual amplitude, A_s , defined at the pivot scale $k = 0.05 \text{ Mpc}^{-1}$ and spectral index, n_s , parameters.

We use MCMC methods as in Chapter 5 to simultaneously sample the 11 cosmological parameters (two adiabatic power-law parameters, five isocurvature power parameters, the baryon density, CDM density, Hubble constant, and optical depth to reionization) and 14 foreground parameters.

6.2 Results




Figure 6.2: Forecast upper limits on the isocurvature power spectrum from *Planck* (grey) compared to *Planck* + SO LAT (red) for CDM isocurvature (left) and neutrino density isocurvature (right), for a set of five wavenumber bins.

We show the expected 95% upper bound on the primordial isocurvature power in the five k -bins in Figure 6.2, for the two types of isocurvature. The CDM density isocurvature 95% upper bounds are shown in the left panel, and the neutrino density isocurvature constraints in the right panel. Apart from the first bin at largest scales, which stays about the same, the upper bounds are reduced by 30% to 78% when the SO Large Aperture Telescope data is added to *Planck*. The greatest improvement occurs for the isocurvature power on small scales where the LAT noise is significantly better than for *Planck*.




Figure 6.3: Forecast posterior distributions for the Λ CDM cosmological parameters, showing how SO LAT data (red) are expected to improve constraints compared *Planck* data alone (black) in the case that the perturbations are not assumed to be perfectly adiabatic. The left panel shows the constraints when binned cold dark matter density isocurvature is allowed to vary, and the right panel show the constraints with neutrino density isocurvature. The left and right panels are very similar to one another because small amounts of isocurvature are preferred.

The expected parameter constraints for the other cosmological parameters are shown in Figure 6.3. Including the LAT improves the constraints for all the parameters, by removing correlations with the isocurvature parameters. Once SO simulated data are included, the binned isocurvature parameters are not strongly degenerate with the cosmological parameters or the foreground parameters.

6.3 Discussion

In this chapter we explored the additional constraining power on primordial isocurvature perturbations that power spectra from SO can offer, in combination with data from *Planck*. In this work we considered only data anticipated from the LAT, which will improve small-scale measurements. We leave for future work the addition of larger-scale data expected from the SO SATs, and improved CMB lensing data derived from the LAT. We found that the SO LAT power spectrum data is expected to significantly tighten constraints on any isocurvature contribution, in the approximate

range $0.01 < k < 0.4 \text{ Mpc}^{-1}$, or alternatively could show up a non-zero isocurvature signal if the underlying model has some departure from pure adiabaticity.

Chapter 7

Conclusion

It's still magic even if you know how it's done.

Terry Pratchett *A Hat Full of Sky*

The cosmic microwave background gives us a remarkably clean window into the early universe, while also carrying information about the CMB photons' journeys from recombination to our telescopes, because some of the photons have been deflected by massive galaxies and clusters, Thomson scattered off electrons during reionization, or inverse-Compton scattered to higher energies off hot electrons or moving electrons in galaxy clusters. CMB experiments that observe at millimeter wavelengths also detect emission from distant dusty galaxies and active galactic nuclei, as well as thermal radiation from dust within our Galaxy and other foreground radiation at those frequencies.

The *Planck* satellite used all-sky measurements of the CMB temperature and polarization in a nine frequency bands ranging from 30 GHz to 857 GHz to study which model of the universe is preferred by the CMB and to measure its parameters to high precision, while also improving our understanding of CMB foregrounds so that they can be marginalized over in the cosmological analysis (Planck Collaboration et al., 2020b). The *Planck* 2018 legacy data release provides cosmic variance limited mea-

surements of the CMB temperature power spectrum to $\ell \approx 2000$, as well as E-mode polarization measurements down to the quadrupole $\ell = 2$, corresponding to large scales which can only be observed from space because of the atmospheric noise encountered on large scales in ground-based observations (Planck Collaboration et al., 2020a). The *Planck* data and likelihoods will be used by current and future cosmologists for years to come, as it serves as both a consistency check and a complement to ground-based CMB data, as well as being used in combination with other cosmological probes to place even tighter constraints on the cosmological parameters. For this reason it is useful to have simple, concise, portable versions of the *Planck* likelihoods that can be easily used by astronomers in different fields.

In Chapter 2 we applied the Massively Optimized Parameter Estimation and Data compression technique (MOPED) to the public *Planck* 2015 temperature likelihood, reducing the dimensions of the data space to one number per parameter of interest. We presented CosMOPED, a lightweight and convenient compressed likelihood code implemented in Python. In doing so we showed that the $\ell < 30$ *Planck* temperature likelihood can be approximated by two Gaussian distributed data points, which allows us to replace the map-based low- ℓ temperature likelihood by a simple Gaussian likelihood. We did not explicitly use the large-scale polarization data in CosMOPED, instead imposing a prior on the optical depth to reionization derived from these data. We showed that the Λ CDM parameters recovered with CosMOPED are consistent with the uncompressed likelihood to within 0.1σ , and tested that a 7-parameter extended model performs similarly well.

In Chapter 3 we presented `Planck-low-py`, a binned low- ℓ temperature and *E*-mode polarization likelihood, as an option to facilitate ease of use of the *Planck* 2018 large-scale data in joint-probe analysis and forecasting. It is written in Python and compresses the $\ell < 30$ temperature and polarization angular power spectra information from *Planck* into two log-normal bins in temperature and seven in polarization.

These angular scales constrain the optical depth to reionization and provide a lever arm to constrain the tilt of the primordial power spectrum. Using `Planck-low-py`, we showed that cosmological constraints on Λ CDM model parameters, and on extensions to Λ CDM including running and isocurvature, agree with those derived with the full `Commander` and `SimA11` likelihoods from the *Planck* legacy release.

The recent limit on the tensor-to-scalar ratio from polarization B-modes observed using the BICEP and Keck Array telescopes at the South Pole (with $r < 0.036$ at 95% confidence) puts pressure on early universe models. We used the publicly-available multi-frequency likelihood to examine various assumptions made in the foreground modeling, including relaxing the assumption of power-law dust, varying the dust temperature, removing the synchrotron index prior, or removing the *Planck* data altogether, and found these to have a negligible impact on r . The error inflation on r due to foreground uncertainty is only $\sim 10\%$. We then estimated foreground-marginalized CMB B-mode bandpower amplitudes and found their distribution to be well fit by a multivariate offset-lognormal distribution, with minimal foreground-induced covariance between bins (only done so far for BICEP3+*Planck*+*WMAP* data). We constructed a marginalized likelihood with no nuisance parameters, serving as a method demonstration for simpler joint-probe cosmological analysis. The tensor-to-scalar ratio inferred from this ‘BK18-lite’ likelihood matches that using the public BK18 likelihood.

Other ground-based CMB experiments such as the Atacama Cosmology Telescope and the South Pole Telescope continue to map the CMB sky at high resolution and with increasing sensitivity, measuring the CMB temperature and polarization on small scales that were not accessible with *Planck*. Upcoming experiments such as the Simons Observatory will continue this work, producing low-noise maps of a large fraction of the sky with resolution of a few arcminutes. This enables us to constrain extensions to Λ CDM that affect scales that were not accessible to *Planck*.

In Chapter 5 we used ACT and SPT data in combination with *Planck* data to study scenarios beyond adiabatic, power law initial conditions in the universe. We first explored an adiabatic model with a broken power law instead of a single power law, to see if this model could account for the preference for scale-invariance ($n_s = 1$) shown by ACT and SPT on small scales and the slightly red-tilted spectrum preferred by *Planck* on large scales ($n_s < 1$). We found that a preferred scale for transitioning between spectral indices emerged, but the improvement in the goodness-of-fit did not justify the inclusion of extra parameters.

We also used ACT and *Planck* data to constrain the level of isocurvature allowed by the data in five wavenumber bins, improving the width of the constraints on small scales with the ACT data. In Chapter 6 we performed forecasts for the upcoming Simon’s observatory’s ability to further improve isocurvature constraints at small scales.

Throughout this dissertation we have explored ways to use CMB satellite and ground-based experiments to better constrain our model of how the universe began and evolved, and ways to make CMB data more accessible to other scientific communities.

Bibliography

Abazajian, K. N., Arnold, K., Austermann, J., et al. 2015, Astroparticle Physics, 63, 66

Abazajian, K. N., Adshead, P., Ahmed, Z., et al. 2016, arXiv e-prints, arXiv:1610.02743

Ade, P., Aguirre, J., Ahmed, Z., et al. 2019, J. Cosmology Astropart. Phys., 2019, 056

Ade, P. A. R., Ahmed, Z., Amiri, M., et al. 2021, Phys. Rev. Lett., 127, 151301

—. 2022, ApJ, 927, 77

Aiola, S., Calabrese, E., Maurin, L., et al. 2020, J. Cosmology Astropart. Phys., 2020, 047

Alam, S., Aubert, M., Avila, S., et al. 2021, Phys. Rev. D, 103, 083533

Albrecht, A., & Steinhardt, P. J. 1982, Phys. Rev. Lett., 48, 1220

Ali, A. M., Adachi, S., Arnold, K., et al. 2020, Journal of Low Temperature Physics, 200, 461

Alonso, D., Sanchez, J., Slosar, A., & LSST Dark Energy Science Collaboration. 2019, MNRAS, 484, 4127

Alsing, J., & Wandelt, B. 2018, MNRAS, 476, L60

Bardeen, J. M. 1980, Phys. Rev. D, 22, 1882

Barkats, D., Aikin, R., Bischoff, C., et al. 2014, ApJ, 783, 67

BICEP2 and Keck Array Collaborations, Ade, P. A. R., Ahmed, Z., et al. 2015, ApJ, 811, 126

BICEP2 Collaboration, Ade, P. A. R., Aikin, R. W., et al. 2014a, ApJ, 792, 62

—. 2014b, Phys. Rev. Lett., 112, 241101

BICEP2 Collaboration, Keck Array Collaboration, Ade, P. A. R., et al. 2018, Phys. Rev. Lett., 121, 221301

BICEP2/Keck Collaboration, Planck Collaboration, Ade, P. A. R., et al. 2015, *Phys. Rev. Lett.*, 114, 101301

Bird, S., Viel, M., & Haehnelt, M. G. 2012, *MNRAS*, 420, 2551

Blas, D., Lesgourgues, J., & Tram, T. 2011a, *J. Cosmology Astropart. Phys.*, 2011, 034

—. 2011b, *J. Cosmology Astropart. Phys.*, 2011, 034

Bond, J. R., & Efstathiou, G. 1984, *ApJ*, 285, L45

Bond, J. R., Jaffe, A. H., & Knox, L. 2000, *ApJ*, 533, 19

Bridle, S. L., Lewis, A. M., Weller, J., & Efstathiou, G. 2003, *MNRAS*, 342, L72

Bucher, M., Moodley, K., & Turok, N. 2000, *Phys. Rev. D*, 62, 083508

Calabrese, E., Hlozek, R. A., Battaglia, N., et al. 2013, *Phys. Rev. D*, 87, 103012

Choi, S. K., Hasselfield, M., Ho, S.-P. P., et al. 2020, *J. Cosmology Astropart. Phys.*, 2020, 045

Clowe, D., Bradač, M., Gonzalez, A. H., et al. 2006, *ApJ*, 648, L109

Das, S., Louis, T., Nolta, M. R., et al. 2014, *J. Cosmology Astropart. Phys.*, 2014, 014

Davis, M. 1997, in *Critical Dialogues in Cosmology*, ed. N. Turok, 13

Davis, M., Huchra, J., Latham, D. W., & Tonry, J. 1982, *ApJ*, 253, 423

Dicke, R. H., Peebles, P. J. E., Roll, P. G., & Wilkinson, D. T. 1965, *ApJ*, 142, 414

Dodelson, S. 2003, *Modern cosmology* (Academic Press)

Dunkley, J., Calabrese, E., Sievers, J., et al. 2013, *J. Cosmology Astropart. Phys.*, 2013, 025

Dutcher, D., Balkenhol, L., Ade, P. A. R., et al. 2021, *Phys. Rev. D*, 104, 022003

Efstathiou, G. 2006, *MNRAS*, 370, 343

Efstathiou, G., & Bond, J. R. 1986, *MNRAS*, 218, 103

—. 1987, *MNRAS*, 227, 33P

Einstein, A. 1915, *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften* (Berlin), Seite 844-847., 844

—. 1917, *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften* (Berlin), 142

Enqvist, K., Kurki-Suonio, H., & Välimäki, J. 2002, Phys. Rev. D, 65, 043002

Eriksen, H. K., Jewell, J. B., Dickinson, C., et al. 2008, ApJ, 676, 10

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, Publications of the Astronomical Society of the Pacific, 125, 306

Friedmann, A. 1922, Zeitschrift fur Physik, 10, 377

Galitzki, N., Ali, A., Arnold, K. S., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10708, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX, ed. J. Zmuidzinas & J.-R. Gao, 1070804

Goodman, J., & Weare, J. 2010, Communications in Applied Mathematics and Computational Science, 5, 65

Gordon, C., & Lewis, A. 2003, Phys. Rev. D, 67, 123513

Graff, P., Hobson, M. P., & Lasenby, A. 2011, MNRAS, 413, L66

Gupta, S., & Heavens, A. F. 2002, MNRAS, 334, 167

Guth, A. H. 1981, Phys. Rev. D, 23, 347

Guth, A. H., & Weinberg, E. J. 1983, Nuclear Physics B, 212, 321

Hamimeche, S., & Lewis, A. 2008, Phys. Rev. D, 77, 103013

Harrison, E. R. 1970, Phys. Rev. D, 1, 2726

Hazumi, M., Ade, P. A. R., Adler, A., et al. 2020, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 11443, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 114432F

Heavens, A., Panter, B., Jimenez, R., & Dunlop, J. 2004, Nature, 428, 625

Heavens, A. F., Jimenez, R., & Lahav, O. 2000, MNRAS, 317, 965

Heavens, A. F., Sellentin, E., de Mijolla, D., & Vianello, A. 2017, MNRAS, 472, 4244

Hensley, B. S., Clark, S. E., Fanfani, V., et al. 2022, ApJ, 929, 166

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

Hoffman, M. D., & Gelman, A. 2011, arXiv e-prints, arXiv:1111.4246

Howlett, C., Lewis, A., Hall, A., & Challinor, A. 2012, J. Cosmology Astropart. Phys., 2012, 027

Hubble, E. 1929, Proceedings of the National Academy of Science, 15, 168

Ijjas, A., & Steinhardt, P. J. 2018, *Classical and Quantum Gravity*, 35, 135004

Kamionkowski, M., & Kovetz, E. D. 2016, *ARA&A*, 54, 227

Kitajima, N., Langlois, D., Takahashi, T., & Yokoyama, S. 2017, *J. Cosmology Astropart. Phys.*, 2017, 042

Knox, L. 1997, *ApJ*, 480, 72

—. 1999, *Phys. Rev. D*, 60, 103516

La Posta, A., Louis, T., Garrido, X., & Hill, J. C. 2021, arXiv e-prints, arXiv:2112.10754

Langlois, D. 1999, *Phys. Rev. D*, 59, 123512

Lee, A., Abitbol, M. H., Adachi, S., et al. 2019, in *Bulletin of the American Astronomical Society*, Vol. 51, 147

Lemaître, G. 1927, *Annales de la Société Scientifique de Bruxelles*, 47, 49

Lemaître, G. 1931, *MNRAS*, 91, 483

Lewis, A. 2013, *Phys. Rev.*, D87, 103529

Lewis, A., & Bridle, S. 2002, *Phys. Rev.*, D66, 103511

Lewis, A., & Challinor, A. 2006, *Physics Reports*, 429, 1

Lewis, A., Challinor, A., & Lasenby, A. 2000, *Astrophys. J.*, 538, 473

Li, Z., Louis, T., Calabrese, E., et al. 2021, arXiv e-prints, arXiv:2112.13839

Linde, A. D. 1982, *Physics Letters B*, 108, 389

Lyth, D. H., & Wands, D. 2003, *Phys. Rev. D*, 68, 103516

Ma, C.-P., & Bertschinger, E. 1995, *ApJ*, 455, 7

MacKay, D. J. C. 2003, *Information theory, inference, and learning algorithms* (Cambridge: Cambridge University Press)

Martinez, V. J., Pons-Borderia, M.-J., Moyeed, R. A., & Graham, M. J. 1998, *MNRAS*, 298, 1212

Mather, J. C., Hauser, M. G., Bennett, C. L., et al. 1992, in *American Institute of Physics Conference Series*, Vol. 245, *American Institute of Physics Conference Series*, ed. H. J. Haubold & R. K. Khanna, 266–278

Misner, C. W. 1969, *Physical Review Letters*, 22, 1071

Moncelsi, L., Ade, P. A. R., Ahmed, Z., et al. 2020, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 11453, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 1145314

Mukhanov, V. F., Feldman, H. A., & Brandenberger, R. H. 1992, *Phys. Rep.*, 215, 203

Ostriker, J. P., & Peebles, P. J. E. 1973, *ApJ*, 186, 467

Panter, B., Jimenez, R., Heavens, A. F., & Charlot, S. 2007, *MNRAS*, 378, 1550

Peebles, P. J. E. 1982, *ApJ*, 263, L1

Peebles, P. J. E., & Yu, J. T. 1970, *ApJ*, 162, 815

Penzias, A. A., & Wilson, R. W. 1965, *ApJ*, 142, 419

Perlmutter, S., Aldering, G., Goldhaber, G., et al. 1999, *ApJ*, 517, 565

Planck Collaboration. 2020, *A&A*, 643, A42

Planck Collaboration, Aghanim, N., Arnaud, M., et al. 2016a, *A&A*, 594, A11

Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016b, *A&A*, 594, A13

Planck Collaboration, Adam, R., Aghanim, N., et al. 2016c, *A&A*, 596, A108

Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2016d, *A&A*, 586, A133

Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020a, *A&A*, 641, A1

—. 2020b, *A&A*, 641, A5

—. 2020c, *A&A*, 641, A6

—. 2020d, *A&A*, 641, A8

Planck Collaboration, Akrami, Y., Arroja, F., et al. 2020e, *A&A*, 641, A10

Planck Collaboration, Akrami, Y., Ashdown, M., et al. 2020f, *A&A*, 641, A11

Prince, H., & Dunkley, J. 2019, *Phys. Rev. D*, 100, 083502

Protopapas, P., Jimenez, R., & Alcock, C. 2005, *MNRAS*, 362, 460

Reichardt, C., Jimenez, R., & Heavens, A. F. 2001, *MNRAS*, 327, 849

Riess, A. G., Filippenko, A. V., Challis, P., et al. 1998, *AJ*, 116, 1009

Riess, A. G., Yuan, W., Macri, L. M., et al. 2021, arXiv e-prints, arXiv:2112.04510

Robertson, H. P. 1935, *ApJ*, 82, 284

Rubin, V. C., & Ford, Jr., W. K. 1970, *ApJ*, 159, 379

Sachs, R. K., & Wolfe, A. M. 1967a, *ApJ*, 147, 73

—. 1967b, *ApJ*, 147, 73

Scaramella, R., Vettolani, G., & Zamorani, G. 1991, *ApJ*, 376, L1

Seljak, U., & Zaldarriaga, M. 1997, *Phys. Rev. Lett.*, 78, 2054

Silk, J. 1968, *ApJ*, 151, 459

Simpson, J. A., & Weiner, E. S. C. 1989, *The Oxford English Dictionary* (Oxford: Clarendon Press)

Smith, R. E., Peacock, J. A., Jenkins, A., et al. 2003, *MNRAS*, 341, 1311

Smoot, G. F., Bennett, C. L., Kogut, A., et al. 1992, *ApJ*, 396, L1

Spergel, D. N., Verde, L., Peiris, H. V., et al. 2003, *ApJS*, 148, 175

SPIIDER Collaboration, Ade, P. A. R., Amiri, M., et al. 2021, *arXiv e-prints*, arXiv:2103.13334

Steinhardt, P. J., & Turok, N. 2002, *Science*, 296, 1436

Sunyaev, R. A., & Zeldovich, Y. B. 1970, *Ap&SS*, 7, 3

—. 1972, *Comments on Astrophysics and Space Physics*, 4, 173

Tegmark, M. 1997, *Phys. Rev. D*, 56, 4514

Tegmark, M., & de Oliveira-Costa, A. 2001, *Phys. Rev. D*, 64, 063001

Tegmark, M., Taylor, A. N., & Heavens, A. F. 1997, *ApJ*, 480, 22

Torrado, J., & Lewis, A. 2020, *arXiv e-prints*, arXiv:2005.05290

Tristram, M., Banday, A. J., Górski, K. M., et al. 2021, *A&A*, 647, A128

Verde, L., Peiris, H. V., Spergel, D. N., et al. 2003, *ApJS*, 148, 195

Walker, A. G. 1937, *Proceedings of the London Mathematical Society*, 42, 90

Zablocki, A., & Dodelson, S. 2016, *Phys. Rev. D*, 93, 083525

Zaldarriaga, M. 1997, *Phys. Rev. D*, 55, 1822

Zeldovich, Y. B. 1972, *MNRAS*, 160, 1P

Zuntz, J., Paterno, M., Jennings, E., et al. 2015, *Astronomy and Computing*, 12, 45

Zwicky, F. 1933, *Helvetica Physica Acta*, 6, 110