



Constraining Primordial Gravitational Waves Using Present and Future CMB Experiments

The Harvard community has made this
article openly available. [Please share](#) how
this access benefits you. Your story matters

Citation	Buza, Victor. 2019. Constraining Primordial Gravitational Waves Using Present and Future CMB Experiments. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.
Citable link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:42029563
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA

Constraining Primordial Gravitational Waves Using Present and Future CMB Experiments

A dissertation presented
by

VICTOR BUZA

to

THE DEPARTMENT OF PHYSICS

in partial fulfillment of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in the subject of

PHYSICS

Harvard University
Cambridge, Massachusetts
May 2019

©2019 – VICTOR BUZA
ALL RIGHTS RESERVED.

Constraining Primordial Gravitational Waves Using Present and Future CMB Experiments

ABSTRACT

Cosmic inflation is our current best theory for what occurred in the universe in the first instances of time. It postulates a brief period of exponential expansion in which quantum fluctuations are magnified to cosmic size and become the seeds for the growth of all structure in the Universe. Inflation makes a number of predictions, the most unique of which is the production of primordial gravitational waves (PGWs). Most of the predictions have since been tested, but the discovery of PGWs has eluded us to this day. The polarized Cosmic Microwave Background (CMB) is a powerful probe of these predictions, including the exciting possible existence of PGWs.

This thesis provides a detailed account of the development of an optimal multi-component spectral-based likelihood analysis framework for joint analyses of heterogeneous multi-frequency CMB datasets, and its subsequent use for joint analysis of BICEP/*Keck*, *Planck* and WMAP CMB polarization data to derive the tightest constraints available on PGWs, parametrized by the tensor-to-scalar ratio.

The manuscript also details the development of a spectral-based Fisher projection framework, specifically targeted towards optimizing tensor-to-scalar parameter constraints in the

presence of galactic foregrounds and gravitational lensing of the CMB, that directly uses information from current BICEP/*Keck* achieved performances, to robustly forecast the science reach of upcoming CMB-polarization endeavors. This methodology allows for rapid iteration over experimental configurations and offers a flexible way to optimize the design of future experiments given a scientific goal. We document the use of this framework to perform forecasts for the next iteration of BICEP/*Keck* instrument – BICEP Array, as well as the next generation ground-based CMB experiment – CMB-S4.

Contents

1	INTRODUCTION AND MOTIVATION	1
1.1	The Cosmic Microwave Background	2
1.2	Inflation	6
1.2.1	Fundamentals and Motivation	6
1.2.2	Matter Fields; Actions and Motions	14
1.2.3	Inflationary Dynamics and Slow Roll Parameters	22
1.2.4	Inflationary Perturbations	25
1.3	CMB Polarization	29
1.4	Thesis Outline	36
2	BICEP/KECK IN A NUTSHELL	38
2.1	The BICEP/Keck Series of Instruments	39
2.2	From Timestreams to Power Spectra	40
2.2.1	Noise Simulations	46
2.2.2	Signal Simulations	47
3	MULTI-COMPONENT LIKELIHOOD ANALYSIS FRAMEWORK	48
3.1	Motivation & Overview	50
3.2	Experiment Description & Real Bandpowers	51
3.3	Multi-component Model	54
3.4	Simulations	60
3.4.1	Signal Simulations	60
3.4.2	Noise Simulations	60
3.5	Semi-analytic Bandpower Covariance Matrix	61
3.6	Likelihood Approximation	61
4	A JOINT ANALYSIS OF BICEP2/<i>Keck</i> AND <i>Planck</i> DATA	65
4.1	Introduction	65
4.2	Maps and Power Spectra	67
4.3	Likelihood Analysis	68
4.3.1	Baseline Likelihood Results	71
4.3.2	Likelihood Variations	72
4.3.3	Likelihood Validation	77
4.4	Discussion and Conclusions	79

5	ADDING BICEP2/ <i>Keck</i> OBSERVATIONS FOR THE 2014 SEASON	81
5.1	Introduction	81
5.2	Maps and Power Spectra	82
5.3	Likelihood Analysis	85
5.3.1	Likelihood Evolution	87
5.3.2	Baseline Likelihood Results	88
5.3.3	Likelihood Variation	91
5.3.4	Likelihood Validation	93
5.4	Discussion and Conclusions	96
6	ADDING BICEP2/ <i>Keck</i> OBSERVATIONS FOR THE 2015 SEASON	97
6.1	Introduction	97
6.2	Maps and Power Spectra	99
6.3	Likelihood Analysis	99
6.3.1	Likelihood Evolution	101
6.3.2	Baseline Likelihood Results	104
6.3.3	Likelihood Variations	107
6.3.3.1	Analysis Variations	108
6.3.3.2	Data Selection Variations	112
6.3.4	Likelihood Validation	113
6.4	Discussion and Conclusions	116
7	PERFORMANCE-BASED FORECASTED CONSTRAINTS ON PRIMORDIAL GRAVITATIONAL WAVES	119
7.1	Fisher Optimization Framework	121
7.1.1	Fisher Formalism	121
7.1.2	Inputs	124
7.1.3	Multicomponent Theory Model	125
7.1.4	Bandpower Covariance Matrix Rescaling	128
7.2	BICEP/ <i>Keck</i> Forecasts	131
7.2.1	Fisher Constraints vs Published BICEP/ <i>Keck</i> Constraints	133
7.2.2	<i>Keck Array</i> high-frequency channel selection	135
7.2.3	BICEP Array Deployment Strategy and Forecasts	140
7.2.4	Inflationary constraints from the BICEP/ <i>Keck</i> program	147
7.3	CMB-S4 Forecasts	150
7.3.1	Experiment Specification	153
7.3.2	Delensing Treatment	156
7.3.3	Optimization and Parameter Constraints	157
7.3.4	Simulations	163
7.3.4.1	Map Noise Simulations	165
7.3.4.2	Foreground Models	165

7.3.4.3	Instrumental Systematics	167
7.3.4.4	Delensing	169
7.3.5	Analysis Methods	169
7.3.6	Results	170
8	CONCLUSION	176
APPENDIX A	GAUSSIAN VARIABLES AND EXPECTATION VALUES	180
A.1	Bandpowers	181
A.2	Bandpower covariance matrix	184
A.2.1	Variance of auto-spectrum	184
A.2.2	Variance of cross-spectrum	185
A.2.3	Covariance for uncorrelated fields	188
A.2.4	Covariance for correlated fields	190
APPENDIX B	BUILDING A SEMI-ANALYTIC BANDPOWER COVARIANCE MATRIX	192
B.1	Introduction	192
B.2	General Semi-analytic Formulation	194
B.3	Validation	202
REFERENCES		210

Listing of figures

1.1	<i>Planck</i> 2018 Map of Temperature Anisotropies	3
1.2	<i>Planck</i> 2018 Temperature Power Spectrum	5
1.3	FRW Conformal Diagram	13
1.4	Current Inflationary Constraints in the $n_s - r$ Plane	30
1.5	Thomson Scattering Generating CMB Polarization	31
1.6	E -mode Polarization	34
1.7	B -mode Polarization	36
1.8	Recent CMB Angular Power Spectrum Measurements	37
2.1	BICEP Observing Field	42
3.1	Likelihood Framework Schematic	52
4.1	<i>Planck</i> 353 GHz T , Q , and U maps	69
4.2	Single- and cross-frequency spectra between BK150 and P353	70
4.3	BKP Baseline Results	71
4.4	BKP Likelihood Variations	73
4.5	BKP Likelihood Results Including Synchrotron	75
4.6	BKP Likelihood Results Including Lensing	77
4.7	BKP Marginalized Likelihoods of Gaussian Simulations	78
4.8	BKP Marginalized Likelihoods of PSM Simulations	79
5.1	BK14 T , Q , U maps at 150 GHz	83
5.2	BK14 T , Q , U maps at 95 GHz	84
5.3	BK14 EE and BB auto- and cross-spectra between 95 & 150 GHz	85
5.4	Full set of BK14 BB auto- and cross-spectra	86
5.5	Evolution of the BKP to the BK14 baseline	88
5.6	BK14 Baseline results of the multicomponent multi-spectral likelihood analysis	89
5.7	BK14 r vs n_s constraints	90
5.8	BK14 Likelihood results for various analysis variations	92
5.9	BK14 Validation on simulations	94
5.10	BK14 Noise uncertainties for the $\ell \sim 80$ BB bandpower	95
6.1	BK15 T , Q , U maps at 95 GHz	100
6.2	BK15 T , Q , U maps at 150 GHz	101
6.3	BK15 T , Q , U maps at 220 GHz	102

6.4	Full set of BK15 BB auto- and cross-spectra	103
6.5	BK15 Effective sky fraction	104
6.6	Evolution of the BK14 to the BK15 baseline	105
6.7	BK15 Baseline results of the multicomponent multi-spectral likelihood analysis .	106
6.8	BK15 r vs n_s constraints	107
6.9	BK15 Likelihood results for various analysis variations	108
6.10	BK15 Likelihood results allowing decorrelation	111
6.11	BK15 Results with free A_L	112
6.12	BK15 Likelihood results for various data selection variations	113
6.13	BK15 Maximum likelihood validation tests	114
6.14	BK15 Validation on simulations with dust decorrelation	115
6.15	CDF tests on lensed- Λ CDM+dust+noise simulations	116
6.16	BK15 Noise uncertainties for the $\ell \sim 80$ BB bandpower	118
7.1	Fisher Framework Schematic	122
7.2	Progression of the BICEP/ <i>Keck</i> experimental program	132
7.3	BKP comparison with Fisher-derived results	136
7.4	BK14 comparison with Fisher-derived results	137
7.5	BK15 comparison with Fisher-derived results	138
7.6	<i>Keck</i> and BICEP3 map comparison	141
7.7	Atmospheric transmission and BICEP/ <i>Keck</i> bandpasses	143
7.8	BICEP Array expected B -mode power spectrum noise uncertainties	145
7.9	BICEP/ <i>Keck</i> projected sensitivity on r	146
7.10	BICEP Array r vs n_s constraints	148
7.11	Schematic timeline of the increase in CMB experiment sensitivity	151
7.12	CMB-S4 Forecasting Loop	154
7.13	Atmospheric Brightness Spectra	155
7.14	Small-aperture CMB-S4 map depth	159
7.15	CMB-S4 Optimized Constraints on r	160
7.16	CMB-S4 Bin-by-bin forecasted tensor constraints for $r=0.01$	161
7.17	CMB-S4 Forecasted r vs n_s constraints for $r = 0.01$	162
7.18	CMB-S4 Forecasted r vs n_s constraints for $r = 0$	163
7.19	CMB-S4 Forecasted r vs n_t constraints for $r = 0.01$ and $r = 0.05$	164
B.1	Direct and Semi-analytic correlation matrices, Model 1	206
B.2	Difference between direct and semi-analytic correlation matrices, Model 1	207
B.3	Direct and Semi-analytic correlation matrices, Model 2	208
B.4	Difference between direct and semi-analytic correlation matrices, Model 2	209

Listing of tables

2.1	BICEP Deployment History	41
7.1	<i>Keck Array</i> high-frequency receiver parameters as used in projections	139
7.2	<i>Keck Array</i> high-frequency receiver projections	139
7.3	BICEP Array receiver parameters as used in projections	142
7.4	SPT-3G delensing projections	144
7.5	CMB-S4 baseline configuration	162
7.6	Summary of Analysis of CMB-S4 Map-Based Simulations	171
7.7	CMB-S4 detection of significance results	172
7.8	CMB-S4 map-based simulation results for simulations containing systematics . .	174

Acknowledgments

The work in this thesis is a culmination of a very long journey. At the moment of this writing, I have been pursuing Physics for almost fifteen years, but the skills that have helped me get here have their roots much further into the past. I grew up on a farm, in the Moldovan village of Danceni, in a country and family that was struggling to make ends meet. My childhood involved a lot of hard manual labor. We did not have an irrigation system, or sewage system, or gas, or any machinery to tend to the animals, vegetables, and fruits that we relied on for food. Everything we did, we did by hand. Pursuing higher education in the US was not only far fetched, it was not even something I could have imagined at the time. What I was fortunate to have, was a mother who was as relentless as she was self-sacrificing. She was the only one around me who believed in the merit of education and always pushed me to learn. Sessions of working outside were interrupted by sessions of learning poetry and mathematics, reading history and Romanian literature, memorizing facts and get quizzed on them. Throughout these all, she kept working. She saw beyond our means to instill values that were foreign to everyone else in our circles while enduring ridicule from people around us who thought studying would land me in a mental facility. She was always the mom that she needed to be for her children. Ever demanding, ever loving, always there. I am forever indebted to my mom for believing in me and pushing me to become the best version of myself

that I could. Were it not for her, I would not have been writing these words, and would likely have shared the fate that unfortunately many other Moldovan children do. She deserves all the gratitude and love in the world, which I am happy to offer her now and always.

I dedicated most of my high-school years to learning physics. I always liked physics, even though I was not always good at it. By high school, I had learned of the existence of scientific powerhouses such as MIT, Princeton, Harvard, Caltech, etc. I also had learned that performing well at National and International Physics Olympiads was a way for someone like me to get there, and I was determined to do so. I am grateful to have had some very caring and dedicated physics teachers in high-school: Tatiana Adreyevna, Eugen Grigoriev, and Igor Evtodiev, you taught me how to think critically, the value and interplay of theory and experiment, and how to take an idea from thought to execution. For that, I thank you.

I arrived at the University of Minnesota, Duluth in 2009. As a public school, UMD offered no financial aid for international students. Letting me take all our family money, which covered just one year of schooling, was just another example of the kind of sacrifices my mother was ready to make for my education. In that year, I got my first taste of real physics. I thank Jonathan Maps, Alec Habig, and John Hiller for allowing me to pursue higher-level classes and the countless hours of discussions that followed – these have exponentially furthered my understanding of nature. I am also eternally grateful for their belief in me, which manifested itself, among many other ways, in their recommendation letters to MIT. They made my dream come true.

At MIT, I had the privilege to learn cosmology from some of the best people in the field:

Max Tegmark and Alan Guth. I am grateful to Max for allowing me to join his group as a sophomore, for teaching me how to think creatively, for the many conversations about life, and for encouraging me to pursue a Ph.D. in observational cosmology. I thank Alan for giving me the chance to learn from him, and for being so generous with his time and knowledge.

Harvard is where I became a proper scientist. The work in this thesis would not have been possible without the contribution of many people. I would like to thank my BICEP/*Keck* collaborators for many years of great scientific explorations. I have learned a tremendous amount from all of them, and have benefited from all of their expertise. I am also grateful to my CMB-S4 collaborators (Lloyd Knox, Raphael Flauger, Tom Crawford, John Carlstrom, Gil Holder, Zeeshan Ahmed, among many others) for expanding my view of cosmology and for teaching me how big projects work behind the scenes.

I would like to thank John Kovac, Clem Pryke, Jamie Bock, and Chao-Lin Kuo for creating a thriving research atmosphere, and for maintaining an incredibly rigorous standard throughout the years. To the Harvard CMB gang: Ben Racine, Marion Dierickx, James Cornelison, Tyler St. Germaine, Denis Barkats, Kirit Karkare, Irene Coyle, thank you for all of your support and banter, and for making me look forward to coming to the office.

I am indebted to Colin Bischoff, who, as a postdoc, overlapped with me for several years, and was instrumental in the development of my analysis skills. I learned a tremendous amount from him, and continue to do so to this day.

It has been an absolute honor to work alongside my advisor John Kovac. I have never met anyone with the same ability as him to state science clearly, and the more I have learned

about the intricacies of pursuing experimental physics, the more I have grown to appreciate the honest, relentless, and raw way in which he practices science. I can only hope that someday I am able to replicate this. John has been a role model for me, not just as a scientist, but as a person as well. He has been unthinkably kind and generous with his support throughout my Ph.D. and has helped me figure out life on numerous occasions. I am happy and humbled to be able to call him a friend.

I am thankful for Lisa Cacciabaudo and Jacob Barandes for putting up with my pestering questions, and for being very helpful every time I have sought their advice. I am also grateful for my committee: John Kovac, Cora Dvorkin, Douglas Finkbeiner, and Christopher Stubbs, for their guidance throughout the years.

My best friend, Eugeniu Plamadeala, deserves special gratitude for being with me through thick and thin for over 15 years. He has seen me fall and get up enough times to know exactly what kind of support to offer. I have always looked up to him, and have always been able to count on him.

Finally, I thank my partner, Nina Maksimova, for helping me get through the darkest moments of my life by injecting joy and lightness, by providing her relentless support, and by occasionally offering the kick up the rear that I needed to keep going. I thank her for the intimacy and home that we have built together – they have allowed me to steady the ship many a time. I thank her for her drive, passion, and quirks – they have shown me a different way to approach life, in good and bad. And I thank her for loving me back, even when I have not made it easy.

To my mom, who has sacrificed everything for her children, often at great expense to herself.
Mom, this is for you.

Pentru mama mea, care a sacrificat totul pentru copiii săi, adesea cu mare detriment pentru
bunăstarea sa. Mami, asta-i pentru tine.

*All you really need to know for the moment is
that the universe is a lot more complicated than
you might think, even if you start from a posi-
tion of thinking it's pretty damn complicated in
the first place.*

Douglas Adams

1

Introduction and Motivation

Cosmic inflation is our current best theory for what occurred in the universe in the first instances of time. It postulates a brief period of exponential expansion in which quantum fluctuations get magnified to cosmic size and become the seeds for the growth of all structure in the Universe. It was first put forward in Guth (1981) to explain the lack of observed magnetic monopoles and to solve the *Flatness* and *Horizon* problems (discussed in detail below)

and has since been a very active field of research. Inflation makes a number of predictions, the most unique of which is the production of primordial gravitational waves (PGWs). Most of the predictions have been since tested, but the discovery of PGWs has eluded us to this day. In this chapter, we will discuss the theory of inflation and its predictions, and explain how the Cosmic Microwave Background (CMB) is a powerful tool to test these predictions, including the exciting possible existence of PGWs.

1.1 THE COSMIC MICROWAVE BACKGROUND

The hot Big Bang model dictates that the early universe was much hotter and denser than today. When the universe was 380,000 years old, it was a hot and dense plasma of electrons, protons and photons. As the universe cooled and expanded, electrons and protons could combine to form neutral hydrogen, in a period called *recombination*. As a result, photons unencumbered by (Thompson) scattering with charged particles could finally free stream. This moment in the history of the universe, which occurred at redshift $z = 1100$, is called *the surface of last scattering*. The photons from this surface continue traveling through the universe and are known as the Cosmic Microwave Background (CMB). As the universe expands further, the CMB photons are redshifted (their wavelength expands with the universe and their energy – $E = hc/\lambda$ – decreases) from temperatures of around $T = 3000K$ to $\mathcal{O}(\text{few K})$ now.

The first observation of the CMB was performed at Bell Laboratories by [Penzias & Wilson \(1965\)](#), in which they measured an isotropic black-body signal with $T = (3.5 \pm 1.0)K$




Figure 1.1: The 2018 *Planck* map of the temperature anisotropies of the CMB, extracted using the SMICA method. The gray outline shows the extent of the confidence mask. *Figure from Planck Collaboration 2018 IV (2018)*

across the sky. The original discovery of this relic black-body radiation solidified the Big Bang model as the status quo explanation for the origin of the universe. Since then, these measurements have been refined to $T = (2.72548 \pm 0.00057) \text{ K}$ (Fixsen (2009)) and it has been discovered that while the CMB temperature is extremely uniform (to 1 part in 10^5), it still possesses anisotropies (as seen in Fig1.2). The history of CMB measurements is rich, and takes place over multiple decades; a comprehensive review of early CMB measurements can be found in Peebles et al. (2009).

Temperature anisotropies were first measured by the COBE satellite (Bennett et al. (1992)) and are an abundant source of information. On the sky, they are best decomposed

in spherical harmonics Y_{lm}

$$\Delta T(\theta, \phi) = \sum_{l=1}^{\infty} \sum_{m=-l}^l a_{lm}^T Y_{lm}(\theta, \phi).$$

Under the assumption that the a_{lm} coefficients are drawn from a Gaussian distribution, the information is entirely contained in the two-point function, or alternatively, the power spectrum. While the mean of the coefficients $\langle a_{lm} \rangle$ is zero, their variance – the power-spectrum – is non-trivial

$$C_l^{XY} = \frac{1}{2l+1} \sum_{m=-l}^l a_{lm}^{X*} a_{lm}^Y. \quad (1.1)$$

Here X and Y represent different possible anisotropy maps, where an example of such a map is represented in Figure 1.2. $X = Y$ corresponds to an auto-power-spectrum, while $X \neq Y$ corresponds to a cross-power-spectrum. This represents the cumulative power in a given multipole $l \sim 180^\circ/\theta$, where θ is the characteristic angular scale. The temperature power spectrum (C_l^{TT}) has been measured by a number of experiments, with the most accurate measurements to date coming from the *Planck* satellite (showcased in Figure 1.2). This spectrum is extremely rich in information, and allows us to constrain multiple cosmological parameters to high precision. The first peak allows one to infer that the universe is close to spatially flat. The second peak indicates substantial amounts of dark baryons, consistent with nucleosynthesis inferences. The third peak yields insight into the physical density of the dark matter. Finally, the damping tail serves as a consistency check on assumptions about recombination




Figure 1.2: *Planck* 2018 temperature power spectrum in red. The base- Λ CDM theoretical spectrum best fit to the likelihoods is plotted in light blue in the upper panel. Residuals with respect to this model are shown in the lower panel. *Figure from Planck Collaboration 2018 VI (2018)*

and initial conditions. For an in-depth review of how the CMB temperature power-spectrum provides all this information, the reader is referred to Dodelson (2003) and Hu & Sugiyama (1995).

Given that we are measuring one surface of one universe, the number of samples we may draw of a given a_{lm} is limited. This fundamental limit places an uncertainty on how well we can measure a given power-spectrum and is called *the cosmic variance limit*

$$\frac{\Delta C_l}{C_l} = \sqrt{\frac{2}{2l+1}}.$$

Our current measurements, presented in Figure 1.2, are cosmic variance limited below $l \sim 2500$.

1.2 INFLATION

We start this sub-chapter by laying out some foundational concepts from Cosmology and General Relativity. We describe the aforementioned *Flatness* and *Horizon* problems and explain how a period of cosmic inflation solves them. We then lay the basis for studying inflationary dynamics by introducing some crucial components from classical field theory and go on to use them in the context of slow-roll inflation. Next, we discuss the production of PGWs during inflation, and follow that with a sub-chapter on how CMB polarization measurements can allow us to test this prediction.

1.2.1 FUNDAMENTALS AND MOTIVATION

Einstein’s field equations relate the geometry of space-time to the distribution of matter and energy in the universe. They can be written in an extremely compact form ([Einstein \(1917\)](#))

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = 8\pi G\mathcal{T}_{\mu\nu},$$

using Einstein notation; G here is Newton’s constant. The following few equations introduce the precise mathematical definitions we need, characterizing the variation of the space-time metric along the various possible directions. For an in-depth discussion of these concepts we

direct the reader to any introductory book on General Relativity. $R^\mu_{\nu\lambda\sigma}$ – the Riemann curvature tensor takes the form:

$$R^\mu_{\nu\lambda\sigma} = \partial_\lambda \Gamma^\mu_{\nu\sigma} - \partial_\sigma \Gamma^\mu_{\nu\lambda} + \Gamma^\mu_{\rho\lambda} \Gamma^\rho_{\nu\sigma} - \Gamma^\mu_{\rho\sigma} \Gamma^\rho_{\nu\lambda},$$

where $\partial_\mu = \frac{\partial}{\partial x^\mu}$. Where the Christoffel symbol – $\Gamma^\mu_{\alpha\beta}$ – is defined as follows

$$\Gamma^\mu_{\alpha\beta} = \frac{1}{2} g^{\mu\lambda} [\partial_\alpha g_{\beta\lambda} + \partial_\beta g_{\alpha\lambda} - \partial_\lambda g_{\alpha\beta}].$$

The Ricci tensor is a contracted version (summed over one index) of the Riemann curvature $R_{\nu\sigma} \equiv R^\mu_{\nu\mu\sigma}$ and the Ricci scalar is the fully-contracted version (summed over all indices) of the Ricci tensor $R = g^{\mu\nu} R_{\mu\nu}$. Finally, the stress-energy tensor, $T_{\mu\nu}$, represents how the energy and pressure of matter is distributed within the space-time.

In this chapter we will focus on a special space-time called the Friedmann-Roberston-Walker (FRW) space-time ([Weinberg \(2008\)](#)), which bears a remarkable resemblance to our observable universe on large scales. However, the formalism we will build in this chapter can be applied to any generic space-time. We start with the observation that while our universe is very lumpy in its matter distribution on small scales, on the largest observable scales it is exceptionally smooth. As we have mentioned, CMB temperature observations indicate that the universe is uniform to 1 part in 10^5 , offering us strong reasons to believe that is no preferred direction (isotropy) and no preferred special point (homogeneity) in the universe. With

the assumptions of isotropy and homogeneity in mind, the most general space-time we can write is the FRW space-time

$$ds^2 = g_{\mu\nu}dx^\mu dx^\nu = -dt^2 + a(t)^2 \left[\frac{dr^2}{1 - Kr^2} + r^2(d\theta^2 + \sin^2\theta d\phi^2) \right], \quad (1.2)$$

where K is the curvature of the space-time, and $a(t)$ is a dimensionless scale factor that parametrizes the relative expansion of the universe.

The same two assumptions also restrict the FRW stress-energy tensor to take the perfect-fluid form $\mathcal{T}_{\mu\nu} = (\rho + p)u_\mu u_\nu + pg_{\mu\nu}$ (Weinberg (2008)), where u_μ is the normalized 4-velocity, and ρ and p are the matter density and pressure. In the frame comoving with the fluid $u_\mu = (1, 0, 0, 0)$ we obtain

$$\begin{aligned} \mathcal{T}_{00} &= \rho, \\ \mathcal{T}_{ij} &= pg_{ij}. \end{aligned} \quad (1.3)$$

For an FRW metric, considering the tt (or 00) component of the Einstein field equations, we obtain

$$3\left(\frac{\dot{a}^2}{a} + \frac{K}{a^2}\right) = 8\pi G\rho.$$

Rearranging this expression we arrive at the well known Friedmann equation

$$H^2 = \frac{8\pi G}{3}\rho - \frac{K}{a^2}. \quad (1.4)$$

We can see that there exists a critical value for the energy density $\rho_{cr} = \frac{3H^2}{8\pi G}$ such that the curvature term vanishes. Therefore, a universe with $\rho = \rho_{cr}$ will have zero curvature $K = 0$; a universe with $\rho > \rho_{cr}$ will have positive curvature $K > 0$; and a universe with $\rho < \rho_{cr}$ will have negative curvature $K < 0$. Alternatively, one can introduce the quantity $\Omega = \frac{\rho}{\rho_{cr}}$, to obtain that $\Omega = 1$ corresponds to a spatially flat universe; $\Omega > 1$ to a positively curved one; and $\Omega < 1$ to a negatively curved one. Rewriting the Friedmann equation in terms of Ω allows us to usefully quantify the deviation of this quantity from unity

$$\frac{\Omega - 1}{\Omega} = \frac{\mathcal{K}}{a^2 \rho}, \quad (1.5)$$

where $\mathcal{K} \equiv \frac{3K}{8\pi G}$. Further dynamical information comes from considering the ij components of the Einstein's field equations

$$-g_{ij} \left(H^2 + \frac{2\ddot{a}}{a} + \frac{K}{a^2} \right) = 8\pi G p g_{ij},$$

or, for the nonzero components of g_{ij}

$$H^2 + \frac{2\ddot{a}}{a} = -8\pi G p - \frac{K}{a^2}.$$

Subtracting Eq.1.4 from above we get

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3p). \quad (1.6)$$

This equation relates the matter density and pressure to the rate of acceleration of the expansion of the universe.

Another important property of the energy-momentum tensor is that it is conserved under covariant differentiation

$$\mathcal{D}_\alpha \mathcal{T}^{\mu\nu} = \partial_\alpha \mathcal{T}^{\mu\nu} + \Gamma_{\alpha\beta}^\mu \mathcal{T}^{\nu\beta} + \Gamma_{\alpha\beta}^\nu \mathcal{T}^{\mu\beta} = 0.$$

For an FRW space-time this yields

$$\mathcal{D}_\nu \mathcal{T}^{t\nu} = \dot{\rho} + 3\frac{\dot{a}}{a}(\rho + p) = 0.$$

Because ρ and p depend only on t in an FRW space-time, it is often useful to consider their ratio, $w \equiv p/\rho$, which only varies with time and must take the same value throughout space at any given moment of time. The relation $p = w\rho$ is known as the equation of state. If w is constant, the solution to the differential equation above yields

$$\rho(t) = \rho_0 a(t)^{-3(1+w)}, \tag{1.7}$$

where ρ_0 is the initial value of the energy density. One can prove that for regular matter the equation of state has $w = 0$, whilst for relativistic matter $w = 1/3$ ([Carroll \(2003\)](#)). For both these cases, the energy density decreases as the universe expands as $\rho(t) = \rho_0 a(t)^{-3}$ and $\rho(t) = \rho_0 a(t)^{-4}$, respectively. This, however, means that for a flat universe, i.e. $\rho = \rho_{cr}$

($\Omega = 1$), there is an unstable equilibrium, i.e. any deviation from unity increases over time as $\frac{\Omega-1}{\Omega} \propto a^{-2}(t)\rho(t)^{-1}$. Current measurements from the [Planck Collaboration 2018 VI \(2018\)](#) tell us $\Omega - 1 = 0.0007 \pm 0.0010$. For the universe to be as flat as it currently is, it would have had to have been orders of magnitude flatter at earlier times!

For instance, taking the measured value cited above for $\Omega - 1$ today, and assuming a universe with negligible curvature that has been dominated by radiation will yield $\Omega - 1 \approx 10^{-20}$ at the time of big-bang nucleosynthesis (BBN), when the universe was roughly $t = 1$ seconds old. Here we see that the universe would have had to have been extremely fine-tuned at its inception in order for us to currently observe a universe as flat as ours. This constitutes the *Flatness Problem*.

Another important problem in cosmology is the *Horizon Problem*. It arises due to the difficulty in explaining the observed homogeneity of causally disconnected regions of space in the absence of a mechanism that sets the same initial conditions everywhere. First, we define the conformal time

$$\tau \equiv \int_0^t \frac{dt'}{a(t')} = \int_0^a \frac{da}{a(aH)},$$

which can be interpreted as a clock that slows down with the expansion of the universe.

Then, the maximum comoving distance light can propagate between t_i and t_f – called the comoving particle horizon – is given by

$$\chi_p(\tau) = c(\tau_f - \tau_i) = \int_{t_i}^{t_f} \frac{cdt}{a(t)} = \int_0^a \frac{cda}{a(aH)}.$$

This is, of course, the fraction of the universe in causal contact. Combining Eq.1.4 and Eq.1.7 we have

$$\frac{1}{aH} = \frac{a^{\frac{1}{2}(1+3w)}}{H_0},$$

which leads to

$$\chi \propto \begin{cases} a, & \text{for radiation,} \\ a^{1/2}, & \text{for matter.} \end{cases}$$

We see here that the comoving horizon grows monotonically with time, i.e. scales that are entering the horizon today were completely outside the horizon at the time of last scattering. The near-homogeneity of the CMB dictates that the universe was homogeneous at the time of CMB decoupling at scales that according to the argument above would have been causally independent. This constitutes the *Horizon Problem*.

Both the *Flatness* and *Horizon problems* arise if we assume types of matter that we can currently see today, i.e. $w \geq 0$. If, on the other hand, we had some type of exotic matter with $w < -1/3$ we note that Eq.1.7 would yield $\rho(t) = \rho_0 a^n$ with $n > -2$ which in turn will allow Eq.1.5 to take the form $(\Omega - 1)/\Omega = \mathcal{K}/a^m$ with $m > 0$. A deviation from unity in Ω will no longer grow over time; in fact, the opposite is true. Note also that Eq.1.6 will yield precisely $\ddot{a} > 0$ when $w < -1/3$. Therefore, we conclude that a possible solution to the *Flatness Problem* is the existence of a period in the evolution of the universe during which it is filled with some type of matter that obeys $w < -1/3$ and during which it undergoes an accelerated expansion.




Figure 1.3: Conformal diagram showing the history of the universe since recombination, assuming no inflationary epoch. The past light cones of two locations in the surface of last scattering do not overlap, and hence were never in causal contact. *Figure adapted from Baumann (2009)*

Similarly, for the *Horizon* problem an $w < -1/3$ will lead to

$$\chi_p(\tau) = \frac{c}{H_0} \int_0^a a^{\frac{1}{2}(-1+3w)} da = \frac{c}{H_0} \int_0^a a^{-n} da = \infty, \text{ for } w < -1/3, \text{ or } n > 1,$$

which is large enough to put the entire observable universe in the causal path of any reference point in space-time.

Cosmic Inflation is a process that satisfies both $w < -1/3$ and $\ddot{a} > 0$, and therefore solves both the *Flatness* and *Horizon problem*.

1.2.2 MATTER FIELDS; ACTIONS AND MOTIONS

We have seen that in order to solve the flatness and horizon problems we require that some of the matter filling the universe behaves according to an equation of state with $w < -1/3$. In this section, we introduce matter fields and the required machinery to study their behavior in curved space-times. We will see how to relate components such as ρ and p to the matter fields and their derivatives. This will allow us to see under what conditions matter might behave such that $w < -1/3$.

First, what is a field and why is it a useful quantity? From classical mechanics we know that for a small number of particles, it is fairly straightforward to write the Lagrangian or the Hamiltonian of the system and solve for the equations of motion. This process becomes cumbersome quickly as we start to increase the number of particles. One has to find another way to encode this information, and one useful way to do so is through densities. These can be densities of matter, energy, momentum, etc. Densities can be further simplified by encoding information in fields.

The reader is already familiar with some types of matter and energy that are represented through fields. Namely, let us consider the *electric field*, which can be used to express information about the electromagnetic energy density. This specific case of the electric field is called a vector field, since it has a magnitude and a direction. One can introduce even simpler structures that don't depend on direction and are single-valued at any point in space and time. These quantities are dubbed scalar fields. Temperature is an example of such a field.

One can measure the temperature at every location in space, repeat the measurements over time, and represent the results with a single scalar field $T(x_\mu)$. By analogy, here we introduce a scalar field ϕ , which we will call the matter field, and it will encode all the information about the matter density in the universe.

To study the behavior of matter fields, we introduce the concept of an action. The action – \mathcal{S} – is an integral over all space and time of the Lagrangian density – \mathcal{L} .

Given that we would like to express the contents of the universe through a single scalar matter field, the action is just sum of the gravitational Einstein-Hilbert action – $\mathcal{S}^{(G)}$ – which is the component that describes the curvature of a given space-time, and the action of the scalar field describing matter – $\mathcal{S}^{(M)}$

$$\mathcal{S} = \mathcal{S}^{(G)} + \mathcal{S}^{(M)} = \int d\text{Vol}(\mathcal{L}^{(G)} + \mathcal{L}^{(M)}). \quad (1.8)$$

The Lagrangian density for the gravitational field in General Relativity takes the form $\mathcal{L} = \frac{1}{16\pi G} \mathcal{R}(x)$ (Weinberg (2008)), where $\mathcal{R}(x)$ is the Ricci scalar – the quantity that contains all the information about the curvature of space-time. In this chapter we are not concerned with this part, and would instead like to focus on constructing the Lagrangian density for matter $\mathcal{L}^{(M)}$. The Lagrangian for a system consists of the kinetic energy minus the potential energy

$$\mathcal{L}^{(M)} = KE_{field} - \mathcal{V}_{field}.$$

We will leave the shape of the potential term arbitrary, pursuing generality, and thus turn our attention to the kinetic term. Consider an analogy with a mechanical system that consists of a ball of unit mass moving along a trajectory $x(t)$. The kinetic term for this system is just $KE_{ball} = \frac{1}{2}\dot{x}^2$. By analogy, a field with trajectory $\phi(x_\mu)$ will have

$$KE_{field} = -\frac{1}{2}g^{\mu\nu}\frac{\partial\phi}{\partial x^\mu}\frac{\partial\phi}{\partial x^\nu} = \frac{1}{2}g^{\mu\nu}\partial\phi_\mu\partial\phi_\nu,$$

where we use $\partial_\mu = \frac{\partial}{\partial x^\mu}$. The last piece we need in order to construct the matter field is the invariant volume of space-time $dVol$. For an Euclidean four-dimensional space-time $dVol = d^4x$. However, we require an invariant notion that works in the context of a generic curved space-time. We can obtain one by including an additional factor that will compensate for the variation of the space-time volume under a coordinate transformation. From General Relativity, we know that this factor is precisely $\sqrt{-g}$, where $g = \text{Det}[g_{\mu\nu}]$. The minus sign is necessary because usually we treat at least one of the metric components as being negative. For this particular manuscript we will choose a mostly positive form of the metric. Hence, the invariant form of the integration measure is $dVol = d^4x\sqrt{-g}$. The action takes the form

$$\mathcal{S} = \int d^4x\sqrt{-g}\left(\frac{1}{16\pi G}\mathcal{R}(x) - \frac{1}{2}g^{\mu\nu}\partial\phi_\mu\partial\phi_\nu - \mathcal{V}(\phi)\right).$$

In order to find the equations of motion, we would like to know how $\mathcal{S}(\phi, \partial_\mu\phi)$ varies with respect to the field ϕ for arbitrary variations $\phi \rightarrow \phi + \delta\phi$, and under what conditions the

action remains stationary. We note right away that $\mathcal{L}^{(G)}$ does not depend on ϕ , so the only portion that varies is the matter portion of the action.

The variation of the action is therefore

$$\delta\mathcal{S} = \int d^4x \sqrt{-g} \left[\left(\frac{\partial\mathcal{L}^{(M)}}{\partial\phi} \right) \delta\phi + \frac{\partial\mathcal{L}^{(M)}}{\partial(\partial_\mu\phi)} \partial_\mu(\delta\phi) \right].$$

First, we can absorb the $\sqrt{-g}$ inside the parentheses as it does not depend on the field ϕ .

Second, we can integrate the second term by parts, which will send $\partial_\mu(\delta\phi)$ to $\delta\phi$ and will introduce a negative sign. We will also introduce a surface term in this process, but this term gives a null contribution as we require that the variation $\delta\phi$ vanishes at infinity. With this in mind, the variation of the action becomes

$$\delta\mathcal{S} = \int d^4x \left[\sqrt{-g} \left(\frac{\partial\mathcal{L}^{(M)}}{\partial\phi} \right) - \partial_\mu \left(\sqrt{-g} \frac{\partial\mathcal{L}^{(M)}}{\partial(\partial_\mu\phi)} \right) \right] \delta\phi.$$

If we require the action to be stationary, i.e. $\delta\mathcal{S} = 0$, for arbitrary $\delta\phi$, we obtain the Euler-Lagrange Equation for a generic curved space-time

$$\frac{\partial\mathcal{L}^{(M)}}{\partial\phi} - \frac{1}{\sqrt{-g}} \partial_\mu \left(\sqrt{-g} \frac{\partial\mathcal{L}^{(M)}}{\partial(\partial_\mu\phi)} \right) = 0.$$

It is not difficult to observe by analogy with the simple ball example that if we take the famil-

iar Euler-Lagrange equations

$$\frac{\partial \mathcal{L}^{(M)}}{\partial \phi} - \frac{d}{dt} \left(\frac{\partial \mathcal{L}^{(M)}}{\partial \dot{\phi}} \right) = 0,$$

and promote the derivative with respect to one coordinate to a derivative with respect to four coordinates ∂_μ , making sure that we take the derivative right for a given metric (i.e. have a factor of $\sqrt{-g}$), we obtain the same Euler-Lagrange equation as we just did from extremizing the action. Now, differentiating the Lagrangian density with respect to ϕ , we obtain

$$\frac{\partial \mathcal{L}^{(M)}}{\partial \phi} = -\frac{\partial \mathcal{V}}{\partial \phi} \equiv -V_{,\phi}.$$

To differentiate with respect to ∂_μ , first relabel the indices in $\mathcal{L}^{(M)}$, in order to avoid confusion about which indices are summed over, and then differentiate

$$\begin{aligned} \frac{\partial \mathcal{L}^{(M)}}{\partial(\partial_\mu \phi)} &= \frac{\partial}{\partial(\partial_\mu \phi)} \left[-\frac{1}{2} g^{\alpha\beta} \partial_\alpha \phi \partial_\beta \phi \right] \\ &= -\frac{1}{2} g^{\alpha\beta} \left[\delta_\alpha^\mu \partial_\beta \phi + \partial_\alpha \phi \delta_\beta^\mu \right] \\ &= -\frac{1}{2} \left[g^{\mu\beta} \partial_\beta \phi + \partial_\alpha \phi g^{\alpha\mu} \right] = -g^{\mu\beta} \partial_\beta \phi. \end{aligned}$$

Where the fourth line comes from relabeling indices in the second term of the third line, and from assuming that the metric is symmetric $g^{\mu\beta} = g^{\beta\mu}$. Combining the above equations, we obtain

$$\frac{1}{\sqrt{-g}} \partial_\mu \left(\sqrt{-g} g^{\mu\nu} \partial_\nu \phi \right) - \mathcal{V}_{,\phi} = 0.$$

The first term of this equation is the covariant d'Alembertian operator, denoted as \square . Notice, for a Minkowski space ($\sqrt{-g} = 1$), we recover the form of the d'Alembertian that we are used to $\square = \partial_\mu \partial^\mu$. With this in mind, the final equation of motion becomes

$$\square\phi - \mathcal{V}_{,\phi} = 0. \quad (1.9)$$

This equation works for any curved space-time within which we can describe matter with a scalar field, and for any choice of potential.

The next step is to calculate $\mathcal{T}_{\mu\nu}$ for a universe filled with a single scalar matter field ϕ for a generic curved space. The most general way to find functional forms for the stress-energy tensors is by varying the action with respect to the metric

$$\mathcal{T}_{\mu\nu} = \frac{-2}{\sqrt{-g}} \frac{\delta \mathcal{S}^{(M)}}{\delta g^{\mu\nu}},$$

where $\mathcal{S}^{(M)} = \int d^4x \sqrt{-g} \mathcal{L}^{(M)}$. Therefore

$$\frac{\delta \mathcal{S}^{(M)}}{\delta g^{\mu\nu}} = \frac{\delta \sqrt{-g}}{\delta g^{\mu\nu}} \mathcal{L}^{(M)} + \sqrt{-g} \frac{\delta \mathcal{L}^{(M)}}{\delta g^{\mu\nu}}.$$

To evaluate this, we need to know the variation of $\sqrt{-g}$ with the metric, for which we use the identity $\delta g = g g^{\mu\nu} \delta g_{\mu\nu} = -g g_{\mu\nu} \delta g^{\mu\nu}$ (Carroll (2003)). Applying the chain rule we find

$$\delta \sqrt{-g} = -\frac{1}{2\sqrt{-g}} \delta g = -\frac{1}{2} \sqrt{-g} g_{\mu\nu} \delta g^{\mu\nu},$$

which allows us to write

$$\frac{\delta \mathcal{S}^{(M)}}{\delta g^{\mu\nu}} = -\frac{1}{2}\sqrt{-g}g_{\mu\nu}\mathcal{L}^{(M)} + \sqrt{-g}\left(-\frac{1}{2}\partial_\mu\phi\partial_\nu\phi\right),$$

and finally

$$\mathcal{T}_{\mu\nu} = \partial_\mu\phi\partial_\nu\phi - g_{\mu\nu}\left[\frac{1}{2}g^{\alpha\beta}\partial_\alpha\phi\partial_\beta\phi + \mathcal{V}(\phi)\right]. \quad (1.10)$$

Just like Eq.1.9, this is extremely general.

Next, we would like to apply the developed formalism to our universe and so we consider the case of a Friedmann-Robertson-Walker cosmology.

As we have seen, the FRW metric takes the following form

$$g^{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & \frac{a^2}{1-Kr^2} & 0 & 0 \\ 0 & 0 & a^2r^2 & 0 \\ 0 & 0 & 0 & a^2r^2\sin^2\theta \end{pmatrix},$$

yielding $\sqrt{-g} = a^3(t)f(r, \theta)$. An FRW cosmology is spatially homogeneous and isotropic,

i.e. $\phi(x^\mu) = \phi(t)$ or $\partial_\nu\phi = \dot{\phi}\delta_{\nu 0}$. This symmetry allows us to significantly simplify the

d'Alembertian of ϕ

$$\begin{aligned}
\Box\phi &= \frac{1}{\sqrt{-g}}\partial_\mu\left(\sqrt{-g}g^{\mu\nu}\partial_\nu\phi\right) \\
&= \frac{1}{a^3(t)f(r,\theta)}\partial_\mu\left[-a^3(t)f(r,\theta)\dot{\phi}\delta^{\mu 0}\right] \\
&= \frac{1}{a^3(t)f(r,\theta)}\partial_t\left[-a^3(t)f(r,\theta)\dot{\phi}\right] \\
&= \frac{1}{a^3(t)}\left[-a^3(t)\ddot{\phi} - 3a^2(t)a\dot{(t)}\dot{\phi}\right] = -\ddot{\phi} - 3H\dot{\phi},
\end{aligned}$$

where we have used the definition for the Hubble parameter $H = \dot{a}(t)/a(t)$. Plugging the result in Eq. 1.9, the equation of motion for ϕ becomes

$$\ddot{\phi} + 3H\dot{\phi} + \mathcal{V}_{,\phi} = 0. \quad (1.11)$$

This equation governs the evolution of $\phi(t)$. In order to find the solution to this equation one needs to know how the Hubble parameter $H = \dot{a}(t)/a(t)$ evolves with time. The behavior of $a(t)$ depends on the composition of the universe, expressed in Eq.1.10. Using the homogeneity of FRW, we obtain for the 00 (or tt) component

$$\begin{aligned}
\mathcal{T}_{00} &= \partial_t\phi\partial_t\phi - g_{tt}\left[\frac{1}{2}g^{tt}\partial_t\phi\partial_t\phi + \frac{1}{2}g^{ij}\partial_i\phi\partial_j\phi + \mathcal{V}(\phi)\right] \\
&= \dot{\phi}^2 - (-1)\left[\frac{1}{2}(-1)\dot{\phi}^2 + \mathcal{V}(\phi)\right] \\
&= \frac{1}{2}\dot{\phi}^2 + \mathcal{V}(\phi).
\end{aligned}$$

Likewise for the ij component

$$\begin{aligned}
\mathcal{T}_{ij} &= \partial_i \phi \partial_j \phi - g_{ij} \left[\frac{1}{2} g^{tt} \partial_t \phi \partial_t \phi + \frac{1}{2} g^{ij} \partial_i \phi \partial_j \phi + \mathcal{V}(\phi) \right] \\
&= 0 - g_{ij} \left[\frac{1}{2} (-1) \dot{\phi}^2 + \mathcal{V}(\phi) \right] \\
&= g_{ij} \left[\frac{1}{2} \dot{\phi}^2 - \mathcal{V}(\phi) \right].
\end{aligned}$$

Combining with equations 1.3 we may write

$$\begin{aligned}
\rho &= \frac{1}{2} \dot{\phi}^2 + \mathcal{V}(\phi), \\
p &= \frac{1}{2} \dot{\phi}^2 - \mathcal{V}(\phi).
\end{aligned} \tag{1.12}$$

Equation 1.11 tells us how the matter field evolves in an FRW space-time and the expressions in 1.12 in turn determine how the space-time itself evolves. These are all the necessary tools to study the dynamics of an FRW cosmology filled with a single matter field.

1.2.3 INFLATIONARY DYNAMICS AND SLOW ROLL PARAMETERS

There are a multitude of ways to drive inflation, one particular way is to assume a period during which the potential energy stored in the fields is much larger than the kinetic energy, i.e. $\mathcal{V}(\phi) \gg \frac{1}{2} \dot{\phi}^2$. One can imagine that the field starts at a given value of $\mathcal{V}(\phi)$ in a shallow convex potential, and slowly rolls to the minimum. This is termed *slow-roll inflation*. This condition yields $p \simeq -\rho$ according to Eq.1.12, or alternatively $w = -1$. In particular using

Eq.1.6 and Eq.1.12 we obtain

$$\frac{\ddot{a}}{a} \simeq \frac{8\pi G}{3} \mathcal{V}(\phi).$$

Meanwhile Eq.1.4 becomes

$$H^2 \simeq \frac{8\pi G}{3} \mathcal{V}(\phi) - \frac{K}{a^2} \simeq \frac{8\pi G}{3} \mathcal{V}(\phi), \quad (1.13)$$

where we again have used $\mathcal{V}(\phi) \gg \frac{1}{2}\dot{\phi}^2$ and the fact that as the universe accelerates, its size, parametrized by the dimensionless scale factor $a(t)$, will grow, and the curvature term will become negligible. We have thus found that under these assumptions $H^2 \simeq \ddot{a}/a$. Taking the time derivative of $H = \dot{a}/a$, we have

$$\dot{H} = \frac{\ddot{a}}{a} - \frac{\dot{a}^2}{a^2} = \frac{\ddot{a}}{a} - H^2. \quad (1.14)$$

Therefore, when $H^2 \simeq \ddot{a}/a$ holds, $\dot{H} \simeq 0$, and the Hubble parameter is constant over the period of inflation. This allows us to simply integrate the expression for $a(t)$

$$H = \frac{\dot{a}}{a} = \frac{1}{a} \frac{da}{dt} \quad \rightarrow \quad \frac{da}{a} = H dt \quad \rightarrow \quad a(t) = a_0 \exp H(t - t_0),$$

to obtain exponential expansion. Another observation is that for $w = -1$, Eq.1.7 leads to $\rho = \rho_0$ constant, i.e. as the universe becomes exponentially larger, its energy density remains constant. At first glance this appears to imply a violation of the conservation of energy; how-

ever, energy is precisely compensated by the negative potential energy of the gravitational field, which grows as there is more and more space for gravity to permeate as the universe expands. In fact, because of this exponential expansion, the universe does not need to obey $w = -1$ everywhere. It is sufficient to have a small patch of it following this equation of state – in very short time the small inflating patch will become most of the universe – and, once inflation starts, it does not need to stop, leading to the possibility of eternal inflation. However, as inflation proceeds, the kinetic term of the field is slowly growing and at some point the equations above will not hold because we are not in the slow-rolling regime anymore. To quantify this effect, we introduce the slow-roll parameters ϵ and η

$$\begin{aligned}\epsilon &\equiv -\frac{\dot{H}}{H^2}, \\ \eta &\equiv \epsilon - \frac{\ddot{\phi}}{H\dot{\phi}}.\end{aligned}\tag{1.15}$$

With some algebraic manipulation, both can be related to the curvature of the potential

([Kaiser \(2011\)](#))

$$\begin{aligned}\epsilon &\simeq \frac{1}{16\pi G} \left(\frac{\mathcal{V}_{,\phi}}{\mathcal{V}} \right)^2, \\ \eta &\simeq \frac{1}{8\pi G} \left(\frac{\mathcal{V}_{,\phi}}{\mathcal{V}} \right).\end{aligned}\tag{1.16}$$

These parameters are small at the beginning of inflation, and, since $\dot{H} \ll 1$ and $|\ddot{\phi}| \ll 1$, they grow steadily throughout inflation. Inflation ends when either of these slow-roll parameters becomes of order 1. From Eq.1.13 we see that $\epsilon = 1$ leads to $\dot{H} = -H^2$ and therefore

implies $\ddot{a} = 0$. The slow-roll parameters are a convenient way to summarize the predictions of an inflationary model and, as we can see, they are related to the potential $\mathcal{V}(\phi)$ and its derivatives. Extracting the values of ϵ and η from cosmological data is therefore equivalent to probing the inflationary potential! In the next section we will see how we can relate these abstract quantities to real observables.

1.2.4 INFLATIONARY PERTURBATIONS

The uncertainty principle is a fundamental property of quantum mechanics. It states that there is a limit to the precision with which physical properties of a particle (such as momentum and position, or energy and time) can be known. This gives rise to brief changes, or quantum fluctuations, in the amount of energy at any and all points in space, allowing for the creation of particle-antiparticle pairs of virtual particles. The seemingly empty vacuum of space is actually bubbling with the creation and annihilation of such pairs. During cosmic inflation, virtual particles near the causal horizon can cross it and become real, then become stretched to astrophysical size, and succeed in making cosmological perturbations, providing the initial seeds for structure formation. In general, the symmetries of an FRW space-time allow for a clean decomposition of the possible types of perturbations into independent scalar, vector, and tensor components. This independence is crucial because it allows for the separate evolution of these perturbations.

During inflation we can define perturbations around the homogeneous background solu-

tions for the matter field $\bar{\phi}(t)$ and the metric $\bar{g}_{\mu\nu}(t)$

$$\phi(t, \vec{x}) = \bar{\phi}(t) + \delta\phi(t, \vec{x}), \quad g_{\mu\nu}(t, \vec{x}) = \bar{g}_{\mu\nu}(t) + \delta g_{\mu\nu}(t, \vec{x}).$$

Scalar perturbations $\delta\phi$ lead to density perturbations $\delta\rho$, which become observable as temperature anisotropies in the CMB. Vector perturbations decay away as $a(t)^{-2}$ (Dodelson (2003)). Tensor perturbations $\delta g_{\mu\nu}(t, \vec{x})$ give rise to primordial gravitational waves, which produce observable distortions in the CMB, as will be described further below.

The variance of the primordial scalar perturbations can be written as (Baumann (2009))

$$\langle \mathcal{R}_{\mathbf{k}} \mathcal{R}_{\mathbf{k}'} \rangle = (2\pi)^3 \delta(\mathbf{k} + \mathbf{k}') \mathcal{P}_{\mathcal{R}}(k), \quad \Delta_s^2 = \Delta_{\mathcal{R}}^2 = \frac{k^3}{2\pi^2} \mathcal{P}_{\mathcal{R}}(k).$$

Δ_s^2 is the dimensionless scalar power spectrum

$$\Delta_s^2 = \Delta_{\mathcal{R}}^2 = \frac{k^3}{2\pi^2} \mathcal{P}_{\mathcal{R}}(k) = \frac{1}{8\pi^2} \frac{1}{\epsilon} \frac{H^2}{M_p^2} \bigg|_{k=aH}, \quad (1.17)$$

where $M_p = 1/\sqrt{8\pi G}$ is the reduced Planck mass and $k = aH$ indicates horizon crossing. The scale dependence of the scalar power spectrum is defined as

$$n_s - 1 \equiv \frac{d \ln \Delta_s^2}{d \ln k}.$$

While a perfectly invariant spectrum has $n_s = 1$, slow-roll inflation predicts $n_s = 1 - 6\epsilon +$

2η (Liddle & Lyth (2000)). With both slow-roll parameters being positive and $\epsilon \sim \eta$, we note that this implies $n_s < 1$. The current observational constraints on n_s were measured by Planck Collaboration 2018 VI (2018) to be $n_s = 0.965 \pm 0.004$.

Similarly, the variance of the primordial tensor perturbations is given by

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

where the factor of 2 in $\Delta_t^2 = 2\Delta_h^2$ arises from there being two polarization modes – h^+ , h^\times – of gravitational waves. The dimensionless tensor power spectrum Δ_t^2 is given by

$$\Delta_t^2 = 2\Delta_h^2 = 2 \frac{k^3}{2\pi^2} \mathcal{P}_h(k) = \frac{2}{\pi^2} \frac{H^2}{M_p^2} \Big|_{k=aH}. \quad (1.18)$$

The scale dependence is defined similarly to the scalar case, modulo the subtraction by 1

$$n_t \equiv \frac{d \ln \Delta_t^2}{d \ln k},$$

indicating that $n_t = 0$ will give an invariant spectrum. Slow-roll inflation predicts $n_t = -2\epsilon$.

The tensor power spectrum and its spectral tilt has not been measured yet. Measuring this spectrum is an active area of research, and what most the work described in this manuscript is dedicated to. In general, as we have seen, measuring n_s and n_t probe inflation directly through the slow-roll parameters, and are therefore a direct probe of the early universe.

Usually the tensor power spectrum is normalized relative to the measured scalar power

spectrum. Therefore, a useful quantity to define is the tensor-to-scalar ratio

$$r \equiv \frac{\Delta_t^2(k)}{\Delta_s^2(k)}.$$

This ratio is a direct measure of the energy scale of inflation (Liddle & Lyth (2000))

$$V^{1/4} \propto \left(\frac{r}{0.01}\right)^{1/4} 10^{16} GeV,$$

with values of $r > 0.01$ corresponding to inflation taking place at grand unified theory (GUT) scales – the energy level above which we believe the electromagnetic force, weak force, and strong force become equal in strength and unify to one force. Probing this level of energy is incredibly interesting, providing rich opportunities of new physics.

Combining equations 1.17, 1.18, and 1.15, we can rewrite the tensor-to-scalar ratio in terms of the matter field

$$r = 16\epsilon = \frac{8}{M_p^2} \left(\frac{\dot{\phi}}{H}\right)^2 = \frac{8}{M_p^2} \left(\frac{d\phi}{dN}\right)^2,$$

where N is the number of e-foldings. Therefore the total field excursion between the end of inflation and when the CMB fluctuations exited the horizon is given by

$$\frac{\Delta\phi}{M_p} = \int_{N_{end}}^{N_{cmb}} dN \sqrt{\frac{r}{8}}.$$

According to Liddle & Lyth (2000), during inflation $r(N)$ varies very slowly, allowing us to

write

$$\frac{\Delta\phi}{M_p} = \mathcal{O}(1) \times \left(\frac{r}{0.01}\right)^{1/2}.$$

For values of $r > 0.01$, the field traverse is larger than the Planck mass – this is called the Lyth Bound. A detection at this level would provide evidence that theories of quantum gravity must accomodate a Planckian field range for the inflaton. The tightest constraints to date on r and n_s are presented in Figure 1.4. The work in this manuscript has contributed considerably towards these state-of-the-art measurements.

It is clear that a possible detection of primordial gravitational waves created by inflation ($r > 0$) would offer a direct and unique window into the early universe, with profound implications for high-energy physics and the quantum nature of gravity (Krauss & Wilczek (2014)). In the next sub-chapter we will look at observable imprints of PGWs on CMB polarization, and will discuss how we can leverage the CMB to constrain inflation.

1.3 CMB POLARIZATION

The CMB radiation is expected to be polarized because of Compton scattering (the relativistic case of Thomson scattering) at the time of decoupling (Dodelson (2003)). Light traveling in the z -direction will have electric and magnetic field oscillating in the $x - y$ plane. Under the assumption that the intensity from the two transverse directions is equal, the scattered light will be unpolarized. If, however, the intensities are not equal, we will obtain a net polarization. Consider the first case shown in Figure 1.5, which represents an unpolarized ray




Figure 1.4: Constraints in the r vs. n_s plane when using *Planck* plus additional data, and when also adding BICEP2/*Keck* data through the end of the 2015 season—the constraint on r tightens from $r_{0.05} < 0.12$ to $r_{0.05} < 0.06$. *Figure from BICEP2 Collaboration et al. (2018)*

incident from the $+\hat{z}$ direction, which has equal intensities in the \hat{x} and \hat{y} directions. This ray is scattered by the electron at the origin. Under the assumption that the deflection takes the ray in the \hat{y} direction, the electric and magnetic field will lie in the $x - z$ plane. However, since we started with no intensity in the \hat{z} direction, the result is a net polarization in the \hat{x} direction (perpendicular to both the incoming and the outgoing directions).

In reality we will have rays from all directions. Consider the next case: isotropic radiation from every direction (monopole), shown in the 2nd panel of Figure 1.5 as two rays incoming from \hat{x} and \hat{z} . In this case, the \hat{x} intensity of the outgoing \hat{y} ray comes from the incoming \hat{z} ray, as before, and the intensity in the \hat{z} comes from the incoming \hat{x} ray. If the intensities are uniform, as assumed, then the result is an unpolarized ray.




Figure 1.5: A temperature quadrupole around a last-scattering electron is required to generate polarization along the line of sight (dot-dashed line); neither a monopole nor a dipole can create polarization. *Figure adapted from Hu & White (1997).*

Now let us consider the simplest anisotropic case: a dipole pattern, shown in the 3rd panel of Figure 1.5 as rays coming from hot and cold spots at $\pm\hat{z}$ and from average temperature spot at \hat{y} . The hot and cold spots will result in an average intensity along the \hat{x} direction for the scattered ray, and the average \hat{y} spot will yield average intensity along \hat{z} . Yet again, we achieve an unpolarized state.

Finally, we consider the case of a quadrupolar pattern, shown in the 4th panel of Figure 1.5. The ray from the hot spot at \hat{z} produces higher intensity along the \hat{x} for the outgoing wave than the ray from the cold spot at \hat{x} does along the \hat{z} . This results in a net polarization. We conclude therefore that one needs a quadrupolar anisotropy to obtain a polarized state.

The last panel of Figure 1.5 depicts polarization in the $x - z$ plane, preferentially in the \hat{x} direction. However, if the incident rays were rotated by $\pi/4$ we would have outgoing polarization along axes that are rotated by $\pi/4$ from the x and z axes. We can therefore depict polarization as a headless vector oriented along the axis for which the intensity is highest and with length corresponding to its magnitude. In the direction perpendicular to propagation, we can write the intensity as a 2D matrix

$$I_{ij} = \begin{pmatrix} T + Q & U \\ U & T - Q \end{pmatrix},$$

yielding $U = I_{12}$, $Q = \frac{1}{2}(I_{11} - I_{22})$ and $T = \frac{1}{2}(I_{11} + I_{22})$. Here T is the temperature field discussed in Section 1.1, and Q and U are parameters that describe polarization. Together, they are three of the four Stokes parameters. The fourth parameter, V , is only non-zero if we have circular polarization, a phenomenon that is expected to occur in the early universe, allowing us to set $V = 0$. The total polarization magnitude and angle are $P = \sqrt{Q^2 + U^2}$ and $\alpha = 1/2 \arctan^{-1}(U/Q)$.

The Stokes parameters are clearly coordinate dependent. Given that it is desirable to work with invariant quantities, we introduce instead two quantities called E-modes and B-modes. We start with a harmonic analysis of $Q \pm iU$, which requires an expansion in terms of spin-2 harmonics

$$(Q \pm iU)(\hat{n}) = \sum_{l,m} a_{\pm 2, lm \pm 2} Y_{lm}(\hat{n}),$$

where $_{\pm 2}Y_{lm}$ are the spin-2 spherical harmonics. We then switch from the $a_{\pm 2,lm}$ basis to

$$a_{E,lm} \equiv -\frac{1}{2}(a_{2,lm} + a_{-2,lm}), \quad a_{B,lm} \equiv -\frac{1}{2i}(a_{2,lm} - a_{-2,lm}),$$

which allows us to define two spin-0 fields

$$E(\hat{n}) = \sum_{l,m} a_{E,lm} Y_{lm}(\hat{n}), \quad B(\hat{n}) = \sum_{l,m} a_{B,lm} Y_{lm}(\hat{n}).$$

These are coordinate independent, but behave differently under parity transformations (E-modes are even and B-modes are odd). Similarly to temperature fluctuations (Eq. 1.1) we can write the power spectrum for polarization as well

$$C_l^{XY} = \frac{1}{2l+1} \sum_{m=-l}^l a_{lm}^{X*} a_{lm}^Y, \quad (1.19)$$

for $X, Y = T, E, B$. It is worth noting that only the auto-spectra TT , EE , BB and the cross-correlation TE survive, while TB and EB vanish due to symmetry (Baumann (2009)), under the assumption that parity is not violated.

We can use the same framework as we did for the Q, U to discuss the patterns produced by E and B -modes. Consider the inhomogeneous density field that produced the observed temperature anisotropies. In particular, pick a density wave that has peaks and troughs as depicted in the first panel of Figure 1.6. This type of wave corresponds to an $m = 0$ quadrupole, which is azimuthally symmetric about the \hat{k} direction, and will yield a scattered




Figure 1.6: A density wave creates polarization in the CMB varying at 0 or 90 with respect to the wavevector, resulting in an E -mode pattern. *Figure from the BICEP2 Collaboration.*

net 0° or 90° polarization, as presented in the second panel of Figure 1.6. We can visually see that this type of pattern is even under parity inversion. The sum of many waves of this kind still produces an E -mode pattern, which is a curl-free polarization component. E -mode polarization was first detected by [Kovac et al. \(2002\)](#) and has since been measured to great precision, as shown in Figure 1.8. Since E modes are produced by the same acoustic oscillations that generate temperature anisotropies, and the velocity of the photon-baryon fluid is maximized for modes that are halfway between compression and rarefaction at recombination, the peaks of the TT and EE spectra are out of phase.

Similarly, consider a tensor wave with an $m \pm 2$ quadrupolar pattern, as depicted in first panel of Figure 1.7. This type of wave will squeeze and stretch space-time as shown, and create a net polarization of $\pm 45^\circ$ with respect to the projected wave-vector ([Kamionkowski et al. \(1997\)](#)). Such polarization patterns are divergence-free and can only be created by tensor modes, which implies that B -mode measurements are not sample-variance restricted as the

E-modes are. Therefore, B-mode polarization is a powerful measure for PGWs (parametrized by the tensor-to-scalar ratio r). Primordial B-modes are expected to peak at degree angular scales, corresponding to the horizon scale at recombination.

In addition to primordial B-modes, there are also so-called “lensing B-modes.” These are E-modes that get distorted into B-modes as photons get deflected by large-scale structure on their way from recombination to now. However, its contribution can be reduced by knowing the cumulative gravitational lensing potential ϕ along the line of sight, and having a faithful E-mode map. Together, the two can be combined to form a lensing B-mode template by lensing the E-mode map with the ϕ field and subtracting this template from the measured B-mode map. This technique is known in the field as *delensing*. Lensing B-modes have been measured by a number of experiments, as shown in Figure 1.8. A detection of PGWs, however, has so far eluded us. PGWs are the only prediction of inflation which has not been confirmed so far. Other predictions: superhorizon fluctuations, gaussian perturbations, adiabatic fluctuations, spatial flatness, and a nearly invariant scalar spectral tilt, have all been tested, most recently in [Planck Collaboration et al. \(2018b\)](#). Detecting PGWs has the potential to definitively confirm inflation, or a process similar to it, as we currently do not have another framework to produce them.




Figure 1.7: A gravitational wave generates polarization in the CMB varying at 45° with respect to the wavevector, resulting in a B -mode pattern. Gravitational waves can also generate E modes and add power to the temperature sky. *Figure from the BICEP2 Collaboration.*

1.4 THESIS OUTLINE

The main focus of this thesis is the development of two major tools that have allowed the CMB field to forge ahead in constraining PGWs with present and future experiments. The first one, presented in Chapter 3, is an optimal multi-component spectral-based likelihood analysis framework for joint analyses of heterogeneous multi-frequency CMB datasets. The second one, presented in Chapter 7, is a spectral-based Fisher projection framework that directly uses information from current achieved performances to robustly forecast the science reach of upcoming CMB-polarization endeavors. Chapter 2 provides an overview of the BICEP/*Keck* series of instruments and analysis pipeline from time ordered data (TOD) to power spectra. Chapters 4-6 describe applications of the multi-component framework to various datasets (BICEP/*Keck* *Planck*/ and WMAP) and document sequential framework up-




Figure 1.8: A compilation of recent CMB angular power spectrum measurements. The plot shows the power spectra of the temperature and E-mode and B-mode polarization signals. Note that for Planck, ACTPol, and SPTpol, the EE points with large error bars are not plotted (to avoid clutter). The dashed line shows the best-fit Λ CDM model to the Planck temperature, polarization, and lensing data. *Figure adapted from Planck Collaboration et al. (2018a).*

dates. Finally, Chapter 8 provides a conclusion and discusses the path forward for the CMB field.

*If you think this Universe is bad, you should see
some of the others.*

Philip K. Dick

2

BICEP/Keck in a Nutshell

BICEP and the *Keck* Array are a series of telescopes, designed from the ground up to specifically target B-mode polarization at degree angular scales, where the signal from PGWs is expected to be brightest. A number of thorough manuscripts have been written about the BICEP/*Keck* instruments, and the BICEP/*Keck* pipeline, in full detail ([James Tolan \(2014\)](#), [Sarah Kernasovskiy \(2014\)](#), [Chin Lin Wong \(2014\)](#), [Kirit Karkare \(2017\)](#)). In this chapter, we

provide only a brief summary of the BICEP/*Keck* program and offer an overview of the necessary pipeline analysis steps that allow us to go from time ordered data to final power spectra.

2.1 THE BICEP/KECK SERIES OF INSTRUMENTS

BICEP2 and the *Keck* array are a set of cryogenically cooled telescopes with refracting, compact, on-axis optics with an aperture of 26.4 cm. Both are sited at the South Pole in Antarctica, taking advantage of the dry atmosphere and stable observing conditions. BICEP2 had 256 dual polarization pixels in the focal plane for a total of 512 antenna-coupled transition edge sensor (TES) bolometers read with a SQUID-based time domain multiplexing system. Each detector had a temperature sensitivity of $\sim 300\mu K_{CMB}\sqrt{s}$, and all detectors observed at 150 GHz. The Keck Array is a series of five BICEP2 like tubes, which started with a combined imaging array of 2500 TES bolometers at 150 GHz, and has since been continuously upgraded to have focal planes at various frequencies. In addition to the all-cold optics, these telescopes have two features which aid greatly in the suppression and characterization of instrumental systematics: first, they are equipped with co-moving absorptive forebaffles, resulting in extremely low far side-lobe response, and second, the entire instrument can be rotated about the line of sight, allowing modulation of polarized signal; since BICEP/*Keck* focal planes all have the dual-polarization pixels oriented in the same direction, and since there is no sky rotation at the South Pole, these instruments need boresight rotation to make Stokes Q/U maps.

The program has had a few iterations over the years, starting with BICEP1 in 2006-2008, continuing with BICEP2 in 2009-2012, the *Keck* array in 2011-Present, BICEP3 in 2015-present, and culminating in BICEP Array which will start deployment in 2019. Table 2.1 summarizes all the receivers deployed since 2010, including BICEP3, which we do not discuss in this text. This manuscript will focus on placing cosmological constraints with data from BICEP2 and the *Keck* array in conjunction with the satellite-based missions WMAP and *Planck*.

Throughout the years BICEP2 and all five Keck Array telescopes observed the same low-foreground patch of the southern sky. In principle, the telescope observes around 600 square degrees of the sky, but factoring in the lower sensitivity around the edges brings the effective sky fraction closer to 400 square degrees. This patch, shown in Figure 2.1, represents a $\sim 1\%$ region of sky, centered at RA $0h$, Dec. 57.5° , and has the advantage of being observable 24/7 throughout the year.

2.2 FROM TIMESTREAMS TO POWER SPECTRA

A brief outline of the BICEP/*Keck* data acquisition and analysis pipeline is presented.

During observations, the telescope scans in azimuth at constant elevation. Once the entire length of the patch is covered (64.2 degrees), the telescope reverses its azimuth direction and scans again. A scan in a single direction is called a *halfscan*. Halfscans are grouped in batches of ~ 100 , performed over the course of 50 minutes, totaling to a full *scanset*. At

	2010	2011	2012
BICEP Mount	B2 (500, 150 GHz)	B2 (500, 150 GHz)	B2 (500, 150 GHz)
<i>Keck</i> Rx0		K3 (500, 150 GHz)	K3 (500, 150 GHz)
<i>Keck</i> Rx1		K2 (500, 150 GHz)	K2 (500, 150 GHz)
<i>Keck</i> Rx2		K1 (500, 150 GHz)	K1 (500, 150 GHz)
<i>Keck</i> Rx3			K4 (500, 150 GHz)
<i>Keck</i> Rx4			K5 (500, 150 GHz)

	2013	2014	2015
BICEP Mount			B3 (1080, 95 GHz)
<i>Keck</i> Rx0	K3 (500, 150 GHz)	K3 (272, 95 GHz)	K3 (272, 95 GHz)
<i>Keck</i> Rx1	K2 (500, 150 GHz)	K2 (500, 150 GHz)	K2 (500, 220 GHz)
<i>Keck</i> Rx2	K1 (500, 150 GHz)	K1 (272, 95 GHz)	K1 (272, 95 GHz)
<i>Keck</i> Rx3	K4 (500, 150 GHz)	K4 (500, 150 GHz)	K4 (500, 220 GHz)
<i>Keck</i> Rx4	K5 (500, 150 GHz)	K5 (500, 150 GHz)	K5 (500, 150 GHz)

	2016	2017	2018
BICEP Mount	B3 (2400, 95 GHz)	B3 (2400, 95 GHz)	B3 (2400, 95 GHz)
<i>Keck</i> Rx0	K3 (500, 210 GHz)	K3 (500, 210 GHz)	K3 (500, 210 GHz)
<i>Keck</i> Rx1	K2 (500, 220 GHz)	K2 (500, 220 GHz)	K2 (500, 220 GHz)
<i>Keck</i> Rx2	K1 (500, 210 GHz)	K1 (500, 210 GHz)	K1 (500, 210 GHz)
<i>Keck</i> Rx3	K4 (500, 220 GHz)	K4 (500, 220 GHz)	K4 (500, 220 GHz)
<i>Keck</i> Rx4	K5 (500, 150 GHz)	K5 (500, 270 GHz)	K5 (500, 270 GHz)

Table 2.1: BICEP/*Keck* Array receivers deployed since 2010. Receiver names, nominal number of optically-coupled detectors, and frequencies are indicated. “B2” refers to BICEP2, “B3” refers to BICEP3, and “K” indicates a particular *Keck* Array cryostat. Table adapted from *Kirit Karkare (2017)*.




Figure 2.1: The BICEP and *Keck Array* observing field shown against the predicted polarization intensity in the Planck Sky Model. This low-foreground patch is often referred to as the *Southern Hole*. *Figure from BICEP2 Collaboration I (2014).*

the end of each scanset, the elevation is updated by 0.25° , and another scanset is performed. Each scanset begins by being centered at an azimuthal coordinate corresponding to $RA=0$, and halfscans are performed about this point. This step is crucial in allowing us to disentangle azimuthally fixed signals, such as ground pick-up, from sky signals, which rotate. Finally, scansets are grouped together in 10 hour long increments called *observation phases*. Observation phases are performed in 2-day increments, driven by the fridge cycle, and are all done at the same deck angle.

As scansets are performed, data from the TES detectors are recorded in a sequence of time-ordered data (TOD's), or timestreams, which also records thermometry and telescope pointing information. The readout multi-channel electronics (MCE) low-pass filters and downsamples the timestreams and passes them to the control computer running *gcp* software.

The data is written to disk in archive files. The first step in the low-level data reduction is deconvolving the temporal response of the MCE and *gcp* filtering. Next, delta function spikes and discontinuous steps in the data are removed in a step called *deglitching* and *destepping*.

Following these procedures, we calibrate our instrument in two steps. The first is a relative calibration process, used to correct the difference between polarization pairs and to combine pairs over the array optimally. Every hour of observation, a small 1° elevation nod (*elnod*) is performed, and the atmospheric temperature gradient is fitted for. This allows for the calculation of the relative detector gains against the median of the entire experiment, and their normalization, in analog to digital units (ADU). The second is an absolute calibration from ADU to real temperature units (μK_{CMB}). This is usually done by comparing the power spectrum of the temperature map with the Planck 143 GHz map and is performed later in the analysis, usually after map-making.

Once the timestreams have been normalized, the timestreams of orthogonal detector pairs are summed and differenced. The pair-summed timestreams are used to make unpolarized measurements of temperature, while the pair-differenced timestreams are used to make timestream Stokes $[Q, U]$ measurements of polarization.

To bin the timestreams into pixels with the appropriate right ascension and declination, pointing information is required. The three axes of the telescope mount encode the pointing direction for a particular observation. A star camera – a near-infrared camera that observes a number of stars with known positions – is then used to translate between these mount encoder positions and actual sky coordinates. This procedure is termed boresight pointing. To

determine the pointing of each detector relative to the boresight pointing, we form cross spectra of each full-year detector map with the Planck temperature map. Detector maps are first made with an approximate location, and then the cross-correlations are used to derive corrections.

Before $[Q, U]$ maps are made, a number of other steps are taken. First, a selection of data cuts is performed. The cuts framework is a complex, multi-level algorithm that identifies and removes faulty timestream data related to a number of possible pathologies. A detailed review of this procedure is detailed in [Sarah Kernasovskiy \(2014\)](#). In short, there are three rounds of cuts. Round 1 cut parameters are calculated and applied for individual halfscans; this round primarily cuts data with glitches, steps, and skipped samples. Round 2 cuts are calculated at the scanset level; this round has three major categories of cuts pertaining to calibration, data quality, and telescope performance. Finally, a number of channels which are deemed faulty are completely cut. After all of these cuts, an average of 50% of the data is retained. Following the cut procedure, a set of timestream filtering operations is performed. First, to remove low frequency $1/f$ atmospheric fluctuations, a third-order polynomial is fit and subtracted from each halfscan. Then, a template of ground fixed signals is constructed for each scanset, and it is removed from the timestreams. This serves to remove azimuthally-fixed signals, which are a potential systematic contamination.

The pair-summed and pair-differenced timestreams are weighted separately based on the variance during a scanset and then binned into a two-dimensional map of the sky. The maps constructed from multiple phases of data are combined using a weighted average. An

inversion from the pair difference map to Stokes $[Q, U]$ maps is then performed, using the pointing information discussed earlier.

Two actions worth noting that are also performed at this step are the deprojection of leading order temperature to polarization leakage terms, due to imperfect differencing between orthogonal pairs of detectors, and the adjustment of the absolute polarization angle to minimize the EB cross spectrum. The latter procedure is termed *self-calibration* and operates under the assumption that parity is not violated, i.e., there is zero on sky TB and EB power. Given this assumption, the absolute polarization angle of the experiment is adjusted until the TB and EB spectra are minimized. See Ref. [BICEP2 Collaboration I \(2014\)](#) for more information.

We apodize to downweight the regions around the edge of the observed area, with the apodization set equal to a map of the weights coadded over the full data set, the Q/U maps are Fourier transformed and converted to the E/B basis in which a possible PGW signal is expected to be most distinct from the standard Λ CDM signal. To suppress E to B leakage, due to the skycut and filtering, we use the matrix purification technique which we have developed [BICEP2 Collaboration I \(2014\)](#), [Keck Array](#) and [BICEP2 Collaborations VII \(2016\)](#) and [James Tolan \(2014\)](#). We then take the variance within annuli (of fixed width $l \sim 35$) of the Fourier plane to estimate the angular power spectra bandpowers.

Next, we create noise-only simulations based on the full data set (see Section 2.2.1) and test their accuracy through standard jackknife tests ([BICEP2 Collaboration I \(2014\)](#)). These jackknives exist to test the real data for systematic contamination, and the noise simulations

are only as accurate as the data used to generate them. These simulations provide an estimate of noise in the power spectra of the real data, so we calculate the power spectra of the noise maps and subtract their mean level from the real data. This method is called *noise debiasing*. While the mean is removed, the variance of the noise is preserved and used as a contribution to the final uncertainty of the measured spectra.

Finally, we correct the measured power spectra for the *filter/beam suppression factor* – the combined effect of polynomial filtering, ground subtraction, deprojection, and the matrix purification have had on each mode in the map. This factor is obtained by passing simulated input skies with power in only one multipole ($\delta(l)$ -function) through all the enumerated steps and obtaining the sky response at each multipole as observed by the pipeline (i.e., *the bandpower window function*). The unnormalized integral of the bandpower window functions are the desired suppression factors. Once these are obtained, real measured bandpowers are divided by the suppression factor to obtain estimated true-sky bandpowers.

2.2.1 NOISE SIMULATIONS

Realistic noise simulations are a necessary component in our analysis. To that end we create *sign-flip noise realizations*. These are created directly from the data. A large number of scansets is produced every year. This method multiplies each scanset by a randomly assigned positive or negative sign, such that the positive and negative halves are equally weighed. When these are added, the random signs lead to the signal component in the map cancelling, leaving only the noise component, which naturally contains noise correlations between de-

tectors. Because of the large number of degrees of freedom associated with the number of scansets, even though the same scansets are used, the 499 simulations we create are largely independent. This method is discussed in detail in [van Engelen et al. \(2012\)](#).

2.2.2 SIGNAL SIMULATIONS

We also generate 499 realizations of lensed and unlensed Λ CDM by resampling timestreams created from simulated input maps and passing them through the full analysis pipeline. The unlensed simulations are useful to empirically determine the purity delivered by the matrix purification algorithm which is used to extract the B-mode signal in the presence of a much stronger E-mode. We also use these simulations for the construction of our bandpower covariance matrix, as detailed in Section 3.5.

*Trying to understand the way nature works
involves a most terrible test of human reasoning
ability. It involves subtle trickery, beautiful
tightropes of logic on which one has to walk in
order not to make a mistake in predicting what
will happen.*

Richard P. Feynman

3

Multi-component Likelihood Analysis

Framework

In this chapter, we describe the development of an optimal multi-component, multi-frequency, spectral-based likelihood analysis framework. The goal of this tool is to allow us to analyze multiple heterogeneous CMB polarization datasets (from BICEP/*Keck Planck*, and WMAP)

in order to disentangle any PGW signal from signals of galactic origin. In addition to the lensing B-mode signal, at low frequencies, charged particles in our galaxy’s magnetic field produce synchrotron radiation; at high frequencies, polarized rotating dust particles create a signal of their own, with intricacies we do not yet fully understand.

We start this chapter with a historical motivation and a bird’s eye overview of the framework. We then introduce all the necessary experimental inputs, the assumed signal theory model, and the likelihood approximations that will be used. An important sub-chapter is dedicated towards the construction of a semi-analytic bandpower covariance matrix from explicit signal and noise simulations, which allows for rescaling to any signal theory, and for reducing Monte Carlo noise given the modest number of simulations. Finally, chapters 4-6 describe the application of this framework, and its evolution, with each iteration of the aforementioned datasets, as they become available.

The framework presented in this manuscript has been described and used to produce the analysis and figures presented in the following published and peer-reviewed articles: *BICEP2/Keck* and *Planck* Collaborations (2015), *Keck Array* and *BICEP2* Collaborations VI (2016), *BICEP2 Collaboration et al.* (2018). Therefore, where it is deemed appropriate, this manuscript follows the cited publications.

3.1 MOTIVATION & OVERVIEW

In March 2014, the BICEP2 team announced the detection of excess degree-scale (multipole range of $30 < l < 150$) B-mode power at 150 GHz over the base lensed- Λ CDM expectation, inconsistent with the null hypothesis at a significance of $> 5\sigma$ [BICEP2 Collaboration I \(2014\)](#). The signal was compatible with one of inflationary origin with $r = 0.20^{+0.07}_{-0.05}$. Cross-correlating against WMAP 23 GHz maps showed that Galactic synchrotron made a negligible contribution to the observed signal. The several polarized dust emission models considered at the time offered estimates of foreground contributions that were subdominant to the signal by a factor of $\sim 5 - 10$. However, these models were not sufficiently constrained by external public data to exclude the possibility of dust emission bright enough to explain the entire excess signal.

A later release by the *Planck* collaboration ([Planck Collaboration Int. XXX \(2016\)](#)), showing measurements of polarized dust over the entire sky, suggested that the polarized fraction of high-galactic latitude dust was higher than the models were predicting. The conclusion was that the entire BICEP2 B-mode signal could be explained by dust. However, while the *Planck* data brought new insights into dust physics and allowed a rough determination of the level of contamination for the BICEP2 CMB polarization experiments, it was not sufficiently constraining to determine the origin of the signal with certainty. The two teams decided that a dedicated joint analysis, incorporating all the pertinent observational details of the two data sets, was necessary in order to assess the dust contributions to the observed

B-mode signal by BICEP2.

To that end, we developed a multi-component, multi-frequency, spectral-based likelihood analysis that allowed for joint consideration of CMB polarization data from multiple experiments. It is worth noting that while this was the first fully-fledged framework for analysis of multiple CMB polarization datasets, which included BICEP2/*Keck* data, previous efforts towards joint analyses existed prior to this one. Figure 3.1 shows a schematic representation of this likelihood framework. In the subsections that follow we will discuss all the modules in detail. We will begin by describing the individual experiments and the preparation of the data to power spectra. Following that, we present the *multi-component* theoretical signal model used to calculate the theory expectation values, the construction of a re-scalable semi-analytic bandpower covariance matrix (BPCM), and the likelihood approximation method used in analyses of BICEP/*Keck* data.

3.2 EXPERIMENT DESCRIPTION & REAL BANDPOWERS

While the framework is entirely flexible and can accommodate any number and flavor of CMB polarization datasets, for the purposes of this manuscript, we will focus on data from four different instruments: BICEP2 *Keck* WMAP and *Planck*. The BICEP2 and *Keck Array* experiments and the necessary analysis steps from maps to power spectra have already been described in Chapter 2. In our analyses, we primarily use the BICEP2/*Keck* combined maps, as described in *Keck Array and BICEP2 Collaborations V (2015)*. The *Planck* maps used for

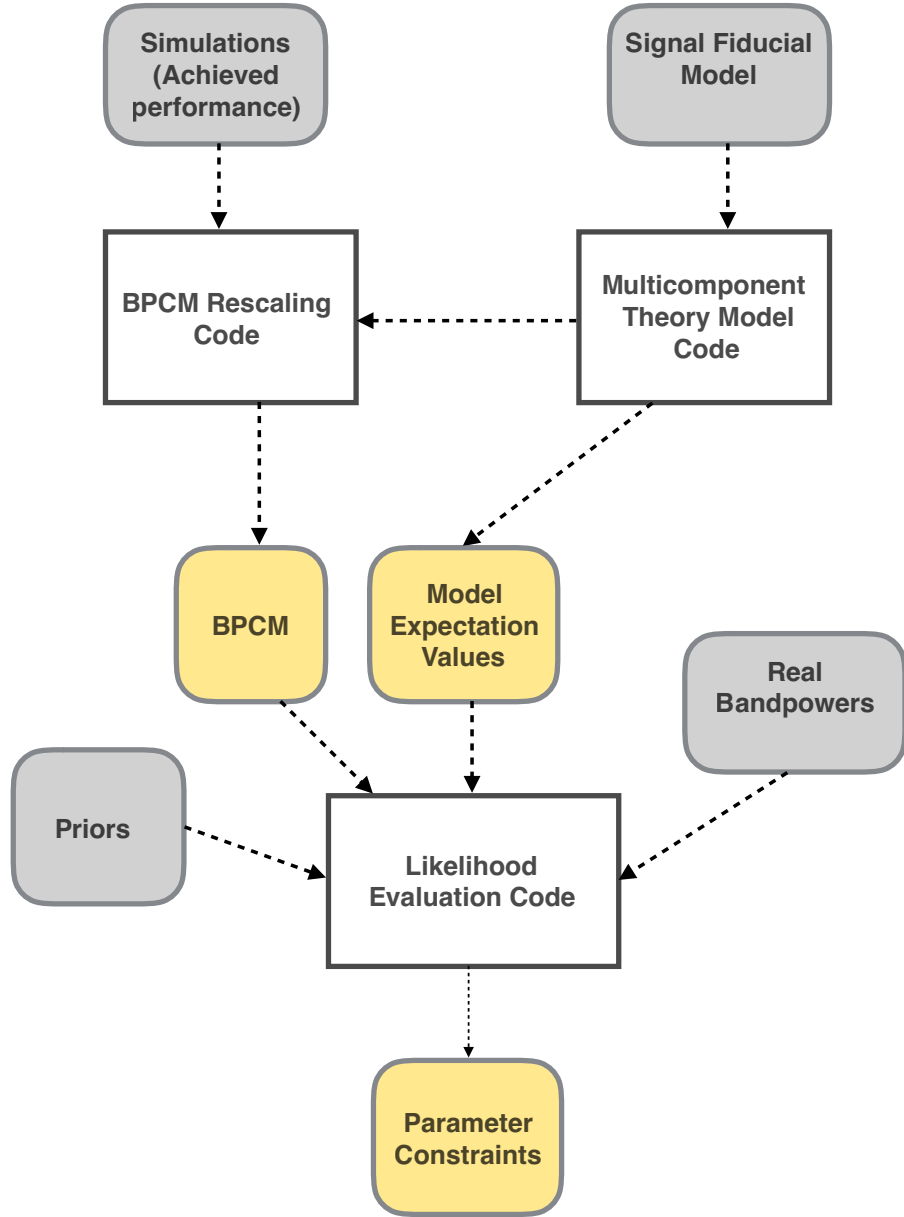


Figure 3.1: Schematic representation of the BICEP/*Keck* multi-component multi-frequency spectral based likelihood framework

cross-correlation with BICEP2/*Keck* are the full-mission polarized maps from *Planck* science release ¹. In more recent analyses we have also folded in the WMAP9 maps at 23 GHz (K-band) and 33 GHz (Ka-band) [Bennett et al. \(2012\)](#).

While the *Planck* maps are filtered only by the instrument beam (the effective beam defined in Refs. [Planck Collaboration IV \(2014\)](#) and [Planck Collaboration VII \(2014\)](#)), the BICEP2/*Keck* maps are also filtered due to the observation strategy and analysis process as described in Chapter 2. In order to facilitate comparison, we, therefore, prepare *Planck* maps “reobserved by BICEP2/*Keck*”. First, we use the HEALPix ² [Górski et al. \(2005\)](#) package to resmooth the *Planck* maps with the BICEP2/*Keck* beam profile. We perform a coordinate rotation of the T , Q , and U maps from Galactic to celestial coordinates using ALTERALM in HEALPix. Next, we pass these through the steps in Chapter 2 to produce maps that include the filtering of modes occurring in the data processing pipeline (these include the polynomial filtering, ground subtraction, and deprojection). The WMAP maps go through an equivalent process.

We convert the maps to power spectra using the methods described in Chapter 2, including the matrix based purification operation to prevent E to B mixing. We generate separate purification matrices to match the filtering of the BICEP/*Keck* maps at each frequency. We take all possible auto-power spectra between these bands. As described in Chapter 2, these are the variance within annuli of the Fourier plane for two copies of the same map. We also

¹<https://pla.esac.esa.int/#maps>

²<http://healpix.sourceforge.net/>

take all possible cross-spectra between bands. The cross-spectrum is equivalent in mathematical definition to the auto-spectrum with the caveat that the two input maps are distinct rather than the same (which is the case for auto-spectra).

The cumulative collection of these auto- and cross- bandpowers represent the set of *real bandpowers* that we use in our analyses to calculate joint likelihoods of the dataset ensemble.

3.3 MULTI-COMPONENT MODEL

The likelihood analysis uses a parametrized model to describe the bandpower expectation values as a combination of cosmological and foreground signals. In our model, we have four distinct signal types: CMB (lensing contribution + tensor contribution), uncorrelated Galactic dust, uncorrelated Galactic synchrotron (sync), and a spatially correlated component between dust and synchrotron. In addition, we also consider dust and synchrotron frequency decorrelation parameters which allow for dust/sync cross-spectral power suppression.

We assume that each signal component is independent, i.e., different signals have zero-cross power. Therefore, for a given spectrum, the code steps through the model components, combines the appropriate amplitude functions for the two experiments contributing to the spectrum and applies the bandpower window functions to get the binned expectation values. Finally, it sums over model components to get the total expectation value for that spectrum. Equation 3.1 outlines the contributions to the BB auto- or cross-spectrum between maps at frequencies ν_1 and ν_2 ($\nu_1 = \nu_2$ for an auto-spectrum) from dust, synchrotron, and the

spatially-correlated component of dust and synchrotron.

$$\begin{aligned} \mathcal{D}_{\ell, BB}^{\nu_1 \times \nu_2} = & A_{dust} \Delta'_d f_d^{\nu_1} f_d^{\nu_2} \left(\frac{\ell}{80} \right)^{\alpha_d} + A_{sync} \Delta'_s f_s^{\nu_1} f_s^{\nu_2} \left(\frac{\ell}{80} \right)^{\alpha_s} + \\ & \pm \epsilon \sqrt{A_{dust} A_{sync}} (f_d^{\nu_1} f_s^{\nu_2} + f_s^{\nu_1} f_d^{\nu_2}) \left(\frac{\ell}{80} \right)^{(\alpha_d + \alpha_s)/2}. \end{aligned} \quad (3.1)$$

The parameter A_{dust} specifies the dust power in units of μK_{CMB}^2 at a pivot frequency 353 GHz and angular scale $\ell = 80$. The parameter A_{sync} specifies synchrotron power in units of μK_{CMB}^2 at a pivot frequency of 23 GHz and angular scale $\ell = 80$. The dust and synchrotron components scale as power laws in ℓ with slopes α_d and α_s , respectively. The parameters α_d and α_s are defined as the ℓ scaling of $\mathcal{D}_\ell \equiv \ell(\ell + 1) C_l / 2\pi$.

The parameter ϵ sets the level of spatial correlation between dust and synchrotron. This correlation coefficient is assumed to be constant across all ℓ , and the ℓ scaling of the correlated component has a slope that is the average of α_d and α_s . If either A_{dust} or A_{sync} are negative, the contribution of the correlated component to the expectation value flips sign. Negative foreground amplitudes are technically nonphysical, but this analytic continuation will become important in later validation steps when we want to explore the entire parameter phase-space.

The additional coefficients f_d^ν and f_s^ν , given by Equations 3.2 and 3.3, describe the scaling of dust and synchrotron power from the pivot frequencies to the measured bandpasses of the maps at frequencies ν_1 and ν_2 . This scaling has two contributions: first, from the foreground SED, and second, from the conversion between units at the target map bandpass and

at the pivot frequency.

The SED model used for dust is a blackbody with temperature $T_d = 19.6\text{K}$ multiplied by a power law with emissivity spectral index β_d [Planck Collaboration Int. XXII \(2015\)](#). This is also known as a greybody spectrum. This simplifying assumption about the behavior of dust holds over the range of frequencies that we consider. The SED model used for synchrotron is defined relative to a Rayleigh-Jeans spectrum as a power law with spectral index β_s . We are required to choose a bandpass convention when we perform the integration of the SED a unit conversion factors over a particular map bandpass. We adopt the same convention as used by [Planck Collaboration IX \(2014\)](#), in which our bandpass functions are defined to be proportional to the response as a function of frequency to a beam-filling source with uniform spectral radiance.

For a greybody signal with spectral index β_d and temperature T_d , the relative scaling for an experiment with bandpass $R(\nu)$ and pivot frequency ν_{pivot} is given by

$$f_d^\nu = \frac{\int d\nu R_1(\nu) \nu^{3+\beta_d} \left(\exp \frac{h\nu}{kT_d} - 1 \right)^{-1}}{\nu_{pivot}^{3+\beta_d} \left(\exp \frac{h\nu_{pivot}}{kT_d} - 1 \right)^{-1}} \times \frac{\nu^4 \exp \frac{h\nu_{pivot}}{kT_{CMB}} \left(\exp \frac{h\nu_{pivot}}{kT_{CMB}} - 1 \right)^{-2}}{\int d\nu R_1(\nu) \nu^4 \exp \frac{h\nu}{kT_{CMB}} \left(\exp \frac{h\nu}{kT_{CMB}} - 1 \right)^{-2}}. \quad (3.2)$$

The first ratio on the right-hand side of the equation is the greybody scaling between the two bandpasses while the second ratio is the conversion between CMB temperature units. For a

synchrotron power-law scaling with spectral index β_s the formula is modified to:

$$f_s^\nu = \frac{\int d\nu R_1(\nu) \nu^{2+\beta_d}}{\nu_{pivot}^{2+\beta_d}} \times \frac{\nu^4 \exp \frac{h\nu_{pivot}}{kT_{CMB}} \left(\exp \frac{h\nu_{pivot}}{kT_{CMB}} - 1 \right)^{-2}}{\int d\nu R_1(\nu) \nu^4 \exp \frac{h\nu}{kT_{CMB}} \left(\exp \frac{h\nu}{kT_{CMB}} - 1 \right)^{-2}}. \quad (3.3)$$

We also consider dust and synchrotron frequency decorrelation. The simplest possible model of a polarized foreground component is that with a fixed spatial pattern on the sky which scales with frequency according to a single SED. In this case, the cross-spectrum between two frequencies is the geometric mean of the respective auto-spectra. In reality, the polarization pattern inevitably varies as a function of frequency, leading to the cross-spectra being suppressed with respect to the geometric mean of the auto-spectra [Planck Collaboration Int. L \(2017\)](#). We refer to this phenomenon as decorrelation.

We define the correlation ratio of the dust

$$\Delta_d = \frac{\mathcal{D}_{80}(217 \times 353)}{\sqrt{\mathcal{D}_{80}(217 \times 217) \mathcal{D}_{80}(353 \times 353)}}, \quad (3.4)$$

where \mathcal{D}_{80} is the dust power at $\ell = 80$. Here $\Delta_d < 1$ corresponds to decorrelation. We scale to other frequency combinations using the factor suggested by [Planck Collaboration Int. L \(2017\)](#)

$$f(\nu_1, \nu_2) = \frac{(\log(\nu_1/\nu_2))^2}{(\log(217/353))^2}, \quad (3.5)$$

Similarly, based on suggestions from [Planck Collaboration Int. L \(2017\)](#) we consider two

possible scalings

$$g(\ell) = \begin{cases} 1 & \text{flat case} \\ (\ell/80) & \text{linear case} \end{cases}. \quad (3.6)$$

The ℓ range in which we have high signal-to-noise measurements is fairly narrow, therefore this choice turns out to make little practical difference.

The scalings above can produce extreme, and non-physical, behavior for frequencies that have a wide separation, or for wide ℓ ranges. We therefore re-map the nominal value using the following function

$$\Delta'_d(\nu_1, \nu_2, \ell) = \exp [\log(\Delta_d) f(\nu_1, \nu_2) g(\ell)]. \quad (3.7)$$

According to this re-mapping, Δ'_d remains in the range 0 to 1 for all values of f and g . This frequency scaling has been shown to correspond to a Gaussian spatial variation in the foreground spectral index ([Vansyngel et al. \(2017\)](#)). In a similar vein, we define the parameter Δ'_s , which describes decorrelation of the synchrotron pattern, but in practice, we do not use this parameter.

We also currently do not include foreground decorrelation parameters in the dust–synchrotron correlated component. In general, we expect a complete foreground model to include all possible correlations between galactic foregrounds at ν_1 and ν_2 , but the current measurements are not sensitive enough to guide us about the form of these correlations. In our analysis, we only consider decorrelation when $\epsilon = 0$.

The foreground contribution to EE is similar, with A_d and A_{sync} scaled by the EE/BB ratios for dust and synchrotron, which are both assumed to be equal to 2 [Planck Collaboration 2018 XI \(2018\)](#), [Krachmalnicoff et al. \(2018\)](#). The model has no EB contributions since we do not expect parity symmetry to be broken. We also do not include $TT/TE/TB$ spectra in the likelihood analysis, nor do we model unpolarized foregrounds.

In addition to foregrounds, we also include CMB lensing and a possible tensor contribution. Using CMB temperature units, the BB spectrum is given by:

$$D_{l,BB}^{\nu_1 \times \nu_2} = \frac{r}{0.1} D_{l,BB}^{tensor} + A_L D_{l,BB}^{lensing}, \quad (3.8)$$

where $D_{l,BB}^{tensor}$ is the BB spectrum produced by a tensor signal with $r = 0.1$, and $D_{l,BB}^{lensing}$ is the expected lensing BB spectrum for Λ CDM. These are obtained using the `CAMB`³ package. Similarly for EE :

$$D_{l,EE}^{\nu_1 \times \nu_2} = \frac{r}{0.1} D_{l,EE}^{tensor} + D_{l,EE}^{lensing}. \quad (3.9)$$

We expect that the second term dominates the EE contribution unless r is very large. Also, the EE lensing component does not vary with A_L as the lensing of B-modes into E-modes is assumed to be negligible compared to the E-mode expectation from Λ CDM.

³<https://camb.info/>

3.4 SIMULATIONS

3.4.1 SIGNAL SIMULATIONS

As described in Section 2.2.2 we generate 499 realizations of lensed and unlensed Λ CDM by timestream resampling from simulated input maps and applying the full analysis pipeline.

We also explicitly simulate simple dust input maps as power-law Gaussian realizations (with amplitude set to $A_d = 3.75 \mu K_{CMB}^2$) and pass these through the timestream sampling and pipeline re-mapping operation. They are then added to the lensed- Λ CDM and noise maps, and taken through to power spectra. We use these when it is important to match the fluctuations present in the real data in detail, such as in the bandpower covariance matrix construction described below.

3.4.2 NOISE SIMULATIONS

For BICEP/*Keck* we use the sign-flip noise simulations describe in Section 2.2.1. To evaluate uncertainties due to Planck instrumental noise, we use 499 noise simulations of each map; these are the standard set of time-ordered data noise simulations projected into sky maps (the FFP8 simulations defined in [Planck Collaboration 2015 XII \(2015\)](#)). For WMAP we use simple inhomogeneous white noise simulations derived from the provided variance maps.

3.5 SEMI-ANALYTIC BANDPOWER COVARIANCE MATRIX

Appendix B presents the full derivation of our semi-analytic bandpower covariance matrix formulation. In short, from the simulated signal-cross-signal, noise-cross-noise and signal-cross-noise spectra, for the full set of experiments under consideration, we can construct the bandpower covariance matrix appropriate for any model containing a set of signal components with given SEDs. When we do this, we set to zero any term which has an expectation value of zero (under the assumption that signal and noise are uncorrelated, and different signals are uncorrelated) to reduce the Monte Carlo error in the resulting covariance matrix given the relatively modest number of realizations used. We also set to zero the covariance between bandpowers that are separated by more than one bin in l , but, importantly, preserve the covariance between the auto- and cross-spectra of the different frequency bands.

3.6 LIKELIHOOD APPROXIMATION

For BICEP2 & *Keck Array* the effective number of degrees of freedom per bandpower (l -bin) is small, leading to significant non-Gaussianity of the likelihood, particularly in the lowest l -bins. Therefore, we choose not to use a Gaussian likelihood approximation, but instead to calculate bandpower likelihoods from our power spectra using the Hamimeche & Lewis likelihood approximation (Hamimeche & Lewis (2008)).

The H-L likelihood takes the following form (equation 49 of [Hamimeche & Lewis \(2008\)](#)):

$$\begin{aligned}
-2 \log \mathcal{L} &= [\mathbf{X}_g]_l^T \left[\mathbf{M}_f^{-1} \right]_{ll'} [\mathbf{X}_g]_{l'} \\
[\mathbf{X}_g]_l &= \text{vecp} \left(\mathbf{C}_{fl}^{1/2} \mathbf{g} [\mathbf{C}_l^{-1/2} \hat{\mathbf{C}}_l \mathbf{C}_l^{-1/2}] \mathbf{C}_{fl}^{1/2} \right)
\end{aligned} \tag{3.10}$$

Here, \mathbf{C}_{fl} are the fiducial model bandpowers, \mathbf{M}_f is the bandpower covariance matrix for the fiducial model, \mathbf{C}_l are the theory bandpower expectation values (different at each point in parameter space), and $\hat{\mathbf{C}}_l$ are the real data bandpowers. Each of these sets of bandpowers is written in a symmetric matrix form, with auto-spectra on the diagonal and cross-spectra off-diagonal. For example, if we use E and B modes from *BK* and *P353*, then the bandpower matrix is 4×4 with the following form (color-coded by spectrum type):

$$\hat{\mathbf{C}}_l = \begin{bmatrix}
\textcolor{red}{BK_E \times BK_E} & \textcolor{green}{BK_E \times BK_B} & \textcolor{red}{BK_E \times P353_E} & \textcolor{green}{BK_E \times P353_B} \\
\textcolor{green}{BK_B \times BK_E} & \textcolor{blue}{BK_B \times BK_B} & \textcolor{green}{BK_B \times P353_E} & \textcolor{blue}{BK_B \times P353_B} \\
\textcolor{red}{P353_E \times BK_E} & \textcolor{green}{P353_E \times BK_B} & \textcolor{red}{P353_E \times P353_E} & \textcolor{green}{P353_E \times P353_B} \\
\textcolor{green}{P353_B \times BK_E} & \textcolor{blue}{P353_B \times BK_B} & \textcolor{green}{P353_B \times P353_E} & \textcolor{blue}{P353_B \times P353_B}
\end{bmatrix}$$

These bandpowers all include the noise bias that we ordinarily subtract out of the real data bandpowers. In the case of the fiducial model or theory bandpowers, the noise bias is added to the calculated expectation values. We expect that only the diagonal entries of this matrix should have significant noise bias. In practice, we allow the off-diagonal noise biases to be non-zero, but we also confirm that they are quite small.

The matrix operation \mathbf{g} applied to $\mathbf{C}_l^{-1/2}\hat{\mathbf{C}}_l\mathbf{C}_l^{-1/2}$ involves finding the eigenvalues of this matrix combination and applying the function $g(x) = \text{sign}(x - 1)\sqrt{2(x - \ln(x) - 1)}$. This function is only real-valued for $x > 0$, meaning that the matrix is positive-definite. The $\mathbf{C}_l^{-1/2}$ terms should always be positive-definite. For the usual H-L case, using all the auto and cross-spectra from a set of maps, positive-definiteness is guaranteed for $\hat{\mathbf{C}}_l$ by a triangle inequality (the amount of power obtained from taking the cross-spectrum of two maps can't exceed the geometric mean of the auto-spectrum power in those maps).

The fiducial model and the associated covariance matrix are fixed during the likelihood calculation, so only the \mathbf{g} varies. The function $g(x)$ is equal to zero for $x = 1$, which occurs when the theory expectation values equal the data bandpowers, i.e., the matrix product $\mathbf{C}_l^{-1/2}\hat{\mathbf{C}}_l\mathbf{C}_l^{-1/2}$ equals the identity matrix. This gives the minimum value for $|\mathbf{X}_g|$ and hence the maximum likelihood. Since the bandpower noise bias terms are added into both the data bandpowers and the theory expectation values, the maximum likelihood solution should not depend strongly on the values used for the noise bias. However, these values do affect the shape of the likelihood, because of the non-linear functional form for $g(x)$. Larger values for the noise bias mean that a particular range of theories will explore a smaller range of x and therefore $g(x)$ will appear more linear.

Note that the unnormalized likelihood values do depend strongly on the noise bias, due to its inclusion in \mathbf{C}_{fl} , but this divides out during any normalization step or in likelihood ratios.

Given a particular parameter space, we sample the likelihood over the entire space us-

ing either a grid evaluation or, for higher-dimensional likelihoods, a Markov-Chain Monte Carlo engine, namely COSMOMC⁴. The COSMOMC module we developed, containing all the inputs and the model described above, is available for download at <http://bicepkeck.org>. In certain cases, when sampling, we use priors on the frequency spectral behaviors of dust and synchrotron emission from previous analyses of WMAP and *Planck* data in other regions of the sky.

⁴<https://cosmologist.info/cosmomc/>

*It doesn't matter how beautiful your theory
is, it doesn't matter how smart you are. If it
doesn't agree with experiment, it's wrong.*

Richard P. Feynman

4

A Joint Analysis of BICEP2/*Keck* and *Planck* Data

4.1 INTRODUCTION

In this chapter, we demonstrate the first application of the likelihood methodology discussed in Chapter 3, to derive cosmological parameter constraints with data from BICEP2 the *Keck*

array and the *Planck* satellite. This was the first joint analysis of these data-sets, meant to allow us to distinguish whether the B-mode signal reported in [BICEP2 Collaboration I \(2014\)](#) was of Galactic or primordial origin. Section 3.1 describes in detail the history and motivation for this investigation. This analysis includes data from BICEP2 (2010-2012) (presented in [BICEP2 Collaboration I \(2014\)](#), [BICEP2 Collaboration II \(2014\)](#), and [BICEP2 Collaboration III \(2015\)](#)), *Planck* PR2 science release (presented in [Planck Collaboration Int. XXX \(2016\)](#)), and the *Keck Array* (2012-2013) (presented in [Keck Array and BICEP2 Collaborations V \(2015\)](#)).

We first introduce the experiments and their corresponding maps and power-spectra. We then define the analysis baseline and present likelihood results from this fiducial analysis choice. Next, we investigate a set of analysis and data variations and show likelihood results for these. Finally, we present a likelihood validation procedure that tests the robustness of the framework under the updated fiducial analysis and offer concluding remarks.

The framework presented in this manuscript has been used to produce all the analysis and figures discussed below. These have been peer-reviewed and published in [BICEP2/Keck and Planck Collaborations \(2015\)](#). Therefore, where it is deemed appropriate, this chapter follows the publication closely. Hereafter we refer to this analysis as BKP.

4.2 MAPS AND POWER SPECTRA

The processing of the maps proceeds as described in Section 3.2 and Chapter 2. We use the BICEP2/Keck combined maps, as described in *Keck Array and BICEP2 Collaborations V* (2015). In certain cases, we also use the BICEP2-only and *Keck*-only maps for cross-checks.

Due to potential systematic effects in the *Planck* maps, the *Planck* auto-spectra are computed as the cross-power spectra of two data-split maps, in which the data are split into two subsets with independent noise. We consider three data split maps: (i) detector-set maps (DS1/DS2), in which the detectors for a particular frequency are separated into two groups, (ii) yearly maps (Y1/Y2), in which the data from two different years are used for the two maps, and (iii) half-ring maps (HR1/HR2), in which each pointing period is divided into halves.

One particular place where this becomes important is in the Hamimeche-Lewis (H-L) approximation described in Section 3.6. As mentioned, usually in H-L using all the auto and cross-spectra from a set of maps guarantees positive-definiteness for $\hat{\mathbf{C}}_l$ (the amount of power obtained from taking the cross-spectrum of two maps can't exceed the geometric mean of the auto-spectrum power in those maps). However, this is not the case for our treatment of the bandpowers here. For example, the $P_{353_B}^{DS1} \times P_{353_B}^{DS2}$ cross-spectrum, which is used as the auto-spectrum for P_{353_B} , and occupies a diagonal position in the bandpower matrix, could fluctuate low due to chance anti-correlation of noise in the first and second halves of the data split. Note that the usual noise bias which is added in the H-L formalism is, in this case, zero since we are essentially operating with a cross-spectrum. To fix this, we follow Appendix C of

Hamimeche & Lewis (2008) and inject a cross-spectrum noise bias

$$N_l^{(H1 \times H2)} = \frac{1}{X} \text{mean} \left(N_l^{(H1)}, N_l^{(H2)} \right)$$

to regularize the $\hat{\mathbf{C}}_l$ matrix, and check that this has little effect on our constraints.

Figure 4.1 shows the resmoothed Planck 353 GHz T, Q, and U maps before and after filtering.

The maps are processed to auto- and cross-power-spectra as described in Section 3.2 and Chapter 2. Figure 4.2 shows the results for BICEP2/*Keck* and *Planck* 353 GHz for TT , TE , EE , and BB . While these are the auto- and cross- spectra for just one pair of experiments, we compute all possible auto- and cross- spectra. See Section II of [BICEP2/*Keck* and *Planck* Collaborations \(2015\)](#) for more details.

4.3 LIKELIHOOD ANALYSIS

We use the framework described in Chapter 3 to perform a joint analysis on all the auto- and cross-spectra. We form our bandpower covariance matrix from simulations with lensed- Λ CDM + noise + tensors (with $r = 0.2$). The fiducial model for the Hamimeche-Lewis approximation is a dust-only model with $r = 0$, $A_s = 0$ and $A_d = 3.6 \mu\text{K}_{CMB}^2$. For the *Planck* single-frequency case, the cross-spectrum of detector-sets (DS1 \times DS2) is used, following [Planck Collaboration Int. XXX \(2016\)](#). The baseline analysis, discussed in Section III.B. of [BICEP2/*Keck* and *Planck* Collaborations \(2015\)](#), is chosen to be the following: the




Figure 4.1: *Planck* 353 GHz T , Q , and U maps before (left) and after (right) the application of BICEP2/*Keck* filtering. The maps have been multiplied by the BICEP2/*Keck* apodization mask. The *Planck* maps are presmoothed to the BICEP2/*Keck* beam profile and have the mean subtracted. The filtering, in particular the third order polynomial subtraction to suppress atmospheric pickup, removes large-angular scale signal along the BICEP2/*Keck* scanning direction (parallel to the right ascension direction in the maps here). *Figure and caption from BICEP2/Keck and Planck Collaborations (2015)*

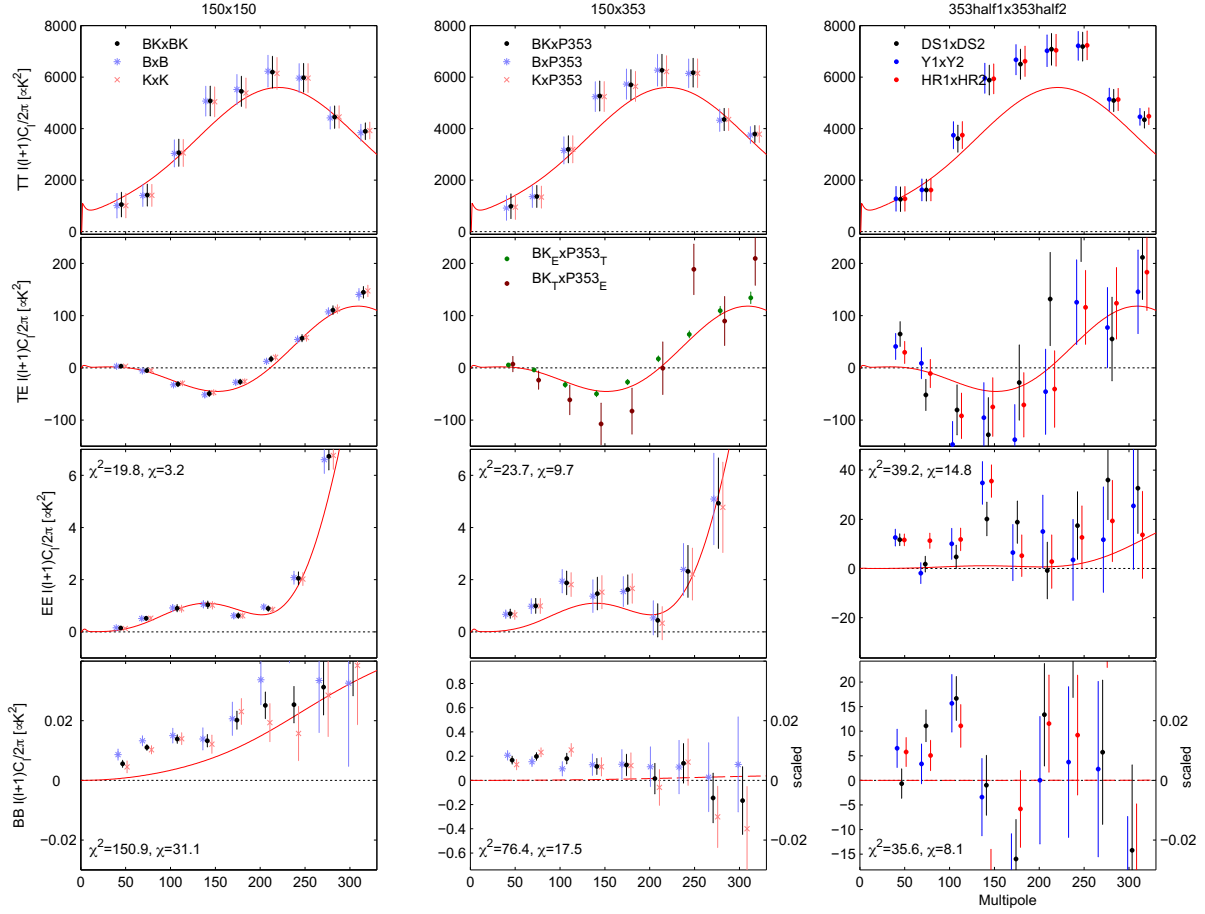


Figure 4.2: Single- and cross-frequency spectra between BICEP2/*Keck* maps at 150 GHz and *Planck* maps at 353 GHz. The red curves show the lensed- Λ CDM expectations. The left column shows single-frequency spectra of the BICEP2, *Keck Array* and combined BICEP2/*Keck* maps. The center column shows cross-frequency spectra between BICEP2/*Keck* maps and *Planck* 353 GHz maps. The right column shows *Planck* 353 GHz data-split cross-spectra. In all cases, the error bars are the standard deviations of lensed- Λ CDM+noise simulations. For EE and BB the χ^2 and χ (sum of deviations) versus lensed- Λ CDM for the nine bandpowers shown is marked at upper/lower left (for the combined BICEP2/*Keck* points and DS1 \times DS2). In the bottom row (for BB) the center and right panels have a scaling applied such that signal from dust with the fiducial frequency spectrum would produce a signal with the same apparent amplitude as in the 150 GHz panel on the left (as indicated by the right-side y -axes). The significant excess apparent in the bottom center panel points to the fact that a substantial amount of the signal detected at 150 GHz by BICEP2 and *Keck Array* could be due to dust. *Figure and caption from BICEP2/Keck and Planck Collaborations (2015)*

frequency selection is restricted to BK150, P217, and P353, the ℓ range is restricted to use the first five bandpowers ($20 < \ell < 200$) of the spectra, the likelihood parameter space is restricted to a three dimensional space r , A_d and β_d , over which we perform a simple grid evaluation.

Using BICEP2/*Keck* and *Planck* cross spectral bandpowers alone, it is not possible to constrain β_d . Therefore, a tight Gaussian prior $\beta_d = 1.59 \pm 0.11$ is imposed. This prior assumes that the SED of dust polarization at intermediate latitudes applies to the high latitude BICEP2/*Keck* field ([Planck Collaboration Int. XXX \(2016\)](#)). We also assume $n_t = 0$ and a scalar pivot scale of $k_s = 0.05 Mpc^{-1}$; all values of r quoted in this chapter are $r_{0.05}$.

4.3.1 BASELINE LIKELIHOOD RESULTS




Figure 4.3: Likelihood results from a lensed- Λ CDM+ r +dust model, fitting BB auto- and cross-spectra taken between maps at 150 GHz, 217, and 353 GHz. The 217 and 353 GHz maps come from *Planck*. The primary results (heavy black) use the 150 GHz combined maps from BICEP2/*Keck*. Alternate curves (light blue and red) show how the results vary when the BICEP2 and *Keck Array* only maps are used. In all cases a Gaussian prior is placed on the dust frequency spectrum parameter $\beta_d = 1.59 \pm 0.11$. In the right panel the two dimensional contours enclose 68% and 95% of the total likelihood. *Figure and caption from BICEP2/Keck and Planck Collaborations (2015)*

Figure 4.3 shows the marginalized 1D and 2D likelihoods of the baseline analysis. The

combined curves (BK+P) in the left and center panels yield: $r = 0.048^{+0.035}_{-0.032}$, $r < 0.12$ at 95% confidence, and $A_d = 3.3^{+0.9}_{-0.8}$. For r the zero-to-peak likelihood ratio is 0.38. Taking $\frac{1}{2}(1 - f(-2 \log L_0/L_{\text{peak}}))$, where f is the χ^2 cdf (for one degree of freedom), we estimate that the probability to get a number smaller than this is 8% if the null model is correct (i.e., if $r = 0$). For A_d the zero-to-peak ratio is 1.8×10^{-6} corresponding to a smaller-than-probability of 1.4×10^{-7} , and a 5.1σ detection of dust power. The maximum likelihood model has parameters $r = 0.05$, $A_d = 3.30 \mu\text{K}^2$ (and $\beta_d = 1.6$).

4.3.2 LIKELIHOOD VARIATIONS

Next, we investigate a number of variations and observe the effects of these choices on the final constraints. The results of these options are presented in Figure 4.4.

- **Choice of *Planck* data split:** we switch the *Planck* split from DS1xDS2 to Y1xY2.

This is the blue line in Figure 4.4. The result hardly changes.

- **Using only BK150 and P353:** dropping the 217 GHz spectra also has little effect.

This corresponds to the red line in Figure 4.4.

- **Using only one auto-spectrum (BK150×BK150) and one cross-spectrum (BK150×P353):** since the statistical weight of the BK150×BK150 and BK150×P353 spectra dominate, excluding the 353 GHz single-frequency spectrum from consideration makes little difference. This corresponds to the yellow line in Figure 4.4.




Figure 4.4: Likelihood results when varying the data sets used and the model priors—see Sec 4.3.2 for details. *Figure from BICEP2/Keck and Planck Collaborations (2015)*

- **Extending the bandpower range:** extending the multipole range to use all 9 bandpowers (corresponding to $20 < \ell < 330$) makes very little difference, since the dominant statistical weight is with the lower bandpowers. This corresponds to the magenta line in Figure 4.4.
- **Including EE spectra:** we can also include the EE spectra. As discussed in Section 3.3, Planck Collaboration Int. XXX (2016) shows that the level of EE from Galactic dust is twice the level of BB . We set $EE/BB = 2$ and find that the constraint on the dust amplitude narrows, while the r constraint changes little. This corresponds to the dashed black line in Figure 4.4.

- **Relaxing the prior on β_d :** we relax the prior on the dust spectral index to $\beta_d = 1.59 \pm 0.33$. This corresponds to the cyan line in Figure 4.4. We note that this pushes the peak of the marginalized 1D r likelihood up. However, as discussed in Section 3.3, it is important to realize that such a strong variation of the spectral index can lead to decorrelation, which is not accounted for in this analysis. Therefore we focus on the original narrow prior. In addition to this prior change, we also try two other prior choices: $\beta_d = 1.3 \pm 0.11$ and $\beta_d = 1.9 \pm 0.11$, for which the peak of the marginalized r likelihood shifts down to $r = 0.021$ and up to $r = 0.073$ respectively.
- **Using a Gaussian determinant likelihood:** as discussed in Section 3.6, by default we use the Hamimeche-Lewis (H-L) likelihood approximation. A usual alternative is to recompute the bandpower covariance matrix \mathbf{C} at each point in parameter space and take a Gaussian determinant form for the likelihood $L = \det(\mathbf{C})^{-1/2} \exp(-(\mathbf{d}^T \mathbf{C}^{-1} \mathbf{d})/2)$, where \mathbf{d} is the deviation of the observed bandpowers from the model expectation values. This corresponds to the green line in Figure 4.4. We note that this change results in a r constraint which peaks slightly lower.
- **Varying the HL fiducial model:** as mentioned in Section 3.6, the HL formalism requires fiducial model expectation values and covariance matrix. Our default fiducial model matches the one for the simulations with dust. In this variation, we change to a model with lensed- Λ CDM+ $r=0.2$. According to Hamimeche & Lewis (2008), changing the fiducial model should cause no ensemble average change in the recovered parame-

ters, though variations for any given realizations do occur. We observe this effect in our analysis, which corresponds to the dashed red line in Figure 4.4.

Two other variations, which are presented separately, are the expansion of the sampled likelihood parameter space to include the synchrotron amplitude (A_{sync}) and to include the lensing amplitude (A_L).




Figure 4.5: Likelihood results for a fit when adding the lower frequency bands of *Planck*, and extending the model to include a synchrotron component. The results for two different assumed degrees of correlation between the dust and synchrotron sky patterns are compared to those for the comparable model without synchrotron (see text for details). *Figure and caption from BICEP2/Keck and Planck Collaborations (2015)*

Adding synchrotron: Here we perform a fit including all the polarized bands of *Planck* and adding a synchrotron component to the base lensed- Λ CDM+noise+ r +dust model. We switch to the Y1 \times Y2 split for this variant since the DS1 \times DS2 data-split is not available for the *Planck* LFI bands. Therefore, we compare to the Y1 \times Y2 case in Figure 4.4 as opposed to

the fiducial analysis. The results are presented in Figure 4.5.

We take synchrotron to have a power law spectrum $\mathcal{D}_\ell \propto \ell^{-0.6}$ [Dunkley et al. \(2009\)](#), with free amplitude A_{sync} . For this particular analysis, the pivot frequency differs from the one described in Section 3.3, in particular A_{sync} is the amplitude at 150 GHz. The frequency spectral index is set to $\beta_s = -3.3$ (the mean value within the BICEP2 field of the MCMC “Model f” spectral index map provided by WMAP [Bennett et al. \(2013\)](#)).

When adding synchrotron and dust separately, we also have to assume a choice about the degree of correlation between the sky patterns of these two foregrounds. Figure 4.5 shows results for the uncorrelated and fully correlated cases. Marginalizing over r and A_d we find $A_{\text{sync}} < 0.0003 \mu\text{K}^2$ at 95% confidence for the uncorrelated case, and much smaller for the correlated case. This synchrotron limit is driven almost entirely by the low-frequency information in the *Planck* 30 GHz band—if we add just this band to the fiducial dataset, we obtain almost identical results. If we instead change the spectral index to $\beta_s = -3.0$ the limit on A_{sync} is doubled for the uncorrelated case and reduced for the correlated.

Varying the lensing amplitude: In our fiducial analysis we fix the lensing amplitude to the ΛCDM expectation value ($A_L = 1$). The *Planck* collaboration constraint at the time was $A_L = 0.99 \pm 0.05$ [Planck Collaboration XVI \(2014\)](#). Here, we use the entire ℓ range available to us, and we allow A_L to vary. The results are presented in Figure 4.6. Marginalizing over r and A_d we find $A_L = 1.13 \pm 0.18$ with a likelihood ratio between zero and peak of 3×10^{-11} . Taking $\frac{1}{2}(1 - f(-2 \log L_0/L_{\text{peak}}))$, where f is the χ^2 cdf (for one degree of freedom),




Figure 4.6: Likelihood results for a fit allowing the lensing scale factor A_L to float freely and using all nine bandpowers. Marginalizing over r and A_d , we find that $A_L = 1.13 \pm 0.18$ and $A_L = 0$ is ruled out with 7.0σ significance. *Figure and caption from BICEP2/Keck and Planck Collaborations (2015)*

this corresponds to a smaller-than probability of 2×10^{-12} , which is equivalent to a 7.0σ detection of lensing in the BB spectrum. We note that this was the most significant direct measurement of lensing in B -mode polarization at the time.

4.3.3 LIKELIHOOD VALIDATION

To validate our likelihood machinery, we perform tests on two types of simulations: first, we use the same lensed- Λ CDM + Gaussian dust + noise simulations as used for the BPCM construction for the fiducial model, and second, we use *Planck* Sky Model (PSM; version 1.7.8) simulations Delabrouille et al. (2013), since these are not necessarily Gaussian, and have much larger variation in the dust amplitude.

We run the fiducial analysis for all 499 Gaussian dust simulations. We present the marginal-




Figure 4.7: Likelihoods for r and A_d , using the fiducial BICEP2/Keck and *Planck* analysis, overplotted on constraints obtained from realizations of a lensed- Λ CDM+noise+dust model with dust power similar to that favored by the real data ($A_d = 3.6 \mu\text{K}^2$). Half of the r curves peak at zero as expected. *Figure and caption from BICEP2/Keck and Planck Collaborations (2015)*

ized 1D constraints on r and A_d for 100 realizations in Figure 4.7. As expected for an unbiased estimator, approximately 50% of the r likelihoods peak above zero. The median 95% upper limit is $r < 0.075$. We find that 8% of the realizations have a ratio L_0/L_{peak} less than the 0.38 observed in the real data, in agreement with the estimate in Sec. 4.3.1.

To prepare the PSM simulations, we take the *Planck* Sky Model evaluated in the *Planck* 353 GHz band, and pull out the 352 $|b| > 35$ deg partially overlapping regions used in *Planck Collaboration Int. XXX (2016)*. We scale the obtained maps to other bands and form spectra as we have done for the real *Planck* maps. We cut out regions with dust power that is much larger than the real data ($A_d < 20 \mu\text{K}^2$), selecting 139 regions. Figure 4.8 presents the resulting marginalized curves. We notice that the r likelihoods broaden as A_d increases, and it is

therefore unsurprising that the fraction of realizations peaking at a value higher than the real data is increased compared to the Gaussian dust simulations with $A_d = 3.6 \mu\text{K}^2$.

We expect on average that 50% of realizations will peak above $r = 0$, and around 8% of realizations to have a L_0/L_{peak} ratio less than the 0.38 observed in the real data. We find 57% and 7%, respectively, consistent with the expected values.




Figure 4.8: Constraints obtained when adding dust realizations from the *Planck* Sky Model version 1.7.8 to the base lensed- Λ CDM+noise simulations. (Curves for 139 regions with peak $A_d < 20 \mu\text{K}^2$ are plotted.) We see that the results for r are unbiased in the presence of dust realizations which do not necessarily follow the $\ell^{-0.42}$ power law or have Gaussian fluctuations about it. *Figure and caption from BICEP2/Keck and Planck Collaborations (2015)*

4.4 DISCUSSION AND CONCLUSIONS

In this chapter, we showed the first application of our multicomponent framework towards a joint analysis of CMB polarization data. We found that the excess B-mode signal detected by BICEP2 contains a large contribution from Galactic dust. The marginalized likelihood on r

peaks at $r = 0.05$, however, the zero-to-peak ratio is only 0.38. As noted, we expect this to happen by chance 8% of the time, which we confirm using Gaussian dust simulations. This significance is too low to be interpreted as a detection of primordial B-modes. We analyzed a number of variations and found the result to be robust under these changes.

In order to further constrain or detect PGWs, additional data with signal-to-noise comparable to that achieved by BICEP2/*Keck* at 150 GHz are required at more than one frequency. The BICEP2/*Keck* noise is much lower in the BICEP2/*Keck* field than the *Planck* noise, however, since dust emission is dramatically brighter at 353 GHz, the *Planck* data is required to detect dust. We do not detect synchrotron in this analysis.

During the 2014 season, two of the *Keck Array* receivers observed in the 95 GHz band, which is closer to the polarized foreground minimum ($\sim 80 - 90$ GHz). In the next Chapter, we will present results which include this data.

*The important thing is not to stop questioning.
Curiosity has its own reason for existence. One
cannot help but be in awe when he contem-
plates the mysteries of eternity, of life, of the
marvelous structure of reality. It is enough if
one tries merely to comprehend a little of this
mystery each day.*

Albert Einstein

5

Adding BICEP2/*Keck* observations for the 2014 Season

5.1 INTRODUCTION

In this chapter, we apply the likelihood methodology discussed in Chapter 3 to derive cosmological parameter constraints when adding new data taken by *Keck Array* in the 2014 season,

including three 150 GHz receivers, and, for the first time, two 95 GHz receivers. We also include, for the first time in our baseline analysis, all seven polarization auto-spectra from *Planck* (30-353 GHz), and add two WMAP frequencies at 23 and 33 GHz.

As in the previous chapter, we begin by introducing the experiments and their corresponding maps and power spectra. After defining the new analysis baseline, we derive likelihood results using this fiducial analysis choice. We then vary the analysis and data selection and demonstrate likelihood results for these. In conclusion, we validate the robustness of our likelihood framework under the updated fiducial analysis and offer a brief discussion.

The work in this manuscript: the likelihood framework, the analysis results presented, and the accompanying figures below, has been peer-reviewed and published in *Keck Array and BICEP2 Collaborations VI* (2016). Therefore, where it is deemed appropriate, this chapter follows the publication closely. Hereafter we refer to this analysis as BK14.

5.2 MAPS AND POWER SPECTRA

We form new maps and power spectra as described in Section 2 and Section 3.2, and similarly to Section 4.2. One difference is that in this analysis we use the Public Release 2 “full mission” maps available from the *Planck* Legacy Archive ¹, and use full auto-spectra for the *Planck* as opposed to single-frequency splits.

Figures 5.1 and 5.2 show the 95 & 150 GHz Q/U maps combining data from BICEP2

¹<http://www.cosmos.esa.int/web/planck/pla> or http://irsa.ipac.caltech.edu/data/Planck/release_2/all-sky-maps

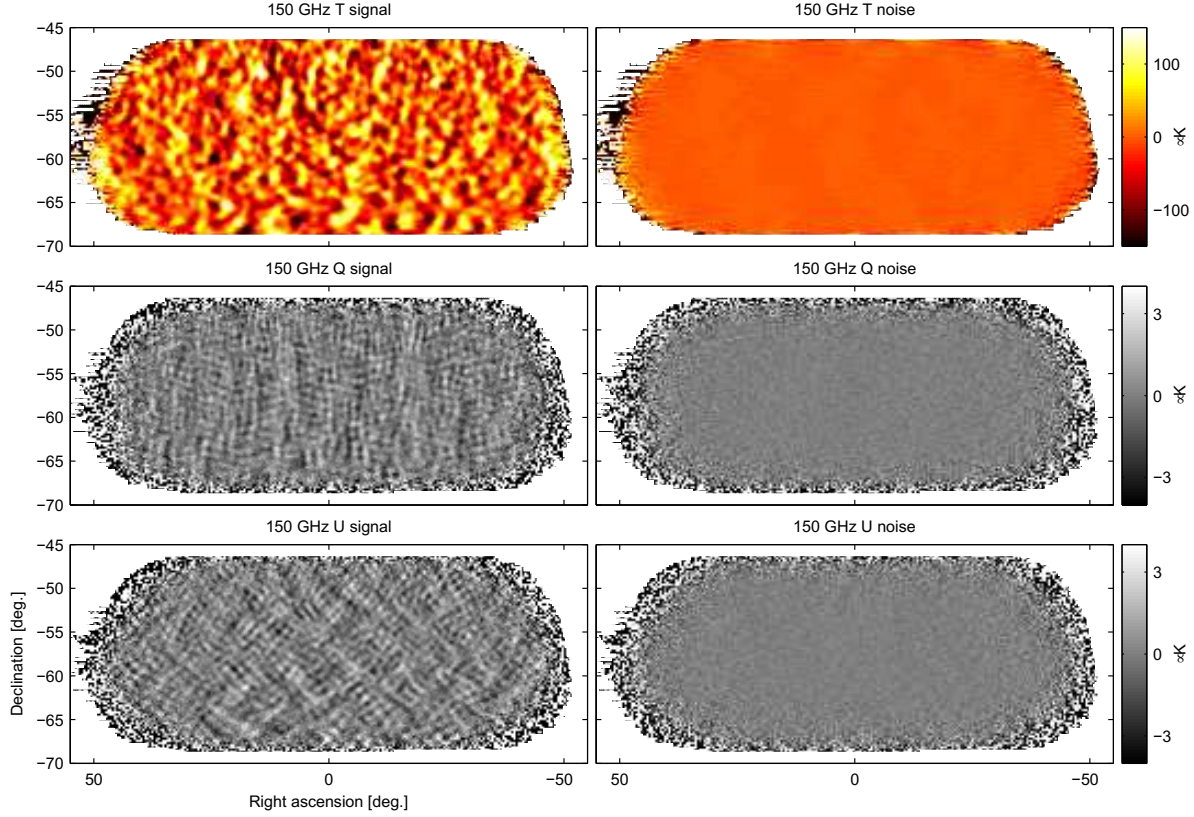


Figure 5.1: T , Q , U maps at 150 GHz using all BICEP2/*Keck* data up to and including the 2014 observing season—we refer to these maps as BK14₁₅₀. The left column shows the pipeline processed signal maps. The right column shows a noise realization. These maps are filtered by the instrument beam (FWHM 30 arcmin), timestream processing, and (for Q & U) deprojection of beam systematics, as described in Chapter 2. The horizontal/vertical and 45 deg structures seen in the Q and U signal maps are expected for an E -mode dominated sky. *Figure from Keck Array and BICEP2 Collaborations VI (2016)*

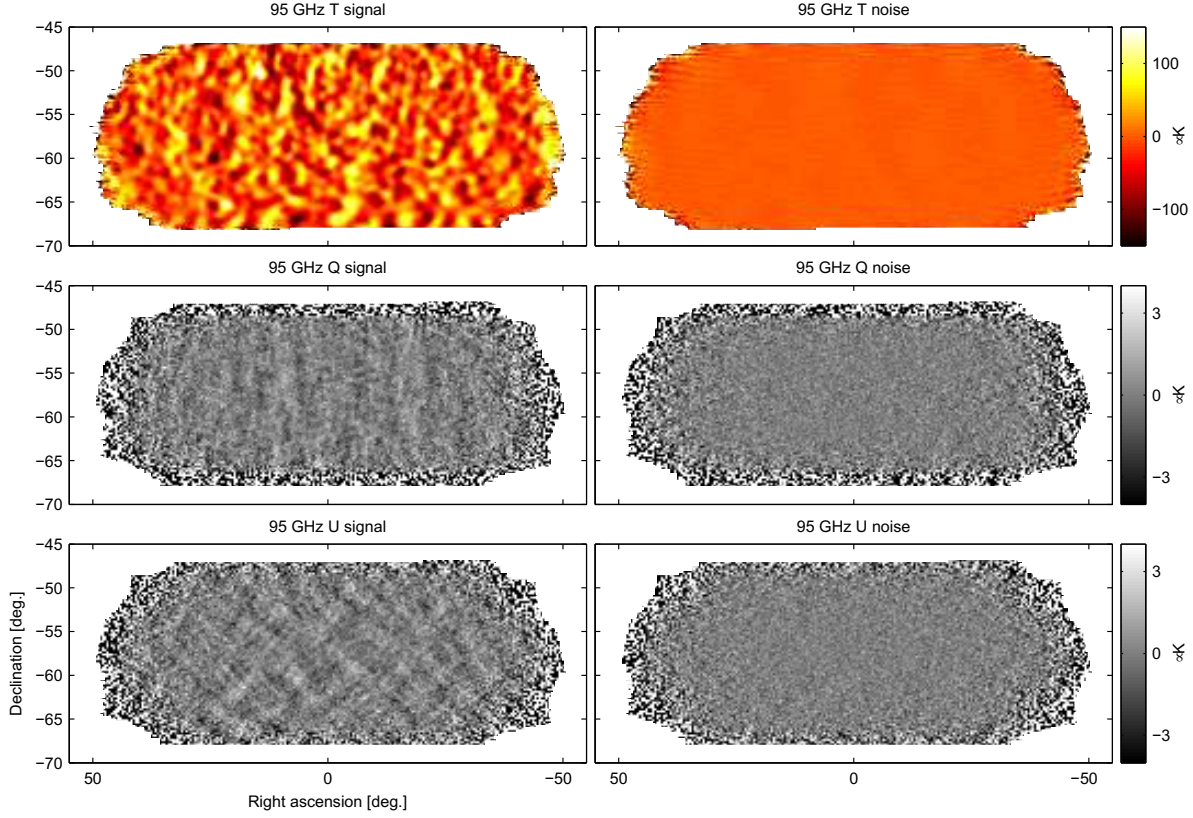


Figure 5.2: T , Q , U maps at 95 GHz using data taken by two receivers of *Keck Array* during the 2014 season—we refer to these maps as $BK14_{95}$. These maps are directly analogous to the 150 GHz maps shown in Figure 5.1 except that the instrument beam filtering is in this case 43 arcmin FWHM. *Figure and caption from Keck Array and BICEP2 Collaborations VI (2016)*

(2010–2012) and *Keck Array* (2012–2014). The 150 GHz maps include three more receiver years over the BKP analysis, improving the Q/U sensitivity from 57 nK deg to 50 nK deg (3.0 μ K arcmin) over an effective area of 395 square degrees. The Q/U sensitivity of the 95 GHz maps, which contain two receiver years of data, is 127 nK deg (7.6 μ K arcmin) over an effective area of 375 square degrees. Save for the difference in noise level and beam size (43 arcmin versus 30 arcmin FWHM), the degree scale structure, dominated by CMB E-modes, is




Figure 5.3: EE and BB auto- and cross-spectra between 95 & 150 GHz using all BICEP2/*Keck* data up to and including that taken during the 2014 observing season—we refer to these spectra as BK14. The points are offset horizontally. The solid black curves show the lensed- Λ CDM theory spectra. The error bars are the standard deviations of the lensed- Λ CDM+noise simulations. The χ^2 and χ (sum of deviations) against lensed- Λ CDM for the lowest five bandpowers are given and can be compared to their expectation value/standard-deviation of 5/3.1 and 0/2.2 respectively. The dashed lines show a lensed- Λ CDM+dust model derived from our previous analysis. *Figure and caption from Keck Array and BICEP2 Collaborations VI (2016)*

near identical in both maps.

In Figure 5.3 we present EE and BB auto- and cross- spectra between the $BK14_{95}$ and $BK14_{150}$ maps. In Figure 5.4 we show all BB auto- and cross-spectra between the BK14 95 & 150 GHz maps and the WMAP and *Planck* bands. These represent the set of baseline bandpowers.

5.3 LIKELIHOOD ANALYSIS

The Likelihood analysis proceeds as in Chapter 4. The evolution of the baseline analysis is presented in the following subsection. For this analysis, we switch to forming our bandpower

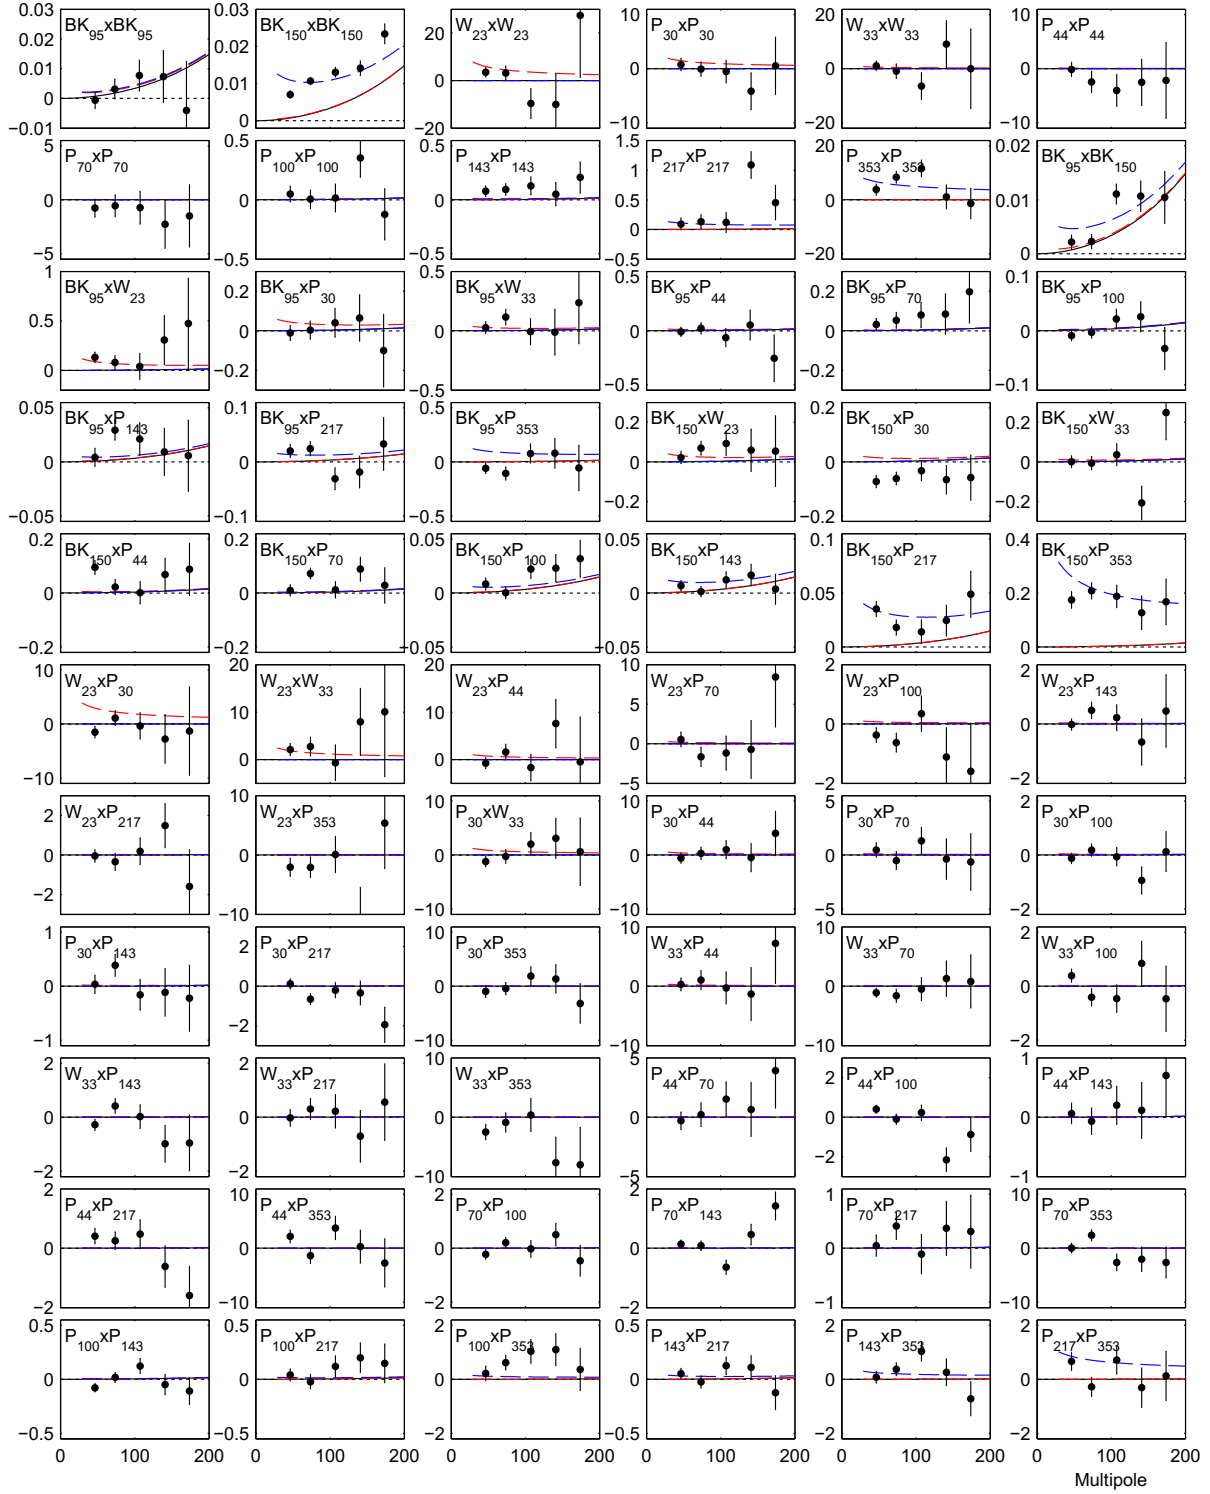


Figure 5.4: *BB* auto- and cross-spectra between the BK14 95 & 150 GHz maps and bands of WMAP and *Planck*. We plot $\ell(\ell+1) C_\ell/2\pi$ (μK^2). The black curves show the lensed- ΛCDM theory spectrum. The error bars are the standard deviations of the lensed- ΛCDM +noise simulations. The blue dashed lines show a baseline lensed- ΛCDM +dust model ($A_{\text{d},353} = 4.3 \mu\text{K}^2$, $\beta_{\text{d}} = 1.6$, $\alpha_{\text{d}} = -0.4$). The red dashed lines show an upper limit lensed- ΛCDM +synchrotron model ($A_{\text{sync},23} = 3.8 \mu\text{K}^2$, $\beta_{\text{s}} = -3.1$, $\alpha_{\text{s}} = -0.6$). *Figure and caption from Keck Array and BICEP2 Collaborations VI (2016)*

covariance matrix from lensed- Λ CDM+noise only simulations (as opposed to simulations that have tensors or dust). This significantly simplifies the data structure that we have to generate. The covariance matrix is then scaled to the desired fiducial model (which is the same as BKP and includes dust) for the H-L likelihood approximation. In this analysis, we expand the parameter space appreciably, and since a grid evaluation becomes prohibitively expensive, we switch to exploring the space with an MCMC, as described in Section 3.6.

5.3.1 LIKELIHOOD EVOLUTION

In this analysis, we extend the baseline from BKP to:

- Use $BK14_{95}$, $BK14_{150}$, all *Planck* frequencies (including auto-spectra), and two WMAP frequencies at 23 and 33 GHz.
- Use the full range of multipoles, i.e., we use bins 1-9 by default ($20 < \ell < 330$).
- Extend the model to include A_{sync} , with an updated central value for the spectral index $\beta_s = -3.1$, and a conservative prior of $\beta_s = -3.1 \pm 0.3$ (Fuskeland et al. (2014)). In this analysis we switch the synchrotron pivot frequency from the 150 GHz considered in BKP, to 23 GHz, as described in Section 3.3.
- Extend the model to allow for dust/synchrotron correlation, and marginalize over a generous range of $0 < \epsilon < 1$.
- Extend the model to marginalize over generous ranges of spatial indices $-1 < \alpha_s < 0$

and $-1 < \alpha_d < 0$ in case true dust and synchrotron deviate from a pure power-law ℓ scaling.

We perform likelihood analyses for each evolution documented above, in order to observe the effects on the marginalized constraints caused by each individual change. The results for this evolution sequence are presented in Figure 5.5. We note that with the exception of the addition of the new *BK* data at 95 and 150 GHz, the rest of the changes have little effect on the r constraints. The addition of WMAP data and full *Planck* auto-spectra, the inclusion of dust/synch correlation, and the updated spectral index, all serve to improve the synchrotron constraints.




Figure 5.5: Evolution of the BKP analysis to the “baseline” analysis as defined for BK14 *Figure from Keck Array and BICEP2 Collaborations VI (2016)*

5.3.2 BASELINE LIKELIHOOD RESULTS

Figure 5.6 shows the results of this baseline analysis. We obtain the following statistics: $r_{0.05} = 0.028^{+0.026}_{-0.025}$, $r_{0.05} < 0.090$ at 95% confidence, $A_{d,353} = 4.3^{+1.2}_{-1.0} \mu\text{K}^2$, and $A_{\text{sync},23} < 3.8 \mu\text{K}^2$ at

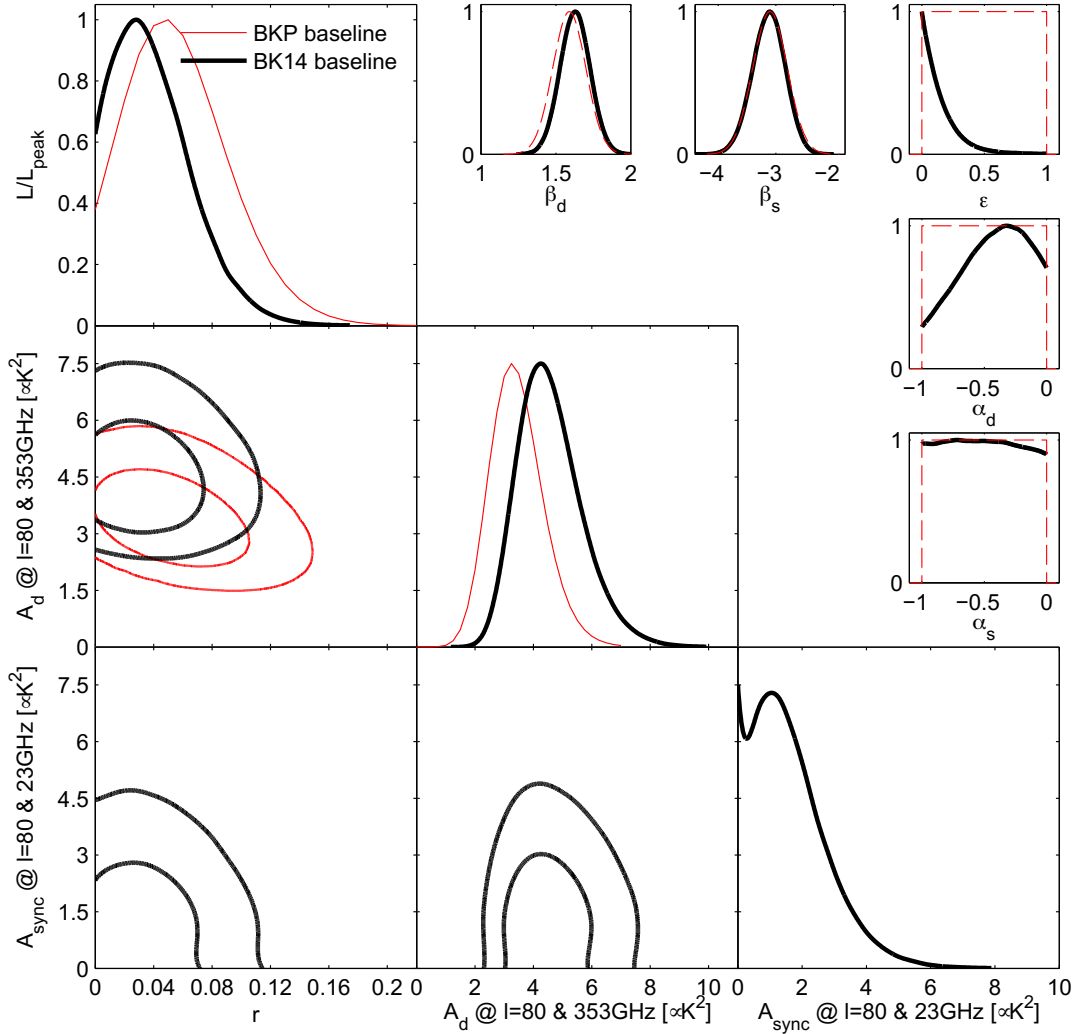


Figure 5.6: Results of the multicomponent multi-spectral likelihood analysis BK14 data. The faint red curves are the primary result from the previous BKP analysis. The bold black curves are the new baseline BK14 results. Differences between these analyses include adding synchrotron to the model, including additional external frequency bands from WMAP & *Planck*, and adding *Keck Array* data taken during the 2014 observing season at 95 & 150 GHz. We see that the peak position of the tensor/scalar ratio curve r shifts down slightly and the upper limit tightens to $r_{0.05} < 0.09$ at 95% confidence. In the β , α and ϵ panels the dashed red lines show the priors placed on these parameters (either Gaussian or uniform). *Figure from Keck Array and BICEP2 Collaborations VI (2016)*




Figure 5.7: Constraints in the r vs. n_s plane when using *Planck* in conjunction with external data + BICEP2/*Keck* data through the end of the 2014 season. The constraint on r tightens from $r_{0.05} < 0.12$ to $r_{0.05} < 0.07$. This figure is adapted from Figure 21 of Ref. [Planck Collaboration 2015 XIII \(2016\)](#)—see there for further details. *Figure from Keck Array and BICEP2 Collaborations VI (2016)*

95% confidence. The zero-to-peak likelihood ration for r is 0.63. Taking $\frac{1}{2}(1 - f(-2 \log L_0/L_{\text{peak}}))$, where f is the χ^2 cdf (for one degree of freedom), the probability to get a ration smaller than the one for the real data is 18% under the null-model assumption. Running the analysis on the lensed- Λ CDM+dust+noise simulations produces a similar number. The zero-to-peak likelihood ratio for A_d indicates a detection of dust at $> 8\sigma$. The maximum likelihood model (including priors) has parameters $r = 0.026$, $A_{d,353} = 4.1 \mu\text{K}^2$, $A_{\text{sync},23} = 1.4 \mu\text{K}^2$, $\beta_d = 1.6$, $\beta_s = -3.1$, $\alpha_d = -0.19$, $\alpha_s = -0.56$, and $\epsilon = 0.00$.

The constraint on r further tightens when including *Planck* temperature data. The analysis of *Planck* full mission TT data in conjunction with external data produces a cosmic variance limited constraint of $r_{0.002} < 0.11$ ($r_{0.05} < 0.12$) at 95% confidence (“*Planck* TT +lowP+lensing+ext” in Equation 39b of Ref. [Planck Collaboration 2015 XIII \(2016\)](#)). The baseline BK14 result is the first B -mode constraint which surpasses those from temperature anisotropies. We use **COSMOMC** to include all the same datasets to reproduce the Ref. [Planck Collaboration 2015 XIII \(2016\)](#)’s result, in the r vs. n_s plane, and show the effect of adding in our BK14 B -mode data. This is presented in Figure 5.7. The joint result produces a constraint of $r_{0.05} < 0.07$ (95%), which represents the best PGW constraint at the time of this analysis. Note that this constraint excludes the ϕ^2 model at many sigma. This particular model very alluring for a significant period of time because it has only one free parameter m , which, once determined by the level of the scalar perturbations, is enough to predict all the other early universe properties. Ruling out this simplistic model suggests that the inflaton field was doing something more complicated as it rolled towards the minimum. For a comprehensive discussion about the implications of the constraints presented in this chapter for various inflationary models, see [Planck Collaboration et al. \(2018b\)](#).

5.3.3 LIKELIHOOD VARIATION

Similar to the BKP analysis, we perform several variations, the results of which are presented in Figure 5.8.

The considered options are as follows:




Figure 5.8: Likelihood results when varying the data sets and the model priors. *Figure from Keck Array and BICEP2 Collaborations VI (2016)*

- **No β_d prior:** we remove the prior on the dust frequency spectral index β_d . This corresponds to the cyan line in Figure 5.8. The marginalized constraint on β_d is approximately Gaussian, with mean/ σ of 1.82/0.26. For this variant, the dust amplitude A_d shifts up, as does the r constraint. The posterior on r is characterized by $r = 0.043^{+0.033}_{-0.031}$ with a zero-to-peak likelihood ratio of 0.44 (10% likely if $r = 0$).
- **Relax the β_s prior:** we relax the prior on the synchrotron frequency spectral index to $-4 < \beta_s < -2$. This corresponds to the green line in Figure 5.8. We can see that this has very little effect
- **Remove the *Planck* low-frequency (LFI) data:** we remove all the *Planck* LFI bands from consideration. This corresponds to the magenta line in Figure 5.8. The peak of the r constraint shifts slightly down, while the A_{sync} constraint peaks strongly away from zero.

- **Remove the WMAP data:** we repeat the exercise above but this time remove low-frequency data from WMAP. This corresponds to the dark yellow line in Figure 5.8. Yet again, this change slightly decreases the zero-to-peak ratio of the r constraint and significantly tightens the A_{sync} constraint.
- **Fix the dust/sync correlation:** recently, Choi & Page (2015) found modest evidence for dust/sync correlation, with $\epsilon = 0.2$. We explore this variant here, represented by the dashed red line in Figure 5.8, and observe that it makes almost no difference to any of the constraints.
- **Add E -mode data:** finally, just as we did for BKP, we add in E -mode data with a fixed EE/BB ratio of 2 for both synchrotron and dust. This corresponds to the dashed black line in Figure 5.8. Assuming this fixed ratio holds, this offers extra constraining power – we see the r curve shift up and narrow slightly, similarly for A_d , while A_s peaks strongly away from zero. Unfortunately, there is no good estimate for a prior on the EE/BB ratios, so we cannot confidently make this extra information part of the baseline analysis.

5.3.4 LIKELIHOOD VALIDATION

Similarly to the BKP analysis, we would like to validate that the likelihood framework recovers unbiased estimates under our new baseline analysis. We test this against the same lensed- Λ CDM + noise + Gaussian dust simulations as in BKP. Since we have switched to using

COSMOMC, we can no longer explore negative amplitudes (negative power leads to a technical issue within COSMOMC). Therefore, in this chapter, we cannot perform full likelihood evaluations as we have done for Section 4.3.3. Instead, we implement a Maximum Likelihood search that uses `minuit`² to find the global likelihood maximum. In fact, since we would like to impose the same priors as we do for the baseline analysis, this is technically a maximum posterior exploration. Figure 5.9 shows the results—the input values are recovered in the mean as expected.

An additional statistic that comes from this particular study is the standard deviation of the recovered ML values for r . Unlike the width of the marginalized posterior, which depends on the peak position of the curve, this new estimate does not suffer from this and is, therefore, a more robust measure of the intrinsic constraining power of the experimental data.




Figure 5.9: Results of validation tests running the likelihood on simulations of a lensed- Λ CDM+dust model ($A_{d,353} = 3.75 \mu\text{K}^2$, $\beta_d = 1.59$ and $\alpha_d = -0.42$). The blue histograms are the recovered ML values with the red line marking their means. The black line shows the input value. In the left panel $\sigma(r) = 0.024$. See Section 5.3.4 for details. *Figure from Keck Array and BICEP2 Collaborations VI (2016)*

²<https://seal.web.cern.ch/seal/snapshot/work-packages/mathlibs/minuit/>

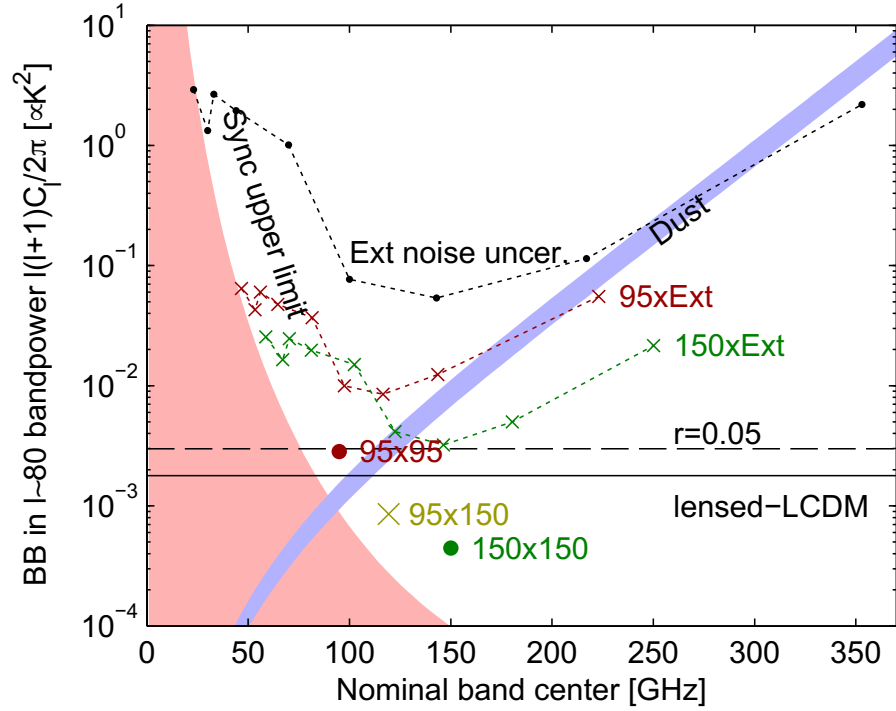


Figure 5.10: Expectation values and noise uncertainties for the $\ell \sim 80$ BB bandpower in the BICEP2/Keck field. The solid and dashed black lines show the expected signal power of lensed- Λ CDM and $r = 0.05$. The blue band shows a dust model consistent with the baseline analysis ($A_d = 4.3^{+1.2}_{-1.0} \mu\text{K}^2$, $\beta_d = 1.6$) while the pink shaded region shows the allowed region for synchrotron ($A_s < 3.8 \mu\text{K}^2$, $\beta_s = -3.1$). The BICEP2/Keck noise uncertainties are shown as large colored symbols, and the noise uncertainties of the WMAP/Planck single-frequency spectra evaluated in the BICEP2/Keck field are shown in black. The red (green) crosses show the noise uncertainty of the cross-spectra taken between 95 (150) GHz and, from left to right, 23, 30, 33, 44, 70, 100, 143, 217 & 353 GHz, and are plotted at horizontal positions such that they can be compared vertically with the dust and sync curves. *Figure and caption from Keck Array and BICEP2 Collaborations VI (2016)*

5.4 DISCUSSION AND CONCLUSIONS

The analysis presented in this chapter is significantly evolved from the one in Chapter 4: we have expanded our foreground parametrization and folded in new data-sets. We have seen that with the addition of extra data we are able to further tighten the constraints on r from $r_{0.05} < 0.12$ at 95% confidence to $r_{0.05} < 0.07$. With this result, B-modes have become the most powerful probe of PGWs, surpassing constraints from temperature data for the first time. These constraints are set to only get better from here. Figure 5.10 compares signal levels and current noise uncertainties at $\ell \sim 80$. During the 2015 season, the *Keck Array* has observed at three frequencies: 95, 150, 220 GHz, and will push these noise levels down, allowing us to improve the separation between foregrounds and a possible PGW signal. In the next chapter, we will add data from this season and describe further updates to our likelihood framework, which help us increase the robustness of our constraints.

Life is a marathon, not a steeplechase.

Iurie Cuza

6

Adding BICEP2/*Keck* Observations for the 2015 Season

6.1 INTRODUCTION

In this chapter, we apply the likelihood methodology discussed in Chapter 3 to derive cosmological parameter constraints when adding new data taken by *Keck Array* in the 2015 sea-

son, including two 95 GHz receivers, a single 150 GHz receiver, and, for the first time, two 220 GHz receivers. Compared to the dataset used in the previous chapter, this analysis doubles the 95 GHz dataset to four receiver-years and adds a new higher frequency band that significantly improves the constraints on the dust contribution compared to those given by using the *Planck* data alone.

From 2010 to 2013, BICEP2 and *Keck Array* jointly recorded a total of 13 receiver-years of data in a band centered on 150 GHz. Two of the *Keck* receivers were switched to 95 GHz before the 2014 season and two more were switched to 220 GHz before the 2015 season. The dataset used in this chapter – the BK15 dataset – thus consists of 4/17/2 receiver-years at 95/150/220 GHz respectively.

As in the previous two chapters, we describe experiments and how their corresponding maps and power-spectra are formed. We proceed by defining the new analysis baseline and present the fiducial likelihood results. Next, we study a set of analysis and data variations and demonstrate likelihood results for these. Finally, we present a likelihood validation procedure that tests the robustness of the framework under the updated analysis choices.

The likelihood analysis presented in this chapter has been peer-reviewed and published in [BICEP2 Collaboration et al. \(2018\)](#), and is a direct result of the work described in this manuscript. Therefore, where it is deemed appropriate, this chapter follows the publication closely. Hereafter we refer to this analysis as BK15.

6.2 MAPS AND POWER SPECTRA

We make maps and power spectra using the same procedures as used for BKP and BK14, described in Chapters 4 and 5. The full set of BICEP/*Keck* $T/Q/U$ maps are presented in Figure 6.1, Figure 6.2, Figure 6.3.

In this analysis, we used the three bands of BICEP2/*Keck* plus the 23 & 33 GHz bands of WMAP and all seven polarized bands of *Planck*. As before, we take all possible auto- and cross-power spectra between these bands. The full set of spectra are shown in Figure 6.4.

Here, for the first time, in Figure 6.5, we choose to show the noise spectra for the three BK15 bands after correction for the filter and beam suppression, and the effective observed sky fraction inferred from the noise fluctuations. The turn up at low- ℓ at least in part due to residual atmospheric $1/f$ in the pair-difference data. We note that it is weakest in the 95 GHz band, where water vapor emission is lowest. These quantities provide a useful synoptic measure of the loss of information due to noise, filtering, and EE/BB separation in the lowest bandpowers.

6.3 LIKELIHOOD ANALYSIS

We perform likelihood analysis using the methods introduced in Chapter 2, which were used in BKP and refined in BK14. Here, as in our previous analysis of BK14 data, the bandpower covariance matrix is derived from 499 simulations of lensed- Λ CDM + noise according to the prescription described in B. We rescale the covariance matrix to the same fiducial model as

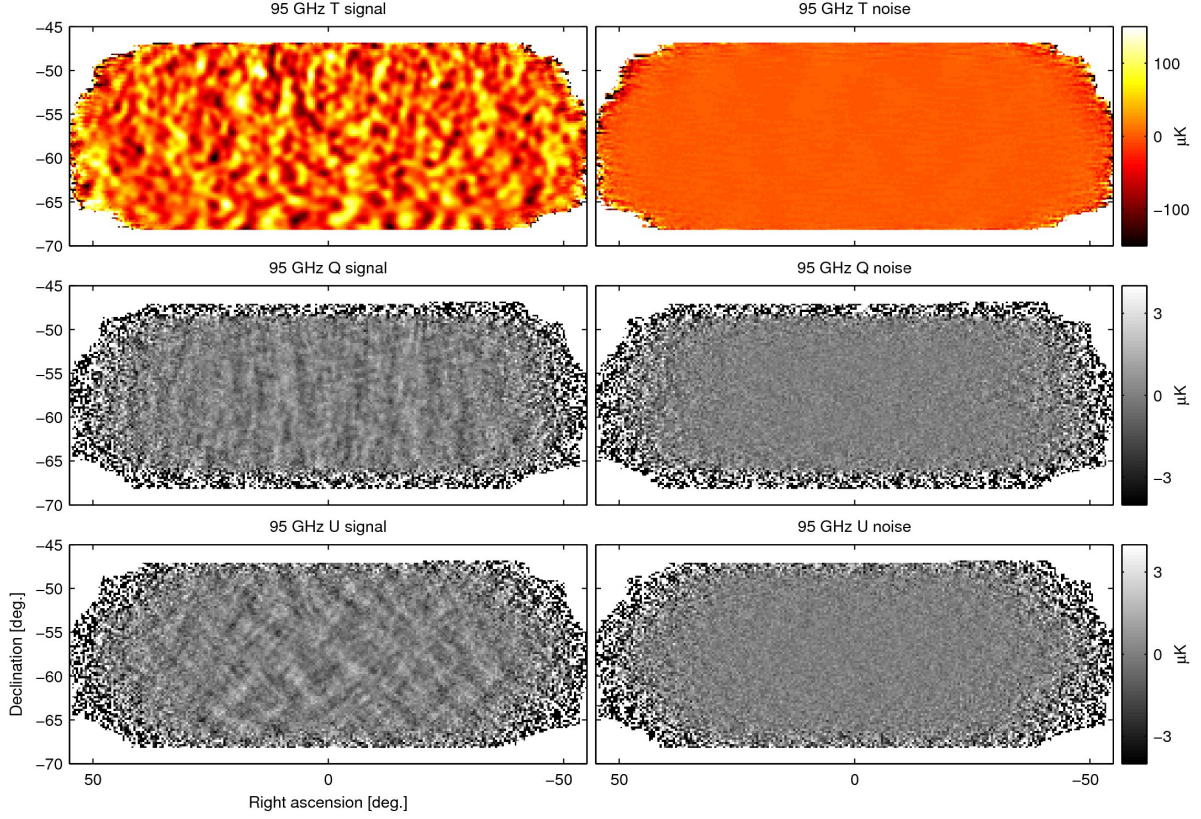


Figure 6.1: T , Q , U maps at 95 GHz using data taken by two receivers of *Keck Array* during the 2014 & 2015 seasons—we refer to these maps as BK15₉₅. The left column shows the real data maps processed through the analysis pipeline. The right column shows a noise realization. These maps are filtered by the instrument beam (FWHM 43 arcmin *Keck Array and BICEP2 Collaborations XI (2018)*), timestream processing, and (for Q & U) deprojection of beam systematics. The horizontal/vertical and 45 deg structures seen in the Q and U signal maps are indicative of an E -mode dominated sky. *Figure and caption from BICEP2 Collaboration et al. (2018)*

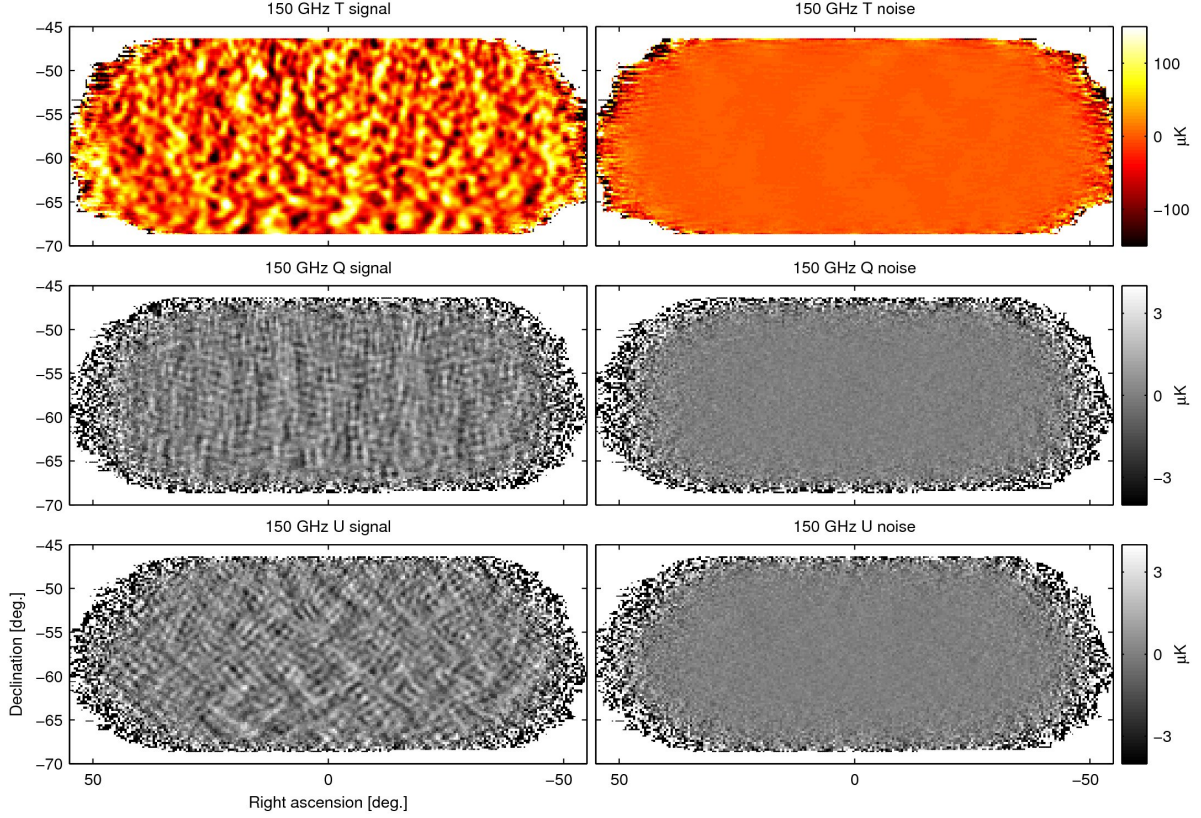


Figure 6.2: T , Q , U maps at 150 GHz using all BICEP2/*Keck* data up to and including that taken during the 2015 observing season—we refer to these maps as BK15₁₅₀. These maps are directly analogous to the 95 GHz maps shown in Figure 6.1 except that the instrument beam filtering is in this case 30 arcmin FWHM *Figure and caption from BICEP2 Collaboration et al. (2018).*

BKP and BK14 ($r = 0$, $A_d = 3.6$, $A_s = 0$, $\beta_d = 1.59$, $\alpha_d = -0.42$, $T_d = 19.6K$)

6.3.1 LIKELIHOOD EVOLUTION

We make only one change to the baseline analysis choices of BK14, expanding the prior on the dust/sync correlation parameter from $0 < \epsilon < 1$ to $-1 < \epsilon < 1$. This choice was made based on the fact the in BK14 the posterior on ϵ was peaking sharply at zero, and while there

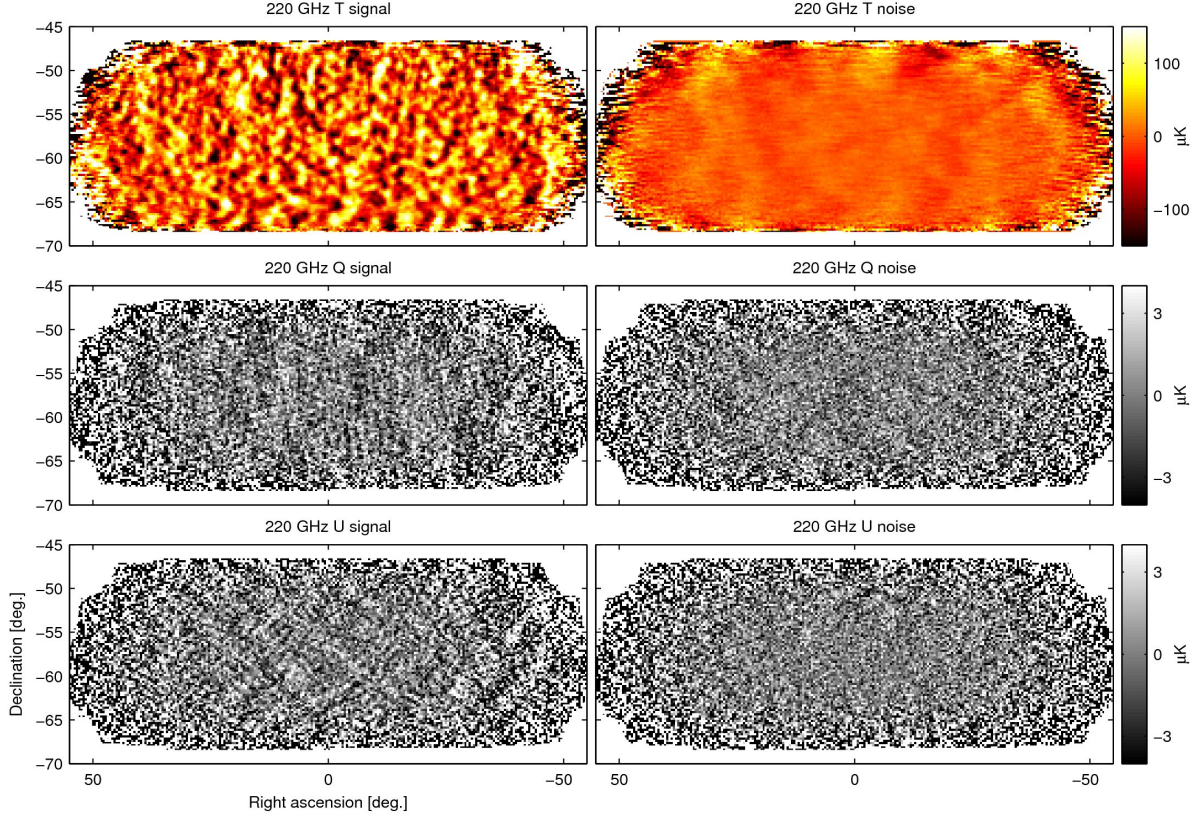


Figure 6.3: T , Q , U maps at 220 GHz using data taken by two receivers of *Keck Array* during the 2015 season—we refer to these maps as $BK15_{220}$. These maps are directly analogous to the 95 GHz maps shown in Figure 6.1 except that the instrument beam filtering is in this case 20 arcmin FWHM. Figure and caption from *BICEP2 Collaboration et al. (2018)*.

is no current theoretical motivation for negative correlation values, anti-correlation of dust and synchrotron spatial patterns is mathematically possible.

In Figure 6.6 we show the evolution of the $BK14$ baseline to the current one. First, we make the change in the ϵ prior, then we update the 95 and 150 GHz selection to include data from *Keck* taken during the 2015 season, and finally, we add in the new *Keck* data at 220 GHz. We note that the net result is a tightening and a slight downward shift of the r poste-




Figure 6.4: The full set of BB auto- and cross-spectra from which the joint model likelihood is derived. The quantity plotted is $100\ell C_\ell/2\pi$ (). Spectra involving BICEP2/Keck data are shown as black points; those using only WMAP/Planck data are shown as blue points. The black lines show the expectation values for lensed- Λ CDM, the red lines show the expectation values of the maximum likelihood lensed- Λ CDM+ r +dust+synchrotron model ($r = 0.02$, $A_{d,353} = 4.7\mu K^2$, $\beta_d = 1.6$, $\alpha_d = -0.58$, $A_{sync,23} = 1.5\mu K^2$, $\beta_s = -3.0$, $\alpha_s = -0.27$, $\epsilon = -0.38$), and the error bars are scaled to that model. Figure and caption from *BICEP2 Collaboration et al. (2018)*




Figure 6.5: *Upper:* The noise spectra of the BK15 maps for 95 GHz (red), 150 GHz (green) and 220 GHz (blue) after correction for signal filtering. No ℓ^2 scaling is applied. *Lower:* The effective sky fraction as calculated from the ratio of the mean noise realization bandpowers to their fluctuation $f_{\text{sky}}(\ell) = \frac{1}{2\ell\Delta\ell} \left(\frac{\sqrt{2}\bar{N}_b}{\sigma(\bar{N}_b)} \right)^2$, i.e. the observed number of B -mode degrees of freedom divided by the nominal full-sky number. The turn-down at low ℓ is due to mode loss to the timestream filtering and matrix purification. *Figure from BICEP2 Collaboration et al. (2018)*

rior.

6.3.2 BASELINE LIKELIHOOD RESULTS

Figure 6.7 shows the results of the BK15 baseline analysis. We obtain the following statis-

tics: $r_{0.05} = 0.020^{+0.021}_{-0.018}$ ($r_{0.05} < 0.072$ at 95% confidence), $A_{\text{d},353} = 4.6^{+1.1}_{-0.9} \mu\text{K}^2$ and

$A_{\text{sync},23} = 1.0^{+1.2}_{-0.8} \mu\text{K}^2$ ($A_{\text{sync},23} < 3.7 \mu\text{K}^2$ at 95% confidence). For r , the zero-to-peak likeli-

hood ratio is 0.66. Taking $\frac{1}{2}(1 - f(-2\log L_0/L_{\text{peak}}))$, where f is the χ^2 CDF (for one degree




Figure 6.6: Evolution of the BK14 analysis to the “baseline” analysis. *Figure from [BICEP2 Collaboration et al. \(2018\)](#)*

of freedom), we estimate that the probability to get a likelihood ratio smaller than this is 18% if, under the null hypothesis. Compared to the baseline result in BK14, the marginalized posterior on r slightly tightens, and shifts down. The posterior on the dust amplitude seems saturated by the sample variance of dust. The posterior on the synchrotron amplitude does not have the double-peak feature anymore due to the new prior on ϵ ; this significantly increases the zero-to-peak likelihood ratio.

The maximum likelihood model (including priors) has parameters $r_{0.05} = 0.020$, $A_{d,353} = 4.7 \mu\text{K}^2$, $A_{sync,23} = 1.5 \mu\text{K}^2$, $\beta_d = 1.6$, $\beta_s = 3.0$, $\alpha_d = 0.58$, $\alpha_s = 0.27$, and $\epsilon = 0.38$.

Using this same baseline, we repeat the r vs n_s analysis as in BK14. Figure 6.8 shows the constraints in the r vs. n_s plane for *Planck* 2015 plus additional data ($r_{0.05} < 0.12$) and when adding in also BK15 ($r_{0.05} < 0.057$). We note that the ϕ model, as well as the entire class of single-field monomial models ϕ^p , now lies entirely outside of the 95% contour. This is a significant reduction of inflationary model parameter space.

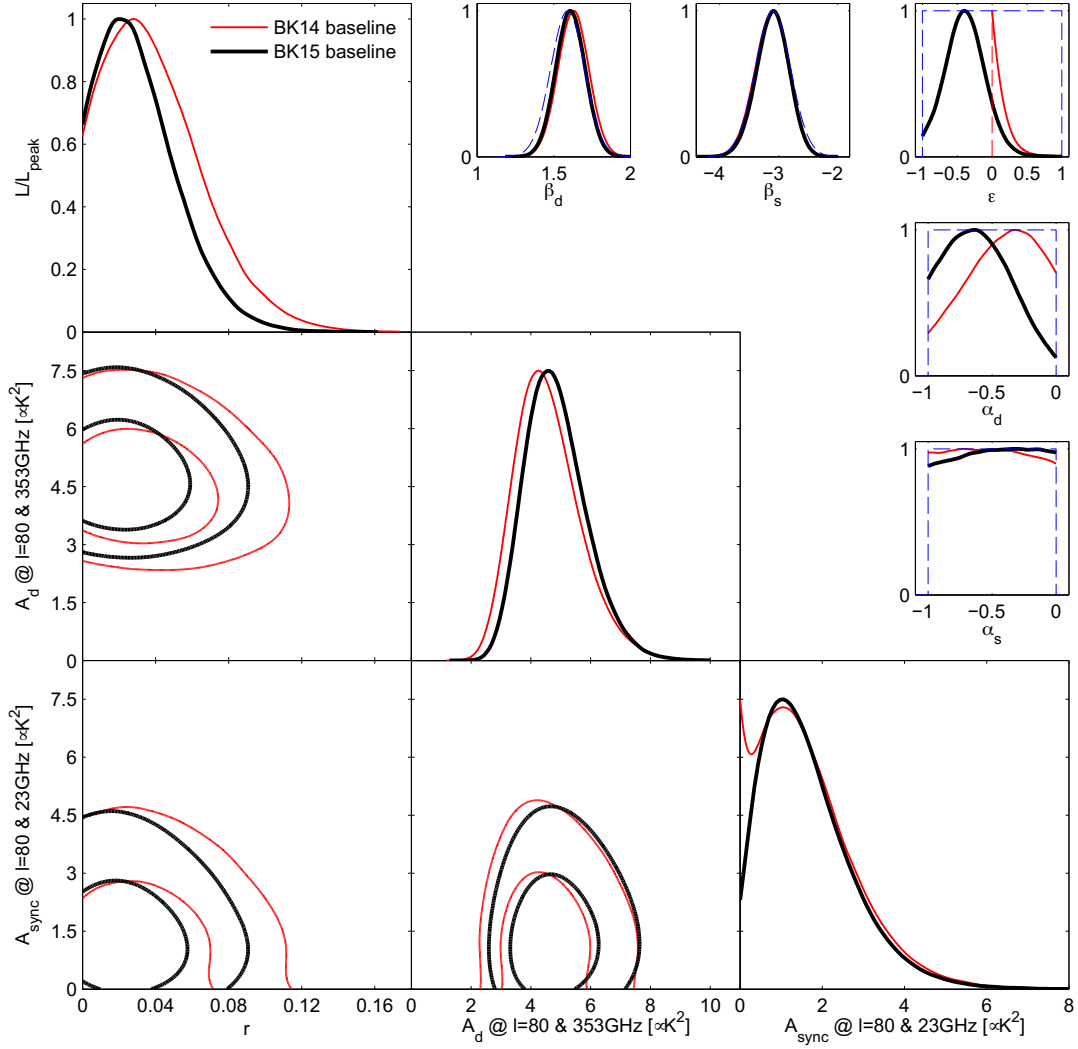


Figure 6.7: Results of a multicomponent multi-spectral likelihood analysis of BICEP2/*Keck*+WMAP/*Planck* data. The red curves are the baseline result from the BK14 analysis. The black curves are the new baseline BK15 results. The upper limit on the tensor-to-scalar ratio tightens to $r_{0.05} < 0.05$ at 95% confidence. In the β , α and ϵ panels the dashed lines show the priors placed on these parameters (either Gaussian or uniform). *Figure from BICEP2 Collaboration et al. (2018)*




Figure 6.8: Constraints in the r vs. n_s plane when using *Planck* plus additional data, and when also adding BICEP2/*Keck* data through the end of the 2015 season—the constraint on r tightens from $r_{0.05} < 0.12$ to $r_{0.05} < 0.06$. This figure is adapted from Figure 21 of Ref. [Planck Collaboration 2015 XIII \(2016\)](#), with two notable differences: switching *lowP* to *lowT* + a τ prior of 0.055 ± 0.009 Ref. [Planck Collaboration XLVI \(2016\)](#) and the exclusion of JLA data and the H_0 prior.

6.3.3 LIKELIHOOD VARIATIONS

As in the previous two chapters, we consider a number of variations to the baseline analysis.

In this particular case, because of the additional constraining power coming from the new bands and extra data at existing frequencies, we can test a larger set of data selection choices.

Therefore, we split the variations into two types: “data set selection” and “analysis selection”.

6.3.3.1 ANALYSIS VARIATIONS

The first set of analysis variations is similar in nature to the alternatives considered in BK14.




Figure 6.9: Likelihood results when varying the analysis choices—see Section 6.3.3 for details. *Figure from BICEP2 Collaboration et al. (2018)*

- **Alternative fiducial models:** we try exchanging the fiducial model used in the H-L approximation from the default model (with $A_{d,353} = 3.6 \mu\text{K}^2$, $A_s = 0$, $r = 0$) used for BKP, BK14, and in the baseline analysis here, to a model with $A_{d,353} = 5 \mu\text{K}^2$, $A_{\text{sync}} = 0$, $r = 0.05$ or one with $A_{d,353} = 5 \mu\text{K}^2$, $A_{\text{sync},23} = 2 \mu\text{K}^2$, $r = 0.05$. These correspond to the blue and red lines in Figure 6.9. We note that the corresponding changes are very small.
- **Relaxing the prior on β_d :** Similarly to BKP and BK14, we remove the prior from *Planck* on β_d . This corresponds to the magenta line in Figure 6.9. In BK14 we observed that this led to an upward shift in r , and a significant broadening of A_d . Here, because of the new data at 220 GHz, the changes are much smaller. The β_d constraint

curve (not shown) is roughly Gaussian, with mean/ σ of 1.65/0.20. We note that with further data additions from the *Keck Array* at high frequencies, we should no longer need the external prior on β_d , thereby removing the uncertainty from the assumption that the average behavior over large sky regions is a good stand-in for dust behavior in our sky patch.

- **Relaxing the prior on β_s :** We extend our prior on the synchrotron spectral index from a Gaussian prior with $\beta_s = -3.1 \pm 0.3$ to a uniform prior with $-4.5 < \beta_s < -2.0$. This corresponds to the green line in Figure 6.9. We note that this choice has hardly any effect on the final marginalized posteriors. This is unsurprising given that we are not yet detecting synchrotron.
- **Fixing ϵ to zero:** We fix the value of the dust/sync correlation to $\epsilon = 0$. This corresponds to the cyan line in Figure 6.9. This choice produces a small downshift in r , as expected from Figure 6.6. We include this variant because we will use it when adding in dust decorrelation as an analysis variant (see Figure 6.10 and the appropriate discussion).
- **Setting a Gaussian prior on ϵ :** We see a similar shift if we instead choose to impose a Gaussian prior on the dust/sync correlation with mean/ σ of 0.48/0.50 ([Planck Collaboration 2018 XI \(2018\)](#)). This corresponds to the yellow line in Figure 6.9.
- **Marginalizing over bandpass uncertainty:** We add to the model a set of param-

eters which describe the uncertainty in the measured BICEP/*Keck* bandpasses. These are quantified by the fractional shift in the band center. We include one parameter for each frequency plus a correlated shift applied to all three channels. For each parameter, we use a Gaussian prior with mean/ σ of 0/0.02. This corresponds to the dashed black line in Figure 6.9. We note that the band-center shifts have little effect on the likelihoods.

A further variation we consider is including dust decorrelation as part of our sampled parameter space. In Section 3.3 we discussed the origin and phenomenology of foreground decorrelation. The motivation for adding this parameter comes from [Planck Collaboration Int. L \(2017\)](#). This particular study claimed the existence of evidence for a cross-spectrum suppression effect, such as the one coming from decorrelation, which increased with ℓ , in particular for cleaner regions of the sky. More recently, [Sheehy & Slosar \(2018\)](#) and [Planck Collaboration 2018 XI \(2018\)](#) re-analyzed the *Planck* data and found no evidence for detection of dust decorrelation.

In general, in order to place robust cosmological constraints, we require a good understanding of foregrounds. Therefore, we would like to search for evidence of decorrelation directly in our observed field (as opposed to studies which analyze large swaths of the sky). To that end, we include this parameter as described in Section 3.3 and explore the effect it has on our baseline posteriors. We consider several choices of prior on the Δ_d parameter: i) Based on Table 1 of [sheehy17](#) and Table 3 of [Planck Collaboration 2018 XI \(2018\)](#) we set a

Gaussian prior with mean/ σ of 0.95/0.05 (truncated above 1), flat with ℓ . ii) A Gaussian prior with mean/ σ of 1.00/0.05, linear with ℓ . iii) A uniform prior 0 to 1, linear with ℓ . The results of this investigation are presented in Figure 6.10. All of these variants result in the r posterior shifting down and peaking at zero. It is important to note that marginalizing over truncated parameter spaces, such as one set by demanding physically meaningful ranges (in this case $r > 0$ and $\Delta_d \leq 1$), can result in systematic downward biases on r even if no decorrelation is present.




Figure 6.10: Likelihood results when allowing dust decorrelation—see Section 6.3.3 for details. *Figure from BICEP2 Collaboration et al. (2018)*

In our baseline analyses, we usually fix the lensing amplitude to its Λ CDM expectation value of ($A_L^{\text{BB}} = 1$). Here we also explore letting the lensing amplitude free. The results are presented in Figure 6.11. Setting a uniform prior, and marginalizing over all other parameters, we obtain $A_L = 1.15^{+0.16}_{-0.14}$ (blue contours). The zero-to-peak likelihood ratio is 1.3×10^{-17} , and the probability of having a lower value is 5.8×10^{-19} , which corresponds to an 8.8σ detection. This is the most significant detection of lensing using B-mode polarization to date. We observe that there is some degeneracy between r and A_L , which pushes the

r posterior down. If we instead impose a prior from Planck, $A = 0.95 \pm 0.04$ (Planck Collaboration 2015 XIII 2016), the recovered r likelihood curve is almost indistinguishable from the baseline case.




Figure 6.11: Likelihood results when allowing the lensing amplitude to be a free parameter. *Figure from BICEP2 Collaboration et al. (2018)*

6.3.3.2 DATA SELECTION VARIATIONS

We also consider a set of data selection variations. The results for these are presented in Figure 6.12. Using the BICEP2/*Keck* data only (magenta) the r posterior peaks at zero, while the A_d posterior broadens slightly. An impressive conclusion is that the constraints on r only degrade by 10-15% when going from the full dataset to just the BICEP2/*Keck* data. This is a massive statement about the sensitivity of the BICEP2/*Keck* data, and our ability to constrain foregrounds without relying on full-sky surveys. Unsurprisingly, the constraints on A_{sync} deteriorates considerably, given that there are no BICEP/*Keck* low-frequency bands. Adding in some low-frequency information in the form of WMAP (green) shifts the

r even more aggressively towards zero, while A_{sync} becomes better constrained. Switching LFI for WMAP (green to yellow) brings r back up a bit and pushes A_{sync} down. Adding high-frequency external information in the form of *Planck* HFI (green to red) pushes r up and leaves A_{sync} unchanged. BICEP/*Keck*+*Planck* (blue) has almost exactly the same r curve as the baseline but a considerably wider A_{sync} curve. We can understand the difference in behaviors for synchrotron when including LFI vs WMAP by noting that in Figure 6.4 the $\text{BK}_{95} \times \text{W}_{23}$ bandpowers are positive while the $\text{BK}_{95} \times \text{P}_{30}$ bandpowers are negative.




Figure 6.12: Likelihood results when varying the data set selection *Figure from BICEP2 Collaboration et al. (2018)*

6.3.4 LIKELIHOOD VALIDATION

Similar to the validation exercise performed in BK14, and presented in Figure 5.9, we carry out full likelihood evaluations of our baseline model for 499 simulations with lensed- Λ CDM + Gaussian dust + noise. The results are shown in Figure 6.13. We note that the constraining power on r , calculated as the standard deviation of the maximum likelihood recovered r

values, increases by 20% when going from BK14 to BK15 (this means a decrease in $\sigma(r)$ from 0.025 to 0.020).




Figure 6.13: Results of a validation test running maximum likelihood search on simulations of a lensed- Λ CDM+dust+noise model with no synchrotron ($A_{d,353} = 3.75\mu K^2$, $\beta_d = 1.6$, $\alpha_d = -0.4$, $A_{sync,23} = 0$). The baseline priors are applied on β_d , β_s , α_d , α_s and ϵ . The blue histograms are the recovered maximum likelihood values with the red lines marking their means and the black lines showing the input values. In the left panel $\sigma(r) = 0.020$. *Figure from [BICEP2 Collaboration et al. \(2018\)](#)*

To check that the framework remains unbiased when running maximum likelihood searches with decorrelation, we repeat the validation exercise described above while including the decorrelation parameter Δ_d . When we add Δ_d in our search, we allow it to take values greater than one. To do this in a symmetrical manner we use

$$\Delta'_d(\nu_1, \nu_2, \ell) = 2 - \exp[\log(2 - \Delta_d) f(\nu_0, \nu_1) g(\ell)]. \quad (6.1)$$

In this exercise we take the linear ℓ scaling. The results of this validation exercise are shown in Figure 6.14.

An additional validation test we perform is to run full COSMOMC likelihood evaluations on




Figure 6.14: Validation tests running the likelihood with the dust decorrelation parameter Δ_d included. *Upper row:* Results for the same lensed- Λ CDM+dust+noise simulations shown in Figure 6.13. *Lower row:* Results for the toy highly decorrelated dust model. The blue histograms are the recovered maximum likelihood values with the red lines marking their means and the black lines showing the input values. *Figure from BICEP2 Collaboration et al. (2018)*

the same ensemble of simulations as above, with the baseline analysis parametrization. The left panel for Figure 6.15 shows the resulting marginalized posteriors on r , while the right panel shows that the CDF of the zero-to-peak likelihood ratios closely follows the simple analytic ansatz $\frac{1}{2}(1 - f(-2 \log L_0/L_{\text{peak}}))$ where f is the χ^2 CDF (for one degree of freedom). We find that 53% of the simulations peak at zero, and 19% have a lower zero-to-peak ratio than the real data. This is powerful empirical evidence that the real data 1D likelihood on r can be taken at face value, assuming that the chosen foreground parameterization accurately reflects reality. When we repeat this exercise but include the decorrelation parameter in the analysis, we find that 72% of the r curves peak at zero.




Figure 6.15: *Left:* Likelihood curves for r when running the baseline analysis on each of the lensed- Λ CDM+dust+noise simulations. We find that 50% of them peak at zero. The real data curve is shown overplotted in heavy black. *Right:* The CDF of the zero-to-peak ratio (red) of the curves shown at right as compared to the simple analytic ansatz (solid black) $\frac{1}{2} (1 - f(-2 \log L_0/L_{\text{peak}}))$ where f is the χ^2 CDF (for one degree of freedom). About one fifth of the simulations offer more evidence for non-zero r than the real data when the true value is actually zero (dashed black). *Figure from BICEP2 Collaboration et al. (2018)*

A number of other validation tests on simulations with third-party foreground models are presented in Appendix E.4. of [BICEP2 Collaboration et al. \(2018\)](#).

6.4 DISCUSSION AND CONCLUSIONS

The previous BK14 analysis yielded the constraint $r_{0.05} < 0.090$ (95%). Adding the *Keck Array* data taken during 2015 we obtain the BK15 result $r_{0.05} < 0.072$. The distributions of maximum likelihood r values in simulations where the true value of r is zero give $\sigma(r_{0.05}) = 0.024$ and $\sigma(r_{0.05}) = 0.020$ for BK14 and BK15 respectively.

Figure 6.16 shows the BK15 noise uncertainties in the $\ell \approx 80$ bandpowers. The new *Keck* 220 GHz band has approximately the same dust sensitivity as *Planck* 353 GHz with just

two receiver-years of operation. Going forward, we expect that if we fold the eight receiver-years of data recorded in the 2016 and 2017 seasons, the noise in the 220×220 & 150×220 spectra will reduce by a factor of 5 & $\sqrt{5}$ respectively. Furthermore, in 2017 alone BICEP3 recorded twice as much data in the 95 GHz band as is included in the current result. Extra sensitivity in this band offers a better handle on synchrotron (through all the cross-spectra with *Planck* and WMAP), as well as a measurement of the CMB in the foreground minimum. The next data release will include all this data and is expected to improve substantially on the current results.

Looking beyond 2017, the BICEP/*Keck* program has planned a number of upgrades which should help to constrain PGWs further. We also expect joint work with SPT-3G towards a demonstration of CMB delensing to become very crucial, as the sample variance from the lensing signal is quickly becoming a dominant bottleneck in our ability to further reduce $\sigma(r)$. Beyond BICEP/*Keck* and SPT-3G, the next generation of ground-based CMB polarization experiment (CMB-S4) is set to take-off within the decade, and yet again improve our ability to tease out a signal of primordial origin, should one exist. To understand the science reach of these upcoming endeavours we turn to the following chapter, where we introduce a new forecasting framework.




Figure 6.16: Expectation values and noise uncertainties for the $\ell \sim 80$ BB bandpower in the BICEP2/Keck field. The solid and dashed black lines show the expected signal power of lensed- Λ CDM and $r_{0.05} = 0.05$ & 0.01 . The blue/red bands show the 1 and 2σ ranges of dust and synchrotron in the baseline analysis including the uncertainties in the amplitude and frequency spectral index parameters ($A_{\text{sync},23}, \beta_s$ and $A_{\text{d},353}, \beta_d$). The BICEP2/Keck auto-spectrum noise uncertainties are shown as large blue circles, and the noise uncertainties of the WMAP/Planck single-frequency spectra evaluated in the BICEP2/Keck field are shown in black. The blue crosses show the noise uncertainty of selected cross-spectra, and are plotted at horizontal positions such that they can be compared vertically with the dust and sync curves. *Figure from BICEP2 Collaboration et al. (2018)*

It never always gets worse.

David Horton

7

Performance-based Forecasted Constraints on Primordial Gravitational Waves

All CMB collaborations face the challenge of faithfully predicting the science reach of their future experiments. Through multiple rounds of published results with *BICEP/Keck* we have learned that the only reliable way to compute forecasts that match our achieved results is to

base them on B-mode noise spectra and covariance matrices derived from on-sky multi-year maps that have passed jackknife resampling tests.

To that end, we developed a spectral-based Fisher projection framework, targeted explicitly towards optimizing tensor-to-scalar parameter constraints in the presence of galactic foregrounds and gravitational lensing of the CMB, that directly uses information from current BICEP/*Keck* achieved performances to forecast the science reach of upcoming CMB-polarization endeavors robustly. This methodology allows for rapid iteration over experimental configurations and offers a flexible way to optimize the design of future experiments given a science goal. This framework is different from others of this kind through its direct implementation of real-life experimental inefficiencies, including, but not limited to: imperfect detector yield, non-uniform detector performance, read-out noise, observing inefficiency, time-stream filtering, beam smoothing, and non-uniform sky coverage.

In this chapter, we describe the development of this forecasting framework and present two applications for future CMB experiments: BICEP Array – the next instrument in the BICEP/*Keck* series, to be deployed in 2019 – and CMB-S4 – the ultimate ground-based CMB polarization experiment, due to start observations in 2027. We will discuss how we have used this tool to compute performance-based forecasts on r for BICEP Array and determine the optimal frequency coverage and survey design for its nominal five-year operation. We will also demonstrate how we have leveraged this work for CMB-S4 to explore the effects of sky coverage, telescope aperture, optics, and beam size, number of detectors, detector sensitivity, number of observing bands, and atmospheric properties, under various foreground assump-

tions. The results of these studies have been instrumental in establishing the current baseline CMB-S4 experiment and have contributed to the publication of the CMB-S4 Science Book [Abazajian et al. \(2016\)](#), the Concept Definition Task Force Report to the Astronomy and Astrophysics Advisory Committee, and the ongoing writing of the Astronomy and Astrophysics Decadal Survey Report.

We finalize by presenting a validation of the forecasts through suites of map-level simulations representing experimental configurations with varying degrees of realism and sky models with increasing levels of foreground complexity.

7.1 FISHER OPTIMIZATION FRAMEWORK

7.1.1 FISHER FORMALISM

At its core, this framework is a Fisherized version of the BICEP/*Keck* likelihood analysis presented in Chapter 3, and our belief in the projections is grounded in that connection to achieved performance and published results. In particular, we emphasize the importance of using map-level signal and noise simulations as a starting point. These simulations are a good description of the maps because they pass jackknife tests derived from them.

Figure 7.1 is a schematic representation of the framework, identifying the user inputs, code modules, and outputs of said modules. This section describes this framework in detail.

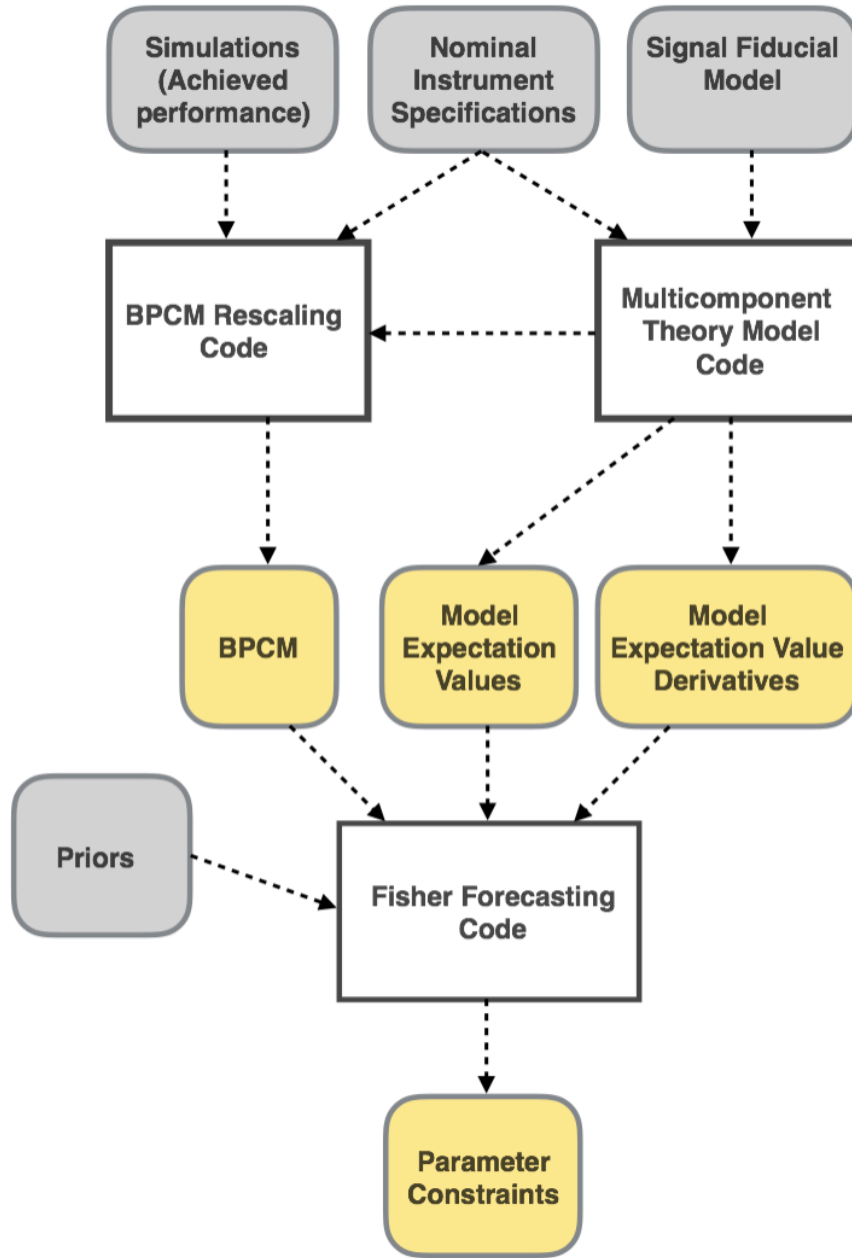


Figure 7.1: This is a schematic representation of our Fisher Machinery. Grey boxes represent user inputs, white boxes represent code modules, and yellow boxes represent code outputs.

Given a likelihood function of the form

$$L(\theta; d) \sim \frac{\exp \left[-\frac{1}{2}(d - \mu(\theta))^T \Sigma(\theta)^{-1} (d - \mu(\theta)) \right]}{\sqrt{\det(\Sigma(\theta))}}, \quad (7.1)$$

where d are the data bandpowers, θ are the theory parameters, $\mu(\theta)$ are the expectation values given the parameters, and $\Sigma(\theta)$ is the bandpower covariance matrix, which can also be a function of the parameters. We can introduce the Fisher Information Matrix

$$F_{ij} = - \left\langle \frac{\partial^2 \log(L(\theta; d))}{\partial \theta_i \partial \theta_j} \right\rangle, \quad (7.2)$$

which is the ensemble average of the log-likelihood curvature, evaluated at the position of the best fit model. This matrix measures how steeply the likelihood falls as we move away from the best fit model, and F^{-1} can be thought of as the best possible covariance matrix for the measurement errors on the parameters θ_i . It can be shown that $\sqrt{(F^{-1})_{ii}}$ is the minimum obtainable standard deviation on the desired parameters [Tegmark \(1997\)](#).

After some algebraic manipulations we can arrive at

$$F_{ij} = \frac{\partial \mu^T}{\partial \theta_i} \Sigma^{-1} \frac{\partial \mu}{\partial \theta_j} + \frac{1}{2} \text{Tr} \left(\Sigma^{-1} \frac{\partial \Sigma}{\partial \theta_i} \Sigma^{-1} \frac{\partial \Sigma}{\partial \theta_j} \right). \quad (7.3)$$

And we calculate our parameter constraints as

$$\sigma_{ii} = \sqrt{(F^{-1})_{ii}}. \quad (7.4)$$

In all the projections below we choose to fix $\Sigma(\theta) = \Sigma$. Eq. 7.3 makes it clear that the construction of the covariance matrix Σ directly impacts the final constraints, and that a misestimation of it could lead to constraints that are far too optimistic. It is with this in mind that we have decided to base our calculations on scaled achieved performances.

7.1.2 INPUTS

This section describes the inputs to the Fisher forecasting code.

Achieved Performance: the code takes information derived from signal and noise simulations of the BICEP/Keck dataset, but could be adapted to use similar information from another experiment. More specifically, we use signal-only, noise-only, and signal x noise bandpower covariance terms, as well as the ensemble-averaged signal and noise bandpowers. These inputs contain information about the actual map noise achieved from multiple receiver-years at $\{95, 150, 220\}$ GHz, including the real penalties for detector yield, distribution of detector performance, weather, and observing efficiency. These inputs also fold in the incomplete mode coverage due to sky coverage, scan strategy, beam smoothing, and filtering in data analysis. For projections, we assume that we can scale down the noise based on increased detector count and integration time and that we can apply beam-size and NET rescaling to estimate performance at other frequencies. We operate under the assumption that detector noise is uncorrelated and we don't suffer heavily from common-mode atmospheric noise.

Instrument Specifications: number of frequencies and bandpass data specifying observing

frequency response for each experiment included in the analysis, beam size, bandpower window functions specifying the response of each observed bandpower to each multipole on the sky, number of detectors per each frequency, and ideal per-detector noise-equivalent temperatures (NET's) per each frequency. The last two items can be used to perform an idealized calculation of the instrument sensitivity. We use these ideal performance numbers only for scaling purposes by comparing to the ideal sensitivity of BICEP/Keck and calculating the appropriate noise scaling factors. These factors are applied to the achieved sensitivities of BICEP/Keck to obtain performance-based sensitivities in our new bands.

Signal Fiducial Model: parameters used to calculate the multicomponent model. Our standard model has fourteen parameters, discussed in the section below.

Priors: if we have prior knowledge of a given parameter θ_i , we can easily introduce this information into the Fisher Matrix by simply adding $P_i = 1/\sigma_i^2$ to the diagonal of the Fisher Matrix, where σ_i is the width of the prior.

7.1.3 MULTICOMPONENT THEORY MODEL

This code module calculates the parametrized model of one or more signal fields, including how the different signals scale to the various observing frequencies. Our model includes four signal components: a CMB component, parametrized by r and the lensing amplitude A_L , and components of dust, synchrotron, and correlated dust and synchrotron, which has a frequency scaling that depends on the relative strength of dust and synchrotron. We assume that the

synchrotron scales as a simple power law in both frequency and l , while for the dust we assume a power law scaling in l , and a greybody spectral energy distribution (SED) for the frequency dependence. In addition, we also consider dust and synchrotron frequency decorrelation parameters which allow for dust/sync cross-spectral power suppression. A detailed description of the full parametric model is presented in Chapter 3. While this model is easily extendable, other models could be substituted here.

The 14 model parameters are:

- r , tensor-to-scalar ratio, at pivot scale $k_t = 0.05$
- A_L , lensing amplitude
- A_{dust} , dust amplitude, in μK_{CMB}^2 , at 353 GHz and $l = 80$
- β_{dust} , dust spectral index
- T_{dust} , dust greybody temperature
- α_{dust} , dust spatial spectral index
- Δ_{dust} , dust frequency decorrelation, at 217×353 GHz and $l = 80$
- EE/BB , ratio for dust
- A_{sync} , synchrotron amplitude, in μK_{CMB}^2 , at 23 GHz and $l = 80$
- β_{sync} , synchrotron spectral index
- α_{sync} , synchrotron spatial spectral index
- Δ_{sync} , synchrotron frequency decorrelation, at 23×33 GHz and $l = 80$
- EE/BB , ratio for synchrotron
- ϵ , synchrotron-dust spatial correlation

Since each signal component is independent, different signals have zero-cross power. Therefore, for a given spectrum, the code steps through the model components, combines the appropriate amplitude functions for the two experiments contributing to the spectrum and applies the bandpower window functions to get the binned expectation values. Finally, it sums over model components to get the total expectation value for that spectrum. In addition, since a Fisher forecast requires knowledge about the response of the model expectation values with respect to the model parameters, the code also outputs the derivatives of the model expectation values.

The Fisher matrix we consider is usually 10-dimensional. The parameters we are constraining are: $\{r, A_{dust}, \beta_{dust}, \alpha_{dust}, \Delta_{dust}, A_{sync}, \beta_{sync}, \alpha_{sync}, \Delta_{sync}, \epsilon\}$. We fix $T_{dust} = 19.6K$ and A_L is varied separately to account for various levels of lensing power, i.e., various levels of delensing efficiencies. The EE/BB ratios are not relevant for the calculations in this manuscript because we are focusing on constraints from B -mode measurements only. The fiducial model for the Fisher forecasting is centered at r of 0 or 0.003, with $A_{dust} = 4.25\mu K_{CMB}^2$ (best-fit value from BK14) and $A_{sync} = 3.8\mu K_{CMB}^2$ (95% upper limit from BK14). The spatial and frequency spectral indices are centered at the preferred Planck values $\beta_{dust} = 1.59$ (with Gaussian prior of width 0.11), $\beta_{sync} = -3.10$ (with a Gaussian prior of width 0.30), $\alpha_{dust} = -0.42$, $\alpha_{sync} = -0.6$, and the dust/sync correlation is centered at $\epsilon = 0$. The central dust decorrelation value is taken to be $\Delta_{dust} = 0.97$ and the synchrotron decorrelation value is assumed to be $\Delta_{sync} = 1$. Unless otherwise stated, the parameters have flat unbounded priors.

7.1.4 BANDPOWER COVARIANCE MATRIX RESCALING

Signal Scaling: the output model expectation values can also be useful in the construction of our bandpower covariance matrix (BPCM). To construct the bandpower covariance matrix components we use lensed- Λ CDM + dust + noise BICEP/Keck simulations (described in Chapter 3). However, because we have the individual signal-only, noise-only, and signal x noise terms, we can record all the BPCM components

- sig = signal-only terms $Cov(S_i \times S_j, S_k \times S_l)$
- noi = noise-only terms $Cov(N_i \times N_j, N_k \times N_l)$
- sn1 = signal×noise terms $Cov(S_i \times N_j, S_k \times N_l)$
- sn2 = signal×noise terms $Cov(S_i \times N_j, N_k \times S_l)$
- sn3 = signal×noise terms $Cov(N_i \times S_j, S_k \times N_l)$
- sn4 = signal×noise terms $Cov(N_i \times S_j, N_k \times S_l)$

And then rescale and combine them to create a bandpower covariance matrix for a new desired multicomponent model. Here, S are signal simulations, N are noise simulations, and the indices i, j, k, l run over fields in the analysis, i.e. all combinations of a map type (T, E, B) and an experiment $(BK_{95}, BK_{150}, \text{etc.})$. As described in Chapter 3, for many combinations of indices, we set certain covariance terms to be identically zero. It is worth noting that having all of these terms allows us to have different numbers of degrees of freedom per bandpower for noise than for signal, a complication that is often ignored in other analyses by setting the noise and signal degrees of freedom to be identical.

While calculating the covariances from the signal and noise simulations, we record the average signal bandpowers from the simulations. For a new signal model, we can calculate the new bandpower expectation values, and rescale the signal components in the bandpower covariance matrix by the appropriate power of the ratio of the recorded average signal bandpowers, and the newly calculated expectation values.

The ability to get a BPCM for any model means only one set of simulations is necessary, and one does not have to run simulations for any and all conceivable scenarios. As mentioned, in all the projections below we choose to fix $\Sigma(\theta) = \Sigma$ and hence we only apply the step above once, i.e., we do not rescale our BPCM at every step of the way.

Noise Scaling: In addition to scaling from one signal model to the other, recording all of the covariance terms allows us to rescale the noise parts as well. Given a dataset for which we have simulations, the noise scaling can go one of two ways. The first allows one to take a frequency present in the dataset and scale down the noise in the BPCM by the desired amount. The second one allows one to add a frequency by taking an existing one, scaling down the noise by the desired amount, and then expanding the BPCM and filling it in with the appropriate variance and covariance terms between this new band and all the existing ones. These two tools allow us to set up a new data structure to explore any combination of frequency bands, with any sensitivity in each band.

We want to base our noise scaling factors on achieved sensitivities rather than ideal performances. To that end, we use the achieved survey weights from BICEP/Keck at $\{95, 150, 220\}$ GHz to obtain projected weights for any channel of a future CMB experiment. The aver-

age achieved per detector-year survey weights are calculated directly from final multi-year BICEP/Keck maps, and therefore include all the non-idealities described in the sections above. These are $w_{BK,achieved}^{per-det-yr} = \{83.3, 58.6, 3.9\} \mu K^{-2}$ (BICEP2 Collaboration et al. (2018)).

Once we have these numbers, we can write the achieved per-detector-year survey weight as follows

$$w_{BK,achieved}^{per-det-yr} = \frac{t_{obs}}{\alpha_{ideal/BK,achieved} NET_{BK,ideal}^2}, \quad (7.5)$$

where $\frac{t_{obs}}{\alpha_{ideal/BK,achieved}}$ is the reality factor that encapsulates the less-than-ideal observing time, receiver performance, cuts, etc. It is the factor that takes us from the ideal scenario to reality. We never have to actively calculate this factor because we operate with ratios of survey weights. To obtain the achieved survey weight at a new frequency, one rescales the achieved survey weight as follows

$$w_{New,achieved} = w_{BK,achieved}^{per-det-yr} \mathcal{N}_{New}^{det-yr} \frac{NET_{BK,ideal}^2}{NET_{New,ideal}^2}, \quad (7.6)$$

where $\mathcal{N}_{New}^{det-yr}$ is the number of detector-years assumed for the new experiment at any particular frequency. The implicit assumption that is made here is that the reality factor for this new frequency is exactly the same as the one from which we are scaling. To not abuse this assumption, the survey weight scaling is always done from the closest frequency that we have an achieved survey weight for, as that is the performance that should guide us.

Therefore, to obtain N_l 's for one channel by scaling the achieved N_l 's of another channel

we have first to scale them by the ratio of the survey weight and then scale by the ratio of the beams squared. Achieved performance inputs fold in the actual BICEP/*Keck* B_l 's, but we rescale based on Gaussian approximations of the beams, which closely follow the BICEP/*Keck* beams. Under the assumption of Gaussian beams, we can write the full relation as

$$N_{l,New} = N_{l,BK} \frac{w_{BK,achieved}}{w_{New,achieved}} \frac{B_{l,New}^2}{B_{l,BK}^2}, \quad (7.7)$$

where $B_{l,\nu}^2 = \exp \frac{-l(l+1)\theta_\nu^2}{8 \log(2)}$, and θ_ν is the FWHM (in radians) of the Gaussian beam. With the N_l 's scalings on hand, we can perform the above mentioned BPCM operations to arrive at a scaled BPCM for a new experiment, that encompasses all the intricacies of reality that BICEP/*Keck* does.

7.2 BICEP/*Keck* FORECASTS

In this section, we present an application of the forecasting framework to the BICEP2/*Keck* series of experiments. As described in Chapter 2, the BICEP2/*Keck* program has seen a number of iterations of instruments. We started with BICEP2, which introduced a new generation of compact refracting polarimeter. BICEP2 was upgraded to *Keck*, which had five copies of BICEP2, each fully optimized for a single frequency band. The third generation receiver – BICEP3 – scaled up the individual *Keck* receiver to have throughput greater than all of *Keck Array*. Finally, BICEP Array being built at the time of this writing, is set to deploy four BICEP3-sized receivers, with each telescope optimized for a different frequency. This progres-

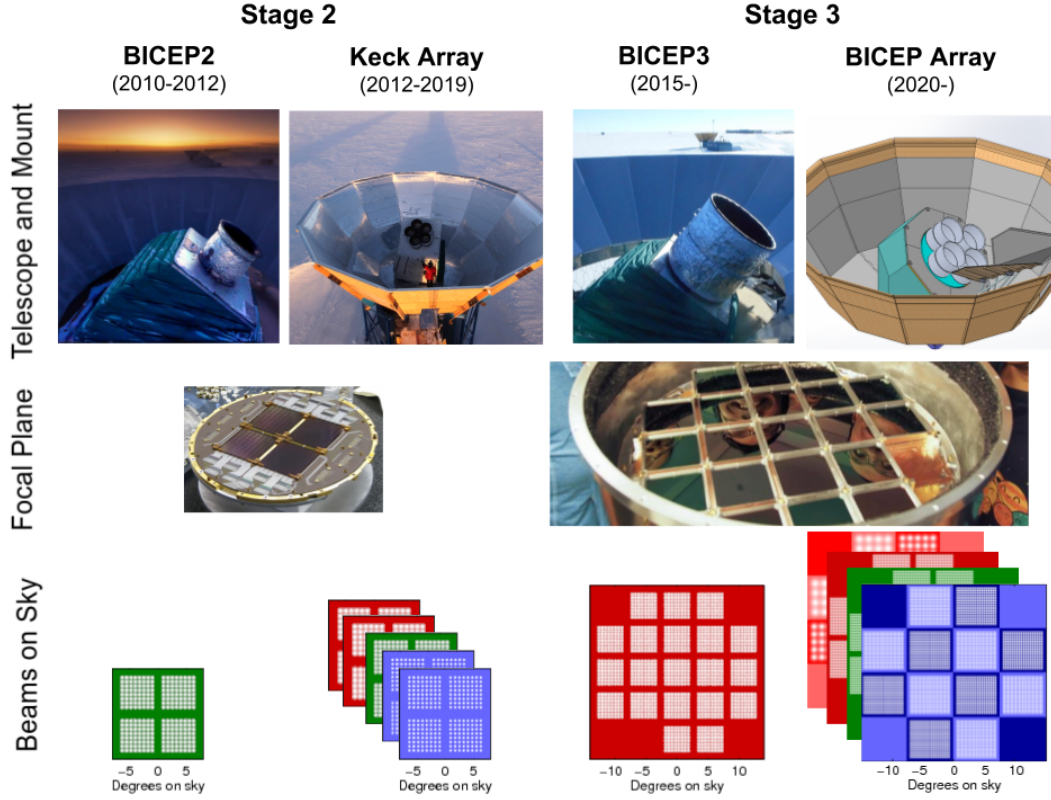


Figure 7.2: The progression of the BICEP/Keck experimental program leading to the BICEP Array. In the bottom row the beam patterns of the focal planes on the sky are shown on a common scale. *Figure from the BICEP/Keck Collaboration.*

sion is presented in Figure 7.2. One of the key features of this progression is multi-frequency observing in order to disentangle a possible PGW signal, which is frequency independent, from signals of galactic origin, which vary with frequency. To fully understand the science reach of these instruments, as well as determine the optimal deployment strategy, we turned to the Fisher framework described above.

7.2.1 FISHER CONSTRAINTS VS PUBLISHED BICEP/*Keck* CONSTRAINTS

The first test we perform is a comparison of Fisher derived constraints on PGWs, parametrized by r , with baseline published BKP, BK14, and BK15 constraints. Section 7.1.2 lists the required inputs to the Fisher framework. In this particular case, because we have available simulations that directly correspond to these specific data-sets, we do not need to rescale our achieved performance or add new bands. Therefore, the bandpower covariance formulation simply follows without requiring any other instrument specifications, as described in Section 3.5 and 7.1.4. The data selection, dimensionality of the parameter space, and prior choice are matched directly to the baseline analyses described in *BICEP2/Keck and Planck Collaborations (2015)*, *Keck Array and BICEP2 Collaborations VI (2016)*, and *BICEP2 Collaboration et al. (2018)*. The fiducial signal models used to calculate model bandpower expectation values, and the derivatives of the expectation values with respect to each parameter, are picked to exactly match the maximum likelihood models derived in the publications cited above. The resulting Fisher constraints are transformed to Fisher ellipses (emulating 68% confidence intervals) and Gaussian 1D probability density functions, as described below, which are then compared to the marginalized 1D and 2D baseline likelihood evaluations.

To obtain Fisher ellipses, for each pair of parameters $\{\theta_i, \theta_j\}$, we take the following sub-

matrix of the inverse of the Fisher information matrix

$$F_{[i,j]}^{-1} = \begin{pmatrix} \sigma_i^2 & \sigma_{ij} \\ \sigma_{ij} & \sigma_j^2 \end{pmatrix},$$

and calculate the ellipse parameters $\{a, b, \alpha\}$ – the semi-major, semi-minor, and projection angle – as follows

$$\begin{aligned} a^2 &= \frac{\sigma_i^2 + \sigma_j^2}{2} + \sqrt{\frac{(\sigma_i^2 - \sigma_j^2)^2}{4} + \sigma_{ij}^2}, \\ b^2 &= \frac{\sigma_i^2 + \sigma_j^2}{2} - \sqrt{\frac{(\sigma_i^2 - \sigma_j^2)^2}{4} + \sigma_{ij}^2}, \\ \tan 2\alpha &= \frac{2\sigma_{ij}}{\sigma_i^2 - \sigma_j^2}. \end{aligned}$$

To form a 1D Gaussian probability density function we calculate

$$P(\theta_i) = \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{(\theta - \bar{\theta})^2}{2\sigma_i^2}\right),$$

where $\bar{\theta}$ is the fiducial model value of the parameter θ .

The comparisons of the Fisher constraints with published *BICEP/Keck* constraints are presented in Figures 7.3, 7.4, and 7.5. In general, we do not expect perfect agreement in the tails because the Fisher ellipses will always be Gaussian by construction, whilst the real contours are manifestly non-Gaussian probability density functions due to the inherent χ^2 distributions of the *BICEP/Keck* bandpowers. On a similar note, the real analysis also uses the

Hammimeche-Lewis likelihood approximation, which predicts non-Gaussian tails. With that in mind, the focus should be on the curvature of the likelihood at the peak, for which we see good agreement between the two results.

7.2.2 *Keck Array* HIGH-FREQUENCY CHANNEL SELECTION

We performed one of the first applications of this forecasting machinery for BICEP/*Keck* in 2015. It was an exercise to determine the optimal high-frequency channel distribution in order to achieve the highest sensitivity on dust. At the time the *Keck Array* was operating with two receivers at 95 GHz, one at 150 GHz and two at 220 GHz. The 2015 season was the first year of observation with the high-frequency camera, and it was apparent that one of the bottlenecks to reducing the variance on r was the lack of high S/N dust maps. The *Planck* 353 GHz map was quickly becoming an insufficient handle on dust. The strategy for 2016 was set to be four receivers at 220 GHz and one at 150 GHz. The decision to swap out the 95 GHz receiver was taken due to the promising performance of BICEP3 (all at 95 GHz). The three options considered for the 2017 season of *Keck Array* were: all five receivers at 220 GHz, three receivers at 220 GHz and two with broadband 250 GHz cameras, and three receivers at 220 GHz and two at 270 GHz. The receiver parameters were calculated using an internal version of `NETlib.py`¹ and are summarized in Table 7.1. Using these parameters, and pinning the Fisher signal model to the BK14 fiducial model, with the exception of setting $r = 0$ for an upper limit calculation, we obtain σ_r results for the three options (presented in Table 7.2).

¹cmb-s4.org/wiki/index.php/New_NET_Calculator_and_Validation

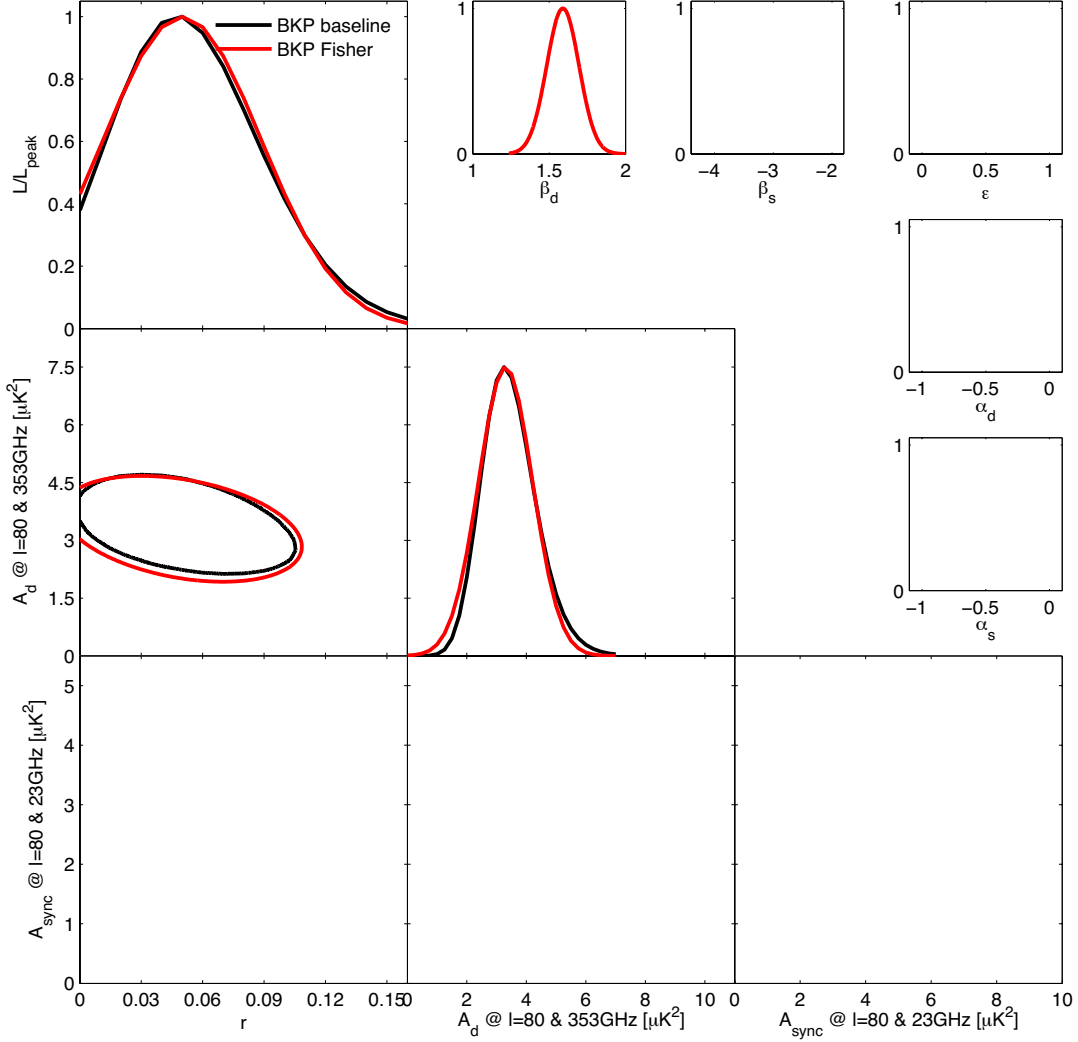


Figure 7.3: Projected constraints compared to published BKP constraints ([BICEP2/Keck and Planck Collaborations \(2015\)](#)). The data selection for the forecasts matches that of the real analysis: *BK13*₁₅₀, P217 and P353. The space constrained by the Fisher calculation matches the BKP parameter space: $r, A_{d,353}, \beta_d$. The fiducial model itself is picked to be the maximum likelihood model derived from the data: $r = 0.05$, $A_d = 3.3 \mu\text{K}^2$, $\beta_d = 1.59$, $\alpha_d = -0.42$, $T_d = 19.6\text{K}$, with a prior on $\beta_d = 1.59 \pm 0.11$ from the full sky *Planck* analysis ([Planck Collaboration 2015 X \(2016\)](#))

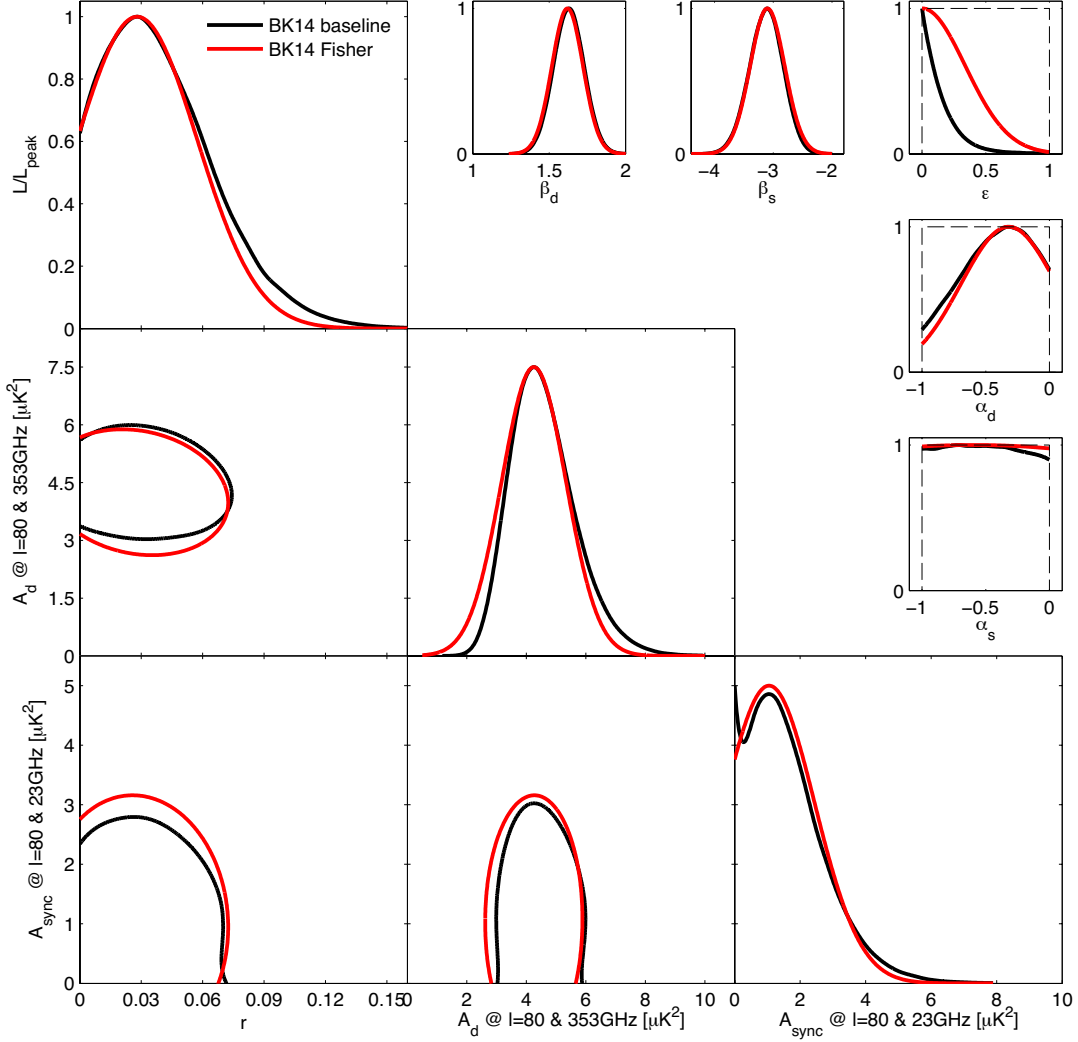


Figure 7.4: Projected constraints compared to published BK14 constraints (*Keck Array* and *BICEP2 Collaborations VI* (2016)). The data selection for the forecasts matches that of the real analysis: $BK14_{95}$, $BK14_{150}$, all seven polarized bands from Planck 30 to 353GHz, and two bands from WMAP at 23 and 33 GHz. The space constrained by the Fisher calculation matches the BKP parameter space: $r, A_{d,353}, \beta_d, \alpha_d, A_{\text{sync},23}, \beta_s, \alpha_s, \epsilon$. The fiducial model itself is picked to be the maximum likelihood model derived from the data: $r = 0.028$, $A_{d,353} = 4.25 \mu\text{K}^2$, $\beta_d = 1.62$, $\alpha_d = -0.42$, $T_d = 19.6\text{K}$, $A_{\text{sync},23} = 1.05 \mu\text{K}^2$, $\beta_s = -3.1$, $\alpha_s = -0.6$, $\epsilon = 0$ with a prior on $\beta_d = 1.59 \pm 0.11$ and $\beta_s = -3.1 \pm 0.3$ from the full sky *Planck* analysis (*Planck Collaboration 2015 X* (2016))

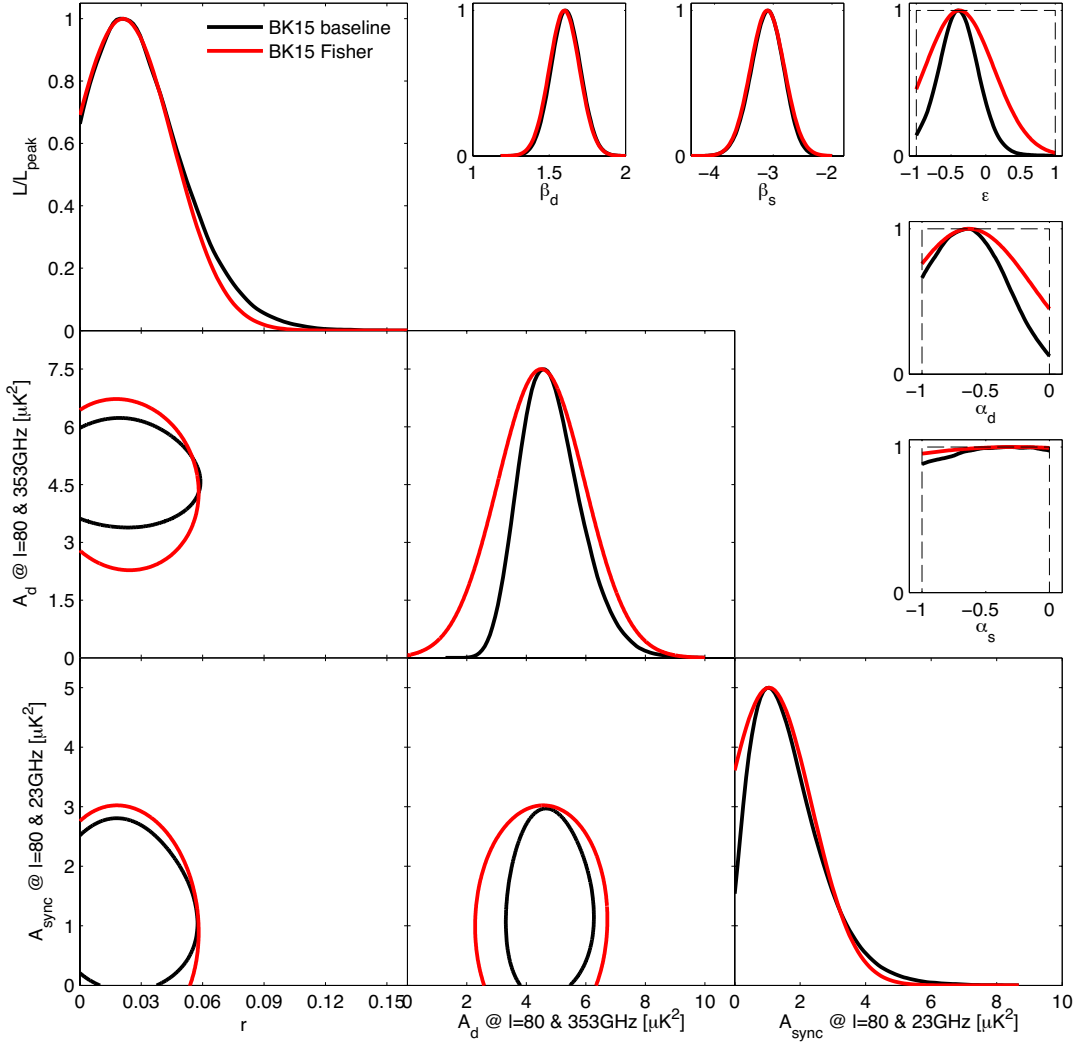


Figure 7.5: Projected constraints compared to published BK14 constraints (BICEP2 Collaboration et al. (2018)). The data selection for the forecasts matches that of the real analysis: $BK15_{95}$, $BK15_{150}$, $BK15_{220}$, all seven polarized bands from Planck 30 to 353GHz, and two bands from WMAP at 23 and 33 GHz. The space constrained by the Fisher calculation matches the BKP parameter space: $r, A_{d,353}, \beta_d, \alpha_d, A_{sync,23}, \beta_s, \alpha_s, \epsilon$. The fiducial model itself is picked to be the maximum likelihood model derived from the data: $r = 0.021$, $A_{d,353} = 4.5 \mu K^2$, $\beta_d = 1.62$, $\alpha_d = -0.63$, $T_d = 19.6 K$, $A_{sync,23} = 1.05 \mu K^2$, $\beta_s = -3.1$, $\alpha_s = -0.27$, $\epsilon = -0.38$ with a prior on $\beta_d = 1.59 \pm 0.11$ and $\beta_s = -3.1 \pm 0.3$ from the full sky *Planck* analysis (Planck Collaboration 2015 X (2016))

Receiver Observing Band (GHz)	Fractional Bandwidth $\Delta\nu/\nu$	Nominal Number of Detectors	Nominal Single Detector NET ($\mu\text{K}_{\text{CMB}}\sqrt{\text{s}}$)	Beam FWHM (arcmin)	Survey Weight Per Year (μK_{cmb}) $^{-2}$ yr $^{-1}$
<i>Keck Array</i>					
220	0.17	512	837	21	2,000
250	0.32	512	756	18	2,500
270	0.17	512	1310	17	800

Table 7.1: *Keck Array* high-frequency receiver parameters as used in sensitivity projections. Boldface numbers are actual/achieved quantities for existing receivers. The non-bolded values in the survey weight column are scaled from the achieved survey weights using only the ratio of the number of detectors, plus, if necessary to change frequency, the ratio of nominal NET values squared.

<i>Keck Array</i> Option	5×220 GHz	3×220 GHz 2×250 GHz	3×220 GHz 2×270 GHz
$\sigma_r(A_L = 1.0) \times 10^3$	8.8	8.4	8.5
$\sigma_r(A_L = 0.5) \times 10^3$	6.8	6.4	6.5
$\sigma_r(A_L = 0.2) \times 10^3$	5.6	5.2	5.3

Table 7.2: *Keck Array* projections under three high-frequency receiver options, for three levels of delensing: $A_L = 1.0$ (no delensing), $A_L = 0.5$ (conservative delensing), $A_L = 0.2$ (optimistic delensing).

Given these results, it was decided that the way forward was the option with three receivers at 220 GHz and two at 270 GHz. While it was certainly the case that the 250 broadband cameras offered a slightly lower σ_r , the technical challenges associated with building a high-frequency broadband camera, as well as the larger level arm offered by the 270 GHz channels, swayed the decision in favor of the 270 GHz solution.

Since then, we have gathered on-sky data from the 220 and 270 GHz *Keck* receivers, and while we have confirmed the projected performance of the 220 channels, the 270 chan-

nels have mostly underperformed by a factor of ~ 2 in survey weight. With that in mind, the decision to deploy two receivers at 270 GHz was reversed, and the *Keck Array* has since been populated with four receivers at 220 GHz and one at 270 GHz. This is reflected in the schedule and sensitivity projections for the entire BICEP/*Keck* program presented in Figure 7.9 (discussed in the next section).

7.2.3 BICEP ARRAY DEPLOYMENT STRATEGY AND FORECASTS

In addition to the *Keck Array*, the BICEP/*Keck* team has been operating BICEP3 – a single telescope at 95 GHz with throughput equal to the entire *Keck Array* (see Figure 7.6), and is soon set to deploy BICEP Array – a set of four BICEP3-like receivers, all at different frequencies. In preparation for this deployment, it was important to understand what optimal observing strategy minimized σ_r . To that end, the team iterated over a set of focal plane designs and arrived at the receiver parameters presented in Table 7.3. The table contains achieved quantities from existing receivers, and quantities scaled from these achieved performances, as discussed in Section 7.1.4. In addition, we also present a visual representation of the BICEP Array bandpasses in comparison with the atmospheric transmission at the South Pole in Figure 7.7. It is worth noting, that while we do not present this in the current manuscript, the 30/40 GHz checkerboard choice was itself a result of a low-frequency forecasting study using this Fisher framework, presented in detail in [Buza \(2017\)](#). The study in question looked at four options: 30 GHz only, 40 GHz only, 30/40 GHz checkerboard, and 30/40 dichroic detectors. It tested these choices under a number of assumptions about the

behaviour of low-frequency foregrounds, in particular, synchrotron radiation, and determined the 30/40 checkerboard solution to yield the lowest synchrotron residuals and therefore the lowest σ_r .




Figure 7.6: A comparison of 95 GHz Q maps between four receiver-years of *Keck* (left) and a single receiver-year (2017) of BICEP3 (right). The increased angular resolution and field-of-view, and the increase in sensitivity can be seen. The observed pattern is due to the dominant E -mode signal, and is detected with high S/N . The colorscale range is $\pm 4 \mu\text{K}$. *Figure from the BICEP/Keck Collaboration.*

We use the receiver parameters presented in Table 7.3 to perform a Fisher optimization on σ_r , over the receiver parameter space, under the assumption that we begin deployment with one receiver, then add two more receivers the following year, and add the fourth and final receiver in the third year. We pick the fiducial model to match the BK14 maximum likelihood model and restrict ourselves to the BICEP2/*Keck* patch for the purposes of this calculation. The optimal deployment strategy offered by this calculation required the following deployment: 30/40 GHz first, followed by the two receivers at 95 GHz and 220/270 GHz and finishing with the 150 GHz receiver. This result makes intuitive sense – while BICEP/*Keck* will have achieved map sensitivities of (2.0, 3.0, 8.0) $\mu\text{K-arcmin}$ at (95, 150, 220) GHz by the

Receiver Observing Band (GHz)	Nominal Number of Detectors	Nominal Single Detector NET ($\mu\text{K}_{\text{CMB}}\sqrt{\text{s}}$)	Beam FWHM (arcmin)	Survey Weight Per Year (μK_{cmb}) ⁻² yr ⁻¹
<i>Keck Array</i>				
95	288	288	43	24,000
150	512	313	30	30,000
220	512	837	21	2,000
270	512	1310	17	800
BICEP3				
95	2560	288	24	213,000
BICEP Array				
< 30	192	260	76	19,500
< 40	300	318	57	20,500
95	4056	288	24	337,400
150	8664	313	15	509,000
< 220	8112	837	11	32,000
< 270	13068	1310	9	21,000

Table 7.3: Receiver parameters as used in sensitivity projections. Boldface numbers are actual/achieved quantities for existing receivers. The non-bolded values in the survey weight column are scaled from the achieved survey weights using only the ratio of the number of detectors, plus, if necessary to change frequency, the ratio of nominal NET values squared.

year 2020 (see Figure 7.9), when BICEP Array is set to deploy its first receiver, the program still lacks high S/N low-frequency measurements, making synchrotron residuals a significant contributor to the overall error on r . Therefore, the optimization prefers the 30/40 GHz channel to be the first. Next, it requires more depth at the frequency where CMB is brightest (95 GHz) and where dust sensitivity is highest (220/270 GHz), leaving the 150 GHz channel last.

While the theoretical calculation can offer a definitive formal solution, experimental constraints do not always allow for the implementation of this solution. For instance, the 220/270 focal plane is set to have over 20,000 detectors, posing a significant challenge which




Figure 7.7: *Left:* Comparison of atmospheric transmission and the bandpasses of BICEP/Keck and BICEP Array. Median atmospheric transmission during the observing season is shown in black, bracketed by the 10th and 90th percentiles. *Right:* Minimally processed timestream pair-sum and pair-diff noise spectra. For 95/150/220 GHz the median performance is shown for a typical mid-winter 2015 scanset; the 270 GHz line shows a successful subset of our prototype 270 GHz detectors for a typical mid-winter 2017 scanset. The stable Antarctic atmosphere enables observations at all of these frequencies that are low-noise across the indicated “science band” from 0.1–1 Hz, corresponding to $25 \lesssim \ell \lesssim 250$. *Figure from the BICEP/Keck Collaboration.*

will require the development of novel manufacturing methods over a lengthy period of time.

Additionally, the formal solution relies on a particular foreground parametrization, which might not be the actual foreground content of the true sky. It is always prudent to not blindly trust the theoretical calculation and build-in experience-based safety factors. With that in mind, it was decided that the 30/40 GHz will stay as the first receiver. The 220/270 GHz receiver will be built last. Given that BICEP3 will have acquired significant amounts of data at 95 GHz, to balance the survey weight distribution the 150 GHz receiver will be deployed at the same time as the 30/40 GHz one, and finally the 95 GHz receiver will be deployed at the same time as the 220/270 GHz receiver. This is the current ideal scenario deployment of

SPT-3G (End of Year)	2019	2020	2021	2022	2023	2024	2025	2026	2027
Effective A_L	0.622	0.496	0.426	0.380	0.347	0.315	0.295	0.280	0.270

Table 7.4: SPT-3G delensing projections derived with an internal linear combination (ILC) method, using SPT-3G90/150/220GHz noise derived from on-sky performance, and an *EB* iterative estimator (described in [Smith et al. \(2012\)](#))

BICEP Array and is presented in Figure 7.9 as the baseline deployment strategy. The nominal plan of operation sees BICEP Array taking data from the start of 2020 to the end of 2023, but it is likely that the program will continue until 2027 when CMB-S4 is set to start making observations. Therefore, Figure 7.9 extends to 2027. Figure 7.8 shows the noise uncertainties in the BICEP/*Keck* field under the nominal operation of BICEP Array (in addition to *Keck* and BICEP3), which shows that we are soon set to measure all three foregrounds (dust, synchrotron, CMB lensing) with high S/N in the maps.

As discussed in Section 6.4, as we continue to acquire data, the noise of our maps decreases, and the dominant residual foreground becomes the sample variance from the lensing signal. This can be removed with high S/N arc-minute-scale measurements of the CMB temperature and polarization, such as the ones from SPT-3G. Therefore, in addition to the BICEP/*Keck* degree-scale measurements, we fold in delensing predictions obtained from the SPT-3G team. These are effective A_L levels (summarized in Table 7.4) derived with an internal linear combination (ILC) method, using SPT-3G 90/150/220GHz noise derived from on-sky performance, and an *EB* iterative estimator (described in [Smith et al. \(2012\)](#)). The results are presented in Figure 7.9.




Figure 7.8: Expectation values and noise uncertainties for the B -mode power spectrum at degree angular scales in the BICEP/Keck field under the nominal operation of BICEP Array. The existing constraints on the SEDs of dust (blue band) and synchrotron (red band) are presented. We also show the levels of the lensing B -mode and two potential PGW signals (black lines). Since CMB units are used, these last are flat with frequency. These can be compared to the noise uncertainties of WMAP/Planck and BICEP/Keck shown as the points for the various observing frequencies. The blue points show the BK15 data set and the red the projected noise levels. *Figure from the BICEP/Keck Collaboration.*




Figure 7.9: Projected sensitivity of the ongoing and planned BICEP/Keck observational program. *Top panel:* Schematic representation of the receiver throughput (Stage 2 vs. 3) at various frequencies by observing season. Colors are labeled in the next panel. *Middle panel:* Map depth at each frequency as a function of time. *Bottom panel:* Sensitivity to r both with and without this proposal, assuming marginalization over the seven parameter foreground model described in the BK14 paper. The raw (no foregrounds) sensitivity is also shown. *Note:* These projections involve direct scalings from published end-to-end analyses and hence include all real-world inefficiencies—the black X’s mark the sensitivities achieved in the BKP, BK14 and BK15 papers [BICEP2/Keck and Planck Collaborations \(2015\)](#), [Keck Array and BICEP2 Collaborations VI \(2016\)](#), [BICEP2 Collaboration et al. \(2018\)](#)

We notice that by the end of 2023, BICEP Array and SPT-3G are set to achieve $\sigma_r \approx 2.0 \times 10^{-3}$ in a non-detection scenario. If we extend the timeline to 2027, this result improves by 10 – 15%. This appears to indicate that there is a saturation of our S/N measurements on the modes available in the BICEP2/*Keck* patch and that a larger sky area is required. While larger sky area increases the noise per mode, it also increases the number of modes available. This is a trade-off that can be formally solved for a particular level of effort. The rule of thumb is that for an $r = 0$ scenario, it is more beneficial to make deep measurements on a smaller patch to obtain a precise handle on the foregrounds such that these can be removed in the process of teasing out a possible small PGW signal. In practice, even with BICEP3 we are already mapping a larger area of sky, and are likely set to map an effective $\sim 3 - 4\%$ of the southern sky with *Keck* and BICEP3; this will also match the SPT-3G survey, which is performing measurements over 1500 square degrees. It is left as a future exercise to design a forecast that combines experiments with different S/N and different on-sky coverage.

7.2.4 INFLATIONARY CONSTRAINTS FROM THE BICEP/*Keck* PROGRAM

As discussed in Chapter 1, measurements on r have direct implications on inflationary models. In this section, we use the forecasts on r presented in the previous section, as well as forecasts on n_s coming from *Planck* and SPT-3G to place constraints in the $r - n_s$ plane, where inflationary predictions are most visually apparent. We show forecasts for a non-detection case $r = 0$ as well as a non-zero PGW signal case $r = 0.015$. The results are presented in Figure 7.10. Below we discuss the implications of these findings.

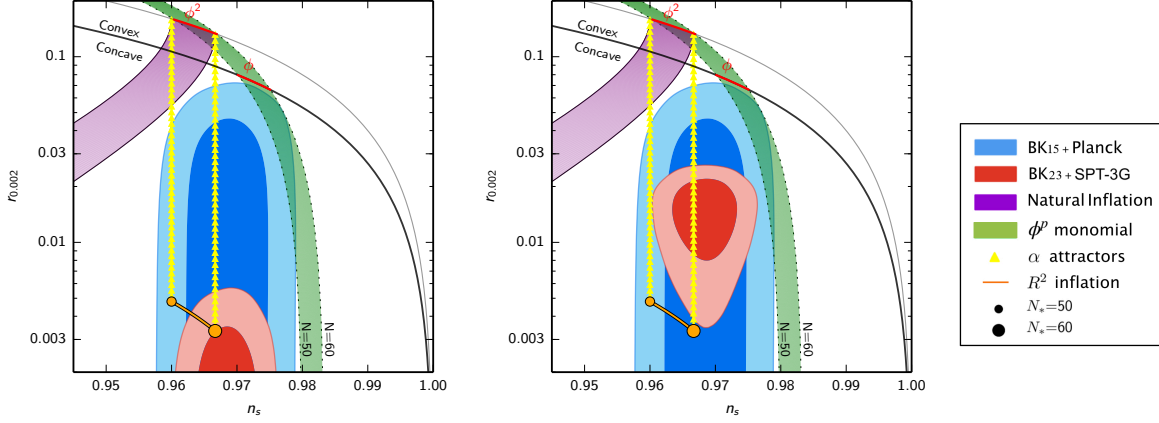


Figure 7.10: Constraints on the inflationary parameters r and n_s from our most recent published analysis (BK15) are shown in blue, as compared to projected constraints with this proposal in red (BK23+SPT-3G), and a range of theoretical models. *Left:* With no detection we would be able to constrain a large class of models. In particular, natural inflation and all of the single-field ϕ^p monomial models would be ruled out or strongly disfavored. This calculation includes the expected modest improvements on $\sigma(n_s)$ from SPT-3G and fixes n_s at the current *Planck* value. *Right:* The scenario where $r = 0.015$. With BK23+SPT-3G we would be able to make a clear detection if r is at this level or above. This would put candidate inflation models in sharp focus and enable detailed characterizations of tensor modes.

First, a detection of primordial B -modes would rule out alternative, non-inflationary, early universe theories which do not predict detectable levels of tensor modes [Khoury et al. \(2001\)](#), [Boyle et al. \(2004\)](#). A signal at the level of $r = 0.01$ or higher would mean that inflation happened at the GUT scale, and the inflaton field range is larger than the Planck mass ($\Delta\phi > M_{pl}$). Therefore, a detection at this level would provide evidence that the theory of quantum gravity must accommodate a super-Planckian field range for the inflaton. This would open up an observational window into quantum gravity and new high-energy theories.

The absence of a signal would also have major implications. In [BICEP2/Keck and Planck](#)

[Collaborations \(2015\)](#) we found that a once compelling inflation model $m^2\phi^2$ is now strongly disfavored. This particular model was alluring because it has only one free parameter m , which, once determined by the level of the scalar perturbations, is enough to predict all the other early universe properties. Ruling out this simplistic model suggests that the inflaton field was doing something more complicated as it rolled towards the minimum. This is a great example of a highly impactful non-detection.

The significance of an upper limit on tensors at $r < 0.005$ (95%) depends on how well motivated the models that would be ruled out by the measurements are. One such model is *natural inflation* [Freese et al. \(1990\)](#), [Adams et al. \(1993\)](#). It produces a flat potential by invoking a pseudo-Nambu-Goldstone boson, commonly occurring elsewhere in particle physics.

Another prominent class of models is chaotic inflation with monomial potentials ϕ^p . All single-field models with $p \geq 2$ are now strongly disfavored by our most recent published analysis [BICEP2 Collaboration et al. \(2018\)](#). Models with $p \leq 1$ have super-Planckian field ranges and can be realized naturally by axions in the context of string theory [Silverstein & Westphal \(2008\)](#), [McAllister et al. \(2010\)](#). In particular, the models with $p = 1, \frac{2}{3}$ are the specific examples of the axion monodromy models. The results in Figure 7.10 imply that both these classes of inflation models would be ruled out.

There also exists a model-insensitive way to appreciate the significance of constraining single-field ϕ^p , summarized in [Abazajian et al. \(2016\)](#). First, we note that the observed departure from scale-invariance $n_s - 1$ is on the order $\sim 1/N_*$, where N_* denotes the number of e -folds between the horizon crossing of the CMB pivot scale and the end of inflation. If the

slow-roll parameters were to scale with $1/N$ during inflation the observed level of deviation would be naturally explained. This assumption leads to only two classes of single-field, slow-roll models, with ϕ^p being one of them.

This means that the upper limit on r placed by BICEP Array and SPT-3G would rule out one of the only two classes of models that naturally explain the observed value of the scalar spectral index. This would leave the remaining class represented by Higgs inflation, the Starobinsky model, and the α attractors with small α 's [Ellis et al. \(2013\)](#), [Kallosh & Linde \(2013\)](#).

7.3 CMB-S4 FORECASTS

In this section, we present the application of our forecasting framework to the development of the CMB-S4 straw-person concept. CMB-S4, due to start observations in 2027, is meant to be the ultimate ground-based CMB polarization experiment. It is set to cross critical thresholds in constraining the B-mode polarization signature of primordial gravitational waves, the number and masses of the neutrinos, possible new light relics, the nature of dark energy, and General Relativity on large scales. To achieve these goals it requires a significant increase in sensitivity, from $\mathcal{O}(2 - 4 \times 10^4)$ detectors in Stage 3 experiments such as BICEP Array, to $\mathcal{O}(5 \times 10^5)$ detectors (see Figure 7.11). Given that Stage 3 detectors are background-limited, and focal planes maximally packed, CMB-S4 will require multiple telescopes, each with a maximally outfitted focal plane of pixels utilizing superconducting, background lim-

ited, CMB detectors. Based on current achieved performances of various types of CMB telescopes, to pursue a high-precision measurement of degree scale B-mode polarization from the ground, one requires small aperture telescopes such as the ones developed in the BICEP/Keck program. It is therefore the default plan for CMB-S4 to include dedicated small aperture telescopes, and to target the recombination bump, with E-mode and B-mode polarization down to $l \sim 30$ in order to measure a possible PGW signal.




Figure 7.11: Schematic timeline showing the expected increase in sensitivity (μK^2) and the corresponding improvement for a few of the key cosmological parameters for Stage-3, along with the threshold-crossing aspirational goals targeted for CMB-S4. *Figure from Abazajian et al. (2016)*

The CMB-S4 concept, described below, evolved from flowing down science requirements

to measurement and instrument requirements, based on our understanding of the impact of astrophysical foregrounds, instrumental systematics, delensing non-idealities, and analysis methodology. Determining all these specifications requires, in principle, full time ordered data simulations. In the absence of these, we rely on semi-analytic tools for the setting of our measurement requirements, as such tools generally offer sufficient speed to allow for optimization. To ensure realism, our general CMB-S4 forecast/simulation approach has been to:

1. Develop a (semi-) analytic spectral forecast that makes use of noise performance that is informed by scaling from actual analyses of real experiments from time-streams to power spectra.
2. Use this forecasting tool to optimize the allocation of detector effort across frequencies, determining certain baseline “checkpoints” in survey definition space.
3. Validate these checkpoint configurations with standardized, version-numbered map-based data challenges. If independent analyses show recovery of science parameters from these challenge maps that does not match analytic forecasts (either in terms of variance or bias), we revise the forecasts accordingly.
4. Iterate between 1 and 3 above, injecting increasing realism in the form of (a) sky model complexity informed by the latest data and modeling efforts, and (b) systematics whose form, parameterization, and likely amplitude is likewise guided by real-world experience.

This loop is represented schematically in Figure 7.12

7.3.1 EXPERIMENT SPECIFICATION

We start our CMB-S4 optimization by assuming eight low-resolution channels at $\{30, 40, 85, 95, 145, 155, 215, 270\}$ GHz and one high-resolution channel at 20 GHz, in four atmospheric windows, see Figure 7.13. The procedure used to come up with the split in each window was to separate the overlapping bands as far as possible while still keeping the calculated per-detector NET within 10-15% of the NET for a detector that spans the full window.

The ideal NET's per detector were calculated with `NETlib.py`² and are assumed to be $\{214, 177, 224, 270, 238, 309, 331, 747, 1281\} \mu K_{CMB} \sqrt{s}$. These NET's are calculated for a 100mK thermal bath, as opposed to 250mK for the Science Book, and are therefore lower. We want to emphasize that these NET numbers are only used to determine the scalings between different channels, and not to calculate sensitivities.

For the low-resolution instruments we pick 0.52m apertures. For the 20 GHz channel, a 0.52m aperture would result in a very broad beam which would dominate the noise at the relevant scales; to circumvent this we place this low frequency channel on a large aperture instrument, and while the scaling of the noise is still done from achieved performance, we choose a more appropriate $l_{knee} = 200$ and $\theta_\nu = 11'$ FWHM, keeping the slope γ the same as for the small-aperture noise (see Eq.7.7). In addition to the low-resolution effort, we assume a separate high-resolution instrument for delensing purposes, described in Section 7.3.4.4. We

²cmb-s4.org/wiki/index.php/New_NET_Calculator_and_Validation

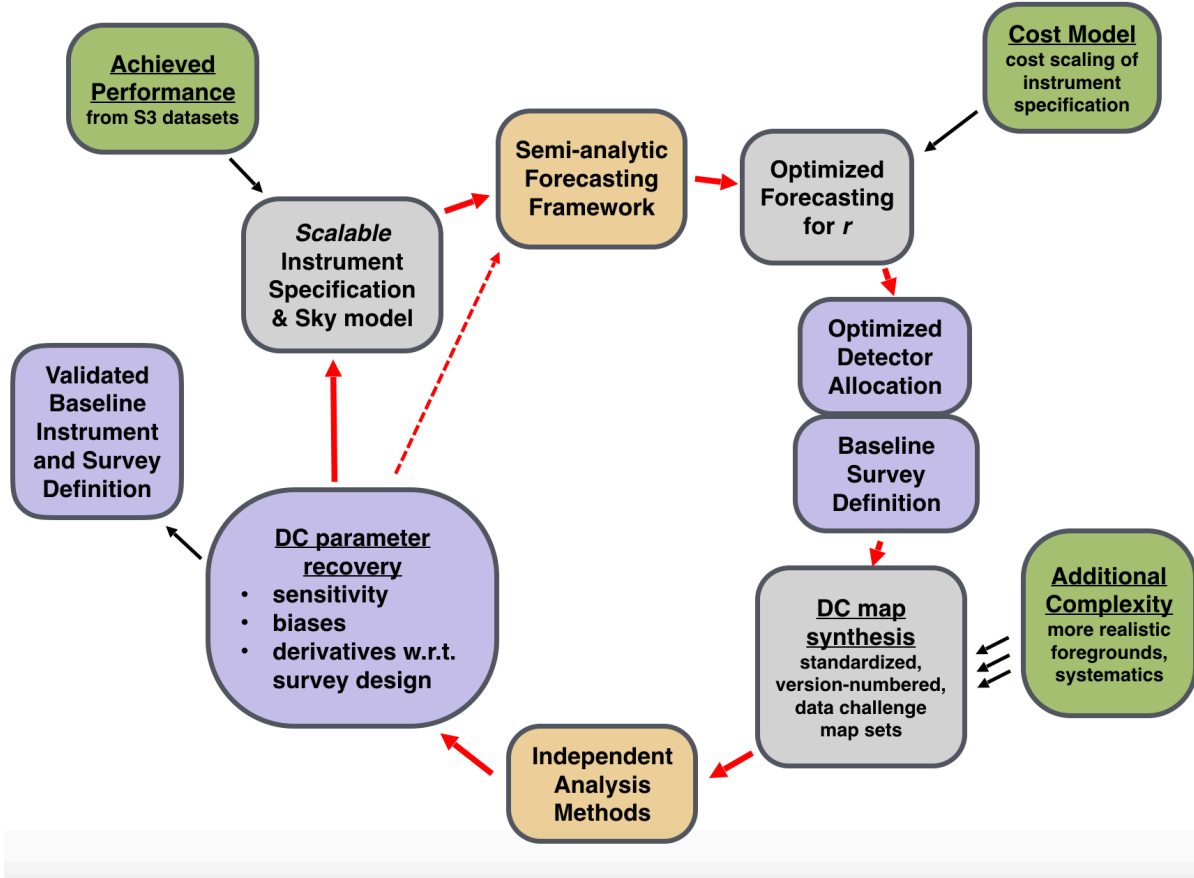


Figure 7.12: Schematic representation of the CMB-S4 Forecasting Loop. We start with achieved performances from Stage 3 datasets, in the form of full covariance matrices and noise spectra, and a set of scalable instrument specifications as well as a fiducial sky model. These are fed as inputs to the Fisher optimization framework, yielding an optimized detector allocation, and a baseline survey definition. Based on this definition, we develop standard Data Challenge noise maps, as well as a suite of signal maps with various degree of complexity. We proceed by analyzing these maps, with multiple independent analysis methods, and check for parameter recovery and presence of biases. If the results suggest a necessary change in survey definition, the process is iterated as needed.




Figure 7.13: This figure shows calculated atmospheric brightness spectra (at zenith) for South Pole at 0.5 mm PWV and Atacama at 1.0 mm PWV (both are near median values). Atmospheric spectra are generated using [Paine \(2017\)](#). The tophat bands are plotted on top of these spectra, with the height of each rectangle equal to the band-averaged brightness temperature using the South Pole spectrum.

also fold in information from two WMAP channels: $\{23, 33\}$ GHz and seven Planck channels: $\{30, 44, 70, 100, 143, 217, 353\}$ GHz, though this extra information is only relevant in the early stages of the optimization.

The optimization unit of effort is equivalent to 500 det-yrs at 150 GHz. For other channels, the number of detectors is calculated as $n_{det,150} \times \left(\frac{\nu}{150}\right)^2$, i.e. assuming comparable focal plane area. The projections run out to a total of 3×10^6 det-yrs, which, if all at 150 GHz, would be equivalent to 500,000 detectors operating for 6 yrs – this seems like a comfortable upper bound for what might be conceivable for S4. S4 scale surveys seem likely to be in the range of 10^6 to 3.0×10^6 det-yrs.

The multipole range is assumed to be $l = [30, 330]$, and using *BICEP/Keck* bandpower

window functions and binning, we obtain 9 bins with nominal centers at l of $\{37.5, 72.5, 107.5, 142.5, 177.5, 212.5, 247.5, 282.5, 317.5\}$.

7.3.2 DELENSING TREATMENT

As mentioned in Section 7.3.1, we assume a separate, high-resolution instrument dedicated to measuring the medium- and small-scale information necessary to construct a template of lensing B modes, so that their effect can be removed. In the semi-analytic optimization process, this instrument is assumed to have 1-arcminute resolution and detector weight at a single frequency. The translation between detector effort and map noise in the delensing instrument is based on the method used for the low-resolution instrument (as described in Sec. 7.1.4 and Eq. 7.8), but with certain non-idealities specific to low-resolution instruments and low- ℓ analysis (such as mode removal and non-uniform coverage) removed. Following the formalism in [Smith et al. \(2012\)](#), we convert the map noise in the delensing survey to a delensing efficiency, or equivalently a fractional residual in lensed B-mode power, specified by setting A_L to the corresponding residual.

The detector effort dedicated to the delensing instrument comes out of the total detector effort budget for the small-area r survey, and the distribution of effort between the low-resolution and delensing instruments is part of the optimization process.

7.3.3 OPTIMIZATION AND PARAMETER CONSTRAINTS

In this section, we try to answer the following question: given a fixed amount of effort, and the instrument specifications offered in the previous sections, what is the optimal effort distribution for foreground cleaning and delensing such that a minimal constraint on r is achieved. To do so we set-up an optimization that calculates the steepest descent through the 10 dimensional space (9 foreground channels + 1 channel for delensing). At each step of the algorithm, the code tries to allocate a fixed amount of effort in each dimension iteratively, rescaling the BPCM accordingly, computing a new Fisher matrix, inverting it to marginalize over all the parameters, and computing the resulting σ_r . Finding the dimension that offers the minimal constraint, the code finally permanently assigns the step of effort to the particular dimension. Though it is generally prohibitive to calculate the full 10D hypercube of σ_r 's, we have validated our approach with the detailed calculation at various points in the optimization. An example of the final path is presented in Figure 7.14.

An additional dimension that is of different nature is the selection of sky fraction. Since we are using BICEP/*Keck* products, which are defined for a particular mask with $\sim f_{sky}^{BK} = 1\%$, we must scale them appropriately for all different fractions. We propagate the effects of sky fraction in the noise spectra and BPCM in two ways: first, we inflate the N_l 's by a factor $\beta = f_{sky}^{S4}/f_{sky}^{BK}$ which boosts the (signal \times noise) and (noise \times noise) terms of the covariance matrix by β and β^2 , to take into account the redistribution of the achieved sensitivity onto a larger patch. Note that the (signal \times signal) component remains unchanged in this step.

Second, we scale down the entire covariance matrix by a factor of β to increase the number of degrees of freedom in the BPCM, accounting for the fact that we are now observing more modes. This procedure scales the signal and noise degrees of freedom independently, preserving the relative effects that filtering, non-uniform coverage, etc have on the covariance structure. In the examples below we pick $f_{sky}^{S4} = 3\%$ as we have done for the CMB-S4 Science Book. The full exercise of optimizing over the f_{sky} dimension is presented in Chapter 2 of [Abazajian et al. \(2016\)](#) as well as [Buza \(2016\)](#).

Upon obtaining the optimized detector count distribution we can calculate the final S4 noise spectra according to Eq.7.7. In order to be able to use these spectra for the design of noise simulations (see 7.3.4.1) it is useful to distill them to a few numbers, to that end we fit them to the following formula

$$N_{l,fit} = \frac{l(l+1)}{2\pi} \frac{\Omega_{pix}}{B_l^2} \left(1 + \left(\frac{l}{l_{knee}} \right)^\gamma \right) \sigma_{map}^2, \quad (7.8)$$

and obtain the map depth σ_{map} , slope γ , and l_{knee} . For the small-aperture data, we find $l_{knee} = 50\text{--}60$ with γ of -2 to -3 , depending on the frequency. The optimized map depths are presented in Figure 7.14.

Figure 7.15 shows the optimized constraints on r as a function of total effort as well as the fraction of effort spent on removing the lensing sample variance and the resulting map rms lensing residual. To reach the desired science goal of $\sigma(r) = 5 \times 10^{-4}$ we note that 1.2×10^6 150 GHz equivalent detector-years are necessary. The optimal distribution of this effort is

summarized in Table 7.5.

Using this distribution, we also show forecasted r - n_s constraints for $r = 0$ and $r = 0.01$ in Figures 7.18 and 7.17, and forecasted r - n_t constraints for $r = 0.01$ and $r = 0.05$ in Figure 7.19. A detailed exploration of the implication of these results is presented in Chapter 2 of [Abazajian et al. \(2016\)](#). Additionally, the discussion presented in Section 7.2.4 is also highly relevant. Finally, Figure 7.16 showcases CMB-S4's ability to constrain the shape of the tensor spectrum given the chosen effort splitting.




Figure 7.14: Optimized map-depth in each of the small-aperture channels as well as in the delensing channel, for an $f_{sky} = 3\%$.




Figure 7.15: Optimized constraints on r as a function of total effort as well as the fraction of effort spent on removing the lensing sample variance and the resulting rms lensing residual.




Figure 7.16: Bin-by-bin forecasted tensor constraints for $r=0.01$, $f_{\text{sky}} = 0.03$, and the default detector effort (10^6 detector years). The boxes denote the forecasted CMB-S4 error bars. Primordial B-mode spectra are shown for two representative values of the tensor-to-scalar ratio: $r=0.001$ and $r=0.01$. The dashed green line shows the Λ CDM expectation for the B modes induced by gravitational lensing of E modes, with the solid line showing the residual lensing power after delensing. The dashed blue and red lines show the dust and synchrotron (current upper limit) model assumed in the forecasting, at the foreground minimum of 95 GHz. The levels of dust and synchrotron are equal to the ones reported in [BICEP2 and Keck Array Collaborations \(2016\)](#). The contribution of dust and synchrotron to the vertical error bars are shown in solid blue and red lines. Since these are calculated from a multi-frequency optimization, the “effective frequency” at which these foreground residuals are defined varies with each bin, allowing the residual lines to go above the input foreground model lines which are defined at a fixed frequency of 95 GHz. Furthermore, due to the low frequency channels having larger beam sizes than the higher frequency ones, in the higher bins, the primordial CMB component will be constrained at a higher effective frequency. Defining the foreground residuals at these effective frequencies will yield a higher amplitude for the dust residual, and a lower amplitude for the synchrotron residual, resulting in the respective shapes of the solid blue and red lines. *Figure from [Abazajian et al. \(2016\)](#)*

Table 7.5: Instrument configuration satisfying the measurement requirements.

Science	Item	Frequency [GHz]									Total
		20	30	40	85	95	145	155	220	270	
$r \dots$	14 x 0.5-m cameras										
	# detectors	...	260	470	17 k	21 k	18 k	21 k	34 k	54 k	168 k
	Angular resolution [FWHM]		77	58	27	24	16	15	11	8.5	
	1 x 6-m telescope										
	# detectors	130	250	500	...	25 k	25 k	...	8.7 k	8.7 k	68 k
	Angular resolution [FWHM]	11	7.0	5.2	...	2.2	1.4	...	1.0	0.8	




Figure 7.17: Forecast of CMB-S4 constraints in the n_s – r plane for a fiducial model with $r = 0.01$. Constraints on r are derived from the expected CMB-S4 sensitivity to the B-mode power spectrum as described in Section 7.1.1. Constraints on n_s are derived from expected CMB-S4 sensitivity to temperature and E-mode power spectra as described in Section 8.10.2 of [Abazajian et al. \(2016\)](#). Also shown are the current best constraints from [BICEP2 and Keck Array Collaborations \(2016\)](#). Chaotic inflation with $V(\phi) = \mu^{4-p}\phi^p$ for $p = 2/3, 1, 2$ are shown as blue lines for $47 < N_* < 57$ (with smaller N_* predicting lower values of n_s). The Starobinsky model and Higgs inflation are shown as small and large filled orange circles, respectively. The lines show the classes of models discussed in Chapter 2 of ([Abazajian et al. \(2016\)](#)). The green band shows the predictions for quartic hilltop models, and the gray band shows the prediction of a sub-class of α -attractor models ([Kallosh & Linde \(2013\)](#)) *Figure from Abazajian et al. (2016)*.




Figure 7.18: Forecast of CMB-S4 constraints in the n_s - r plane for a fiducial model with $r = 0$. Constraints on r are derived from the expected CMB-S4 sensitivity to the B-mode power spectrum as described in Section 7.1.1. Constraints on n_s are derived from expected CMB-S4 sensitivity to temperature and E-mode power spectra as described in Section 8.10.2 of [Abazajian et al. \(2016\)](#). Also shown are the current best constraints from [BICEP2 and Keck Array Collaborations \(2016\)](#). The Starobinsky model and Higgs inflation are shown as small and large filled orange circles. The lines show the classes of models discussed in Chapter 2 of [Abazajian et al. \(2016\)](#) that naturally explain the observed value of the scalar spectral index for different characteristic scales in the potential, $M = M_P/2$, $M = M_P$, $M = 2M_P$, and $M = 5M_P$. Longer dashes correspond to larger values of the scale M . *Figure from [Abazajian et al. \(2016\)](#)*

7.3.4 SIMULATIONS

Using simulations to optimize the design of a CMB experiment inevitably involves a trade-off between the degree of detail that the simulations are able to capture and the computational (and human) cost of generating and analyzing them. This trade-off includes the choice of domain in which the simulation is generated, ranging from the most detailed but most expensive time domain through the map domain to the most simplified but most flexible spectral




Figure 7.19: Forecasts for joint constraints on the tensor to scalar ratio r , and the tensor spectral index, n_t assuming fiducial values of $r = 0.01$ (left) or $r = 0.05$ (right). The pivot scale is set to $k_t = 0.0099 \text{ Mpc}^{-1}$ to break the degeneracy. The forecasts assume $f_{\text{sky}} = 3\%$ and $f_{\text{sky}} = 40\%$ for the fiducial values of $r = 0.01$ and $r = 0.05$, respectively.

domain. Inclusion of additional detail can help to validate general results, to explore their sensitivity to assumptions about foreground models, sky coverage, and instrumental noise and systematics, and in more mature stages of design can inform specific instrument and survey strategy choices.

Here we review the methods used to explore parameter space for the small-area survey, including map level noise simulations, sky models, and observation strategy. We also describe our approach to modeling instrumental systematics, the delensing survey, and the analysis methods.

We also use these simulations to validate the spectral domain forecasts for configurations where the approaches are directly comparable.

7.3.4.1 MAP NOISE SIMULATIONS

As in the Science Book, we have mostly used a sky area of 3%, with some simulations also done for 1% and 10%. We use 7.8 to obtain the desired noise prescription and then generate Gaussian noise realizations at each band. Small-aperture cameras have a very wide instantaneous field of view and hence the observed sky region necessarily has a large edge taper. For the nominal 3% sky coverage simulations, we assumed a circular sky patch with full coverage at $r < 12$ deg and “relative hits” tapering to zero with a cosine-squared shape for $12 \text{ deg} < r < 27 \text{ deg}$. The noise realizations are divided by the square-root of this coverage pattern such that the noise “blows up around the edge” as it does in real maps.

Future timestream simulations will include an explicit scan strategy on the sky and will produce more realistic sky coverage patterns, but for the moment we regard the above as a reasonable compromise between idealism and reality.

7.3.4.2 FOREGROUND MODELS

To make simulated sky maps we add realizations of lensed CMB both without and with an r component to models of the Galactic foregrounds. So far we have run simulations with seven foreground models:

0. Simple Gaussian realizations of synchrotron and dust with power-law angular power spectra at amplitudes set to match the observations in the BICEP/Keck field, and simple uniform SEDs (power law for synchrotron, greybody for dust).

1. The PySM³ model **a1d1f1s1**, where the letters refer to anomalous microwave emission, dust, free-free and synchrotron respectively, and the numbers are the base models described in [Thorne et al. \(2016\)](#).
2. The PySM model **a2d4f1s3**, where the models have been updated to variants that are also described in [Thorne et al. \(2016\)](#). Note that these include 2% polarized AME, a curvature of the synchrotron SED, and a two-temperature model for dust.
3. The PySM model **a2d7f1s3**, where the dust model has been updated to a sophisticated physical model of dust grains as described in [Hensley \(2015\)](#). This model is interesting in that it does not necessarily conform to the greybody SED.
4. The dust model in 3 is replaced by a model of polarized dust emission that incorporates HI column density maps as tracers of the dust intensity structures, and a phenomenological description of the Galactic magnetic field as described in [Ghosh et al. \(2017\)](#). The model is expanded beyond that described in the paper to produce a modest amount of decorrelation of the dust emission pattern as a function of frequency motivated by the analysis of Planck data in [Planck Collaboration \(2017\)](#).
5. A toy model where the dust decorrelation suggested in Figure 3 of [Planck Collaboration \(2017\)](#) is taken at face value ($\Delta_{80}^{217 \times 353} = 0.85$, at $\ell = 80$) and scaled to other frequencies using the functional form given in appendix B of [Vansyngel et al. \(2017\)](#),

³https://github.com/bthorne93/PySM_public

with a linear scaling in ℓ . While such a model is not ruled out by current data it appears to be very hard to produce such strong decorrelation in physics-based models.

We also note that [Sheehy & Slosar \(2017\)](#) have re-analyzed the same Planck data and, while they find that the high level of decorrelation in this model is consistent with the data, their best fit to that same data has no decorrelation.

6. A model based on MHD simulations ([Kritsuk et al. 2017](#)) of the Galactic magnetic field, which naturally produces non-Gaussian correlated dust and synchrotron emission.

Models 1 to 4 use the actual large-scale modes of the real sky as measured above the noise in the Planck data. This means that these models are intrinsically “single-realization,” and this must be borne in mind when interpreting the results. The PySM models fill in the small-scale structure with power-law Gaussian extrapolations, while models 4 and 6 naturally produce non-Gaussian small-scale structure. However, all of these models are consistent with current data, and we should be careful not to necessarily associate nominal sophistication with greater probability to more closely reflect reality.

7.3.4.3 INSTRUMENTAL SYSTEMATICS

As discussed, control of instrumental systematics is a critical design consideration. However, predicting and modeling these effects realistically is a difficult task that is dependent on actual instrument and survey design details, and in any case their impact on an actual result comes not through the modeled effects but through unmodeled residuals. For this study we

have taken the first steps in simulating various generic classes of additive systematic by injecting additional noise-like components into the maps, and then re-analyzing them without knowledge of what was put in. We have experimented with components that are both correlated and uncorrelated across frequency bands, and which have white, $1/\ell$, and white + $1/\ell$ spectra, at varying levels compared to single-frequency map noise or, for correlated cases, combined map noise. Examples of mechanisms that might produce map residuals within this class, after modeling them and either correcting or filtering their leading-order effects, include bandpass mismatches, beam and pointing variations, calibration variations, cross-talk effects, half-wave-plate leakages, ground pickup, and readout irregularities.

Other classes of systematics can be simulated by manipulating the reanalysis procedure only. Examples of such effects include bandpass, polarization angle, calibration, and beam shape uncertainties.

For assessing the impact of instrumental systematics on measurement requirements, in determining both the required survey depths and the maximum allowable levels of systematic effects in the final single-frequency survey maps, our general procedure is to feed parameterizations of various systematic effects into semi-analytic forecasts to judge at what levels classes of systematics introduce parameter biases or additional uncertainties that are significant compared to science targets for those parameters.

7.3.4.4 DELENSING

We have generated high-resolution simulated maps on which we can run explicit lensing reconstruction and then include that information in the analysis. However, that process is not yet converged, and so for the present we approximate delensing by scaling down the Λ CDM lensing signal by the appropriate factor, as described in Section 7.3.2.

7.3.5 ANALYSIS METHODS

To make simulated maps the noise realizations described in Section 7.3.4.1 are added to the sky models described in Section 7.3.4.2. For each realization one then has a stack of multi-frequency $I/Q/U$ maps containing non-uniform noise, foregrounds and signal, and the challenge is to re-analyze them to recover the parameter of interest (in this case r). This can be done by different teams using different methods, and could be done in a blind manner, although we have not done this yet.

So far we have experimented with two methods. The first is a map-based ILC cleaning method (e.g., [Eriksen et al. 2004](#)), which seeks the linear combination of maps that minimizes the remaining CMB signal, followed by a marginalization over residual foregrounds. This method has the advantage that it does not need to know the bandpasses of the frequency channels.

The second method is an evolution of the parametric multi-component fit to the ensemble of auto- and cross-spectra as used for the BICEP/Keck analysis to date ([BICEP2](#), [Keck](#)

Array and Planck Collaborations 2015, BICEP2 and Keck Array Collaborations 2016). This method fits the observed bandpowers to a model composed of the lensing expectation plus dust and synchrotron contributions and a possible r component. Dust and synchrotron each have an amplitude (A_d and A_s), a spatial spectral parameter (α_d and α_s), and a frequency spectral parameter (β_d and β_s). We also allow dust/synchrotron correlation (ϵ), and decorrelation of the foreground patterns over frequency (Δ_d and Δ_s).

Both of these analysis methods are only close to optimal when the foreground behavior is close to uniform across the observing field. For analysis of larger fields, algorithms that fit, for example, the frequency spectral indices individually in (large) pixels, will be required.

7.3.6 RESULTS

Table 7.6 summarizes the results of the analysis for simulations of the CDT Configuration (1.18×10^6 150-GHz-equivalent detector-years) and residual lensing power $A_L = 0.1$. The results from the parametric analysis naturally depend on whether a marginalization over decorrelation is performed, while the ILC does not attempt to capture the effects of decorrelation on the recovery of r and $\sigma(r)$. This is evidenced by the large bias for the ILC method for Model 5 when compared to the parametric analysis that directly accounts for a possible decorrelation (last column). In general, we see that for $r = 0$ the simple Gaussian foreground Model 0 gives $\sigma(r) \sim 5 \times 10^{-4}$, exactly as expected from the semi-analytic code. As we progress to the more complex foreground models, $\sigma(r)$ is generally in the range $5\text{--}8 \times 10^{-4}$.

The level of biases is generally at $\lesssim 0.3\sigma$ for all the models. These simulations are sets

of 500 realizations, so the statistical uncertainty on the bias is $\approx 0.04\sigma$. However, the strong decorrelation in Model 5 as well as the high-significance detection of decorrelation in Model 4 does significantly increase $\sigma(r)$. While the parametric method is able to account for the decorrelation, by construction information is lost, and in fact if one believed in such a scenario, re-optimization to concentrate the sensitivity at closer-in frequencies would be called for.

Table 7.6: Results of two analysis methods applied to map-based simulations assuming the strawperson configuration and our suite of sky models. All simulations assume an instrument configuration including a (high-resolution) 20 GHz channel, a survey of 3% of the sky with $\sim 1.2 \times 10^6$ 150-GHz-equivalent detector-years, and $A_L = 0.1$.

r value	Sky model	ILC		Parametric (no decorr.)		Parametric (incl. decorr.)	
		$\sigma(r) \times 10^{-4}$	r bias $\times 10^{-4}$	$\sigma(r) \times 10^{-4}$	r bias $\times 10^{-4}$	$\sigma(r) \times 10^{-4}$	r bias $\times 10^{-4}$
0	0	5.4	0.3	4.4	0.2	5.6	0.3
	1	8.6	1.1	4.7	6.9	6.4	5.3
	2	8.0	-2.3	4.8	4.0	6.6	2.0
	3	6.8	0.5	4.7	6.1	6.7	0.7
	4	7.5	5.0	7.8	42.0	8.1	-5.8
	5 ^a	16	18	33.6	341.8	12.8	-0.3
	6	5.7	-0.4	4.8	0.6	6.5	1.8
0.003	0	7.2	-4.0	6.2	0.3	8.1	0.4
	1	10	0.2	6.5	7.0	8.5	5.5
	2	10	-2.8	6.5	4.1	8.0	2.1
	3	7.4	-1.4	6.6	6.8	8.7	1.1
	4	10	5.8	9.9	49.7	10.3	-4.0
	5 ^a	20	20	36.1	351.7	14.1	-0.5
	6	8.0	-0.4	7.1	1.6	8.6	2.6

^a An extreme decorrelation model—see Section 7.3.4.2. The parametric analysis includes a decorrelation parameter. No attempt is made in the ILC analysis to model decorrelation. The middle column shows the parametric analysis when we don’t include the decorrelation parameter.

Table 7.7 shows results on detection significance for the strawperson configuration for

sky Model 6. For $r = 0$, the 95% upper limit is about $2.1\sigma(r)$. The value of the tensor-to-scalar ratio for which we expect a 5σ detection is expected after 4 years of operation is $r = 0.004$. For a tensor-to-scalar ratio of $r = 0.003$, the median detection significance after 4 years is expected to be 4σ . If a detection were to be emerging at this point, extending the run time to 8 years would be justified to reach a 5σ detection.

While for given assumptions $\sigma(r)$ can be precisely forecast, the achieved detection level for r depends on the realization of the B -mode field in the observed patch of sky and the instrument noise. Therefore we can only forecast a distribution of detection levels. For a tensor-to-scalar ratio of $r = 0.003$ and the strawperson concept with 8 years of observing we expect to achieve more than 3σ detection with a probability of 0.99, more than 4σ with a probability of 0.93, more than 5σ with a probability of 0.53, and more than 6σ with a probability of 0.14. For simplicity we focus on $\sigma(r)$, and on median detection levels as well as median 95% confidence upper limits to state the typical outcome.

Table 7.7: Results on detection significance for the strawperson concept selected for CMB-S4, using the two analysis methods.

r value	Duration	Sky model	ILC		Parametric (incl. decorrelation)	
			95% CL UL	Det. Sig.	95% CL UL	Det. Sig.
0	4 years	6	1.0×10^{-3}	...	1.0×10^{-3}	...
0.003	4 years	6	...	4.0	...	4.2
	8 years	6	...	5.1	...	5.6

The numbers in Table 7.6 clearly show dependence on the foreground model used in

the simulation. If the actual foregrounds are substantially different from any of these cases, then the biases could be larger. To get some understanding of how large the biases could be, and what instrument modifications might help to reduce them, we have also looked at ILC biases in the extreme case that the foreground residuals are not modeled, but simply absorbed into the estimated B-mode power spectrum. Doing so with simulations based on Sky Model 6 increases the magnitude of the bias on r from 0.4×10^{-4} to 1.3×10^{-4} . The dominant contribution to the bias comes from synchrotron residuals, which motivated placing one lower-frequency channel on the large-aperture. Doing this with the 20 GHz channel improves the angular resolution from $76.6'$ to $11'$, while also increasing the noise power ℓ_{knee} from 50 to 200. We found that this change reduced the magnitude of the bias from 4.0×10^{-4} to 1.3×10^{-4} .

Table 7.8 summarizes the results of analysis of simulations including additive systematic effects, in different combinations of uncorrelated and correlated contamination with varying spectra, added on top of foreground model 3. The levels of systematic contamination for these simulations were chosen to predict biases on r of 1×10^{-4} in semi-analytic forecasts. We can see that the different combinations explored increase biases on r by amounts that typically vary from $0.5\text{--}1.5 \times 10^{-4}$ for the two different analyses, over the different cases. We find that to restrict bias on r to this level, the sum of additive contamination effects needs to be controlled to 3–7% of the single-frequency survey noise, or (in the case of correlated systematics) 6–11% of the total combined noise levels. Such percentages are consistent with the upper limits currently achieved for residual additive systematic contamination compared to survey

noise by small-aperture experiments (e.g., [BICEP2 and Keck Array Collaborations 2016](#)). Assuming CMB-S4 will include a sustained effort to continue to control, understand, and model systematic effects down to levels limited by survey noise, these percentages provide reasonable benchmark requirements.

Table 7.8: Map-based simulation results for simulations containing systematics. Simulations are as in Table 7.6 for sky model 3 and $r = 0$, with additive systematic effects in varying combinations, the amplitudes of which are specified as percentages of survey noise.

Systematic	Uncorrected		Corrected		ILC		Parametric	
	A [%]	B [%]	A [%]	B [%]	$\sigma(r) \times 10^{-4}$	$r \text{ bias} \times 10^{-4}$	$\sigma(r) \times 10^{-4}$	$r \text{ bias} \times 10^{-4}$
None	0	0	0	0	5.3	0.0	7.2	0.0
Uncorrelated white	3.3	0	0	0	6.0	0.84	8.0	0.63
Uncorrelated $1/\ell$. .	0	6.8	0	0	5.0	0.99	7.0	0.85
Correlated white . .	0	0	5.8	0	6.3	1.2	7.3	1.41
Correlated $1/\ell$	0	0	0	10.5	5.2	1.0	6.7	0.97
Uncorr. white + $1/\ell$	1.6	3.5	0	0	5.6	0.89	7.5	0.76
Corr. white + $1/\ell$.	0	0	2.9	5.3	5.5	0.98	6.9	1.04
Both, white + $1/\ell$	0.8	1.7	1.5	2.6	5.6	1.1	7.9	0.98

Results of simulating systematic errors in the determination of bandpasses vary by analysis method. The construction of the ILC method makes it largely insensitive to such errors. The parametric analysis, which includes parametrized models of the frequency spectra of different foregrounds, shows biases on r at the 1×10^{-4} level for uncorrelated random deviations in bandcenter determination of 0.8%, or for correlated deviations of 2%, which we adopt as reasonable benchmark requirements to accommodate a variety of both blind and astrophysical foreground modeling approaches.

In the current simulations, the treatment of delensing is still somewhat crude. We have

explored alternate methods of translating detector effort to map noise, including scaling directly from the noise in a fielded, high-resolution experiment (SPTpol), but with some assumed modifications in per-detector sensitivity. The scaling using this alternate method is slightly more pessimistic than the default scaling, and we adopt this more pessimistic scaling for the numbers in the Science and Measurement Requirements section (Section 2). We have also included multiple frequency bands in the delensing instrument in Section 2, because we cannot conclusively rule out the possibility that non-Gaussianity in small-scale Galactic foregrounds will cause a bias in a single-frequency delensing survey. The level of delensing assumed in the optimization code is reproduced with the configuration in Section 2 if the information from the two “CMB channels” (95 and 145 GHz) can be combined optimally for CMB sensitivity. If instead we assume we have to combine frequency bands to explicitly project out a dust component and a synchrotron component, the delensing efficiency degrades by 5–10%. Effort is currently underway to include delensing with fully non-Gaussian small-scale foregrounds in the Data Challenges. Separate simulations with one model of fully non-Gaussian dust (Vansyngel et al. 2017) indicate that biases to delensing are negligible (also see Challinor et al. 2017), but a larger parameter space of foreground models will need to be explored.

8

Conclusion

Over the coming decade we should see the field of CMB polarization mature considerably.

A number of collaborations are working arduously to pursue not just discovering primordial gravitational waves, but also expanding our knowledge about the small scales by learning more about the masses of the neutrinos, possible new light relics, the nature of dark energy, cluster science, and lensing science, among many others. In Chapter 7, as well as in the con-

cluding sections of Chapters 4, 5, 6, we have touched on the future of the BICEP/Keck program as well as the next generation ground-based CMB experiment CMB-S4. In addition to these, the Atacama Cosmology Telescope (ACT), the South Pole Telescope (SPT), the Simons Observatory (SO), the Simons Array (SA), the “Searching For Echoes of Inflation” (SPIDER) telescope, the Cosmology Large Angular Scales Surveyor (CLASS), and possible other experiments, are set to continue developing new technologies and analysis methods which will notably change the landscape of cosmology.

For detecting primordial B-modes, one of the main challenges is the aforementioned lensing B-mode contribution from the weak lensing of E-modes as the CMB photons travel to us. For instrument noise levels below $5 \mu\text{K-arcmin}$, this lensing signal becomes the dominant source of contamination and its sample variance significantly worsens our constraining power on PGWs. Unlike galactic foregrounds, observing at multiple frequencies is not sufficient to remove this signal. However, its contribution can be reduced by knowing the cumulative gravitational lensing potential ϕ along the line of sight, and having a faithful E-mode map. Together, the two can be combined to form a lensing B-mode template by lensing the E-mode map with the ϕ field and subtracting this template from the measured B-mode map. This technique is known in the field as *delensing*.

Though the theoretical foundation for delensing in pursuit of PGWs was laid out over a decade ago, in practice this analysis has never been done before. In addition, a proper delensing analysis will need to account for galactic foregrounds at large scales (for PGWs) and small scales (for the ϕ tracer). The likelihood framework described in this manuscript is able

to do the former by including a lensing template as an extra virtual frequency band, and taking all the appropriate cross-correlations. There is ongoing effort from the BICEP/*Keck* and SPT collaborations to use this framework to demonstrate delensing of the BICEP/*Keck* maps by forming a lensing potential from the Cosmic Infrared Background (CIB) data from Planck and an E-mode map from a combination of SPTpol, BK and Planck maps.

Going forward, as CMB data becomes more sensitive, the fractional value of delensing with the CIB will become increasingly smaller, and we will need better data sets from which to derive large-scale structure tracers. Since the gravitational lensing deflections of the CMB are only on the order of 2 arcmin, we may reconstruct ϕ with high fidelity using high sensitivity, high angular resolution CMB polarization maps – specifically, in the case of BICEP/*Keck*, using SPT-3G maps. The work described in Chapter 7 shows that after four years of observations with SPT-3G, we can reduce the effective lensing power by a factor of more than three, and, in conjunction with BICEP Array, achieve a constraint of $\sigma(r) \sim 0.005$. Over the coming years, we should expect to see work towards build delensing pipelines to perform foreground cleaning at small scales, produce unbiased tracers, and demonstrate that a large reduction in lensing is possible. This will also serve as necessary groundwork for CMB-S4, for which the calculations in Chapter 7 predict that we will need to reduce the lensing power by 90% in order to achieve $\sigma(r) \sim 0.0005$.

In the near future, first with Stage 3 experiments and later with CMB-S4, the CMB community will encounter new data processing and analysis challenges that will stem from significantly increasing detector counts and experimental complexity and pushing the de-

sired science thresholds to unprecedented levels. Raising the number of detectors to $10^5 - 10^6$ brings a new hurdle to data cuts and systematic studies – the data volume is large enough that current methods are insufficient, and constructing multiple time-ordered-data simulations may be prohibitive. Significant work will be needed to automate these processes and perform end-to-end simulations for the aforementioned experiments.

Finally, as the CMB field evolves, improving parameter constraints will require combining different data sets from both CMB and large-scale structure experiments. For instance, the CMB-S4 baseline includes both small-aperture telescopes (foreground mitigation for r) and large-aperture telescopes (delensing for r); BICEP/*Keck*/SPT are collaborating; so are other experiments in the field, such as Simons Array/Simons Observatory/AdvACT. Furthermore, cross-correlations of CMB data with DES, DESI, LSST, and Euclid will prove invaluable in constraining properties of dark energy, neutrino masses, General Relativity, lensing and cluster science, and beyond.



Gaussian variables and expectation values

This Appendix includes some useful calculations with Gaussian variables. The expressions presented here are reproduced (with permission) from a set of notes put together by Colin Bischoff. In this document, we will be working exclusively with Gaussian fields with zero mean. We will denote the fields with lower case Roman letters, a, b, c, \dots , and variances $\sigma_a^2, \sigma_b^2, \sigma_c^2, \dots$

The normalized probability density function for a Gaussian variable is given by

$$P(a) = \frac{1}{\sqrt{2\pi\sigma_a^2}} \exp\left(-\frac{a^2}{2\sigma_a^2}\right). \quad (\text{A.1})$$

$P(a)$ is an even function, so we immediately know that $\langle a^\zeta \rangle = 0$ for odd values of ζ . For even values of ζ , $\langle a^\zeta \rangle$ is non-zero. The following identities will be frequently used throughout this document:

$$\langle a^0 \rangle = \int_{-\infty}^{\infty} P(a) da = 1 \quad (\text{A.2})$$

$$\langle a^2 \rangle = \int_{-\infty}^{\infty} a^2 P(a) da = \sigma_a^2 \quad (\text{A.3})$$

$$\langle a^4 \rangle = \int_{-\infty}^{\infty} a^4 P(a) da = 3\sigma_a^4 \quad (\text{A.4})$$

A.1 BANDPOWERS

To measure power, we just take the square of the field, i.e. a^2 . We can see from Equation (A.3) that the expectation value of this power estimate $\langle a^2 \rangle = \sigma_a^2$, as desired.

We can also measure the cross-power between two fields by taking their product, i.e. ab .

If fields a and b are uncorrelated, then the expectation value factorizes: $\langle ab \rangle = \langle a \rangle \langle b \rangle = 0$.

On the other hand, if a and b are partially correlated, then we can rewrite them in terms of

uncorrelated fields x , y , and z .

$$a = x + y$$

$$b = \lambda x + z$$

where

$$\sigma_a^2 = \sigma_x^2 + \sigma_y^2$$

$$\sigma_b^2 = \lambda^2 \sigma_x^2 + \sigma_z^2.$$

Then we can calculate the expectation value of the cross-power as

$$ab = \lambda x^2 + xz + \lambda xy + yz$$

$$\langle ab \rangle = \lambda \sigma_x^2.$$

For a bandpower, we will average the power estimates for many modes of the field. For this document, I will use an abstract but simplified concept of modes for the field. In actual data analysis, we start with a Gaussian map described in the pixel basis, then change basis to Fourier space (statistics should still be Gaussian), and calculate bandpowers by averaging over annuli. The number of points to average over depends on the Fourier plane gridding, but these points are correlated with one another, leading to a smaller effective number of

modes that is determined by the apodization mask and filtering. I will be working in the basis of these effective degrees-of-freedom, which I will take to be k independent modes, indexed by subscripts i or j . As a further simplification, I will assume that all modes of the field have identical variance (relaxing this assumption would just lead to more bookkeeping, with the i subscript attached to σ_a^2).

Now the auto-spectrum bandpower for field a , and its expectation value, can be written as

$$\mathcal{C}_{aa} = \frac{1}{k} \sum_{i=1}^k a_i^2 \tag{A.5}$$

$$\langle \mathcal{C}_{aa} \rangle = \frac{1}{k} \sum_i \langle a_i^2 \rangle \tag{A.6}$$

$$= \frac{1}{k} \sum_i \sigma_a^2 \tag{A.7}$$

$$= \sigma_a^2. \tag{A.8}$$

The expectation value operation, denoted by angle brackets, is an average over independent realizations of the field, not an average over modes i (those are iterated using the summation symbol).

A.2 BANDPOWER COVARIANCE MATRIX

A.2.1 VARIANCE OF AUTO-SPECTRUM

First, let's calculate the variance of \mathcal{C}_{aa} .

$$\text{Var}(\mathcal{C}_{aa}) = \langle \mathcal{C}_{aa}^2 \rangle - \langle \mathcal{C}_{aa} \rangle^2 \quad (\text{A.9})$$

$$\langle \mathcal{C}_{aa}^2 \rangle = \frac{1}{k^2} \left\langle \left(\sum_i a_i^2 \right) \left(\sum_j a_j^2 \right) \right\rangle. \quad (\text{A.10})$$

The product of these two summations contains k^2 terms, which we can divide up into k terms with $i = j$ and $k(k-1)$ terms with $i \neq j$.

$$\langle \mathcal{C}_{aa}^2 \rangle = \frac{1}{k^2} \left[\sum_i \langle a_i^4 \rangle + \left(\sum_i \langle a_i^2 \rangle \right) \left(\sum_{j \neq i} \langle a_j^2 \rangle \right) \right] \quad (\text{A.11})$$

$$= \frac{1}{k^2} [3k\sigma_a^4 + k(k-1)\sigma_a^4] \quad (\text{A.12})$$

$$= \left(\frac{2}{k} + 1 \right) \sigma_a^4 \quad (\text{A.13})$$

$$\text{Var}(\mathcal{C}_{aa}) = \left(\frac{2}{k} + 1 \right) \sigma_a^4 - \sigma_a^4 \quad (\text{A.14})$$

$$= \frac{2}{k} \sigma_a^4. \quad (\text{A.15})$$

We see that the bandpower variance is proportional to the square of the power and inversely proportional to the number of modes observed.

A.2.2 VARIANCE OF CROSS-SPECTRUM

Next, consider the variance of a cross-spectrum bandpower, for the case where fields a and b are uncorrelated.

$$Var(\mathcal{C}_{ab}) = \langle \mathcal{C}_{ab}^2 \rangle - \langle \mathcal{C}_{ab} \rangle^2 \quad (\text{A.16})$$

$$\langle \mathcal{C}_{ab}^2 \rangle = \frac{1}{k} \left\langle \left(\sum_i a_i b_i \right) \left(\sum_j a_j b_j \right) \right\rangle \quad (\text{A.17})$$

$$= \frac{1}{k^2} \left[\sum_i \langle a_i^2 b_i^2 \rangle + \left(\sum_i \langle a_i b_i \rangle \right) \left(\sum_{j \neq i} \langle a_j b_j \rangle \right) \right] \quad (\text{A.18})$$

$$= \frac{1}{k^2} [k \sigma_a^2 \sigma_b^2] \quad (\text{A.19})$$

$$= \frac{1}{k} \sigma_a^2 \sigma_b^2 \quad (\text{A.20})$$

$$Var(\mathcal{C}_{ab}) = \frac{1}{k} \sigma_a^2 \sigma_b^2 \quad (\text{A.21})$$

$$= \frac{1}{2} \sqrt{Var(\mathcal{C}_{aa}) Var(\mathcal{C}_{bb})} \quad (\text{A.22})$$

We see that the variance of the ab cross-spectrum is one half the geometric mean of the variances for the aa and bb auto-spectra.

If fields a and b are totally correlated, then we can write $b = \lambda a$ and calculate the cross-

spectrum variance.

$$\mathcal{C}_{bb} = \lambda^2 \mathcal{C}_{aa} \tag{A.23}$$

$$\text{Var}(\mathcal{C}_{bb}) = \lambda^4 \text{Var}(\mathcal{C}_{aa}) \tag{A.24}$$

$$\mathcal{C}_{ab} = \lambda \mathcal{C}_{aa} \tag{A.25}$$

$$\text{Var}(\mathcal{C}_{ab}) = \lambda^2 \text{Var}(\mathcal{C}_{aa}) \tag{A.26}$$

$$= \sqrt{\text{Var}(\mathcal{C}_{aa}) \text{Var}(\mathcal{C}_{bb})} \tag{A.27}$$

We see that for the full correlated case, there is no longer any factor of one half between the variance of the cross-spectrum and the geometric mean of the auto-spectra variances.

For the most general case with partial correlation between a and b , we will switch to uncorrelated fields x, y, z , as in Section A.1.

FIELD a AUTO-SPECTRUM:

$$C_{aa} = \frac{1}{k} \sum_i (x_i^2 + 2x_i y_i + y_i^2) \quad (\text{A.28})$$

$$\langle C_{aa} \rangle = \sigma_x^2 + \sigma_y^2 \quad (\text{A.29})$$

$$\langle C_{aa}^2 \rangle = \frac{1}{k^2} \left\langle \left(\sum_i x_i^2 + 2x_i y_i + y_i^2 \right) \left(\sum_j x_j^2 + 2x_j y_j + y_j^2 \right) \right\rangle \quad (\text{A.30})$$

$$\begin{aligned} &= \frac{1}{k^2} \left[\sum_i (\langle x_i^4 \rangle + 4 \langle x_i^3 y_i \rangle + 2 \langle x_i^2 y_i^2 \rangle + 4 \langle x_i^2 y_i^2 \rangle + 4 \langle x_i y_i^3 \rangle + \langle y_i^4 \rangle) + \right. \\ &\quad \left. \sum_i \sum_{j \neq i} (\langle x_i^2 x_j^2 \rangle + 4 \langle x_i^2 x_j y_j \rangle + 2 \langle x_i^2 y_j^2 \rangle + 4 \langle x_i y_i x_j y_j \rangle + 4 \langle x_i y_i y_j^2 \rangle + \langle y_i^2 y_j^2 \rangle) \right] \quad (\text{A.31}) \end{aligned}$$

$$= \frac{1}{k^2} [3k\sigma_x^4 + 6k\sigma_x^2\sigma_y^2 + 3k\sigma_y^4 + k(k-1)\sigma_x^4 + 2k(k-1)\sigma_x^2\sigma_y^2 + k(k-1)\sigma_y^4] \quad (\text{A.32})$$

$$= \left(\frac{2}{k} + 1 \right) (\sigma_x^2 + \sigma_y^2)^2 \quad (\text{A.33})$$

$$\text{Var}(C_{aa}) = \frac{2}{k} (\sigma_x^2 + \sigma_y^2)^2 \quad (\text{A.34})$$

FIELD b AUTO-SPECTRUM:

$$C_{bb} = \frac{1}{k} \sum_i (\lambda^2 x_i^2 + 2\lambda x_i z_i + z_i^2) \quad (\text{A.35})$$

$$\langle C_{bb} \rangle = \lambda^2 \sigma_x^2 + \sigma_z^2 \quad (\text{A.36})$$

$$\text{Var}(C_{bb}) = \frac{2}{k} (\lambda^2 \sigma_x^2 + \sigma_z^2)^2 \quad (\text{A.37})$$

CROSS-SPECTRUM BETWEEN a AND b :

$$\mathcal{C}_{ab} = \frac{1}{k} \sum_i (\lambda x_i^2 + x_i z_i + \lambda x_i y_i + y_i z_i) \quad (\text{A.38})$$

$$\langle \mathcal{C}_{ab} \rangle = \lambda \sigma_x^2 \quad (\text{A.39})$$

$$\text{Var}(\mathcal{C}_{ab}) = \frac{1}{k} [2\lambda^2 \sigma_x^4 + \sigma_x^2 \sigma_z^2 + \lambda^2 \sigma_x^2 \sigma_y^2 + \sigma_y^2 \sigma_z^2] \quad (\text{A.40})$$

$$= \frac{1}{2} \sqrt{\text{Var}(\mathcal{C}_{aa}) \text{Var}(\mathcal{C}_{bb})} + \frac{1}{k} \lambda^2 \sigma_x^4 \quad (\text{A.41})$$

$$= \frac{1}{2} \sqrt{\text{Var}(\mathcal{C}_{aa}) \text{Var}(\mathcal{C}_{bb})} \left(1 + \frac{\langle \mathcal{C}_{ab} \rangle^2}{\langle \mathcal{C}_{aa} \rangle \langle \mathcal{C}_{bb} \rangle} \right) \quad (\text{A.42})$$

When fields a and b are uncorrelated, $\langle \mathcal{C}_{ab} \rangle = 0$ and Equation (A.42) reduces to the expression given in Equation (A.22). In the other limit, when the fields are totally correlated, $\langle \mathcal{C}_{ab} \rangle^2 = \langle \mathcal{C}_{aa} \rangle \langle \mathcal{C}_{bb} \rangle$ and we obtain the expression from Equation (A.27).

A.2.3 COVARIANCE FOR UNCORRELATED FIELDS

The covariance between \mathcal{C}_{aa} and \mathcal{C}_{bb} is zero if fields a and b are totally uncorrelated. We can see this because, for that case,

$$\text{Cov}(\mathcal{C}_{aa}, \mathcal{C}_{bb}) = \langle \mathcal{C}_{aa} \mathcal{C}_{bb} \rangle - \langle \mathcal{C}_{aa} \rangle \langle \mathcal{C}_{bb} \rangle \quad (\text{A.43})$$

$$= \langle \mathcal{C}_{aa} \rangle \langle \mathcal{C}_{bb} \rangle - \langle \mathcal{C}_{aa} \rangle \langle \mathcal{C}_{bb} \rangle \quad (\text{A.44})$$

$$= 0. \quad (\text{A.45})$$

Next, consider the covariance between \mathcal{C}_{aa} and \mathcal{C}_{ab} , again for a and b uncorrelated.

$$Cov(\mathcal{C}_{aa}, \mathcal{C}_{ab}) = \langle \mathcal{C}_{aa} \mathcal{C}_{ab} \rangle - \langle \mathcal{C}_{aa} \rangle \langle \mathcal{C}_{ab} \rangle \quad (\text{A.46})$$

$$\langle \mathcal{C}_{aa} \rangle = \sigma_a^2 \quad (\text{A.47})$$

$$\langle \mathcal{C}_{ab} \rangle = 0 \quad (\text{A.48})$$

$$\langle \mathcal{C}_{aa} \mathcal{C}_{ab} \rangle = \frac{1}{k^2} \left\langle \left(\sum_i a_i^2 \right) \left(\sum_j a_j b_j \right) \right\rangle \quad (\text{A.49})$$

$$= \frac{1}{k^2} \left[\sum_i \langle a_i^3 b_i \rangle + \sum_i \sum_{j \neq i} \langle a_i^2 a_j b_j \rangle \right] \quad (\text{A.50})$$

$$= 0. \quad (\text{A.51})$$

So the (unsurprising) result for uncorrelated fields is that the bandpower covariance matrix is diagonal.

A.2.4 COVARIANCE FOR CORRELATED FIELDS

If we work out the case of partially correlated a and b , we get the following covariance between auto-spectra:

$$Cov(\mathcal{C}_{aa}, \mathcal{C}_{bb}) = Cov(\mathcal{C}_{xx} + 2\mathcal{C}_{xy} + \mathcal{C}_{yy}, \lambda^2 \mathcal{C}_{xx} + 2\lambda \mathcal{C}_{xz} + \mathcal{C}_{zz}) \quad (\text{A.52})$$

$$= \lambda^2 Var(\mathcal{C}_{xx}) \quad (\text{A.53})$$

$$= \frac{2}{k} \lambda^2 \sigma_x^4 \quad (\text{A.54})$$

$$= \sqrt{Var(\mathcal{C}_{aa}) Var(\mathcal{C}_{bb})} \left(\frac{\langle \mathcal{C}_{ab} \rangle^2}{\langle \mathcal{C}_{aa} \rangle \langle \mathcal{C}_{bb} \rangle} \right). \quad (\text{A.55})$$

The covariance between an auto-spectrum and the cross-spectrum is

$$Cov(\mathcal{C}_{aa}, \mathcal{C}_{ab}) = Cov(\mathcal{C}_{xx} + 2\mathcal{C}_{xy} + \mathcal{C}_{yy}, \lambda \mathcal{C}_{xx} + \mathcal{C}_{xz} + \lambda \mathcal{C}_{xy} + \mathcal{C}_{yz}) \quad (\text{A.56})$$

$$= \lambda Var(\mathcal{C}_{xx}) + 2\lambda Var(\mathcal{C}_{xy}) \quad (\text{A.57})$$

$$= \frac{2}{k} \lambda \sigma_x^2 (\sigma_x^2 + \sigma_y^2) \quad (\text{A.58})$$

$$= [Var(\mathcal{C}_{aa})^3 Var(\mathcal{C}_{bb})]^{1/4} \left(\frac{\langle \mathcal{C}_{aa} \rangle \langle \mathcal{C}_{ab} \rangle}{\sqrt{\langle \mathcal{C}_{aa} \rangle^3 \langle \mathcal{C}_{bb} \rangle}} \right). \quad (\text{A.59})$$

In the most general case, we have four fields, a, b, c, d , with an arbitrary covariance structure given by \mathbf{C} (note: referring here to the *field* covariance, not bandpower covariance). We can solve this by diagonalizing the covariance matrix, $\mathbf{C} = \mathbf{P}^T \mathbf{D} \mathbf{P}$, where \mathbf{P} contains eigenvectors that let us transform to an alternate basis of four fields, w, x, y, z , with diagonal co-

variance matrix **D**. Using this orthogonal basis, we can solve

$$Cov(\mathcal{C}_{ab}, \mathcal{C}_{cd}) = \frac{1}{2} [Var(\mathcal{C}_{aa}) Var(\mathcal{C}_{bb}) Var(\mathcal{C}_{cc}) Var(\mathcal{C}_{dd})]^{1/4} \left(\frac{\langle \mathcal{C}_{ac} \rangle \langle \mathcal{C}_{bd} \rangle + \langle \mathcal{C}_{ad} \rangle \langle \mathcal{C}_{bc} \rangle}{\sqrt{\langle \mathcal{C}_{aa} \rangle \langle \mathcal{C}_{bb} \rangle \langle \mathcal{C}_{cc} \rangle \langle \mathcal{C}_{dd} \rangle}} \right). \quad (\text{A.60})$$

It takes a lot of algebra to get to Equation (A.60), but you can pretty much guess at it by inspection of Equations (A.42, A.55, A.59). The product of auto-spectrum variances (term in square brackets) ensures that if you rescale any of the four fields, the covariance will rescale properly. The correlations between fields are encoded in the ratios between expectation values of cross-spectra and the geometric mean of the auto-spectrum expectation values (term in parentheses).

B

Building a Semi-analytic Bandpower Covariance Matrix

B.1 INTRODUCTION

This appendix describes the development of a framework to build a bandpower covariance matrix (BPCM) that relies on analytic reasoning and signal-only/noise-only simulations. We

assume that the signal is independent of noise, thereby throwing out random signal-noise correlations that come into a direct estimation of the covariance matrix. Early in the development of the multi-component framework, we have seen that the results depended a lot on the off-diagonal elements in each ℓ -block of the band power covariance matrix. We saw that by zeroing out certain elements, under the assumptions that they should be driven to zero in the limit of an infinite simulation set, we got different results than if we left those elements in. However, this zero-ing was approximate, and it was required to take a closer look at what each of the elements in BPCM looked like in the case of a limited number of simulations.

The framework is capable of taking in N experiments with separate signal and noise simulations; however, in our case, we will distinguish between independently simulated experiments (i.e., experiments like BICEP2, BICEP1) and dependently simulated experiments (i.e., experiments simulated by re-observing BICEP2, like $WMAP_K$). For each independent experiment, the inputs are signal expectation values, signal variance, noise expectation values, and noise variance. For each dependent experiment, the inputs are just noise expectation values and noise variance. As usual, we zero BPCM terms that differ by more than one ℓ -bin.

In addition to that, the new framework allows one to input multiple types of signal (like dust or sync) into the BPCM, by scaling the already available simulations to reflect any choice of underlying model (as seen in sect;2). This lets us perform our usual H-L calculation and test how dependent it is on the choice of fiducial model, or alternatively it lets us perform a dynamic likelihood calculation by evaluating the BPCM at each point in the phase space we are exploring, and using a simple χ^2 likelihood.

B.2 GENERAL SEMI-ANALYTIC FORMULATION

Let the letters a, b, c, d denote the different experiments in our analysis, and the subscripts $i = 1, 2, 3$ denote the various signal components: CMB, dust, sync, for each of the aforementioned experiments. Then, the total signal for experiment a is the sum of all components $\sum_i^N a_i$, where N is the number of signal components, and each a_i is a Gaussian field with mean zero. Furthermore, we will assume that all signals are independent of noise, and each experiment has different noise n_{expt} . Then, a generic covariance between two spectra is represented by:

$$\begin{aligned}
& Cov\{Expt(A) \times Expt(B), Expt(C) \times Expt(D)\} = \\
& = \langle (\sum_i a_i + n_a)(\sum_i b_i + n_b)(\sum_i c_i + n_c)(\sum_i d_i + n_d) \rangle + \\
& - \langle (\sum_i a_i + n_a)(\sum_i b_i + n_b) \rangle \langle (\sum_i c_i + n_c)(\sum_i d_i + n_d) \rangle = \\
& = \langle (\sum_i a_i \sum_i b_i + n_a \sum_i b_i + n_b \sum_i a_i + n_a n_b)(\sum_i c_i \sum_i d_i + n_c \sum_i d_i + n_d \sum_i c_i + n_c n_d) \rangle + \\
& - \langle \sum_i a_i \sum_i b_i + n_a n_b \rangle \langle \sum_i c_i \sum_i d_i + n_c n_d \rangle = \\
& = \langle \sum_i a_i \sum_i b_i \sum_i c_i \sum_i d_i \rangle + \langle n_c n_d \rangle \langle \sum_i a_i \sum_i b_i \rangle + \langle n_b n_d \rangle \langle \sum_i a_i \sum_i c_i \rangle + \\
& + \langle n_b n_c \rangle \langle \sum_i a_i \sum_i d_i \rangle + \langle n_a n_d \rangle \langle \sum_i b_i \sum_i c_i \rangle + \langle n_a n_c \rangle \langle \sum_i b_i \sum_i d_i \rangle + \\
& + \langle n_a n_b \rangle \langle \sum_i c_i \sum_i d_i \rangle + \langle n_a n_b n_c n_d \rangle - \langle \sum_i a_i \sum_i b_i \rangle \langle \sum_i c_i \sum_i d_i \rangle + \\
& - \langle n_c n_d \rangle \langle \sum_i a_i \sum_i b_i \rangle - \langle n_a n_b \rangle \langle \sum_i c_i \sum_i d_i \rangle - \langle n_a n_b \rangle \langle n_c n_d \rangle = \\
& = \langle \sum_i a_i \sum_i b_i \sum_i c_i \sum_i d_i \rangle - \langle \sum_i a_i \sum_i b_i \rangle \langle \sum_i c_i \sum_i d_i \rangle + \langle n_a n_b n_c n_d \rangle + \\
& - \langle n_a n_b \rangle \langle n_c n_d \rangle + \langle n_b n_d \rangle \langle \sum_i a_i \sum_i c_i \rangle + \langle n_b n_c \rangle \langle \sum_i a_i \sum_i d_i \rangle + \\
& + \langle n_a n_d \rangle \langle \sum_i b_i \sum_i c_i \rangle + \langle n_a n_c \rangle \langle \sum_i b_i \sum_i d_i \rangle
\end{aligned}
\tag{B.1}$$

We would like to simplify the difference between the first two terms of the last equality.

Taking a closer look at the first term, and making use of the assumption that two different types of signal are uncorrelated, that is that $\langle a_i a_j \rangle = \langle b_i b_j \rangle = \langle a_i b_j \rangle = \dots = 0$ for $i \neq j$, we find

$$\langle \sum_i a_i \sum_i b_i \sum_i c_i \sum_i d_i \rangle = \sum_i \langle a_i b_i c_i d_i \rangle + \sum_{i \neq j} \langle a_i b_i \rangle \langle c_j d_j \rangle \quad (\text{B.2})$$

$$+ \sum_{i \neq j} \langle a_i c_i \rangle \langle b_j d_j \rangle + \sum_{i \neq j} \langle a_i d_i \rangle \langle b_j c_j \rangle \quad (\text{B.3})$$

Similarly, the second term expands to

$$\langle \sum_i a_i \sum_i b_i \rangle \langle \sum_i c_i \sum_i d_i \rangle = \sum_i \langle a_i b_i \rangle \langle c_i d_i \rangle + \sum_{i \neq j} \langle a_i b_i \rangle \langle c_j d_j \rangle$$

And the difference between the two yields

$$\sum_i \text{Cov}\{a_i b_i, c_i d_i\} + \sum_{i \neq j} \langle a_i c_i \rangle \langle b_j d_j \rangle + \sum_{i \neq j} \langle a_i d_i \rangle \langle b_j c_j \rangle$$

Which brings the total answer to

$$\begin{aligned} & \sum_i \text{Cov}\{a_i b_i, c_i d_i\} + \text{Cov}\{n_a n_b, n_c n_d\} + \langle n_b n_d \rangle \langle \sum_i a_i c_i \rangle + \langle n_b n_c \rangle \langle \sum_i a_i d_i \rangle + \\ & + \langle n_a n_d \rangle \langle \sum_i b_i c_i \rangle + \langle n_a n_c \rangle \langle \sum_i b_i d_i \rangle + \sum_{i \neq j} \langle a_i c_i \rangle \langle b_j d_j \rangle + \sum_{i \neq j} \langle a_i d_i \rangle \langle b_j c_j \rangle \end{aligned}$$

We now have to think how to go from this Gaussian field representation to a bandpower rep-

resentation, after all, the benefit of this semi-analytic BPCM construction is that we could build it from signal-only, and noise-only aps, which contain bandpower information.

To construct a bandpower, one has to sum over the map modes. One might notice that until now there was no mention of different map modes; this is because the above case is, in fact, a special case in which we have only one map mode. In the case of multiple map modes, and different modes from spectra to spectra, one has to take an extra sum. Denoting by the Greek letters α, β the map mode number, and leaving the Roman letters i, j to represent different types of signal, the line

$$< (\sum_i a_i + n_a)(\sum_i b_i + n_b)(\sum_i c_i + n_c)(\sum_i d_i + n_d) > \quad (\text{B.4})$$

$$- < (\sum_i a_i + n_a)(\sum_i b_i + n_b) > < (\sum_i c_i + n_c)(\sum_i d_i + n_d) > \quad (\text{B.5})$$

becomes

$$< \sum_{\alpha} \{ (\sum_i a_{i\alpha} + n_{a\alpha})(\sum_i b_{i\alpha} + n_{b\alpha}) \} \sum_{\beta} \{ (\sum_i c_{i\beta} + n_{c\beta})(\sum_i d_{i\beta} + n_{d\beta}) \} > \quad (\text{B.6})$$

$$- < \sum_{\alpha} \{ (\sum_i a_{i\alpha} + n_{a\alpha})(\sum_i b_{i\alpha} + n_{b\alpha}) \} > < \sum_{\beta} \{ (\sum_i c_{i\beta} + n_{c\beta})(\sum_i d_{i\beta} + n_{d\beta}) \} > \quad (\text{B.7})$$

We will have all the same terms as in the calculation above, except that now we have to be careful about different map modes. To make it clear how this changes the formulation, let's

look at the terms that contributed to the Signal Covariance term in the boxed equation

$$< \sum_i a_i \sum_j b_j \sum_m c_m \sum_p d_p > - < \sum_i a_i \sum_j b_j > < \sum_m c_m \sum_p d_p >$$

this difference will now look like

$$\begin{aligned}
& < \sum_{i,\alpha} a_{i\alpha} \sum_{j,\alpha} b_{j\alpha} \sum_{m,\beta} c_{m\beta} \sum_{p,\beta} d_{p\beta} > - < \sum_{i,j,\alpha} a_{i\alpha} b_{j\alpha} > < \sum_{m,p,\beta} c_{m\beta} d_{p\beta} > = \\
& = < \sum_{i,j,\alpha} a_{i\alpha} b_{j\alpha} \sum_{m,p,\beta} c_{m\beta} d_{p\beta} > - < \sum_{i,j,\alpha} a_{i\alpha} b_{j\alpha} > < \sum_{m,p,\beta} c_{m\beta} d_{p\beta} > = \\
& = \sum_{i,\alpha,\beta} < a_{i\alpha} b_{i\alpha} c_{i\beta} d_{i\beta} > + \sum_{i \neq j, \alpha, \beta} < a_{i\alpha} c_{i\beta} > < b_{j\alpha} d_{j\beta} > + \sum_{i \neq j, \alpha, \beta} < a_{i\alpha} d_{i\beta} > < b_{j\alpha} c_{j\beta} > + \\
& - \sum_{i,\alpha,\beta} < a_{i\alpha} b_{i\alpha} > < c_{i\beta} d_{i\beta} > = \\
& = \sum_i Cov\{ \sum_{\alpha} a_{i\alpha} b_{i\alpha}, \sum_{\beta} c_{i\beta} d_{i\beta} \} + \sum_{i \neq j, \alpha, \beta} < a_{i\alpha} c_{i\beta} > < b_{j\alpha} d_{j\beta} > + \sum_{i \neq j, \alpha, \beta} < a_{i\alpha} d_{i\beta} > < b_{j\alpha} c_{j\beta} > = \\
& = \begin{cases} \sum_i Cov\{ \sum_{\alpha} a_{i\alpha} b_{i\alpha}, \sum_{\beta} c_{i\beta} d_{i\beta} \} + \sum_{i \neq j, \gamma} < a_{i\gamma} c_{i\gamma} > < b_{j\gamma} d_{j\gamma} > + \\ & + \sum_{i \neq j, \gamma} < a_{i\gamma} d_{i\gamma} > < b_{j\gamma} c_{j\gamma} >, \text{if there is mode overlap} \\ 0, \text{if there is no mode overlap} \end{cases}
\end{aligned} \tag{B.8}$$

where γ is the set of overlapping modes and where in the last step we took into account that the map modes are independent, so that $< a_{i\alpha} b_{i\alpha} > = 0$ if $\alpha \neq \beta$.

Now, for a cross term of the type

$$< n_b n_d > < \sum_i a_i \sum_j c_j >$$

We will get

$$< \sum_{\alpha} n_{b\alpha} \sum_{\beta} n_{d\beta} > < \sum_{i,\alpha} a_{i\alpha} \sum_{j,\beta} c_{j\beta} > = \begin{cases} \sum_{i,\gamma} < n_{b\gamma} n_{d\gamma} > < a_{i\gamma} c_{i\gamma} >, & \text{mode overlap} \\ 0, & \text{no mode overlap} \end{cases}$$

where γ is the set of overlapping modes. Notice however, that there is only one sum over modes for the cross terms, so we can't just form two band-powers. Instead, we can see that for each i we effectively have

$$\sum_{\gamma} < n_{b\gamma} n_{d\gamma} > < a_{i\gamma} c_{i\gamma} > = \frac{C_l^{a_i, c_i} C_l^{n_b, n_d}}{\sqrt{k_l^{a_i, c_i} k_l^{n_b, n_d}}} = \frac{1}{2} \sqrt{\text{Var}(a_i c_i) \text{Var}(n_b n_d)}$$

where we have used that k is the number of degrees of freedom, and is given by

$$k_l^{a_i, c_i} = 2 \frac{(C_l^{a_i, c_i})^2}{\text{Var}(a_i c_i)}$$

So then, the entire result will look like

$$\begin{aligned}
& \sum_i Cov\{\sum_\alpha a_{i\alpha}b_{i\alpha}, \sum_\beta c_{i\beta}d_{i\beta}\} + Cov\{\sum_\alpha n_{a\alpha}n_{b\alpha}, \sum_\beta n_{c\beta}n_{d\beta}\} + \sum_{i,\gamma} \langle n_{b\gamma}n_{d\gamma} \rangle \langle a_{i\gamma}c_{i\gamma} \rangle + \\
& + \sum_{i,\gamma} \langle n_{b\gamma}n_{c\gamma} \rangle \langle a_{i\gamma}d_{i\gamma} \rangle + \sum_{i,\gamma} \langle n_{a\gamma}n_{d\gamma} \rangle \langle b_{i\gamma}c_{i\gamma} \rangle + \sum_{i,\gamma} \langle n_{a\gamma}n_{c\gamma} \rangle \langle b_{i\gamma}d_{i\gamma} \rangle + \\
& + \sum_{i \neq j, \gamma} \langle a_{i\gamma}c_{i\gamma} \rangle \langle b_{j\gamma}d_{j\gamma} \rangle + \sum_{i \neq j, \gamma} \langle a_{i\gamma}d_{i\gamma} \rangle \langle b_{j\gamma}c_{j\gamma} \rangle
\end{aligned} \tag{B.9}$$

Or dropping the Greek letter indices, and taking a combination of Roman letters to mean a bandpower (rather than Gaussian modes), we have:

$$\begin{aligned}
& \sum_i Cov\{a_i b_i, c_i d_i\} + Cov\{n_a n_b, n_c n_d\} + \\
& + \frac{1}{2} \sum_i \sqrt{Var(a_i c_i) Var(n_b n_d)} + \frac{1}{2} \sum_i \sqrt{Var(a_i d_i) Var(n_b n_c)} + \\
& + \frac{1}{2} \sum_i \sqrt{Var(b_i c_i) Var(n_a n_d)} + \frac{1}{2} \sum_i \sqrt{Var(b_i d_i) Var(n_a n_c)} + \\
& + \frac{1}{2} \sum_{i \neq j} \sqrt{Var(a_i c_i) Var(b_j d_j)} + \frac{1}{2} \sum_{i \neq j} \sqrt{Var(a_i d_i) Var(b_j c_j)}
\end{aligned}$$

So how do we compute this? We have signal simulations for a particular model, so a simple way to get these signal components for other models is to rescale the signal simulations in each ℓ bin. An example of this scaling for one of the terms is presented below

$$Var(a_i b_i) = \left(\frac{\langle a_i b_i \rangle}{\langle a_{simulations} b_{simulations} \rangle} \right)^2 Var(a_{simulations} b_{simulations})$$

Adopting the shorthand $a_{simulations} = a_s$ and the scaling introduced above, the final result looks like

$$\begin{aligned}
& Cov\{a_s b_s, c_s d_s\} \sum_i \frac{\langle a_i b_i \rangle \langle c_i d_i \rangle}{\langle a_s b_s \rangle \langle c_s d_s \rangle} + Cov\{n_a n_b, n_c n_d\} + \\
& + \frac{1}{2} \sqrt{Var(a_s c_s) Var(n_b n_d)} \sum_i \frac{\langle a_i c_i \rangle}{\langle a_s c_s \rangle} + \frac{1}{2} \sqrt{Var(a_s d_s) Var(n_b n_c)} \sum_i \frac{\langle a_i d_i \rangle}{\langle a_s d_s \rangle} + \\
& + \frac{1}{2} \sqrt{Var(b_s c_s) Var(n_a n_d)} \sum_i \frac{\langle b_i c_i \rangle}{\langle b_s c_s \rangle} + \frac{1}{2} \sqrt{Var(b_s d_s) Var(n_a n_c)} \sum_i \frac{\langle b_i d_i \rangle}{\langle b_s d_s \rangle} + \\
& + \frac{1}{2} \sqrt{Var(a_s c_s) Var(b_s d_s)} \sum_{i \neq j} \frac{\langle a_i c_i \rangle \langle b_j d_j \rangle}{\langle a_s c_s \rangle \langle b_s d_s \rangle} + \\
& + \frac{1}{2} \sqrt{Var(a_s d_s) Var(b_s c_s)} \sum_{i \neq j} \frac{\langle a_i d_i \rangle \langle b_j c_j \rangle}{\langle a_s d_s \rangle \langle b_s c_s \rangle}
\end{aligned}
\tag{B.10}$$

All the bandpowers with index i are obtained from *like_getexpvals.m* which takes model parameters and evaluates bandpower expectation values of the model for all specified fields, experiments, and ℓ bins. While all the signal and noise bandpowers, as well as variances and covariances of these (elements denoted by subscript s) are taken from the raw aps.

For each element of the BPCM, we exclude certain terms based on the fact that each of the $a, b, c, d, n_a, n_b, n_c, n_d$ are independent Gaussian fields with zero mean. For an explicit demonstration of this calculation of each element, please refer to the more extensive posting. For pedagogy, I will demonstrate the calculation of one of the elements here, namely $Cov\{B2 \times B2, B2 \times B1\}$. We take for our BPCM model the simulations fiducial model, which is just a (tensor+lensing) model and thus has only one signal component (CMB). Also, in this case, $a = b = d$. Hence:

$$\begin{aligned}
Cov\{B2 \times B2, B2 \times B1\} = Cov\{a_s a_s, a_s c_s\} &+ 0 + \frac{1}{2} \sqrt{Var(a_s c_s) Var(n_a n_a)} + 0 \\
&+ \frac{1}{2} \sqrt{Var(a_s c_s) Var(n_a n_a)} + 0 + 0 + 0
\end{aligned} \tag{B.11}$$

All the other elements are evaluated similarly.

B.3 VALIDATION

The implementation of the *semi-analytic* bandpower covariance matrix above is based on existing signal and noise simulations but can also be rescaled to reflect a multi-component model different from the one used to generate the simulations. This procedure involves one key simplifying assumption, which is that the various signal fields all have similar covariance structure, i.e., they populate similar Fourier modes as our CMB signal simulations. To make sure that everything is being calculated correctly, we generate a set of toy model simulated

bandpowers and contrast the direct covariance matrix construction to the semi-analytic one. The example we present here is one with three experiments (BICEP1 100 GHz, BICEP2 150 GHz, and WMAP7 23 GHz), and is the first application of this framework we have performed:

First, we outline the procedure for the BPCM construction for this particular case.

- We run a set of scripts to calculate $signal \times signal$, $noise \times noise$, and $signal \times noise$ angular power spectra (aps) for all experiment combinations. The aps generating scripts use *reduc_makeaps* but then combine and reorganize the results. In the end, we have one set of files that contain single experiment results (3 aps per file: $S \times S$, $N \times N$, and $S \times N$). A second set of files includes one file for each pairwise combination of experiments, with 10 aps per file. In this step, for this example, we specify the signal simulations to mean lensed- Λ CDM+r=0.2 because that is a good fit to the BICEP2 bandpowers. However, we are free to use any signal simulations as long as we record the model that describes the simulations. For instance, to test the rescaling, we also simulate bandpowers with lensed- Λ CDM+r=0.05 and $A_{dust} = 3.75 \mu K_{CMB}$.

An important note here is that we are using the BICEP2 signal simulations as a standard for WMAP signal simulations. This is justified to the extent that we believe the WMAP maps (after beam correction) include unbiased estimates of every mode on the sky. If that is the case, then the only signal filtering that occurs is what we apply with our reobserving matrix, and the BICEP2 signal simulations describe how the reobserva-

tion process affects sky signals.

- The next step is to calculate covariances from these aps. One goal of the semi-analytic construction is to avoid extra noise in the covariance matrix from terms that would average to zero in the large number of simulations limit, so we are selective about which covariances to calculate. The output of this process is a set of *six* covariance matrices.
 - Covariance matrix between *signal* \times *signal* bandpowers. For this matrix, we calculate and keep all terms.
 - Covariance matrix between *noise* \times *noise* bandpowers. If we assume that the noise is independent for all experiments, then we only keep terms that are the variance of noise bandpowers (i.e. $Cov(N_i \times N_i, N_i \times N_i)$ or $Cov(N_i \times N_j, N_i \times N_j)$ but not $Cov(N_i \times N_i, N_j \times N_j)$). However, we also allow for covariance across ℓ bins, so this matrix includes some other diagonals in addition to the main diagonal.
 - Four covariance matrices for *signal* \times *noise* bandpowers. These consist of terms with the form $Cov(S_i \times N_i, S_k \times N_l)$, $Cov(S_i \times N_j, N_k \times S_l)$, $Cov(N_i \times S_j, S_k \times N_l)$, and $Cov(N_i \times S_j, N_k \times N_l)$, respectively. These terms are non-zero only if the two noise terms come from the same experiment/field and the two signal terms are correlated (i.e., same field, but potentially different experiments). For the case where we include both E and B -mode signals, we assume that the E and B signal fields are uncorrelated, and the same for the E and B noise fields.

- With those six covariance matrices in hand, we can calculate the bandpower covariance matrix for an arbitrary model by applying appropriate scale factors and summing over signal components, as shown in Equation B.10. As a first step, here we calculate the semi-analytic bandpower covariance matrix for the same model that was used in the simulations (i.e., lensed- Λ CDM+r=0.2). This is done by simply summing the six components, without any rescaling. Then we use this matrix to rescale to the model which includes dust as well.

For the toy model simulated bandpowers, we assume degrees-of-freedom per band taken from the BICEP/*Keck* aps (but rounded to integer values). We simulate a single signal field, with variance equal to the bandpower expectation values for lensed- Λ CDM+r=0.2; all three experiments (B2, B1 100 GHz, and WMAP 23 GHz) see a perfectly correlated signal. Separately, we do the same for the model with dust. Each experiment gets its own independent noise field, with the same number of degrees-of-freedom and variance equal to the appropriate noise-only bandpowers.

The toy model also includes correlations between adjacent ℓ bins. To induce correlation, we generate a series of gaussian random numbers for each signal and noise component in each ℓ bin, then replace the first n_i numbers in each bin with the last n_i numbers from the previous bin, rescaled to the appropriate noise level and degrees of freedom of the new bin. Since we know that the real datasets have $\sim 10\%$ overlap between ℓ bins, we choose n_i using the formula $n_i = (\nu_i + \nu_{i+1})/2 \times 0.1$.




Figure B.1: Direct (left) and semi-analytic (right) correlation matrices with logarithmic color scale. The matrix is symmetric with 6 spectra times 5 ell bins; see the bottom left for the layout of each 6×6 block. For both cases, the correlations are calculated using the diagonal elements of the direct covariance matrix for normalization.

We then calculate a million realizations of the signal and noise fields, for each model, then derive bandpowers, and store them in files that can be read by the semi-analytic `bpcm` code. At the same time, we add the signal and noise fields and calculate traditional signal+noise bandpowers. These bandpowers are used to compute the *direct* bandpower covariance matrix.

Figure B.1 shows correlation matrices calculated from the direct and semi-analytic covariance matrices (both for the lensed- Λ CDM+r=0.2 model). To calculate a correlation matrix, we divide the i, j covariance element by the square root of $A_{ii}A_{jj}$ obtained from the direct bandpower covariance matrix (A). This means that the diagonal elements of the direct correlation matrix are equal to one by construction, but this is not exactly true for the semi-analytic correlation matrix. The color scale is logarithmic.

Figure B.2 shows the difference between the two correlation matrices, now on a linear




Figure B.2: Difference between direct and semi-analytic correlation matrices. Linear color scale. Both covariance matrices are derived for the lensed- Λ CDM+r=0.2 model.

color scale. There is no visible structure in the difference map for the correlation matrices.

Taking the standard deviation across all entries gives a value of $0.961e-3$, which is very close to the sample variance expected from 10^6 simulations: $1/\sqrt{10^6}$.

Figure B.3 shows a different set of correlation matrices. The direct covariance matrix is calculated for the model with lensed- Λ CDM+r=0.05 and $A_{dust} = 3.75 \mu K_{CMB}$. The semi-analytic covariance matrices is calculated from the same simulations as before (i.e., the lensed- Λ CDM+r=0.2 model) and rescaled to the new model.




Figure B.3: Direct (left) and semi-analytic (right) correlation matrices with logarithmic color scale. The direct covariance matrix is calculated for the model with lensed- Λ CDM+r=0.05 and $A_{dust} = 3.75 \mu K_{CMB}$. The semi-analytic covariance matrices is calculated from the same simulations as before (i.e., the lensed- Λ CDM+r=0.2 model) and rescaled to the new model.

Figure B.4 shows the difference of these two correlation matrices, on a linear color scale. Yet again, there is no visible structure in the difference map for the correlation matrices, and the entries are within the sample variance expected from 10^6 simulations: $1/\sqrt{10^6}$.

This is a conclusive test that both the original formulation of the semi-analytic matrix (without rescaling) and the rescaling process, is working as expected.




Figure B.4: Difference between direct and semi-analytic correlation matrices. Linear color scale. The direct covariance matrix is calculated for the model with lensed- Λ CDM+r=0.05 and $A_{dust} = 3.75 \mu\text{K}_{CMB}$. The semi-analytic covariance matrices is calculated from the same simulations as before (i.e., the lensed- Λ CDM+r=0.2 model) and rescaled to the new model.

References

- Abazajian, K. N., Adshead, P., Ahmed, Z., Allen, S. W., Alonso, D., Arnold, K. S., Bacigalupi, C., Bartlett, J. G., Battaglia, N., Benson, B. A., Bischoff, C. A., Borrill, J., Buza, V., Calabrese, E., Caldwell, R., Carlstrom, J. E., Chang, C. L., Crawford, T. M., Cyr-Racine, F.-Y., De Bernardis, F., de Haan, T., di Serego Alighieri, S., Dunkley, J., Dvorkin, C., Errard, J., Fabbian, G., Feeney, S., Ferraro, S., Filippini, J. P., Flauger, R., Fuller, G. M., Gluscevic, V., Green, D., Grin, D., Grohs, E., Henning, J. W., Hill, J. C., Hlozek, R., Holder, G., Holzapfel, W., Hu, W., Huffenberger, K. M., Keskitalo, R., Knox, L., Kosowsky, A., Kovac, J., Kovetz, E. D., Kuo, C.-L., Kusaka, A., Le Jeune, M., Lee, A. T., Lilley, M., Loverde, M., Madhavacheril, M. S., Mantz, A., Marsh, D. J. E., McMahon, J., Meerburg, P. D., Meyers, J., Miller, A. D., Munoz, J. B., Nguyen, H. N., Niemack, M. D., Peloso, M., Peloton, J., Pogosian, L., Pryke, C., Raveri, M., Reichardt, C. L., Rocha, G., Rotti, A., Schaan, E., Schmittfull, M. M., Scott, D., Sehgal, N., Shandera, S., Sherwin, B. D., Smith, T. L., Sorbo, L., Starkman, G. D., Story, K. T., van Engelen, A., Vieira, J. D., Watson, S., Whitehorn, N., & Kimmy Wu, W. L. (2016). CMB-S4 Science Book, First Edition. *ArXiv e-prints*.
- Abazajian, K. N. et al. (2016). CMB-S4 Science Book, First Edition.
- Adams, F. C., Bond, J. R., Freese, K., Frieman, J. A., & Olinto, A. V. (1993). Natural inflation: Particle physics models, power law spectra for large scale structure, and constraints from COBE. *Phys. Rev.*, D47, 426–455.
- BICEP2/Keck and Planck Collaborations (2015). Joint Analysis of BICEP2/Keck Array and Planck Data. *Physical Review Letters*, 114(10), 101301.
- Keck Array and BICEP2 Collaborations V (2015). BICEP2/Keck Array V: Measurements of B-mode Polarization at Degree Angular Scales and 150 GHz by the Keck Array. *Astrophys. J.*, 811, 126.
- Keck Array and BICEP2 Collaborations VI (2016). BICEP2/Keck Array VI: Improved Constraints on Cosmology and Foregrounds from BICEP2 and Keck Array Cosmic Mi-

- microwave Background Data with Inclusion of 95 GHz Band. *Physical Review Letters*, 116(3), 031302.
- Keck Array and BICEP2 Collaborations XI (2018). BICEP2/Keck Array XI: Beam Characterization and Temperature-to-Polarization Leakage in the BK15 Dataset.
- Keck Array and BICEP2 Collaborations VII (2016). BICEP2/Keck Array. VII. Matrix Based E/B Separation Applied to Bicep2 and the Keck Array. *Astrophys. J.*, 825, 66.
- Baumann, D. (2009). TASI Lectures on Inflation.
- Bennett, C., Larson, D., Weiland, J., Jarosik, N., Hinshaw, G., et al. (2012). Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results.
- Bennett, C. L., Larson, D., Weiland, J. L., Jarosik, N., Hinshaw, G., Odegard, N., Smith, K. M., Hill, R. S., Gold, B., Halpern, M., Komatsu, E., Nolte, M. R., Page, L., Spergel, D. N., Wollack, E., Dunkley, J., Kogut, A., Limon, M., Meyer, S. S., Tucker, G. S., & Wright, E. L. (2013). Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. *Astrophys. J. Suppl. Ser.*, 208, 20.
- Bennett, C. L., Smoot, G. F., Hinshaw, G., Wright, E. L., Kogut, A., de Amici, G., Meyer, S. S., Weiss, R., Wilkinson, D. T., Gulkis, S., Janssen, M., Boggess, N. W., Cheng, E. S., Hauser, M. G., Kelsall, T., Mather, J. C., Moseley, Jr., S. H., Murdock, T. L., & Silverberg, R. F. (1992). Preliminary separation of galactic and cosmic microwave emission for the COBE Differential Microwave Radiometer. *ApJ*, 396, L7–L12.
- BICEP2 and Keck Array Collaborations (2016). Improved Constraints on Cosmology and Foregrounds from BICEP2 and Keck Array Cosmic Microwave Background Data with Inclusion of 95 GHz Band. *Phys. Rev. Lett.*, 116, 031302.
- BICEP2 Collaboration, Keck Array Collaboration, Ade, P. A. R., Ahmed, Z., Aikin, R. W., Alexander, K. D., Barkats, D., Benton, S. J., Bischoff, C. A., Bock, J. J., Bowens-Rubin, R., Brevik, J. A., Buder, I., Bullock, E., Buza, V., Connors, J., Cornelison, J., Crill, B. P., Crumrine, M., Dierickx, M., Duband, L., Dvorkin, C., Filippini, J. P., Fliescher, S., Grayson, J., Hall, G., Halpern, M., Harrison, S., Hildebrandt, S. R., Hilton, G. C., Hui, H., Irwin, K. D., Kang, J., Karkare, K. S., Karpel, E., Kaufman, J. P., Keating, B. G., Kefeli, S., Kernasovskiy, S. A., Kovac, J. M., Kuo, C. L., Larsen, N. A., Lau, K.,

- Leitch, E. M., Lueker, M., Megerian, K. G., Moncelsi, L., Namikawa, T., Netterfield, C. B., Nguyen, H. T., O'Brient, R., Ogburn, R. W., Palladino, S., Pryke, C., Racine, B., Richter, S., Schillaci, A., Schwarz, R., Sheehy, C. D., Soliman, A., St. Germaine, T., Staniszewski, Z. K., Steinbach, B., Sudiwala, R. V., Teply, G. P., Thompson, K. L., Tolan, J. E., Tucker, C., Turner, A. D., Umiltà, C., Vieregg, A. G., Wandui, A., Weber, A. C., Wiebe, D. V., Willmert, J., Wong, C. L., Wu, W. L. K., Yang, H., Yoon, K. W., & Zhang, C. (2018). Constraints on Primordial Gravitational Waves Using Planck, WMAP, and New BICEP2/Keck Observations through the 2015 Season. *Physical Review Letters*, 121(22), 221301.
- BICEP2, Keck Array and Planck Collaborations (2015). Joint Analysis of BICEP2/Keck Array and Planck Data. *Phys. Rev. Lett.*, 114, 101301.
- Boyle, L. A., Steinhardt, P. J., & Turok, N. (2004). The Cosmic gravitational wave background in a cyclic universe. *Phys. Rev.*, D69, 127302.
- Buza, V. (2016). Fisher projections for cmb-s4, v3. http://users.physics.harvard.edu/~buza/20150331_fisher/.
- Buza, V. (2017). Bicep-array performance-based low-frequency forecasting study. http://bicep.rc.fas.harvard.edu/bkcmb/analysis_logbook/analysis/20170404_BA_low/.
- Carroll, S. (2003). *Spacetime and Geometry: An Introduction to General Relativity*. Pearson.
- Challinor, A. et al. (2017). Exploring cosmic origins with CORE: gravitational lensing of the CMB. *arXiv e-prints*.
- Chin Lin Wong (2014). *Beam Characterization and Systematics of the BICEP2 and Keck Array Cosmic Microwave Background Polarization Experiments*. PhD thesis, Harvard University.
- Choi, S. K. & Page, L. A. (2015). Polarized galactic synchrotron and dust emission and their correlation. *J. Cosmol. Astropart. Phys.*, 12, 020.
- BICEP2 Collaboration I (2014). BICEP2 I: Detection of B-Mode Polarization at Degree Angular Scales by BICEP2. *Physical Review Letters*, 112(24), 241101.

- BICEP2 Collaboration II (2014). BICEP2 II: Experiment and Three-Year Data Set. *Astrophys. J.*, 792, 62.
- BICEP2 Collaboration III (2015). BICEP2 III: Instrumental Systematics. *ArXiv e-prints*.
- Delabrouille, J., Betoule, M., Melin, J.-B., Miville-Deschênes, M.-A., Gonzalez-Nuevo, J., Le Jeune, M., Castex, G., de Zotti, G., Basak, S., Ashdown, M., Aumont, J., Baccigalupi, C., Banday, A. J., Bernard, J.-P., Bouchet, F. R., Clements, D. L., da Silva, A., Dickinson, C., Dodu, F., Dolag, K., et al. (2013). The pre-launch Planck Sky Model: a model of sky emission at submillimetre to centimetre wavelengths. *Astron. Astrophys.*, 553, A96.
- Dodelson, S. (2003). *Modern Cosmology*. Academic Press.
- Dunkley, J., Amblard, A., Baccigalupi, C., Betoule, M., Chuss, D., Cooray, A., Delabrouille, J., Dickinson, C., Dobler, G., Dotson, J., Eriksen, H. K., Finkbeiner, D., Fixsen, D., Fos-alba, P., Fraisse, A., Hirata, C., Kogut, A., Kristiansen, J., Lawrence, C., Magalhães, A. M., Miville-Deschenes, M. A., et al. (2009). Prospects for polarized foreground removal. *AIP Conf. Proc.*, 1141, 222–264.
- Einstein, A. (1917). Cosmological Considerations in the General Theory of Relativity. *Königlich Preußische Akademie der Wissenschaften*.
- Ellis, J., Nanopoulos, D. V., & Olive, K. A. (2013). No-Scale Supergravity Realization of the Starobinsky Model of Inflation. *Phys.Rev.Lett.*, 111, 111301.
- Eriksen, H. K., Banday, A. J., Gorski, K. M., & Lilje, P. B. (2004). Foreground removal by an internal linear combination method: Limitations and implications. *Astrophys. J.*, 612, 633–646.
- Fixsen, D. J. (2009). THE TEMPERATURE OF THE COSMIC MICROWAVE BACKGROUND. *The Astrophysical Journal*, 707(2), 916–920.
- Freese, K., Frieman, J. A., & Olinto, A. V. (1990). Natural inflation with pseudo - Nambu-Goldstone bosons. *Phys.Rev.Lett.*, 65, 3233–3236.
- Fuskeland, U., Wehus, I. K., Eriksen, H. K., & Næss, S. K. (2014). Spatial Variations in the Spectral Index of Polarized Synchrotron Emission in the 9 yr WMAP Sky Maps. *Astrophys. J.*, 790, 104.

- Ghosh, T., Boulanger, F., Martin, P. G., Bracco, A., Vansyngel, F., Aumont, J., Bock, J. J., Doré, O., Haud, U., Kalberla, P. M. W., & Serra, P. (2017). Modelling and simulation of large-scale polarized dust emission over the southern Galactic cap using the GASS Hi data. *Astron. Astrophys.*, 601, A71.
- Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. (2005). HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *Astrophys. J.*, 622, 759–771.
- Guth, A. H. (1981). The Inflationary Universe: A Possible Solution to the Horizon and Flatness Problems. *Phys.Rev.*, D23, 347–356.
- Hamimeche, S. & Lewis, A. (2008). Likelihood analysis of CMB temperature and polarization power spectra. *Phys. Rev. D*, 77(10), 103013.
- Hensley, B. (2015). *On the nature of interstellar grains*. PhD thesis, Princeton University.
- Hu, W. & Sugiyama, N. (1995). Toward understanding CMB anisotropies and their implications. *Phys. Rev.*, D51, 2599–2630.
- Hu, W. & White, M. (1997). A CMB polarization primer. , 2, 323–344.
- James Tolan (2014). *Testing Inflationary Cosmology with BICEP2 and the Keck Array*. PhD thesis, Stanford University.
- Kaiser, D. (2011). Cosmic inflation: An informal primer. MIT-CTP.
- Kallosh, R. & Linde, A. (2013). Universality Class in Conformal Inflation. *JCAP*, 1307, 002.
- Kamionkowski, M., Kosowsky, A., & Stebbins, A. (1997). A Probe of Primordial Gravity Waves and Vorticity. *Phys. Rev. Lett.*, 78, 2058–2061.
- Khoury, J., Ovrut, B. A., Steinhardt, P. J., & Turok, N. (2001). The Ekpyrotic universe: Colliding branes and the origin of the hot big bang. *Phys. Rev.*, D64, 123522.
- Kirit Karkare (2017). *Multifrequency Beam Systematics and Measurements of B-Mode Polarization with the BICEP/Keck Array Cosmic Microwave Background Experiments*. PhD thesis, Harvard University.

- Kovac, J. M., Leitch, E. M., Pryke, C., Carlstrom, J. E., Halverson, N. W., & Holzapfel, W. L. (2002). Detection of polarization in the cosmic microwave background using DASI. *Nature*, 420, 772–787.
- Krachmalnicoff, N., Carretti, E., Baccigalupi, C., Bernardi, G., Brown, S., Gaensler, B. M., Haverkorn, M., Kesteven, M., Perrotta, F., Poppi, S., & Staveley-Smith, L. (2018). The S-PASS view of polarized Galactic Synchrotron at 2.3 GHz as a contaminant to CMB observations. *ArXiv e-prints*.
- Krauss, L. M. & Wilczek, F. (2014). Using cosmology to establish the quantization of gravity. *Phys. Rev. D*, 89(4), 047501.
- Kritsuk, A. G., Ustyugov, S. D., & Norman, M. L. (2017). The structure and statistics of interstellar turbulence. *New Journal of Physics*, 19(6), 065003.
- Liddle, A. R. & Lyth, D. H. (2000). *Cosmological Inflation and Large-Scale Structure*. Cambridge University Press.
- McAllister, L., Silverstein, E., & Westphal, A. (2010). Gravity Waves and Linear Inflation from Axion Monodromy. *Phys.Rev.*, D82, 046003.
- Paine, S. (2017). The am atmospheric model (v. 9.2).
- Peebles, P. J. E., Page, Jr., L. A., & Partridge, R. B. (2009). *Finding the Big Bang*.
- Penzias, A. A. & Wilson, R. W. (1965). A Measurement of Excess Antenna Temperature at 4080 Mc/s. *Astrophys. J.*, 142, 419–421.
- Planck Collaboration (2017). Planck intermediate results. L. Evidence of spatial variation of the polarized thermal dust spectral energy distribution and implications for CMB B-mode analysis. *Astron. Astrophys.*, 599, A51.
- Planck Collaboration, Akrami, Y., Arroja, F., Ashdown, M., Aumont, J., Baccigalupi, C., Ballardini, M., Banday, A. J., Barreiro, R. B., Bartolo, N., Basak, S., Battye, R., Benabed, K., Bernard, J. P., Bersanelli, M., Bielewicz, P., Bock, J. J., Bond, J. R., Borrill, J., Bouchet, F. R., Boulanger, F., Bucher, M., Burigana, C., Butler, R. C., Calabrese, E., Cardoso, J. F., Carron, J., Casaponsa, B., Challinor, A., Chiang, H. C., Colombo, L. P. L., Combet, C., Contreras, D., Crill, B. P., Cuttaia, F., de Bernardis, P., de Zotti, G., Delabrouille, J., Delouis, J. M., Désert, F. X., Di Valentino, E., Dickinson, C.,

Diego, J. M., Donzelli, S., Doré, O., Douspis, M., Ducout, A., Dupac, X., Efstathiou, G., Elsner, F., Enßlin, T. A., Eriksen, H. K., Falgarone, E., Fantaye, Y., Fergusson, J., Fernandez-Cobos, R., Finelli, F., Forastieri, F., Frailis, M., Franceschi, E., Frolov, A., Galeotta, S., Galli, S., Ganga, K., Génova-Santos, R. T., Gerbino, M., Ghosh, T., González-Nuevo, J., Górski, K. M., Gratton, S., Gruppuso, A., Gudmundsson, J. E., Hamann, J., Handley, W., Hansen, F. K., Helou, G., Herranz, D., Hivon, E., Huang, Z., Jaffe, A. H., Jones, W. C., Karakci, A., Keihänen, E., Keskitalo, R., Kiiveri, K., Kim, J., Kisner, T. S., Knox, L., Krachmalnicoff, N., Kunz, M., Kurki-Suonio, H., Lagache, G., Lamarre, J. M., Langer, M., Lasenby, A., Lattanzi, M., Lawrence, C. R., Le Jeune, M., Leahy, J. P., Lesgourgues, J., Levrier, F., Lewis, A., Liguori, M., Lilje, P. B., Lilley, M., Lindholm, V., López-Caniego, M., Lubin, P. M., Ma, Y. Z., Macías-Pérez, J. F., Maggio, G., Maino, D., Mandolese, N., Mangilli, A., Marcos-Caballero, A., Maris, M., Martin, P. G., Martínez-González, E., Matarrese, S., Mauri, N., McEwen, J. D., Meerburg, P. D., Meinhold, P. R., Melchiorri, A., Mennella, A., Migliaccio, M., Millea, M., Mitra, S., Miville-Deschênes, M. A., Molinari, D., Moneti, A., Montier, L., Morgante, G., Moss, A., Mottet, S., Münchmeyer, M., Natoli, P., Nørgaard-Nielsen, H. U., Oxborrow, C. A., Pagano, L., Paoletti, D., Partridge, B., Patanchon, G., Pearson, T. J., Peel, M., Peiris, H. V., Perrotta, F., Pettorino, V., Piacentini, F., Polastri, L., Polenta, G., Puget, J. L., Rachen, J. P., Reinecke, M., Remazeilles, M., Renzi, A., Rocha, G., Rosset, C., Roudier, G., Rubiño-Martín, J. A., Ruiz-Granados, B., Salvati, L., Sandri, M., Savelainen, M., Scott, D., Shellard, E. P. S., Shiraishi, M., Sirignano, C., Sirri, G., Spencer, L. D., Sunyaev, R., Suur-Uski, A. S., Tauber, J. A., Tavagnacco, D., Tenti, M., Terenzi, L., Toffolatti, L., Tomasi, M., Trombetti, T., Valiviita, J., Van Tent, B., Vibert, L., Vielva, P., Villa, F., Vittorio, N., Wandelt, B. D., Wehus, I. K., White, M., White, S. D. M., Zaccari, A., & Zonca, A. (2018a). Planck 2018 results. I. Overview and the cosmological legacy of Planck. *arXiv e-prints*, (pp. arXiv:1807.06205).

Planck Collaboration, Akrami, Y., Arroja, F., Ashdown, M., Aumont, J., Baccigalupi, C., Ballardini, M., Banday, A. J., Barreiro, R. B., Bartolo, N., Basak, S., Benabed, K., Bernard, J. P., Bersanelli, M., Bielewicz, P., Bock, J. J., Bond, J. R., Borrill, J., Bouchet, F. R., Boulanger, F., Bucher, M., Burigana, C., Butler, R. C., Calabrese, E., Cardoso, J. F., Carron, J., Challinor, A., Chiang, H. C., Colombo, L. P. L., Combet, C., Contreras, D., Crill, B. P., Cuttaia, F., de Bernardis, P., de Zotti, G., Delabrouille, J., Delouis, J. M., Di Valentino, E., Diego, J. M., Donzelli, S., Doré, O., Douspis, M.,

Ducout, A., Dupac, X., Dusini, S., Efstathiou, G., Elsner, F., Enßlin, T. A., Eriksen, H. K., Fantaye, Y., Fergusson, J., Fernandez-Cobos, R., Finelli, F., Forastieri, F., Frailis, M., Franceschi, E., Frolov, A., Galeotta, S., Galli, S., Ganga, K., Gauthier, C., Génova-Santos, R. T., Gerbino, M., Ghosh, T., González-Nuevo, J., Górski, K. M., Gratton, S., Gruppuso, A., Gudmundsson, J. E., Hamann, J., Handley, W., Hansen, F. K., Herranz, D., Hivon, E., Hooper, D. C., Huang, Z., Jaffe, A. H., Jones, W. C., Keihänen, E., Keskitalo, R., Kiiveri, K., Kim, J., Kisner, T. S., Krachmalnicoff, N., Kunz, M., Kurki-Suonio, H., Lagache, G., Lamarre, J. M., Lasenby, A., Lattanzi, M., Lawrence, C. R., Le Jeune, M., Lesgourgues, J., Levrier, F., Lewis, A., Liguori, M., Lilje, P. B., Lindholm, V., López-Caniego, M., Lubin, P. M., Ma, Y. Z., Macías-Pérez, J. F., Maggio, G., Maino, D., Mandolesi, N., Mangilli, A., Marcos-Caballero, A., Maris, M., Martin, P. G., Martínez-González, E., Matarrese, S., Mauri, N., McEwen, J. D., Meerburg, P. D., Meinhold, P. R., Melchiorri, A., Mennella, A., Migliaccio, M., Mitra, S., Miville-Deschênes, M. A., Molinari, D., Moneti, A., Montier, L., Morgante, G., Moss, A., Münchmeyer, M., Natoli, P., Nørgaard-Nielsen, H. U., Pagano, L., Paoletti, D., Partridge, B., Patanchon, G., Peiris, H. V., Perrotta, F., Pettorino, V., Piacentini, F., Polastri, L., Polenta, G., Puget, J. L., Rachen, J. P., Reinecke, M., Remazeilles, M., Renzi, A., Rocha, G., Rosset, C., Roudier, G., Rubiño-Martín, J. A., Ruiz-Granados, B., Salvati, L., Sandri, M., Savelainen, M., Scott, D., Shellard, E. P. S., Shiraishi, M., Sirignano, C., Sirri, G., Spencer, L. D., Sunyaev, R., Suur-Uski, A. S., Tauber, J. A., Tavagnacco, D., Tenti, M., Toffolatti, L., Tomasi, M., Trombetti, T., Valiviita, J., Van Tent, B., Vielva, P., Villa, F., Vittorio, N., Wandelt, B. D., Wehus, I. K., White, S. D. M., Zacchei, A., Zibin, J. P., & Zonca, A. (2018b). Planck 2018 results. X. Constraints on inflation. *arXiv e-prints*, (pp. arXiv:1807.06211).

Planck Collaboration 2015 X (2016). Planck 2015 results. X. Diffuse component separation: Foreground maps. *Astron. Astrophys.*, 594, A10.

Planck Collaboration 2015 XII (2015). Planck 2015 results. XII. Full Focal Plane simulations. *ArXiv e-prints*.

Planck Collaboration 2015 XIII (2016). Planck 2015 results. XIII. Cosmological parameters. *Astron. Astrophys.*, 594, A13.

Planck Collaboration 2018 IV (2018). Planck 2018 results. IV. Diffuse component separation. *ArXiv e-prints*.

- Planck Collaboration 2018 VI (2018). Planck 2018 results. VI. Cosmological parameters. *ArXiv e-prints*.
- Planck Collaboration 2018 XI (2018). Planck 2018 results. XI. Polarized dust foregrounds. *ArXiv e-prints*.
- Planck Collaboration Int. L (2017). Planck intermediate results. L. Evidence of spatial variation of the polarized thermal dust spectral energy distribution and implications for CMB B-mode analysis. *Astron. Astrophys.*, 599, A51.
- Planck Collaboration Int. XXII (2015). Planck intermediate results. XXII. Frequency dependence of thermal emission from Galactic dust in intensity and polarization. *Astron. Astrophys.*, 576, A107.
- Planck Collaboration Int. XXX (2016). Planck intermediate results. XXX. The angular power spectrum of polarized dust emission at intermediate and high Galactic latitudes. *Astron. Astrophys.*, 586, A133.
- Planck Collaboration IV (2014). Planck 2013 results: LFI Beams. *Astron. Astrophys.*, 571, A4.
- Planck Collaboration IX (2014). Planck 2013 results: HFI spectral response. *Astron. Astrophys.*, 571, A9.
- Planck Collaboration VII (2014). Planck 2013 results: HFI time response and beams. *Astron. Astrophys.*, 571, A7.
- Planck Collaboration XLVI (2016). Planck intermediate results. XLVI. Reduction of large-scale systematic effects in HFI polarization maps and estimation of the reionization optical depth. *Astron. Astrophys.*, 596, A107.
- Planck Collaboration XVI (2014). Planck 2013 results. XVI. Cosmological parameters. *Astron. Astrophys.*, 571, A16.
- Sarah Kernasovskiy (2014). *MEASURING THE POLARIZATION OF THE COSMIC MICROWAVE BACKGROUND WITH THE KECK ARRAY AND BICEP2*. PhD thesis, Stanford University.
- Sheehy, C. & Slosar, A. (2017). No evidence for dust B-mode decorrelation in Planck data. *ArXiv e-prints*.

- Sheehy, C. & Slosar, A. (2018). No evidence for dust B-mode decorrelation in Planck data. *Phys. Rev. D*, 97(4), 043522.
- Silverstein, E. & Westphal, A. (2008). Monodromy in the CMB: Gravity Waves and String Inflation. *Phys.Rev.*, D78, 106003.
- Smith, K. M., Hanson, D., LoVerde, M., Hirata, C. M., & Zahn, O. (2012). Delensing CMB Polarization with External Datasets. *JCAP*, 1206, 014.
- Tegmark, M. (1997). How to make maps from CMB data without losing information. *Astrophys. J.*, 480, L87–L90.
- Thorne, B., Dunkley, J., Alonso, D., & Naess, S. (2016). The Python Sky Model: software for simulating the Galactic microwave sky. *arXiv e-prints*.
- van Engelen, A. et al. (2012). A measurement of gravitational lensing of the microwave background using South Pole Telescope data. *Astrophys. J.*, 756, 142.
- Vansyngel, F., Boulanger, F., Ghosh, T., Wandelt, B. D., Aumont, J., Bracco, A., Levrier, F., Martin, P. G., & Montier, L. (2017). Statistical simulations of the dust foreground to cosmic microwave background polarization. *Astron. Astrophys.*, 603, A62.
- Vansyngel, F., Boulanger, F., Ghosh, T., Wandelt, B. D., Aumont, J., Bracco, A., Levrier, F., Martin, P. G., & Montier, L. (2017). Statistical simulations of the dust foreground to cosmic microwave background polarization. *Astron. Astrophys.*, 603, A62.
- Weinberg, S. (2008). *Cosmology*. OUP Oxford.