
Precision Cosmology from Small-Scale CMB Observations

by

Hidde T. Jense

A Thesis submitted to Cardiff University
for the degree of Doctor of Philosophy

December 12, 2024

Abstract

Recent developments in cosmological observables are leading to an era of precision cosmology. At the forefront are observations of the Cosmic Microwave Background (CMB), which are capable of testing entire cosmological models by themselves. I present an overview of the development of CMB observations with legacy space-, and modern ground-based observatories, as well as the work done on the creation of new, modern Bayesian likelihood codes capable of constraining cosmological parameters from these observations. Additionally, developments in machine learning have allowed for accelerated inference by using neural network emulators. I present work on the development of a complete software suite with the `CosmoPower` emulator framework, opening new avenues for the prescription, creation, and application of new high-accuracy emulators that are applicable for these precise measurements. Finally, I bring all this work together in an overview of the work done on the upcoming data release (DR6) of the Atacama Cosmology Telescope (ACT). This work entails the development and testing of the ACT DR6 likelihood for both standard and extended cosmologies on a suite of purpose-built simulations. When published, the future ACT DR6 release will include small-scale temperature and polarization data of the CMB that will provide the most stringent cosmological tests from the CMB alone to date.

Publications

First Author Publications

- **H. T. Jense**, I. Harrison, A. Spurio Mancini et al., 2024; *A complete framework for cosmological emulation and inference with CosmoPower*, arXiv:2405.07903

Co-Author Publications

- F. J. Qu, K. M. Surrao, B. Bolliet, J. C. Hill, B. D. Sherwin and **H. T. Jense**, 2024; *Accelerated inference on accelerated cosmic expansion: New constraints on axion-like early dark energy with DESI BAO and ACT DR6 CMB lensing*, arXiv:2404.16805
- B. Bolliet, A. Spurio Mancini, J. C. Hill, M. S. Madhavacheril, **H. T. Jense**, E. Calabrese, and J. Dunkley 2023; *High-accuracy emulators for observables in Λ CDM, N_{eff} , Σm_ν , and w cosmologies*, arXiv:2303.01591
- Z. Li, T. Louis, E. Calabrese., and **H. T. Jense** et al. 2021; *The Simons Observatory: a new open-source power spectrum pipeline applied to the Planck Legacy Data*, arXiv:2112.13839

Other Publications

- S. Giardiello et al. 2024; *The Simons Observatory: impact of bandpass, polarization angle and calibration uncertainties on small-scale power spectrum analysis*, arXiv:2403.05242
- S. Shaikh, I. Harrison et al. 2023; *Cosmology from Cross-Correlation of ACT-DR4 CMB Lensing and DES-Y3 Cosmic Shear*, arXiv:2309.04412
- W. Coulton et al. 2023; *The Atacama Cosmology Telescope: High-resolution component-separated maps across one-third of the sky*, arXiv:2307.01258
- M. S. Madhavacheril et al. 2023; *The Atacama Cosmology Telescope: DR6 Gravitational Lensing Map and Cosmological Parameters*, arXiv:2304.05203
- F. J. Qu et al. 2023; *The Atacama Cosmology Telescope: A Measurement of the DR6 CMB Lensing Power Spectrum and its Implications for Structure Growth*, arXiv:2304.05202

Contents

Abstract	iii
Publications	v
1 Introduction and the Cosmic Microwave Background	1
1.1 A Cosmological History of the Universe	2
1.2 The Λ CDM model	6
1.3 The Cosmic Microwave Background	8
1.3.1 Initial Conditions	9
1.3.2 Decoupling, the Sound Horizon and Acoustic Peaks	11
1.3.3 Recombination and the Dark Ages	13
1.3.4 The Epoch of Reionization	14
1.3.5 Lensing of the CMB	16
1.3.6 Other Probes of Cosmology	17
1.4 The CMB Power Spectrum	19
1.4.1 Temperature Anisotropies	19
1.4.2 Polarization Anisotropies	20
1.4.3 Modeling the Power Spectrum	24
1.5 Measurements of the Power Spectrum	25
1.6 Testing the Λ CDM model with the CMB	28
1.6.1 Shape of the Universe	28
1.6.2 Non-flat Primordial Power Spectra	28
1.6.3 Neutrino Mass	29
1.6.4 Number of Relativistic Species	29
1.6.5 Dark Energy	30
1.7 Leading experiments for CMB Measurements	30
1.7.1 Early Space and Balloon Missions	30
1.7.2 Precision Cosmology from Space	31
1.7.3 Precision Cosmology from the Ground	34
2 Cosmology from the CMB	41
2.1 General Likelihood Implementation	43
2.1.1 Likelihood Function	43
2.1.2 The Data Vector	43
2.1.3 Binning	45
2.2 Full Spectrum Model	46

2.2.1	Astrophysical Foreground Model	46
2.2.2	Instrumental Nuisance Parameters	51
2.3	Markov Chain Monte Carlo Sampling	53
2.3.1	Maximum Likelihood Estimation and Model Selection	57
2.4	Current Constraints on Λ CDM	58
3	Likelihood Software for new CMB Analyses	61
3.1	Likelihood Design	63
3.2	A Re-implementation of the <i>WMAP</i> 9-Year Likelihood	65
3.3	Multi-Frequency Spectra Likelihoods	69
3.3.1	An Overview of the <code>MFLike</code> likelihood	69
3.4	<code>MFLike</code> applied to <i>Planck</i> legacy data	72
3.4.1	Likelihood Function	73
3.4.2	Tests of Parameters	75
3.5	<code>MFLike</code> Development and Validation for the Simons Observatory	77
3.5.1	Using <code>MFLike</code> on SO simulations	80
3.6	A CMB-only Likelihood	80
3.6.1	The Likelihood for the CMB-only dataset	85
3.6.2	Recovery of fiducial CMB bandpowers and cosmology	85
4	Emulators for High-Precision Cosmology	89
4.1	Limitations of Conventional Cosmological Codes	90
4.2	Cosmological Emulators	91
4.2.1	Neural Network design	91
4.2.2	Reusability	94
4.3	Implementation	95
4.3.1	Emulated Quantities	96
4.3.2	Cosmological Models	100
4.3.3	Training Data	102
4.3.4	Network Design and Training	103
4.3.5	Principal Component Analysis	104
4.3.6	Accuracy of Emulated Observables	106
4.4	Packaging Prescription	108
4.4.1	Generating Training Data	113
4.4.2	Choices for Emulator design and Training	117
4.5	Comparison of Recovered Cosmology	117
4.5.1	Simulated data vectors	117
4.5.2	Results	119
5	New Results from the <i>Atacama Cosmology Telescope</i>	123
5.1	The <i>Atacama Cosmology Telescope</i>	124
5.2	The DR6 Power Spectra	124
5.2.1	Cross Spectra	124
5.2.2	Theory Model	127
5.2.3	Inclusion of external datasets	130
5.3	Likelihood Validation	131

5.3.1	Simulated Data Vector	131
5.3.2	Recovery of Fiducial Parameters	132
5.4	Parameter Null Tests on Λ CDM	132
5.4.1	Null tests for Beyond- Λ CDM	134
5.5	Forecasts for Constraints on Cosmological Parameters	136
5.5.1	Constraints on the Λ CDM Model	136
5.5.2	Constraints on Beyond- Λ CDM Models	139
6	Conclusion	141
6.1	Thesis Overview	141
6.2	Developments in Cosmological Emulator Frameworks	141
6.3	Future Cosmological Constraints from CMB Experiments	142

2 Oboes

2 Clarinets in A

1 English Horn

1st & 2nd Bassoon

Violoncellos

pp

p

p


p

p

pp

The image shows a musical score for five instruments: 2 Oboes, 2 Clarinets in A, 1 English Horn, 1st & 2nd Bassoon, and Violoncellos. The score is written in 6/8 time and consists of five staves. The Oboe and Clarinet parts have dynamic markings of *p* (piano). The English Horn part has dynamic markings of *p*. The Bassoon part has dynamic markings of *p*. The Violoncello part has dynamic markings of *pp* (pianissimo). The score includes various musical notations such as notes, rests, and slurs.

Introduction and the Cosmic Microwave Background



In this introduction, I provide an overview of the necessary physics of background cosmology, the Cosmic Microwave Background (CMB), the standard model of cosmology (Λ CDM), and measurements of Λ CDM with the CMB.

The structure of this chapter is as follows:

- In section 1.1, I will give a broad overview of non-perturbed background cosmology and the expansion history of the universe;
- In section 1.2, I explain the Λ CDM model that is currently considered the “standard model” of cosmology;
- In section 1.3, I explain how this Λ CDM model gives rise to the Cosmic Microwave Background (CMB), and how the history of the universe is imprinted in features of the CMB;
- In section 1.4, I explain the power spectrum of the CMB, which is the main probe we aim to measure with experiments;
- In section 1.5, I give an idea of what goes into measuring this power spectrum;

- In section 1.6, I give an overview of extensions of the Λ CDM model, and what is currently at the forefront of CMB research;
- In section 1.7, I give an overview of important past, present, and future experiments that observe(d) the CMB.

The base material has been sourced from a variety of resources, primarily from

- Daniel Baumann’s lecture notes on *Advanced Cosmology*, sourced from his webpage <http://cosmology.amsterdam/education/advanced-cosmology/>, last opened August 2024;
- Scott Dodelson, *Modern Cosmology, Second Edition*, 2020, Academic Press.

with any further references given in the text.

1.1 A Cosmological History of the Universe

The entire history of our universe begins, of course, with Albert Einstein. In 1916 he published his *Theory of General Relativity* (GR), which forms the basis for our best understanding of gravitation. In it, gravity is described not as an attraction between masses, but by a curvature of the spacetime through which objects move as a consequence of mass.

The relation between mass-energy and curvature is summarized in the *Einstein field equations*,

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu}. \quad (1.1)$$

Here, $R_{\mu\nu}$ is the Ricci tensor, and $R = R^\mu{}_\mu$ is its contraction, the Ricci scalar. $g_{\mu\nu}$ is the metric tensor, and $T_{\mu\nu}$ the stress-energy tensor. The scalar Λ is the vacuum energy density, which appears in the equations as a degree of freedom.

In the years that followed Einstein’s discovery, several people, notably Friedmann, Lemaître, Robertson, and Walker, independently worked on a solution to the field equations of GR that we now know as the *FLRW metric*

$$ds^2 = -c^2 dt^2 + a(t)^2 \left[\frac{dr^2}{1 - \kappa r^2} + r^2 d\theta^2 + r^2 \sin^2 \theta d\phi^2 \right]. \quad (1.2)$$

In this FLRW metric, space can homogeneously expand or contract with a **scale factor** $a(t)$. Spacetime also has a **curvature** $\kappa = +1, 0$ or -1 , where these three values refer to an *open, flat* or *closed* spatial geometry respectively. This metric only describes

the curvature of spacetime. For us to inspect any properties of spacetime, and to learn how it affects objects and their kinematics, we need to describe the right hand side of Equation 1.1 by means of a stress-energy tensor

$$T_{\mu\nu} = (\rho + p)u_\mu u_\nu - pg_{\mu\nu}. \quad (1.3)$$

Here, we have made the assumption that our Universe is filled by a perfect fluid with a density ρ and an isotropic pressure p , and u^μ is the four-velocity of an observer moving relative to the fluid. For an observer comoving with the fluid, $u^\mu = (1, 0, 0, 0)$, and the fluid only retains its diagonal term

$$T_{\mu\nu} = \text{diag}(-\rho, \vec{p}). \quad (1.4)$$

We can now combine our metric from Equation 1.2 and our stress-energy tensor Equation 1.4 into our field equations Equation 1.1 and find

$$H^2 \equiv \frac{\dot{a}^2}{a^2} = \frac{8\pi G\rho}{3} - \frac{\kappa}{a^2} \quad (1.5)$$

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

These equations describe the **Hubble parameter** H , which is the time evolution of the scale parameter a as a function of the **density** ρ , **pressure** p , and curvature κ of the universe. The parameter H measures the rate at which two comoving points are receding from each other relative to their comoving distance. This is often expressed as *Hubble's law*, $v = Hd$, with v the recession velocity and d the physical distance between two comoving points. For a distant object, this recession velocity is often measured through its redshift z . In cosmology, this redshift is not necessarily due to a true velocity, and thus we define cosmological redshift in terms of the scale factor as $z + 1 \equiv 1/a$.

Since general solutions for Equation 1.5 are hard to find, one often looks at the limiting behaviour of these equations as a function of the relation between the density and pressure of the contents of the universe. A common way to do this is to set $\kappa = 0$, and then describe the **equation of state** for these contents as:

$$p = w\rho \quad (1.7)$$

Where w is the **parameter of state**, usually taken to be a constant, that equates the relation between pressure and density of the contents of the universe. Assuming a single component with a known parameter of state, solving the Friedmann equations

gives a simple expression for the scale factor,

$$a(t) = a_0 t^{\frac{2}{3(w+1)}}, \quad (1.8)$$

where a_0 is a constant of proportionality, commonly taken such that $a_0 = a(t_0) = 1$, where t_0 is present time*.

As the universe is filled with many different components, it is useful to look at a variety of important ones, and how their properties affect the expansion rate. Such components can primarily be classified by how their pressure relates to their density.

Matter is the component of all non-relativistic, massive particles, for which the pressure is negligible compared to their density, i.e. $p = 0$ and hence $w = 0$. Substituting this into Equation 1.8, we get

$$a(t) = a_0 t^{2/3}. \quad (1.9)$$

Radiation is the component of both relativistic massive particles that are dominated by their kinetic energy, or massless particles like photons. Their pressure is a third of their density, so $w = 1/3$, and

$$a(t) = a_0 t^{1/2}. \quad (1.10)$$

Then there is **dark energy**, which follow from the special case where $w = -1$, or where $p = -\rho$. For this value, this solution breaks down and one needs to solve the Friedmann equations directly to find that

$$a(t) \sim e^t \quad (1.11)$$

Such behaviour is usually associated with the dark energy component, denoted with the symbol Λ . This component is usually understood as an intrinsic energy density of space (also known as “vacuum energy”), that causes space to expand over time.

The behaviour of all these different models is shown together in Figure 1.1.

If applied to our Universe, however, one finds that a single component is not quite accurate enough, as we already know of the existence of at least two components in our Universe (those being matter and radiation). However, one can write down their solution as a linear combination of different solution to Equation 1.8, and solve the time evolution for the scale factor for a universe with multiple components. To do this, it is easier to rewrite Equation 1.5 by taking:

*Throughout this work, a subscript zero will always be used for “at present time”, e.g. X_0 is “the value of X at present time.”

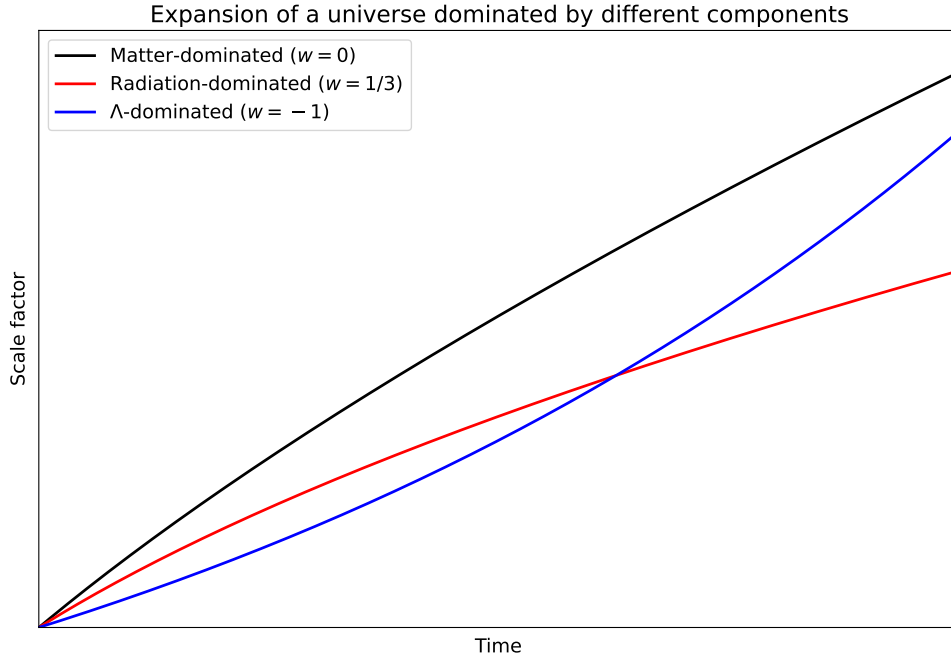


Figure 1.1. The expansion rate of a universe dominated by **(black)** matter ($a \propto t^{2/3}$), **(red)** radiation ($a \propto t^{1/2}$), and **(blue)** a cosmological constant ($a \propto e^t$).

$$\rho_c = \frac{3H^2}{8\pi G} \quad (1.12)$$

And observing that if $\rho = \rho_c$, $\kappa = 0$. This density measure ρ_c is known as the **critical density** and it the density for which a universe with flat geometry is expanding at a specific rate H . In fact, it becomes easy to rescale our density with respect to the critical density

$$\Omega \equiv \frac{\rho}{\rho_c} \quad (1.13)$$

We can now recognize that a flat universe is a universe at which $\Omega = 1$ — namely a universe at which the density is exactly the critical density and thus has flat spatial curvature. With this we can write the first Friedmann equation more simply as:

$$\frac{H^2}{H_0^2} = \Omega_{0,r}a^{-4} + \Omega_{0,m}a^{-3} + \Omega_{0,\Lambda} + \Omega_{0,\kappa}a^{-2} \quad (1.14)$$

Where $\Omega_{0,i}$ is the present-day density of a specific component, with r for radiation, m for matter, and Λ for dark energy. There is also a term with $\Omega_{0,\kappa} \equiv 1 - \Omega_{0,r}$, which is the density of curvature, which for a flat universe would be zero and disappear as a term. By taking a known set of $\Omega_{0,i}$ parameters, one can integrate this equation over time and follow the expansion history of the universe according to GR. This is shown in




Figure 1.2. The expansion history for a flat universe with $\Omega_{m,0} = 0.3$, $\Omega_{\Lambda,0} = 0.7$ and $\Omega_{r,0} \sim 10^{-4}$ (these numbers were chosen to approximate our best measurement for our own Universe), calculated by numerically integrating Equation 1.14. The **top** figure shows the scale factor as a function of time, while the **bottom** figure shows the density fraction of the various components over time. The vertical dashed lines show two epochs, first where $\Omega_r = \Omega_m$ (known as *matter-radiation equality*), and later where $\Omega_m = \Omega_\Lambda$. In these three regimes, the scale factor expands as in Equation 1.8 and Equation 1.11.

Figure 1.2.

Because these different components dominate the universe at different stages, the evolution of the universe will be different throughout its history. The parameterization of Equation 1.14 also only gives an overview of the *average* density across the universe, which assumes the universe is perfectly homogenous. In the next section, section 1.2, I will go through our best understanding of this evolution in more detail, and give a qualitative overview of how inhomogeneities behave over time.

1.2 The Λ CDM model

In 1927, Georges Lemaître independently derived the FLRW metric and, coupled with observations of the recession of distant galaxies, proposed an expansion of the Universe (Lemaître, 1927). More importantly, Lemaître argued that since the Universe is expanding, it must have expanded from a smaller state in the past. This proposal led to Edwin Hubble's observations in 1929 where he showed that distant galaxies are indeed receding, at a rate known as the **Hubble Parameter**, which at present is measured at

$H_0 \approx 67 - 72$ km/s/Mpc. In 1964, the discovery of the cosmic microwave background or CMB (Penzias and Wilson, 1965) – which I will describe in more detail in section 1.3 – provided evidence of an initial hot, dense state of the Universe, which expanded over time into its current distribution of galaxies within a cosmic web of large scale filaments.

Observations of the CMB showed that the Universe was incredibly homogeneous and smooth early on. If small perturbations exist, then these will over time accumulate more mass and grow into the large-scale structure of the Universe. Early observations showed however, that the Universe at present day is highly clustered, while the early-Universe CMB is far too smooth to seed the perturbations that can grow into this clustered structure. The inability at the time to relate the relatively homogeneous nature of the CMB with the highly clustered nature of matter at present day gave rise to the idea that cold dark matter (CDM) has to dominate the mass component of the Universe. In addition, the relatively flat shape of spacetime and the apparent thermal equilibrium of the early Universe gave rise to the formulation of the cosmic inflation paradigm (see subsection 1.3.1).

These observations led to the present-day concordance model for cosmology[†], known as the **Λ CDM model**. This model makes only a few basic assumptions:

- General Relativity is a correct description of gravity on cosmological scales;
- The Universe is homogeneous and isotropic on cosmological scales;
- There is a dark matter component that was *cold* at matter-radiation equality (cold dark matter, or CDM);
- There is a component driving the accelerated expansion of the Universe similar to a *cosmological constant* Λ .

Λ CDM cosmology allows us to summarize the entire cosmological history of the Universe in only six basic parameters, which is the number needed to fit to cosmological observables. While more free parameters, such as the spatial curvature, or mass of neutrinos, can be introduced as additional degrees of freedom, measurements of these parameters have not yet yielded any deviations from basic assumptions or external measurements. For these two given examples, we have not detected a deviation from the simple assumption of flat spatial curvature, nor have we made a nonzero measurement of the mass of neutrinos from cosmological probes. There is no one set of parameters

[†]It is called a “concordance” model because, when it was devised, it reached a consensus between the cosmology described by these different observational probes.

that must be used to fit for — instead, there is some freedom of choice, and the remaining parameters can be exactly derived from the original six.

A common set of six parameters chosen are:

1. The Hubble parameter H_0 ;
2. The present-day physical density of baryonic matter $\Omega_{b,0}h^2$ and cold dark matter $\Omega_{c,0}h^2$, where $h \equiv H_0/100 \text{ km/s/Mpc}$ is a normalized present-day Hubble parameter;
3. The amplitude A_s and spectral index n_s of the primordial power spectrum;
4. The optical depth at reionization τ_{reio} .

These first three parameters have been described above already. The latter three will be described more in detail in subsection 1.3.1 and subsection 1.3.4, later on. An additional parameter, θ_* , will be mentioned in subsection 1.3.2. Because this parameter can be fully derived from the other six, it is not a free parameter, but it is often used in analysis instead of the Hubble parameter H_0 .

1.3 The Cosmic Microwave Background

The **Cosmic Microwave Background** or **CMB** is one of the key observables in cosmology. It was first predicted to exist in the 1940s by Alpher and Herman (Alpher, 1948). Based on the assumption that the Universe is expanding, and that the Universe expanded in the past from a smaller state, we can argue that the Universe must have been denser and hotter in the past. In the distant past, the density of matter must have been so great that any heat in the form of radiation could not dissipate due to optical depth being far greater than unity. At some point however, the expansion of the Universe caused the density of matter to drop so low that optical depth dropped far enough for this radiation to radiate out. This is the first epoch at which the Universe became transparent to radiation — at any earlier point in time the Universe must have been opaque to radiation and we cannot make any observations from photons. Looking back at this point in the Universe's history shows an all-sky background of photons, radiating away from an opaque layer of matter: this primordial radiation is what we call the CMB.

The CMB was detected first in 1965 by Arno Penzias and Robert Wilson, during attempts to remove any sources of antenna noise for long-distance radio communications (Penzias and Wilson, 1965). It was originally formed as the thermal emission of the

primordial plasma, with the photons last interacting with this plasma via Thomson scattering. The point at which these photons last scattered is called the **surface of last scattering**, which is observed at present day to be around $z_* = 1090.97 \pm 0.89$ (Bennett et al., 2013).

At present, the CMB radiation forms an all-sky black-body emission at a temperature $T_{\text{CMB}} = 2.755 \text{ K}$ (Fixsen, 2009), which corresponds to a temperature around decoupling of $T_* = T_{\text{CMB}}(1 + z_*) \approx 3000 \text{ K}$. At this observed present-day temperature, the black-body emission peaks around a frequency of 162 GHz, or a wavelength of 8 millimetres. Coincidentally and usefully, this peak is right at the transition point where Galactic synchrotron radiation (which dominates at lower frequencies) and Galactic thermal dust emission (which dominates at higher frequencies) cross-over, and the sum of their emission is at a minimum. Thus Galactic contamination is at a minimum in the range where the CMB is brightest. Additionally, at this wavelength, the Earth's atmosphere is highly transparent, with a gap between the absorption spectra of water vapor and carbon dioxide. This makes it possible to observe the CMB from the surface of the Earth.

Despite any assumptions we have made of homogeneity and isotropy, the universe itself is not perfectly homogeneous. Any form of primordial inhomogeneity causes tiny fluctuations in the primordial plasma, and thus the CMB forms an all-sky field of fluctuations in temperature and polarization. These fluctuations of the order of $\Delta T_{\text{CMB}} \leq \pm 300 \mu\text{K}$ are measurable across the sky, meaning inhomogeneities existed of the order of $\Delta T/T \sim O(10^{-5})$. An all-sky measurement of these anisotropies as measured by the *Planck* satellite is shown in Figure 1.3. In addition to this temperature anisotropy component, the CMB also has a *polarized* component, which too inhibits its own anisotropies. I will delve into more details about this in subsection 1.4.2.

To explain how primordial features and cosmological physics affected the CMB, both at and after the surface of last scattering, I will summarize the thermal history of the universe in this section.

1.3.1 Initial Conditions

The Λ CDM model assumes that the universe started existing as a tiny, hot state of matter, radiation, and other components. The state must have been small enough that the entire universe was a quantum system, and any later fluctuations in density must have been seeded by the quantum fluctuations in this field. This was immediately followed, within about 10^{-30} seconds of this state's existence, of a period of rapid, exponential expansion during which the universe expanded by a factor




Figure 1.3. A map of the all-sky CMB temperature anisotropy as measured by *Planck* (Planck Collaboration, 2020a). Note that while the mean temperature $T_{\text{CMB}} = 2.756\text{K}$, the amplitude of the measured fluctuations are no more than $\Delta T_{\text{CMB}} = \pm 300\mu\text{K}$, meaning the CMB is extremely homogeneous across the sky.

$\sim e^{60}$. This rapid expansion is known as **inflation**. In addition to creating the initial seeds for inhomogeneities, inflation also explains the relative flatness of the universe, $|\Omega_0 - 1| \ll 1$, as any overdensity got smoothed out over a larger volume.

As a result of inflation, these inhomogeneities were blown up to form a universe filled with overdensities - regions which contained more mass than the average density in the universe. The initial conditions follow a power spectrum of overdensities, often called the **primordial power spectrum**, parameterized as

$$P(k) = A_s \left(\frac{k}{k_*} \right)^{n_s - 1}, \quad (1.15)$$

where k is the wavenumber of an overdensity, k_* is an arbitrary **pivot scale**, A_s is the **fluctuation amplitude** at that pivot scale, and n_s is the **spectral index** of the power law. This power law describes the initial state of scalar (density) perturbations in the universe. In addition, the initial conditions should include some *tensor* perturbations in the form of some primordial gravitational waves. These tensor perturbations follow a similar power law, but with an initial *tensor* fluctuation amplitude A_t and spectral index n_t .

The origin of these Gaussian initial conditions is assumed to be the **inflation epoch** of the universe, during which the universe rapidly expanded from a small, hot, primordial plasma, which were Gaussian before reheating. The exact physics of inflation determine the possible values for A_s , n_s , A_t , and n_t .

Current measurements from the CMB anisotropies in temperature and polarization have given us good measurements of the scalar spectral index parameters A_s and n_s (as we will see in later chapters, namely chapter 3 and chapter 5). Because we do not have a good measurement of it yet, a value that is of interest in the context of cosmology, is the **tensor-to-scalar ratio** $r \equiv A_t/A_s$, which measures the amount of primordial tensor fluctuations relative to the amount of scalar fluctuations. Current constraints put this value at $r < 0.036$ (Ade and et al., 2021).

1.3.2 Decoupling, the Sound Horizon and Acoustic Peaks

Because the primordial power spectrum of overdensities is nonzero, the universe has some initial anisotropies in it, which will evolve over time. Overdense regions of dark matter will collapse under gravity, while overdense regions of radiation will attempt to diffuse away towards underdense regions. Baryons, which are attracted to gravity, will attempt to follow the dark matter into overdense regions, but so long as

the baryonic density is high enough, they will — unlike the dark matter — experience a radiation pressure that forces them outwards into underdense regions. Thus initially, these overdense pockets of baryons are forced outwards against gravitational infalling.

For a universe that is also changing in scale over time, the concept of a physical size measure needs to account for a time-varying distance measure. The **particle horizon** is the maximal distance that a signal can travel between some initial time and a later time:

$$r(t) = \int_0^t \frac{dt}{a(t)}. \quad (1.16)$$

In the early universe however, the high density of baryons causes photons to remain coupled to them, meaning a photon cannot travel at a speed faster than the sound speed c_s . As such, any overdensity in the primordial plasma grows in size as

$$r_s(t) = \int_0^t \frac{c_s(t)dt}{a(t)}, \quad (1.17)$$

where $c_s(t) = 1/\sqrt{3}$ while photons remain coupled to baryons. As these overdensities grow in physical size, the number density of baryons drops, and as a result photons can diffuse out of these overdensities, lowering the outward pressure and thus the sound speed $c_s(t) \lesssim 1/\sqrt{3}$. At the point where photons decouple from baryons, we are left with a shell-shaped overdensity of baryons.

In response to the disappearance of this thermal radiation pressure, these baryon overdensities will suddenly be left with only a gravitational pull inwards towards the dark matter overdensity in the middle, and start collapsing. Because they will contract again in physical size, this will increase the number density of baryons again, which means the radiation pressure increases, forcing the shell outwards again. The baryonic shell will oscillate around a size known as the **sound horizon** size, $r_* \equiv r_s(t_*)$, which is a universal quantity that can be computed as the scale of the sound horizon at the time t_* when radiation finally fully decouples from baryonic matter. A similar quantity, the sound horizon at drag epoch $r_d \equiv r_s(t_d)$, is the maximal distance that baryons can travel due to radiation pressure, which is slightly higher due to baryons traveling outward before they reach their largest scales under gravitational infalling. Baryons in the large scale structure (LSS) of the universe tend to exist in hollow, overdense shells with a radius r_d . These overdensities are known as **baryon acoustic oscillations** or BAO features, where baryons undergo oscillating periods of compression and rarefaction.

On the sky, fluctuations of a size r_* form patterns at an angular scale θ_* , known as the **angular scale at decoupling**. Because the exact size of θ_* is dependent on the background evolution of the primordial plasma, the value of θ_* can conversely be used as a measurement of the universe's cosmological properties.

1.3.3 Recombination and the Dark Ages

As the universe expands in size, the matter density drops because the same number of particles now occupy a larger volume, and at the same time the temperature drops because thermal energy similarly occupies a larger volume of space. Because of this, we reach the **epoch of recombination**[‡], where the temperature is low enough for the electrons and protons in the primordial plasma to combine and form neutral hydrogen.

In Figure 1.4, I show the reionization history of the universe as described by the Λ CDM model. The figure shows the redshift of the universe z , with high redshift being the early universe and zero redshift the present day, against the ionization fraction $\chi_e(z)$, being the number of free electrons per atomic nucleus. For a universe that consists fully of hydrogen, this number will vary between zero (no free electrons) and one (one free electron per hydrogen nucleus). However, Λ CDM assumes some brief early universe nucleosynthesis, and thus a small amount of primordial helium should exist. Because a helium atom contains two electrons per nucleus, we should see more than one electron per atomic nucleus if the universe is fully ionized. Observations indicate that the universe consisted for about 25% Helium by mass (Aver et al., 2015). In Figure 1.4 we see, that at very high redshift, the universe is completely ionized at $\chi_e \approx 1.17$. The epoch of recombination then happens between $10^3 \lesssim z_{\text{rec}} \lesssim 10^4$, where we see a stepwise drop in the ionization fraction: this is because the different atomic properties of helium and hydrogen mean that these two elements will recombine with their (in helium's case: first and second) electrons at different temperatures and thus at different redshifts. Once the universe has cooled down enough, all electrons will exist in neutral hydrogen and helium, and thus the ionization fraction $\chi_e = 0$ after recombination, as seen at $z \lesssim 10^3$.

The period following this epoch of recombination is called the **dark ages**, as it represents a period during which neutral hydrogen and helium filled the entire universe in a state where the gas was too warm and insufficiently clustered to form stellar objects. This period lasting from about $z \sim 10^3$ to the reionization period is poorly understood, as there was no physical interaction between the CMB photons and neutral gas. One possible observation for probing this moment in the universe's evolution would be the

[‡]although it is called *recombination*, matter was not in a combined prior to this epoch.

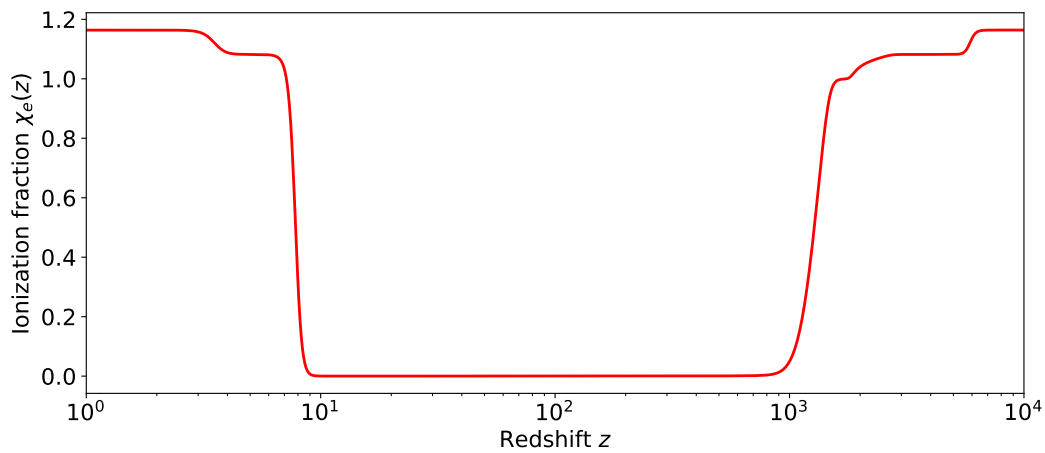


Figure 1.4. The ionization fraction of the universe χ_e as a function of redshift z . The universe was ionized at first (high redshift), then underwent a process of **recombination** (drop in ionization fraction at $z_{\text{rec}} \sim 1000$), only to be **reionized** at a later stage (increase in ionization fraction at $z_{\text{reio}} = 7.68 \pm 0.79$ (Planck Collaboration, 2020d)). The main decrease and later increase in ionization fraction is due to the ionization and later recombination of hydrogen, which forms the dominant atomic component in the universe. The complicated structure of these increases and decreases in ionization fraction is due to the existence of helium (J. N. Lockyer, Esq, 1868), which recombines/reionizes at different energies than hydrogen. These two different transition epochs are explained in subsection 1.3.3 and subsection 1.3.4 respectively.

21 cm hyperfine transition of hydrogen, which, when accounting for redshift, could be visible in the 15–200 MHz band on Earth.

1.3.4 The Epoch of Reionization

While the physics of recombination suggests the universe must have transitioned to a neutral state, current observations indicate the opposite is true, and that the universe is in a highly ionized state. This suggests that the universe must have undergone a second transition back from recombination to an ionized state.

By observing the absorption spectra of distant quasars, we can probe the redshift at which the universe transitioned from this neutral state back to an ionized state. Neutral hydrogen absorbs strongly in the Lyman- α regime, and thus the period between the epoch of recombination up to the present day must include a period of strong absorption of Ly α emission. This period of strong absorption is known as the *Gunn-Peterson trough*, and measurements indicate the transition to the current ionized state must have happened before $z \gtrsim 6$ (Becker et al., 2001). This transition from the early, neutral state to the current, ionized state is known as the **epoch of reionization**. This epoch is shown in Figure 1.4, as the increase in ionization fraction at $z \sim 8$. Because hydrogen and helium ionize at slightly different energy levels, we see a second increase in the ionization fraction at a slightly later $z \sim 5$, which is due to the second reionization of helium. After this, at $z \ll z_{\text{reio}}$, the universe is once again in a fully ionized state at $\chi_e \approx 1.17$.

Measurements of the Lyman- α forest in the spectrum of quasars puts an upper limit on the redshift when this transition must have occurred, but due to the highly sensitive nature of this measurement, it is hard to use this for a detailed measurement of the entire epoch.

For observations of the CMB, the exact physics of reionization are not of interest, since the impact they have on the exact shape of the CMB is minimal. What is of interest however, is *when* exactly reionization happened, and *how long* it lasted. This is often expressed in terms of the *redshift of reionization* z_{reio} as well as a duration Δz . Since the exact model for reionization does not matter much at current experimental sensitivity, it is commonly modeled as a simple transition from low to high ionization fraction

$$\chi_e(z) = \frac{1 + \tanh((z - z_{\text{reio}})/\Delta z)}{2}. \quad (1.18)$$

Current constraints on the epoch of reionization are limited, so often the parameter $\Delta z \approx 0.5$ is kept fixed at this somewhat arbitrary value. The CMB alone has a very weak dependence on this parameter Δz , and thus its value cannot be directly inferred

from it. However, as described later in this section, it will be possible to infer this from secondary anisotropies through the Sunyaev-Zel'dovich effect.

Since the universe is not fully transparent to radiation at this point, the CMB photons are slightly affected by the **optical depth at reionization** τ_{reio} , found by the *Planck* mission to be $\tau_{\text{reio}} = 0.0544 \pm 0.0073$. This optical depth is equivalent to a redshift $z_{\text{reio}} = 7.50 \pm 0.82$. Any increase in optical depth will suppress anisotropies in the CMB through absorption, meaning the overall amplitude of anisotropy measurements is inversely related to the optical depth at reionization.

It is believed that the epoch of reionization was caused by the formation of the first, massive, population of stars (see e.g. Zaroubi (2012); Klessen and Glover (2023)). As the neutral hydrogen gas clustered to form the initial large scale structure of the universe around $z \sim 8$, dense pockets of gas formed the first, high-mass stars.

The formation of these ionized pockets in the universe creates an additional secondary effect on the CMB known as the **Sunyaev-Zel'dovich effect** (SZ effect). CMB photons traveling along the line-of-sight through pockets of ionized gas will scatter off them through Compton scattering. With the ionized gas in the universe forming a complicated large-scale structure in the universe after clustering, we differentiate between the **thermal** (tSZ) and **kinematic** (kSZ) SZ effects.

The tSZ effect is created when a CMB photon is scattered by electrons in a thermal distribution within a galaxy cluster. The tSZ effect can be used to detect galaxy clusters from CMB map observations, study their gas profiles, and probe the baryon physics within them (see e.g. Hilton et al. (2021)).

Electrons outside a galaxy cluster can experience bulk velocity through gravitational infalling towards the cluster mass. This bulk motion can scatter the CMB and will generate the kSZ effect, which is a much smaller effect than the tSZ effect.

The nature of reionization physics and cluster formation, and their interaction, have an effect on the physical distribution of electrons through the large scale structure of the universe, and thus on the details of the kSZ model. The kSZ effect is often modeled in two components: an instantaneous *patchy* kSZ effect that is generated as the CMB photons travel through the bubbles of reionized material that form during reionization. Secondly, there is the *late-time* kSZ effect, that is formed by the CMB photons scattering off the bulk motion of electrons in the large scale structure after the epoch of reionization.

Because the kSZ effect is caused by the distribution of electrons, which itself is generated during the epoch of reionization, the physics of this epoch will leave an

imprint on the kSZ effect itself. Precision measurements of the kSZ effect will be able to probe reionization physics, and parameters such as the duration of reionization, Δz , or the clustering of matter σ_8 , better than observations of the CMB alone (e.g. MacCrann et al. (2024); Hadzhiyska et al. (2024)).

1.3.5 Lensing of the CMB

Because of the inhomogeneous distribution of mass throughout the universe, photons from the CMB are deflected by the presence of overdense regions such as galaxy clusters. This in effect creates a **lensing** effect of the CMB, where matter overdensities along the line of sight can warp an observer's view of the CMB. A correct model for the CMB must therefore include a model for the deflection of the CMB photons through lensing, and either include this in the final observation model, or include some paradigm to remove the lensing effects on the CMB.

In practice, the deflection of photons through lensing by the LSS is small, and lensing effects are often approximated to first order. Assuming we observe the sky in some direction \vec{n} , any mass near the line of sight will generate some small gravitational potential field $\phi(\vec{n})$, which causes a deflection of nearby photons, meaning we end up seeing the CMB that was really at position $\vec{n} + \nabla\phi(\vec{n})$. If we try to measure some unlensed function $F(\vec{n})$ along the sky, then, to first order, we will really see a *lensed* version F^L of this. This function F^L is to first order

$$F^L(\vec{n}) = F^u(\vec{n} + \nabla\phi(\vec{n})) \approx F^u(\vec{n}) + \nabla\phi(\vec{n}) \cdot \nabla F^u(\vec{n}) + O(\nabla^2\phi), \quad (1.19)$$

where F^u is the *unlensed* version of F . The value of the lensing potential $\phi(\vec{n})$ is critical in modeling and measuring the CMB. Because the second term on the right hand side is proportional to ∇F , any small-scale effect in the CMB (which changes more rapidly across the sky and thus has a higher value of ∇F) will be more affected by lensing effects. Thus, as we aim to make more precise measurements of small-scale features in the CMB, including the effects of lensing become more important.

In addition, as we will see later in subsection 1.4.2, any polarized component of the CMB will be affected by lensing, as lensing will warp the view of the Stokes Q and U modes.

The correct modeling of lensing involves a correct modeling of the clustering of matter across the universe. This involves predicting how matter is initially distributed, how this distribution evolves over cosmic time, and how this evolving distribution correlates with itself and with the CMB photons to deflect their paths towards the observer.

1.3.6 Other Probes of Cosmology

While for this thesis, the CMB is the main cosmological probe of interest, there are other observables that in one way or another constrain cosmological models and parameters. As the CMB probes the universe at high redshift ($z_{\text{CMB}} \approx 1050$), combinations of CMB experiments with probes that measure different redshifts can help break degeneracies that exist in each experiment separately, or constrain parameters in a different way that makes the joint constraint much tighter than the individual ones.

One of the more traditional measurements for cosmology is the use of *distance ladder measurements*, whereby the present-day Hubble parameter is measured by probing the redshift-distance relation $cz = H_0 d$ directly – relating the redshift z , Hubble parameter H_0 and distance d to an object. Key to this procedure is a method to accurately infer the distance to an object, for which the *cosmic distance ladder* is used: a stepwise procedure to measure the distance to nearby objects, using this to measure the distance to moderately far-away objects, which are then used to calibrate the measurement of distances to far-away objects. The *SH0ES* project uses the *Hubble Space Telescope* and *Gaia* to measure the distance to nearby Cepheid variable stars to calibrate the distance to supernovae in nearby galaxies, which are then used to measure the distance to supernovae in distant galaxies, out to a redshift of $z \leq 0.01$. The value reported in Riess et al. (2022), at $H_0 = (73.04 \pm 1.04)$ km/s/Mpc is famously in disagreement with the value found from CMB measurements by *Planck* (see later in Table 2.1), at a 5σ discrepancy between them. This disagreement, which persists between nearby measurements of the Hubble parameter being high, and high-redshift measurements being lower, is commonly known as the *Hubble tension* and is an open problem in cosmology.

Baryon Acoustic Oscillation (BAO) measurements involve probing the distribution of baryonic matter across distance and redshift. Experiments which probe the BAO scale aim to measure the distribution of galaxies, and their angular ($\Delta\theta$) and redshift (Δz) distances, comparing it to the drag scale r_d (see subsection 1.3.2) – since the latter is a fixed size, this gives a measurement of how large this drag scale appears on the sky as a function of its redshift, known as the angular diameter distance $D_L(z)$. Measuring these indirectly probes the evolution rate of the universe as a function of redshift, $H(z)$, which traces the background evolution of the universe. Measurements such as those from the recent *DESI* results (i.e. DESI Collaboration (2024)) constrain the matter density Ω_m and present-day expansion rate H_0 of the universe.

Another way to probe the matter distribution in the universe is to look at the dark matter distribution, or rather, means to infer it. Since the clustered nature of dark

matter creates an uneven gravitational potential throughout the universe, photons that travel from a distant object to us are deflected from their path due to this potential. *Weak gravitational lensing* of galaxies seeks to probe the gravitational potential by estimating the deflection of photons coming from distant galaxies to us at present day. The weak lensing of these photons creates a deformed image of this galaxy, and the correlation of the shapes and/or sizes of these galaxies probes the dark matter distribution between the galaxy and a present-day observer. The *Dark Energy Survey* (DES) and *Kilo-Degree Survey* (KiDS-1000) performed an analysis of the cosmic shear, measuring the matter clustering parameter $S_8 \equiv \sigma_8 \sqrt{\Omega_m/0.3} = 0.790^{+0.018}_{-0.014}$ (Dark Energy Survey Collaboration and Kilo-Degree Survey Collaboration, 2023).

Finally, an upcoming probe that is of interest to cosmology is the use of *Gravitational Waves*. For one, standard sirens, which are astrophysical sources of gravitational waves of which the amplitude of the wave at the source is known precisely, could be used as an alternative to supernovae in distance ladder measurements. The use of standard sirens in these kinds of measurements could shed light on the prospects of the Hubble tension by testing the calibrations of distance ladder measurements using Cepheid variables. Alternatively, the observation of optical counterparts to gravitational wave events could be used to measure the distance to galaxies without the need for distance ladder measurements. Although initial results have been published from these kinds of measurements (e.g. Hotokezaka et al. (2018); Mukherjee et al. (2020)), the lack of sufficient events with optical counterparts have yet to make this method competitive with supernova or CMB measurements of the Hubble parameter. A second use for gravitational waves is the *Cosmic Gravitational Wave Background*, which would be a source of background gravitational waves originating from the physics of inflation (see subsection 1.3.1). A measurement of the amplitude or shape of the power spectrum of these gravitational waves would measure the tensor mode power spectrum, which would be a parameter directly related to inflation models.

1.4 The CMB Power Spectrum

When making models or measurements of the CMB, we often work in harmonic space rather than real space. Predictions about the CMB are statistical in nature, and working in harmonic space (akin to working in Fourier space) is a natural way of making measurements of things that are spatially-averaged on a sphere.

1.4.1 Temperature Anisotropies

The CMB appears on the sky as some temperature map $T(\vec{n})$, with \vec{n} the direction vector. Under the assumption that this temperature follows a Gaussian distribution, we often care about the variation $\Theta^T(\vec{n}) = \delta T(\vec{n})/\bar{T}$, with \bar{T} the mean CMB temperature. This function can be decomposed into a spherical harmonic decomposition as

$$\Theta^T(\vec{n}) = \sum_{\ell m} a_{\ell m}^T Y_{\ell m}(\vec{n}) \quad (1.20)$$

with $Y_{\ell m}(\vec{n})$ the spherical harmonics functions, characterized by its parameters ℓ, m , and its amplitude coefficients $a_{\ell m}^T$. This decomposition is akin to a Fourier transform, but on a spherical surface (the sky) instead of a flat plane. Computing these coefficients $a_{\ell m}$ is therefore done with a spherical harmonic transform, which is the spherical surface equivalent of a Fourier transform:

$$a_{\ell m} = \int Y_{\ell m}^*(\vec{n}) \Theta(\vec{n}) d\vec{n}, \quad (1.21)$$

with * indicating the complex conjugate, and the integral over the surface of the sphere.

If we assume the temperature anisotropies to be statistically isotropic, we should care only about statistical correlations in Θ^T . Thus, we care about the isotropic two-point correlation function

$$\langle \Theta^T(\vec{n}) \Theta^T(\vec{n}') \rangle = \sum_{\ell} \frac{2\ell + 1}{4\pi} C_{\ell}^{TT} P_{\ell}(\cos \theta), \quad (1.22)$$

where $\cos \theta \equiv \vec{n} \cdot \vec{n}'$, P_{ℓ} are the Legendre polynomials, and C_{ℓ}^{TT} are the coefficients of the **angular power spectrum** (“angular” because they only depend on the angle θ between the two directions of observation). Since this correlation function measures the autospectrum of temperature with itself, these coefficients will be denoted as C_{ℓ}^{TT} . If the CMB is Gaussian, then these coefficients will capture all information about the CMB, and all higher-order odd correlation functions $\langle \Theta^3 \rangle$, $\langle \Theta^5 \rangle$, etc. will be zero. In this case, the coefficients of the angular power spectrum will be

$$C_{\ell} = \sum_m a_{\ell m}^* a_{\ell(-m)}, \quad (1.23)$$

with the sum over all possible values of m . A theory curve for the TT power spectrum is shown in blue (the top-left figure) in Figure 1.5, based on the results from the *Planck* 2018 data release (see more details in later chapters). Marked in this

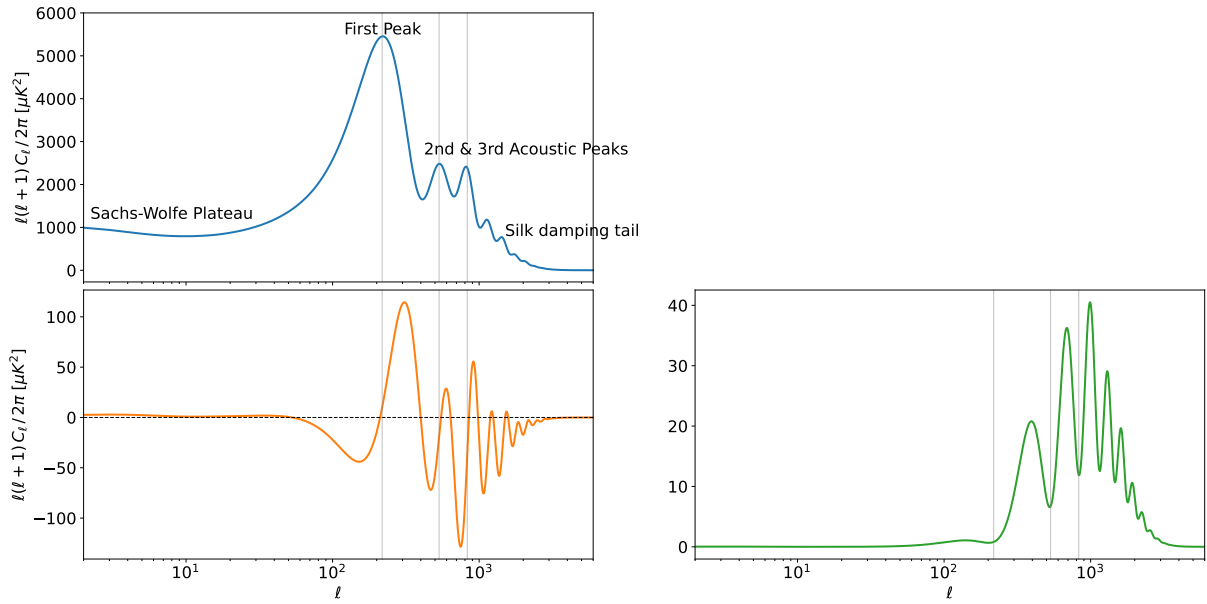


Figure 1.5. The theoretical temperature autospectra C_ℓ^{TT} (**blue**) and C_ℓ^{EE} (**green**), as well as the cross-spectrum C_ℓ^{TE} . The spectra are here plotted as $\mathcal{D}_\ell \equiv \ell(\ell + 1)C_\ell/2\pi$ for a fiducial cosmology similar to the *Planck* 2018 best fit (Planck Collaboration, 2020d). Several important features in the TT autospectrum are labeled. The vertical grey lines indicate the approximate locations of the first three acoustic peaks in the TT autospectra, which correspond to dips in the EE autospectrum and zero-crossings in the TE cross-spectrum.

figure at the acoustic peaks, which are, as mentioned before, generated by baryon-radiation interactions around the time of decoupling (see subsection 1.3.2). These peaks correspond to the scales of maximum compression or rarefaction of the plasma oscillations. Smaller scale features are less prominent, as photon diffusion causes them over time to be “washed out” and be suppressed. This effect is known as *Silk damping*, and is the main reason for the decay in amplitude for peaks at higher ℓ .

If primordial non-Gaussianities are observed in the CMB, they will point to the existence of non-Gaussianities in the primordial power spectrum. These would be seen in the existence of a bispectrum

$$B(k_1, k_2, k_3) \sim f_{NL} [P(k_1)P(k_2) + P(k_2)P(k_3) + P(k_1)P(k_3)], \quad (1.24)$$

where f_{NL} is a free parameter that expresses the amplitude of the bispectrum B with respect to the primordial power spectrum P . Current constraints from the CMB put f_{NL} at $f_{NL}^{\text{local}} = -0.9 \pm 5.1$; $f_{NL}^{\text{equil}} = -26 \pm 47$; $f_{NL}^{\text{ortho}} = -38 \pm 24$ (Planck Collaboration, 2019). This points to the CMB being highly Gaussian. As such, we approximate our baseline results to a Gaussian field, while we continue to explore non-Gaussian contributions.

1.4.2 Polarization Anisotropies

The CMB anisotropies are caused by Thomson scattering within the primordial fluid. Because Thomson scattering has different emission coefficients for polarizations perpendicular and tangential to the scattering angle, if there is a local quadrupole in the CMB temperature anisotropy, this directional dependence can generate a polarized anisotropy. As such, in the presence of a local quadrupole, the CMB will naturally have a polarized anisotropic component as well, arising from the anisotropic properties of the initial density fluid. If there is a local hot-cold quadrupole in the CMB temperature field, then Thomson scattering will generate a polarized anisotropy with polarization angles orthogonal to this quadrupole anisotropy – i.e. the hot-hot axis of the local temperature quadrupole will be orthogonal to the hot-hot axis of the local polarization anisotropy.

Under the assumption of global isotropy, the polarization of the CMB should only consist of components which are rotationally invariant. This is often done by looking at the curl-free and divergence-free components of the polarization, also known as the **E- and B-modes** of the CMB. See Figure 1.6 for a comparison of the difference between E- and B-modes.

E-modes appear naturally as a CMB component due to the fact that Thomson scattering is polarization-dependent. Local temperature quadrupole anisotropies in the primordial plasma will cause charged particles to diffuse from hot to cold regions, and the scattering of CMB photons off of these particles will induce a small polarized E-mode component.

B-modes are not generated via this process. Instead, any B-modes visible must originate from tensor perturbations in the early universe, and thus be related to some physics of inflation. Such B-modes are called **primordial B-modes**, and a measurement of the large scale angular B-mode autospectrum would allow for a measurement of the tensor-to-scalar ratio r , as described in subsection 1.3.1.

Physically, polarization as measured by a telescope is often expressed in terms of the Stokes Q and U parameters, indicating the horizontal-vertical, and orthogonal-diagonal components of a photon's polarization mode. Similarly to how we decompose the temperature anisotropy into the spherical harmonics, we decompose the polarization fields together into two polarization spherical harmonics

$$a_{\ell m}^E \pm ia_{\ell m}^B = - \int_{\pm 2} Y_{\ell m}^*(\vec{n}) [Q \pm iU](\vec{n}) d\vec{n}, \quad (1.25)$$

where $_{\pm 2}Y_{\ell m}$ are the spin-2 spherical harmonics, as opposed to the spin-0 spherical harmonics that are used in the temperature field. For this reason, the polarization modes sometimes called spin-2 components of the CMB.




Figure 1.6. A sketch showing a comparison of (**blue, left**) a curl-free E–mode pattern and (**red, right**) a divergence-free B–mode pattern. Notice that both E– and B–modes are invariant under a uniform rotation. Rotating each polarized photon in an E–mode by 45 degrees will yield a B–mode, and vice-versa.

Gravitational lensing of polarized photons can cause a deflection of the photon path, as mentioned in subsection 1.3.5. This happens through the same manner in which the CMB temperature map gets deflected

$$[Q^L \pm iU^L](\vec{n}) = [Q^u \pm iU^u](\vec{n} + \nabla\phi(\vec{n})) \quad (1.26)$$

Based on Figure 1.6, rotating each polarized photon of an E–mode by 45 degrees will create B–modes. Hence, the addition of lensing effects will lead to some E–modes of the CMB to form **lensed B–modes**, as opposed to the *primordial* B–modes mentioned above.

For this thesis, only the E-mode polarization, its autospectrum C_ℓ^{EE} , and its cross-spectrum C_ℓ^{TE} with temperature are relevant. The full B–mode spectrum is quite weak in amplitude (on the order of $O(10^{-4})$ of the TT power spectrum amplitude), and B–modes are not easy to measure. The experiments treated in this work do not have a signal-dominated measurement of the B–modes, and thus I shall not discuss them in this thesis.

Similar to the temperature anisotropy, there is a polarization anisotropy field $\Theta^E(\vec{n}) = \delta E(\vec{n})/\bar{E}$, with \bar{E} the mean CMB temperature of the E-mode polarization field. Akin to Equation 1.22, combining our measurement of the temperature field and the

polarization field, we get three measures of correlation

$$\begin{aligned}
 \langle \Theta^T(\vec{n})\Theta^T(\vec{n}') \rangle &= \sum_{\ell} \frac{2\ell + 1}{4\pi} C_{\ell}^{TT} P_{\ell}(\cos \theta), \\
 \langle \Theta^T(\vec{n})\Theta^E(\vec{n}') \rangle &= \sum_{\ell} \frac{2\ell + 1}{4\pi} C_{\ell}^{TE} P_{\ell}(\cos \theta), \\
 \langle \Theta^E(\vec{n})\Theta^E(\vec{n}') \rangle &= \sum_{\ell} \frac{2\ell + 1}{4\pi} C_{\ell}^{EE} P_{\ell}(\cos \theta).
 \end{aligned} \tag{1.27}$$

Here, the numbers C_{ℓ}^{TT} measure the amplitudes of the **autocorrelation** of the temperature field, and C_{ℓ}^{EE} measure the autocorrelation of the E-mode polarization field. The numbers C_{ℓ}^{TE} measure the **cross-correlation** of the temperature and E-mode polarization field. An example of these functions is shown in Figure 1.5, with the TE-cross spectrum in orange (lower left) and the EE-autospectrum in green (lower right).

The polarization spectra contain similar peaks to the temperature spectra. We saw before that the peaks in the temperature autospectrum were created by the compression and rarefaction of the primordial plasma. Contrary to this, the peaks in the polarization autospectrum are generated by the motion of plasma with respect to the photons. If we approximate these acoustic oscillations as harmonic oscillators, then the velocity will be zero at maximum compression and rarefaction, while it reach a maximum inbetween two such moments. As such, the EE power spectrum is approximately 90 degrees out of phase with the temperature power spectrum. This is shown in Figure 1.5 by the vertical markers at the first three acoustic peaks in the temperature autospectrum, which correspond to minima in the polarization autospectrum. At these points, the two spectra are out of phase, and the TE cross-spectrum is close to zero.

1.4.3 Modeling the Power Spectrum

Computing what the CMB power spectrum should look like given a set of cosmological parameters is no easy task. It involves computing the background evolution of the universe, the evolution of overdensities of baryonic matter, dark matter, radiation, and other components such as neutrinos within this background, and ultimately deriving how a photon that scatters at the surface of last scattering reaches our present-day instruments while traveling through this evolving universe. This ultimately comes down to solving a series of coupled non-linear differential equations that describe the time-evolution of the matter-radiation fluid in the universe. Solving

these equations give rise to a set of **transfer functions** $\Delta_\ell^X(k)$ that explain how a Fourier mode k is transformed into the ℓ spherical harmonic of some observable X – where for the CMB there are $X = \{T, E\}$.

Once these transfer functions $\Delta_\ell^X(k)$ are computed, we can compute the correlation of our observed quantities by integrating over them and our primordial power spectrum as:

$$C_\ell^{XY} = 4\pi \int \Delta_\ell^X(k) \Delta_\ell^Y(k) P(k) d \ln k, \quad (1.28)$$

where the factor 4π comes from the fact that we're integrating over the surface of a sphere.

The effects of lensing on the CMB, as explained briefly in subsection 1.3.5, then need to be taken into account as an additional modification of the C_ℓ functions. This is often done by computing a lensing power spectrum $C_\ell^{\phi\phi}$, defined as the spherical harmonics of the lensing potential ϕ across the sky, which is derived from the redshift-dependence of the power spectrum of matter clustering. Lensing effects on the CMB power spectrum can then be computed as the sums across multiple ℓ ranges of combinations of C_ℓ^{XY} and $C_\ell^{\phi\phi}$.

Computing CMB power spectra with camb

Computing what these transfer functions, and the inclusion of lensing effects, look like is a complicated endeavour. Software packages, such as the code `camb` (Lewis et al., 2000)[§], exist to compute observables such as the CMB power spectrum. Given a set of cosmological parameters, `camb` can compute the expected observed CMB power spectrum. It does so by first computing the background evolution of the universe by solving the Friedmann equations I described in section 1.1. These background evolution results can already be used for computing cosmological constraints from BAO measurements or distance ladder measurements. To compute the CMB power spectrum however, `camb` needs to compute the evolution of a perturbation within the universe, and then propagate this to the line-of-sight scattering of a photon that travels from the surface of last scattering to the present, which yields these transfer functions $\Delta_\ell(k)$. Integrating over these transfer functions then allows one to compute the *unlensed* CMB power spectrum, as shown in Equation 1.28. Including the effects of lensing involves sums over different ℓ modes for the unlensed power spectrum and the power spectrum

[§]<https://github.com/cmbant/CAMB> and <https://camb.info/>

of the lensing potential, including cross-terms for the E and B polarization modes which get coupled due to these lensing effects. When all these computations have finished, the user is left with an array of power spectra, C_ℓ^{XY} , which contain models for the temperature and polarization power spectra given the requested input parameters up to some accuracy target.

Another advantage to using a code like `camb` is the fact that these code are often highly modular: certain components such as the treatment of recombination physics, or the parameterization of the primordial power spectrum, are all handled by individual modules that can be swapped out for different codes if desired. For example, the treatment of dark energy, which by default is assumed to have an equation of state with $w = -1$, can be swapped out for a component where w is redshift-dependent.

Later, in chapter 4 I will describe my work on novel techniques to speed up these computations using neural network emulators.

1.5 Measurements of the Power Spectrum

Measurements of the CMB power spectrum take several steps. First, the CMB is measured as a field on the sky, with temperature T and polarization anisotropies Q, U as a function of position \vec{n} . Once the CMB is mapped across the sky, this function gets decomposed into the spherical harmonic coefficients $a_{\ell m}$. These are then summed over to compute the power spectrum C_ℓ s.

For an ideal measurement of the CMB, without any instrumental noise or other external factors, our measurement of the power spectrum C_ℓ s is limited by the fact that we have only one sky to measure, and thus our measured value of the power spectrum will randomly scatter around the theoretical prediction due with a Gaussian error. This ideal error, known as the **cosmic variance limit**, is given by

$$\Sigma_{\ell_1, \ell_2}^{XY, ZW} = \frac{\delta_{\ell_1, \ell_2}}{f_{\text{sky}}(2\ell + 1)} \left[(C_\ell^{XY} C_\ell^{ZW}) + (C_\ell^{XZ} C_\ell^{YW}) \right], \quad (1.29)$$

where $\Sigma_{\ell_1, \ell_2}^{XY, ZW}$ measures the covariance between $C_{\ell_1}^{XY}$ and $C_{\ell_2}^{ZW}$, and f_{sky} is the fraction of the sky that is observed with this ideal experiment. Note that the inclusion of the Kronecker delta function δ_{ℓ_1, ℓ_2} means that this covariance matrix is diagonal, meaning there is no covariance between different modes. Figure 1.7 shows a comparison of the theoretical TT power spectrum, compared with a cosmic variance-limited measurement.

To limit the impact of scatter by different C_ℓ s, the measured data points of the




Figure 1.7. The cosmic variance limit (CVL) of the TT auto-spectrum. The **blue line** shows the theoretical prediction for the power spectrum, the **grey dots** represent a random sample of potential measurements of the C_ℓ^{TT} coefficients, while the **red dots** are these random samples averaged over a moving window of $\Delta\ell = 50$. For increasingly smaller angular scales (higher ℓ), the scatter decreases and approaches the mean, proportional to $\sigma_\ell^{TT} \sim (2\ell + 1)^{-1/2}$. We notice that by taking the moving average over the random samples, we recover the power spectrum, but lose some granularity in the low- ℓ regime.

CMB are often *binned*. This binning procedure simply takes a group of adjacent ℓ -modes in a predetermined range, and averages over them to estimate the mean and error of a binned data point. The effect of binning is also shown in Figure 1.7 in the red data points and errorbars.

For real experimental observations, we will end up with a covariance matrix that has a much more complex structure than this. These additional terms are due to (un)correlated noise from the instrument itself, and mode coupling due to the lensing of the CMB; the latter of these two dominates the off-diagonal contributions of the covariance matrix, especially in the blocks for the polarized modes of the CMB.

On top of the terms mentioned here, the covariance will also include terms coming from foreground emission in the microwave regime, lensing effects of the CMB, and correlations between adjacent bins. A common way to estimate the full covariance

matrix is to either analytically estimate each individual term and add them up, or to make multiple simulations of the sky and compute the covariance from monte carlo estimation (see e.g. Atkins et al. (2023)). The result is a dense covariance matrix that includes any kinds of factors known for the covariance of different modes and cross-spectra.

With the covariance of the C_ℓ 's known, it becomes a question of how to measure the model parameters, and how to propagate the error on the measurement of the CMB to the error on these parameters. The methodology of Gaussian likelihood modeling and MCMC parameter sampling is discussed in detail in section 2.1.

1.6 Testing the Λ CDM model with the CMB

The Λ CDM model makes several assumptions that are not necessarily true, but close enough to constraints that some fiducial value or model is often left fixed. Such assumptions can be relaxed. For example one could instead of keeping the mass of neutrinos fixed to some value obtained from particle physics experiments, try to measure the mass of neutrinos from the CMB and/or other cosmological probes directly. Or one could see if dark energy has a different equation of state from $w = -1$, that we assumed before. Such models are considered *extensions* of Λ CDM and are at the forefront of cosmological research. I will go through a variety of the most important extension models.

1.6.1 Shape of the Universe

We assumed earlier in section 1.1 that the universe is flat, meaning that $\Omega_{\kappa,0} = 0$. This is not a requirement of the universe, and the value of $\Omega_{\kappa,0}$ is a free parameter. Early experiments in cosmology, as I will later touch on in subsection 1.7.1, were aimed at measuring this value precisely, and it was soon discovered that at the very least, $\Omega_{\kappa,0} \approx 0$. For an evolving universe, 0 is a metastable equilibrium value, meaning that any tiny deviation from zero will cause the universe to start deviating from flatness more over time. The fact that the universe is very close to flat was considered a problem for cosmology in the 70s, until in the 80s the inflation paradigm (see subsection 1.3.1) was proposed as a solution to among other things this flatness problem.

At present, the flatness condition is implicitly assumed in Λ CDM, as so far all measurements are consistent with a flat universe. A non-flat universe with $\Omega_{\kappa,0}$ as a seventh free parameter is an extension model that is still of interest for measurement.

1.6.2 Non-flat Primordial Power Spectra

The initial conditions for the universe assume a primordial plasma with a density distribution following a nearly flat ($n_s \approx 1$) power law, as described earlier in subsection 1.3.1. This scale-invariant power spectrum was originally predicted from the inflation paradigm, which itself was proposed to address the existing horizon and flatness problems in cosmology at the time.

The exact physics of inflation however, are not known to date, and any kind of field theory assumptions for inflation do not fit within the Standard Model of particle physics. These physics can be probed by probing the shape of the primordial power spectrum. The most stringent two-parameter model for the primordial power spectrum $P(k)$ can be relaxed in numerous ways through extension models. One example is by adding a simple **running of the spectral index**, n_{run} , which quantifies how much the spectral index n_s is varying across different values of k .

1.6.3 Neutrino Mass

Neutrinos are light, weakly-interacting particles that are difficult to measure directly with particle physics experiments. Measurements of neutrino oscillations indicate that neutrinos have some non-zero mass, but attempts at measuring this mass have not been successful so far. Because of their light masses, neutrinos form a relativistic particle contribution to the radiation density of the early universe, whilst forming a dark matter contribution at late times – their interaction strength being too weak to act as baryonic matter. Because the nature of this transition depends on the mass of the neutrinos, there is a possibility for inferring the total mass between the three neutrino eigenstates, Σm_ν , from cosmological measurements.

Current constraints from particle physics put the sum of the neutrino masses at $\Sigma m_\nu > 0.06$ eV, assuming a normal ordering of the mass eigenstates with a minimal mass among the two lowest mass states.

Cosmological probes, such as measurements of the CMB or of BAO features, give a constraint on the upper bound of Σm_ν . Current constraints from the Dark Energy Spectroscopic Instrument (DESI) survey on galaxy clustering and BAO shapes, give an upper bound of $\Sigma m_\nu < 0.072$ eV at 95% credible interval (DESI Collaboration, 2024).

1.6.4 Number of Relativistic Species

In the early universe, when the primordial plasma has a temperature higher than the mass of a baryonic particle, that particle behaves more like radiation than matter as

the two remain coupled together. Neutrinos, which have a mass much less than most other baryons, provide a contribution to radiation that is often parametrized via the number N_{eff} , known as the **effective number of relativistic species**, as

$$\rho_r = \rho_\gamma + \rho_\nu = \rho_\gamma \left[1 + N_{\text{eff}} \frac{7}{8} \left(\frac{4}{11} \right)^{4/3} \right], \quad (1.30)$$

where the numerical factor $7/8(4/11)^{4/3}$ comes from particle physics and the number of spin states that neutrinos can occupy. The main prediction for N_{eff} comes from particle physics, where the Standard Model predicts $N_{\text{eff}} = 3.044$, under the assumption of three light neutrinos, interacting via the electroweak interaction. Deviations from this number can point to the existence of particles or physics beyond the Standard Model, whether an additional boson or fermion, or an interaction term (Bennett et al., 2020).

The value for N_{eff} can be probed through the background expansion rate of the universe, the abundances of primordial elements, and radiation perturbations. Deviations from the standard value of N_{eff} can be measured through the damping tail and position of the acoustic peaks in temperature and E-mode polarization (Hou et al., 2013).

1.6.5 Dark Energy

The dark energy component of the universe, described earlier in section 1.1, is a component that is assumed to have a **parameter of state** $w = -1$. This assumption can be relaxed in a variety of ways. A common model is to add some free parameter w we can fit for, giving us the $w\text{CDM}$ model.

Alternatively, some cosmological probes can at times fit for a time-varying function, often parametrized as

$$w(a) = w_0 + w_a(1 - a) \quad (1.31)$$

with w_0 and w_a free parameters to model for, and a the scale parameter of the universe. This extension is often called the w_0w_a extension.

Measurements of the CMB are not good at constraining the nature of dark energy, and constraints on the w_0w_a cosmology from the CMB alone are poor. At the scales where the CMB is affected by varying values of w_0w_a , constraints on measurements of the CMB are dominated by cosmic variance. Adding these degrees of freedom opens up parameter degeneracies between w_0w_a and other parameters. As such, the CMB alone cannot constrain this kind of model for dark energy well.

1.7 Leading experiments for CMB Measurements

1.7.1 Early Space and Balloon Missions

When it became apparent in the 1970s that early-universe physics in the universe determine the nature of anisotropies of the CMB, experiments to measure these anisotropies became of interest.

Initial efforts were quite limited. It was not until the launch of the *Cosmic Background Explorer* (COBE) in 1989 that the first precision measurements of the all-sky anisotropies were performed. COBE measured the sky anisotropy on angular scales $> 7^\circ$, which is greater than the horizon size at decoupling (see subsection 1.3.2). The smoothness of the CMB at this scale was essential in showing that the early universe was in thermal equilibrium between matter and radiation. Results presented in 1992 showed agreement with the initial scale-invariant inflation model predictions ($n_s \approx 1$), as well as anisotropies of a scale that could seed the structure-growth needed for the present-day large scale structure clustering (Smoot et al., 1992).

Almost a decade later, the next great breakthrough came with the *Balloon Observations Of Millimetric Extragalactic Radiation and Geophysics* (BOOMERanG) and *Millimeter Anisotropy eXperiment IMaging Array* (MAXIMA) missions. The two experiments were both balloon-borne telescopes that aimed to map out the first peak in the CMB anisotropy power spectrum, with the intention of testing the flatness of the universe. The BOOMERanG data set showed that the first acoustic peak in the CMB power spectrum peaked at $\ell = 197 \pm 6$, which was the first indication that the universe was indeed flat in shape (de Bernardis and et al., 2000), indicating that $0.88 < \Omega_0 < 1.12$ at the 95% confidence level. The concurrent MAXIMA mission gave a measurement of the anisotropies down to $\ell \sim 785$ (Hanany et al., 2000). This measurement gave rise to the first values for the baryon density of $\Omega_b h^2 = 0.030 \pm 0.005$ (Bond and et al., 2000), which was the first measurement of this value independent of primordial nucleosynthesis estimates.

These measurements of the CMB anisotropies showed that the CMB was a critical probe of cosmological parameters, one sufficiently constraining on the Λ CDM model even by itself. Soon after the success of the BOOMERanG and MAXIMA missions, space-based observatories were launched and started mapping the full-sky CMB to make a complete measurement of the power spectra.

1.7.2 Precision Cosmology from Space

Wilkinson Microwave Anisotropy Probe

The **Wilkinson Microwave Anisotropy Probe** (*WMAP*) was launched in 2001, measuring the CMB in both temperature and polarization until 2010. It created a full-sky map of the CMB with a 13 arcminute resolution. The final, 9-year temperature and polarization results provided the first precision measurements of the CMB anisotropies over the entire sky (Hinshaw et al., 2013). By observing in five different frequency bands between 23 and 94 GHz, it allowed for mapping foreground contaminations such as synchrotron radiation, free-free emission, and astrophysical dust.

The *WMAP* collaboration published their results in incremental releases. The first release, in 2003, provided a measurement of all six cosmological parameters based on the Λ CDM model, as well as some constraints of extension models. Although the constraints of *WMAP* at this time were not enough to provide strong evidence on Λ CDM as a model, the first measurement of the redshift of reionization from *WMAP* ruled out warm dark matter as a dark matter candidate (Spergel et al., 2003). Combined constraints with external probes gave constraints on the flatness of the universe, the equation of state for dark energy, and an upper limit on the mass and amount of neutrinos in the universe. The later five-year results provided the first statistically significant detection of cosmic neutrinos with a high-confidence measurement of $N_{\text{eff}} > 0$ (Dunkley et al., 2009).

Later, improved measurements of the CMB power spectra in temperature and E-mode polarization yielded an independent, complete six-parameter measurement of Λ CDM, as well as constraints of extensions of this model. With the release of the 9-year data release in 2012, the flatness of the universe was constrained to $1 - \Omega_0 = -0.0027 \pm 0.0039$, giving a major improvement over the *BOOMERanG* experiment over a decade earlier (Hinshaw et al., 2013).

To this day, the *WMAP* measurement of the first acoustic peak remains useful for joint analyses with small-scale measurements of the CMB. In section 3.2, I will delve more into the analysis of this data, and my work on a modern reimplementation of the likelihood software for use in modern joint analyses.

Planck

The *Planck* spacecraft succeeded the *WMAP* probe, launched in 2009 and started its all-sky survey of the microwave background between 30 and 857 GHz later that year. The all-sky measurements by *Planck* provided a large-scale measurement of the CMB that complemented the start of observations by the ground-based telescopes ACT and




Figure 1.8. An overview of modern precision measurements of the CMB power spectra. The largest angular scales for temperature and E-mode polarization are currently best measured by *Planck* (blue). The *Atacama Cosmology Telescope* (ACT, red) and *South Pole Telescope* (SPT, yellow) are ground-based observatories aiming to measure the intermediate- and small-scale anisotropies. The POLARBEAR (green) and BICEP/Keck (purple) experiments are ground-based observatories aiming to make the first precision measurements of the B-mode polarization spectra. The dashed line shows the best-fitting cosmology based on the *Planck* data, and the dash-dotted line shows the primordial B-mode signal expected from a tensor-to-scalar ratio $r = 0.1$. Figure by Erminia Calabrese (2023).

SPT (see subsection 1.7.3), which focused on small-scale features. *Planck* exceeded *WMAP* in sensitivity and resolution, both in its measurements of the CMB temperature and polarization, giving improved all-sky maps of the CMB. The wider frequency range of *Planck* allowed it to map out Galactic thermal dust emission, which is an important contaminant in the microwave sky, and dominates at higher frequencies than the CMB.

The latest legacy results from *Planck*, published five years after its decommissioning, collated the most extensive multi-frequency maps of the full microwave sky; presented the most accurate large-scale measurement of the CMB anisotropies in temperature and polarization; constrained five of the six parameters of the Λ CDM model to sub-percent level accuracy; and showed no hints of beyond- Λ CDM physics by itself (Planck Collaboration, 2020c).

The CMB power spectra as measured by *Planck* in the blue points, and the best-fitting cosmology in the dashed grey line, are shown in Figure 1.8, compared to measurements of the CMB by newer, ground-based observatories.

To this day, results from the *Planck* measurement of the CMB anisotropies remain a state-of-the-art, independent measurement of the Λ CDM model, and is often used as a baseline for other measurements, from either the CMB or other probes, to compare against. The stringent upper limits of *Planck* on non-standard models for inflation, primordial non-Gaussianity, and isocurvature modes have vastly constrained the parameter space for explorations of these kinds of theoretical models. Limits on extensions to Λ CDM, tighter than those measured by *WMAP*, have shown that *Planck* is highly consistent with, and supportive of, the Λ CDM model. A summary of the constraints on Λ CDM and its extensions by *Planck* are shown later in Table 2.1. By itself, *Planck* constrained the spatial curvature of the universe much closer to flatness, provided a 7σ measurement of the existence of light neutrinos, showed that the primordial power spectrum is extremely flat, and gave an upper limit to the tensor-to-scalar ratio. Each of these measurements provided compelling arguments in favor of the Λ CDM model.

As part of this thesis, in section 3.4, I will go through the methodology that was used to analyse *Planck*'s measurement of the small-scale CMB temperature and E-mode spectra, as well as my work in our recreation of the full analysis using the *Simons Observatory* pipeline.

1.7.3 Precision Cosmology from the Ground

When measuring smaller scales in the CMB, larger instruments are needed. As it is easier to build larger instruments on the ground than it is to send them out into space, the smaller scales of the CMB are more accessible from the ground. On the ground, larger dishes can be built, serviced, and upgraded when needed, which is impossible to do in deep space – both *WMAP* and *Planck* orbited far from Earth. Spacecraft thus need to be sent in one mission with a finite lifespan and thus observing time.

The CMB blackbody function peaks at 162 GHz, which coincidentally coincides with a gap in the absorption spectrum of the Earth’s atmosphere. As such, ground-based experiments for the CMB are a viable option to investigate the small-scales, complementary to the large-scale observations that are more accessible with full-sky observations from space. Locations with low water vapour in the atmosphere, such as the Atacama Desert in Chile and the Antarctic region at the South Pole are well-suited for these, to reduce additional contamination from the atmosphere.

Initial explorations of the small-scale temperature and polarization modes of the CMB came with the Antarctic *Degree Angular Scale Interferometer (DASI)* instrument, which in 2001 successfully reported a measurement of the second and third acoustic peaks in the temperature autospectrum (Halverson et al., 2002). Later, this was followed up with a detection of the E-mode polarization spectrum, as well as the TE-cross spectrum measurement (Leitch et al., 2005).

At present, the *South Pole Telescope (SPT)* experiment, located at the Amundsen-Scott South Pole Station in the Antarctic, is observing the microwave sky in temperature and E-mode polarization across a $\sim 1500\text{deg}^2$ patch of the sky in three bands at 95, 150, and 220 GHz (Carlstrom and et al., 2011). The low humidity at the south pole, especially during the Antarctic winter, allows for minimal atmospheric contamination in sub-millimetre observations. Since finishing construction in 2007, SPT has undergone multiple upgrades to its camera, with the most recent *SPT-3G* camera being installed in 2017. Its most recent data release, presented in Balkenhol et al. (2023), contained signal-dominated measurements of the CMB power spectra between $300 \leq \ell \leq 3000$ in TE and EE, as well as between $750 \leq \ell \leq 3000$ in TT. The constraints from SPT, which provide smaller error bars than those of *Planck* and go to higher multipoles, showed great agreement with *Planck* across the entire data range, both at the power spectrum level and at the parameter level.

Complementary to this, the *Background Imaging of Cosmic Extragalactic Polarization*

(*BICEP*) and *Keck Array* experiments are telescopes aiming at measuring the B-mode autospectrum. *BICEP/Keck* have been observing the polarized microwave sky from the Amundsen-Scott South Pole Station since 2006, aiming to beat the noise levels and Galactic dust foregrounds as much as possible. A measurement of the large-scale B-mode polarization autospectrum would give insight into the physics of cosmic inflation by probing the tensor-to-scalar ratio r , as described in subsection 1.3.1. Figure 1.8 shows the most up-to-date results from the *BICEP2/Keck* experiments in purple. So far, the experiments, along with *WMAP* and *Planck* data, have put the most stringent upper limit of $r < 0.036$ at 95% confidence level (Ade and et al., 2021).

The Atacama Cosmology Telescope

The **Atacama Cosmology Telescope** (ACT) was a 6-metre telescope located on Cerro Toco in the Atacama Desert in the north of Chile, at an elevation of 5190 metres. Built in 2007 and seeing first light at the end of that year, ACT has gone through several generations of cameras. The latest camera, AdvACT, has observed roughly half of the sky from 2017 until decommissioning in 2022, at 30, 40, 95, 150, and 220 GHz frequencies in both temperature and polarization, the last three of which are to be used for the upcoming Data Release 6 (DR6) analysis, and constitute a majority of the data volume.

The main goals of ACT were to make a measurement of the CMB power spectra in temperature and E-mode polarization at small, arcminute-size, scales. These measurements, which are complementary to the large angular scales probed by the previous *WMAP* and partially concurrent *Planck* observatories, aimed to be a stringent test of the Λ CDM model at small scales. The initial results from the first ACT camera, MBAC, gave a measurement of the small-scale temperature power spectrum at 148 and 218 GHz which was consistent with the cosmology measured by *WMAP* (Dunkley et al., 2011). Additionally, observations of the CMB by ACT alone gave the first evidence supporting dark energy (Sherwin et al., 2011), and lensing of the CMB (Das et al., 2011)

The following instrument, the ACTPol camera, was designed to measure the temperature and polarization anisotropies at 98 and 150 GHz. The DR4 maps and power spectra were published in 2020, giving new, small-scale measurements of the polarization modes that went beyond the scales accessible by *Planck* (Choi et al., 2020). These measurements were compared to the earlier *WMAP* and *Planck* results, and found to be in good consistency with pre-existing results and with Λ CDM. The addition of these precision small-scale features gave improved constraints on the running of the spectral index, the number of relativistic species, and the primordial Helium

fraction (Aiola et al., 2020).

At the small scale anisotropies ($\ell > 3000$) which ACT is targeting, the microwave sky is dominated by signals of astrophysical origins, the full list of the most important components is explained later in subsection 2.2.1. The ACT power spectrum contains measurements of these components in different frequencies across 40% of the sky in the range $600 < \ell < 8000$, the smallest of which are arcminute scale in size.

The ACT DR4 CMB power spectrum measurement is shown in Figure 1.8 in the red data points, along with the lensing power spectrum from ACT DR6. The additional constraining power of ACT over *Planck* comes from the smaller scales ($\ell > 2500$) which were not measurable with *Planck*, as well as the improved measurement of the E-mode polarizations over *Planck*.

The sensitivity of ACT, which improves over *Planck* at small scales, has made substantial contributions to cluster cosmology by providing high-quality catalogs of optically-confirmed clusters across over ten thousand square degrees of the sky (Hilton et al., 2021). Since ACT was making continuous scans of the sky in the microwave regime, ACT data has been used in the search for transient sources (Naess and et al., 2021a), while also putting upper limits on the existence of a super-Earth planet (“Planet 9”) within the Solar system (Naess and et al., 2021b).

Some of the more recent results of ACT include the DR6 lensing maps (Qu et al., 2023; Madhavacheril et al., 2023). The upcoming results from the ACT DR6 power spectrum analysis will give new constraints on the Λ CDM cosmological parameters at the same level of, but independent from, *Planck*. Additional new measurements of astrophysical small-scale anisotropies will give additional constraining power on the epoch of reionization from the SZ effect (Calabrese et al., 2014).

Analysing the ACT DR6 primary power spectrum data has been a major part of my thesis, and in chapter 5, I will cover this analysis in detail.

Simons Observatory

The **Simons Observatory** (SO) is a CMB experiment under construction on Cerro Toco in the Atacama desert in the north of Chile, with some components already finished and some components due to begin overstations during the next year. SO is aiming to be the next generation in CMB experiments, building on the developments in ground-based measurements of the microwave sky and targeting a variety in scientific objectives for

CMB cosmology and intergalactic astrophysics (The Simons Observatory collaboration, 2019).

SO will be located in the Atacama Desert at an altitude of 5200 metres on the Cerro Toco plateau in Chile, built near the former site of ACT. The full SO experiment will be split between a single Large Aperture Telescope (LAT) and a collection of Small Aperture Telescopes (SATs).

The SATs will be a set of at least four 0.5-metre refracting telescopes, aimed at measuring the large-scale ($\ell \leq 200$) primordial B-mode signal. Some of the SATs have finished construction in 2023 and have begun collecting data as of writing. The nominal plan for SO will involve six SATs, of which two of those coming from the United Kingdom, and extensions of the SO Project being planned with international contributions from SO:UK (United Kingdom), SO:JP (Japan).

The LAT will be a single telescope with a 6 metre primary mirror, similar in size to ACT. The LAT receiver will have seven optics tubes, spanning frequency bands between 27 to 280 GHz, with space for more optics tubes in the future. Each optics tube will have three detector arrays, one for each of intensity and two linear polarization detectors. One low-frequency (LF) tube will make measurements in two bands centred at 27 and 39 GHz; four mid-frequency (MF) tubes will have bands centred at 93 and 145 GHz; and two high-frequency (HF) tubes will have bands at 225 and 280 GHz. The LAT will attain arcminute angular resolution, with a field of view of approximately 7.8° in diameter (Galitzki and et al., 2018).

The SO LAT survey will cover a large area of the sky, at roughly 40% of the sky with uniform coverage. The sky area is chosen to provide optimal overlap with other survey experiments, such as the *Dark Energy Survey* (DES, Dark Energy Survey Collaboration (2005)), the *Dark Energy Spectroscopic Instrument* (DESI, DESI Collaboration (2016, 2022)), the *Vera C. Rubin Observatory* (formerly known as LSST, The LSST Dark Energy Science Collaboration et al. (2021)), and the *Euclid* satellite (Euclid Collaboration et al., 2024).

Going beyond the precision levels of current CMB experiments, the *Simons Observatory* will provide new, highly-accuracy small-scale measurements of the CMB. This will provide the most precise, independent test of the Λ CDM model from the CMB alone since *Planck*. In addition, the small-scale features of the CMB will provide insights into constraints of beyond- Λ CDM physics. In the small scale damping tail of the CMB, effects such as the number of relativistic species N_{eff} , will be precisely measured as a test of the standard model complementary to particle physics. Measurements of the Hubble parameter H_0 will reveal whether the so-called Hubble tension is real

or an artefact of *Planck* systematics. Lensing effects on the CMB will constrain the mass of neutrinos down to $\sigma(\Sigma m_\nu) = 0.09$ eV (The Simons Observatory collaboration, 2019), and combinations with external probes will take us to a regime where we can potentially make a discovery. The large sky area of the LAT survey will provide additional astrophysical science for Galactic and extragalactic components (Hensley et al., 2022). These include Galactic thermal dust emission, extragalactic Sunyaev-Zel'dovich emission, distant radio-loud galaxies, and the cosmic infrared background (CIB).

In section 3.5, I will discuss the work I have done in the development of the SO analysis software. This work is essential for the future analysis of the primary CMB power spectrum analysis with SO LAT data.

Future Experiments for CMB Cosmology

Beyond the experiments mentioned above, future experiments are investigated and science cases are explored for the next generation of CMB measurements.

The *LiteBIRD* mission is a proposed space observatory aiming to measure the large-scale polarization modes of the CMB, with the purpose of measuring the low- ℓ B-mode polarization peak generated by inflationary gravitational waves (LiteBIRD Collaboration, 2022). Current, mostly ground-based, constraints on the inflationary features in the B-mode autospectrum have found an upper limit in the tensor-to-scalar ratio of $r < 0.032$ (Ade and et al., 2021), see also subsection 1.3.1. The plans for the future *LiteBIRD* mission is a space-based observatory, which will be able to access the full-sky polarization modes, with the objective of estimating the lensing B-mode polarization, in order to subtract them off the full B-mode map and thus infer the primordial B-modes only. The mission proposal for *LiteBIRD* aims to constrain the tensor-to-scalar ratio $r \sim O(10^{-3})$ (Namikawa et al., 2023).

Concurrently, the CMB-S4 observatory is a planned, so-called *Stage-IV* observatory[¶], aimed to consist of several large aperture telescopes, improving over previous designs to contain hundred of thousands of detectors each, alongside a series of small aperture telescopes aimed at measuring the degree-scale B-mode polarization (Abazajian and et al., 2016).

[¶]The various “stages” of CMB experiments are often discerned through the number of detectors per single telescope: from hundreds for Stage-I, to thousands for Stage-II, to tens of thousands for Stage-III, and beyond that for Stage-IV.

This application of Bayes' theorem can be seen as a way to update our knowledge of θ : we start with some prior idea about what θ might be, and after measuring some data X we get a new, updated, conditional probability $\mathbb{P}(\theta|X, M)$. The evidence $\mathbb{P}(X|M) = \int \mathcal{L}(X|\theta, M)\pi(\theta) d\theta$ is the likelihood marginalized over the full parameter space of the given model. We often have the freedom to vary our parameter θ while keeping our data X fixed: as such, I shall write the conditional probability of our data matching our model as the **likelihood function** $\mathcal{L}(\theta|X, M)$. Finally, the evidence $\mathbb{P}(X|M)$ exists as a normalization factor over the data X , and is independent of θ . I will implicitly drop this factor, since we will be mostly considering parameter estimation with a single model, and as such the evidence acts only as a normalisation constant and can be ignored.

In this chapter, I will go through the steps involved in the application of Bayesian Inference within cosmology from cosmic microwave background observations.

- In section 2.1, I will lay out what the general structure of the likelihood function $\mathcal{L}(\theta|X)$ is, and how we forward model the primary CMB signal;
- Following that, in subsection 2.2.1, I will go through the parametrization of the secondary emission in the form of astrophysical foregrounds, which contain information about the astrophysics that adds a secondary term to the microwave signal, and form a nuisance component of our full signal;
- In subsection 2.2.2, I will explain the basic forward modeling we employ for instrumental systematics, such as calibration uncertainties or bandpass mismatch, and how they affect our signal model;
- Finally, in section 2.3, I will explain the concept of Markov Chain Monte Carlo (MCMC) sampling and its use in sampling the posterior $\mathbb{P}(\theta|X, M)$ from the likelihood $\mathcal{L}(\theta|X)$.

The following paper forms parts of this chapter:

- **The Simons Observatory: impact of bandpass, polarization angle and calibration uncertainties on small-scale power spectrum analysis**, *S. Giardiello, M. Gerbino, L. Pagano, et al., (2024)*;

2.1 General Likelihood Implementation

2.1.1 Likelihood Function

We seek the function $\mathcal{L}(X|\theta, M)$, which is the likelihood \mathcal{L} of our data X given some forward modeling based on parameters θ and a model M . For our data, this is whatever measurement of the signal C_ℓ s we measure. As described before in the introduction, this will be some combination of (1) the CMB C_ℓ s coming from our power spectrum model given a cosmology; (2) any additional contribution from secondaries, which are based on astrophysical effects; (3) any modification of these combined signals due to the systematics of our instrumental properties. I will denote our theoretical model for the total signal as $C_\ell^{\text{th}}(\theta, M)$, and the signal we measured as \hat{C}_ℓ^{d} .

For large enough ℓ , the microwave power spectrum is fully Gaussian*. This means we can write our likelihood function as:

$$\log \mathcal{L}(\hat{C}_\ell^{\text{d}}|\theta, M) = -\frac{1}{2} [\hat{C}_\ell^{\text{d}} - C_\ell^{\text{th}}(\theta, M)]^T \Sigma^{-1} [\hat{C}_\ell^{\text{d}} - C_\ell^{\text{th}}(\theta, M)] - \frac{1}{2} \log \det \Sigma, \quad (2.2)$$

where $C_\ell^{\text{th}}(\theta, M)$ is the theoretical prediction of the signal. The $\frac{1}{2} \ln \det \Sigma$ factor is a constant term that normalizes the Gaussian likelihood such that it integrates to unity. Since we often only compare likelihood ratios or log-likelihood differences, such terms commonly cancel out. I will therefore implicitly drop this constant term in future expressions of the likelihood function.

Because the CMB C_ℓ s span a large range of orders of magnitude, we find that it is more efficient to work in $\mathcal{D}_\ell \equiv \ell(\ell + 1)C_\ell/2\pi$ space. Since the difference is merely a numerical, constant prefactor in our data and theory vectors, and covariance matrix, we can trivially write our likelihood as:

$$\log \mathcal{L}(\hat{\mathcal{D}}_\ell^{\text{d}}|\theta, M) = -\frac{1}{2} [\hat{\mathcal{D}}_\ell^{\text{d}} - \mathcal{D}_\ell^{\text{th}}(\theta, M)]^T \Sigma^{-1} [\hat{\mathcal{D}}_\ell^{\text{d}} - \mathcal{D}_\ell^{\text{th}}(\theta, M)], \quad (2.3)$$

where now the covariance matrix Σ includes the constant $\ell(\ell + 1)/2\pi$ factor to translate from C_ℓ to \mathcal{D}_ℓ space.

*The CMB is very Gaussian above $\ell \approx 100$, although in Planck Collaboration (2016), the *Planck* analysis found that for $\ell = 30 \sim 100$, the Gaussian shape is reasonably close as well with minimal bias on their parameter constraints.

2.1.2 The Data Vector

For inference from small-scale CMB features, we are interested in both the temperature and E-mode polarization of the CMB. As introduced before in subsection 1.4.2, this means we are interested in the TT and EE autospectra, as well as the TE cross-spectrum.

Our full data and model vectors \mathcal{D}_ℓ will thus consist of three sets of data: namely the TT, TE, and EE spectra. These three are commonly concatenated, one after another, into one longer vector. Whenever I write \mathcal{D}_ℓ , I implicitly mean the concatenation of \mathcal{D}_ℓ^{TT} , \mathcal{D}_ℓ^{TE} , \mathcal{D}_ℓ^{EE} , or any kind of permutation of these three.

The covariance matrix Σ for such a data vector may also include correlations between these signals. For example, any noise in the temperature measurement will affect both the TT and TE power spectra, thus the covariance matrix will contain information in its TT-TE block to describe the correlated noise property that these different spectra have.

Apart from measuring both temperature and E-mode polarization separately, we also measure both across multiple frequencies. As such, most likelihood functions we consider will consist of a full data vector \mathcal{D}_ℓ with multiple instances of the TT, TE, and EE power spectra, for each different pair of frequencies that have been measured. While the CMB component of these cross-spectra is assumed to be the same, the terms coming from astrophysical sources will vary with frequency in a manner that is different from a blackbody spectrum.

In our data vector, we denote this by explicitly denoting both whether a term is a temperature T or polarization mode E, and which frequency either term is measured at. If component X is measured at some frequency ν_1 and component Y is measured at some frequency ν_2 , then the cross-spectrum of these two is $\mathcal{D}_{\ell,\nu_1\nu_2}^{XY}$. This cross-spectrum is potentially different from a measurement of component X measured at ν_2 and Y at ν_1 .

As a simple illustrative example, for a theoretical experiment that measures the CMB in temperature and E-mode polarization, both at 90 and 150 GHz, the full data vector could have a structure similar to:

$$\mathcal{D}_{\ell,\nu_1\nu_2} = \left\{ \mathcal{D}_{\ell,90\times 90}^{TT}, \mathcal{D}_{\ell,90\times 90}^{TE}, \mathcal{D}_{\ell,90\times 90}^{EE}, \right. \\ \left. \mathcal{D}_{\ell,90\times 150}^{TT}, \mathcal{D}_{\ell,90\times 150}^{TE}, \mathcal{D}_{\ell,150\times 90}^{TE}, \mathcal{D}_{\ell,90\times 150}^{EE}, \right. \\ \left. \mathcal{D}_{\ell,150\times 150}^{TT}, \mathcal{D}_{\ell,150\times 150}^{TE}, \mathcal{D}_{\ell,150\times 150}^{EE} \right\}, \quad (2.4)$$

where it should be noted that $\mathcal{D}_{\ell,90 \times 150}^{TE} \neq \mathcal{D}_{\ell,150 \times 90}^{TE}$, since one of them measures temperature at 90 GHz and polarization at 150 GHz, while the other one measures temperature at 150 GHz and polarization at 90 GHz. Sometimes this may be indicated by writing $\mathcal{D}_{\ell,150 \times 90}^{TE}$ as $\mathcal{D}_{\ell,90 \times 150}^{ET}$ instead (note the swapping of TE instead of 90×150).

2.1.3 Binning

Since our measurement of each \mathcal{D}_ℓ is, even at cosmic variance limit, going to scatter around the theoretical \mathcal{D}_ℓ for a given cosmology, we expect that our observational power spectrum will not be a perfectly smooth curve. As mentioned before in section 1.5, we compress the data into a binned measurement.

By averaging our signal over adjacent ℓ modes, we can bring down the error of our signal, improving our signal-to-noise ratio, at the cost of losing a bit of the finer structure in the features of the signal. This process of averaging the signal over several different ℓ samples is called **binning**.

Given a collection of samples \mathcal{D}_ℓ , we write the binned vector \mathcal{D}_b as

$$\mathcal{D}_b = \sum_{\ell} w_{b\ell} \mathcal{D}_\ell \quad (2.5)$$

with $w_{b\ell}$ the window function of bin b . Similarly, we bin the covariance matrix the same way was

$$\Sigma_{b_1, b_2}^{XY, ZW} = \sum_{\ell_1, \ell_2} w_{b_1 \ell_1} w_{b_2 \ell_2} \Sigma_{\ell_1, \ell_2}^{XY, ZW}, \quad (2.6)$$

where now we estimate the covariance between binned \mathcal{D}_b modes instead of between separate \mathcal{D}_ℓ modes. In the last chapter, Figure 1.7 showed the power spectrum with the binned data points in red.

For CMB experiments, estimating what the covariance matrix looks like is a multi-step process. Analytical estimates can be made by assuming the cosmic variance limit and uncorrelated white noise, but these do not take covariances between different bins into account. A different way to estimate the covariance matrix is to create a map of the sky according to a fiducial CMB model and some model for foregrounds. This map is then polluted with a model for the instrumental noise of the experiment, before being passed through the same pipeline that will later compute the power spectrum from the data. The resulting power spectrum from this simulated map will contain a random, noisy evaluation of the CMB based on some simulation. By computing the power spectra of many different random evaluations of the CMB, foregrounds, and noise, the covariance matrix can be computed as the covariance of these different spectra. The

challenge in this process comes mainly from correctly simulating the instrumental noise properties on these maps. Atkins et al. (2023) gives an overview of this noise simulation for the *Atacama Cosmology Telescope* as part of the upcoming DR6 data release, which I will later talk about in chapter 5.

2.2 Full Spectrum Model

The main two things we need to model the power spectrum is some theory prediction and a way in which we model how a theoretical signal would be measured by the instrument.

The theoretical observable curve will include a component $\mathcal{D}_\ell^{\text{CMB}}$ from the CMB. This curve is usually computed from some cosmological prediction with an Einstein-Boltzmann code, as described in section 1.4. In addition, there is a component $\mathcal{D}_{\ell,\nu_1\nu_2}^{\text{fg}}$ coming from a number of astrophysical foregrounds, such as thermal emission from dust, radio emission from distance galaxies, or the interaction between the CMB and hot plasma in galaxy clusters. I will go into details for the full description of these relevant components later in subsection 2.2.1. Since these components do not necessarily follow a blackbody spectrum, they will have a different frequency dependence than the CMB power spectrum, and thus are indexed by an additional ν_1, ν_2 index for describe that.

Then there are a variety of ways in which the theory curve gets modified for instrumental systematics. For the most part, imperfect signal transmission and polarization filters will simply multiply the entire theory signal with some number, while a mismatch between the modeled frequency and the measured frequency can cause shifts in components which have a frequency dependent model. Asymmetries in the beam profile can create incorrect polarization anisotropies as the temperature signal “leaks” into the polarization signal, and uncercainties in the polarization angle can create a confusion between a measured E- and B-mode signal. A complete overview of the components relevant for my analyses is given in subsubsection 2.2.2.

We write our full model vector as

$$\mathcal{D}_b^{\text{th}}(\theta, M) = w_{b\ell} \left[c_{ij}^{XY} \left(\mathcal{D}_\ell^{\text{CMB}} + \mathcal{D}_{\ell,\nu_1\nu_2}^{\text{fg}} \right) \right] \quad (2.7)$$

where $w_{b\ell}$ are our window functions (see subsection 2.1.3), c_{ij}^{XY} is a calibration factor for a given cross-spectrum, and $\mathcal{D}_{\ell,\nu_1\nu_2}^{\text{fg}}$ is an additional component from astrophysical foreground emission. I have described modeling of the CMB power spectrum before in subsection 1.4.3. I will lay out the component of the foreground model in subsection 2.2.1, and the instrumental nuisance factors in subsubsection 2.2.2.

2.2.1 Astrophysical Foreground Model




Figure 2.1. A comparison of the observed TT and pure CMB signal at 150×150 GHz for an SO-like experiment where most of the Galactic dust is masked. The CMB signal for a fiducial cosmology is shown in **grey**, while the full foreground signal is shown in **black**. The sum of the two components is shown in the **black dashed** line. Note that the CMB only dominates the signal up to an $l \sim 3000$, and above this the astrophysical foreground emission takes over. For experiments which include signals above $l = 1000 \sim 2000$, correctly modeling the contamination by astrophysical foregrounds becomes imperative.

At around scales of $l \sim 3000$, the microwave background is dominated by astrophysical foregrounds (see Figure 2.1). This means that the full signal \mathcal{D}_ℓ needs to be modeled as

$$\mathcal{D}_{\ell, \nu_1 \nu_2} = \mathcal{D}_\ell^{\text{CMB}} + \mathcal{D}_{\ell, \nu_1 \nu_2}^{\text{SZ}} + \mathcal{D}_{\ell, \nu_1 \nu_2}^{\text{CIB}} + \mathcal{D}_{\ell, \nu_1 \nu_2}^{\text{dust}} + \mathcal{D}_{\ell, \nu_1 \nu_2}^{\text{PS}} \quad (2.8)$$

where now $\mathcal{D}_{\ell, \nu_1 \nu_2}$ depends also on the frequencies $\nu_1 \nu_2$ of the cross-spectrum, $\mathcal{D}_\ell^{\text{CMB}}$ is the CMB signal, which is blackbody and thus does not depend on frequency, $\mathcal{D}_{\ell, \nu}^{\text{SZ}}$ is the Sunyaev-Zel'dovich effect, $\mathcal{D}_{\ell, \nu}^{\text{CIB}}$ is the cosmic infrared background (CIB), $\mathcal{D}_{\ell, \nu}^{\text{dust}}$ is galactic dust, and $\mathcal{D}_{\ell, \nu}^{\text{PS}}$ is the contribution from radio point sources (e.g. Active Galactic Nuclei). For the polarization components TE and EE, only the dust and point source terms are important. Additional terms, such as galactic synchrotron emission,

are ignored in this case, as, for the small-scale analyses treated in this work, they are below the detection limit for the frequency range of current ground-based telescopes.

I will follow Choi et al. (2020) in describing the various components here. A plot with models for all foreground components in temperature is shown in Figure 2.2.

The Sunyaev–Zel’dovich Effect

The Sunyaev–Zel’dovich (SZ) effect is a secondary anisotropy caused by inverse Compton scattering of CMB photons off high-energy, charged particles (often electrons) in galaxy clusters. Often for CMB observations, two components are considered: the *thermal* SZ effect and the *kinematic* SZ effect, also denoted tSZ and kSZ respectively. While both are the same effect (inverse Compton scattering of the CMB photons), they are caused by electrons in different regimes. I have given an overview in the SZ before in subsection 1.3.4.

The tSZ effect is caused by electrons in hot gas in a thermal distribution within a cluster. When a CMB photon travels through a cluster, the thermal distribution of the electrons inside tends to scatter the photon as it is traveling through. This scattering causes local a spectral distortion of the CMB blackbody, as low-energy photons are scattered upwards to higher energies. At frequencies below the peak frequency of the CMB (162 GHz), an SZ cluster appears like a small “cold spot” on the map, since there is a deficiency in low-energy photons, while at higher frequencies, these clusters appear as “hot spots”, since there is an increase in higher-energy photons (Mroczkowski et al., 2019).

Outside of primary CMB power spectrum analysis, maps of the tSZ effect can be compared to X-ray emission of the intracluster gas to determine the baryonic mass content of the cluster (Hilton et al., 2021). A model for the tSZ power spectrum can be motivated by cosmological simulations, and is given by:

$$\mathcal{D}_{\ell, \nu_1 \nu_2}^{\text{tSZ}} = A_{\text{tSZ}} \mathcal{D}_{\ell, 0}^{\text{tSZ}} \frac{f(\nu_1) f(\nu_2)}{f(\nu_0)^2}, \quad (2.9)$$

where $f(\nu) = x \coth(x/2) - 4$, with $x \equiv h\nu/k_B T_{\text{CMB}}$, a free amplitude A_{tSZ} , and a template $\mathcal{D}_{\ell, 0}$ normalized to unity at some ℓ_0 and ν_0 (Battaglia et al., 2012). These pivot points are often taken to be $\ell_0 = 3000$ and $\nu_0 = 150\text{GHz}$.

The kSZ effect is caused by galaxy clusters having a peculiar velocity with respect to the Hubble flow. It is modeled as:

$$\mathcal{D}_\ell^{\text{kSZ}} = A_{\text{kSZ}} \mathcal{D}_{\ell,0}^{\text{kSZ}}, \quad (2.10)$$

where $\mathcal{D}_{\ell,0}^{\text{kSZ}}$ is a template spectrum derived from cosmological simulations (Battaglia et al., 2010).

Measurements of the exact nature of the kSZ power spectrum are of interest for cosmological research as well. As described earlier in subsection 1.3.4, the physics of the Epoch of Reionization impact the structure of the cosmic web, which leads to a potentially measurable effect in the kSZ power spectrum. By precisely measuring the kSZ power spectrum at small scales, a measurement of reionization physics, which is normally only loosely constrained by the large scale EE power spectrum, can give new insights in a moment in the Universe's early history that is poorly understood so far (Calabrese et al., 2014).

The Cosmic Infrared Background

The Cosmic Infrared Background (CIB) is the thermal dust emission from star-forming galaxies at high redshift. This component contains two parts, a poisson term (CIBp) and a clustered term (CIBc), which are modeled as

$$\mathcal{D}_{\ell,\nu}^{\text{CIBp}} = A_p \frac{\ell(\ell+1)}{\ell_0(\ell_0+1)} \left(\frac{\mu(\nu_1, \beta_p) \mu(\nu_2, \beta_p)}{\mu(\nu_0, \beta_p)^2} \right), \quad (2.11)$$

$$\mathcal{D}_{\ell,\nu}^{\text{CIBc}} = A_c \mathcal{D}_{\ell,0}^{\text{CIBc}} \frac{\ell(\ell+1)}{\ell_0(\ell_0+1)} \left(\frac{\mu(\nu_1, \beta_c) \mu(\nu_2, \beta_c)}{\mu(\nu_0, \beta_c)^2} \right), \quad (2.12)$$

with $\mu(\nu, \beta) \equiv \nu^\beta B_\nu(T_d) g(\nu)$ a modified blackbody function for some effective dust temperature $T_d = 9.6\text{K}$. This dust temperature, while not perfectly known, is not well constrained in the frequency range observed by most CMB experiments, as it is degenerate with the spectral index β . For the clustered component, the power spectrum shape $\mathcal{D}_{\ell,0}^{\text{CIBc}}$ is often taken to be a template in $\ell < 3000$ coming from high-frequency measurements of *Planck* (Planck Collaboration, 2014), and extrapolated as $\mathcal{D}_{\ell,0}^{\text{CIBc}} \propto \ell^{\alpha_c}$ with $\alpha_c = 0.8$ above that (Addison et al., 2012). The spectral indices β_p, β_c are related but not necessarily equal. Current constraining power is consistent with $\beta_p = \beta_c$ being fixed.

In addition, there is a cross-correlation between the clustered CIB and the tSZ effect. This correlation is included either by the creation of a joint tSZ-CIBc template, or by simply subtracting a term of the form

$$\mathcal{D}_{\ell,\nu}^{\text{tSZ}\times\text{CIB}} = -2\xi\sqrt{A_{\text{tSZ}}A_{\text{CIBc}}}\mathcal{D}_{\ell,0}^{\text{tSZ-CIBc}}\frac{f(\nu_1)\mu(\nu_2,\beta_c) + f(\nu_2)\mu(\nu_1,\beta_c)}{f(\nu_0)} \quad (2.13)$$

with A_{tSZ} , A_c the amplitudes of the tSZ and clustered CIB spectra, ξ the correlation parameter, and $\mathcal{D}_{\ell,0}^{\text{tSZ-CIBc}}$ the spectrum template (Addison et al., 2012). The functions $f(\nu)$ and $\mu(\nu, \beta)$ the same as those mentioned above.

Galactic Dust

Galactic dust is an important, non-isotropic foreground component. The center and plane of the Milky Way galaxy dominates the microwave emission, especially at higher frequencies. For small angular scales, observing the sky away from the galactic plane is an option, but for large angular scales, observing the galactic plane is needed. For the latter, the *Planck* satellite used its 353 GHz channel to map out the galactic foreground, which was then used to identify areas of strong emission, which are then used to mask out regions where galactic dust emission dominates the microwave emission.

Ground-based experiments rely on these high-frequency measurements from space to mask the Galaxy. For example, in the ACT DR4 analysis, the *Planck* Galactic mask was used to cut out the sky regions containing dominant galactic foregrounds. The remaining area is expected to contain emissions of only $O(1\%)$ of the CMB power in the 90 and 150 GHz channels of ACT. This residual dust was then modeled as a power law in harmonic space:

$$\mathcal{D}_{\ell,\nu}^{\text{dust}} = A_{\text{dust}}\left(\frac{\ell}{\ell_0}\right)^{\alpha_d+2}\mathcal{F}_\nu^d g_1(\nu)^2, \quad (2.14)$$

with A_{dust} a free scaling amplitude for dust, $\ell_0 = 500$ an arbitrary, pre-determined pivot scale, and α_d the spectral index, found to be -2.6 for TT, -2.4 for TE and -2.6 for EE (Planck Collaboration, 2020b; Choi et al., 2020). The remaining two terms, \mathcal{F}_ν^d is a conversion factor between the Planck modified blackbody dust model in antenna temperature, and $g_1(\nu)$ is a conversion factor from antenna to CMB temperature. Here, the antenna temperature is a measurement of the total amount of power received by the detector, which the factor \mathcal{F}_ν^d converts from the modified blackbody temperature at 353 GHz by *Planck* to the temperature measured by the ACT detector at a frequency ν . The factor $g_1(\nu)$ then converts this from a signal with a given power at a frequency ν to a temperature fluctuation relative to the CMB temperature at that frequency.

Radio Sources

Distant, unresolved radio sources provide an additional Poisson-like noise component. Each experiment will have a detection limit for bright sources, which can be removed at the map level by either masking or subtracting a template. A large number of radio sources will however, be too faint and remain unresolved and cannot be removed in this manner. The power from this latter component is modeled as a power law:

$$\mathcal{D}_{\ell,\nu}^{\text{radio}} = A_s \left(\frac{\nu}{\nu_0} \right)^{\beta_s} \frac{\ell(\ell+1)}{\ell_0(\ell_0+1)} \quad (2.15)$$

where $\nu_0 = 150\text{GHz}$, $\ell_0 = 3000$ are arbitrary pivot scales, and A_s and β_s are the amplitude and spectral index of the model.

It is usually assumed that β_s is the same in both temperature and polarization, $\beta_s^T = \beta_s^E$, but this is not necessarily true, and with the prospect of improved measurements of the polarization power spectra, this assumption can be scrutinized.

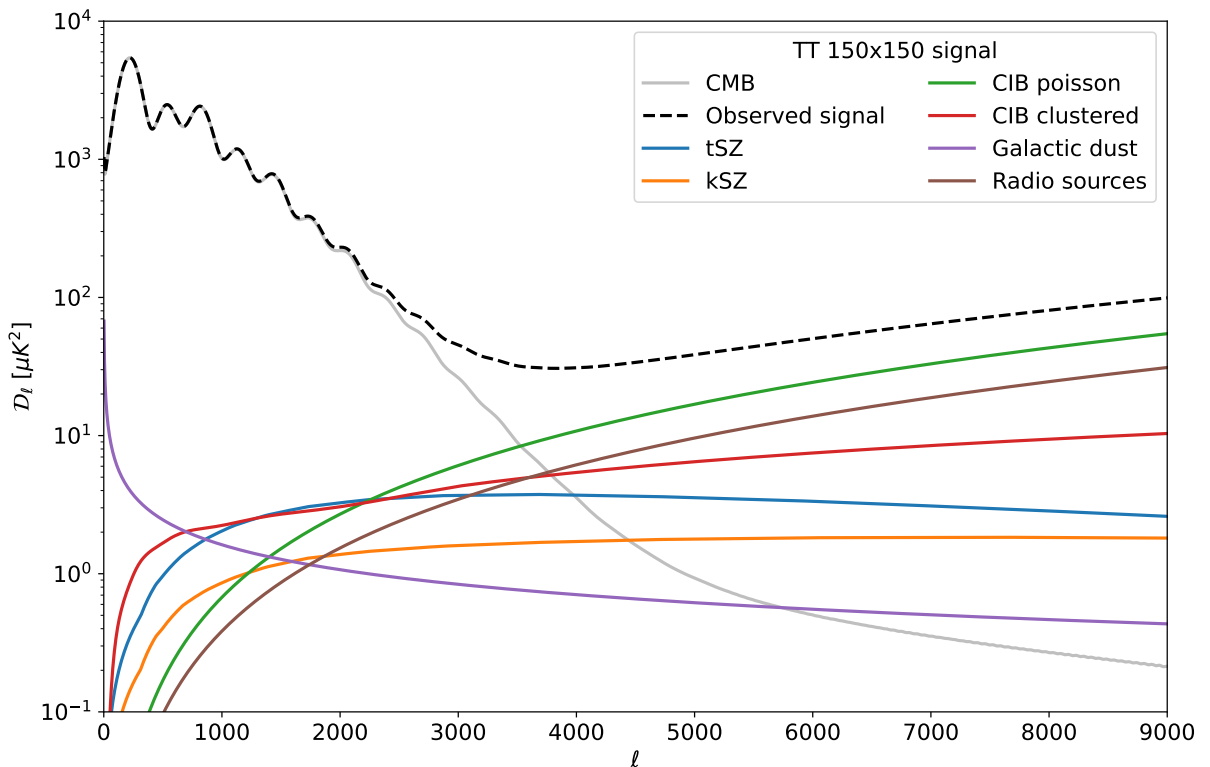


Figure 2.2. A fiducial model for the expected contamination of the TT signal at 150×150 GHz for an SO-like experiment. The pure CMB signal (grey) dominates the signal below $\ell \sim 3000$ (see also Figure 2.1), while above that the various astrophysical foregrounds become a dominant contribution to the full TT signal.

2.2.2 Instrumental Nuisance Parameters

Instrumental noise and miscalibration factors can affect our measurement of the CMB and astrophysical foregrounds. To some extent, corrections for these factors can be applied to the power spectrum measurement prior to likelihood modelling, but uncertainties of what the exact value of these factors are have to be marginalized over to obtain proper posterior parameter constraints.

For most experiments, we need to correct for a variety of instrumental systematics. These systematics include things like gain calibration, polarization efficiencies, or mismatch in the bandpass center. A summary of the impact of these uncertainties on cosmological measurements is given in Giardiello et al. (2024). I provide an overview of how some of these parameters are included in the model. These will be part of the data vectors that we will fit for in chapter 5.

Calibration and Polarization Efficiency Factors

Calibration and polarization efficiency factors arise from imperfect transmission of the signal throughout the instrument. Since any instrument will absorb a small amount of signal as it passes through the optics, the true signal will be slightly stronger than reported by the detector. In addition, the full CMB map needs to be calibrated with respect to the CMB dipole to correct for a contribution due to the the motion of the Sun through the Galactic plane.

Each of these parameters is applied as a multiplicative factor to the model, as:

$$\begin{aligned}\mathcal{D}_{\ell,\nu_1\nu_2}^{TT,ij} &\rightarrow \mathcal{D}_{\ell,\nu_1\nu_2}^{TT,ij} / \text{cal}_{\text{dip}}^2 c_{\nu_1}^i c_{\nu_2}^j, \\ \mathcal{D}_{\ell,\nu_1\nu_2}^{TE,ij} &\rightarrow \mathcal{D}_{\ell,\nu_1\nu_2}^{TE,ij} / \text{cal}_{\text{dip}}^2 c_{\nu_1}^i c_{\nu_2}^j p_{\nu_2}^j, \\ \mathcal{D}_{\ell,\nu_1\nu_2}^{EE,ij} &\rightarrow \mathcal{D}_{\ell,\nu_1\nu_2}^{EE,ij} / \text{cal}_{\text{dip}}^2 c_{\nu_1}^i p_{\nu_1}^i c_{\nu_2}^j p_{\nu_2}^j,\end{aligned}$$

where i and ν_1 refer to the first detector array and its frequency, and j and ν_2 refer to the second detector array and its frequency in the cross-spectrum. We note that every component is multiplied by the global dipole calibration gain factor, each component of a given array d at a given frequency ν is multiplied by the calibration factor c_{ν}^d , and each polarized component of that array is multiplied by p_{ν}^d .

Bandpass Integration and Mismatch

The optical path of the instrument is frequency-dependent, and the transmission of the signal will depend on its frequency ν . We treat the signal as if it is a singular frequency,

but in reality the detector will absorb a whole range of frequencies, giving a transmission bandpass $F(\nu)$. For the most part, this means that for a foreground component, its actual signal strength will be proportional to its **integrated signal**:

$$\mathcal{D}_{\ell, \nu_1 \nu_2}^{XY, fg} \propto \int d\nu_1 \int d\nu_2 F(\nu_1) F(\nu_2) f_{\text{SED}}^{fg}(\nu_1, \nu_2) \mathcal{D}_{\ell}^{XY, fg}, \quad (2.16)$$

where $f_{\text{SED}}^{fg}(\nu_1, \nu_2)$ is the spectral energy distribution (SED) of the foreground (described above), and the integrals are performed over the entire bandpass. Since we are working in blackbody temperature units, the CMB signal is not frequency-dependent and thus the CMB \mathcal{D}_{ℓ} s will not depend on the bandpass $F(\nu)$.

If this bandpass profile is measured but there is a mismatch in its central frequency, then the amplitudes of all foregrounds will shift around due to their differing SEDs with respect to the central frequency. As such, if the true center of the bandpass of detector d is not ν but actually $\nu + \Delta_d^\nu$, then the signal strength needs to be modeled as:

$$\mathcal{D}_{\ell, \nu_1 \nu_2}^{XY, fg} \propto \int d\nu_1 \int d\nu_2 F(\nu_1 + \Delta_{d_1}^{\nu_1}) F(\nu_2 + \Delta_{d_2}^{\nu_2}) f_{\text{SED}}^{fg}(\nu_1 + \Delta_{d_1}^{\nu_1}, \nu_2 + \Delta_{d_2}^{\nu_2}) \mathcal{D}_{\ell}^{XY, fg}, \quad (2.17)$$

where now all the bandpass profiles and SEDs are shifted with their respective bandpass shift parameters Δ_d^ν .

2.3 Markov Chain Monte Carlo Sampling

Now that we have some likelihood function $\mathcal{L}(X|\theta, M)$ and some priors $\pi(\theta)$, we seek to evaluate the full posterior $\mathbb{P}(\theta|X, M)$ (see Equation 2.1). However, evaluating the full posterior is difficult due to the high dimensionality and complexity of our parameter space – note that we have at least 6 cosmological, about 15 astrophysical, and commonly of the order of 5 to 15 instrumental parameters. Not only would we need to find a maximum a posteriori (MAP) for our data, but computing the uncertainties on our parameters requires somehow computing the Hessian of the posterior likelihood function, which may not be computationally feasible – and that is assuming our parameter constraints are relatively Gaussian.

To this end, we need to employ some methodology to find the posterior constraints on our parameters. A commonly used process for this is **Markov chain Monte Carlo** (MCMC) sampling, a stochastic process that seeks to find a stable limiting equilibrium distribution that closely follows the posterior distribution. While in theory an MCMC process will only recover the posterior distribution in the limit of infinite samples, it is often more than feasible to run an MCMC sampling process for a finite amount of time,

and under certain tests and checks the resulting posterior density function (PDF) can be assumed to be “close enough” that it is not statistically significantly different from the true PDF.

For MCMC sampling, there are multiple publicly available softwares for cosmological sampling. We make use of the sampling software `cobaya` Lewis (2013); Torrado and Lewis (2019, 2021)[†], a large suite of `python` software for Bayesian analysis. Among other things, `cobaya` comes with a full MCMC sampler and a χ^2 -minimizer to quickly explore the best-fitting region in parameter space, and allows accessible interfacing to custom-made theory and likelihood software. In addition, `cobaya` contains the necessary interfacing for parallel computing, meaning that multiple chains can be run at the same time, while allowing their samplers to exchange intermediate results with each other to improve estimation of posterior probability statistics from multiple runners at the same time during the sampling process.

There are a variety of MCMC algorithms that can do this, but we will employ the **Metropolis–Hastings algorithm** (Metropolis et al., 1953; Hastings, 1970). The general process of this Metropolis–Hastings algorithm is as follows:

1. Pick a starting point θ_1 , this is where our chain starts sampling;
2. Draw a sample $y \sim \mathcal{G}(y|\theta_i)$;
3. Evaluate the acceptance ratio $\alpha = \mathbb{P}(y|X)/\mathbb{P}(\theta_i|X)$;
4. Compute a random number $r \sim \mathcal{U}(0, 1)$, if $r \leq \alpha$, then $\theta_{i+1} = y$, else $\theta_{i+1} = \theta_i$;
5. Repeat steps 2–5 until we consider our chain sufficiently long or converged.

In this method, the samples $\{\theta_1, \dots, \theta_n\}$ are the **accepted samples** that form our chain, the point y is called our **candidate** for each step in the chain, α is the **acceptance ratio**, \mathcal{G} is our **proposal distribution**, $\mathbb{P}(\theta|X)$ is our **target distribution**, and $\mathcal{U}(0, 1)$ is a uniform distribution on the interval $[0, 1]$.

We seek a target distribution $\mathbb{P}(\theta|X, M)$, which we find should follow our likelihood function $\mathcal{L}(X|\theta, M)$ and our prior distribution $\pi(\theta)$, with X our data and θ our parameter vector. Since we have described before methods of computing the *log-likelihood* function, it is also common to transform our acceptance test from a uniform distribution to an exponential one, by observing that the test whether

[†]<https://github.com/CobayaSampler/cobaya>

$$\alpha \equiv \frac{\mathcal{L}(X|y)\pi(y)}{\mathcal{L}(X|\theta_i)\pi(\theta_i)} \stackrel{?}{\geq} r \sim \mathcal{U}(0, 1), \quad (2.18)$$

is equivalent to checking

$$\log \alpha = [\log \mathcal{L}(X|y) + \log \pi(y)] - [\log \mathcal{L}(X|\theta_i) + \log \pi(\theta_i)] \stackrel{?}{\geq} \log r \sim \text{Exp}(1), \quad (2.19)$$

where $\text{Exp}(1)$ is the standard exponential distribution. This can be a more efficient way of computing the acceptance rate if computing the log-likelihood function is calculated.

After sufficient sampling, the probability density function $\mathbb{P}(\theta|X, M)$ can be estimated from the distribution of the accepted samples in the chain $\{\theta_1, \dots, \theta_n\}$. Once this probability density function is known, summary statistics of this can be derived, such as the mean and credible intervals of the parameters. The purpose of MCMC sampling is thus to estimate the PDF and any derived statistics from samples taken of the PDF, and MCMC sampling with the Metropolis–Hastings algorithm is a method to properly sample the PDF.

Figure 2.3 shows a simple, pedagogical example of a 2-dimensional probability distribution that is sampled with the Metropolis–Hastings algorithm. The figure shows a two-dimensional plane, with a two-dimensional multivariate normal probability distribution in blue, that is sampled with the Metropolis–Hastings algorithm in the red points. Notice how the chain was started far away from the center of the PDF, but after a relatively short amount of steps, the chain converged towards the center of the PDF and started sampling a region of high probability. After sampling 200 data points, the resulting estimate of the PDF is shown in red, which follows a similar shape (but slightly underestimates the scale) of the true PDF.

Proposal Updating and Burn-in

Since for a Markov chain model, any steps taken can only rely on the current state of our chain, the distribution $\mathcal{G}(y)$ can only rely on our current chain step θ_i , hence $\mathcal{G} = \mathcal{G}(y|\theta_i)$. It is common for this distribution to pick a multivariate Gaussian distribution with some covariance Σ that we know or assume is close to our posterior distribution covariance.

Updating the proposal distribution is a useful tool in converging to a proper posterior distribution. A method often employed is to start with some guess for the




Figure 2.3. An example of a 2-dimensional posterior distribution sampled with a markov chain monte carlo method. The 1 and 2σ confidence levels of the **blue ellipses** show the true likelihood distribution of the parameters. The **red points** shows the 200 random samples drawn by following the Metropolis-Hastings algorithm. The chain started in the lower right corner at $(10, -8)$, far outside the posterior distribution, and within 17 accepted samples, the chain reached the 95%-confidence level of the posterior distribution. The proposal distribution was an uncorrelated multivariate Gaussian. After removing the first one-third of the chain, the inferred posterior is shown in the **red ellipses**.

proposal distribution, and, after a period of sampling, updating this guess based on the chain of accepted samples so far.

Without knowledge of the posterior distribution, it is difficult to know what to use as a proposal distribution. In some cases, the approximate shape of the proposal can be guessed, for example it is known that for CMB observations, the Hubble parameter H_0 and cold dark matter density $\Omega_c h^2$ are quite correlated, so any proposal that captures that correlation will be better than one that does not. It might also be the case that sub-blocks of the covariance matrix are known, for example if the covariance of the foreground parameters is known for one cosmological model, then it can still be used for inference with a different cosmological model, as the foregrounds are usually not majorly correlated with the cosmological parameters.

When no proposal distribution is given, `cobaya` has the option to make an educated guess the shape of the proposal distribution based on the known covariance of parameters from cosmological observables. Alternatively, `cobaya` uses a fully diagonal covariance matrix, of which the entries are given by the user. For these cases, it is often useful to underestimate the proposal distribution, as a proposal distribution that is too wide will lead to a poor acceptance rate initially, due to a lot of samples being rejected.

This does mean that initially, when our guess distribution is poor, our chain might not be sampling proportionately from the true posterior distribution. As a result, it is common practice to let our chains run for a while first so they can start exploring parameter space first, and then prune part of the chain at the starting so-called **burn-in period** and consider only the later part of the chain in posterior evaluation. It is common practice to discard a fraction of samples from the beginning of the chain, typically $0.3 \sim 0.5$.

Convergence Checking

An MCMC is **converged** when the distribution of the accepted samples $\{\theta_i\}$ is sufficiently close to the posterior distribution $P(\theta|X)$. Since of course the latter is not known a priori, it raises a question on how to determine whether our distribution has converged or not.

One method would be to visually inspect the posteriors of the chain and, if they are not visually smooth, continue running the chain for longer. Another method might be to sample the posterior probability with multiple chains, and having them start in different regions of parameter space, so that the initial exploration and burn-in phase will sample different parts of the posterior probability density function.

`cobaya` uses the method described in Lewis (2013), where a single statistic, the **Gelman-Rubin statistic** or $R - 1$ is computed over part of the entire chain. This value represents the largest orthonormalized variance between chains with respect to the means. The $R - 1$ statistic decreases as the chain reaches convergence, and a target value of $R - 1 \sim O(10^{-2})$ is a reasonably good aim for convergence. As a lower value indicates a better estimate of the posterior distribution and the covariance of the parameters, whenever the $R - 1$ statistic drops, `cobaya` also updates the proposal distribution \mathcal{G} of the chain to a multivariate Gaussian with a covariance matrix of the chain up to this point.

2.3.1 Maximum Likelihood Estimation and Model Selection

Often, we are also interested in the Maximum Likelihood Estimate (MLE) $\hat{\theta}$ of our model. Mathematically speaking, the MLE is simply the set of parameters that maximizes the likelihood function $\mathcal{L}(\theta|X, M)$. We are interested in the MLE because it gives the values of the parameters that give the best fit of the model to our data. Based on this MLE, we seek to find a model selection criterion, which is simply a score we assign to a model that, when compared to the score of a different model, provides us with a means to see how well these two different models compare to each other. A good model selection criterion will choose a model that finds the right balance between the complexity of a model and the goodness of the fit to the data.

We can estimate the MLE by running a χ^2 -minimizer (or a log-likelihood maximizer). Finding such a minimum or maximum can be tricky, and much work has been done in developing algorithms that can find good estimates without needing to compute the derivatives of the log-likelihood function, while avoiding getting stuck in local minima or maxima. Within `cobaya`, the `BOBYQA` algorithm is used to find the minimum χ^2 , or the MLE, of a log-likelihood function (Cartis et al., 2018, 2021; Powell, 2009).

There is no simple model selection criterion that unambiguously proves one model is better than another. The posterior odds ratio, $\mathbb{P}(M_1|X)/\mathbb{P}(M_2|X)$, is the quantity that we are interested in, but computing this requires knowledge of the evidence $\mathbb{P}(X|M)$, which is often time-consuming to estimate directly for a given model. As such, we rely on approximate methods to perform this task. For cosmological and astrophysical model selection, we chose to use the **Bayesian Information Criterion** (BIC), which was found to be a good choice for selection different parameterizations (Liddle, 2004), and is defined[‡] as:

$$\text{BIC}(M) = 2 \log \left(\mathcal{L}(\hat{\theta}|X, M) \right) - |M| \log N, \quad (2.20)$$

where M is the model that describes our theory $\mathcal{D}_b^{\text{th}}(\theta, M)$, $\mathcal{L}(\hat{\theta}|X, M)$ is the MLE of the model, $|M|$ is the number of parameters describing M (thus the number of parameters in θ), and N is the number of data points in our measurement (Bhat and Kumar, 2010).

If we compare the BIC of two different models, then the model with the highest BIC is the model that fits the data best. Notice that for a model with a lower MLE, the BIC may still end up being higher if the number of parameters is lower than that for a

[‡]Some definitions of the BIC have a factor 2 difference, or (as is the case for Liddle (2004)) may differ in sign convention.

model with a higher MLE. The second term in the definition of the BIC acts as a penalty for adding extra parameters to the model.

2.4 Current Constraints on Λ CDM

Here I give an example of how we can take a state-of-the-art measurement of the CMB, along with the functions and methods described above, to get estimates of cosmological parameters for the Λ CDM model.

In 2018, the *Planck* constraints on Λ CDM parameters, as derived from their analysis of the TT, TE, and EE CMB power spectra, were published in Planck Collaboration (2020d). These constraints were obtained from the all-sky CMB in T and E polarization, with measurement of the TT, TE, and EE power spectra between $30 \leq \ell \leq 2508$. The large-scale ($\ell \leq 29$) temperature analysis was performed by the `Commander` likelihood, which includes non-Gaussian corrections needed for these large scales (Planck Collaboration, 2020a). An additional separate analysis of the EE power spectrum in the range $2 \leq \ell \leq 30$, called the `lowE` analysis, was included, which is needed to constrain the optical depth at reionization τ_{reio} .

Table 2.1 gives an overview of the constraints on Λ CDM cosmology from the *Planck* 2018 analysis. It shows the constraints on the six parameters that were varied in their cosmological model, as well as six derived parameters that were computed from these constraints. The data observed by the *Planck* satellite was well described by Λ CDM, with good constraints on all six parameters from this model: apart from τ_{reio} , all parameters were measured with 1% precision or less. To date, this measurement of the Λ CDM model from *Planck* is among the most stringent constraints of the entire Λ CDM model from the CMB alone. By itself, *Planck* provided high precision measurements of all six parameters of the Λ CDM model, while also putting tight constraints on the possibilities for extension models.

In the absence of any constraints on the parameters beyond the six of Λ CDM, we often keep these beyond- Λ CDM parameters fixed to some fiducial parameter that we base on either an assumption or other measurements of physics. As such, unless otherwise specified, we keep the spatial curvature Ω_k , the running of the spectral index $dn_s/d \ln k$, and the tensor-to-scalar ratio r all fixed to zero. We also keep the sum of the neutrino masses Σm_ν to 0.06 eV based on the lower bound found in neutrino oscillation experiments, and the effective number of relativistic species N_{eff} to 3.044 based on calculations from the Standard Model of particle physics. Despite the good fit of Λ CDM to the *Planck* data, there is still room for exploration of several extensions of the Λ CDM model. In the lower part of Table 2.1, I summarize several extensions

Parameter	<i>Planck</i> constraint	
$\Omega_b h^2$	$(2.233 \pm 0.015) \times 10^{-2}$	
$\Omega_c h^2$	$(11.98 \pm 0.12) \times 10^{-2}$	
$100\theta_{\text{MC}}$	1.04077 ± 0.00047	
τ_{reio}	$(5.40 \pm 0.74) \times 10^{-2}$	
n_s	0.9652 ± 0.0042	
$\log(10^{10} A_s)$	3.043 ± 0.014	
H_0	$(67.37 \pm 0.54) \text{ km/s/Mpc}$	
$100\theta_*$	1.04097 ± 0.00046	
z_{reio}	7.50 ± 0.82	
Ω_m	0.321 ± 0.013	
σ_8	0.8118 ± 0.0089	
S_8	0.840 ± 0.024	
Ω_κ	0	$-0.044^{+0.033}_{-0.034}$
Σm_ν	0.06 eV	$< 0.257 \text{ eV}$
N_{eff}	3.044	$2.92^{+0.36}_{-0.37}$
$dn_s/d \ln k$	0	-0.006 ± 0.013
r	0	< 0.107

Table 2.1. (Top half): The best-fitting parameters for Λ CDM as found by Planck Collaboration (2020d). The parameters in this case are: the observed densities in baryonic matter $\Omega_b h^2$ and cold dark matter $\Omega_c h^2$; the approximate angular scale at last scattering, θ_{MC} ; the optical depth at reionization τ_{reio} ; the power law index n_s and initial amplitude A_s of primordial fluctuations. **(Middle):** Various other parameters that can be derived from the first six, and their inferred constraints as found by *Planck*: the present-day Hubble parameter H_0 in units of km/s/Mpc; the angular scale at last scattering, θ_* ; the redshift of the midpoint of reionization z_{reio} ; the matter density Ω_m ; and the matter clustering parameters σ_8 and S_8 . **(Bottom half):** Extensions to the Λ CDM model, their assumed fixed values in Λ CDM, and their constraints in extended cosmologies (see section 1.6 for descriptions of these models): the density of curvature Ω_κ ; the sum of the mass-eigenstates of neutrinos Σm_ν in units of eV; the effective number of relativistic species N_{eff} ; the running of the primordial spectral index $dn_s/d \ln k$; and the tensor-to-scalar ratio r .

to the Λ CDM model, their assumed (fixed) values, and constraints found by *Planck* in extended cosmological analyses. I have described these extension models before in section 1.6.

These *Planck* constraints showed here were derived from the `plik` likelihood function. In the next chapter, specifically in section 3.4, I will delve into details of this `plik` likelihood function, as part of our work in the re-analysis of the *Planck* data with the *Simons Observatory* pipeline.

Likelihood Software for new CMB Analyses

In this chapter, I discuss the development of modern likelihoods for CMB data. This involves a discussion of legacy likelihood design, as well as modern, state-of-the-art likelihood implementations. For the most part, this chapter will be parallel to the mathematical description I provided in chapter 2, and will instead focus on the software implementation side of these likelihoods.

I will describe likelihood implementations for data from the *WMAP*, *Planck*, ACT, and SO experiments. I have previously described these experiments, and their observations, in more detail in section 1.7.

- In section 3.1, I will give an overview of the requirements for likelihood software, and a schematic overview of their desired functionality;
- In section 3.2, I will present the single-spectrum *WMAP* likelihood and my modern reimplementation of the *WMAP* v5 likelihood in `pyWMAP`;
- In section 3.3, I will present the `MFLike` likelihood for the Simons Observatory, as an example of modern likelihood design;
- In section 3.4, I will show how we adapted `MFLike` for our reanalysis of the *Planck* PR3 data, and recovered the same cosmological constraints.
- In section 3.5, I will present results from applying the `MFLike` likelihood on simulated data vectors for the Simons Observatory, and methods we employed to validate the likelihood function;
- Finally, I will end in section 3.6 with the design for the foreground marginalization procedure to compress the multifrequency data into a smaller CMB-only dataset,

and the design for the likelihood for this data, and I will provide an example of this procedure applied to a simulated SO data vector.

The following papers form parts of this chapter:

- **The Simons Observatory: impact of bandpass, polarization angle and calibration uncertainties on small-scale power spectrum analysis**, *S. Giardiello, M. Gerbino, L. Pagano, et al., (2024)*;

My contributions for this paper consist of aiding in the development of the `CosmoPower` framework (which I describe in more details in chapter 4) and `MFLike` likelihood software (which I describe later in this chapter in subsection 3.3.1 and section 3.5), and their applications within the inference framework used in this paper. I also contributed as a member of the SO collaboration on the development of the SO analysis software used within these analyses. I adapted parts of subsection 3.3.1 and section 3.5 after this paper.

- **The Atacama Cosmology Telescope: DR6 Gravitational Lensing Map and Cosmological Parameters**, *M. S. Madhavacheril et al., (2024)*;
- **The Atacama Cosmology Telescope: A Measurement of the DR6 CMB Lensing Power Spectrum and Its Implications for Structure Growth**, *F. J. Qu et al., (2024)*;

My contributions for these papers consist of my implementation of the `pyWMAP` software (which I describe later in this chapter section 3.2), which was used for the ACT +WMAP analysis within these papers. I also contributed as a member of the ACT collaboration on the development of the ACT analysis software used within these analyses.

- **The Simons Observatory: a new open-source power spectrum pipeline applied to the Planck Legacy Data**, *Z. Li, T. Louis, E. Calabrese., and H. T. Jense et al., (2021)*;

My contributions for this paper involve writing the *Planck* likelihood within the `MFLike` software. My work forms sections 5 and 6 in this paper, including figures 11, 12, and 13. I adapted section 3.4 after the corresponding sections from this paper, and included the figures as Figure 3.4, Figure 3.5, and Figure 3.6 respectively.

3.1 Likelihood Design

Cosmological data releases from CMB experiments include making publicly available software required for likelihood analyses. Over the past two decades, a full analysis was often done with bespoke software for each experiment, specifically designed for the instrument which data was being analysed. Elements were often hard-coded or difficult to modify for explorations beyond the baseline case.

For example, in 2012, the *WMAP* mission released its final, 9-year data release. To compute the best-fitting cosmology and parameter constraints, the accompanying `WMAPv5` likelihood was released. This likelihood was written as a *Fortran* module containing the necessary functions to compute the covariance matrix and high- ℓ ($\ell > 32$) spectra differences, as well as the low- ℓ ($\ell < 32$) likelihood from the pixel data and noise estimates. This likelihood required to be manually compiled to be incorporated in the sampling software `CosmoMC`, which was also written in *Fortran*, and could be cross-linked to theoretical Boltzmann codes such as `camb` by including them in the compilation unit as well. After compilation, parameters and run settings could be passed to `CosmoMC` for the particular cosmology one was interested in by modifying a simple plain text setup file.

However, for performing different data cuts, for example if one was interested in the constraints from the temperature-only data from *WMAP*, or if one wanted to compare if the low- ℓ pixel-likelihood yielded the same or similar results as the low- ℓ Gibbs sampled bandpowers, one would have to modify the underlying source code manually and recompile the entire unit to re-run this experiment for comparisons.

When running one likelihood, this would not be necessarily difficult or tedious, and most software engineers would hopefully release human-readable and well-documented code that was easily modified to suit one's needs. For extended explorations when many runs are needed, and considering a general user, including students and other cosmologists, this approach becomes difficult and prone to mistakes or limitations.

While some aspects of the likelihood are indeed experiment-specific, many components are common to all datasets and to be able to combine multiple experiments it is imperative to streamline modifications, combinations, and choices of options. This need for multi-experiment analysis has driven us to a more modular, readable, and accessible design. We have transitioned from writing experiment-specific, single-purpose, *Fortran* likelihood software to modern `python` implementations, which interface with modern sampling software such as `cobaya`. These new sampling software are capable of interfacing with multiple likelihood and theory codes under variable

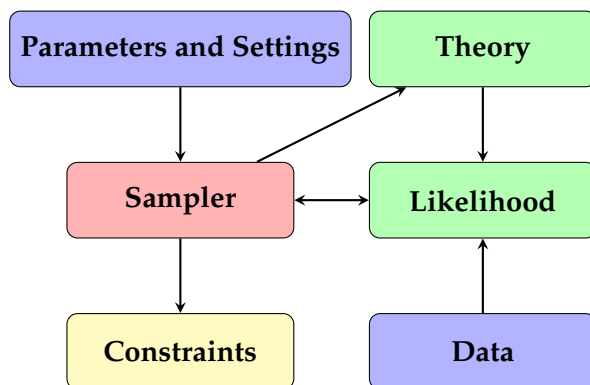


Figure 3.1. A schematic overview of the components of a modern likelihood sampling interface. The **sampler** software is at the center: it loads in any collection of **parameters and settings** and interfaces with the needed components of the code. The **theory** codes compute any kind of intermediate observables (for example: computing the CMB power spectra from cosmological parameters, or modeling some kind of foreground model from astrophysical and/or cosmological parameters). The **likelihood** function then takes in the theory observables and other parameters from the sampler, and returns the $\log \mathcal{L}$ of the model given the **data**. The sampler then provides results which allow one to compute parameter **constraints** and results.

parameters and settings, without the need for the end-user to directly edit the underlying code bases. The new `python` software packages can be installed on the user-side with minimal hassle. An example of how a modern likelihood-to-parameter framework works is shown in Figure 3.1.

My work during my research has included reimplementing legacy likelihoods in modern analysis frameworks, and developing new likelihoods for current and future experiments. We sought to update and recreate existing legacy software, such as the *WMAP* (see section 3.2) and *Planck* (see section 3.4) likelihoods – which are still of interest for future analyses – in modern, up-to-date environments. Alongside this we have been building the infrastructure for current (ACT DR6, see chapter 5) and future (SO, see section 3.5) ground-based observations.

We reflect back on chapter 2, notably on Equation 2.3, where I prescribed the general likelihood function we wish to implement for our dataset. For most of the work in this thesis, we focus on the $\ell > 30$ regime of the power spectra, where the CMB is highly Gaussian in nature. As such, the Gaussian likelihood function given in Equation 2.3 is applicable to our statistical model for the likelihood.

For a single-frequency dataset, our instrument measures the CMB at a single frequency in temperature and E-mode polarization. The full data vector will have the form $\hat{D}_b^d = \{\hat{D}_b^{TT}, \hat{D}_b^{TE}, \hat{D}_b^{EE}\}$, where b is the bin of each datapoint, and we have the

error for and covariance between each data point, described by a covariance matrix Σ_{b_1, b_2} .

Our log-likelihood will be a function $\log \mathcal{L}(\hat{\mathcal{D}}_b^d | \theta)$, where $\theta = \{\Omega_b h^2, \Omega_c h^2, \theta_{MC}, n_s, \log A_s, \tau_*, \dots\}$ is the set of cosmological parameters, plus any kind of potential nuisance parameters that we need to include. This function takes these parameters and computes the log-likelihood of the data given the model. For this, such a likelihood commonly takes the following steps:

1. Compute the theoretical CMB vector $\mathcal{D}_\ell^{\text{CMB}} = \{\mathcal{D}_\ell^{TT, \text{CMB}}, \mathcal{D}_\ell^{TE, \text{CMB}}, \mathcal{D}_\ell^{EE, \text{CMB}}\}$ from the cosmological parameters with an Einstein-Boltzmann code;
2. Bin each spectrum with a given window function w as $\mathcal{D}_b^{XY, \text{CMB}} = \sum_\ell w_{b\ell} \mathcal{D}_\ell^{XY, \text{CMB}}$;
3. Include any additional systematics, such as calibration uncertainty, that modifies the spectrum (see subsection 2.2.2); I will denote it here as a generic function $\mathcal{D}_b^{\text{th}}(\theta) = f(\mathcal{D}_b^{\text{CMB}}, \theta)$, which includes any other parameters θ that we consider in our model;
4. Compare the forward theory model with the data and the covariance matrix, and calculate the likelihood $\log \mathcal{L} = -\frac{1}{2} [\hat{\mathcal{D}}_b^d - \mathcal{D}_b^{\text{th}}(\theta)] \Sigma^{-1} [\hat{\mathcal{D}}_b^d - \mathcal{D}_b^{\text{th}}(\theta)]$.

The exact nature of the function f will depend on the specific experiment setup, and the treatment of foreground nuisances. However, for different experiments, each of these steps can include modules which are shared between experiments, and the software for doing this capable of doing these steps in order for multiple experiments at the same time.

We can see these actions enumerated here reflected back in Figure 3.1. Computing the theoretical CMB forward model vector is performed by the **theory** code. This theoretical model is passed on to the **likelihood** code, which adds additional systematics effects and other model components to the full theoretical model vector, and compares it with the observed **data** to compute the log-likelihood function. The **sampler** code is then responsible for sampling the log-likelihood function across parameter space, and infer any **constraints** from the results.

3.2 A Re-implementation of the WMAP 9-Year Likelihood

The *Wilkinson Microwave Anisotropy Probe* (WMAP) was an all-sky satellite mission that observed the CMB for nine years between 2001 and 2010. The 9-year data release, presented in Hinshaw and et al. (2013); Bennett and et al. (2013), included all-sky maps

in five bands between 20 and 95 GHz. It provided a measurement of the CMB power spectrum in temperature up to $\ell \approx 1200$, and in the temperature-polarization cross-spectrum up to $\ell \approx 1000$.

Despite its age and the release of more precise data, the creation of a `python` implementation of the *WMAP* likelihood has useful benefits. Primarily, for ground-based observatories, the limited sky coverage provides limited measurement of the larger angular scales, importantly the first acoustic peak in the range $2 < \ell < 400$ which is essential in constraining the baryon density $\Omega_b h^2$ and the primordial spectral index n_s . Because *WMAP* provides an independent, signal-dominated measurement of this part of the power spectrum, the usage of *WMAP* as opposed to *Planck* allows for an independent measurement of the CMB as a means to verify the *Planck* results. In addition, *WMAP* is noise-dominated in the regime where the signal from current ground-based telescopes begins, meaning that the covariance between *WMAP* and ground-based data is negligible. This allows for the inclusion of *WMAP* data without having to cut the data from ground-based telescopes, nor the inclusion of a deeper re-analysis of the covariance between *WMAP* and newer data. Finally, the *WMAP* signal is dominated by the CMB signal, with limited contamination from foregrounds. A single amplitude for the SZ emission in temperature is used for the analysis. This means that the parameter volume does not increase much by combining with *WMAP*.

My reimplementaion of the *WMAP* likelihood, released as `pyWMAP`, was used in the ACT DR6 lensing analysis, see Madhavacheril et al. (2023). This implementation closely follows the original *WMAP* v5 likelihood, presented in Bennett and et al. (2013), but does not include the low- ℓ components in TE, or EE (data below $\ell = 24$). The full code for `pyWMAP` is publicly available online at <https://github.com/HTJense/pyWMAP>.

For the high- ℓ part, the parameter estimation follows Hinshaw et al. (2003), where they use the Gaussian estimation of the power spectrum \mathcal{D}_ℓ s. The binned power spectrum is assumed to have the form

$$\mathcal{D}_b = w_{b\ell} \mathcal{D}_\ell + N_b, \quad (3.1)$$

where N_b is a binned uncorrelated noise component. We did not re-estimate the noise curve, but instead re-used their estimate of

$$N \approx N_0 + N_1 \left(\frac{450}{\ell} \right), \quad (3.2)$$

where N_0 and N_1 are fixed coefficients. These coefficients were found by the *WMAP* team by assuming the measurements of one autospectrum in two different years would be the same, apart from its noise measurement. They found the coefficients of N by assuming two different years of measurements would contain the same CMB power spectrum but independent noise power spectra. Thus, by cross-correlating different years, one can cancel out the noise and be retained with only the CMB power spectrum, and thus the noise curve would be,

$$\langle N \rangle = \frac{1}{2} \sum_{j_1, j_2} [\mathcal{D}_\ell^{j_1 j_1} - \mathcal{D}_\ell^{j_1 j_2}], \quad (3.3)$$

where the indices j stand for the different years of observations, and $\mathcal{D}_\ell^{j_1 j_1}$ stands for the autospectrum of one year with itself, while $\mathcal{D}_\ell^{j_1 j_2}$ stands for the cross-spectrum of one year of observations with a different year of observations (Hinshaw et al., 2007). The values for N_0 and N_1 were then found by fitting Equation 3.2 to this estimator over the range $33 \leq \ell \leq 1024$.

We then compute the covariance matrix by estimating the uncertainty per \mathcal{D}_ℓ as

$$\Delta \mathcal{D}_\ell = \frac{1}{f_{\text{sky}}} \sqrt{\frac{2}{2\ell + 1}} (\mathcal{D}_\ell^{\text{fid}} + N_\ell), \quad (3.4)$$

where $\mathcal{D}_\ell^{\text{fid}}$ is a fiducial model for the CMB and foreground power spectrum. This gives a good estimate for cosmic variance, instrumental noise, and mode coupling. We see that this equation is similar to our cosmic-variance limit given in Equation 1.29, but with an added term for the uncorrelated noise. An additional component needs to be added to the covariance for the effects of beam deconvolution and uncertainties in point-source subtraction, which is a term I did not re-estimate for my work and follows the same treatment as (Hinshaw et al., 2003). This adds a slightly off-diagonal component to the covariance matrix.

In the regime where *WMAP* measures the microwave background, at $\ell < 1200$, the signal is dominated by the CMB. Separate cleaning is done to remove thermal dust emission from the galaxy, and apart from that, the only remaining contamination comes from the Sunyaev-Zel'dovich effect. To model this, a single template spectrum with a freely-varying relative amplitude is added to the theory curve, as

$$\mathcal{D}_\ell^{\text{th}} = \mathcal{D}_\ell^{\text{CMB}} + A_{\text{SZ}} \mathcal{D}_\ell^{\text{SZ}} \quad (3.5)$$

with $0 \leq A_{\text{SZ}} \leq 2$ for the amplitude of the SZ component relative to the fiducial value. It was found that marginalizing over the SZ amplitude affects the cosmological




Figure 3.2. The marginalized Λ CDM cosmological parameters (**blue**) recovered with my `python WMAP` likelihood implementation versus the (**black**) marginals from the original `WMAPv5` likelihood constraints. I show that my likelihood implementation perfectly recovers the parameter constraints from the *WMAP* measurements of the CMB.

parameters by $\lesssim 0.4\sigma$, while the amplitude of the SZ template itself was unconstrained.

For the low- ℓ temperature analysis, I reimplemented the analysis of the Gibbs sampling method which is described in (Dunkley et al., 2009). This analysis uses a Blackwell–Rao (BR) estimator for computing the likelihood, based on some initial estimate of the data $\mathcal{D}_{\ell s}$.

For my reimplement of the *WMAP* likelihood, we verified that the resulting log-likelihood as well as some posterior parameter constraints matches those of the original *WMAP* analysis. We verified that for a fiducial CMB power spectrum, the chi-square of the code is sufficiently close to the original code, confirming that $\Delta\chi^2 \approx 2 \times 10^{-4}$ (where the absolute $\chi^2 \approx 7557.96$) for the theoretical power spectrum that was packaged with the original *WMAP* v5 likelihood. In addition, we also verified that individual components of the likelihood resulted in a chi-square value that was similarly

close to the original code. The remaining differences can be down to computational differences, such as a different algorithm being used for matrix inversion between the two codes. In Figure 3.2, I compare the 1D posteriors of the constraints recovered with my *WMAP* likelihood versus the original *WMAP* likelihood ran with the same settings.

3.3 Multi-Frequency Spectra Likelihoods

As mentioned in the previous chapters, observations of the CMB power spectra get contaminated by astrophysical foregrounds, and this combined signal gets affected by instrumental noise properties. This means that present and future observations of the CMB will require methods to separate these nuisances from the CMB signal and extract cosmological information from the T and E mode data.

With the advent of new, high-precision observations from telescopes such as the Simons Observatory, there is an ongoing drive to go from older single-purpose likelihoods to more modern, open, and modular designs that allow for better access to underlying data products and analyses.

As mentioned before in subsection 2.1.2, when observing the microwave sky in multiple frequencies, multiple cross-spectra between different frequency pairs are possible. This means that the full data vector will no longer just be single TT, TE, and EE spectra observations, but instead each spectrum will appear multiple times in the data vector, each for different cross-frequency pairs.

In this section, I will go through the implementation of `MFLike`, a Gaussian **M**ulti**F**requency **L**ikelihood software that is under development for the future SO LAT analysis (see subsection 1.7.3 for a summary of SO and the LAT), and has been applied for a re-analysis of legacy *Planck* data (section 3.4), as well as the analysis of the upcoming ACT DR6 data (chapter 5).

With the design philosophy of modularity and interoperability, `MFLike` consists of multiple components, with the intention that, for example, the computation of power spectra for astrophysical foregrounds exists in a standalone library that interfaces with not just the multifrequency likelihood, but with other components of the SO (or other experiment) pipelines as well. I will discuss each of the components relating to the `MFLike` software suite, parallel to the discussion of the likelihood function described in the previous chapter.

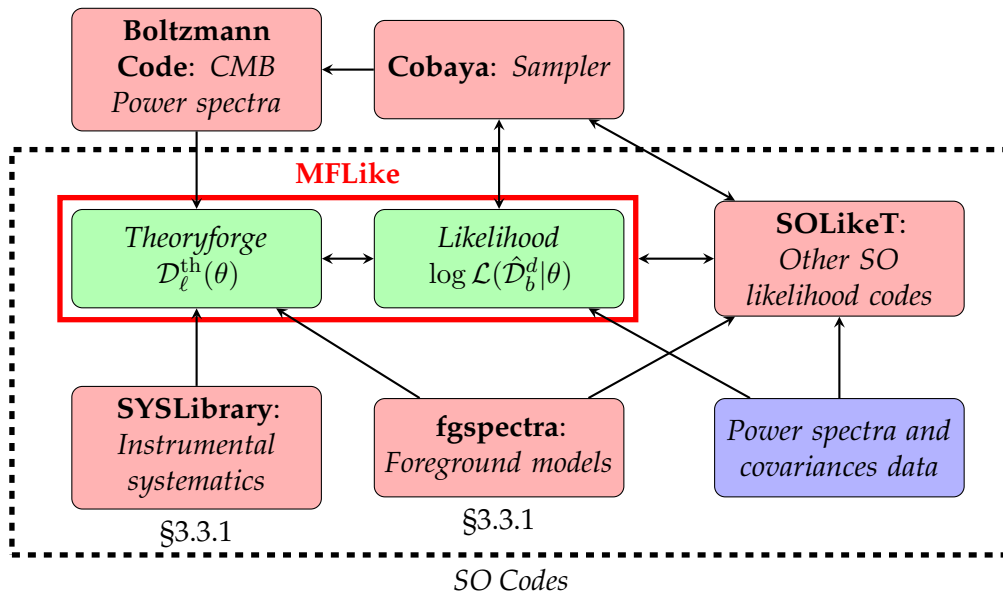


Figure 3.3. An overview of the various components of the SO analysis pipeline (**black, dashed box**), the main components of MFLike (**red, solid box**), and how each of these interface with each other and various external components. The full package of MFLike consists of the *Theoryforge* (**green**), which takes in the various theoretical components and builds a model for the data vector, and the *Likelihood* (**green**), which compares the theory model to the data and computes the log-likelihood function. The *Theoryforge* interfaces with the *Boltzmann code* to obtain theoretical CMB power spectra, as well as the SO softwares *fgspectra* for power spectra of astrophysical foregrounds, and *SYSLibrary* for instrumental systematics modeling. The combined power spectrum model is then passed to the *Likelihood*, which compares it to the *Data* (**blue**) computes the log-likelihood function, passing this to *cobaya*. The SO analysis pipeline also contains codes for computing log-likelihoods from other probes (such as CMB lensing, or cross-covariance with galaxy clustering), and each of these codes can interface with each other to compute joint constraints on cosmology.

3.3.1 An Overview of the MFLike likelihood

A schematic overview of all components that MFLike interfaces with is given in Figure 3.3. The main code base of MFLike is split into two components: the *Theoryforge*, which computes the theory model given the parameters, and the *Likelihood*, which computes the log-likelihood given the theory model. The software package formed of these two components allows one to compute cosmological constraints from multifrequency CMB data. The full code for MFLike is available online at https://github.com/simonsobs/LAT_MFLike.

The reasoning behind this high level of modularity and interfacing, is the attempt at building a catalog of softwares that can handle data and models beyond just SO results.

fgspectra: Foreground Modeling

`fgspectra` is the SO library for computing multifrequency power spectra for astrophysical foreground models. The full code is available online at <https://github.com/simonsobs/fgspectra>.

Mathematically, `fgspectra` seeks to model any astrophysical foreground component as any function of the frequency-dependent part of the function $f(\nu)$; the ℓ -dependent part of the function C_ℓ ; and whatever other parameters θ the foreground model depends on. Comparing the generic idea with the various components described in subsection 2.2.1, we see that most spectra are implemented as a simple factorized cross-spectrum,

$$\mathcal{D}_{\ell,\nu_1,\nu_2}(\theta) = f(\nu_1, \theta)f(\nu_2, \theta)C_\ell(\theta), \quad (3.6)$$

with $f(\nu, \theta)$ the spectral dependency of the component, which may depend on additional parameters θ . In fact, the only component that is commonly not a simple factorized cross-spectrum is the joint tSZ-CIB model, which contains a correlation term and thus has a slightly more complicated form that includes Equation 2.13.

By writing all this in generic form rather than implementing a fixed-pipeline foreground model, `fgspectra` permits easy configuration of foreground models. With the LAT's aim at measuring high-accuracy multi-frequency cross-spectra, this will allow for investigation of the exact modeling of astrophysical foregrounds, for example for the testing models of correlations or decorrelations between different components.

At present, the publicly available library of `fgspectra` contains templates capable of describing all components described previously in subsection 2.2.1. These include ℓ space templates for tSZ, kSZ, and CIB components; frequency-dependency

terms for tSZ, modified blackbody, and power laws; and cross-spectrum models for simple factorized cross-spectra, and cross-spectra with correlations (e.g. tSZxCIB). Beyond currently used models, a wide variety of additional models is available as well, such as the ability to tilt a template in ℓ space to explore mis-modeling systematics, templates for synchrotron radiation and free-free emission, and cross-spectra for de-correlation between different components.

SYSLibrary: Instrumental Systematics

SYSLibrary is the SO library for computing multifrequency power spectrum modifications based on instrumental systematics models. The full SYSLibrary package comes with implementations for per-spectrum calibration factors, polarization angle miscalibration, and T to E leakage. The full code is available online at <https://github.com/simonsobs/syslibrary>.

For this thesis, only per-spectrum calibration is considered, and it is used for both a per-array gain calibration and polarization efficiency factor, as described before in subsection 2.2.2.

Full Theory Computation

The full theory computation is now based on a combination of (1) the theoretical CMB power spectra obtained from the Boltzmann code, (2) the astrophysical foreground model obtained from `fgspectra`, and (3) any modifications to this model due to the model applied by SYSLibrary. The full model is the sum of the CMB power spectrum, plus the foreground model integrated over the bandpass, both calibrated for gain and polarization efficiencies, and then binned according to the window functions.

MFLike is written to interface with the cosmological sampling software `cobaya`. As a consequence, it is easy to recover constraints on cosmological, astrophysical, and other nuisance parameters from the MFLike log-likelihood function through MCMC sampling (see section 2.3).

As two examples of the use of MFLike, I provide in section 3.4 an overview of our work on the re-analysis of the *Planck* 2018 data with MFLike. Afterwards, in section 3.5, I show our work within the SO collaboration on validating the MFLike likelihood on simulated data vectors for the SO LAT.

3.4 MFLike applied to *Planck* legacy data

The original *Planck* plik likelihood is the likelihood implementation for the high- ℓ multipoles from the *Planck* satellite (Planck Collaboration, 2020b). Since the CMB is sufficiently Gaussian in the $\ell > 30$ regime, it follows a simple Gaussian implementation of the likelihood function as previously described in section 2.1.

For the work we did in Li et al. (2021), we sought to reproduce the *Planck* 2018 analysis using the SO analysis pipeline. For my work, I adapted the MFLike likelihood to read the *Planck* data products and model the noise properties of *Planck*, while adapting the foreground nuisance models from `fgspectra` to include models for and interface with the *Planck* data. The full code for MFLike-plik is publicly available online at https://github.com/simonsobs/LAT_MFLike/tree/mflike-plik.

3.4.1 Likelihood Function

I developed MFLike-plik to adapt MFLike for the high- ℓ ($\ell > 30$) part of the TT, TE, and EE spectra of *Planck*. The datavector of the spectra, \mathcal{D}_b^d , and its covariance matrix Σ , were obtained by SO collaborators re-evaluating the *Planck* maps with the SO pipeline, while the likelihood was implemented by me as a Gaussian likelihood, with the log-likelihood function given in Equation 2.2. For the *Planck* likelihood, this forward model for the cross-frequency binned spectrum is written as:

$$\mathcal{D}_{b,\nu_1\nu_2}^{XY,\text{th}}(\theta) = w_{bl}c_{\nu_1}^Xc_{\nu_2}^Y \left[\mathcal{D}_\ell^{XY,\text{CMB}}(\theta_1) + \mathcal{D}_{\ell,\nu_1\nu_2}^{XY,\text{fg}}(\theta_2) + \mathcal{D}_{\ell,\nu_1\nu_2}^{XY,\text{sys}} \right] \quad (3.7)$$

where $\mathcal{D}_\ell^{XY,\text{CMB}}$ is a CMB theory prediction for a given set of cosmological parameters θ_1 , and the other two terms account for frequency-dependent foregrounds and residual systematics, respectively.

Foregrounds: $\mathcal{D}_{\ell,\nu_1\nu_2}^{XY,\text{fg}}(\theta_2)$ is the secondary signal from galactic and extragalactic emission and depends on a set of 20 foreground parameters θ_2 . We describe this model in detail in subsection 2.2.1. In temperature this term includes thermal and kinematic Sunyaev-Zel'dovich (SZ) emission, dusty star-forming and radio galaxies appearing as point sources (ps), a clustering term for the cosmic infrared background (CIB), a correlation between thermal SZ and CIB (tSZxCIB), and thermal dust emission from our galaxy (dust). In polarization the galactic dust is modeled for both TE and EE. For our implementation, we adopted the public SO library `fgspectra` that builds cross-spectra predictions for foreground components at different frequencies for a given set of model parameters.

Parameter	Description	Prior
A_{tSZ}	Thermal SZ amplitude.	≥ 0
A_{kSZ}	Kinematic SZ amplitude.	≥ 0
A_{217}^{CIB}	CIB amplitude at 217 GHz.	≥ 0
ξ	tSZ-CIB correlation scale.	$0 \leq \xi \leq 1$
A_{100}^{PS}	Unresolved radio sources in TT at 100x100.	≥ 0
A_{143}^{PS}	Unresolved radio sources in TT at 143x143.	≥ 0
$A_{143 \times 217}^{\text{PS}}$	Unresolved radio sources in TT at 143x217.	≥ 0
A_{217}^{PS}	Unresolved radio sources in TT at 217x217.	≥ 0
A_{100}^{dust}	Galactic dust amplitude in TT at 100x100.	$(8.6 \pm 2.0) \mu K^2$
A_{143}^{dust}	Galactic dust amplitude in TT at 143x143.	$(10.6 \pm 2.0) \mu K^2$
$A_{143 \times 217}^{\text{dust}}$	Galactic dust amplitude in TT at 143x217.	$(23.5 \pm 8.5) \mu K^2$
A_{217}^{dust}	Galactic dust amplitude in TT at 217x217.	$(91.9 \pm 20.0) \mu K^2$
$A_{100}^{\text{dust,TE}}$	Galactic dust amplitude in TE at 100x100.	$(0.130 \pm 0.042) \mu K^2$
$A_{100 \times 143}^{\text{dust,TE}}$	Galactic dust amplitude in TE at 100x143.	$(0.130 \pm 0.036) \mu K^2$
$A_{100 \times 217}^{\text{dust,TE}}$	Galactic dust amplitude in TE at 100x217.	$(0.46 \pm 0.09) \mu K^2$
$A_{143}^{\text{dust,TE}}$	Galactic dust amplitude in TE at 143x143.	$(0.207 \pm 0.072) \mu K^2$
$A_{143 \times 217}^{\text{dust,TE}}$	Galactic dust amplitude in TE at 143x217.	$(0.69 \pm 0.09) \mu K^2$
$A_{217}^{\text{dust,TE}}$	Galactic dust amplitude in TE at 217x217.	$(1.938 \pm 0.54) \mu K^2$
c_{100}	Calibration of the 100 GHz spectra relative to 143 GHz.	1.0002 ± 0.0007
c_{217}	Calibration of the 217 GHz spectra relative to 143 GHz.	0.99805 ± 0.00065

Table 3.1. The nuisance parameter model for the `MFLike` `plik` analysis we presented in Li et al. (2021). We also imposed the additional prior $A_{\text{tSZ}} + 1.6A_{\text{kSZ}} = (9.5 \pm 3.0) \mu K^2$.

Because the *Planck* masks are frequency-dependent, the residual dust emission is not the same in each cross-spectrum. As such, the frequency scaling described in subsection 2.2.1 is not appropriate. Instead, the dust is modeled independently per cross-frequency, and the amplitude of each template is constrained with a prior per cross-frequency which is computed individually. We adapted the standard dust template in `fgspectra` but used a custom frequency-dependent component to independently vary these amplitudes. A full overview of all nuisance parameters and their priors is given in Table 3.1.

Systematics: $\mathcal{D}_{\ell,\nu_1\nu_2}^{XY,\text{sys}}$ is a term correcting the model for residual levels of systematic effects. We added a module in `MFLike-plik` to include the frequency-dependent templates for beam leakage, sub-pixel noise corrections, and correlated noise corrections to the model vector with fixed amplitudes. We reused the model from Planck Collaboration (2020b) and did not re-estimate these templates in our analysis.

Calibration and polarization efficiencies: The full theory vector is then calibrated and corrected for polarization efficiencies c_ν^X before comparing to the data. As in Planck Collaboration (2020b), the model is calibrated assuming a fixed temperature calibration equal to unity at 143GHz, and polarization efficiencies are also fixed to 1.021, 0.966, and 1.04 for 100, 143, and 217GHz, respectively. The 100 and 217GHz temperature calibration factors are added as nuisance parameters as a third component to the θ vector and varied with the same Gaussian priors used in Planck Collaboration (2020b).

3.4.2 Tests of Parameters

We directly compared our re-estimated spectra and covariance products with the *Planck* PR3 products. We looked at the outputs of the cosmological pipeline and tested how cosmological parameters derived from our new inputs compare with those of Planck Collaboration (2020d) for Λ CDM and a select few single-parameter extension models.

To do this, I coupled our likelihood to `cobaya` (see section 2.3), extending the *Planck* parameter configuration files already existing in `cobaya` to call `MFLike-plik` for the high- ℓ spectra. We also include the pre-existing PR3 likelihood implementation for the low- ℓ *Planck* data in temperature and EE polarization, called “lowl” and “lowE” respectively.

The results of our Λ CDM run are shown in Figure 3.4, with the *Planck* PR3 results in orange and our re-analyzed `S0mode` results in blue. We find excellent agreement for

the six basic Λ CDM parameters between our inputs and Planck Collaboration (2020d), with all parameters agreeing to within 0.1σ .

The distributions of the foreground parameters are shown in Figure 3.5. In this case, all parameters are consistent with those reported in Planck Collaboration (2020d) within 0.1σ , except for the amplitudes of thermal dust in TE at frequencies higher than 100 GHz, which show small shifts of $\leq 1\sigma$. Since the dust amplitudes in TT and TE were varied with a Gaussian prior imposed, but the EE amplitudes were remained fixed, we attribute the small shifts in TE parameters to the small differences in the EE spectra at low multipoles we recovered, where the dust contamination is significant, at 143 and 217 GHz. We tested that these shifts were indeed coming from spectrum differences, rather than differences in the covariance matrix, by estimating parameters using the PR3 spectra together with our new *SOMode* covariance matrix. This combination resulted in no parameter shifts in the foreground parameters.

We also explore a standard set of Λ CDM extensions to validate our products on multiple theories and parametrizations. We estimate the constraints on the effective number of relativistic species N_{eff} , the running of the spectral index n_{run} , the amplitude of lensing smoothing in the power spectra A_L , and the spatial curvature Ω_κ . In all cases, we recover the Planck 2018 results (as seen in Planck Collaboration (2020d)) to within 0.1σ as shown in Figure 3.6.

3.5 MFLike Development and Validation for the Simons Observatory

As mentioned before, the MFLike likelihood is designed for use with the to-be-completed *Simons Observatory Large Aperture Telescope* (SO LAT), which is under construction in the Atacama Desert in Chile. The basic framework has been shown in subsection 3.3.1, and here I will focus on summarising its status for deployment on SO data.

The LAT will observe the microwave sky in temperature and polarization at 93, 145, and 225GHz. The nominal plan for sky coverage is to observe $f_{\text{sky}} \approx 40\%$, which gives a signal-dominated and cosmic-variance limited power spectrum at $\ell \leq 3000$ in temperature and $\ell \leq 2000$ in polarization (The Simons Observatory collaboration, 2019), and new observations spanning up to a much higher $\ell \sim 9000$. Cosmology, foreground cleaning, and astrophysical science from small-scale, multifrequency data will become a major component of the SO science products. Within the SO *Power Spectrum* analysis working group, I have contributed to the development of MFLike for

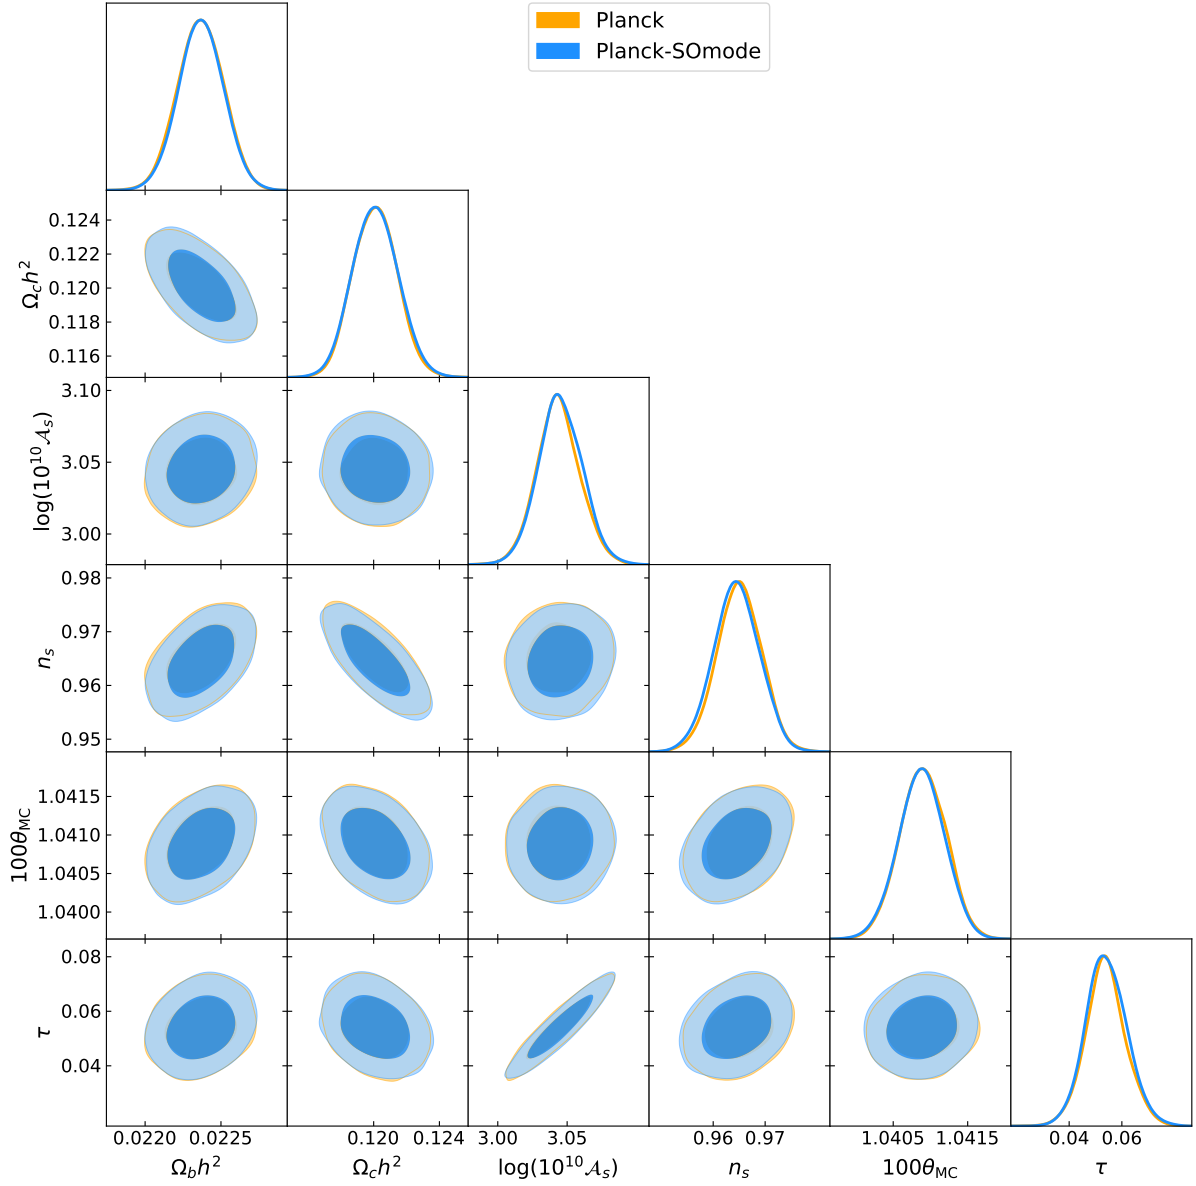


Figure 3.4. The cosmological parameters found with the MFLike-plik likelihood (**blue**) and the original plik results (**orange**). Both constraints make use of the external original likelihood for low- ℓ temperature and polarization data, which is needed to constrain τ (Li et al., 2021).

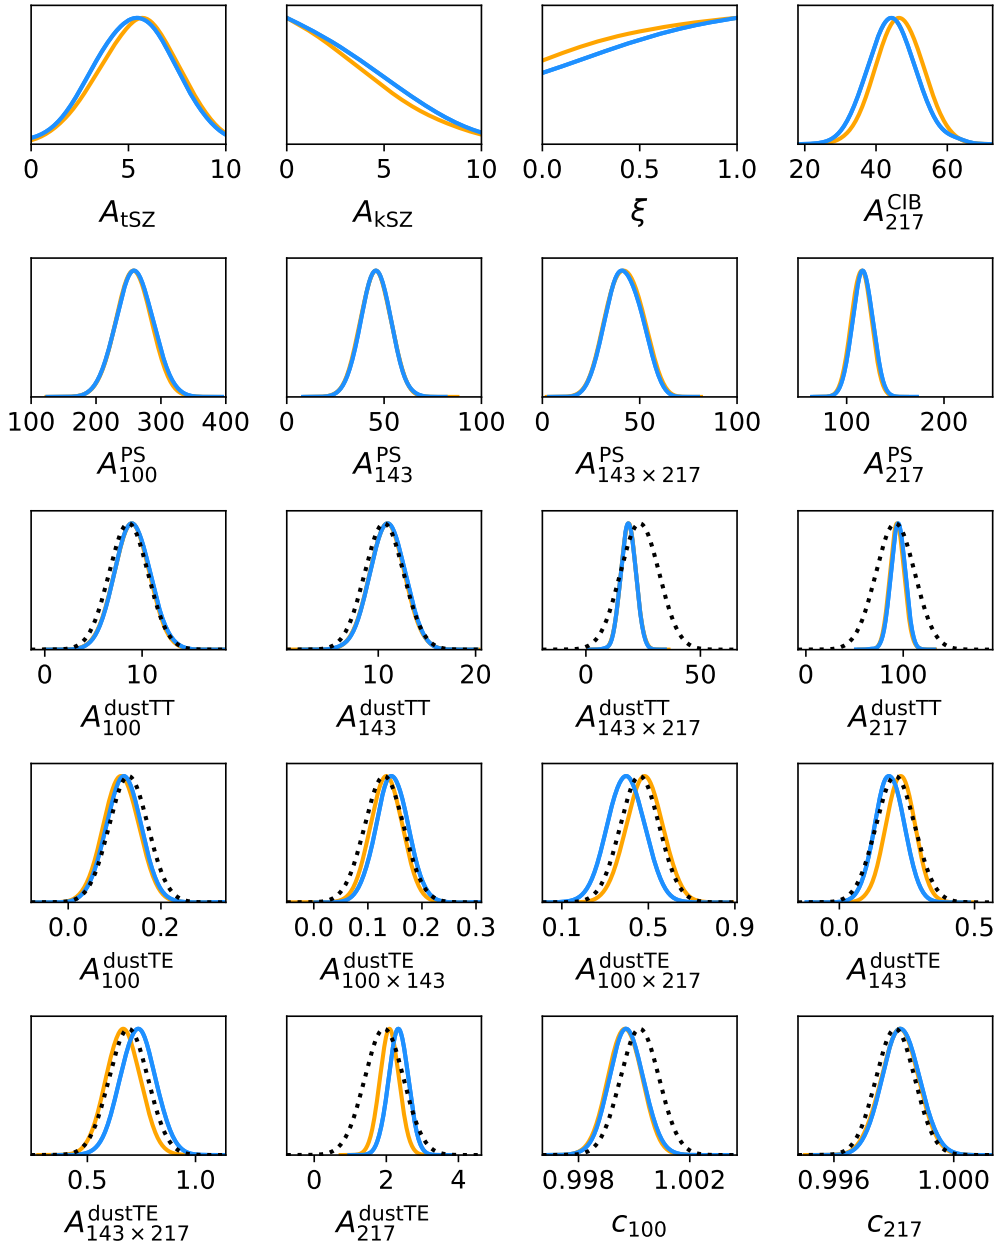


Figure 3.5. The foreground and calibration parameters found with the MFLike-plik likelihood (**blue**) versus the original plik results (**orange**). The black dotted lines illustrate the gaussian priors imposed on the dust and calibration amplitudes. An additional joint prior is used for the tSZ and kSZ amplitudes and is not shown here. The two pipelines give consistent results, with small shifts only in some of the TE dust amplitudes, which we attribute to the small differences found in the EE spectra at low multipoles (Li et al., 2021).

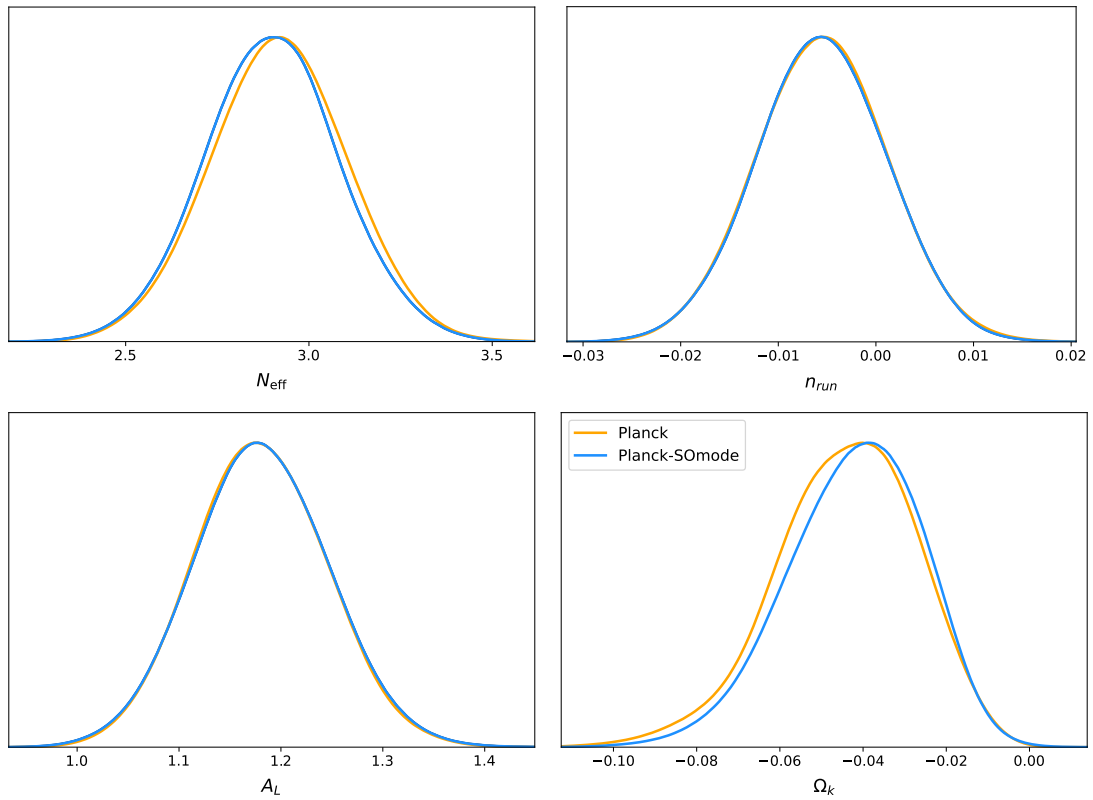


Figure 3.6. The Λ CDM extensions found with the MFLike-`plik` likelihood (blue) and the original `plik` results (orange). We show the constraints on (top left) the effective number of relativistic species N_{eff} , (top right) the running of the spectral index n_{run} , (lower left) the amplitude of lensing smoothing in the power spectra A_L , (lower right) and the spatial curvature Ω_k (Li et al., 2021).

the LAT power spectrum analysis. I have also contributed to the SO *Likelihood and Theory* package `SOLikeT`^{*}, which will make SO data easily available for cross-probe analysis, interfacing `MFLike` for power spectra to other observables such as CMB lensing, cross-correlations with cosmic shear (see e.g. Shaikh et al. (2023)), halo modeling, or galaxy cluster counts.

3.5.1 Using `MFLike` on SO simulations

After building the basic suite of software composing `MFLike`, and validating it on *Planck*, within the SO collaboration we have started extensive tests on SO simulations. As an important step to validate the functioning of the likelihood, we run the `MFLike` software on simulated data vectors for SO LAT observations. Simulated data vectors were generated from Gaussian sky realizations of a fiducial set of parameters. The simulated power spectra \hat{D}_ℓ^d and covariance were generated using `PSPipe`[†], the same software that will be used for the real data analysis when the LAT comes online. These simulated data vectors on Gaussian scales between $30 \leq \ell \leq 9000$ are then passed to `MFLike` for modeling and measurements.

At the likelihood level, we model the full multifrequency microwave signal as a combination of CMB and foreground components. The full foreground model is the same as the one described in subsection 2.2.1, with SZ, CIB, Galactic dust, and radio point source components in temperature and E-mode polarization.

The full suite of simulated data vectors is available publicly as part of the `MFLike` software.

I show the recovery of the fiducial cosmological parameters for all 100 simulated data vectors in Figure 3.7. For this figure, I ran the `cobaya` MCMC sampler, interfaced with the `MFLike` likelihood and a CMB theory code `CosmoPower` (see chapter 4), to constrain the cosmological and astrophysical foreground parameters. I did not consider any instrumental systematics in this analysis, and I refer to Giardiello et al. (2024) for a recent extensive analysis of SO simulations, including the impact of various systematics and the importance of external characterization of instrumental properties. For Λ CDM cosmological parameters, I show a good recovery of the fiducial input parameters compared to the mean of the 100 simulations.

^{*}<https://github.com/simonsobs/SOLikeT>

[†]<https://github.com/simonsobs/PSPipe>




Figure 3.7. Recovery of the fiducial cosmological parameters of a set of simulated data vectors for SO LAT. These data vectors are Gaussian realizations of a fiducial cosmology and foreground model, and do not consider any instrumental systematics for the LAT (see Giardiello et al. (2024) for a complete overview of the effect of systematic uncertainties in these simulations). I imposed a Gaussian prior on the optical depth $\tau_{\text{reio}} = (5.44 \pm 0.58) \times 10^{-2}$, since this parameter cannot be constrained by the LAT alone. Various **grey** distributions are the posterior distributions obtained from the 100 different Gaussian realizations of the data. The **black** distribution is the mean of all of these 100 posteriors, with the vertical line showing the mean of this distribution. The **red** vertical line shows the fiducial input value, and the captions show the difference between the fiducial input value and the recovered mean over 100 simulations. Overall, I make a good recovery of the fiducial cosmology, with minimal discrepancies or biases ($< \pm 0.25\sigma$) on any parameter.

3.6 A CMB-only Likelihood

For purely cosmological analysis, it is more practical to work with a CMB-only dataset. For this, some kind of subtraction of the foreground contamination is necessary. This can be done at the map level, where techniques such as component separation are used to differentiate the maps in linear space to different maps in component emission (e.g. Coulton et al. (2023)).

As a different method, we employed **foreground marginalization** at the power spectrum level to independently marginalize over the foreground contamination and produce a clean, CMB-only power spectrum for cosmological analysis. These CMB-only power spectra can then be exploited with a simple CMB-only likelihood (as done for ACT, *Planck*, and SPT (Dunkley et al., 2013; Planck Collaboration, 2016, 2020b; Choi et al., 2020)). To do this, we start from `MFLike` and sample at the same time the foreground and nuisance parameters, as well as the individual CMB bandpowers without assuming any cosmological model. For the latter step, we employ the Gibbs sampling method described in (Dunkley et al., 2013), which goes as follows.

We recall from Equation 3.7, that our data vector has the form

$$\mathcal{D}_{b,\nu_1\nu_2}(\theta) = \mathcal{D}_{b,\nu_1\nu_2}^{\text{CMB}} + \mathcal{D}_{b,\nu_1\nu_2}^{\text{sec}}(\theta_2), \quad (3.8)$$

where $\mathcal{D}_\ell^{\text{sec}}$ refers to any secondaries (foreground contamination and instrumental uncertainties) in our signal, which depend on the parameters θ_2 of our model for the secondary contamination, and we do not assume any model for the CMB power spectrum based on the cosmological parameters θ_1 . Since our data is binned, we write the spectrum in bandpowers b , as $\mathcal{D}_b = \sum_\ell w_{b,\ell} \mathcal{D}_\ell$. Since the CMB is blackbody, we know that for any ν_1, ν_2 , the CMB \mathcal{D}_b s will be the same. We can thus estimate the mean and covariance of our CMB \mathcal{D}_b s efficiently by means of a **mapping matrix** \mathbf{M} that maps the n_{cmb} elements of our CMB data vector to the n_{spec} elements of our observed data vector. This gives us:

$$\mathcal{D}_b(\theta) = \mathbf{M} \mathcal{D}_b^{\text{CMB}} + \mathcal{D}_b^{\text{sec}}(\theta_2), \quad (3.9)$$

where \mathbf{M} helps us map our blackbody CMB \mathcal{D}_b s to our multi-frequency data \mathcal{D}_b s. To this end, \mathbf{M} is a rectangular matrix that is equal to 1 if $\mathcal{D}_b^{\text{CMB}}$ contains CMB bandpowers that feed into \mathcal{D}_b , or 0 otherwise. We need to correct this however, for the fact that our CMB bandpowers are calibrated with our calibration factors c_ν^d and polarization efficiencies p_ν^d (see subsection 2.2.2). If a theoretical power spectrum is multiplied by $c_{\nu_1}^i c_{\nu_2}^j$ in our data vector, then we need to include this factor in our

mapping matrix. Schematically, if we consider the cross-spectra for arrays A and B , our mapping matrix will now have the following shape:

$$\underbrace{\begin{pmatrix} (c_\nu^A)^2 & 0 & 0 \\ c_\nu^A c_\nu^B & 0 & 0 \\ (c_\nu^B)^2 & 0 & 0 \\ 0 & c_\nu^A c_\nu^A p_\nu^A & 0 \\ 0 & c_\nu^A c_\nu^B p_\nu^B & 0 \\ 0 & c_\nu^A p_\nu^A c_\nu^B & 0 \\ 0 & c_\nu^B c_\nu^B p_\nu^B & 0 \\ 0 & 0 & (c_\nu^A p_\nu^A)^2 \\ 0 & 0 & c_\nu^A p_\nu^A c_\nu^B p_\nu^B \\ 0 & 0 & (c_\nu^B p_\nu^B)^2 \end{pmatrix}}_{\mathbf{M}} \underbrace{\begin{pmatrix} \mathcal{D}_b^{TT} \\ \mathcal{D}_b^{TE} \\ \mathcal{D}_b^{EE} \end{pmatrix}}_{\mathcal{D}_b^{\text{CMB}}} = \begin{pmatrix} \mathcal{D}_b^{TT, A \times A} \\ \mathcal{D}_b^{TT, A \times B} \\ \mathcal{D}_b^{TT, B \times B} \\ \mathcal{D}_b^{TE, A \times A} \\ \mathcal{D}_b^{TE, A \times B} \\ \mathcal{D}_b^{TE, B \times A} \\ \mathcal{D}_b^{TE, B \times B} \\ \mathcal{D}_b^{EE, A \times A} \\ \mathcal{D}_b^{EE, A \times B} \\ \mathcal{D}_b^{EE, B \times B} \end{pmatrix}. \quad (3.10)$$

To estimate our $\mathcal{D}_b^{\text{CMB}}$, we want to marginalize over all secondary parameters θ_2 , which gives us a posterior distribution

$$\mathbb{P}(\mathcal{D}_b^{\text{CMB}} | C_{b, \nu_1 \nu_2}) = \int \mathbb{P}(\mathcal{D}_b^{\text{CMB}}, \theta_2 | C_{b, \nu_1 \nu_2}) \pi(\theta_2) d\theta_2. \quad (3.11)$$

For this, Gibbs sampling is an efficient way to map out the joint distribution $\mathbb{P}(\mathcal{D}_b^{\text{CMB}}, \theta_2 | C_{b, \nu_1 \nu_2}^{\text{sec}})$. Gibbs sampling is a generalized form of the Metropolis-Hastings MCMC algorithm, where some joint distribution $\mathbb{P}(X_1, X_2)$ can be sampled even if one only knows the distributions $\mathbb{P}(X_1 | X_2)$ and $\mathbb{P}(X_2 | X_1)$. In our case, we can sample $\mathbb{P}(\theta_2 | \mathcal{D}_b^{\text{CMB}})$ with the Metropolis-Hastings algorithm, while computing $\mathbb{P}(\mathcal{D}_b^{\text{CMB}} | \theta_2)$ by direct sampling of the CMB bandpowers from the data vector.

We get a full, multi-frequency log-likelihood as:

$$-2 \ln \mathcal{L} = (\mathbf{M} \mathcal{D}_b^{\text{CMB}} + C_{b, \nu_1 \nu_2}^{\text{sec}} - C_{b, \nu_1 \nu_2})^T \boldsymbol{\Sigma}^{-1} (\mathbf{M} \mathcal{D}_b^{\text{CMB}} + C_{b, \nu_1 \nu_2}^{\text{sec}} - C_{b, \nu_1 \nu_2}), \quad (3.12)$$

For fixed θ_2 , and thus fixed $C_{b, \nu_1 \nu_2}^{\text{sec}}$, the conditional distribution for the CMB bandpowers, $\mathbb{P}(\mathcal{D}_b^{\text{CMB}} | \theta_2, C_{b, \nu_1 \nu_2})$ is given by:

$$-2 \ln \mathbb{P}(\mathcal{D}_b^{\text{CMB}} | \theta_2, C_{b, \nu_1 \nu_2}) = (\mathcal{D}_b^{\text{CMB}} - \hat{\mathcal{D}}_b^{\text{res}})^T \mathbf{Q}^{-1} (\mathcal{D}_b^{\text{CMB}} - \hat{\mathcal{D}}_b^{\text{res}}) \quad (3.13)$$

where $\hat{\mathcal{D}}_b^{\text{res}}$ is the mean of of the residual of the data minus the foregrounds, and \mathbf{Q} is the covariance of the CMB \mathcal{D}_b s. We can find these by taking the derivative

of Equation 3.12 with respect to the CMB bandpowers, assuming the CMB is a pure, Gaussian distribution with uniform flat priors, giving us:

$$\mathbf{Q} = \mathbf{M}^T \boldsymbol{\Sigma}^{-1} \mathbf{M}, \quad (3.14)$$

and

$$\hat{\mathcal{D}}_b^{\text{res}} = \mathbf{Q}^{-1} \left[\mathbf{M}^T \boldsymbol{\Sigma}^{-1} \left(\mathcal{D}_{b,\nu_1\nu_2} - \mathcal{D}_{b,\nu_1\nu_2}^{\text{sec}} \right) \right]. \quad (3.15)$$

We can draw a random sample from this Gaussian distribution by taking the Cholesky decomposition of the covariance matrix, $\mathbf{Q} = \mathbf{L}\mathbf{L}^T$, and drawing a random vector

$$\mathcal{D}_b^{\text{CMB}} = \hat{\mathcal{D}}_b^{\text{res}} + \mathbf{L}y \quad (3.16)$$

where $y \sim \mathcal{N}(0, 1)$ is a vector drawn from a standard normal distribution.

Moving back to Gibbs sampling, we can now sample the joint posterior of $\mathcal{D}_b^{\text{CMB}}$ and θ_2 by alternating sampling steps between one and the other. This changes our MCMC sampling procedure, as described in section 2.3, to the following Gibbs sampling procedure:

1. Pick a starting point $\theta_{2,1}$ and compute the bandpowers $\mathcal{D}_{b,1}^{\text{CMB}}$, this is where our chain starts sampling;
2. At our current point $\theta_{2,i}$ and $\mathcal{D}_{b,i}^{\text{CMB}}$, evaluate the posterior likelihood $\mathbb{P}(\mathcal{D}_{b,i}^{\text{CMB}}, \theta_{2,i})$;
3. Draw a new sample $\theta'_2 \sim \mathcal{G}(\theta'_2 | \theta_{2,i}, \mathcal{D}_b^{\text{CMB}})$;
4. At this point, compute the new CMB bandpowers $\mathcal{D}_b^{\text{CMB}'}$ from our data and our sample $\vec{\theta}'_2$, and add them to the chain $\mathcal{D}_{b,i+1}^{\text{CMB}} = \mathcal{D}_b^{\text{CMB}'}$;
5. Check if the log-posterior ratio $\alpha = \mathbb{P}(\mathcal{D}_b^{\text{CMB}'}, \theta'_2) / \mathbb{P}(\mathcal{D}_{b,i}^{\text{CMB}}, \theta_{2,i})$ is accepted;
6. If the point is accepted, then $\theta_{2,i+1} = \theta'_2$, else $\theta_{2,i+1} = \theta_{2,i}$;
7. Repeat steps 2–6 until we consider our chain sufficiently long or converged;

To avoid full degeneracy between the bandpower amplitudes and the calibration parameters, we often choose to fix the dipole calibration to unity and one polarization efficiency parameter to its best-fitting value we recovered from the full MCMC run. We then need to remember to include the dipole calibration and a single polarization

efficiency parameter in our final likelihood. Since it only involves two additional parameters, as opposed to a full model of two dozen astrophysical and instrumental nuisance parameters, this leads to minimal impact on parameter constraints.

3.6.1 The Likelihood for the CMB-only dataset

We are now left with a set of samples for the CMB bandpowers, and can build a likelihood to estimate the cosmological parameters from these bandpowers. To do this, we first compute the mean and covariance from the spectra, assuming they follow a Gaussian distribution, $\hat{\mathcal{D}}_b$ and Σ .

The CMB-only likelihood now needs only an input theoretical CMB power spectrum, $\mathcal{D}_\ell^{\text{CMB}}$, as well as a reduced set of nuisance parameters. Because we need to fix some calibration parameters to avoid full degeneracies, we need to include them in the CMB-only likelihood as an additional nuisance to marginalize over.

This means our entire simplified likelihood function becomes

$$\log \mathcal{L} = -\frac{1}{2} \left[\hat{\mathcal{D}}_b - c^{XY} w_{b\ell} \mathcal{D}_\ell^{XY, \text{th}}(\theta) \right] \Sigma^{-1} \left[\hat{\mathcal{D}}_b - c^{XY} w_{b\ell} \mathcal{D}_\ell^{XY, \text{th}}(\theta) \right]^T \quad (3.17)$$

with $w_{b\ell}$ the window functions of the CMB bandpowers, c^{XY} the nuisance parameters we did not yet marginalize over, and $\mathcal{D}_\ell^{XY, \text{th}}(\theta)$ the theoretical power spectrum. This has reduced the total number of parameters for $\log \mathcal{L}$ from ~ 40 down to ~ 8 , as we no longer need to marginalize over all the foreground and various instrumental systematic parameters: all of this information has been captured in the covariance matrix Σ . Additionally, the full multifrequency cross-spectrum information has been compressed by about a factor 10, as we no longer have multiple measurements of each CMB power spectrum, but instead only have one spectrum for TT, one for TE, and one for EE. This also compresses the size of the full covariance matrix by a factor 100.

3.6.2 Recovery of fiducial CMB bandpowers and cosmology

I ran a `python` implementation of the above algorithm, using a custom interface with the `MFLike` likelihood code, on one of the simulated data vectors for SO LAT. I extracted the CMB bandpowers in TT, TE, and EE up to $\ell = 7000$, and assumed the CMB bandpowers were zero above this. This is done to avoid degeneracies between the CMB and foreground components, and the assumption that the CMB is zero can safely be made in the regime where the CMB bandpowers are $O(10\%)$ of the weighted




Figure 3.8. The CMB bandpowers (**points**) for a SO simulated data vector compared to the fiducial power spectrum (**black line**). The subfigures show the TT (**blue, top left**), TE (**orange, lower left**), and EE (**green, lower right**) power spectra in the full $30 \leq \ell \leq 7000$ range where we are able to extract a CMB-only signal from LAT data. The extracted CMB bandpowers are consistent with the fiducial power spectrum with minimal bias to the measurement. Although the mean bandpowers for the TT and EE autospectra appear to deviate from the fiducial spectrum at high ℓ ($\ell > 5000$), this is due to simulation-specific scatter, the logarithmic scale used in the plot, as well as degeneracies with the foregrounds; it does not affect the parameter measurement, as shown in Figure 3.9.

cross-spectrum error. Because the simulations do not yet include any instrumental systematics modeling, I opted to leave all instrumental nuisance parameters fixed to their fiducial inputs. This means we do not need any additional calibration parameters later in the CMB-only likelihood either. In the future, when we do need to model for these systematics, we will pick one polarized array (preferably the one that makes the most sensitive measurement of the CMB), and fix its polarization efficiency to its maximum likelihood estimation during the extraction step, alongside the overall gain calibration. These two additional parameters then need to be marginalized over in the CMB-only likelihood.

I show a comparison of the fiducial CMB power spectra and the recovered

CMB bandpowers in Figure 3.8. We see an excellent recovery of the fiducial power spectrum, with no visible biases in the extracted CMB bandpowers. The final recovered bandpowers fall within $O(1 - 10\mu K^2)$ of the theoretical power spectrum, with a high fidelity recovery of the CMB signal in the $1000 \leq \ell \leq 5000$ range.

We can now recover cosmological parameters from this data through a CMB-only analysis. I wrote a likelihood as described in subsection 3.6.1 for the extracted bandpowers from this simulation, and compared it to the fiducial parameters as well as the constraints from the multifrequency analysis with `MFLike`. The bandpower windows used in this likelihood were the same as those in `MFLike`. Only the CMB signal, $\mathcal{D}_\ell^{XY,th}(\theta)$ needs to be modeled, and no additional systematics need to be marginalized over for now.

I show the results of running a CMB-only likelihood in Figure 3.9. Overall, we see that the foreground marginalization procedure recovers a power spectrum which contains the same cosmological information as the multifrequency data, while needing over a factor 10 less samples to converge. The difference in constraints in the cosmological parameters compared to the fiducial parameters is purely down to sample variance, as is shown earlier in Figure 3.7. The foreground marginalization procedure imposes no bias in the cosmological constraints while offering a greater speed in convergence for Bayesian analysis.

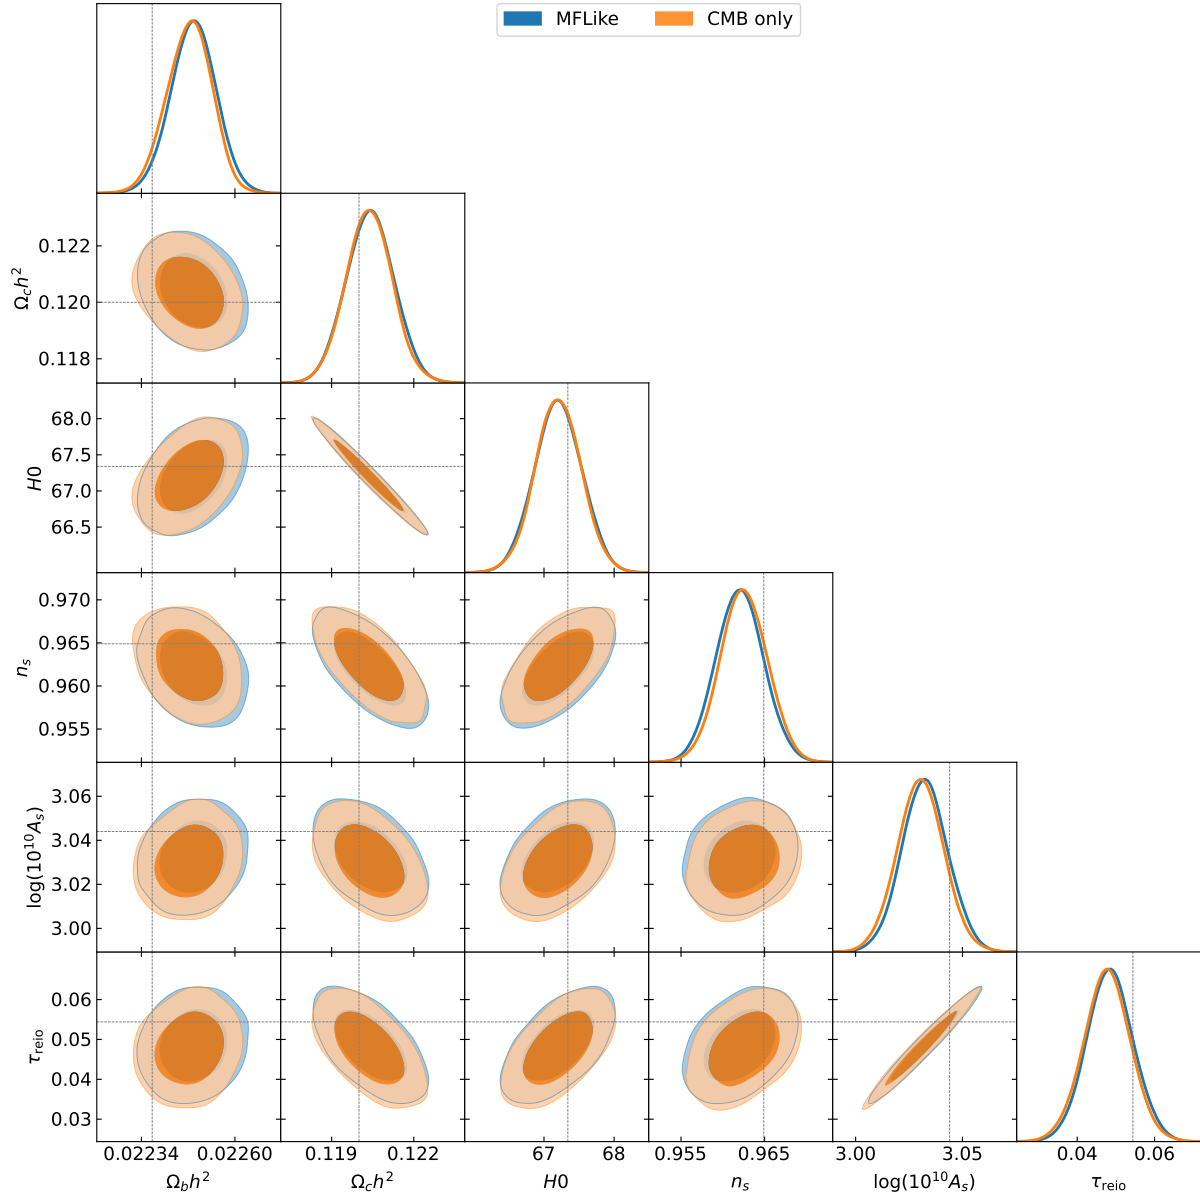



Figure 3.9. A comparison of the recovered cosmological parameters from a single SO LAT simulation, computed with both a full MFLike analysis (**blue**), and a foreground-marginalized CMB-only analysis (**orange**), with the fiducial cosmology shown in the black vertical lines. We see excellent agreement between the two results. Performing a foreground-marginalization shows no significant loss in constraining power, while the number of required samples to reach a fully converged MCMC analysis was reduced from ~ 370000 to ~ 26000 .

Emulators for High-Precision Cosmology



In this chapter, I will discuss modern developments in the deployment of *cosmological emulators* in Bayesian inference. With the development of new, small-scale observatories comes an ever-increasing pressure on cosmological theory codes to provide more precise computations of theory models. This is leading to an exponential increase in the need for computer power, as these theory models have non-linear complexity in ℓ scaling, accuracy increases, and beyond- Λ CDM physics modeling. Emulators provide a natural solution to this problem, and part of my work has been to create and characterize neural network emulators with the `python` package `CosmoPower`, as well as to create new software interfaces for `CosmoPower` to connect it with `cobaya` and other Bayesian inference frameworks for widespread adoption of emulator usage.

In this chapter, I will go through the need for and function of neural network emulators. The structure of this chapter is as follows:

- I will delve into the limitations of current, traditional, Einstein-Boltzmann codes in section 4.1, and argue the need for emulators;
- I will explain the concept of *emulators* in section 4.2, explaining what an emulator is, and the basic mathematics behind neural network (NN) emulators;

- In section 4.3, I will go through our work on the software `CosmoPower`, its implementation of NN emulators for cosmological inference, and the framework we have developed for the creation and characterization; I will also show our set of (beyond)- Λ CDM emulators that I have created for our work;
- In section 4.4, I will show our design concept for human-readable and machine-parsable prescription files, which are a simple prescription for using the pipeline described in section 4.3 without the need for any bespoke software development by the end-user;
- I will conclude with displaying the accuracy and speed of our new suite of emulators in section 4.5, showing that our emulators are accurate enough to recover cosmic variance-limited cosmological parameters without any biases.

The following papers were used for this chapter:

- **A complete framework for cosmological emulation and inference with CosmoPower**, Jense, Harrison, Calabrese, Spurio Mancini, Bolliet, Dunkley, and Hill (2024);

This paper, which I authored and is based on my work, forms the main body of this chapter. Most figures, notably Figure 4.2, 4.3, 4.5, 4.6, 4.7, 4.8, are my work and included directly from the paper. Sections 4.3 through 4.5 are adapted from this paper.

- **CosmoPower: emulating cosmological power spectra for accelerated Bayesian inference from next-generation surveys**, Spurio Mancini, Piras, Alsing, Joachimi, and Hobson (2022);

While I have not personally contributed to this paper, work in this chapter continues from this paper, and parts of section 4.2 are based on this paper. I explicitly refer to this paper when I use content from it.

- **High-accuracy emulators for observables in Λ CDM, N_{eff} , Σm_ν , and w cosmologies**, Bolliet, Spurio Mancini, Hill, Madhavacheril, Jense, Calabrese, and Dunkley (2023);

I have co-authored this paper and part of the methodology from this chapter is adapted from this work. I explicitly refer to this paper when I use content from it.

4.1 Limitations of Conventional Cosmological Codes

We have seen throughout the previous chapters how we compute models for cosmological probes and fit these to observational data to recover constraints on cosmological parameters. We recall that for the CMB, this is done by computing a set of transfer functions $\Delta_\ell^X(k)$ and using it to relate the primordial power spectrum $P(k)$ to the observed quantity C_ℓ^{XY} as

$$C_\ell^{XY} = 4\pi \int \Delta_\ell^X(k) \Delta_\ell^Y(k) P(k) d\ln k \quad (4.1)$$

However, this computed quantity does not take into account *lensing* effects, and we can only observe the lensed CMB with telescopes. Lensing effects cause a coupling between different ℓ modes of the CMB, and especially at high multipoles, these effects can become an important observational effect. Because of this, if we observe the CMB up to some ℓ , we actually need to compute our theoretical model to a *higher* ℓ_{\max} in order to accurately including the effects in the lensing tail of the CMB. At present this is done with conventional codes that model the CMB power spectrum, so-called Einstein-Boltzmann codes, which can accurately model the CMB power spectrum to high precision, provided one has the computation power.

These requirements however, yield a series of difficult computational limitations on our theoretical model for the CMB C_ℓ s:

- With more precise data, we need to increase our computational precision in order to accurately constrain cosmological parameters;
- With more small-scale data, we need to compute our observables to higher multipoles, which require more accurate transfer functions $\Delta_\ell^X(k)$ that reach regimes where our integration step size over k needs to be shorter;
- As we reach higher multipoles, the effects of CMB lensing become more important, which couples our already high ℓ modes to even higher ℓ multipoles, meaning we need to compute our theory curve to *even higher* multipoles than before.

4.2 Cosmological Emulators

A **Cosmological emulator** is any kind of computational system that recreates (“emulates”) the mapping between cosmological parameters θ and observable quantities C without relying on the full theoretical computation of classical Boltzmann solvers.

4.2.1 Neural Network design

Following the design set out in the original `CosmoPower` paper Spurio Mancini et al. (2022), we employ Neural Network (NN) emulators. A neural network is a set of densely connected neurons that relate the inputs to the outputs via a network of weighted connections. See Figure 4.1 for a schematic overview.




Figure 4.1. A schematic overview of a Neural Network. It consists of several *neurons* (circles) that are connected via *weights* (lines). This particular example network consists of four *layers*, with the first *input layer* having three neurons (representing here the cosmological input parameters $\Omega_b h^2$, $\Omega_c h^2$, and H_0), the two *hidden layers* having five neurons each, and the final *output layer* having 2507 neurons (it thus emulates a mapping $\mathbb{R}^3 \rightarrow \mathbb{R}^{2507}$, computing CMB multipoles $C_\ell(\Omega_b h^2, \Omega_c h^2, H_0)$ for $2 \leq \ell \leq 2508$). The layers are *densely connected*, meaning that between consecutive layers A and B, each neuron in A is connected to each neuron in B. The value of a specific node is determined by the activation function f of that node and the weighted inputs $W\vec{x}$ of all connected nodes in the previous layer.

From a mathematical point of view, we have a series of *input* values $\vec{x}^1 = (x_1^1, x_2^1, \dots, x_n^1)$, sets of *weights* W^1, \dots, W^{l-1} (which are matrices) and *biases* $\vec{\beta}^1, \dots, \vec{\beta}^{l-1}$ (which are also vectors), and some activation function $f(\vec{x})$. We define the relation

between layer i and layer $i + 1$ as:

$$\vec{x}^{i+1} = f(W^i \vec{x}^i + \vec{\beta}^i), \quad (4.2)$$

where ultimately the output we are interested in is the final value \vec{x}^l .

There is a lot of freedom in the choice of the number of layers (the value of l), the number of hidden nodes in each layer (the shapes of the matrices W and the vectors $\vec{\beta}$), and the type of activation function (the function f). The exact type of activation function is a free choice and is not necessarily easy to decide a priori. A popular choice is the so-called **Rectified Linear Unit (ReLU)** function

$$f(\vec{x}) = \max(0, \vec{a} \odot \vec{x}), \quad (4.3)$$

where \odot signifies element-wise multiplication, and the \max function is operating element-wise. This function has the advantage of being simple, having a very simple derivative.

In our case, we will end up using a more general activation function, that can interpolate between a linear activation and the ReLU function. We implement the *weighted swish function* as I later define in Equation 4.5.

Training

The question then becomes to optimize the network, such that the mapping $\vec{\theta} \rightarrow C_\ell^{XY}$ is accurate enough for our analysis. This is done via **machine learning**, where the network is initially set to a random state, and iteratively tweaked to reproduce pre-generated training data. The goal is to minimize the different between the predicted observables $\hat{C}_\ell^{XY}(\vec{\theta})$ and the true observable $C_\ell^{XY}(\vec{\theta})$. This difference is often expressed in terms of the **loss function** \mathcal{E} , which in our case we chose to be the mean square difference

$$\mathcal{E} \equiv \sum_\ell \left(\frac{\hat{C}_\ell - \mu_\ell}{\sigma_\ell} - \frac{C_\ell - \mu_\ell}{\sigma_\ell} \right)^2 = \sum_\ell \left(\frac{\hat{C}_\ell - C_\ell}{\sigma_\ell} \right)^2, \quad (4.4)$$

where each \hat{C}_ℓ and C_ℓ is normalized with respect to the mean μ_ℓ and standard deviation σ_ℓ of the training data. For other emulated quantities, for example the matter power spectrum $P(k)$ or any kind of derived parameter like σ_8 , we use the same loss function and normalize the outputs with respect to the mean and standard deviation of the training dataset.

We generate a large number of training samples, which consists of pairs of $\vec{\theta}$ and C_ℓ^{XY} that are sampled across our parameter space $\{\vec{\theta}\}$. To ensure an even distribution

of our parameter space, we employ Latin Hypercube Sampling (LHS) for an optimal distribution of input parameters. Each training pass consists of the following:

1. We group the training data into different *batches* for improved stability, setting aside a small ($\sim 10\%$) subset of our samples for validation testing;
2. We pass our batches through the network and evaluate the predicted $\hat{C}_\ell^{XY} = N(\vec{\theta})$ (where N is the function that is the entirety of our neural network);
3. We tweak our hyperparameters W^i , $\vec{\beta}^i$, and $\vec{\gamma}^i$ to decrease loss function \mathcal{E} , across the entire batch;
4. After passing through each of our training batches, we evaluate the network at our validation batch once and report on the *validation loss* \mathcal{E}_{val} ;
5. If the validation error stagnates and does not decrease for a number of training steps in a row, we stop this training pass and move on to the next.

Assuming our input training data set is large enough, the number of training steps is large enough, and our training goes well, we will slowly converge onto a set of hyperparameters that accurately reproduce our mapping $C_\ell^{XY}(\vec{\theta})$.

The crucial step in this is step 3, where we want to tweak our hyperparameters such that the loss function \mathcal{E} goes down with each training step. This is commonly done via some optimization algorithm, such as stochastic gradient descent. Since our emulator is purely described via differentiable function $f(\vec{x})$ and linear operations between the weights, biases, and input values, it is mathematically possible and computationally feasible to evaluate the gradients $d\mathcal{E}/dW^i$, $d\mathcal{E}/d\vec{\beta}^i$, and $d\mathcal{E}/d\vec{\gamma}^i$. These are then passed to the optimization algorithm, which then tweaks the hyperparameters in a manner that should improve the accuracy of the network.

The actual accuracy of a network is often not expressed in terms of the loss \mathcal{E} but rather the *validation loss* \mathcal{E}_{val} , which is the loss over a set of data points that are not used in the set passed to the optimizer. This separate validation dataset is used to avoid overfitting of the model, which is a real risk when the original dataset is used for validation.

4.2.2 Reusability

A large problem in the use of cosmological emulators is the upfront investment cost of training a neural network. A single cosmological analysis requires running an MCMC chain with $O(10^5)$ samples, which would require a similar number of

evaluations of an Einstein-Boltzmann code. Preparing an emulator for Stage-IV accuracy takes a similar number of training samples, which need to be pre-generated with a classical Einstein-Boltzmann code. Then, the neural network needs to be trained on these samples, which is an additional investment cost on creation *before* one gets to running the MCMC. However, once one has a sufficiently stable and accurate emulator, it no longer needs to be trained again and can be used for many analyses without having to run the underlying Einstein-Boltzmann code again. For this reason, reusability is key in emulators, as an emulator only needs to be used twice for it to be more efficient than an Einstein-Boltzmann code.

4.3 Implementation

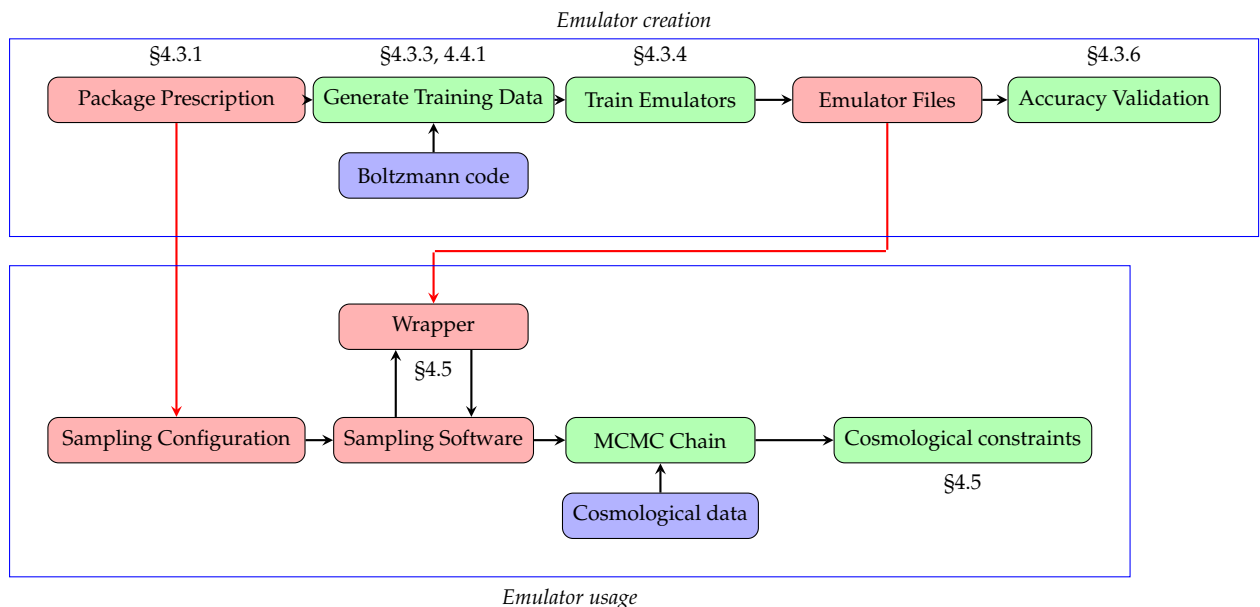


Figure 4.2. An overview of the workflow with `CosmoPower`: To create a new emulator (**top blue box**), we write a *packaging prescription*, use that to *generate training data*, and from that *train emulators* which outputs several *emulator files*, for which we can easily generate plots which *validate the accuracy* of the emulators. This packaging prescription and set of emulator files are then shared with the end-user (**red arrows**), who wants to use the emulators (**bottom blue box**): the prescription is put inside the *sampling configuration file*, which is given to our *software wrappers*, which provide the user with an *MCMC chain* that can be used to find *cosmological constraints*. The various labels refer to the sections where these individual steps are described in this chapter.

In this section I describe the details of our emulators: what is emulated and with which inputs, how an emulation is performed and how the emulators are validated. This serves both as a full description of the emulators that I released with my work,

Parameter	Range	Default Value
$\Omega_b h^2$	[0.015, 0.03]	–
$\Omega_c h^2$	[0.09, 0.15]	–
$\ln(10^{10} A_s)$	[2.5, 3.5]	–
n_s	[0.85, 1.05]	–
τ_{reio}	[0.02, 0.20]	–
H_0 [km/s/Mpc]	[40, 100]	–
z_{pk}	[0, 20]	–
$\log T_{\text{AGN}}$	[7.3, 8.3]	7.8
N_{eff}	[1.5, 5.5]	3.044
Σm_ν [eV]	[0, 0.5]	0.06
w_0	[−1.5, 0]	−1.0
w_a	[−2, 1.5]	0.0

Table 4.1. Table of parameter ranges over which we trained our emulators. Compare this with the textual specification in Figure 4.11. The **top** section of the table refers to the background cosmology parameters used in almost all emulators. The **middle** section of the table contains the redshift and baryonic feedback parameter $\log T_{\text{AGN}}$ used only in the $P(k, z)$ and $\sigma_8(z)$ emulators, with their default values from `camb` used in the CMB and background evolution emulators. The **bottom** section of the table shows the ranges of the single-/two-parameter extension model emulators, and their default values taken in the base Λ CDM case. Each emulator takes in the first six parameters, and one or two extension parameters, with the exception for $C_\ell^{\phi\phi}$, and background quantities, which do not rely on τ_{reio} .

and as guidelines on the creation of new emulators packaged and usable in the same way (e.g. for extended cosmological models). By *emulator* I mean a ‘black box’ code which is capable of ingesting a set of cosmological parameters $\vec{\theta}$ and outputting a set of predictions for the summary statistics of a set of observables $\{\vec{d}_1(\vec{\theta}), \vec{d}_2(\vec{\theta}), \dots, \vec{d}_N(\vec{\theta})\}$ which are indistinguishable (within a given tolerance) from the set which would have been produced by a code which explicitly implements numerical models of the physics relating the \vec{d} and $\vec{\theta}$. As the emulation works effectively as an interpolation of the quantities \vec{d} between known points, I rely on the fact that the \vec{d} vary smoothly with respect to the input parameters.

To meet the requirements for Stage-IV analyses, we use the `camb` accuracy settings suggested by McCarthy et al. (2022); Hill et al. (2022) as adequate for convergence of the likelihood value obtained from data with this level of precision, summarised in Figure 4.4.




Figure 4.3. An overview of the accuracy reached by a trained C_ℓ^{TT} emulator given the number of training spectra used to train the emulator, for an increasing number of input parameters. The solid coloured lines and point represent the 68% error of a C_ℓ^{TT} emulator trained with N_S samples (the dotted shaded lines and points show the similar behaviour observed at 99%), averaged over the entire ℓ -range. We show the full-size emulators generated with $N_S = 100000$ for Λ CDM and $N_S = 120000$ for extended models, as well as emulators with a smaller training set to show how accuracy scales with N_S and input parameters. We train emulators for Λ CDM (6 parameters, **blue**), $+\Sigma m_\nu$ (7 parameters, **orange**), and $+N_{\text{eff}} + \Sigma m_\nu$ (8 parameters, **green**), each on a random smaller subset of the full training dataset, scaling the training batch size proportional to the size of the subset. We show how the mean emulation error decreases as the number of training spectra increases, and increases as we increase the complexity of the parameter space. We note, however, that scaling of emulators accuracy with number of input parameters is non linear, the nature and impact on the emulated quantity of the specific parameter will matter for this behaviour.

Quantity	Range	Emulator
C_ℓ^{TT}	$2 \leq \ell \leq 10000$	NN of log-spectra
C_ℓ^{TE}	$2 \leq \ell \leq 10000$	NN+PCA of spectra
C_ℓ^{EE}	$2 \leq \ell \leq 10000$	NN of log-spectra
C_ℓ^{BB}	$2 \leq \ell \leq 10000$	NN of log-spectra
$C_\ell^{\phi\phi}$	$2 \leq \ell \leq 10000$	NN+PCA of log-spectra
$P_{\text{lin}}(k, z)$	$10^{-4} \leq k \leq 50$	NN of log-spectra
$P_{\text{NL}}(k, z)$	$10^{-4} \leq k \leq 50$	NN of log-spectra
$P_{\text{NL}}/P_{\text{lin}}(k, z)$	$10^{-4} \leq k \leq 50$	NN of spectra ratio
$H(z)$	$0 \leq z \leq 20$	NN of evolution
$\sigma_8(z)$	$0 \leq z \leq 20$	NN of evolution
$D_A(z)$	$0 \leq z \leq 20$	NN of evolution
<i>derived parameters</i>	–	NN of value of derived parameters

Table 4.2. Emulated quantities, ranges of scales covered and type of emulator employed for each of them.

```

1 lmax: 10000
2 kmax: 10.0
3 k_per_logint: 130
4 nonlinear: True
5 lens_potential_accuracy: 8
6 lens_margin: 2050
7 lAccuracyBoost: 1.2
8 min_l_logl_sampling: 6000
9 DoLateRadTruncation: False
10 recombination_model: CosmoRec

```

Figure 4.4. Accuracy settings for `camb`, based on the settings earlier suggested in McCarthy et al. (2022); Hill et al. (2022). This can be compared with Figure 4.11, which contains the parameters over which the training data is to be generated.

4.3.1 Emulated Quantities

In Table 4.2 we show the full list of quantities output by the Einstein-Boltzmann code (`camb` v.1.5.2) which we emulate. As output observables we generate the CMB temperature, polarisation and lensing potential angular power spectra; linear and non-linear matter power spectra (and their ratio); and a limited set of background expansion and derived perturbation quantities also output by the Einstein-Boltzmann code.

We compute the CMB angular power spectra in the multipole range $2 \leq \ell \leq 10000$ in each of TT, TE, EE, and BB combinations for different cosmological models. In the basic configurations, we use as inputs for our emulators the six cosmological parameters of the standard Λ CDM model: $\Omega_b h^2$, $\Omega_c h^2$, H_0 , n_s , $\log(10^{10} A_s)$, and τ_{reio} . We add additional model parameters to these for separate emulators for Λ CDM extension models as explained below. When not explicitly varied, neutrinos are described by fixing $N_{\text{eff}} = 3.044$, corresponding to the contribution from the three Standard Model neutrino species, with one of them carrying a total 0.06 eV mass.

We also emulate the CMB lensing potential $\phi\phi$ power spectrum in the same multipole range. For this we use the same parameter inputs except the optical depth to reionization, given that the lensing potential power spectrum does not depend on it.

In addition to emulators for these CMB observables, we also generated emulators of several non-CMB observables. Firstly, we created an emulator for the matter power spectrum $P(k, z)$. Similarly to how the CMB power spectrum describes the distribution of CMB photons on the sky, the matter power spectrum describes the spatial clustering of matter at wavenumbers k and redshift z . For the matter power spectrum, we compute the linear matter power spectrum $P_{\text{lin}}(k)$ for five input parameters: $\Omega_b h^2$, $\Omega_c h^2$, $\ln(10^{10} A_s)$, n_s , and H_0 , plus again the extra parameters for the extension models. For all matter power spectra we also treat the redshift z as an input parameter, resulting in an emulator function which acts as $P_{\text{lin}}(k, \vec{\theta})$, where $\vec{\theta}$ includes the redshift. For the non-linear matter power spectrum, we emulate both the $P_{\text{NL}}(k, z)$ spectrum itself and the non-linear boost $P_{\text{NL}}/P_{\text{lin}}(k, z) - 1$. For the emulators included in our work, we emulate the 2020 version of HMCcode described in Mead et al. (2021). We sample the wavenumber k at 500 points in the range $10^{-4} \leq k \leq 50 \text{ Mpc}^{-1}$ with logarithmic spacing. Note that we compute $P(k)$ up to $k = 100 \text{ Mpc}^{-1}$ for improved accuracy.

We additionally emulated several background evolution quantities. For these, we use redshift in the range $0 \leq z \leq 20$, sampled at 5000 equally-spaced points, as the modes along which we evaluate the redshift-evolution of the Hubble parameter $H(z)$, the angular diameter distance $D_A(z)$, and the clustering parameter $\sigma_8(z)$ for the

five input parameters $\Omega_b h^2$, $\Omega_c h^2$, $\ln(10^{10} A_s)$, n_s , and H_0 , plus the extension model parameters where relevant. Adding these background quantities to our emulator packages allows for additional cosmological constraints from e.g. BAO measurements. Additional background quantities, such as $f\sigma_8(z) \equiv -(1+z)d\sigma_8/dz$, can also be easily computed from these quantities with minimal overhead or loss of accuracy.

Finally, we compute ten derived parameters, namely:

1. The angular acoustic scale θ_* at the surface of last scattering;
2. The matter clustering parameter σ_8 ;
3. The primordial helium fraction Y_{He} ;
4. The redshift z_{reio} of reionization, defined as the midpoint of reionization described by a simple hyperbolic tangent;
5. The optical depth $\tau_{r,\text{end}}$ at the end of recombination;
6. The redshift z_* at the surface of last scattering;
7. The sound horizon scale r_* at the surface of last scattering;
8. The redshift z_d at the baryon drag epoch;
9. The sound horizon scale r_d at the baryon drag epoch;
10. The effective number of relativistic species N_{eff} .

It is common to use θ_* , the angular scale when optical depth is unity, or the approximate parameter θ_{MC} , as a sampled parameter in MCMC analyses of CMB data due to its lower level of covariance with other parameters than H_0 . We noted that `camb` and `CLASS` use different points at which to evaluate the angular scale (with `CLASS` defining θ_s as the angular scale at maximum visibility, which is close to but not the same as θ_* , which is used in `camb`). We have commented on this in prior work on Stage-IV emulator creation with `CLASS`, described in Bolliet et al. (2023). To maintain cross-compatibility between our emulators, and to remain consistent with our earlier work, we therefore use H_0 as an input, and not θ_* . Including these derived parameters as emulators allow us to recover the posterior distributions on these quantities, either directly storing their computed values while sampling the chain, or afterwards by post-processing a converged MCMC chain.

Quantity	Parameters
$C_\ell^{TT}, C_\ell^{TE}, C_\ell^{EE}, C_\ell^{BB}$	$(\Omega_b h^2, \Omega_c h^2, \ln(10^{10} A_s), n_s, H_0, \tau_{\text{reio}})$
$C_\ell^{\phi\phi}$	$(\Omega_b h^2, \Omega_c h^2, \ln(10^{10} A_s), n_s, H_0)$
$H(z), D_A(z)$	$(\Omega_b h^2, \Omega_c h^2, \ln(10^{10} A_s), n_s, H_0)$
<i>derived parameters</i>	$(\Omega_b h^2, \Omega_c h^2, \ln(10^{10} A_s), n_s, H_0, \tau_{\text{reio}})$
$P_L(k, z_{\text{pk}})$	$(\Omega_b h^2, \Omega_c h^2, \ln(10^{10} A_s), n_s, H_0, z_{\text{pk}})$
$P_{NL}(k, z_{\text{pk}}), P_{NL}/P_L(k, z_{\text{pk}}) - 1$	$(\Omega_b h^2, \Omega_c h^2, \ln(10^{10} A_s), n_s, H_0, z_{\text{pk}}, \log T_{\text{AGN}})$
$\sigma_8(z)$	$(\Omega_b h^2, \Omega_c h^2, \ln(10^{10} A_s), n_s, H_0, \log T_{\text{AGN}})$

Table 4.3. The input parameters used for each emulator for the Λ CDM case. For the extension models, all extension parameters are used across all emulators.

4.3.2 Cosmological Models

We provide emulators for CMB power spectra from the Λ CDM model with parameters $(\Omega_b h^2, \Omega_c h^2, \ln(10^{10} A_s), n_s, H_0, \tau_{\text{reio}})$ defined above as well as the following four extended models:

1. Λ CDM + N_{eff} : varying the effective number of relativistic species N_{eff} ;
2. Λ CDM + Σm_ν : varying the sum of neutrino masses Σm_ν ;
3. Λ CDM + $N_{\text{eff}} \Sigma m_\nu$: varying both the number and mass sum of neutrinos;
4. Λ CDM + $w_0 w_a$: varying the dark energy equation of state described with two parameters w_0 and w_a .

Each of these four extension models is emulated separately, with the extension parameters used as additional inputs. These parameters, except τ_{reio} , are also used as inputs for the redshift evolutions $H(z)$ and $D_A(z)$. We also do not use τ_{reio} for the lensing potential $C_\ell^{\phi\phi}$.

The matter power spectrum is often computed as the *linear* and *non-linear* power spectra, the latter of which includes terms for baryonic interactions, which tend to be model-dependent. For the linear $P_L(k, z_{\text{pk}})$, non-linear $P_{NL}(k, z_{\text{pk}})$ and non-linear boost $P_{NL}/P_L(k, z_{\text{pk}}) - 1$ emulators, we use all the same input parameters except τ_{reio} , and include an extra parameter z_{pk} for the redshift of the power spectrum. For the non-linear models and $\sigma_8(z)$, we need to include the parameter of the baryonic feedback model, which is the AGN temperature $\log T_{\text{AGN}}$ that appears in HMCode (Mead et al., 2021).

All input parameter choices are summarized in Table 4.3.

For the remaining model choices, which set a primordial helium fraction set from BBN consistency using PRIMAT (Pitrou et al., 2018), recombination from the CosmoRec

code (Chluba and Thomas, 2010; Chluba et al., 2010), and reionization modeled with a simple hyperbolic tangent with a redshift width $\Delta z = 0.5$. Most of these options are the default settings in `camb`. We only changed the recombination code to `CosmoRec`, whereas the `camb` default is to use the older `RECFAST` code.

4.3.3 Training Data

Training of emulators involves creating a set of output data \vec{d} at a finite sample of known parameter values $\vec{\theta}$ using the code to be emulated (i.e. the Einstein-Boltzmann code here). These data will then subsequently used in subsection 4.3.4 for the neural network to learn an approximate (but high accuracy) mapping between input and output. Training data must be generated at a high enough resolution in the input parameters that we can smoothly interpolate between outputs. The training data only need to be generated once, to train the emulator, and do not need to be generated using the computationally intensive numerical code again in any subsequent inference.

Following Spurio Mancini et al. (2022) and Bolliet et al. (2023), we generate $N_S = 10^5$ sets of output spectra as training data, of which 20% will be used for validating the network accuracy, and the rest for training. Our parameter space is shown in Table 4.1. We employ Latin Hypercube (LHC) sampling for ensuring our parameter space is evenly sampled. For extended models, we choose to generate slightly more spectra at $N_S = 1.2 \times 10^5$, to compensate for the expanded parameter space. To demonstrate the need for this and to provide some guidance on how to select N_S , we show a comparison of the mean prediction error versus the size of the training dataset in Figure 4.3, for a varying number of input parameters. The figure shows that there is not a simple linear scaling with the number of parameters. Although increasing the number of parameters always requires a larger training set to reach the desired target accuracy, the physical nature and range of variation of the specific additional parameter will impact the results. For example, if we extend Λ CDM varying N_{eff} or Σm_ν , we observe different behaviours, even if in both cases it is only one additional input parameter (7 input parameters compared to 6 for Λ CDM). We explain this by noting that cosmological observables have different responses to different parameters, according to the physics signature they are tracking. For example, the CMB TT spectrum will exhibit a strong dependence on N_{eff} – changing both the peak position and amplitude at all scales, but less so on Σm_ν which will primarily appear at scales dominated by lensing. Hence in Figure 4.3 the Λ CDM + N_{eff} case requires more training than Λ CDM + Σm_ν . When we expand further the model to Λ CDM + $N_{\text{eff}}\Sigma m_\nu$ (8 input parameters compared to 6 for Λ CDM), we observe a very similar behaviour to the 7-parameter case Λ CDM + N_{eff} , because we

have already covered most of the strongly-varying training region. We conclude that to achieve the desired convergence of the emulators, the user will need to monitor the behaviour of their specific model and perform some exploratory studies of how the emulators depend on the model parameters.

4.3.4 Network Design and Training

Following Spurio Mancini et al. (2022), we implement the emulators as dense neural networks, with four hidden layers of 512 neurons each. Each emulator takes the normalised parameters as input, and maps it to normalised spectra. We use the activation function from Spurio Mancini et al. (2022):

$$f(\vec{x}) = \left[\vec{\gamma} + \left(1 + e^{-\vec{\beta} \odot \vec{x}}\right)^{-1} \odot (1 - \vec{\gamma}) \right] \odot \vec{x}, \quad (4.5)$$

which is sometimes referred to as the *weighted swish function*, with $\vec{\beta}$ and $\vec{\gamma}$ hyperparameters we can optimize. For the optimizer, we re-use the Adam optimizer. The input and output quantities are normalised with respect to mean and standard deviations of the respective ranges. For most quantities, as detailed in Table 4.2 we emulate the logarithm of the spectrum, as the high dynamic range of these values makes it easier for the emulator to reconstruct the log-values. We employ the same method for the background quantities $H(z)$, $\sigma_8(z)$, and $D_A(z)$, where we reconstruct the logarithm of the redshift evolution.

For the C_ℓ^{TE} emulator, the resulting raw spectra include zero-crossings which make emulating the log-spectra impossible. Because the unscaled spectra still contain a high dynamic range in values, we followed Spurio Mancini et al. (2022) in first decomposing the spectra with a Principal Component Analysis (PCA) and then subsequently emulating the sets of PCs. Similar to before, we decompose the C_ℓ^{TE} spectra into 512 PCs. Even though they remain completely positive, we also decompose the $C_\ell^{\phi\phi}$ spectra into 64 PCs. We find that this is more effective at emulating the $\phi\phi$ spectra, which we explain with the reduced dimensionality of the information contained in the $\phi\phi$ spectra. We construct scree plots, showing the eigenvalues associated with each PC in the decomposition, and identify the “elbow” at which higher PC numbers no longer carry significant weight and can be discarded. We did some investigating into some robust guidelines for this decomposition, which is summarised in subsection 4.3.5. With this setup, our emulator design for the CMB spectra remains fully consistent with the original emulators from Spurio Mancini et al. (2022).

For the matter power spectra $P_{\text{lin}}(k, z_{\text{pk}})$ and $P_{\text{NL}}(k, z_{\text{pk}})$, we choose to emulate $\log P_{\text{lin}}(k, z_{\text{pk}})$ and the non-linear boost $P_{\text{NL}}/P_{\text{lin}}(k, z_{\text{pk}}) - 1$ for best performance. These

quantities are functions of two parameters, the wavenumber k and redshift z_{pk} . Similar to previous emulators we have developed, we use k as the one-dimensional grid along which we sample our spectra, and use z_{pk} as an additional input for our $P(k)$ emulators.

The time it takes to train an emulator depends on many factors, including the size of the dataset, the number of inputs and outputs of the network, the hardware performance, as well as some slightly random factors due to the inherently stochastic nature of training. At 10^5 training samples for a network, we find it takes $O(1h)$ to train a C_ℓ network on a GPU. If no GPU hardware or the required software is available, then the emulators can alternatively be trained on a CPU, which for the same case still only takes $O(10h)$ to perform. For the generation of training data, which can take $O(10s)$ for a relatively simple model like Λ CDM, that means the only upfront cost is $O(10days)$ for generating training data – or less, since this program is highly parallelizable.

Of interest to a user of a pre-trained emulator, is the time spent on generating the initial training sample compared to evaluating a pre-trained emulator. For the accuracy settings and models presented here, it takes `camb` about 12 seconds to compute either the CMB spectra on a cluster. This is compared to evaluating the emulator, which took about 11 *milliseconds* on an average end-user laptop without GPU-acceleration. It should also be noted that for more complicated cosmological models, the computation performed by `camb` can be slower, while the emulator evaluation speed does not increase unless the architecture needs to be significantly changed.

4.3.5 Principal Component Analysis

The use of Principal Component Analysis (PCA) can be worthwhile in improving the accuracy of the emulator by compressing the full data into a smaller number of free components. While the reduction in freedom in the output is reduced and has therefore less capacity to accurately recover the original spectra, the reduced dimensionality of the output vector means that the emulator can more efficiently train on this reproduction.

The choice of whether or not to use PCA is not trivial, and there is no simple test that can conclusively show that the use of PCA compression is guaranteed to be beneficial before training an emulator. While for some cases, like C_ℓ^{TE} , the use of a PCA is needed due to the zero-crossing of the observed quantity, it may not be obvious *a priori* that the use of a PCA can improve it for other quantities as well.

It was observed in Spurio Mancini et al. (2022) that the $C_\ell^{\phi\phi}$ emulator improved in accuracy when employing PCA compression. We observed that this can be explained by making a *scree* plot, which is a line plot of the eigenvalues of all retained PCA components. We show a scree plot of the training data for our Λ CDM emulators




Figure 4.5. A scree plot, showing the *unexplained variance* of a PCA compression for the various CMB quantities, as a function of the number of retained principal components. The “scree” of each line is the flat plateau of each line. We observed that for $C_\ell^{\phi\phi}$, this scree lies around 64 principal components (vertical red line), and hence a PCA compression of 64 retained components is effective for a $C_\ell^{\phi\phi}$ emulator. Conversely, however, observing the scree for C_ℓ^{BB} at around 100 principal components, we expected the same to see for this quantity. We attribute the lack of an improvement in emulation for this quantity to the presence of important features which shift in ℓ -space for that quantity, which would not be retained by our implementation of PCA.

in Figure 4.5. By observing where this line flattens out (the “scree” of the line), one can estimate the amount of components that need to be retained in the PCA. For the $C_\ell^{\phi\phi}$ spectra, we found that this scree appears around 60 components, which means around 64 components should be sufficient to accurately decompose the 10000 ℓ modes of the spectra without loss of information. Similarly, a scree plot showed that a few hundred components should be sufficient for C_ℓ^{TE} .

However, a scree plot is not necessarily conclusive. We observed that the C_ℓ^{BB} are also dense enough that about 200 PCA components should be capable of accurately recovering them. Upon training such an emulator however, we found that direct emulation of C_ℓ^{BB} was more accurate than one that employed PCA compression. We think this is due to the fact that the BB spectra contain features which vary in ℓ under certain parameter variations, and hence cannot be properly accounted for in PCA compression. Since our regular emulators were shown to be more than accurate for physical analysis, we did not do an in-depth analysis of this discrepancy. Further investigation, or a different type of information compaction that does allow for horizontal shifts in ℓ -space, can perhaps allow for more accurate emulators in the future.

4.3.6 Accuracy of Emulated Observables

To assess the accuracy of our emulators we perform a number of comparisons between the observables emulated and those calculated directly with `camb`. This allows us to understand if we have reached the theoretical calculation accuracy required for Stage-IV analyses.

In Figure 4.6 we report the difference between direct `camb` outputs and emulated observables for Λ CDM, showing contours corresponding to the fraction of our training spectra (across the full parameter space) which lie within a given level of agreement with the emulated values. All the CMB spectra reach sub-percent accuracy (note that the TE higher values are numerical artefacts due to diving for a signal crossing 0); the matter power spectrum is accurate at the few percent level for very large range of wave numbers.

For the CMB observables, as done in previous works we can also compute the difference relative to (or ‘in units of’) a specific experiment’s sensitivity which tracks the noise for each observable N_ℓ^{XY} with

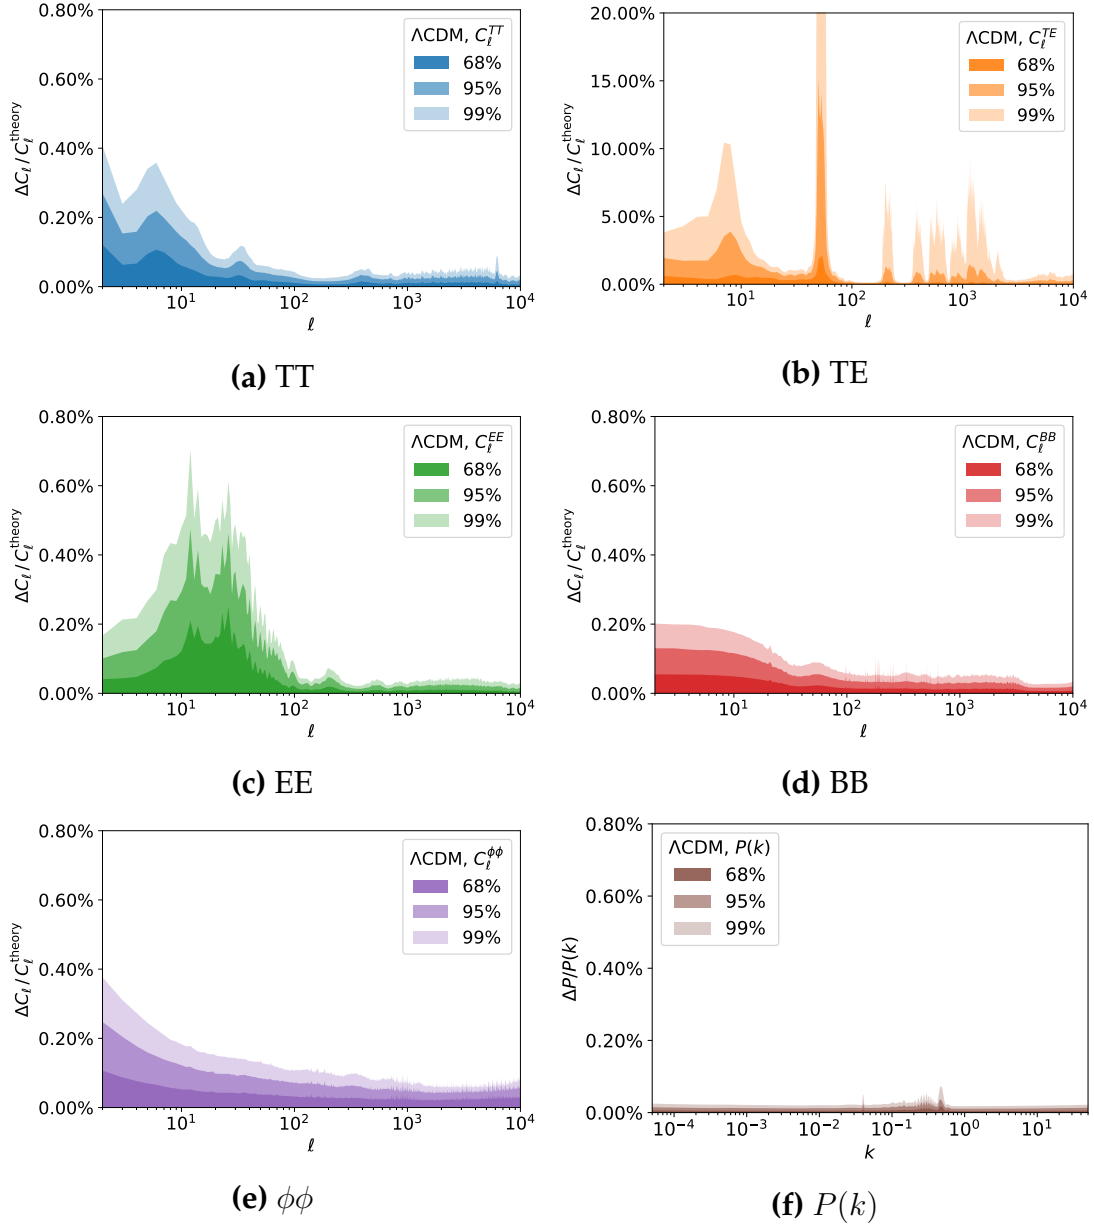


Figure 4.6. A validation graph generated from our trained networks for Λ CDM. We show the error in the reconstructed CMB power spectrum in TT (blue, top-left), TE (orange, top-right), EE (green, centre-left), BB (red, centre-left), $\phi\phi$ (purple, bottom-left), and $P(k)$ (brown, bottom-right) relative to the `camb` theory curve. The bands show the 68/95/99% contours (from darkest to lightest shades). Note the different scale for TE, for which errors get blown up due to the zero-crossings of the input power spectrum.

$$\left(\sigma_\ell^{XY}\right)^2 = \frac{1}{f_{\text{sky}}(2\ell + 1)} \left[\left(C_\ell^{XX} C_\ell^{YY}\right) + \left(C_\ell^{XY} C_\ell^{YX}\right) \right]. \quad (4.6)$$

where, for the cosmic variance limit, $f_{\text{sky}} = 1$ and $N_\ell^{XY} = 0$ for all XY . For SO, we use $f_{\text{sky}} = 0.4$ and we use the publicly available N_ℓ^{TT} , N_ℓ^{EE} , and N_ℓ^{BB} noise curves, presented in The Simons Observatory collaboration (2019).

We show this accuracy of our emulators relative to the cosmic variance-limited experimental noise in figures 4.7, 4.8 for Λ CDM, and Λ CDM + Σm_ν respectively. Similar figures for the other extension models are presented in Jense et al. (2024).

Since we create several emulators that emulate the same quantity in different ways, we are also able to compare their emulations to learn more about the effect that the setup has on the emulation quality. Of note we made a test of the recovery of:

- The non-linear matter power spectrum, which we emulate both directly as $P_{\text{NL}}(k)$ and through the linear and boost emulators $P_L(k)(1 + P_{\text{boost}}(k))$;
- The clustering parameter σ_8 at redshift zero, which is emulated by the derived parameters emulator and by the redshift evolution $\sigma_8(z)$ emulator.

Since the matter power spectrum emulators are all trained across similar training data, testing this comparison was relatively easy. We simply passed the validation input parameters for the non-linear emulator through the linear and boost emulators, and then computed the non-linear power spectrum from these two emulated quantities, and compared it with the same recovery from the non-linear emulator. We show a comparison of these two accuracy quantities in Figure 4.9. We note that although Spurio Mancini et al. (2022) suggested emulating the non-linear boost, we find that direct emulation of the nonlinear power spectrum is slightly better at emulation, particularly in the $10^{-2} < k < 10^1$ range.

We compare the recovery of the present-day matter clustering $\sigma_8(z = 0)$ in Figure 4.10, where we evaluate the derived parameter emulator and the zero redshift bin of the $\sigma_8(z)$ emulator. Because the latter emulator uses baryonic feedback parameters as an additional input, but the former does not, we generated training samples across the fixed value of $\log T_{\text{AGN}} = 7.8$ (see Table 4.1). We found that for this data recovery, the $\sigma_8(z)$ emulator was better at computing this quantity than the derived parameter emulator. This result may seem surprising at first, but can be explained relatively straightforward with the fact that the full set of training data for the $\sigma_8(z)$ emulator contains more data, as even data at $z > 0$ can help the emulator pick up information about the $z = 0$ bin.

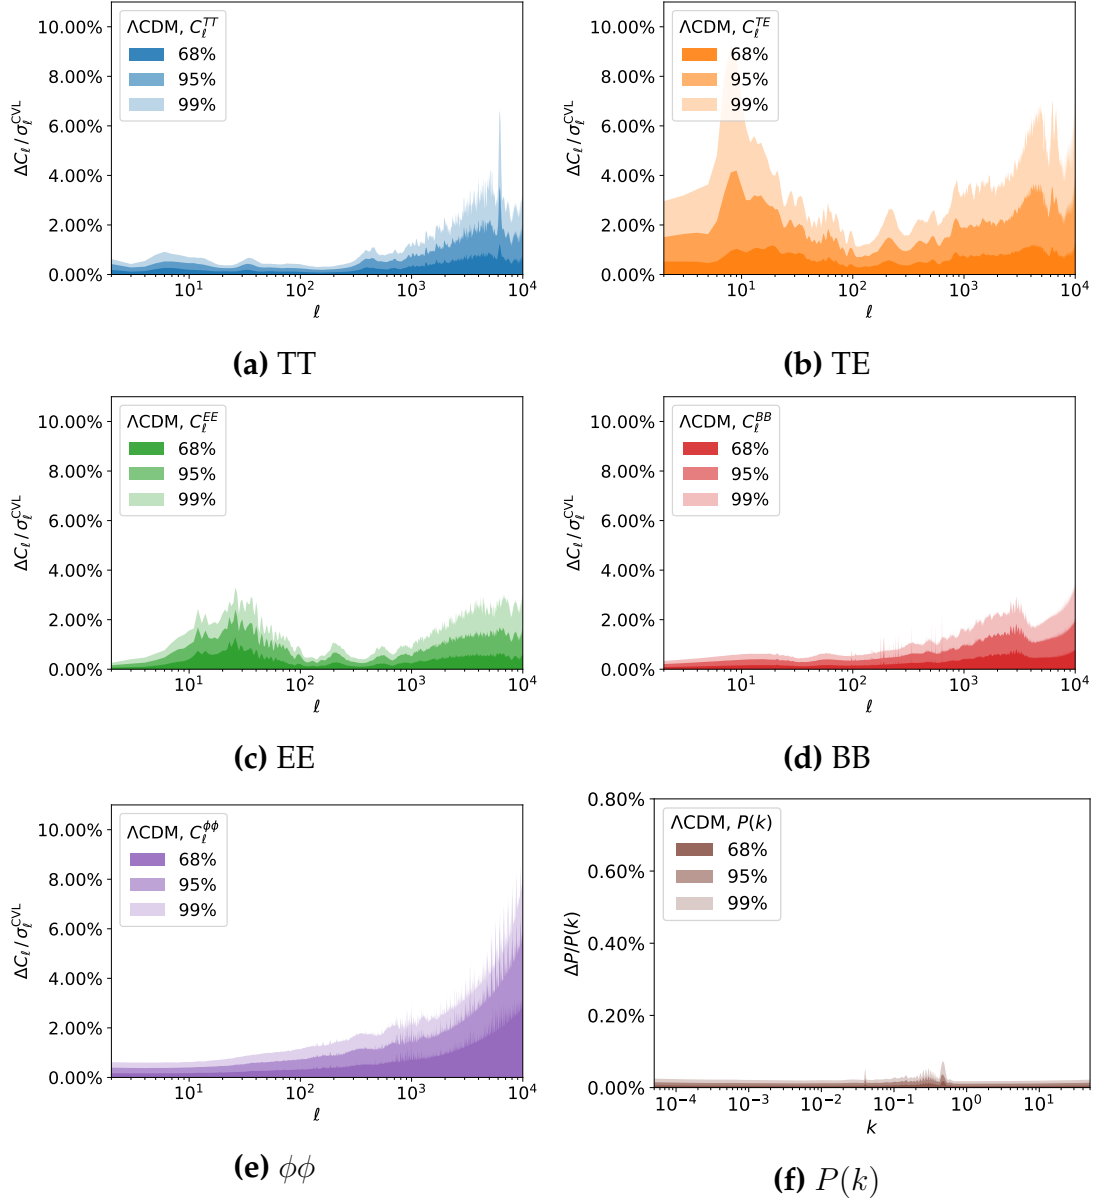


Figure 4.7. A validation graph generated from our trained networks for Λ CDM. We show the error in the reconstructed CMB power spectrum in TT (blue, top-left), TE (orange, top-right), EE (green, centre-left), BB (red, centre-left), $\phi\phi$ (purple, bottom-left), and $P(k)$ (brown, bottom-right) relative to the experimental noise of a cosmic variance-limited experiment. The bands show the 68/95/99% contours (from darkest to lightest shades).

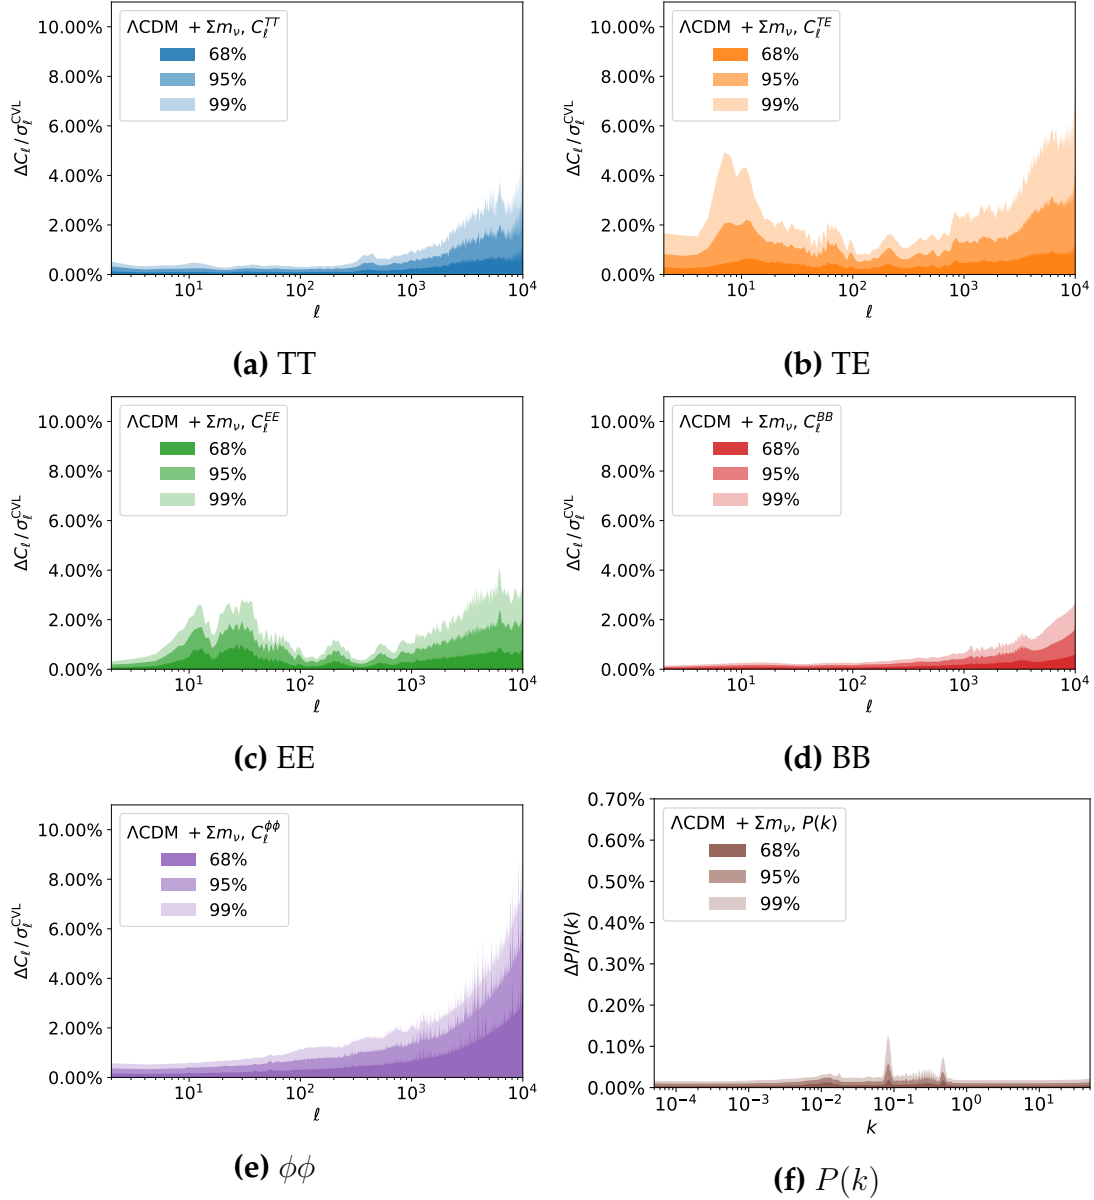


Figure 4.8. A validation graph generated from our trained networks for $\Lambda\text{CDM} + \Sigma m_\nu$. We show the error in the reconstructed CMB power spectrum in TT (blue, top-left), TE (orange, top-right), EE (green, centre-left), BB (red, centre-left), $\phi\phi$ (purple, bottom-left), and $P(k)$ (brown, bottom-right) relative to the experimental noise of a cosmic variance-limited experiment. The bands show the 68/95/99% contours (from darkest to lightest shades).

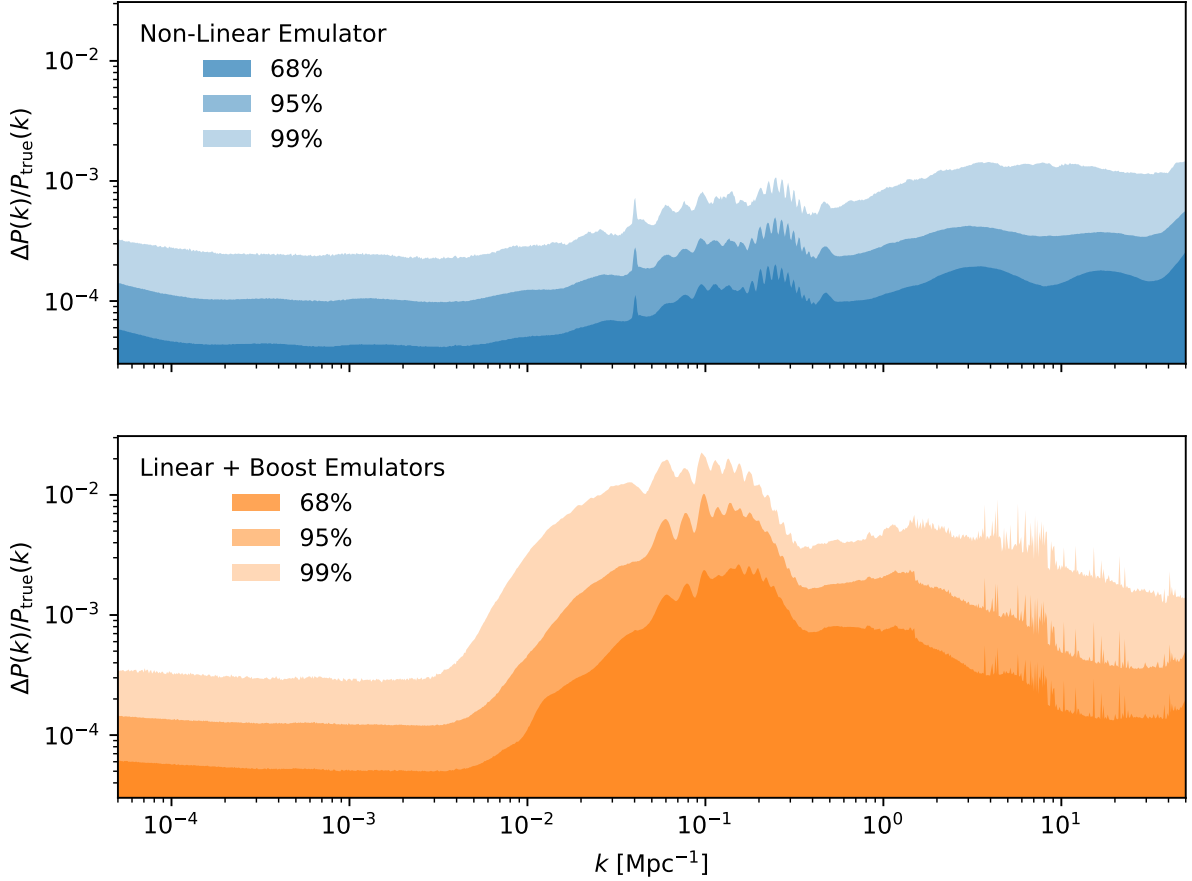


Figure 4.9. A comparison of the error in the emulation of the non-linear matter power spectrum by direct emulation (**top panel, blue**), and through the combination of linear and non-linear boost emulators (**bottom panel, orange**). We note that at low k , the emulators are equally effective at emulation, but at high k , the direct emulator is better at emulating the non-linear matter power spectrum than the combination of the linear and boost emulators, although the difference decreases again in the non-linear tail of the power spectrum ($k > 10^0$). This suggests that with our setup, the boost emulator is not necessarily better at emulating the non-linear power spectrum, and direct emulation is sufficient for sub-percent ($\ll 10^{-2}$) inaccuracy.




Figure 4.10. A comparison of the recovery of the derived parameter σ_8 (**blue**) and the evaluation of the full $\sigma_8(z)$ emulator at $z = 0$ (**red**) from validation samples. Since the former emulator does not take the baryonic feedback model into account, we fixed the relevant parameter $\log T_{\text{AGN}} = 7.8$ across this testing dataset. Although the validation data used is more similar to the training data for the derived parameter than the redshift evolution emulator, the latter is better at emulating the quantity. We attribute this to the fact that the space of training spectra for the $\sigma_8(z)$ may be better suited for recovering this quantity, although we note that we were not able to test if this recovery holds up as well when varying the baryonic feedback model.

4.4 Packaging Prescription

As part of this work, alongside new emulators we designed a packaging prescription for `CosmoPower` emulators. This prescription is both human- and machine-readable and serves as a description of what the emulator is capable of and its full design specifications. We updated the `CosmoPower` software package to include a full parser for the packaging prescription.

To create and train a new emulator, the packaging prescription is designed to guide both the author and a later user through the process of considering what quantities are emulated, how, and to what accuracy.

In this section, we describe the main steps of creating an emulator, namely: (1) describing the input parameters and output data, and generating the training spectra with the Einstein-Boltzmann code, (2) detailing the specifications of the emulator and the training parameters, and performing the training process, and (3) testing the validation of emulators. We follow the creation of the emulators we specified in section 4.3, and describe how the packaging prescription of these emulators is setup, as well as alternative options and choices available for the user.

4.4.1 Generating Training Data

As mentioned before, `CosmoPower` uses LHC sampling, which allows for an evenly spaced grid of sampling points that are sufficiently distributed that the entire parameter space is covered with minimal variation in sampling density. In Figure 4.11 we show how to prescribe the LHC grid in the prescription file.

The `emulated_code` block of the packaging contains information about the Einstein-Boltzmann code being emulated, in particular the name and version number. If a customized version of a code is used, it is possible to manually specify the import path with the `boltzmann_path` keyword. The `inputs` keyword is the list of named parameters which will be varied as inputs to the Einstein-Boltzmann code. `extra_args` contains code parameters which embody any model choices or approximation and accuracy settings.

The `samples` block specifies the `Ntraining` training spectra to be generated. The packaging prescription recognises four different types of parameters in the `parameters` block:

1. Sampled parameters, these are the parameters that the LHC is created over, and are defined with a minimum-maximum pair for the range over which the LHC is sampled, e.g. `ombh2: [0.015, 0.03];`

```

1 emulated_code:
2   name: camb
3   version: "1.5.2"
4   inputs: [ombh2, omch2, As, ns, H0, tau]
5   extra_args:
6     <...>
7
8 samples:
9   Ntraining: 100000
10  parameters:
11    ombh2: [0.015, 0.03]
12    omch2: [0.09, 0.15]
13    # We want to sample on log(10^10 As), but our
14    # Boltzmann code takes As as an input.
15    logA: [2.5, 3.5]
16    As: "lambda logA: 1.e-10 * np.exp(logA) "
17    tau: [0.02, 0.20]
18    ns: [0.85, 1.05]
19    H0: [40.0, 100.0]
20    # Parameters computed by the Boltzmann code
21    thetastar:
22    sigma8:
23    YHe:
24    zrei:
25    taurend:
26    zstar:
27    rstar:
28    zdrag:
29    rdrag:
30    Neff:

```

Figure 4.11. Code snippet for sampling and parameters block, compare this with Table 4.1. In the example here, we setup the aforementioned six parameters to sample over, add an intermediate parameter A_s , and add the nine parameters which are derived directly from the Boltzmann code, in this case `camb`. Note that `camb` expects the primordial amplitude A_s to be provided, but it is far more common to sample over $\ln(10^{10} A_s)$ instead. By defining the `logA` parameter and marking the `As` parameter as a derived parameter from that, we can perfectly accomplish this. At the bottom we show the nine parameters we derive from the Boltzmann code - in this case, they are computed by `camb`. It is possible to use any of the parameters defined in this block as an input to the networks, including the parameters derived from the Boltzmann code. The `extra_args` block would include any accuracy settings, as seen in Figure 4.4.

2. Derived parameters, these are parameters that are trivially derived from other sampled parameters, and are defined with a text string prescribing a python lambda function equation to derive them directly, e.g. As : `"lambda logA: 1.e-10 * np.exp(logA)";`
3. Fixed parameters, these are simply defined by writing a single numerical value that the parameter is set to, e.g. `mnu: 0.06;`
4. Computed parameters, these are parameters that we cannot easily compute ourselves, but the Boltzmann code can, and these are defined by simply leaving an empty tag in the parameter list. These parameters are specified by variable names available to `CosmoPower` at the spectra generation stage via the python interfaces of the Einstein-Boltzmann codes being emulated, e.g. `YHe: .`

Any of these types of parameters can be used as an input to a network, and any of the first three types can be used as an input for the Einstein-Boltzmann code. It is for example possible to create an LHC over a range of Hubble parameter H_0 , while using the angular scale θ_* , as computed by the Einstein-Boltzmann code, as an input for the emulators.

The `networks` block specifies the neural networks to be created using the training data. It is possible to specify multiple networks, each under a quantity heading, which each have their own set of network properties specified as further blocks and keywords. When creating `CosmoPower` networks, the current list of quantities to which can be emulated is defined and described as follows:

- `Cl/xy`: referring to (lensed) CMB angular power spectra C_ℓ^{XY} with X, Y any combination of T/E/B ($C_\ell^{TT}, C_\ell^{TE}, C_\ell^{EE}, C_\ell^{TB}, C_\ell^{EB},$ and C_ℓ^{BB});
- `Cl/pp`: CMB lensing potential spectrum for $\phi\phi$, there are also options available for cross-spectra with primary CMB via `Cl/pt`, `Cl/pe`, and `Cl/pb`;
- `Pk/lin` and `Pk/nonlin`: Matter power spectrum for linear $P_{\text{lin}}(k, z)$ and non-linear $P_{\text{nl}}(k, z)$;
- `Pk/nlboost`: The non-linear boost $(P_{\text{NL}}/P_{\text{lin}} - 1)(k, z)$ defined as the non-linear boost to the linear matter power spectrum;
- `Hubble`, `Omegab`, `Omegac`, `Omegam`, `sigma8` and `DA`: The redshift-evolving quantities $H(z), \Omega_b(z), \Omega_c(z), \Omega_m(z), \sigma_8(z),$ and $D_A(z)$.

```

1 networks:
2   - quantity: "Cl/tt"
3     inputs: [ombh2, omch2, logA, ns, H0, tau]
4     type: NN
5     log: True
6     modes:
7       label: 1
8       range: [2,10000]
9     n_traits:
10      n_hidden: [512, 512, 512, 512]
11    training:
12      validation_split: 0.1
13      learning_rates: [1.e-2, 1.e-3, 1.e-4, 1.e-5, 1.e-6, 1.e-7]
14      batch_sizes: [1000, 2000, 5000, 10000, 20000, 50000]
15      gradient_accumulation_steps: 1
16      patience_values: 100
17      max_epochs: 1000

```

Figure 4.12. Code snippet for network block. We setup a network that emulates $\log_{10}(C_\ell^{TT})(\theta)$ with our six input parameters $\theta = \{\Omega_b h^2, \Omega_c h^2, \log(10^{10} A_s), n_s, h, \tau\}$ and ℓ between 2 and 10000. The network is a fully connected dense neural network with 4 hidden layers of 512 neurons each. Our training block defines the fraction of example spectra used for validation estimation, the learning rates of each learning step, the batch size over which we average, any gradient accumulation steps, patience values, and maximum number of training epochs.

It is also possible to specify derived quantities. This network will automatically use all parameters from the `parameter` block that are computed by the Einstein-Boltzmann code as outputs. So, when we specify a `derived` network in our emulators similar to our C_ℓ^{TT} emulator, we create an emulator that emulates the computation of the nine quantities mentioned in subsection 4.3.1 (which are the nine parameters we listed in Figure 4.11).

In Figure 4.12 we show an example for the `network` block of an emulator trained on primary CMB C_ℓ^{TT} data for $2 \leq \ell \leq 10000$. We discuss the choices made in this block in more detail in section 4.3.

Once the packaging file has been set up with the sections specified above, it becomes easy to generate training data for networks by calling:

```
python -m cosmopower generate <yamlfile>
```

In addition, the `--resume` flag can be used to increase more samples for an already existing set of data points, if it is found afterwards that the training set size is not large enough for training to result in good recovery of spectra from the emulator. When resuming the generation of samples, any pre-existing LHC will be used (if compatible with the given prescription) and any pre-existing samples are not re-generated. This

can be used for continuing a run that was cancelled or stopped before, adding new quantities that were not computed earlier, or increasing the number of samples beyond the LHC that was generated beforehand.

We store the generated training data in `hdf5` files, which are optimised for large, table-like datasets, and allow for both fast read-write access and good data compression. We also include the option to automatically split the data into multiple files, to prevent memory issues from opening a too large a single file at once. For our Λ CDM emulators, this means that we generate about 4 GB worth of training spectra per emulator, split across ten files.

4.4.2 Choices for Emulator design and Training

Because of the large amount of freedom in choosing the hyperparameters for the emulator design and training, it can be hard to determine what settings are optimal for a good training pass. In addition, the impact of certain decisions can wildly vary from either minimal to substantial. As a result, we cannot provide clear guidance on what settings to use but there are a few rules of thumb that can be used when determining the training settings which we recommend:

- The validation split should be about 10-20%;
- Each iteration, the learning rate should go down and the batch size should go up – this makes the emulator learn at more precise steps as it gets closer to a local optimum;
- If a learning iteration reaches the maximum number of epochs instead of a patience value, that means it could have learned for longer, and it hasn't fully optimised yet - try to increase the batch size or learning rate for this iteration or an earlier one.

`CosmoPower` keeps track of the validation loss for every epoch, and saves this to a plain text file for post-training analysis and diagnosis of training issues.

4.5 Comparison of Recovered Cosmology

We now demonstrate that we can use our emulators in parameter inference analysis, generating posterior samples using Monte Carlo chains with wrappers we created for mcmc sampling softwares, interfacing with the same network packaging. We created two wrappers, one for `cobaya` and one for the similar software `CosmoSIS` (Zuntz et al., 2015). In order to utilise all of the output quantities we do this for a set of

Parameter	Fiducial Value
$\Omega_b h^2$	2.2383×10^{-2}
$\Omega_c h^2$	12.011×10^{-2}
H_0	67.32 km/s/Mpc
n_s	0.966
$\log(10^{10} A_s)$	3.0448
τ	5.43×10^{-2}
$\log T_{\text{AGN}}$	7.8
Σm_ν	0.12eV

Table 4.4. The fiducial parameters used for generating the smooth data vector. The first six parameters refer to the cosmology, while $\log T_{\text{AGN}}$ is the baryonic feedback parameter used in the non-linear model of `camb`. The last parameter is specific for the extension model we tested, with a neutrino mass for the inverted hierarchy to ensure that we could recover a closed posterior for our $+\Sigma m_\nu$ emulators. The remaining accuracy settings are the same as in Figure 4.4.

observables: primary CMB, CMB lensing, galaxy weak lensing, and galaxy clustering. Note that this allows for quick and easy cross-validation of the results from using different Einstein-Boltzmann codes between different inference packages (e.g., `CLASS` and `camb` in `cobaya`). This is particularly important because leading cosmology collaborations adopt different combinations of these codes while releasing results which we want to compare and combine.

4.5.1 Simulated data vectors

For full validation, it is important to check that not only the emulators recover the cosmological observables to high accuracy, but also that there is no inherent bias when using our emulators for estimation of the final cosmological parameters. To do this, we can generate simulated data for the observables we emulate with a theoretical covariance matrix and perform a parameter inference analysis on them using the wrappers described above.

Cosmic-variance-limited CMB data

For our testing purposes, we generate a smooth data vector with cosmic-variance-limited noise (such that our conclusions apply to all current and future experiments). This data vector contains data from a fiducial cosmology (see Table 4.4) for the CMB power spectra C_ℓ^{TT} , C_ℓ^{TE} , and C_ℓ^{EE} , as well as the lensing potential spectrum $C_\ell^{\phi\phi}$. For the CMB data vector, the cosmic-variance-limited noise model is similar to Equation 4.6,

with $N_\ell^{XX} = N_\ell^{XY} = 0$ for all combinations of XX and XY . We constrain our analysis to the multipole range $2 \leq \ell \leq 6000$. To explore the parameter space we add a log-likelihood function as a simple Gaussian chi-square distribution:

$$\log \mathcal{L} = -\frac{1}{2} \sum_{\ell} \left(\frac{C_\ell^{\text{pred}} - C_\ell^{\text{data}}}{\sigma_\ell} \right)^2. \quad (4.7)$$

Since the data vector is smooth, we expect to recover the exact input parameters with a final $\chi^2 = 0$. We found that for the fiducial cosmological parameters, `CosmoPower` was able to recover the same χ^2 value down to $\Delta\chi^2 = 0.005$.

4.5.2 Results

Figure 4.13 shows the recovered contours of `cobaya+CosmoPower` versus the `cobaya+camb` posteriors from a CMB cosmic-variance-limited dataset. We show that we can reproduce the `camb` best-fit cosmology and posterior distribution to $< 0.1\sigma$ of the cosmic variance limit error bars in both inference codes. Figure 4.14 shows the same result within `cobaya` for the $+\sum m_\nu$ emulator as an example for an extended model.

For a simple Λ CDM model and the cosmic-variance-limited CMB data, we found that a `camb` chain took ~ 10 hours, while for `CosmoPower` it takes only ~ 20 minutes to run to convergence. Most of this speed-up comes from the fact that at this level of accuracy, an evaluation of a `camb` power spectrum takes $\sim 20s$ to compute, while the same computation takes `CosmoPower` $\sim 0.1s$, at which point computing any non-trivial likelihood function becomes the limiting factor. When going to beyond- Λ CDM models, the time it takes to run a `camb` chain will go up due to the increased complexity or accuracy requirements from the computations. For `CosmoPower` however, the pre-trained emulators do not require more complicated computation when running these chains, and as such the time it takes a `CosmoPower` chain to converge will only increase slightly due to the larger parameter space that needs to be explored.

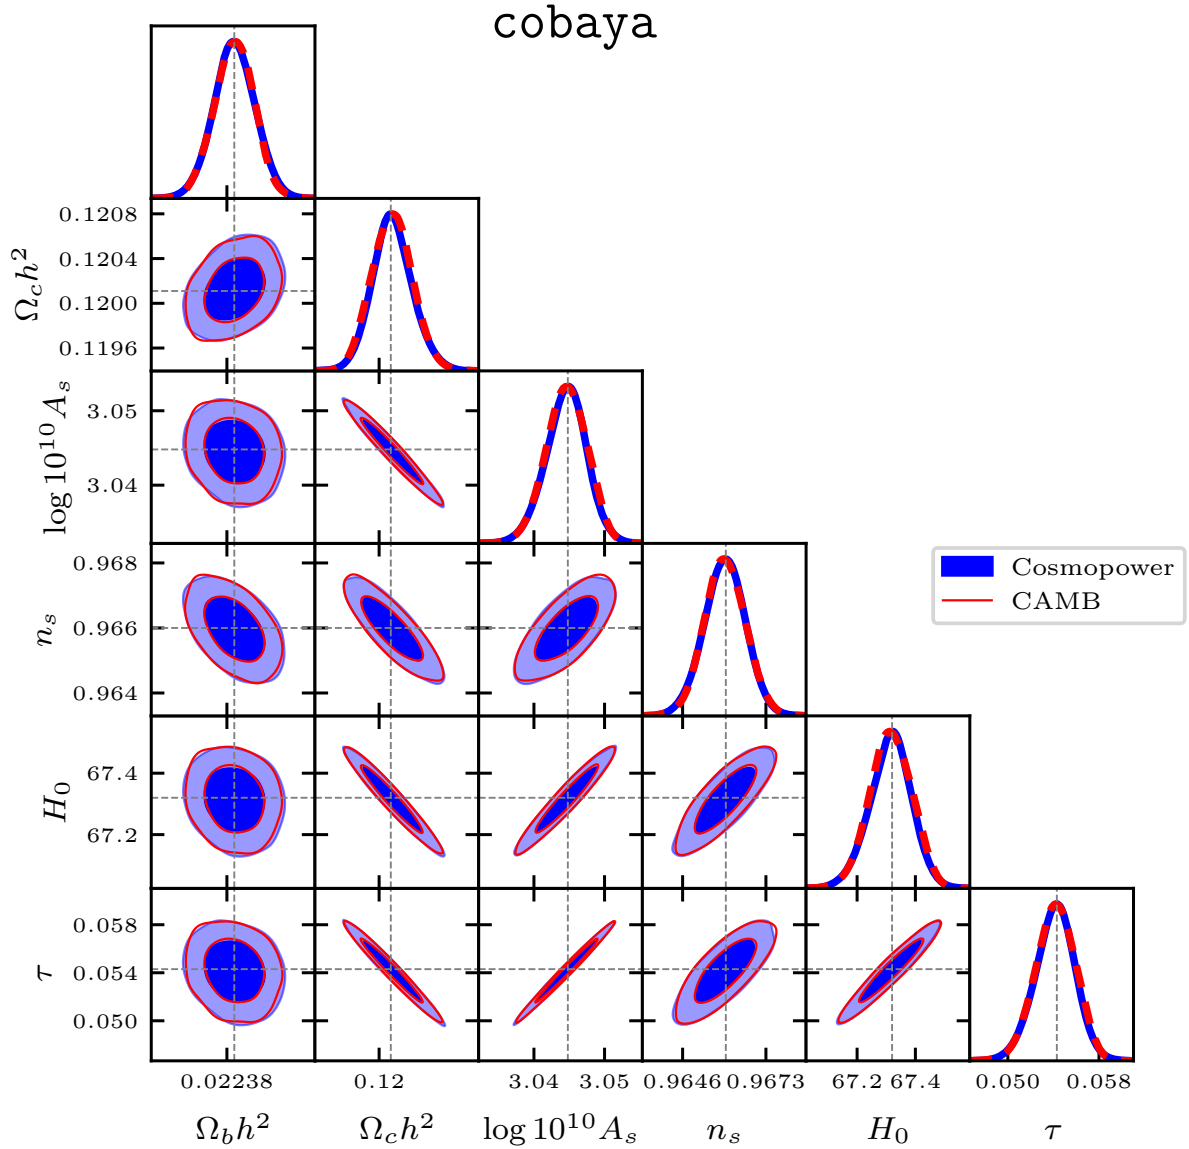


Figure 4.13. To illustrate that we can estimate posteriors in `cobaya` we show the same 68% and 95% confidence levels for Λ CDM parameters from CMB cosmic-variance-limited power spectra, obtained from a full MCMC run done for `CosmoPower` (**blue**) or with the `camb` (**red**) for `cobaya`. The grey dashed lines show the fiducial parameters, and we note a good recovery of these.

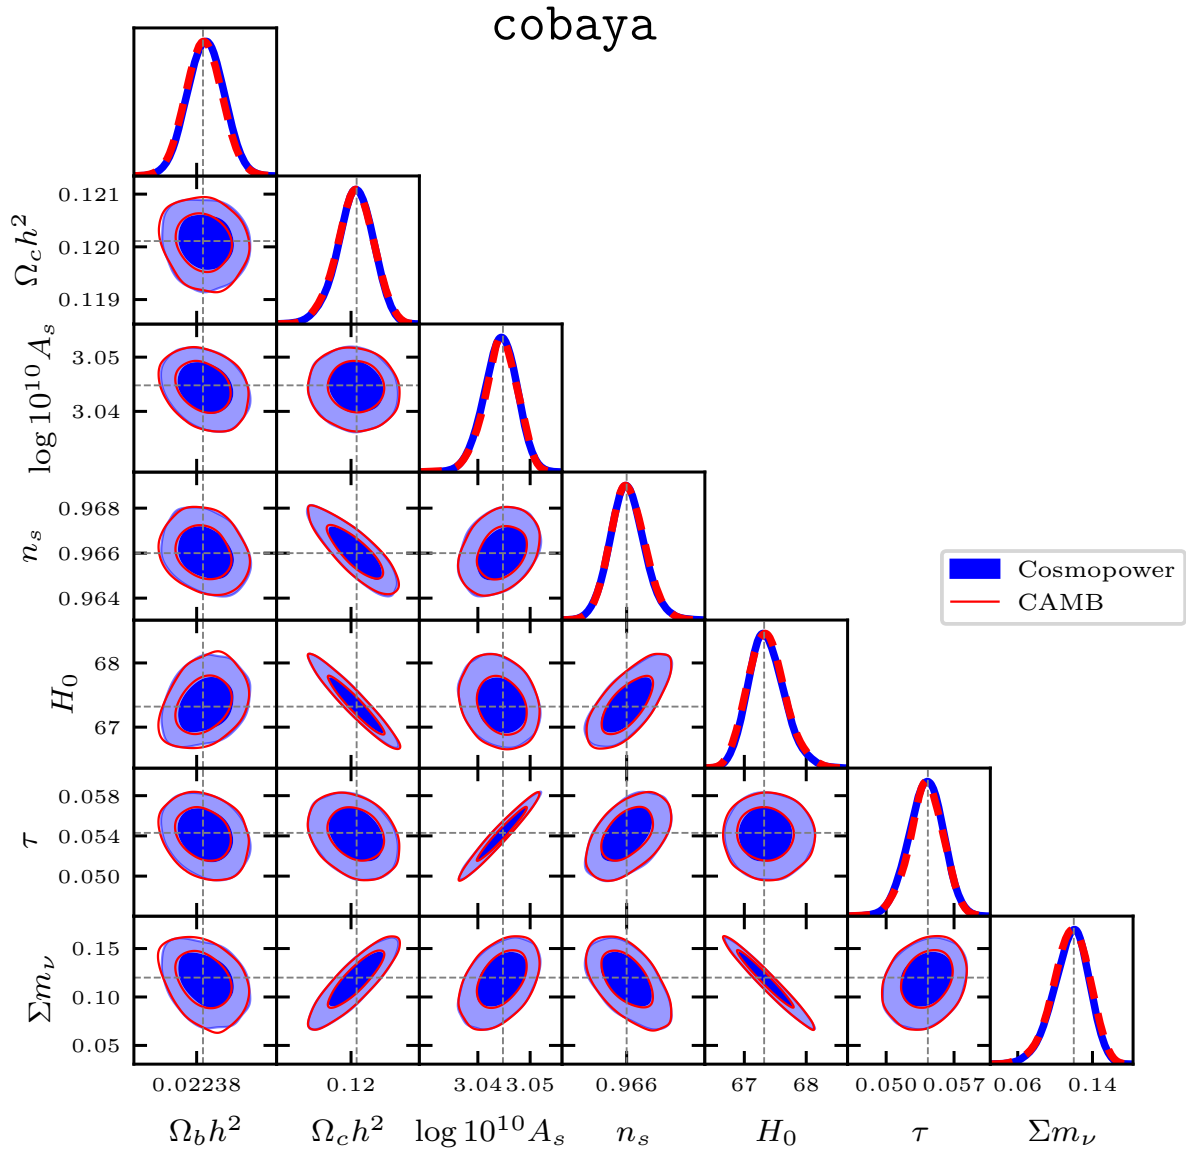


Figure 4.14. The same figure as Figure 4.13, but for the Λ CDM + Σm_ν extension. We note a good recovery of the input parameters, including for the extension parameter.




Figure 4.15. A comparison of the log-posterior (**top panel**) and fractional difference in the log-posterior (**bottom panel**) for a chain run with `camb` and each of the sample points re-evaluated with `CosmoPower`. The chain was run on a mock galaxy-lensing sample using `CosmoSIS`. Note how the `CosmoPower` emulator was able to recover the same log-posterior to within 0.005% of the `camb` log-posterior on the same sample points.

New Results from the *Atacama Cosmology Telescope*

In this chapter, I discuss work I have done on the data analysis for Data Release 6 of the Atacama Cosmology Telescope (ACT DR6). Work on ACT DR6 will be published in a series of collaboration papers in the near future. Here, I summarize my contributions to these future publications.

The makeup of this chapter is as follows:

- In section 5.1, I will go through the specifications of the ACT instrument, a broad overview of the sky area observed by ACT.
- Following in section 5.2, I will give an overview of the observed power spectra in ACT DR6, what we are aiming to measure with it, and what external datasets we will combine with for increased constraints.
- I will summarize the various validation and null tests in section 5.3, and the recovery of Λ CDM parameters in section 5.4.
- Finally, in section 5.5, I will go over the cosmological constraints from the ACT DR6 power spectrum.

All this work has, within the collaboration, been done on data coming from ACT. However, as this work is not yet finalized and the data has not yet been published, I will describe the whole pipeline and expected results, using a simulation of ACT DR6 data as an example. Any constraints from the ACT DR6 CMB power spectrum are forecasts for the upcoming data release.

Some sections regarding the ACT instrument, or the ACT DR6 data are taken from public ACT papers:

- The introduction to section 5.1 is adapted from a similar description of the ACT DR6 lensing dataset in Madhavacheril et al. (2023);
- The details about the polarized arrays used in the ACT DR6 data, given in Table 5.1, is taken from Atkins et al. (2023).

All figures and forecasts presented in this chapter are my own, created as part of my contributions within the ACT collaboration. While the work presented here is part of the upcoming ACT DR6 release and papers, none of the text used here is directly adapted from these papers.

5.1 The Atacama Cosmology Telescope

The Atacama Cosmology Telescope (ACT) was a 6-meter off-axis Gregorian telescope, located at an elevation of 5190 m on Cerro Toco in the Atacama Desert in Chile. Operational from 2007 until 2022, it measured the microwave sky at arcminute resolution at millimeter wavelengths. Since 2016, the telescope was equipped with the Advanced ACTPol (AdvACT) receiver containing arrays of superconducting transition-edge sensor bolometers, sensitive to both temperature and polarization at frequencies centered roughly at 30, 40, 97, 149 and 225 GHz (Fowler et al., 2007; Thornton and et al., 2016); we denote these bands as f_{030} , f_{040} , f_{090} , f_{150} and f_{220} . For the Data Release 6 (DR6) analysis, we use the night-time temperature and polarization AdvACT data collected from 2017 to 2021 covering the CMB-dominated frequency bands f_{090} and f_{150} of detector arrays PA5 and PA6, as well as the small-scale temperature data from f_{220} of PA4 for improved constraints on astrophysical foregrounds. The full size of this survey covers roughly 40% of the sky at arcminute resolution. See Table 5.1 for information on the instrumentation data.

5.2 The DR6 Power Spectra

5.2.1 Cross Spectra

Following the description of the PAs from Table 5.1, ACT DR6 will contain measurements of the power spectrum in temperature and E-mode polarization in three different frequency bands, namely 90, 150, and 220 GHz. The first two of these frequency bands contain both temperature and the E-mode polarization and are measured by two different PAs, PA5 and PA6, while the 220 GHz band is only measured in temperature and only by the PA4 detector array. I will denote the individual measurements of the

Detector Array	Frequency (GHz)	Polarization	Beam (arcmin)
PA4	220	No	1.0
	182 – 277		
PA5	90	Yes	2.0
	77 – 112	Yes	1.4
	150		
	124 – 172		
PA6	90	Yes	2.0
	77 – 112	Yes	1.4
	150		
	124 – 172		

Table 5.1. The different polarized arrays (PAs) that are used in ACT DR6, and the frequencies and modes observed by each, and their beam sizes. The frequency column cites the nominal central frequency of the array, and the full range of the bandpass of that array. The beam size refers to the average beam full width at half maximum (Atkins et al., 2023).

CMB by the different detector arrays in frequency channels as `pa4_f220`, `pa5_f090`, etc.

We end up with 5 measurements of the CMB temperature and 4 of the E-mode polarization. To minimize the impact of the large mode contamination (from e.g. the atmosphere, or ground pickup), the cross-spectra with PA4/PA5/PA6 are cut at $\ell < 1000/600/800$ in temperature respectively, and at $\ell = 500$ in polarization. We use the `pa4_f220` channel only in TT cross-spectra, meaning that we only have TE and EE data in 90 and 150 GHz. We thus have a total of 15 measurements of multifrequency CMB TT autospectra, 16 TE cross-spectra, and 10 EE autospectra.

I show a comparison of the forecasted sensitivity of ACT DR6 with the *Planck* measurement of the CMB in Figure 5.1. This figure shows a comparison of the ACT DR6 forecasted CMB-only power spectra compared to the *Planck* 2018 CMB-only power spectra from Planck Collaboration (2020b). These ACT DR6 power spectra have been foreground-marginalized using the Gibbs sampling procedure explained before in section 3.6, extracting the CMB bandpowers up to $\ell \leq 6000 \sim 6500$ in both temperature and polarization. We found that this ℓ value is optimal for getting the lowest variance measurement of the foreground-marginalized temperature data – as changing this value to a lower/higher value leads to the Gibbs sampling procedure not properly disentangling the CMB bandpowers from the kSZ foreground emission (which is also blackbody). At $\ell = 6000 \sim 6500$, the CMB emission is well enough into the noise floor of the measurement, and and blackbody signal beyond this point can be safely assumed

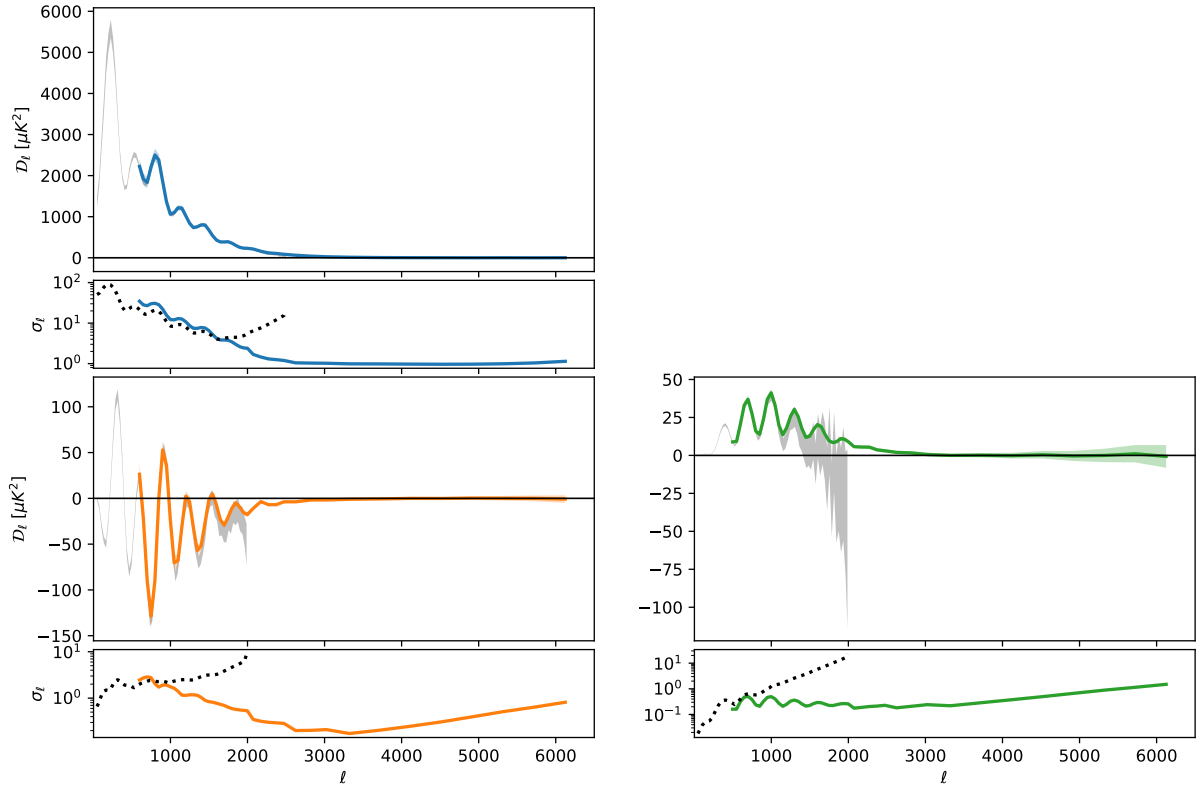


Figure 5.1. A comparison of the forecasted sensitivity of ACT DR6 and the sensitivity of *Planck* in TT (**blue, top-left**), TE (**orange, lower left**), and EE (**green, lower right**). Each corner shows a forecast of a nuisance-marginalized CMB-only power spectrum in the top figure (in these figures, the errors shown are the 5σ confidence intervals for visibility), and a comparison of the errorbar on these measurements in the lower figure (note the log scale on the latter). The bands are the binned power spectra from the *Planck* in grey and the ACT DR6 simulation in blue/orange/green. The ACT DR6 CMB-only forecast is obtained by using the foreground marginalization procedure described before in section 3.6 on a single simulation of the ACT DR6 power spectra, while the *Planck* CMB-only power spectra are described in section 3.5.2 of Planck Collaboration (2020b). We see that ACT DR6 is comparable to *Planck* in TT in the range $600 \leq \ell \leq 2000$, and better at smaller scales and in polarization mode power spectra. We expect ACT DR6 to have μK^2 -sensitivity in TT after combining all cross-spectra, and sub- μK^2 -sensitivity in combined polarization cross-spectra, all down to $\ell = 6000 \sim 6500$.

to be fully from the kSZ effect.

5.2.2 Theory Model

We use the specifications from section 2.1 for the ACT DR6 power spectrum model. For the CMB signal itself, the likelihood code `MFLike` can interface with a CMB theory code, either an Einstein-Boltzmann code such as `camb` or `CLASS`, or our cosmological emulators for `CosmoPower` as described in chapter 4. This interfacing with theory codes is shown earlier in Figure 3.1, and the surrounding text. We model the CMB up to $\ell = 9000$, using high-precision settings for `camb` or `CLASS` to ensure proper modelling of the small-scale features in the CMB power spectra. Our cosmological emulators presented earlier follow settings described in Figure 4.4, and accuracy settings of comparable level were used whenever we used `camb` or `CLASS` directly.

In the remainder of this section, I will go through the details of which part of the model is used in ACT DR6. For details on the model itself, I refer back to the aforementioned sections for an in-depth discussion of these components. The full set of parameters and any priors we set upon them are given in Table 5.2.

Astrophysical Foregrounds

For all frequencies, we apply the *Planck* cosmology mask to our maps first. For the TT spectra, we model the thermal and kinematic Sunyaev–Zel’dovich effects, the clustered and poisson components of the Cosmic Infrared Background, the cross-correlation between tSZ and CIB, residual diffuse Galactic dust emission, and residual unresolved point sources. For polarization components, we only model the last two components. Contrary to the previous DR4 data release, we do not employ an external prior to the amplitude of the CIB templates, since we now include the $\ell < 220$ data which helps constrain it directly. In the unmasked area, the diffuse Galactic foreground emission is on the order of 1% of the CMB in power for TT, TE, and EE. For the power spectra, we employ an astrophysical foreground model as described earlier in subsection 2.2.1. For the diffuse Galactic dust emission, we measure the level of emission separately by cross-correlating our maps with the *Planck* 353 GHz map, and use this as a prior on the amplitude of our model. These priors are imposed on the parameters within `cobaya`. I list the fiducial mean and the forecasted constraints on the TT, TE, and EE dust amplitudes in the last column of Table 5.2. This leaves us with 14 parameters for our model, which is summarized in Table 5.2.

With the inclusion of new, high- ℓ measurements in temperature, we have the ability to differentiate between different astrophysical foreground models. Of interest

Parameter	Description	Prior
A_{tSZ}	Thermal SZ amplitude.	≥ 0
α_{tSZ}	Thermal SZ template shape mismatch.	
A_{kSZ}	Kinematic SZ amplitude.	≥ 0
A_c	Clustered CIB amplitude.	≥ 0
β_c	Clustered CIB spectral index.	
ξ	tSZ-CIB correlation scale.	$0 \leq \xi \leq 0.2$
A_p	Poisson CIB amplitude.	≥ 0
β_p	Poisson CIB spectral index.	
A_s^{TT}	Unresolved radio sources in TT.	
β_s	Spectral index of the radio templates.	≤ 0
A_g^{TT}	Galactic dust amplitude in TT.	$(8.83 \pm 0.32) \mu K^2$
A_s^{TE}	Unresolved radio sources in TE.	
A_g^{TE}	Galactic dust amplitude in TE.	$(0.43 \pm 0.03) \mu K^2$
A_s^{EE}	Unresolved radio sources in EE.	> 0
A_g^{EE}	Galactic dust amplitude in EE.	$(0.165 \pm 0.017) \mu K^2$

Table 5.2. The 15 parameters of the ACT DR6 foreground model, and their priors. See subsection 2.2.1 for a description of the model for these parameters. If not mentioned, we impose an uninformative, wide, uniform prior on the parameter – for most parameters, simply a wide, non-negative prior is used. The additional parameter α_{tSZ} was added to quantify a potential mismatch in the tSZ template shape. Due to the lack of sufficient constraining power, we often consider $\beta_c = \beta_p$, reducing the number of free parameters to 14.

Parameter	Description	Prior
c_{ACT}	Dipole calibration	1 ± 0.003
$c_{\text{pa4}}^{\text{f220}}$ $c_{\text{pa5}}^{\text{f090}}$ $c_{\text{pa5}}^{\text{f150}}$ $c_{\text{pa6}}^{\text{f090}}$ $c_{\text{pa6}}^{\text{f150}}$	Per-frequency array gain calibration	1 ± 0.013 1 ± 0.0016 1 ± 0.0020 1 ± 0.0018 1 ± 0.0024
$p_{\text{pa5}}^{\text{f090}}$ $p_{\text{pa5}}^{\text{f150}}$ $p_{\text{pa6}}^{\text{f090}}$ $p_{\text{pa6}}^{\text{f150}}$	Per-frequency array polarization efficiency	
$\Delta_{\text{pa4}}^{\text{f220}}$ $\Delta_{\text{pa5}}^{\text{f090}}$ $\Delta_{\text{pa5}}^{\text{f150}}$ $\Delta_{\text{pa6}}^{\text{f090}}$ $\Delta_{\text{pa6}}^{\text{f150}}$	Per-frequency array bandpass shift	(0 ± 3.6) GHz (0 ± 1.0) GHz (0 ± 1.3) GHz (0 ± 1.2) GHz (0 ± 1.1) GHz

Table 5.3. The 15 nuisance parameters for our model of the ACT DR6 instrument. See subsection 5.2.2 for a description of the model for these parameters. We impose Gaussian priors on the gain calibrations and bandpass shifts based on external calibration with respect to the *Planck* temperature maps, and measurements of the instrumental bandpass. The polarization efficiencies are free to vary within relatively uninformative flat priors (between $0.9 < p < 1.1$).

to us is the difference between the templates for the thermal Sunyaev-Zel’dovich effect as presented in Battaglia et al. (2010, 2012) and Kim and The AGORA Collaboration (2014). To this end, we added an additional parameter, α_{tSZ} , that affects the shape of the thermal SZ template, modifying Equation 2.9 as

$$\mathcal{D}_{\ell, \nu_1 \nu_2}^{\text{tSZ}} = A_{\text{tSZ}} \mathcal{D}_{\ell, 0}^{\text{tSZ}} \left(\frac{\ell}{\ell_0} \right)^{\alpha_{\text{tSZ}}} \frac{f(\nu_1) f(\nu_2)}{f(\nu_0)^2}, \quad (5.1)$$

where the additional $(\ell/\ell_0)^{\alpha_{\text{tSZ}}}$ term effectively tilts the high- ℓ tail of the tSZ template, if it does not properly match the data. We used the tSZ template from Battaglia et al. (2012) as our baseline, and let the shape parameter α_{tSZ} vary freely. A value of $\alpha_{\text{tSZ}} = 0.0$ would thus prefer the templates from Battaglia et al. (2012), while a value of roughly $\alpha_{\text{tSZ}} = -0.3$ would indicate a tSZ model that is steeper at high- ℓ values and is closer to the one from Kim and The AGORA Collaboration (2014).

Instrumental Systematics

The ACT DR6 instrumental systematics model is made up of 15 parameters. Of these, there are three parameters per frequency array: one overall gain calibration, one polarization efficiency, and one bandpass shift. The only exception is the array PA4, which is unpolarized, and thus does not have a polarization efficiency. On top of these fourteen, there is one additional parameter for an overall calibration with respect to the CMB dipole, which we obtain by calibrating our data to *Planck*.

We obtain priors on the calibration from calibrating our maps to *Planck*, and on the bandpass shifts from measurements of the instrument. We impose only flat, uninformative priors on our per-frequency array polarization efficiencies.

Because our overall calibration to the dipole is obtained from *Planck*, joint runs combining ACT and *Planck* data must use a joint dipole calibration amplitude parameter $c_{\text{ACT}} = A_{\text{Planck}}$.

5.2.3 Inclusion of external datasets

We include several external datasets to the DR6 data. These either involve measurements of the CMB at scales not covered by ACT, or the inclusion of other measurements, such as lensing of the CMB, by the same or different experiments. For the purposes of this thesis, I will focus mainly on external measurements of the large angular scales of the CMB.

Since the optical depth at reionization is constrained by the $\ell < 10$ bump in polarization data, when not combining with any other data, we employ a prior coming from the `SROLL2` analysis of the *Planck* large-scale polarization data (Pagano et al., 2020). We also analyze the ACT DR6 data with a direct inclusion of the `SROLL2` dataset.

To measure the first acoustic peak independently of *Planck*, we combine with the temperature and polarization spectra from *WMAP*. These spectra are noise-dominated in the parts where the ACT data begins, so we ignore the covariance between the two datasets when doing this combination. We ignore the $\ell < 23$ polarization data from *WMAP*, and instead employ either a prior on τ or use the `SROLL2` data. We use the `pyWMAP` implementation of the `WMAPv5` dataset for `cobaya`, which I implemented and described before in section 3.2.

When including the *Planck* data on the scales that are common between *Planck* and ACT, we employ a cut on the *Planck* temperature data to only use the $\ell < 1000$ part of *Planck*, and only the $\ell > 1000$ data of ACT. Since the polarization measurement of *Planck* is noise-dominated in the regime where the ACT data begins, we use the full polarization data from both instruments in this combination. We compare inclusion of

both the `plik` and `NPIPE` analysis of the *Planck* data ("Planck Collaboration", 2020). We use the `plik` implementation which I implemented and described before in section 3.4.

CMB Lensing

Lensing of the CMB by mass clusters has been measured by ACT in Madhavacheril et al. (2023); Qu et al. (2023). CMB lensing indirectly probes the growth of structure at intermediate redshifts of $z \sim 0.5 - 5$, thereby providing additional information on the matter density Ω_m and clustering parameter σ_8 . We include the lensing potential power spectrum $C_\ell^{\phi\phi}$ as measured by ACT in Qu et al. (2023). We employ the baseline `act` model when not using any *Planck* data, or the `act_planck` baseline model when we do include such data (Carron et al., 2022). These datasets are implemented in the `act_dr6_lenslike` package in `cobaya`, allowing easy inclusion of this data for our analyses.

5.3 Likelihood Validation

A crucial first step was to validate the likelihood and analysis pipeline. For this, we needed a series of mock datasets where we know the fiducial parameter value, while also including a noise model that is a realistic representation of the true ACT noise. We then needed to validate that we recovered the input parameters of our simulations, to confirm that our pipeline is working correctly. When that was finished, we did the same null tests that we planned to do for ACT DR6 on our data vector. This would give us an idea of the distribution and confidence intervals of the null results from the actual data.

The methodology we used for performing the validation and null tests on simulated data is similar to the one described in section 6.3.1 of Choi et al. (2020).

5.3.1 Simulated Data Vector

We generated a set of 1600 Gaussian CMB realizations of the full sky. These realizations contain an unlensed CMB component based on the Planck 2018 best-fitting model (Planck Collaboration, 2020d), which are then modified to account for the lensing contribution.

Foreground emission from extragalactic sources is simulated with Gaussian random fields on the full sky as well. The amplitudes of these spectra are fixed at a fiducial value, which was known in advance, and was based on previous ACT results (Choi et al., 2020).

Because foreground emission is slightly non-Gaussian, but our model is Gaussian, we also tested the recovery of the parameters from a single, non-Gaussian realization based on the Agora simulation (Kim and The AGORA Collaboration, 2014). We used a separate model for the SZ emission and measurement of the dust contamination in these simulations. The purpose of this test was to ensure that our foreground model captures enough complexity that our cosmological parameter recovery is not biased by non-Gaussian foreground emission.

5.3.2 Recovery of Fiducial Parameters

In total, we had access to 1600 Gaussian CMB realizations. Because of the time it took to analyse these, we restricted our analysis to a random selection of 100 simulated data vectors for full MCMC analysis, and a selection of 400 simulated data vectors for best-fit analysis, which was faster for certain null tests but required more measurements for an accurate estimate of the mean and variance. To speed up the analysis of 400 data vectors, we used `CosmoPower`, as described in chapter 4 and Jense et al. (2024).

The settings and data cuts on these simulations will be equivalent to what we intend to set for ACT DR6. The only difference is that the simulations will have idealized instrumental properties, meaning they include perfect gain calibration and polarization efficiencies, and no bandpass shifts. Despite this, we still marginalize over these parameters in our analyses.

In addition to recovering the fiducial parameters for our full dataset, we also validated that our likelihood recovers the fiducial parameters from just the temperature and polarization data separately, that is, for only the TT, TE, and EE components separately, as well as for merely the TE/EE components together.

We established that for our likelihood to be unbiased, our recovered parameter mean should not deviate from the fiducial cosmology by more than $0.2\sigma/\sqrt{N}$, where N is the number of simulations we tested. I show the parameter recovery from a subset of 100 simulations in Figure 5.2.

5.4 Parameter Null Tests on Λ CDM

We established a series of null tests for our power spectra. To ensure that our data was consistent with itself, we sought to split the dataset into distinct subsets that were not sensitive to changes in cosmology, and test that these subsets were consistent in recovering the same cosmology. This involves splitting the full dataset into two subsets A and B , and check that for each parameter p , the parameter difference

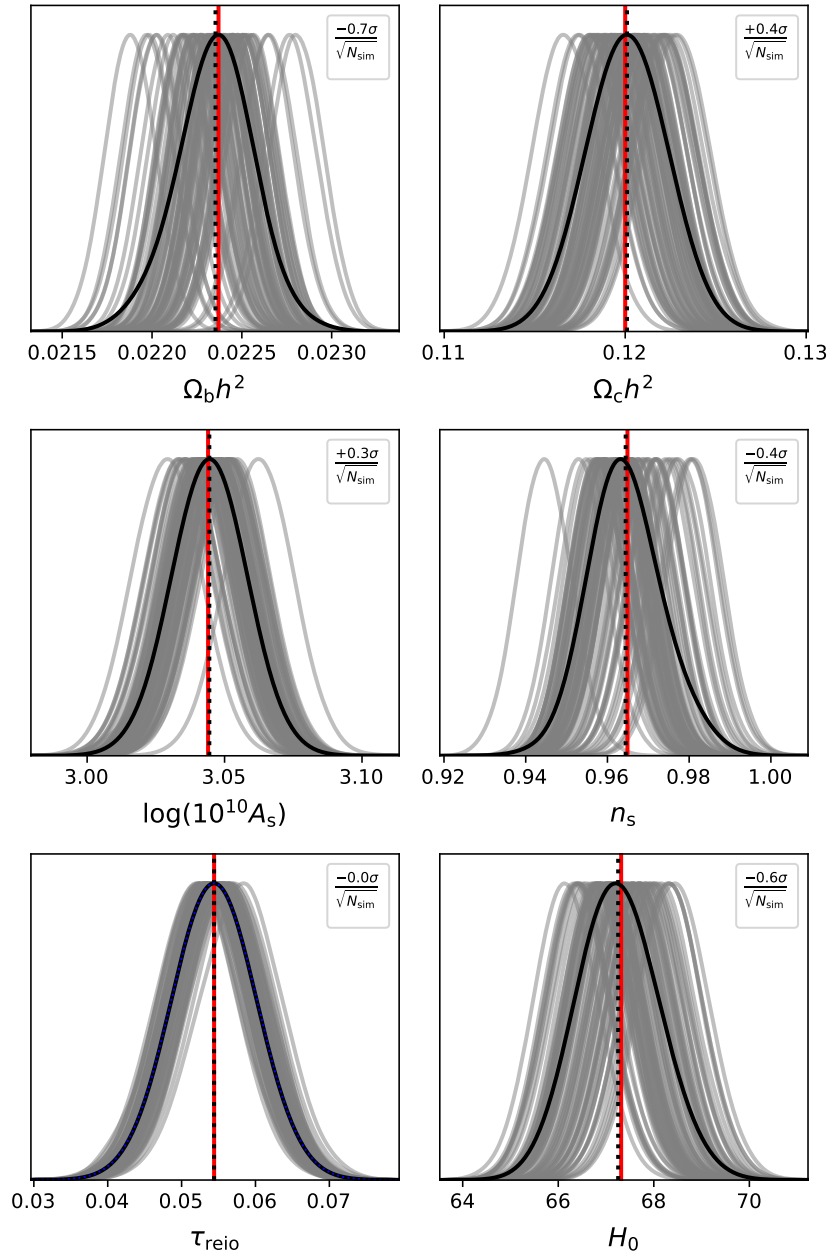


Figure 5.2. A plot of the parameter recovery of 100 simulations of the ACT DR6 data. Each of the posterior parameter distributions of each of the 100 simulations is individually shown in **grey**, and we can see that they scatter around the fiducial parameter, shown in **red**. Additionally, the mixture distribution (viz. the “average distribution”) is shown in **black**, and the recovered mean of all distributions is shown in the **dotted black** line. The optical depth at reionization τ_{reio} is constrained by a prior, shown in **blue**. We see that there appears to be a slight bias in some parameters, as some of the parameter means are $0.2 - 1\sigma/\sqrt{N}$ away from the fiducial mean: we confirmed that this is due to the relatively small size of the subset of simulations we used, and this apparent bias vanishes when we increase the number of simulations used.

$(p_A - p_B) \sim \mathcal{N}(0, \sigma_p)$. For N independently measured cosmological parameters, we compute an N -dimensional χ^2 value

$$\chi^2 = (p_A - p_B) (\Sigma_A + \Sigma_B)^{-1} (p_A - p_B)^T \quad (5.2)$$

where Σ_A, Σ_B are the covariances of the parameters p in datasets A and B respectively. For the ACT cosmology, we have five independently measured cosmological parameters, with only the optical depth of reionization τ_{reio} not measured directly by ACT. Hence we compute this parameter-difference χ^2 value over the remaining five Λ CDM parameters.

Measuring this parameter-level χ^2 value gives us an impression of whether our parameter recovery is biased or not: for an unbiased parameter recovery, the distribution of the computed χ^2 values should tend towards a χ^2 distribution as we increase the number of simulations.

For ACT DR6, there are a variety of splits we could attempt, such as splitting per array (i.e. $A = \text{pa5}, B = \text{pa6}$), and per frequency (i.e. $A = \text{f090}, B = \text{f150}$). Because a per-array split could be biased by instrument-based choices, since for example our ℓ -cuts are different per array, a per-array split could be biased by the fact that one array uses less data than another one. Therefore, we opted to analyze only frequency splits for our null tests. These splits do not include the `pa4_f220` data, but this data is mainly important for constraining astrophysical foregrounds, which are not of interest for these null tests.

I show the results of a parameter null in Figure 5.3. For this figure, I ran 1000 χ^2 -minimizers over the separate `f090` and `f150` data cuts and recorded the bestfit cosmological parameters for each of these minima. In the figure, I show the two-dimensional scatter between each parameter difference - that is, each point shows how much two different parameters differ from each other for a single simulation across two data splits. I show that across all parameters, the minima scatter around zero difference. This means that our likelihood has no parameter bias when performing different data splits, showing a good recovery of parameters across simulations. In addition, the scatter of these simulations gives us an expected confidence region in the five-dimensional parameter space for our data nulls.

5.4.1 Null tests for Beyond- Λ CDM

Since ACT DR6 will provide new insights in small-scale CMB cosmology for the purpose of precision testing of the Λ CDM model, we want to ensure that we have an




Figure 5.3. A 90-150 null for Λ CDM. In each subplot, each datapoint shows the parameter difference between the 90-only and 150-only constraint. We do not include τ_{reio} in this plot as it is not measured by our data and needs a prior to constrain. The black dashed lines show where we assume the center of the distribution should be, centred on zero, while the grey ellipses show the 1 and 2σ confidence ellipses of the parameter differences. We see that the distributions of parameter differences is consistent with zero, showing no bias coming from our data split.

unbiased recovery of beyond- Λ CDM parameters coming from this data. Parameters such as the effective number of relativistic species, N_{eff} , affect the CMB in the small-scale damping tail of the power spectrum. We made similar null tests on our simulated data vectors to see that we accurately recover the fiducial value of N_{eff} with our array and frequency splits. At the same time, opening up additional parameters should not affect our recovery of the original 6 Λ CDM parameters. The recovery of these parameters is shown in Figure 5.4, where we see that opening up an additional beyond- Λ CDM parameter has no impact on our parameter recovery null, while the additional parameter itself is also recovered well.

5.5 Forecasts for Constraints on Cosmological Parameters

Here I present forecasts on constraints on cosmological parameters from ACT DR6 simulations. ACT DR6 will improve over *Planck* at the small scales, and these measurements will improve our constraints on Λ CDM, while combinations with large scales from *WMAP* or *Planck* will improve over the *Planck*-alone constraints that make up Λ CDM constraints today.

We make forecasts for the combination of ACT DR6 alone compared to *Planck* alone, as well as the combination of ACT DR6 + *Planck* together (with the cuts on ACT and *Planck* data described before in subsection 5.2.3). In all cases, the *Planck* data also includes the low- ℓ T and E likelihoods.

To obtain forecasts for ACT DR6, I ran the CMB extraction formalism over a single simulation, extracting the foreground-marginalized CMB power spectrum up to $\ell \lesssim 6000$ in TT, TE, and EE, confirming that there was no bias on the nuisance parameters compared to the multifrequency likelihood. I then ran a chain over this CMB-only data vector, including an additional marginalization over the overall calibration and polarization efficiency parameters. To make a forecast for ACT + *Planck*, I include only the *Planck* data at $\ell < 1000$ and the ACT temperature CMB-only data at $\ell > 1000$ (we do not cut any polarization data from either dataset), ignoring any covariance between the two datasets. Since the ACT simulation is different from the *Planck* data, I only cite the standard deviation of the resulting posterior parameter constraints, showing how much of an improvement the inclusion of ACT DR6 data will have over *Planck* alone.

5.5.1 Constraints on the Λ CDM Model

ACT DR6 will provide a measurement of 5 out of the 6 Λ CDM parameters, with only the optical depth at reionization, τ_{reio} , depending on large-scale polarization




Figure 5.4. A 90-150 null for $\Lambda\text{CDM} + N_{\text{eff}}$. Each datapoint shows the parameter difference between the 90-only and 150-only constraint. We do not include τ_{reio} in this plot as it is not measured by our data range. The black dashed lines show where we assume the center of the distribution should be, centred on zero. We see that the distribution of parameter differences is consistent with zero, showing no bias coming from our data split.

that is outside the regime of ACT observations. When not combining with external measurements, we impose the Gaussian prior $\tau_{\text{reio}} = (5.44 \pm 0.73) \times 10^{-2}$, coming from Pagano et al. (2020).

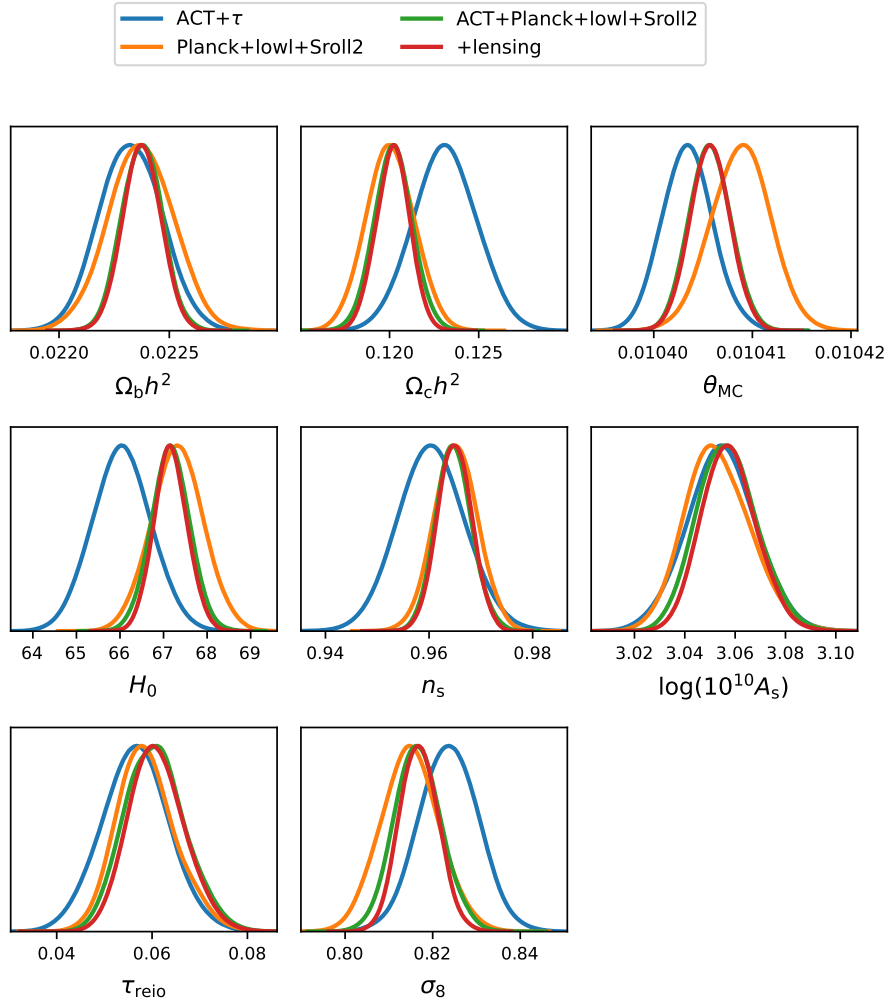
I show the constraints from a measurement of Λ CDM, comparing ACT and *Planck*, and their combination, in Figure 5.5. ACT alone is less constraining than *Planck* on most parameters, mostly because of the lack of large scale features – missing out on the first two acoustic peaks, and the reionization bump in the EE power spectrum. However, the combination of ACT and *Planck* will give constraints on cosmology from both large- and small-scale CMB features, giving constraints much tighter than have ever been obtained from the CMB alone so far: inclusion of the ACT data will improve constraints on five out of the six parameters by 10–55%. The inclusion of high-fidelity, precise measurements of the TE and EE mode polarization power spectra will also greatly improve the constraints on Λ CDM and extension models, as described hereafter.

5.5.2 Constraints on Beyond- Λ CDM Models

The small-scale data of ACT DR6 will give a signal-dominated measurement of the CMB up to $\ell \lesssim 6000$. This new data will give us tests of the Λ CDM model beyond the constraints of *Planck*. There are many models beyond- Λ CDM that are of interest, but to compress it into a single example here, I show the constraints on the single-parameter extension for Λ CDM + N_{eff} here, which is a good proxy for any extension that predominantly affects the small-scale features of the CMB power spectrum.

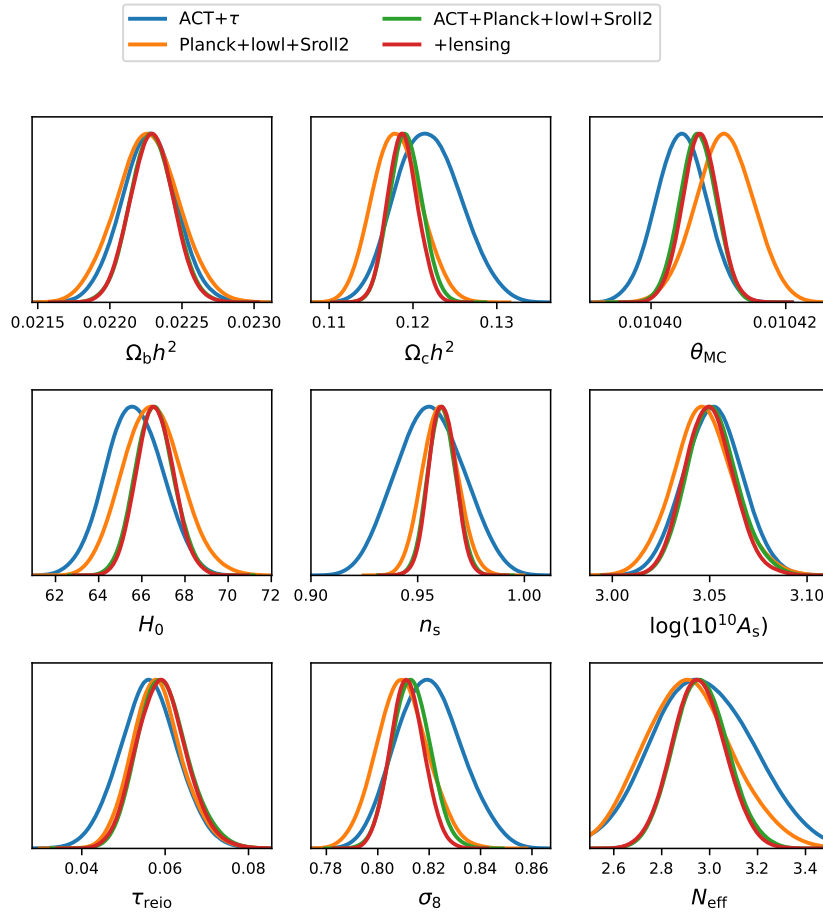
I give an overview of the constraining power on extension models from the ACT DR6 data compared to and combined with *Planck* in Figure 5.6. The table in this figure shows that the extra degree of freedom from the additional parameter blows up most of the errorbars, when compared to Figure 5.5. The inclusion of ACT DR6 not only further constrains the additional parameter, N_{eff} , but it also more tightly constraints the six parameters of Λ CDM within the freedom of this extension.

The forecasts presented here provide an outlook towards the upcoming ACT DR6 data release, including forecasts on the constraining power from the new, precise, small-scale measurement of the CMB anisotropies. These new measurements will provide a new, constraining test for Λ CDM and its extensions.



Parameter	ACT + τ	Planck	ACT +Planck	+lensing
$\Omega_b h^2$	$\pm 1.4 \times 10^{-4}$	$\pm 1.5 \times 10^{-4}$	$\pm 0.90 \times 10^{-4}$	$\pm 0.88 \times 10^{-4}$
$\Omega_c h^2$	$\pm 1.8 \times 10^{-3}$	$\pm 1.2 \times 10^{-3}$	$\pm 1.1 \times 10^{-3}$	$\pm 0.94 \times 10^{-3}$
H_0 [km/s/Mpc]	± 0.67	± 0.60	± 0.44	± 0.39
θ_{MC}	$\pm 2.6 \times 10^{-6}$	$\pm 4.7 \times 10^{-6}$	$\pm 2.1 \times 10^{-6}$	$\pm 2.0 \times 10^{-6}$
n_s	$\pm 6.6 \times 10^{-3}$	$\pm 4.2 \times 10^{-3}$	$\pm 3.5 \times 10^{-3}$	$\pm 3.3 \times 10^{-3}$
$\log(10^{10} A_s)$	$\pm 1.4 \times 10^{-2}$	$\pm 1.4 \times 10^{-2}$	$\pm 1.2 \times 10^{-2}$	$\pm 1.1 \times 10^{-2}$
τ_{reio}	–	$\pm 5.4 \times 10^{-3}$	$\pm 5.4 \times 10^{-3}$	$\pm 5.4 \times 10^{-3}$
σ_8	$\pm 6.8 \times 10^{-3}$	$\pm 8.9 \times 10^{-3}$	$\pm 5.4 \times 10^{-3}$	$\pm 4.4 \times 10^{-3}$

Figure 5.5. Image: A comparison of the constraining power of ACT DR6 (with a prior on τ_{reio} , **blue**), *Planck* (using the `plik` TT/TE/EE+lowl+Sroll12 data, **orange**), ACT +*Planck* (combined with a simple ℓ cut as described in subsection 5.2.3, **green**) and ACT +*Planck* +lensing (also including the ACT DR6 lensing likelihood, **red**). Note that the ACT DR6 simulated data vector is a single Gaussian realization, and thus scatters around a different mean from the other two datasets, which are real data. The usage of *Planck* refers to both the high- ℓ as well as low- ℓ (lowl+Sroll12) datasets. **Table:** A comparison of the (forecasted) error bars from each of these four combinations for Λ CDM.



Parameter	ACT + τ	<i>Planck</i>	ACT + <i>Planck</i>	+lensing
$\Omega_b h^2$	$\pm 1.9 \times 10^{-4}$	$\pm 2.2 \times 10^{-4}$	$\pm 1.5 \times 10^{-4}$	$\pm 1.5 \times 10^{-4}$
$\Omega_c h^2$	$\pm 3.6 \times 10^{-3}$	$\pm 2.9 \times 10^{-3}$	$\pm 1.9 \times 10^{-3}$	$\pm 1.8 \times 10^{-3}$
H_0 [km/s/Mpc]	± 1.3	± 1.4	± 0.90	± 0.85
θ_{MC}	$\pm 3.7 \times 10^{-6}$	$\pm 4.2 \times 10^{-6}$	$\pm 2.7 \times 10^{-6}$	$\pm 2.7 \times 10^{-6}$
n_s	$\pm 16.0 \times 10^{-3}$	$\pm 8.2 \times 10^{-3}$	$\pm 6.4 \times 10^{-3}$	$\pm 6.1 \times 10^{-3}$
$\log(10^{10} A_s)$	$\pm 1.5 \times 10^{-2}$	$\pm 1.6 \times 10^{-2}$	$\pm 1.3 \times 10^{-2}$	$\pm 1.3 \times 10^{-2}$
τ_{reio}	–	$\pm 5.4 \times 10^{-3}$	$\pm 5.4 \times 10^{-3}$	$\pm 5.4 \times 10^{-3}$
σ_8	$\pm 1.3 \times 10^{-2}$	$\pm 1.0 \times 10^{-2}$	$\pm 0.75 \times 10^{-2}$	$\pm 0.66 \times 10^{-2}$
N_{eff}	± 0.20	± 0.19	± 0.11	± 0.11

Figure 5.6. Image: A comparison of the constraining power for an extended model of cosmology, of ACT DR6 (with a prior on τ_{reio} , **blue**), *Planck* (using the plik TT/TE/EE+lowl+Sroll2 data, **orange**), ACT +*Planck* (combined with a simple ℓ cut as described in subsection 5.2.3, **green**) and ACT +*Planck* +lensing (also including the ACT DR6 lensing likelihood, **red**). Note that the ACT DR6 simulated data vector is a single Gaussian realization, and thus scatters around a different mean from the other two datasets, which are real data. The usage of *Planck* refers to both the high- ℓ as well as low- ℓ (lowl+Sroll2) datasets. **Table:** A comparison of the (forecasted) error bars from each of these three datasets for Λ CDM + N_{eff} . The usage of *Planck* refers to both the high- ℓ as well as low- ℓ (lowl+Sroll2) datasets.

CosmoPower package. These emulators can speed up by orders of magnitude the estimation of posteriors on cosmological and nuisance parameters from experimental data and hence enable investigation of models which extend the fiducial Λ CDM cosmology and the checking of the robustness of any conclusions made to a plethora of modelling choices. By creating a specification for packaging and distributing such emulators and providing wrappers for their use in popular inference packages we hope to improve efficiency and reproducibility in cosmological studies, by allowing appropriate emulators to be widely used by many different studies once they have been trained. This kind of reproducibility across platforms will also assist in combining different data sets to improve statistical constraining power and investigate more models in more detail.

I have used the framework to produce a suite of emulators of quantities calculated by `camb` v1.5.2: CMB primary angular power spectra C_ℓ^{TT} , C_ℓ^{TE} , C_ℓ^{EE} , C_ℓ^{BB} ; CMB lensing power spectra $C_\ell^{\phi\phi}$; linear and non-linear matter power $P(k)_{\text{lin}}$, $P(k)_{\text{NL}}$ and a variety of background and derived quantities. I have demonstrated the accuracy of the emulators at both the spectrum level and the parameter-recovery level to accuracy appropriate for Stage-IV data (and beyond to the cosmic variance limit for the CMB spectra).

6.3 Future Cosmological Constraints from CMB Experiments

The software I have contributed to will play a crucial role in the upcoming ACT DR6 data release, and the future *Simons Observatory* data analysis. I have shown a wide range of tests we have performed on simulated data vectors for the ACT DR6 data release, which we hope to publish in the near future. The work I presented here includes forecasts for the constraining power on the CMB-only data, the unbiased recovery of cosmological parameters from Gaussian simulations of the ACT DR6 power spectra, and forecasts for Λ CDM and beyond- Λ CDM parameter constraints with ACT DR6, both alone and combined with *Planck*.

When published, the ACT DR6 power spectra will provide the most stringent test for Λ CDM that we have to date. They will not only provide a test for Λ CDM independent from *Planck*, serving as a separate validation of the *Planck* results, but the joint ACT and *Planck* data will provide a new test of cosmological and physical models from the CMB.

Once this data are available for the public, work will continue in training people to use the then publicly-available likelihood software to perform their own analyses. On top of the many new results coming from that, this will also provide a platform

to thoroughly test `MFLike` for the future release of SO data in the years to come. The continuing analysis of the impact of astrophysical and instrumental nuisances – which will be of high importance for SO.

The creation of emulators for (beyond)- Λ CDM cosmology will increase in importance with the future prospect of cross-probe analysis. The continuing pressure on computational resources with the expansion of observational windows and improvement of constraining power will lead to greater benefits in the use of emulators. Our work in the development of a coherent suite of software for the usage of `CosmoPower` emulators for the development and application of emulators will open avenues for new use of emulators in single- and multi-probe cosmological analysis.

All this work lays important foundations for the future of cosmological analysis. Looking back at only the previous twenty years of scientific breakthroughs made in cosmology, it is intriguing and amazing to think what the next ten or twenty years will hold for us.

Bibliography

- K. N. Abazajian and et al. Cmb-s4 science book, first edition, 2016. URL <https://arxiv.org/abs/1610.02743>.
- G. E. Addison, J. Dunkley, and D. N. Spergel. Modelling the correlation between the thermal Sunyaev Zel'dovich effect and the cosmic infrared background. *Monthly Notices of the Royal Astronomical Society*, 427(2):1741–1754, 12 2012. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2012.21664.x. URL <https://doi.org/10.1111/j.1365-2966.2012.21664.x>.
- P. Ade and et al. Improved Constraints on Primordial Gravitational Waves using Planck, WMAP, and BICEP/Keck Observations through the 2018 Observing Season. *Physical Review Letters*, 127(15), Oct. 2021. ISSN 1079-7114. doi: 10.1103/physrevlett.127.151301. URL <http://dx.doi.org/10.1103/PhysRevLett.127.151301>.
- S. Aiola, E. Calabrese, L. Maurin, S. Naess, B. L. Schmitt, and et al. The Atacama Cosmology Telescope: DR4 maps and cosmological parameters. *Journal of Cosmology and Astroparticle Physics*, 2020(12):047, dec 2020. doi: 10.1088/1475-7516/2020/12/047. URL <https://dx.doi.org/10.1088/1475-7516/2020/12/047>.
- R. A. Alpher. A neutron-capture theory of the formation and relative abundance of the elements. *Phys. Rev.*, 74:1577–1589, Dec 1948. doi: 10.1103/PhysRev.74.1577. URL <https://link.aps.org/doi/10.1103/PhysRev.74.1577>.
- Z. Atkins, A. J. Duivenvoorden, W. R. Coulton, F. J. Qu, and et al. The atacama cosmology telescope: Map-based noise simulations for dr6, 2023. URL <https://arxiv.org/abs/2303.04180>.
- E. Aver, K. A. Olive, and E. D. Skillman. The effects of He I λ 10830 on helium abundance determinations. *Journal of Cosmology and Astroparticle Physics*, 2015(07):011–011, July 2015. ISSN 1475-7516. doi: 10.1088/1475-7516/2015/07/011. URL <http://dx.doi.org/10.1088/1475-7516/2015/07/011>.

- L. Balkenhol, D. Dutcher, A. Spurio Mancini, A. Doussot, K. Benabed, S. Galli, and et al. Measurement of the CMB temperature power spectrum and constraints on cosmology from the SPT-3G 2018 TT , TE , and EE dataset. *Phys. Rev. D*, 108:023510, Jul 2023. doi: 10.1103/PhysRevD.108.023510. URL <https://link.aps.org/doi/10.1103/PhysRevD.108.023510>.
- N. Battaglia, J. R. Bond, C. Pfrommer, J. L. Sievers, and D. Sijacki. Simulations of the sunyaev-zel'dovich power spectrum with active galactic nucleus feedback. *The Astrophysical Journal*, 725(1):91, nov 2010. doi: 10.1088/0004-637X/725/1/91. URL <https://dx.doi.org/10.1088/0004-637X/725/1/91>.
- N. Battaglia, J. R. Bond, C. Pfrommer, and J. L. Sievers. On the cluster physics of sunyaev-zel'dovich and x-ray surveys. ii. deconstructing the thermal sz power spectrum. *The Astrophysical Journal*, 758(2):75, sep 2012. doi: 10.1088/0004-637X/758/2/75. URL <https://dx.doi.org/10.1088/0004-637X/758/2/75>.
- R. H. Becker, X. Fan, R. L. White, M. A. Strauss, V. K. Narayanan, R. H. Lupton, J. E. Gunn, J. Annis, N. A. Bahcall, J. Brinkmann, A. J. Connolly, I. Csabai, P. C. Czarapata, M. Doi, T. M. Heckman, G. S. Hennessy, Ž. Ivezić, G. R. Knapp, D. Q. Lamb, T. A. McKay, J. A. Munn, T. Nash, R. Nichol, J. R. Pier, G. T. Richards, D. P. Schneider, C. Stoughton, A. S. Szalay, A. R. Thakar, and D. G. York. Evidence for Reionization at $z \sim 6$: Detection of a Gunn-Peterson Trough in a $z = 6.28$ Quasar. *The Astronomical Journal*, 122(6):2850–2857, Dec. 2001. ISSN 0004-6256. doi: 10.1086/324231. URL <http://dx.doi.org/10.1086/324231>.
- C. L. Bennett and et al. Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. *The Astrophysical Journal Supplement Series*, 208(2):20, Sept. 2013. ISSN 1538-4365. doi: 10.1088/0067-0049/208/2/20. URL <http://dx.doi.org/10.1088/0067-0049/208/2/20>.
- C. L. Bennett, D. Larson, J. L. Weiland, N. Jarosik, G. Hinshaw, N. Odegard, K. M. Smith, R. S. Hill, B. Gold, M. Halpern, E. Komatsu, M. R. Nolte, L. Page, D. N. Spergel, E. Wollack, J. Dunkley, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. L. Wright. Nine-year wilkinson microwave anisotropy probe (wmap) observations: Final maps and results. *The Astrophysical Journal Supplement Series*, 208(2):20, Sept. 2013. ISSN 1538-4365. doi: 10.1088/0067-0049/208/2/20. URL <http://dx.doi.org/10.1088/0067-0049/208/2/20>.
- J. J. Bennett, G. Buldgen, M. Drewes, and Y. Y. Wong. Towards a precision calculation of the effective number of neutrinos n_{eff} in the standard model: the qed equation of

-
- state. *Journal of Cosmology and Astroparticle Physics*, 2020(03):003–003, Mar. 2020. ISSN 1475-7516. doi: 10.1088/1475-7516/2020/03/003. URL <http://dx.doi.org/10.1088/1475-7516/2020/03/003>.
- H. Bhat and N. Kumar. On the derivation of the bayesian information criterion. 01 2010.
- B. Bolliet, A. Spurio Mancini, J. C. Hill, M. Madhavacheril, H. T. Jense, E. Calabrese, and J. Dunkley. High-accuracy emulators for observables in Λ CDM, N_{eff} , Σm_ν , and w cosmologies. 2023.
- J. R. Bond and et al. CMB Analysis of Boomerang & Maxima & the Cosmic Parameters Ω_{tot} , $\Omega_b h^2$, $\Omega_c d m h^2$, Ω_{Λ} , n_s , 2000.
- E. Calabrese, R. Hložek, and et al. Precision epoch of reionization studies with next-generation CMB experiments. *Journal of Cosmology and Astroparticle Physics*, 2014(08): 010–010, Aug. 2014. ISSN 1475-7516. doi: 10.1088/1475-7516/2014/08/010. URL <http://dx.doi.org/10.1088/1475-7516/2014/08/010>.
- J. E. Carlstrom and et al. The 10 Meter South Pole Telescope. *Publications of the Astronomical Society of the Pacific*, 123(903):568, apr 2011. doi: 10.1086/659879. URL <https://dx.doi.org/10.1086/659879>.
- J. Carron, M. Mirmelstein, and A. Lewis. CMB lensing from Planck PR4 maps. *Journal of Cosmology and Astroparticle Physics*, 2022(09):039, Sept. 2022. ISSN 1475-7516. doi: 10.1088/1475-7516/2022/09/039. URL <http://dx.doi.org/10.1088/1475-7516/2022/09/039>.
- C. Cartis, J. Fiala, B. Marteau, and L. Roberts. Improving the Flexibility and Robustness of Model-Based Derivative-Free Optimization Solvers, 2018. URL <https://arxiv.org/abs/1804.00154>.
- C. Cartis, L. Roberts, and O. Sheridan-Methven. Escaping local minima with local derivative-free methods: a numerical investigation. *Optimization*, 71(8):2343–2373, Feb. 2021. ISSN 1029-4945. doi: 10.1080/02331934.2021.1883015. URL <http://dx.doi.org/10.1080/02331934.2021.1883015>.
- J. Chluba and R. M. Thomas. Towards a complete treatment of the cosmological recombination problem: Towards a complete treatment of recombination. *Monthly Notices of the Royal Astronomical Society*, pages no–no, Dec. 2010. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2010.17940.x. URL <http://dx.doi.org/10.1111/j.1365-2966.2010.17940.x>.

- J. Chluba, G. M. Vasil, and L. J. Dursi. Recombinations to the rydberg states of hydrogen and their effect during the cosmological recombination epoch: Recombinations to the rydberg states of hi. *Monthly Notices of the Royal Astronomical Society*, 407(1): 599–612, July 2010. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2010.16940.x. URL <http://dx.doi.org/10.1111/j.1365-2966.2010.16940.x>.
- S. K. Choi, M. Hasselfield, S.-P. P. Ho, B. Koopman, M. Lungu, and et al. The Atacama Cosmology Telescope: a measurement of the Cosmic Microwave Background power spectra at 98 and 150 GHz. *Journal of Cosmology and Astroparticle Physics*, 2020(12): 045–045, Dec. 2020. ISSN 1475-7516. doi: 10.1088/1475-7516/2020/12/045. URL <http://dx.doi.org/10.1088/1475-7516/2020/12/045>.
- W. R. Coulton, M. S. Madhavacheril, A. J. Duivenvoorden, J. C. Hill, and I. A.-C. et al. The atacama cosmology telescope: High-resolution component-separated maps across one-third of the sky. 2023.
- Dark Energy Survey Collaboration. The Dark Energy Survey. 2005. URL <https://arxiv.org/abs/astro-ph/0510346>.
- Dark Energy Survey Collaboration and Kilo-Degree Survey Collaboration. Des y3 + kids-1000: Consistent cosmology combining cosmic shear surveys. *The Open Journal of Astrophysics*, 6, Oct. 2023. ISSN 2565-6120. doi: 10.21105/astro.2305.17173. URL <http://dx.doi.org/10.21105/astro.2305.17173>.
- S. Das, B. D. Sherwin, and et al. Detection of the power spectrum of cosmic microwave background lensing by the atacama cosmology telescope. *Phys. Rev. Lett.*, 107:021301, Jul 2011. doi: 10.1103/PhysRevLett.107.021301. URL <https://link.aps.org/doi/10.1103/PhysRevLett.107.021301>.
- P. de Bernardis and et al. A flat Universe from high-resolution maps of the cosmic microwave background radiation. *Nature*, 404(6781):955–959, Apr. 2000. ISSN 1476-4687. doi: 10.1038/35010035. URL <http://dx.doi.org/10.1038/35010035>.
- DESI Collaboration. The DESI Experiment Part I: Science, Targeting, and Survey Design, 2016. URL <https://arxiv.org/abs/1611.00036>.
- DESI Collaboration. Overview of the Instrumentation for the Dark Energy Spectroscopic Instrument. *The Astronomical Journal*, 164(5):207, Oct. 2022. ISSN 1538-3881. doi: 10.3847/1538-3881/ac882b. URL <http://dx.doi.org/10.3847/1538-3881/ac882b>.

- DESI Collaboration. Desi 2024 vi: Cosmological constraints from the measurements of baryon acoustic oscillations, 2024. URL <https://arxiv.org/abs/2404.03002>.
- J. Dunkley, E. Komatsu, M. R.olta, D. N. Spergel, D. Larson, G. Hinshaw, and L. P. et al. Five-year wilkinson microwave anisotropy probe observations: Likelihoods and parameters from the wmap data. *The Astrophysical Journal Supplement Series*, 180 (2):306, feb 2009. doi: 10.1088/0067-0049/180/2/306. URL <https://dx.doi.org/10.1088/0067-0049/180/2/306>.
- J. Dunkley, R. Hlozek, J. Sievers, and et al. The Atacama Cosmology Telescope: Cosmological Parameters from the 2008 Power Spectrum. *The Astrophysical Journal*, 739(1):52, sep 2011. doi: 10.1088/0004-637X/739/1/52. URL <https://dx.doi.org/10.1088/0004-637X/739/1/52>.
- J. Dunkley, E. Calabrese, J. Sievers, and et al. The atacama cosmology telescope: likelihood for small-scale cmb data. *Journal of Cosmology and Astroparticle Physics*, 2013 (07):025, jul 2013. doi: 10.1088/1475-7516/2013/07/025. URL <https://dx.doi.org/10.1088/1475-7516/2013/07/025>.
- Euclid Collaboration, Y. Mellier, and et al. Euclid. i. overview of the euclid mission, 2024. URL <https://arxiv.org/abs/2405.13491>.
- D. J. Fixsen. The temperature of the cosmic microwave background. *The Astrophysical Journal*, 707(2):916–920, Nov. 2009. ISSN 1538-4357. doi: 10.1088/0004-637x/707/2/916. URL <http://dx.doi.org/10.1088/0004-637X/707/2/916>.
- J. W. Fowler, M. D. Niemack, S. R. Dicker, and et al. Optical design of the Atacama Cosmology Telescope and the Millimeter Bolometric Array Camera. *Applied Optics*, 46(17):3444, May 2007. ISSN 1539-4522. doi: 10.1364/ao.46.003444. URL <http://dx.doi.org/10.1364/AO.46.003444>.
- N. Galitzki and et al. The simons observatory: instrument overview. July 2018. doi: 10.1117/12.2312985. URL <http://dx.doi.org/10.1117/12.2312985>.
- S. Giardiello, M. Gerbino, L. Pagano, D. Alonso, B. Beringue, B. Bolliet, E. Calabrese, G. Coppi, J. Errard, G. Fabbian, I. Harrison, J. C. Hill, H. T. Jense, B. Keating, A. L. Posta, M. Lattanzi, A. I. Lonappan, G. Puglisi, C. L. Reichardt, and S. M. Simon. The simons observatory: impact of bandpass, polarization angle and calibration uncertainties on small-scale power spectrum analysis, 2024.

- B. Hadzhiyska, S. Ferraro, B. R. Guachalla, E. Schaan, and et al. Evidence for large baryonic feedback at low and intermediate redshifts from kinematic Sunyaev-Zel'dovich observations with ACT and DESI photometric galaxies, 2024. URL <https://arxiv.org/abs/2407.07152>.
- N. W. Halverson, E. M. Leitch, C. Pryke, J. Kovac, J. E. Carlstrom, W. L. Holzapfel, M. Dragovan, J. K. Cartwright, B. S. Mason, S. Padin, T. J. Pearson, A. C. S. Readhead, and M. C. Shepherd. Degree Angular Scale Interferometer First Results: A Measurement of the Cosmic Microwave Background Angular Power Spectrum. *The Astrophysical Journal*, 568(1):38–45, Mar. 2002. ISSN 1538-4357. doi: 10.1086/338879. URL <http://dx.doi.org/10.1086/338879>.
- S. Hanany, P. Ade, A. Balbi, J. Bock, J. Borrill, A. Boscaleri, P. de Bernardis, P. G. Ferreira, V. V. Hristov, A. H. Jaffe, A. E. Lange, A. T. Lee, P. D. Mauskopf, C. B. Netterfield, S. Oh, E. Pascale, B. Rabii, P. L. Richards, G. F. Smoot, R. Stompor, C. D. Winant, and J. H. P. Wu. Maxima-1: A measurement of the cosmic microwave background anisotropy on angular scales of 10[arcmin]–5°. *The Astrophysical Journal*, 545(1):L5–L9, Dec. 2000. ISSN 0004-637X. doi: 10.1086/317322. URL <http://dx.doi.org/10.1086/317322>.
- W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 04 1970. ISSN 0006-3444. doi: 10.1093/biomet/57.1.97. URL <https://doi.org/10.1093/biomet/57.1.97>.
- B. S. Hensley, S. E. Clark, V. Fanfani, N. Krachmalnicoff, G. Fabbian, D. Poletti, G. Puglisi, G. Coppi, J. Nibauer, R. Gerasimov, N. Galitzki, S. K. Choi, and et al. The simons observatory: Galactic science goals and forecasts. *The Astrophysical Journal*, 929(2):166, Apr. 2022. ISSN 1538-4357. doi: 10.3847/1538-4357/ac5e36. URL <http://dx.doi.org/10.3847/1538-4357/ac5e36>.
- J. C. Hill, E. Calabrese, and et al. Atacama cosmology telescope: Constraints on prerecombination early dark energy. *Phys. Rev. D*, 105:123536, Jun 2022. doi: 10.1103/PhysRevD.105.123536. URL <https://link.aps.org/doi/10.1103/PhysRevD.105.123536>.
- M. Hilton, C. Sifón, S. Naess, M. Madhavacheril, M. Oguri, E. Rozo, E. Rykoff, and et al. The Atacama Cosmology Telescope: A Catalog of >4000 Sunyaev-Zel'dovich Galaxy Clusters. *The Astrophysical Journal Supplement*, 253(1):3, Mar. 2021. doi: 10.3847/1538-4365/abd023.

- G. Hinshaw and et al. Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results. *The Astrophysical Journal Supplement Series*, 208(2):19, Sept. 2013. ISSN 1538-4365. doi: 10.1088/0067-0049/208/2/19. URL <http://dx.doi.org/10.1088/0067-0049/208/2/19>.
- G. Hinshaw, D. N. Spergel, L. Verde, R. S. Hill, and S. S. M. et al. First-year wilkinson microwave anisotropy probe (wmap) observations: The angular power spectrum. *The Astrophysical Journal Supplement Series*, 148(1):135, sep 2003. doi: 10.1086/377225. URL <https://dx.doi.org/10.1086/377225>.
- G. Hinshaw, M. R. Nolta, C. L. Bennett, and et al. Three-year wilkinson microwave anisotropy probe (wmap) observations: Temperature analysis. *The Astrophysical Journal Supplement Series*, 170(2):288, jun 2007. doi: 10.1086/513698. URL <https://dx.doi.org/10.1086/513698>.
- G. Hinshaw, D. Larson, E. Komatsu, D. N. Spergel, C. L. Bennett, J. Dunkley, M. R. Nolta, M. Halpern, R. S. Hill, N. Odegard, L. Page, K. M. Smith, J. L. Weiland, B. Gold, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, E. Wollack, and E. L. Wright. Nine-year wilkinson microwave anisotropy probe (wmap) observations: Cosmological parameter results. *The Astrophysical Journal Supplement Series*, 208(2): 19, sep 2013. doi: 10.1088/0067-0049/208/2/19. URL <https://dx.doi.org/10.1088/0067-0049/208/2/19>.
- K. Hotokezaka, E. Nakar, O. Gottlieb, S. Nissanke, K. Masuda, G. Hallinan, K. P. Mooley, and A. T. Deller. A hubble constant measurement from superluminal motion of the jet in gw170817, 2018. URL <https://arxiv.org/abs/1806.10596>.
- Z. Hou, R. Keisler, L. Knox, M. Millea, and C. Reichardt. How massless neutrinos affect the cosmic microwave background damping tail. *Physical Review D*, 87(8), Apr. 2013. ISSN 1550-2368. doi: 10.1103/physrevd.87.083008. URL <http://dx.doi.org/10.1103/PhysRevD.87.083008>.
- J. N. Lockyer, Esq. Notice of an observation of the spectrum of a solar prominence. *Proceedings of the Royal Society of London*, 17:91–92, 1868.
- H. T. Jense, I. Harrison, E. Calabrese, A. Spurio Mancini, B. Bolliet, J. Dunkley, and J. C. Hill. A complete framework for cosmological emulation and inference with cosmopower. *ArXiv e-Prints*, 2024.

BIBLIOGRAPHY

- J.-h. Kim and The AGORA Collaboration. The AGORA High-resolution Galaxy Simulations Comparison Project. *ApJS*, 210(1):14, Jan. 2014. doi: 10.1088/0067-0049/210/1/14.
- R. S. Klessen and S. C. O. Glover. The first stars: formation, properties, and impact, 2023. URL <https://arxiv.org/abs/2303.12500>.
- Larry Wasserman. *All of Statistics: A Concise Course in Statistical Inference*. Springer Science and Business Media, Inc., 2004.
- E. M. Leitch, J. M. Kovac, N. W. Halverson, J. E. Carlstrom, C. Pryke, and M. W. E. Smith. Degree angular scale interferometer 3 year cosmic microwave background polarization results. *The Astrophysical Journal*, 624(1):10–20, May 2005. ISSN 1538-4357. doi: 10.1086/428825. URL <http://dx.doi.org/10.1086/428825>.
- G. Lemaître. Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques. *Annales de la Société Scientifique de Bruxelles*, 47:49–59, Jan. 1927.
- A. Lewis. Efficient sampling of fast and slow cosmological parameters. *Physical Review D*, 87(10), May 2013. ISSN 1550-2368. doi: 10.1103/physrevd.87.103529. URL <http://dx.doi.org/10.1103/PhysRevD.87.103529>.
- A. Lewis, A. Challinor, and A. Lasenby. Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models. *ApJ*, 538(2):473–476, Aug. 2000. doi: 10.1086/309179.
- Z. Li, T. Louis, E. Calabrese, H. Jense, D. Alonso, J. R. Bond, S. K. Choi, J. Dunkley, G. Fabbian, X. Garrido, A. H. Jaffe, M. S. Madhavacheril, P. D. Meerburg, U. Natale, and F. J. Qu. The simons observatory: a new open-source power spectrum pipeline applied to the planck legacy data. 2021.
- A. R. Liddle. How many cosmological parameters? *Monthly Notices of the Royal Astronomical Society*, 351(3):L49–L53, July 2004. ISSN 1365-2966. doi: 10.1111/j.1365-2966.2004.08033.x. URL <http://dx.doi.org/10.1111/j.1365-2966.2004.08033.x>.
- LiteBIRD Collaboration. Probing cosmic inflation with the litebird cosmic microwave background polarization survey. *Progress of Theoretical and Experimental Physics*, 2023(4), Nov. 2022. ISSN 2050-3911. doi: 10.1093/ptep/ptac150. URL <http://dx.doi.org/10.1093/ptep/ptac150>.

- N. MacCrann, F. J. Qu, T. Namikawa, B. Bolliet, H. Cai, E. Calabrese, S. K. Choi, O. Darwish, S. Ferraro, Y. Guan, J. C. Hill, M. Hilton, R. Hložek, D. Kramer, M. S. Madhavacheril, K. Moodley, N. Sehgal, B. D. Sherwin, C. Sifón, S. T. Staggs, H. Trac, A. V. Engelen, and E. M. Vavagiakis. The Atacama Cosmology Telescope: Reionization kSZ trispectrum methodology and limits, 2024. URL <https://arxiv.org/abs/2405.01188>.
- M. S. Madhavacheril, F. J. Qu, B. D. Sherwin, N. MacCrann, and Y. L. et al. The atacama cosmology telescope: Dr6 gravitational lensing map and cosmological parameters. 2023.
- F. McCarthy, J. C. Hill, and M. S. Madhavacheril. Baryonic feedback biases on fundamental physics from lensed cmb power spectra. *Phys. Rev. D*, 105:023517, Jan 2022. doi: 10.1103/PhysRevD.105.023517. URL <https://link.aps.org/doi/10.1103/PhysRevD.105.023517>.
- A. J. Mead, S. Brieden, T. Tröster, and C. Heymans. `hmcodes-2020`: improved modelling of non-linear cosmological power spectra with baryonic feedback. *Monthly Notices of the Royal Astronomical Society*, 502(1):1401–1422, Jan. 2021. ISSN 1365-2966. doi: 10.1093/mnras/stab082. URL <http://dx.doi.org/10.1093/mnras/stab082>.
- N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 06 1953. ISSN 0021-9606. doi: 10.1063/1.1699114. URL <https://doi.org/10.1063/1.1699114>.
- T. Mroczkowski, D. Nagai, K. Basu, J. Chluba, J. Sayers, and et al. Astrophysics with the Spatially and Spectrally Resolved Sunyaev-Zeldovich Effects. A Millimetre/Submillimetre Probe of the Warm and Hot Universe. *Space Science Reviews*, 215(1):17, Feb. 2019. doi: 10.1007/s11214-019-0581-2.
- S. Mukherjee, A. Ghosh, M. J. Graham, C. Karathanasis, M. M. Kasliwal, I. M. Hernandez, S. M. Nissanke, A. Silvestri, and B. D. Wandelt. First measurement of the hubble parameter from bright binary black hole gw190521, 2020. URL <https://arxiv.org/abs/2009.14199>.
- S. Naess and et al. The Atacama Cosmology Telescope: Detection of Millimeter-wave Transient Sources. *The Astrophysical Journal*, 915(1):14, June 2021a. ISSN 1538-4357. doi: 10.3847/1538-4357/abfe6d. URL <http://dx.doi.org/10.3847/1538-4357/abfe6d>.

BIBLIOGRAPHY

- S. Naess and et al. The Atacama Cosmology Telescope: A Search for Planet 9. *The Astrophysical Journal*, 923(2):224, Dec. 2021b. ISSN 1538-4357. doi: 10.3847/1538-4357/ac2307. URL <http://dx.doi.org/10.3847/1538-4357/ac2307>.
- T. Namikawa, A. I. Lonappan, and et al. Litebird science goals and forecasts: Improving sensitivity to inflationary gravitational waves with multitracer delensing, 2023. URL <https://arxiv.org/abs/2312.05194>.
- L. Pagano, J.-M. Delouis, S. Mottet, J.-L. Puget, and L. Vibert. Reionization optical depth determination from planck hfi data with ten percent accuracy. *Astronomy & Astrophysics*, 635:A99, Mar. 2020. ISSN 1432-0746. doi: 10.1051/0004-6361/201936630. URL <http://dx.doi.org/10.1051/0004-6361/201936630>.
- A. A. Penzias and R. W. Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *ApJ*, 142:419–421, July 1965. doi: 10.1086/148307.
- C. Pitrou, A. Coc, J.-P. Uzan, and E. Vangioni. Precision big bang nucleosynthesis with improved Helium-4 predictions. *Submitted to Phys. Rept.*, 2018.
- Planck Collaboration. Planck2013 results. xxx. cosmic infrared background measurements and implications for star formation. *Astronomy & Astrophysics*, 571:A30, Oct. 2014. ISSN 1432-0746. doi: 10.1051/0004-6361/201322093. URL <http://dx.doi.org/10.1051/0004-6361/201322093>.
- Planck Collaboration. Planck 2015 results: XI. CMB power spectra, likelihoods, and robustness of parameters. *Astronomy & Astrophysics*, 594:A11, Sept. 2016. ISSN 1432-0746. doi: 10.1051/0004-6361/201526926. URL <http://dx.doi.org/10.1051/0004-6361/201526926>.
- Planck Collaboration. Planck 2018 results. ix. constraints on primordial non-gaussianity, 2019.
- Planck Collaboration. Planck 2018 results. iv. diffuse component separation. *Astronomy & Astrophysics*, 641:A4, 2020a. doi: 10.1051/0004-6361/201833881. URL <https://doi.org/10.1051/0004-6361/201833881>.
- Planck Collaboration. Planck 2018 results: V. cmb power spectra and likelihoods. *Astronomy & Astrophysics*, 641:A5, Sept. 2020b. ISSN 1432-0746. doi: 10.1051/0004-6361/201936386. URL <http://dx.doi.org/10.1051/0004-6361/201936386>.

- Planck Collaboration. Planck 2018 results: I. Overview and the cosmological legacy of Planck. *Astronomy & Astrophysics*, 641:A1, Sept. 2020c. ISSN 1432-0746. doi: 10.1051/0004-6361/201833880. URL <http://dx.doi.org/10.1051/0004-6361/201833880>.
- Planck Collaboration. Planck 2018 results. vi. cosmological parameters. *Astronomy & Astrophysics*, 641:A6, sep 2020d. doi: 10.1051/0004-6361/201833910. URL <https://doi.org/10.1051%2F0004-6361%2F201833910>.
- "Planck Collaboration". Planck intermediate results - LVII. Joint Planck LFI and HFI data processing. *A&A*, 643:A42, 2020. doi: 10.1051/0004-6361/202038073. URL <https://doi.org/10.1051/0004-6361/202038073>.
- M. J. D. Powell. The BOBYQA algorithm for bound constrained optimization without derivatives, 2009.
- F. J. Qu, B. D. Sherwin, M. S. Madhavacheril, D. Han, and K. T. C. et al. The Atacama Cosmology Telescope: A Measurement of the DR6 CMB Lensing Power Spectrum and its Implications for Structure Growth. 2023.
- A. G. Riess, W. Yuan, L. M. Macri, D. Scolnic, D. Brout, S. Casertano, D. O. Jones, Y. Murakami, G. S. Anand, L. Breuval, T. G. Brink, A. V. Filippenko, S. Hoffmann, S. W. Jha, W. D'arcy Kenworthy, J. Mackenty, B. E. Stahl, and W. Zheng. A comprehensive measurement of the local value of the hubble constant with $1\text{kms}^{-1}\text{Mpc}^{-1}$ uncertainty from the hubble space telescope and the sh0es team. *The Astrophysical Journal Letters*, 934(1):L7, July 2022. ISSN 2041-8213. doi: 10.3847/2041-8213/ac5c5b. URL <http://dx.doi.org/10.3847/2041-8213/ac5c5b>.
- S. Shaikh, I. Harrison, A. van Engelen, and G. A. M. et al. Cosmology from cross-correlation of act-dr4 cmb lensing and des-y3 cosmic shear. 2023.
- B. D. Sherwin, J. Dunkley, S. Das, and et al. Evidence for dark energy from the cosmic microwave background alone using the atacama cosmology telescope lensing measurements. *Phys. Rev. Lett.*, 107:021302, Jul 2011. doi: 10.1103/PhysRevLett.107.021302. URL <https://link.aps.org/doi/10.1103/PhysRevLett.107.021302>.
- G. F. Smoot, C. L. Bennett, A. Kogut, E. L. Wright, and et al. Structure in the COBE Differential Microwave Radiometer First-Year Maps. *Astrophysical Journal Letters*, 396:L1, Sept. 1992. doi: 10.1086/186504.

D. N. Spergel, L. Verde, H. V. Peiris, E. Komatsu, M. R.olta, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. First-year wilkinson microwave anisotropy probe (wmap)* observations: Determination of cosmological parameters. *The Astrophysical Journal Supplement Series*, 148(1):175, sep 2003. doi: 10.1086/377226. URL <https://dx.doi.org/10.1086/377226>.

A. Spurio Mancini, D. Piras, J. Alsing, B. Joachimi, and M. P. Hobson. *CosmoPower*: emulating cosmological power spectra for accelerated bayesian inference from next-generation surveys. *Monthly Notices of the Royal Astronomical Society*, 511(2):1771–1788, jan 2022. doi: 10.1093/mnras/stac064. URL <https://doi.org/10.1093%2Fmnras%2Fstac064>.

The LSST Dark Energy Science Collaboration, R. Mandelbaum, T. Eifler, R. Hložek, T. Collett, E. Gawiser, D. Scolnic, and et al. The LSST Dark Energy Science Collaboration (DESC) Science Requirements Document, 2021. URL <https://arxiv.org/abs/1809.01669>.

The Simons Observatory collaboration. The simons observatory: science goals and forecasts. *Journal of Cosmology and Astroparticle Physics*, 2019(02):056, feb 2019. doi: 10.1088/1475-7516/2019/02/056. URL <https://dx.doi.org/10.1088/1475-7516/2019/02/056>.

R. J. Thornton and et al. The Atacama Cosmology Telescope: The Polarization-Sensitive ACTPol Instrument. *The Astrophysical Journal Supplement Series*, 227(2):21, Dec. 2016. ISSN 1538-4365. doi: 10.3847/1538-4365/227/2/21. URL <http://dx.doi.org/10.3847/1538-4365/227/2/21>.

J. Torrado and A. Lewis. Cobaya: Bayesian analysis in cosmology. <https://ascl.net/1910.019>, Oct. 2019.

J. Torrado and A. Lewis. Cobaya: code for bayesian analysis of hierarchical physical models. *Journal of Cosmology and Astroparticle Physics*, 2021(05):057, May 2021. ISSN 1475-7516. doi: 10.1088/1475-7516/2021/05/057. URL <http://dx.doi.org/10.1088/1475-7516/2021/05/057>.

S. Zaroubi. *The Epoch of Reionization*, page 45–101. Springer Berlin Heidelberg, Sept. 2012. ISBN 9783642323621. doi: 10.1007/978-3-642-32362-1_2. URL http://dx.doi.org/10.1007/978-3-642-32362-1_2.

J. Zuntz, M. Paterno, E. Jennings, D. Rudd, A. Manzotti, S. Dodelson, S. Bridle, S. Sehrish, and J. Kowalkowski. Cosmosis: Modular cosmological parameter estimation. *Astronomy and Computing*, 12:45–59, Sept. 2015. ISSN 2213-1337. doi: 10.1016/j.ascom.2015.05.005. URL <http://dx.doi.org/10.1016/j.ascom.2015.05.005>.