The early Universe from CMB B-modes: observational challenges

Susanna Azzoni

St Cross College, University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity Term 2023

Abstract

The extreme conditions of density and temperature make the Universe around the initial moment of the Big-Bang, an extraordinary laboratory for probing poorly understood physics, mixing quantum and general relativity phenomena. The physics of the early Universe can be studied through the Cosmic Microwave Background (CMB), the oldest observable "light" in the Universe. The CMB temperature has been mapped recently with unprecedented precision by the Planck satellite. However, its full potential in testing fundamental physics lies in measurements of its polarization at high resolution and sensitivity. This will be unlocked for the first time with current and future ground-based and space-based observatories. The main goal is to discover the imprint of inflation, when the Universe expanded by an incredibly large factor during a fraction of a second around the time of the Big Bang. Inflation models predict the production of stochastic gravitational waves which distorted spacetime while propagating through the Universe, and generated an observable parity-violating signature in the polarization of the CMB, the so-called B-mode polarization. This nano-Kelvin signal is quantified by the tensor-to-scalar ratio parameter r. Current limits are given by the BICEP/Keck ground experiments, and the goal of the community is to achieve an accuracy at the level of 10^{-3} within a few years. The potential of primordial B-modes to help identify what prompted inflation, whilst also investigating physics at grand unification energy scales, makes the search for this faint CMB polarization signal one of the most compelling goals of modern cosmology.

However, with the capabilities of next-generation experiments, the large leap in detector count and experiment complexity introduce challenges to current methods. Some of these challenges are intrinsically computational: to explore the enormous datasets, we must develop sufficiently reliable and efficient data analysis pipelines to extract features from the observations. Other challenges are algorithmic, including the selection of optimal ways to disentangle the CMB signal of interest from foregrounds and to mitigate systematic errors. This in turn imposes stringent requirements on the instrument design and calibration. Tackling these challenges is the focus of my research.

This thesis presents several methods for foreground cleaning and mitigation of instrument systematics, that allow for unbiased measurements of the cosmic microwave background (CMB) B-mode polarisation power spectrum, that can be used for cosmological parameter estimation. These methods have been devised for the Simons Observatory (SO) and LiteBIRD satellite and are adaptable to a variety of CMB experiments. I present two new methods for component separation in the presence of spatial variations in Galactic foreground spectral properties and their impact on parameter inference: a moment expansion method and a hybrid method, which is an extension of the former. Both these methods account for spatial variability while recovering unbiased constraints on r in the presence of complex foregrounds and noise. Through the development of these methods, I also contributed significantly to the SO cross-spectrum analysis pipeline, which became SO's main Small Aperture Telescope forecast likelihood pipeline. This work led to the discovery that the moments-based method is currently the most promising avenue for SO. I also present a study on the precision of angular position measurement of the half-wave plate polarisation modulation unit for LiteBIRD, and developed a simulation pipeline to mitigate the propagation of HWP systematic effects to final constraints. This study is part of an instrument systematic mitigation campaign that is crucial for the success of the next-generation CMB experiments.

The early Universe from CMB B-modes: observational challenges

Susanna Azzoni

St Cross College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity Term 2023

Acknowledgements

The guidance and constant support of my supervisors David Alonso and Tomotake Matsumura have been invaluable to my personal and academic growth, and I owe them a debt of gratitude. Through my collaboration with them, I have gained a profound understanding of the qualities that distinguish exceptional scientists.

Additionally, I am deeply grateful for the unwavering support of my family and friends, who have been a consistent source of strength throughout my DPhil. I dedicate this work to my parents, who have always demonstrated their support for me and have never stopped believing in me, even during times when I doubted myself. Luca deserves a special recognition for making the past few years even more fulfilling than I anticipated. I cannot imagine being as happy or productive without the constant encouragement that Luca provides me with every day. I would also like to express my gratitude to Rossana and Jack, for setting an example of compassion and dedication throughout my life. To my lifelong friends, Corinne, Valentina, Roberta, Andrea, Olimpia, Niccolò, and Giulia, who have always supported me from far away during my journey.

I must also mention Jakob, who played a crucial role in my journey to becoming a PhD student at Oxford. Without him, I may not have taken the leap to pursue this opportunity. My gratitude extends to the numerous people in the BIPAC and the friends who have shared this journey with me over the past three years. While there are too many of you to mention by name, your support and presence have made the past years more meaningful.

I would also like to thank the numerous individuals who have contributed to my research, including Max Abitbol, Guillaume Patanchon, Nicoletta Krachmalnicoff, Josquin Errard, Kevin Wolz, and Carlo Baccigaluppi. Moreover, I am thankful to the Simons Observatory and LiteBIRD collaborations, and I extend my appreciation to Erminia Calabrese, Hirokazu Ishino, Jens Chluba for their invaluable input and discussions. I also express my gratitude to Jonathan Patterson for smoothly running the Glamdring Cluster hosted by the University of Oxford, where a lot of the data processing for my research was performed.

I sincerely thank the examiners of my DPhil viva, Pedro Ferreira and Carlo Contaldi, for their patience in reading through this thesis, valuable feedback and rigorous evaluation of my research.

Last but not least, I must acknowledge COVID-19 for adding a new dimension to my academic journey. Without its constant presence during the last 3 years, I may have taken for granted the joy of in-person interactions and the beauty of remote research.

Declaration

I declare that no part of this thesis has been accepted, or is currently being submitted, for any degree or diploma or certificate or any other qualification in this University or elsewhere.

Chapter 3 is based on work presented in Azzoni, S., et al., published May 2021 in the Journal of Cosmology and Astroparticle physics (JCAP): "A minimal power-spectrum-based moment expansion for CMB B-mode searches" [1]. I led the project, including the development of the method and produced all the results for this Chapter. This work led to a collaboration paper that I co-authored that can be found in Ref. [2].

Chapter 4 is based on work presented in Wolz, K., Azzoni, S., et al., which was submitted to the Astronomy and Astrophysics Journal (A&A) and is currently under review: "The Simons Observatory: pipeline comparison and validation for large-scale B-modes" [3]. I led the development and testing of the moment expansion for pipeline A and developed code for the cross-pipeline comparison.

Chapter 5 is based on work presented in Azzoni, S., et al., published March 2023 in JCAP: "A hybrid map-C_ℓ component separation method for primordial CMB B-mode searches" [4]. I led the project, including the development of the method and produced all the results for this Chapter.

Chapter 6 is based on work presented in Azzoni, S et al., which is currently under review: "Half-Wave-Plate rotational instabilities propagated through an end-to-end pipeline for LiteBIRD". I led the project, including the development of the method and produced all the results for this Chapter. This work led to collaboration papers that I co-authored that can be found in Refs. [5–9].

Susanna Azzoni (May 2023)

Abstract

The extreme conditions of density and temperature make the Universe around the initial moment of the Big-Bang, an extraordinary laboratory for probing poorly understood physics, mixing quantum and general relativity phenomena. The physics of the early Universe can be studied through the Cosmic Microwave Background (CMB), the oldest observable "light" in the Universe. The CMB temperature has been mapped recently with unprecedented precision by the Planck satellite. However, its full potential in testing fundamental physics lies in measurements of its polarization at high resolution and sensitivity. This will be unlocked for the first time with current and future ground-based and space-based observatories. The main goal is to discover the imprint of inflation, when the Universe expanded by an incredibly large factor during a fraction of a second around the time of the Big Bang. Inflation models predict the production of stochastic gravitational waves which distorted spacetime while propagating through the Universe, and generated an observable parity-violating signature in the polarization of the CMB, the so-called B-mode polarization. This nano-Kelvin signal is quantified by the tensor-to-scalar ratio parameter r. Current limits are given by the BICEP/Keck ground experiments, and the goal of the community is to achieve an accuracy at the level of 10^{-3} within a few years. The potential of primordial B-modes to help identify what prompted inflation, whilst also investigating physics at grand unification energy scales, makes the search for this faint CMB polarization signal one of the most compelling goals of modern cosmology.

However, with the capabilities of next-generation experiments, the large leap in detector count and experiment complexity introduce challenges to current methods. Some of these challenges are intrinsically computational: to explore the enormous datasets, we must develop sufficiently reliable and efficient data analysis pipelines to extract features from the observations. Other challenges are algorithmic, including the selection of optimal ways to disentangle the CMB signal of interest from foregrounds and to mitigate systematic errors. This in turn imposes stringent requirements on the instrument design and calibration. Tackling these challenges is the focus of my research.

This thesis presents several methods for foreground cleaning and mitigation of instrument systematics, that allow for unbiased measurements of the cosmic microwave background (CMB) B-mode polarisation power spectrum, that can be used for cosmological parameter estimation. These methods have been devised for the Simons Observatory (SO) and LiteBIRD satellite and are adaptable to a variety of CMB experiments. I present two new methods for component separation in the presence of spatial variations in Galactic foreground spectral properties and their impact on parameter inference: a moment expansion method and a hybrid method, which is an extension of the former. Both these methods account for spatial variability while recovering unbiased constraints on r in the presence of complex foregrounds and noise. Through the development of these methods, I also contributed significantly to the SO cross-spectrum analysis pipeline, which became SO's main Small Aperture Telescope forecast likelihood pipeline. This work led to the discovery that the moments-based method is currently the most promising avenue for SO. I also present a study on the precision of angular position measurement of the half-wave plate polarisation modulation unit for LiteBIRD, and developed a simulation pipeline to mitigate the propagation of HWP systematic effects to final constraints. This study is part of an instrument systematic mitigation campaign that is crucial for the success of the next-generation CMB experiments.

Contents

List of Figures

List of Abbreviations

CMB . . . . . Cosmic Microwave Background.

EM tensor . . Energy-Momentum tensor.

SVT . . . . . . Scalar-Vector-Tensor.

GR . . . . General Relativity.

GUT . . . . . . Grand Unified Theory.

K-G . . . . . . Klein-Gordon.

QFT . . . . . . Quantum field theory.

BAO . . . . . . Baryon Acoustic Oscillations.

PGW . . . . . Primordial gravitational wave.

TOD . . . . . . Time-ordered-data.

ISM . . . . . . Interstellar medium.

GMF . . . . . Galactic magnetic field.

LSS . . . . . . Large-scale structure.

SZ . . . . Sunyaev-Zel'dovich.

tSZ . . . . Thermal SZ effect.

kSZ . . . . . . Kinematic SZ effect.

ISW . . . . . . Integrated Sachs-Wolfe.

RMS . . . . . . Root-mean-square.

PSD . . . . . . Power spectral density.

SED . . . . . . Spectral Energy Distribution.

Life can only be understood backwards;
but it must be lived forwards.

— Soren Kierkegaard

Chapter 1

Cosmology from the CMB

Contents

The Cosmic Microwave Background (CMB) is a remnant radiation from the Big Bang, and has been redshifted over time to become a microwave background radiation. The CMB offers a direct way to validate models that explain how everything we see in the Universe originates from quantum mechanics. By examining the subtle correlations in the background variations (*anisotropies*), which arise from the interplay of gravitational and quantum physics at high energies, we can gain insights into the unification of these fundamental forces. Additionally, the correlations that arise at later times in the CMB encode valuable information about the distribution of both ordinary and dark matter in the Universe, as well as the properties of neutrinos.

In the late 1930s, Georges Lemaître proposed the concept of the CMB as a remnant of the Big Bang. Despite initial skepticism from the scientific community, George Gamow and Ralph Alpher's groundbreaking research in 1948 supported the Big Bang theory by demonstrating the agreement between observed element abundances and their model [10]. That same year, Alpher and Robert Herman made the first quantitative prediction of background radiation, estimating its temperature to be around 5 K. In the early 1960s, Robert Dicke, Jim Peebles, and Yakov Zel'dovich independently made similar predictions. In 1978 the Nobel Prize was awarded to Arno Penzias and Robert Wilson for the discovery of the CMB. While working on a radiometer for satellite communication experiments at Bell Laboratories in 1965, they unexpectedly observed an excess temperature of 4.2 K in their antenna readings, which was attributed to the echo of the ancient Universe. Since its discovery, the CMB has provided significant insights into the fundamental aspects of space and time, bridging the gap between our understanding of physics at the smallest scales and highest energies to the largest scales in the Universe. CMB studies have achieved remarkable success by combining theory, phenomenology, and advanced experiments. The isotropy of the CMB, observed with incredible precision of nearly one part in a hundred thousand, played a crucial role in the development of inflation and cold dark matter theories in the 1980s. In 1992, NASA's COBE satellite, equipped with the DMR (Differential Microwave Radiometer) [11] and FIRAS (Far Infrared Absolute Spectrophotometer) [12] instruments, measured the CMB's black-body spectrum and confirmed the presence of anisotropies, laying the foundation for subsequent higher-resolution measurements. The significant discoveries made by COBE were recognized with the Nobel Prize in 2006.

In the subsequent decade, various ground and balloon-based instruments were employed to measure the temperature variations in the CMB at different angular scales, ranging from degrees to arcminutes. These measurements have provided us with a wealth of information about the early stages of the Universe. Starting in 1998, several experiments such as the Microwave Anisotropy Telescope (MAT or TOCO) [13], BOOMERanG [14], MAXIMA [15], and Archeops [16], played a crucial role in revealing acoustic peaks in the angular power spectrum of the CMB. These findings supported the predictions of inflation and provided evidence for an accelerating Universe, as observed through type Ia supernovae (SNe). The discovery of secondary peaks in the CMB angular power spectrum was made by experiments like DASI [17], BOOMERanG [14], and VSA (Very Small Array) [18] in 2001 and 2002. In 2002, the Cosmic Background Imager (CBI) experiment published the CMB power spectrum up to a multipole moment of $\ell = 3500$, allowing for the first measurement of the CMB damping tail [19]. Similar-scale measurements were also conducted by ACBAR [20]. In 2001, NASA launched the Wilkinson Microwave Anisotropy Probe (WMAP) satellite, which observed the sky for nine years [21]. WMAP played a pivotal role in ushering in a new era of precision cosmology. It provided highly accurate measurements of key cosmological parameters that characterise the composition, geometry, size, and age of our Universe. The most recent comprehensive data on CMB anisotropies comes from the Planck satellite [22], launched by the European Space Agency (ESA) in 2009. Planck aimed to make definitive measurements of the primary temperature anisotropies in the CMB, enabling precise tests of the standard cosmological model. These early measurements from Planck also confirmed the universal baryon density, consistent with calculations from big bang nucleosynthesis, and provided evidence for the non-baryonic nature of dark matter.

In recent years, scientists have shifted their focus towards studying the polarisation of the CMB, which can be divided into two modes: E-mode and B-mode. The dominant mode, E-mode polarisation, is primarily caused by density fluctuations in the early Universe. The first detection of E-mode polarization was made by the DASI experiment in 2002 [23], followed by subsequent observations by CAPMAP [24], CBI, WMAP, and Planck. By studying E-mode polarisation, researchers have gained insights into the matter distribution in the Universe and the effects of gravitational lensing on the CMB. On the other hand, the subtle B-mode polarisation in the CMB arises from gravitational waves generated during the inflationary epoch, a period of rapid expansion in the early Universe. While the study of the CMB has already revolutionised cosmology by providing significant insights into the early Universe, the detection of primordial B-mode polarisation represents the next frontier, as it would provide valuable information about the origins of the Universe. Specifically, detecting CMB B-modes would offer evidence for inflation and allow us to constrain the energy scale of inflation, providing insights into high-energy physics and quantum gravity. Such a discovery would be a remarkable advancement in cosmology, offering a deeper understanding of the origin and evolution of our Universe. The study of primordial B-mode polarisation in the CMB is currently an active area of research, with various experiments underway to detect these subtle polarisation patterns. These experiments include BICEP[1]

[1] bicepkeck.org

(BICEP1, BICEP2, KECK Array, BICEP3, BICEP Array), ACT[2] (ACT, ACTpol, AdvACT), Polarbear[3] (Polarbear-1, Polarbear-2a, Simons Array), and SPT[4] (SPT, SPTpol, SPT-3G), among others. Additionally, forthcoming experiments such as the Simons Observatory, CMB-S4, and LiteBIRD are planned to enhance the sensitivity and resolution of CMB measurements.

The aim of this chapter is to provide a comprehensive overview of our current understanding of the Universe, focusing on the CMB. Section 1.1 provides an overview of the "smooth" Universe, establishing the foundational concepts required to comprehend its evolution and to understand the subsequent sections. Section 1.2 delves into the Universe's perturbations, which are crucial for modelling the Universe's evolution and understanding the anisotropies that are observed today. Section 1.3 investigates the Hot Big Bang model, which is the most widely accepted cosmological model to date, whose essential components in the progression from the early Universe to the current era are reviewed. The problems of this model are discussed in Section 1.4, leading to the theoretical foundations of inflation theories. Finally, Section 1.5 provides an in-depth examination of the CMB, discussing its properties, and the relation with primordial gravitational waves (PGWs). The physical significance of the discovery of PWGs and how they can provide critical insights into the Universe's early stages is discussed in the final Section 1.6.

1.1 The smooth Universe

Throughout centuries of study, it has become clear that our position in the Universe is not special, in that we are not at its centre. This has led to certain assumptions about the nature of our universe. Specifically, if we are insignificant and the Universe is also insignificant everywhere, then at any given time, the Universe should look the same everywhere. This uniformity is known as a *homogeneous* space-time.

Building on this concept, homogeneity implies that the Universe looks the same in all directions from any given point in space, creating an *isotropic* space-time[5]. When the Universe exhibits both homogeneity and isotropy, it satisfies the *Cosmological Principle*.

As it stands, our Universe on large enough scales is believed to conform to this principle. In other words, the properties of the Universe at large scales do

[2] act.princeton.edu

[3] bolo.berkeley.edu/polarbear/

[4] pole.uchicago.edu/

[5] Homogeneity and isotropy are related but distinct concepts. If a manifold is isotropic around *any* one point, then it must also be homogeneous. However, if a Universe is only isotropic around one point, it is not necessarily homogeneous (and vice versa).

not depend on the position of the observer or the direction it looks[6]. These assumptions have been supported by a range of observational evidence, including the Cosmic Microwave Background (CMB) radiation and the large scale structure of the Universe. In particular, the CMB is a nearly uniform background radiation that permeates the entire Universe. It has been observed to be almost perfectly isotropic, with small fluctuations in temperature that are consistent with the predictions of the standard cosmological model. Also, the distribution of galaxies that we observe on the largest scales, far away from us, appears to be randomly dispersed with a uniform density in all directions. The observed homogeneity and isotropy of the Universe on large scales allow us to infer what the early Universe must have looked like and how it evolved over time to form the structures we see today. This is described in this Section, which draws upon several references [25–30].

1.1.1 Cosmological geometric structure

The goal is now to construct a metric which would describe such a smooth Universe. Obviously, the universe is neither homogenous nor isotropic over time (i.e. it looks different in the past and in the future). It is possible however to "slice" spacetime into constant-time spaces[7], or *hypersurfaces*, $d\Sigma_3$ [29], for which the geometry can be described.

1.1.1.1 The FLRW metric

The symmetries imposed by the Cosmological Principle define implicitly a specific hypersurface that is homogeneous and isotropic, or maximally symmetric [8]. This can be described using the Friedmann-Lemaître-Robertson-Walker (*FLRW*) metric, which is the most general spacetime model that satisfies the cosmological principle.

The distance between two neighbouring points (or line element) of the FLRW spacetime metric is given by [28]:

$$d\tau^2 \equiv -g_{\mu\nu}(x)dx^\mu dx^\nu = dt^2 - a^2(t)(d\Sigma_3)^2, \tag{1.1.1.1}$$

where $g_{\mu\nu}(x)$ is the metric tensor, describing the geometry of spacetime, and x^μ are the coordinates of a point in spacetime (or event). The expression above describes the spacetime in the *comoving coordinates*[9], which are defined in such a way that

[6]Isotropy implies rotationally symmetry, and homogeneity implies symmetry under translations.

[7]The slicing process is also known as *foliation* [31].

[8]It is worth noting that a different slicing of the same spacetime will not necessarily return hypersurfaces that are homogeneous and isotropic.

[9]The comoving coordinates relate to the physical coordinates as $x^i_{\rm phys} = a(t)x^i$

the spatial coordinates are the labels of freely falling[10] fluid elements that fill the Universe and the time coordinate t is the time measured by the imaginary clocks carried by each fluid element.

The right-hand side of Equation 1.1.1.1 consists of two terms: the time interval term dt^2, and the spatial term, including the *scale factor* $a(t)$, which describes the overall scale of the Universe as it evolves over time. Since the Universe is homogeneous and isotropic, the scale factor does not depend on location or direction, and has no physical unit. The spatial part of the line element represents the metric of the 3D spatial hypersurface at a given time, and is given by

$$(d\Sigma_3)^2 = \left[d\boldsymbol{x}^2 + \kappa\frac{(\boldsymbol{x}\cdot d\boldsymbol{x})^2}{1-\kappa\boldsymbol{x}^2}\right] \equiv \gamma_{ij}\ \mathrm{d}x^i\ \mathrm{d}x^j \tag{1.1.1.2}$$

i.e. a sum of the square of the spatial interval ($d\boldsymbol{x}^2$) and a correction term for the spatial curvature, with κ being the curvature parameter of the 3D spatial slices. The space-like metric tensor is defined as

$$\gamma_{ij} \equiv \delta_{ij} + \kappa\frac{x_i x_j}{1-\kappa\left(x_k x^k\right)}, \tag{1.1.1.3}$$

where $x^i, (i = 1, 2, 3)$ are the spatial comoving coordinates, and the Kronecker δ_{ij} is a tensor with 1's on the diagonal and 0's elsewhere.

The 3D metric γ_{ij}, being maximallly symmetric, has six Killing vectors, i.e. vector fields that preserve the geometric properties of the spacetime (homogeneity and isotropy). This has two qualitative implications:

1. The spatial curvature[11] must be the same at every point in space, and hence only a function of time, $R^{(3)}(t) \equiv 6\kappa/a(t)^2$.

2. The curvature κ must be constant everywhere, allowing for a negatively curved (hyperbolic, open) Universe ($\kappa < 0$), a positively curved (spherical, closed) Universe ($\kappa > 0$) or a flat (Euclidean) Universe ($\kappa = 0$)[12].

To ensure that the spatial interval between two points is positive everywhere, including at $\boldsymbol{x} = 0$, it is necessary to have $\dot{a} > 0$, i.e. we have an *expanding Universe*.

The space-like metric varies depending on the selection of coordinates used. Spherical polar coordinates[13] (r, θ, ϕ) are advantageous because they clearly exhibit the symmetries of the space, such that the FLRW metric is

$$d\tau^2 = dt^2 - a^2(t)\left[\frac{dr^2}{1-\kappa r^2} + r^2\left(d\theta^2 + \sin^2\theta d\phi^2\right)\right] \tag{1.1.1.4}$$

[10]This is why the dt^2 component have a value of -1 in the comoving coordinates.
[11]This can be defined here as the *Ricci scalar*, calculated from the spatial part of the metric
[12]If $\kappa = 0$ and $a(t) = 1$ we have the *Minkowski metric*.
[13]Using the transformation $d\boldsymbol{x}^2 = dr^2 + r^2 d\Omega, \quad d\Omega \equiv d\theta^2 + \sin^2\theta d\phi^2$.

and the metric γ_{ij} becomes diagonal, with elements

$$\gamma_{rr} = \frac{1}{1-\kappa r^2}, \quad \gamma_{\theta\theta} = r^2, \quad \gamma_{\phi\phi} = r^2 \sin^2\theta, \quad \gamma_{00} = -1. \tag{1.1.1.5}$$

To simplify γ_{rr}, it is useful to redefine the radial coordinate $\mathrm{d}\chi \equiv \mathrm{d}r/\sqrt{1-\kappa r^2}$, such that Eq. 1.1.1.4 becomes

$$d\tau^2 = dt^2 - a^2(t)\left[d\chi^2 + \left\{\begin{array}{c}\sin^2\chi\\ \chi^2\\ \sinh^2\chi\end{array}\right\}\left(d\theta^2+\sin^2\theta d\phi^2\right)\right], \quad \left\{\begin{array}{c}\kappa=+1\\ \kappa=0\\ \kappa=-1\end{array}\right\} \tag{1.1.1.6}$$

where we have rescaled $a(t)$ such that $\kappa = \pm 1$ or 0. For simplicity, we will assume that the Universe is flat from this point forward, i.e. $\kappa = 0$, as most observational evidence supports this scenario. It will sometimes be useful to express the FLRW metric in *conformal time* η, which relates to *physical* or *cosmological time*, given by $dt = a d\eta$. The four-dimensional metric in Eq. 1.1.1.6 becomes

$$d\tau^2 = -g_{\mu\nu}(x)dx^\mu dx^\nu = a^2(\eta)\left(d\eta^2 - |d\boldsymbol{x}|^2\right) \tag{1.1.1.7}$$

where we have used the shorthand $|d\boldsymbol{x}|^2 = \left[d\chi^2 + \chi^2(d\theta^2 + \sin^2\theta\, d\varphi^2)\right]$.

1.1.2 Cosmological evolution

1.1.2.1 Matter and Energy

The FLRW metric contains the free function $a(t)$, that can be determined from the contents of the Universe using General Relativity. To this end, we want to find the FLRW metric as a solution of *Einstein's field equations*:

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = 8\pi G T_{\mu\nu} + \Lambda g_{\mu\nu}, \tag{1.1.2.1}$$

where $G_{\mu\nu}$ is the Einstein tensor, $g_{\mu\nu}$ is the metric tensor as defined in Eq. 1.1.1.7 describing the geometry of spacetime, $R_{\mu\nu}$ is the Ricci tensor and $R \equiv g_{\mu\nu}R^{\mu\nu}$ is the Ricci scalar, both related to the curvature of spacetime (quantified by the Riemann curvature tensor $R^\mu_{\nu\alpha\beta}$, where $R_{\mu\nu} \equiv R^\alpha_{\mu\alpha\nu}$), G stands for Newton's Constant of Gravitation, $T_{\mu\nu} = g^{\mu\lambda}T_{\lambda\nu}$ is the stress-energy or *energy-momentum (EM) tensor*, which describes the distribution of matter and energy within spacetime, Λ is the cosmological constant.

Einstein's equations state that the curvature of spacetime ("geometric" first equality) is directly related to the distribution of matter and energy within it ("matter" second equality). In other words, the presence of matter and energy in a region of spacetime will cause that region to curve, and the degree of curvature will depend on the amount and distribution of matter and energy.

A usual choice is that of a perfect fluid to describe the components that comprise the energy content of the Universe, for which the EM tensor can be written in terms of the energy distribution (ρ) and momentum density or pressure (p) as follows:

$$T_{\mu\nu} = (\rho + p)u_\mu u_\nu - p g_{\mu\nu}, \tag{1.1.2.2}$$

where u_μ is the fluid's 4-velocity of the matter, and the Einstein summation convention is used. It follows that, in the absence of matter, the EM tensor is zero. In the presence of matter, the EM tensor determines the curvature of spacetime, which in turn affects the motion of matter and the propagation of light.

We take the fluid to be made of multiple components i, such that $T_{\mu\nu} = \sum_i T^i_{\mu\nu}$, and for each component we relate the pressure and energy density using the *equation of state*. For most cosmological fluids, e.g. ideal "pressureless" matter (baryons, cold dark matter, or any other non-relativistic particles), radiation (photons, massless neutrinos, or any other massless particles) and vacuum energy (Λ), takes the simple form:

$$p_i = w_i \rho_i\,, \quad \left\{ \begin{array}{cc} w = 0 & \text{matter} \\ w = 1/3 & \text{radiation} \\ w = -1 & \text{vacuum energy or } \Lambda \end{array} \right\} \tag{1.1.2.3}$$

where w_i represents the equation of state parameter for the i-th component, and in the latter case we have considered Λ to be a component of the EM tensor, $T^{(\Lambda)}_{\mu\nu} = \Lambda g_{\mu\nu}/(8\pi G)$[14]. In comoving coordinates, $u_\mu \propto (1, 0, 0, 0)$, such that in the inertial frame moving with the matter the EM tensor is

$$T^\mu_\nu = \mathrm{diag}(\rho, -p, -p, -p), \tag{1.1.2.4}$$

where each element only depends on time, and all non-diagonal components have to vanish by the requirement of isotropy and homogeneity[15].

[14] In the case of Λ, a departure from $w_\Lambda = -1$ may indicate that dark energy is more intricate than a basic cosmological constant. Due to this reason, this parameter has gained significant importance in the quest for dark energy.

[15] It is worth noting that in general, the EM tensor can have off-diagonal components, corresponding to anisotropic stresses. Additionally, T^0_i can measure the heat flux, in the fluid context proportional to the fluid velocity.

1.1.2.2 The Friedmann Equations

Replacing the perfect-fluid $T_{\mu\nu}$ reduces Einstein's equations 1.1.2.1 to two independent differential equations for $a(\eta)$, known as the *Friedmann equations*:

$$\left(\frac{\dot{a}}{a}\right)^2 + \kappa = \mathcal{H}^2 + \kappa = \frac{8\pi G}{3} a^2 \rho + \frac{1}{3}\Lambda a^2 \tag{1.1.2.5}$$

$$\left(\frac{\dot{a}}{a}\right)^{\cdot} = \dot{\mathcal{H}} = -\frac{4\pi G}{3} a^2(\rho + 3p) + \frac{1}{3}\Lambda a^2 = \left(\frac{\ddot{a}}{a}\right) - \left(\frac{\dot{a}}{a}\right)^2 \tag{1.1.2.6}$$

where $\dot{} \equiv d/d\eta$, $\rho = -T^0_0, p = T^i_i$, and we define the Hubble parameter in conformal time as $\mathcal{H} = \dot{a}/a$. It can be inferred from the second Friedmann equation that any species having an equation of state with $w < -1/3$ would result in a *positive acceleration*. The *Hubble parameter* describes the expansion rate of the Universe and is defined by:

$$H = \frac{da/dt}{a} = \frac{\dot{a}}{a^2} = \mathcal{H}/a\,, \tag{1.1.2.7}$$

whose quantity is 1/(time). From the above definition of H, one can calculate the age of the Universe, which gives $H(t)dt = da/a$. If we know H as a function of a instead of t, we obtain $t = \int da/aH(a)$.

A consequence of the Friedmann equations, due to the contracted Bianchi identity, is *energy-momentum conservation* ($\nabla_\nu T^{\mu\nu} = 0$)[16], which implies the *continuity equation*:

$$\dot{\rho} = -3\left(\frac{\dot{a}}{a}\right)(\rho + p) = -3(1+w)\mathcal{H} \tag{1.1.2.8}$$

This can be integrated to obtain the time-dependent evolution of different components

$$\rho(a) = \rho_0\, a^{-3(1+w)} \quad \propto \left\{ \begin{array}{cc} a^{-3} & \text{matter} \\ a^{-4} & \text{radiation} \\ a^0 & \text{vacuum energy} \end{array} \right\} \tag{1.1.2.9}$$

Using the solution in Eq. 1.1.2.9, we obtain a compact form for the first Friedmann equation:

$$H^2 = H_0^2 \sum_i \Omega_i\, a^{-3(1+w_i)}, \tag{1.1.2.10}$$

in terms of dimensionless cosmological density parameters $\Omega_{i,0}$, given by the ratio of a component density to the critical density required for a flat Universe, defined from Eq. 1.1.2.5 as $\rho_c = 3H/8\pi G$, evaluated at present time η_0. We define

$$\Omega_{\rho_i} = \left(\frac{\rho_i}{\rho_c}\right)_{\eta=\eta_0}, \quad \Omega_\Lambda = \left.\frac{\Lambda}{3H^2}\right|_{\eta=\eta_0}, \quad \Omega_\kappa = \left.\frac{-\kappa}{H^2}\right|_{\eta=\eta_0}; \tag{1.1.2.11}$$

[16] The energy-momentum conservation is obtained by taking the covariant derivative of the Friedmann equations and using the contracted Bianchi identity.

where the first equation correspond to the general expression for the density of matter ρ_m (both dark and baryonic) and radiation ρ_r, while the other two correspond to dark energy (Λ) and curvature (κ), respectively. These components contribute to the *total energy density* of the Universe, and $\Omega = \Omega_m + \Omega_r + \Omega_\Lambda = 1 - \Omega_\kappa$. Substituting these into Eq. 1.1.2.10, we can rewrite the Friedmann equation in a compact form:

$$H^2 = H_0^2 \left[\Omega_{r,0} a^{-4} + \Omega_{m,0} a^{-3} + \Omega_{\kappa,0} a^{-2} + \Omega_\Lambda\right] \tag{1.1.2.12}$$

The scaling factors of radiation (a^{-4}), matter (a^{-3}), and vacuum energy (a^0) have two significant qualitative implications. Firstly, during earlier times, the impact of curvature is essentially negligible as its contribution increases more slowly compared to matter and radiation. Secondly, for the majority of its existence, the Universe was ruled by a single component, starting with radiation, followed by matter, and vacuum energy.

From the Friedmann equations in 1.1.2.5-1.1.2.6, or by solving for a in 1.1.2.12, where $\kappa = \Lambda = 0$ and $w =$ constant, one concludes that the scale factor behaves like a power law, given by $a \propto \eta^{\frac{2}{1+3w}} \propto t^{\frac{2}{3(1+w)}}$ for $w \neq -1$ and $a \propto -\eta^{-1} \propto e^{Ht}$ for $w = -1$. The time dependence of the scale factor for different components is then:

$$a \propto \left\{ \begin{array}{ccc} \eta^2 & \propto t^{2/3} & \text{for } w = 0, \text{ matter} \\ \eta & \propto t^{1/2} & \text{for } w = 1/3, \text{ radiation} \\ \eta^{-1} & \propto e^{-Ht} & \text{for } w = -1, \text{ dark energy} \end{array} \right\}. \tag{1.1.2.13}$$

Consequently, assuming that a increases with time, in the early stages ($a << 1$), the Universe experienced a radiation-dominated phase where $a \propto t^{1/2}$, followed by a deceleration of expansion during matter domination where $a \propto t^{2/3}$, and it is currently accelerating towards exponential expansion in the Λ-dominated phase[17].

Based on current observations from CMB, supernovae, and galaxy surveys, the estimated values for various cosmological parameters are summarized below:

- $H_0 = 67.66 \pm 0.42$ km/s/Mpc — *Hubble constant*[18];
- $\Omega_m = 0.3147 \pm 0.0074$ — *matter density*, of which:
  - $\Omega_b = 0.02233 \pm 0.00015$ — *baryon*
  - $\Omega_{DM} = 0.2607 \pm 0.0045$ — *dark matter*;
- $\Omega_\Lambda = 0.6844 \pm 0.0084$ — *dark energy density*;
- $\Omega_\kappa = -0.0057 \pm 0.0035$ — *curvature*;
- $\Omega_r = 0.0001361 \pm 0.0000033$ — *radiation density*.

[17] It is worth noting that the powers of t in both radiation- and matter-domination are shallower than t, reflecting the decelerating nature of the expansion.

[18] The reported value comes from CMB measurements by the Planck Collaboration [32]. This value presents, however, a discrepancy with alternative measurement techniques using local objects (the so-called *Hubble tension*). Findings from observations of cepheids, using the Hubble Space Telescope Key Project, give a higher value of $H_0 = 73.24 \pm 1.74$ km/s/Mpc [33]. Similarly, the observations of type Ia supernovae indicate $H_0 = 68.1 \pm 1.28$ km/s/Mpc [34], and the study of time-delays in quasars suggests $H_0 = 73.3^{+1.7}_{-1.8}$ km/s/Mpc.

1.2 The perturbed Universe

Although the Universe is generally thought to be homogeneous and isotropic, observations have revealed that it is not perfectly so. Inhomogeneities in cosmology increase in size due to the attractive force of gravity. This means that in the past, these inhomogeneities were much smaller. As a result, for a significant part of their existence, these inhomogeneities can be treated as linear perturbations. However, the linear treatment becomes invalid on smaller scales in our recent history. This poses a challenge for reconstructing the primordial inhomogeneities from large-scale structures. Nonetheless, the linear treatment is suitable for describing the fluctuations of the CMB during the time of last scattering. Therefore, the CMB serves as the most effective observational tool for studying primordial inhomogeneities. To learn more about the relativistic theory of cosmological perturbations, which is briefly introduced in this chapter, readers are encouraged to refer to standard reviews in the literature [see, e.g. 25, 26, 28–30, 35–39].

1.2.1 Perturbed geometry

To study the effects on the geometry of the Universe, we make the assumption that the inhomogeneities only cause small deviations from homogeneity, which can be analysed using first-order perturbation theory. To this end, we define a perturbed geometry $g_{\mu\nu} = \bar{g}_{\mu\nu} + \delta g_{\mu\nu}$ where $\bar{g}_{\mu\nu}$ is the background (or unperturbed) FLRW metric, and the small perturbations around it can be defined as $\delta g_{\mu\nu} = a^2 h_{\mu\nu}$, where $|h_{\mu\nu}| \ll 1$. The perturbed metric in conformal time is, therefore:

$$d\tau^2 = (\bar{g}_{\mu\nu} + a^2 h_{\mu\nu}) dx^\mu dx^\nu. \tag{1.2.1.1}$$

1.2.1.1 Gauge Problem

In practice, constructing $\bar{g}$ and $\bar{T}$ from the physical fields g and T is not trivial. First, the metric perturbations are not uniquely defined, but depend on the choice of spacetime coordinates (or *gauge choice*). A wise choice of the gauge will be so that the averaged fields $\bar{g}$ and $\bar{T}$ satisfy Friedmann's equations. This is known as the *gauge freedom problem* in general relativity.

One way to deal with this is to describe perturbations using a *gauge-invariant*[19] general-relativistic formalism. Another way to avoid gauge freedom involves *fixing*

[19]Gauge transformations are a type of coordinate transformations that do not change the physical observables of a system. Instead, they transform the mathematical description of the system into an equivalent one, which is physically indistinguishable. A gauge-invariant formalism then implies that the expression is invariant under gauge transformations, which means that it gives the same physical result in different coordinate systems.

the gauge, and rely on computing only pure observables. It is worth noting that the metric tensor is not directly observable. What is observable are quantities like electromagnetic intensity, redshift, and angular size. The gauge problem therefore becomes irrelevant if the quantities of interest are expressed in terms of the true observables.

SVT decomposition For practical purposes, in what follows we will consider only the case of a flat FLRW background spacetime, as defined in Eq. 1.1.1.7, where we have defined the perturbed background as $g_{\mu\nu} = \bar{g}_{\mu\nu} + a^2 h_{\mu\nu}$. We consider now a fixed admissible Friedmann background ($\bar{g}$), We can decompose perturbations according to their behaviour under spatial rotations: the *Scalar-Vector-Tensor (SVT) decomposition.* The perturbed term can be parametrized in terms of SVT as follows:

$$h_{\mu\nu}dx^\mu dx^\nu = -2Ad\eta^2 - 2B_i d\eta dx^i + 2H_{ij}dx^i dx^j \tag{1.2.1.2}$$

where A corresponds to the scalar perturbation, B_i is the vector perturbation, and the H_{ij} is the tensor perturbation.

The SVT decomposition is especially useful due to the fact that the Einstein equations for scalars, vectors and tensors do not mix at linear order and can therefore be treated separately (see Appendix A.1 for more details). We obtain a general expression for the perturbed line element:

$$\begin{split} d\tau^2 = a^2[(1+2A)d\eta^2 + 2(\partial_i B + \hat{B}_i)d\eta\, dx^i + \\ (-(1-2C)\eta_{ij} + 2\partial_i\partial_j E + \partial_i \hat{F}_j + \partial_j \hat{F}_i + \hat{E}_{ij})dx^i\, dx^j]. \end{split} \tag{1.2.1.3}$$

Again, here η denotes conformal time, and x^i denotes comoving spatial coordinates. The metric perturbations comprise 10 degrees of freedom and can then be described by:

$$\begin{aligned} &4 \text{ Scalars}: && A, B, C, E && \text{(corresponding to 4 degrees of freedom)},\\ &2 \text{ Vectors}: && \hat{B}_i, \hat{F}_i && \text{(corresponding to 4 degrees of freedom)},\\ &1 \text{ Tensor}: && \hat{E}_{ij} && \text{(corresponding to 2 degrees of freedom)}, \end{aligned}$$

Gauge transformations If we consider a coordinate transformation by an infinitesimal change, of the form

$$x^\mu \to \tilde{x}^\mu \equiv x^\mu + \xi^\mu(\eta, \boldsymbol{x}) \tag{1.2.1.4}$$

and apply SVT decomposition to the gauge perturbation ξ^μ, we can split it into scalar and vector components: $(\xi^0, \xi^i) = (\xi^0, \partial_i\zeta + \xi^i_\perp)$, where (ξ^0, ζ) are scalars, and $\xi^i_\perp$ is a transverse vector.

Applying the appropriate transformation to each of the scalars, vectors, and tensors appearing in the perturbed line element in Eq. 1.2.1.3, we obtain the transformation laws[20]:

$$\begin{aligned}
&\text{Scalars}: \quad \tilde{A} = A - a^{-1}(a\xi^0)', \quad \tilde{B} = B + \zeta' - \xi^0, \quad \tilde{C} = C + \mathcal{H}\xi^0, \quad \tilde{E} = E + \zeta. \\
&\text{Vectors}: \quad \tilde{B}_i = B_i + \xi'_{\perp,i}, \quad \tilde{F}_i = F_i + \xi_{\perp,i}. \\
&\text{Tensors}: \quad \tilde{E}_{ij} = E_{ij}.
\end{aligned}$$

We find that coordinate transformations depend on two scalar and two vector degrees of freedom. These gauge modes can then be used to null out two of the scalar and two of the vector perturbation modes from the perturbed metric's degrees of freedom. In other words, only two scalar perturbation modes are gauge-invariant, and vector perturbations can be described with a single transverse vector.

Gauge choice The equations above can be used to obtain *gauge-invariant variables*, i.e. combinations of perturbations that are independent of the choice of coordinate system. With a convenient choice of gauge transformations, we define two gauge-invariant scalar variables (Ψ, Φ), one gauge-invariant transverse vector (V_i), and we retain the tensor metric perturbations, which are gauge-invariant by construction:

$$\begin{aligned}
&\Psi = A - a^{-1}(a(B - E'))', \quad \Phi = C + \mathcal{H}(B - E'), \\
&V_i = B_i - E'_i, \quad h_{ij} \equiv E_{ij}
\end{aligned} \tag{1.2.1.5}$$

where Ψ, Φ are the so-called *Bardeen potentials*. These gauge-invariant variables do not change under gauge transformation, and can therefore be considered as the "real" spacetime perturbations.

An alternative approach to the gauge problem consists in *fixing the gauge*, and relying on the computation of real observables, which do not depend on the gauge. A convenient choice is that of the so-called *"Newtonian" gauge*. This is defined by setting $B = E = 0$,

$$d\tau^2 = a^2(\eta)\left((1+2\Psi)d\eta^2 - (1-2\Phi)\left[\delta_{ij} + h_{ij}\right]dx^i dx^j\right). \tag{1.2.1.6}$$

[20] See Appendix A.2 for more details.

which includes both scalar and tensor perturbations, and where the component of the metric can be written as

$$g_{00}(\boldsymbol{x},\eta) = 1 + 2\Psi(x,\eta), \tag{1.2.1.7}$$

$$g_{0i}(\boldsymbol{x},\eta) = 0, \tag{1.2.1.8}$$

$$g_{ij}(\boldsymbol{x},\eta) = a^2(\eta)\,(\delta_{ij} + h_{ij})\,[1 - 2\Phi(\boldsymbol{x},\eta)]. \tag{1.2.1.9}$$

It is worth noting that vector perturbations are not included in the equations above. Here, we focus only on the degrees of freedom that are relevant to simple models for the initial evolution of the Universe, which do not predict the generation of vector perturbations. It can also be proven that, if vector perturbations were included, subsequent evolution after inflation would lead to their decay [29].

1.2.2 Perturbed Einstein field equations

By substituting the Newtonian gauge metric into Einstein's field equations 1.1.2.1, we can obtain a set of equations that describe the behaviour of the metric potentials. Specifically, we will focus on linear terms in the metric potentials, and we obtain:

$$\nabla^2\Phi - 3\mathcal{H}(\dot{\Phi} + \mathcal{H}\Psi) = 4\pi G a^2\,\delta T^0_0, \tag{1.2.2.1}$$

$$\partial_i(\dot{\Phi} + \mathcal{H}\Psi) = 4\pi G a^2\,\delta T^0_i, \tag{1.2.2.2}$$

$$\ddot{\Phi} + \mathcal{H}(2\Phi + \Psi)^{\cdot} + (2\dot{\mathcal{H}} + \mathcal{H}^2)\Psi + \frac{1}{3}\nabla^2(\Psi - \Phi) = -\frac{4\pi}{3} G a^2\,\delta T^i_i, \tag{1.2.2.3}$$

$$\partial_i\partial_j(\Psi - \Phi) = 8\pi G a^2\,\delta T^i_j \;\; (i \neq j). \tag{1.2.2.4}$$

The last equation describes the evolution of the difference between the two metric potentials, namely Ψ and Φ. This difference is sourced by the off-diagonal components of the perturbed EM tensor, which represent the anisotropic stress of the fluid, and describe the difference in pressure in different directions. However, in the standard cosmological model, the only fluid species that have significant anisotropic stress are neutrinos, a type of fundamental particles that have a tiny but non-zero mass, and they interact with other particles through the weak force, making them difficult to detect. The effect of the anisotropic stress of neutrinos on the evolution of the universe is very small, and therefore, assuming that $\Phi = \Psi$ is often an excellent approximation.

1.2.3 Perturbed matter

As mentioned before, matter in cosmology is commonly described as a perfect fluid. Its homogeneous part is described by the energy-momentum (EM) tensor defined earlier in Eq. 1.1.2.4.

EM tensor perturbations The density and pressure are defined as the density and pressure measured by an observer that is co-moving with the fluid ($\rho \equiv T_{\mu\nu}u^\mu u^\nu$, $p \equiv T_{\mu\nu}(g^{\mu\nu} - u^\mu u^\nu)/3$). In the homogeneous case, $\bar{u}_\mu = (a, 0, 0, 0)$, and therefore any spatial component of u_μ denotes a small coordinate velocity of the fluid, which can be treated as a linear perturbation. In detail, we write the perturbed 4-velocity as

$$u^\mu = \left(\frac{1-\Psi}{a}, \frac{\mathbf{v}}{a}\right), \tag{1.2.3.1}$$

where the 0-th component is fixed by the normalisation $u^\mu u_\nu = 1$. The perturbations to the energy-momentum tensor can then be written as:

$$\delta T^0_0 = \bar{\rho}\delta, \quad \delta T^0_i = (\bar{\rho} + \bar{p})v^i, \quad \delta T^i_j = -\delta p\, \delta^i_j. \tag{1.2.3.2}$$

where the v_i is the bulk velocity. Substituting this into Einstein's equations, we obtain:

$$k^2\Psi + 3\mathcal{H}(\dot{\Psi} + \mathcal{H}\Psi) = -4\pi G a^2\, \bar{\rho}\, \delta, \tag{1.2.3.3}$$

$$k^2(\dot{\Psi} + \mathcal{H}\Psi) = 4\pi G a^2\, (\bar{\rho} + \bar{p})\theta, \tag{1.2.3.4}$$

$$\ddot{\Psi} + 3\mathcal{H}\dot{\Psi} + (2\dot{\mathcal{H}} + \mathcal{H}^2)\Psi = 4\pi G a^2\, c_s^2 \bar{\rho}\, \delta, \tag{1.2.3.5}$$

where c_s is the sound speed, defined as $c_s^2 \equiv \delta p/\delta\rho$. Eq. 1.2.3.3 is the generalized Poisson's equation, where the second term on the right-hand side represents relativistic corrections, which are important only on scales comparable to the horizon, namely $k \lesssim \mathcal{H}$ ($\eta k \lesssim 1$).

Density perturbations During matter domination ($a \propto \eta^2$, $\mathcal{H} = 2/\eta$), assuming no pressure ($c_s^2 = 0$), Eq. 1.2.3.5 has solution $\psi = C_1 + C_2/\eta^5$, i.e. the gravitational potential Ψ does not evolve. Plugging this into Eq. 1.2.3.3, neglecting the decaying mode, and solving for *density perturbations* $\delta \equiv \delta\rho/\rho$, we obtain

$$\delta = \left[-\frac{(k\eta)^2}{6} - 2\right]\Psi. \tag{1.2.3.6}$$

This, on sub-horizon scales, is equivalent to $\delta \propto ak^2\Psi$, while on super-horizon scales, where relativistic effects come into play, the overdensity is frozen at $\delta \simeq -2\Psi$[21].

Although Ψ is constant on super-horizon scales $k \ll aH$ ($k\eta \gg 1$) when the equation of state is constant, it undergoes minor fluctuations during epoch transitions, like inflation to radiation domination, and then to matter domination.

[21]However, the behaviour on large scales is not gauge-invariant.

In contrast, it is possible to define a quantity which stays constant on super-horizon scales *at all times* (including during the Universe's transitions between different epochs):

$$\mathcal{R} \equiv -\Psi - \frac{\mathcal{H}(\dot{\Psi} + \mathcal{H}\Psi)}{4\pi G a^2(\bar{\rho} + \bar{p})}. \tag{1.2.3.7}$$

This is known as the *curvature perturbation* on comoving hypersurfaces, and is defined above in the Newtonian gauge.

Energy-momentum conservation From the conservation of the EM tensor, $\nabla_\mu T^m u_\nu = 0$, the continuity and Euler's equations can be derived:

$$\nu = 0: \quad \dot{\delta} = -\left(1 + \frac{\bar{p}}{\bar{\rho}}\right)\left(\frac{1}{a}\partial_i v^i - 3\dot{\Phi}\right) - 3\mathcal{H}\left(\frac{\delta p}{\bar{\rho}} - \frac{\bar{p}}{\bar{\rho}}\delta\right), \tag{1.2.3.8}$$

$$\nu = i: \quad \dot{v^i} = -\left(\mathcal{H} + \frac{\dot{\bar{p}}}{\bar{\rho} + \bar{p}}\right) v^i - \frac{1}{a}\left[\frac{1}{\bar{\rho} + \bar{p}}\left(\partial^i \delta p - \partial_j \delta^{ij}\right) - \partial^i \Psi\right]. \tag{1.2.3.9}$$

1.3 The Hot Big Bang model

The Universe is considered to be adiabatically expanding. Using the cosmological FLRW model of a flat Universe, we can express the time dependence with the density, using the Friedmann equations, as

$$t = \sqrt{\frac{3}{32\pi G\rho}} \quad \text{(radiation domination).} \tag{1.3.0.1}$$

$$t = \sqrt{\frac{1}{6\pi G\rho}} \quad \text{(matter domination)} \tag{1.3.0.2}$$

The Universe was radiation dominated in the early phases, as indicated by Eq. 1.1.2.9. For a photon gas, Stefan-Boltzmann's law tells us that $\rho_{\text{radiation}} \propto T^4$, which indicates that the Universe gets hotter and denser the further back in time we go. According to the *standard Hot Big-Bang model*, extrapolating sufficiently back in time ($t \to 0$), the Universe can be found in a state of arbitrarily high temperature and density ($\rho \to \infty$), where the laws of physics as we know them break down. This hot, dense, and highly energetic state is the so-called "Big Bang", at which point the temperature of the Universe is estimated using the Planck temperature, $T_p = \sqrt{\hbar c^5/Gk^2} \approx 1.4 \times 10^{32}$ K. Accordingly, the Universe then expanded rapidly, cooling and evolving to its current state over billions of years.

To understand the thermal history of the universe, it is crucial to compare the rate of expansion (H) to the *rate of interaction* can be expressed as:

$$\Gamma \equiv n \langle \sigma v \rangle \tag{1.3.0.3}$$

where n is the number density of particles, v is the relative velocity, and σ is the interaction cross-section, which depends on energy.

When $\Gamma \gg H$, particle interactions occur much faster than the universe's expansion, so local thermal equilibrium is reached before the expansion has a significant impact. As the universe cools, the rate of interactions may decline more rapidly than the expansion rate. When $\Gamma \sim H$, particles break away from the thermal bath. A given species i will decouple from the thermal bath when the interaction rate drops below the rate of expansion, $\Gamma_i \ll H$. It is worth noting that different particle types may have varying interaction rates and thus decouple at different times.

This allows us to define a brief thermal history of the Universe, through the main events of the Hot Big-Bang model, which are broadly described in the following section. This is largely based on the following references: [30, 36, 40–43].

1.3.1 Thermal equilibrium

The very early universe expanded very rapidly, but also had a very high temperature and density, yielding particles to constantly collide and exchange energy, at a much faster rate than the typical expansion rate, i.e. $\Gamma \gg H$. Therefore, we assume that the initial conditions were close to thermal equilibrium, meaning that the particles in the universe were in a state of balance with each other and their temperature and density were uniform throughout. This concept was initially explored by Gamow, Alpher, and Herman in [10, 44–46].

Thermodynamics. If we further assume that we can treat the simple fluids of interest as *perfect gases*[22], we can easily recover their thermodynamics. We start with a box of volume $V = L^3$ and expand the fields inside into periodic waves with harmonic boundary conditions.

In this system, the number of allowed quantum states per unit volume of k space, where k represents the momentum of a particle, is determined by the *density of states*:

$$dN = g\frac{V}{(2\pi)^3}d^3k \tag{1.3.1.1}$$

where g is a degeneracy factor that takes into account the possible states of the system (such as spin), and the term $\frac{V}{(2\pi)^3}d^3k$ represents the volume of the shell of

[22]Gases at low pressures and high temperatures, where intermolecular interactions are weak, can be approximated as "perfect" or "ideal" gases. Their behaviour is described by the *ideal gas law*, $PV = nRT$, where P is the pressure, V is the volume, n is the number of particles, R is the gas constant, and T is the absolute temperature.

states between k and $k + dk$ in 3D momentum space. The probability of finding a particle in a specific quantum state with energy ϵ in a system that is in thermal equilibrium is then quantified by its *occupation number* (or *distribution function*):

$$\langle f_{\mathrm{eq}} \rangle = \left[e^{(\epsilon-\mu)/k_B T} \pm 1 \right]^{-1} \tag{1.3.1.2}$$

where μ is the chemical potential of the system, k_B is the Boltzmann constant, T is the temperature, and the $\pm$ sign is for fermions and bosons, respectively.

In this system, the number of particles, N, will naturally adjust to reach its equilibrium state. At equilibrium, the system is expected to be stationary with respect to small changes in N. In other words, the number of particles in the system will remain relatively constant unless there is a significant change to the system's temperature, pressure, or volume. Accordingly, for a thermal radiation background in equilibrium the *chemical potential* μ is always zero. This is because μ appears in the *first law of thermodynamics* as the change in energy associated with a change in particle number,

$$dE = TdS - PdV + \mu dN \tag{1.3.1.3}$$

When μ is zero, it implies that the energy of the system is independent of the number of particles, and the system has reached equilibrium.

Using the expressions above, we can recover two important quantities that allow us to determine the properties of matter at high temperatures in the early universe: the number density and the energy density of particles in thermal equilibrium. The expression for the thermal equilibrium background *number density* of particles is given by:

$$n = \frac{1}{V} \int f dN = g \frac{1}{(2\pi\hbar)^3} \int_0^\infty \frac{4\pi p^2 dp}{e^{\epsilon(p)/kT} \pm 1}, \tag{1.3.1.4}$$

where the integration has been carried out in momentum space, with $\epsilon = \sqrt{m^2c^4 + p^2c^2}$. There are two interesting limits of this expression:

1. In the *relativistic regime*, $T \gg m$, relativistic particles (radiation) in thermal equilibrium behave as if they are massless. The thermal equilibrium background number density of particles can be described by the following expression:

$$n = \left(\frac{kT}{c}\right)^3 \frac{4\pi g}{(2\pi\hbar)^3} \int_0^\infty \frac{y^2 dy}{e^y \pm 1}. \tag{1.3.1.5}$$

2. In the *non-relativistic regime*, $T \ll m$, after the temperature drops below the mass of the particles, the ± 1 term in the occupation number can be ignored. As a result, the number is suppressed by a dominant factor of $\exp\left(-mc^2/kT\right)$, indicating that the background "turns on" at approximately $kT \sim mc^2$. At this threshold energy, photons and other species in equilibrium can produce particle-antiparticle pairs.

The **energy density** of the background can be easily obtained by multiplying the integrand by the energy of each mode $\epsilon(p)$:

$$u = \rho c^2 = g \frac{1}{(2\pi\hbar)^3} \int_0^\infty \frac{4\pi p^2 dp}{e^{\epsilon(p)/kT} \pm 1} \epsilon(p) \tag{1.3.1.6}$$

1.3.1.1 Relativistic particles (radiation)

Thermodynamics. For $T \gtrsim 100\,\mathrm{GeV}$, all known particles are relativistic, and particle masses can be ignored. Therefore we use temperature T as a dimensionful scale to express other quantities. In this limit, $\epsilon(p) = pc$, Eq. 1.3.1.6 above becomes

$$u = \frac{\pi^2}{30(\hbar c)^3} g (kT)^4 \quad \text{(bosons)} \tag{1.3.1.7}$$

For Fermions, we can express their occupation number in terms of the Bosonic occupation number at two different temperatures, with the following trick: $1/\left(e^x + 1\right) = 1/\left(e^x - 1\right) - 2/\left(e^{2x} - 1\right)$. This means that a gas of Fermions can be thought of as a mixture of Bosons at two different temperatures. Using this, and knowing that boson number density and energy density scale as $n_B \propto T^3$ and $u_B \propto T^4$, we find that the number density and energy density of Fermions are related to those of Bosons as $n_\mathrm{F} = (3/4) n_\mathrm{B}$ and $u_\mathrm{F} = (7/8) u_\mathrm{B}$.

To fully characterize the thermal background, it's also important to determine its entropy. This takes an interesting form in the ultrarelativistic limit, where the energy density follows the usual black-body scaling law $u \propto T^4$, and the pressure is $P = u/3$. Thus, the entropy density, $s \equiv S/V$, of the system can be expressed as

$$s = \frac{4}{3} \frac{u_B}{T} = \frac{2\pi^2 k}{45(\hbar c)^3} g (kT)^3 \quad \text{(bosons)}, \tag{1.3.1.8}$$

where the left-hand side is obtained by substituting Eq. 1.3.1.7 for u_B, and for fermions we then have $s_F = (7/8) s_B$. The equation above shows that $s_B \propto T^3$, as is the boson number density in an ultrarelativistic background, $n_B \propto T^3$, as mentioned before. This leads to the interesting observation that entropy density is proportional to the number density: $s \propto n$. This observation justifies the use of the term *"entropy per baryon"* to describe the ratio of the number density of photons in the universe to the number density of baryons. The large value of this ratio, which is on the order of 10^9, indicates that most of the entropy in the universe is associated with photons rather than baryons[23].

Degrees of freedom From the energy density $u = \rho c^2$ in the relativistic limit, given by Eq. 1.3.1.7, we find the *total energy density of radiation*:

$$\rho_r c^2 = \frac{\pi^2}{30(\hbar c)^3}\left(\sum_{i\in \text{ bosons}} g_i (kT)^4 + \frac{7}{8}\sum_{j\in \text{ fermions}} g_j (kT)^4\right) \equiv g_* \frac{\pi^2}{30}(kT)^4, \quad (1.3.1.9)$$

where g_* counts the total number of relativistic degrees of freedom. A similar equation can be written for the total entropy density of relativistic particles in the Universe:

$$s_r = \left[2\pi^2 k/45(\hbar c)^3\right] g_*^S (kT)^3 \quad (1.3.1.10)$$

If all relativistic species are in thermal equilibrium, then $g_* = g_*^S$. In our Universe this is the case until $t \sim 1$ s, when neutrinos decouple from the thermal bath (see below).

1.3.1.2 Time and temperature

In a thermal system i with temperature T_i that is expanding adiabatically, *the entropy is conserved*, so we get:

$$d\left(g_{*,i}^S a^3 T_i^3\right) = 0 \quad \Rightarrow \quad T_i \propto (g_{*,i}^S)^{-1/3} a^{-1} \quad (1.3.1.11)$$

i.e. the entropy is $s_i \propto T_i^3 \propto a^{-3}$. If there are multiple independent thermal systems, as is the case after decoupling, this law applies to each system separately, where the temperatures of each system may differ.

CMB blackbody spectrum If the photons in the early universe originated from a thermal equilibrium distribution, we would expect to see a blackbody spectrum. This can be explained as follows.

Each photon with frequency ν has an energy of $E_\gamma = p = h\nu$, and the total energy received over an area dA with photon direction (momentum) in solid angle $d\Omega$ is given by:

$$E_\gamma f_\gamma d^3\mathbf{p} d^3\mathbf{x} = E_\gamma f_\gamma p^2 dp d\Omega_\mathbf{p} d^3\mathbf{x} = E^3 f_\gamma dE d\Omega_\mathbf{p} dA c dt \quad (1.3.1.12)$$

[23] This observation also justifies the assumption of *adiabaticity*, which is based on the idea that the entropy of a system remains constant during an adiabatic process. In the case of the universe, this assumption is reasonable because most of the entropy is associated with photons, which are expected to be in thermal equilibrium throughout most of the universe's history.

Here, f_γ is the equilibrium distribution function for photons, given by Eq.1.3.1.2, where $\epsilon = E_\gamma$, and photons are bosons with $\mu = 0$ and two spins (polarisations). The intensity (power per unit area per unit solid angle per unit frequency) observed is the blackbody spectrum, given by:

$$B(\nu) = \frac{2h\nu^3}{c^2}\frac{1}{e^{\frac{h\nu}{kT}} - 1} \tag{1.3.1.13}$$

Observations by FIRAS fit this spectrum to high accuracy with a temperature of $T_{\gamma 0} = 2.725$ K, which tells us that the photons were very close to equilibrium at some point in the early universe. As the universe expands adiabatically, the temperature of cosmic radiation decreases while the blackbody spectrum of photons is preserved. It is worth noting that the spectrum peaks at a given $T_{\rm peak}$, corresponding to a frequency $\nu_{\rm peak}$. Since the photon frequency gets redshifted, the temperature gets redshifted as well, and we assume that:

$$T(z) = T_\gamma\,(1 + z), \tag{1.3.1.14}$$

This is justified by noting that photon energies scale as $E \propto 1/a$ and assuming that the typical energy in blackbody radiation is $\sim kT$. The same result can also be demonstrated by preserving entropy (Eq. 1.3.1.11). Since in that case $s \propto a^{-3}$, and we saw that $s \propto T^3$ (Eq. 1.3.1.10), we find that $T_i \propto a^{-1}$, assuming $g^S_{*,i}$ remains constant.

1.3.2 Non-relativstic matter and freeze-out

In general, if a species has $\Gamma_i \gg H$ and interactions can maintain equilibrium, the occupation number $f_{\rm eq}$ will follow the equilibrium distribution described in Eq. 1.3.1.2. In these conditions, a particle species that interacts with the photons will have the same temperature as the photons: $T_i = T_\gamma$. However, when a species i becomes completely decoupled ($\Gamma_i \ll H$), its particles will travel freely, and the distribution function will *freeze-out* in the form it had at decoupling, i.e. the particles have effectively ceased to interact, reaching a state of chemical equilibrium, while their momenta gradually redshift as the universe expands. This preserves a "snapshot" of the properties of the universe from the time when the particle was last in thermal equilibrium. The expansion of the Universe then cools down the containing matter, such that particle species can become non-relativistic ($T_{Di} \gg m_i$), and T_{Di} is the *decoupling or freeze-out temperature*, characteristic for each particle species i.

In this limit, the energy density, pressure, number density and entropy density of particles in thermal equilibrium are given by:

$$\rho_i = m_i n_i \tag{1.3.2.1}$$

$$p_i = n_i T \tag{1.3.2.2}$$

$$n_i = g_i m_i \left(\frac{m_i T}{2\pi}\right)^{\frac{3}{2}} \exp\left(-\frac{m_i - \mu_i}{T}\right) \tag{1.3.2.3}$$

$$s_F = \frac{(m_i + T)\, n_i}{T}. \tag{1.3.2.4}$$

1.3.2.1 Neutrino decoupling

In cases where a massive particle decouples while still relativistic ($T_D \gg m$), its distribution function remains frozen in the form of the distribution function f_{eq} of massless particles, and T_D is the decoupling or freeze-out temperature, characteristic for each particle species. Later, when the temperature of the thermal bath drops below the particle's mass, it becomes non-relativistic, with energy $E \simeq m$. The distribution function and number density of the particles still correspond to the frozen-in form of relativistic particles, while the energy density becomes that of non-relativistic particles, as in Eq. 1.3.2.1. This is what happens in the case of massive neutrinos.

At the later stages of the Big Bang, energies are such that only light particles survive in equilibrium: photons (γ), neutrinos (ν) and positron-electron pairs e^+e^-. Neutrinos are coupled to the thermal bath via weak interaction processes like

$$\nu_e + \bar{\nu}_e \leftrightarrow e^+ + e^-, \tag{1.3.2.5}$$

$$e^- + \bar{\nu}_e \leftrightarrow e^- + \bar{\nu}_e \tag{1.3.2.6}$$

The cross section for these types of interactions can be estimated as

$$\sigma \simeq G_F^2 E^2 \sim G_F^2 T^2, \tag{1.3.2.7}$$

where G_F is the Fermi constant $\left(G_F = \pi\alpha_W / \left(\sqrt{2} m_W^2\right) = 1.1664 \times 10^{-5}\mathrm{GeV}^{-2}\right)$. Using the formula given in Eq. 1.3.0.3, we can calculate the rate of interaction to be $\Gamma \sim G_F^2 T^5$. However, as the temperature decreases, the rate of interaction decreases much faster than the Hubble rate, which is approximately $H \sim T^2/M_{\mathrm{pl}}$. This can be expressed as:

$$\frac{\Gamma}{H} \sim \left(\frac{T}{1\mathrm{MeV}}\right)^3 \tag{1.3.2.8}$$

indicating that neutrinos decouple at around 1MeV, corresponding to a time of roughly 1 s.

1.3.2.2 Electron-positron annihilation

As the temperature drops below $T = 10^{9.7}$ K, the electron-positron pairs (e^+e^-) annihilate. Although the annihilations could potentially produce pairs of photons or neutrinos due to electron interactions through either the electromagnetic or the weak force, in reality, weak reactions freeze out earlier, at around $T \simeq 10^{10}$ K. Consequently, the electron-positron annihilation increases the number of photons compared to neutrinos, while conserving entropy. The entropy of an $e^\pm + \gamma$ gas, which is proportional to the number density, is easily calculated as each particle species has $g = 2$ (polarisation or spin). The total entropy is, therefore, given by

$$s\left(\gamma + e^+ + e^-\right) = \frac{11}{4}s(\gamma) \tag{1.3.2.9}$$

By equating this to the photon entropy at a new temperature, we obtain the factor by which the photon temperature (T_γ) is enhanced relative to that of the neutrinos (T_ν). This factor reveals the presence of a *neutrino background* (CνB) with a temperature of

$$T_\nu = \left(\frac{4}{11}\right)^{1/3} T_{\gamma 0} = 1.945 \text{ K}, \tag{1.3.2.10}$$

given the CMB temperature today $T_{\gamma 0} = 2.725$ K.

1.3.3 Matter-radiation equality

The total matter density and radiation today is

$$\rho_{m0} = \frac{3H_0^2\Omega_m}{8\pi G} = 1.88 \times 10^{-32}\Omega_m h^2 \text{ kg cm}^{-3} \tag{1.3.3.1}$$

$$\rho_{r0} = g_* \frac{\pi^2}{30}\frac{(k_B T_{\gamma 0})^4}{c^5\hbar^3} = 7.8 \times 10^{-37} \text{ kg cm}^{-3} \tag{1.3.3.2}$$

Here we used the definition that the Hubble parameter today is $H_0 = h100$ km s^{-1}Mpc^{-1}.

Using the fact that $\rho_r/\rho_m \propto a_0/a = 1 + z$, it follows that the redshift corresponding to when the matter-radiation contributions to energy density were equal (*matter-radiation equality*) is given by:

$$1 + z_{\rm eq} = 2.4 \times 10^4 \Omega_m h^2 \tag{1.3.3.3}$$

Evidence suggests $\Omega_m h^2 \sim 0.133$, so $z_{\rm eq} \sim 3200\,(T_{eq} \sim 0.75\,\text{eV})$. Prior to this redshift the universe is radiation dominated, afterwards it is matter dominated until dark energy becomes important at low redshift $(z \sim 1)$.

1.3.4 Recombination and the CMB

The process of recombination in the early universe involves the capture of electrons by protons to form neutral hydrogen atoms, and it is essential for the CMB radiation to be released.

Saha equilibrium. After electron positron annihilation, the concentration of electrons can be written as $n_{\rm e} = X_{\rm e} n_{\rm B}$ where $n_{\rm B}$ is the concentration of baryons[24] and $X_{\rm e}$ is the *ionization fraction.* At high temperatures, deep in the radiation-dominated epoch, electrons and nuclei could not yet effectively combine to form neutral hydrogen ($X_{\rm e} > 1$). When the concentration of electrons equals the concentration of protons, i.e. $n_{\rm e} = n_{\rm B}$, the Universe is globally neutral. As the temperature gets low enough, after matter-radiation equality, the abundance of free electrons fell steeply, since electrons and nuclei started to combine to form neutral hydrogen:

$$p + e^- \leftrightarrow H + \gamma \tag{1.3.4.1}$$

This time is called *recombination*[25]. In order to establish when recombination took place, we assume that hydrogen is formed in its lowest energy level during recombination.

We further assume that the neutral and ionised hydrogen, as well as the free electrons, are in thermal equilibrium, which is true if the recombination rate is much higher than the expansion rate, $\Gamma_{rec} \gg H$. In this situation, since electrons, protons and hydrogen are non-relativistic, their concentrations ($n_{\rm e}$, $n_{\rm p}$, $n_{\rm H}$, respectively) follow from the equilibrium distribution in Eq. 1.3.2.3. Since we are in a situation of chemical equilibrium, and recalling that $\mu_\gamma = 0$, we can relate the various chemical potentials according to the order of the reaction, i.e. $\mu_{\rm p} + \mu_{\rm e} = \mu_{\rm H}$. To remove the dependence on the chemical potentials, we consider the following ratio

$$\left(\frac{n_{\rm H}}{n_e n_p}\right)_{\rm eq} = \frac{g_{\rm H}}{g_e g_p}\left(\frac{m_{\rm H}}{m_e m_p}\frac{2\pi}{T}\right)^{3/2} e^{(m_p+m_e-m_{\rm H})/T} \tag{1.3.4.2}$$

$$= \left(\frac{m_{\rm e} T}{2\pi}\right)^{-3/2} e^{E_0/T}, \tag{1.3.4.3}$$

where in the last equality we have defined the degrees of freedom $g_p = g_e = 2$ and $g_H = 4$[26]. Here we have used $m_{\rm H} \approx m_p$ in the prefactor, but not it in the exponential.

[24] We are going to assume that at this point all the baryons in the Universe are in the form of protons, neglecting the smaller but significant helium fraction.

[25] The name can be misleading, as this is the first time electrons and nuclei combined.

[26] The spins of the electron and proton in a hydrogen atom can be aligned or anti-aligned, giving one singlet state and one triplet state, so $g_H = 1 + 3 = 4$

where we have defined the absolute value of the *binding energy of the hydrogen atoms* $E_0 = m_\mathrm{e} + m_\mathrm{p} - m_\mathrm{H} = 13.26\mathrm{eV}$, corresponding to the lowest energy level.

The equation above can be rewritten in terms of the ionisation fraction $X_e \equiv n_e/n_B$, assuming $n_e = n_p$, and expressing the baryon density as $n_B = \eta \cdot n_\gamma = \eta \cdot (2\zeta(3) \cdot T^3)/\pi^2$, and $\eta = 5.5 \cdot 10^{-10}\,(\Omega_B h^2/0.020)$. At this point, over 90% of the nuclei are protons, so the total baryon number density can be approximated as $n_B \approx n_p + n_\mathrm{H} = n_e + n_\mathrm{H}$. Therefore we find:

$$\frac{1-X_e}{X_e^2} = \frac{n_\mathrm{H}}{n_e^2} n_B = \frac{n_H}{n_e^2}\frac{n_B}{n_\gamma} n_\gamma = \frac{2\zeta(3)}{\pi^2}\frac{n_B}{n_\gamma}\left(\frac{2\pi T}{m_e}\right)^{3/2} \exp\left(\frac{E_0}{T}\right). \tag{1.3.4.4}$$

This is the *Saha equation* for the equilibrium ionization fraction, which determines X_e as a function of temperature. Solving for T as a function of X_e (which must be done numerically), and remembering that $1 + z = T/T_{\gamma 0}$, we obtain an estimate for the redshift at which recombination happens, $z^{\mathrm{eq}}_{\mathrm{rec}} \sim 1300$.

Two-photon decay. The Saha equation is a useful approximation only at the beginning of recombination, when there are few enough photons in the high-energy tail of the blackbody distribution. However, as recombination proceeds, the capture of a free electron by a proton to form a neutral hydrogen atom, $e + p \to H + \gamma$, becomes less efficient. That is, if the electron is captured directly to the ground state, the photon emitted during this process has sufficient energy to immediately ionize another neutral hydrogen atom, resulting in no net recombination.

Instead, recombination mostly proceeds through a two-step process. First, the electron is captured to an excited state, which can be any energy level above the ground state, before relaxing down to the 2p state. At this point, there is a bottleneck in the process because the 2p state has a long lifetime and is metastable, meaning that the electron can remain in this state for a relatively long time, which slows down the recombination process. To complete the recombination process, the electron in the 2p state needs to transition to the ground state 1s of a hydrogen atom (Lyman-α transition). This gives a resonance Lyman-α photon, which can immediately excite the 1s state of another atom.

The whole hydrogen recombination therefore can occur through the emission of two photons, in a process known as *two-photon decay.* The two photons each carry away some of the energy released during the transition and have energies that are much lower than the Lyman-α photon, resulting in a gradual redshifting of the distribution of resonance photons, which eventually allows them to escape, leaving a net recombination. Once the hydrogen atom is in the ground state, it can no longer be ionized, and so it remains neutral[27].

The two-step recombination process leads therefore to a slower rate of recombination than predicted by the Saha equation. A full non-equilibrium calculation, accounting for the recombination rates to different excited states, and the transitions between them, finds recombination to occur at around $z_{\rm rec}^{\rm non\text{-}eq} \sim 1090$, later than predicted by the Saha equation ($z^{\rm eq} > z^{\rm non\text{-}eq}$).

Last scattering. Observationally, recombination is important as it marks the first time that photons could travel freely. In the early highly-ionized Universe, photons moved in a random walk due to Thomson scattering, making the universe opaque. As the universe continued to expand, the mean free path of photons increased, outside the horizon. As a result, photons decoupled from matter and were able to travel almost freely, making the Universe essentially transparent. The photons that emanated at this point, known as *last scattering surface*, form the CMB that we observe today.

In reality, however, recombination did not happen instantaneously. Hence, when we look at a given point on the last scattering surface, we are actually seeing a mixture of photons coming from different times during recombination. The *visibility function*, $g(\eta)$, denotes the probability of a photon's last scattering as a function of time η, and is given by:

$$g(\eta) \equiv -\dot{\tau}(\eta) e^{-\tau(\eta)}, \quad \text{where} \quad \tau(\eta) \equiv \int_{\eta}^{\eta_0} \Gamma\left(\eta'\right) d\eta'. \tag{1.3.4.5}$$

Above we have defined $\tau(\eta)$, the *optical depth* between time η and today, where $\Gamma\left(\eta'\right) = a n_e \sigma_T$. The factor $e^{-\tau}$ then defines the fraction of photons from time η that we actually receive rather than being scattered, and $g(\eta)$ is sharply peaked around the time of recombination, when the optical depth drops to almost zero, as the mean-free-path of photons was very long.

After recombination, most photons travelled through the inhomogeneous Universe almost unimpeded. These photons left imprints of fluctuations in the radiation temperature, gravitational potentials, and bulk velocity of the radiation at the location of their last scattering, creating the CMB anisotropies that we observe today (see Sec. 1.5 for more details).

[27]The ionization of the plasma cannot be maintained by electromagnetic interactions alone, as is assumed in the thermal-equilibrium approach. Instead, the simultaneous emission of two photons conserve energy and angular momentum, allowing an electron to make a transition from an excited state to a lower energy state without violating the laws of physics, e.g. allowing for the $2S \to 1S$ transition, which is otherwise strictly forbidden at first order because it violates the selection rules for one-photon emission.

1.4 Inflation

The widely accepted Hot Big-Bang model predicts that the Universe began with an initial singularity where $a \to 0$, based on the evolution of a radiation-dominated Universe. However, some predictions of the model lead to a number of problems when confronted with cosmological data. The concept of inflation was initially introduced to address these issues, as described below. Incidentally, it soon became apparent that inflation could also provide a plausible explanations for the observed fluctuations in both the geometry of spacetime and the distribution of matter.

Inflation can be broadly defined as a phase of accelerated expansion of the Universe, which is characterised by $\ddot{a} > 0$, where a is the scale factor. Based on current cosmological observations, it is possible that our present Universe is undergoing such an expansion. However, when referring to *inflation*, we are interested in an inflationary phase that occurred in the very early Universe and at different energy scales. References [28, 30, 36, 47] form the primary basis for this section's discussion of inflation.

1.4.1 Problems of the Big Bang model

Observations of the CMB have highlighted two primary issues that motivate the need for inflation. These are summarised in what follows, together with a brief explanation of how inflation provides plausible solutions to them.

The horizon problem. This arises when we consider the large-scale homogeneity and isotropy of the observable universe. In particular, the CMB is emitted at $z_d \sim 1100$. At this point, the *comoving horizon* measures

$$\chi_{\rm h} = \int_{z_d}^{\infty} \frac{dz'}{H(z')} \simeq 250\,\text{Mpc}. \tag{1.4.1.1}$$

Meanwhile, the distance from the last-scattering surface to us is

$$\chi_{\rm lss} = \int_{0}^{z_d} \frac{dz'}{H(z')} \simeq 14\,\text{Gpc}. \tag{1.4.1.2}$$

As a result, the angle subtended by the horizon is roughly $\theta_{\rm h} = \chi_{\rm h}/\chi_{\rm lss} \sim 1^\circ$. Given this, a CMB map ought to reveal about $\mathcal{O}(10^4)$ causally-disconnected patches. Nonetheless, despite this, the temperature of the CMB is uniform to within 1 part in 10^5 all over the sky.

The flatness problem. This refers to the discrepancy between the observed value of the curvature parameter Ω_κ and the expected value based on the evolution of the Universe. The curvature parameter at any given point in the Universe's history, is defined in Eq. 1.1.2.11 in conformal time, and can be expressed equivalently as:

$$\Omega_\kappa(t) = \frac{\kappa}{(aH)^2} = \frac{\kappa}{\dot{a}^2}\,. \tag{1.4.1.3}$$

In a decelerating Universe, $\dot{a}$ decreases over time, causing $|\Omega_k|$ to increase. However, observations of the CMB indicate that the Universe is nearly flat, which means that the value of $|\Omega_{k,0}|$ is very small, with an upper limit of around $\lesssim 10^{-3}$, which suggests that in the past it must have been much smaller (e.g., $|\Omega_k| \lesssim 10^{-17}$ during the nucleosynthesis epoch). This raises the question of why the initial conditions were so precisely tuned to create a Universe that is nearly Euclidean.

Solutions If we assume that the universe expands according to a power law, $a(t) \propto t^\alpha$, then the size of the horizon at a given time t is given by the integral:

$$\chi_H = \int_0^t \frac{dt'}{a(t')}. \tag{1.4.1.4}$$

If $\alpha < 1$, this integral is finite when evaluated from $t = 0$, which means that the size of the horizon at any given time is finite, and there are regions of the universe that are outside of each other's causal horizon. This is the root of the horizon problem.

To make the size of the horizon arbitrarily large (and the integral above diverge), we need $\alpha > 1$, in which case we find:

$$\ddot{a} \propto \alpha(\alpha - 1) > 0 \tag{1.4.1.5}$$

i.e. the expansion is accelerating. Therefore, an epoch of accelerated expansion is necessary to solve the horizon problem.

In addition to solving the horizon problem, an epoch of accelerated expansion can also help with the flatness problem. The curvature of the universe is determined by $\Omega_\kappa \propto \dot{a}^{-2}$. An accelerated expansion will increase $\dot{a}$, which in turn will decrease Ω_κ. This allows Ω_κ to take arbitrarily large values at sufficiently early times, which is necessary for consistency with observations.

1.4.2 Mechanism for inflation.

The theory of inflation stands out because the rapid expansion can transform the initial vacuum quantum fluctuations into "macroscopic" cosmological perturbations at a larger scale (as highlighted in the seminal works in [48–51]). Therefore, inflation yields natural initial conditions that align with present-day observations.

According to the Friedmann equations, acceleration of the Universe can only occur if the equation of state satisfies the condition $p < -\frac{1}{3}\rho$. A simple example of an equation of state that satisfies this condition is a cosmological constant, which corresponds to a cosmological fluid with $p = -\rho$. However, a strict cosmological constant would result in a never-ending exponential inflation, which is not compatible with a radiation or matter era. Alternatively, a dynamical field that eventually "switches off", such as a *scalar field*, can also satisfy this condition and provide a mechanism for inflation, which we will now discuss in more detail.

1.4.2.1 Scalar field

The Lagrangian is a mathematical function that describes the dynamics of a physical system. In the case of a scalar field ϕ, the Lagrangian density is given by

$$\mathcal{L}_\phi = \frac{1}{2}\partial_\mu\phi\partial^\mu\phi - V(\phi), \tag{1.4.2.1}$$

where the first term describes the kinetic energy of the field, while the second term, $V(\phi) = m^2\phi^2$, describes its potential energy density. In the case of the scalar field, the stress-energy tensor is given by [43]

$$\begin{aligned} T^{\mu\nu} &= -2\frac{\partial\mathcal{L}_m}{\partial g_{\mu\nu}} - g^{\mu\nu}\mathcal{L}_m \\ &= \partial^\mu\phi\partial^\nu\phi - g^{\mu\nu}\left(\frac{1}{2}\partial\nu'\phi\partial^{\nu'}\phi - V\right) \end{aligned} \tag{1.4.2.2}$$

In an isotropic and homogeneous Universe, $\phi \equiv \phi(t)$, i.e. the scalar field only depends on time. In this case, the stress-energy tensor $T^{\mu\nu}$ can be written in the same form as that of a perfect fluid (see Eq. 1.1.2.4), from which we can determine the scalar field effective energy density and pressure, as $\rho = T^0{}_0$ and $T^i{}_j = -p\delta^i{}_j$ for $i, j = 1, 2, 3$. In other words, the pressure is the same in all directions, while the density can vary. We find:

$$\rho = \frac{1}{2}\dot{\phi}^2 + V(\phi), \quad p = \frac{1}{2}\dot{\phi}^2 - V(\phi) \tag{1.4.2.3}$$

where $\dot{\phi}$ represents the time derivative of the scalar field. From the first Friedmann equation 1.1.2.5, and using the definition of ρ above, we obtain:

$$H^2 = \frac{8\pi G}{3}\left[\frac{1}{2}\dot{\phi}^2 + V(\phi)\right], \tag{1.4.2.4}$$

where the second and third term correspond to kinetic and potential energy.

It is useful to define the *equation of motion for the scalar field*, known as the *Klein-Gordon* (K-G) equation, in a Universe undergoing expansion is given by:

$$\ddot{\phi} + 3H\dot{\phi} + V'(\phi) = 0 \tag{1.4.2.5}$$

where $\dot{} \equiv d/dt$ and $V'(\phi)$ is the derivative of the potential energy density with respect to ϕ. The K-G equation is similar to that of a damped harmonic oscillator, with the force given by $-V'(\phi)$ and the expansion rate H acting as the damping term[28].

For inflation to occur, specific requirements must be met by the scalar field. In order to be deemed successful, an inflationary model must satisfy two critical conditions: first, the ability to facilitate accelerated expansion, and second, a prolonged but finite duration of that expansion. Here, we present a definition of a successful inflationary model for a scalar field.

1.4.2.2 Slow-roll

Conditions for inflation. In order to obtain inflation, i.e. $\ddot{a} > 0$, the pressure defined above in Eq. 1.4.2.3 must be negative, $p < 0$. In the case of a scalar field, this can be achieved only if

$$\dot{\phi}^2 \ll V. \tag{1.4.2.6}$$

This leads to the idea of *"slow-roll" inflation*, whereby a scalar field (*"inflaton"*) starts from an initial value away from the minimum of the potential, with a small velocity, i.e. its kinetic energy is much smaller than its potential energy. In this limit, we have $p = -\rho$, which corresponds to the equation of state of a Universe dominated by a vacuum energy Λ, i.e. Eq. 1.1.2.3 with $w = -1$, which then leads to *exponential expansion* in the very early Universe, induced by a scalar field (*"inflaton"*) in a potential $V(\phi)$. From this, we also find that $\rho + 3p < 0$, hence the Friedmann Eq. 1.1.2.6 entails

$$\ddot{a} \propto -(\rho + 3p) > 0, \tag{1.4.2.7}$$

[28]This can be compared to a ball rolling on a hill with friction, where the shape of the hill corresponds to the potential and the friction corresponds to the expansion rate. If the field starts away from the minimum of the potential, it will roll down the hill towards the minimum.

i.e. at the beginning of inflation the scalar field drives an *accelerated expansion*.

In order to have a long-lasting acceleration, the velocity of the field must change very slowly compared to the Universe's expansion rate:

$$|\ddot{\phi}| \ll 3H\dot{\phi} \tag{1.4.2.8}$$

In other words, as the field rolls down the potential towards the minimum, releasing the energy that drives the exponential expansion, it gains kinetic energy and speeds up. However, the friction term slows down its descent, allowing inflation to last for a long time.

With the slow-roll conditions defined in Eqs. 1.4.2.8-1.4.2.8, the K-G and Friedmann equations (Eq. 1.4.2.5 and 1.4.2.4, respectively) reduce to the basic *slow-roll equations*:

$$3\dot{\phi} = -\frac{1}{H}V' \tag{1.4.2.9}$$

$$H^2 = \frac{8\pi G}{3}V = \frac{1}{3M_P^2}\left(V\right), \tag{1.4.2.10}$$

where we have used the reduced Planck mass $M_P = m_P/\sqrt{8\pi} = 1/\sqrt{8\pi G}$.

This mechanism allows for inflation to occur without requiring fine-tuning of initial conditions. Moreover, it enables a smooth transition out of inflation ("graceful exit"), as described below.

Graceful exit. Once the scalar field reaches the minimum of the potential, it oscillates around it, and the kinetic energy becomes larger than the potential energy, $\dot{\phi}^2/2 > V(\phi)$, such that the oscillations are damped by the $3H\dot{\phi}$ friction term. Eventually, the oscillation frequency $\propto m$ becomes larger than H. Assuming that the minimum of the potential can be approximated as a quadratic function, the K-G equation reads:

$$\ddot{\phi} + m^2\phi = 0 \quad \Rightarrow \quad \phi \propto \cos(mt + \alpha) \tag{1.4.2.11}$$

In this regime, the pressure associated with the fluid, given by Eq. 1.4.2.3, will average to zero over many oscillations. As $p \simeq 0$, the fluid behaves, on average, as non-relativistic matter, hence the Universe expands as $a \propto t^{2/3}$, which means that $H = \frac{\dot{a}}{a}$ decreases with time as $H \propto t^{-1/3}$, ending the exponential expansion.

Reheating. In the simple inflationary model described above, the universe starts in a potential-dominated state, and inflates until the field falls enough that the kinetic energy becomes important. This stage is associated with a highly speculative process called "reheating", which averted an empty Universe after inflation. Accordingly, at the end of inflation, the inflaton begins to oscillate around its minimum and decays, releasing energy in the form of particles, which initially interacted strongly with one another, causing the Universe to rapidly heat up, averting an empty Universe after inflation.

Parametrization. Inflationary models can be parameterized by expanding the potential in a Taylor series, i.e. in increasingly higher derivatives of V, using the *slow-roll parameters*. The first two are defined as:

$$\varepsilon_V(\phi) = -\frac{\dot{H}}{H^2} = 1 - \frac{\ddot{a}a}{\dot{a}^2}, \quad \eta_V(\phi) = \frac{d\log\varepsilon}{d\log a}. \tag{1.4.2.12}$$

where $\varepsilon_V(\phi) \ll 1$ in order to have exponential expansion, and $\eta_V(\phi) \ll 1$ in order for inflation to last for a long time. For a K-G field, these correspond to

$$\varepsilon_V(\phi) = 3\frac{\dot{\phi}^2/2}{\dot{\phi}^2/2 + V} \simeq \frac{M_{\rm Pl}}{2}\left(\frac{V'}{V}\right)^2, \tag{1.4.2.13}$$

$$\eta_V(\phi) = 2\varepsilon - \frac{2\ddot{\varphi}}{H\dot{\varphi}} \simeq M_{\rm Pl}^2\frac{V''}{V}, \tag{1.4.2.14}$$

The expression above provides a straightforward interpretation: $\varepsilon_V(\phi)$ measures the slope of the potential energy density during inflation, while $\eta_V(\phi)$ measures its curvature. The slow-roll parameters therefore provide a way to parameterize different inflationary models and to compare them with observations. In particular, measurements of the CMB radiation can provide constraints on the values of these parameters and thus on the details of the inflationary epoch.

1.4.2.3 Quantifying the expansion.

The exponential expansion of the Universe is quantified by the number of *e-foldings*, denoted by N, the number of times the Universe has doubled in size. This relates to the scale factor before and after inflation as:

$$\frac{a(t_f)}{a(t_i)} = e^N \quad \Rightarrow \quad N = \int_{t_i}^{t_f} H dt, \tag{1.4.2.15}$$

where t_i and t_f are the initial and final times of inflation, respectively, and we used the fact that $N = \ln(a_f/a_i)$ to write $dN = d\ln a$, and $d\ln a = Hdt$. This result can be rewritten using the slow-roll equations 1.4.2.9-1.4.2.10,

$$Hdt = H\frac{dt}{d\phi}d\phi = H\frac{d\phi}{\dot{\phi}} = -\frac{3H^2 d\phi}{V_{,\phi}} = -\frac{1}{M_P^2}\frac{V}{V_{,\phi}}d\phi. \tag{1.4.2.16}$$

Therefore, we arrive at

$$N = -\frac{1}{M_P^2}\int_{\phi_i}^{\phi_f} \frac{V}{V_{,\phi}} d\phi. \tag{1.4.2.17}$$

Typically, a value of $N \gtrsim 40-60$ is sufficient to solve the horizon and flatness problems, with exactness dependent on the reheating temperature.

1.4.3 Perturbations in the primordial Universe

Soon after inflation was proposed, it was discovered that, besides solving the horizon and curvature problems, it could also provide a natural explanation for the origin of the small inhomogeneities we observe in an otherwise homogeneous Universe. Metric perturbations are sourced by quantum field fluctuations around the vacuum state, which are converted into "macroscopic" cosmological perturbations by the accelerated expansion. In order to calculate the fluctuations produced by metric perturbations, their power spectrum, and then their subsequent evolution, we need to study the inflaton field at the quantum level.

1.4.3.1 Inflaton quantum fluctuations

We discuss how to obtain the quantization of the free scalar field ϕ, for different geometries (perturbed or not) below. We will see that this involves the calculation of its momentum conjugate, π_ϕ, which can be obtained from the Lagrangian density $\mathcal{L}_\phi$ in Eq. 1.4.2.1 as:

$$\pi_\phi \equiv \frac{\partial \mathcal{L}_\phi}{\partial \dot{\phi}} = \dot{\phi}. \tag{1.4.3.1}$$

The unperturbed metric case. We start by allowing for spatial variations of the scalar field (inhomogeneities) of the form $\phi(\mathbf{x},\eta) = \phi(\eta) + \delta\phi(\mathbf{x},\eta)$, and assuming that the geometry is homogeneous, i.e. the FLRW metric is unperturbed. Here, we consider *perturbations that are well inside the horizon*, where the Hubble damping term can be neglected, and define how they evolve in the slow roll approximation.

We start by defining the *action* S of the scalar field, which relates to the Lagrangian density $\mathcal{L}_\phi$ as:

$$S = \int d^4x \sqrt{-g}\mathcal{L}_\phi. \tag{1.4.3.2}$$

where $\mathbf{g}$ is the metric. For a perturbed $\phi(\mathbf{x},\eta)$, this can be estimated by $S = S_0[\bar{\phi}] + S_2[\bar{\phi},\delta\phi] + \ldots$, where $S_0[\bar{\phi}]$ is the classical background action, and S_2 is the leading quadratic part due to the perturbation[29]. Given the unperturbed conformal time FLRW metric, $\sqrt{-g} = a^4$, from S_2[30] we can estimate the second-order part of the Lagrangian for the fluctuations in an expanding spacetime as:

$$\mathcal{L}_2 = \frac{a^4}{2}\partial_\mu \delta\phi \partial^\mu \delta\phi = \frac{a^2}{2}\left[\left(\dot{\delta\phi}\right)^2 - (\boldsymbol{\nabla}\delta\phi)^2\right], \tag{1.4.3.3}$$

for which $\pi_{\delta\phi}$ can be evaluated:

$$\pi_{\delta\phi} \equiv \frac{\partial \mathcal{L}_2}{\partial(\dot{\delta\phi})} = a^2\dot{\delta\phi} \tag{1.4.3.4}$$

The fluctuations will evolve following the K-G equation, defined in 1.4.2.5. In the limit described above, whereby $k \gg \mathcal{H}$, the K-G for the perturbed scalar field in Fourier space can be written as:

$$\ddot{\delta\phi} + 2\mathcal{H}\dot{\delta\phi} + k^2\delta\phi = 0 \tag{1.4.3.5}$$

where $k = |\mathbf{k}^2|$ is the comoving wave vector. Changing variables to the canonical $u_\mathbf{k} \equiv a\,\delta\phi_\mathbf{k}$, where a is the scale factor, the equation reads:

$$\ddot{u}_\mathbf{k} + \left(k^2 - \frac{\ddot{a}}{a}\right)u_\mathbf{k} = 0. \tag{1.4.3.6}$$

i.e. the so-called *Mukhanov-Sasaki equation*. On small scales, the second term in parentheses can be ignored, and the field evolves as a harmonic oscillator:

$$\delta\phi_\mathbf{k} = A_\mathbf{k}\frac{e^{\pm ik\eta}}{a}. \tag{1.4.3.7}$$

It is necessary to evaluate a physical criterion to determine the *initial condition amplitude* $A_\mathbf{k}$.

Quantisation and vacuum fluctuations. To quantise a field means that it can be represented as a set of quantum harmonic oscillators, each having a discrete set of energy levels. Even if we assume that the system is in its lowest possible energy state (or *vacuum*), there is an inherent uncertainty to what its true value is at all points in space, due to the uncertainty principle[31], i.e. $\Delta\phi(\mathbf{x})\Delta\pi(\mathbf{x}) \geq \frac{\hbar}{2}$, meaning that there are still fluctuations in the field amplitude.

[29] There is no term linear in $\delta\phi$ since it must vanish for $\phi(t)$ to minimise the action, so the leading perturbation term is obtained by Taylor expanding the potential: $S_2 = \int d^4x a^4 \frac{1}{2}\left\{\partial_\mu(\delta\phi)\partial^\mu(\delta\phi) - V_{,\phi\phi}(\bar{\phi})(\delta\phi)^2\right\}$.

[30] Effectively, we neglect $V_{,\phi\phi}$ to obtain $\mathcal{L}_2$.

[31] The uncertainty principle states that the more precisely we know the position of a particle, the less precisely we can know its momentum, and vice versa, i.e. $\Delta x\Delta p \geq \hbar/2$. This applies to fields as well, since a field can be thought of as a collection of particles, or "excitations", which can be created or destroyed, for which there is an inherent uncertainty to what their true value is at all points in space.

We can determine the minimum fluctuation amplitude, or "vacuum fluctuation", of a canonical scalar field using quantum field theory (QFT). Accordingly, $u_{\mathbf{k}}$ can be quantised[32], and the associated quantum harmonic operator $\hat{u} = a\hat{\phi}$ can be written in terms of ladder operators $\hat{b}_{\mathbf{k}}$ and $\hat{b}^{\dagger}_{\mathbf{k}}$, such that:

$$\hat{u}_{\mathbf{k}}(\eta) = u_k(\eta)\hat{b}_{\mathbf{k}} + u^*_k(\eta)\hat{b}^{\dagger}_{\mathbf{k}}, \tag{1.4.3.8}$$

where $\hat{b}$ and $\hat{b}^{\dagger}$ are the annihilation and the creation operator, respectively. Then $u_k(\eta)$ is the positive-energy solution to the K-G equation of motion in Fourier space, in Eq. 1.4.3.6:

$$u_k(\eta) = \frac{e^{-ik\eta}}{\sqrt{2k}}, \tag{1.4.3.9}$$

where the factor of $1/\sqrt{2k}$ is inserted so that the operators $\hat{b}$ and $\hat{b}^{\dagger}$ satisfy the usual harmonic oscillator canonical commutation relation $\left[\hat{b}_{\mathbf{k}}, \hat{b}^{\dagger}_{\mathbf{k}'}\right] = (2\pi)^3\delta(\mathbf{k} - \mathbf{k}')$.

Power spectrum of fluctuations. This choice of $u_k(\eta)$ above is equivalent to the choice of the vacuum state $|0\rangle$, which satisfies $\hat{b}_{\mathbf{k}}|0\rangle = 0$. The *vacuum expectation value* of the field is therefore $\langle \hat{u}_{\mathbf{k}} \rangle \equiv \langle 0|\hat{u}|0\rangle = 0$, while its variance is given by the Legendre transform of the correlation function of the field:

$$\left\langle 0 \left| \hat{u}(\eta, \mathbf{k})\hat{u}^{\dagger}(\eta, \mathbf{k}') \right| 0 \right\rangle = \frac{2\pi^2}{k^3}\mathcal{P}_u(k)\delta\left(\mathbf{k} - \mathbf{k}'\right). \tag{1.4.3.10}$$

This defines the *power spectrum* of the field u:

$$\mathcal{P}_u(k) = \frac{k^3}{2\pi^2}\left|u_k(\eta)\right|^2 \tag{1.4.3.11}$$

Since u is related to the inflaton fluctuation by $u = a\delta\phi$, the power spectrum of $\delta\phi$ is

$$\mathcal{P}_{\delta\phi}(k) = \frac{k^3}{2\pi^2}\left|\frac{u_k(\eta)}{a(\eta)}\right|^2, \tag{1.4.3.12}$$

which gives, in the limit where wavelengths are smaller than the Hubble radius ($k \gg \mathcal{H}$), one recovers $\mathcal{P}_{\delta\phi}(k) = (k/2\pi a)^2$.

In the opposite limit ($k \ll \mathcal{H}$), we can obtain an approximate result for the power spectrum. Since the field equation 1.4.3.5 is still approximately accurate if

[32] Once the canonical conjugate variables, ϕ and π, are defined, we can use them to define the corresponding quantum operators, $\hat{\phi}$ and $\hat{\pi}$, which in the quantum limit satisfy the commutation relations: $[\hat{\phi}(\mathbf{x}), \hat{\pi}(\mathbf{x}')] = i\hbar\delta(\mathbf{x} - \mathbf{x}')$. In practice, we use Fourier transform of the field, which allows us to work with the commutation relation for each mode separately. This leads to the creation and annihilation operators for each mode, which can be used to create and destroy "excitations" of the field. Therefore this allows us to quantise the field by representing it as a collection of harmonic oscillators, each with a discrete set of energy levels.

H is evaluated at horizon crossing, and assuming $\mathcal{P}_{\delta\phi}(k) \approx \mathcal{P}_{\delta\phi}(k,\eta)$, the power spectrum outside the Hubble radius is

$$\mathcal{P}_{\delta\phi}(k) \simeq \left(\frac{H}{2\pi}\right)^2 \tag{1.4.3.13}$$

Since H varies only slowly during inflation, the spectrum is *approximately scale-invariant*. Similarly, once well outside the horizon, $\delta\phi$ remains approximately constant, leading to a constant power spectrum.

The perturbed metric case. In reality, from Einstein equations, we know that perturbations in the matter distribution will be coupled to the metric perturbations. Considering metric perturbations in the Newtonian gauge (Eq. 1.2.1.7) up to second order, the perturbed action is:

$$S\left[\bar{\phi}+\delta\phi, \bar{g}_{\mu\nu}+h_{\mu\nu}\right] = S_0\left[\bar{\phi}, \bar{g}_{\mu\nu}\right] + S_1\left[\delta\phi, h_{\mu\nu}; \bar{\phi}, \bar{g}_{\mu\nu}\right] + S_2\left[\delta\phi, h_{\mu\nu}; \bar{\phi}, \bar{g}_{\mu\nu}\right] \tag{1.4.3.14}$$

Again, the first-order term S_1 vanishes when using FLRW equations of motion. The second-order term S_2 can then be used to quantise the perturbations, as for the perturbations of a scalar field described in the previous section. As described in Sec. 1.2, we avoid the description of vector modes and consider only scalar and tensor perturbations in the linear regime, which allows us to treat them separately in the following sections.

1.4.3.2 Scalar perturbations

Scalar perturbations arise from the scalar field perturbation $\delta\phi$ and the A, B, C and E modes of the metric perturbations. Similarly to the case of vacuum fluctuations, where we introduced u, it is possible to define a single variable to write the second-order action for scalar perturbations:

$$v \equiv a\left(\delta\phi - \frac{\dot{\phi}}{\mathcal{H}}C\right), \tag{1.4.3.15}$$

whose corresponding action takes the simple form:

$$S = \frac{1}{2}\int d\eta d^3x \left[\dot{v}^2 - \partial_i v \partial^i v + \frac{\ddot{z}}{z}v^2\right] \tag{1.4.3.16}$$

where we have defined $z \equiv a\dot{\phi}/\mathcal{H}$. It can be shown that v is gauge-invariant, and is proportional to the *comoving curvature perturbation* defined in Eq. 1.4.3.17, which for a single scalar field reads:

$$\mathcal{R} = -C + \frac{\mathcal{H}}{\dot{\phi}}\delta\phi. \tag{1.4.3.17}$$

This is a linear combination of scalar field and metric perturbations.

In the slow-roll approximation, ϕ and H evolve more slowly than a, therefore $\ddot{z}/z \simeq \ddot{a}/a$, and the results of the previous section obtained for u can be applied to v. Therefore the normalised function corresponding to the *Bunch-Davies vacuum* is approximately given by $v_k \simeq e^{-ik\eta}/\sqrt{2k}$. Following similar steps as above to recover Eq. 1.4.3.13, and given $v = -z\mathcal{R}$, the power spectrum for $\mathcal{R}$ is found to be:

$$\mathcal{P}_{\mathcal{R}}(k) = \frac{k^3}{2\pi^2}\left|\frac{v_k}{a}\right|^2 \overset{\text{slow-roll}}{=} \frac{1}{24\pi^2}\left(\frac{V}{m_P^4 \varepsilon_V}\right)_{k=aH}, \tag{1.4.3.18}$$

where the latter equality is valid in the slow-roll regime and for $k \leq \mathcal{H}$, and we used the inflaton potential V and the first slow-roll parameter ε_V defined in Eq. 1.4.2.13. The reduced Planck mass $m_P \equiv \hbar c/\sqrt{8\pi G}$ was introduced.

The corresponding dimensionless quantity, namely the *dimensionless primordial power spectrum* $\Delta^2_{\mathcal{R}}(k)$, is often preferred to $P_{\mathcal{R}}$, and is defined as:

$$\langle \mathcal{R}^*_{\mathbf{k}} \mathcal{R}_{\mathbf{k}'} \rangle \equiv \frac{2\pi^2}{k^3}\Delta^2_{\mathcal{R}}(k)\,(2\pi)^3 \delta^D(\mathbf{k} - \mathbf{k}'). \tag{1.4.3.19}$$

i.e. the variance of $\mathcal{R}_{\mathbf{k}}$ at the time of horizon exit (after which, $\mathcal{R}$ is constant). From this, we find the relation between $\Delta^2_{\mathcal{R}}$ and $\mathcal{P}_{\mathcal{R}}$:

$$\Delta^2_{\mathcal{R}}(k) = \frac{k^3}{2\pi^2} P_{\mathcal{R}}(k). \tag{1.4.3.20}$$

The dimensionless power spectrum represents the amount of variance in the field per logarithmic interval of wavenumber k. Combining Eqs. 1.4.3.7 and 1.4.3.17, we find

$$\Delta^2_{\mathcal{R}}(k) = \frac{k^3}{2\pi^2}\frac{H^2}{\dot{\varphi}^2}\frac{1}{2ka^2}\bigg|_{k=aH} = \frac{1}{2m_P^2 \varepsilon_V}\left(\frac{H}{2\pi}\right)^2\bigg|_{k=aH}. \tag{1.4.3.21}$$

During the slow-roll phase, the quantities H and ε_V remain nearly constant. As a result, one of the significant predictions of inflation emerges: an *almost scale-invariant primordial power spectrum*. However, the slight variations in ε_V and H towards the end of inflation introduce some level of scale dependence into the primordial power spectrum. In fact, the amplitude of a perturbation relies on the values of H and $\dot{\phi}$ at the point where its scale exits the Hubble radius, which can differ slightly for different scales. To quantify the level of scale invariance, we define the *scalar spectral index* n_s as:

$$n_s - 1 \equiv \frac{d \ln \mathcal{P}_{\mathcal{R}}(k)}{d \ln k} \overset{\text{slow-roll}}{=} -2\varepsilon_V - \eta_V, \tag{1.4.3.22}$$

and so we expect departures from $n_s = 1$ to be small[33]. This definition implies a power-law spectrum, which can be modelled as:

$$\ln \mathcal{P}_{\mathcal{R}}(k) = \ln A_s + (n_s - 1) \ln \left(\frac{k}{k_0}\right) + \frac{\alpha_s}{2}\left(\ln \left(\frac{k}{k_0}\right)\right)^2 + \ldots \tag{1.4.3.23}$$

for some pivot scale k_0, where A_s is the amplitude and α_s is the *running* of n_s, given by $\alpha_s \equiv dn_s/d\ln k$. We currently measure $n_s = 0.9665 \pm 0.0038$ and $\alpha_s = -0.0041 \pm 0.0067$ [32], which is incompatible with a perfectly scale-invariant power spectrum, and therefore compatible with an early phase of inflation.

1.4.3.3 Tensor perturbations (primordial gravitational waves)

The metric tensor perturbations h_{ij}, also known as *primordial gravitational waves* (PGWs), would provide a significant probe of the early Universe, if they are ever detected in the future. In the inflationary paradigm, these PGWs are created from vacuum quantum fluctuations, just like scalar perturbations.

Just like ϕ, the tensor perturbation h_{ij} also satisfies the Mukhanov-Sasaki equation, i.e. we can define a variable $w = ah_{ij}$, such that:

$$\ddot{w}_{\mathbf{k}} + \left(k^2 - \frac{\ddot{a}}{a}\right) w_{\mathbf{k}} = 0. \tag{1.4.3.24}$$

As for scalar perturbations, the tensor power spectrum as a function of the inflation potential V can be expressed as:

$$\mathcal{P}_h(k) = \frac{2}{3\pi^2}\left(\frac{V}{m_P^4}\right)_{k=aH}, \tag{1.4.3.25}$$

and we also define the tensor spectral index:

$$n_t \equiv \frac{d\ln \mathcal{P}_h(k)}{d\ln k} \overset{\text{slow-roll}}{=} -2\epsilon_V. \tag{1.4.3.26}$$

Thus, the primordial tensor power spectrum can be parametrised as

$$\ln \mathcal{P}_h(k) = \ln A_t + n_t \ln \left(\frac{k}{k_0}\right) + \frac{\alpha_t}{2}\left(\ln \left(\frac{k}{k_0}\right)\right)^2 + \ldots \tag{1.4.3.27}$$

where the tensor power amplitude, A_t, which is related to the scalar amplitude by the *tensor-to-scalar ratio*, r. This is defined as the ratio of scalar and tensor

[33] It is worth noting that a spectrum with $n_s \simeq 1$ ensures that the amplitude of perturbations remains small both on large scales to maintain the homogeneity of the Universe, and on small scales to uphold the achievements of big-bang nucleosynthesis.

modes power spectra, and under the slow-roll approximation is directly proportional to the slow-roll parameter ϵ_V :

$$r \equiv \frac{\mathcal{P}_h}{\mathcal{P}_\mathcal{R}} = 16\epsilon_V = -8n_t. \tag{1.4.3.28}$$

The tensor-to-scalar ratio r is of central importance for cosmology, as this is almost the unique parameter setting the amplitude of tensor modes in the primordial Universe. As detailed in the next chapter, the detection of primordial B-modes in the CMB would be a smoking gun for inflation theories.

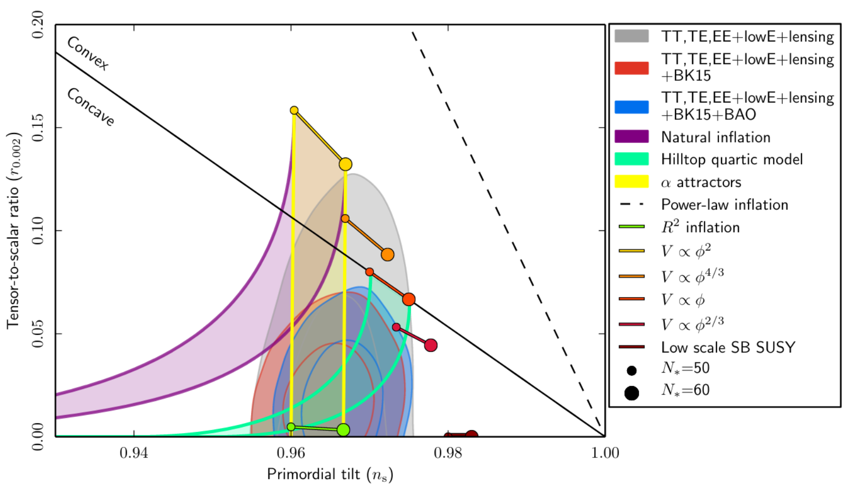

Figure 1.1: Observational constraints on n_s and r from different data sets, including Planck 2018 results (TT, EE, TE, lowE, and lensing), BICEP/Keck 2014 release (BK14) and BAO surveys such as 6dF and SDSS. The figure depicts the 95% and 68% confidence level contours. The observational constraints are compared to various inflationary models and their parameters, where V is the potential, N_* is the number of e-folds. Credit: [52].

The two parameters, the scalar spectral index n_s and the tensor-to-scalar ratio r, thus depend on inflationary models and their precise measurement allow discriminating between models. The estimation of the scalar spectral index already allowed to disfavour models that predict $n_s = 1$ as well as $n_s > 1$. The joint constraints on inflation models coming from n_s and r are shown in Fig. 1.1, and more information can be found in [52].

1.5 Cosmic Microwave Background

1.5.1 An almost perfect blackbody

The cosmic microwave background (CMB) radiation is a thermal emission that pervades the universe. It is the afterglow of the Big Bang, when the universe was hot and dense. In the FLRW background, the temperature of the CMB is assumed to be uniform in all directions, which means that the temperature anisotropies are zero. As the universe expanded and cooled, the radiation was stretched to longer wavelengths, resulting in a peak in the microwave region of the electromagnetic spectrum, and we can measure an almost perfect blackbody spectrum at a temperature of 2.725 K.

While the CMB appears to be isotropic on large scales, a more detailed examination reveals that it actually contains tiny temperature fluctuations, or anisotropies, on the order of just a few parts per million, approximately $\mathcal{O}\left(10^{-5}\right)$, which are dependent on the direction of observation. As explained in section 1.4.3, these anisotropies can be linked to the primordial fluctuations caused by metric perturbations. Essentially, the fluctuations in the metric caused changes in the energy of photons as they propagated through space, resulting in observable temperature variations in the CMB.

1.5.2 Dipole anisotropy

These CMB anisotropies can be decomposed into different modes, each corresponding to a different angular scale (parametrised by the *multipole moment* ℓ) on the sky. The zero-order mode, or *monopole* ($\ell = 0$), corresponds to the largest scales, and increasing modes correspond to smaller scales. At the first-order mode, or *dipole* ($\ell = 1$), the temperature changes relative to the monopole are relatively small and can be described by a linear relationship with direction.

If we are moving with a 3-velocity $\mathbf{v}$ relative to the CMB, then the energy of a photon detected by us at energy E and moving in direction $\hat{\mathbf{e}}$ is different from its energy in the rest frame of the CMB. Then a photon with 4-momentum $p^\mu = E(1, \hat{\mathbf{e}})$ that we detect moving along direction $\hat{\mathbf{e}}$ has energy $E = p_\mu u^\mu$, where $u^\mu = \gamma(1, -\mathbf{v})$[34] is the 4-velocity of the CMB rest-frame, i.e. the frame where there is no photon momentum density, hence dipole. The energy in the CMB frame is:

$$E = E_\gamma(1 + \mathbf{v} \cdot \hat{\mathbf{e}}) \tag{1.5.2.1}$$

where the photon's energy is $E_\gamma = h_P \nu$.

The photons are described by the Planck distribution function which, in terms of the energy and direction seen by us, becomes:

$$f\left(p^\mu\right) = \frac{1}{e^{-E_\gamma(1+\mathbf{v}\cdot\hat{\mathbf{e}})/k_B T_{\mathrm{CMB}}-1}} \tag{1.5.2.2}$$

[34] The Lorentz factor $\gamma = 1/\sqrt{1 - |\mathbf{v}|^2}$ is a quantity expressing how much the measurements of time, length, and other physical properties change for an object while that object is moving.

Although this still looks like a blackbody along any direction, as described in Sec. 1.3.1.2, the observed temperature varies over the sky as:

$$T(\hat{\mathbf{e}}) = \frac{T_{\rm CMB}}{\gamma(1+\mathbf{v}\cdot\hat{\mathbf{e}})} \approx T_{\rm CMB}(1-\mathbf{v}\cdot\hat{\mathbf{e}}) \tag{1.5.2.3}$$

This variation is a *dipole anisotropy*, induced by the Doppler term $e_i v^i$. The measured dipole indicates that the solar system barycentre has a speed of 3.68×10^5 m s^{-1} relative to the CMB. Equation 1.5.2.3 shows that relative motion produces quadrupole anisotropies and alters the monopole temperature, but only at $\mathcal{O}\left(|\mathbf{v}|^2\right)$. It is worth noting that the large dipole must be subtracted to get at the $\mathcal{O}\left(10^{-5}\right)$ anisotropies from recombination.

1.5.3 Scalar perturbations to temperature anisotropies

We shall explore the effect of metric fluctuations on the CMB in the Newtonian gauge, defined in Eq. 1.2.1.6. Scalar perturbations to the metric are sourced by density fluctuations (and vice versa). Thanks to SVT decomposition, we can focus on scalar perturbations in isolation, before treating tensor perturbations in the next section.

To fully understand the factors that cause temperature fluctuations in the CMB, it is necessary to solve the Boltzmann equation, which involves calculating the evolution of the photon distribution function, while also accounting for the effects of Thomson scattering before and during recombination [see, e.g. 26]. However, if we make two simplifying assumptions, we can still obtain a reasonably accurate explanation of temperature fluctuations. First, we assume "instantaneous recombination", which implies that the CMB photons were emitted precisely at the moment of recombination when the ionization fraction crossed a particular threshold. Although this assumption may not be entirely accurate, its impact on the results is negligible, as recombination occurs rapidly. Secondly, rather than calculating the evolution of the photon distribution function through the Boltzmann equation, we will instead measure the change in frequency of a single photon as it travels. We can estimate the temperature fluctuation as follows:

$$\frac{\delta T}{\bar{T}} = \frac{\delta \nu}{\bar{\nu}}, \tag{1.5.3.1}$$

where we have used the relationship between the frequency and temperature of a blackbody radiation spectrum, given by 1.5.2.2, where the peak frequency is proportional to the temperature as $\nu \propto k_B T/h$.

1.5.3.1 Photon propagation in a perturbed universe

The equation of motion for a photon in a given spacetime is determined by the *geodesic equation*:

$$\frac{d^2x^\mu}{d\lambda^2} + \Gamma^\mu_{\alpha\beta}\frac{dx^\alpha}{d\lambda}\frac{dx^\beta}{d\lambda} = 0\,, \tag{1.5.3.2}$$

where $x^\mu(\lambda)$ is the photon's trajectory, parametrised by the affine parameter λ, and $\Gamma^\mu_{\alpha\beta}$ are the Christoffel symbols, which can be calculated from the perturbed metric in Eq. 1.2.1.7 as:

$$\Gamma^\mu_{\alpha\beta} = \frac{1}{2}g^{\mu\lambda}\left(\partial_\alpha g_{\lambda\beta} + \partial_\beta g_{\lambda\alpha} - \partial_\lambda g_{\alpha\beta}\right). \tag{1.5.3.3}$$

where we have used the notation $\partial_\mu \equiv \partial/\partial x^\mu$.

The photon 4-momentum can be used to define λ as

$$p^\mu = dx^\mu/d\lambda, \tag{1.5.3.4}$$

and can be described in terms of the energy E and direction cosines e^i of its propagation direction[35], as seen by an observer at rest in the coordinates:

$$p^\mu = \epsilon a^{-2}[(1-\Psi),(1+\Phi)\hat{\mathbf{e}}] \tag{1.5.3.5}$$

where we have defined $\epsilon \equiv Ea$, the *comoving energy* that is constant if measured by an observer at rest in the Newtonian gauge, $E = p^\nu u_\nu$ and the observer's 4-velocity $u^\nu = (1-\Psi)\delta_{0\nu}/a$. Using the definition in Eq. 1.5.3.5 to calculate Eq. 1.5.3.4 we obtain

$$\begin{aligned} &\frac{d\eta}{d\lambda} = p^0 = \frac{\epsilon}{a^2}(1-\Psi), \quad \frac{dx^i}{d\lambda} = \frac{\epsilon}{a^2}(1+\Phi)e^i \\ \Rightarrow \quad &\frac{dx^i}{d\eta} = (1+\Phi+\Psi)e^i \end{aligned} \tag{1.5.3.6}$$

If we parametrise Eq. 1.5.3.2 with conformal η, we find that a photon in vacuum travels along the path

$$(1-\Psi)\frac{\epsilon}{a^2}\frac{dp^\mu}{d\eta} + \Gamma^\mu_{\nu\rho}p^\nu p^\rho = 0 \tag{1.5.3.7}$$

and it travels on a null curve, travels i.e. $g_{\mu\nu}p^\mu p^\nu = 0$[36]. Solving the geodesic equation for the perturbed metric to first order allows us to determine the energy

[35] $\delta_{ij}e^ie^j = 1$, and e^i corresponds to the 3-vector $\hat{\mathbf{e}}$.

[36] More generally, the 4-momentum satisfies $g_{\mu\nu}p^\mu p^\nu = -m^2$, where the left-hand side is null in the case of massless photons.

shift, which involves solving for p^0. To do so, we need to compute the components $\Gamma^0_{\alpha\beta}$, using the relations in 1.2.1.7. The only non-zero component of $g^{\mu\lambda}$ is the inverse of the time-time component g_{00}. If both lower indices in Eq. 1.5.3.3 are spatial, the first two terms in brackets vanish ($g_{0i} = 0$) and only the last term survives, and we obtain

$$\begin{aligned}\Gamma^0_{ij} &= \frac{1-2\Psi}{2}\frac{\partial}{\partial t}\left[\delta_{ij}a^2(1+2\Phi)\right] \\ &= \delta_{ij}a^2[\mathcal{H} + 2\mathcal{H}(\Phi-\Psi) + \dot{\Phi}].\end{aligned} \tag{1.5.3.8}$$

Including this and the relation in Eq. 1.5.3.6 in the 0-component of Eq. 1.5.3.7, we get

$$\begin{aligned}\frac{1}{\epsilon}\frac{d\epsilon}{d\eta} - 2\mathcal{H} - \frac{d\Psi}{d\eta} + \mathcal{H} + \dot{\Psi} + 2e^i\partial_i\Psi \\ + \delta_{ij}e^ie^j[\mathcal{H} - \dot{\Phi} - 2\mathcal{H}(\Phi+\Psi)][1+2(\Phi+\Psi)] = 0,\end{aligned} \tag{1.5.3.9}$$

where $d/d\eta$ is the derivative along the path of the photon, such that to first order:

$$\frac{d\Psi}{d\eta} = \dot{\Psi} + e^i\partial_i\Psi \tag{1.5.3.10}$$

$$\frac{de^i}{d\eta} = -\left(\delta^{ij} - e^ie^j\right)\partial_j(\phi+\psi). \tag{1.5.3.11}$$

The formula 1.5.3.11 describes *gravitational lensing*, which involves the alteration of the photon direction due to the gradient of the potential $\Phi + \Psi$ perpendicular to the line of sight. Since CMB temperature anisotropies are variations in temperature across different directions in the sky, in the background model the direction remains unchanged, meaning there are no temperature variations. However, the CMB temperature still depends on direction, but only in a first-order manner, which means that the temperature changes are small and can be approximated by a linear relationship with direction. Therefore, knowing the constant direction in the background model is enough to compute the CMB temperature anisotropies.

We can then rewrite the Eq.1.5.3.9 to find the expression for the evolution of the comoving energy along the photon path, in the presence of metric perturbations.

$$\frac{1}{\epsilon}\frac{d\epsilon}{d\eta} = -\frac{d\Psi}{d\eta} + (\dot{\Phi} + \dot{\Psi}). \tag{1.5.3.12}$$

On the right-hand side of this expression, the first term indicates that, while ϵ remains constant in the background, it is affected by changes in Ψ along the path and by the time-dependent evolution of the potentials (the second term)[37]. This can

[37]The latter term in the RHS of Eq. 1.5.3.12 becomes especially significant in the late stages of the CMB, when dark energy dominates, and around the time of last scattering, because the universe was not yet fully matter-dominated.

be easily integrated along the unperturbed trajectory, and we obtain the comoving energy today $\epsilon(\eta_0)$ in terms of the energy at conformal time η, corresponding to the point of emission on the perturbed last scattering surface:

$$\epsilon\left(\eta_0\right)=\epsilon(\eta)\left[1+\Psi(\eta)-\Psi_0+\int_\eta^{\eta_0} d\eta\left(\dot{\Psi}+\dot{\Phi}\right)\right], \tag{1.5.3.13}$$

which implies that in the unperturbed scenario, the cosmological redshift reduces the energy of photons as $E \propto 1/a$.

The frequency of a photon as observed by an observer with 4-velocity u^μ is given by $h_P\nu = p^\mu u_\mu$. When the observer has a peculiar velocity $\mathbf{v}$, their 4-velocity becomes $u^\mu = a^{-1}(1-\psi, \mathbf{v})$, and the frequency can be related to ϵ as $h_P\nu = a^{-1}\epsilon(1+\hat{\mathbf{n}}\cdot\mathbf{v})$, where $\hat{\mathbf{n}} = -\hat{\mathbf{e}}$. This leads to the perturbed *redshift relation* between the emitted (ν) and observed ($\nu_0 \equiv \nu(\eta_0)$) frequencies:

$$\frac{\nu_0}{\nu}=\frac{1}{1+z}=\frac{1}{a}\left[1+\left(\Psi-\Psi_0\right)+\hat{\mathbf{n}}\cdot\left(\mathbf{v}-\mathbf{v}_0\right)+\int_\eta^{\eta_0} d\eta\,(\dot{\Phi}+\dot{\Psi})\right], \tag{1.5.3.14}$$

where the three perturbation terms have simple physical interpretations in the Newtonian gauge:

- $\Psi(\eta)-\Psi_0$, the difference in potentials between the point of emission and reception, determines the *gravitational redshift*[38], i.e. the net redshift of a photon as it climbs out of the potential well (losing energy), or blueshift as it falls into the well (gaining energy).

- $\hat{\mathbf{n}}\cdot(\mathbf{v}-\mathbf{v}_0)$, the relative velocity between the source and the observer constitutes the standard *Doppler term*. The Doppler shift is imprinted by the high velocities of the baryons at recombination, which were still undergoing acoustic oscillations, leading to temperature fluctuations on all angular scales.

- $\int_\eta^{\eta_0} d\eta\,(\dot{\Phi}+\dot{\Psi})$ is the *Integrated Sachs-Wolfe* (ISW) effect, due to the net blue or red-shift of photons as they pass through evolving gravitational potentials along the line of sight. The gravitational potential only evolves during radiation and Λ domination, giving rise to the "early" and "late" ISW effects. In a universe dominated by matter, the gravitational potentials remain static, resulting in $\dot{\Phi}=\dot{\Psi}=0$, causing the ISW effect to disappear. This occurs because the blue and red-shifts cancel out when the photon falls into and climbs out of the potentials.

[38] Also known as the Sachs-Wolfe contribution.

Using Equation 1.5.3.1, we can relate the observed temperature fluctuation today to the fluctuation at recombination as:

$$\left.\frac{\delta T}{T}\right|_0 = \left(\frac{\delta T}{T} + \Psi - \hat{\mathbf{n}}\cdot\mathbf{v}\right)_{\rm rec} + \int_{\eta_{\rm rec}}^{\eta_0} d\eta(\dot{\Phi} + \dot{\Psi}). \tag{1.5.3.15}$$

Here we have omitted terms that depend on the gravitational potential and the peculiar velocity at the observer, as they only contribute to the monopole and dipole term. Furthermore, we can use the Stefan-Boltzmann law $\rho_\gamma \propto T^4$ to relate the temperature fluctuation at recombination to the radiation overdensity as $\delta T/T = \delta_\gamma/4$. Combining these results, we arrive at the following expression:

$$\frac{\delta T}{T}(\hat{\mathbf{n}}) = \left(\frac{\delta_\gamma}{4} + \Psi - \hat{\mathbf{n}}\cdot\mathbf{v}\right)(\eta_{\rm rec}, \chi_{\rm rec}\hat{\mathbf{n}}) + \int_{\eta_{\rm rec}}^{\eta_0} d\eta\,(\dot{\Phi} + \dot{\Psi})(\eta, \chi\hat{\mathbf{n}}), \tag{1.5.3.16}$$

where $\chi = \eta_0 - \eta$ is the comoving distance along the unperturbed photon trajectory.

1.5.3.2 Angular power spectrum

The temperature fluctuation $\delta T/T$ is defined on a sphere, and its decomposition is best achieved using spherical harmonics:

$$\frac{\delta T}{T}(\hat{\mathbf{n}}) = \sum_{\ell m} a_{\ell m}\, Y_{\ell m}(\hat{\mathbf{n}}), \quad a_{\ell m} = \int d\hat{\mathbf{n}}\, \frac{\delta T}{T}(\hat{\mathbf{n}})\, Y^*_{\ell m}(\hat{\mathbf{n}}). \tag{1.5.3.17}$$

For a statistically isotropic field, such as $\delta T/T$, the angular power spectrum C_ℓ is defined in a manner similar to the power spectrum $\mathcal{P}(k)$:

$$\langle a_{\ell m} a^*_{\ell' m'}\rangle = C_\ell\, \delta_{\ell\ell'}\delta_{mm'}. \tag{1.5.3.18}$$

Our objective is to establish a connection between the C_ℓ of the CMB and the primordial power spectrum $\Delta^2_{\mathcal{R}}(k)$, as defined in Eq. 1.4.3.21. We examine a quantity $f(\hat{\mathbf{n}})$ defined on the sphere, which is obtained through integrating along a radial lightcone over a three-dimensional quantity $F(\mathbf{x}, \eta)$:

$$f(\hat{\mathbf{n}}) = \int^{\eta_0} d\eta\, q_f(\chi) F(\chi\hat{\mathbf{n}}, \eta), \tag{1.5.3.19}$$

where $q_f(\chi)$ is a radial kernel. If F is linearly related to the primordial curvature perturbation $\mathcal{R}_\mathbf{k}$ by a transfer function $T_F(k, \eta)$, where k is the wave number and η is the conformal time:

$$F_\mathbf{k}(\eta) = T_F(k, \eta)\, \mathcal{R}_\mathbf{k}. \tag{1.5.3.20}$$

then the harmonic coefficients of f are related to $\mathcal{R}_\mathbf{k}$ through:

$$f_{\ell m} = \int \frac{dk}{2\pi^2} k^2\, \Delta^f_\ell(k) i^\ell \int d\hat{\mathbf{n}}_\mathbf{k}\, Y^*_{\ell m}(\hat{\mathbf{n}}_\mathbf{k})\, \mathcal{R}_\mathbf{k}, \tag{1.5.3.21}$$

where $\hat{\mathbf{n}}_{\mathbf{k}} \equiv \mathbf{k}/k$, and

$$\Delta_\ell^f(k) = \int d\chi \, q_v(\chi) \, T_F(k,\eta) \, j_\ell(k\chi). \tag{1.5.3.22}$$

We can demonstrate this by decomposing F into its Fourier coefficients and using the plane-wave expansion:

$$e^{i\mathbf{k}\cdot\mathbf{x}} = 4\pi \sum_{\ell m} i^\ell \, j_\ell(k\chi) \, Y_{\ell m}(\hat{\mathbf{n}}) \, Y^*_{\ell m}(\hat{\mathbf{n}}_{\mathbf{k}}), \tag{1.5.3.23}$$

where $j_\ell(x)$ is the spherical Bessel function of order ℓ. To obtain the angular power spectrum of f, we square the coefficients $f_{\ell m}$ and compute the average over multiple realizations of the primordial metric fluctuations. Using Eq. 1.4.3.19 for the primordial power spectrum, we arrive at the following expression:

$$C_\ell^f = 4\pi \int \frac{dk}{k} \, |\Delta_\ell^f(k)|^2 \, \Delta^2_{\mathcal{R}}(k). \tag{1.5.3.24}$$

Generally, a projected quantity like the temperature fluctuation in the CMB can be a sum of multiple contributions, denoted as $f(\hat{\mathbf{n}}) = \sum_\alpha f_\beta(\hat{\mathbf{n}})$, and the previous expression can be easily extended to:

$$C_\ell^f = 4\pi \int \frac{dk}{k} \left(\sum_{\alpha\beta} \Delta_\ell^\alpha(k) \Delta_\ell^\beta(k) \right) \Delta^2_{\mathcal{R}}(k). \tag{1.5.3.25}$$

It is important to note that while the above analysis is applicable to the δ_γ, Sachs-Wolfe, and ISW terms in equation Eq. 1.5.3.16, we need to adjust the definition of Δ_ℓ^f for Doppler-like terms with the following form:

$$g(\hat{\mathbf{n}}) = \int^{\eta_o} d\eta \, q_g(\chi) \, \hat{\mathbf{n}} \cdot \mathbf{v}. \tag{1.5.3.26}$$

In this case, the angular transfer function reads:

$$\Delta_\ell^g(k) = -\int d\chi \, q_g(\chi) \, \chi \, T_\theta(k,\eta) \frac{j'_\ell(k\chi)}{k\chi}, \tag{1.5.3.27}$$

where T_θ is the transfer function for the velocity divergence, and $j'_\ell(x) \equiv dj_\ell(x)/dx$.

1.5.4 Tensor perturbations to polarisation anisotropies

As we have seen in Sec. 1.4.3, other than scalar perturbations, there are other types of gravitational perturbations in the early Universe, in particular tensor perturbations or primordial gravitational waves (PGWs), described in Sec. 1.4.3.3. These imprint specific patterns in the CMB polarisation, which can be decomposed into two distinct modes: E-modes, which arise from scalar perturbations, and B-modes, which arise from tensor perturbations. Here we discuss how these polarisation anisotropies are generated. The information presented in this section draws from [26, 53–56].

1.5.4.1 Origin of CMB polarisation

The CMB polarisation arises when the primordial photons are scattered by hot electrons on their way to us. *Thomson scattering* is the process responsible for this phenomenon, whereby incoming radiation is absorbed and re-emitted by an electron in a different direction. The scattered radiation is polarised in a direction perpendicular to the original polarisation, with the intensity of the scattered radiation peaking in the direction normal to the incident polarisation.

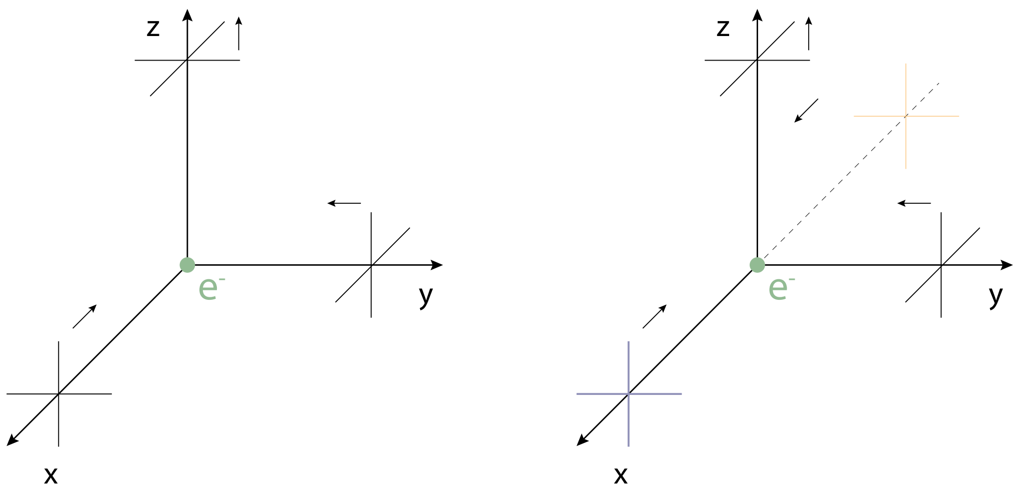

Figure 1.2: Thomson scattering of incoming isotropic (left) and dipole (right) radiation. The black lines in the diagram represent radiation with an average intensity, whereas the violet (thicker) and orange (thinner) lines correspond to radiation that is hotter or colder than average, respectively. Despite these variations in intensity, the net effect of Thomson scattering is the emission of outgoing unpolarised light, which holds true for both isotropic and dipole radiation.

Suppose a plane wave of light is moving towards the origin along the x-axis, with electric and magnetic fields oscillating in the y-z plane. If the intensity in the transverse directions (y, z) is equal, the light is unpolarised. For example, consider an electron being bombarded by unpolarised isotropic radiation (monopole) from all directions. Figure 1.2 represents a simplified version of this scenario, where only two directions ($+\hat{x}$ and $+\hat{y}$) are shown. Since the incoming amplitudes from both directions have equal intensity, the outgoing radiation along the x and y-axes will have equal intensities, resulting in unpolarised outgoing radiation along the $\hat{z}$ direction. Therefore, *isotropic radiation does not produce polarisation.*

Incoming dipole radiation also does not produce polarisation. For example, when incoming radiation is hotter (colder) than average from the $+x(+y)$ direction and colder than average from the $-x(-y)$ direction, the outgoing intensity along the x-axis comes from the ±y incident radiations that have an average temperature.

The two rays from the $\pm x$ directions also produce an outgoing intensity that is neither hot nor cold along the y-axis because it is just equal to the average intensity of the incoming rays in along the $\pm x$ direction. Since these have the same average intensity, the intensity of the outgoing wave along the x and y axes are equal, resulting in outgoing unpolarised radiation.

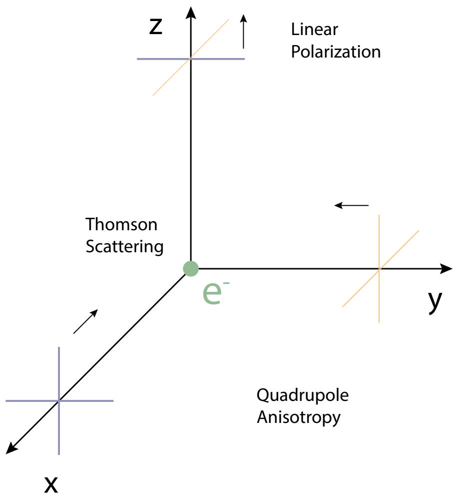

Figure 1.3: When radiation undergoes Thomson scattering with a quadrupole anisotropy, it results in linear polarisation. The thicker violet lines correspond to hot radiation, while the thinner orange lines indicate cold radiation.

Polarised radiation is produced when the incident radiation field has a quadrupole variation in intensity (i.e. temperature). This can happen when the average incoming radiation from one direction is hotter or colder than it is from another direction (see Fig. 1.3 for an example). Thomson scattering can generate polarisation in the cosmic background radiation if there is a quadrupole anisotropy of the photon flux at one point on the last scattering surface. However, this can only occur during a short period around recombination[39], so the amplitude of the polarisation is expected to be smaller than the temperature anisotropies. Late re-ionisation can enhance the polarisation at large scales, but the polarisation signal is still expected to be small.

1.5.4.2 Formalism

After explaining how Thomson scattering generates polarisation, our goal is now to develop a formula that describes the Q and U components of polarisation for the scattered radiation by summing over all incident rays. To do this, we first define the polarisation vectors for the outgoing beam of light. Specifically, we

[39]Prior to recombination, the density of electrons is so high that the mean free path of photons is too small to generate a quadrupole moment. Conversely, after recombination, the electron density is too low to result in significant Thomson scattering.

let $\hat{\epsilon}_1$ be perpendicular to the scattering plane, and $\hat{\epsilon}_2$ be in the scattering plane. We can use the Thomson scattering cross-section to describe the process of an incident wave with linear polarisation $\hat{\epsilon}'$ being scattered into a wave with linear polarisation $\hat{\epsilon}$. This cross-section is given by

$$\frac{d\sigma}{d\Omega_{n'}} = \frac{3\sigma_T}{8\pi} \left|\hat{\epsilon}' \cdot \hat{\epsilon}\right|^2 \tag{1.5.4.1}$$

Here, σ_T is the total Thomson cross-section, and $d\Omega_{n'}$ is the differential solid angle in the direction of the incoming ray $\hat{n}'$.

Let's begin with an unpolarised incident plane wave with intensity I'. This means that the Q', U', and V' polarisations of the incident wave are all zero. The Q polarisation of the outgoing wave can be obtained by taking the difference between the cross-sections for photons polarised in two different directions, $\hat{\epsilon}_1$ and $\hat{\epsilon}_2$. Using this, the following equation is obtained [54]:

$$Q = \frac{3\sigma_T}{8\pi} \left(\sum_{j=1}^{2} \left|\hat{\epsilon}_1 \cdot \hat{\epsilon}'_j\left(\hat{n}'\right)\right|^2 - \sum_{j=1}^{2} \left|\hat{\epsilon}_2 \cdot \hat{\epsilon}'_j\left(\hat{n}'\right)\right|^2 \right) \tag{1.5.4.2}$$

where $\hat{n}'$ is the direction of the scattered radiation, and $\hat{\epsilon}'_j$ represents the polarisation vector of the scattered radiation. To simplify things, we can assume that the $\hat{\hat{z}}$ axis is along the propagation direction of the scattered light. In this case, we can choose the two outgoing polarisation axes to be $\hat{\epsilon}_1 = \hat{x}$ and $\hat{\epsilon}_2 = \hat{y}$. By integrating over all directions of $\hat{n}'$, we arrive at the following expression:

$$\begin{aligned} Q(\hat{z}) &= \frac{3\sigma_T}{8\pi} \int d\Omega_{n'} I'\left(\hat{n}'\right) \sum_{j=1}^{2} \left(\left|\hat{x} \cdot \hat{\epsilon}'_j\left(\hat{n}'\right)\right|^2 - \left|\hat{y} \cdot \hat{\epsilon}'_j\left(\hat{n}'\right)\right|^2 \right) && (1.5.4.3) \\ &= -\frac{3\sigma_T}{8\pi} \int d\Omega_{n'} I'\left(\hat{n}'\right) \sin^2\theta' \cos\left(2\phi'\right). && (1.5.4.4) \end{aligned}$$

where in the last equality the dot products have been simplified by expressing $\hat{\epsilon}'_1$ and $\hat{\epsilon}'_2$ in terms of Cartesian coordinates, i.e. $\hat{\epsilon}'_1\left(\theta', \phi'\right) = \left(\cos\theta' \cos\phi', \cos\theta' \sin\phi', -\sin\theta'\right)$, $\hat{\epsilon}'_2\left(\theta', \phi'\right) = \left(-\sin\phi', \cos\phi', 0\right)$.

Similarly, we can obtain the expression for $U(\hat{z})$. The polarisation's U-component is directly proportional to the difference between the cross-sections for outgoing photons polarised in the $(\hat{x} + \hat{y})/\sqrt{2}$ and $(\hat{x} - \hat{y})/\sqrt{2}$ directions. Consequently, the U-component is derived as follows:

$$\begin{aligned} U(\hat{z}) =& \frac{3\sigma_T}{8\pi} \int d\Omega_{n'} I'\left(\hat{n}'\right) \sum_{j=1}^{2} \left(\left|\left(\frac{\hat{x}+\hat{y}}{\sqrt{2}}\right) \cdot \hat{\epsilon}'_j\left(\hat{n}'\right)\right|^2 - \left|\left(\frac{\hat{x}-\hat{y}}{\sqrt{2}}\right) \cdot \hat{\epsilon}'_j\left(\hat{n}'\right)\right|^2 \right) \\ =& -\frac{3\sigma_T}{8\pi} \int d\Omega_{n'} I'\left(\hat{n}'\right) \sin^2\theta' \sin\left(2\phi'\right). \end{aligned} \tag{1.5.4.5}$$

The polarisation in the $\hat{z}$-direction will then have the form

$$Q(\hat{z}) - iU(\hat{z}) = -\frac{3\sigma_T}{8\pi} \int d\Omega_{n'} \sin^2\theta' e^{2i\phi'} I'(\theta', \phi') \tag{1.5.4.6}$$

We can further expand the incident intensity by spherical harmonics

$$I'(\theta', \phi') = \sum_{\ell m} a'_{\ell m} Y_{\ell m}(\theta', \phi') \tag{1.5.4.7}$$

Using the orthonormality of the $Y_{\ell m}$ and considering that the integrand in Eq. 1.5.4.6 is proportional to $Y_{22}(\theta', \phi') = \sqrt{15/32\pi} e^{2i\phi'} \sin^2\theta'$, Eqs. 1.5.4.6-1.5.4.7 lead to

$$(Q - iU)(\hat{z}) = \frac{3\sigma_T}{2\pi} \sqrt{\frac{2\pi}{15}} a'_{22} \tag{1.5.4.8}$$

Thus, *polarisation is generated along the outgoing $\hat{z}$-axis provided that the a'_{22} quadrupole moment of the incoming radiation is non-zero.* To determine the outgoing polarisation in a direction making an angle β with the z-axis, the same incoming radiation $I'(\hat{z})$ field must be expanded in the new coordinate system

$$I'(\hat{z}) = \tilde{I}'(\hat{n}_{rot}) = \sum_{\ell m} \tilde{a}'_{\ell m} Y_{\ell m}(\hat{n}_{rot}), \tag{1.5.4.9}$$

where $\tilde{.}$ refers to the quantities in the frame $\hat{n}_{\text{rot}} = R(\beta)\hat{z}$. The rotated multipole coefficients are

$$\tilde{a}'_{\ell m} = \int d\hat{n}_{rot} Y^*_{\ell m}(\hat{n}_{\text{rot}}) \tilde{I}'(\hat{n}) = \sum_{m'=-m}^{m} \mathcal{D}^{\ell *}_{mm'}(R(\beta)) a'_{\ell m'}, \tag{1.5.4.10}$$

where $\mathcal{D}^{\ell}_{mm'}$ is the Wigner D-symbol. We see that only a^*_{2m} components of incident radiation contribute to $\tilde{a}'_{22}$ which generates the polarisation in the $\hat{n}_{\text{rot}}$ direction. Hence the polarisation in the new direction is

$$(Q - iU)(\hat{n}_{rot}) = \frac{3\sigma_T}{2\pi} \sqrt{\frac{2\pi}{15}} \sum_{m=-2}^{2} \mathcal{D}^{2*}_{2m}(\hat{n}_{rot}) a'_{2m} \tag{1.5.4.11}$$

Thus, *the total scattered radiation in any direction of sky is polarised only if there exists a non-zero quadrupole moment in the incident radiation field.*

In the previous passage, we derived the polarisation state $(Q - iU)$ of the scattered radiation by first obtaining the polarisation in the $\hat{z}$ direction in Eq. 1.5.4.8 and then using the Wigner $\mathcal{D}$ and the rotation matrix to find the general spatial dependence for polarisation in a general direction $\hat{n}$ (Eq. 1.5.4.11).

For a given incident radiation I', it is possible to show [see, e.g. 53] that a scalar perturbation would induce a variation with the following angular dependence:

$$\delta I(\theta, \varphi) \propto \cos^2\theta - \frac{1}{3}, \tag{1.5.4.12}$$

thus the incident radiation induces only $m = 0$ contributions to the sum in Eq. 1.5.4.11. Thus, only E modes are present in the spherical harmonic transform of the polarisation. The angular dependence caused by a tensor wave can instead be derived as:

$$\delta I(\theta, \varphi) \propto \left(1 - \cos^2\theta\right)\cos 2\varphi. \tag{1.5.4.13}$$

The incident radiation induces $m \neq 0$ modes in Eq. 1.5.4.11, due to the dependency on φ in the Equation above, resulting in both E- and B-modes in the spherical harmonic transform of the polarisation. For the full derivation and further discussion, refer to [53, 56].

1.5.5 Secondary anisotropies

The CMB also carries information about the large-scale structure (LSS) of the Universe and the physical processes that occurred after the CMB was emitted. As the CMB radiation propagates, it interacts with the matter and radiation it encounters, leaving imprints on the CMB that are known as *secondary anisotropies*. Secondary anisotropies can arise from a variety of physical processes, both gravitational and scattering in nature, and can provide valuable information about the distribution of matter in the universe, the properties of dark matter and dark energy, and the history of the universe from the time of the CMB emission to the present day. Since the work presented in this thesis mainly focuses on CMB primary anisotropies, and in particular B-mode polarisation, we only briefly describe the effects that induce secondary anisotropies in what follows.

1.5.5.1 Gravitational lensing

The LSS in the Universe can bend the path of CMB photons as they travel towards us. This effect is called *gravitational lensing*. This effect does not directly modify the intensity of the CMB, but it leads to a reshuffling of its observed angular coordinates, causing CMB patterns to appear distorted. For polarisation, lensing can convert E-modes to B-modes, giving rise to the so-called "lensing B-modes". As intensity is conserved, the induced changes in the statistical properties of the CMB fluctuations are second-order effects. Nonetheless, they can be detected with high significance in CMB data, making gravitational lensing a powerful tool for studying the LSS and evolution of the Universe. For more information, refer to [e.g. 26, 57, 58].

On the CMB temperature fluctuations, gravitational lensing remaps the unlensed temperature distribution. The lensed fluctuations, $\delta T'$, relate to the unlensed, δT, via the lensing potential ϕ_L, which depends on the gravitational potential along the line of sight. The displacement associated with the gradient of ϕ changes the photon paths, and a first-order Taylor expansion yields the expression:

$$\delta T'(\boldsymbol{\theta}) = \delta T(\boldsymbol{\theta} - \delta\boldsymbol{\theta}) \simeq \delta T(\boldsymbol{\theta}) - \nabla_\theta \phi_L(\boldsymbol{\theta}) \cdot \nabla_\theta \delta T(\boldsymbol{\theta}), \tag{1.5.5.1}$$

which in Fourier space, for an angular wave vector $\mathbf{l}$, reads:

$$\delta T'_{\mathbf{l}} = \delta T_{\mathbf{l}} + \int \frac{d^2\mathbf{l}'}{(2\pi)^2} \mathbf{l}' \cdot (\mathbf{l} - \mathbf{l}') \, \Phi_{L,\mathbf{l}'} \, \delta T_{\mathbf{l}-\mathbf{l}'}. \tag{1.5.5.2}$$

Since the LSS that contributes to the lensing potential ϕ_L is mostly at low redshifts ($z \lesssim 10$), while CMB photons were emitted at higher redshifts ($z \sim 1100$), then ϕ_L and δT are not strongly correlated with each other. However, the lensing effect can induce correlations between different Fourier modes of the temperature field. In fact, for a fixed ϕ_L, the second term in Eq. 1.5.5.1 breaks the statistical isotropy of the unlensed temperature fluctuations, inducing correlations between different Fourier modes. The correlator of two lensed temperature modes with different wave vectors $\mathbf{l} \neq \mathbf{l}'$ can be expanded to lowest order in ϕ_L as follows:

$$\langle \delta T'_{\mathbf{l}} \delta T'^{*}_{\mathbf{l}'} \rangle_{\phi_L} = (\mathbf{l}\, C^T_\ell - \mathbf{l}' C^T_{\ell'}) \cdot \Phi_{L,\mathbf{l}-\mathbf{l}'} \, (\mathbf{l} - \mathbf{l}'), \tag{1.5.5.3}$$

where the notation $\langle \cdots \rangle_{\Phi_L}$ denotes ensemble averaging for a fixed ϕ_L. To estimate $\Phi_{L,\mathbf{L}}$ for a given mode $\mathbf{L}$, the products of lensed temperature modes $\delta T'_{\mathbf{l}} \delta T'^{*}_{\mathbf{l}'}$ can be combined such that $\mathbf{l} - \mathbf{l}' = \mathbf{L}$, using the *quadratic estimator*:

$$\hat{\Phi}_{L,\mathbf{L}} = \int \frac{d^2\mathbf{l}}{(2\pi)^2} \delta T_{\mathbf{l}} \, \delta T_{\mathbf{l}-\mathbf{L}} \, g(\mathbf{l}, \mathbf{L}). \tag{1.5.5.4}$$

The optimal kernel $g(\mathbf{l}, \mathbf{L})$ can be determined by minimising the variance of the resulting lensing potential map [58]. Lensing reconstruction produces maps of the lensing convergence at the surface of last scattering, which contain information about matter density fluctuations integrated along the line of sight from $z = 1100$. By cross-correlating these maps with other tracers of matter fluctuations at specific redshifts, we can study the growth history from recombination and recover the redshift dependence of structure growth. This is a powerful tool for studying the nature of dark energy and modified gravity.

1.5.5.2 Scattering effects

Sunyaev-Zel'dovich (SZ) effect. The SZ effect describes the scattering of CMB photons by ionised gas in galaxy clusters. There are two components of the SZ effect:

1. The *thermal SZ effect* (tSZ) is caused by the random motion of electrons in the hot gas of galaxy clusters, resulting in a distortion of the CMB spectrum away from a perfect blackbody. This effect can be measured by analyzing the temperature of the CMB at different frequencies, which allows for the separation of the SZ signal from other astrophysical sources.

2. The *kinematic SZ effect* (kSZ) is due to the bulk motion of the gas in galaxy clusters along the line of sight, causing a shift in the temperature of the CMB. This effect does not modify the shape of the CMB spectrum, but it can still be detected by analysing the temperature fluctuations in the CMB.

Both the thermal and kinematic SZ effects are important probes of the properties of gas in galaxy clusters and the large-scale structure of the Universe. Ongoing and planned experiments are dedicated to measuring these effects with high precision, which will provide valuable insights into the nature of the Universe's evolution.

Integrated Sachs-Wolfe (ISW) effect. This effect arises from the fluctuating gravitational potential of the LSS. As photons pass through regions of varying density and gravitational potential, their energy is affected. In the presence of accelerating expansion of the universe, photons travelling through less dense regions lose energy, while those passing through denser regions gain energy. This results in a small temperature shift in the CMB.

Reionization. This refers to the process by which the neutral hydrogen gas in the early universe was ionized by the first stars and galaxies that formed. Reionization can be thought of as the opposite of recombination, when the ionized gas that resulted from the Big Bang cooled and recombined into neutral hydrogen gas.

Reionization occurred at a redshift $z_{\mathrm{reio.}} \approx 8$, when the first generations of galaxies began to radiate ultraviolet and X-ray photons that heated and ionized their surroundings and eventually the entire intergalactic medium. As for recombination (see Eq. 1.3.4.5), we can define the *optical depth to reionization*, denoted by $\tau_{\mathrm{reio.}}$, which measures the probability that a photon travelling through the Universe will interact with an electron before reaching us. It is defined as a line-of-sight integral along the free electron fraction from reionization, $X_e(z)$:

$$\tau_{\text{reio.}} \equiv n_H \sigma_T \int_0^{z_{\max}} dz X_e(z) \frac{(1+z)^2}{H(z)}, \tag{1.5.5.5}$$

where n_H is the total number density of hydrogen nuclei today.

The effect of reionization on the CMB temperature power spectrum is a suppression of power at small scales (large multipoles) due to the scattering of CMB photons by the ionized gas. This suppression factor is given by $e^{-2\tau_{\text{reio.}}}$. In addition, reionization creates a bump in the polarisation power spectrum at large scales (small multipoles), which scales in power as $\tau^2_{\text{reio.}}$. At scales $\ell \gtrsim 20$, the effect of reionization is heavily correlated with the A_s parameter, which determines the overall amplitude of the primordial power spectrum. Therefore, precise measurements of the E-mode polarisation are crucial for improving our constraints on $\tau_{\text{reio.}}$.

1.6 Physical Significance of PGW discovery

Should the predictions of the leading models concerning the Hot Big Bang's origin prove accurate, CMB polarisation experiments could uncover the mark of primordial gravitational waves (PGWs) in the CMB's polarisation pattern. This finding carries enormous potential for opening up a new avenue to the primordial universe, providing evidence for the quantisation of gravity, uncovering new physics at the level of grand unified theories, and offering insights into nature's symmetries.

Constrain models of inflation or rule it out. Simple inflationary models have the ability to create measurable tensor perturbations. The current upper limit on the tensor-to-scalar ratio r has already ruled out certain slow-roll inflation models [59]. However, there is a strong motivation [32, 52] to further lower the current limits as certain large-field plateau models predict a non-negligible value of $r \approx 0.003$ [60]. The forthcoming generation of CMB experiments aims to either detect or exclude a signal at the $r = 0.001$ level at a few σ significance using the B-mode amplitude. A signal at this level would provide evidence for certain inflationary models [61, 62], such as $r \propto 1/N^2$ models (e.g. Higgs or R^2 inflation) [63, 64].

Conversely, cosmologies that do not rely on inflationary models propose a bounce as an alternative to the big bang singularity, which would result in a seamless transition from contraction to expansion [65]. The gradual contraction phase preceding the bounce would smooth and flatten the universe, generating nearly scale-invariant super-horizon perturbations of quantum origin that would eventually give rise to the structure in the post-bounce universe [66]. Detectable tensor perturbations are not expected in these models, meaning that a detection of primordial B-modes would serve as evidence against them.

Besides the tensor-to-scalar ratio, exploring the large scale polarisation signal can be utilised to investigate the constraints on the tensor tilt, n_T. These can be used to test consistency of different inflation models, whereby simple inflationary models predict $r = -8n_T$.

Non-standard correlations. The information could be used to detect the existence of non-standard correlations such as non-zero TB and EB, produced by early or late occurrences, such as chiral gravitational waves and Faraday rotation [67–71].

Energy scale of the early Universe. For slow roll inflation, $\dot{\phi}^2 \ll V$ so Eqns. 1.4.2.9,1.4.2.10,1.4.2.4,1.4.3.25,1.4.3.28 together lead to a relation between the energy scale of inflation $V^{1/4}$ and the tensor-to-scalar ration r:

$$V^{1/4} \approx \left(\frac{3\pi^2}{2} r \mathcal{P}_\mathrm{s}\right)^{1/4} M_\mathrm{Pl} \approx \left(\frac{r}{0.001}\right)^{1/4} 10^{16}\mathrm{GeV} \tag{1.6.0.1}$$

where in the last equality we have used the fiducial values at $r = 0.001, k_0 = 0.05\mathrm{Mpc}^{-1}$ and $\ln(10^{10} A_s) = 3.089$ (the value given in [52]), $M_\mathrm{Pl} = 2.43 \times 10^{18}\mathrm{GeV}$. Thus, the energy scale observed during inflation aligns with the energy level predicted by Grand Unified Theories (GUTs), only a few orders of magnitude below the Planck scale. This has enormous implications for high-energy particle physics. At present, the unification of gauge couplings and the lower limits on proton lifetime provide the only hints of physics at such high energies, which are beyond the capabilities of human-made ground particle colliders [72].

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

— Richard Feynman

Chapter 2

CMB Observations

Contents

2.1 CMB Observables

In the following Chapter, we will explore the temperature and polarisation anisotropies as the two primary observables of the Cosmic Microwave Background (CMB). We will specifically focus on the unique indicators of primordial gravitational waves (PGWs) in these observables and highlight why polarisation presents a highly encouraging approach for detecting them.

2.1.1 Stokes parameters and coherency matrix

The polarisation of the CMB can be described statistically in terms of the two components of an electromagnetic wave with frequency ν propagating in the z-direction.

$$E_x(\mathbf{n}) = A_x(t)\cos\left(\nu t - \theta_x(t)\right) \tag{2.1.1.1}$$

$$E_y(\mathbf{n}) = A_y(t)\cos\left(\nu t - \theta_y(t)\right). \tag{2.1.1.2}$$

where A_x, A_y are complex amplitudes of the transverse electric field $\vec{E}$. The polarisation of the CMB arises when these two components are correlated. All the statistical information is encoded in the *coherency matrix*:

$$C = \begin{pmatrix} \left\langle |A_x|^2 \right\rangle & \left\langle A_x A_y^* \right\rangle \\ \left\langle A_x^* A_y \right\rangle & \left\langle |A_y|^2 \right\rangle \end{pmatrix} = \frac{1}{2}\begin{pmatrix} I+Q & U-\mathrm{i}V \\ U+\mathrm{i}V & I-Q \end{pmatrix}, \tag{2.1.1.3}$$

and I, Q, U, V are the *Stokes parameters.* These are defined locally by time-averaging over the amplitude and phases of the electromagnetic wave in each direction $\mathbf{n}$:

$$I = \left\langle A_x^2 \right\rangle + \left\langle A_y^2 \right\rangle \tag{2.1.1.4}$$

$$Q = \left\langle A_x^2 \right\rangle - \left\langle A_y^2 \right\rangle \tag{2.1.1.5}$$

$$U = \left\langle 2A_x A_y \cos\left(\theta_x - \theta_y\right) \right\rangle \tag{2.1.1.6}$$

$$V = \left\langle 2A_x A_y \sin\left(\theta_x - \theta_y\right) \right\rangle \tag{2.1.1.7}$$

where $\langle \cdots \rangle$ represent expectation values. These parameters describe the intensity I (or temperature T) and polarisation, both linear (Q, U) and circular (V), of the CMB in each direction $\mathbf{n}$ in the sky.

2.1.2 Spherical and spin-2 harmonics decomposition

As we measure temperature and polarisation on the sky, which is the surface of a sphere, we need to use the spherical harmonics and spin-2 tensor harmonics to decompose them.

2.1.2.1 Temperature

The temperature anisotropy of the CMB is a scalar field, so it is invariant under rotation and has zero parity (*spin-0*). Hence it can be expanded using the standard plane-wave decomposition (*Fourier transform*) on the celestial sphere:

$$T(\hat{n}) = \sum_{\ell m} a^T_{\ell m} Y_{\ell m}(\hat{n}) \tag{2.1.2.1}$$

where $Y_{\ell m}(\hat{n})$ are the spherical harmonics, and $\hat{n} = (\theta, \phi)$ is the unit vector along the line of sight, corresponding to a given pixel. Since $T(\hat{n})$ is real, the decomposition coefficients have to satisfy the reality condition $a^*_{\ell m} = (-1)^m a_{\ell -m}$.

2.1.2.2 Polarisation

In the case of polarisation, it is convenient to work with a complex linear combination $Q+iU$. For counter-clock wise rotations by an angle ψ,

$$(Q \pm iU)'(\hat{n}) = e^{\mp 2i\psi}(Q \pm iU)(\hat{n}) \tag{2.1.2.2}$$

This property shows that $Q \pm iU$ has the spin of ± 2 (or ∓ 2, depending on one's definition). This means we cannot decompose $Q \pm iU$ into the ordinary plane-wave decomposition, we need the *spin-2 decomposition* [73, 74]. Expanding in spin-weighted harmonics ${}_sY_{lm}(\hat{n})$ (see Appendix A.3 for more details), we get the spin-2 quantities

$$(Q \pm iU)(\hat{n}) = \sum_{\ell m} a_{\pm 2,\ell m \pm 2}Y_{\ell m}(\hat{n}) \tag{2.1.2.3}$$

where $a^*_{-2,\ell m} = a_{2,\ell -m}$. For Q, U, we can use spin raising/lowering operators ($\eth\,/\,\bar{\eth}$) to obtain spin-0 quantities, similarly to T:

$$\bar{\eth}^2(Q+iU)(\hat{n}) = \sum_{\ell m}\left[\frac{(\ell+2)!}{(\ell-2)!}\right]^{1/2} a_{2,\ell m}Y_{\ell m}(\hat{n}) \tag{2.1.2.4}$$

$$\eth^2(Q-iU)(\hat{n}) = \sum_{\ell m}\left[\frac{(\ell+2)!}{(\ell-2)!}\right]^{1/2} a_{2,\ell m}Y_{\ell m}(\hat{n}) \tag{2.1.2.5}$$

where $\eth = -\left(\partial_x + i\partial_y\right)$, and $\bar{\eth} = -\left(\partial_x - i\partial_y\right)$.

Combining equations 2.1.2.3, 2.1.2.4 we obtain the expansion coefficients:

$$\begin{aligned} a_{2,\ell m} &= \int d\Omega\, {}_2Y^*_{\ell m}(\hat{n})(Q+iU)(\hat{n}) \\ &= \left[\frac{(\ell+2)!}{(\ell-2)!}\right]^{-1/2} \int d\Omega Y^*_{\ell m}(\hat{n})\bar{\eth}^2(Q+iU)(\hat{n}) \end{aligned} \tag{2.1.2.6}$$

$$\begin{aligned} a_{-2,\ell m} &= \int d\Omega\, {}_{-2}Y^*_{\ell m}(\hat{n})(Q-iU)(\hat{n}) \\ &= \left[\frac{(\ell+2)!}{(\ell-2)!}\right]^{-1/2} \int d\Omega Y^*_{\ell m}(\hat{n})\eth^2(Q-iU)(\hat{n}) \end{aligned} \tag{2.1.2.7}$$

These are useful to define the linear combination:

$$\begin{aligned} a^E_{\ell m} &= -\left(a_{2,\ell m} + a_{-2,\ell m}\right)/2 \\ a^B_{\ell m} &= i\left(a_{2,\ell m} - a_{-2,\ell m}\right)/2 \end{aligned} \tag{2.1.2.8}$$

We can use these to define two spin-0 quantities E, B [73]:

$$E(\hat{n}) = -\frac{1}{2}\left[\bar{\eth}^2(Q+iU) + \eth^2(Q-iU)\right] = \sum_{\ell m}\left[\frac{(\ell+2)!}{(\ell-2)!}\right]^{1/2} a^E_{\ell m}Y_{\ell m}(\hat{n}) \tag{2.1.2.9}$$

$$B(\hat{n}) = \frac{i}{2}\left[\bar{\eth}^2(Q+iU) - \eth^2(Q-iU)\right] = \sum_{\ell m}\left[\frac{(\ell+2)!}{(\ell-2)!}\right]^{1/2} a^B_{\ell m}Y_{\ell m}(\hat{n}), \tag{2.1.2.10}$$

where the pattern of E remains unchanged under parity transformation (electric parity), while B changes sign (magnetic parity)[1] [72].

2.1.3 Statistics of the CMB

Using the expressions above, we now want to extract information about the statistical properties of the $a_{\ell m}$'s of the CMB perturbations. As predicted by inflation, the primordial perturbations are expected to be highly Gaussian and their evolution is well approximated by linear theory at least until last scattering. Thus, we anticipate that the CMB anisotropies will also follow *Gaussian statistics*, and they are so small that the universe is globally isotropic. If we treat the expansion components in Eqns. 2.1.2.1 and 2.1.2.3 as independent Gaussian variables, their distribution can be completely characterised by the mean and the variance[2]. Their average value vanishes $\left\langle a^T_{\ell m}\right\rangle = \left\langle a^E_{\ell m}\right\rangle = \left\langle a^B_{\ell m}\right\rangle = 0$, and their covariances are given by:

$$\left\langle a^{T*}_{\ell' m'} a^T_{\ell m}\right\rangle = C^{TT}_\ell \delta_{\ell\ell'}\delta_{m'm} \tag{2.1.3.1}$$

$$\left\langle a^{E*}_{\ell' m'} a^E_{\ell m}\right\rangle = C^{EE}_\ell \delta_{\ell'\ell}\delta_{m'm} \tag{2.1.3.2}$$

$$\left\langle a^{B*}_{\ell' m'} a^B_{\ell m}\right\rangle = C^{BB}_\ell \delta_{\ell'\ell}\delta_{m'm} \tag{2.1.3.3}$$

$$\left\langle a^{T*}_{\ell' m'} a^E_{\ell m}\right\rangle = C^{TE}_\ell \delta_{\ell'\ell}\delta_{m'm} \tag{2.1.3.4}$$

$$\left\langle a^{B*}_{\ell' m'} a^E_{\ell m}\right\rangle = \left\langle a^{B*}_{\ell' m'} a^T_{\ell m}\right\rangle = 0 \tag{2.1.3.5}$$

where C^{XY}_ℓ are the angular *power spectra*[3], defined as the ensemble average of the coefficients over all realisations of the sky. Here, δ is the Kronecker symbol, which is introduced since the $a_{\ell m}$'s are independent random variables, i.e. the covariance is diagonal in ℓ and m, and it is m-independent due to isotropy. The cross correlations in 2.1.3.5 vanish if the parity is respected in the average (since B has opposite parity to T and E). Therefore, we only need four power spectra to describe the statistics of the CMB, $C^{TT}_\ell, C^{EE}_\ell, C^{BB}_\ell, C^{TE}_\ell$. The angular power spectra can be related to the 2-point correlation function in real space, $\xi(\theta)$, as described in more detail in Appendix A.4.

[1] Similarly to the electric ("gradient") and magnetic ("curl") fields.

[2] A Gaussian distribution is fully characterised by a mean (μ) and a variance $\left(\sigma^2\right)$. This is because all higher odd moments of the distribution are 0, and even moments can be written in terms of the variance (*Wick's theorem*).

[3] More specifically, when $X = Y$, the ensemble average for a single field is called the "auto" power spectrum, and for two different fields ($X \neq Y$) is the "cross" power spectrum.

2.1.3.1 Power spectrum estimation

All the information contained in the CMB perturbations is therefore summarised in their power spectra, which measure the amount of power at each angular scale in the CMB fluctuations, and are rotationally invariant quantities. Ideally, the ensemble average of the $a_{\ell m}$'s should be measured on an infinite number of sky realisations. In practice, this is not possible (we can only observe a single realisation of the CMB field).

Full sky. A *statistical estimator* of the observed C_ℓ is used instead[4], by calculating the mean on m modes, i.e. $m = 2\ell + 1$ modes for each ℓ:

$$\hat{C}_\ell^{XY} = \frac{1}{2\ell+1} \sum_m a_{\ell m}^X a_{\ell m}^{Y*}, \tag{2.1.3.6}$$

the sum of the square of the normal Gaussian variables $a_{\ell m}$ with $\nu = 2\ell + 1$ *degrees of freedom* are χ^2_ν distributions, whose expectation values $\langle \hat{C}_\ell \rangle \equiv C_\ell$ (i.e. the estimator is *unbiased*).

The uncertainty in the estimated C_ℓ's, due to the fact that only a finite number of modes can be observed, is given by the *cosmic variance*:

$$\mathrm{Cov}\left[\hat{C}_\ell^{XY}, \hat{C}_\ell^{ZW}\right] = \frac{2}{(2\ell+1)} \sum_{\ell'=\ell_{\min}}^{\ell_{\max}} C_{\ell'}^{XY} C_{\ell'}^{ZW} \tag{2.1.3.7}$$

which can be derived using the *Wick's theorem*[5]. The cosmic variance therefore decreases as the number of observed ℓ modes increases. For high ℓ modes, where many m modes can be observed, the cosmic variance is not a major source of uncertainty. However, for low ℓ modes, corresponding to large angular scales, the cosmic variance can be a dominant source of uncertainty [75].

The power spectrum of the CMB fluctuations C_ℓ can be estimated numerically from the power spectrum of the fluctuations generated by inflation $\mathcal{P}(k)$ using one of the publicly available codes: `CMBFAST`, `CAMB`, `CLASS` [76–78].

[4] In the C_ℓ estimator, the ensemble average can be replaced with an average over different positions and directions. According to the *ergodic hypothesis*, different regions that are far apart in the sky are statistically independent of each other and can be seen as different realisations of the same stochastic process. As we can only observe the CMB field at a specific location and time, we are left with an average over different directions or values of m.

[5] The Wick's theorem states that for a Gaussian field, any N-point (N-even) statistics can be written as a function of the 2-point correlation.

Partial sky. In practice, an observed map $\tilde{a}$ is typically modelled as the product of the true map a and a weight map w:

$$\tilde{a}(\hat{\mathbf{n}}) = w(\hat{\mathbf{n}})a(\hat{\mathbf{n}}). \tag{2.1.3.8}$$

While we will focus on scalar fields like T for simplicity, the techniques used can be extended to fields of any spin, including spin-2 such as Q and U, as discussed in detail in [79]. The weight map can be interpreted as a mask that identifies which pixels have been observed ($w = 1$) and which have not ($w = 0$), as well as a local weight that inversely scales regions of high noise.

The observed field $\tilde{a}$ can be expressed as a convolution of the spherical harmonic coefficients of the true field m and the weight map w. This results in mode coupling between the power spectra of the observed and true fields. This is quantified with the "coupled pseudo-C_ℓ" $\tilde{C}_\ell^{ab}$. The relation between $\tilde{C}_\ell^{ab}$ and the true underlying power spectrum C_ℓ^{ab} is given by

$$\left\langle \tilde{C}_\ell^{ab} \right\rangle = \sum_{\ell'} M_{\ell\ell'}^{vw} C_{\ell'}^{ab}, \tag{2.1.3.9}$$

where the coupling between different ℓ modes is captured by the mode-coupling matrix (MCM) $M_{\ell\ell'}^{vw}$, which depends only on the weight maps of the two fields being correlated. The MCM relates the observed coupled power spectrum to the true underlying power spectrum via a matrix multiplication. The MCM can be computed analytically and efficiently using the orthogonality of the Wigner-3j symbols, as demonstrated in [80]. Explicit expressions for the MCM of different combinations of spin-0 and spin-2 fields can be found in [79]. In the limit of full-sky data, the MCM reduces to the identity matrix, and departures from this limit give rise to a statistical off-diagonal coupling between neighbouring multipoles. However, in many practical cases, the MCM is non-invertible, meaning that the coupled pseudo-C_ℓ cannot be turned into an unbiased estimator of the true power spectrum C_ℓ^{ab}.

The pseudo-C_ℓ estimator then proceeds along the following three steps:

1. *Binning.* To regularise the mode-coupling matrix (MCM), we employ binning of the power spectra into bandpowers that are represented by weighted sums of different ℓ values as follows:

$$\tilde{C}_q^{ab} = \sum_{\ell \in q} B_q^\ell \tilde{C}_\ell^{ab} \tag{2.1.3.10}$$

where q represents a bandpower and B_q^ℓ is the weight associated with the ℓ-th multipole moment in that bandpower. The binned MCM is then defined as:

$$\mathcal{M}_{qq'}^{vw} \equiv \sum_{\ell \in q} \sum_{\ell' \in q'} B_q^\ell M_{\ell\ell'}^{vw}, \tag{2.1.3.11}$$

which can be invertible for sufficiently broad bandpowers and serves as a regularisation technique for the MCM.

2. *Inverting the MCM.* Inverting the binned MCM, and removing the mode-coupled noise power spectrum $\tilde{N}q^{ab}$, returns the decoupled bandpowers:

$$\hat{C}_q^{ab} = \sum_{q'} (\mathcal{M}^{vw})^{-1}_{qq'} \left(\tilde{C}_{q'}^{ab} - \tilde{N}_{q'}^{ab} \right), \tag{2.1.3.12}$$

3. *Bandpower convolution.* $\hat{C}_q^{ab}$ can be related to the theoretical power spectrum $C\ell^{ab}$ by convolving it with the bandpower window functions $\mathcal{F}q\ell^{vw}$:

$$C_q^{ab} = \sum_{\ell} \mathcal{F}_{q\ell}^{vw} C_\ell^{ab}, \tag{2.1.3.13}$$

where

$$\mathcal{F}_{q\ell}^{vw} = \sum_{q'} (\mathcal{M}^{vw})^{-1}_{qq'} \sum_{\ell' \in q'} w_q^{\ell'} M_{\ell'\ell}^{uv}. \tag{2.1.3.14}$$

are obtained through a fast matrix-vector multiplication involving the binned MCM and the mode-coupling matrix $M_{\ell'\ell}^{uv}$, which describes how different modes of the power spectrum are coupled due to finite pixel size and partial sky coverage.

When analysing spin-2 quantities, a mask can not only lead to coupling between ℓ modes, but also to coupling between the two polarisation modes, E- and B-modes. To avoid the E-mode variance from leaking into the B-modes, this effect should be removed at the map level, a procedure known as *B-mode purification* [79, 81, 82].

2.1.4 Observations

The CMB is studied by measuring its temperature and polarisation anisotropy maps, which can be decomposed into their respective harmonic modes (see Sec. 2.1.2), and then cross-correlating the different modes one obtains the power spectrum, as explained in Sec 2.1.3. Each power spectrum provides distinct information about the physics of the CMB. The different effects that generate primary and secondary anisotropies can be observed at different angular scales, producing characteristic patterns in the power spectra. These patterns are briefly examined in what follows.

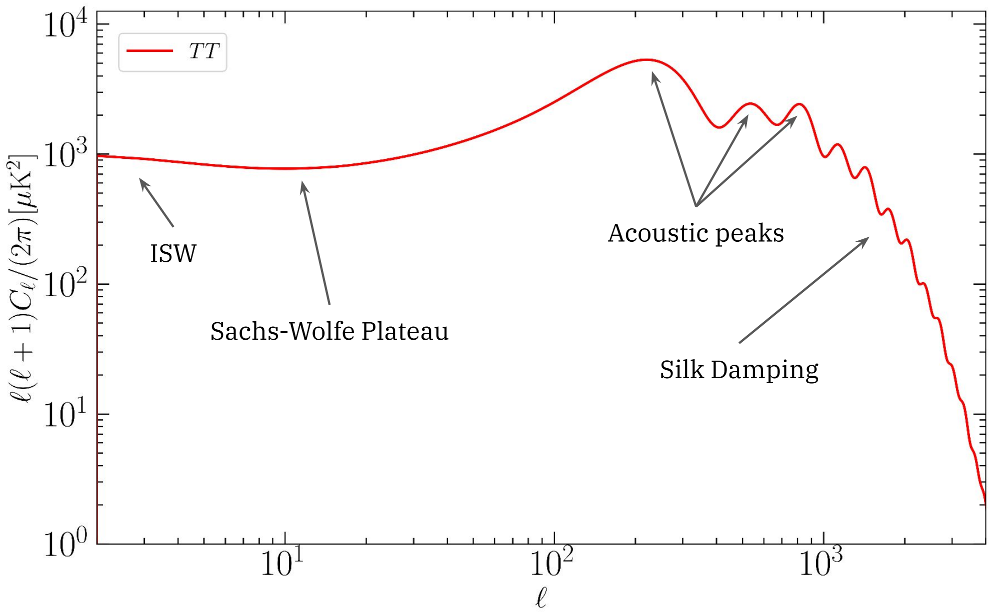

Figure 2.1: The CMB temperature power spectrum calculated with `CAMB`, as a function of the multipole ℓ, which relates to the angular scale as $\ell \approx \pi/\theta$. Different regimes are enlightened, which are discussed in detail in Sec. 2.1.4.1.

2.1.4.1 Temperature

The CMB temperature anisotropies have an amplitude that is about 10^5 times smaller than the average background temperature, which reflects the remarkable homogeneity of the primordial Universe.

In the temperature power spectrum C_ℓ^{TT}, displayed in Fig. 2.1, three main regions can be distinguished:

1. *Sachs-Wolfe region* – In the low ℓ region ($\ell \lesssim 100$) of the power spectrum, corresponding to large angular scales, the anisotropies are dominated by the early-time Sachs-Wolfe effect, which occurs during the radiation-dominated era before recombination. In other words, the *Sachs-Wolfe plateau* represents the angular scales where the density perturbations in the early Universe were still in the linear regime, and the gravitational potential fluctuations were not yet affected by the growth of structures. At $\ell \lesssim 5$, there is a small upturn in the power spectrum shape, known as the *Integrated Sachs-Wolfe (ISW) rise*, due to the effect of time-varying gravitational fields at very large scales. In particular, varying dark energy equation of state could be probed through the ISW rise. However, measurement uncertainties in this region of the spectrum are dominated by cosmic variance.

2. *Acoustic peaks*: at intermediate angular scales ($100 \lesssim \ell \lesssim 1000$), the anisotropies are mainly due to acoustic oscillations of the photon-baryon fluid before the decoupling of matter and radiation. These oscillations generate peaks and troughs in the temperature power spectrum. The position and height of the peaks are sensitive to the geometry, expansion rate, and composition of the Universe.

3. *Silk damping tail*: At small angular scales ($\ell \gtrsim 1000$), the photon-baryon fluid is tightly coupled and diffusion of photons smooths out the small-scale fluctuations, and suppresses the power spectrum, leading to the so-called Silk damping tail. Another effect that contributes to the damping is the finite thickness of the last scattering surface, which arises because the decoupling of photons from matter is not an instantaneous process. As a result, a given observed mode can have contributions from physical oscillation modes that have decoupled at different moments, leading to a loss of coherence in the temperature fluctuations. At very high multipole moments ($\ell \gtrsim 4000$), the acoustic peaks become indistinguishable due to the dominance of secondary anisotropies (see Sec. 1.5.5), due to e.g. the kSZ effect, the Rees-Sciama effect and point sources [6].

2.1.4.2 Polarisation

The amplitude of the polarisation anisotropies is much smaller than that of the temperature anisotropies. If the latter have an amplitude of about one part in 100,000, or 10^{-5}, the E-mode polarisation anisotropies have an amplitude of about one part in a million, or 10^{-6}, and the lensing B-mode polarisation anisotropies are even smaller, at most one part in ten million, or $10^{-7} - 10^{-8}$. The polarisation spectra are shown in Fig. 2.2.

E-modes The EE power spectrum is generated by the same physical process as the temperature anisotropies, i.e. the density fluctuations in the primordial Universe. This means that the global structure of the E-modes power spectrum is similar to that of the temperature power spectrum. However, the phase shift between the velocity field and the density field causes the E-modes to be out of phase with the temperature modes in the harmonic domain.

[6] The Rees-Sciama effect is caused by arises from the time variation of the gravitational potential along the line of sight due to non-linear growth (unlike ISW), while point sources refer to extragalactic sources, such as radio galaxies and quasars, that can produce non-negligible temperature fluctuations in the CMB.

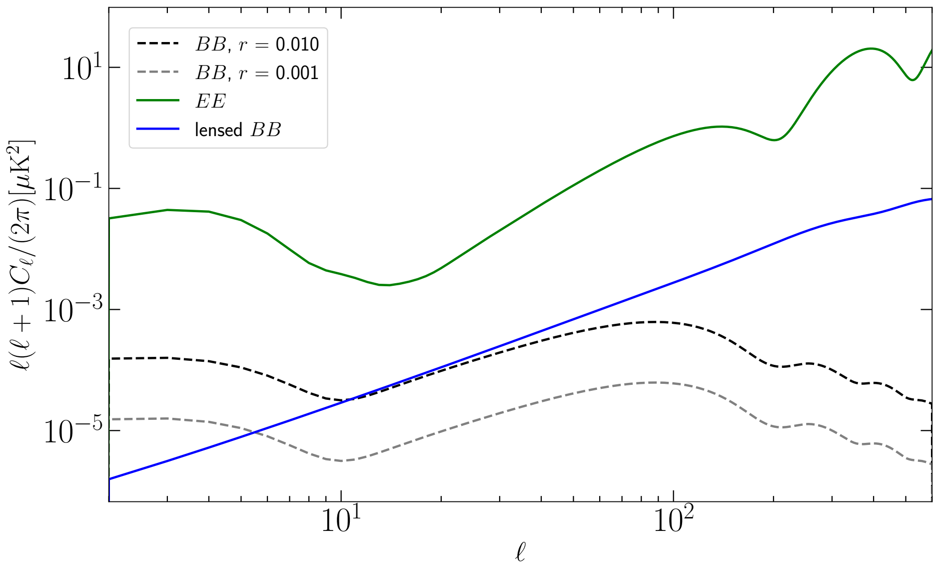

Figure 2.2: E-mode and B-mode power spectra calculated with `CAMB`, described in Sec. 2.1.4.2. The amplitude of the former is 10^3 times larger than the amplitude of the latter. The lensing BB spectrum (blue) can be distinguished from the primordial BB spectrum (dashed lines), plotted for different values of the tensor-to-scalar ratio r.

Another important feature of the E-modes power spectrum is the reionisation bump at low ℓ, which arises due to inverse differential Thompson scattering during reionisation. The amplitude and position of this bump can be used to measure the optical depth to reionisation, τ, and the redshift of reionization, z_{reion}, respectively[7]. The precise measurement of these parameters can help us better understand the epoch of reionization. It is worth noting that the primordial B-modes power spectrum is also expected to exhibit a similar structure due to reionisation at low ℓ.

B-modes Here, we focus on the power spectrum of primordial B-modes, which are distinct from lensing B-modes that will be discussed in the subsequent section. The amplitude of the power spectrum of primordial B-modes is directly proportional to the tensor-to-scalar ratio r (see Eq. 1.4.3.28).

The shape of the power spectrum for large angular scales of primordial B-modes is similar to that of E-modes, with a visible reionisation bump at low ℓ, as shown in Figure 2.2. In order to align with the predictions of a range of admissible inflationary models (see Sec. 1.6), forthcoming experiments have set a goal to measure r at a

[7]Note that a - degenerated - measurement of τ can be obtained thanks to the amplitude of the power spectra that scales like $A_s e^{-2\tau}$. Moreover, the position of the reionization peak relates to the redshift of reionisation as $\ell_{\text{reion.}} \propto \sqrt{z_{\text{reion.}}}$.

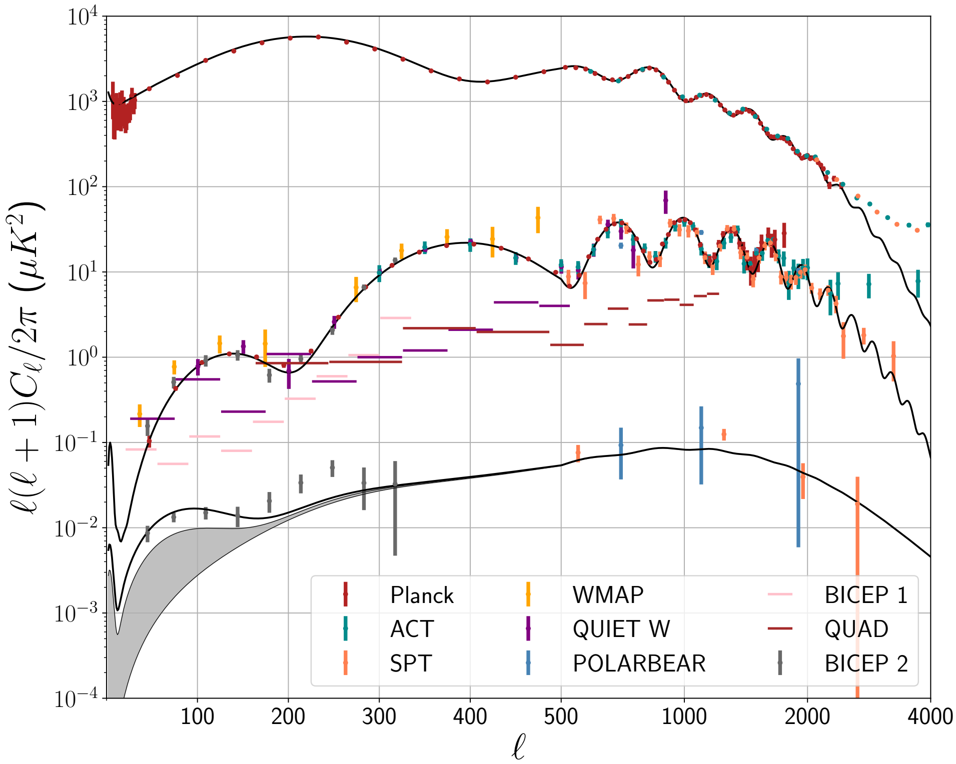

Figure 2.3: Theory power spectra (TT, EE, BB) that are derived from Planck (2013) parameters and calculated using `CAMB` [77] including the gravitational lensing contribution (in black), compared with the data coming from several experiments over the years, including Planck [22], ACT [83], SPT [84], WMAP [21], QUIET [85], POLARBEAR [86], QUAD [87], BICEP1 [88] and BICEP2/Keck Array [89]. The line of the B-mode power spectrum includes $r = 0.1$ tensor modes, and overlaps with the fits to the BICEP2 data.

substantially reduced value of 0.001, which is within the capabilities of realistic detectors. Thus, this discovery would be a significant milestone in cosmology, confirming one of the key predictions of inflation.

Ground experiments mostly focus on small patches of the sky to measure the primordial B-modes recombination peak at $\ell \sim 80$, which corresponds to less than a one-degree-wide sky patch. However, measuring the power spectrum of B-modes at larger angular scales, is only possible from space.

2.1.4.3 Measurements

Several experiments have measured the CMB temperature and polarisation fields over the years, and are described briefly in this Section. In Figure 2.3, some of these measurements are compared with the TT, EE and BB spectra calculated using CAMB, including gravitational lensing and tensor modes (the latter, only for the B-mode spectrum).

Temperature In the 1970s, the existence of CMB temperature anisotropies was predicted, but it was not until 1992 that they were first measured by the satellite COBE [12], specifically on angular scales larger than 10°. Subsequently, in 2001, two balloon-borne experiments, BOOMERanG (Balloon Observations Of Millimetric Extragalactic Radiation ANd Geophysics) [90] and MAXIMA (Millimeter Anisotropy eXperiment IMaging Array) [15], detected the first acoustic peak, while secondary peaks were measured by BOOMERanG and DASI (Degree Angular Scale Interferometer) [17]. These detections were a significant milestone in modern cosmology and lent support to the ΛCDM model as the standard cosmological model. Despite their importance, these experiments only covered less than 1% of the sky.

Since then, NASA's WMAP (Wilkinson Microwave Anisotropies Probe) mission (2001-2010) [21] and ESA's Planck mission (launched in 2009) [22] have greatly enhanced full-sky measurements, improving angular resolution and temperature sensitivity. This progress has facilitated a better understanding of the primordial universe and the phenomena occurring between the observer and the last scattering surface.

More recent experiments such as the Atacama Cosmology Telescope (including different generations: ACT, ACTpol, AdvACT) [83, 91–93] and the South Pole Telescope (SPT) [84] have further pushed the boundaries of CMB observations. Presently, temperature anisotropy measurements are capable of reaching the cosmic variance limit on scales as small as a few arcminutes, thereby offering valuable insights into the early universe's structure and evolution.

Polarisation The detection of polarisation E-modes dates back to 2002, when the DASI experiment [17], a ground based interferometer located at the South Pole Amundsen-Scott research station, first observed the TE and EE signal. Following DASI, other experiments have claimed EE detection, such as CAPMAP [24], an evolution of the PIQUE system [94] and CBI (Cosmic Background Imager), a radio interferometric receiver working in the band $26-36$GHz and covering a wide range of multipoles ($300 \lesssim \ell \lesssim 3500$), which reported a 8.9σ detection of EE on small angular scales [95], and the WMAP satellite, which measured TE correlations with high ℓ resolution in many σ as of the first data release of 2003 [96], probing the low multipoles sensitive to the CMB photon's optical depth τ. Subsequently, observations by WMAP and Planck [52] have allowed a comprehensive understing across a wide range of angular scales. However, there is still room for improvement at both very large and small scales, where the cosmic variance limit has not yet been reached. The situation is more complicated for B-modes. As previously discussed, these are expected to arise from two sources: lensing B-modes due to weak gravitational lensing (see Sec. 1.5.5.1), and primordial B-modes resulting from a hypothetical inflationary process (see Sec. 1.4.3.3). Lensing B-modes were first observed in 2013 by SPT [97], confirmed by the POLARBEAR experiment in 2014 [86], and have since been detected by various other CMB polarisation experiments, including BICEP2 (Background Imaging of Cosmic Extragalactic polarisation) [98], Keck Array [99], SPTpol [100], combined BICEP2/Keck Array, Planck and WMAP [89, 101, 102], ACTPOL [103] and SPIDER (Spectro-Polarimetric Imaging of the Degree-Scale CMB Anisotropy) [104]. The current state-of-the-art constrain on primordial B-modes was determined by the BICEP2/Keck Array experiment, in conjunction with WMAP and Planck data sets, which determined the upper limit on the value of r to be 0.032 at a 95% confidence level [59] when evaluated at a pivot scale of 0.05 Mpc^{-1}. However, an unequivocal detection of primordial B-modes remains a challenge, as they have not yet been observed by any experiment. This is due in part to the difficulty of distinguishing between CMB emission and polarised foregrounds on the line-of-sight. This is discussed in detail in the following Chapters.

2.2 Upcoming CMB B-modes observatories

One of the primary science targets of upcoming CMB observatories is the detection of the primordial B-mode polarisation pattern in CMB radiation at large scales, which are produced by gravitational waves in the primordial universe. Thus, an unequivocal measurement of r, or a stringent upper bound, would be able to greatly constrain the landscape of theories on the early Universe, as discussed in Chapter 1. These CMB B-modes observatories include ground-based and space-based telescopes.

Ground-based telescopes, such as SPT, BICEP and Keck Array (located at the South Pole), ACT, POLARBEAR, Simons Array and the upcoming Simons Observatory and CMB-S4 (located at the Atacama desert in Chile), can take advantage of larger apertures compared to satellites (since launching larger telescopes into space is prohibitively expensive). Moreover, ground-based experiments can concentrate their sensitivity on a small portion of the sky for an extended period, providing high resolution. Additionally, the availability of instruments on the ground provides practical advantages such as extended project lifetimes that are not limited by the availability of cryogens or fuel, phased construction, and are easy to upgrade.

There are also space-based observatories like the Planck satellite mission operated by the European Space Agency and the upcoming LiteBIRD mission by the Japanese Aerospace Exploration Agency (JAXA). These instruments have synergies and complementarities to those on the ground [105]. Nonetheless, in space instruments experience 5 to 10 times lower detector noise compared to their ground-based counterparts and can detect the full frequency range of interest without the constraints of atmospheric windows or opacity, which are prohibitive for ground-based experiments above $\sim$ 300GHz. Furthermore, the L2 Lagrange point's outstanding thermal stability reduces thermal drifts and strengthens systematic error mitigation in resulting datasets. Space missions also benefit from their ability to observe the entire sky, which is not feasible from any single terrestrial location.

Near-future experiments such as the BICEP Array [106] target a detection at the level of $r \sim 0.01$, while in the following decade, next-generation projects such as LiteBIRD [107] and CMB-S4 [69] aim at $r \sim 0.001$. The focus of this thesis is on two specific experiments, which are discussed in greater detail below.

2.2.1 Simons Observatory

The Simons Observatory (SO) targets the detection of primordial gravitational waves at the level of $r \sim 0.01$ [108], and its performance at realizing this goal is the main focus of this paper. SO is a ground-based experiment, located at the Cerro Toco site in the Chilean Atacama desert, which will observe the microwave sky in six frequency channels, from 27 to 280 GHz, due to start in 2023. SO consists of two main instruments. On the one hand, a Large Aperture Telescope (LAT) with a 6 m diameter aperture will target small-scale CMB physics, secondary anisotropies and the CMB lensing signal. Measurements of the latter will serve to subtract lensing-induced B-modes from the CMB signal to retrieve primordial B-modes [using a technique known as "delensing", see 109]. On the other hand, multiple Small Aperture Telescopes (SATs) with 0.5 m diameter apertures will make large-scale, deep observations of $\sim$ 10% of the sky, with the main aim of constraining the primordial B-mode signal, peaking on scales $\ell \sim 80$ (the so-called "recombination bump"). See [108] for an extended discussion on experimental capabilities.

2.2.2 LiteBIRD

The Lite (Light) satellite for the study of B-mode polarisation and Inflation from cosmic background Radiation Detection, or LiteBIRD, is a satellite mission that has been selected by the Japan Aerospace Exploration Agency (JAXA) to study the polarisation CMB and search for evidence of cosmic inflation. The satellite is planned to be launched in the late 2020s and will be placed in a halo orbit around the L2 Sun-Earth Lagrange point, which is a stable point in space where the gravitational forces of the Sun and Earth balance out. This position allows for a stable observing platform with minimal interference from the Earth and its atmosphere. LiteBIRD will observe the CMB over a three-year mission and will cover a frequency range of $34 - 448$GHz using 15 different frequency bands. The satellite will have a resolution of $0.5°$ at 100GHz. With these capabilities, LiteBIRD can detect B-modes or alternatively set an upper limit of $r < 0.002$ at 95 % confidence.

2.3 Data Analysis Pipeline

Telescopes scan the sky, using a given number of detectors as a "paint brush" and the scanning motion as paint strokes to paint a larger image. The direct output of such a telescope is not a single value per pixel, but instead a time series of values for each detector, called the *time-ordered data* (TOD). Based on this TOD and information about the telescope's scanning pattern, it is possible to reconstruct an image of the sky in a process called "*map making*".

2.3.1 Time-Ordered Data (TOD)

Before obtaining an image using a scanning telescope, it is necessary to create an accurate model that describes how the telescope transforms signals from the sky into time-ordered data (TOD), and then invert this model to reconstruct the original sky image. However, dealing with infinite-resolution skies or time-streams with arbitrary response functions is impractical, so the data needs to be discretized to enable numerical analysis. Therefore, when working with TODs, several simplifying assumptions are made:

- The TOD is discrete rather than continuous.
- The sky is modelled as a set of discrete pixels, which is a good approximation as long as the pixels are smaller than our angular resolution. We employ the HEALPix [110] scheme[8], where the resolution is parametrised by N_{side}, such that the number of pixels covering the entire sky is $N_{pix} = 12 \times N_{side}^2$.

[8]The Hierarchical Equal Area isoLatitude Pixelization (HEALPix) scheme organizes the centres of the pixels in isolatitude rings with equal azimuthal spacing, allowing the use of Fast Fourier Techniques (FFT) in the spherical harmonic transformations. This results in a faster calculation of spherical harmonics by a factor of $\sqrt{N_{\rm pix}}$, compared to traditional layouts. Additionally, the scheme assumes equal pixel area (i.e., equal flux density), and the "hierarchical" property makes changing the resolution easier. The angular size of each pixel is given by $\Omega_{pix} = 4\pi/N_{pix}$.

- The sky is assumed to be constant during each TOD, which implies that we can use sky-fixed coordinates such as galactic or equatorial coordinates.

- The response of the telescope is assumed to be linear. This linearity requirement is a standard design requirement for most detectors since non-linear responses can be challenging to handle.

Together, these assumptions allow us to model the TOD as a linear combination of signal and noise as a function of time t:

$$\vec{d}_t = s_t + \vec{n}_t \tag{2.3.1.1}$$

and the observed data vector has dimension ($N_{\text{det}} \times N_{\text{samp}}$), where N_{det} is the number of detectors and N_{samp} the number of data samples[9]. For practical purposes, here we start by presenting a simplistic data reduction pipeline.

2.3.1.1 Signal

The signal for a single detector is given by $s_{dt} = P_{tp}\vec{m}_{dp}$ [111], where the subscripts indicate the dependence on the detector d (corresponding to a polarisation angle, which is different for different detectors), the time t, and the pixel $p(\hat{n})$. Here, $P_{tp} \equiv \mathsf{P}$ is the *pointing matrix*, which is a rectangular matrix of dimension ($N_{\text{pix}} \times N_{\text{samp}}$) that relates the sky Stokes parameters at each pixel, represented as a map vector $\vec{m}$, with the sampling time. The pointing matrix specifies which pixel in the map was observed during each TOD sample[10].

For a CMB polarisation experiment, we aim to reconstruct the temperature map, $\vec{I}$, as well as the polarisation maps, $\vec{Q}$, and $\vec{U}$. The map will therefore have dimension ($N_{\text{pix}} \times 3$), and its components can be simply written as $m_p = [I_p, Q_p, U_p]$ (see Sec. 2.1.1 for a discussion on Stokes parameters). The projection in time domain of the three Stokes parameters, for a single detector, can be written as:

$$\vec{d} = P[\vec{I} + \vec{Q}\cos(2\psi) + \vec{U}\sin(2\psi)] + \vec{n} \tag{2.3.1.2}$$

where ψ is the detector polarisation angle expressed in a given sky coordinate system, and we dropped the subscripts.

[9]The number of data samples is determined by the sampling frequency of the detector.
[10]Note that many TOD samples may correspond to the same pixel in a sky map.

2.3.1.2 Noise

The majority of time-ordered data (TOD) collected by a scanning telescope consists of noise, with different sources of noise having characteristic spectral distributions that dominate specific frequency ranges. Thus, it is best to apply the noise model in the Fourier domain. The frequency domain version of the noise, $N_{dd'\nu\nu'}$, can be obtained from the time domain version, $N_{dd'tt'}$, using the discrete Fourier transform $\mathcal{F}_{t\nu} = \mathrm{e}^{-\frac{2\pi i\nu}{N_{\mathrm{samp}}}}/\sqrt{N_{\mathrm{samp}}}$. This is expressed by the equation:

$$N_{dd'\nu\nu'} \equiv \left\langle n_{d\nu} n^{\dagger}_{d'\nu'} \right\rangle = \mathcal{F}_{\nu t} N_{dd'tt'} \mathcal{F}^{-1}_{t'\nu'} \tag{2.3.1.3}$$

The noise is assumed to be Gaussian, with covariance given by:

$$N_{dd'\nu\nu'} = \sqrt{\phi_d(\nu)\phi_{d'}(\nu)} C_{dd'\nu} \delta_{\nu\nu'} \tag{2.3.1.4}$$

where $C_{dd'\nu}$ is the detector correlations, $\phi_d(\nu)$ is the noise power spectrum for detector d. In the case of stationary noise, the matrix $N_{dd'\nu\nu'}$ can be modelled as diagonal in Fourier space ($N_{\nu\nu'} = \delta_{\nu\nu'} N_\nu$), with each detector having a noise realization $\vec{n}$ independently drawn from a Gaussian distribution with zero mean and variance σ^2_{det}. However, time-dependent thermal variations across the focal plane and atmospheric emission (for ground-based experiments) can lead to correlations between detectors and also between different timestreams, resulting in a non-diagonal $C_{dd'\nu}$. This results in a noise power spectrum of the form:

$$\phi(\nu) \equiv \sigma_0^2 \left[1 + \left(\frac{\nu}{\nu_{\mathrm{knee}}}\right)^{\alpha}\right] \tag{2.3.1.5}$$

where σ_0 is the white noise level, ν_{knee} is the knee frequency, and α is the slope.

2.3.1.3 Realistic scenario

When dealing with actual data, more components need to be included in P when modelling the signal. A more general expression to model the measured power at time t is:

$$d_t = gH_t \otimes \left[n_t + \frac{1}{2}\int d\nu\, W(\nu) \int d\Omega\, s(t,\nu,\hat{n},r)\right] \tag{2.3.1.6}$$

where $\otimes$ is the Kronecker product, H_t is the time transfer function, n is the noise, g is the responsivity or the conversion factor between Watt and Volt, $W(\nu)$ is the bandpass as a function of frequency ν, and the unit vector $\hat{n}$ gives the pointing position in sky. The integrated signal $s(t,\nu,\hat{n},r)$ can be expressed as:

$$\begin{aligned} s(t,\nu,\hat{n},r) =& B_{\parallel}(r,\nu,\hat{n}) * [I(\nu,\hat{n}) + \epsilon(\nu)(Q(\nu,\hat{n})\cos 2\psi(\nu) + U(\nu,\hat{n})\sin 2\psi(\nu))] \\ &+ B_{\perp}(r,\nu,\hat{n}) * [I(\nu,\hat{n}) - \epsilon(\nu)(Q(\nu,\hat{n})\cos 2\psi(\nu) + U(\nu,\hat{n})\sin 2\psi(\nu))] \end{aligned} \tag{2.3.1.7}$$

where $*$ indicates convolution, $B_{\parallel}$ ($B_{\perp}$) is the co-polar (cross-polar) beam at the detector coordinate position r in the focal plane, $\epsilon(\nu)$ is the polarisation efficiency independent of the beam and only dependent on the properties of the detector, $\psi(\nu)$ is the detector polarisation orientation angle including sinuous antenna wobbling.

2.3.2 Map making

With our signal and noise model in hand, we then want to reconstruct a map of the sky. There are different ways of solving $d = Pm + n$ for m. Given the data model, it is possible to either directly solve for an unbiased map (as in maximum likelihood mapmaking or destriping), or to measure and correct for the bias in a biased estimator (as in filter+bin mapmaking). These common methods are briefly described below.

- *Filter and bin*: the TOD is filtered before obtaining a binned map, where "binning" means averaging the TOD samples which "observe" a given pixel and returns a map comprised of these averages [85, 112]. It is then common to characterize the intrinsic bias of the method using observation matrices or simulations.

- *Destriping*: the noise is split into a correlated and an uncorrelated part, and the correlated noise is modelled as a series of slowly changing degrees of freedom to be solved for jointly with the sky image itself [113, 114].

- *Maximum-likelihood*: these methods find the generalized least-squares solution for the map [85, 111, 115–118]. The main limitation of these methods is that the covariance matrix is often poorly conditioned, making it slow for iterative solution methods to recover the large scales.

2.3.2.1 Maximum-Likelihood methods

For the covariance of the data we assume this expression [111]:

$$\langle d_t \rangle = P_{tp} m_p \tag{2.3.2.1}$$

where $\langle \cdots \rangle$ is the average over noise realisations. We assume that the noise is Gaussianly distributed, such that the likelihood of m given the data stream is given by

$$\mathcal{L}(m) = \frac{\mathrm{e}^{-\frac{1}{2}(d-Pm)^T N^{-1}(d-Pm)}}{\sqrt{|2\pi N|}}. \tag{2.3.2.2}$$

The minimum variance map is then found by maximising $\mathcal{L}(m)$ with respect to the signal over all possible signals:

$$\frac{d\mathcal{L}(\hat{m})}{d\hat{m}} = 0 \quad \Rightarrow \quad \hat{m} = \left(P^T N^{-1} P\right)^{-1} P^T N^{-1} d. \tag{2.3.2.3}$$

where m is unbiased ($\langle \hat{m} \rangle = m$) and Gaussian distributed with covariance $(P^T N^{-1} P)$. Making maps from TODs is therefore an inverse problem, where the covariance needs to be inverted to recover the maps. In practice, the covariance can be seen as an operator where the pointing P projects the pixels values into time-ordered samples, N^{-1} performs inverse-variance weighting of the data and P^T projects the data back onto a map. In other words, we can construct a covariance matrix that acts on the TOD to return an estimate for the map, which is unbiased and optimally noise-weighted.

2.3.2.2 Iterative solvers

In practice, the matrices involved are often too large to solve by direct inversion, due to the large number of samples in the TOD ($\mathtt{N}$ has dimension $N_{\text{samp}} \times N_{\text{samp}}$)

One option is to resort to *conjugate gradient techniques* [see, e.g. 119, 120] to calculate the inverse covariance without ever explicitly building the inverse matrix. These techniques are iterative and several iterations must be run to arrive at an acceptable estimate for the map. One example is implemented in Chapter 6, where more details on this technique are provided.

2.3.3 Likelihood and parameter inference

Our objective is to obtain measurements of various parameters, including cosmological ones, from maps. Specifically, we aim to assess the *posterior distribution* of the parameters (θ) based on the given data ($\mathbf{d}$), represented by $\mathcal{P}(\theta \mid \mathbf{d})$.

Bayesian methods have become increasingly popular, particularly for analysing noisy CMB data, and will be used in throughout this thesis to obtain constraints[11].

[11] Generally, both Bayesian and frequentist methods can be used to estimate cosmological parameters from CMB data. One advantage of the Bayesian approach is that it allows for the incorporation of prior knowledge about the parameters, such as constraints from other cosmological probes. This can lead to more accurate and precise parameter estimates, particularly when the data is limited or noisy (as is generally the case with CMB data). However, the choice of prior can have a significant impact on the results, and care must be taken to choose an appropriate prior. Frequentist methods, on the other hand, do not require the specification of a prior distribution, but can be limited by the assumptions of the model and the need for large sample sizes. Additionally, they do not provide a direct measure of uncertainty, but rely on confidence intervals to give a range of plausible parameter values.

Bayes' theorem can be employed to establish the relationship between the probability of the data given the parameters $\mathcal{P}(\mathbf{d} \mid \theta)$ (the likelihood), and the posterior of the parameters based on the data, given by:

$$\mathcal{P}(\theta \mid \mathbf{d}) = \frac{\mathcal{P}(\mathbf{d} \mid \theta)\mathcal{P}(\theta)}{\mathcal{P}(\mathbf{d})} \tag{2.3.3.1}$$

where $\mathcal{P}(\theta)$ is the probability of our model (known as the *prior*).

2.3.3.1 Pixel-based Likelihood

In the case of the CMB, maps $\mathbf{m} = (I, Q, U)$ exhibit Gaussian fluctuations, whose probability density function $\mathcal{P}(\mathbf{m} \mid C_\ell)$ is given by:

$$\mathcal{L}\left(C_\ell\right) = \mathcal{P}\left(\mathbf{m} \mid C_\ell\right) = \frac{1}{|2\pi\mathsf{C}|^{1/2}} \exp\left(-\frac{1}{2}\mathbf{m}^{\mathrm{T}}\mathsf{C}^{-1}\mathbf{m}\right) \tag{2.3.3.2}$$

where $\mathbf{m}$ is a data vector of length $\mathrm{N_{pix}}$ that contains the pixels of the map, and C is a $(\mathrm{N_{pix}} \times \mathrm{N_{pix}})$ covariance matrix that encodes the correlation between pixels. Specifically, the covariance is given by

$$\mathsf{C}_{ij} = \langle m_i m_j^\top \rangle = \mathsf{S}(\theta) + \mathsf{N}, \tag{2.3.3.3}$$

where N is the pixel space noise covariance matrix and $\mathsf{S}(\theta)$ is the signal covariance that depends on the parameters θ that we want to sample. This is related to the (theoretical) power spectrum C_ℓ through the following expression (for temperature T):

$$\mathsf{S} = \langle \Delta T_i \Delta T_j \rangle = \frac{1}{4\pi} \sum_\ell (2\ell + 1) C_\ell P_\ell \tag{2.3.3.4}$$

assuming isotropy of the signal, where P_ℓ is the Legendre polynomial.

There are certain benefits to working in pixel-space. For instance, handling masks can be done with ease. However, when it comes to computing the likelihood for a given set of θ parameters, we must determine the determinant and inverse of C, which is a time-consuming process that scales as $\mathcal{O}(N_{\mathrm{pix}}^3)$. This approach may be viable for low resolutions, but it becomes unfeasible at higher resolutions, whereby alternative solutions may be necessary.

2.3.3.2 Harmonic-space likelihood

An alternative approach to using the maps directly is to estimate the power spectrum ($\hat{C}_\ell$) from the maps and use that as our "data".

In a noiseless full-sky experiment, where only temperature is considered, we can compute $\hat{C}_\ell$ using the estimator in Eq. 2.1.3.6. Accordingly, we can define a quantity $\hat{Y} = \sum_{m=-\ell}^{\ell} \left|a_{\ell m}/\sqrt{C_\ell}\right|^2$, which follows a χ^2 distribution with $\nu = 2\ell+1$ degrees of freedom:

$$\mathcal{P}\left(\hat{Y}_\ell \mid C_\ell\right) = \frac{\hat{Y}_\ell^{\nu/2-1}}{\Gamma(v/2)2^{v/2}} \exp\left(-\frac{\hat{Y}_\ell}{2}\right) \tag{2.3.3.5}$$

where C_ℓ is the theoretical spectrum.

This means that the estimated (or observed) $\hat{C}_\ell$, being a multiple of $\hat{Y}_\ell$, has a Gamma Γ distribution with ν degrees of freedom:

$$\mathcal{P}\left(\hat{C}_\ell \mid C_\ell\right) \propto C_\ell^{-1} \left(\frac{\hat{C}_\ell}{C_\ell}\right)^{v/2-1} \exp\left(-\frac{\nu}{2}\frac{\hat{C}_\ell}{C_\ell}\right) \tag{2.3.3.6}$$

This expression corresponds to the full-sky likelihood $\mathcal{L}(C_\ell)$, given the fiducial power spectrum, expressed in terms of $\hat{C}_\ell$. The $\hat{C}_\ell$ distribution has a mean of C_ℓ and a variance of $2C_\ell^2/\nu$. The maximum of the distribution is $(\nu-2)/\nu\, C_\ell$, which is not the same as the mean of the distribution, resulting in a skewed distribution of observed $\hat{C}_\ell$. However, in the limit $\nu \to \infty$ (at high ℓ), the central limit theorem ensures that the distribution tends towards a Gaussian distribution with the same mean and variance. It is important to note that the variance of the distribution represents the cosmic variance.

If we include polarisation as well, this formalism extends to give a Wishart distribution (i.e. a multi-dimensional generalization of the Γ distribution) for $\hat{\mathbf{C}} = (\hat{C}^{\mathrm{TT}}\, \hat{C}^{\mathrm{TE}}\, \hat{C}_l^{\mathrm{EE}}\, \hat{C}_l^{\mathrm{BB}})^T$. For $\nu \to \infty$, the distribution of the $\hat{C}_\ell$'s tends to a Gaussian distribution, due to the central limit theorem [75]. In this approximation, the likelihood can be written as in Eq. 2.3.3.6:

$$-\ln \mathcal{L}(\hat{\mathbf{C}} \mid \mathbf{C}(\theta)) = \frac{1}{2}[\hat{\mathbf{C}} - \mathbf{C}(\theta)]^T \mathsf{C}^{-1}[\hat{\mathbf{C}} - \mathbf{C}(\theta)] + \text{ const} \tag{2.3.3.7}$$

where $\hat{\mathbf{C}}$ is the data, $\mathbf{C}(\theta)$ is the model which depends on the parameters θ we want to determine, and C^{-1} is the inverse covariance matrix, which can be estimated with a fixed fiducial set of parameters.

However, on large scales, the small number of available modes invalidates the central limit theorem and, as quadratic functions of Gaussian fields, power spectra exhibit non-Gaussian features in their likelihoods. To account for this effect, one solution is provided by the non-Gaussian likelihood developed by Hamimeche and Lewis (HL) [121], that coincides with the exact likelihood in full sky and is a good approximation in the cut-sky regime. This likelihood requires an estimate of the covariance matrix of the full set of power spectra, which can be obtained from simulations (e.g. see Chapters 3-5.) The HL likelihood additionally requires an estimate of the fiducial power spectra, as well as the noise power spectrum.

2.3.3.3 Sampling technique

Expressing the posterior distribution in an analytical form can be difficult, so we need to find a numerical representation for $\mathcal{P}(\boldsymbol{\theta} \mid \mathbf{d})$ to marginalize over a subset of our parameters or determine statistical summaries of the distribution. This can be accomplished by sampling $N \gg 1$ values, $\{\boldsymbol{\theta}_i : i = 1, \ldots, N\}$, from the distribution, and using them to estimate the expectation value of any function of the parameters, $f(\boldsymbol{\theta})$ as

$$\langle f(\boldsymbol{\theta})\rangle = \int \mathrm{d}\boldsymbol{\theta} f(\boldsymbol{\theta})\mathcal{P}(\boldsymbol{\theta} \mid \mathbf{d}) \approx \frac{1}{N}\sum_{i=1}^{N} f\left(\boldsymbol{\theta}_i\right). \tag{2.3.3.8}$$

The natural question now arises as to how one should generate such samples. One way to approach this problem is using the *Markov Chain Monte Carlo (MCMC)* techniques. MCMC is a powerful class of algorithms for sampling from complex, high-dimensional distributions. Unlike other numerical techniques, MCMC algorithms do not require explicit knowledge of the posterior distribution or any functional form of the target distribution. Instead, they build a Markov chain of samples, where each sample is generated based on the previous one, ensuring that the chain converges to the target distribution.

The fundamental idea behind MCMC is to construct a chain that has the target distribution as its stationary distribution. The chain is constructed by proposing a new sample, $\boldsymbol{\theta}'$, from a proposal distribution $q(\boldsymbol{\theta}'|\boldsymbol{\theta})$, which depends on the current state $\boldsymbol{\theta}$. Then, the proposed sample is accepted or rejected based on a probability ratio called the *Metropolis-Hastings acceptance probability*:

$$A\left(\boldsymbol{\theta}' \mid \boldsymbol{\theta}\right) = \min\left(1, \frac{\mathcal{P}\left(\boldsymbol{\theta}' \mid \mathbf{d}\right) q\left(\boldsymbol{\theta} \mid \boldsymbol{\theta}'\right)}{\mathcal{P}(\boldsymbol{\theta} \mid \mathbf{d}) q\left(\boldsymbol{\theta}' \mid \boldsymbol{\theta}\right)}\right) \tag{2.3.3.9}$$

where $\mathcal{P}\left(\boldsymbol{\theta} \mid \mathbf{d}\right)$ is the posterior distribution we wish to sample from, and $q(\boldsymbol{\theta}'|\boldsymbol{\theta})$ is the proposal distribution used to generate a new sample. The proposed sample is accepted with probability $A(\boldsymbol{\theta}'|\boldsymbol{\theta})$, or otherwise, the current sample is repeated, ensuring that the Markov chain has the desired target distribution as its stationary distribution.

MCMC algorithms have several advantages over other numerical methods. Firstly, they can sample from complex, high-dimensional distributions where other techniques may fail. Secondly, they provide a way to compute posterior expectations and probabilities, such as marginal distributions or posterior predictive distributions, without explicit knowledge of the normalizing constant. Lastly, they can be easily parallelized, making them suitable for large-scale problems[12]. This is why they are widely used in the context of CMB data analysis, and are adopted throughout this thesis, where we will often use an affine invariant ensemble sampler, as implemented in the `emcee` package [122].

2.4 Instrument model

The characterization and understanding of instrumental effects is crucial in the interpretation of astronomical data. Accurate modelling of the instrument response is necessary to disentangle the astrophysical signal from the instrumental systematic effects. Here, we discuss various formalisms used to describe the instrumental response, including the Jones formalism and the Mueller formalism. We also introduce a general framework to understand instrument systematics and their impact on measurements, and how the formalisms can be used to model and correct for these effects.

In particular, we introduce an optical component that is common to many CMB polarisation dedicated instruments, that will be characterised in more detail in Ch. 6: the half-wave plate (HWP). The HWP plays a critical role in manipulating and restoring the polarisation of incoming radiation, and is especially useful when dealing with a multi-frequency focal plane.

2.4.1 Jones formalism

The propagation of an incident radiation $\vec{E}$ through a receiver can be described by a *Jones matrix* $\mathbf{J}$ [123], such that the "output" electric field $\vec{E}_o$, after passing through the receiver, is

$$\overrightarrow{E_o} = \mathbf{J}\overrightarrow{E} \tag{2.4.1.1}$$

[12]It is worth mentioning that MCMC algorithms can be computationally expensive, especially for high-dimensional problems, and require careful tuning of the proposal distribution to ensure efficient exploration of the target distribution. Given that the number of parameters being sampled in this thesis is restricted, MCMC is the most suitable approach, and we make sure that convergence is achieved, to ensure that the chain has fully explored the target distribution.

where $\mathbf{J}$ is a 2×2 complex matrix. It describes how the instrument linearly transforms the two-dimensional vector representing the incoming radiation field $\vec{E}$ into the two-dimensional vector of the outgoing field $\vec{E}_o$.

For an instrument with several components, the Jones matrix is the product of the Jones matrices for each component. For example, the ideal Jones matrix for an instrument in which the incident radiation passing through a rotating HWP is [124]

$$\begin{aligned}
\mathbf{J}_{\mathrm{hwp}} &= \mathbf{J}_{\mathrm{rot}}^{\mathrm{T}}\,\mathbf{J}_{\mathrm{hwp}}\,\mathbf{J}_{\mathrm{rot}} \\
&= \begin{pmatrix} \cos(\omega t) & -\sin(\omega t) \\ \sin(\omega t) & \cos(\omega t) \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \begin{pmatrix} \cos(\omega t) & \sin(\omega t) \\ -\sin(\omega t) & \cos(\omega t) \end{pmatrix} \\
&= \begin{pmatrix} \cos(2\omega t) & \sin(2\omega t) \\ \sin(2\omega t) & -\cos(2\omega t) \end{pmatrix},
\end{aligned} \tag{2.4.1.2}$$

where ω is the angular velocity of the half-wave plate, $\mathbf{J}_{\mathrm{rot}}$ the rotation matrix and $\mathbf{J}_{\mathrm{hwp}}$ the ideal Jones matrix of the half-wave plate.

After passing through the receiver, a pair of ideal orthogonal linear detectors measure the power as

$$S_i = \tfrac{1}{2}(I \pm Q\cos(4\omega t) \pm U\sin(4\omega t)) \tag{2.4.1.3}$$

To model systematic errors within a polarisation-sensitive instrument, the Jones matrix can be described by introducing diagonal terms, i.e. the complex gain parameters g_x and g_y and non-diagonal terms [125], i.e. the complex coupling parameters e_x and e_y associated to the orthogonal polarisations

$$\mathbf{J} = \begin{pmatrix} 1 - g_x & e_x \\ e_y & 1 - g_y \end{pmatrix}. \tag{2.4.1.4}$$

For example, systematic errors arising from the half-wave plate can be modelled by a Jones matrix for the half-wave plate $\mathbf{J}_{\mathrm{hwp}}$, such that the electric field $\vec{E}$ propagated through the half-wave plate becomes

$$\vec{E}_i = \mathbf{J}_{\mathrm{rot}}^{\mathrm{T}}\mathbf{J}_{\mathrm{hwp}}\mathbf{J}_{\mathrm{rot}}\vec{E} \tag{2.4.1.5}$$

2.4.2 Mueller formalism

While the Jones matrix expresses the transformation of the electric field in the x and y directions, the *Mueller matrix* describes how the different polarisation states transform. Therefore, they can be used to "manipulate" partial polarised states, described using the Stokes parameters [see, for practical examples, 7, 126–130].

The Mueller matrix is a 4×4 matrix and can be written as the direct product of the 2×2 Jones matrices.

We calculate the tensor product of the outgoing field given by Eq. 2.4.1.1

$$\vec{E}_o^* \otimes \vec{E}_o = \mathbf{J}\vec{E} \otimes \mathbf{J}^* \vec{E}^* \tag{2.4.2.1}$$

In general, the direct product $\left\langle \vec{E} \otimes \vec{E}^* \right\rangle$ gives the vector $\vec{C}$ defined as

$$\vec{C} = \begin{bmatrix} C_{xx} \\ C_{yy} \\ C_{yx} \\ C_{yy} \end{bmatrix} = \left\langle \vec{E} \otimes \vec{E}^* \right\rangle = \left\langle \begin{bmatrix} E_x \\ E_y \end{bmatrix} \otimes \begin{bmatrix} E_x^* \\ E_y^* \end{bmatrix} \right\rangle = \begin{bmatrix} \langle E_x E_x^* \rangle \\ \left\langle E_x E_q^* \right\rangle \\ \langle E_y E_x^* \rangle \\ \left\langle E_y E_y^* \right\rangle \end{bmatrix} \tag{2.4.2.2}$$

With Eq. 2.4.2.1, one can write the transmission of the electric field through an instrument, described by its Jones matrix, as

$$\vec{C}_o = \left\langle \mathbf{J}\vec{E} \otimes \mathbf{J}^* \vec{E}^* \right\rangle = \langle \mathbf{J} \otimes \mathbf{J}^* \rangle \left(\vec{E} \otimes \vec{E}^* \right\rangle = \langle \mathbf{J} \otimes \mathbf{J}^* \rangle \, \vec{C} \tag{2.4.2.3}$$

Using Eq. 2.4.2.2, the polarisation state of this electric field can be described by the *Stokes vector* $\vec{S}$, given by:

$$\vec{S} = \begin{bmatrix} I \\ Q \\ U \\ V \end{bmatrix} = \begin{bmatrix} C_{xx} + C_{yy} \\ C_{xx} - C_{yy} \\ C_{xy} + C_{yx} \\ \mathrm{i}\,(C_{xy} - C_{yx}) \end{bmatrix} \tag{2.4.2.4}$$

From Eq. 2.4.2.2, we obtain the expression of the Stokes vector as $\vec{S} = \mathbf{A}\vec{C}$, where

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & -1 \\ 0 & 1 & 1 & 0 \\ 0 & i & -i & 0 \end{bmatrix} \tag{2.4.2.5}$$

Substituting Eq. 2.4.2.3 into $\vec{S}$, it follows that the outgoing Stokes vector $\vec{S}_{out}$ can be written as

$$\vec{S}_{out} = \mathbf{A}\left(\mathbf{J} \otimes \mathbf{J}^*\right)\mathbf{A}^{-1}\vec{S} = \mathbf{M}\vec{S}_{in} \tag{2.4.2.6}$$

where we have defined the Mueller matrix $\mathbf{M} = \mathbf{A}\left(\mathbf{J} \otimes \mathbf{J}^*\right)\mathbf{A}^{-1}$ that describes how the Stokes parameters transform:

$$\mathbf{M} = \begin{pmatrix} M_{II} & M_{IQ} & M_{IU} & M_{IV} \\ M_{QI} & M_{QQ} & M_{QU} & M_{QV} \\ M_{UI} & M_{UQ} & M_{UU} & M_{UV} \\ M_{VI} & M_{VQ} & M_{VU} & M_{VV} \end{pmatrix}. \tag{2.4.2.7}$$

For an instrument with several components, the Mueller matrix $\mathbf{M}$, is the product of the Mueller matrices for each component. Describing the instrument in terms of Mueller matrices is not only useful for characterizing the polarisation properties of the instrument, but it can also be used to characterize some instrument systematics.

2.4.3 Instrument systematics

Systematics refer to any bias or error that affects the measurements of an instrument [125], and can arise from various sources, such as instrument imperfections, environmental effects, or measurement techniques. These induce deviations of the measured parameter values from the unknown true values. When the deviation from the true value is assumed to be small, it can be approximated by a first-order expansion. In terms of Eq. 2.3.1.6, the uncertainties caused by these deviations can be written as:

$$d' = (g + \delta g)\,(H_t + \delta H_t) \otimes \left[n + \delta n + \frac{1}{2}\int d\nu (W + \delta W) \int d\Omega (p + \delta p)\right] \quad (2.4.3.1)$$

where the different δ's stand for different systematics: gain δg, transfer functions δH_t, noise models δn and band-pass function δW. Following the expression in 2.3.1.7, δp can be expanded as:

$$\begin{aligned}\delta p =& \delta B_{\parallel}[I + \epsilon(Q\cos 2\psi + U \sin 2\psi)] + B_{\perp}[I - \epsilon(Q\cos 2\psi + U\sin 2\psi)]\\ &+ B_{\parallel}[I + \epsilon(Q\cos 2\psi + U \sin 2\psi)] + \delta B_{\perp}[I - \epsilon(Q\cos 2\psi + U\sin 2\psi)]\\ &+ B_{\parallel}[\delta I + \epsilon(\delta Q\cos 2\psi + \delta U \sin 2\psi)] + B_{\perp}[\delta I - \epsilon(\delta Q\cos 2\psi + \delta U\sin 2\psi)]\\ &+ B_{\parallel}[I + \delta\epsilon(Q\cos 2\psi + U \sin 2\psi)] + B_{\perp}[I - \delta\epsilon(Q\cos 2\psi + U\sin 2\psi)],\end{aligned} \quad (2.4.3.2)$$

where each line corresponds to a different systematic error: in order, we have the co-polar beam, the cross-polar beam, the polarisation efficiency of the detector, and the polarisation orientation of the detector. I, as part of the Simons Observatory and LiteBIRD collaborations, have already started implementing requirements for some of the systematics mentioned above, [e.g., see 5–9, 131].

Additionally, the uncertainties in the Stokes parameters can be further studied as:

$$\delta S = \delta(RMR) S_{\text{in}} + RMR\delta S_{\text{in}} \quad (2.4.3.3)$$

where S is the vector of the Stokes parameters. The first term includes systematic errors in the HWP Mueller matrix, which, in the case of normal incidence, can be broken down into the rotation matrix and a Mueller matrix defined relative to the fixed axis of the HWP (although this decomposition is not feasible for oblique incidence). The uncertainties in the HWP Mueller matrix will be studied in Chapter 6, with a particular focus towards the polarisation efficiency and rotation terms of the matrix.

The uncertainties in the input Stokes parameters, δS_{in}, are due to the uncertainties present in the foreground spectrum shapes and components. These are analysed separately in the following Section.

2.5 Foregrounds

polarised radiation is not only generated by the CMB, but also by other astrophysical sources, such as dust and synchrotron emission from our own galaxy, as well as radio emission from extragalactic sources. s

2.5.1 Galactic foregrounds

Arguably the largest source of systematic uncertainty for both space and ground experiments will be contamination by the astrophysical polarised signal from the Milky Way. Fortunately, our Galaxy is transparent to CMB photons. However, the Milky Way's sub-millimeter emission, particularly near its galactic plane, can overshadow the CMB. Additionally, some components in the Milky Way's emission are polarised. The polarised Galactic emission is dominated by synchrotron and polarised thermal dust. The combination of both dominates the CMB polarisation signal on the scales of interest by at least a factor ~ 5 over the whole frequency range. Devising methods to separate this contamination from the CMB in a robust and efficient manner is therefore of central importance in this quest.

2.5.1.1 Thermal dust

The thermal dust emission is the most dominant galactic foreground at frequencies greater than 100GHz. This emission originates from the solid phase of the interstellar medium (ISM), composed of elements such as Carbon, Silicon, Magnesium, Iron, and Oxide, which predominantly form silicate and amorphous carbon molecules. These molecules[13] aggregate into dust grains that range in size from 0.005 to 0.25 μm [134] and have an asymmetric elongated shape, making them electric dipoles that partially align with the magnetic field of the ISM [135], giving rise to its polarisation. The dust grains align themselves with their long axes parallel to the magnetic field, resulting in a net polarisation signal along the direction of the magnetic field. Stellar photons heat these dust grains, which absorb the UV/optical light, and then release radiation in the mid-infrared to sub-millimeter-wavelength range (via internal vibrational and rotational modes).

This polarisation signal is typically modelled using the Davis-Greenstein formalism [135], which assumes that the alignment efficiency of the dust grains depends on the angle between the magnetic field direction and the grain's symmetry axis.

[13] Molecules making up dust include graphite (C), enstatite ($(\mathrm{Mg},\mathrm{Fe})\mathrm{SiO_3}$), olivine ($(\mathrm{Mg},\mathrm{Fe})_2\mathrm{SiO_4}$), silicon carbide (SiC), iron (Fe), and magnetude $\mathrm{Fe_3O_4}$ [132, 133]

The *polarisation fraction* of the thermal dust emission depends on the dust temperature, the magnetic field strength, and the properties of the dust grains. Locally, the polarisation fraction p and angle ϕ can be defined as:

$$p \equiv \frac{\sqrt{Q^2+U^2}}{T}, \quad \phi = \frac{1}{2}\arctan\left(\frac{U}{Q}\right). \tag{2.5.1.1}$$

This is typically highest at frequencies between 100 and 300 GHz, where the thermal dust emission dominates over other foregrounds. The Planck mission has constrained the upper limit $p = 19.8\% \pm 0.07$ [136]. In addition, an indicator of the level of organization in the orientation of the magnetic field in the plane of the sky is the *polarisation angle dispersion*:

$$\Delta\phi^2(\mathbf{n},\delta) \equiv \int d^2\Omega'\delta\left(|\mathbf{n}'|-\delta\right)\left(\phi(\mathbf{n})-\phi\left(\mathbf{n}+\mathbf{n}'\right)\right), \tag{2.5.1.2}$$

i.e. the integral taken over all possible directions $\mathbf{n}'$ of the difference in polarisation angle between two points separated by a distance δ on the sky. This quantity can be used to study the structure of the galactic magnetic field (GMF) in the plane of the sky [136, 137]. In particular, measurements of the dispersion angle reveal that dust is organized in filaments, which in turn can be used to probe the structure of the underlying magnetic fields[14].

A simple model for the dust emission at a given observation frequency ν is a *modified black-body*[15] spectrum, defined at reference frequency ν_0:

$$I_\nu^{\text{dust}} = A_d\left(\frac{\nu}{\nu_0}\right)^{\beta_d-2}\frac{B_\nu\left(T_d\right)}{B_{\nu_0}\left(T_d\right)} \tag{2.5.1.3}$$

where A_d is the dust amplitude, β_d is the dust spectral index, T_d is the dust temperature. The Planck satellite has measured the mean value across the sky of the spectral index and temperature for dust, $\beta_d = 1.51 \pm 0.01$ and $T_d \approx 19.6$ K ± 1.3 K, respectively. It should be noted that these values are obtained solely from total intensity measurements data, while the spectral index for polarisation was estimated at $\beta_d = 1.53 \pm 0.02$, with significant variations across the sky [138, 139].

Unfortunately, the polarisation properties of many foreground sources, including dust, are not well understood, and there is ongoing research to try to better characterize these sources. Advanced models are being developed to accurately replicate the physical effects that cause dust polarisation. These models are becoming

[14]The interaction between GMFs and foregrounds is an active area of research in galactic science. The ultimate goal is to develop a foreground model based on the three-dimensional structure of the underlying magnetic fields rather than on template observations for various foregrounds.

[15]Also referred to as a "grey-body" radiation with a modified Planck function.

increasingly sophisticated, with some incorporating multiple dust populations with varying characteristics, [140–142] while others use a statistical approach by coupling dust populations with the turbulent interstellar medium and galactic magnetic fields [143]. The importance of these refined models has intensified as scientists search for B-modes. Ultimately, advanced ISM models are essential to achieve accurate component separation.

Frequency decorrelation The Planck satellite's observations reveal a decorrelation between different frequency channels. The $B-$ mode cross-power spectra of thermal dust between frequency channels ν_1 and ν_2 can be modelled as

$$\ell(\ell+1)C_\ell^{BB}\left[\nu_1,\nu_2\right] = A_\ell^2\left(\frac{\nu_1\nu_2}{\nu_0^2}\right)^{\beta_d-2}\frac{B_{\nu_1}}{B_{\nu_0}}\frac{B_{\nu_2}}{B_{\nu_0}}f_d\left(\Delta_d,\nu_1,\nu_2\right), \tag{2.5.1.4}$$

where the *decorrelation parameter* Δ_d is defined in the function

$$f_d\left(\Delta_d,\nu_1,\nu_2\right) \equiv \exp\left(-\Delta_d\left(\ln\frac{\nu_1}{\nu_2}\right)^2\right). \tag{2.5.1.5}$$

Alternatively, one can define the correlation ratio

$$\mathcal{R}_\ell^{XY} = \frac{C_\ell^{XY}\left[\nu_1,\nu_2\right]}{\sqrt{C_\ell^{XY}\left[\nu_1,\nu_1\right]C_\ell^{XY}\left[\nu_2,\nu_2\right]}}. \tag{2.5.1.6}$$

According to references [139, 144], the Planck data shows no dust B-mode decorrelation between frequency channels, which is a crucial aspect for detecting primordial B-modes. This finding supports the possibility of measuring $r > 0.01$ with current upper-levels. However, it's important to note that frequency decorrelation can occur due to variations of the spectral indices across the sky, as well as along the line of sight, if not properly accounted for in the analysis [102].

2.5.1.2 Synchrotron

Synchrotron emission refers to the radiation produced by charged relativistic particles, typically electrons, that are accelerated in magnetic fields present in our own Galaxy. As the charged particles spiral around the magnetic field lines and lose energy, they emit synchrotron radiation.

The magnetic fields in our own Galaxy are primarily generated by Type Ib and Type II supernova remnants at small scales, while the Galactic magnetic field is believed to result from the shearing effect of differential rotation at larger scales. There are two main sources of synchrotron emission: electrons trapped in the magnetic fields of discrete supernova remnants and diffuse emission from cosmic ray electrons spread throughout the Galaxy. In supernova remnants, the magnetic field strength typically reaches around 75μG, while the diffuse Galactic magnetic field is about an order of magnitude fainter, ranging from $1-5\mu$G.

The energy spectrum of cosmic ray electrons is described by a relativistic electron number density distribution $N(E) \sim E^{-p}$. Since $N(E)$ and the magnetic field strength vary across the Galaxy, the resulting synchrotron emission exhibits a wide range of spectral behaviours. As a result, the observed morphology and spectral parameters of synchrotron sky maps change significantly with frequency and with the line of sight.

The power-law that describes optically thin synchrotron emission has a spectral index, denoted as α, which is related to the index of the electron-energy distribution, referred to as p, by the equation $\alpha = -(p+3)/2$. At radio frequencies in the few GHz range, discrete supernova remnants exhibit a spectral index of approximately $\alpha \sim -0.5$ (corresponding to $p \sim 2$), despite having stronger magnetic field strengths. However, they only contribute around 10% of the total synchrotron emission of the Galaxy at 1.5 GHz, [145]. On the other hand, the diffuse component of synchrotron emission, which accounts for more than 90% of the observed emission, has a direction-dependent spectral index that typically ranges from $-1.1 < \alpha < -0.5$ $(2.0 < p < 3.2)$ [146].

If synchrotron flux density scales as a power law also the equivalent brightness temperature follows a power-law:

$$I_\nu^{\mathrm{sync.}} = A_{\mathrm{sync.}} \left(\frac{\nu}{\nu_0}\right)^{\beta_{\mathrm{sync.}}}, \tag{2.5.1.7}$$

where the spectral index $\beta_{\mathrm{sync.}}$ determines the steepness of the change. At a frequency of 1 GHz, the spectral index is approximately -2.7 and becomes steeper at higher frequencies, reaching approximately -3.0 [147]. In the frequency range of 2.3 to 33 GHz, the spectral index can decrease to around -3.2 for polarisation [148]. Additionally, the spectral index is not constant across different directions in the sky, with values ranging from -4.4 to -2.5. Its spectral and spatial variability makes the modelling and removal of synchrotron emission in CMB B-mode experiments particularly challenging. A careful characterisation of the GMF and its spectral properties is required to accurately separate the synchrotron foreground from the CMB B-mode signal.

Synchrotron emission is subject to self-absorption, which occurs when a photon emitted by synchrotron radiation interacts with a charged particle and is absorbed. This phenomenon is prevalent in regions where synchrotron emission is optically thick, particularly at low frequencies. Self-absorption is independent of the electron-energy distribution index, p, and in the optically thick regime, the absorption coefficient scales as $I \sim \nu^{5/2}$.

When a single charge emits synchrotron radiation, the resulting wave is elliptically polarised. However, when considering a distribution of particles that varies smoothly with pitch angle (the angle between the particle's velocity and magnetic field), the elliptical component cancels out. Thus, the radiation becomes partially linearly polarised [149]. In the absence of non-uniform magnetic fields, the polarisation fraction of the total polarised intensity, can be described again by Eq. 2.5.1.1, and is typically higher at lower frequencies below 70 GHz, being the dominant radio sky emission below 40 GHz [147]. For non-uniform magnetic fields, p is related to β_s, by the expression

$$p = \frac{(3\beta_s + 3)}{(3\beta_s + 1)} \approx 0.7 \tag{2.5.1.8}$$

However, in practice, the degree of observed polarisation is generally less than 20%, as non-uniform magnetic field directions along the line of sight can reduce the polarisation [150].

However, synchrotron emission can also be depolarised due to Faraday rotation, the rotation of polarised light proportional to the magnetic field component that is parallel to the direction of propagation. The degree of Faraday rotation depends on the strength and orientation of the magnetic field, as well as the density and temperature of the plasma [151].

2.6 Focus of this thesis

Even though the only source of primordial large-scale B-modes at linear order are tensor fluctuations, in practice a measurement is complicated by several factors. Some of these are addressed in this thesis. First, diffuse Galactic foregrounds have significant polarised emission, and in particular foreground components such as synchrotron radiation and thermal emission from dust produce B-modes with a significant amplitude, as discussed in Sec. 2.5. Component separation methods, which exploit the different spectral energy distributions (SED) of the CMB and foregrounds to separate the different components, are thus of vital importance [152, 153]. To this end, several different component separation methods have been developed and discussed in Chapters 3, 4, 5. Practical implementations of these methods must also be able to carry out this separation in the presence of instrumental noise and systematic effects [e.g. 131, 154]. The treatment of specific classes of instrument systematic effects is also a central topic of this thesis, and is treated in detail in Chapter 6.

Chapter 3

A minimal power-spectrum-based moment expansion for CMB B-mode searches

Contents

3.1 Introduction

The primordial B-mode signal is extremely faint compared to other sources of B-modes. On the one hand, gravitational lensing by the intervening large-scale structure generates a B-mode contribution which, on small scales, is similar to a white-noise component with amplitude $\sigma_N = 5\,\mu K$ arcmin [155]. On the other hand, the most important astrophysical source is the contamination from Galactic polarised foregrounds [98]. In particular, polarised synchrotron dominates the sky emission at low frequencies ($\nu \lesssim 40\,\mathrm{GHz}$), while thermal dust emission is more relevant at higher frequencies ($\nu \gtrsim 150\,\mathrm{GHz}$). The combination of both, in any case, dominates over the CMB B-mode signal, including the lensing contribution, over the whole frequency range on degree scales. The separation of the multi-frequency data into different components is therefore a crucial part in the analysis of CMB B-mode data [98, 101, 156–159]. Manifestly optimal map-based component separation methods have been designed and applied to existing satellite datasets [160–165]. However, their implementation on high-resolution datasets is computationally challenging due to the larger number of modes. Although approximate parametric map-level methods exist (e.g. [166–168]), that can deal with large numbers of pixels more efficiently, their implementation on ground-based data is further complicated by the presence of complex filtering and inhomogeneous non-white noise introducing non-trivial correlations between pixels. Partly for this reason, multi-frequency C_ℓ-based approaches, where the signal is modelled directly at the level of the cross-frequency power spectra, and where some of these complications are easier to deal with (e.g. through the use of transfer functions [80]), have been developed and used for ground-based data [101, 102, 169, 170].

The main drawback of these methods, in their simplest incarnation, is the difficulty to account for the expected spatial variability of foreground spectral properties. Although the main effects of this spatial variability, in the form of frequency decorrelation [144, 171], can be effectively taken into account in specific cases (e.g. uncorrelated spectral index variations [143]), developing a framework to account for this variability in a general scenario will be useful in the analysis of data from ongoing and future ground-based experiments given their higher sensitivity. In this work, we will make use of the so-called moment expansion formalism, introduced in [172–175], to derive the simplest extension to the standard power spectrum-level parametrizations of foreground spatial variability, and quantify the ability of this method to obtain unbiased constraints on the tensor-to-scalar ratio from existing data and upcoming experiments. This problem was already addressed by [176], where the most general moment expansion was presented and applied to *Planck* data. In this work, we instead study the simplest version of this expansion, and evaluate its performance for next-generation ground-based CMB data using foreground simulations with varying levels of realism.

This chapter is structured as follows. Section 3.2 describes the basic elements of power-spectrum-based component separation methods, the main effects of spatially-varying foreground spectral properties, and introduce the moment expansion method we will use to account for those. This section also includes a more accurate, non-perturbative calculation of the impact of foreground spatial variations using a formalism similar to that used to estimate the effects of gravitational lensing on the primary CMB power spectrum. Section 3.3 describes the ingredients of the synthetic sky simulations used to validate the method. This validation is described in detail in Section 3.4, which presents the limits of applicability of the leading-order moment expansion and its performance on simulations with varying degrees of foreground complexity. After validating the method, using simulated observations mimicking the expected performance of the Simons Observatory (SO), we apply it to the public data from the BICEP2/*Keck* collaboration and obtain constraints on r marginalised over foreground spatial variations. We summarize and discuss our results in Section 3.5.

3.2 Formalism

3.2.1 Sky model

We model the polarised sky signal at position $\hat{\mathbf{n}}$ and frequency ν, $\mathbf{m}_\nu(\hat{\mathbf{n}}) \equiv (Q_\nu(\hat{\mathbf{n}}), U_\nu(\hat{\mathbf{n}}))$, as a sum of components of the form

$$\mathbf{m}_\nu(\hat{\mathbf{n}}) = \sum_c \mathbf{T}_c(\hat{\mathbf{n}})\, S^c_\nu(\vec{\beta}_c(\hat{\mathbf{n}})), \tag{3.2.1.1}$$

where $\mathbf{T}_c(\hat{\mathbf{n}})$ is the amplitude of component c at a pivot frequency ν_0^c, and S^c_ν is its frequency spectrum (normalized to $S^c_{\nu_0} = 1$). $\vec{\beta}_c(\hat{\mathbf{n}})$ is a set of parameters describing the spectrum, which can vary as a function of sky position.

We will consider three components, which are described in detail in Ch. 2:

- **CMB**: in antenna temperature units, the spectrum of CMB temperature anisotropies is
$$S^{\rm CMB}_\nu = e^x \left(\frac{x}{e^x - 1}\right)^2, \quad x = \frac{h\nu}{k_B T_{\rm CMB}}, \tag{3.2.1.2}$$
where h is the Planck constant, k_B is the Boltzmann constant, and $T_{\rm CMB} = 2.7255\,{\rm K}$ is the CMB monopole temperature [177]. The CMB spectrum is isotropic and is not normalized at any pivot frequency.

- **Thermal dust:** dust grains in the interstellar medium are heated by stellar radiation, producing emission on microwave frequencies. The alignment of elongated dust grains with the Galactic magnetic field (GMF) produces a linear polarisation perpendicular to both the magnetic field and the direction of propagation, making dust the most relevant foreground for B-mode searches on frequencies $\nu \gtrsim 150\,\mathrm{GHz}$ (for more detail, refer to Sec. 2.5). Thermal dust emission is well-characterized by a modified black-body (MBB) spectrum of the form [138]

$$S_\nu^{\mathrm{D}} = \left(\frac{\nu}{\nu_0^{\mathrm{D}}}\right)^{\beta_{\mathrm{D}}} \frac{B_\nu(T_{\mathrm{D}})}{B_{\nu_0^{\mathrm{D}}}(T_{\mathrm{D}})}, \tag{3.2.1.3}$$

where β_{D} and T_{D} are the dust spectral index and temperature, and

$$B_\nu(T) = \frac{2h\nu^3}{c^2}\left[\exp\left(\frac{h\nu}{kT}\right) - 1\right]^{-1} \tag{3.2.1.4}$$

is the Planck black-body spectrum. We will consider spatial variations in β_{D}, which takes values $\beta_{\mathrm{D}} \sim 1.6$. The restricted frequency range available to most ground-based experiments, including the SO ($\nu \lesssim 280\,\mathrm{GHz}$), makes B-mode studies almost insensitive to the value of T_{D}, and therefore we fix it to $T_{\mathrm{D}} = 19.6\,\mathrm{K}$ here. Note that departures from a pure MBB law are anticipated by dust models (e.g. [178, 179]).

- **Synchrotron:** Galactic synchrotron emission is caused by the interaction of high-energy cosmic ray electrons with the GMF [149]. Synchrotron is strongly polarised, and is characterized by a smooth power-law spectrum tracing the energy distribution of cosmic ray electrons (for more detail, refer to Sec. 2.5). The synchrotron spectrum used here is therefore

$$S_\nu^{\mathrm{S}} = \left(\frac{\nu}{\nu_0^{\mathrm{S}}}\right)^{\beta_{\mathrm{S}}}, \tag{3.2.1.5}$$

where β_{S} is the synchrotron spectral index, which takes values $\beta_{\mathrm{S}} \sim -3$.

The spatially-varying degrees of freedom of the sky model are therefore the amplitudes of the three components ($\mathbf{T}_{\mathrm{CMB}}(\hat{\mathbf{n}})$, $\mathbf{T}_{\mathrm{D}}(\hat{\mathbf{n}})$, $\mathbf{T}_{\mathrm{S}}(\hat{\mathbf{n}})$) and the two spectral indices ($\beta_{\mathrm{D}}(\hat{\mathbf{n}})$, $\beta_{\mathrm{S}}(\hat{\mathbf{n}})$), where we will use the notation 'D' to indicate dust parameters and 'S' for synchrotron.

A full characterization of these variables would ideally require a pixel-based component separation approach [e.g. 158, 161, 162, 167] where they are constrained in each pixel individually. We will touch upon these approaches in more detail in

Chapter 5. Unfortunately this approach would significantly degrade the uncertainties in the final B-mode constraints, due to the limited frequency coverage of ground-based experiments [158, 167] (e.g. 6 bands in the case of SO). Although "pooling" approaches have been proposed where single spectral indices are constrained in larger sky regions to reduce the number of degrees of freedom, the degradation of final constraints can still be significant [168], and the resulting component-separated amplitude maps often retain features associated with the choice of pooling. Moreover, especially in the case of ground-based data, map-based analyses can be extremely complex, given the various filtering operations that the final frequency maps are subjected to, which can lead to important non-linear biases in the component-separated maps [180]. On the other hand there exist well-tested methods to account for these effects when computing power spectra between different frequency maps (e.g. the use of cross-split correlations to avoid imperfections in modelling the noise bias, or the use of transfer functions or observation matrices to correct for the impact of map-level filtering operations [159, 170, 181]). For this reason, most cosmological analyses from ground-based data have used a "multi-frequency" likelihood where the analysis is based on modelling the full set of cross-frequency power spectra [89, 102, 159]. This will motivate us propagating the map-based model described here to the power spectrum level in Section 3.2.2.

We can make further progress by assuming that the spatial variations in the spectral parameters β_c^i across the mapped footprint are small, and using the "moments-based" expansion [175]. Expanding $\beta_c^i = \bar{\beta}_c^i + \delta\beta_c^i(\hat{\mathbf{n}})$, where $\bar{\beta}_c^i$ is the mean of β_c^i across the map, and $\delta\beta_c^i$ describes its spatial fluctuations, we can expand Eq. 3.2.1.1 to second order in $\delta\beta_c^i$ as

$$\mathbf{m}_\nu(\hat{\mathbf{n}}) = \sum_c \left[\mathbf{T}_c(\hat{\mathbf{n}})\bar{S}_\nu^c + \mathbf{T}_c(\hat{\mathbf{n}})\,\delta\beta_c^i(\hat{\mathbf{n}})\,\partial_i\bar{S}_\nu^c + \frac{1}{2!}\mathbf{T}_c(\hat{\mathbf{n}})\,\delta\beta_c^i(\hat{\mathbf{n}})\,\delta\beta_c^j(\hat{\mathbf{n}})\,\partial_i\partial_j\bar{S}_\nu^c + \mathcal{O}(\delta\beta^3)\right], \tag{3.2.1.6}$$

where the first order Taylor expansion is equivalent to the "Delta-map" method developed in [182, 183]. In Eq. 3.2.1.6 we have used the shorthand

$$\bar{S}_\nu^c = S_\nu^c(\bar{\beta}), \quad \partial_i\bar{S}_\nu^c \equiv \left.\frac{\partial S_\nu^c}{\partial\beta_c^i}\right|_{\beta=\bar{\beta}}, \quad \partial_i\partial_j\bar{S}_\nu^c \equiv \left.\frac{\partial^2 S_\nu^c}{\partial\beta_c^i\partial\beta_c^j}\right|_{\beta=\bar{\beta}}. \tag{3.2.1.7}$$

It is worth noting that the derivatives of the spectrum take a specific form in the case of spectral indices, which enter the spectrum as $(\nu/\nu_0)^\beta$:

$$\frac{\partial^n S_\nu^c}{\partial\beta^n} = \left[\log\left(\frac{\nu}{\nu_0^c}\right)\right]^n S_\nu^c. \tag{3.2.1.8}$$

In Eq. 3.2.1.1 and in what follows we have used Einstein's notation with respect to summation over repeated indices identifying spectral parameters (i, j etc.).

3.2.2 C_ℓ-based cleaning method

Let us now propagate the formalism above to the power spectrum between two frequency maps

$$\langle a^{\nu}_{\ell m} a^{\nu' *}_{\ell' m'} \rangle = \delta_{\ell\ell'}\delta_{mm'} C^{\nu\nu'}_{\ell}, \tag{3.2.2.1}$$

where $a^{\nu}_{\ell m}$ are the spherical harmonic coefficients associated with map $m_\nu(\hat{\mathbf{n}})$, which we have assumed is statistically isotropic. Note that, for simplicity, we will assume that all fields involved ($m_\nu(\hat{\mathbf{n}})$, $T_c(\hat{\mathbf{n}})$) are spin-0 fields, even though our main application, the polarisation of the CMB, is a spin-2 field. The generalization of our main result to fields of arbitrary spin is presented in Sec. 3.2.3, where we justify that neglecting the spin-2 nature is a good approximation in our analysis[1]. We will thus apply these results directly to the pseudo-scalar B-mode component.

Using Eq. 3.2.1.6, the power spectrum $C^{\nu\nu'}_\ell$ is, up to second order in $\delta\beta$, given by:

$$C^{\nu\nu'}_\ell = C^{\nu\nu'}_\ell|_{0\times0} + C^{\nu\nu'}_\ell|_{0\times1} + C^{\nu\nu'}_\ell|_{1\times1} + C^{\nu\nu'}_\ell|_{0\times2}, \tag{3.2.2.2}$$

where

$$C^{\nu\nu'}_\ell|_{0\times0} \equiv \sum_{cc'} \bar{S}^c_\nu\, \bar{S}^{c'}_{\nu'}\, C_\ell(T_c, T_{c'}), \tag{3.2.2.3}$$

$$C^{\nu\nu'}_\ell|_{0\times1} \equiv \sum_{cc'} \partial_i\bar{S}^c_\nu\, \bar{S}^{c'}_{\nu'}\, C_\ell(T_c\delta\beta^i_c, T_{c'}) + (\nu\leftrightarrow\nu'), \tag{3.2.2.4}$$

$$C^{\nu\nu'}_\ell|_{1\times1} \equiv \sum_{cc'} \partial_i\bar{S}^c_\nu\, \partial_j\bar{S}^{c'}_{\nu'}\, C_\ell(T_c\delta\beta^i_c, T_{c'}\delta\beta^j_{c'}), \tag{3.2.2.5}$$

$$C^{\nu\nu'}_\ell|_{0\times2} \equiv \frac{1}{2}\sum_{cc'} \partial_i\partial_j\bar{S}^c_\nu\, \bar{S}^{c'}_{\nu'}\, C_\ell(T_c\delta\beta^i_c\delta\beta^j_c, T_{c'}) + (\nu\leftrightarrow\nu'), \tag{3.2.2.6}$$

where $C_\ell(a,b)$ denotes the power spectrum between fields a and b, and $(\nu\leftrightarrow\nu')$ implies the same term swapping the roles of ν and ν'.

The 0×0 term, which is exact in the absence of spatial variations of spectral indices, is the basis for the fiducial power-spectrum level cleaning methods used by CMB collaborations [89, 102], and it involves the modelling of $N_c(N_c+1)/2$ power spectra for N_c different components. Assuming one single free spectral parameter for each component, the 0×1, 1×1 and 0×2 terms would in general imply modelling an additional N_c^2, $N_c(N_c+1)/2$ and N_c^2 different power spectra respectively (for a total of $N_c(3N_c+1)$ spectra). These individual contributions are associated to distinct spectral responses in terms of the combination of spectrum derivatives that accompany them, and it should therefore be possible to separate

[1]This is further reinforced by the good agreement of our formalism with Gaussian simulations presented in Section 3.4.2.

all contributions cleanly given a large enough number of frequency channels N_ν (leading to $N_\nu(N_\nu+1)/2$ distinct cross-frequency correlations). For example, on a bandpower-by-bandpower basis, with two foreground sources, associated with a single spectral parameter each, as well as the CMB, this approach would involve modelling 15 different cross-spectra (compared with 4 spectra if the higher-order terms are neglected). An experiment with 6 frequency channels would be able to measure 21 different cross-frequency spectra, and should therefore be able to separate the 15 different contributions, albeit at a significant cost in statistical uncertainties [176]. In the presence of more than one free spectral parameter per component, the number of independent spectra to model would increase rapidly.

With the objective of ameliorating the complexity of this model, we will make the following three simplifying assumptions:

1. In the 1×1 and 0×2 terms, we will assume that different components are uncorrelated, i.e. terms like $C_\ell(T_{\rm S}\delta\beta_{\rm S}, T_{\rm D}\delta\beta_{\rm D})$ are zero. We know that this assumption is wrong at some level. polarised dust and synchrotron are associated to the same GMF, and there is evidence that they are correlated on large scales. We will indeed account for this correlation in the 0×0 term but ignore it in the higher-order ones, effectively treating the cross-component correlation coefficient as another perturbative parameter.

2. We will assume Gaussian statistics for all fields involved. This automatically cancels the 0×1 term, which only contains 3-point functions. Furthermore this allows us to express all four-point functions in Eqs. 3.2.2.5 and 3.2.2.6 as products of two-point functions. Galactic foregrounds are well-known to be non-Gaussian, although the assumption of Gaussianity may be a better approximation in polarisation than intensity [143]. Nevertheless, the rationale is the same as before, treating the non-Gaussian foreground terms as higher-order in the perturbative expansion.

3. We will assume that amplitudes T_c and spectral index fluctuations $\delta\beta_c$ are uncorrelated, therefore ignoring terms of the form $C_\ell(T_c, \delta\beta_c)$. This is also not generally true, since variations in spectral indices are likely to trace the same structures (e.g. dust filaments) that generate the foreground signals. As before, we ignore these correlations, treating the associated correlation coefficients as additional perturbative parameters that would make them higher-order in the expansion.

These assumptions yield the simplest possible description of the multi-frequency power spectrum at leading order in the spatial variation of the foreground spectral parameters. For the specific model used here (synchrotron and dust with free spectral indices), the three surviving contributions to Eq. 3.2.2.2 read:

$$C_\ell^{\nu\nu'}|_{0\times0} = \bar{S}^{\rm D}_\nu \bar{S}^{\rm D}_{\nu'} C_\ell^{\rm DD} + \bar{S}^{\rm S}_\nu \bar{S}^{\rm S}_{\nu'} C_\ell^{\rm SS} + \left(\bar{S}^{\rm D}_\nu \bar{S}^{\rm S}_{\nu'} + \bar{S}^{\rm S}_\nu \bar{S}^{\rm D}_{\nu'}\right) C_\ell^{\rm SD}, \tag{3.2.2.7}$$

$$C_\ell^{\nu\nu'}|_{1\times1} = \sum_{c\in\{\rm D,S\}} \partial_\beta \bar{S}^c_\nu\, \partial_\beta \bar{S}^c_{\nu'} \sum_{\ell_1\ell_2} \frac{(2\ell_1+1)(2\ell_2+1)}{4\pi} \begin{pmatrix} \ell & \ell_1 & \ell_2 \\ 0 & 0 & 0 \end{pmatrix}^2 C_{\ell_1}^{cc} C_{\ell_2}^{\beta_c}, \tag{3.2.2.8}$$

$$C_\ell^{\nu\nu'}|_{0\times2} = \sum_{c\in\{\rm D,S\}} \frac{1}{2}\left[\bar{S}^c_\nu\, \partial^2_\beta \bar{S}^c_{\nu'} + \bar{S}^c_{\nu'}\, \partial^2_\beta \bar{S}^c_\nu\right] C_\ell^{cc} \sigma^2_{\beta_c}, \tag{3.2.2.9}$$

where we have used the shorthand

$$C_\ell^{cc} \equiv C_\ell(T_c, T_c), \quad C_\ell^{\beta_c} \equiv C_\ell(\beta_c, \beta_c), \quad \sigma^2_{\beta_c} \equiv \sum_\ell \frac{2\ell+1}{4\pi} C_\ell^{\beta_c}. \tag{3.2.2.10}$$

At the cost of generality, the method is therefore significantly simpler, requiring only the modelling of two additional power spectra, $C_\ell^{\beta_{\rm D}}$ and $C_\ell^{\beta_{\rm S}}$. We have explored the impact of the simplifying assumptions used here through the analysis of idealized and realistic simulations, as well as precursor data, as described in Section 3.4.

We parametrize the different ingredients described above by extending the model used by [102] for the 0×0 contribution. The amplitude power spectra are modelled as power laws of the form

$$\frac{\ell(\ell+1)}{2\pi} C_\ell^{cc} = A_c \left(\frac{\ell}{\ell_0}\right)^{\alpha_c}, \tag{3.2.2.11}$$

with the dust-synchrotron cross-correlation parametrized through a scale-independent correlation coefficient: $C_\ell^{\rm SD} = \epsilon_{\rm DS}\sqrt{C_\ell^{\rm DD} C_\ell^{\rm SS}}$. Likewise, the power spectrum of $\delta\beta_c$ is parametrised as

$$C_\ell^{\beta_c} = B_c \left(\frac{\ell}{\ell_0}\right)^{\gamma_c}, \tag{3.2.2.12}$$

with $\ell_0 = 80$ in all cases. Finally, the CMB B-mode power spectrum is parametrized as

$$C_\ell^{\rm CMB} = A_{\rm lens} C_\ell^{\rm lens} + r\ C_\ell^{\rm tens}\Big|_{\rm r=1}, \tag{3.2.2.13}$$

where $C_\ell^{\rm lens}$ and $C_\ell^{\rm tens}|_{\rm r=1}$ are templates for the B-mode power spectrum caused by gravitational lensing and by primordial tensor fluctuations with tensor-to-scalar ratio $r = 1$ respectively. The model therefore has 13 free parameters:

$$\{r, A_{\rm lens}, A_{\rm D}, \alpha_{\rm D}, \beta_{\rm D}, B_{\rm D}, \gamma_{\rm D}, A_{\rm S}, \alpha_{\rm S}, \beta_{\rm S}, B_{\rm S}, \gamma_{\rm S}, \epsilon_{\rm SD}\}. \tag{3.2.2.14}$$

3.2.3 Spin-s generalization of the first-order expansion

The expressions in Eqs. 3.2.2.7, 3.2.2.8 and 3.2.2.9 can be easily generalized to the case of spin-s quantities (as is the case for the spin-2 CMB polarisation field). Using the same formalism and notation presented in [79], the map $\mathbf{m}_\nu$ and amplitudes $\mathbf{T}_c$ in Eq. 3.2.1.6 are promoted to spin-s fields with both Q and U components in real space, and E and B-mode components in harmonic space:

$$\begin{aligned} m_\nu &\longrightarrow \mathbf{m}_\nu \equiv (m_\nu^Q, m_\nu^U), \\ a_{\ell m} &\longrightarrow \mathbf{a}_{\ell m} \equiv (a_{\ell m}^E, a_{\ell m}^B), \end{aligned}$$

while the spectral index variations $\delta\beta_c$ remain scalar, real-valued fields. The power spectrum in Eq. 3.2.2.1 between any two spin-s field then becomes a 2×2 matrix containing the four correlations between their E and B components:

$$\langle \mathbf{a}_{\ell m} \mathbf{b}^\dagger_{\ell' m'} \rangle \equiv \delta_{\ell\ell'}\delta_{mm'} \mathsf{C}_\ell^{ab} \equiv \delta_{\ell\ell'}\delta_{mm'} \begin{pmatrix} C_\ell^{a^E b^E} & C_\ell^{a^E b^B} \\ C_\ell^{a^B b^E} & C_\ell^{a^B b^B} \end{pmatrix}. \tag{3.2.3.1}$$

Following the same techniques used in the derivation of the standard pseudo-C_ℓ estimator (see e.g. [79]), it is easy to show that the expressions in Eqs. 3.2.2.7 and 3.2.2.9 for the 0×0 and 0×2 contributions remain formally unchanged:

$$\mathsf{C}_\ell^{\nu\nu'}|_{0\times 0} = \bar{S}_\nu^{\rm D} \bar{S}_{\nu'}^{\rm D}\, \mathsf{C}_\ell^{\rm DD} + \bar{S}_\nu^{\rm S} \bar{S}_{\nu'}^{\rm S}\, \mathsf{C}_\ell^{\rm SS} + \left(\bar{S}_\nu^{\rm D} \bar{S}_{\nu'}^{\rm S} + \bar{S}_\nu^{\rm S} \bar{S}_{\nu'}^{\rm D}\right) \mathsf{C}_\ell^{\rm SD}, \tag{3.2.3.2}$$

$$\mathsf{C}_\ell^{\nu\nu'}|_{0\times 2} = \sum_{c\in\{\rm D,S\}} \frac{1}{2} \left[\bar{S}_\nu^c\, \partial_\beta^2 \bar{S}_{\nu'}^c + \bar{S}_{\nu'}^c\, \partial_\beta^2 \bar{S}_\nu^c\right] \mathsf{C}_\ell^{cc} \sigma_{\beta_c}^2, \tag{3.2.3.3}$$

while the 1×1 term in Eq. 3.2.2.8 becomes

$$\mathsf{C}_\ell^{\nu\nu'}|_{1\times 1} = \sum_{c\in\{\rm D,S\}} \partial_\beta \bar{S}_\nu^c\, \partial_\beta \bar{S}_{\nu'}^c \sum_{\ell_1\ell_2} \frac{(2\ell_1+1)(2\ell_2+1)}{4\pi} \begin{pmatrix} \ell & \ell_1 & \ell_2 \\ s & -s & 0 \end{pmatrix}^2 C_{\ell_2}^{\beta_c}\, \hat{\mathsf{d}}_{S_\ell} \mathsf{C}_{\ell_1}^{cc} \hat{\mathsf{d}}_{S_\ell}^\dagger, \tag{3.2.3.4}$$

where $S_\ell \equiv \ell + \ell_1 + \ell_2$, and we have defined the matrix

$$\hat{\mathsf{d}}_n \equiv \frac{1}{2} \begin{pmatrix} 1+(-1)^n & -i[1-(-1)^n] \\ i[1-(-1)^n] & 1+(-1)^n \end{pmatrix}. \tag{3.2.3.5}$$

Thus we see that the main additional effect of the spectral index variation on the spin-2 polarised foregrounds is the mixing of E and B modes in the 1×1 term. This is similar to the generation of CMB lensing B modes from primordial E modes, which motivates the analogy with CMB lensing used in Sec. 3.2.4, or to the leakage between E and B in the presence of a sky mask. In the case of the featureless foreground power spectra considered here, this effect is degenerate with the unknown amplitude of the spectral index variations, and therefore we find the spin-0 approximation described in the previous section accurate enough for our main analysis, as will be demonstrated in Sec. 3.4.2.

3.2.4 Non-perturbative calculation and CMB lensing

If the spectral properties of a given component were perfectly homogeneous, the maps of that component at different frequencies would be simply rescaled versions of the same field. The presence of spatially-varying spectral parameters, however, causes additional perturbations on this field at different frequencies. The situation is similar to the basic description of the lensed CMB: an unperturbed field (in this case the primordial CMB fluctuation) is perturbed by a non-linear modification caused by another field (the lensing deflection). Furthermore, the minimal model explored here, in which both foreground amplitudes and spectral index fluctuations are treated as uncorrelated Gaussian random fields, makes the analogy between both phenomena almost exact. Because of this, we can use some of the methods developed within the context of CMB lensing to improve on the moment expansion method presented above. In particular, this section presents a full, non-perturbative calculation of the multi-frequency power spectrum for a given component in analogy to the calculation of the lensed CMB power spectrum.

Let us start by considering a single component c with amplitude $T_c(\hat{\mathbf{n}})$ at a pivot frequency ν_0, and a spectrum of the form $S^c_\nu(\beta_c) = (\nu/\nu_0)^{\beta_c} F^c_\nu$, where F_ν is an arbitrary function of frequency normalised to $F^c_{\nu_0} = 1$, and $\beta_c(\hat{\mathbf{n}}) = \bar{\beta}_c + \delta\beta_c(\hat{\mathbf{n}})$ is the component's spatially-varying spectral index with mean $\bar{\beta}_c$. At a frequency ν, the component's sky emission is

$$T_{c,\nu}(\hat{\mathbf{n}}) = S^c_\nu(\beta_c(\hat{\mathbf{n}}))\, T_c(\hat{\mathbf{n}}) = \bar{S}^c_\nu\, T_c(\hat{\mathbf{n}})\, e^{x_\nu\, \delta\beta(\hat{\mathbf{n}})}, \tag{3.2.4.1}$$

where $\bar{S}^c_\nu \equiv S^c_\nu(\bar{\beta}_c)$, and we have defined $x_\nu \equiv \log(\nu/\nu_0)$.

The multi-frequency correlation function of the perturbed field is defined as

$$\xi^{\nu\nu'}(\theta) \equiv \langle T_{c,\nu}(\hat{\mathbf{n}})\, T_{c,\nu'}(\hat{\mathbf{n}}')\rangle = \bar{S}^c_\nu\, \bar{S}^c_{\nu'}\, \langle T_c(\hat{\mathbf{n}})\, T_c(\hat{\mathbf{n}}')\rangle \left\langle e^{x_\nu\, \delta\beta(\hat{\mathbf{n}}) + x_{\nu'}\, \delta\beta(\hat{\mathbf{n}}')}\right\rangle, \tag{3.2.4.2}$$

where $\cos\theta \equiv \hat{\mathbf{n}} \cdot \hat{\mathbf{n}}'$ and, in the second equality, we have assumed that T_c and $\delta\beta_c$ are statistically independent.

The second expectation value can be calculated analytically using the following well-known result for Gaussian variables:

$$\langle e^y\rangle \equiv \int_{-\infty}^{\infty} dy \frac{e^{-y^2/(2S^2)}}{\sqrt{2\pi S^2}} e^y = e^{S^2/2}, \tag{3.2.4.3}$$

where S^2 is the variance of y. Applying this result to the last term in Eq. 3.2.4.2 we obtain

$$\xi^{\nu\nu'}(\theta) = \bar{S}^c_\nu\, \bar{S}^c_{\nu'}\, \xi^{cc}(\theta)\, \exp\left[(x_\nu^2 + x_{\nu'}^2)\frac{\sigma^2_{\beta_c}}{2} + x_\nu x_{\nu'}\, \xi^\beta(\theta)\right], \tag{3.2.4.4}$$

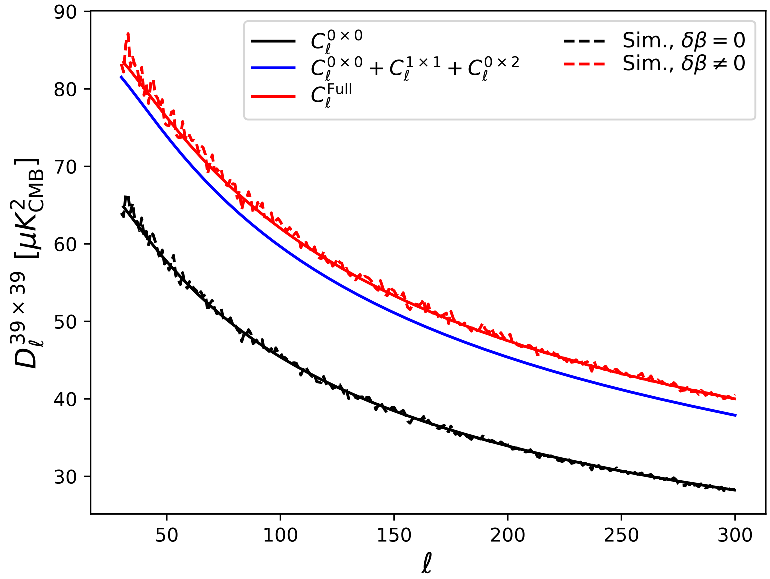

Figure 3.1: Average dust power spectrum at 39 GHz from 100 Gaussian simulations with spectral index variation $\sigma_{\beta_{\rm D}}=0$ (dashed black) and $\sigma_{\beta_{\rm D}}=0.16$ (dashed red) with a pivot $\nu_0^{\rm D}=353\,$GHz. The solid black line shows the theoretical prediction assuming no spectral index variation, while the solid blue line shows the leading-order moment expansion used in this work. The solid red line shows the exact calculation described in Equations 3.2.4.4, 3.2.4.6 and 3.2.4.7.

where we have defined the unperturbed correlation function and the spectral index correlation function:

$$\xi^{cc}(\theta)\equiv\langle T_c(\hat{\mathbf{n}})T_c(\hat{\mathbf{n}}')\rangle\,,\quad \xi^{\beta}(\theta)\equiv\langle\delta\beta_c(\hat{\mathbf{n}})\,\delta\beta_c(\hat{\mathbf{n}}')\rangle\,. \tag{3.2.4.5}$$

Correlation functions and power spectra of scalar fields are related to each other through:

$$\xi^X(\theta)=\sum_{\ell=0}^{\infty}\frac{2\ell+1}{4\pi}L_\ell(\cos\theta)\,C_\ell^X\simeq\int_0^\infty\frac{d\ell\,\ell}{2\pi}J_0(\ell\theta)\,C_\ell^X, \tag{3.2.4.6}$$

$$C_\ell^X=2\pi\int_0^\pi d(\cos\theta)\,L_\ell(\cos\theta)\,\xi^X(\theta)\simeq2\pi\int_0^\infty d\theta\,\theta\,J_0(\ell\theta)\,\xi^X(\theta), \tag{3.2.4.7}$$

where the second equality in each line is valid in the flat-sky approximation.

Given a model for the power spectra of the foreground amplitude and spectral index variations C_ℓ^{cc} and $C_\ell^{\beta_c}$, the multi-frequency power spectrum can be computed at all orders in $\delta\beta$ through a 3-step process:

1. Calculate $\xi^{cc}(\theta)$ and $\xi^\beta(\theta)$ (and $\sigma_\beta^2\equiv\xi^\beta(0)$) from C_ℓ^{cc} and C_ℓ^β using Eq. 3.2.4.6.

2. Calculate $\xi^{\nu\nu'}(\theta)$ from $\xi^{cc}(\theta)$ and $\xi^\beta(\theta)$ using Eq. 3.2.4.4.

3. Calculate $C_\ell^{\nu\nu'}$ from $\xi^{\nu\nu'}(\theta)$ using Eq. 3.2.4.7.

The Hankel transforms translating between correlation functions and power spectra (Eqs. 3.2.4.6 and 3.2.4.7) can be calculated using computationally efficient methods (e.g. [184]).

Figure 3.1 shows the improvement of this method over the lowest-order moment expansion (Eq. 3.2.2.2). The figure shows the dust power spectrum at 39 GHz for $\sigma_{\beta_{\rm D}} = 0.16$ calculated from the average of 100 Gaussian simulations (dashed red), the zeroth-order approximation $C_\ell^{\nu\nu'}|_{0\times 0}$ (black), the lowest-order moment expansion (Eq. 3.2.2.2, blue), and the full calculation described here (solid red). The non-perturbative calculation is able to recover the simulated power spectrum exactly.

Two interesting limits can be explored in Eq. 3.2.4.4. First, for small spectral index variations, the first term in the Taylor expansion of the exponential factor in this equation leads to two contributions, proportional to σ_β^2 and $\xi^\beta(\theta)$. These correspond to the real-space versions of the moment expansion terms $C_\ell^{\nu\nu'}|_{0\times 2}$ and $C_\ell^{\nu\nu'}|_{1\times 1}$ respectively. Secondly, in the limit of uncorrelated spectral index variations ($C_\ell^{\beta_c} = \Omega_{\rm pix}\sigma^2_{\beta_c}$, where $\Omega_{\rm pix}$ is the pixel size), the frequency decorrelation is scale-independent, and given by [143]

$$\frac{C_\ell^{\nu\nu'}}{\sqrt{C_\ell^{\nu\nu}C_\ell^{\nu'\nu'}}} = \exp\left[-\frac{1}{2}\sigma^2_{\beta_c}\Omega_{\rm pix}\log^2\left(\frac{\nu}{\nu'}\right)\right], \tag{3.2.4.8}$$

recovering the parametrisation in terms of decorrelation parameter used by [102].

Even with the use of fast methods for Hankel transforms, implementing the full model slows down the computation of the likelihood used here significantly, and therefore all our results use the moment expansion. As we have shown in section 3.4.2, the accuracy of the expansion is sufficient for the range of frequencies and sensitivities explored here. Furthermore, although the full calculation yields unbiased results, its applicability is fairly limited to the case of Gaussian spectral index variations in two dimensions that are statistically uncorrelated with the foreground amplitudes. The inaccuracies associated with these assumptions are likely to be more important than the differences with the moment expansion calculation. That being said, the intuition gained from the CMB lensing analysis could be useful for other applications, such as employing lensing reconstruction techniques to multi-frequency maps in order to recover spectral index maps (see e.g. [185]).

Parameter	Prior	Bounds
r	Top-hat	[-1, 1]
$A_{\rm lens}$	Top-hat	[0, 2]
A_D	Top-hat	[0, ∞]
α_D	Top-hat	[-1, 0]
β_D	Gaussian	1.6 $\pm$ 0.5
γ_D	Top-hat	[-6, -2]
B_D	Top-hat	[0, 10]
A_S	Top-hat	[0, ∞]
α_S	Top-hat	[-1, 0]
β_S	Gaussian	-3 $\pm$ 0.6
B_S	Top-hat	[0, 10]
γ_S	Top-hat	[-6, -2]
ϵ_{SD}	Top-hat	[-1, 1]

Table 3.1: Summary of the priors used in the analysis. Note that all amplitude parameters (except r) have physically-motivated positivity priors.

3.2.5 Power spectrum likelihood

To derive constraints on the free parameters of the model, we will use a multi-frequency power spectrum likelihood (see Sec. 2.3.3 for more details). In this case, the data vector is the full matrix of cross-frequency power spectra $C_\ell^{\nu\nu'}$, with the corresponding theory prediction described in the previous section. As discussed in Sec. 2.3.3, the central limit theorem becomes invalid at large scales causing power spectra to exhibit non-Gaussian features in their likelihoods. To account for this effect, we use the non-Gaussian likelihood developed by [121] (HL hereon). This likelihood requires an estimate of the covariance matrix of the full set of power spectra. In order to accurately account for the effects of incomplete sky coverage and EB leakage, we estimate this covariance matrix from a set of 500 Gaussian simulations, generated as described in Section 3.3.2. The HL likelihood additionally requires an estimate of the fiducial power spectra, as well as the noise power spectrum. We produce the former from the fiducial set of parameters used to generate the simulations (see Section 3.3.2), and the latter by averaging the power spectra of the 500 noise realisations generated for these simulations. Atmospheric noise and various systematics will likely limit the largest scales that can be reliably used by ground-based experiments, and therefore, in addition to the use of the realistic noise curves described in Section 3.3.4, we use the restricted scale range $30 \leq \ell \leq 300$ where B-mode signal from the recombination bump is concentrated [102].

The posterior distribution is given by the product of this likelihood and a set of priors. The priors have been chosen to be wide enough that the parameters are constrained by the data in most cases. They are summarised in Table 3.1. Note that we impose a physically motivated priors on all power spectrum amplitudes (A_c, B_c) forcing them to be positive. An exception to this is r, which we allow to be negative in order to detect possible negative biases. While negative r values do not make sense physically, they may result from e.g. volume effects caused by choosing specific priors on other parameters that we have marginalised over. Opening the prior on r to negative values allows us to monitor these unwanted effects, offering a simple robustness check.

We sample this posterior distribution using the affine-invariant Monte-Carlo Markov chain ensemble sampler `emcee` [122]. The chains were started around the maximum likelihood point, found via Powell's minimisation scheme [186] as implemented in `SciPy` [187]. The likelihood evaluations take about 10 milliseconds per iteration without the moment terms, and order 100ms with them, due to the evaluation of the sum in the 1×1 term 3.2.3.4. Further optimisations were not necessary as a converged MCMC chain can be completed in a few days.

The full pipeline that was implemented for this study, `BBPipe`, is publicly available[2], of which the more updated version is included in the submodule `BBPower`[3]

3.3 Simulations

In order to test the validity of the moment expansion method described in the previous section, we test it on a suite of sky simulations. These simulations were created using the publicly available code `BBSims`[4]. These simulations include the most relevant sky components, as discussed in Section 3.2.1, with varying degrees of realism in order to explore the impact of the assumptions of the method regarding foreground properties on its performance. The simulations also incorporate the contribution from instrumental noise and limited sky coverage for a SO-like experiment, as described in Section 3.3.4.

3.3.1 CMB

In all simulations, the CMB contribution was generated as a Gaussian random field drawn from the power spectrum in Eq. 3.2.2.13, given by the *Planck* 2018 best-fit ΛCDM parameters. Our baseline model does not include any primordial tensor signal ($r = 0$) but incorporates lensing power in the BB spectra ($A_{\rm lens} = 1$).

[2] See github.com/simonsobs/BBPipe.
[3] See github.com/simonsobs/BBPower.
[4] See github.com/susannaaz/BBSims.

3.3.2 Gaussian foreground simulations

We generate a large suite of "Gaussian" simulated skies. For these, we simulate sky maps for the amplitude and spectral index variation maps ($\mathbf{T}_c(\hat{\mathbf{n}})$ and $\delta\beta_c(\hat{\mathbf{n}})$) as Gaussian random fields governed by power spectra following the power-law models in Eqs. 3.2.2.11 and 3.2.2.12. We then add the mean spectral indices $\bar{\beta}_c$ to $\delta\beta_c(\hat{\mathbf{n}})$ and use the templates of the Python Sky Model software (`PySM`[5], [188]) to generate observed sky maps in the six SO frequency channels (see Section 3.3.4). These maps are generated using the `HEALPix`[6] pixelisation scheme [110] with resolution parameter $N_{\rm side} = 256$. The pixel resolution ($\delta\theta \sim 0.2°$) is enough to cover the ℓ range relevant for our analysis. Amplitude and spectral index maps were generated as uncorrelated Gaussian fields.

The aim of these Gaussian simulations is twofold. First, by using the exact same model assumed by the cleaning method (Gaussian fields, uncorrelated indices and amplitudes), we can test the validity of the leading-order expansion in $\delta\beta_c$ for different levels of spectral index variation, and compare it with the full, non-perturbative result. Secondly, the HL likelihood used here (see Section 3.2.5) requires an estimate of the power spectrum covariance. In order to fully incorporate the effects of inhomogeneous noise, mode-coupling and E/B mixing in the covariance matrix, we use these Gaussian simulations to compute it.

To estimate covariance matrices, we generate two suites of 500 Gaussian simulations. Both suites were generated with constant spectral indices ($B_{\rm D} = B_{\rm S} = 0$) but with different values for the amplitude power spectrum parameters. The first suite used the best-fit foreground parameters found by the BICEP2/*Keck* collaboration [102]:

$$A_{\rm D} = 5\,\mu{\rm K}^2, \quad \alpha_{\rm D} = -0.42, \quad A_{\rm S} = 2\,\mu{\rm K}^2, \quad \alpha_{\rm S} = -0.6,$$

and was used in the analysis of the method's performance as a function of spectral index variation amplitude. The second suite used the foreground parameters that best fit the dust and synchrotron template maps used in the realistic set of simulations described in Section 3.3.3, and were used to both validate the method against realistic simulations as described in Section 3.4.3, and in the simulation challenge described in Section 3.4.4. The corresponding parameter values are

$$A_{\rm D} = 28\,\mu{\rm K}^2, \quad \alpha_{\rm D} = -0.16, \quad A_{\rm S} = 1.6\,\mu{\rm K}^2, \quad \alpha_{\rm S} = -0.93. \tag{3.3.2.1}$$

[5] See github.com/bthorne93/PySM_public
[6] http://healpix.sf.net

All the Gaussian simulations presented here used the same values for the constant spectral indices ($\beta_{\rm D} = 1.6$, $\beta_{\rm S} = -3$) and the same pivot frequencies ($\nu_0^{\rm D} = 353\,{\rm GHz}$ and $\nu_0^{\rm S} = 23\,{\rm GHz}$). The dust-synchrotron correlation coefficient was set to $\epsilon_{\rm SD} = 0$.

To study the performance of the moment expansion method for different levels of spectral index variation, we generate a number of additional simulations with spectral index maps generated as Gaussian realisations of $C_\ell^{\beta_c}$ in Eq. 3.2.2.12 with varying values for the amplitude B_c. Instead of varying B_c directly, we generate maps of $\delta\beta_c(\hat{\mathbf{n}})$ with an arbitrary amplitude and then renormalize them to enforce a given per-pixel standard deviation $\sigma(\beta_c) \equiv \sqrt{\langle \delta\beta_c^2 \rangle}$. Therefore our results will be presented in terms of $\sigma(\beta_c)$ as a more meaningful parameter, rather than B_c. Unless otherwise stated, we fix the spectral tilt γ_c to the arbitrary values $(\gamma_{\rm D}, \gamma_{\rm S}) = (-3.5, -2.5)$. Note that we will also study the impact of the value of γ_c on the results, since this parameter regulates the distribution of spectral index fluctuations on different scales.

In summary, the different types of simulations include the two suites used to generate the covariance matrix discussed above, Gaussian foreground realisations with varying spectral index amplitudes described in Section 3.4.2, "realistic" foreground simulations introduced in the following Section and described in Section 3.4.3 and simulations with varying degrees of foreground complexity used in the "simulation challenge" described in Section 3.4.4.

3.3.3 Realistic foreground simulations

In order to test the validity of the assumptions adopted by our simplified moment expansion (Gaussianity, independence of spectral index and amplitudes, and between components), we produce an additional set of simulations with higher level of realism.

- **`d1s1` model**: We use the foreground templates provided by `PySM`. Specifically, we use the dust amplitude map and spectral index map assumed by the `d1` model, coming from *Planck* high frequency observations, as well as the synchrotron amplitude map included in the `s1` model, coming from *WMAP* 23 GHz maps. The synchrotron spectral index map provided with `s1` was originally derived from a combination of the Haslam 408 GHz map [189, 190] and the WMAP 23 GHz map [191, 192]. The map presents very mild fluctuations in $\beta_{\rm S}$, and only on very large scales. The level of variation and the overall value of $\beta_{\rm S}$ in this model has been shown by [148] to be too low using data from the SPASS experiment.

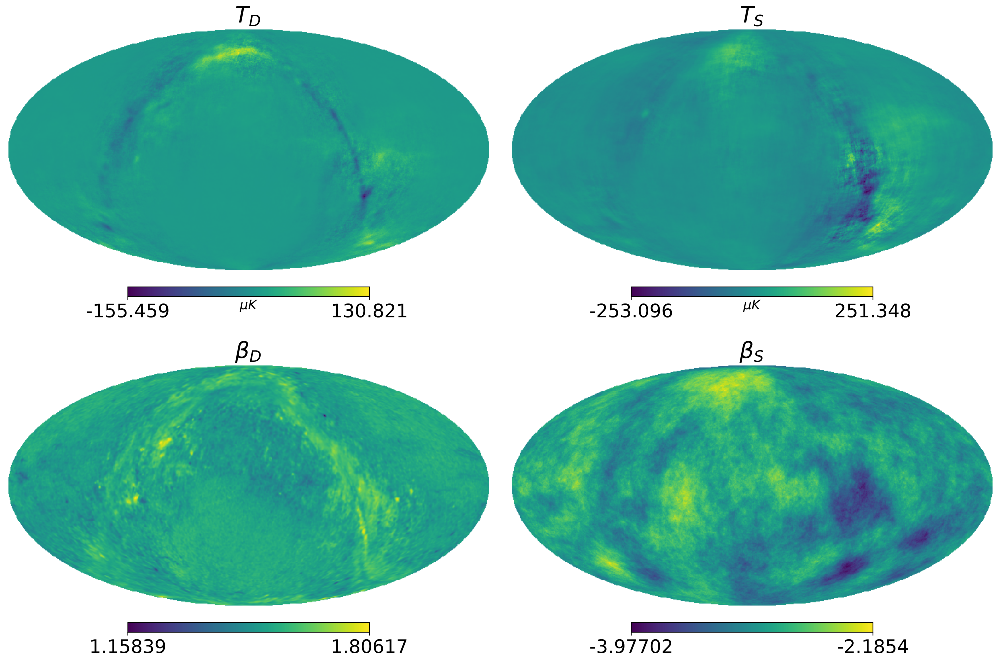

Figure 3.2: Dust and synchrotron polarised amplitude maps (*top row*) and spectral index maps (*bottom row*).

- `d1sm` **model**: To increase the complexity of the synchrotron contribution we instead generate an alternative spectral index map generated by re-scaling the `s1` map and extending it to smaller scales using a power-law spectral index power spectrum $C_\ell^{\beta_{\rm S}}$ matching the measurements of [148]. The dust spectral index in `d1` corresponds to the estimate of $\beta_{\rm D}(\hat{\mathbf{n}})$ from the *Planck* data using the `Commander` component separation code [193].

Figure 3.2 shows the polarised amplitude and spectral index maps used in these simulations. The RMS fluctuation around the mean of the spectral index maps are

$$\sigma_{\beta_{\rm D}} = 0.04, \quad \sigma_{\beta_{\rm S}} = 0.22. \tag{3.3.3.1}$$

It is worth noting that the level of realism of these simulations is similar to those used to quantify the performance of future B-mode facilities in e.g. [108, 194, 195].

Three main aspects of these simulations that will challenge the assumptions made by our method are:

1. The spectral index and amplitude maps are *non-Gaussian*.
2. The spectral index maps are based on existing observations, and therefore should be realistically correlated with the polarised amplitudes on large scales.
3. The spectral index fluctuations are not necessarily distributed according to a power-law power spectrum on all scales.

Frequency (GHz)	FWHM (arcmin)	Noise (baseline) (μK-arcmin)	$\ell_{\rm knee}$ –	$\alpha_{\rm knee}$ –
27	91	35	15	-2.4
39	63	21	15	-2.4
93	30	2.6	25	-2.5
145	17	3.3	25	-3.0
225	11	6.3	35	-3.0
280	9	16	40	-3.0

Table 3.2: Summary of the beam Full Width at Half Maximum (FWHM) apertures, baseline and goal sensitivity levels for each band of the SO Small Aperture Telescope (SAT) from [108]. The correlated noise power spectrum is parametrised as $N_\ell = N_{\rm white}[(\ell/\ell_{\rm knee})^{\alpha_{\rm knee}} + 1]$.

3.3.4 Instrumental effects

All the simulations described in the previous two sections include instrumental noise designed to mimic an experiment following the specifications of the SO as described in [108] (SO19 hereon). The simulations are generated in the six frequency bands covered by SO, centred around $\nu = 27, 39, 93, 145, 225$ and 280 GHz. The three most relevant effects are instrument beam, scale-dependent noise, and inhomogeneous sky coverage. For simplicity, we do not include the effects of bandpass convolution and instead assume delta-function bandpasses centred at the frequencies listed above.

We generate noise realisations following the same two-step process described in SO19. First, we generate noise power spectra using the noise calculator released with the data supplement of SO19, of the form:

$$N_\ell = N_{\rm white}\left[\left(\frac{\ell}{\ell_{\rm knee}}\right)^{\alpha} + 1\right], \tag{3.3.4.1}$$

where $N_{\rm white}$ is the white noise variance (in units of $\mu K^2_{\rm CMB}$ srad), while $\ell_{\rm knee}$ and $\alpha_{\rm knee}$ parametrise the scale at which $1/f$ noise starts dominating, and the steepness with which it rises on large scales, respectively. We assume the baseline noise level, an optimistic knee scale $\ell_{\rm knee}$, and a total of 5 years of observation. Details can be found in the SO19 paper and in Table 3.2. We then generate homogeneous noise maps for the six frequency channels as Gaussian realisations of these power spectra. Finally, we scale these maps inversely with the square-root of the hits count map included in the data supplement (and shown in Figure 3.3) to account for the inhomogeneous sky coverage.

The signal maps are convolved with a Gaussian beam with a Full Width at Half Maximum (FWHM) aperture given by $\sigma_{\rm FWHM} = 1.22\lambda/D$, assuming diffraction-limited optics. Here $D = 42\,$cm is the diameter of the SO Small Aperture Telescopes,

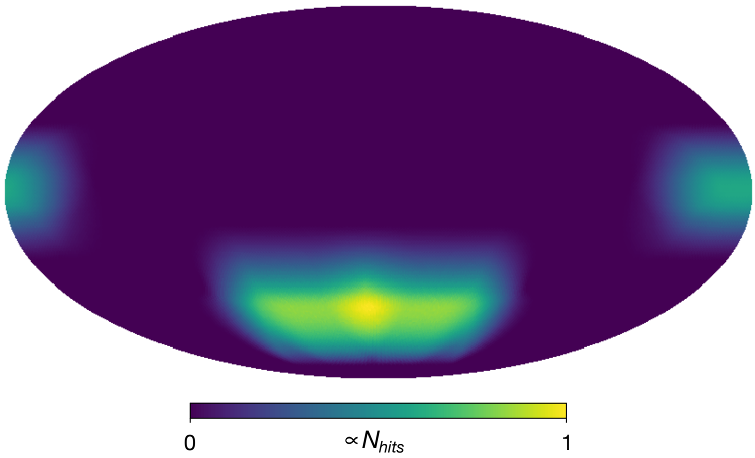

Figure 3.3: Sky mask used in the analysis, proportional to the map of hit counts in Equatorial coordinates used for the SO SATs.

$\lambda = \nu/c$ is the central wavelength of the band. Note that, for each simulation, we generate four independent noise realisations, each with a noise amplitude twice as large as the five-year SO sensitivity. This allows us to use these four realisations as independent data splits when estimating power spectra in order to avoid modelling the noise bias as part of the likelihood (see Section 3.2.5).Signal and noise maps are then added in the observed footprint and saved to file.

3.3.5 Power spectrum measurement

We extract the full set of multi-frequency B-mode power spectra $C_\ell^{\rm BB}$ from each simulation as follows. Each simulation consists of 24 pairs of (Q, U) maps, corresponding to the 6 frequency channels and 4 data splits. We compute the cross-spectrum between all pairs of maps using a pseudo-C_ℓ estimator as implemented in `NaMaster`[7] [79]. We use a differentiable sky mask constructed by smoothing the hits map provided with the SO19 data supplement with a 1° FWHM beam and applying a "C1" apodization with a 5° width (see [108]) to the resulting map. The simulated map of hit counts in Equatorial coordinates for the SATs is displayed in Figure 3.3.

Once all unique auto- and cross-spectra between the 24 different maps have been calculated, we produce coadded power spectra for every pair of frequencies (ν_1, ν_2) by averaging over all power spectra involving maps at those frequencies corresponding to different data splits. By doing this, we isolate the contribution from the inhomogeneous noise bias to the auto-correlations, which are not used to generate the final coadded spectra. Discarding the auto-correlations leads to some

[7] https://github.com/LSSTDESC/NaMaster

loss of sensitivity. However, we find this loss to be negligible when comparing the final constraints on r (presented in the next section) with the official SO forecasts [108]. All power spectra were measured in a set of equispaced bandpowers with width $\Delta\ell = 10$. Of these, only the 27 bandpowers in the range $30 < \ell < 300$ were used in the likelihood analysis, making the total size of the data vector $N_{\rm data} = 567$.

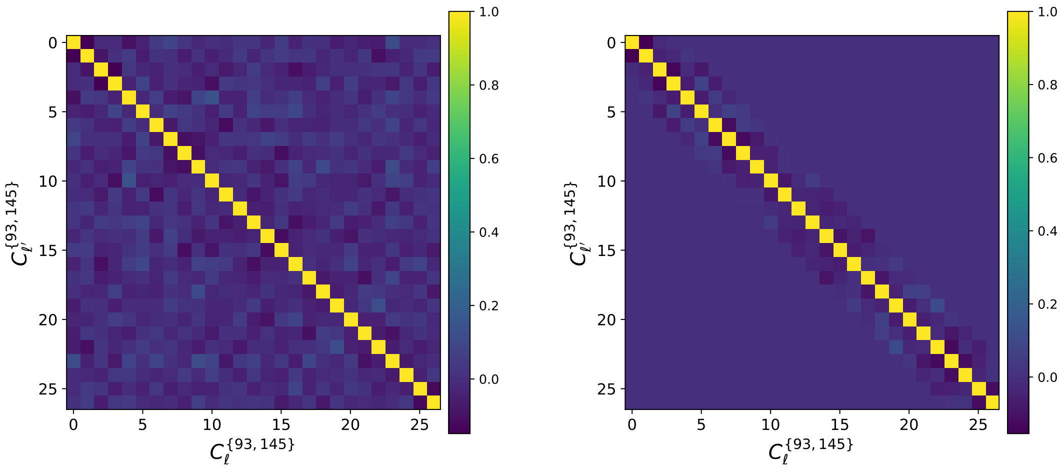

Figure 3.4: *Left:* Correlation matrix obtained from the full covariance of 500 Gaussian simulations at $30 \leq \ell \leq 300$. Here we have selected the block corresponding to the BB cross spectrum $C_\ell^{\nu\times\nu'}$, where $\nu = 93$ GHz and $\nu' = 143$ GHz. The associated covariance is strongly dominated by the diagonal terms. *Right:* All off-diagonal elements of the covariance above the third superdiagonal are set to zero in order to reduce the statistical noise due to the finite number of simulations.

As noted in Section 3.3.2, we repeat this process for two suites of 500 Gaussian simulations to generate the covariance matrices used in the likelihood analysis. We find each block of the resulting covariance involving two pairs of frequencies to be strongly dominated by its diagonal elements. Therefore, to reduce the statistical noise in the covariance due to the finite number of simulations, we set all off-diagonal elements in each block to zero, except for the diagonal and the first three superdiagonals. This is evident in Figure 3.4, which shows the correlation matrix derived from the scaled covariance with unit values along its diagonal.

In this analysis we avoid leakage of E to B modes induced by the partial sky coverage in order to minimise the computational cost of using B-mode purification when calculating the power spectrum covariance. The following section 3.3.6 shows that, in the range of scales used here, this approximation does not affect our results. This was also verified in [108].

Any residual mode-coupling effects in the pseudo-C_ℓ estimator are taken into account analytically when evaluating the theory predictions for these measured power spectra [79].

3.3.6 E-B mixing

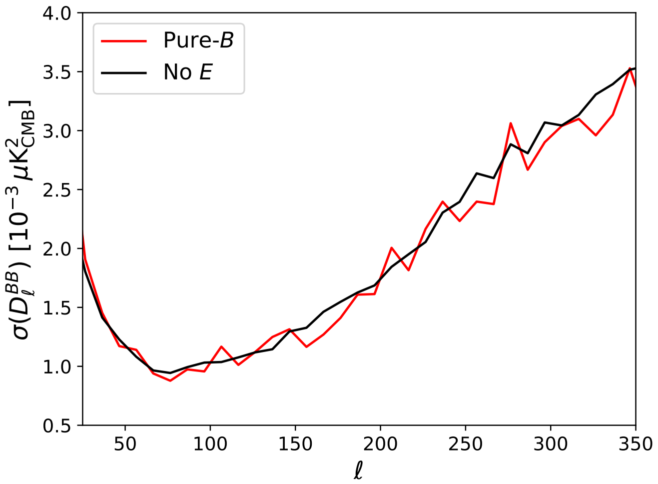

Figure 3.5: Error on the B-mode power spectrum D_ℓ calculated at 95 GHz for 500 simulations. The ideal case when E-modes are not included (*black*) is compared with the result of B-mode purification (*red*). The two cases show commensurate levels of uncertainty on the power spectrum at scales $30 \leq \ell \leq 300$.

The CMB E-mode component of the signal is significantly larger than the B-modes. In the optimal case, the polarisation field can be linearly decomposed into a E and B part, and all the B fields are orthogonal to all the E fields (i.e. pure) and vice versa. However, when dealing with partial-sky observations, E and B modes in the cut sky become linear combinations of the true E and B modes [196, 197].

The leakage of E to B modes contributes to the variance of the B-mode power spectrum estimator, limiting our chances to detect the tiny primordial B-mode signal. Therefore, rather than removing the leakage in harmonic space (i.e. the standard pseudo-C_ℓ approach), in cut-sky regimes it is preferable to define operators at the map level such that the cut-sky harmonic coefficients will only depend on the corresponding E/B field (i.e. B-mode purification formalism [81]). This solution will therefore mitigate the leakage.

In our analysis, we simulate optimal measurements of the B-mode power spectrum. That is, we generate two suites of 500 simulations (see Section 3.3.2) without including E-modes, which allows us to bypass the computationally expensive purification step when using `NaMaster` to compute the covariance. However, this means that we may be underestimating the uncertainty of the power-spectrum estimates.

In order to quantify this, Figure 3.5 compares the error bars of the power-spectrum estimates in the ideal case (corresponding to the case where E-modes are not included) in black, with those of the purified B-mode spectra (i.e. where $E-B$ mixing is mitigated) in red. We find comparable levels of uncertainty for $30 \leq \ell \leq 300$, where the lower angular scales are dominated by cosmic variance and the higher scales by the ℓ^2 factor. The performance of the moment expansion method is therefore not affected by the absence of E-modes at the scales of interest.

3.4 Results

3.4.1 Convergence of the model

Before testing the method on simulated data, let us first gain some intuition on its behaviour by studying the convergence of the power-spectrum-level moment expansion. The moment expansion will converge if each higher-order term in the series is monotonically decreasing [175]. Using the formalism described in Section 3.2.1, and assuming that all spectral parameters are spectral indices, Eq. 3.2.1.6 becomes a geometric series of the form

$$S_\nu^c(\beta(\hat{\mathbf{n}})) = \bar{S}_\nu^c \left(1 + x + x^2 + \dots\right) \tag{3.4.1.1}$$

where $x = \log\left(\nu/\nu_0^c\right) \delta\beta_c$. The moment expansion will therefore converge as long as $x \lesssim 1$. Thus, if we want the method to converge on the frequency interval $[\nu_1, \nu_2]$, there is a certain maximum $\delta\beta$ that the expansion can tolerate. Figure 3.6 shows the convergence bound on $\delta\beta$ for $\nu_1 = 30\,\mathrm{GHz}$ (blue) and $\nu_2 = 300\,\mathrm{GHz}$ (red) as a function of the pivot frequency ν_0. The maximum $\delta\beta$ for the full region is achieved at $\nu_0 = \sqrt{\nu_1\nu_2} \simeq 95\,\mathrm{GHz}$ and corresponds to $|\delta\beta|_{\max} = 2/\log(\nu_2/\nu_1) \simeq 0.87$.

Thus, a good choice of pivot frequency can improve both the convergence of the model and minimise foreground-related biases for a given level of spectral index variability. Note, however, that the choice that maximises the allowed variability of $\delta\beta$ for a converged expansion is not necessarily the optimal choice that minimises foreground biases on r for a finite order in the moment expansion. For example, taking the lowest-order expansion used here, and considering a model consisting of dust, synchrotron and CMB, the choice of pivot frequency for dust should balance the need to describe dust accurately in the frequency channels where it dominates the emission (280 and 220 GHz for SO), while providing a reasonable extrapolation of it on the CMB-sensitive frequencies (145 and 93 GHz for SO), regardless of the convergence of the model at all orders in the expansion in the

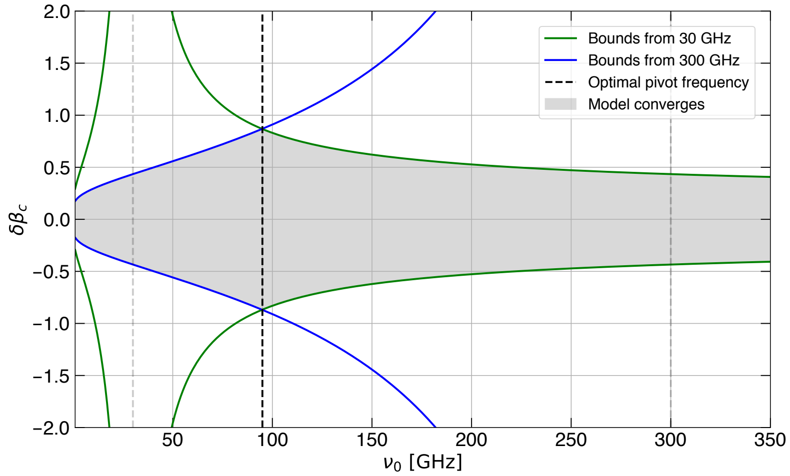

Figure 3.6: Convergence regions of the moment expansion method.

furthest, synchrotron-dominated frequency channels (39 and 27 GHz for SO). We will study the impact of the choice of ν_0 in Section 3.4.2.

Another convergence-related aspect of the particular model used here is the choice of power-law to describe the power-spectrum of spectral index variations $C_\ell^{\beta_c}$. Over a given range of scales $\ell_{\rm min} < \ell < \ell_{\rm max}$, the standard deviation of $\delta\beta_c$ is given by

$$\sigma^2(\beta_c) = \sum_{\ell=\ell_{\rm min}}^{\ell_{\rm max}} \frac{2\ell+1}{4\pi} C_\ell^{\beta_c} = \frac{B_c}{4\pi\ell_0^{\gamma_c}} \left[2\zeta(-\gamma_c - 1) + \zeta(-\gamma_c) - 3\right], \tag{3.4.1.2}$$

where, in the second equality, ζ is the Riemann "zeta" function, we have assumed the power-law model used here (Eq. 3.2.2.12), and used $\ell_{\rm min} = 2$, $\ell_{\rm max} = \infty$. The standard deviation therefore diverges for $\gamma \geq -2$. The models used here, supported by current measurements of the synchrotron and dust spectral index [198], satisfy this constraint. However, as before, ultimately we only need this lowest-order expansion to describe the data on a limited range of scales, in which case these convergence constraints need not be strictly imposed.

3.4.2 Gaussian simulations

We evaluate the performance of our method with a set of Gaussian simulations generated following the prescriptions described in Section 3.3.2. We will study this performance as a function of the level of spectral index variation, parametrised by the standard deviation σ_β rather than the amplitude parameters B_c.

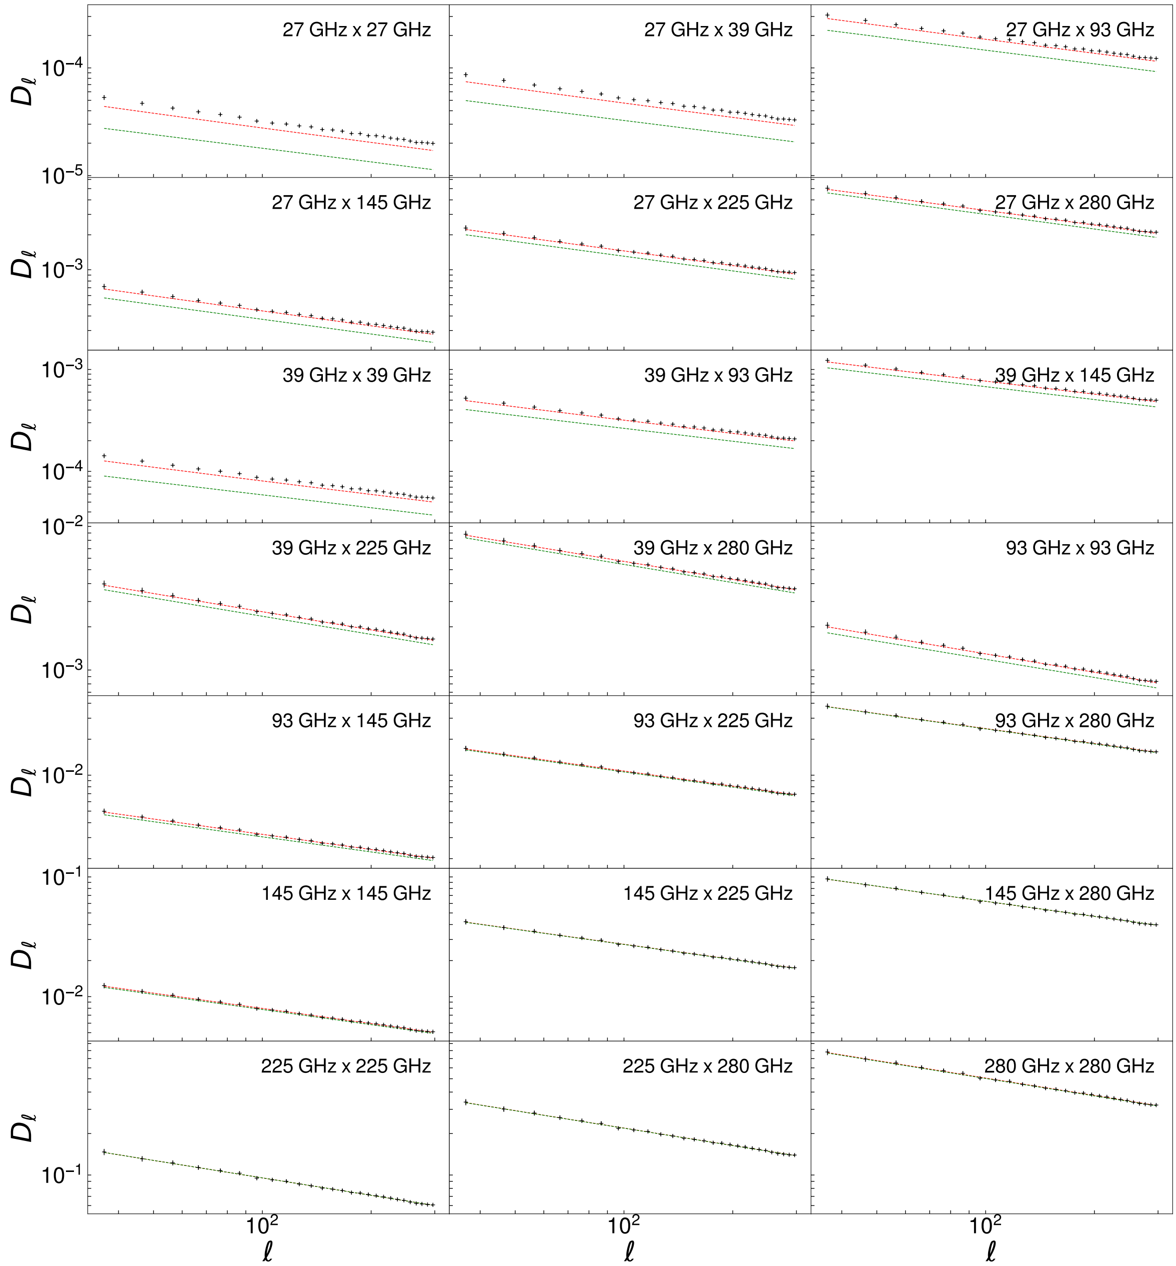

Figure 3.7: Power spectra of the average of ten Gaussian dust-only simulations, assuming a spatially varying spectral index with $\sigma(\beta_\mathrm{D}) = 0.3$ (*black*). The lines show the prediction without accounting for this spatial variation (*green*), and with the inclusion the moment expansion terms (*red*). The order-0 prediction clearly underestimates the power spectrum on relevant frequencies (e.g. 93×93 GHz), likely leading to a bias on r. The inclusion of the order-2 terms recovers the simulated data accurately up to effective frequencies $\sqrt{\nu\nu'} \simeq 60\,\mathrm{GHz}$. At the lowest frequencies (27, 39 GHz), where dust is subdominant to synchrotron and CMB emission, the order-2 expansion underpredicts the power spectrum significantly.

3.4.2.1 Power spectrum predictions

First, we use these simulations to quantify the ability of the lowest-order moment expansion to describe Gaussian simulated data containing all of the higher order terms. The relevance of these terms will determine the bias on r associated with the method, which should increase with increasing σ_β. Figure 3.7 displays the multi-frequency power spectra $C_\ell^{\nu\nu'}$ in the range $30 \leq \ell \leq 300$ used here. The spectral index variation is $\sigma(\beta_D) = 0.3$ with $\gamma_D = -3.5$ at $\nu_0 = 220$ GHz. These simulations contained only dust, the dominant foreground source for SO, with input parameters $A_{\rm D} = 5\ \mu{\rm K}^2$, $\alpha_{\rm D} = -0.42$, $\beta_{\rm D} = 1.6$, in order to study the recovery of this particular component across the whole frequency range. Here the green line shows the predicted dust spectrum assuming a homogeneous spectral index, equal to the true mean $\beta_{\rm D}$ used in the simulations. The red line then shows the result of adding the 1×1 and 0×2 terms.

We can see that the order-0 prediction clearly underestimates the power spectrum on relevant frequencies (such as the 93×93 GHz combination), likely leading to a bias on r. However, in this particular example, considering only dust, the inclusion of the order-2 terms is able to recover the simulated data accurately up to effective frequencies $\sqrt{\nu\nu'} \simeq 60\,{\rm GHz}$. At the lowest frequencies (27, 39 GHz), where the emission is dominated by synchrotron, the order-2 expansion again underpredicts the power spectrum noticeably.

The 1×1 term in the moment expansion (Eq. 3.2.2.8) involves a convolution of C_ℓ^{cc} and $C_\ell^{\beta_c}$. The associated sum over multipoles should in principle cover all integer $\ell_{1,2}$, however, in practice, it is only evaluated up to a maximum multipole $\ell_{\rm max}$. Figure 3.8 shows, as a dashed red line, the predicted dust auto-spectrum at $\nu = 93\,{\rm GHz}$ using $\ell_{\rm max} = 384$, compared with the average of ten Gaussian simulations with $\sigma_{\beta_{\rm D}} = 0$ (black points) and $\sigma_{\beta_{\rm S}} = 0.2$ (red points). The figure shows clearly that the moment expansion calculation becomes numerically inaccurate on scales $\ell > \ell_{\rm max}$, showing that the impact of the choice of $\ell_{\rm max}$ is limited to multipoles similar or larger than that scale. Thus, when implementing the moment expansion, it is sufficient to choose $\ell_{\rm max}$ to be slightly larger than the maximum multipole used in the analysis.

3.4.2.2 Constraints on r

In order to propagate these results to final constraints on the tensor-to-scalar ratio, we proceed with the multi-frequency power spectrum likelihood analysis described in Section 3.2.5. We start by producing a suite of 500 Gaussian simulations with constant spectral indices and $A_{\rm D} = 5\,\mu{\rm K}^2,\ \alpha_{\rm D} = -0.42,\ A_{\rm S} = 2\,\mu{\rm K}^2,\ \alpha_{\rm S} = -0.6$ as described in Section 3.3.2. These simulations are used throughout this analysis to compute the power spectrum covariance matrix and to validate our method. We also used these simulations to validate our pipeline and implementation of the [121] likelihood, by making sure that we are able to recover the input foreground and CMB parameters in a subset of the suite.

In the presence of spatially-varying spectral indices, a significant bias in the tensor-to-scalar ratio can arise if the corresponding effect in the multi-frequency spectra is not taken into account in the model. In the case of Gaussian simulated data, higher B_c values generally correspond to higher biases on r. This is illustrated in Figure 3.9, which shows the recovered best-fit values of r and their 1σ uncertainty for simulations run with increasingly larger values of σ_β (equal for both synchrotron and dust). For each value of σ_β, a set of 20 simulations were generated, with the same seeds in each set, and the figure in the upper panel shows the individual results from each simulation as well as the average over simulations to minimise the impact of sample variance. All simulations were run with an input $r = 0$, and therefore the bias on r is directly given by its mean measured value. Results are

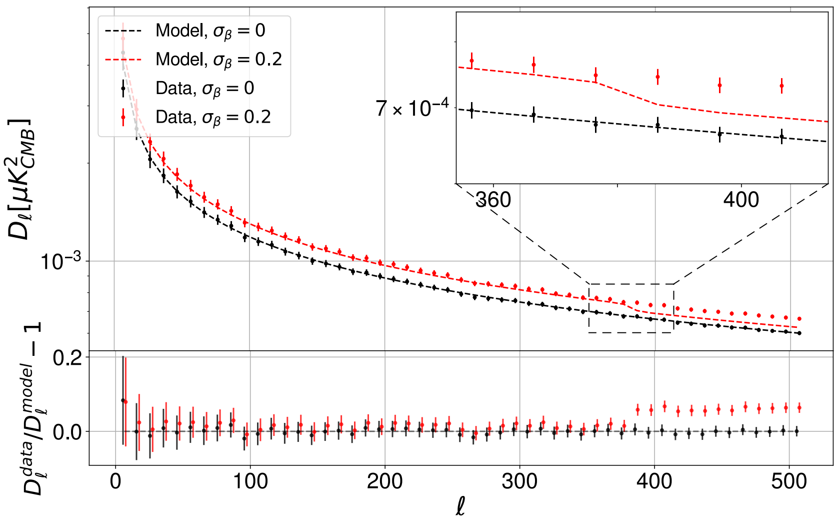

Figure 3.8: *Upper panel:* Dust power spectrum at 93 GHz. The *red* and *black* points show the average of 10 Gaussian simulations with $\sigma_{\beta_{\rm D}} = 0$ and 0.2 respectively. The *dashed red* line shows the theoretical prediction with the moment expansion method using a maximum multipole $\ell_{\rm max} = 384$ when convolving C_ℓ^{cc} and $C_\ell^{\beta_c}$ in the 1×1 term (Eq. 3.2.2.8). As shown in the inset, the impact of this choice is limited to scales $\ell \gtrsim \ell_{\rm max}$, and thus this value must simply be chosen to lie outside the range of scales used in the analysis. *Lower panel:* Fractional difference between the model and the simulated power spectra. The black error bars are recovered from the 10 realisations.

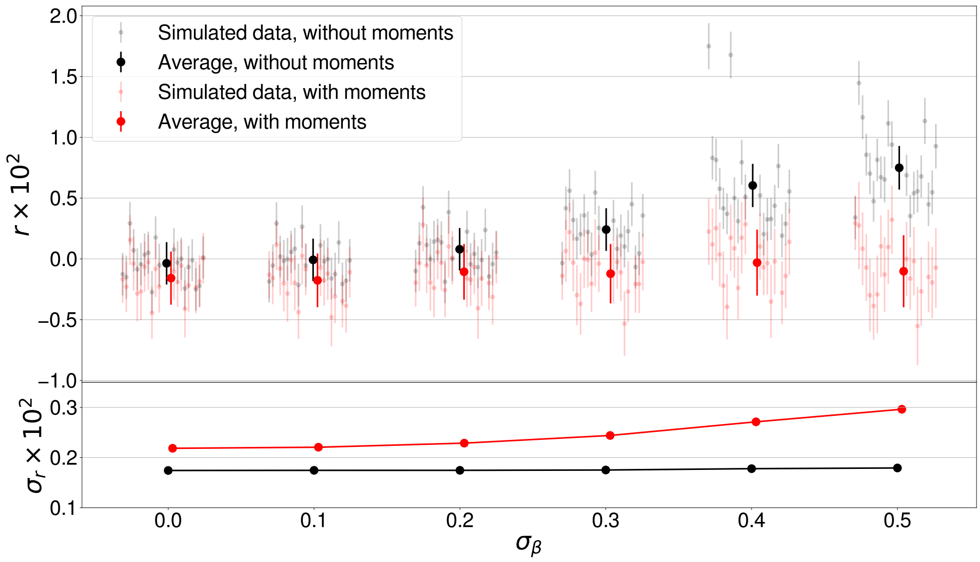

Figure 3.9: *Upper panel:* Best-fit values of r for twenty realisations of the sky calculated at different values of σ_{β_c} (the same for synchrotron and dust). Results are shown for a model assuming constant spectral indices (*black*) and using the moment expansion method (*red*). The position of each simulation in the x axis is shifted slightly from its true σ_{β_c} for clarity. The larger, solid dots, at the centre of each σ_{β_c} value show the mean and standard deviation of each suite of simulations. *Lower panel:* Statistical uncertainty $\sigma(r)$ averaged over the ten realisations in the case of constant spectral indices (*black*) and using the moment method (*red*). The moment expansion is able to correct the bias on r for all values of σ_{β_c} considered, at the cost of increased final uncertainties with respect to a model with constant spectral indices (which themselves increase monotonically with σ_{β_c}).

shown for the final constraints on r found assuming constant spectral indices (black dots) and using the moment expansion method to account for their spatial variation (red dots). We see in the bottom panel of 3.9 that the statistical uncertainty on r consistently increases with higher values of σ_{β_c} when accounting for spectral index variation. The upper panel figure shows that the bias on r when ignoring the spatial variation of β_c grows with σ_{β_c}, becoming of the same order as the statistical uncertainties ($\sigma(r) \simeq 0.002$) for $\sigma_{\beta_c} \sim 0.25$. Accounting for the spectral index variation through our minimal moment expansion consistently corrects this bias, making it compatible with zero for the full range of σ_{β_c} studied here, which encompasses the range of variation allowed by current data [193].

It is worth noting that the mean of the posterior distribution of r for data with no (or mild) spectral index variation analysed using the moment expansion is consistently biased low by $\Delta r \simeq -\sigma(r) \simeq -0.002$. This is due to the parameter degeneracy between r and the spectral index amplitudes B_c, coupled with the

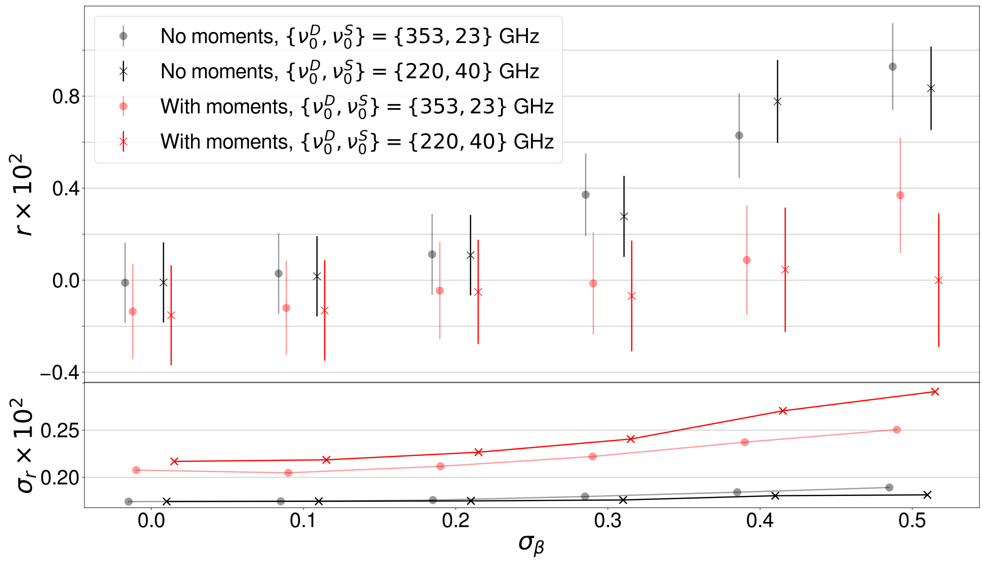

Figure 3.10: *Upper panel:* Mean best-fit values of r for ten sky simulations with varying levels of spectral index variation parametrised by σ_β, without accounting for moments (*black*) and with moments correction (*red*). Results are shown when the analysis is done using 40 GHz and 220 GHz as pivot frequencies for synchrotron and dust respectively (*crosses*), and when using 23 GHz and 353 GHz (*shaded circles*). The simulations were generated using the latter values as pivot frequencies for the foreground amplitude maps. Using pivot frequencies that coincide with channels close to the experiment's central bands allows the moment expansion to recover unbiased results on r for the whole range of σ_{β_c}, whereas the alternative pivot frequencies lead to a bias for large σ_{β_c}. *Bottom panel:* Statistical uncertainty on r in the same cases.

positivity prior $B_c \geq 0$ imposed on the latter. A positive B_c increases the amplitude of the corresponding component's power spectra in the central, CMB-sensitive frequencies (93 and 145 GHz), which the model can compensate through a slightly negative r. We find, however, that the best-fit r value reported here is not significantly biased. This is a well-known effect that arises also when parametrising spectral index variation through a frequency decorrelation parameter Δ_c when imposing a physical prior $\Delta_c \leq 1$ (see Appendix F of [102]). The impact of physically-motivated priors should therefore be considered when interpreting the results of these analyses. We leave a more detailed study of this effect for future work.

The previous results were found with a model using pivot frequencies $\nu_0^{\rm D} = 220$ GHz and $\nu_0^{\rm S} = 40$ GHz. This choice of pivot frequencies is motivated by the fact that SO has two pairs of foreground monitor channels at 27/39 GHz and 220/285 GHz for synchrotron and dust respectively. Since the allowed variability of β_c is larger when the pivot frequencies are closer to the centre frequencies of the experiment's band, choosing the pivot frequencies to lie in the corresponding monitor channel lying closer to the CMB-sensitive bands ensures an accurate description of the foregrounds at the monitor frequencies and a reasonable extrapolation in the intermediate bands. An application of this method to instruments with more frequency bands, such as future space-borne missions [199, 200], will likely require a more careful analysis of the optimal pivots. The relevance of this choice is shown in Figure 3.10, where we compare the results from our previous analysis with the same method using pivot frequencies at $\nu_0^{\rm D} = 353\,{\rm GHz}$ and $\nu_0^{\rm S} = 23\,{\rm GHz}$, motivated by the location of sensitive bands in previous experiments [192, 201]. When using these more distant pivot frequencies, we find that the moment expansion is not able to recover unbiased results for the highest value of σ_{β_c} explored here.

It is worth noting that the results above were obtained for simulations in which the spectral tilt of the spectral index fluctuations were $\gamma_{\rm D} = \gamma_{\rm S} = -3$. This corresponds to fairly steep power spectra, which induce substantial large-scale variations of the spectral indices. Since we analyse cut-sky simulations, the effective mean spectral indices in our sky patch are not necessarily centred at the fiducial values ($\bar{\beta}_{\rm D} = 1.6$, $\bar{\beta}_{\rm S} = -3$) we assumed to generate the simulated data. In fact, due to the large scale power associated with $\delta\beta_c$, the variation of the effective mean spectral index in the analysed footprint can easily be as large as the σ_β assumed in our simulations. In order to avoid this, we use large Gaussian priors on $\beta_{\rm D}$ and $\beta_{\rm S}$ (0.5 and 0.6 respectively). We have verified that the results presented here hold also for simulations with flatter spectral tilts $\gamma_c = -2.1$, in line with current measurements of the synchrotron spectral index [148].

The posterior distributions for one realisation of a Gaussian simulation with $\sigma_{\beta_{\rm D}} = 0.2$, $\sigma_{\beta_{\rm S}} = 0.3$, $A_{\rm D} = 28\mu{\rm K}^2$, $A_{\rm S} = 1.6\mu{\rm K}^2$ is shown in Figure 3.11. The left panel shows the results for the tensor-to-scalar ratio and the average spectral indices with and without the inclusion of the additional moments parameters in blue and orange, respectively. The posterior distribution of the moment amplitudes and spectral tilts, displayed in the right panel of Figure 3.11, shows that $B_{\rm D}$ and $B_{\rm S}$ are not significantly larger than zero compared with their error bars. The power spectrum-based approach used here is therefore unable to detect the spectral index variation from the Gaussian simulated data, with a level of spectral index variation compatible with existing measurements of $\beta_{\rm S}$ and $\beta_{\rm D}$.

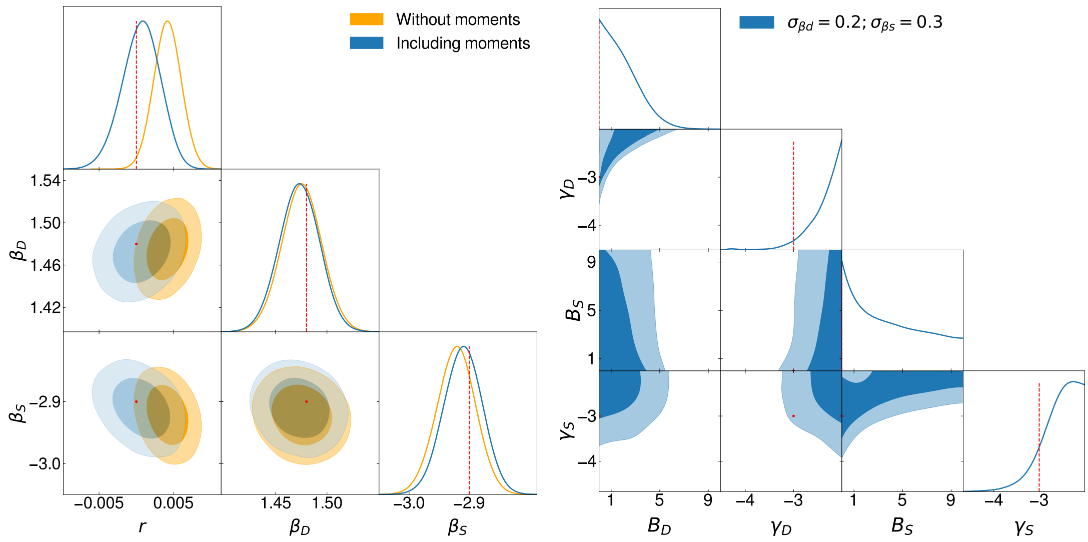

Figure 3.11: *Left panel:* posterior distribution for r, $\beta_{\rm D}$ and $\beta_{\rm S}$ in Gaussian simulations with spatially-varying indices ($\sigma_{\beta_{\rm D}} = 0.2$ and $\sigma_{\beta_{\rm S}} = 0.3$). Results are shown for a model assuming constant spectral indices (*orange* contours), and using the moment expansion method (*blue* contours). The input parameters used to run the simulation are shown as red points and dashed lines. *Right panel:* posterior distribution for the moment expansion parameters. Note that the values of B_c displayed in our analysis are 6 orders of magnitude larger for clarity.

3.4.3 Realistic simulations

Having validated the method for simulations that follow the implemented theoretical model, we now turn to simulations containing "realistic" foreground spectral index maps and amplitudes that do not adhere to our model assumptions (Gaussian fields, uncorrelated indices and amplitudes), as described in Section 3.3.3. As discussed in Section 3.3.3, although the level of complexity increases compared to the Gaussian simulations, these simulations contain a comparatively low level of spectral index variation for both dust and synchrotron.

In Figure 3.12 the power spectrum of the "realistic" amplitude maps for dust and synchrotron is compared with power-law fits. We find that, within the range of scales considered here, the power spectra of the `PySM` maps are well described by power laws with amplitudes $(A_{\rm D}, A_{\rm S}) = (27.7,\ 1.6)\mu{\rm K}^2$ and tilts $(\alpha_{\rm D}, \alpha_{\rm S}) =$ (-0.16, -0.93). Thus, any bias on r resulting from the analysis of these simulations can be attributed to the spatially-varying spectral indices, and not to an incorrect modelling of the scale dependence of foreground amplitudes. We use these values to generate another suite of 500 Gaussian simulations (described in Section 3.3.2), which we use to estimate the power spectrum covariance. It is worth noting that the foreground power spectra exhibit clear departures from a perfect power law on scales larger than those used here [193], but this does not affect our results.

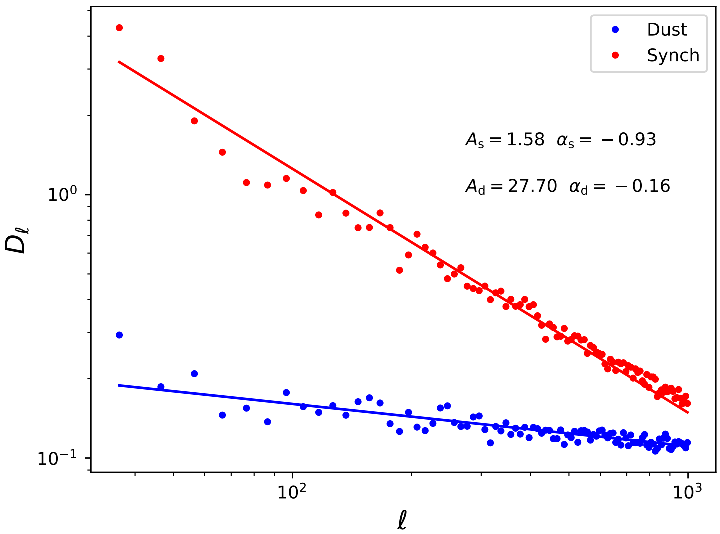

Figure 3.12: Power spectra of the `PySM` polarised synchrotron and dust maps (*red* and *blue* dots respectively) in the sky patch used in our analysis. The solid lines show the best-fit power-law spectra with amplitudes (A_D, A_S) = (27.70, 1.58)μK^2 and tilts (α_D, α_S) = (-0.16, -0.93).

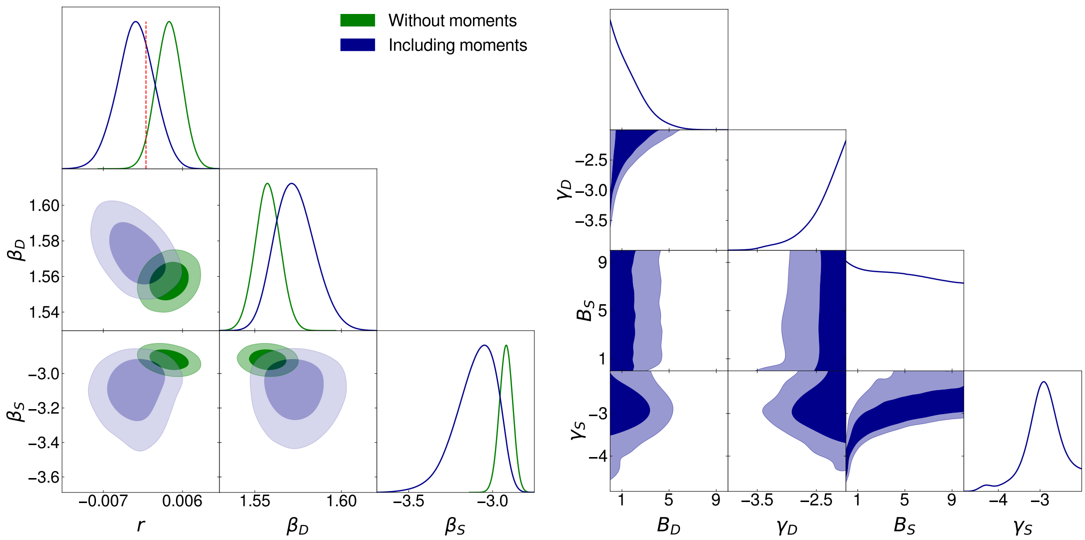

Figure 3.13: *Left panel:* posterior distribution of the multi-frequency power spectrum likelihood for r, β_D and β_S from the realistic simulations. The green contours show the baseline result assuming constant spectral indices, while the blue curves show the constraints using the moment expansion method. *Right panel*: distribution of the moment expansion parameters. No significant detection of spatially-varying indices is found with this method. Note that all values of B_c displayed in our analysis are 6 orders of magnitude larger for clarity.

Ten different realisations of these realistic simulations were generated using the same foreground amplitude and spectral index maps, but varying the CMB and noise components. These were then analysed through our pipeline with and without the inclusion of the additional moments parameters to recover the posterior distribution for r. The results for one realisation are shown in Figure 3.13. The baseline result of the r posterior distribution is centred at 0.0029 ± 0.0023 averaged over the ten realisations. The higher complexity of these simulations introduces a bias on r, which is comparable to the bias found with the Gaussian simulations with a spectral index variation $\sigma_\beta = 0.3$.

After including the higher-order foreground terms and marginalising over the 4 additional parameters, the constraints on r are $r = 0.0005 \pm 0.0028$ averaged over simulations. Thus, the bias on r is corrected by over 1σ if we include the extra foregrounds parameters (the posterior of these is shown in the right panel of Figure 3.13). The inclusion of moments induces a small increase in the uncertainty on r, with $\sigma(r) = 0.0028 \pm 0.0004$ averaged over simulations. This is visible in the corresponding wider posterior distribution of r in Fig. 3.13. The spectral index parameters, on the other hand, seem to absorb some of the additional parameter freedom, with $\beta_{\rm D}$ increasing its posterior standard deviation by $\sim 30\%$.

3.4.4 Simulation challenge

As a final validation test for our implementation of the moment expansion method, we have carried out a "simulation challenge" to determine the ability of the method to absorb a variety of foreground parametrisation. A set of 12 different simulations were generated independently and then analysed blindly without knowledge of the contents of each simulation. In order to quantify the performance of the method, the analysis was carried out with and without the moment expansion. Covariance matrices were estimated using the second suite of Gaussian simulations described in Section 3.3.2. The simulations combined the levels of complexity encoded in the Gaussian and realistic simulations described in the previous sections with additional ingredients.

- Simulations were generated with Gaussian amplitudes, with $A_{\rm D} = 28\,\mu{\rm K}^2$, $\alpha_{\rm D} = -0.16$, $A_{\rm S} = 1.6\,\mu{\rm K}^2$, $\alpha_{\rm S} = -0.93$ (labelled "G" in Table 3.3) and with the realistic amplitude templates included in `PySM` (labelled "P" in Table 3.3).

- Two simulations were run with the realistic `PySM` amplitude templates for synchrotron and dust and with the addition of spinning dust (AME) with a 2% polarisation fraction (labelled "P+A").

- The thermal dust contribution was propagated in frequency using the modified black-body spectrum in Eq. 3.2.1.3 (labelled "MBB") as well as the model by [179] (labelled "H&D").

- Foreground spectral indices were generated as Gaussian fields (labelled by their standard deviation σ_β), as well as using the more complex templates described in Section 3.4.3.

- The underlying value of r was varied between $r = 0$ and $r = 0.01$ in different simulations.

- Finally, two simulations were run using the statistical model described in [143] (labelled "VS" in the table). In this model, the three-dimensional structure of the GMF is described by a finite number of layers (we use $N_{\rm layer} = 7$ layers). The coherent component of the magnetic field is the same in all layers, while its turbulent part is generated as a Gaussian random field. Once the direction of the GMF is determined, maps of the Q and U Stokes parameters are generated by scaling the dust intensity map found by *Planck* [193]. As an additional level of complexity, and in an attempt to describe the three-dimensional distribution of the dust spectral index, we associate each layer with a different Gaussian realisation of $\delta\beta_{\rm D}$ with standard deviation $\sigma_{\beta_{\rm D}} = 0.13$ (the combined RMS variation for 7 layers is $\sigma_{\beta_{\rm D}} \simeq 0.35$).

The results for the different simulations run as part of this challenge are summarised in Table 3.3. Highlighted in red are the results of simulations in which a the best-fit value of r was found to be more than 2σ away from the input value. We find that in most cases where foregrounds introduce a bias on r at this level using the standard method, the moment expansion method is able to recover unbiased results at the same level.

The results of the analysis for the first four simulations (rows 1-4 in Table 3.3) are consistent with the results presented in Section 3.4.2. When no spectral index variation is introduced ($\sigma_\beta = (0, 0)$), the final results both with and without moments are compatible with the true values of r within their respective $\sigma(r)$, with a modest widening in the final constraints when moments are included. When Gaussian spectral index variation is introduced ($\sigma_\beta = (0.2, 0.3)$), the baseline result exhibits a bias up to the 3σ level, which is reduced by 2σ using the moment expansion method, with a $\sim 30\%$ increase in uncertainty.

Similarly, the $1 - 2\sigma(r)$ level bias induced by the `PySM` templates for the spectral indices (rows 5 and 6) is reduced by $\gtrsim 85\%$ of $\sigma(r)$ compared to the baseline result,

Simulation		No moments		With moments	
Description; $(\sigma_{\beta_D}, \sigma_{\beta_S})$	$r_{\rm true}$	$r_{\rm fit} \pm \sigma(r)$	χ^2	$r_{\rm fit} \pm \sigma(r)$	χ^2
G, MBB; $\sigma_\beta = (0, 0)$	0	-0.0013 ± 0.0021	0.8	-0.0024 ± 0.0024	0.8
G, MBB; $\sigma_\beta = (0, 0)$	0.01	0.0116 ± 0.0022	0.8	0.0099 ± 0.0025	0.8
G, MBB; $\sigma_\beta = (0.2, 0.3)$	0	0.0088 ± 0.0023	0.9	0.0038 ± 0.0035	0.8
G, MBB; $\sigma_\beta = (0.2, 0.3)$	0.01	0.0158 ± 0.0025	0.9	0.0098 ± 0.0035	0.9
P, MBB; $\sigma_\beta = \texttt{PySM}$	0	0.0051 ± 0.0022	0.9	0.0036 ± 0.0026	0.9
P, MBB; $\sigma_\beta = \texttt{PySM}$	0.01	0.0130 ± 0.0023	0.9	0.0104 ± 0.0027	0.9
P+A, MBB; $\sigma_\beta = \texttt{PySM}$	0	0.0071 ± 0.0022	1.1	0.0041 ± 0.0025	1.1
P+A, MBB; $\sigma_\beta = \texttt{PySM}$	0.01	0.0148 ± 0.0023	1.1	0.0127 ± 0.0026	1.1
G, H&D; $\sigma_\beta = (0, 0)$	0	0.0058 ± 0.0026	1.1	0.0003 ± 0.0037	1.1
G, H&D, $\sigma_\beta = (0, 0)$	0.01	0.0122 ± 0.0024	1.1	0.0055 ± 0.0038	1.1
P, H&D; $\sigma_\beta = \texttt{PySM}$	0	0.0052 ± 0.0025	1.1	0.0001 ± 0.0033	1.1
P, H&D; $\sigma_\beta = \texttt{PySM}$	0.01	0.0120 ± 0.0024	1.1	0.0069 ± 0.0034	1.1
P, VS; $\sigma_\beta = (0.13, N.A.)$	0	0.0114± 0.0024	1.0	-0.0036 ± 0.0036	1.0
P, VS; $\sigma_\beta = (0.13, N.A.)$	0.01	0.0184 ± 0.0025	1.0	0.0029 ± 0.0034	1.0

Table 3.3: Results from the simulation channel, where the χ^2 is evaluated per degree of freedom. The values marked in red show the results with a bias $|r_{\rm fit} - r_{\rm true}| \geq 2\sigma(r)$.

with a 17% degradation in $\sigma(r)$. These results are unaffected by the addition of spinning dust (AME) to the foreground components (rows 7 and 8).

When the H&D dust SED is used in the simulations (rows 9 to 12), the deviations from the original MBB spectra introduce a bias at the $1 - 2\sigma$ level. The additional freedom in the dust model due to the inclusion of the moment parameters corrects the bias whilst also widening the final constraints by $\sim 40\%$.

Finally, the more complex non-Gaussian VS simulations induce a large bias on r, which the moment expansion is able to correct at the cost of increasing the final uncertainties by $\sim 50\%$. We find, however, that for an input $r = 0.01$, the moments method underestimates it by $\simeq 2.1\sigma$. After finding this, we verified that the same result is reproduced on a second realisation of the VS model, and therefore believe that this is not due to a statistical fluke. The cause of this bias is not clear. It could be due to the additional complexity of the VS simulations (non-Gaussianity, additional spectral index variation and frequency decorrelation). It could also be that the simulated data in this case differs significantly from the model used to construct the covariance matrix and the resulting likelihood is ill behaved. We leave a more thorough analysis of the performance of the lowest-order moment expansion used here on data with this level of complexity for future work.

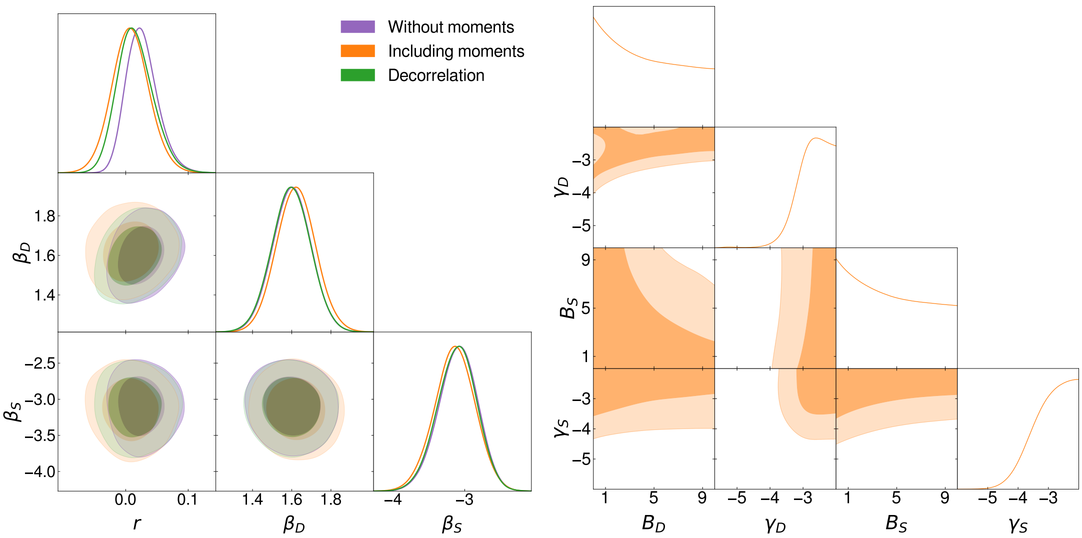

Figure 3.14: *Left*: Posterior distribution of r, and spectral indices β_D and β_S using our analysis pipeline on the publicly available BK15X data. The green curves show the baseline moments-less case and agrees with the BK15X published posterior distributions, with r peaking at 0.023. Marginalising over foregrounds spatial variation shifts the r curve closer to zero (orange) without a significant impact on the posteriors. The results using the decorrelation parameter method, as done in [102], are displayed in purple. *Right*: Distribution of the additional SED parameters corresponding to the moments of the parameter distributions.

It is worth noting that none of the simulations analysed here included the effects of polarised extragalactic sources, which may be relevant for future B-mode searches [202].

3.4.5 BICEP2/Keck Array data

In order to further validate the moment expansion method, as well as to explore the sensitivity of current B-mode constraints to the spatial variability of foreground spectral indices, we have applied the method to the latest publicly available data from the BICEP2/*Keck* collaboration [102] (BK15X hereon).

The BK15X dataset is fully described in [102]. The power spectrum data contain cross-correlations between 12 frequency bands, including 3 BICEP2/*Keck* bands (at 95, 150 and 220 GHz), 7 *Planck* bands (30, 44, 70, 100, 143, 217 and 353 GHz) [201] and 2 low-frequency WMAP bands (23 and 33 GHz) [192]. The analysis is performed over the BICEP2/*Keck* footprint, covering approximately 400 deg^2. All the information concerning cross-correlations between different frequencies and polarisation bands, their covariance matrix, all frequency bandpass transmission curves, and bandpower window functions are publicly available[8]. All power spectra

[8]See http://bicepkeck.org/bk15_2018_release.html.

have been measured in a set of 9 equi-spaced bandpowers which cover the multipole range $\ell \lesssim 400$. The total size of the data vector is $N_{\rm data} = 2700$, of which 702 elements correspond to B-mode-only correlations.

In order to validate the implementation of our multi-frequency foreground model, we first reproduced the fiducial BK15X results by running our component separation pipeline on the B-mode data using the same parameter priors used by BK15X. The results are displayed in green in Figure 3.14. We recover the published r posterior distribution [102] almost perfectly, and other foreground parameter constraints (e.g. on $\beta_{\rm S}$, $\beta_{\rm D}$) agree to better than 5% with those presented in BK15X.

The orange contours in the same figure show the constraints after marginalising over the four additional moment expansion parameters. We observe a small shift in the r posterior mean towards smaller values, accompanied by a broadening of the distribution by $\sim 24\%$. The right panel of Figure 3.14 shows the posterior distribution of the moment amplitudes and spectral tilts. We find no evidence of spectral index variations in the BK15X data.

The original BK15X analysis studied the impact of spatially-varying spectral indices by introducing two frequency decorrelation parameters [102] with different levels of scale dependence. As we showed in Sec. 3.2.4, a constant decorrelation parameter is equivalent to a non-perturbative moment expansion in the specific case of scale-independent (i.e. uncorrelated) spectral index variations[9]. In order to compare the impact of both parametrisations on the final constraints on r, we reproduced the decorrelation results found by BK15X with our pipeline. The results, shown as purple contours in Fig. 3.14, display a similar downward shift in the r posterior, while its width is slightly smaller than the full moment expansion. Thus, although parametrising the impact of spatially-varying indices in terms of frequency decorrelation captures one of the most important effects in the data vector, the moment expansion is able to effectively marginalise over additional freedom in the scale dependence of these spatial variations. Although the relevance of this additional freedom is small in current datasets, it may prove to be important when more sensitive data become available.

[9] Note that, although the analysis of BK15X allowed for different forms of scale dependence for the decorrelation parameter, these do not map directly onto a model for the spatial fluctuations in the foreground spectral indices. We have only considered the constant decorrelation case here.

3.5 Conclusion

A detailed characterisation of Galactic polarised emission is necessary to disentangle CMB B-modes from other sources of polarised emission. In this context, multi-frequency power-spectrum-based component separation pipelines have been used by ground-based experiments to derive the current state-of-the-art constraints on the tensor-to-scalar-ratio r from B-modes [102]. C_ℓ-based methods provide several advantages when handling ground-based B-mode data, since they are computationally less challenging than pixel-based techniques, and allow for a straightforward treatment of correlated noise, complicated map filtering, and certain systematics [203].

This method, however, has major drawbacks when applied to data with higher sensitivity over wider patches of the sky. In particular, the characterisation of spatially-varying foreground spectra, being difficult to model at the power spectrum level, poses major challenges. We have addressed this issue here by designing a power-spectrum-based component separation approach, based on existing moment expansion methods. In order to curb the number of new free parameters that rapidly appear in the standard series expansion, which would degrade the final constraints on r significantly, we impose three strong assumptions on the model: spectral index variations are Gaussianly distributed, foreground amplitudes and spectral index variations are uncorrelated, and spectral index variations of different foreground sources are uncorrelated. Departures from these assumptions are thus ignored as higher-order terms in the expansion. The resulting parametric model has four additional free parameters: the amplitudes and slopes of the power spectra describing the dust and synchrotron spectral index fluctuations.

In order to quantify the performance of this method, as well as the impact of its assumptions on the final r constraints, we have made use of a suite of sky simulations with increasing degrees of realism. These include Gaussian foreground simulations following the same assumptions of the model (Section 3.4.2), as well as more realistic simulations based on various models proposed in the literature (Sections 3.4.3 and 3.4.4). These simulations assumed an instrumental setup similar to that expected of the SO Small-Aperture Telescopes ($\sim 10\%$ of the sky with a $\sim 2\mu$K arcmin white-noise level), accounting for both inhomogeneous sky coverage and non-white noise. Finally, we have applied this method to the B-mode data made publicly available by the BICEP2/*Keck* collaboration.

Overall we find that the method is able to correct the bias to the tensor-to-scalar ratio induced by spectral index variations for most realistic foreground models.

From the Gaussian simulation suite we find that the leading-order expansion is able to cope with spectral index variations at the level of $\sigma_\beta \lesssim 0.5$, compatible with existing measurements of the polarised foreground spectral indices [148, 193]. We also find that a judicious choice of the pivot frequencies defining the amplitude and spectral indices for a given foreground source can improve the performance of the method. Although we do not attempt an exact derivation of the optimal pivots, we follow the rule of thumb of using a pivot frequency corresponding to the foreground monitor channel closest to the foreground minimum.

The additional freedom in the foreground model due to the four moment parameters results in a moderate widening of the final constraints on r. In the case of an SO-like dataset, this corresponds to a $\sim 30-50\%$ increase in $\sigma(r)$, with a more moderate increase of $\sim 20\%$ for the current BICEP2/*Keck* data. As shown in Sec. 3.2.4, the moment expansion used here is a generalisation of the frequency decorrelation parameter used in the BK15X analysis, and we find that both methods have a similar effect on the posterior distribution for those data. Although the levels of spectral index variability explored here can induce a bias on the measurement of r at the level of 1 or 2σ, we find that it is not possible to significantly detect the effects of this variation on the foreground multi-frequency power spectra for SO-like sensitivities.

The model used here is based on a leading-order expansion of the cross-frequency power spectra with respect to the spectral index variations. We show in Sec. 3.2.4 that the impact of spatially-varying indices on the power spectrum can be calculated exactly at all orders using methods developed in the context of CMB lensing reconstruction. Although the applicability of other lensing-inspired techniques in the context of non-Gaussian foregrounds may be limited, it could be an interesting avenue to pursue in the future.

Given this option to include new foreground parameters, one natural question to ask is whether these parameters are necessary to describe a given set of real data. One could perform simple model selection tests using information criteria or use a evidence-based methods [204] to judge whether moment parameters are required by the data. Given that the results here show only a fairly moderate increase in $\sigma(r)$, we leave this analysis to future work. It should be noted that such an exploration is warranted even on the standard B-mode analyses in the literature (e.g. to study whether, from a model selection perspective, frequency decorrelation parameter needs to be included in the model [102]).

Our study has been limited to the analysis of primordial B-modes from ground-based facilities targeting the recombination bump on scales $30 \lesssim \ell \lesssim 300$. Its applicability to space missions targeting the reionization bump on larger scales, and covering a wider range of frequencies [199, 205], may be be limited, and the use of pixel-based methods is likely more appropriate. Nevertheless, we expect that the methodology presented here, as well as its potential extensions, some already explored in the literature [176], will be useful in the analysis of future ground-based observatories, such as the SO [108] or CMB Stage-4 [69], which will require the characterisation of spatially-varying foreground spectra, and marginalisation over them in order to achieve reliable constraints on r.

Chapter 4

The Simons Observatory: pipeline comparison and validation for large-scale B-modes

Contents

4.1 Introduction

In this work, we aim at validating three independent B-mode analysis pipelines. The first, "pipeline A", encompasses the pipelines explored in Chapter 3: the baseline cross-C_ℓ and the moment-based extension. The development, testing, and analysis of the latter, which is the main focus of the results presented in this Chapter, was carried out by the author of this thesis. We compare the performance of all three pipelines regarding a potential r measurement by the Simons Observatory's SATs, and evaluate the capability of the survey to constrain $\sigma(r = 0) \leq 0.003$ in the presence of foreground contamination and instrumental noise. To that end, we produce sky simulations encompassing different levels of foreground complexity, CMB with different values of r and different amounts of residual lensing contamination, and various levels of the latest available instrumental noise.

We feed these simulations through the analysis pipelines and test their performance, quantifying the bias and statistical uncertainty on r as a function of foreground and noise complexity. The three pipelines are briefly described in Section 4.2, together with a summary of the statistical tests adopted for to compare the models of pipeline A. Section 4.3 presents the simulations used in the analysis, and compares them with the simulations used in the previous Chapter 3. In Section 4.4, we present our forecasts for r. Section 4.4.2 shows preliminary results on a set of new, complex foreground simulations. In Section 4.5 we summarise and draw our conclusions.

4.2 Methods

In this section we present the three component separation pipelines, that adopt complementary approaches widely used in the literature: power-spectrum-based parametric cleaning ([89, 102]), Needlet Internal Linear Combination (NILC) blind cleaning [206–208], and map-based parametric cleaning [209]. In the following, these are denominated pipelines A, B and C, respectively. The cleaning algorithms operate on different data spaces (harmonic, needlet and pixel space) and vary in their cleaning strategy (parametric, meaning that we assume an explicit model for the frequency spectrum of the foreground components, or blind, meaning that we do not model the foregrounds or make any assumptions on how their frequency spectrum should be). Hence, they do not share the same set of method-induced systematic errors. This will serve as an important argument in favour of claiming robustness of our inference results.

4.2.1 Pipelines

- **Pipeline A: Cross-C_ℓ cleaning**. Pipeline A is based on a multi-frequency power-spectrum-based component separation method, which is equivalent to the baseline "moment-less" pipeline described in Chapter. 3, where the data vector is the full set of cross-power spectra between all frequency maps, $C_\ell^{\nu\nu'}$, and the model is equivalent to the $C_\ell^{\nu\nu'}|_{0\times0}$ term in Eq. 3.2.2.3. The power spectra are measured using a pseudo-C_ℓ approach with B-mode purification as implemented in `NaMaster` [79] (unlike the previous Chapter 3, where the purification step was bypassed (see Sec. 3.3.6). The full pipeline is publicly

$A_{\rm lens}$	r	A_d	α_d	β_d	A_s	α_s	β_s	ϵ_{ds}	B_s	γ_s	B_d	γ_d
TH	TH	TH	TH	G	TH	TH	G	TH	TH	TH	TH	TH
1.0	0.0	25	0.0	1.54	2.0	−1.0	−3.0	0.0	0.0	−4.0	5.0	−4.0
1.0	0.1	25	0.5	0.11	2.0	1.0	0.3	1.0	10.0	2.0	5.0	2.0

Table 4.1: Parameter priors for pipeline A, considering both the C_ℓ-fiducial model and the C_ℓ-moments model. The first line displays the full list of the model parameters. The second line specifies the prior types, either Gaussian (G) or top-hat (TH), considered distributed symmetrically around the centre value (specified in the second line). The last line displays the standard deviation (in the case of G priors) or the half width (for TH priors).

available[1]. We will refer to results using this model as "C_ℓ-fiducial". This terminology is used to distinguish this pipeline from the expansion provided by the moment-based parameterisation, described in Ch. 3, used to describe the spatial variation of β_d and β_s. We will refer to results using this method as "C_ℓ-moments", or "A + moments". Table 4.1 lists the priors on the parameters. It is worth noting that this model is equivalent to the fiducial model in the absence of spatial variations of the spectral parameters.

- **Pipeline B: NILC cleaning**. The second pipeline is based on the blind Internal Linear Combination (ILC) method, which assumes no information on foregrounds whatsoever, and instead only assumes that the observed data contains one signal of interest (the CMB) plus noise and contaminants [210]. The method assumes a model for the data $\mathbf{d}_\nu$ that can be defined in either pixel or harmonic space as $\mathbf{d}_\nu = a_\nu \mathbf{s} + \mathbf{n}_\nu$, where a_ν is the spectrum of the CMB, $\mathbf{s}$ is its amplitude, and $\mathbf{n}_\nu$ is the contamination in channel ν, which includes the foregrounds and instrumental noise. The method aims at reconstructing a map of the CMB component $\tilde{\mathbf{s}}$ as a linear combination of the data with a set of weights w_ν allowed to vary across the map,

$$\tilde{\mathbf{s}} = \sum_\nu w_\nu \mathbf{d}_\nu = \mathbf{w}^T \hat{\mathsf{d}}\,, \tag{4.2.1.1}$$

where both $\mathbf{w}$ and $\hat{\mathsf{d}}$ are $N_\nu \times N_{\rm pix}$ matrices,with $N_{\rm pix}$ being the number of pixels. Optimal weights are found by minimising the variance of $\tilde{\mathbf{s}}$. The result is given by:

$$\mathbf{w}^T = \frac{\mathbf{a}^T \hat{\mathsf{C}}^{-1}}{\mathbf{a}^T \hat{\mathsf{C}}^{-1} \mathbf{a}}, \tag{4.2.1.2}$$

[1]See github.com/simonsobs/BBPower.

where **a** is the black-body spectrum of the CMB (i.e. a vector filled with ones if maps are in thermodynamic temperature units) and $\hat{\mathsf{C}} = \langle \hat{\mathsf{d}}\, \hat{\mathsf{d}}^T \rangle$ is the frequency-frequency covariance matrix per pixel of the observed data. Note that we assume no correlation between pixels for the optimal weights. In this analysis, the Needlet ILC method (NILC) [206–208] has been used. NILC uses localisation in pixel and harmonic space by finding different weights **w** for a set of harmonic filters, called "needlet windows". These windows are defined in harmonic space $h_i(\ell)$ for $i = 0, ..., n_{\rm windows} - 1$ and must satisfy the constraint $\sum_{i=0}^{n_{\rm windows}-1} h_i(\ell)^2 = 1$ in order to preserve the power of the reconstructed CMB.

- **Pipeline C: map-based cleaning**. Our third pipeline is a map-based parametric pipeline and is described in more detail in Sections 2.3.3,5.2.1,5.2.1.2. Pipeline C also offers the option to marginalise over a dust template, with a similar approach to earlier methods [211, 212]. The model of the component-separated CMB map include an additional dust term:

$$C_\ell^{\rm CMB} = C_\ell^{\rm CMB}(r,\, A_{\rm lens}) + A_{\rm dust} \tilde{C}_\ell^{\rm dust}. \tag{4.2.1.3}$$

  In principle one could add synchrotron or other terms in Equation 4.2.1.3 but we limit ourselves to dust as it turns out to be the largest contamination. In the remainder of this chapter, we will refer to this method as "C + dust marginalisation".

4.2.2 Model comparison for pipeline A

The $\hat{\chi}^2_{\rm min}$ goodness of fit is a statistical test used to evaluate how well a theoretical model fits to a set of observed data points, and is described in this Section. This test will be used in Sec. 4.4 to evaluate the goodness of the fit for the C_ℓ-fiducial and the moment expansion models of pipeline A. We also make use of a refined model comparison tool that quantifies the preference toward either model, considering every individual simulation seed: the *Akaike Information Criterion* (AIC, [213]). This is also described below.

4.2.2.1 $\hat{\chi}^2_{\rm min}$ goodness of fit

Aside from studying the data products at the end of the B-mode analysis, namely the clean CMB power spectra and the marginalised r posterior, we can take a step back and ask how well the input data are described by our theoretical models. With pipelines B and C, the input data are multi-frequency maps that are not directly compared to a likelihood model, instead a map-based cleaning step produces clean CMB maps that are fitted on the power spectrum level with a cosmological model, and therefore the goodness of the model fit with pipelines B and C will not be considered. With pipeline A, this question has a more straightforward answer: here, the input data are multi-frequency power spectra directly fitted by a C_ℓ-likelihood model describing CMB, Galactic synchrotron and dust. Another interesting aspect that can be addressed in the same context is the comparison of the goodness of the fit between the C_ℓ-fiducial and the C_ℓ-moments model. We will therefore focus on pipeline A and A + moments. The $\hat{\chi}^2_{\rm min}$ statistic is defined as

$$\hat{\chi}^2_{\rm min} = (\mathbf{d}-\mathbf{t})^T \mathsf{C}^{-1}(\mathbf{d}-\mathbf{t})\,, \tag{4.2.2.1}$$

where $\mathbf{d}$ is the data vector, $\mathbf{t}$ is the best-fit theory prediction, and C is the data covariance. From the definition of the likelihood in Eq. 2.3.3.6, it is easily found that $-2\ln\hat{\mathcal{L}} \propto \chi^2$, therefore, maximising the likelihood function is equivalent to minimising the χ^2 statistic. It can be shown that for Gaussian distributed data, the minimum-χ^2 statistic follows a χ^2 distribution with $N-P$ degrees of freedom, where N is the number of data points and P is the number of parameters (see e.g. [214]). In our case ($N-P=567$), this is approximately a Gaussian distribution with mean $N-P$ and variance $2(N-P)$. If we use Gaussian foregrounds, our likelihood is accurate since a) the model is known, b) the power spectrum data are very close to Gaussian (owing to the central limit theorem), and c) the covariance is accurate, since it was estimated from the corresponding set of simulations. If our guess of the covariance was an over- or under-estimation of the data covariance, then the $\hat{\chi}^2_{\rm min}$ would peak at lower, or higher values, respectively. The $\hat{\chi}^2_{\rm min}$ is calculated and discussed in Sec. 4.4.3 for the models of pipeline A (fiducial or moment-based expansion, described in Sec. 4.2.1).

4.2.2.2 Akaike Information Criterion (AIC)

While the $\hat{\chi}^2_{\rm min}$ statistic evaluates the goodness of the fit for each model independently, we also make use of a refined model comparison tool: the *Akaike Information Criterion* (AIC, [213]). The AIC is a statistical measure used to compare the goodness of fit of different models to a given dataset, and is defined as:

$$\mathrm{AIC} = \hat{\chi}^2_{\rm min} + 2k. \tag{4.2.2.2}$$

where k is the number of free parameters in the model. The AIC penalises the model for having more free parameters, which can lead to overfitting of the data

Frequency	FWHM	Baseline $N_{\rm white}$	Goal $N_{\rm white}$	Pessimistic $\ell_{\rm knee}$	Optimistic $\ell_{\rm knee}$	$\alpha_{\rm knee}$
27	91	46	33	30	15	-2.4
39	63	28	22	30	15	-2.4
93	30	3.5	2.5	50	25	-2.5
145	17	4.4	2.8	50	25	-3.0
225	11	8.4	5.5	70	35	-3.0
280	9	21	14	100	40	-3.0

Table 4.2: Instrument and noise specifications used to produce the simulations in this work. Note that these levels correspond to homogeneous noise, while our default analysis assumes noise maps weighted according to the SAT hits map.

and a decrease in the model's ability to make accurate predictions for new data. Thus, a simpler model is preferred over a more complex one, provided that it fits the data almost as well, and the best model is the one with the lowest AIC value. The difference in AIC values between two models provides a measure of their relative likelihood, with a higher difference indicating strong evidence in favour of the model with the lower AIC value. This is used in Sec. 4.4.3 to quantify the preference toward either model of pipeline A, considering every individual simulation seed. Specifically, we compute the difference $\Delta\mathrm{AIC} = 2\Delta k + \hat{\chi}^2_{\rm min}(\text{moments}) - \hat{\chi}^2_{\rm min}(\text{fiducial})$, where $\Delta k = 4$ is the number of excess parameters of the C_ℓ-moments over C_ℓ-fiducial, penalising the former for its higher complexity. Following e.g. [215], the number $\exp(\Delta\mathrm{AIC}/2)$ can be interpreted as the relative model odds $\mathcal{P}(C_\ell - \text{fiducial})/\mathcal{P}(C_\ell - \text{moments})$.

4.3 Simulations

A set of dedicated simulations was produced to test our data analysis pipelines and compare results. The simulated maps include cosmological CMB signal, Galactic foreground emission as well as instrumental noise.

4.3.1 Instrumental specifications and noise

We simulate polarised Stokes Q and U sky maps as observed by the SO-SAT telescopes. All maps are simulated using the `HEALPix` pixelation scheme [110] with resolution parameter $N_{\rm side} = 512$.

We model the SO-SAT noise power spectra as

$$N_\ell = N_{\rm white}\left[1 + \left(\frac{\ell}{\ell_{\rm knee}}\right)^{\alpha_{\rm knee}}\right], \tag{4.3.1.1}$$

where $N_{\rm white}$ is the white noise component while $\ell_{\rm knee}$, and $\alpha_{\rm knee}$ describe the contribution from $1/f$ noise. Following [108] (hereinafter SO2019), we consider different possible scenarios: "baseline" and "goal" levels for the white noise component, and "pessimistic" and "optimistic" correlated noise. The values of $(N_{\rm white}, \ell_{\rm knee}, \alpha_{\rm knee})$ associated with the different cases are reported in Table 4.2. Note that the values of $N_{\rm white}$ correspond to noise on a sky fraction of $f_{\rm sky} = 10\%$ and 5 years of observation time, as in SO2019. Differently from SO2019, we cite polarisation noise levels corresponding to a uniform map coverage, accounting for the factor of ~ 1.3 difference compared to Table 1 in SO2019. We simulate noise maps as Gaussian realisations of the N_ℓ power spectra. In our main analysis, we use noise maps with pixel weights computed from the SO-SAT hits map (see Figure 3.3) and refer to this as "inhomogeneous noise". In Section 4.4.1, all results assume inhomogeneous noise. Note that, although inhomogeneous, the noise realisations used here lack some of the important anisotropic properties of realistic $1/f$ noise, such as stripes due to the scanning strategy. Thus, together with the impact of other time-domain effects (e.g. filtering), we leave a more thorough study of the impact of instrumental noise properties for future work.

In all simulations, the CMB contribution was generated as described in Sec. 3.3.1.

4.3.2 Foregrounds

We use four sets of models to simulate the sky emission:

- **Gaussian foregrounds**: thermal dust emission and synchrotron radiation are simulated as Gaussian realisations of power law EE and BB power spectra, as described in Sec. 3.3.2. This idealised model has been used to validate the different pipelines and to build approximate signal covariance matrices from 500 random realisations.

- `d0s0` **model**: the `d0s0` `PySM` model includes the same templates for both dust and synchrotron emission specified in Sec. 3.3.3 for `d1s1`, but the SEDs are considered to be uniform across the sky.

- `d1s1` **model**: this is equivalent to the `d1s1` model described in Sec. 3.3.3, which includes spatial variability for spectral parameters.

- `dmsm` **model**: for synchrotron, this is equivalent to the `dmsm` model described in Sec. 3.3.3, while dust is a simplified version, where we smoothed the β_d and T_d templates at an angular resolution of 2 degrees, in order to down-weight the contribution of instrumental noise fluctuations in the original `PySM` maps.

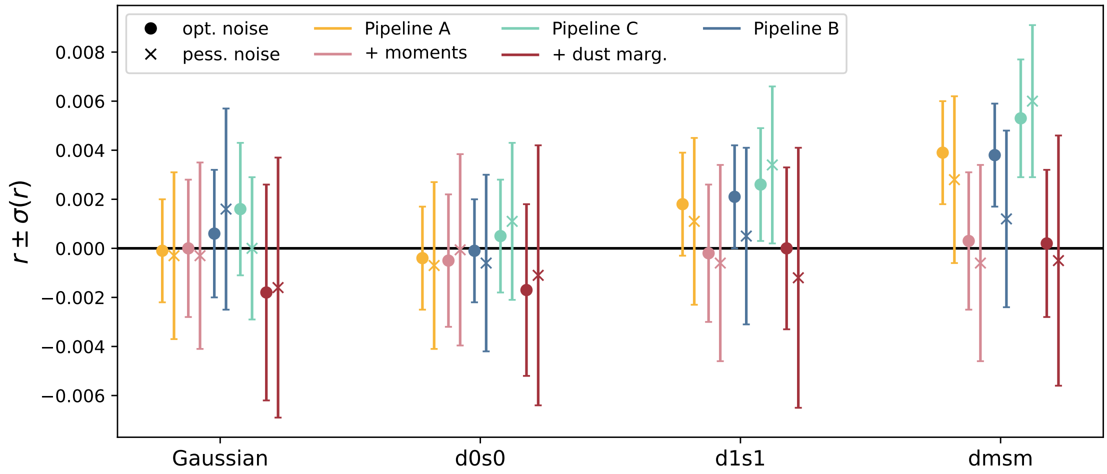

Figure 4.1: Compilation of the mean r with its (16, 84)% credible interval as derived from 500 simulations, applying the three nominal pipelines (plus extensions) to four foreground scenarios of increasing complexity. We assume a fiducial cosmology with $r = 0$ and $A_{\rm lens} = 1$, inhomogeneous noise with goal sensitivity and optimistic $1/f$ noise component (dot markers), and inhomogeneous noise with baseline sensitivity and pessimistic $1/f$ noise component (cross markers). Note that the NILC results for Gaussian foregrounds are based on a smaller sky mask

As for the CMB simulations, the multi-frequency foreground maps at the SO reference frequencies have been convolved with Gaussian beams, to reach the expected angular resolution. We assume delta-like frequency bandpasses in order to accelerate the production of these simulations, although all pipelines are able to handle finite bandpasses. Therefore this approximation should not impact the performance of any of the pipelines presented here.

4.4 Results

Simulations were generated for 4 different noise and foreground models, respectively, (see Section 4.3), for a total of 16 different foreground-noise combinations. For the main analysis, we consider a fiducial CMB model with $A_{\rm lens} = 1$ (no delensing) and $r = 0$ (no primordial tensor fluctuations). In addition, we explored three departures from the fiducial CMB model, with input parameters ($A_{\rm lens} = 0.5$, $r = 0$), ($A_{\rm lens} = 1$, $r = 0.01$) and ($A_{\rm lens} = 0.5$, $r = 0.01$). Here we report the results found for all these cases.

4.4.1 Constraints on r

Let us start by examining the final constraints on r obtained by each pipeline applied to 500 simulations. These results are summarised in Figure 4.1, which shows the mean r and (16, 84)% credible intervals found by each pipeline as a function of the input foreground model (labels on the x axis). Results are shown for 5 pipeline setups: pipeline A using the C_ℓ-fiducial model (red) and using the C_ℓ-moments model (yellow), pipeline B (blue), pipeline C (green) and including the marginalisation over A_d (cyan). For each pipeline, we show two points with error bars. The dot markers and smaller error bars correspond to the results found in the best-case instrument scenario (goal noise level, optimistic $1/f$ component), while the cross markers and larger error bars correspond to the baseline noise level and pessimistic $1/f$ component.

We will first discuss the nominal pipelines A, B and C without considering any extensions. We find that for the simpler Gaussian and `d0s0` foregrounds, the nominal pipelines obtain unbiased results, as expected. Pipeline B shows a slight positive bias for Gaussian foregrounds, which can be traced back to the pixel covariance matrix used to construct the NILC weights. The more complex `d1s1` foregrounds lead to a $\sim 1\sigma$ bias in the goal and optimistic noise scenario. The `dmsm` foregrounds lead to a noticeable increase of the bias of up to $\sim 2\sigma$, seen with pipeline C in all noise scenarios and with pipeline A in the goal-optimistic case, and slightly less with pipeline B. The modifications introduced in the `dmsm` foreground model include a larger spatial variation in the synchrotron spectral index β_s with respect to `d1s1`, and are a plausible reason for the increased bias on r.

Remarkably, we find that, in their simplest incarnation, all pipelines achieve comparable statistical uncertainty on r, ranging from $\sigma(r) \simeq 2.1 \times 10^{-3}$ to $\sigma(r) \simeq 3.6 \times 10^{-3}$ (a 70% increase), depending on the noise model. Changing between the goal and baseline white noise levels results in an increase of $\sigma(r)$ of $\sim 20 - 30\%$. Changing between the optimistic and pessimistic $1/f$ noise has a similar effect on the results from pipelines A and B, although $\sigma(r)$ does not increase by more than 10% when changing to pessimistic $1/f$ noise for pipeline C. These results are in reasonable agreement with the forecasts presented in [108].

Now let us have a look at the pipeline extensions, A + moments and C + dust marginalisation. Notably, in all noise and foreground scenarios, the two extensions are able to reduce the bias on r to below 1σ. For the Gaussian and `d0s0` foregrounds, we consistently observe a small negative bias (at the $\sim 0.1\sigma$ level for A + moments and $< 0.5\sigma$ for C + dust marginalisation). This bias may be caused by the introduction of extra parameters that are prior dominated, like the dust template's amplitude in the absence of residual dust contamination, or the moment parameters in the absence of varying spectral indices of foregrounds. If those extra parameters are weakly degenerate with the tensor-to-scalar ratio, the marginal r

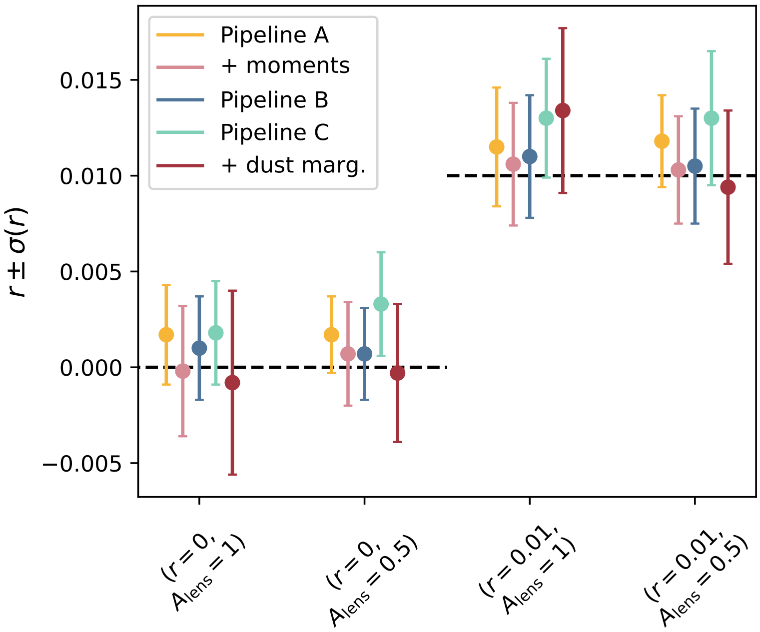

Figure 4.2: Results on r and the (16, 84)% credible interval derived from 500 simulations, applying the three nominal pipelines plus extensions, assuming input models including primordial B-modes and 50% delensing efficiency. We assume baseline noise level with optimistic $1/f$ component and the `d1s1` foreground template.

posterior will shift according to the choice of the prior on the extra parameters. The observed shifts in the tensor-to-scalar ratio and their possible relation with these volume effects will be investigated in a future work. For the more complex `d1s1` and `dmsm`, both pipeline extensions effectively remove the bias observed in the nominal pipelines, achieving a $\sim 0.5\sigma$ bias and lower.

The statistical uncertainty $\sigma(r)$ increases for both pipeline extensions, although by largely different factors. While C + dust marginalisation yields $\sigma(r)$ between 3.0×10^{-3} and 5.9×10^{-3}, the loss in precision for A + moments is significantly smaller, with $\sigma(r)$ varying between 2.7×10^{-3} and 4.0×10^{-3} depending on the noise scenario, an average increase of $\sim 25\%$ compared to pipeline A. In any case, within the assumptions made regarding the SO noise properties, it should be possible to detect a primordial B-mode signal with $r = 0.01$ at the 2–3σ level with no delensing. The impact of other effects, such as time domain filtering or anisotropic noise may affect these forecasts, and will be studied in more detail in the future.

This analysis was repeated for input CMB maps generated assuming $r = 0$ or 0.01 and $A_{\rm lens} = 0.5$ or 1. For simplicity, in these cases we considered only the baseline white noise level with optimistic $1/f$ noise, and the moderately complex `d1s1` foreground model. The results can be found in Figure 4.2. A 50% delensing efficiency results in a reduction in the final $\sigma(r)$ by 25-30% for pipelines A and B,

$\sim$10-20% for A + moments, and 0-33% for C + dust marginalisation. The presence of primordial B-modes with a detectable amplitude increases the contribution from cosmic variance to the error budget, with $\sigma(r)$ growing by up to 40% if $r = 0.01$, in agreement with theoretical expectations. Using C + dust marginalisation, and considering no delensing, we even find $\sigma(r)$ decreasing, hinting at the possible breaking of the degeneracy between r and $A_{\rm dust}$. We conclude that all pipelines are able to detect the $r = 0.01$ signal at the level of $\sim 3\sigma$. As before, we observe a 0.5-1.2σ bias on the recovered r that is eliminated by both the moment expansion method and the dust marginalisation method.

4.4.2 More complex foregrounds: `d10s5`

During the completion of this work, the new `PySM3` Galactic foreground models were made publicly available. In particular, in these new models, templates for polarised thermal dust and synchrotron radiation have been updated including the following changes:

1. Large-scale thermal dust emission is based on the GNILC maps [164], which present a lower contamination from CIB emission with respect to the `d1` model, based on `Commander` templates.

2. For both thermal dust and synchrotron radiation small scales structures are added by modifying the logarithm of the polarisation fraction tensor[2].

3. Thermal dust spectral parameters are based on GNILC products, with larger variation of β_d and T_d parameter at low resolution compared to the `d1` model. Small-scale structure is also added as Gaussian realisations of power-law power spectra.

4. The new template for β_s includes information from the analysis of S-PASS data [148], in a similar way as the one of the `sm` model adopted in this work. In addition, small-scale structures are present at sub-degree angular scales.

These modifications are encoded in the models called `d10` and `s5` in the updated version of `PySM`. Although these models are still to be considered preliminary, both in terms of their implementation details in `PySM`[3] and in general, being based on datasets that may not fully involve the unknown level of foreground complexity, we

[2] See pysm3.readthedocs.io/en/latest/preprocess-templates.

[3] The `PySM` library is currently under development, with beta versions including minor modifications of the foreground templates being realised regularly. In this part of our analysis we make use of `PySM v3.4.0b3`.

decided to dedicate an extra section to their analysis. For computational speed, we ran the five pipeline set-ups on a reduced set of 100 simulations containing the new `d10s5` foregrounds template, CMB with a standard cosmology ($r = 0$, $A_{\rm lens} = 1$) and inhomogeneous noise in the goal-optimistic scenario. The resulting marginalised posterior mean and (16, 84)% credible intervals on r, averaged over 100 simulations, are:

$$\begin{aligned}
r &= 0.0194 \pm 0.0021 && \text{(pipeline A)} \\
\mathbf{r} &= \mathbf{0.0025 \pm 0.0031} && (\textbf{A + moments}) \\
r &= 0.0144 \pm 0.0023 && \text{(pipeline B)} \\
r &= 0.0220 \pm 0.0026 && \text{(pipeline C)} \\
r &= -0.0015 \pm 0.0051 && \text{(C + dust marg.)}
\end{aligned} \tag{4.4.2.1}$$

Note that the respective bias obtained with pipelines A, B and C are at 9, 6 and 8σ, quadrupling the bias of the `dmsm` foreground model. Crucially, this bias is reduced to $\sim 0.8\sigma$ with the A + moments pipeline, with a 40% increase in $\sigma(r)$ compared to pipeline A, and $\sim 0.3\sigma$ with the C + dust-marginalisation pipeline, with a 95% increase in $\sigma(r)$ compared to pipeline C. This makes A + moments the unbiased method with the lowest statistical error, and was therefore highlighted above.

4.4.3 Model comparison for pipeline A

The $\hat{\chi}^2_{\rm min}$ described in Sec 4.2.2 for 100 simulated data realisations containing standard cosmology ($r = 0$, $A_{\rm lens} = 1$), inhomogeneous noise in the goal-optimistic scenario, is calculated and discussed for the models of pipeline A (fiducial or moment-based expansion, described in Sec. 4.2.1), and in four foreground cases (Gaussian, `d0s0`, `d1s1`, `dmsm`).

As shown in Figure 4.3, the C_ℓ-fiducial model achieves minimum χ^2 values of $\approx 600 \pm 40$. Although this is an increase of $\Delta\chi^2 \sim 30$ with respect to the less complex foreground simulations, the associated probability to exceed is $p \sim 0.16$ (assuming our null distribution is a χ^2 with $N_{\rm data} - N_{\rm parameters} = 567$ degrees of freedom), therefore it would not be possible to identify the presence of a foreground bias by virtue of the model providing a bad fit to the data. On the other hand, A + moments achieves best-fit χ^2 values of $\approx 531 \pm 33$, with a $< 5\%$ decrease in $\Delta\chi^2$ with respect to less complex foreground simulations, indicating an improved fitting accuracy.

We also find that the empirical distribution of the $\hat{\chi}^2_{\rm min}$ statistic matches the theoretical expectation not only for Gaussian foregrounds, but also the non-Gaussian cases. This shows that our estimate of the covariance matrix is appropriate, even for the more complex foreground models.

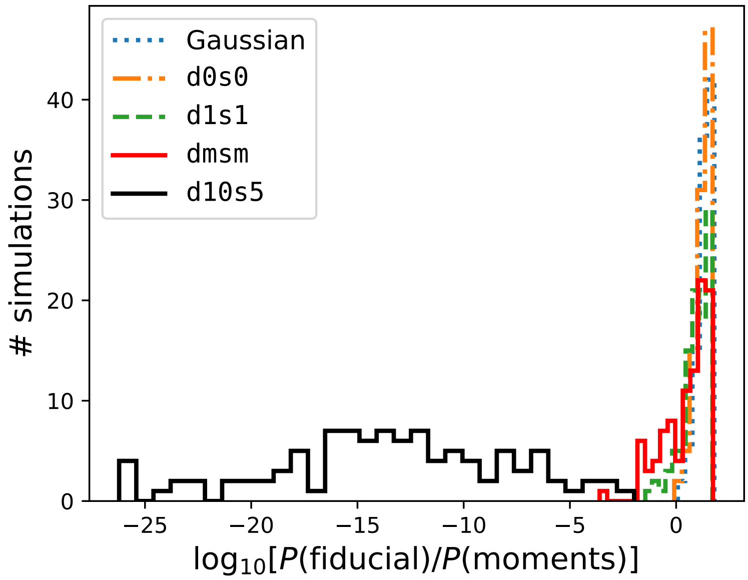

Figure 4.3: Empirical distribution of the AIC-based relative model odds (see Sec 4.2.2 for details) between the C_ℓ-fiducial and the C_ℓ-moments model, evaluated for 100 simulations and generated with five different Galactic foreground templates, including the new `PySM` foreground model `d10s5`. We find strong preference for the C_ℓ-moments model in the `d10s5` foreground scenario, and only then.

The relative model odds between C_ℓ-fiducial and C_ℓ-moments vary between 10^{-26} and 10^{-2}, clearly favouring C_ℓ-moments. Out of 100 `d10s5` simulations, 99 yield model odds below 1% and 78 below 10^{-5}. As opposed to the less complex foreground simulations (`d0s0`, `d1s1` and `dmsm`), `d10s5` gives strong preference to using the moment expansion in the power spectrum model. Note that the AIC-based model odds are computed from the differences of χ^2 values that stem from the same simulation seed, while the χ^2 analysis only considers the models individually, which explains why the former makes for a much more powerful model comparison test.

These results consider only the most optimistic noise scenario. Other cases would likely lead to larger uncertainty and, as a consequence, lower relative biases. In this regard, it is highly encouraging to see two pipeline extensions being able to robustly separate the cosmological signal from Galactic synchrotron and dust emission with this high-level complexity. This highlights the importance of accounting for and marginalising over residual foreground contamination due to frequency decorrelation for the level of sensitivity that SO and other next-generation observatories will achieve. These results should also be accompanied by a comprehensive set of robustness tests able to identify signatures of foreground contamination in the data. This will form the basis of a future work.

4.5 Conclusions

In this chapter, we have presented three different component separation pipelines designed to place constraints on the amplitude of cosmological B-modes on polarised maps of the SO Small Aperture Telescopes. The pipelines are based on multi-frequency C_ℓ parametric cleaning (Pipeline A), blind Needlet ILC cleaning (Pipeline B), and map-based parametric cleaning (Pipeline C). We have also studied extensions of pipelines A and C that marginalise over additional residual foreground contamination, using a moment expansion or a dust power spectrum template, respectively. We have tested and compared their performance on a set of simulated maps containing lensing B-modes with different scenarios of instrumental noise and Galactic foreground complexity. The presence of additional instrumental complexity, such as time-domain filtering, or anisotropic noise, are likely to affect our results. The impact of these effects will be more thoroughly studied in future work.

We find the inferred uncertainty on the tensor-to-scalar ratio $\sigma(r)$ to be compatible between the three pipelines. While the simpler foreground scenarios (Gaussian, `d0s0`) do not bias r, spectral index variations can cause an increased bias of 1–2σ if left untreated, as seen with more complex foreground scenarios (`d1s1`, `dmsm`). Modelling and marginalising over the spectral residuals is vital to obtain unbiased B-mode estimates. The extensions to pipelines A and C are thus able to yield unbiased estimates on all foreground scenarios, albeit with a respective increase in $\sigma(r)$ by $\sim 20\%$ (A + moments) and $> 30\%$ (C + dust marginalisation). These results are in good agreement with the forecasts presented in [108]. After testing on simulations with an $r = 0.01$ cosmology, we conclude that under realistic conditions, SO should be able to detect a $r = 0.01$ signal at $\sim$ 2–3σ after 5 years of observation.

We have also carried out a preliminary analysis of new, more complex, foreground models recently implemented in `PySM3`, in particular the `d10s5` foreground template. The much higher level of spatial SED variation allowed by this model leads to a drastic increase in the bias on r by up to $\sim 9\sigma$, when analysed with the nominal pipelines A, B and C. Fortunately, this bias can be reduced to below 1σ when using A + moments and C + dust marginalisation. These extensions lead to a 40% and 95% degradation of the error bars, respectively. Our results highlight the importance of marginalising over residuals caused by frequency decorrelation for SO-like sensitivities. Although we have not analysed `d10s5` to the same depth as the other foreground models presented here, it is encouraging to confirm that we have the tools at hand to obtain robust, unbiased constraints on the tensor-to-scalar ratio in the presence of such complex Galactic foregrounds.

Chapter 5

A hybrid map-C_ℓ component separation method for primordial CMB B-mode searches

Contents

5.1 Introduction

In the context of foreground component separation techniques, a generalisation of the C_ℓ*-based* "multi-frequency power spectrum likelihood" approach [101, 102, 169, 170, 216] was proposed in Chapter 3, using a minimal moment expansion of the foreground SEDs to account for both the amplitude and scale dependence of the effect induced by spatial variation of spectral parameters, within a consistent physical model. Chapters 3-4 have shown that the method can significantly improve over the vanilla approach, being able to mitigate the foreground bias for realistic sky models over areas much larger than that covered by BICEP-*Keck*. This minimal expansion can be fully generalised, including all potential correlations between different spectral parameter variations and amplitudes, and avoiding any assumptions about their non-Gaussianity. This has been shown [2, 217]to yield unbiased results given a sufficiently large number of frequency channels. However, this methodology suffers from two main shortcomings: first, it implicitly treats foregrounds as statistically isotropic fields by compressing their information into the power spectrum, ignoring couplings between different angular modes, and approximating the power spectrum covariance as that of a Gaussian field. Secondly, by adopting a simple functional form for the power spectrum of the foreground amplitudes, which are the dominant sky components, the model may lead to biases if the real sky deviates from this functional form on the range of scales explored. This is particularly relevant for experiments targeting the low-ℓ regime, but it may be important at any scale given enough sensitivity.

At the other end of the spectrum, parametric *map-based methods* aim to describe the sky at the map level, separating the CMB component and using it to constrain cosmological parameters [161, 162, 164, 167, 176, 218–220]. An example was provided in Ch. 4, described under the name of "pipeline B". Note that, while these methods mostly exploit the frequency information to separate components at the map level, recently studies have been carried out on techniques to separate out Galactic foreground purely based on their spatial statistics [185, 221]. Map-based methods have several advantages: first, they are manifestly optimal for any specified sky model[1]. Secondly, since all components are modelled at the map level, it is natural and simple to account for spatial variations in the foreground spectra. Finally, foreground amplitudes are reconstructed in each pixel without assuming a specific model for their spatial correlations, thus avoiding potential biases inherent in assuming a functional form for the foreground power spectra. This approach is not without shortcomings, however: first, while it is easy to incorporate spatially varying spectra in the model, controlling the scales over which spectral parameters vary is not straightforward. The most conservative approach, allowing for uncorrelated variations over all pixels in the map, can lead to significant degradation in the final cosmological constraints, especially if only a small number of frequency channels are available [158, 167, 222]. Several methods have been proposed to select larger regions over which spectral parameters can be held constant or develop strong correlations, leading to rich and complex Bayesian forward modelling schemes [167, 168, 223]. The computational complexity of the resulting frameworks, particularly if

[1] Although C_ℓ methods will always achieve equivalent constraints if all components are Gaussian.

simultaneously sampling the CMB power spectra, is another disadvantage of these methods. While computing time is often a fair price to pay for improved or more robust constraints, it is often necessary to carry out multiple analyses of the same data before a clear understanding of all relevant sources of systematic uncertainty is achieved. Thus, fast, lightweight methods are extremely useful when confronted with new data. Finally, incorporating the impact of real-world imperfections in the data (e.g. filtering, gain or polarisation angle variations) is often not straightforward in map-based methods, or may increase their computational complexity significantly.

In this chapter, we explore an alternative method combining features from map-based and C_ℓ-based analyses. In this hybrid approach, foreground amplitudes are determined at the map level, free of assumptions regarding their statistical correlations, assuming homogeneous spectral properties. We then subtract the contribution from these foregrounds from the data, and model the residual foreground contamination (due to imperfect subtraction or spatial SED variations), at the power spectrum level. The methods is thus robust to most assumptions (Gaussianity, power spectrum form) about the dominant foreground amplitudes, while limiting the degradation in the final constraints caused by spectral index variations by imposing those assumptions on the smaller foreground residuals, while retaining information about their frequency dependence. We will study the applicability of this method for next-generation ground-based B-mode experiments targetting large sky areas ($f_{\rm sky} \gtrsim 0.1$) in the presence of inhomogeneous, correlated instrumental noise, for foreground models with varying degrees of complexity. We will also study the potential limitations of the method when employed on larger sky fractions, with consequently larger foreground contamination, quantifying its applicability for future space-based missions.

The chapter is structured as follows. Section 5.2 presents the method in detail. The simulations used to validate it are described in Section 5.3, including the instrumental setup assumed. The method is then validated in Section 5.4, where we explore its performance as a function of foreground complexity and sky fraction. The main conclusions of this study are summarised in Section 5.5.

5.2 Formalism

The component separation method proposed here consists of two stages:

1. We find a set of well-educated constant spectral indices $\beta_{\rm S}$ and $\beta_{\rm D}$ and clean out the spatially-constant part of the foregrounds at the map-level using the methods described in [158, 166–168, 224]. In practice, this step returns a filter

Q that de-projects the linear combinations of the data that follow the best-fit SEDs of the foreground sources under consideration (dust and synchrotron in this case). This step is described in Section 5.2.1.

2. We use the results of the previous step to measure the auto- and cross-power spectra of the map residuals at all frequencies after removing the constant-SED foreground components. Then, under the assumption that the residual foreground contamination is small, we can model its contribution to the multi-frequency power spectrum assuming a relatively simple model, with the frequency dependence given by the derivatives of the foreground SEDs with respect to the spectral indices. As done in Chapter 3, we assume a power law model for the scale dependence of the foreground residual power spectra. This step is described in Section 5.2.2.

5.2.1 Map-domain cleaning

5.2.1.1 Map-level parametrisation

At the map level, we model the sky signal using matrix notation as:

$$\mathbf{m} = \mathsf{S}\,\mathbf{T} + \mathbf{n}. \tag{5.2.1.1}$$

where

- $\mathbf{m}$ is the map vector containing the measured polarised sky signal $(Q_\nu(\hat{n}), U_\nu(\hat{n}))$ at angular coordinates $\hat{n}$ in each observing frequency ν.
- S is the so-called mixing matrix, with dimension $N_\nu \times N_c$, where N_c is the number of components, and N_ν is the number of frequencies. This matrix encodes the emission law $S^c_\nu(\beta_c(\hat{n}))$ of each component. We assume the same SEDs for both Q and U.
- $\mathbf{T}$ is a vector of size N_c containing the amplitude of each component for a given pixel.
- $\mathbf{n}$ is a vector containing the instrumental noise at different frequencies for a given pixel. Its covariance matrix will be labelled $\mathsf{N} \equiv \langle \mathbf{n}\mathbf{n}^T \rangle$.

Note that, implicitly, we have assumed that all quantities in Eq. 5.2.1.1, including the foreground spectra, depend on sky coordinates $\hat{n}$.

We consider three components, with a given spectral dependence, as described in Sec. 3:

1. **Thermal dust:** as described in Sec. 2.5, dust is the most relevant foreground for B-mode searches on frequencies $\nu \gtrsim 100\,\mathrm{GHz}$. Thermal dust emission is thought to be well-characterised by a modified black-body spectrum, of the form given by Eq. 3.2.1.3, parametrised by β_D and T_D, the dust spectral index and temperature, respectively.

2. **Synchrotron:** as described in Sec. 2.5, synchrotron is strongly polarised, and is characterised by a smooth power-law spectrum tracing the energy distribution of cosmic ray electrons. Thus, the synchrotron spectrum used here is modelled as a power law, a described by Eq. 3.2.1.5, and parametrised by β_S, the synchrotron spectral index.

3. **CMB:** the spectrum of CMB temperature anisotropies is expressed in antenna temperature units, and described by Eq. 3.2.1.2.

We will consider spatial variations in β_D, which takes values $\beta_\mathrm{D} \sim 1.6$. The restricted frequency range available to most ground-based experiments, including the Simons Observatory ($\nu \lesssim 280\,\mathrm{GHz}$) makes B-mode studies almost insensitive to the value of T_D, and therefore we fix it to $T_\mathrm{D} = 19.6\,\mathrm{K}$ here. β_S is the synchrotron spectral index, which takes values $\beta_\mathrm{S} \sim -3$.

We can separate the spectra S into two parts: a spatially-independent "mean" $\bar{\mathsf{S}}$ and a perturbation $\Delta\mathsf{S}$, such that Eq. 5.2.1.1 becomes

$$\mathbf{m} = \bar{\mathsf{S}}\mathbf{T} + \Delta\mathsf{S}\mathbf{T} + \mathbf{n}. \tag{5.2.1.2}$$

In the first term, $\bar{\mathsf{S}}$ represents our best guess of the mean spectra of all components. Here, we will determine best-fit constant spectral indices $(\beta_\mathrm{S}, \beta_\mathrm{D})$ using the map-based spectral likelihood described in Section 5.2.1.2, and $\bar{\mathsf{S}}$ is given by evaluating the spectra above at these best-fit values. The residual $\Delta\mathsf{S}$ includes the spatially-dependent part of the foreground spectra, as well as any difference between the estimated best-fit spectral indices and their true spatial average. The subsequent treatment of $\Delta\mathsf{S}$ does not change as long as these differences are small.

5.2.1.2 Map-level spectral likelihood

The procedure to obtain the best-fit spectral indices at the map level follows the methods described in [158, 166, 167]. Given the model in Eq. 5.2.1.1, the posterior marginalised distribution for $\boldsymbol{\beta} \equiv (\beta_\mathrm{D}, \beta_\mathrm{S})$ is given by

$$-2\ln\mathcal{L}_\beta = \left(\bar{\mathsf{S}}^T\mathsf{N}^{-1}\mathbf{m}\right)^T\left(\bar{\mathsf{S}}^T\mathsf{N}^{-1}\bar{\mathsf{S}}\right)^{-1}\left(\bar{\mathsf{S}}^T\mathsf{N}^{-1}\mathbf{m}\right), \tag{5.2.1.3}$$

where N is the noise covariance matrix. As described in [158], this results from analytically marginalising over the linear amplitude parameters $\mathbf{T}$, and imposing a Jeffreys-like prior that corrects for volume effects in the spectral index parameters due to their non-linear nature in the model. We obtain the best-fit spectral indices $\boldsymbol{\beta}_{\rm BF}$, defining $\bar{\mathsf{S}}$, as the values that maximise this likelihood, $\boldsymbol{\beta}_{\rm BF} \equiv {\rm argmax}(\mathcal{L}_\beta)$, assuming that they are homogeneous across the sky. Furthermore, we use a Fisher-matrix approach to estimate the uncertainty on the spectral parameters. The inverse covariance of $\boldsymbol{\beta}_{\rm BF}$ is therefore approximated as the Hessian of the log-posterior at the best fit:

$$({\rm Cov}_\beta)^{-1}_{ij} \simeq -\left.\frac{\partial^2 \ln \mathcal{L}_\beta}{\partial\beta_i\partial\beta_j}\right|_{\beta=\beta_{\rm BF}} \tag{5.2.1.4}$$

This approximation is accurate for sufficiently constraining data [157, 225] (as is the case of SO), and allows us to avoid sampling the full two-dimensional parameter space (although this can be done if necessary)[2].

A number of technical aspects regarding the implementation of this method in practice must be noted at this stage. First, we assume that the data is provided in the form of splits, i.e. a set of maps containing the same sky but different noise realisations (e.g. corresponding to observations at different times). We use differences between these splits, containing only noise, to estimate the noise covariance matrix N. Moreover, we assume, for simplicity, that the noise covariance matrix is diagonal in pixel space (i.e. noise in different pixels are uncorrelated). In this work we also assume that the noise covariance is diagonal in frequency space, although accounting for off-diagonal correlations would pose no additional difficulty to the method. It is worth emphasizing that, even though our implementation of the method assumes uncorrelated noise for computational efficiency, this is never the case in real observations, nor in the simulations used here. As we will see, this does not spoil the validity of the method, since this first step only requires us to find spectral index values that are sufficiently close to the spatial mean over the observed footprint. Note that we do account for inhomogeneity in the noise component, using the hits count map. After calculating N, we coadd all data splits into a single set of maps, which we then process as described above to find $\bar{\mathsf{S}}$.

[2] The pipeline implemented is based on the public code available at github.com/damonge/BFoRe.

5.2.1.3 Map-level residuals

After determining the best-fit spectral indices and the corresponding mixing matrix $\bar{\mathsf{S}}$, the corresponding best-fit amplitudes can be found analytically by solving Eq. 5.2.1.2 as a least-squared problem:

$$\mathbf{T}_{\rm BF} = \left(\bar{\mathsf{S}}^T\mathsf{N}^{-1}\bar{\mathsf{S}}\right)^{-1}\bar{\mathsf{S}}^T\mathsf{N}^{-1}\mathbf{m}. \tag{5.2.1.5}$$

It is worth emphasizing that, at this stage, we have not imposed any priors on the statistical correlations of foreground amplitudes, and instead treated them as an independent parameter in each pixel. This is one of the key differences with the usual power-spectrum-based methods, where a functional form for the foreground power spectra is often assumed.

From $\mathbf{T}_{\rm BF}$ in Eq. 5.2.1.5 we can get a best-fit estimate of the foreground contribution to our observations as $\mathbf{m}^{\rm FG}_{\rm BF} = \bar{\mathsf{S}}\mathsf{P}\mathbf{T}_{\rm BF}$, where $\mathsf{P}\equiv{\rm diag}(1,1,...,1,0)$ is an $N_c\times N_c$ diagonal matrix that selects only the non-CMB components of $\mathbf{T}_{\rm BF}$ (we have assumed the CMB to be the last component). We can then subtract $\mathbf{m}^{\rm FG}_{\rm BF}$ from the data to obtain the residual maps:

$$\mathbf{r}\equiv\mathbf{m}-\mathbf{m}^{\rm FG}_{\rm BF} = \mathsf{Q}\mathbf{m}, \tag{5.2.1.6}$$

where the matrix Q is

$$\mathsf{Q}\equiv\mathbb{1}-\bar{\mathsf{S}}\mathsf{P}\left(\bar{\mathsf{S}}^T\mathsf{N}^{-1}\bar{\mathsf{S}}\right)^{-1}\bar{\mathsf{S}}^T\mathsf{N}^{-1}. \tag{5.2.1.7}$$

The matrix Q is a projector ($\mathsf{Q}^2=\mathsf{Q}$), and satisfies the nice property that $\mathsf{Q}\bar{\mathsf{S}} = \bar{\mathsf{S}}(\mathbb{1}-\mathsf{P})$, where $\mathbb{1}-\mathsf{P}$ selects only the CMB component. Note also that, since the spatial and frequency dependence of N is factorisable, Q is constant across the sky. Finally, it is straightforward to show that ${\rm Tr}(\mathsf{Q}) = N_\nu - N_c + 1$. Thus, in our case, with $N_c=3$, Q projects the N_ν-dimensional dataset onto a subspace of dimension $N_\nu-2$, eliminating two independent modes corresponding to synchrotron and dust assuming constant spectral indices (i.e. $\mathsf{Q}\,\bar{\mathsf{S}}^{\rm D}=\mathsf{Q}\,\bar{\mathsf{S}}^{\rm S}=0$).

Substituting our model for $\mathbf{m}$ (Eq. 5.2.1.1) into Eq. 5.2.1.6, and using the properties of Q, we obtain

$$\mathbf{r} = \mathsf{Q}\left(\bar{\mathsf{S}}\mathbf{T}+\Delta\mathsf{S}\mathbf{T}+\mathbf{n}\right), \tag{5.2.1.8}$$

$$= \mathbf{m}_{\rm CMB}+\mathsf{Q}\Delta\mathsf{S}\mathbf{T}+\mathsf{Q}\mathbf{n}, \tag{5.2.1.9}$$

where $(\mathbf{m}_{\rm CMB})_\nu(\hat{n})\equiv S^{\rm CMB}_\nu T_{\rm CMB}(\hat{n})$ is the CMB component of the maps. Since the CMB row of $\Delta\mathsf{S}$ is zero (given that the CMB spectrum is constant), the second term in the equation above only contains foreground sources.

Assuming that the fluctuations in the spectral indices with respect to their constant best-fit values are small, we can use a linear expansion of $\Delta\mathsf{S}$ around $\boldsymbol{\beta}_{\rm BF}$ (as in Chapter 3):

$$\Delta S^c_\nu(\hat{n}) = \partial_{\beta_c}\bar{S}^c_\nu\,\delta\beta_c(\hat{n}). \tag{5.2.1.10}$$

where $\partial_{\beta_c} \equiv \partial/\partial\beta_c$, and we do not implicitly sum over c. Then, we can rewrite Eq. 5.2.1.9 as:

$$r_\nu(\hat{n}) = S^{\rm CMB}_\nu T_{\rm CMB}(\hat{n}) + \tilde{S}^c_\nu\tilde{T}_c(\hat{n}) + \tilde{n}_\nu(\hat{n}), \tag{5.2.1.11}$$

where the modified spectra and amplitudes are:

$$\tilde{S}^c_\nu \equiv Q^{\nu'}_\nu\,\partial_{\beta_c}\bar{S}^c_{\nu'}, \quad \tilde{T}_c(\hat{n}) \equiv \delta\beta_c(\hat{n})T_c(\hat{n}), \tag{5.2.1.12}$$

and we implicitly sum over ν'. The pipeline that was implemented for this part of the study, `BBHybrid`, is publicly available[3]

It is worth noting that, if the spectral indices are truly constant across the sky, and if the inferred value of $\boldsymbol{\beta}$ is equal to the true one, then $\mathbf{r}$ in Eq. 5.2.1.6 is a set of multi-frequency maps, each containing the same copy of the CMB. Otherwise each map in Eq. 5.2.1.6 contains a spectral residual, with a frequency dependence given by the first-order expansion of each component's SED. We will then model and marginalise over these residuals at the power spectrum level, as described in the next section.

5.2.2 C_ℓ likelihood on deprojected spectral modes

To separate the CMB contribution at the level of the power spectrum, we use a multi-frequency power spectrum likelihood, similar to that used in the latest analysis carried out by the BICEP-*Keck* collaboration [89, 102]. In this case, the data vector is the full matrix of cross-power spectra $C_\ell^{\nu\nu'}$ between all pairs from a set of frequency maps[4]. The model used to describe this data vector, and the procedure used to estimate it from the data are described here.

[3]See github.com/susannaaz/BBHybrid.
[4]The code is available at github.com/simonsobs/BBPower.

5.2.2.1 Power spectrum model

Here, the map model described above gets propagated through to power spectrum. We start from a linear model of the form:

$$d_\nu(\hat{n}) = \mathcal{S}^c_\nu \mathcal{T}_c(\hat{n}) + n_\nu, \tag{5.2.2.1}$$

where $d_\nu(\hat{n})$ is a set of frequency maps, $\mathcal{T}_c(\hat{n})$ is a set of component amplitude maps, and $\mathcal{S}^c_\nu$ is a *constant* matrix characterising the spectra of the different components. Here we will explore two possible approaches:

1. In the **baseline** case, d_ν is the raw set of frequency maps, and $\mathcal{S}^c_\nu$ and $\mathcal{T}_c$ are the SEDs and amplitude maps of dust, synchrotron, and the CMB.

2. In the **hybrid** case, d_ν is the set of residual maps described in the previous section, and $(\mathcal{S}^c_\nu, \mathcal{T}_c)$ are given by the modified spectra and amplitudes $(\tilde{S}^c_\nu, \tilde{T}_c)$ in Eq. 5.2.1.12, complemented with the CMB SED and amplitude map.

Since we assume the mixing matrix $\mathcal{S}^c_\nu$ to be spatially constant, the theoretical prediction for the multi-frequency power spectrum of the map-level model in Eq. 5.2.2.1 is simply given by

$$C^{\nu\nu'}_\ell = \sum_{c,c'} \mathcal{S}^c_\nu \mathcal{S}^{\nu'}_{c'} C^{cc'}_\ell, \tag{5.2.2.2}$$

where $C^{cc'}_\ell$ is the cross-power spectrum between the amplitude maps of components c and c'.

Following the model used by [102], we parametrise $C^{cc'}_\ell$ for the different components described above as follows.

- Foreground auto-correlations are modelled as power-laws of the form:

$$\frac{\ell(\ell+1)}{2\pi} \mathsf{C}^{cc}_\ell = A_c \left(\frac{\ell}{\ell_0}\right)^{\alpha_c}, \tag{5.2.2.3}$$

  with $c \in \{\mathrm{D}, \mathrm{S}\}$, $\ell_0 \equiv 80$.

- The dust-synchrotron cross-correlation is parametrised through a scale-independent correlation coefficient ϵ_{DS}:

$$C^{\mathrm{DS}}_\ell = \epsilon_{\mathrm{DS}} \sqrt{C^{\mathrm{DD}}_\ell C^{\mathrm{SS}}_\ell} \tag{5.2.2.4}$$

- The CMB B-mode power spectrum is parametrised as

$$C_\ell^{\rm CMB} = A_{\rm lens}\, C_\ell^{\rm lens} + r\ \left. C_\ell^{\rm tens}\right|_{\rm r=1}\,, \tag{5.2.2.5}$$

  where $C_\ell^{\rm lens}$ and $C_\ell^{\rm tens}|_{\rm r=1}$ are templates for the B-mode power spectrum caused by gravitational lensing and by primordial tensor fluctuations with tensor-to-scalar ratio $r = 1$ respectively. We assume no correlation between the CMB and foreground components.

In both the "baseline" and "hybrid" approaches, the model has a total of 9 free parameters: $\boldsymbol{\theta} = \{A_{\rm lens}, r, A_{\rm D}, \alpha_{\rm D}, \beta_{\rm D}, A_{\rm S}, \alpha_{\rm S}, \beta_{\rm S}, \epsilon_{\rm DS}\}$. However, since the bulk of the foreground contamination is removed at the map level in the hybrid approach, we will fix $\epsilon_{\rm DS} = 0$ in that case, and explore the impact of freeing up this parameter. Note that, although the best-fit $\beta_{\rm D}$ and $\beta_{\rm S}$ have been determined in the first step of the "hybrid" method, we allow for them to vary in the power spectrum likelihood, with a prior given by the uncertainties calculated in that step.

5.2.2.2 Power spectrum likelihood

Given a set of measured power spectra $\hat{C}_\ell^{\nu\nu'}$, and their theoretical prediction outlined in the previous section, construct a Gaussian likelihood of the form

$$\chi^2 \equiv -2\ln p(\hat{\mathsf{C}}_\ell|\boldsymbol{\theta}) = \sum_{a,b}[C_\ell^{ij} - \hat{C}_\ell^{ij}]\,(\boldsymbol{\Sigma}^{-1})_{ij\ell,mn\ell'}[C_{\ell'}^{mn}(\theta) - \hat{C}_{\ell'}^{mn}], \tag{5.2.2.6}$$

where $\boldsymbol{\Sigma}$ is the covariance matrix of the power spectrum measurements. Note that, strictly speaking, power spectra of Gaussian fields follow a Wishart distribution, which can be approximated as proposed by [121], but we verified that the Gaussian approximation is sufficiently accurate on the scales used here ($30 \leq \ell \leq 300$). This approximation becomes inaccurate on larger scales, where the small number of available modes invalidates the central limit theorem and, as quadratic functions of Gaussian fields, power spectra follow a markedly non-Gaussian distribution. To account for this effect, we employ the method of [121] ("HL" hereon) when including scales $\ell < 30$ (see Section 5.4.3).

As described in Ch. 3, we estimate the cross-correlation of the data splits introduced in Section 5.2.1.1 in order to avoid modelling and subtracting the noise bias. We also make use of 100 Gaussian simulations, described in Section 5.3.2, to estimate the covariance matrix $\boldsymbol{\Sigma}$. In order to avoid numerical noise in the covariance, we set all off-diagonal elements coupling two ℓs separated by more than 2 bandpowers to zero. When applying our method to simulations with larger sky areas, in Section 5.4.3, we will approximate the covariance using the "Knox formula" [226]:

$$\Sigma_\ell^{a,b} \equiv \langle \Delta\hat{\mathsf{C}}_\ell^{a,b} \Delta\hat{\mathsf{C}}_\ell^{c,d} \rangle = \delta_{\ell\ell'} \frac{\mathsf{C}_\ell^{a,c}\mathsf{C}_\ell^{b,d} + \mathsf{C}_\ell^{a,d}\mathsf{C}_\ell^{b,c}}{(2\ell+1)\Delta\ell\, f_{sky}}, \tag{5.2.2.7}$$

where $\Delta\ell$ is the number of multipoles in the bandpower labelled with multipole ℓ, and $f_{\rm sky}$ is the usable sky fraction.

We measure power spectra using bandpowers of width $\Delta\ell = 10$. By default, we will only include scales in the range $30 \leq \ell \leq 300$, appropriate for a ground-based experiment, affected by atmospheric noise, targeting a B-mode detection from the recombination bump.

5.2.2.3 Residual C_ℓs and covariance

In both the "baseline" and "hybrid" approaches, the parameter inference starts from the full set of cross-correlations between all raw frequency maps, $\hat{\mathsf{C}}_\ell$. In the "baseline" approach, these are compared directly with the model in Section 5.2.2.1. In the "hybrid" approach, however, we must first construct the power spectra of the residual maps.

The residual maps $\mathbf{r}$ are related to the raw maps $\mathbf{m}$ via Eq. 5.2.1.9, where Q is the projector matrix defined in Eq. 5.2.1.7, and estimated in the first step of the method from the map-level likelihood. The power spectrum of the residual maps, C_ℓ^r is thus related to the power spectrum of the raw maps simply via

$$\mathsf{C}_\ell^r = \mathsf{Q} \cdot \mathsf{C}_\ell \cdot \mathsf{Q}^T. \tag{5.2.2.8}$$

Although it would be tempting to simply apply the modelling and likelihood described in the preceding sections to $\hat{\mathsf{C}}_\ell^r$ calculated as above from the raw power spectrum measurements, this is, unfortunately, not possible. Since Q projects out the two independent modes corresponding to synchrotron and dust with constant spectral indices, the rank of C_ℓ^r is $N_\nu - 2$, and not N_ν, and therefore it has null eigenvalues. This will also be the case for the power spectrum covariance, which is therefore not invertible. This makes the Gaussian likelihood in Eq. 5.2.2.6 numerically unstable.

To solve this problem, rather than making use of pseudo-inverses, we explicitly construct a new basis for our data in the subspace defined by the image of Q, and express the data and covariance in this lower-dimensional space. Since Q is a projector, and $\mathrm{Tr}(\mathsf{Q}) = N_\nu - 2$, we can diagonalise it as

$$\mathsf{Q} = \mathsf{B}\,\mathsf{D}\,\mathsf{B}^{-1}, \tag{5.2.2.9}$$

where B is a matrix containing the eigenvectors of Q as columns, and $\mathsf{D} = \mathrm{diag}(1, ..., 1, 0, 0)$ is the eigenvalue matrix. We have arbitrarily selected the eigenvector order such that the two null eigenvectors are at the end. Let $\tilde{\mathsf{D}}$ be the matrix that results from removing the last two null rows of D, and define the matrix $\mathsf{R} \equiv \tilde{\mathsf{D}}\mathsf{B}^{-1}$. R has dimensions $(N_\nu - 2) \times N_\nu$, and preserves the important property of Q that $\mathsf{R}\bar{\mathsf{S}}^{\mathrm{D}} = \mathsf{R}\bar{\mathsf{S}}^{\mathrm{D}} = 0$. Thus R projects general N_ν-dimensional vectors onto the $(N_\nu - 2)$-dimensional image of Q, explicitly reducing the dimensionality of the resulting vector. In order to obtain non-singular residual power spectra and covariances, we therefore use R instead of Q.

In summary, given the $N_\nu \times N_\nu$ matrix of multi-frequency power spectra $C_\ell^{\nu\nu'}$, and its covariance matrix, we construct the $(N_\nu - 2) \times (N_\nu - 2)$ residual power spectrum matrix

$$\tilde{C}_\ell^{\alpha\beta} \equiv R_{\alpha\nu} R_{\beta\nu'} C_\ell^{\nu\nu'}, \tag{5.2.2.10}$$

and its covariance matrix

$$\mathrm{Cov}\left(\tilde{C}_\ell^{\alpha\beta}, \tilde{C}_{\ell'}^{\gamma\delta}\right) = R_{\alpha\nu_1} R_{\beta\nu_2} R_{\gamma\nu_3} R_{\delta\nu_4} \, \mathrm{Cov}\left(C_\ell^{\nu_1\nu_2} C_\ell^{\nu_3\nu_4}\right). \tag{5.2.2.11}$$

Note that the transformation in Eq. 5.2.2.10 is applied to both the input power spectrum data, and to the theoretical prediction described in Section 5.2.2.1.

5.2.3 Summary of the method

The hybrid method proposed here can be summarised as follows.

1. Starting from a set of N_ν frequency maps $\mathbf{m}$, we obtain the best-fit spectral index parameters $\boldsymbol{\beta}_{\mathrm{BF}}$, the associated spectra, and the projector Q that removes the best-fit dust and synchrotron components from the data, using the map-level spectral likelihood of Section 5.2.1.2. We also construct the $(N_\nu - 2) \times N_\nu$ matrix
$$\mathsf{R} \equiv \tilde{\mathsf{D}}\, \mathsf{B}^{-1}, \tag{5.2.3.1}$$
where B is the matrix of eigenvectors of Q, and $\tilde{\mathsf{D}}$ selects only the rows of B^{-1} corresponding to the non-zero eigenvalues of Q.

2. We estimate the full set of power spectra between all pairs of frequencies, $\hat{C}_\ell^{\nu\nu'}$, from the same maps, and their covariance matrix.

3. We use R to deproject the constant dust and synchrotron modes from $\hat{C}_\ell^{\nu\nu'}$ and its covariance, according to Eqs. 5.2.2.10 and 5.2.2.11.

4. We calculate the theoretical prediction for these power spectra according to Eq. 5.2.2.2, using simple power-law models to describe the scale dependence of the foreground residuals, and the derivative of the foreground SEDs to describe their frequency dependence (see Section 5.2.1.3). We apply the R matrix to this prediction using again Eq. 5.2.2.10.

5. We use a Gaussian likelihood to obtain parameter constraints by comparing the data against this theoretical prediction. When inferring parameter constraints, by default we vary the CMB parameters $\{A_{\rm lens}, r\}$, the foreground power spectrum parameters $\{A_c, \alpha_c\}$ ($c \in \{\mathrm{D}, \mathrm{S}\}$), and the spectral indices $\{\beta_{\rm D}, \beta_{\rm S}\}$, imposing a prior on the latter based on the map-level constraints.

The main differences with respect to the "baseline" power-spectrum-based approach are that the latter skips steps 1 and 3 above, applies the power-law modelling for the full foreground power spectra (instead of the residuals), assuming constant spectral indices, and therefore does not make use of the R matrix to project onto the space of residuals. The assumption of constant spectral indices can be relaxed at the power spectrum level by implementing a moment expansion method, as done in Ch. 3, or by allowing for frequency decorrelation [143]. The main shortcomings of these purely C_ℓ-based methods, which the hybrid approach is able to address, are the need to assume a model for the scale dependence of the foreground amplitudes, and to either ignore spectral index variations, or to assume a model for their statistics (e.g. Gaussian, uncorrelated with amplitude variations, or with a flat scale dependence). In turn, the hybrid approach only assumes a model for the scale dependence for the much smaller foreground residuals, limiting the number of parameters needed to model them.

It must be noted that, as described above, the method proposed is not rigorous from a statistical point of view, as it does not propagate the uncertainties of all parameters self-consistently. Instead of splitting the parameter inference into two stages (map-based and C_ℓ-based), a more principled approach, similar to [161], would simultaneously sample the linear foreground and CMB amplitudes, the constant spectral indices, the power spectrum of the foreground residuals, and the CMB power spectrum parameters, in a joint map-level likelihood, effectively treating the foreground residuals as a two-dimensional Gaussian process. While fully self-consistent, this kind of approach is significantly more computationally demanding, although not intractable. Instead, we first use the maps to obtain tight constraints on the average spectral indices, and then reuse the data at the power spectrum level to marginalise over the foreground residuals. As we will show, however, in practice we find that the method proposed here is able to produce reliable constraints when applied to simulated data. The main reason for this is the fact that the information contained in the maps, for the instrumental sensitivities explored here, allows for a very precise measurement of the mean spectral indices in the first step. This first step is therefore akin to obtaining an external, high-precision measurement of $\boldsymbol{\beta}$, which can be used to deproject the constant foreground amplitudes from the data, rendering the procedure fully self-consistent.

5.3 Simulations

In order to test the validity of the hybrid component separation method described in the previous section, we test it on a suite of sky simulations. These simulations include the most relevant sky components, as discussed in Section 5.2. In all simulations, the CMB contribution was generated as a Gaussian random field drawn from the power spectrum in Eq. 5.2.2.5. We use fiducial values ($A_{\rm lens} = 1$, $r = 0$), unless otherwise stated. The simulations also incorporate the contribution from instrumental noise and the effects of a limited sky coverage as described in Section 5.3.1. We generate simulations of the polarised Stokes Q and U sky maps in six frequency bands as described in Table 3.2. These maps are generated using the `HEALPix` pixelisation scheme [110] with resolution parameter $N_{\rm side} = 256$.

5.3.1 Instrumental effects

All the simulations produced in this work include a model of the instrumental noise which replicates the characteristics of the Simons Observatory (SO), as described in [108] (SO19 hereon). The most relevant effects are the scale dependence and inhomogeneity of the noise, and the limited sky coverage, which are briefly described below (refer to Sec. 3.3.4 for more details).

- **Scale-dependent noise**. The noise is modelled using a power spectrum N_ℓ described as in Eq. 3.3.4, for which more details are provided in Table 3.2.
- **Sky footprint and inhomogeneous noise**. Except in Section 5.4.3, we will assume a sky footprint that follows the hits count map in Fig. 3.3. We use this mask to compute all power spectra (except in Section 5.4.3) using a pseudo-C_ℓ algorithm with B-mode purification as implemented in `NaMaster`[5] [79].

[5] github.com/LSSTDESC/NaMaster.

Note that we have not included the effects of an instrumental beam in these simulations in order to simplify the analysis. The impact of beams is straightforward to incorporate in the power spectrum analysis, and thus does not impact this method. At the map-level stage, maps should be pre-smoothed to a common beam before obtaining the best-fit spectral parameters. Again, this should not invalidate any of the results presented here.

As mentioned in Section 5.2.1, for each sky simulation, we generate 4 independent noise maps, each with a noise amplitude twice as large as the target sensitivity. This allows us to use these 4 realisations as independent data splits. Therefore, each simulation consists of 24 pairs of (Q, U) maps, corresponding to the 6 frequency channels and 4 data splits. In the map-domain cleaning described in Section 5.2.1, the data vector which enters the spectral likelihood in Eq. 5.2.1.3 is given by

$$\hat{\mathbf{m}} = \frac{1}{N_{\rm splits}} \sum_{i}^{N_{\rm splits}} (\mathbf{s} + \mathbf{n}_i), \tag{5.3.1.1}$$

i.e. the map of the coadded splits. When computing the multi-frequency power spectrum matrix that forms the basis of the C_ℓ-cleaning stage, we start by computing all BB auto- and cross-correlations between different pairs of maps, for a total of 300 different BB power spectra between all $N_{\rm splits} = 4$ splits and $N_\nu = 6$ frequency bands. In order to avoid modelling the noise bias as part of the likelihood described in Sec. 5.2.2), we then generate a coadded power spectrum matrix by averaging over all power spectra combining different data splits for each frequency pair. This results in a collection of 21 distinct coadded power spectra, which form the independent elements of the 6×6 multi-frequency power spectrum matrix.

5.3.2 Foreground simulations

The simulations described below were created using the publicly available code `BBSims`[6]. We generate two classes of foreground simulations: gaussian and "more realistic". The latter comprises three models, which we label as `d1sm`, HD and VS, that introduce different aspects of foreground complexity and include the foreground templates provided by `PySM`. These models include "realistic" non-Gaussian amplitude and spectral index maps that are based on existing observations, and described below.

[6]See github.com/susannaaz/BBSims.

- **Gaussian foreground simulations.** We generate a large suite of "Gaussian" foreground realisations, following the same procedure described in Sec 3.3.2. For these, we simulate sky maps for the amplitude ($\mathbf{T}_c(\hat{n})$) of each component as Gaussian random fields governed by power spectra following the power-law model in Eq. 5.2.2.3. The aim of these Gaussian simulations is twofold. First, we employ them to compute the power-spectrum covariance as described in Section 5.2.2.2. To estimate covariance matrices, we generate a suite of 100 Gaussian simulations with constant spectral indices. Secondly, we use these Gaussian simulations to study the performance of the method proposed here with different levels of spectral index variation. We parametrise this variation in terms of the standard deviation of $\sigma(\beta_c)$ (see Eq. 3.4.1.2) as in Ch 3, with increasing values $\sigma(\beta_c) = \{0, 0.1, 0.2, 0.3\}$ for both synchrotron and dust. We fix the spectral tilt γ_c to $(\gamma_{\rm D}, \gamma_{\rm S}) = (-3.5, -2.5)$. We then add the mean spectral indices $\bar{\beta}_c$ to $\delta\beta_c(\hat{n})$ and use `PySM`[7], [188]) to generate observed sky maps in the six frequency channels of Table 3.2. All the Gaussian simulations presented here use the same values for the constant spectral indices ($\bar{\beta}_{\rm D} = 1.6$, $\bar{\beta}_{\rm S} = -3$) and the same pivot frequencies ($\nu_0^{\rm D} = 353\,{\rm GHz}$, $\nu_0^{\rm S} = 23\,{\rm GHz}$). The dust-synchrotron correlation coefficient was set to $\epsilon_{\rm SD} = 0$. We use the foreground parameters that best fit the dust and synchrotron template maps used in the realistic set of simulations, described in Section 5.3.2, within the footprint used here, and listed in Eq. 3.3.2.1.

- **`d1sm` model.** A set of 20 realisations is produced, where for the dust amplitude and spectral index map we adopt the templates included in the `d1` model; similarly, the synchrotron amplitude follows the `s1` model, while its spectral index map is created by re-scaling the `s1` map and extending it to smaller scales using a power-law spectral index power spectrum $C_\ell^{\beta_{\rm S}}$ consistent with the measurements of [148]. For more details on the `d1sm` model, please refer to Sec 3.3.3. The level of realism included in these simulations is compatible with the one used to assess the performance of forthcoming B-mode experiments in e.g. [108, 194, 195]. Assuming free amplitudes at the first step of our pipeline, as described in Sec. 5.2.1, should allow us to account for the higher complexity of these models. The foreground amplitude maps used in these simulations reproduce the properties of the power spectrum used to generate the Gaussian realisations described in the previous section within the footprint use here, and within the scales ($\ell \in [30, 300]$) we use in our main analysis. It should be

[7] github.com/bthorne93/PySM_public.

pointed out that, on scales larger than those analysed here, the foreground power spectra present obvious departures from a perfect power law [193], but this does not affect our analysis. Therefore, our model of the scale dependence of foreground amplitudes described in Section 5.2.2 will not induce any bias on r for these simulations; the bias will arise because of the spatially-varying spectral indices.

- **Hensley & Draine (HD).** additional sets of 20 simulations The thermal dust contribution is modelled using the specifications provided by [179], which we label "HD" hereafter. This is based on a micro-physical model of dust grains, taking into account the strength of the local radiation field, the grain compositions (carbonaceous, and silicate with varying degrees of iron abundance) and the temperature distribution depending on grain size. In these simulations, the synchrotron SED still corresponds to the `sm` model described above.

- **Vansyngel (VS).** An additional suite of 20 simulations was run using the statistical model described in [143] (labelled "VS"). This model takes into account the spin orientation of the dust grains, in order to simulate the diffuse polarised emission from galactic dust. The three-dimensional structure of the GMF is described by a finite number of layers (we use $N_{\rm layer} = 7$ layers). The coherent component of the magnetic field is the same in all layers, while its turbulent part is generated as a Gaussian random field. Once the direction of the GMF is determined, maps of the Q and U Stokes parameters are generated by scaling the dust intensity map found by *Planck* [193]. As an additional level of complexity, and in an attempt to describe the three-dimensional distribution of the dust spectral index, we associate each layer with a different Gaussian realisation of $\delta\beta_{\rm D}$ with standard deviation $\sigma_{\beta_{\rm D}} = 0.13$ (the combined RMS variation for 7 layers is $\sigma_{\beta_{\rm D}} \simeq 0.35$). In these simulations, the synchrotron SED still corresponds to the `sm` model.

5.4 Results

We now report on the performance of the hybrid component separation methodology proposed here, as a function of different aspects of foreground complexity. We will present the best-fit and standard deviation on r, obtained with the *hybrid* method, and compare it with the results found using the *baseline* C_ℓ-based method, and with its *moments*-based extension, described in Ch 3.

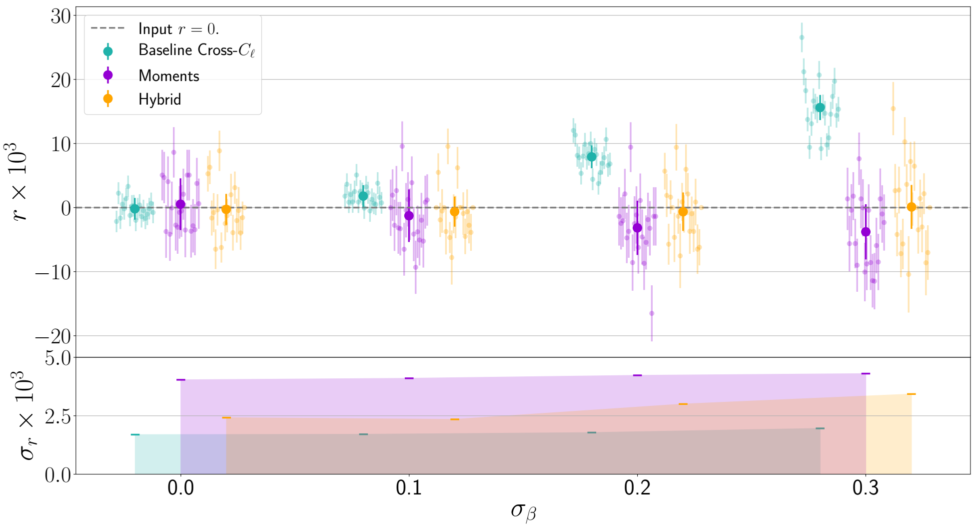

Figure 5.1: *Upper panel:* Best-fit values of r for 20 realisations of the sky calculated at different values of σ_{β_c} (the same for synchrotron and dust). Results are shown for the baseline cross-C_ℓ pipeline assuming constant spectral indices (*turquoise*), adding moments (*purple*) and using the hybrid method (*orange*). The position of each simulation in the x axis is shifted slightly from its true σ_{β_c} for clarity. The larger, solid dots, at the centre of each σ_{β_c} value show the mean and standard deviation of each suite of simulations. *Lower panel:* Statistical uncertainty $\sigma(r)$ averaged over the twenty realisations in the case of constant spectral indices (*turquoise*), adding moments (*purple*) and using the hybrid method (*orange*). Both the hybrid and the moments method are able to correct the bias on r for all values of σ_{β_c} considered, at the cost of increased final uncertainties with respect to a model with constant spectral indices (which themselves increase monotonically with σ_{β_c}) .

5.4.1 Gaussian foreground simulations

We use 20 Gaussian simulations with varying values of σ_β, representative of the range allowed by current data [193], to explore the impact of increased spatial variability of the spectral indices.

The results are shown in Fig. 5.1 as a function of σ_β (equal for synchrotron and dust). The upper panel shows the best-fit and standard deviation of β for each simulation, for the baseline, moments, and hybrid methods in turquoise, purple, and orange, respectively. The darker, larger points, show the overall mean and standard deviation of all simulations for each value of σ_β. The lower panel then shows the standard deviation obtained by each method. The mean measured value of r effectively correspond to the bias we get at different levels of spectral variation, since all the simulations were run with an input $r = 0$.

We find that, as σ_β increases, the bias on r obtained by the baseline method grows up to $\delta r = 0.017$ for the highest value of $\sigma_\beta = 0.3$ explored. Since the

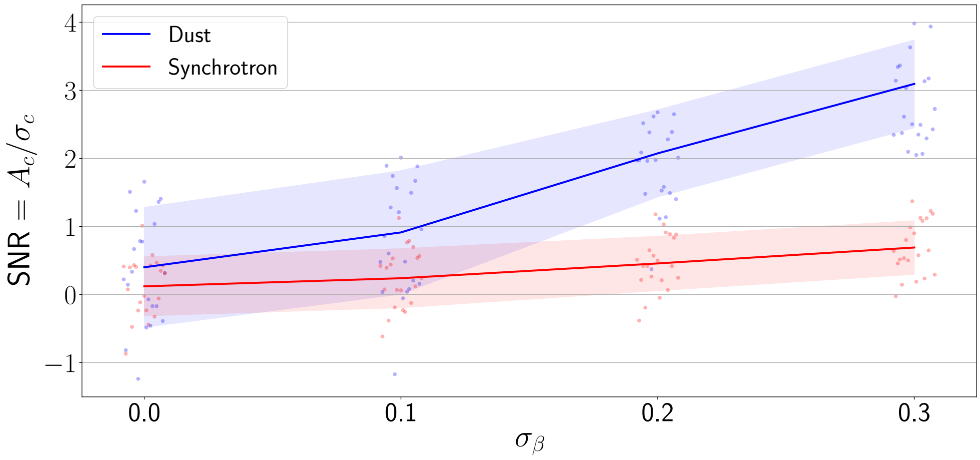

Figure 5.2: Signal-to-Noise Ratio (SNR) of SNR of the amplitudes of the foreground residual maps, dust A_d^{BB} (in blue) and synchrotron A_s^{BB} (in red), using the hybrid method for the Gaussian foreground simulations with increasing spectral index variation ($\sigma_\beta = 0 - 0.3$, the same for dust and synchrotron). For both components, the dots correspond to the SNR of each realisation, the line is the mean and the shaded area covers the 1σ standard deviation. We find that, at values of $\sigma_\beta > 0.1$, it is possible to significantly detect the effect of the spatial variation of the spectral parameters, particularly the dust spectral index, on the foregrounds multi-frequency power spectra for SO-like sensitivities.

baseline method achieves uncertainties $\sigma(r) \simeq 0.002$, this corresponds to a large, $\sim 8\sigma$ parameter bias. The hybrid method is then able to correct for this bias in all cases, at the cost of a slight increase in the final parameter uncertainty, with $\sigma(r)$ growing to $\sim$ 0.0024-0.0033 (a $\sim$ 20-50% increase). The moments expansion method is also able to correct this bias, as was shown in Chapter 3, although at a higher cost in terms of $\sigma(r)$, which rises by $\sim 50 - 70\%$.

At SO-like sensitivities, we are able to detect the impact of spatially-varying foreground SEDs, particularly in the case of dust. Figure 5.2 shows the Signal-to-Noise Ratio (SNR) of the foreground amplitudes using the hybrid method for the Gaussian foreground simulations with increasing spectral index variation ($\sigma_\beta = 0 - 0.3$, the same for dust and synchrotron). The SNR is calculated for each foreground component (c) as the ratio of the amplitude (A_d^{BB} for dust, A_s^{BB} for synchrotron) to the respective σ obtained from the posterior distribution. We find that, at values of $\sigma_\beta > 0.1$, it is possible to significantly detect the effect of the spatial variation of the dust spectral index for SO-like sensitivities, while a scatter at the level of $\sigma_\beta \gtrsim 0.3$ is required to detect variations in the synchrotron spectral index.

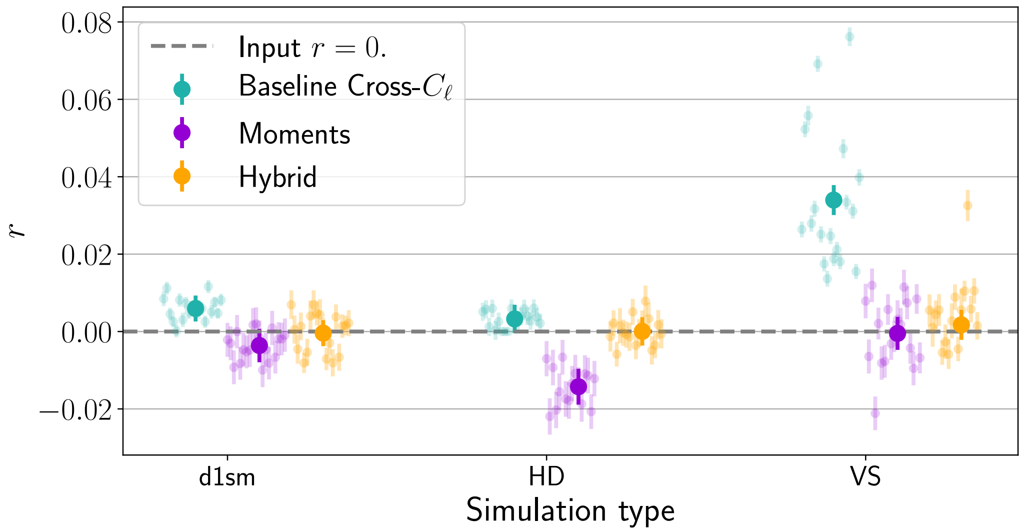

Figure 5.3: Best-fit values of r for 20 sky simulations for different types of realistic simulations, displayed on the x-axis (`d1sm`, HD, VS). The orange dots show the baseline cross-C_ℓ results assuming constant spectral indices, while the purple and orange dots show the results of the hybrid method and the moments method, respectively. The larger, solid dots at the centre of each simulation type show the mean of each suite of simulation. The hybrid method is able to correct the bias on r for all the different types of realistic simulations that appears when assuming constant spectral indices. The moments method recovers unbiased results for the `d1sm` and VS simulations, but introduces a negative 3σ bias in the HD case.

5.4.2 Realistic foreground simulations

Having used the Gaussian simulations to explore the performance of the method as a function of spatial spectral index variability, we now make use of the realistic foreground simulations described in Section 5.3.2 to validate the method in the presence of other aspects of foreground complexity. The `d1sm` simulation includes non-Gaussian and statistically inhomogeneous foreground amplitudes, and realistic spectral index maps. The HD simulations make use of a dust spectral law that is not a modified black-body in detail. Finally, the VS simulations account for the three-dimensional variability in the dust spectral index, and include an enhanced level of non-Gaussianity in the dust amplitude. The results are shown in Fig. 5.3 for the three different models (shown in the x axis).

For `d1sm`, the mean best-fit r from 20 realisations using the baseline method is centered at 0.0059 ± 0.0023, corresponding to a $\sim 2.5\sigma$ bias. With the hybrid method, this reduces to $r = -0.0004 \pm 0.003$. The bias is eliminated at the cost of increasing uncertainties by $\sim 25\%$, as expected from the results found with the Gaussian simulations. Similarly, the moments method recovers unbiased results with a $\sim$ 40% increase in $\sigma(r)$.

The conclusions are similar when analysing the HD simulations. The different dust SED used in this model introduces a bias at the $1-2\sigma$ level when analysed with the baseline method, obtaining $r = 0.0033 \pm 0.002$. The hybrid method corrects this bias (we obtain a mean $r = 4.3 \cdot 10^{-5}$), with a widening of the final constraints by $\sim 40\%$. The analysis with the moments method introduces instead a negative 3σ bias. This must be due to the foreground bias departing from the theoretical model used here. The two main assumptions are that the variations in $\beta_{\rm D}$ are small enough that the residual spectrum can be described by the Taylor expansion of the modified black-body SED, and that the residuals follow a power-law scale dependence. The additional complexity in the way the non-Gaussian foreground amplitude maps are generated might contribute to either of these causes. It could also be that the simulated data in this case differs significantly from the Gaussian assumption used to construct the covariance matrix and the resulting likelihood is ill behaved.

The significantly more complex VS simulations introduce a much larger bias on r, at the level of 16σ, when using the baseline method. As in the `d1sm` case, both the hybrid and the moments method eliminate this bias, with an increase in uncertainties by $\sim 30\%$ and $\sim 40\%$ respectively. This is in agreement with the findings of Ch. 3. Note that one of the key advantages of the hybrid method is that it does not rely on any assumption regarding the Gaussianity of the spectral index variations, unlike the moments method, and therefore it provides more flexibility to handle the additional complexity of the VS model. The impact of this additional complexity, however, does not seem to hamper the moments expansion method significantly by comparison with the hybrid approach explored here.

5.4.3 Larger sky fractions

The ability of the hybrid method to remove the bulk of the foreground contamination at the map level, without assuming a model for their spatial correlations, or for the correlation between them and the foreground residuals resulting from spatially-varying spectra, could potentially allow us to obtain constraints on sky areas significantly larger than the footprint explored so far. This is relevant in the context of future space missions, such as LiteBIRD [199, 205, 228], or when B-mode constraints become cosmic-variance dominated due to imperfect delensing. We have therefore generated and analysed a new set of simulations covering larger sky areas.

In this case, we generate three new sky masks three masks by setting different thresholds on the polarised intensity ($P = \sqrt{Q^2 + U^2}$) of the Planck 353 GHz map [227], smoothed on a scale of 1°. The thresholds are defined to yield observable sky fractions $f_{\rm sky} = \{0.3, 0.6, 0.8\}$. The resulting sky masks, avoiding the Galactic

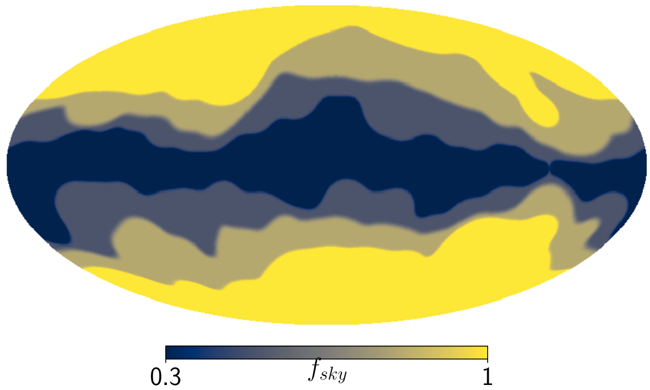

Figure 5.4: Masks applied to the simulations in Section 5.4.3, which exclude the Galactic plane from the estimation of the data. Three masks are created by setting different threshold on the polarisation intensity ($P = \sqrt{Q^2 + U^2}$) of the Planck 353 GHz map [227], to obtain sky cuts $f_{\rm sky} = 0.3$ (dark blue), 0.6 (lighter blue) and 0.8 (grey). The yellow area corresponds to the observed footprint in galactic coordinates.

plane, are displayed in Fig. 5.4. These are smoothed with a 10° FWHM beam and made differentiable with the application of a "C1" apodization with a scale of 5° to the resulting map [79]. The result is similar to the sky mask used in [139]. The larger sky fractions cover regions where foreground contamination is higher, and where spectral index variation is potentially more damaging. For example, in Figure 5.5 we find that the mean SNR of the amplitudes of the foreground residual maps recovered with the hybrid method in the `d1sm` model is $\mathrm{SNR} = 0.2, 2.1, 8.9$ (for dust) and $\mathrm{SNR} = 1.2, 2.7, 7.9$ (for synchrotron) in the $f_{\rm sky} = 0.3$, 0.6, and 0.8 footprints, respectively.

In order to avoid the computational cost of re-calculating the power spectrum covariance matrix from hundreds of simulations for each sky fraction, here we use the approximate formula in Eq. 5.2.2.7. Note that this approximation is not accurate enough for actual data analysis, which should involve extensive simulations, but it serves the needs of this forecasting exercise, as illustrated in [108]. The power spectra introduced in the formula contain both signal and noise contributions. Although the approximation accounts for the impact of incomplete sky observations in reducing the number of available independent modes, it does not account for the effects of mode-mode coupling, or the mixing of E and B modes. To get around this limitation, we have also modified the simulations used for this part of the analysis. We remove any E-to-B leakage from the simulations by hand, by transforming the original Q and U maps into E and B, setting the E-mode component to zero, and transforming back to Q and U. Although this degrades the level of realism used in the previous sections, it allows us to account for the main

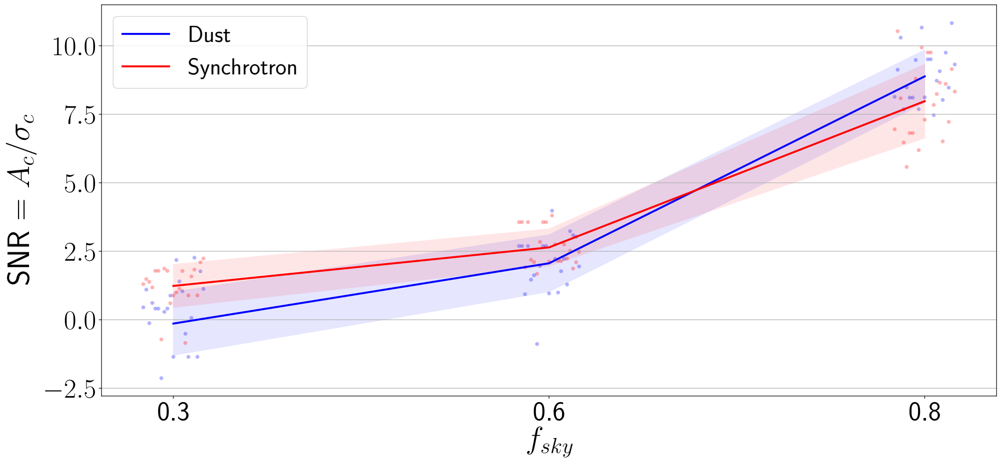

Figure 5.5: Signal-to-Noise Ratio (SNR) of the amplitudes of the foreground residual maps, dust A_d^{BB} (in blue) and synchrotron A_s^{BB} (in red), using the hybrid method for the `d1sm` foreground simulations with increasing observed sky fractions ($f_{\rm sky}0.3 - 0.8$). For both components, the dots correspond to the SNR of each realisation, the line is the mean and the shaded area covers the 1σ standard deviation. We find that at increasingly larger sky fractions, the foreground contamination is higher, and we can significantly detect the effect of the spatial variation of the spectral parameters on the foregrounds multi-frequency power spectra at $f_{\rm sky} > 30\%$.

impact of foreground complexity: the large dynamic range of foregrounds on large scale areas, and the departure from perfect power-law behaviour on large scales. For simplicity, in these simulations we assume the same noise power spectrum used before (Eq. 3.3.4.1), with a flat sky coverage (i.e. the noise properties are homogeneous throughout the observed footprint).

Fig. 5.6 shows the constraints on r found with the "baseline", "moments", and "hybrid" methods, as a function of sky fraction, for the `d1sm` foreground model. Over the 30% footprint, all methods obtain unbiased constraints on r. This, however, changes rapidly as we increase the sky area. At $f_{\rm sky} = 0.6$ and 0.8, the baseline method yields $r = 0.004 \pm 0.0008$ and $r = 0.013 \pm 0.0009$, corresponding to a $\sim 5\sigma$ and a $\sim 16\sigma$ bias, respectively. Remarkably, the hybrid method is able to absorb this bias in all cases ($r = 0.00001 \pm 0.0011$ and -0.0007 ± 0.0012 for 60% and 80% $f_{\rm sky}$ respectively), even when using 80% of the sky. The validity of this result, obviously depends on the level of realism of the foreground simulations used here, which may not be sufficiently accurate close to the Galactic plane. However, the result highlights the robustness of the method to high levels of foreground contamination, in comparison with standard approaches. Similarly, the moments method is able to recover unbiased results, recovering $r = 0.0 \pm 0.0013$

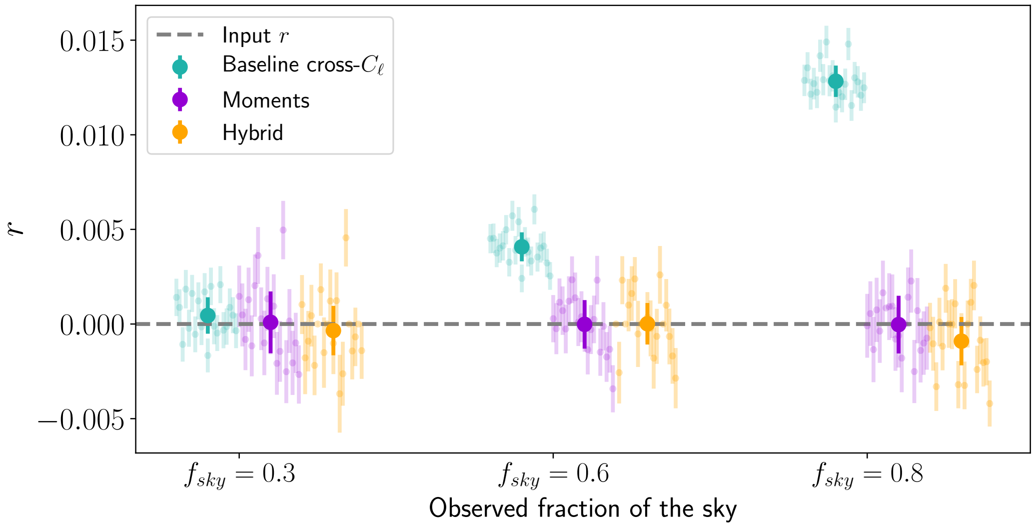

Figure 5.6: Constraints on r obtained by the "baseline", "moments", and "hybrid" methods (in turquoise, purple and orange), as a function of sky fraction, for the `d1sm` foreground model. With sky areas wider than 30% ($f_{\rm sky} > 0.3$), the baseline pipeline yields biased results. Both the hybrid method and the moments method are able to eliminate this bias, with an increase in $\sigma(r)$ by $30 - 40\%$ (hybrid) and by $50 - 80\%$ (moments) with respect to the baseline constraints.

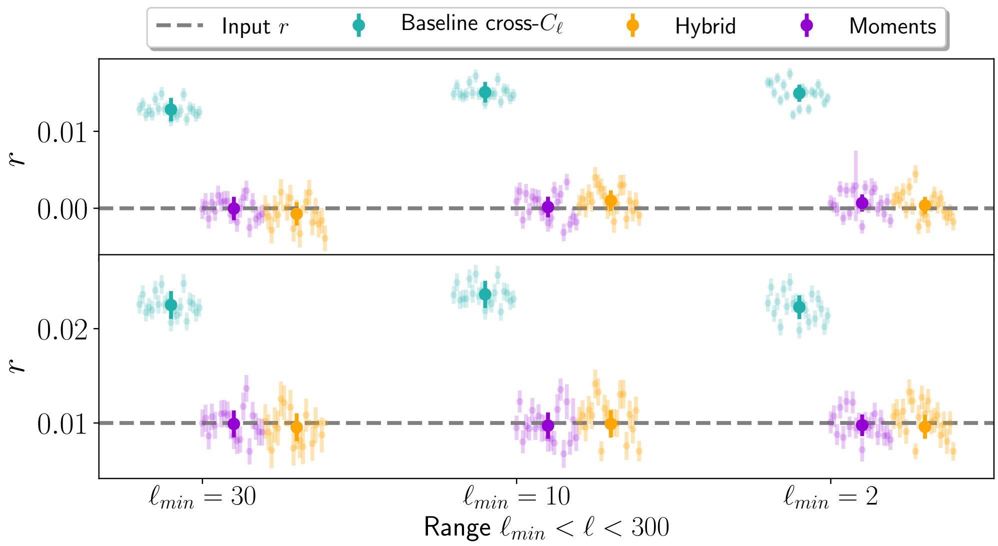

Figure 5.7: Best-fit values of r for 20 sky simulations observing 80% of the sky ($f_{\rm sky} = 0.8$) across different ranges of angular scales (from $\ell_{\rm min}$ up to $\ell = 300$) for input $r = 0$ (*upper panel*) and $r = 0.01$ (*lower panel*). The $x-$axis displays increasing values of angular scales (lower $\ell_{\rm min}$). The red dots show the baseline cross-C_ℓ results, while the green dots show the results of the hybrid method. The larger, solid dots at the centre of each suite of simulation show the mean across simulations. The hybrid method is able to correct the bias on r even at higher angular scales down to $\ell_{\rm min} = 2$.

and -0.0001 ± 0.0015 for 60% and 80% $f_{\rm sky}$, but with a $20-30\%$ increase in $\sigma(r)$ compared to the hybrid method constraints.

The availability of a larger sky area allows us to constrain primordial B-modes using larger scales. Fig. 5.7 shows the constraints obtained when using the $f_{\rm sky}=0.8$ footprint as a function of the minimum multipole, $\ell_{\rm min}$, included in the analysis. The two panels in this figure show the results for simulations with an input $r=0$ (top), and for $r=0.01$ (bottom), which thus allows us to study the performance of the method in the presence of a detectable B-mode signal. For our default $\ell_{\rm min}=30$, we obtain unbiased constraints on r with the hybrid approach in both cases ($r=-0.0007\pm0.0012$ and $r=0.009\pm0.0013$), whereas the baseline approach recovers a biased r at the same level as that displayed in Fig. 5.6. When including larger angular scales ($\ell_{\rm min}<30$), we use the HL non-Gaussian likelihood [121], as described in Sec. 5.2.2.2. When extending the scale range to $\ell_{\rm min}=10$, the hybrid method obtains $r=0.0005\pm0.0013$ (input $r=0$) and $r=0.0099\pm0.0013$ (input $r=0.01$). When including even larger scales, down to $\ell_{\rm min}=2$, the hybrid method recovers again unbiased results, i.e. $r=0.0002\pm0.0011$ (input $r=0$) and $r=0.0096\pm0.0012$ (input $r=0.01$), compared to the 19σ bias introduced by the baseline cross-C_ℓ pipeline. Similarly, the moments method recovers the input r value on all scales, with a $\sim10\%$ increase in $\sigma(r)$ compared to the hybrid method, and introducing a small negative $<0.2\sigma$ bias.

5.5 Conclusion

Here, we investigate a hybrid approach that combines characteristics from both map-based and C_ℓ-based studies. In the first step of the method, we determine foreground amplitudes at the map level without making any assumptions about their statistical correlations and assuming spatially uniform spectral characteristics. We next take this spatially-homogeneous foreground contribution away from the original data, and the residual foreground contamination (due to inaccurate subtraction or spatial SED fluctuations) is modelled at the power spectrum level. This allows us to retain the same C_ℓ-level assumptions, imposing a model for the scale dependence of the foreground amplitudes, while also limiting the degradation in the final constraints brought on by spectral index variations, by applying those assumptions to the smaller foreground residuals (instead of the full foreground spectra) and retaining information about their frequency dependence. Therefore this method allows us to overcome the limitations imposed by C_ℓ approaches (Gaussianity, power spectrum form), limiting the number of parameters needed to model the smaller foreground residuals.

We have applied this method to simulated data including realistic foreground models and instrumental noise. We have shown that we are able to obtain unbiased constraints on r in the presence of realistic (and often pessimistic) foreground complexity. We compare this with the results of a “baseline” fiducial cross-C_ℓ analysis (analogue to the method used to derive the current state-of-the-art constraints on the tensor-to-scalar-ratio r from B-modes [89]), and with the extended “moments” approach described in Chapter 3, which is able to further mitigate contamination from foreground residuals. Generally, we find that the hybrid method outperforms the baseline one in terms of bias, and the moments method in terms of $\sigma(r)$ (and bias for some of the most complicated foreground models).

In the simpler foreground scenario, which includes no spatial variation of the spectral parameters (corresponding to the Gaussian simulations with $\sigma_\beta = 0$), we obtain unbiased tensor-to-scalar ratio r with all the three pipelines, with a respective increase in $\sigma(r)$ by $\sim 30\%$ (moments) and $< 10\%$ (hybrid) compared to the baseline approach. In Section 5.4.1, we find that increasing spectral index variations ($\sigma_\beta = 0.1 - 0.3$) can introduce a bias on r of up to 8σ with the baseline method, which both the hybrid and the moments methods are able to correct for at the cost of an increase in statistical uncertainty (the former by $\sim 20 - 50\%$, the latter by $\sim 50 - 70\%$).

We have also explored more realistic foreground simulations based on various models proposed in the literature (Section 5.4.2) with different levels of complexity. Overall, we find that the hybrid method is able to correct the bias on r for all the different types of realistic simulations that appears when assuming constant spectral indices, at the cost of a limited $\sigma(r)$ increase by $< 30\%$. The moments analysis recovers unbiased results for most of these simulations. However, the reliance of the moments method on simplifying assumptions (Gaussianity of the spectral index variations, uncorrelation among spectral indices and foreground amplitudes) limits its ability to recover unbiased results for one type of complex foreground model explored here.

Unlike the moments method, the hybrid approach has the advantage of not relying on any assumptions about the shape of the spectral index fluctuations, and proceeds by removing the bulk of the foreground contamination at the map level, without assuming a model for their spatial correlations, or for the correlation between them and the foreground residuals resulting from spatially-varying spectra. This allows us to obtain constraints on increasingly larger sky areas and at larger angular scales in Section 5.4.3. Here we find that the results remain largely unbiased, even for experiments covering $\sim 80\%$ of the sky, suggesting the potential application of this method to future satellite missions such as LiteBIRD.

Chapter 6

Half-Wave-Plate rotational instabilities propagated through an end-to-end pipeline for LiteBIRD

Contents

6.1 Introduction

Here we study the propagation of a specific class of instrumental systematic effects that bias the final constraints on r. More specifically, we focus on the effects of the rotational variability of the Half Wave Plate (HWP), a polarisation modulator that will be deployed by future CMB experiments, such as the phase-A satellite mission LiteBIRD [199, 200] (see Sec. 2.2.2).

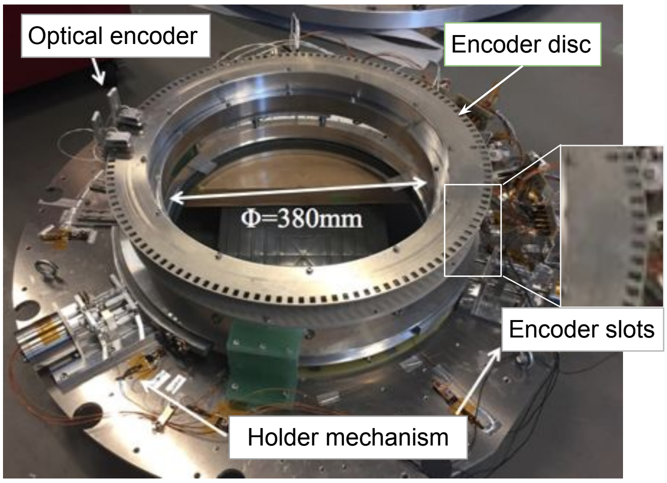

Figure 6.1: LiteBIRD Half-Wave Plate (HWP) and encoder mechanism.

The HWP angle is measured using an optical encoder, which is illustrated in Fig. 6.1. However, intrinsically imperfect instruments will lead to inconsistencies between the reconstructed HWP angle and the underlying true angle. Through the analysis of the HWP phase instabilities, the required instrument calibration can be determined, including the precision needed for the encoder measurement to prevent introducing bias into the cosmological analysis. Failure to calibrate the polarisation angle accurately can lead to significant systematic errors, resulting in leakage of E-modes into B-modes and biasing the estimation of the tensor-to-scalar ratio. This is described below in Sec. 6.2, where a general formalism for the treatment of instrument systematics is presented before describing the specific case of interest.

To propagate the systematic errors to cosmology, we have developed a stand-alone end-to-end simulation pipeline adapted to LiteBIRD, described in Sec. 6.2.4. Through the latter, we simulate CMB-only observations at single frequencies (140 GHz) to emphasise the systematic effects on the B-mode power spectrum C_ℓ^{BB}. We analyze the effects of a possible mismatch between the measured HWP properties (used in the time-ordered data stage) and the actual profiles (used in the map-making step of the pipeline).

Sec. 6.3 describes the ingredients of the synthetic simulations of the systematic effects introduced by the HWP rotational instabilities, which are propagated to the analysis pipeline. The systematic effects are first quantified in terms of a bias on C_ℓ^{BB} with respect to the ideal "no-systematics" case, and then propagated to the tensor-to-scalar ratio r. At the present stage, we impose the requirement on the specification of individual systematics to result in 1% of the expected LiteBIRD sensitivity (i.e. $\delta r \sim 10^{-5}$). The results are reported in Sec. 6.4.

6.2 Formalism

However, introducing a new element in the optics chain of the instrument unavoidably introduces new instrumental systematic effects. To benefit from HWP advantages without loosing too much because of new systematics, we therefore need to carefully model HWP-induced effects.

As introduced in Sec. 2.4, the HWP is used to modulate the polarised sky signal, and introduces several advantages:

1. the HWP modulates the polarised sky signal at higher frequencies than the $1/f$ noise component of the atmosphere, reducing noise correlations in the time-domain data. This high-frequency modulation also mitigates low-frequency sources of contamination, such as long-time noise drifts;

2. The HWP can also minimise beam, bandpass, and gain systematics by allowing separate polarised sky maps for each detector separately. As a result, the impact of mismatched detector properties in a detector pair is reduced;

3. the higher frequency sampling of the sky polarisation vector results in more independent measurements of each pixel, known as cross-linking, which further reduces systematics.

However, the introduction of a new element in the optical chain of the instrument inevitably creates new instrumental systematic effects. Therefore, careful modeling of the HWP-induced effects is necessary to benefit from its advantages while minimising new sources of systematics. To do so, we adopt the Mueller description of the instrument, introduced in Sec. 2.4.2.

6.2.1 The HWP Mueller matrix

As an example, we use Eq.2.4.2.6 to define the relation between the Stokes parameters of the light incident on the detectors to those of the light incoming from the sky as

$$\begin{pmatrix} I(t,\nu,\hat{n}) \\ Q(t,\nu,\hat{n}) \\ U(t,\nu,\hat{n}) \end{pmatrix} = \mathbf{M}_{\mathrm{OPT}}(\theta)\mathbf{M}_{\mathrm{H}}(\theta,\alpha,\nu) \begin{pmatrix} I_{\mathrm{in}}(t,\nu,\hat{n}) \\ Q_{\mathrm{in}}(t,\nu,\hat{n}) \\ U_{\mathrm{in}}(t,\nu,\hat{n}) \end{pmatrix} \tag{6.2.1.1}$$

where I, Q, U are the Stokes parameters after the HWP, $M_{\mathrm{OPT}}(\theta)$ is the Mueller matrix of the light transfer in the optical system and $M_{\mathrm{H}}(\theta,\alpha,\nu)$ is the Mueller matrix of the HWP, where θ is the incident angle with respect to the aperture, $\alpha = \omega_{HWP}(t)t$ is the HWP azimuthal rotation angle and $\omega_{HWP}(t)$ is the HWP revolution rate. For this discussion, we consider any other optical elements other than the HWP to be described by the identity matrix. Furthermore, here we assume no input V term, since the HWP is the first element in the LiteBIRD's optical system.

The Stokes parameters, contained in the vector $\vec{S}$, propagated through a rotating HWP become

$$\vec{S}_{\rm out} = \mathbf{R}(-\theta)\mathbf{M}_{\rm hwp}\mathbf{R}(\theta)\vec{S}_{\rm in}\,, \tag{6.2.1.2}$$

where the first and final operators on $S_{\rm in}$ are rotation matrices

$$\mathbf{R}(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos 2\theta & -\sin 2\theta & 0 \\ 0 & \sin 2\theta & \cos 2\theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{6.2.1.3}$$

that describe the rotation of the Stokes vector into and out of the HWP.

The HWP is used to modulate the polarisation, and in the ideal case the input polarisation is flipped with respect to the HWP's extraordinary axis. In terms of the Stokes parameters, this operation is equivalent to inverting the +U polarisation component to −U (and +V to −V). Therefore, the HWP Mueller matrix in the ideal scenario reduces to the identity matrix

$$\mathbf{M}^{\rm ideal}_{\rm hwp} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \tag{6.2.1.4}$$

6.2.2 HWP systematics

Consider a measured polarised signal $[Q \pm iU]'$, that is the true sky signal $[Q \pm iU]$ modified by a systematic effect. In general, one can define the vector $\delta\vec{S}$ that describes the uncertainties in the Stokes parameters in terms of the Mueller matrix for a given systematic error $\mathbf{M}_{\rm sys}$ [229],

$$\delta\vec{S} = \delta\mathbf{M}_{\rm sys} \cdot \vec{S}_{\rm in} = \begin{pmatrix} \delta_{II} & \delta_{IQ} & \delta_{IU} & \delta_{IV} \\ \delta_{QI} & \delta_{QQ} & \delta_{QU} & \delta_{QV} \\ \delta_{UI} & \delta_{UQ} & \delta_{UU} & \delta_{UV} \\ \delta_{VI} & \delta_{VQ} & \delta_{VU} & \delta_{VV} \end{pmatrix} \cdot \vec{S}_{\rm in} \tag{6.2.2.1}$$

where $\delta\vec{S} = \vec{S}^{\rm meas} - \vec{S}^{\rm sky}$, and $\vec{S}^{\rm meas}$ is the vector of measured Stokes parameters and $\vec{S}^{\rm sky}$ is the vector of the Stokes parameters coming from the sky.

In a more realistic scenario, different systematics of the HWP will affect different terms in $\mathbf{M}_{\rm hwp}$ in Eq. 6.2.1.4, which will become non-zero. These systematic effects are related to several HWP properties:

- δ_{II}: Reduction of intensity
- δ_{QQ}, δ_{UU}: polarisation efficiency
- δ_{QU}, δ_{UQ}: polarisation rotation
- δ_{IQ}, δ_{IU}: Instrumental polarisation

In this work, we consider an ideal HWP where the only systematic effects is due to its rotational instabilities, i.e. where the HWP angle is shifted by $\delta\alpha$. In this case, the measured polarised signal $[Q \pm iU]'$ corresponds to the true sky signal $[Q \pm iU]$ modified by a polarisation angle rotation of $\delta\alpha$, such that:

$$[Q \pm iU]' = e^{-2i\delta\alpha}[Q \pm iU] \tag{6.2.2.2}$$

In this case, the error matrix then reduces to the rotation matrix:

$$\delta\mathbf{M}^{\mathrm{HWP}}_{\mathrm{sys}} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & \varepsilon & \rho & 0 \\ 0 & -\rho & \varepsilon & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & \cos 2\delta\alpha & \sin 2\delta\alpha & 0 \\ 0 & -\sin 2\delta\alpha & \cos 2\delta\alpha & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{6.2.2.3}$$

where only the polarisation efficiency and rotation terms (i.e. the central block of the HWP matrix) are affected. The complex term ε changes the amplitude of polarisation, and ρ mixes both Q and U [125, 229]. This becomes clear when propagating to power spectra.

6.2.3 Propagation of systematics to power spectra

Generally, one can propagate Eq. 6.2.2.1 to harmonic space. Here, we define a general expression that relates the observed power spectra to the Mueller matrix of the HWP systematics.

Starting from the spin-2 polarisation, defined in Eq. 2.1.2.3 we recover an expression for Q and U in terms of the decomposition coefficients:

$$Q = \sum_{\ell m} \left[(a_{2,\ell m\, 2}Y_{\ell m}(\hat{n})) + (a_{-2,\ell m\, -2}Y_{\ell m}(\hat{n})) \right] \cdot \frac{1}{2} \tag{6.2.3.1}$$

$$U = \sum_{\ell m} \left[(a_{2,\ell m\, 2}Y_{\ell m}(\hat{n})) - (a_{-2,\ell m\, -2}Y_{\ell m}(\hat{n})) \right] \cdot \frac{1}{2i}. \tag{6.2.3.2}$$

From Eq. 2.1.2.8 we have:

$$a_{2,\ell m} = -ia^{B}_{\ell m} - a^{E}_{\ell m}; \quad a_{-2,\ell m} = ia^{B}_{\ell m} - a^{E}_{\ell m} \tag{6.2.3.3}$$

Combining the equations above, and re-arranging the terms, we can recover the Stokes Q, U in terms of $a^E_{\ell m}, a^B_{\ell m}$

$$Q = \sum_{\ell m} \left[a^E_{\ell m} \left(-_2Y_{\ell m}(\hat{n}) +_{-2} Y_{\ell m}(\hat{n}) \right) + a^B_{\ell m} i \left(-_2Y_{\ell m}(\hat{n}) +_{-2} Y_{\ell m}(\hat{n}) \right) \right] \cdot \frac{1}{2} \tag{6.2.3.4}$$

$$U = \sum_{\ell m} \left[a^E_{\ell m} \left(-_2Y_{\ell m}(\hat{n}) +_{-2} Y_{\ell m}(\hat{n}) \right) + a^B_{\ell m} \left(-i_2Y_{\ell m}(\hat{n}) -_{-2} Y_{\ell m}(\hat{n}) \right) \right] \cdot \frac{1}{2}$$

This allows us to write a matrix of the spherical harmonics:

$$M_{a_{\ell m}} = \begin{pmatrix} M_{II} & M_{IQ} & M_{IU} \\ M_{QI} & M_{QQ} & M_{QU} \\ M_{UI} & M_{UQ} & M_{UU} \end{pmatrix} \tag{6.2.3.5}$$

$$= \begin{pmatrix} Y_{\ell m}(\hat{n}) & 0 & 0 \\ 0 & -\frac{1}{2}\left({}_2Y_{\ell m}(\hat{n}) +_{-2} Y_{\ell m}(\hat{n})\right) & -\frac{i}{2}\left({}_2Y_{\ell m}(\hat{n}) -_{-2} Y_{\ell m}(\hat{n})\right) \\ 0 & \frac{i}{2}\left({}_2Y_{\ell m}(\hat{n}) -_{-2} Y_{\ell m}(\hat{n})\right) & -\frac{1}{2}\left({}_2Y_{\ell m}(\hat{n}) +_{-2} Y_{\ell m}(\hat{n})\right) \end{pmatrix},$$

such that the full vector of $a_{\ell m}$'s in terms of the Stokes parameters is:

$$\begin{pmatrix} a^{\mathsf{T}}_{\ell m} \\ a^E_{cm} \\ a^B_{\ell m} \end{pmatrix} = \int M_{a_{\ell m}} \begin{pmatrix} I \\ Q \\ U \end{pmatrix} d\hat{n}\,, \tag{6.2.3.6}$$

which combined with Eq. 6.2.3.5 gives

$$\begin{aligned} a^{\mathsf{T}}_{\ell m} &\sim M_{II} T \\ a^E_{\ell m} &\sim M_{QQ} Q + M_{QU} U \\ a^B_{\ell m} &\sim M_{UQ} Q + M_{UU} U = -M_{QU} Q + M_{QQ} U \end{aligned} \tag{6.2.3.7}$$

This is useful to define a general expression that relates the matrix of Fourier coefficients $M_{a_{\ell m}}$ of the observed CMB power spectra, defined in Eq. 6.2.3.5, to the Mueller matrix of the systematics $\delta M_{\rm sys}$:

$$\begin{pmatrix} \tilde{a}^{\mathsf{T}}_{\ell m} \\ \tilde{a}^E_{cm} \\ \tilde{a}^B_{\ell m} \end{pmatrix} = \int M_{a_{\ell m}} \delta M_{\rm sys} \begin{pmatrix} I \\ Q \\ U \end{pmatrix} d\hat{n} \tag{6.2.3.8}$$

$$= \begin{pmatrix} M_{II} & M_{IQ} & M_{IU} \\ M_{QI} & M_{QQ} & M_{QU} \\ M_{UI} & M_{UQ} & M_{UU} \end{pmatrix} \begin{pmatrix} \delta_{II} I + \delta_{IQ} Q + \delta_{IU} U \\ \delta_{QI} I + \delta_{QQ} Q + \delta_{QU} U \\ \delta_{UI} I + \delta_{UQ} Q + \delta_{UU} U \end{pmatrix} \tag{6.2.3.9}$$

The approach described above is general and can be used to describe a wide variety of systematic effects.

In the specific case of interest, defined by Eq. 6.2.2.3, we can easily recover the modified Fourier coefficients that are affected by the systematic in terms of the observed $a_{\ell m}$'s in Eq. 6.2.3.7:

$$\begin{aligned} \tilde{a}^E_{\ell m} &= \varepsilon a^E_{\ell m} + \rho a^B_{\ell m} \\ \tilde{a}^B_{\ell m} &= \varepsilon a^B_{\ell m} - \rho a^E_{\ell m}. \end{aligned} \tag{6.2.3.10}$$

We can then recover the error on the affected power spectra δC_ℓ^{XY}, using the relation in 2.1.3.6, where the relevant expectation values are given by

$$\left\langle \tilde{a}^{E*}_{\ell m}\tilde{a}^{E}_{\ell m}\right\rangle = \varepsilon^2 \left\langle a^{E*}_{\ell m}a^{E}_{\ell m}\right\rangle + \rho^2 \left\langle a^{B*}_{\ell m}a^{B}_{\ell m}\right\rangle \tag{6.2.3.11}$$

$$\left\langle \tilde{a}^{B*}_{\ell m}\tilde{a}^{B}_{\ell m}\right\rangle = \varepsilon^2 \left\langle a^{B*}_{\ell m}a^{B}_{\ell m}\right\rangle + \rho^2 \left\langle a^{E*}_{\ell m}a^{E}_{\ell m}\right\rangle \tag{6.2.3.12}$$

$$\left\langle \tilde{a}^{E^*}_{\ell m}\tilde{a}^{B}_{\ell m}\right\rangle = -\varepsilon\rho \left\langle a^{E*}_{\ell m}, a^{E}_{\ell m}\right\rangle + \varepsilon\rho \left\langle a^{B*}_{l_m}, a^{B}_{\ell m}\right\rangle \tag{6.2.3.13}$$

where we have made use of the fact that $\left\langle a^{E*}_{\ell m}a^{B}_{\ell m}\right\rangle = \left\langle a^{B*}_{\ell m}a^{E}_{\ell m}\right\rangle = 0$ due to parity violation.

The measured spectra are then given by the original power spectra with an additional error term induced by the HWP systematic, $C_\ell^{XY} + \delta C_\ell^{XY}$, such that

$$C_\ell^{\mathrm{EE,meas}} = C_\ell^{EE} + \varepsilon^2 C_\ell^{EE} + \rho^2 C_\ell^{BB} \tag{6.2.3.14}$$

$$C_\ell^{BB,\mathrm{meas}} = C_\ell^{BB} + \varepsilon^2 C_\ell^{BB} + \rho^2 C_\ell^{EE} \tag{6.2.3.15}$$

$$C_\ell^{EB,\mathrm{meas}} = -\varepsilon\rho C_\ell^{EE} + \varepsilon\rho C_\ell^{BB} \tag{6.2.3.16}$$

where $C_\ell^{EE/BB,\text{ meas}}$ are the E- and B-mode power spectra including systematic effects, $C_\ell^{EE/BB}$ the underlying power spectra. Thus, we find that the measured spectra are given by the original ones with an additional error term induced by the HWP systematic. The equations 6.2.3.14-6.2.3.16 above clearly show that the efficiency term affects the amplitude of the E/B-mode power spectra, and the rotation term yields a E-B mixing.

It is worth noting at this point that, using the equations above, it is straightforward to forecast the propagation of the errors to the tensor-to-scalar ratio r, in the case of a *global rotation angle error*, i.e. $\delta\alpha(t) \equiv \delta\alpha$ is constant. Consider the relation between the systematics parameters and the uncertainty in the polarisation angle $\delta\alpha$, implicitly defined in Eq. 6.2.2.3, whereby:

$$\rho = \sin(2\delta\alpha), \quad \varepsilon = \cos(2\delta\alpha), \tag{6.2.3.17}$$

and the equations for the measured power spectra above. Noting that the inflationary parameter r measures the ratio of B-mode power to E-mode power, we can roughly estimate the bias on r:

$$\delta r \sim \sin^2(2\delta\alpha). \tag{6.2.3.18}$$

Therefore, any miscalibration of HWP phase given by a global $\delta\alpha = \text{const}$ needs to be $\lesssim 1$ degree for $\delta r \lesssim 10^{-3}$ (or $\lesssim 0.01$ for 1% of the LiteBIRD systematic budget). In the case of more realistic HWP angle errors, it is more complicated to analytically predict how they would propagate to the final cosmological parameters. Thus we resort to simulation and propagate them through an end-to-end analysis pipeline, described below.

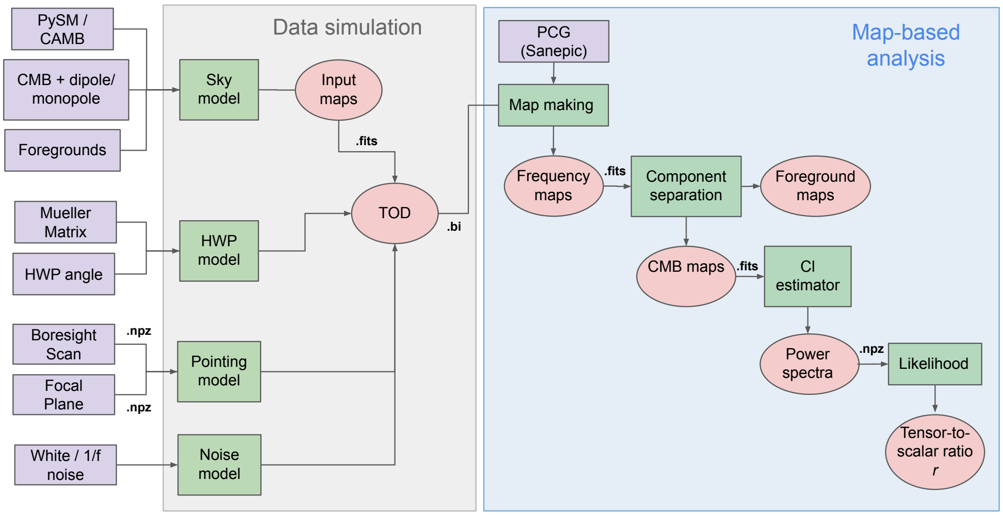

Figure 6.2: Schematics of the end-to-end analysis pipeline `STOMPer`, used in this study to propagate the HWP systematics at the TOD level to the cosmological parameter r.

6.2.4 Analysis pipeline

A realistic analysis of the systematics induced by the HWP angle errors requires a framework that goes from the time-ordered data (TOD), to cleaned maps and finally to the cosmological parameter r. To this end, the end-to-end analysis pipeline `STOMPer`[1] has been developed.

6.2.4.1 From TOD to sky maps

As described in Sec. 2.3.1, the time-ordered data (TOD) is represented as a sum of signal and noise

$$\mathbf{d} = \mathsf{P}\mathbf{m} + \mathbf{n} \tag{6.2.4.1}$$

where $\mathbf{d}$ is a vector, P is the pointing matrix that relates the sky pixel Stokes parameters represented as a vector $\mathbf{m}$ with the sampling time, and $\mathbf{n}$ is the noise vector. What we want to do is to estimate the vector $\mathbf{m}$ given the observed data $\mathbf{d}$ with the noise characteristics.

The noise is assumed to be Gaussian, and for the purpose of this work it will be sufficient to consider the idealised case of uncorrelated detectors at the focal plane, such that the noise matrix can be modeled as diagonal[2]. The noise realisation

[1] github.com/susannaaz/STOMPer

[2] Although the map-maker is suitable to work with non-diagonal noise covariance matrices of the form $N_{ij} = \langle n(t_i)\, n(t_j)\rangle$ as well.

$\mathbf{n}$ is then independently drawn for each detector from a Gaussian distribution with zero mean and variance σ^2_{det} .

The map is then reconstructed through map-making (see Sec. 2.3.2). To this end, a Preconditioned Conjugate Gradient (PCG) method, based upon [119], was implemented. This method aims to minimise the χ^2 of the data:

$$\chi^2 = \mathbf{n}^\intercal \mathsf{N}^{-1}\mathbf{n} = (\mathbf{d} - \mathsf{P}\mathbf{m})^\intercal \mathsf{N}^{-1}(\mathbf{d} - \mathsf{P}\mathbf{m}) \tag{6.2.4.2}$$

such that $\partial\chi^2/\partial m = 0$ to obtain the minimum variance map, with solution:

$$\mathbf{m} = \left(\mathsf{P}^\intercal \mathsf{N}^{-1}\mathsf{P}\right)^{-1} \mathsf{P}^\intercal \mathsf{N}^{-1}\mathbf{d} \tag{6.2.4.3}$$

The linear system is solved iteratively by means of the PCG to invert the covariance without ever explicitly building the matrix $(\mathsf{P}^\intercal \mathsf{N}^{-1}\mathsf{P})$. This is done by defining a “preconditioner”, i.e. a best estimate of the matrix to invert, by computing exactly its diagonal. To do so, we assume the following:

- *Gaussianity*: this ensures that only the second-order moments are needed to describe the noise statistics, which is then fully described by cross- and auto-power spectra of the noise.
- *Stationary noise*: this implies that the matrix is *circulant*, i.e. the matrix to invert is diagonal in Fourier space.
- *Continuous data stream*: there are no gaps in the data.

This allows us to apply the matrix $(\mathsf{P}^\intercal \mathsf{N}^{-1}\mathsf{P})^{-1}\, \mathsf{P}^\intercal \mathsf{N}^{-1}$ directly to $\mathbf{d}$ to produce the map, treating it effectively as an operator. The PCG requires several iterations to arrive at an acceptable estimate of the map. The number of required iterations strongly depends on the noise model and on which maximum scale we aim to recover in the map. In our case, we set $\partial\chi^2/\partial m = 10^{-16}$, and the map is reconstructed at $N_{\text{side}} = 256$, requiring $\mathcal{O}(10)$ iterations.

A schematics of the full analysis pipeline is displayed in Fig.6.2. This end-to-end simulation pipeline has been adapted to LiteBIRD but can be easily repurposed for other experiments. In this preliminary study, we simulate single-frequency, CMB-only observations to emphasise the systematic effects on the B-mode power spectrum C_ℓ^{BB}. We analyse the effects of a possible mismatch between the measured HWP properties (included in the TOD simulation step of the pipeline) and the actual profiles.

6.2.4.2 Propagation of errors to tensor-to-scalar ratio r

In LiteBIRD, we impose the requirement on each individual instrumental systematic effect contributing to the total systematic error, to result in less than 1% of the lensing B-mode power spectrum [8].

The LiteBIRD mission aims to measure the tensor-to-scalar ratio with a total systematic error of $\delta r \leq 0.57 \times 10^{-3}$, which is estimated using the likelihood function defined as:

$$-2 \ln \mathcal{P}\left(\hat{C}_\ell^{BB} \mid r\right) = (2\ell+1)\left[\frac{\hat{C}_\ell^{BB}}{C_\ell^{BB}} + \ln C_\ell^{BB} - \frac{2\ell-1}{2\ell+1}\ln \hat{C}_\ell^{BB}\right]. \tag{6.2.4.4}$$

The data vector is generated as $\hat{C}_\ell^{BB} = r_{\text{input}}\, C_\ell^{GW} + C_\ell^{\text{lens}} + N_\ell^{BB}$ and the fiducial model is given by $C_\ell^{BB} = rC_\ell^{GW} + C_\ell^{\text{lens}} + N_\ell^{BB}$, where C_ℓ^{GW} is the primordial gravitational wave spectrum with $\tau = 0.06$ and C_ℓ^{lens} is the lensing contribution. The white noise spectrum is defined as

$$N_\ell^{BB} = \left(\frac{\pi}{10800}\frac{w_p^{-\frac{1}{2}}}{\mu K \text{ arcmin}}\right)^2, \tag{6.2.4.5}$$

and the total likelihood function is given by $\ln \mathcal{L}(r) = \sum_{\ell=2}^{200} \ln \mathcal{P}\left(\hat{C}_\ell^{BB} \mid r\right)$, from which we recover δr to 1 sigma level by solving $\int_0^{\delta r} \mathcal{L}(r)dr / \int_0^{\infty} \mathcal{L}(r)dr = 0.68$. In the white noise, $w_p^{-1/2}$ is estimated to be $7.6\mu K$, given the sensitivity of the LiteBIRD detector array design, which achieves $\delta r = 0.57 \times 10^{-3}$. Based on the above, we set the accuracy requirement for measuring the HWP properties to be

$$\delta r < 5.7 \cdot 10^{-6}, \tag{6.2.4.6}$$

which corresponds to 1% of the expected LiteBIRD sensitivity on r.

6.3 Simulations

The scanning strategy implemented in the analysis pipeline follows that of LiteBIRD, as described in [8, 199] and illustrated in Figure 6.3. The simulated scan pattern and telescope response thus go into the pointing P. For the maps $\mathbf{m}$, we simulate the set of the Stokes parameters, starting from CMB maps at 140 GHz, with the aim of eventually including foregrounds. The maps are then used to produce 1 year TOD simulations for a single detector.

Here we describe the simulations of the HWP angle error $\delta\alpha$. We start from a simplistic (and unrealistic) assumption that the angle error is constant over time, which incidentally corresponds to the worst-case scenario in terms of bias to r, in Sec. 6.3.1, to more realistic models for the angle error, going from a random jitter in Sec. 6.3.2 to a drift over time in Sec. 6.3.3.

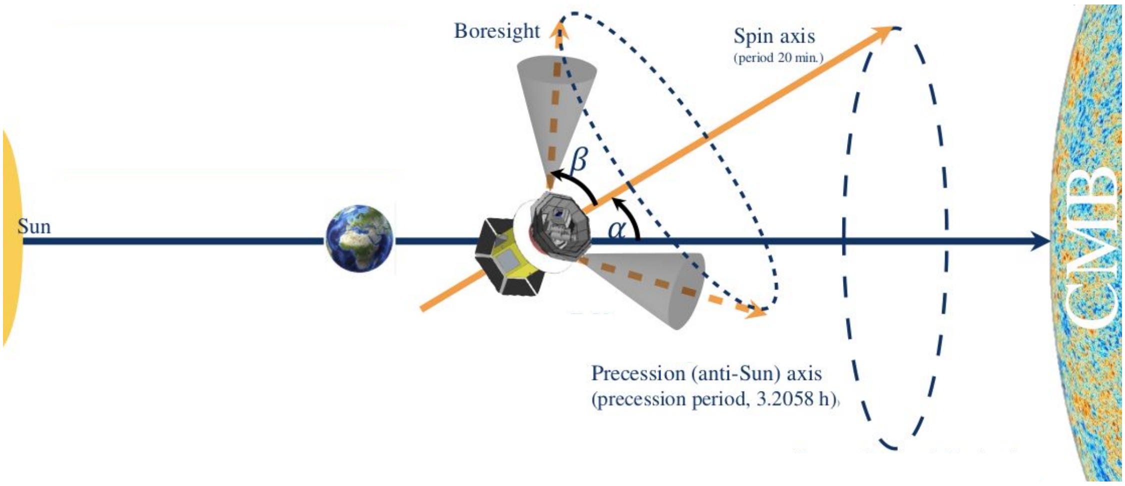

Figure 6.3: LiteBIRD scanning strategy. During the scan, the spacecraft will rotate at a rate of 0.1rpm around its spin axis. The spin axis will also rotate, but at a much slower pace, ranging from 90 minutes to almost a day in current studies. The spin axis will be angled at $\alpha = 45°$ to the Sun-L2 axis, while the angle β between the boresight and the spin axis is $50°$ [199].

6.3.1 Constant offset

We start by considering a constant offset over time of the measured HWP angle, i.e. a HWP phase lag that is fixed throughout the mission. We use these simulations to compare the results of the analysis pipeline with the formalism described in Sec. 6.2.3.

Propagating the Stokes parameters to angular power spectra, we can analytically recover the error on the observed E- and B-mode power spectra:

$$\delta C_\ell^{BB} = \cos^2(2\delta\alpha) C_\ell^{BB} + \sin^2(2\delta\alpha) C_\ell^{EE} \tag{6.3.1.1}$$

$$\delta C_\ell^{EE} = \cos^2(2\delta\alpha) C_\ell^{EE} + \sin^2(2\delta\alpha) C_\ell^{EE} \tag{6.3.1.2}$$

where $\delta C_\ell \equiv C_\ell^{\text{meas}} - C_\ell$ and we have used the expression for the measured spectra defined in Eqns. 6.2.3.14,6.2.3.15, including the definitions in Eq. 6.2.3.17 to relate the uncertainty in the polarisation angle $\delta\alpha(t)$ to the systematics parameters.

We include a global polarisation rotation systematic in our simulations in Sec. 6.4.1, and propagate it through the analysis pipeline. The aim of this exercise is to see if we can recover the power-spectrum model.

6.3.2 Random jitter

We consider a random jitter of the HWP angle, i.e. to simulate $\delta\alpha$ we draw random samples from a normal Gaussian distribution with probability density

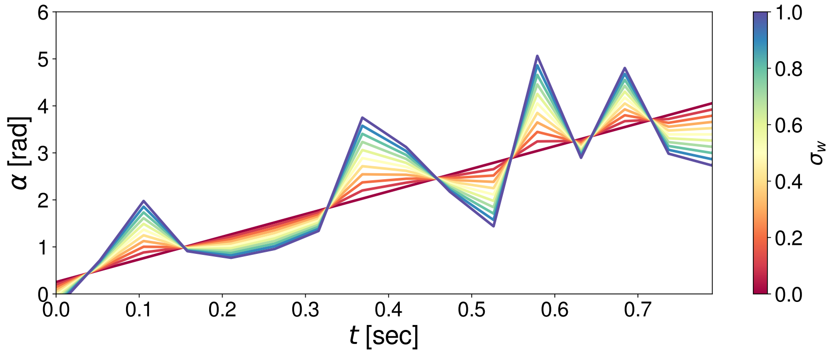

Figure 6.4: The HWP angle α position over a short period of time, including the angle error $\delta\alpha$ simulated as a random jitter and determined for increasing standard deviations σ_w (for a single realisation). The angle increases over time as it tracks the actual HWP position as it rotates, so the "slope" is given by the rotation rate.

$p(x) = e^{-\frac{(x-\mu)^2}{2\sigma_w^2}} / \sqrt{2\pi\sigma_w^2}$, where μ is the mean and σ_w the standard deviation. In Fourier space, this corresponds to a white noise spectrum of the form

$$P(f) = \sigma_w^2 \,. \tag{6.3.2.1}$$

In our simulations, we set $\mu = 0$ and we increase σ_w in order to get the requirements for 1% of the total systematic budget. This is illustrated in Fig. 6.4, where the HWP angle, for different values of the angle error σ_w, increases as a function of time. We parametrise the angle error $\delta\alpha$ with the root mean square (RMS) of the TOD of the angle variation per bolometer sample, corresponding to

$$\sigma_{\rm tod}^{\delta\alpha} \;=\; \sqrt{\int \; P(f) \; df}. \tag{6.3.2.2}$$

which is a function of the single model parameter $\sigma_{\rm tod}^{\delta\alpha}\,(\sigma_w)$ and has angle units.

6.3.3 Drift between measurements

A more realistic model of the error includes a drift in the HWP rotational variation, whereby the angle error is accumulated over time in between two measurements, until the encoder "re-calibrates" by measuring the HWP position. In this case, $\delta\alpha(t)$ is sampled from a Gaussian with zero mean and variance given by the power spectrum

$$P(f) = \sigma_{fk}^2 \left[1 + \left(\frac{f}{f_{\rm knee}}\right)^{\alpha}\right]^{-1}. \tag{6.3.3.1}$$

where the angle error is parametrised by the knee frequency f_{knee} of the spectrum and again the RMS of the TOD of the angle variation per bolometer sample, corresponding to $\sigma_{\rm tod}^{\delta\alpha}$ (see Eq. 6.3.2.2). We set $\alpha = 2$, such that the model of the angle over time corresponds to a random walk. The aim is then to constrain f_{knee}, the frequency corresponding to the inverse of the separation in time between two measurements made by the encoder.

To simulate this effect realistically, for different values of the model parameters we iterate over several steps:

1. **Generate a HWP phase time-stream with no corrections from the encoder**. This corresponds to a drifting signal over 1 year, with data points at sampling frequency $f_{\rm samp}$. We implement it as a random walk in real space, which corresponds to a $1/f^2$ power spectrum. This allows us to avoid Fourier domain transforms, hence reducing the overall computational cost.

2. **Determine the data points measured by the encoder, including some random error in time**. The encoder data is sampled at a higher frequency than that of the bolometer/data sampling frequency (~ 10 times higher). We simulate the HWP angle data acquired by the encoder using a spline fit.

3. **Calibrate**. We subtract the spline fit from the angle data to isolate the residual HWP angle error $\delta\alpha(t)$, which will need to be introduced in the full TOD generator code where the data is sampled at the bolometer sampling frequency $f_{\rm bol} = 19$Hz.

4. **Downsample the residual $\delta\alpha$ to $f_{\rm bol}$**. To do so, we bin the angle residual at every $1/f_{\rm bol}$ Hz, and we extrapolate the value at each bin using linear interpolation. These values make up the error to the HWP phase TOD $\delta\alpha$ that are added to the HWP angle in the full TOD generator code.

Given the LiteBIRD instrument specifications [8], we know the HWP is spinning at 46 rpm, and so its frequency is $f_{\rm HWP} = 0.766$ Hz. From this, we define the period of time that elapses in between two encoder measurements,

$$T_{\rm meas} = \frac{1}{f_{\rm HWP}} \cdot \frac{1}{N_{\rm holes}} \tag{6.3.3.2}$$

where $N_{\rm holes}$ is the number of slots in the encoder used to gather the HWP angle measurements. The frequency at which the angle starts drifting in Eq. 6.3.3.1 is given by

$$f_{\rm knee} = \frac{1}{T_{\rm meas}}, \tag{6.3.3.3}$$

and the frequency at which data are sampled in the laboratory is approximately 10 times higher, i.e. $f_{\rm samp} = 10 \times f_{\rm knee}$. We then use this model to test the effect of varying degrees of angle drift, by varying the number of encoder slots (and hence f_{knee}).

The aim of this exercise is three-fold. First, we can constrain the HWP drift with f_{knee}. Secondly, we obtain the HWP requirements for 1% of the total systematic budget, in terms of the model parameters, $f_{\mathrm{knee}}, \sigma^{\delta\alpha}_{\mathrm{tod}}$. Third, we can further obtain requirements for the number of holes that need to be included in the encoder in order to have accurate HWP angle measurements. Currently, $N_{\mathrm{holes}} = 128$ slots are included, and we want to test whether this number can be decreased without biasing the cosmology. We therefore include different values of N_{holes} in our analysis in Section 6.4.3, and summarised in Table 6.1.

6.4 Results

We now report on the results of the simulations run through the analysis pipeline described in Section 6.3. We will present the best-fit and standard deviation on r, as a function of different levels of HWP angle error, for increasingly realistic models of the error.

6.4.1 Constant offset

As described in Sec. 6.3.1, we use CMB-only simulations with different constant values of $\delta\alpha$, corresponding to a constant HWP angle offset over time, which is introduced at the TOD level. The simulations are then propagated through the analysis pipeline, where in the map-making stage we assume $\delta\alpha = 0$, i.e. an unbiased HWP. From the produced maps, we calculate the power spectra and compare these with the power-spectrum model that was derived in Eqns. 6.3.1.2,6.3.1.1. The result for one of the input systematic error, corresponding to $\delta\alpha = 0.18°$ is reported in Fig. 6.5, where the "measured" power spectra, corresponding to the ones calculated from the map-maker output maps, which include the systematic

Table 6.1: Tabulated values of the model parameters for the angle drift, assuming different values for the number of encoder slots.

ω_{HWP} [rpm]	N_{holes}	T_{meas} [s]	f_{knee} [Hz]	f_{samp} [Hz]
46	128	0.010	98.05	980.48
46	40	0.033	30.64	306.40
46	10	0.131	7.67	76.6

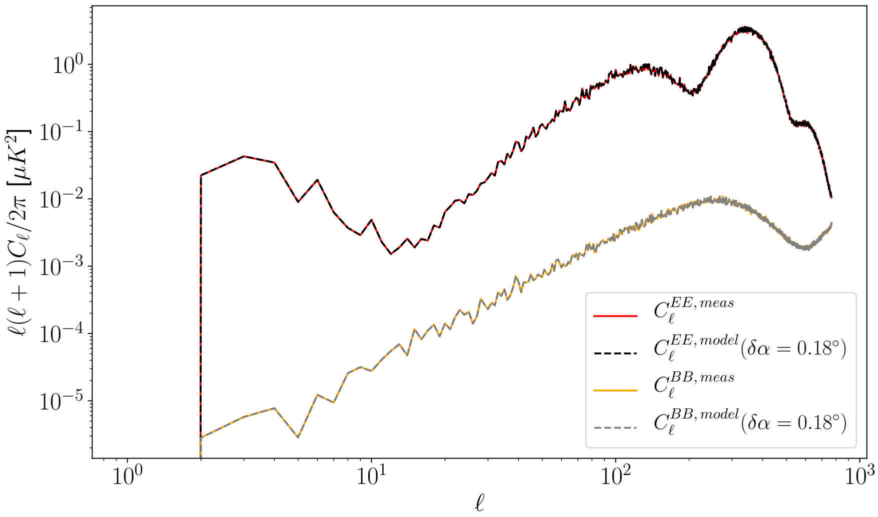

Figure 6.5: The measured power spectra of the polarisation maps that are produced from the biased TOD simulations, for a fixed HWP angle error $\delta\alpha$ is compared with the model. We find that, for the systematic error considered here, the power spectra of the biased maps are well described by the analytic model with $\delta\alpha = 0.18°$, which corresponds to the input systematic error included in the TOD simulation.

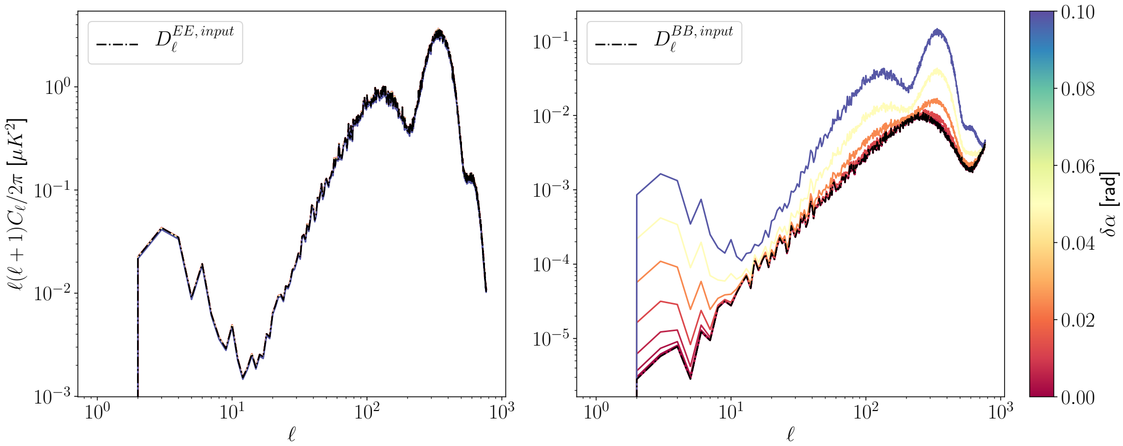

Figure 6.6: E- and B-mode power spectra (left and right, respectively) of the recovered output maps for increasing constant values of the HWP angle error $\delta\alpha$. With an increase in error values, the output spectra progressively deviate from the power spectrum of the input map (black dashed lines) exhibiting clear EB-mixing. This effect is particularly noticeable in the B-mode spectrum for higher $\delta\alpha$'s, due to the larger amplitude of the E-mode spectrum.

bias, overlap with the model of Eqns. 6.3.1.2,6.3.1.1. The dotted lines show the best-fit model with error $\delta\alpha = 0.18°$, which is consistent with the value that was inputted in the TOD simulation.

The results for different values of $\delta\alpha$ is illustrated in Fig. 6.6, which shows the E- and B-mode power spectra calculated from the map-maker's output maps as a

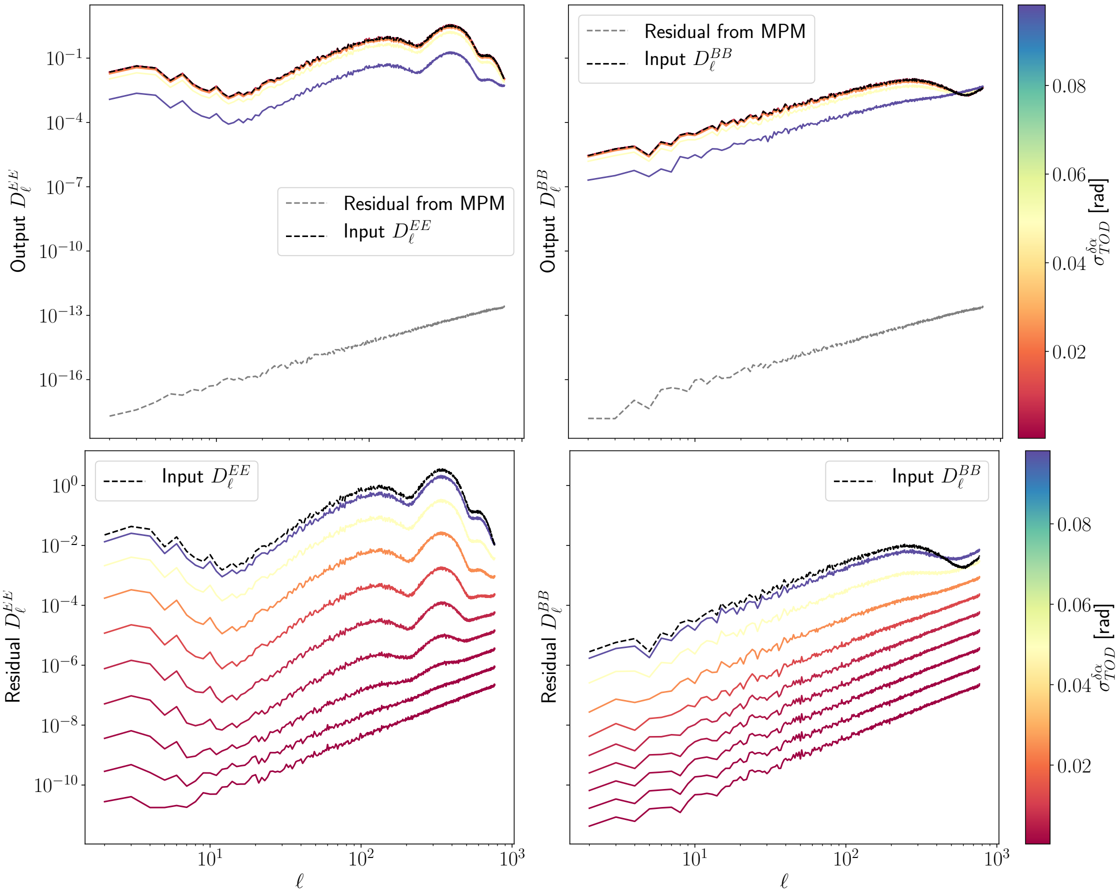

Figure 6.7: E and B-mode power spectra of the recovered output maps (top panel) and the residual maps (below) for different values of the random jitter angle error $\sigma_{TOD}^{\delta\alpha}$. We see that, for higher values of the angle error, the output spectra increasingly depart from the power spectrum of the input map (black dashed line), resulting in an increase in amplitude of the residual spectra and a net depolarisation effect.

function of the angle error. As expected, we find that higher values of the angle error correspond to more biased power spectra. There is a very dominant $E-B$ leakage, thus for higher $\delta\alpha$'s the BB output spectra (on the right) resemble the input EE spectrum and shifts away from the input BB, depicted by the black dashed line, indicating that the ρ term dominates the power spectrum model in Eq. 6.2.3.15.

6.4.2 Random jitter

As described in Section 6.3.2, the error on the HWP angle $\delta\alpha$ can be characterised as a random jitter, parametrised by $\sigma_{TOD}^{\delta\alpha}$. By increasing the level of random jitter, we seek to ascertain the degree of bias that the HWP angle error $\delta\alpha$ introduces to the tensor-to-scalar ratio r.

The results are shown in Fig. 6.7, which illustrates the E- and B-mode power spectra calculated from the maps that are produced by the map-maker, and from

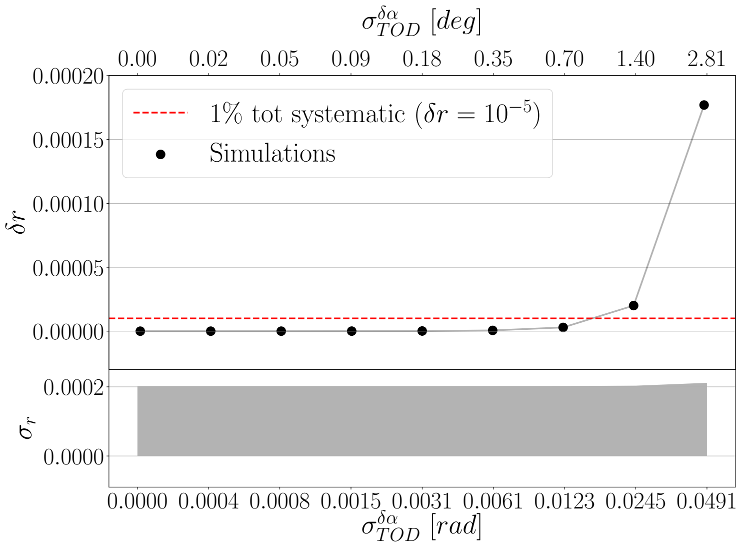

Figure 6.8: *Top panel:* recovered δr from simulations as a function of the angle jitter error. We find that at $\sigma_{TOD}^{\delta\alpha} = 0.7°$, the bias $\delta r = 2.94 \cdot 10^{-6}$, and values of $\sigma_{TOD}^{\delta\alpha} > 0.9°$ significantly bias the tensor to scalar ratio, as δr exceeds the 1% of the total systematic budget (i.e. $\delta r > 10^{-5}$, illustrated by the horizontal dashed red line). *Bottom panel:* σ_r, the error on the recovered best fit δr, slightly increases at higher values of $\sigma_{TOD}^{\delta\alpha}$.

their difference with the ideal input map, as a function of the angle error $\sigma_{TOD}^{\delta\alpha}$. As expected, we find that higher values of the angle error correspond to more biased power spectra, thus the corresponding residual increases. The output C_ℓ's mainly scale in amplitude compared to the input C_ℓ in black, inducing a net depolarisation. However, they only negligibly vary their shape, which indicates that the EB-mixing effect is negligible. Thus, the angle uncertainty averages over sky, and ρ term in Eqns. 6.2.3.14,6.2.3.15 cancels out, while the ϵ term dominates. This is consistent with the result that $C_\ell^{EB,\,\mathrm{meas}}$ is compatible with zero. This happens because the angle error is randomised, and so the uncertainty builds up over time and eventually averages over the sky. Therefore, the random jitter of the angle has the effect of reducing the amplitude of the polarisation maps by a global factor.

The bias on the tensor-to-scalar ratio δr is illustrated in Figure 6.8 as a function of $\sigma_{TOD}^{\delta\alpha}$. We find that values $\sigma_{TOD}^{\delta\alpha} > 0.9$ degrees significantly bias the tensor to scalar ratio, as δr exceeds the 1% of the total systematic budget, illustrated by the horizontal dashed red line. For example, for $\sigma_{TOD}^{\delta\alpha} = [0.7, 1.4, 2.8]$ degrees we recover $\delta r = [2.9 \cdot 10^{-6}, 2 \cdot 10^{-5}, 1.8 \cdot 10^{-4}]$. To reiterate, these results are valid under the assumption that the angle error is purely random. A more accurate physical model for the angle error, including a drift over time, is analysed in the following section.

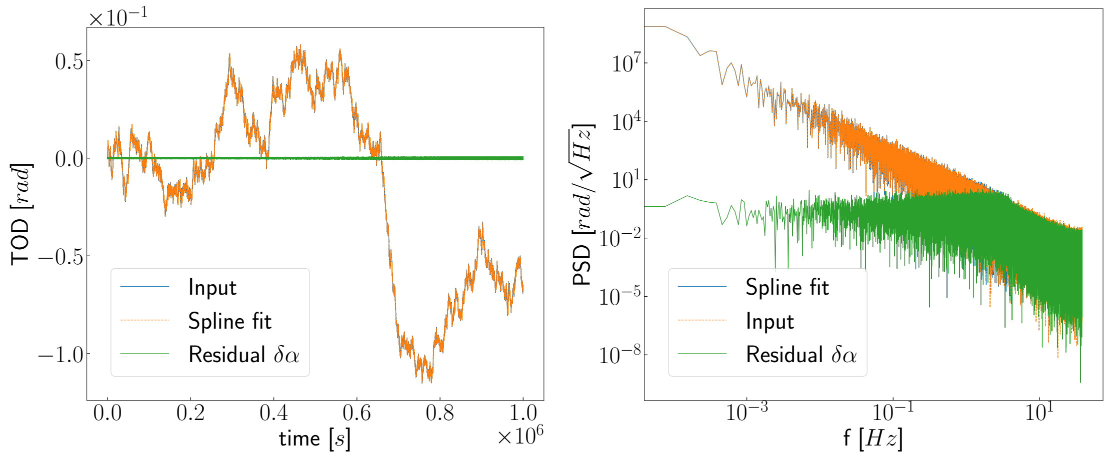

Figure 6.9: *Left:* The TOD of the HWP angle is illustrated for the "input" data generated as a complete drift over time with no corrections from the encoder (orange), its spline fit (blue), and the residual between the two (green) which determines the angle error $\delta\alpha$. *Right:* The corresponding power spectral density (PSD) of each component is shown.

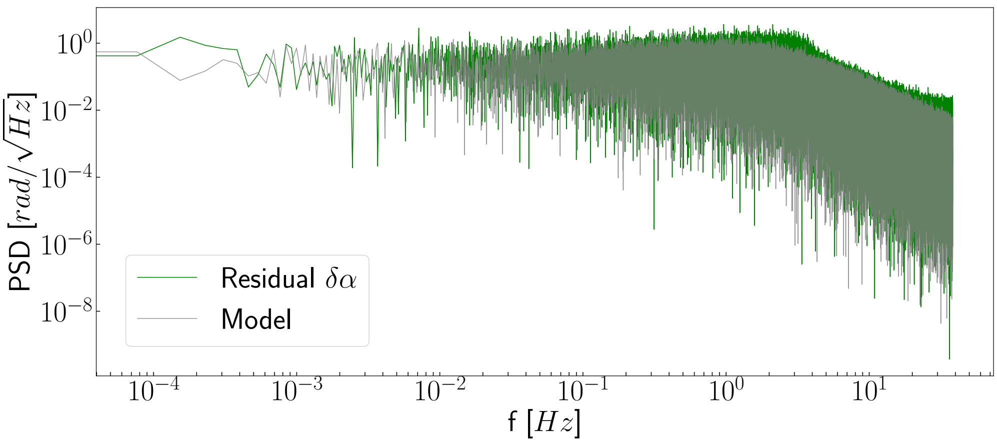

Figure 6.10: The PSD of the downsampled residual (green) is well fitted by the overlapping grey line, which is the best-fit model of Eq. 6.3.3.1 with parameters $\sigma^{\delta\alpha}_{TOD} = 0.0002$ rad/$\sqrt{\text{Hz}}$ and $f_k = 0.0001$ Hz. The residual (green) is dominated by white noise for $f < f_k$.

6.4.3 Drift between measurements

As explained in Section 6.3.3, the error associated with the HWP angle, $\delta\alpha(t)$, can be represented as a drift that is characterised by the model parameters $\sigma^{\delta\alpha}_{TOD}$ and f_k. Using the procedure described in Section 6.3.3, simulations are produced for different values of the model parameters, of which Figure 6.9 illustrates an example. The error $\delta\alpha(t)$ is obtained from the difference between the input data shown in orange and the spline fit shown in blue, of which the the PSD is obtained in the right figure.

The PSD of the residual $\delta\alpha$ is compared with the model of Eq. 6.3.3.1 in the

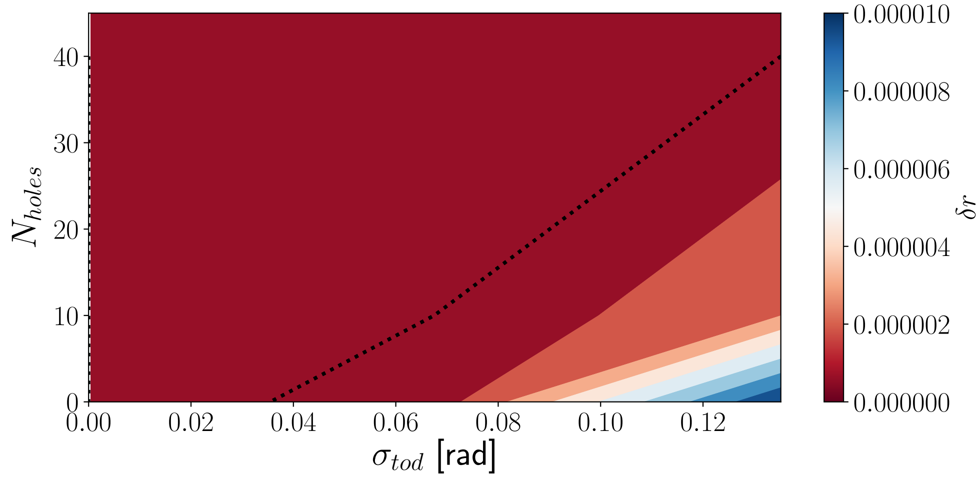

Figure 6.11: Recovered δr for increasing values of $\sigma^{\delta\alpha}_{TOD}$ and N_{holes}. Note that each value of the number of slots in the encoder N_{holes} is directly related to the knee frequency, according to the values in Table 6.1). For low values of N_{holes}, the bias on r increases more rapidly, and the black dotted line corresponds 1% of the total LiteBIRD systematic budget.

Fig. 6.10, where the best-fit model with parameters $\sigma^{\delta\alpha}_{TOD} = 0.0002$ rad/$\sqrt{\text{Hz}}$ and $f_k = 0.0001$ Hz is observed to be a good fit for the data. The PSD of the residual $\delta\alpha$ is approximately flat at $f < f_k$, thus for $f_k \to \infty$ ($N_{\text{holes}} \gg 1$) the spectrum is dominated by white noise. Given the relations in Eqns. 6.3.3.2-6.3.3.3, we have $f_k \propto N_{\text{holes}}$, thus for high number of slots in the encoder we expect the results to be compatible with those of the random jitter angle error presented in the previous Section.

The residual $\delta\alpha(t)$ is included in the HWP angle defined in the TOD simulation step of the analysis pipeline described in Sec 6.2.4. The systematic error for different values of $\sigma^{\delta\alpha}_{TOD}$ and N_{holes}, corresponding to a specific f_{knee}, following the relations 6.3.3.2-6.3.3.3 (the values are displayed in Table 6.1), gets propagated through the pipeline to recover the corresponding bias on the tensor-to-scalar ratio δr. The results are illustrated in the contour plot of Fig. 6.11. We find that for lower N_{holes}, the bias δr increases more rapidly. Higher δr values are obtained for lower f_{knee} and higher $\sigma^{\delta\alpha}_{TOD}$, corresponding to the bottom right corner in the Figure. The 1% of the total LiteBIRD systematic budget corresponds to the black dotted line. Interestingly, we find that a $\delta r > 10^{-5}$ is achieved for values of $N_{\text{holes}} > 40$ with a drift corresponding to $\sigma_{\text{tod}} \gtrsim 0.13$ rad (or $\gtrsim 7.4$ degrees). This suggests that the number of encoder slots currently adopted in LiteBIRD for measuring the HWP angle can widely be relaxed. For $N_{\text{holes}} = 10$, the 1% of the total bias δr allowed by LiteBIRD corresponds to $\sigma_{\text{tod}} \sim 3.4$ degrees.

6.5 Conclusions

In this Chapter, we have investigated the impact of half-wave-plate (HWP) rotational instabilities on measurements of the Cosmic Microwave Background. By studying the HWP phase instabilities we can address the issue of the encoder polarization angle uncertainties for LiteBIRD. The miscalibration of the polarization angle can be a major source of the systematic error, yielding a mixing of E-modes and B-modes which biases the estimation of the the tensor-to-scalar ratio r. Therefore, the question we have been trying to address is: how well do we need to measure the HWP angles in order to prevent certain systematics? The answer has ultimately allowed us to set the instrumental requirements for the HWP angle encoder.

We have utilized the HWP Mueller matrix formalism to model the effects of HWP systematics, as described in Sections 6.2.1,6.2.2. This allowed us to build a general formalism to describe the propagation of different instrumental systematics to power spectra, described in Sec. 6.2.3. For this study, we have built an end-to-end analysis pipeline, which includes the transformation of time-ordered data (TOD) to sky maps and the propagation of errors to the tensor-to-scalar ratio r, was described in Sec. 6.2.4.

With this, we propagate the simulations of HWP angle systematics to final constraints. The systematic effect was simulated as originating from increasingly realistic error of the HWP rotation angle, starting from a constant offset, to a random jitter, and finally to a drift between measurements, as described in Sec. 6.3. Our results have shown that these effects can introduce significant biases in the power spectra, with constant offset having the largest impact. However, the more realistic models for the angle error reveal that the effect on the final cosmological constrain is negligible, suggesting that the number of slots that are physically allowed in the LiteBIRD encoder mechanism can effectively be relaxed.

It is worth noting that this study was performed using CMB-only simulations at single frequencies (140 GHz). Whilst this allows us to emphasize the systematic effects on C_ℓ^{BB}, a more thorough study including foreground components and multi-frequency observations should be considered. In this setup, three HWPs, each corresponding to a dedicated telescope (low-, mid-, and high-frequency) should be included. This will be investigated in a future work.

Furthermore, the HWP rotational instabilities may introduce an effective $1/f$ signal for the polarization. The tail from this effective $1/f$ signal would appear at four times the HWP rotation frequency in the modulated timestreams. The coupling of the large modulated signal at $4f_{HWP}$, where the monopole sits, with other systematic effects such as non-linearities or time-varying gains, could induce artefacts in the data. This analysis would allow us to set some requirements for the instrumental polarization, whilst also mitigating the focal plane temperature variations induced by the gain. We leave a more thorough study of the impact of HWP rotational instabilities coupled with other systematic effects for future work.

"The Answer to the ultimate question of Life,
The Universe and Everything is...42!"

— Douglas Adams, The Hitchhiker's Guide to the Galaxy

Chapter 7

Conclusions

The potential of primordial CMB B-modes to open a new window into the physics of the very early Universe makes the search for this faint CMB polarization signal one of the most compelling goals of modern cosmology. A strict upper constraint on the large-scale amplitude of tensor fluctuations has the potential to further support or disprove various inflationary models [61].

However, as explained in Chapter 2, their detection is complicated by several analysis and technical challenges. Many sources of systematic uncertainty contaminate the faint primordial $B-$modes with the creation of "spurious" $B-$mode signal. Among these, the most significant astrophysical source is coming from Galactic polarized foregrounds [98]. In particular, the combination of polarized synchrotron and thermal dust emission predominates over the CMB B-mode signal on degree angular scales over the whole frequency range covered by current and forthcoming CMB experiments.

Therefore, one of the most important steps in the analysis of CMB B-mode data is the separation of the multi-frequency data into its different components. In principle, the separation of CMB radiation and foregrounds can be performed either in real space, dealing with maps of the Q/U Stokes parameters, or in Fourier space, processing the E/B modes power spectra. Different CMB experiments have followed either approach, revealing their advantages and caveats. Whilst achieving satisfactory results on simulated data, current methods are facing additional challenges put forth by the increasing sensitivity of the next-generation CMB experiments, where the precision of the different approaches will become crucial.

In Chapters 3-5, I have explored new methods to do component separation in the presence of spatial variations in the spectral properties of Galactic foregrounds and their impact on parameter inference. In Chapter 3, I led the development of a power-spectrum-based component separation approach, based on a moment expansion method. The method is able to account for the spatial variability while recovering unbiased constraints on r in the presence of complex foregrounds and noise. The expansion of this method has also been developed in the context of space-based experiments in a collaborative effort within LiteBIRD [2].

To achieve these results, I led significant contributions to the Simons Observatory (SO) cross-spectrum analysis pipeline `BBPipe`, and I implemented and tested the submodule `BBPower` in the case of higher levels of Galactic foregrounds complexity. This has become the main Small Aperture Telescope forecast likelihood pipeline adopted by SO. This analysis pipeline has been employed within a collaborative effort, where we produce the forecasts of different component separation methods, in order to reach an agreement on unified frameworks for forecasting, as presented in Ch. 4. This work has revealed that the moments-based method is currently the most promising avenue for SO, being the only pipeline able to extract unbiased constraints in the presence of realistic foregrounds and correlated and inhomogeneous noise, with a small degradation in statistical uncertainties.

In Chapter 5, I have extended this methodology through the development of a two-step hybrid method, which makes use of both map-based and power-spectrum-based techniques, and found that it achieves tighter constraints while remaining unbiased, even on larger sky fractions (up to $f_{sky} = 80\%$). This method addresses the intrinsic shortcomings of the method described in the previous chapter, and has the potential to adapt to the higher sensitivity needs of different forthcoming CMB experiments, both ground- and space-based.

Furthermore, an excellent control of instrumental systematic contamination is needed to image the low primordial B-mode polarization signature. For example, any mismatches in the polarized beams or detector gains will lead to contamination of the measured polarization by the much stronger unpolarised signal, referred to as I-P leakage [131]. These issues can be mitigated by including additional polarization modulation schemes into the instrument design, which motivates the use of Half Wave Plates (HWPs) by several ground-based experiments. For the first time on a CMB satellite mission, the HWP will also be a key element of LiteBIRD, with the ability to overcome low frequency noise and to remove many systematic effects, but which will potentially introduce important new systematic effects to model.

The detailed study of HWP systematics and their impact on cosmological parameters is therefore crucial for the success of the next-generation CMB experiments. For this purpose, in Chapter 6 I led a study on the precision of measurement of the HWP angular position in the context of LiteBIRD. I implemented a stand-alone simulation pipeline, `STOMPER`, going from time-ordered-data (TOD) to maps through a preconditioned conjugate gradient (PCG) map-making algorithm to mitigate the propagation of HWP systematic effects to final constraints. This study is used as a reference within the collaboration and was recently published in the collaboration's article [228] and my first-author publication is in preparation.

This thesis presents robust data analysis methods and pipelines that will empower the upcoming generation of CMB polarization experiments to acquire competitive constraints. With increased observational power comes increased sensitivity to astrophysical and systematic sources of noise, which most dominantly limit our ability to extract information from the data. The methods developed here aid in characterizing and evaluating both sources of noise, allowing us to explore an uncharted region of parameter space. This will enable us to answer fundamental questions about the early Universe and potentially provide evidence for the quantization of gravity, while uncovering new physics at the level of grand unified theories.

Although in this thesis we develop several complementary frameworks to analyse CMB polarization data at different stages, CMB polarization studies often require consistent data analysis starting from raw unprocessed TOD to final cosmological parameters. Thus, it would be desirable to develop more comprehensive integrated data analysis techniques, whereby a set of modules are iterated between to accomplish the analysis goal, as opposed to modular, non-integrated pipelines. This approach supports simultaneous treatment of different instrumental systematics, noise processing and foreground errors. Most of the current data analysis pipelines are divided into stages: analysis of the raw data to clean up the instrumental effects; map-making; foreground cleaning; likelihood analysis. Whilst modular approaches are more intuitive, their division is far from optimal, because the projections of different systematics on the final maps are correlated. Moreover, the astrophysical component maps can be affected differently from frequency to frequency by a single systematic effect, thus leaving a residual on the CMB maps after component separation. This introduces the need for fully integrated pipelines for end-to-end simulations. Many examples of integrated data processing pipelines have demonstrated the advantages of this approach, for example in the context of Planck, [227, 230–233], where fitting for multiple systematics jointly maximizes the signal-to-noise ratio (SNR). The tools developed in this thesis, in the context of LiteBIRD and SO, should be expanded in the future to include different experiments, and classes of systematics, both in time-domain and pixel-space, that are of interest for forthcoming CMB surveys. An integrated analysis will be able to make use of the full statistical power of the experiment, where maximising the SNR is of the essence, whilst also allowing a detailed control over systematics. The joint analysis of different datasets, which will be crucial for the first data releases of future experiments, would also favour this approach.

Future work should also address the need for consistent analysis pipelines which account for the volume of likelihood space for non-linear parameters. The effects of priors and marginalization on the current Bayesian constraints on the tensor-to-scalar ratio r and other parameters can be highly influential. For example, the results of Chapter 3 revealed a parameter degeneracy between r and the amplitudes of the spectral indices βs. This is a well-known effect that arises when imposing physically-motivated priors on the parametrization used for spectral index variation [102]. One common solution is to use Jeffreys' priors for the spectral parameters [161] implemented in Bayesian MCMC analyses, although the comparison with other likelihood sampling techniques has not been explored in this context. A contingent exploration would be that of optimal power spectrum estimators, which is currently limited to the Pseudo-C_ℓ method in this thesis.

Although primarily applied to the analysis of CMB data in this thesis, many of the techniques introduced can be applicable to other areas of astrophysics. For example, the "discarded" higher order terms of the moments method in Chapter 3, that are of no use to recover r, potentially encode valuable information on the foreground components that can be used for galactic science studies. Similarly, a frequency-dependent residual as that obtained with the hybrid method in Chapter 5 has the potential to open up the exploration of many spatially varying components, for example to extract the kSZ signal accurately and to study the large-scale structure of the Universe. Another example is the formalism introduced in the study of the HWP angle errors introduced in Chapter 6, which is general enough that it can be used to study parity-violating EB physics induced by any angle error. The primary objective of the techniques used throughout this thesis is to detect correlations in data that are plagued by noise and may possess unknown characteristics. This enables the extraction of a large amount of information from regions of parameter space that have not been thoroughly explored. By accomplishing this goal, there is a tantalising prospect of obtaining competitive constraints on fundamental physics through innovative approaches.

Appendices

Appendix A

A.1 SVT decomposition

The perturbed geometry can be parametrised in terms of SVT as follows:

$$h_{\mu\nu}dx^\mu dx^\nu = -2Ad\eta^2 - 2B_i d\eta dx^i + 2H_{ij}dx^i dx^j \tag{A.1.0.1}$$

where A corresponds to the scalar perturbation, B_i is the vector perturbation, and the H_{ij} is the tensor perturbation.

We therefore consider each term separately below.

- **Scalars** - A scalar function can generally be decomposed into a "background" mode $\bar{A}$ and a perturbation term δA,

$$A = \bar{A} + \delta A \tag{A.1.0.2}$$

- **Vectors** - By Helmholtz's theorem, any 3-vector can be split into the gradient of a scalar B and and the curl of a vector B_i (i.e. a divergenceless, or "transverse" vector, such that $\partial^i \hat{B}_i = 0$):

$$B_i = \partial_i B + \hat{B}_i. \tag{A.1.0.3}$$

- **Tensors** - The tensor perturbation, defined by H_{ij}, is traceless ($H^i_i = 0$) and transverse ($\nabla^i H_{ij} = 0$). Similarly to vectors, any rank-2 symmetric tensor can be written as

$$H_{ij} = 2C\delta_{ij} + 2\partial_{\langle i}\partial_{j\rangle}E + 2\partial_{\langle i}\hat{F}_{j)} + 2\hat{E}_{ij}, \tag{A.1.0.4}$$

  where the first two terms correspond to a scalar quantity, the third to a vector, the fourth to a tensor, and we define

$$\begin{aligned} \partial_{\langle i}\partial_{j\rangle}E &\equiv \left(\partial_i\partial_j - \frac{1}{3}\delta_{ij}\nabla^2\right)E \\ \partial_{(i}\hat{F}_{j)} &\equiv \frac{1}{2}\left(\partial_i\hat{F}_j + \partial_j\hat{F}_i\right) \end{aligned} \tag{A.1.0.5}$$

  where ˆ indicates transverse quantities, as is the vector $\hat{\mathbf{F}}$ ($\partial_i F_i = 0$).

A.2 Gauge transformations for linear perturbations

- **Scalars.** Consider a scalar function $\phi(x)$ which, in coordinates x^μ is composed of a "background" mode $\bar{\phi}(x)$ and a perturbation $\delta\phi$, i.e.:

$$\phi(x) = \bar{\phi}(x) + \delta\phi(x). \tag{A.2.0.1}$$

Since ϕ is a scalar, in the new coordinates $\tilde{\phi}(\tilde{x}) = \phi(x)$, therefore:

$$\tilde{\phi}(\tilde{x}) = \phi(x) = \phi(\tilde{x} - \xi) = \bar{\phi}(\tilde{x}) + \delta\phi(\tilde{x}) - \xi^\mu \partial_\mu \bar{\phi}, \tag{A.2.0.2}$$

where we have ignored all terms of higher order than 1 in $\delta\phi$ or ξ. In the new coordinates, the perturbation is therefore

$$\tilde{\delta\phi} = \delta\phi - \xi^\mu \delta_\mu \bar{\phi}. \tag{A.2.0.3}$$

An example of such a scalar is the energy density ρ, which can be expressed as $\rho = U^\mu U^\nu T_{\mu\nu}$.

- **Vectors.** As before, contravariant vector $V^\mu(x) = \bar{V}^\mu(x) + \delta V^\mu$ [1]. In this case, we need to take into account the transformation law for vectors, and expanding to linear order:

$$\begin{aligned}
\tilde{V}^\mu(\tilde{x}) &= \frac{\partial \tilde{x}^\mu}{\partial x^\nu} V^\nu(x) \\
&= (\delta^\mu_\nu + \partial_\nu \xi^\mu)\left(\bar{V}^\nu(\tilde{x}) + \delta V^\nu - \xi^\rho \partial_\rho \bar{V}^\nu\right) \\
&= \bar{V}^\mu(\tilde{x}) + \delta v^\mu - \xi^\nu \partial_\nu \bar{V}^\mu + \bar{V}^\nu \partial_\nu \xi^\mu,
\end{aligned} \tag{A.2.0.4}$$

and thus the perturbation is modified as

$$\delta\tilde{V}^\mu = \delta V^\mu - \xi^\nu \partial_\nu \bar{V}^\mu + \bar{V}^\nu \partial_\nu \xi^\mu. \tag{A.2.0.5}$$

Examples of contravariant vectors include displacement vectors, velocity vectors, and momentum vectors in classical mechanics, as well as four-momentum vectors and electromagnetic four-potentials in special and general relativity.

[1] The term "contravariant" refers to the fact that the components of the vector transform in the opposite way to the basis vectors of the coordinate system. In other words, if the basis vectors transform covariantly (i.e., with the Jacobian matrix), then the components of a contravariant vector transform contravariantly (i.e., with the inverse of the Jacobian matrix).

- **Tensors.** For a rank-2 covariant tensor, such as the metric,

$$\tilde{g}_{\mu\nu}(\tilde{x}) = \frac{\partial x^\rho}{\partial \tilde{x}^\mu}\frac{\partial x^\sigma}{\partial \tilde{x}^\nu} g_{\mu\nu}(x). \tag{A.2.0.6}$$

To derive the transformation law for linear perturbations of the metric tensor, we start by considering a small perturbation of the metric tensor, $\delta g_{\mu\nu}(x)$, and a small diffeomorphism transformation, $\xi^\mu(x)$. We then consider the transformed metric tensor, $\tilde{g}_{\mu\nu}(\tilde{x})$, where $\tilde{x}^\mu = x^\mu + \xi^\mu(x)$. Using the chain rule for partial derivatives, we can express the transformed metric in terms of the original metric and the perturbations as follows:

$$\begin{aligned}\tilde{g}_{\mu\nu}(\tilde{x}) &= \frac{\partial x^\rho}{\partial \tilde{x}^\mu}\frac{\partial x^\sigma}{\partial \tilde{x}^\nu} g_{\rho\sigma}(x+\xi)\\ &= g_{\rho\sigma}(x+\xi)\frac{\partial x^\rho}{\partial (x^\mu+\xi^\mu)}\frac{\partial x^\sigma}{\partial (x^\nu+\xi^\nu)}\\ &= g_{\rho\sigma}(x+\xi)\left(\delta^\rho_\mu + \partial_\mu\xi^\rho\right)\left(\delta^\sigma_\nu + \partial_\nu\xi^\sigma\right)\\ &= g_{\mu\nu}(x) + \delta g_{\mu\nu}(x) - \xi^\rho\partial_\rho g_{\mu\nu}(x) - g_{\mu\rho}\partial_\nu\xi^\rho - g_{\rho\nu}\partial_\mu\xi^\rho + \mathcal{O}(\xi^2),\end{aligned} \tag{A.2.0.7}$$

where we have expanded the transformed metric to linear order in both the perturbation and the diffeomorphism transformation. This gives us the desired transformation law for linear metric perturbations:

$$\tilde{\delta g}_{\mu\nu} = \delta g_{\mu\nu} - \xi^\rho\partial_\rho\bar{g}_{\mu\nu} - \bar{g}_{\mu\rho}\partial_\nu\xi^\rho - \bar{g}_{\rho\nu}\partial_\mu\xi^\rho, \tag{A.2.0.8}$$

Inspecting Eqs. A.2.0.3, A.2.0.5, and A.2.0.8, a general transformation law is easy to infer for tensor of arbitrary rank and index position.

A.3 Spin-weighted decomposition

The conventional spherical harmonics on the surface of a sphere are insufficient for expanding polarisation. Instead, the appropriate mathematical framework for representing the angular distribution of the polarization of the CMB on the celestial sphere is given by the spin-weighted harmonics $sY\ell m(\hat{n})$. These functions constitute a comprehensive and orthonormal basis set on the sphere:

$$\int d\Omega_s Y^*_{\ell' m'}(\theta,\phi)_s Y_{\ell m}(\theta,\phi) = \delta_{\ell'\ell}\delta_{m'm} \tag{A.3.0.1}$$

$$\sum_{\ell,m} {}_sY^*_{\ell m}(\theta,\phi)_s Y_{\ell m}\left(\theta',\phi'\right) = \delta\left(\phi-\phi'\right)\delta\left(\cos\theta-\cos\theta'\right). \tag{A.3.0.2}$$

The parity relation for these functions is

$${}_sY_{\ell m} \rightarrow (-1)^l {}_{-s}Y_{\ell m} \tag{A.3.0.3}$$

A pair of raising and lowering operators, namely $\eth$ and $\bar{\eth}$, are available to modify the spin-weight of a function. These operators raise and lower the spin-weight of a function, i.e.

$$(\eth_s f)' = e^{-i(s+1)\psi} (\eth_s f), \tag{A.3.0.4}$$

$$\left(\bar{\eth}_s f\right)' = e^{-i(s-1)\psi} \left(\bar{\eth}_s f\right), \tag{A.3.0.5}$$

where the prime refers to a quantity in a frame rotated ψ from the original frame. The explicit form of these operators is

$$\eth_s f(\theta,\phi) = -\sin^s\theta \left[\partial_\theta + \frac{i}{\sin\theta}\partial_\phi\right] \sin^{-s}\theta_s f(\theta,\phi), \tag{A.3.0.6}$$

$$\bar{\eth}_s f(\theta,\phi) = -\sin^{-s}\theta \left[\partial_\theta - \frac{i}{\sin\theta}\partial_\phi\right] \sin^{s}\theta_s f(\theta,\phi). \tag{A.3.0.7}$$

Using the raising and lowering operators, one can relate the spin- s spherical harmonics and the usual spherical harmonics

$${}_sY_{\ell m} = \left[\frac{(\ell-s)!}{(\ell+s)!}\right]^{1/2} \eth^s Y_{\ell m} \tag{A.3.0.8}$$

for $0 \leq s \leq \ell$, and

$${}_sY_{\ell m} = \left[\frac{(\ell+s)!}{(\ell-s)!}\right]^{1/2} (-1)^s \bar{\eth}^{-s} Y_{\ell m} \tag{A.3.0.9}$$

for $-\ell \leq s \leq 0$. Hence, using these, we can obtain

$${}_{\pm 2}Y_{\ell m} = \left[\frac{(\ell-s)!}{(\ell+s)!}\right]^{1/2} \left[\partial_\theta^2 - \cot\theta\partial_\theta \pm \frac{2i}{\sin\theta}\left(\partial_\theta - \cot\theta\right)\partial_\phi - \frac{1}{\sin^2\theta}\partial_\phi^2\right] Y_{\ell m}. \tag{A.3.0.10}$$

Useful properties of the spin-weighted spherical harmonics include

$${}_sY^*_{\ell m} = (-1)^{m+s}{}_{-s}Y_{\ell -m}, \tag{A.3.0.11}$$

$$\eth_s Y_{\ell m} = [(\ell-s)(\ell+s+1)]^{1/2}{}_{s+1}Y_{\ell m}, \tag{A.3.0.12}$$

$$\bar{\eth}_s Y_{\ell m} = -[(\ell+s)(\ell-s+1)]^{1/2}{}_{s-1}Y_{\ell m}, \tag{A.3.0.13}$$

$$\bar{\eth}\eth_s Y_{\ell m} = -(\ell-s)(\ell+s+1)_s Y_{\ell m}. \tag{A.3.0.14}$$

A.4 C_ℓ's and 2-point correlation function

Using the Legendre polynomials

$$\sum_m Y^*_{\ell m}(\mathbf{n}_i)\, Y_{\ell m}(\mathbf{n}_j) = \frac{2\ell+1}{4\pi} P_\ell\left(\hat{\mathbf{n}}_i \cdot \hat{\mathbf{n}}_j\right) \tag{A.4.0.1}$$

and the addition theorem, we can establish a connection between the angular power spectrum and the 2-point correlation function in real space as follows:

$$\xi(\theta) = \langle \Theta_i \Theta_j \rangle = \sum_{\ell} \frac{2\ell+1}{4\pi} C_\ell P_\ell \left(\hat{\mathbf{n}}_i \cdot \hat{\mathbf{n}}_j \right) \tag{A.4.0.2}$$

Due to the isotropy of the universe, the two-point correlation function only depends on the angular separation θ between two points on the sky, and not on the specific orientation of the separation.

The two-point correlation function, $\xi(\theta)$, is less informative than the angular power spectrum, C_ℓ, because it lacks distinguishing features, such as the prominent acoustic peaks. Additionally, $\xi(\theta)$ suffers from correlated errors. However, it can be useful for certain calculations, such as replacing the convolution of two C_ℓ values with the product of two $\xi(\theta)$ values in some cases, as demonstrated by [234]:

References

[1] S. Azzoni et al. "A minimal power-spectrum-based moment expansion for CMB B-mode searches". In: *Journal of Cosmology and Astroparticle Physics* 2021.5, 047 (May 2021), p. 047. arXiv: 2011.11575 [astro-ph.CO].

[2] L. Vacher et al. "Moment expansion of polarized dust SED: A new path towards capturing the CMB B-modes with LiteBIRD". In: *Astronomy and Astrophysics* 660, A111 (Apr. 2022), A111. arXiv: 2111.07742 [astro-ph.CO].

[3] K. Wolz et al. "The Simons Observatory: pipeline comparison and validation for large-scale B-modes". In: *arXiv e-prints*, arXiv:2302.04276 (Feb. 2023), arXiv:2302.04276. arXiv: 2302.04276 [astro-ph.CO].

[4] S. Azzoni et al. "A hybrid map-C_ℓ component separation method for primordial CMB B-mode searches". In: *Journal of Cosmology and Astroparticle Physics* 2023.3, 035 (Mar. 2023), p. 035. arXiv: 2210.14838 [astro-ph.CO].

[5] T. Hasebe et al. "Sensitivity Modeling for LiteBIRD". In: *Journal of Low Temperature Physics* (Dec. 2022).

[6] J. Hubmayr et al. "Optical Characterization of OMT-Coupled TES Bolometers for LiteBIRD". In: *Journal of Low Temperature Physics* 209.3-4 (Nov. 2022), pp. 396–408.

[7] Ryota Takaku et al. "Impact of the effective thickness from anti-reflective sub-wavelength structures in achromatic half-wave plate design". In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*. Ed. by Jonas Zmuidzinas and Jian-Rong Gao. Vol. 12190. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. Aug. 2022, 121902K, 121902K.

[8] E. Allys et al. "Probing cosmic inflation with the LiteBIRD cosmic microwave background polarization survey". In: *Progress of Theoretical and Experimental Physics* 2023.4, 042F01 (Apr. 2023), 042F01.

[9] Nadia Dachlythra et al. "The Simons Observatory: Beam characterization for the Small Aperture Telescopes". In: *arXiv e-prints*, arXiv:2304.08995 (Apr. 2023), arXiv:2304.08995. arXiv: 2304.08995 [astro-ph.IM].

[10] R. A. Alpher, H. Bethe, and G. Gamow. "The Origin of Chemical Elements". In: *Physical Review* 73.7 (Apr. 1948), pp. 803–804.

[11] G. F. Smoot et al. "Structure in the COBE Differential Microwave Radiometer First-Year Maps". In: *Astrophysical Journal, Letters* 396 (Sept. 1992), p. L1.

[12] D. J. Fixsen et al. "Cosmic Microwave Background Dipole Spectrum Measured by the COBE FIRAS Instrument". In: *Astrophysical Journal* 420 (Jan. 1994), p. 445.

[13] A. D. Miller et al. “A Measurement of the Angular Power Spectrum of the Cosmic Microwave Background from L = 100 to 400”. In: *Astrophysical Journal, Letters* 524.1 (Oct. 1999), pp. L1–L4. arXiv: astro-ph/9906421 [astro-ph].

[14] P. de Bernardis et al. “Multiple Peaks in the Angular Power Spectrum of the Cosmic Microwave Background: Significance and Consequences for Cosmology”. In: *Astrophysical Journal* 564.2 (Jan. 2002), pp. 559–566. arXiv: astro-ph/0105296 [astro-ph].

[15] S. Hanany et al. “MAXIMA-1: A Measurement of the Cosmic Microwave Background Anisotropy on Angular Scales of 10’-5°”. In: *Astrophysical Journal, Letters* 545.1 (Dec. 2000), pp. L5–L9. arXiv: astro-ph/0005123 [astro-ph].

[16] A. Benoit et al. “The cosmic microwave background anisotropy power spectrum measured by Archeops”. In: *Astronomy and Astrophysics* 399 (Mar. 2003), pp. L19–L23. arXiv: astro-ph/0210305 [astro-ph].

[17] N. W. Halverson et al. “Degree Angular Scale Interferometer First Results: A Measurement of the Cosmic Microwave Background Angular Power Spectrum”. In: *Astrophysical Journal* 568.1 (Mar. 2002), pp. 38–45. arXiv: astro-ph/0104489 [astro-ph].

[18] Paul F. Scott et al. “First results from the Very Small Array - III. The cosmic microwave background power spectrum”. In: *Monthly Notices of the RAS* 341.4 (June 2003), pp. 1076–1083. arXiv: astro-ph/0205380 [astro-ph].

[19] T. J. Pearson et al. “The Anisotropy of the Microwave Background to l = 3500: Mosaic Observations with the Cosmic Background Imager”. In: *Astrophysical Journal* 591.2 (July 2003), pp. 556–574. arXiv: astro-ph/0205388 [astro-ph].

[20] C. L. Kuo et al. “High-Resolution Observations of the Cosmic Microwave Background Power Spectrum with ACBAR”. In: *Astrophysical Journal* 600.1 (Jan. 2004), pp. 32–51. arXiv: astro-ph/0212289 [astro-ph].

[21] Kendrick M. Smith, Oliver Zahn, and Olivier Doré. “Detection of gravitational lensing in the cosmic microwave background”. In: *Physical Review D* 76.4, 043510 (Aug. 2007), p. 043510. arXiv: 0705.3980 [astro-ph].

[22] Planck Collaboration et al. “Planck early results. I. The Planck mission”. In: *Astronomy and Astrophysics* 536, A1 (Dec. 2011), A1. arXiv: 1101.2022 [astro-ph.IM].

[23] J. M. Kovac et al. “Detection of polarization in the cosmic microwave background using DASI”. In: *Nature* 420.6917 (Dec. 2002), pp. 772–787. arXiv: astro-ph/0209478 [astro-ph].

[24] D. Barkats et al. “First Measurements of the Polarization of the Cosmic Microwave Background Radiation at Small Angular Scales from CAPMAP”. In: *Astrophysical Journal, Letters* 619.2 (Feb. 2005), pp. L127–L130. arXiv: astro-ph/0409380 [astro-ph].

[25] Steven Weinberg. *Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity*. New York, NY: Wiley, 1972. URL: https://cds.cern.ch/record/100595.

[26] Scott Dodelson. *Modern Cosmology*. 2003.

[27] James B. Hartle. *Gravity : an introduction to Einstein’s general relativity*. 2003.

[28] Steven Weinberg. *Cosmology*. 2008.

[29] Ruth Durrer. *The Cosmic Microwave Background*. Cambridge University Press, 2008.

[30] Daniel Baumann. *Cosmology lecture notes*. URL: http://cosmology.amsterdam/education/cosmology/.

[31] Eric Gourgoulhon. *3+1 Formalism in General Relativity*. Vol. 846. 2012.

[32] Planck Collaboration et al. "Planck 2018 results. VI. Cosmological parameters". In: *Astronomy and Astrophysics* 641, A6 (Sept. 2020), A6. arXiv: 1807.06209 [astro-ph.CO].

[33] Adam G. Riess et al. "Large Magellanic Cloud Cepheid Standards Provide a 1% Foundation for the Determination of the Hubble Constant and Stronger Evidence for Physics beyond ΛCDM". In: *Astrophysical Journal* 876.1, 85 (May 2019), p. 85. arXiv: 1903.07603 [astro-ph.CO].

[34] M. Betoule et al. "Improved cosmological constraints from a joint analysis of the SDSS-II and SNLS supernova samples". In: *Astronomy and Astrophysics* 568, A22 (Aug. 2014), A22. arXiv: 1401.4064 [astro-ph.CO].

[35] Hideo Kodama and Misao Sasaki. "Cosmological Perturbation Theory". In: *Progress of Theoretical Physics Supplement* 78 (Jan. 1984), p. 1.

[36] V. F. Mukhanov, H. A. Feldman, and R. H. Brandenberger. "Theory of cosmological perturbations". In: *Physics Reports* 215.5-6 (June 1992), pp. 203–333.

[37] Chung-Pei Ma and Edmund Bertschinger. "Cosmological perturbation theory in the synchronous and conformal Newtonian gauges". In: *Astrophys. J.* 455 (1995), pp. 7–25. arXiv: astro-ph/9506072.

[38] Ruth Durrer. "Cosmological Perturbation Theory". In: ed. by E. Papantonopoulos. Vol. 653. 2005, p. 31.

[39] Daniel Baumann. "TASI Lectures on Primordial Cosmology". In: *arXiv e-prints*, arXiv:1807.03098 (July 2018), arXiv:1807.03098. arXiv: 1807.03098 [hep-th].

[40] Edward W. Kolb and Michael S. Turner. *The early universe*. Vol. 69. 1990.

[41] John A Peacock. *Cosmological Physics*. Cambridge University Press, 1999.

[42] John Peacock. *Advanced Cosmology*. URL: http://www.roe.ac.uk/japwww/teaching/cos5.html.

[43] Antony Lewis. *Advanced Cosmology*. URL: https://cosmologist.info/teaching/EU/.

[44] G. Gamow. "Expanding Universe and the Origin of Elements". In: *Physical Review* 70.7-8 (Oct. 1946), pp. 572–573.

[45] Ralph A. Alpher and Robert C. Herman. "On the Relative Abundance of the Elements". In: *Physical Review* 74.12 (Dec. 1948), pp. 1737–1742.

[46] Ralph A. Alpher and Robert C. Herman. "Remarks on the Evolution of the Expanding Universe". In: *Physical Review* 75.7 (Apr. 1949), pp. 1089–1095.

[47] Andrew R Liddle and David H Lyth. *Cosmological Inflation and Large-Scale Structure*. Cambridge: Cambridge Univ. Press, 2000. URL: https://cds.cern.ch/record/452061.

[48] V. F. Mukhanov and G. V. Chibisov. “Quantum fluctuations and a nonsingular universe”. In: *Soviet Journal of Experimental and Theoretical Physics Letters* 33 (May 1981), p. 532.

[49] A. H. Guth and S. -Y. Pi. “Fluctuations in the New Inflationary Universe”. In: *Physical Review Letters* 49.15 (Oct. 1982), pp. 1110–1113.

[50] A. A. Starobinsky. “Dynamics of phase transition in the new inflationary universe scenario and generation of perturbations”. In: *Physics Letters B* 117.3-4 (Nov. 1982), pp. 175–178.

[51] S. W. Hawking. “The development of irregularities in a single bubble inflationary universe”. In: *Physics Letters B* 115.4 (Sept. 1982), pp. 295–297.

[52] Planck Collaboration X. “Planck 2018 results. X. Constraints on inflation”. In: *Astronomy and Astrophysics* 641, A10 (Sept. 2020), A10. arXiv: 1807.06211 [astro-ph.CO].

[53] Uroš Seljak and Matias Zaldarriaga. “Signature of Gravity Waves in the Polarization of the Microwave Background”. In: *Physical Review Letters* 78.11 (Mar. 1997), pp. 2054–2057. arXiv: astro-ph/9609169 [astro-ph].

[54] Yasin Memari. *Polarization of the Cosmic Microwave Background Radiation*. 2007. URL: https://www.roe.ac.uk/ifa/postgrad/pedagogy/2007_memari.pdf.

[55] Anthony Challinor. *Large-scale structure formation*. 2009. URL: https://cosmologist.info/teaching/EU/ADC_Structure_formation2.pdf.

[56] Marc Kamionkowski and Ely D. Kovetz. “The Quest for B Modes from Inflationary Gravitational Waves”. In: *Annual Review of Astron and Astrophysis* 54 (Sept. 2016), pp. 227–269. arXiv: 1510.06042 [astro-ph.CO].

[57] M. Bartelmann and P. Schneider. “Weak gravitational lensing”. In: *Physics Reports* 340.4-5 (Jan. 2001), pp. 291–472. arXiv: astro-ph/9912508 [astro-ph].

[58] Antony Lewis and Anthony Challinor. “Weak gravitational lensing of the CMB”. In: *Physics Reports* 429.1 (June 2006), pp. 1–65. arXiv: astro-ph/0601594 [astro-ph].

[59] M. Tristram et al. “Improved limits on the tensor-to-scalar ratio using BICEP and Planck data”. In: *Physical Review D* 105.8, 083524 (Apr. 2022), p. 083524. arXiv: 2112.07961 [astro-ph.CO].

[60] A. A. Starobinsky. “A new type of isotropic cosmological models without singularity”. In: *Physics Letters B* 91.1 (Mar. 1980), pp. 99–102.

[61] Jérôme Martin et al. “The best inflationary models after Planck”. In: *Journal of Cosmology and Astroparticle Physics* 2014.3, 039 (Mar. 2014), p. 039. arXiv: 1312.3529 [astro-ph.CO].

[62] Jérôme Martin, Christophe Ringeval, and Vincent Vennin. “Encyclopædia Inflationaris”. In: *Physics of the Dark Universe* 5 (Dec. 2014), pp. 75–235. arXiv: 1303.3787 [astro-ph.CO].

[63] A. A. Starobinskiǐ. “Spectrum of relict gravitational radiation and the early state of the universe”. In: *Soviet Journal of Experimental and Theoretical Physics Letters* 30 (Dec. 1979), p. 682.

[64] Fedor Bezrukov and Mikhail Shaposhnikov. "The Standard Model Higgs boson as the inflaton". In: *Physics Letters B* 659.3 (Jan. 2008), pp. 703–706. arXiv: 0710.3755 [hep-th].

[65] Anna Ijjas and Paul J. Steinhardt. "Bouncing cosmology made simple". In: *Classical and Quantum Gravity* 35.13, 135004 (July 2018), p. 135004. arXiv: 1803.01961 [astro-ph.CO].

[66] Aaron M. Levy, Anna Ijjas, and Paul J. Steinhardt. "Scale-invariant perturbations in ekpyrotic cosmologies without fine-tuning of initial conditions". In: *Physical Review D* 92.6, 063524 (Sept. 2015), p. 063524. arXiv: 1506.01011 [astro-ph.CO].

[67] Peter A. R. Ade et al. "POLARBEAR constraints on cosmic birefringence and primordial magnetic fields". In: *Physical Review D* 92.12, 123509 (Dec. 2015), p. 123509. arXiv: 1509.02461 [astro-ph.CO].

[68] Carlo R. Contaldi. "Imaging cosmic polarization rotation". In: *International Journal of Modern Physics D* 25.11, 1640014 (June 2016), p. 1640014.

[69] Kevork N. Abazajian et al. "CMB-S4 Science Book, First Edition". In: *arXiv e-prints*, arXiv:1610.02743 (Oct. 2016), arXiv:1610.02743. arXiv: 1610.02743 [astro-ph.CO].

[70] Planck Collaboration et al. "Planck 2015 results. XVII. Constraints on primordial non-Gaussianity". In: *Astronomy and Astrophysics* 594, A17 (Sept. 2016), A17. arXiv: 1502.01592 [astro-ph.CO].

[71] BICEP2 Collaboration et al. "BICEP2 / Keck Array IX: New bounds on anisotropies of CMB polarization rotation and implications for axionlike particles and primordial magnetic fields". In: *Physical Review D* 96.10, 102003 (Nov. 2017), p. 102003. arXiv: 1705.02523 [astro-ph.CO].

[72] Anthony Challinor and Hiranya Peiris. "Lecture notes on the physics of cosmic microwave background anisotropies". In: *Cosmology and Gravitation: XIII Brazilian School on Cosmology and Gravitation (XIII BSCG)*. Ed. by Mario Novello and Santiago Perez. Vol. 1132. American Institute of Physics Conference Series. May 2009, pp. 86–140. arXiv: 0903.5158 [astro-ph.CO].

[73] E. T. Newman and R. Penrose. "Note on the Bondi-Metzner-Sachs Group". In: *Journal of Mathematical Physics* 7.5 (May 1966), pp. 863–870.

[74] J. N. Goldberg et al. "Spin-s Spherical Harmonics and ð". In: *Journal of Mathematical Physics* 8.11 (Nov. 1967), pp. 2155–2161.

[75] W. J. Percival and M. L. Brown. "Likelihood techniques for the combined analysis of CMB temperature and polarization power spectra". In: *Monthly Notices of the RAS* 372.3 (Nov. 2006), pp. 1104–1116. arXiv: astro-ph/0604547 [astro-ph].

[76] Uros Seljak and Matias Zaldarriaga. "A Line-of-Sight Integration Approach to Cosmic Microwave Background Anisotropies". In: *Astrophysical Journal* 469 (Oct. 1996), p. 437. arXiv: astro-ph/9603033 [astro-ph].

[77] Antony Lewis, Anthony Challinor, and Anthony Lasenby. "Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models". In: *Astrophysical Journal* 538.2 (Aug. 2000), pp. 473–476. arXiv: astro-ph/9911177 [astro-ph].

[78] Julien Lesgourgues. “The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview”. In: *arXiv e-prints*, arXiv:1104.2932 (Apr. 2011), arXiv:1104.2932. arXiv: 1104.2932 [astro-ph.IM].

[79] David Alonso et al. “A unified pseudo-C_ℓ framework”. In: *Monthly Notices of the RAS* 484.3 (Apr. 2019), pp. 4127–4151. arXiv: 1809.09603 [astro-ph.CO].

[80] Eric Hivon et al. “MASTER of the Cosmic Microwave Background Anisotropy Power Spectrum: A Fast Method for Statistical Analysis of Large and Complex Cosmic Microwave Background Data Sets”. In: *Astrophysical Journal* 567.1 (Mar. 2002), pp. 2–17. arXiv: astro-ph/0105302 [astro-ph].

[81] Kendrick M. Smith. “Pseudo-C_ℓ estimators which do not mix E and B modes”. In: *Physical Review D* 74.8, 083002 (Oct. 2006), p. 083002. arXiv: astro-ph/0511629 [astro-ph].

[82] J. Grain, M. Tristram, and R. Stompor. “Polarized CMB power spectrum estimation using the pure pseudo-cross-spectrum approach”. In: *Phys. Rev. D* 79 (12 June 2009), p. 123515. URL: https://link.aps.org/doi/10.1103/PhysRevD.79.123515.

[83] Sudeep Das et al. “Detection of the Power Spectrum of Cosmic Microwave Background Lensing by the Atacama Cosmology Telescope”. In: *Physical Review Letters* 107.2, 021301 (July 2011), p. 021301. arXiv: 1103.2124 [astro-ph.CO].

[84] B. A. Benson et al. “SPT-3G: a next-generation cosmic microwave background polarization experiment on the South Pole telescope”. In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*. Ed. by Wayne S. Holland and Jonas Zmuidzinas. Vol. 9153. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. July 2014, 91531P, 91531P. arXiv: 1407.2973 [astro-ph.IM].

[85] QUIET Collaboration et al. “Second Season QUIET Observations: Measurements of the Cosmic Microwave Background Polarization Power Spectrum at 95 GHz”. In: *Astrophysical Journal* 760.2, 145 (Dec. 2012), p. 145. arXiv: 1207.5034 [astro-ph.CO].

[86] Polarbear Collaboration et al. “A Measurement of the Cosmic Microwave Background B-mode Polarization Power Spectrum at Sub-degree Scales with POLARBEAR”. In: *Astrophysical Journal* 794.2, 171 (Oct. 2014), p. 171. arXiv: 1403.2369 [astro-ph.CO].

[87] P. Ade et al. “First Season QUaD CMB Temperature and Polarization Power Spectra”. In: *Astrophysical Journal* 674.1 (Feb. 2008), pp. 22–28. arXiv: 0705.2359 [astro-ph].

[88] D. Barkats et al. “Degree-scale Cosmic Microwave Background Polarization Measurements from Three Years of BICEP1 Data”. In: *Astrophysical Journal* 783.2, 67 (Mar. 2014), p. 67. arXiv: 1310.1422 [astro-ph.CO].

[89] BICEP2 Collaboration et al. “Improved Constraints on Cosmology and Foregrounds from BICEP2 and Keck Array Cosmic Microwave Background Data with Inclusion of 95 GHz Band”. In: *Physical Review Letters* 116.3, 031302 (Jan. 2016), p. 031302. arXiv: 1510.09217 [astro-ph.CO].

[90] P. de Bernardis et al. “A flat Universe from high-resolution maps of the cosmic microwave background radiation”. In: *Nature* 404.6781 (Apr. 2000), pp. 955–959. arXiv: astro-ph/0004404 [astro-ph].

[91] Nick Hand et al. “Evidence of Galaxy Cluster Motions with the Kinematic Sunyaev-Zel’dovich Effect”. In: *Physical Review Letters* 109.4, 041101 (July 2012), p. 041101. arXiv: 1203.4219 [astro-ph.CO].

[92] Sigurd Naess et al. “The Atacama Cosmology Telescope: CMB polarization at 200 $< \ell <$ 9000”. In: *Journal of Cosmology and Astroparticle Physics* 2014.10 (Oct. 2014), pp. 007–007. arXiv: 1405.5524 [astro-ph.CO].

[93] B. J. Koopman et al. “Advanced ACTPol Low-Frequency Array: Readout and Characterization of Prototype 27 and 39 GHz Transition Edge Sensors”. In: *Journal of Low Temperature Physics* 193.5-6 (Dec. 2018), pp. 1103–1111. arXiv: 1711.02594 [astro-ph.IM].

[94] E. J. Wollack et al. “An Instrument for Investigation of the Cosmic Microwave Background Radiation at Intermediate Angular Scales”. In: *Astrophysical Journal* 476.2 (Feb. 1997), pp. 440–457. arXiv: astro-ph/9601196 [astro-ph].

[95] A. C. S. Readhead et al. “Polarization Observations with the Cosmic Background Imager”. In: *Science* 306.5697 (Oct. 2004), pp. 836–844. arXiv: astro-ph/0409569 [astro-ph].

[96] A. Kogut et al. “First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Temperature-Polarization Correlation”. In: *Astrophysical Journal, Supplement* 148.1 (Sept. 2003), pp. 161–173. arXiv: astro-ph/0302213 [astro-ph].

[97] K. T. Story et al. “A Measurement of the Cosmic Microwave Background Gravitational Lensing Potential from 100 Square Degrees of SPTpol Data”. In: *Astrophysical Journal* 810.1, 50 (Sept. 2015), p. 50. arXiv: 1412.4760 [astro-ph.CO].

[98] BICEP2 Collaboration et al. “Detection of B-Mode Polarization at Degree Angular Scales by BICEP2”. In: *Physical Review Letters* 112.24, 241101 (June 2014), p. 241101. arXiv: 1403.3985 [astro-ph.CO].

[99] BICEP2 and Keck Array Collaborations et al. “BICEP2/Keck Array V: Measurements of B-mode Polarization at Degree Angular Scales and 150 GHz by the Keck Array”. In: *Astrophysical Journal* 811.2, 126 (Oct. 2015), p. 126. arXiv: 1502.00643 [astro-ph.CO].

[100] R. Keisler et al. “Measurements of Sub-degree B-mode Polarization in the Cosmic Microwave Background from 100 Square Degrees of SPTpol Data”. In: *Astrophysical Journal* 807.2, 151 (July 2015), p. 151. arXiv: 1503.02315 [astro-ph.CO].

[101] BICEP2/Keck Collaboration et al. “Joint Analysis of BICEP2/Keck Array and Planck Data”. In: *Physical Review Letters* 114.10, 101301 (Mar. 2015), p. 101301. arXiv: 1502.00612 [astro-ph.CO].

[102] BICEP2 Collaboration et al. “Constraints on Primordial Gravitational Waves Using Planck, WMAP, and New BICEP2/Keck Observations through the 2015 Season”. In: *Physical Review Letters* 121.22, 221301 (Nov. 2018), p. 221301. arXiv: 1810.05216 [astro-ph.CO].

[103] Thibaut Louis et al. "The Atacama Cosmology Telescope: two-season ACTPol spectra and parameters". In: *Journal of Cosmology and Astroparticle Physics* 2017.6, 031 (June 2017), p. 031. arXiv: 1610.02360 [astro-ph.CO].

[104] P. A. R. Ade et al. "A Constraint on Primordial B-modes from the First Flight of the SPIDER Balloon-borne Telescope". In: *Astrophysical Journal* 927.2, 174 (Mar. 2022), p. 174.

[105] Engineering National Academies of Sciences and Medicine. *Pathways to Discovery in Astronomy and Astrophysics for the 2020s*. 2021.

[106] Howard Hui et al. "BICEP Array: a multi-frequency degree-scale CMB polarimeter". In: *Proc. SPIE Int. Soc. Opt. Eng.* 10708 (2018). arXiv: 1808.00568 [astro-ph.IM].

[107] M. Hazumi et al. "LiteBIRD: A Satellite for the Studies of B-Mode Polarization and Inflation from Cosmic Background Radiation Detection". In: *Journal of Low Temperature Physics* 194.5-6 (Mar. 2019), pp. 443–452.

[108] Peter Ade et al. "The Simons Observatory: science goals and forecasts". In: *Journal of Cosmology and Astroparticle Physics* 2019.2, 056 (Feb. 2019), p. 056. arXiv: 1808.07445 [astro-ph.CO].

[109] Toshiya Namikawa et al. "Simons Observatory: Constraining inflationary gravitational waves with multitracer B -mode delensing". In: *Physical Review D* 105.2, 023511 (Jan. 2022), p. 023511. arXiv: 2110.09730 [astro-ph.CO].

[110] K. M. Górski et al. "HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere". In: *Astrophysical Journal* 622.2 (Apr. 2005), pp. 759–771. arXiv: astro-ph/0409513 [astro-ph].

[111] Max Tegmark. "How to Make Maps from Cosmic Microwave Background Data without Losing Information". In: *Astrophysical Journal, Letters* 480.2 (May 1997), pp. L87–L90. arXiv: astro-ph/9611130 [astro-ph].

[112] K. K. Schaffer et al. "The First Public Release of South Pole Telescope Data: Maps of a 95 deg^2 Field from 2008 Observations". In: *Astrophysical Journal* 743.1, 90 (Dec. 2011), p. 90. arXiv: 1111.7245 [astro-ph.CO].

[113] E. Keihänen et al. "Making cosmic microwave background temperature and polarization maps with MADAM". In: *Astronomy and Astrophysics* 510, A57 (Feb. 2010), A57. arXiv: 0907.0367 [astro-ph.CO].

[114] D. Sutton et al. "Fast and precise map-making for massively multi-detector CMB experiments". In: *Monthly Notices of the RAS* 407.3 (Sept. 2010), pp. 1387–1402. arXiv: 0912.2738 [astro-ph.CO].

[115] G. de Gasperis et al. "ROMA: A map-making algorithm for polarised CMB data sets". In: *Astronomy and Astrophysics* 436.3 (June 2005), pp. 1159–1165. arXiv: astro-ph/0502142 [astro-ph].

[116] C. M. Cantalupo et al. "MADmap: A Massively Parallel Maximum Likelihood Cosmic Microwave Background Map-maker". In: *Astrophysical Journal, Supplement* 187.1 (Mar. 2010), pp. 212–227. arXiv: 0906.1775 [astro-ph.CO].

[117] A. Traficante et al. "Data reduction pipeline for the Hi-GAL survey". In: *Monthly Notices of the RAS* 416.4 (Oct. 2011), pp. 2932–2943. arXiv: 1106.0698 [astro-ph.IM].

[118] Rolando Dünner et al. "The Atacama Cosmology Telescope: Data Characterization and Mapmaking". In: *Astrophysical Journal* 762.1, 10 (Jan. 2013), p. 10. arXiv: 1208.0050 [astro-ph.IM].

[119] G. Patanchon et al. "SANEPIC: A Mapmaking Method for Time Stream Data from Large Arrays". In: *The Astrophysical Journal* 681.1 (July 2008), pp. 708–725. URL: https://doi.org/10.1086/588543.

[120] J. Papež, L. Grigori, and R. Stompor. "Solving linear equations with messenger-field and conjugate gradient techniques: An application to CMB data analysis". In: *Astronomy and Astrophysics* 620, A59 (Nov. 2018), A59. arXiv: 1803.03462 [astro-ph.CO].

[121] Samira Hamimeche and Antony Lewis. "Likelihood analysis of CMB temperature and polarization power spectra". In: *Physical Review D* 77.10, 103013 (May 2008), p. 103013. arXiv: 0801.0554 [astro-ph].

[122] Daniel Foreman-Mackey et al. "emcee: The MCMC Hammer". In: *Publications of the ASP* 125.925 (Mar. 2013), p. 306. arXiv: 1202.3665 [astro-ph.IM].

[123] R. M. A. Azzam, N. M. Bashara, and Stanley S. Ballard. "Ellipsometry and Polarized Light". In: *Physics Today* 31.11 (Jan. 1978), p. 72.

[124] Sean A. Bryan, Thomas E. Montroy, and John E. Ruhl. "Modeling dielectric half-wave plates for cosmic microwave background polarimetry using a Mueller matrix formalism". In: *Applied Optics* 49.32 (Nov. 2010), p. 6313. arXiv: 1006.3359 [astro-ph.IM].

[125] Daniel O'Dea, Anthony Challinor, and Bradley R. Johnson. "Systematic errors in cosmic microwave background polarization measurements". In: *Monthly Notices of the RAS* 376.4 (Apr. 2007), pp. 1767–1783. arXiv: astro-ph/0610361 [astro-ph].

[126] T. Essinger-Hileman et al. "Systematic effects from an ambient-temperature, continuously rotating half-wave plate". In: *Review of Scientific Instruments* 87.9, 094503 (Sept. 2016), p. 094503. arXiv: 1601.05901 [astro-ph.IM].

[127] Maria Salatino et al. "Studies of systematic uncertainties for Simons Observatory: polarization modulator related effects". In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*. Ed. by Jonas Zmuidzinas and Jian-Rong Gao. Vol. 10708. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. July 2018, 1070848, p. 1070848. arXiv: 1808.07442 [astro-ph.IM].

[128] Kunimoto Komatsu et al. "Demonstration of the broadband half-wave plate using the nine-layer sapphire for the cosmic microwave background polarization experiment". In: *Journal of Astronomical Telescopes, Instruments, and Systems* 5, 044008 (Oct. 2019), p. 044008. arXiv: 1905.13520 [astro-ph.IM].

[129] Adriaan J. Duivenvoorden et al. "Probing frequency-dependent half-wave plate systematics for CMB experiments with full-sky beam convolution simulations". In: *Monthly Notices of the RAS* 502.3 (Apr. 2021), pp. 4526–4539. arXiv: 2012.10437 [astro-ph.CO].

[130] S. Giardiello et al. "Detailed study of HWP non-idealities and their impact on future measurements of CMB polarization anisotropies from space". In: *Astronomy and Astrophysics* 658, A15 (Feb. 2022), A15. arXiv: 2106.08031 [astro-ph.CO].

[131] Maximilian H. Abitbol et al. "The Simons Observatory: gain, bandpass and polarization-angle calibration requirements for B-mode searches". In: *Journal of Cosmology and Astroparticle Physics* 2021.5, 032 (May 2021), p. 032. arXiv: 2011.02449 [astro-ph.CO].

[132] A. P. Jones. "Depletion patterns and dust evolution in the interstellar medium". In: *Journal of Geophysics Research* 105.A5 (May 2000), pp. 10257–10268. arXiv: astro-ph/9907066 [astro-ph].

[133] B. T. Draine. "Interstellar Dust Models: Extinction, Absorption and Emission". In: *EAS Publications Series*. Ed. by F. Boulanger et al. Vol. 35. EAS Publications Series. Jan. 2009, pp. 245–268.

[134] J. S. Mathis, W. Rumpl, and K. H. Nordsieck. "The size distribution of interstellar grains." In: *Astrophysical Journal* 217 (Oct. 1977), pp. 425–433.

[135] Jr. Davis Leverett and Jesse L. Greenstein. "The Polarization of Starlight by Aligned Dust Grains." In: *Astrophysical Journal* 114 (Sept. 1951), p. 206.

[136] Planck Collaboration et al. "Planck intermediate results. XIX. An overview of the polarized thermal emission from Galactic dust". In: *Astronomy and Astrophysics* 576, A104 (Apr. 2015), A104. arXiv: 1405.0871 [astro-ph.GA].

[137] Roger H. Hildebrand et al. "Dispersion of Magnetic Fields in Molecular Clouds. I". In: *Astrophysical Journal* 696.1 (May 2009), pp. 567–573. arXiv: 0811.0813 [astro-ph].

[138] Planck Collaboration et al. "Planck intermediate results. XXII. Frequency dependence of thermal emission from Galactic dust in intensity and polarization". In: *Astronomy and Astrophysics* 576, A107 (Apr. 2015), A107. arXiv: 1405.0874 [astro-ph.GA].

[139] Planck Collaboration et al. "Planck 2018 results. XI. Polarized dust foregrounds". In: *arXiv e-prints*, arXiv:1801.04945 (Jan. 2018), arXiv:1801.04945. arXiv: 1801.04945 [astro-ph.GA].

[140] Douglas P. Finkbeiner, Marc Davis, and David J. Schlegel. "Extrapolation of Galactic Dust Emission at 100 Microns to Cosmic Microwave Background Radiation Frequencies Using FIRAS". In: *Astrophysical Journal* 524.2 (Oct. 1999), pp. 867–886. arXiv: astro-ph/9905128 [astro-ph].

[141] Aaron M. Meisner and Douglas P. Finkbeiner. "Modeling Thermal Dust Emission with Two Components: Application to the Planck High Frequency Instrument Maps". In: *Astrophysical Journal* 798.2, 88 (Jan. 2015), p. 88. arXiv: 1410.7523 [astro-ph.GA].

[142] B. T. Draine and Brandon S. Hensley. "On the Shapes of Interstellar Grains: Modeling Infrared Extinction and Polarization by Spheroids and Continuous Distributions of Ellipsoids". In: *Astrophysical Journal* 910.1, 47 (Mar. 2021), p. 47. arXiv: 1710.08968 [astro-ph.GA].

[143] F. Vansyngel et al. "Statistical simulations of the dust foreground to cosmic microwave background polarization". In: *Astronomy and Astrophysics* 603, A62 (July 2017), A62. arXiv: 1611.02577 [astro-ph.CO].

[144] Christopher Sheehy and Anže Slosar. “No evidence for dust B -mode decorrelation in Planck data”. In: *Physical Review D* 97.4, 043522 (Feb. 2018), p. 043522. arXiv: 1709.09729 [astro-ph.CO].

[145] P. Biermann. “On the radio continuum flux from the disks of spiral galaxies.” In: *Astronomy and Astrophysics* 53.2 (Dec. 1976), pp. 295–303.

[146] U. Lisenfeld and H. J. Völk. “On the radio spectral index of galaxies”. In: *Astronomy and Astrophysics* 354 (Feb. 2000), pp. 423–430. arXiv: astro-ph/9912232 [astro-ph].

[147] Planck Collaboration et al. “Planck 2015 results. XXV. Diffuse low-frequency Galactic foregrounds”. In: *Astronomy and Astrophysics* 594, A25 (Sept. 2016), A25. arXiv: 1506.06660 [astro-ph.GA].

[148] N. Krachmalnicoff et al. “S-PASS view of polarized Galactic synchrotron at 2.3 GHz as a contaminant to CMB observations”. In: *Astronomy and Astrophysics* 618, A166 (Oct. 2018), A166. arXiv: 1802.01145 [astro-ph.GA].

[149] George B. Rybicki and Alan P. Lightman. *Radiative Processes in Astrophysics.* 1986.

[150] A. Waelkens et al. “Simulating polarized Galactic synchrotron emission at all frequencies. The Hammurabi code”. In: *Astronomy and Astrophysics* 495.2 (Feb. 2009), pp. 697–706. arXiv: 0807.2262 [astro-ph].

[151] David Alonso, Pedro G. Ferreira, and Mario G. Santos. “Fast simulations for intensity mapping experiments”. In: *Monthly Notices of the RAS* 444.4 (Nov. 2014), pp. 3183–3197. arXiv: 1405.1751 [astro-ph.CO].

[152] Jacques Delabrouille and J. -F. Cardoso. “Diffuse source separation in CMB observations”. In: *International Summer School on Data Analysis in Cosmology.* Feb. 2007. arXiv: astro-ph/0702198.

[153] S. M. Leach et al. “Component separation methods for the PLANCK mission”. In: *Astronomy and Astrophysics* 491.2 (Nov. 2008), pp. 597–615. arXiv: 0805.0269 [astro-ph].

[154] P. Natoli et al. “Exploring cosmic origins with CORE: Mitigation of systematic effects”. In: *Journal of Cosmology and Astroparticle Physics* 2018.4, 022 (Apr. 2018), p. 022. arXiv: 1707.04224 [astro-ph.CO].

[155] Lloyd Knox and Yong-Seon Song. “Limit on the Detectability of the Energy Scale of Inflation”. In: *Physical Review Letters* 89.1, 011303 (July 2002), p. 011303. arXiv: astro-ph/0202286 [astro-ph].

[156] M. Remazeilles et al. “Sensitivity and foreground modelling for large-scale cosmic microwave background B-mode polarization satellite missions”. In: *Monthly Notices of the RAS* 458.2 (May 2016), pp. 2032–2050. arXiv: 1509.04714 [astro-ph.CO].

[157] Josquin Errard et al. “Robust forecasts on fundamental physics from the foreground-obscured, gravitationally-lensed CMB polarization”. In: *Journal of Cosmology and Astroparticle Physics* 2016.3, 052 (Mar. 2016), p. 052. arXiv: 1509.06770 [astro-ph.CO].

[158] David Alonso et al. "Simulated forecasts for primordial B -mode searches in ground-based experiments". In: *Physical Review D* 95.4, 043504 (Feb. 2017), p. 043504. arXiv: 1608.00551 [astro-ph.CO].

[159] S. Adachi et al. "A Measurement of the Degree Scale CMB B-mode Angular Power Spectrum with POLARBEAR". In: *arXiv e-prints*, arXiv:1910.02608 (Oct. 2019), arXiv:1910.02608. arXiv: 1910.02608 [astro-ph.CO].

[160] H. K. Eriksen et al. "Power Spectrum Estimation from High-Resolution Maps by Gibbs Sampling". In: *Astrophysical Journal, Supplement* 155.2 (Dec. 2004), pp. 227–241. arXiv: astro-ph/0407028 [astro-ph].

[161] H. K. Eriksen et al. "Joint Bayesian Component Separation and CMB Power Spectrum Estimation". In: *Astrophysical Journal* 676.1 (Mar. 2008), pp. 10–32. arXiv: 0709.1058 [astro-ph].

[162] J. Dunkley et al. "Five-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Bayesian Estimation of Cosmic Microwave Background Polarization Maps". In: *Astrophysical Journal* 701.2 (Aug. 2009), pp. 1804–1813. arXiv: 0811.4280 [astro-ph].

[163] Mathieu Remazeilles, Jacques Delabrouille, and Jean-François Cardoso. "Foreground component separation with generalized Internal Linear Combination". In: *Monthly Notices of the RAS* 418.1 (Nov. 2011), pp. 467–476. arXiv: 1103.1166 [astro-ph.CO].

[164] Planck Collaboration et al. "Planck 2018 results. IV. Diffuse component separation". In: *Astronomy and Astrophysics* 641, A4 (Sept. 2020), A4. arXiv: 1807.06208 [astro-ph.CO].

[165] Mathieu Remazeilles, Aditya Rotti, and Jens Chluba. "Peeling off foregrounds with the constrained moment ILC method to unveil primordial CMB B-modes". In: *arXiv e-prints*, arXiv:2006.08628 (June 2020), arXiv:2006.08628. arXiv: 2006.08628 [astro-ph.CO].

[166] Radek Stompor et al. "Maximum likelihood algorithm for parametric component separation in cosmic microwave background experiments". In: *Monthly Notices of the RAS* 392.1 (Jan. 2009), pp. 216–232. arXiv: 0804.2645 [astro-ph].

[167] Radek Stompor, Josquin Errard, and Davide Poletti. "Forecasting performance of CMB experiments in the presence of complex foreground contaminations". In: *Physical Review D* 94.8, 083526 (Oct. 2016), p. 083526. arXiv: 1609.03807 [astro-ph.CO].

[168] R. D. P. Grumitt, Luke R. P. Jew, and C. Dickinson. "Hierarchical Bayesian CMB component separation with the No-U-Turn Sampler". In: *Monthly Notices of the RAS* 496.4 (June 2020), pp. 4383–4401. arXiv: 1910.14170 [astro-ph.CO].

[169] J. Dunkley et al. "The Atacama Cosmology Telescope: likelihood for small-scale CMB data". In: *Journal of Cosmology and Astroparticle Physics* 2013.7, 025 (July 2013), p. 025. arXiv: 1301.0776 [astro-ph.CO].

[170] Steve K. Choi et al. "The Atacama Cosmology Telescope: A Measurement of the Cosmic Microwave Background Power Spectra at 98 and 150 GHz". In: *arXiv e-prints*, arXiv:2007.07289 (July 2020), arXiv:2007.07289. arXiv: 2007.07289 [astro-ph.CO].

[171] Planck Collaboration et al. "Planck intermediate results. L. Evidence of spatial variation of the polarized thermal dust spectral energy distribution and implications for CMB B-mode analysis". In: *Astronomy and Astrophysics* 599, A51 (Mar. 2017), A51. arXiv: `1606.07335 [astro-ph.CO]`.

[172] Max Tegmark. "Removing Real-World Foregrounds from Cosmic Microwave Background Maps". In: *Astrophysical Journal* 502.1 (July 1998), pp. 1–6. arXiv: `astro-ph/9712038 [astro-ph]`.

[173] Matias Zaldarriaga, Steven R. Furlanetto, and Lars Hernquist. "21 Centimeter Fluctuations from Cosmic Gas at High Redshifts". In: *Astrophysical Journal* 608.2 (June 2004), pp. 622–635. arXiv: `astro-ph/0311514 [astro-ph]`.

[174] Mário G. Santos, Asantha Cooray, and Lloyd Knox. "Multifrequency Analysis of 21 Centimeter Fluctuations from the Era of Reionization". In: *Astrophysical Journal* 625.2 (June 2005), pp. 575–587. arXiv: `astro-ph/0408515 [astro-ph]`.

[175] Jens Chluba, James Colin Hill, and Maximilian H. Abitbol. "Rethinking CMB foregrounds: systematic extension of foreground parametrizations". In: *Monthly Notices of the RAS* 472.1 (Nov. 2017), pp. 1195–1213. arXiv: `1701.00274 [astro-ph.CO]`.

[176] Anna Mangilli et al. "Dust moments: towards a new modelling of the galactic dust emission for CMB B-modes analysis". In: *arXiv e-prints*, arXiv:1912.09567 (Dec. 2019), arXiv:1912.09567. arXiv: `1912.09567 [astro-ph.CO]`.

[177] D. J. Fixsen. "The Temperature of the Cosmic Microwave Background". In: *Astrophysical Journal* 707.2 (Dec. 2009), pp. 916–920. arXiv: `0911.1955 [astro-ph.CO]`.

[178] B. T. Draine and Brandon Hensley. "Magnetic Nanoparticles in the Interstellar Medium: Emission Spectrum and Polarization". In: *Astrophysical Journal* 765.2, 159 (Mar. 2013), p. 159. arXiv: `1205.7021 [astro-ph.GA]`.

[179] Brandon S. Hensley and B. T. Draine. "Thermodynamics and Charging of Interstellar Iron Nanoparticles". In: *Astrophysical Journal* 834.2, 134 (Jan. 2017), p. 134. arXiv: `1611.08607 [astro-ph.GA]`.

[180] Planck Collaboration et al. "Planck 2013 results. XV. CMB power spectra and likelihood". In: *Astronomy and Astrophysics* 571, A15 (Nov. 2014), A15. arXiv: `1303.5075 [astro-ph.CO]`.

[181] BICEP2 Collaboration et al. "BICEP2/Keck Array. VII. Matrix Based E/B Separation Applied to Bicep2 and the Keck Array". In: *Astrophysical Journal* 825.1, 66 (July 2016), p. 66. arXiv: `1603.05976 [astro-ph.IM]`.

[182] Nobuhiko Katayama and Eiichiro Komatsu. "Simple Foreground Cleaning Algorithm for Detecting Primordial B-mode Polarization of the Cosmic Microwave Background". In: *Astrophysical Journal* 737.2, 78 (Aug. 2011), p. 78. arXiv: `1101.5210 [astro-ph.CO]`.

[183] Kiyotomo Ichiki et al. "Delta-map method of removing CMB foregrounds with spatially varying spectra". In: *Progress of Theoretical and Experimental Physics* 2019.3, 033E01 (Mar. 2019), 033E01. arXiv: `1811.03886 [astro-ph.CO]`.

[184] Steven G. Murray and Francis J. Poulin. "hankel: A Python library for performing simple and accurate Hankel transformations". In: *Journal of Open Source Software* 4.37 (2019), p. 1397. URL: https://doi.org/10.21105/joss.01397.

[185] Oliver H. E. Philcox, Blake D. Sherwin, and Alexand er van Engelen. "Detection and removal of B-mode dust foregrounds with signatures of statistical anisotropy". In: *Monthly Notices of the RAS* 479.4 (Oct. 2018), pp. 5577–5595. arXiv: 1805.09177 [astro-ph.CO].

[186] Michael JD Powell. "An efficient method for finding the minimum of a function of several variables without calculating derivatives". In: *The computer journal* 7.2 (1964), pp. 155–162.

[187] Pauli Virtanen et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". In: *Nature Methods* 17 (2020), pp. 261–272.

[188] Ben Thorne et al. "The Python Sky Model: software for simulating the Galactic microwave sky". In: *Monthly Notices of the Royal Astronomical Society* 469.3 (2017), pp. 2821–2833.

[189] CGT Haslam et al. "A 408 MHz all-sky continuum survey. I-Observations at southern declinations and for the North Polar region". In: *Astronomy and Astrophysics* 100 (1981), pp. 209–219.

[190] CGT Haslam et al. "A 408 MHz all-sky continuum survey. II-The atlas of contour maps". In: *Astronomy and Astrophysics Supplement Series* 47 (1982), p. 1.

[191] J Dunkley et al. "Prospects for polarized foreground removal". In: *AIP Conference Proceedings*. Vol. 1141. 1. American Institute of Physics. 2009, pp. 222–264.

[192] C. L. Bennett et al. "Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results". In: *Astrophysical Journal, Supplement* 208.2, 20 (Oct. 2013), p. 20. arXiv: 1212.5225 [astro-ph.CO].

[193] Planck Collaboration et al. "Planck 2015 results. X. Diffuse component separation: Foreground maps". In: *Astronomy and Astrophysics* 594, A10 (Sept. 2016), A10. arXiv: 1502.01588 [astro-ph.CO].

[194] Kevork Abazajian et al. "CMB-S4 Science Case, Reference Design, and Project Plan". In: *arXiv e-prints*, arXiv:1907.04473 (July 2019), arXiv:1907.04473. arXiv: 1907.04473 [astro-ph.IM].

[195] The CMB-S4 Collaboration et al. "CMB-S4: Forecasting Constraints on Primordial Gravitational Waves". In: *arXiv e-prints*, arXiv:2008.12619 (Aug. 2020), arXiv:2008.12619. arXiv: 2008.12619 [astro-ph.CO].

[196] Antony Lewis, Anthony Challinor, and Neil Turok. "Analysis of CMB polarization on an incomplete sky". In: *Physical Review D* 65.2, 023505 (Jan. 2002), p. 023505. arXiv: astro-ph/0106536 [astro-ph].

[197] Emory F. Bunn. "Efficient decomposition of cosmic microwave background polarization maps into pure E, pure B, and ambiguous components". In: *Physical Review D* 83.8, 083003 (Apr. 2011), p. 083003. arXiv: 1008.0827 [astro-ph.CO].

[198] Planck Collaboration et al. "Planck 2015 results. IX. Diffuse component separation: CMB maps". In: *Astronomy and Astrophysics* 594, A9 (Sept. 2016), A9. arXiv: 1502.05956 [astro-ph.CO].

[199] Masashi Hazumi et al. “Litebird: A satellite for the studies of b-mode polarization and inflation from cosmic background radiation detection”. In: *Journal of Low Temperature Physics* 194.5-6 (2019), pp. 443–452.

[200] H. Sugai et al. “Updated Design of the CMB Polarization Experiment Satellite LiteBIRD”. In: *Journal of Low Temperature Physics* 199.3-4 (Jan. 2020), pp. 1107–1117. arXiv: 2001.01724 [astro-ph.IM].

[201] Planck Collaboration et al. “Planck 2015 results. I. Overview of products and scientific results”. In: *Astronomy and Astrophysics* 594, A1 (Sept. 2016), A1. arXiv: 1502.01582 [astro-ph.CO].

[202] G. Puglisi et al. “Forecasting the Contribution of Polarized Extragalactic Radio Sources in CMB Observations”. In: *Astrophysical Journal* 858.2, 85 (May 2018), p. 85. arXiv: 1712.09639 [astro-ph.CO].

[203] Maximilian H. Abitbol et al. “The Simons Observatory: Bandpass and polarization-angle calibration requirements for B-mode searches”. In: *arXiv e-prints*, arXiv:2011.02449 (Nov. 2020), arXiv:2011.02449. arXiv: 2011.02449 [astro-ph.CO].

[204] F. Feroz, M. P. Hobson, and M. Bridges. “MULTINEST: an efficient and robust Bayesian inference tool for cosmology and particle physics”. In: *Monthly Notices of the RAS* 398.4 (Oct. 2009), pp. 1601–1614. arXiv: 0809.3437 [astro-ph].

[205] T. Matsumura et al. “Mission Design of LiteBIRD”. In: *Journal of Low Temperature Physics* 176.5-6 (Sept. 2014), pp. 733–740. arXiv: 1311.2847 [astro-ph.IM].

[206] J. Delabrouille et al. “A full sky, low foreground, high resolution CMB map from WMAP”. In: *Astronomy and Astrophysics* 493.3 (Jan. 2009), pp. 835–857. arXiv: 0807.0773 [astro-ph].

[207] Soumen Basak and Jacques Delabrouille. “A needlet internal linear combination analysis of WMAP 7-year data: estimation of CMB temperature map and power spectrum”. In: *Monthly Notices of the RAS* 419.2 (Jan. 2012), pp. 1163–1175. arXiv: 1106.5383 [astro-ph.CO].

[208] Soumen Basak and Jacques Delabrouille. “A needlet ILC analysis of WMAP 9-year polarization data: CMB polarization power spectra”. In: *Monthly Notices of the RAS* 435.1 (Oct. 2013), pp. 18–29. arXiv: 1204.0292 [astro-ph.CO].

[209] Davide Poletti and Josquin Errard. In: *In prep.* (2023).

[210] C. L. Bennett et al. “First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Foreground Emission”. In: *Astrophysical Journal, Supplement* 148.1 (Sept. 2003), pp. 97–117. arXiv: astro-ph/0302208 [astro-ph].

[211] Josquin Errard and Radek Stompor. “Characterizing bias on large scale CMB B -modes after Galactic foregrounds cleaning”. In: *Physical Review D* 99.4, 043529 (Feb. 2019), p. 043529. arXiv: 1811.00479 [astro-ph.CO].

[212] LiteBIRD Collaboration. “Probing Cosmic Inflation with the LiteBIRD Cosmic Microwave Background Polarization Survey”. In: *Progress of Theoretical and Experimental Physics* (Nov. 2022). ptac150. URL: https://doi.org/10.1093/ptep/ptac150.

[213] H. Akaike. "A new look at the statistical model identification". In: *IEEE Transactions on Automatic Control* 19.6 (1974), pp. 716–723.

[214] M. Maltoni and T. Schwetz. "Testing the statistical compatibility of independent data sets". In: *Phys. Rev. D* 68 (2003), p. 033020. arXiv: `hep-ph/0304176`.

[215] E.J. Wagenmakers and S. Farrell. "AIC model selection using Akaike weights". In: *Psychon Bull Rev.* 11.1 (2004), pp. 192–6.

[216] P. A. R. Ade et al. "Improved Constraints on Primordial Gravitational Waves using Planck, WMAP, and BICEP/Keck Observations through the 2018 Observing Season". In: *Physical Review Letters* 127.15, 151301 (Oct. 2021), p. 151301. arXiv: `2110.00483 [astro-ph.CO]`.

[217] Mathieu Remazeilles, Aditya Rotti, and Jens Chluba. "Peeling off foregrounds with the constrained moment ILC method to unveil primordial CMB B modes". In: *Monthly Notices of the RAS* 503.2 (May 2021), pp. 2478–2498. arXiv: `2006.08628 [astro-ph.CO]`.

[218] Brandon S. Hensley and Philip Bull. "Mitigating Complex Dust Foregrounds in Future Cosmic Microwave Background Polarization Experiments". In: *Astrophysical Journal* 853.2, 127 (Feb. 2018), p. 127. arXiv: `1709.07897 [astro-ph.CO]`.

[219] Lisa McBride, Philip Bull, and Brandon S. Hensley. "Characterizing line-of-sight variability of polarized dust emission with future CMB experiments". In: *arXiv e-prints*, arXiv:2207.14213 (July 2022), arXiv:2207.14213. arXiv: `2207.14213 [astro-ph.CO]`.

[220] Léo Vacher et al. "High precision modeling of polarized signals: moment expansion method generalized to spin-2 fields". In: *arXiv e-prints*, arXiv:2205.01049 (May 2022), arXiv:2205.01049. arXiv: `2205.01049 [astro-ph.CO]`.

[221] Niall Jeffrey et al. "Single frequency CMB B-mode inference with realistic foregrounds from a single training image". In: *Monthly Notices of the RAS* 510.1 (Feb. 2022), pp. L1–L6. arXiv: `2111.01138 [astro-ph.CO]`.

[222] Josquin Errard and Radek Stompor. "Characterizing bias on large scale CMB B -modes after Galactic foregrounds cleaning". In: *Physical Review D* 99.4, 043529 (Feb. 2019), p. 043529. arXiv: `1811.00479 [astro-ph.CO]`.

[223] Giuseppe Puglisi et al. "Improved galactic foreground removal for B-mode detection with clustering methods". In: *Monthly Notices of the RAS* 511.2 (Apr. 2022), pp. 2052–2074. arXiv: `2109.11562 [astro-ph.CO]`.

[224] F. Stivoli et al. "Maximum likelihood, parametric component separation and CMB B-mode detection in suborbital experiments". In: *Monthly Notices of the RAS* 408.4 (Nov. 2010), pp. 2319–2335. arXiv: `1004.4756 [astro-ph.CO]`.

[225] J. Errard, F. Stivoli, and R. Stompor. "Publisher's Note: Framework for performance forecasting and optimization of CMB B-mode observations in the presence of astrophysical foregrounds [Phys. Rev. D 84, 063005 (2011)]". In: *Physical Review D* 84.6, 069907 (Sept. 2011), p. 069907. arXiv: `1105.3859 [astro-ph.CO]`.

[226] Lloyd Knox. “Cosmic Microwave Background Anisotropy Observing Strategy Assessment”. In: *Astrophysical Journal* 480.1 (May 1997), pp. 72–78. arXiv: astro-ph/9606066 [astro-ph].

[227] Planck Collaboration et al. “Planck intermediate results. LVII. Joint Planck LFI and HFI data processing”. In: *Astronomy and Astrophysics* 643, A42 (Nov. 2020), A42. arXiv: 2007.04997 [astro-ph.CO].

[228] LiteBIRD Collaboration et al. “Probing Cosmic Inflation with the LiteBIRD Cosmic Microwave Background Polarization Survey”. In: *arXiv e-prints*, arXiv:2202.02773 (Feb. 2022), arXiv:2202.02773. arXiv: 2202.02773 [astro-ph.IM].

[229] M. -A. Bigot-Sazy et al. “Self-calibration: an efficient method to control systematic effects in bolometric interferometry”. In: *Astronomy and Astrophysics* 550, A59 (Feb. 2013), A59. arXiv: 1209.4905 [astro-ph.IM].

[230] E. Keihänen and M. Reinecke. “ArtDeco: a beam-deconvolution code for absolute cosmic microwave background measurements”. In: *Astronomy and Astrophysics* 548, A110 (Dec. 2012), A110. arXiv: 1208.1399 [astro-ph.IM].

[231] Planck Collaboration et al. “Planck 2018 results. III. High Frequency Instrument data processing and frequency maps”. In: *Astronomy and Astrophysics* 641, A3 (Sept. 2020), A3. arXiv: 1807.06207 [astro-ph.CO].

[232] Planck Collaboration et al. “Planck intermediate results. XLVI. Reduction of large-scale systematic effects in HFI polarization maps and estimation of the reionization optical depth”. In: *Astronomy and Astrophysics* 596, A107 (Dec. 2016), A107. arXiv: 1605.02985 [astro-ph.CO].

[233] J. -M. Delouis et al. “SRoll2: an improved mapmaking approach to reduce large-scale systematic effects in the Planck High Frequency Instrument legacy maps”. In: *Astronomy and Astrophysics* 629, A38 (Sept. 2019), A38. arXiv: 1901.11386 [astro-ph.CO].

[234] Gayoung Chon et al. “Fast estimation of polarization power spectra using correlation functions”. In: *Monthly Notices of the RAS* 350.3 (May 2004), pp. 914–926. arXiv: astro-ph/0303414 [astro-ph].