Inflation : phenomenological study and LiteBIRD space mission preparation

Gilles Weymann-Despres

université PARIS-SACLAY

Inflation: phenomenological study and LiteBIRD space mission preparation

Inflation : étude phénoménologique et préparation de la mission spatiale LiteBIRD

Thèse de doctorat de l'Université Paris-Saclay

École doctorale n° 576, Particules, hadrons, énergie et noyau : instrumentation, imagerie, cosmos et simulation (PHENIICS)
Spécialité de doctorat: Science des Astroparticules et Cosmologie
Graduate School : Physique. Référent : Faculté des sciences d'Orsay

Thèse préparée dans l'unité de recherche **IJCLab (Université Paris-Saclay, CNRS)**, sous la direction de **Sophie HENROT-VERSILLÉ**, Directeur de Recherche

Thèse soutenue à Paris-Saclay, le 13 septembre 2024, par

Gilles WEYMANN-DESPRES

Composition du jury

Membres du jury avec voix délibérative

Achille STOCCHI Professeur des universités, IJCLab, Université Paris-Saclay	Président
Anthony J. BANDAY Directeur de recherche, IRAP, Université Paul Sabatier Toulouse III	Rapporteur & Examinateur
Juan-Francisco MACIAS-PÉREZ Directeur de recherche, LPSC, Université Grenoble Alpes	Rapporteur & Examinateur
Joanna DUNKLEY Professeur, Princeton University	Examinatrice
Tomotake MATSUMURA Professeur associé, IPMU, The University of Tokyo	Examinateur
Nathalie PALANQUE-DELABROUILLE Directrice de recherche, LBNL, University of California	Examinatrice
Vivian POULIN-DÉTOLLE Chargé de recherche, LUPM, Université de Montpellier	Examinateur

THESE DE DOCTORAT
NNT : 2024UPASP074

université PARIS-SACLAY
ÉCOLE DOCTORALE
Particules, hadrons, énergie et noyau : instrumentation, imagerie, cosmos et simulation (PHENIICS)

Titre: Inflation : étude phénoménologique et préparation de la mission spatiale LiteBIRD
Mots clés: Cosmologie, Inflation, LiteBIRD, Phénoménologie, Fond diffus cosmologique, Supersymétrie

Résumé:

Cette thèse est consacrée à l'étude de l'inflation cosmologique, une phase d'expansion accélérée de l'univers primordial qui reste, à ce jour, spéculative. L'observable central de ce travail est le fond diffus cosmologique (CMB), la plus ancienne lumière encore visible aujourd'hui, dont l'étude statistique permet d'inférer des informations cruciales sur la cosmologie.

Nous entamons cette étude sur un volet expérimental, avec la préparation du satellite LiteBIRD. Au milieu de la prochaine décennie, ce dernier mesurera la polarisation du CMB à grande échelle avec une précision inédite, permettant ainsi de contraindre la présence d'ondes gravitationnelles primordiales générées durant l'inflation. Pour obtenir une telle sensibilité et éviter tout potentiel effet systématique, une maîtrise parfaite de l'instrument et de l'analyse de donnée est indispensable. Dans ce cadre, nous présentons notre première implémentation du modèle de l'instrument dans une base de données dédiée, ainsi que les outils nécessaires à la production de certains paramètres instrumentaux. Nous avons notamment produit les quaternions qui encodent les informations de pointage et d'orientation de chaque détecteur, et implémenté les faisceaux, les bandes passantes, le modèle de bruit de l'instrument et la spécification du système de lecture.

De plus, nous avons mis en place un pipeline complet pour analyser les cartes de polarisation que LiteBIRD fournira. Nous avons testé ce pipeline sur des simulations de l'instrument présentant diverses complexités. L'analyse se décompose en trois étapes. La première étape est la séparation des composantes afin de nettoyer les cartes des avant-plans. Nous avons optimisé une méthode agnostique qui ne nécessite pas de connaissances préalables sur les propriétés des avant-plans. La deuxième étape consiste à estimer les spectres à partir des cartes nettoyées et masquées, pour laquelle nous avons implémenté et testé diverses méthodes non biaisées et quasi optimales. Enfin, nous avons évalué la performance de plusieurs fonctions de vraisemblance pour inférer les paramètres cosmologiques. En plus de contraindre la présence d'ondes gravitationnelles primordiales, cette analyse permettra d'affiner notre compréhension de l'époque de la réionisation, liée au puissant rayonnement émis par la première génération d'étoiles.

La troisième partie de cette thèse se concentre sur une étude phénoménologique de l'inflation, en particulier sur un modèle d'inflation qui s'inscrit dans un cadre de physique des particules : le modèle supersymétrique minimal. En collaboration avec des cosmologistes, théoriciens et physiciens des particules, nous avons montré que les données existantes du satellite Planck sont suffisamment précises pour que les erreurs systématiques dans les prédictions du modèle dominent le budget d'erreur dans un exercice d'inférence. Ces erreurs systématiques théoriques sont dues à la non-inclusion des corrections radiatives et à une compréhension imparfaite de la fin de l'inflation. Nous avons donc incorporé les corrections nécessaires et identifié des points dans l'espace des paramètres qui satisfont à la fois les contraintes observationnelles de la physique des particules (comme la masse du Higgs et les recherches directes de SUSY au LHC) et de la cosmologie (comme la fraction de matière noire dans l'univers et les propriétés des perturbations observées par Planck). Ce travail démontre la possibilité d'unifier la description de la physique des particules et de la cosmologie dans un seul modèle cohérent, ouvrant ainsi la voie à une exploration complète de ce cadre.

université PARIS-SACLAY

ÉCOLE DOCTORALE
Particules, hadrons, énergie et noyau : instrumentation, imagerie, cosmos et simulation (PHENIICS)

Title: Inflation: phenomenological study and preparation of the LiteBIRD space mission
Keywords: Cosmology, Inflation, LiteBIRD, Phenomenology, Cosmic microwave background, Supersymmetry

Abstract: This thesis is devoted to the study of cosmological inflation, a phase of accelerated expansion in the early universe that remains speculative to this day. The central observable for this study is the cosmic microwave background (CMB), the oldest light still visible today, whose statistical study enables cosmological inference.

We first approach the study from an experimental perspective, focusing on the preparation of the LiteBIRD satellite. Set to launch in the middle of the next decade, LiteBIRD will measure the large-scale polarisation of the CMB with unprecedented precision, allowing for stringent constraints on the presence of primordial gravitational waves generated during inflation. To achieve the required sensitivity and minimise systematic effects, we must ensure precise control of both the instrument and data analysis. As part of this effort, we have implemented the instrument model in a dedicated database, along with the tools necessary to produce key instrumental parameters. This includes generating quaternions that encode each detector's pointing and orientation information, as well as implementing beam models, bandpasses, the noise model, and the specification of the readout system.

Furthermore, we have developed a complete pipeline for analysing the polarisation maps that LiteBIRD will deliver. We have tested this pipeline on realistic simulations of the instrument with various levels of complexity. The analysis pipeline consists of three stages. The first stage involves component separation to remove foreground contamination from the maps. We optimise an agnostic method that does not rely on prior knowledge of the foreground properties. The second stage focuses on estimating power spectra from the cleaned and masked maps. To this end, we have implemented and tested various unbiased and quasi-optimal methods. Finally, we assess the performance of different likelihood functions to infer cosmological parameters. In addition to constraining primordial gravitational waves, this analysis will enhance our understanding of the epoch of reionisation, which is due to the intense radiation from the first generation of stars.

In the third section of the thesis, we focus on a phenomenological study of inflation, particularly on a model of inflation situated within a particle physics framework: the minimal supersymmetric model. In collaboration with cosmologists, theorists, and particle physicists, we demonstrate that the existing data from the Planck satellite are already precise enough that systematic errors in the model's predictions dominate the error budget in an inference context. These theoretical systematics arise from the non-inclusion of radiative corrections and an incomplete understanding of the end of inflation. We have included the necessary corrections and identified points in parameter space that satisfy both the observational constraints of particle physics (such as the Higgs mass and direct SUSY searches at the LHC) and cosmology (including the dark matter fraction in the universe and the properties of scalar perturbations as observed by Planck). Our work demonstrates the feasibility of unifying particle physics and cosmology descriptions within a single self-consistent model, paving the way for a comprehensive exploration of the inflationary MSSM or other high-energy physics models.

Acknowledgement

First and foremost, I would like to express my deepest gratitude to Sophie for her exceptional supervision throughout my thesis. None of this would have been possible without her skills, kindness, motivation, and involvement.

I also want to extend my thanks to Tony Banday and Juan Macias-Pérez for agreeing to review this manuscript, as well as to my jury members, Achille Stocchi, Jo Dunkley, Tomotake Matsumura, Nathalie Palanque-Delabrouille, and Vivian Poulin-Détolle.

I also thank all my collaborators. I want to give special recognition to Gilbert and Vincent; your scientific insight, pedagogy, patience, and humility have been invaluable to me. You dedicated so much of your time, and I have learned immensely from you. It has been an honour to work with you, and I hope we will continue to collaborate in the future. Thanks to Laurent for the passionate discussions on supersymmetry and our daily corridor-mate interactions. Thanks also to Anja, Dirk, and Richard. I also want to acknowledge the entire LiteBIRD collaboration, and particularly the E-mode group for our weekly exchanges. The interactions with various wonderful people have been incredibly rewarding; I hope you all recognise yourselves, even if I can't mention each of you by name.

An excellent working environment was essential to the success of this thesis, and I am deeply grateful to the IJCLab CMB team for these three years. Every member has been scientifically and personally central and made this thesis possible. First, thanks to Matt, from whom I have learned a great deal; your advice and support throughout have been indispensable. Stéphane, your impressive knowledge, pedagogy, and humility have been crucial: thank you for patiently answering my questions without judgment and sharing your expertise. It is so valuable and pleasant to have scientists like you around. Thanks to Xavier for consistently delivering jokes at just the right moment. I am very grateful to Thibaut, who has been a major support; not only are you a scientist I deeply respect, but you have also been a supportive and humorous friend. Thank you for your intensity, your noisy office entries, and everything else. My deepest thanks go to Adrien, with whom I have shared countless discussions, laughs, and beers. I am thrilled to have another year of this ahead! Thanks, Giacomo, for sharing the office and being my housemate for four months; I owe you a lot (*eg.* those Domino's pizza, my Moka, our passionate discussions, ...). Big thanks also to Merry. The three of you have been genuinely wonderful friends throughout

these three years. Thanks also to Camille and Léon, my first fellow internship students. A big thank you as well to Bérénice, Noémie, and Thomas for the great times we have shared lastly.

I also want to acknowledge my other cosmological friends: Arianna, Magdy, Clément, Baptiste, Léo, Louise, Thomas, and Hugo. Thanks also to the wonderful people I met at Ecole de Gif, Ecole Euclid, Tokyo, Elba, and more recently in Noirmoutier.

My close friends have been a constant source of support, and I am profoundly grateful. Thanks to my housemates Arsène, Corentin, Fab, Sam, Stann, Giac, and Baptiste for the dinners, climbing sessions, runs, crêpes, evenings, and weekends. Your positivity has been essential over the past few years. I also want to infinitely thank Lucas, Noémie, and Simon, who have been another major supportive pillar during my thesis. Thank you for the countless laughs, lunches, beers, coffees, parties, and more. I will certainly miss your presence next year. Thanks also to Rebecca and Theo, who shared many of these moments with us.

Lastly, I want to deeply thank my lifelong friends: Hugo; Mathieu, Yanis, Charline, Lucas, Nelly, Guillaume, Louis, Maelle, and Arnaud; Bastien, Noémie, Rémi, Swann, and Zoé; Clémentine, and Pierre; Mathieu, Yvain, and Romain. I apologise for listing your names in this manner; I hope you know how much each of you means to me and how grateful I am to have you all around; there isn't enough space here to recount all the memories we share. I also thank my dear friends from Padua, with whom I had so much fun partying and traveling during six extraordinary months; I hope to see you again soon. The same applies to my fellow NPAC mates.

Finally, I want to thank my precious family: Maman, Papa and Simon, Arsène, my aunts and uncles, my cousins, as well as Mamie, Papi and Grand-mère, for their continuous support throughout my life. A special thought goes to Grand-père, to whom this thesis is dedicated.

Last but certainly not least, I want to thank you, Sruthi, for being who you are and for your unwavering support during these years and in the ones to come.

Table of contents

Conventions **i**

I Introduction to cosmology, CMB and inflation 1

1 The standard model of Cosmology **5**
1.1 An expanding homogeneous Universe . . . 6
1.2 A slightly perturbed Universe . . . 12
1.3 Picture at the recombination . . . 22

2 Cosmic Microwave Background **27**
2.1 Photon anisotropy today . . . 28
2.2 General formalism . . . 31
2.3 Statistics of the maps . . . 37
2.4 Current observational status . . . 41

3 Inflation as the cherry on the top **49**
3.1 A solution to conceptual issues of the cosmological model . . . 50
3.2 Single-field slow-roll inflation . . . 54
3.3 Slow-roll inflationary perturbations . . . 57

4 Challenges and target for future CMB experiments **63**
4.1 Effects on the late line-of-sight . . . 64
4.2 Instrumental effects . . . 72
4.3 Estimating or forecasting cosmological constraints . . . 74
4.4 Current and future observational targets . . . 78
4.5 Next-generation CMB experiments . . . 83

II The LiteBIRD mission and its Instrument Model 87

5 LiteBIRD overview 91
5.1 Requirements . . . 91
5.2 Characteristics and design . . . 92
5.3 Data-analysis processes and analyses . . . 98

6 LiteBIRD Instrument Model 101
6.1 IMo definition and organisation . . . 102
6.2 Low-level instrumental quantities . . . 105
6.3 High-level derived quantities . . . 112
6.4 Detectors characteristics and mapping . . . 120

Conclusion 129

III LiteBIRD data analysis preparation 131

Preliminary: simulating a microwave sky as seen by detectors 135

7 Component separation and foreground cleaning 143
7.1 Methods overview . . . 144
7.2 Harmonic Internal Linear Combination . . . 146
7.3 HILC for LiteBIRD: setup and products . . . 149
7.4 Optimisation and results . . . 153

8 Power-spectrum estimation 165
8.1 Methods overview . . . 166
8.2 Comparison of cross-spectra in a white noise case . . . 171
8.3 Application to LiteBIRD and results . . . 177

9 Likelihood investigation and parameter estimation 187
9.1 Preliminary τ and r uncertainty estimate . . . 188
9.2 Likelihoods approximations overview . . . 189
9.3 Application to the white-noise case and results . . . 199
9.4 Application to LiteBIRD and results . . . 210

Conclusion 223

IV Slow-roll inflation phenomenology & MSSM inflation 225

10 Slow-roll predictions and BSM particle-physics 229
10.1 Predicting the primordial power spectra in slow-roll approximation . . . 230
10.2 Going beyond the particle standard model . . . 234

11 MSSM-inflation:
Towards a coherent description of particle physics and cosmology **243**
11.1 Slow-rolling the MSSM . . . 245
11.2 Analysis framework . . . 252
11.3 Slow-roll conditions, initial conditions, fine-tuning . . . 256
11.4 Inflationary constraints in the parameter space . . . 262
11.5 Combining HEP and cosmological constraints . . . 269

Conclusion **275**

Conclusions and perspectives **279**

Appendices **283**
A Detectability of the slow-roll tensor tilt . . . 283
B IMo tree architecture . . . 285
C Foreground complexity in L2-050 and H3-402 . . . 289
D Comments on the HL transform . . . 291
E Illustration of the $\ln R_{\rm rad}$ calculation from Γ_ϕ . . . 293
F Other eMSSM cosmological observables . . . 295

Synthèse en français **299**

Conventions

By convention, we will work with natural units imposing $c = h = k_{\rm B} = \sqrt{8\pi G} M_{\rm Pl} = 1$ where c is the speed of the light in vacuum, h the Planck constant, k_B the Boltzmann constant, G the gravitational constant and $M_{\rm Pl}$ the Planck mass. These factors can be recovered at any discussion stage by dimensional analysis.

In this work, we limit our use of bold mathematical symbols to vectors or matrices in pixel indices if not specified otherwise. Otherwise, we stick to the Einstein notation (or use explicit sum symbols).

Part I

Introduction to cosmology, CMB and inflation

This first part aims to introduce the key concepts, theoretical background, and context required for the subsequent discussions. First, we explain how to model our expanding and slightly perturbed Universe. Next, we discuss the properties and important tools related to the Cosmic Microwave Background (CMB), which is the central observable in this thesis. We then motivate and introduce the concept of cosmological inflation, a speculative era of accelerated expansion in the very early Universe. The investigation of its experimental and phenomenological aspects forms the main focus of this thesis. Finally, we provide an overview of the various observational challenges that CMB experiments must address in order to put constraints on cosmology and inflation.

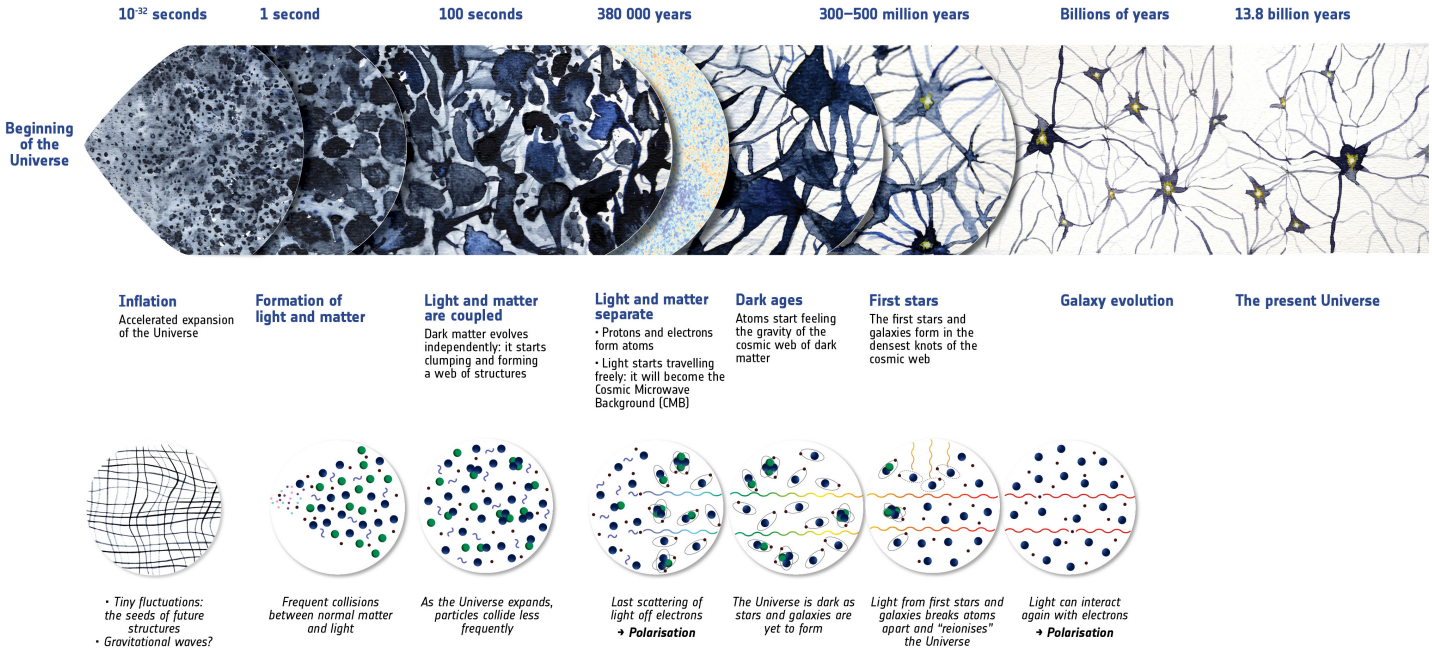

Figure I.1: Credit: *European Space Agency* [1]. Artistic summary of our modern understanding of the Universe's history. Our Universe can be described at leading order as a homogeneous and expanding spacetime filled with various matter and radiation fluids. To describe its structure and the presence of a slightly anisotropic microwave radiation called *cosmic microwave background* emitted around 380 000 years after the beginning of the Universe, one needs to track the evolution of small perturbations in these fluids. This can be done at linear order, which provides predictions in remarkable agreement with the observations. The question of the primordial origin of these perturbations can be answered by cosmological inflation, a speculative phase of accelerated expansion of the very early Universe, a tiny fraction of a second after its beginning.

CHAPTER 1

The standard model of Cosmology

In this first chapter, we will introduce the standard description of cosmology, starting from the description of the homogeneous and expanding background, before introducing small perturbations to this background that will be able to explain the inhomogeneous Universe as we know it.

Contents

1.1 An expanding homogeneous Universe 6
1.1.1 Basic properties and notations 6
1.1.2 Friedman equations 7
1.1.3 The Universe's energy budget and history 9
1.2 A slightly perturbed Universe 12
1.2.1 Scalar-vector-tensor decomposition 13
1.2.2 Scalar perturbations 13
1.2.3 Vector perturbations 21
1.2.4 Tensor perturbations 21
1.3 Picture at the recombination 22
1.3.1 Large scales: Sachs-Wolfe effect 23
1.3.2 Intermediate scales: oscillations in the tight coupling limit . . . 23
1.3.3 Small scales: diffusion damping 25

1.1. An expanding homogeneous Universe

Cosmology is based on the *Cosmological Principle*, which states that the Universe is homogeneous on large scales of around 100 Mpc. This principle is justified by the observed *isotropy* of the Universe combined with the *Copernican Principle*, which asserts that the place in which we live is not special in the Universe.

A second observational characteristic of our Universe, discovered by Edwin Hubble in 1929, is that it is expanding. Two distant galaxies move away from each other, and the content of the Universe is constantly diluted. This first chapter will introduce the definition and notations that are required for describing such a universe.

1.1.1. Basic properties and notations

To track the relative evolution of distances in the Universe (for instance, the *distance* d between two galaxies from *cosmic time* t to today, t_0), we introduce the *scale factor*

$$a(t) \equiv \frac{d(t)}{d(t_0)}, \tag{1.1}$$

which is one today by convention.

The wavelength of a photon travelling in an expanding background is stretched out proportionally to the scale factor. Equivalently, the thermodynamic temperature diminishes proportionally with the scale factor. This photon energy dilution along its travel from emission time t to today t_0 is called the redshift, z:

$$1 + z \equiv \frac{\lambda(t_0)}{\lambda(t)} = \frac{1}{a(t)} = \frac{T(t)}{T(t_0)}. \tag{1.2}$$

The redshift is commonly used as a time variable instead of t since it is directly observable. To relate both parameterisations, one needs to determine the function $a(t)$, which has been a major challenge for cosmologists during the last century. The so-called *Big-Bang* is the limit $z \to \infty$ *i.e.* $t \to 0$, when the Universe was asymptotically infinitely hot and dense.

A fundamental object for describing a spacetime geometry is the *metric*, the tensor that allows computing the infinitesimal distance $\mathrm{d}s^2 = g_{\mu\nu}\mathrm{d}x^\mu \mathrm{d}x^\nu$. For example, for Euclidean space, this prescription for computing distances reduces to the Pythagorean theorem. The simplest useful metric in cosmology called the Friedmann-Lemaitre-Roberson-Walker (FLRW) metric, describes the geometry of a homogeneous and expanding Universe:

$$\mathrm{d}s^2 = -\mathrm{d}t^2 + a^2(t)\left[\frac{\mathrm{d}r^2}{1-\mathcal{K}r^2} + r^2\left(\mathrm{d}\theta^2 + \sin^2\theta \mathrm{d}\phi^2\right)\right] \tag{1.3}$$

with t the cosmic time, r the *comoving radial coordinate*, θ and ϕ the *comoving angular coordinates*, and $\mathcal{K}$ the *curvature*, which is -1 for an *open*, 0 for a *flat*, or $+1$ for a *closed* Universe. Two distance conventions can be chosen, L the *physical distance* and χ the *comoving distance*, linked by:

$$L(t) = a(t)\chi, \tag{1.4}$$

where χ is defined from the path of a radial light ray for which $\mathrm{d}s^2 = -\mathrm{d}t^2 + a^2(t)\mathrm{d}\chi^2 = 0$:

$$\chi = \int_{t_e}^{t_0} \frac{\mathrm{d}t}{a(t)}. \tag{1.5}$$

Another name and associated definition for the same quantity (up to its dimension) is the *conformal time* η, defined as $\mathrm{d}\eta \equiv \mathrm{d}t/a(t)$. In this thesis, the notation "$\dot{}$" means differentiating with respect to the cosmic time ($\dot{O} \equiv \frac{\mathrm{d}O}{\mathrm{d}t}$), while "$'$" means differentiating with respect to the conformal time ($O' \equiv \frac{\mathrm{d}O}{\mathrm{d}\eta}$).

Differentiating Eq. 1.4 with respect to the time, the velocity of a galaxy in the Hubble flow with respect to us can be derived:

$$v = \frac{\mathrm{d}L}{\mathrm{d}t} = \dot{a}(t)\chi = H(t)L, \tag{1.6}$$

where $H(t)$ is the so-called *Hubble rate* and is defined as

$$H(t) \equiv \frac{\dot{a}(t)}{a(t)}. \tag{1.7}$$

Evaluating this relation today (time t_0) allows us to obtain the *Hubble law*, first verified for nearby galaxies in 1929:

$$v = H_0 L, \tag{1.8}$$

where $H_0 \equiv H(t_0)$ is commonly called the *Hubble constant*. Let's also define the *reduced Hubble parameter* h defined such that

$$H_0 = 100\ h\mathrm{km/sec/Mpc}. \tag{1.9}$$

Finally, one can rewrite the comoving distance Eq. 1.5 as an integral on the redshift:

$$\chi = \int_1^{a_e} \frac{\mathrm{d}a}{a}\frac{\mathrm{d}t}{\mathrm{d}a} \tag{1.10}$$

$$= \int_1^{a_e} \frac{\mathrm{d}a}{H(a)a^2} \tag{1.11}$$

$$= \int_0^{z_e} \frac{\mathrm{d}z}{H(z)} \tag{1.12}$$

where the last step comes from Eq. 1.2 which implies $\frac{\mathrm{d}a}{\mathrm{d}z} = -a^2$.

1.1.2. Friedman equations

Now that we have established these useful properties and notations, let us introduce the *Einstein equations* that are necessary to describe the dynamic of space-time metrics. The first fundamental object that enters them is the *Einstein tensor* $G_{\mu\nu}$ that describes the curvature of space-time, and the *energy-momentum tensor* $T_{\mu\nu}$ that describes the energy content of the Universe. The Einstein tensor is a function of the metric itself:

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} \tag{1.13}$$

where $R_{\mu\nu}$ is the *Ricci tensor* (and R its determinant, the *Ricci scalar*), constructed from the *Christoffel symbols* defined as

$$\Gamma^{\rho}_{\mu\nu} \equiv \frac{1}{2}g^{\rho\lambda}(\partial_\nu g_{\lambda\mu} + \partial_\mu g_{\lambda\nu} - \partial_\lambda g_{\mu\nu}), \tag{1.14}$$

in the following manner, denoting $A_{,i} \equiv \frac{\partial A}{\partial x^i}$,

$$R_{\mu\nu} \equiv 2\Gamma^{\rho}_{\mu[\nu,\rho]} + 2\Gamma^{\rho}_{\lambda[\rho}\Gamma^{\lambda}_{\nu]\mu}, \tag{1.15}$$

where the brackets denote an anti-symmetrisation over the indices.

The second ingredient required to track the evolution history of the Universe is the energy-momentum tensor $T_{\mu\nu}$ that describes the Universe's energy content. In general, $T_{\mu\nu}$ is defined from the action S of a given species as

$$T_{\mu\nu} \equiv -\frac{2}{\sqrt{-g}}\frac{\partial \mathcal{S}}{\partial g_{\mu\nu}}. \tag{1.16}$$

In cosmology, the various components can be described at the largest scales by a perfect fluid with no shear stresses or viscosity. Moreover, we stick to the homogeneity assumption to start with. In this case, the stress-energy tensor becomes

$$T_{\mu\nu} = \rho u_\mu u_\nu + \frac{p}{a^2}\, g_{\mu\nu}|_{\mu\neq 0,\nu\neq 0}\,, \tag{1.17}$$

where u_μ is the 4-velocity of the fluid, p is its pressure and ρ its density. Then, the Einstein equation reads:

$$G_{\mu\nu} = 8\pi G T_{\mu\nu}, \tag{1.18}$$

If one uses the FLRW metric Eq. 1.3 to derive $G_{\mu\nu}$, the Einstein's equations governing the evolution of a homogeneous and expanding Universe are called the *Friedmann equation* and the *Raychaudhuri equation*, which can be respectively written as

$$H^2 = \frac{8\pi G}{3}\rho - \frac{\mathcal{K}}{a^2}, \tag{1.19}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3P), \tag{1.20}$$

where ρ and P are the *density* and the *pressure* of a cosmological fluid, linked by the *equation of state* $\frac{P}{\rho} = \omega$, where ω is a parameter depending on the fluid. It is also possible to rewrite Eqs. 1.19 and Eq. 1.20 indexing the fluids of the Universe by i and introducing

$$\Omega_i \equiv \frac{\rho_i^{(0)}}{\rho_{\rm crit}^{(0)}}, \tag{1.21}$$

where $\rho_i^{(0)}$ is the density of i today and $\rho_{\rm crit}^{(0)} \equiv \frac{3H_0^2}{8\pi G}$ (the *critical density* of the Universe today):

$$H^2 = H_0^2 \sum_i \Omega_i (1+z)^{3(1+\omega_i)} - \frac{\mathcal{K}}{a^2}, \tag{1.22}$$

$$\frac{\ddot{a}}{a} = H_0^2 \sum_i \Omega_i \frac{1+3\omega}{2}. \tag{1.23}$$

Combining Eqs. 1.22, 1.23, one obtains the *continuity equation* valid for each component X,

$$\dot{\rho}_X + 3H(\rho_X + p_X) = 0, \tag{1.24}$$

which can be solved for ρ using the equation of state for a stable fluid of parameter ω_X:

$$\rho_X(z) = \rho_{X,0}(1+z)^{3(1+\omega_X)}. \tag{1.25}$$

Consequently, according to Eq. 1.19, if a species X dominates, the scale factor will evolve according to

$$a(t) \propto \begin{cases} t^{\frac{2}{3(1+\omega_X)}} & \text{for } \omega_X \neq -1, \\ \exp\left(\sqrt{\frac{\rho_{X,0}}{3}}\frac{t}{M_\mathrm{P}}\right) & \text{for } \omega_X = -1, \end{cases} \tag{1.26}$$

where $M_\mathrm{P} \equiv \frac{1}{\sqrt{8\pi G}}$ is the *Planck mass*, and $\rho_{X,0}$ an integration constant. Hence, knowing the ω parameters and current densities of each component of the Universe, one can extract its evolution.

1.1.3. The Universe's energy budget and history

For simplicity, every non-relativistic species interacting with photons are called **baryons**, including electrons and nucleons. It is a pressureless fluid, hence, with the simple equation of state

$$P_\mathrm{b} = \omega_\mathrm{b}\rho_\mathrm{b} = 0. \tag{1.27}$$

During a matter-dominated era, Eq. 1.26 leads to

$$1/z = a \propto t^{2/3} \propto \eta^2. \tag{1.28}$$

The **photon**, γ, is a stable and relativistic ($T \gg m$) boson interacting mainly via a non-relativistic scattering (called *Thomson* or *Compton*) with the electrons. Its equation of state is

$$P_\gamma = \omega_\gamma\rho_\gamma = \rho_\gamma/3. \tag{1.29}$$

Hence, during a radiation-dominated era, Eq. 1.26 leads to

$$1/z = a \propto t^{1/2} \propto \eta. \tag{1.30}$$

The photon decouples from the primordial plasma when the electrons combine with the hydrogen nuclei to give neutral hydrogen atoms ($pe \leftrightarrow H\gamma$). Prior to this event, the photons and electrons were tightly coupled and kinetic equilibrium was enforced because of the numerous scattering events. To find the temperature at which this event happens, one needs to know the binding energy of the hydrogen and the ratio baryon/photon $R = n_\mathrm{b}/n_\gamma$. The distributions of the elements all behave like classical dilute gas, following the *Boltzmann distribution*:

$$f(E_s) = \exp\left(-\frac{E_s - \mu_s}{T_s}\right), \tag{1.31}$$

where μ_s is the chemical potential, E_s the energy and T_s the temperature of the component s. To find the associated density, one integrates the distribution over the phase space,

$$n_s = g_s e^{\mu_s/T} \int \frac{\mathrm{d}^3p}{(2\pi)^3} e^{-E_s(p)/T}, \tag{1.32}$$

where g_s is the number of freedom of the species s. We also denote $n_s^{(0)}$ the density for a perfect chemical equilibrium with $\mu_s = 0$. Then, Eq. 1.32 can be rewritten as

$$n_s^{(0)} = \begin{cases} g_s \left(\frac{m_s T}{2\pi}\right)^{3/2} e^{-m_s/T} & m_s \gg T, \\ g_s \frac{T^3}{\pi^2} & m_s \ll T, \end{cases} \tag{1.33}$$

depending on whether the species is relativistic or not. Let's call x_e the *ionisation fraction*, defined as the ratio between the electron density, n_e, and the hydrogen nuclei density, $n_{\rm H} + n_{\rm H^+}$, where $n_{\rm H}$ is the fraction of neutral hydrogen $n_{\rm H^+}$ of ionised hydrogen:

$$x_e \equiv \frac{n_e}{n_{\rm H} + n_{\rm H^+}}. \tag{1.34}$$

As long as only hydrogen atoms can be ionised, $n_{\rm H^+} = n_e$ and this fraction lies between zero and one[1]. In this particular regime, the densities of neutral hydrogen atoms $n_{\rm H}$, proton n_p and electrons n_e evolve around perfect chemical equilibrium while preserving the *Saha equation*,

$$\frac{n_e n_p}{n_{\rm H}} = \frac{n_e^{(0)} n_p^{(0)}}{n_{\rm H}^{(0)}}. \tag{1.35}$$

The evolution of x_e can be followed by inserting Eqs. 1.33 and 1.34 into Eq. 1.35:

$$\frac{x_e^2}{1 - x_e} = \frac{1}{n_e + n_{\rm H}} \left[\left(\frac{m_e T}{2\pi} \right)^{3/2} e^{-[m_e + m_p - m_{\rm H}]/T} \right]. \tag{1.36}$$

As long as only hydrogen atoms can be ionised, the evolution of the ionisation fraction can be solved for, showing that the Universe transitioned from ionised to neutral in a short amount of time. The temperature at which this transition happens is called the *recombination* temperature. When the Universe's temperature drops below this value, the mean free path of the photons immediately increases because the free electrons on which they scatter become rarer and rarer as they combine to form hydrogen. Usually, decoupling is defined as equality between the horizon size and the mean free path of the species considered. The so-called *photon decoupling* happens very shortly after the recombination. Given the current observational constraints that we will introduce in Sec. 2.4, these two events happen around $z \simeq 1100$ in the standard model of cosmology. The decoupled photons then travel freely through the older Universe. At leading order, this leads to an isotropic black-body radiation of temperature[2] T_0 called the *Cosmic Microwave Background* (CMB), whose photon temperature is distributed as $B_\nu(T_0)$ where B_ν is the Planck law:

$$B_\nu(T) = \frac{2\nu^3}{e^{\frac{\nu}{T}} - 1}. \tag{1.38}$$

The **neutrinos**, ν, are fermions, composed of three species (ν_e, ν_μ, ν_τ). At early times, they all were lighter than the thermal bath temperature; hence, they were relativistic species. As for the photons, the resulting equation of state is

$$P_\nu = \omega_\nu \rho_\nu = \rho_\nu / 3. \tag{1.39}$$

The neutrinos also interact with electrons through $e\nu \leftrightarrow e\nu$ and $ee \leftrightarrow \nu\nu$, but the associated interaction rate is much lower than for photons. Hence, it did not wait for the recombination

[1] While it could become greater than one, *eg.* in presence of ionised Helium atoms.

[2] The peak frequency ($\nu_{\rm peak}$ in GHz) or wavelength ($\lambda_{\rm peak}$ in mm) can be linked to T_0 by differentiating Eqs. 1.38 with respect to ν or λ and equating to zero:

$$\nu_{\rm peak} / \left(58.79 \text{ GHz/K}\right) = \left(5.10 \text{ mm} \cdot \text{K}\right) / \lambda_{\rm peak} = T_0. \tag{1.37}$$

According to measurements, [2] and see Sec. 2.4, $T_0 = 2.725$ K so $\nu_{\rm peak} = 160.2$ GHz and $\lambda_{\rm peak} = 1.87$mm.

and the associated rarefaction of electrons to decouple from its surroundings. Comparing the standard model interaction rate with the Hubble radius in the proper cosmological era allows one to derive the *neutrino decoupling temperature* ($T_\nu \simeq$ 2-3 MeV in the standard model [3]). The neutrinos then free stream through the Universe, and the massive ones transition at late time to non-relativistic species (inducing a modification of Eq. 1.39 and slowing down in the formation of large-scale structures).

To describe the observations, another non-relativistic fluid is required, which only interacts through the metric, the **cold-dark matter**, which follows a matter equation of state:

$$P_{\rm c} = \omega_{\rm c}\rho_{\rm c} = 0. \tag{1.40}$$

The difference with baryons is that cold-dark matter does not interact with any of the other species. As we will see, the abundant presence of such a fluid is unavoidable in cosmology. The question of its fundamental nature is still an open question of physics. For instance, it could be a dark sector particle that decoupled from the visible sector well before the electroweak transition. The typical mass and interaction rate with standard model particles of this dark sector is constrained by dark matter direct searches and particle physics colliders.

The last component to introduce is a **dark-energy** fluid, Λ:

$$P_\Lambda = \omega_\Lambda\rho_\Lambda = -\rho_\Lambda. \tag{1.41}$$

When homogeneous and stable, this kind of fluid is equivalent to introducing a *cosmological constant* $g_{\mu\nu}\Lambda$ on the right-hand side of the Einstein equation (consequently $\rho_\Lambda \equiv \frac{\Lambda}{8\pi G}$). During a *de Sitter Universe*, *i.e.* dominated by such a dark-energy fluid, Eq. 1.26 leads to

$$1/z = a \propto \exp\left(Ct\right) \propto -\exp\left(-C\eta\right), \tag{1.42}$$

with C some integration constant.

Such an equation of state is uncommon: it means the fluid has a negative pressure. Nevertheless, it is usually invoked to explain the observed current acceleration phase of the Universe, alongside (up to some small corrections) a speculative era of accelerated expansion of the very early Universe, *inflation*, which is one of the main topics of the thesis. As for dark matter, neither the fundamental nature of current dark energy nor inflation is known.

We summarise in Fig. 1.1 the evolution of the various fluids throughout the history of the Universe. The numerical values chosen for this plot are compatible with the current cosmological constraints, which we will introduce in Sec. 2.4. The upper panel illustrates the evolution of the densities normalised to the total density today. The x-axis tracks $1/a$ in logarithmic scale, so going towards the left of the Figure means going back in time. The cosmic time in years is specified on the upper axis. The cold-dark matter, baryon, neutrino (3 massive species), and dark energy are represented in blue, light blue, red, orange and green, respectively. One recover the equations of state Eqs. 1.40, 1.27, 1.39, 1.29, 1.41 applied to Eq. 1.32.

Especially, one can identify three eras:

- A **radiation-dominated era** during which the scale factor evolves with Eq. 1.30, which ends at a redshift of $z_{\rm eq} \simeq 3400$ corresponding approximately to 50000 years after the Big-Bang. This transition is illustrated by a first dashed grey vertical line (in this case, $z_{\rm eq} = 3398$ and $t_{\rm eq} = 51100$ years).

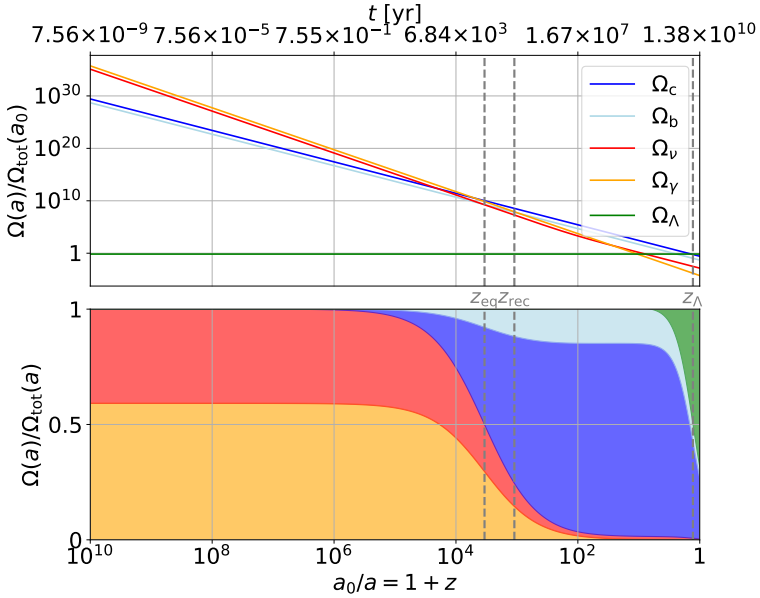

Figure 1.1: TOP: fraction of the various fluid species at each (cosmic) time (or scale factor or redshift), normalised to the current total density. BOTTOM: Proportion plot representing the same quantities normalised to the total density at each time. Anticipating on Sec. 2.4, the chosen numerical values are compatible with the current observational constraints and correspond to the fiducial we introduce later on in Tab. 4.1.

- A **matter-dominated era**, during which the scale factor evolves with Eq. 1.28. CMB is emitted during this era, at a redshift of $z_{\rm rec} \simeq 1090$ and time of 380 000 years, as symbolised by a second dashed grey vertical line. This ends at a redshift of $z_\Lambda \simeq 0.3$, lasting between 100 000 and 10 billion years after the Big-Bang. The end of this phase is displayed by the third grey dashed line (in this case, $z_\Lambda = 0.29$, which corresponds to $t_\Lambda = 10.3$ Gyrs).
- A **Λ-dominated era**, during which the scale factor evolves with Eq. 1.42, which we currently stand in, 13.8 billion years after the Big-Bang.

These three eras are also evidenced in the lower panel. The y-axis represents the same density quantity as above, in a proportion plot, normalised to the total density at each considered time.

1.2. A slightly perturbed Universe

In addition to this dynamic but homogeneous Universe, it is necessary to describe the structures present within it and their evolution. More precisely, each species i is distributed according to a distribution $f_i(\vec{x}, \vec{p}, t)$ that depends on both position and momentum. This distribution obeys the Einstein and Boltzmann relativistic equations, which describe how perturbations in various species evolve, interact with each other, and interplay with the metric through gravity. In the case of a perfectly homogeneous Universe, this formalism reduces to the usual Friedman and Raychaudhuri equations Eqs. 1.19 and 1.20. In this section, we will slightly deviate from the FLRW metric by introducing small scalar perturbations, typically less than 10^{-4} at the

largest and *linear* scales (the ones we will focus on). Then, we will present how one derives the Einstein-Boltzmann system at linear order in these perturbations.

1.2.1. Scalar-vector-tensor decomposition

To model our slightly inhomogeneous Universe, we introduce in the most generic manner small fluctuations in the FLRW metric Eq. 1.3 (where we set $\mathcal{K} = 0$, choice further motivated in Sec. 2.4) [4, 5]:

$$\mathrm{d}s^2 = a(\eta)^2 \left\{ -(1+2\Psi)\mathrm{d}\eta^2 + 2B_i \, \mathrm{d}x^i \, \mathrm{d}\eta + [\delta_{ij} + h_{ij}] \, \mathrm{d}x^i \, \mathrm{d}x^j \right\} \tag{1.43}$$

Ψ is a 3-scalar, function of space and time. ω_i is a 3-vector, function of space and time, but itself can be decomposed according to the Helmholtz decomposition theorem in a longitudinal (scalar) and a (divergent free) transverse (vector) part: $B_i = -\partial_i B - \bar{B}_i$ where $\partial_i \bar{B}_i = 0$. In the same way, the ten degrees of freedom of the symmetric 3-tensor h_{ij}, function of space and time, can be decomposed in scalar, vector and tensor parts: $h_{ij} = 2\Phi\delta_{ij} + 2\partial_i\partial_j E + 2\partial_{(i}\bar{E}_{j)} + 2\bar{E}_{ij}$, with $\partial_i \bar{E}^{ij} = 0$ and $\bar{E}^i_i = 0$. Because of spatial isotropy, each perturbation type evolves independently of the others at linear order [6]. This is the so-called *decomposition theorem*.

The above decomposition is very generic but also redundant because of gauge freedom [6–8]. A choice of gauge is required to keep an adequate number of gauge invariant variables: the four scalar, four vector, and ten tensor functions in the above actually reduce for each type to two independent physical degrees of freedom (*dof*'s for short). A popular gauge choice, the so-called conformal Newtonian gauge [7], results in the following independent *dof*'s: Ψ and Φ the scalar potentials (2 *dof*'s), $\omega_i \equiv \bar{B}_i$ the transverse vector potential (3-1=2 *dof*'s), and $\gamma_{ij} \equiv \bar{E}_{ij}$ both transverse and traceless spatial metric (6-3-1=2 *dof*'s), leading to the perturbed element line [7, 9]:

$$\mathrm{d}s^2 = a^2(\eta) \left\{ -(1+2\Psi)\mathrm{d}\eta^2 + 2\omega_i \, \mathrm{d}x^i \, \mathrm{d}\eta + [(1+2\Phi)\delta_{ij} + 2\gamma_{ij}] \, \mathrm{d}x^i \, \mathrm{d}x^j \right\}. \tag{1.44}$$

When $\Phi = \Psi = 0$, this metric simply reduces to the FRLW one (which we have written in spherical coordinates and with respect to cosmic-time in Eq. 1.3). We will focus on the evolution of these perturbations in the next three sections: the scalar Ψ and Φ in Sec. 1.2.2, the vector ω_i in Sec. 1.2.3 and the tensor γ_{ij} in Sec. 1.2.4. We will also introduce the associated *super-horizon* initial conditions, delaying the question of their fundamental (inflationary) origin to Sec. 3.3.1.

1.2.2. Scalar perturbations

First, we present the derivation at linear order, focusing on scalar perturbations of the metric that originate and are originated by density fluctuations, as done for the first time in [10]. According to the decomposition theorem, this computation can be done by fixing the non-scalar *dof*'s to zero and only consider the scalar parts of the perturbed metric Eq. 1.44. The scalar perturbed metric in the conformal gauge reduces to [5]:

$$\mathrm{d}s^2 = -(1+2\Psi)\mathrm{d}t^2 + a^2(1+2\Phi)\delta_{ij}\mathrm{d}x^i\mathrm{d}x^j. \tag{1.45}$$

1.2.2.1 Boltzmann equations

To describe how the angular distribution function evolves with scattering, we use Boltzmann's equation, in which we drop all second-order terms in the perturbation. In its schematic version,

this equation reads [5]

$$\frac{Df}{Dt} = C[f], \tag{1.46}$$

where f is the distribution of the considered species and C is the collision term. The distribution may depend on the position, momentum, and time, so the total derivative of the distribution can be decomposed as

$$\frac{Df(t,x_i,p_i)}{Dt} = \frac{\partial f}{\partial t} + \frac{\partial f}{\partial \vec{x}}\frac{\mathrm{d}\vec{x}}{\mathrm{d}t} + \frac{\partial f}{\partial \vec{p}}\frac{\mathrm{d}\vec{p}}{\mathrm{d}t} \tag{1.47}$$

$$= \frac{\partial f}{\partial t} + \frac{\partial f}{\partial x^i}\frac{\mathrm{d}x^i}{\mathrm{d}t} + \frac{\partial f}{\partial p}\frac{\mathrm{d}p}{\mathrm{d}t} + \frac{\partial f}{\partial \hat{p}^i}\frac{\mathrm{d}\hat{p}^i}{\mathrm{d}t}, \tag{1.48}$$

where $\hat{p}^i$ is the unit vector describing the direction of the momentum and p its amplitude. We will need Boltzmann's equations for each Universe component: photons, baryons, dark matter and neutrinos. Let's focus, for instance, on the derivation of the Boltzmann's equation for **photons**. It is the most complete case since it requires a relativistic treatment and includes a collision term with baryons. Let's write P the 4-momentum. the perturbed FLRW gives for its square $P^2 = -(1+2\Psi)(P^0)^2 + p^2$ which also vanishes for massless particles. Then, P^0 at first order in Ψ simply is written as

$$P^0 = p(1-\Psi). \tag{1.49}$$

We can write $P^\mu = \frac{\mathrm{d}x^\mu}{\mathrm{d}\lambda}$ with λ parametrizing the particle path. Because of Eq. 1.49 and $P^i = p\hat{p}^i\frac{1-\Phi}{a}$, the spatial velocity is written as

$$\frac{\mathrm{d}x^i}{\mathrm{d}t} = \frac{\mathrm{d}x^i}{\mathrm{d}\lambda}\frac{\mathrm{d}\lambda}{\mathrm{d}t} = \frac{P^i}{P^0} = \frac{\hat{p}^i}{a}(1+\Psi-\Phi). \tag{1.50}$$

The photon distribution is Planckian and anisotropic at first order, $\frac{\partial f}{\partial \hat{p}^i} = 0$, and $\frac{\partial f}{\partial x^i} = 0$. Moreover, from the the *geodesic equation* that describes how particles move in space-time, one can obtain

$$\frac{\mathrm{d}p}{\mathrm{d}t} = -p\left(H + \frac{\partial \Phi}{\partial t} + \frac{\hat{p}^i}{a}\frac{\partial \Psi}{\partial x^i}\right). \tag{1.51}$$

Finally, inserting these into Eq. 1.47 that the total derivative of the distribution is written as

$$\frac{Df(t,x,p)}{Dt} = \frac{\partial f}{\partial t} + \frac{\partial f}{\partial x^i}\frac{\hat{p}^i}{a} - p\frac{\partial f}{\partial p}\left(H + \frac{\partial \Phi}{\partial t} + \frac{\hat{p}^i}{a}\frac{\partial \Psi}{\partial x^i}\right). \tag{1.52}$$

The photon distribution can be viewed as a small deviation around a *Bose-Einstein distribution*,

$$f(\vec{x},p,\hat{p},t) = \left(\exp\left\{\frac{p}{T(t)\left[1+\Theta(\vec{x},\hat{p},t)\right]}\right\} - 1\right)^{-1}, \tag{1.53}$$

where $\Theta = \frac{\delta T}{T}$ is the perturbation in photon temperature assumed to only depend on the momentum p. Expanding it for small perturbations Θ,

$$f \simeq f^{(0)} - p\frac{\partial f^{(0)}}{\partial p}\Theta \quad \text{with} \quad f^{(0)} \equiv \left[\exp\left(\frac{p}{T}\right) - 1\right]^{-1}. \tag{1.54}$$

Eventually, inserting Eq. 1.54 inside Eq. 1.52, one gets the left-hand side of the Boltzmann's equation:

$$\frac{Df}{Dt} = \underbrace{\frac{\partial f^{(0)}}{\partial t} - Hp\frac{\partial f^{(0)}}{\partial p}}_{\text{0th order}} \underbrace{- p\frac{\partial f^{(0)}}{\partial p}\left(\frac{\partial \Theta}{\partial t} + \frac{\hat{p}^i}{a}\frac{\partial \Theta}{\partial x^i} + \frac{\partial \Phi}{\partial t} + \frac{\hat{p}^i}{a}\frac{\partial \Psi}{\partial x^i}\right)}_{\text{1st order}}. \tag{1.55}$$

The integrated right-hand side collision term, which comes from Compton scattering ($\gamma(p) + e^-(q) \leftrightarrow \gamma(p') + e^-(q')$) is given by

$$\begin{aligned} C[f(\vec{p})] =& \frac{1}{p}\int \frac{\mathrm{d}^3 q}{(2\pi)^3 2E_e(q)} \int \frac{\mathrm{d}^3 q'}{(2\pi)^3 2E_e(q')} \int \frac{\mathrm{d}^3 p'}{(2\pi)^3 2E_e(p')} |\mathcal{M}|^2 (2\pi)^4 \\ & \times \delta^3[\vec{p} + \vec{q} - \vec{p}' - \vec{q}']\delta[E(p) + E(q) - E(p') - E(q')]\{f_e(\vec{q}')f(\vec{p}') - f_e(\vec{q})f(\vec{q})\}. \end{aligned} \tag{1.56}$$

The Dirac functions enforce the energy momentum conservation of the Compton scattering. Thanks to the Feynman rules, the *scattering matrix* $\mathcal{M}$ can be approximated computed as $|\mathcal{M}|^2 = 8\pi\sigma_T m_e^2$, with σ_T the Thomson cross-section and m_e the mass of the electron. After a few delta functions manipulations and an integration over q', one gets

$$\begin{aligned} C[f(\vec{p})] =& \frac{\pi}{4m_e^2 p}\int \frac{\mathrm{d}^3 q\, f_e(\vec{q})}{(2\pi)^3} \int \frac{\mathrm{d}^3 p'}{(2\pi)^3 p'} 8\pi\sigma_T m_e^2 \\ & \times \left\{\delta[p - p'] + \frac{(\vec{p} - \vec{p}')\vec{q}}{m_e}\frac{\partial \delta(p - p')}{\partial p'}\{f(\vec{p}') - f(\vec{p})\}\right\}. \end{aligned} \tag{1.57}$$

Eventually, after replacing f with Eq. 1.54 and integrating over the momentum space of p', one gets

$$C[f(\vec{p})] = -p\frac{\partial f^{(0)}}{\partial p} n_e \sigma_T [\Theta_0 - \Theta(\hat{p}) + \hat{p}.\vec{v_\mathrm{b}}], \tag{1.58}$$

where Θ_ℓ is the ℓ-th multipole[3] of Θ, n_e is the electron density, and v_b the mean velocity of baryons. Finally, the photon Boltzmann's equation to first order in perturbation reads:

$$\underbrace{\frac{\partial f^{(0)}}{\partial t} - Hp\frac{\partial f^{(0)}}{\partial p}}_{\text{0th order}} \underbrace{- p\frac{\partial f^{(0)}}{\partial p}\left[\frac{\partial \Theta}{\partial t} + \frac{\hat{p}^i}{a}\frac{\partial \Theta}{\partial x^i} + \frac{\partial \Phi}{\partial t} + \frac{\hat{p}^i}{a}\frac{\partial \Psi}{\partial x^i}\right]}_{\text{1st order}} = \underbrace{-p\frac{\partial f^{(0)}}{\partial p} n_e \sigma_T [\Theta_0 - \Theta(\hat{p}) + \hat{p}.\vec{v_\mathrm{b}}]}_{\text{1st order}}. \tag{1.60}$$

On the one hand, since $\frac{\partial f^{(0)}}{\partial t} = \frac{\partial f^{(0)}}{\partial T}\frac{\mathrm{d}T}{\mathrm{d}t} = -\frac{\mathrm{d}T}{\mathrm{d}t}\frac{p}{T}\frac{\partial f^{(0)}}{\partial p}$, the leading-order simply reduces to

$$\left(\frac{-\mathrm{d}T/\mathrm{d}t}{T} - \frac{\mathrm{d}a/\mathrm{d}t}{a}\right)\frac{\partial f^{(0)}}{\partial p} = 0, \tag{1.61}$$

which directly implies the background relation of proportionality between the temperature and the inverse scale factor given in Eq. 1.2.

[3] also called Legendre moment and defined as

$$\Theta_\ell(k, \eta) \equiv \frac{1}{(-i)^\ell}\int_{-1}^{1} \frac{d\mu}{2} P_\ell(\mu)\Theta(\mu, k, \eta), \tag{1.59}$$

where P_ℓ is the ℓ-th Legendre polynomial.

On the other hand, the linear order equation gives:

$$\frac{\partial \Theta}{\partial t} + \frac{\hat{p}^i}{a}\frac{\partial \Theta}{\partial x^i} + \frac{\partial \Phi}{\partial t} + \frac{\hat{p}^i}{a}\frac{\partial \Psi}{\partial x^i} = n_e \sigma_T \left(\Theta_0 - \Theta(\hat{p}) + \hat{p}\cdot\vec{v_{\rm b}}\right). \tag{1.62}$$

One introduces the *optical depth* at a conformal time η as

$$\tau(\eta) \equiv \int_{\eta}^{\eta_0} {\rm d}\tilde{\eta}\, n_e \sigma_T a, \tag{1.63}$$

where $n_e\sigma_T$ is its derivative with respect to η that quantifies the number of scattering events per η intervals. In Fourier space, with respect to the conformal time, Eq. 1.62 reads

$$\Theta' + ik\mu\Theta + \Phi' + ik\mu\Psi = \tau'[\Theta_0 - \Theta + \mu v_{\rm b}]. \tag{1.64}$$

where k is the wave vector or mode, and $\mu \equiv \hat{p}\hat{k}$.

Our description of the collision term has been simplistic so far. When one accounts for the angular dependence of Compton scattering, a dependence on the quadrupole, which is typically non-zero, arises in the collision term, and Eq. 1.64 becomes

$$\Theta' + ik\mu\Theta + \Phi' + ik\mu\Psi = \tau'[\Theta_0 - \Theta + \mu v_{\rm b} - \frac{1}{2}P_2(\mu v_{\rm b})\Theta_2]. \tag{1.65}$$

with the second Legendre polynomial being given by $P_2(x) = (3x^2-1)/2$.

Additionally, on top of their temperature, photons carry two additional degrees of freedom through their polarisation. Only linear polarisation is generated from scalar perturbations through Thomson scattering of a quadrupole. Fig. 1.2 illustrates schematically this mechanism: due to a hotter incident beam from along x as compared to the one along y, the outgoing radiation beam has a smaller intensity along the x than y resulting in a linearly polarised output beam. Hence, at linear order, the polarisation amplitude of the scattered photon only depends on the photon quadrupole Θ_2 seen by the last scattering electron:

$$\Theta_P(\hat{n}, \vec{k}) \propto \left[1 - (\hat{n}, \vec{k})^2\right] \Theta_2(k). \tag{1.66}$$

From this source term, one can write down the Boltzmann equation that drives the evolution of the polarisation amplitude [12]:

$$\Theta_P' + ik\mu\Theta_P = -\tau' \left[-\Theta_P + \frac{3}{4}\left(1-\mu^2\right)\Pi\right], \tag{1.67}$$

with $\Pi = \Theta_2 + \Theta_{P,2} + \Theta_{P,0}$. The left-hand side free-streaming term is similar to the photon temperature one, while the right-hand side collision terms consist of the source term given by Eq. 1.66 added to a loss term. As for the temperature photon case, the collision term is weighted by the number of collisions per η, τ'. In the next section, we see that this average linear polarisation leads to parity-even patterns of polarisation in the CMB, called E-modes. The coupling of the temperature field to the polarisation field also modifies the photon temperature Boltzmann equation Eq. 1.65, and leads to the following final form:

$$\Theta' + ik\mu\Theta + \Phi' + ik\mu\Psi = \tau'[\Theta_0 - \Theta + \mu v_{\rm b} - \frac{1}{2}P_2(\mu)\Pi]. \tag{1.68}$$

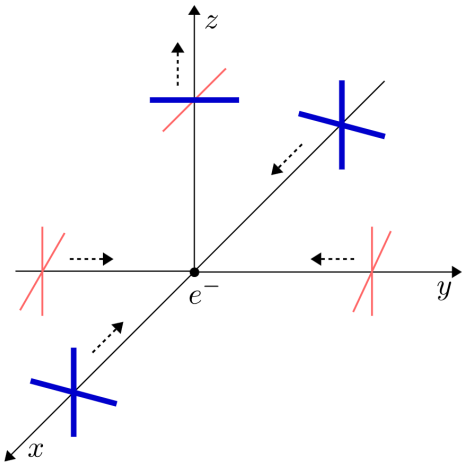

Figure 1.2: From [11]. A temperature quadrupolar anisotropy felt by the last scattering particle leads to a polarised scattered beam.

We have now provided an overview of the complete derivation of the Boltzmann equation for photons. We will be briefer for the other species.

As long as the **neutrinos** are relativistic, their Boltzmann's equation is the same as in the photon case ($\mathcal{N}$, the density contrast of neutrinos, playing the role of Θ). However, one should drop the collision term because these weakly-interacting particles are decoupled from the plasma:

$$\mathcal{N}' + ik\mu\mathcal{N} = -\Phi' - ik\mu\Psi. \tag{1.69}$$

One can do the same procedure for **dark matter**. One recovers that $n_{\rm c}^{(0)} \propto a^{-3}$ at leading order (consistently with Eqs. 1.26 and 1.40), as well as two equations at first order for each zeroth and first moments of the following Boltzmann's equation:

$$\delta_{\rm c}' + ikv_{\rm c} = -3\Phi', \tag{1.70}$$

$$v_{\rm c}' + \frac{a'}{a}v_{\rm c} = -ik\Psi, \tag{1.71}$$

where $\delta_{\rm c}$ is the perturbation in the dark matter density *i.e.* is such that $n_{\rm c} = n_{\rm c}^{(0)}\,[1+\delta(\vec{x},t)]$, and $v_{\rm c}$ is the velocity of dark matter.

For **baryons**, one gets the same equations, but with a Compton scattering term appearing in the velocity equation:

$$\delta_{\rm b}' + ikv_{\rm b} = -3\Phi', \tag{1.72}$$

$$v_{\rm b}' + \frac{a'}{a}v_{\rm b} = -ik\Psi + \frac{\tau'}{R}[v_{\rm b} + 3i\Theta_1], \tag{1.73}$$

where $\delta_{\rm b}$ is the perturbation in the dark matter density *i.e.* is such that $n_{\rm b} = n_{\rm b}^{(0)}(1+\delta_{\rm b}(\vec{x},t))$, and $R = \frac{3\rho_{\rm b}}{4\rho_\gamma}$ is the baryon-to-photon ratio.

These equations describe how scalar perturbations impact the various species present in the Universe. For instance, the photon inhomogeneities derived here will be imprinted on the so-called *last scattering surface* of electrons at the recombination, leading to CMB anisotropies that

we describe in the next section, while the inhomogeneities in the baryon distributions will evolve and lead to the large-scale structures in the sky.

Let's mention that we did not present a Boltzmann equation for **dark energy**, assuming it behaves like a (homogeneous) cosmological constant, which is so far in agreement with observations.

1.2.2.2 Einstein equations

We also need the Einstein equations that describe how fluctuations in the particle distributions interplay with metric perturbations. To write down them, we start from the perturbed FLRW metric Eq. 1.3 to evaluate the linear order Christoffel symbols Eq. 1.14,

$$\begin{aligned}
&\Gamma^0{}_{00} = \dot{\Psi},\\
&\Gamma^0{}_{0i} = \Gamma^0{}_{i0} = \nabla\Psi = ik_i\Psi,\\
&\Gamma^0{}_{ij} = \delta_{ij}a^2\big[H + 2H(\Phi-\Psi) + \dot{\Phi}\big],\\
&\Gamma^i{}_{00} = \frac{ik^i}{a^2}\Psi,\\
&\Gamma^i{}_{j0} = \Gamma^i_{0j} = \delta_{ij}\Big(H + \dot{\Phi}\Big),\\
&\Gamma^i{}_{jk} = i\Phi(\delta_{ij}k_k + \delta_{ik}k_j - \delta_{jk}k_i).
\end{aligned} \tag{1.74}$$

The Riemann tensor Eq. 1.15 can be expressed from these as

$$\begin{aligned}
R_{00} &= \Gamma^\alpha{}_{00,\alpha} - \Gamma^\alpha{}_{0\alpha,0} + \Gamma^\alpha{}_{\beta\alpha}\Gamma^\beta{}_{00} - \Gamma^\alpha{}_{\beta0}\Gamma^\beta{}_{0\alpha}\\
&= -\frac{3}{a}\ddot{a} - \frac{k^2}{a^2}\Psi - 3\ddot{\Phi} + 3H(\dot{\Psi} - \dot{\Phi}),\\
R_{ij} &= \delta_{ij}\Bigg[\big(2a^2H^2 + a\ddot{a}\big)(1 + 2\Phi - 2\Psi) + a^2H(6\dot{\Phi} - \dot{\Psi}) + a^2\ddot{\Phi} + k^2\Phi\Bigg]\\
&\quad + k_ik_j(\Phi + \Psi),
\end{aligned} \tag{1.75}$$

while at linear order, its determinant, the Ricci scalar, reads:

$$\mathcal{R}^{(1)} = -12\Psi\left(H^2 + \frac{\ddot{a}}{a}\right) + \frac{2k^2}{a^2}\Psi + 6\ddot{\Phi} - 6H(\dot{\Psi} - 4\dot{\Phi}) + \frac{4k^2}{a^2}\Psi. \tag{1.76}$$

With these ingredients, one can derive the Einstein equations Eq. 1.18 at first order. First, we focus on the time-time component. The left-hand side part is straightforwardly obtained from the Einstein tensor definition Eq. 1.13. The right-hand-side part requires evaluating the stress-energy tensor for every species.

Only keeping the leading order term, the *time-time component* of the Einstein equation reads:

$$-3H\dot{\Phi} + 3\Psi H^2 - \frac{k^2\Phi}{a} = -4\pi G[\rho_{\rm m}\delta_{\rm m} + 4\rho_{\rm r}\Theta_{\rm r,0}], \tag{1.77}$$

with $\rho_{\rm m}\delta_{\rm m} \equiv \rho_{\rm c}\delta_{\rm c} + \rho_{\rm b}\delta_{\rm b}$ and $\rho_{\rm r}\Theta_{{\rm r},n} \equiv \rho_\gamma\Theta_n + \rho_\nu\Theta_n$. In terms of conformal time, $\dot{a} \to \frac{a'}{a}$, so Eq. 1.77 becomes

$$k^2\Phi + 3\frac{a'}{a}\left(\Phi' - \Psi\frac{a'}{a}\right) = 4\pi Ga^2[\rho_{\rm m}\delta_{\rm m} + 4\rho_{\rm r}\Theta_{\rm r,0}]. \tag{1.78}$$

This is the relativistic equivalent of the Poisson equation in an expanding background.

The second non-redundant Einstein equation is the *longitudinal trace-less space-space* part of the Einstein equation (that can be obtained by applying $\left(\hat{k}_i\hat{k}^j - \frac{1}{3}\delta_i^j\right)$ to Eq. 1.18). After developing all spatial terms entering in Einstein's equation, one obtains the following relation:

$$k^2(\Phi + \Psi) = -32\pi G a^2[\rho_{\rm r}\Theta_{\rm r,2}], \tag{1.79}$$

which indicates that the two scalar potentials are equal and opposite to the radiation quadrupole, with a dominant contribution from neutrinos before the recombination, given that the photons were tightly coupled to baryons at this time.

1.2.2.3 Initial conditions

Eqs. 1.68 & 1.67 (photons density), Eq. 1.69 (neutrinos), Eqs. 1.72 & 1.73 (baryons), Eqs. 1.70 & 1.71 (dark matter), and Eqs. 1.78 & 1.79 (Einstein's scalar equations) form a equation system that should be solved jointly to determine the evolution of the scalar perturbation in the primordial plasma.

To do so, one should fix its initial conditions, requiring rewriting the system at early times, *i.e.* when the modes of interest were *super-horizon*, meaning $k\eta \ll 1$. In this regime, we notice that for any quantity q, $\dot{q} \simeq \frac{q}{\eta} \ll q$. Every perturbation has much larger wavelengths than the sound horizon, so the whole Universe is homogeneous. Hence, at early times, the monopole dominates: $\Theta_0 \gg \Theta_n$ and $\mathcal{N}_0 \gg \mathcal{N}_n$ for any $n > 0$.

Then, Eqs. 1.68 and 1.69 become

$$\Theta_0' + \Phi' = 0, \tag{1.80}$$
$$\mathcal{N}_0' + \Phi' = 0, \tag{1.81}$$

while Eqs. 1.71 and 1.71 become:

$$\delta_{\rm c}' + 3\Phi' = 0, \tag{1.82}$$
$$\delta_{\rm b}' + 3\Phi' = 0, \tag{1.83}$$

Since they are comparable to $\Theta_1 \ll \Theta_0$, we consider that $v_{\rm c} = v_{\rm b} = 0$ at early times[4].

Hereafter, as illustrated in Fig. 1.1 and corroborated by observations, we assume that we lie in a radiation-dominated era at early times *i.e.* $\rho_{\rm m}\delta_{\rm m} \ll \rho_{\rm r}\Theta_{\rm r,0}$. We have mentioned in Sec. 1.1 that under such a condition, the conformal time is $\eta \propto t^{1/2} \propto a$, so $\frac{\dot{a}}{a} = \frac{1}{\eta}$. Then, Eq. 1.78 becomes

$$\begin{aligned}\frac{\Phi'}{\eta} - \frac{\Psi}{\eta^2} &= \frac{16\pi G\rho a^2}{3}\Theta_{\rm r,0} \\ &= \frac{2}{\eta^2}\Theta_{\rm r,0},\end{aligned} \tag{1.85}$$

[4]More precisely, one can show from Eq. 1.73 that the superhorizon behaviour for the dipole is:

$$\Theta_1 = \mathcal{N}_1 = \frac{iv_{\rm b}}{3} = \frac{iv_{\rm c}}{3} = -\frac{k}{6aH}\Phi. \tag{1.84}$$

where the last step follows from the Friedmann equation and the fact that the energy budget is dominated by $\rho_{\rm r}$. Multiplying this equation by η^2, differentiating *wrt.* η and using Eq. 1.81 leads to

$$\Phi''\eta + \Phi' - \Psi' = -2\Phi'. \tag{1.86}$$

Neglecting $\Theta_{r,2}$ in Eq. 1.79 implies that $\Phi = -\Psi$, and Eq. 1.86 becomes

$$\Phi''\eta + 4\Phi' = 0. \tag{1.87}$$

We keep the non-decaying solution where Φ is constant with respect to η.

An important assumption we make from now on is the adiabaticity of the initial conditions, *i.e.* perturbations for all species at a given scale η_i share the same spatial phase: $\forall i, j, \delta(n_i/n_j) = 0$ [13]. This assumption is natural if the process creating the perturbations treats all species similarly, as in the simplest inflation models, as we will see in Chapter 3. The coherent perturbation among the species is called the *curvature perturbation* $\mathcal{R}$, and the associated density contrasts satisfy:

$$\frac{1}{4}\delta_\gamma = \frac{1}{4}\delta_\nu = \frac{1}{3}\delta_{\rm b} = \frac{1}{3}\delta_{\rm c}. \tag{1.88}$$

Hence, at initial time η_i, the following relations are satisfied

$$\Theta_0(\eta_i) = \mathcal{N}_0(\eta_i). \tag{1.89}$$

Consequently, Eq. 1.85 becomes

$$\Phi(\eta_i) = 2\Theta_0(\eta_i). \tag{1.90}$$

Combining Eqs. 1.81 and 1.82 leads to

$$\dot{\delta}_{\rm c} = \dot{\delta}_{\rm b} = 3\dot{\Theta}_0, \tag{1.91}$$

which we can integrate for adiabatic perturbations $\frac{n_{\rm c}}{n_\gamma} = \frac{n_{\rm c}^{(0)}}{n_\gamma^{(0)}} \frac{1+\delta}{1+3\Theta_0} = \text{cst}$ (similarly for baryons):

$$\delta = \delta_{\rm b} = 3\Theta_0. \tag{1.92}$$

To summarise, the adiabatic initial conditions for the scalar perturbations are entirely described by the curvature perturbation:

$$\mathcal{R} = \delta_{\rm b} = \delta_{\rm c} = 3\Theta_0 = 3\mathcal{N}_0 = \frac{3}{2}\Phi = -\frac{3}{2}\Psi. \tag{1.93}$$

This gauge-invariant quantity remains constant on super-horizon scales in the absence of *entropy perturbations* (*i.e.* non-adiabatic), justifying its importance: it can be used to propagate the perturbations throughout the super-horizon regime from a primordial era to the subsequent cosmological eras. Super-horizon-scale curvature perturbations are *frozen*.

The power spectrum of these initial conditions, $\mathcal{P}_\mathcal{R}(k)$, defined as

$$\left\langle \mathcal{R}(\vec{k})\mathcal{R}^*\left(\vec{k}'\right)\right\rangle\Big|_{\eta_i} \equiv (2\pi)^3\delta^{(3)}\left(\vec{k} - \vec{k}'\right)\mathcal{P}_\mathcal{R}(k) \tag{1.94}$$

is a key observable of cosmology: it provides the seed distribution for all inhomogeneities in our Universe. This spectrum is often parametrised as a Taylor series expansion around some *pivot scale* k_* [14, 15]:

$$\mathcal{P}_\mathcal{R}(k) = A_\mathrm{s}\left(\frac{k}{k_*}\right)^{n_\mathrm{s}-1+\frac{1}{2}\alpha_\mathrm{s}\ln(k/k_*)+\frac{1}{6}\beta_\mathrm{s}\ln^2(k/k_*)+\ldots} \tag{1.95}$$

where A_s is called the scalar *amplitude*, n_s the *spectral index* or *tilt* , α_s the *running*, β_s *running of the running*.

1.2.3. Vector perturbations

Vector perturbations cannot be simply associated with density fluctuations. Even if they can be produced in particular scenarios, they rapidly decay in most of the early Universe models, especially inflationary ones [7] on which we focus in Chapter 3. Hence, in the following, we assume $\omega_i = 0$ in Eq. 1.44.

1.2.4. Tensor perturbations

Contrary to scalar gravitational perturbations, gravitational waves are not directly coupled to the density of the various species. However, they induce an additional contribution to the photon's anisotropies, both in terms of temperature and of polarisation [9, 16–18]. Working with Θ_P, which only describes the polarisation amplitude, was enough for Thomson scattering-induced polarisation. Instead, for gravitational waves, one has to follow the evolution of all Stokes parameters that fully specify both polarisation and temperature.

1.2.4.1 Boltzmann equation

This treatment leads to the following equation for the photon tensor distribution function:

$$\Theta^\mathrm{T}(k,\mu,\phi) = \Theta^\mathrm{T}_+(k,\mu)\left(1-\mu^2\right)\cos(2\phi) + \Theta^\mathrm{T}_\times(k,\mu)\left(1-\mu^2\right)\sin(2\phi) \tag{1.96}$$

where μ, similarly to the scalar case, is the angle between the direction of propagation of the gravitational wave and the Fourier mode considered, and where Θ^T_+ and $\Theta^\mathrm{T}_\times$ satisfy

$$\frac{d\Theta^\mathrm{T}_t}{d\eta} + ik\mu\Theta^\mathrm{T}_t + \frac{1}{2}h'_t = \tau'\left[\Theta^\mathrm{T}_t - \frac{1}{10}\Theta^\mathrm{T}_{t,0} - \frac{1}{7}\Theta^\mathrm{T}_{t,2} - \frac{3}{70}\Theta^\mathrm{T}_{t,4}\right] \tag{1.97}$$

and where $\Theta^\mathrm{T}_{t,\ell}$ is the ℓ-th multipole of the anisotropy with $t = +$ or $\times$. Similar equations can be derived for the general polarisation state. This will have a key impact on the photons emitted by the last scattering surface.

1.2.4.2 Einstein equations

The Einstein equations applied to the purely-tensor metric given by Eq. 1.44 and in absence of anisotropic stress leads to the following evolution equation for the two tensor *dof*'s of γ_{ij} that we generically denote h_+ and $h_\times$:

$$h''_t + 2\frac{a'}{a}h'_t + k^2 h_t = 0 \tag{1.98}$$

where $t = +$ or $\times$. This is the equation describing the propagation of gravitational waves in the primordial Universe. In particular, for a gravitational wave propagating in the z direction, h_+ and $h_\times$ appear in the metric through

$$\gamma_{ij} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{1.99}$$

1.2.4.3 Initial conditions

As for the scalar curvature perturbation, initial conditions are required for γ_{ij}, *i.e.* h_+ and $h_\times$ through the prescription of an initial power spectrum $\mathcal{P}_h(k)$:

$$(2\pi)^3\delta^{(3)}\left(\vec{k}-\vec{k}'\right)\mathcal{P}_h(k) \equiv \frac{1}{2}\left[\left\langle h_+(\vec{k})h_+^*\left(\vec{k}'\right)\right\rangle\Big|_{\eta_i} + \left\langle h_\times(\vec{k})h_\times^*\left(\vec{k}'\right)\right\rangle\Big|_{\eta_i}\right] \tag{1.100}$$

$$= \frac{1}{4}\left\langle \gamma_{ij}(\vec{k})\gamma_{ij}^*\left(\vec{k}'\right)\right\rangle\Big|_{\eta_i} \tag{1.101}$$

As for scalar, this spectrum can be parametrised as a Taylor series expansion around some *pivot scale* k_* [14, 15]:

$$P_h(k) = A_\mathrm{t}\left(\frac{k}{k_*}\right)^{n_\mathrm{t}+\frac{1}{2}\alpha_\mathrm{t}\ln(k/k_*)+\frac{1}{6}\beta_\mathrm{t}\ln^2(k/k_*)+\ldots} \tag{1.102}$$

where A_t is called the tensor *amplitude*, n_t the *spectral index* or *tilt* , α_t the *running*, β_t *running of the running*. We also define at the *pivot scale* the *tensor-to-scalar ratio* r as

$$r \equiv \frac{A_\mathrm{t}}{A_\mathrm{s}}. \tag{1.103}$$

A priori, $\mathcal{P}_h(k)$ is independent of $P_\mathcal{R}(k)$. These are two key central objects to this work. One of the most popular explanations for the mechanism that produced them is cosmological inflation, an era of accelerated expansion of the very early Universe before the radiation era. Its simplest descriptions allow us to predict the shapes of $\mathcal{P}_h(k)$ and of $P_\mathcal{R}(k)$ (see Chapter 3) and to explain the observed Universe (see Sec. 2.4), and hence initial conditions to the Einstein-Boltzmann system.

So far, we have a set of equations describing the evolution of species' density contrasts sourced by scalar gravitational potential inhomogeneities, their initial condition, and the effect of tensor modes for the photon distribution. Now, we use these equations to describe how were the photon distribution inhomogeneities at last scattering surface.

1.3. Picture at the recombination

In order to determine the properties of the photon emitted by the last scattering surface, we solve the Einstein-Boltzmann system for $\Theta(k, \eta_\mathrm{rec})$, where η_rec is the conformal time at the recombination. In Fourier space, approximate solutions can be obtained in the *large-scale* [19], *intermediate-scale* [20] and *small-scale* regimes. These three regimes correspond respectively to $k\eta \ll 1$, $k\eta \sim 1$ and $k\eta \gg 1$ with $\eta \sim \eta_\mathrm{rec}$.

1.3.1. Large scales: Sachs-Wolfe effect

Similarly to the initial condition situation we encountered previously, the largest scales are the ones still outside the horizon at the recombination, so $k\eta \ll 1$. Hence, we can insert Eq. 1.92 in Eq. 1.78 to obtain

$$3\frac{a'}{a}\left(\Phi' + \frac{a'}{a}\Phi\right) = 4\pi G a^2 \rho_{\rm c}\delta\left(1 + \frac{4}{3y}\right), \tag{1.104}$$

with $y \equiv \frac{a}{a_{\rm eq}} = \frac{\rho_{\rm c}}{\rho_{\rm r}}$. Since $\frac{\rm d}{{\rm d}\eta} = \frac{{\rm d}y}{{\rm d}\eta}\frac{\rm d}{{\rm d}y} = aHy\frac{\rm d}{{\rm d}y}$, one can rewrite the previous equation as

$$y\frac{{\rm d}\Phi}{{\rm d}y} + \Phi = \frac{3y+4}{6(y+1)}\delta_{\rm c}. \tag{1.105}$$

Putting aside δ, differentiating with respect to y, and using Eq. 1.82, we get the following second order differential equation for Φ:

$$\frac{{\rm d}^2\Phi}{{\rm d}y^2} + \frac{21y^2+54y+32}{2y(y+1)(3y+4)}\frac{{\rm d}\Phi}{{\rm d}y} + \frac{\Phi}{y(y+1)(3y+4)} = 0, \tag{1.106}$$

which we can solve to obtain a general expression for Φ, which can evaluate in particular for the (matter-dominated) recombination, when $y \to \infty$:

$$\begin{aligned}\Phi(y) &= \frac{\Phi(0)}{10}\frac{1}{y^3}\left(16\sqrt{1+y} + 9y^3 + 2y^2 - 8y - 16\right)\\ &\xrightarrow{y\to\infty} \frac{9}{10}\Phi(0).\end{aligned} \tag{1.107}$$

Integrating Eq. 1.79 and fixing the constant thanks to Eq. 1.90 leads to writing the photon density contrast at the recombination as

$$\Theta(k,\eta_{\rm rec}) = -\Phi(k,\eta_{\rm rec}) + \frac{3}{2}\Phi(k,\eta_i) \tag{1.108}$$

$$= -\Phi(k,\eta_{\rm rec}) + \frac{10}{9}\frac{3}{2}\Phi(k,\eta*) \tag{1.109}$$

$$= \frac{2}{3}\Phi(k,\eta_{\rm rec}). \tag{1.110}$$

Especially one can express a useful quantity for the next step, $(\theta_0 + \Psi)$. Neglecting $\Theta_{\rm r,2}$ in Eq. 1.79, it reads

$$(\Theta_0 + \Psi)(k,\eta_{\rm rec}) = -\frac{1}{6}\delta(k,\eta_{\rm rec}) \tag{1.111}$$

which follows from integrating Eq. 1.82 between η_i and $\eta_{\rm rec}$ and fixing the initial condition with Eq. 1.93. This large-scale-regime equation leads to the so-called *Sachs-Wolfe* anisotropies of the CMB that we present in Sec. 2.1.

1.3.2. Intermediate scales: oscillations in the tight coupling limit

To understand what happens at intermediate scales, we consider the limit in which $\tau \gg 1$ before the recombination and horizon size mode (*i.e.* $k\eta \sim 1$) [5, 20]. To obtain an expression for the photon density contrast multipole, multiply the photon Boltzmann Eq. 1.65 by the Legendre polynomial of order ℓ and integrate over the incident angle μ:

$$\Theta'_\ell + \frac{k}{(-i)^{\ell+1}}\int_{-1}^{1}\frac{{\rm d}\mu}{2}\mu P_\ell(\mu)\Theta(\mu) = \tau'\Theta_\ell. \tag{1.112}$$

Using the property of Legendre polynomials that $(\ell+1)P_{\ell+1}(\mu) = (2\ell+1)\mu P_\ell(\mu) - \ell P_{\ell-1}(\mu)$, we get:

$$\Theta'_\ell - \frac{k\ell}{2\ell+1}\Theta_{\ell-1} + \frac{k(\ell+1)}{2\ell+1}\Theta_{\ell+1} = \tau'\Theta_\ell. \tag{1.113}$$

We can apply this equation for the two first moments:

$$\Theta'_0 + k\Theta_1 = -\Phi', \tag{1.114}$$

$$\Theta'_1 - \frac{k\Theta_0}{3} = \frac{k\Psi}{3} + \tau'\Big[\Theta_1 - \frac{iv_\mathrm{b}}{3}\Big]. \tag{1.115}$$

To express the term v_b, let's also rewrite equation Eq. 1.73 at first order as

$$v_\mathrm{b} \simeq -3i\Theta_1 + \frac{R}{\dot\tau}\Big(-3i\Theta_1 - \frac{a'}{a}3i\Theta_1 + ik\Psi\Big), \tag{1.116}$$

which, combined with Eq. 1.115 yields

$$\Theta'_1 + \frac{\dot a}{a}\frac{R}{1+R}\Theta_1 - \frac{k}{3(1+R)}\Theta_0 = \frac{k\Psi}{3}. \tag{1.117}$$

We can now differentiate with respect to time 1.114 and replace Θ'_1 by its expression Eq. 1.117, to obtain a second order differential equation for Θ_0:

$$\Theta''_0 + \frac{\dot a}{a}\frac{R}{1+R}\dot\Theta_0 + k^2c_\mathrm{s}^2\Theta_0 = F(k,\eta), \tag{1.118}$$

with

$$F(k,\eta) \equiv -\frac{k^2}{3}\Psi - \frac{a'}{a}\frac{R}{1+R}\Phi' - \Phi'', \tag{1.119}$$

and c_s is called the *sound speed* in the plasma, defined as

$$c_\mathrm{s} \equiv \sqrt{\frac{1}{3(1+R)}}. \tag{1.120}$$

One can rewrite this last result as

$$\left\{\frac{\mathrm{d}^2}{\mathrm{d}\eta^2} + \frac{\dot R}{1+R}\frac{\mathrm{d}}{\mathrm{d}\eta} + k^2c_\mathrm{s}^2\right\}(\Theta_0+\Phi) = \frac{k^2}{3}\left(\frac{1}{1+R}\Phi - \Psi\right). \tag{1.121}$$

This last second order differential equation has a solution which can be found with Green's functions, approximating in the relativistic limit $R \to \infty$ where $kc_\mathrm{s} \simeq \frac{k}{\sqrt{3}}$:

$$\Theta_0(\eta) + \Phi(\eta) = [\Theta_0(0) + \Phi(0)]\cos(kr_\mathrm{s}) + \frac{k}{\sqrt{3}}\int_0^\eta \mathrm{d}\tilde\eta[\Phi(\tilde\eta) - \Psi(\tilde\eta)]\sin[k(r_\mathrm{s}(\eta) - r_\mathrm{s}(\tilde\eta))]. \tag{1.122}$$

with the sound horizon

$$r_\mathrm{s}(\eta) \equiv \int_0^\eta \mathrm{d}\tilde\eta c_\mathrm{s}(\tilde\eta). \tag{1.123}$$

If the first term in Eq. 1.122 dominates over the second one, one understands that this effect will yield peaks in Fourier space at $k = n\pi/r_\mathrm{s}$ with n integer.

Finally, at intermediate scales, the dipole is not negligible compared to the monopole; we can do a similar derivation as before but keeping in equations Θ_1 instead of Θ_0, and solve the resulting second-order equation:

$$\Theta_1(\eta) = \frac{1}{\sqrt{3}}[\Theta_0(0) + \Phi(0)]\sin(kr_\mathrm{s}) - \frac{k}{3}\int_0^\eta \mathrm{d}\tilde{\eta}[\Phi(\tilde{\eta}) - \Psi(\tilde{\eta})]\cos[k(r_\mathrm{s}(\eta) - r_\mathrm{s}(\tilde{\eta}))]. \tag{1.124}$$

The first term yields an out-of-phase contribution with respect to the monopole one. We shall see that Eq. 1.122 and Eq. 1.124 yields in the CMB the *acoustic oscillation* anisotropies, which are characterised by peaks in the power spectrum.

1.3.3. Small scales: diffusion damping

We need to account for the photon diffusion on the last scattering surface to describe inhomogeneities in the photon distribution at small angular scales [21]. To do so, one needs this time to account for the quadrupole in Eq. 1.113. At small scales, radiation pressure overtakes gravity, and gravitational-potential terms can be neglected. The resulting set of equations for the three first moments of Θ and for v_b reads

$$\Theta_0' + k\Theta_1 = 0, \tag{1.125}$$

$$\Theta_1' + k\left(\frac{2}{3}\Theta_2 - \frac{1}{3}\Theta_0\right) = \tau'\left(\Theta_1 - \frac{iv_\mathrm{b}}{3}\right), \tag{1.126}$$

$$\Theta_2' - \frac{2k}{5}\Theta_1 = \frac{9}{10}\tau'\Theta_2, \tag{1.127}$$

$$3i\Theta_1 + v_\mathrm{b} = \frac{R}{\tau'}\left[v_\mathrm{b}' + \frac{a'}{a}v_\mathrm{b}\right]. \tag{1.128}$$

In the tight coupling limit, all quantities oscillate so they can be written as proportional to $\exp^{i\int \mathrm{d}\eta\omega}$ with $\omega \simeq kc_\mathrm{s}$. Moreover, we are dealing with high frequencies at the smallest scales, $\dot{v}_\mathrm{b} = i\omega v_\mathrm{b} \gg \frac{\dot{a}}{a}v_\mathrm{b}$, which allows one to solve for v_b in Eq. 1.128:

$$v_\mathrm{b} = -3i\Theta_1\left(1 - \frac{i\omega R}{\tau'}\right)^{-1} \simeq -3i\Theta_1\left[1 + \frac{i\omega R}{\tau'} - \left(\frac{\omega R}{\tau'}\right)^2\right], \tag{1.129}$$

and to relate the quadrupole to the dipole and to the monopole with Eqs. 1.127 and 1.125:

$$\Theta_2 = -\frac{4k}{9\tau'}\Theta_1 = \frac{4i\omega}{9\tau'}\Theta_0, \tag{1.130}$$

Inserting these into Eq. 1.126, one obtains the following dispersion relation:

$$\omega^2 = k^2c_\mathrm{s}^2 + 2\omega\delta\omega, \tag{1.131}$$

where the sound speed c_s and the first order correction to the frequency $\delta\omega$ read:

$$c_\mathrm{s} = \frac{1}{\sqrt{3(1+R)}} \tag{1.132}$$

$$\delta\omega = -\frac{ik^2}{2(1+R)\dot{\tau}}\left(c_\mathrm{s}^2R^2 + \frac{8}{27}\right). \tag{1.133}$$

From this dispersion relation, one can rewrite more precisely the time dependence of the monopole and dipole

$$\Theta_0, \Theta_1 \sim e^{ik\int \mathrm{d}\eta c_\mathrm{s}(\eta')} e^{-\frac{k^2}{k_D^2(\eta')}}, \tag{1.134}$$

where the damping wavenumber is defined in full generality as

$$\frac{1}{k_D^2(\eta)} \equiv \int_0^\eta \frac{\mathrm{d}\tilde{\eta}}{6(1+R)n_e\sigma_T a(\tilde{\eta})} \left(\frac{R^2}{1+R} + \frac{8}{9}\right). \tag{1.135}$$

This results in a blurring at the smallest scales that corresponds to anisotropies $k > k_\mathrm{D}$, which can be viewed intuitively as the effect of the non-zero thickness of the last scattering surface. This effect is called *diffusion* or *Silk damping* [21].

Conclusion

In the two first sections of this chapter, Sec. 1.1 and Sec. 1.2, we have introduced the main equations to describe first-order perturbations on top of an expanding background. This description of our Universe already allows us to understand most of our cosmological observations. Then, in Sec. 1.3, we have depicted the Universe at the recombination, at the time of the photon decoupling, in three regimes. In the next chapter, we will focus on a particular cosmological observable, the *cosmic microwave background* that we have already briefly mentioned in Sec. 1.1. The CMB consists of photons emitted by the last scattering surface when the Universe transitioned from being opaque to transparent.

CHAPTER 2

Cosmic Microwave Background

In this chapter, we will first discuss how perturbations in the photon distribution at the recombination, $\eta_{\rm rec}$, lead to anisotropies in the CMB. After polarised photons are emitted in all directions by particles at the last-scattering surface, they travel freely in the Universe at first order, with only a few interactions with electrons. Subsequently, we will introduce the general formalism of polarised CMB and the statistics of its maps. Finally, we will provide an overview of the current constraints on the Λ-CDM model coming from the CMB.

Contents

2.1 Photon anisotropy today 28
2.1.1 From temperature contrast 28
2.1.2 From scalar polarisation 30
2.1.3 From tensor perturbations 31
2.1.4 From the reionisation 31
2.2 General formalism 31
2.2.1 Statistics of a function on a sphere 32
2.2.2 Temperature anisotropies 32
2.2.3 Polarisation, E- and B- modes 33
2.2.4 Scalar and tensor power spectra 34
2.3 Statistics of the maps 37
2.3.1 Real-space probability density function 37
2.3.2 Power spectrum estimate of a single sky 39
2.3.3 Pixelisation of the sky 40
2.4 Current observational status 41
2.4.1 First measurement, COBE, Archeops and WMAP 41

2.4.2 Impact of the Planck mission on the Λ-CDM model 42
2.4.3 SPT, ACT and BICEP 45
2.4.4 Current constraints on initial conditions 45

2.1. Photon anisotropy today

First, we discuss how the inhomogeneities in the photon temperature contrast at the recombination that we have introduced in Sec. 1.3, the scalar polarisation, the tensor perturbation and the Universe reionisation lead to anisotropies for us observers.

2.1.1. From temperature contrast

To describe the transition from tightly-coupled photons to quasi *free-streaming* ones, one needs to rewrite the photon Boltzmann equation, Eq. 1.65, as

$$e^{-ik\mu\eta+\tau}\frac{\mathrm{d}}{\mathrm{d}\eta}\left(\Theta e^{ik\mu\eta-\tau}\right)=\tilde{S}, \tag{2.1}$$

where $\tilde{S}$ is the *source term* that reads

$$\tilde{S}\equiv-\dot{\Phi}-ik\mu\Psi-\dot{\tau}\left[\Theta_0+\mu v_\mathrm{b}-\frac{1}{2}P_2(\mu)\Pi\right]. \tag{2.2}$$

Θ can be solved for by multiplying by the exponential factor and integrating between two times, η_i before the recombination and η_0:

$$\begin{aligned}\Theta(\eta_0)&=\Theta(\eta_i)e^{ik\mu(\eta_i-\eta_0)}e^{-\tau(\eta_i)}+\int_{\eta_i}^{\eta_0}\mathrm{d}\eta\tilde{S}(\eta)e^{ik\mu(\eta-\eta_0)-\tau(\eta)}\\&=\int_0^{\eta_0}\mathrm{d}\eta\tilde{S}(\eta)e^{ik\mu(\eta-\eta_0)-\tau(\eta)}\end{aligned} \tag{2.3}$$

where we have neglected the first term because $\eta_i\simeq 0$ and $\tau(\eta_i)\gg 1$. Using the identity

$$\int_{-1}^{1}\frac{\mathrm{d}\mu}{2}P_\ell(\mu)e^{ik\mu(\eta-\eta_0)}=\frac{1}{(-i)^\ell}j_\ell(k(\eta-\eta_0)), \tag{2.4}$$

where j_ℓ is the Bessel function of order l, the ℓ-th multipole of the photon distribution can be rewritten as

$$\Theta_\ell(k,\eta_0)=\int_0^{\eta_0}\mathrm{d}\eta S(k,\eta)j_\ell(k(\eta-\eta_0)), \tag{2.5}$$

with

$$S(k,\eta)\equiv e^{-\tau}\left[-\dot{\Phi}-\dot{\tau}\left(\Theta_0+\frac{\Pi}{4}\right)\right]+\frac{\mathrm{d}}{\mathrm{d}\eta}\left[e^{-\tau}\left(\Psi-\frac{iv_\mathrm{b}\dot{\tau}}{k}\right)\right]-\frac{3}{4k^2}\frac{\mathrm{d}^2}{\mathrm{d}\eta^2}\left(e^{-\tau}\dot{\tau}\Pi\right). \tag{2.6}$$

One introduces the visibility function as

$$g(\eta)\equiv-\dot{\tau}e^{-\tau}, \tag{2.7}$$

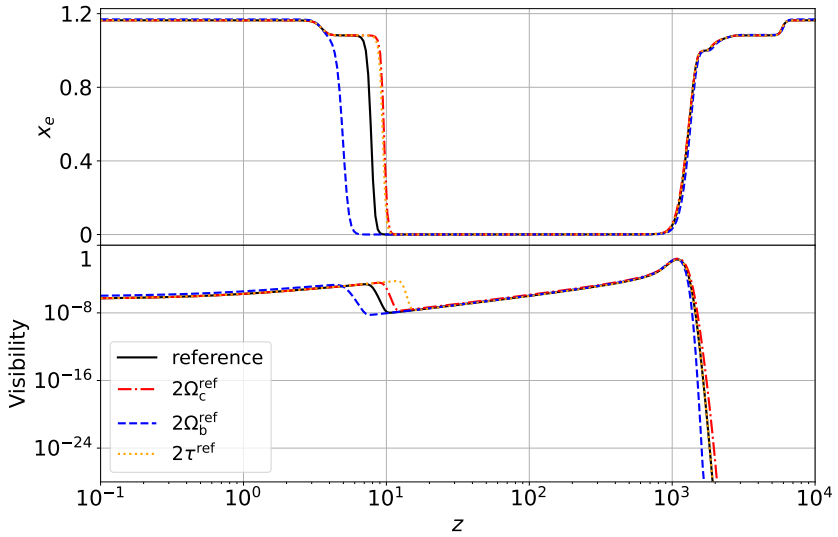

Figure 2.1: Visibility (upper panel) and ionisation fraction (lower panel) as functions of the redshift. Most electrons last scatter at the recombination, around $z \simeq 1100$ ($\eta \simeq 275$ Mpc). In this case, consistently with its definition Eq. 1.34, x_e can be greater than one because of the presence of ionised Helium atoms.

which can is the probability density that a photon is last scattered at η. The x_e, opacity and visibility functions (defined respectively in Eqs. 1.34, 1.63 and 2.7) lead to express the visibility as a function of the fraction of ionised electrons in the Universe during each era. We represent these two functions, respectively, on the upper and lower panel of Fig. 2.1, for a reference illustrated in black that corresponds to the cosmological model best fit, in red when the fraction of cold-dark matter is twice and in blue when the fraction of cold-dark matter is twice. From this figure, one sees that most photons last scatter around the recombination, which corresponds to $z_{\rm rec} \simeq 1100$ with little dependence on the underlying cosmological parameters.

One also notices a second bump in the visibility function that corresponds to a *reionisation* of the Universe [22–24] around $z \simeq 8$. We further comment on the effect of this era at the end of the section. Neglecting the small polarisation terms, one can write down various terms that contribute to the anisotropies seen today. The source function reads

$$S(k,\eta) \simeq g(\eta)\Big[\Theta_0(k,\eta) + \Psi(k,\eta)\Big] + \frac{\mathrm{d}}{\mathrm{d}\eta}\left(\frac{iv_{\rm b}(k,\eta)g(\eta)}{k}\right) + e^{-\tau}\Big[\dot{\Psi} - \dot{\Phi}\Big], \tag{2.8}$$

which, inserted in Eq. 2.5, gives

$$\begin{aligned}\Theta_\ell(k,\eta_0) \simeq &\int_0^{\eta_0} \mathrm{d}\eta\; g(\eta)\Big[\Theta_0(k,\eta) + \Psi(k,\eta)\Big] j_\ell(k(\eta_0-\eta)) \\ &- \int_0^{\eta_0} \mathrm{d}\eta\; g(\eta)\frac{iv_{\rm b}(k,\eta)}{k}\frac{\mathrm{d}}{\mathrm{d}\eta} j_\ell(k(\eta_0-\eta)) \\ &+ \int_0^{\eta_0} \mathrm{d}\eta\; e^{-\tau}\Big[\dot{\Psi}(k,\eta) - \dot{\Phi}(k,\eta)\Big] j_\ell(k(\eta_0-\eta)).\end{aligned} \tag{2.9}$$

Fig. 2.1 illustrates that the visibility function is very sharp around the recombination, and it is the case relative to $(\Theta_0 + \Psi)$ or j_ℓ. Hence, to focus first on the effect of the **recombination**

on anisotropies, we approximate the visibility by a Dirac function at $\eta_{\rm rec}$, $g(\eta) \simeq \delta(\eta - \eta_{\rm rec})$. Then, using the relation $\frac{{\rm d}j_\ell}{{\rm d}x} = j_{\ell-1} - \frac{\ell+1}{x} j_\ell$ and the relation between dipole and baryon velocity Eq. 1.84, the previous expression for the photon multipole at η_0 becomes

$$\begin{aligned}\Theta_\ell(k,\eta_0) \simeq & \Big[\Theta_0(k,\eta_{\rm rec}) + \Psi(k,\eta_{\rm rec})\Big] j_\ell(k(\eta_0 - \eta_{\rm rec})) \quad \textbf{(SW) and (AO)} \\ & + 3\Theta_1(k,\eta_{\rm rec}) \left(j_{\ell-1}[k(\eta_0 - \eta_{\rm rec})] - \frac{(\ell+1) j_\ell[k(\eta_0 - \eta_{\rm rec})]}{k(\eta_0 - \eta_{\rm rec})} \right) \quad \textbf{(D)} \\ & + \int_0^{\eta_0} {\rm d}\eta \; e^{-\tau} \Big[\dot{\Psi}(k,\eta) - \dot{\Phi}(k,\eta)\Big] j_\ell(k(\eta - \eta_0)). \quad \textbf{(ISW)} \end{aligned} \tag{2.10}$$

Each term can be explained physically:

(SW) the **Sachs-Wolfe term**, due to the temperature *monopole* at the recombination. As already noticed in Eq. 1.108, it is the sum of two contributions of inverse sign: the effect of the photon underdensity $\Theta_0 < 0$ (*resp.* overdensity $\Theta_0 > 0$) itself, added to the loss $\Psi < 0$ (*resp.* gain $\Psi > 0$) in energy because of the gravitational redshift caused by the climbing (*resp.* falling) of the gravitational potential. As shown by Eq. 1.111, the latter effect is stronger than the former, and the cold CMB photons come from overdensities at the recombination: $\delta T/T \simeq -\frac{1}{6}\delta\rho/\rho$.

(AO) the monopole **acoustic oscillations** that yield anisotropies at specific scales, see Eq. 1.122.

(D) the **Doppler term**, due to temperature *dipole* at the recombination, arising because of the Doppler effect due to the oscillating motion of the last scattering surface. It also shows acoustic oscillation patterns that are out of phase with respect to monopole ones; see Eq. 1.124.

(ISW) **Integrated Sachs-Wolfe term**. The photon energy is modified because of gravitational redshift when the gravitational potentials evolve, which happens mainly during the radiation era (*early ISW*) and dark energy era (*late ISW*).

2.1.2. From scalar polarisation

On top of this picture, the additional temperature and linear polarisation perturbations due to Thomson scattering, which appears in Eq. 2.6, also lead to anisotropies in temperature and in polarisation. One can solve the source term for polarisation similarly to what we have presented for temperature. For instance, in standard cosmological models, [25] shows analytically that the CMB polarisation is proportional to the temperature dipole at the recombination,

$$\Theta_P\left(\eta_{\rm rec}\right) \simeq -0.17\left(1-\mu^2\right) \Delta\eta_{\rm rec} k \Theta_1\left(\eta_{\rm rec}\right) \tag{2.11}$$

demonstrating that the polarised signal arises from the velocity field of baryons (see Eq. 1.84). Thus, the unpolarised regions of the CMB coincide with extrema in potential, hence extrema in temperature (while the regions of maximum polarisation amplitude lie in the steepest regions of the gravitational potential, where the baryon velocity is maximum and the temperature is the average one). This also has an incidence on the typical polarisation patterns due to scalar perturbations, which have to be *curl-free*; see Sec. 2.2.4 in which we formally introduce these so-called *E-modes*.

2.1.3. From tensor perturbations

We have also seen that tensor perturbations also impact the photon emission from the last scattering surface; see Eq. 1.97. They lead to anisotropies in temperature and polarisation, and they constitute the only way to produce parity-odd polarisation patterns with curl. We precise what it implies in terms of angular power spectra in Sec. 2.2.4.

2.1.4. From the reionisation

Additionally, the first stars appear during the late-time Universe, typically around $z \simeq 8$ [22–24], and inject energy into the surrounding medium. Consequently, the hydrogen atoms that became neutral after the recombination reionise. This is illustrated on the upper panel of Fig. 2.1, where the transition from neutral to ionised Universe is done around $z \simeq 8$ depending on the underlying cosmology. This effect depends strongly on the astrophysical history of the reionisation, which is not currently well known. Fig 2.1 is based on a simple hyperbolic tangent transition with two degrees of freedom, that one can parameterise fixing the reionisation duration $\Delta z = 0.5$ and the total opacity along the line of sight $\tau \equiv \tau(\eta_0)$ [26]. Indeed, the presence of new free electrons in the light of sight of the photons induces a sudden increase in the visibility function.

As a first consequence, since all primordial-anisotropy terms come with a $e^{-\tau}$ damping factor, these are blurred by the new reionisation scattering events at all scales. Secondly, new polarisation anisotropies are sourced by Thomson scattering of the quadrupole as seen by the reionisation electrons [27], exactly by the same mechanism as the one described in Eq. 1.67. Because the reionisation has occurred lately, these are only visible in the CMB largest scales. Using trigonometry, one can estimate the maximum angle on the today sky associated with anisotropies sourced at the reionisation as the ratio between the physical distance to the recombination and the angular diameter distance of the reionisation $\theta_{\rm reio} \simeq \frac{\int_{z_{\rm reio}}^{z_{\rm rec}} {\rm d}z/H(z)}{\int_0^{z_{\rm reio}} {\rm d}z/H(z)}$. In a matter-dominated Universe, the Friedmann equation Eq. 1.22 approximately yields this angle as a function of $z_{\rm reio}$, and independently of the other cosmological parameters: $\theta_{\rm reio} \simeq \left[1-(1+z_{\rm reio})^{-1/2}\right]\left[(1+z_{\rm reio})^{-1/2} - z_{\rm rec}^{-1/2}\right] \simeq 0.2$ for $z_{\rm reio} = 8$ which translates to $\ell_{\rm max} \simeq \pi/\theta_{\rm reio} \simeq 16$ for $z_{\rm reio} = 8$.

We have now introduced all we need to describe a perturbed Universe and the resulting cosmic microwave background emitted by the recombination. These are the elementary blocks of the so-called cosmological Λ-*CDM model*, which we introduce in Eq. 2.4.

2.2. General formalism

We have seen that CMB anisotropies carry a wealth of information about cosmology. More precisely, they provide information through their statistics, *i.e.*, their angular distribution over the sky. We have already used the multipole expansion of the photon temperature contrast. In this section, we further introduce the tools needed to describe the statistics of the CMB in temperature and polarisation: spherical harmonic and coefficients, power spectra, Stokes parameters, E- and B- polarisation. In a last discussion, Sec. 2.2.4, we present the predictions on the power-spectra done by the standard model of cosmology we have been introducing so far.

2.2.1. Statistics of a function on a sphere

We have seen that the perturbed Universe can be viewed as a random field with uncorrelated Fourier modes. Moreover, the most standard realisations of inflation, especially slow-roll inflation, predict fluctuations that are very nearly Gaussian (see Chapter 3). Since the CMB probes fluctuations mostly in the linear regime, this Gaussian property is preserved, and the CMB looks very Gaussian.

Any function on a sphere can be decomposed in a spherical harmonics basis,

$$f(\hat{n}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} f_{\ell m} Y_{\ell m}(\hat{n}) \tag{2.12}$$

where ℓ is the multipole, $f_{\ell m}$ are the non-local spherical harmonics coefficients, while $Y_{\ell m}$ are the basis functions. They satisfy

$$Y^*_{\ell m} = (-1)^m Y_{\ell -m}, \tag{2.13}$$

and they are orthonormal:

$$\int \mathrm{d}\hat{n} Y_{\ell m}(\hat{n}) Y^*_{\ell' m'}(\hat{n}) = \delta_{\ell \ell'} \delta_{m m'}. \tag{2.14}$$

Consequently, the harmonic coefficients of a real spherical function satisfy

$$f^*_{\ell m} = (-1)^m f_{\ell -m}, \tag{2.15}$$

and can be written as

$$f_{\ell m} = \int d\hat{n} f(\hat{n}) Y^*_{\ell m}(\hat{n}). \tag{2.16}$$

For a Gaussian field, the spherical coefficients are such that $\langle f_{\ell m} \rangle = 0$. For a statistically isotropic field, they additionally satisfy:

$$\langle f_{\ell m} f^*_{\ell' m'} \rangle = C_\ell \delta_{\ell \ell'} \delta_{m m'}. \tag{2.17}$$

C_ℓ is called the *angular power spectrum* of f. The $\langle\rangle$ represents the ensemble average.

2.2.2. Temperature anisotropies

For instance, let's assume that an observer at coordinates $\vec{x}$ and conformal time η observes the CMB temperature in a direction $\hat{n}$, $T(\vec{x}, \eta, \hat{n})$. He sees a photon contrast defined as

$$\Theta(\vec{x}, \eta, \hat{n}) \equiv \frac{T(\vec{x}, \eta, \hat{n}) - T(\eta)}{T(\eta)} = \int \frac{\mathrm{d}k^3}{(2\pi)^3} e^{i \vec{k} \vec{x}} \Theta(\vec{k}, \hat{n}, \eta), \tag{2.18}$$

where $T(\eta)$ is the average CMB temperature at time conformal η, and where $\Theta_\ell(k)$ is typically given by Eq. 2.5. This temperature contrast can be written in the $Y_{\ell m}(\hat{n})$ base

$$\Theta(\vec{x}, \eta, \hat{n}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a^T_{\ell m}(\vec{x}, \eta) Y_{\ell m}(\hat{n}), \tag{2.19}$$

where the temperature spherical coefficients read

$$a^T_{\ell m}(\vec{x}, \eta) = \int \mathrm{d}\Omega Y^*_{\ell m}(\hat{n}) \Theta(\vec{x}, \eta, \hat{n}), \tag{2.20}$$

and have the Gaussian properties given by Eq 2.17. From these equations, analytical derivation can be performed. For instance, the low-ℓ C^{TT}_ℓ, which is dominated by the Sachs-Wolfe effect, is derived in [5]. Hereafter, for conciseness, we consider the field today on Earth and drop the arguments $\vec{x}, \eta$.

2.2.3. Polarisation, E- and B- modes

To describe the polarisation state of an electromagnetic radiation [28], one often uses the four Stokes parameters I, Q, U, V, where I in our case coincides with Θ as defined in Eq. 2.18, and Q and U represent the two degrees of freedom of linear polarisation in an arbitrary base. The circular polarisation, V, is ignored in cosmology because no usual physical mechanism is able to generate it. Hereafter we denote the polarisation amplitude as P:

$$P = \sqrt{Q^2 + U^2} \tag{2.21}$$

Contrary to I or P, which are invariant by a rotation ψ around $\hat{n}$, Q and U transform as

$$\begin{aligned} Q' &= Q\cos 2\psi + U \sin 2\psi, \\ U' &= -Q \sin 2\psi + U \cos 2\psi. \end{aligned} \tag{2.22}$$

From Q and U, one can build two spin-2 quantities $(Q \pm iU)$ that transform under rotation as:

$$(Q \pm iU)'(\hat{n}) = e^{\mp 2i\psi}(Q \pm iU)(\hat{n}). \tag{2.23}$$

Similarly to the temperature, these quantities can be decomposed in spherical harmonics:

$$(Q \pm iU)(\hat{n}) = \sum_{\ell m} a_{\pm 2,\ell m}\ {}_{\pm 2}Y_{\ell m}(\hat{n}), \tag{2.24}$$

where ${}_{\pm 2}Y_{\ell m}(\hat{n})$ are the spin-± 2 harmonic functions and where the associated coefficients read

$$a_{\pm 2,\ell m} = \int d\Omega\ {}_{\pm 2}Y^*_{\ell m}(\hat{n})(Q \pm iU)(\hat{n}). \tag{2.25}$$

The so-called E- and B- modes are another way of completely specifying linear polarisation. Their spherical coefficients can be expressed as a linear combination of the spin-± 2 coefficients:

$$\begin{aligned} a^E_{\ell m} &= -\left(a_{2,\ell m} + a_{-2,\ell m}\right)/2, \\ a^B_{\ell m} &= i\left(a_{2,\ell m} - a_{-2,\ell m}\right)/2. \end{aligned} \tag{2.26}$$

As for temperature, E- and B- modes are usually statistically isotropic Gaussian fields over the sphere. If one defines the following vector:

$$\mathbf{A}_{\ell m} = \begin{pmatrix} a^{\mathrm{T}}_{\ell m} \\ a^{\mathrm{E}}_{\ell m} \\ a^{\mathrm{B}}_{\ell m} \end{pmatrix}, \tag{2.27}$$

for most standard cosmological models, it follows a multivariate Gaussian distribution involving C_ℓ's. In full generality, its covariance matrix is written as

$$\mathbf{C}^{XY}_{\ell\ell'} = \left\langle \left(\mathbf{A}^X_{\ell m} - \left\langle \mathbf{A}^X_{\ell m}\right\rangle\right)\left(\mathbf{A}^Y_{\ell' m'} - \left\langle \mathbf{A}^Y_{\ell' m'}\right\rangle\right)^{\mathrm{H}} \right\rangle, \tag{2.28}$$

where the H exponent stands for the Hermitian transpose. As for temperature only, each scale evolves independently so $C_{\ell \neq \ell'} = 0$. One can reduce the information in the covariance matrix to a more convenient n_ℓ-shaped vector of 3×3 matrices and $\mathbf{A}_{\ell m} \sim \mathcal{N}_3(0, \mathbf{C}_\ell)$ with

$$\mathbf{C}_\ell = \begin{pmatrix} C^{TT}_\ell & C^{TE}_\ell & C^{TB}_\ell \\ C^{TE}_\ell & C^{EE}_\ell & C^{EB}_\ell \\ C^{TB}_\ell & C^{EB}_\ell & C^{BB}_\ell . \end{pmatrix} \tag{2.29}$$

where the power spectra C_ℓ^{XY} are defined similarly as Eq. 2.17

$$\left\langle a_{\ell m}^{X} a_{\ell' m'}^{Y*} \right\rangle = \delta_{\ell\ell'}\delta_{mm'} C_\ell^{XY}. \tag{2.30}$$

where $X, Y \in T, E, B$.

Hereafter, instead of representing Q and U maps, we will sometimes represent the E and B maps, defined with Eq. 2.12, hence simply treating E and B-modes' as independent scalar fields.

Fundamentally, $E-$ and $B-$modes are of opposite parity [29]: under the transformation $\hat{n} \to -\hat{n}$, the E modes remain unchanged, whereas the B modes change sign. This results in characteristic patterns for these two. Around a maximum of polarisation amplitude (equivalently in temperature), E modes are *curl-free* while B modes are *divergence-free*, as illustrated in Fig. 2.2. This property explains the naming of the two modes by analogy with electromagnetic. Near a maximum in polarisation amplitude, E-modes trace whether the polarisation vectors point towards ($E < 0$) or circle around ($E > 0$) the maximum, while B-modes are rotated by 45° with respect to E-modes (forming a basis), $B < 0$ corresponds to polarisation vectors turning in the clockwise sense, while it is the contrary for $B > 0$.

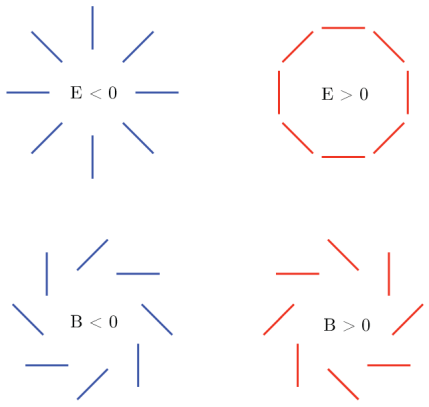

Figure 2.2: From [29]. TOP: Purely E-modes polarisation pattern BOTTOM: Purely B-modes polarisation pattern.

2.2.4. Scalar and tensor power spectra

Beyond this important property, the reason why this decomposition is crucial for CMB is that cosmological E- and B- modes are not sourced by the same physics. Especially, scalar perturbations are only able to generate E-mode patterns. Gravitation is a radial force so the velocity field of baryons is irrotational near gravitational potentials. Since we have seen with Eq. 2.11 that Thomson scattering polarisation results from moving baryons, scalar-induced polarisation consists in E-modes at first order in perturbation theory. One can show it formally by differentiating the cosmological predictions for the spectra, see [28]. The scalar contribution to

these various spectra is written generically as

$$\begin{aligned} C_\ell^{TT,EE\ (\mathrm{S})} &= (4\pi)^2 \int k^2 dk P_{\mathcal{R}}(k) \left[\Delta_\ell^{T,E\ (\mathrm{S})}(k)\right]^2 \\ C_\ell^{TE,\ (\mathrm{S})} &= (4\pi)^2 \int k^2 dk P_{\mathcal{R}}(k) \Delta_\ell^{T,\ (\mathrm{S})}(k) \Delta_\ell^{E,\ (\mathrm{S})}(k) \\ C_\ell^{BB,EB,TB\ (\mathrm{S})} &= 0. \end{aligned} \tag{2.31}$$

where $P_{\mathcal{R}}(k)$ is the primordial scalar power spectrum introduced in Eq. 1.94, and $\Delta_\ell^{T,E\ (\mathrm{S})}(k)$ are called the *transfer functions* for scalar temperature or E modes, that can be obtained from integrating the Boltzmann equations Eqs. 1.67 and 1.67. Instead, the tensor modes lead to the following contributions

$$\begin{aligned} C_\ell^{TT,EE,BB\ (\mathrm{T})} &= (4\pi)^2 \int k^2 dk \mathcal{P}_h(k) \left[\Delta_\ell^{T,E,B\ (\mathrm{T})}(k)\right]^2 \\ C_\ell^{TE,\ (\mathrm{T})} &= (4\pi)^2 \int k^2 dk \mathcal{P}_h(k) \Delta_\ell^{T,\ (\mathrm{T})}(k) \Delta_\ell^{E,\ (\mathrm{T})}(k) \\ C_\ell^{EB,TB\ (\mathrm{T})} &= 0. \end{aligned} \tag{2.32}$$

where $\mathcal{P}_h(k)$ is the primordial tensor power spectrum introduced in Eq. 1.100, and $\Delta_\ell^{T,E,B\ (\mathrm{T})}(k)$ are the transfer function for tensor temperature, E- or B- modes, that can be obtained from integrating the Boltzmann equations Eqs. 1.97. The complete expressions are given in [28]. The most important conclusion is that B modes are only generated from tensor perturbations.

One also deduces that in most standard cosmological scenarios, Eq. 2.29 reduces to

$$\mathbf{C}_\ell = \begin{pmatrix} C_\ell^{TT} & C_\ell^{TE} & 0 \\ C_\ell^{TE} & C_\ell^{EE} & 0 \\ 0 & 0 & C_\ell^{BB}. \end{pmatrix}. \tag{2.33}$$

We represent these various theoretical power spectra in Fig. 2.3, for which we have adopted a fiducial cosmology compatible with the current observations apart from the tensor-to-scalar ratio, which we take equal to 0.1 for this particular illustration. The upper-left panel illustrates TT, the upper-right: EE, the lower-left: BB and the lower-right: TE. $C_\ell^{(\mathrm{S})}$ are illustrated in blue while $C_\ell^{(\mathrm{T})}$ are illustrated in orange.

Notice that the lensing, which we introduce later on in Sec. 4.1.1, is not included. Even though constraints on tensor modes can come from measurements of TT, TE, and BB, the preferred channel to detect them is BB.

One can recognise some of the effects we have qualitatively described earlier:

- The effect of Thomson rescattering at the **reionisation** that we have introduced in Sec. 4.4.2.1 is clearly identified in EE, TE and BB as a bump between $\ell = 2$ and $\ell = 10$ (it could go up to $\ell = 25$ for other fiducial cosmologies). It scales with τ^2 in EE and BB and with τ in TE. Apart from these bumps, τ is almost entirely degenerated with A_{s} through the relation $A_{\mathrm{s}} e^{-2\tau}$, since both of these parameters act as a global amplitude for the spectra.
- The **Sachs-Wolfe** effect, introduced in Eq. 1.111 and appearing in the first term of Eq. 2.10, leads to the other features below $\ell = 50$ in TT, TE and EE, especially an almost scale invariant D_ℓ^{TT} with amplitude $\frac{8}{25} A_{\mathrm{s}}$ [11]. A detection of the Sachs-Wolfe effect is a first but non definitive hint for superhorizon scale fluctuations.

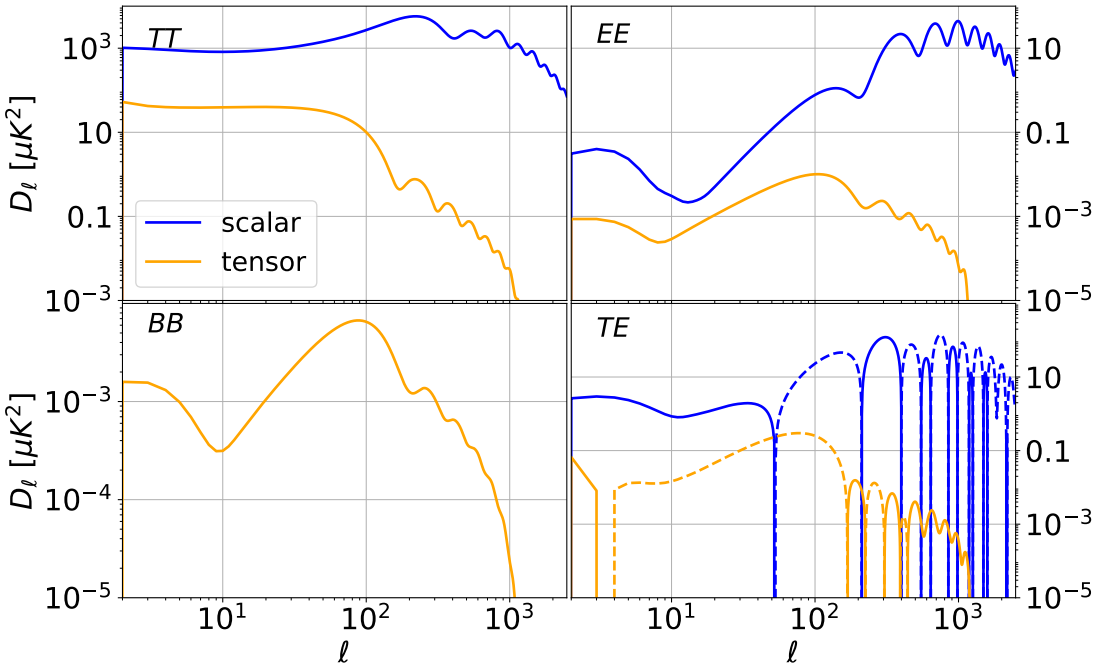

Figure 2.3: Tensor and scalar power spectra for $r = 0.1$ and the fiducial in 4.1. UPPER-LEFT: TT, UPPER-RIGHT: EE, LOWER-LEFT: BB LOWER-RIGHT: TE. Scalar-sourced power spectra are illustrated in blue while tensor-sourced power ones are illustrated in orange.

- Instead, the **negative** TE between $50 < \ell < 150$ is a distinctive signature of superhorizon fluctuations at decoupling [30, 31]: the few non-standard mechanisms that are able to generate temperature fluctuations at the largest scales in the CMB without invoking superhorizon perturbations also predict positive C_ℓ^{TE} in this multipole region. We comment more on the detection of this signature in Sec. 2.4 and on the resulting *horizon arguments* for inflation in Sec. 3.1.1.
- The scalar-induced **acoustic oscillations**, due to the monopole and the Doppler terms in Eq. 2.10, see also Eqs. 1.122 and 1.124, are responsible for the peaks visible in all scalar spectra. They cannot simply be identified to the cosine and sine peaks in Fourier space because of the presence of the Bessel functions in the integral in Eq. 2.10. Peaks in EE and TT are out-of-phase because EE results from the temperature dipole, which is itself out-of-phase with the monopole, which dominates the contribution in TT.
- The **tensor spectra** only arises because of the gravitational waves spectrum, which presents its own acoustic oscillation peaks.
- The **integrated Sachs-Wolfe** effect that occurred during the radiation era leads to increasing the amplitude of the spectra at the same range of multipole, peaking around $\ell = 100$. The late-time Sachs-Wolfe effect will induce the spectral enhancement at very large scales in the next billion years.
- The **diffusion damping** that we have introduced in Sec. 1.3.3, which is due to the thickness of the last scattering surface, is the reason why the anisotropies are erased at small scales and why the spectra fall off at large ℓ. The temperature spectrum is the most

affected.

- One can show from Eqs. 2.31 and 2.32 that at first order, even after the various tight-coupled effects, the spectra conserve the information of the primordial spectra. In particular, the scalar (*resp.* tensor) spectra are proportional to A_s (*resp.* A_t) and roughly have a global tilt of n_s (*resp.* n_t).

2.3. Statistics of the maps

So far, we have only discussed theoretical predictions. In this section, we explain how the statistical definitions introduced in Sec. 2.2 can be adapted to the study of a measured CMB map (a single pixelised sky that contains limited information). We start introducing the pixel probability density function, which can be expressed as a function of the C_ℓ. We then introduce the first and most natural way to estimate a power spectrum on a full (but unique) sky before introducing a correction to this estimator that accounts for the finite size of the pixels.

2.3.1. Real-space probability density function

This discussion aims to derive a complete expression of the statistics of the CMB in map space, including polarisation. We have three Stokes degrees of freedom per pixel to treat: I, Q and U. Assuming no noise, the n pixels of a CMB map form a vector $\mathbf{m}$ of shape $3n$, which follows a multivariate Gaussian distribution,

$$f_{C_\ell}(\mathbf{m}) = \frac{1}{\sqrt{|2\pi\mathbf{M}|}} \exp\left(-\frac{1}{2}\mathbf{m}^T\mathbf{M}^{-1}\mathbf{m}\right) \tag{2.34}$$

where $\mathbf{M} = \langle \mathbf{m}\ \mathbf{m}^T \rangle$ is the associated covariance matrix of shape $(3n, 3n)$. This covariance matrix depends on the cosmology through C_ℓ, which typically should include polarised spectra, but also on the referential convention used to define the polarisation component of $\mathbf{m}$. Pixel per pixel, this covariance matrix reduces to n^2 (3, 3) matrices:

$$M_{ij} = \mathbf{R}(\alpha_{ij})\, \mathcal{S}(\hat{n}_i \cdot \hat{n}_j)\, \mathbf{R}(\alpha_{ji})^T, \tag{2.35}$$

where the rotation matrix R of shape $(3, 3)$ accomplishes a rotation into a global reference frame where the reference directions are meridians. In the IQU convention, this change of referential is operated by

$$\mathbf{R}(\alpha_{ij}) \equiv \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos 2\alpha_{ij} & \sin 2\alpha_{ij} \\ 0 & -\sin 2\alpha_{ij} & \cos 2\alpha_{ij} \end{pmatrix}, \tag{2.36}$$

where α_{ij} is the angle between the axis of the meridian passing by i, and the axis of the great circle containing i and j, which computation is further detailed in [32]. The matrix $\mathcal{S}$ is the signal covariance matrix of shape $(3, 3)$ and only depends on cosmology through the angular separation between the two pixels $\hat{n}_i \cdot \hat{n}_j$:

$$\mathcal{S}(\hat{n}_i \cdot \hat{n}_j) \equiv \begin{pmatrix} \langle I_i I_j \rangle & \langle I_i Q_j \rangle & \langle I_i U_j \rangle \\ \langle I_i Q_j \rangle & \langle Q_i Q_j \rangle & \langle Q_i U_j \rangle \\ \langle I_i U_j \rangle & \langle U_i Q_j \rangle & \langle U_i U_j \rangle \end{pmatrix}, \tag{2.37}$$

Its coefficients can be written as a combination of theoretical power spectra and Legendre polynomials[1]. For instance, using Eqs. 2.12, 2.17 with $f(\hat{n}) = I(\hat{n}) = \Delta T(\hat{n})$ and the corresponding harmonic coefficients $a^T_{\ell m}$, one obtains

$$\begin{aligned}
\langle I_i I_j \rangle &\equiv \langle I\left(\hat{n}_i\right) I\left(\hat{n}_j\right) \rangle && (2.38)\\
&= \sum_{\ell m} \sum_{\ell' m'} \underbrace{\langle a^T_{\ell m} a^{T*}_{\ell' m'} \rangle}_{C^{TT}_\ell \delta_{\ell\ell'} \delta_{mm'}} Y_{\ell m}(\hat{n}_i) Y^*_{\ell' m'}\left(\hat{n}_j\right) && (2.39)\\
&= \sum_{\ell} \sum_{m} Y_{\ell m}(\hat{n}_i) Y^*_{\ell m}\left(\hat{n}_j\right) C^{TT}_\ell && (2.40)\\
&= \sum_{\ell} \frac{2\ell+1}{4\pi} P^{(0)}_\ell\left(\hat{n}_i \cdot \hat{n}_j\right) C^{TT}_\ell && (2.41)
\end{aligned}$$

Introducing $z \equiv \hat{n}_i \cdot \hat{n}_j$, the coefficients in Eq. 2.37 write:

$$\begin{aligned}
\langle I_i I_j \rangle (z) &\equiv \sum_\ell \frac{2\ell+1}{4\pi} P^{(0)}_\ell(z) C^{TT}_\ell,\\
\langle I_i Q_j \rangle (z) &\equiv -\sum_\ell \frac{2\ell+1}{4\pi} F^{10}_\ell(z) C^{TE}_\ell,\\
\langle I_i U_j \rangle (z) &\equiv -\sum_\ell \frac{2\ell+1}{4\pi} F^{10}_\ell(z) C^{TB}_\ell,\\
\langle Q_i Q_j \rangle (z) &\equiv \sum_\ell \frac{2\ell+1}{4\pi} \left(F^{12}_\ell(z) C^{EE}_\ell - F^{22}_\ell(z) C^{BB}_\ell\right),\\
\langle U_i U_j \rangle (z) &\equiv \sum_\ell \frac{2\ell+1}{4\pi} \left(F^{12}_\ell(z) C^{BB}_\ell - F^{22}_\ell(z) C^{EE}_\ell\right),\\
\langle Q_i U_j \rangle (z) &\equiv \sum_\ell \frac{2\ell+1}{4\pi} \left(F^{12}_\ell(z) + F^{22}_\ell(z)\right) C^{EB}_\ell,
\end{aligned} \qquad (2.42)$$

These relations applied to $z = 0$ allow one to derive the variance of the fields:

$$\langle II \rangle = \sum_\ell \frac{2\ell+1}{4\pi} C^{TT}_\ell \qquad (2.43)$$

$$\langle QQ + UU \rangle = \sum_\ell \frac{2\ell+1}{4\pi} \left(C^{EE}_\ell + C^{BB}_\ell\right). \qquad (2.44)$$

These expressions for the field variance represent a motivation for plotting

$$D_\ell \equiv \frac{\ell(\ell+1)}{2\pi} C_\ell \qquad (2.45)$$

rather than C_ℓ, choice often made. Indeed, D_ℓ gives the contribution to the variance per logarithmic range in ℓ, since $\sum_\ell \frac{2\ell+1}{4\pi} \simeq \int \mathrm{d}\log\ell \frac{\ell(\ell+1)C_\ell}{2\pi}$.

[1] $P^{(m)}_v$ is the Legendre polynomial of order m and degree v

Evaluating the polarisation terms of Eq. 2.42 requires:

$$
\begin{aligned}
F^{10}(z) &= 2\frac{\frac{\ell z}{(1-z^2)}P^{(0)}_{\ell-1}(z) - \left(\frac{\ell}{1-z^2} + \frac{\ell(\ell-1)}{2}\right)P^{(0)}_{\ell}(z)}{[(\ell-1)\ell(\ell+1)(\ell+2)]^{1/2}}, \\
F^{12}(z) &= 2\frac{\frac{(\ell+2)z}{(1-z^2)}P^{(2)}_{\ell-1}(z) - \left(\frac{\ell-4}{1-z^2} + \frac{\ell(\ell-1)}{2}\right)P^{(2)}_{\ell}(z)}{(\ell-1)\ell(\ell+1)(\ell+2)}, \\
F^{22}(z) &= 4\frac{(\ell+2)P^{(2)}_{\ell-1}(z) - (\ell-1)zP^{(2)}_{\ell}(z)}{(\ell-1)\ell(\ell+1)(\ell+2)\,(1-z^2)}.
\end{aligned}
\tag{2.46}
$$

We discuss how Eq. 2.34 is modified in the presence of pixelisation in Sec. 2.3.3 and in the presence of an instrumental noise and beam in Sec. 4.2.

2.3.2. Power spectrum estimate of a single sky

With a single and pixelised realisation of the sky, the relations presented in Sec. 2.2.1 slightly change. Since we only observe a *unique Universe*, we only have access to a *single realisation* of the field on the sphere, hence the ensemble average of Eq. 2.16 is not directly accessible. One can only estimate the power spectrum from the $2\ell+1$ available coefficients per multipole, thanks to the ergodicity hypothesis,

$$
\hat{C}^{XY}_{\ell} = \frac{1}{\nu_\ell}\sum_m a^X_{\ell m} a^{Y*}_{\ell m},
\tag{2.47}
$$

where we define the number of degrees of freedom per ℓ as

$$
\nu_\ell = 2\ell + 1.
\tag{2.48}
$$

This limited information contained in the sky constrains the possibility of estimating the power spectrum. Any estimate will come with a minimal probability density function, whose compactness around the true value is limited by the finite number of modes per estimated scale. In particular, Eq. 2.47 is optimal when considering an entire sky, in which case its covariance matrix can be computed analytically, see [33]. In particular, the diagonal terms can be derived in a straightforward way [34], developping the expectation value of the product of 4 random variables with the Wick's theorem and using the parity relation $a^*_{\ell m} = a_{\ell -m}$:

$$
\begin{aligned}
\sigma^2(C_\ell) = \langle(\hat{C}_\ell - C_\ell)^2\rangle &= \left\langle\left(\frac{1}{\nu_\ell}\sum_m a_{\ell m}a^*_{\ell m}\right)^2\right\rangle - 2C_\ell\left\langle\frac{1}{\nu_\ell}\sum_m a_{\ell m}a^*_{\ell m}\right\rangle + C_\ell^2 \\
&= \frac{1}{{\nu_\ell}^2}\sum_{mm'}\langle a_{\ell m}a^*_{\ell m}a_{\ell m'}a^*_{\ell m'}\rangle - 2C_\ell^2 + C_\ell^2 \\
&= \frac{1}{{\nu_\ell}^2}\left[\sum_{mm'}\langle a_{\ell m}a^*_{\ell m}\rangle\langle a_{\ell m'}a^*_{\ell m'}\rangle + \langle a_{\ell m}a_{\ell m'}\rangle\langle a^*_{\ell m}a^*_{\ell m'}\rangle + \langle a_{\ell m}a^*_{\ell m'}\rangle\langle a_{\ell m'}a^*_{\ell m}\rangle\right] - C_\ell^2 \\
&= \frac{1}{{\nu_\ell}^2}\left[\sum_{mm'} C_\ell^2 + C_\ell^2\delta_{m-m'} + C_\ell^2\delta_{mm'}\right] - C_\ell^2 \\
&= \frac{1}{{\nu_\ell}^2}\left[{\nu_\ell}^2 C_\ell^2 + 2\nu_\ell C_\ell^2\right] - C_\ell^2 \\
&= \frac{2}{\nu_\ell}C_\ell^2.
\end{aligned}
\tag{2.49}
$$

More generally, if one defines $\tilde{\mathbf{C}}_\ell$ as

$$\tilde{\mathbf{C}}_\ell \equiv \begin{pmatrix} C_\ell^{TT} \\ C_\ell^{TE} \\ C_\ell^{EE} \\ C_\ell^{BB} \end{pmatrix}, \tag{2.50}$$

one can show from 2.33 that in full-sky, it has a covariance matrix Ξ that reads:

$$\mathbf{\Xi}_\ell \equiv \mathrm{Cov}\left(\tilde{\mathbf{C}}_\ell\right) = \frac{2}{2\ell+1} \begin{pmatrix} \left(C_\ell^{TT}\right)^2 & C_\ell^{TT} C_\ell^{TE} & \left(C_\ell^{TE}\right)^2 & 0 \\ C_\ell^{TT} C_\ell^{TE} & \left[C_\ell^{TT} C_\ell^{EE} + \left(C_\ell^{TE}\right)^2\right]/2 & C_\ell^{TE} C_\ell^{EE} & 0 \\ \left(C_\ell^{TE}\right)^2 & C_\ell^{TE} C_\ell^{EE} & \left(C_\ell^{EE}\right)^2 & 0 \\ 0 & 0 & 0 & \left(C_\ell^{BB}\right)^2 \end{pmatrix}. \tag{2.51}$$

This can also be shown thanks to Wick's theorem. Instead, we will prove this result in Sec. 9 while discussing the full distribution of the estimator.

We will introduce how to include the effect of a partial sky in Sec. 4.1.2 and the effect of instrumental beam and noise in Sec. 4.2. For now, we investigate how the C_ℓ-estimate Eq. 2.47 translates on a pixelised sky.

2.3.3. Pixelisation of the sky

A common choice for pixelisation is called the *Hierarchical EqualArea and iso-Latitude Pixelisation* (*HEALPix*) (see [35, 36]). The lowest resolution corresponds to the case in which 4 pixels are around the equator, the north and south poles, making 12 pixels in total. A higher resolution is achieved from a given resolution by subdividing each pixel into four new pixels. Hence, the number of pixels in a map at a given resolution is $12N_{\rm side}^2$ where $N_{\rm side}$, the number of divisions along the side of a base-resolution pixel, is a power of 2.

On top of that, we include the pixelisation of the sky and Eq. 2.16 becomes:

$$\hat{a}_{\ell m} = \frac{4\pi}{N_{\rm pix}} \sum_{p=0}^{N_{\rm pix}-1} Y^*_{\ell m}(p)\, f(p)\,. \tag{2.52}$$

For a pixelised signal, $f(p)$ is the average within each pixel p of the underlying signal:

$$f(p) = \frac{N_{\rm pix}}{4\pi} \int_{\hat{n}\in p} \mathrm{d}\hat{n} f(\hat{n}), \tag{2.53}$$

This averaging induces a pixel effect in the estimators with respect to the true quantities that we introduce here. Eq. 2.12 becomes

$$f(p) = \sum_{\ell=0}^{\ell_{\rm max}} \sum_{m=-\ell}^{\ell} a_{\ell m} W_{\ell m}^{\rm pix}(p) Y_{\ell m}(p), \tag{2.54}$$

where the summation is stopped at $\ell_{\rm max}$, and where $W_{\ell m}^{\rm pix}(p)$ is the spherical harmonic transform of the pixel p. Instead of using this function of ℓ, m, and p, we usually introduce the effective m-averaged pixel window function, $W_\ell^{\rm pix}$, its average over the pixels m's for each angular scale.

$$W_\ell^{\rm pix} \equiv \frac{1}{N_{\rm pix}} \frac{4\pi}{2\ell+1} \sum_{p=0}^{N_{\rm pix}-1} \sum_{m=-\ell}^{\ell} \left| W_{\ell m}^{\rm pix}(p) \right|^2. \tag{2.55}$$

Under the assumption that the exact structure of the pixel can be ignored and that the pixels are not too different (so averaging the window function over the pixel makes sense), this pixel window function acts like an additional beam in the map that is due to the pixelisation. One can correct for this effect in the estimator Eq. 2.47 as one corrects the experiment beam window functions (see Sec. 8.1.2):

$$\hat{C}_\ell^{\rm unpix} = \frac{1}{W_\ell^{\rm pix\ 2}} \hat{C}_\ell^{\rm pix}. \tag{2.56}$$

In the case of the analysis of a full sky map with $N_{\rm side}$ resolution, Eq. 2.47, combined to Eq. 2.52 corrected for the pixel effect, provides an unbiased and optimal estimator of the power spectrum. We further discuss the limitations of this estimator in Sec. 8.

In Python, the module `Healpy` [37] implements spherical transformations with the Healpix pixelisation, for temperature and polarisation signals. We will use as a shortcut the appellation `anafast`, which is the `Healpix` name for Eqs. 2.52 and 2.47.

2.4. Current observational status

In this section, we overview the current CMB observational status by introducing the past experiments, and the current constraints on the so-called Λ-*CDM model* and on the initial conditions.

2.4.1. First measurement, COBE, Archeops and WMAP

The discovery of the CMB dates back to 1964 when Arno Penzias and Robert Wilson detected a persistent microwave signal, whose physical nature was realised soon after. Following this discovery, ground-based and balloon experiments have been characterising the CMB more precisely. However, these early efforts were limited by atmospheric interference and instrument sensitivity.

In 1989, the COBE (Cosmic Background Explorer) satellite was launched by NASA and provided the first detailed measurements of the CMB's spectrum and anisotropies. Onboard, FIRAS (Far Infrared Absolute Spectrophotometer) provided the most precise CMB black body spectrum measurement [2], illustrated in Fig. 2.4. In this figure, the error bar size is multiplied

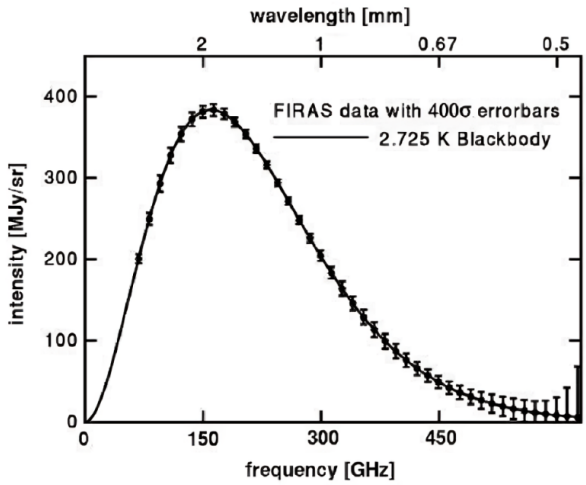

Figure 2.4: CMB spectrum as measured by COBE-FIRAS. The errorbars are multiplied by 400, and the Planck law best fit is represented in plain line. From [2].

by 400 in order to make them visible. The plain line corresponds to the best fit of the spectrum by the Planck law, confirming that the CMB follows a nearly perfect black-body spectrum with a temperature of approximately 2.725 K. This measurement is still today the best spectroscopic measurement of the CMB. The first detections of CMB anisotropies were made by the COBE DMR (Differential Microwave Radiometer) instrument, which confirmed the existence of temperature fluctuations in the CMB at the level of $\langle \delta T/T \rangle \simeq 10^{-5}$.

In 2002, the Archeops balloon [38]detected for the first time the first acoustic peak, confirming the Universe's flatness and allowing for an estimate of the cosmological constant (when combined with other measurements). The Wilkinson Microwave Anisotropy Probe (WMAP) was launched in 2001 by NASA and delivered in 2003 the most precise measurement of the first acoustic peak and cosmological parameters to date [39, 40].

2.4.2. Impact of the Planck mission on the Λ-CDM model

The Planck satellite, launched in 2009 by ESA, mapped the entire sky with unprecedented sensitivity and resolution. It provided the most detailed measurements of the CMB, allowing exquisite measurements of temperature anisotropies and of E-modes. In Fig. 2.5, we illustrate the final CMB maps produced by Planck [41]. The left panel represents the temperature map. The line-of-sights that are colder than the average temperature are represented in blue (they correspond to the dense regions at the recombination; see term **(A)** of Eq. 2.10 for an explanation), while the red spots illustrate the warm line-of-sights. In the right panel, a bar represents the polarisation vector field at each point of the map, which is superimposed with the smoothed temperature map. One verifies that, as mentioned earlier, the extrema in temperature corresponds to unpolarised regions. Moreover, E-mode polarisation patterns (that trace the velocity field at the recombination) can be seen around the temperature/polarisation extrema (see [42] in which this is even better illustrated through the stacking of hot spots in E and T maps).

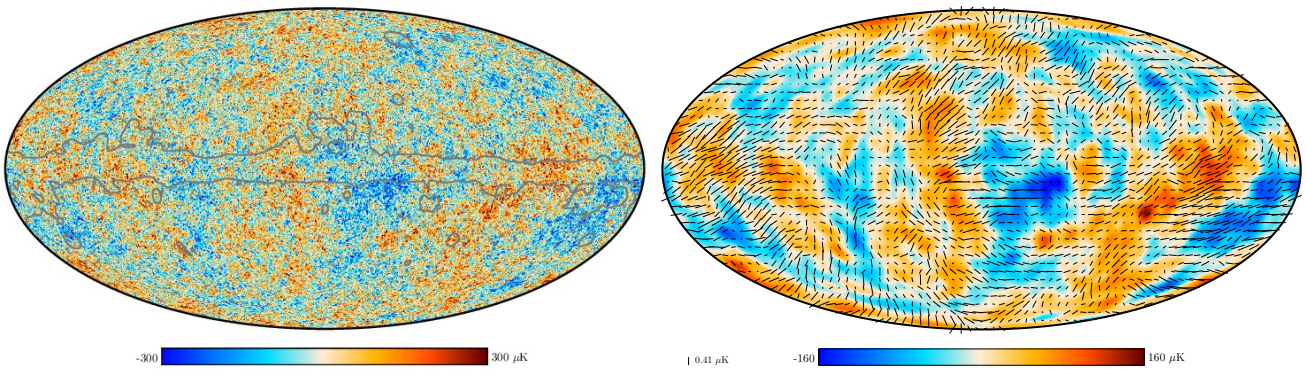

Figure 2.5: CMB sky measured by *Planck*, from [41]. Left: temperature map. Right: polarisation field as rods of varying length superimposed on the temperature map. The masked Galactic plane is delineated by a grey line in the temperature map, region that has been inpainted for visual purposes. The units are thermodynamic CMB temperature (with respect to the average temperature), which we formally introduce in Sec. 4.1.2.

The dotted blue points in Fig. 2.6 represent the Planck power spectra D_ℓ^{TT}, D_ℓ^{EE} and D_ℓ^{BB} on

the upper panel, D_ℓ^{TE} on the middle panel and $D_\ell^{\phi\phi}$ on the lower panel, which we will comment in Sec. 4.1.1. Still to date, the Planck data are the most constraining CMB data for Λ-CDM model and most of its extensions. To mention only one of the features whose detection have been confirmed by WMAP (cyan data points) and Planck, we can mention the detection of the negative correlation between temperature and polarisation between $50 < \ell < 150$, in the TE panel. As emphasised earlier, this is a characteristic signature of superhorizon fluctuations at the recombination.

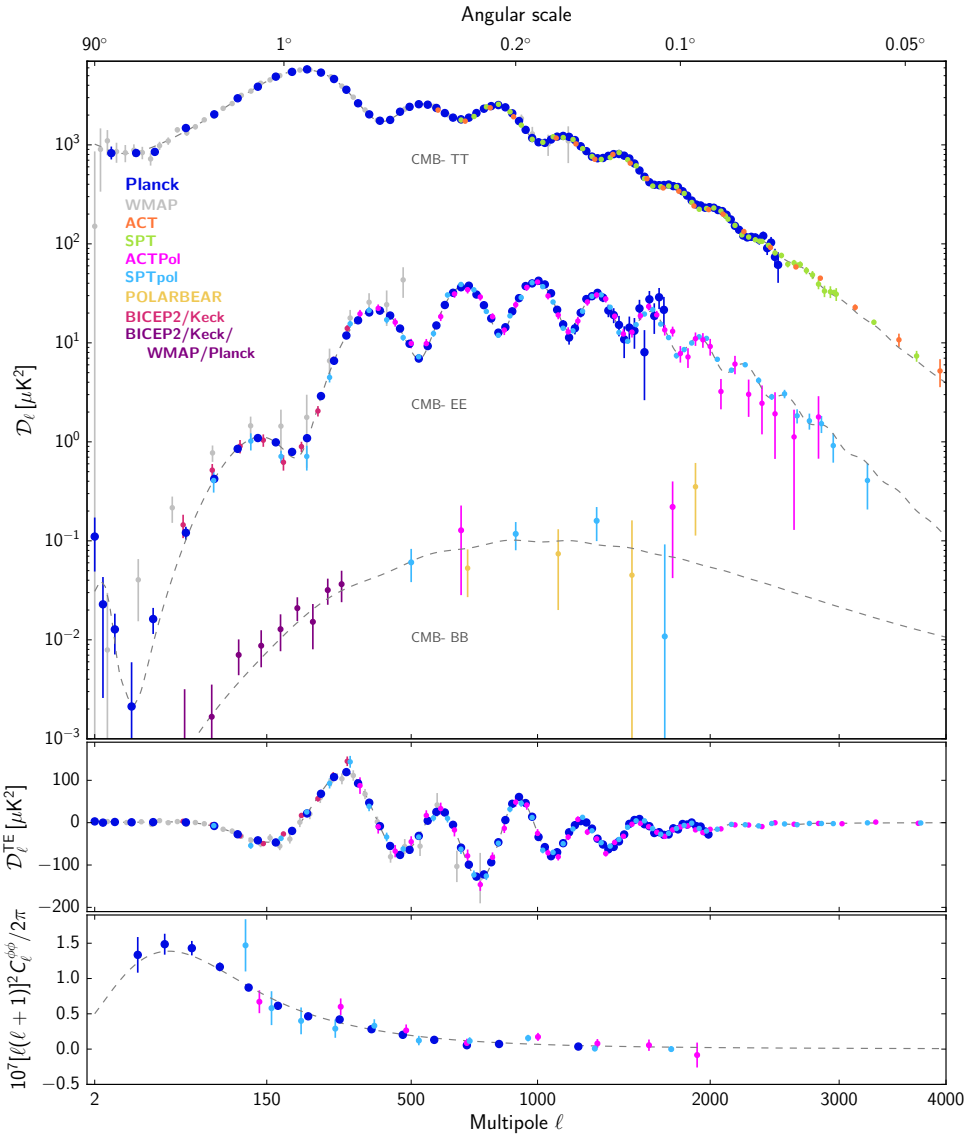

Figure 2.6: From Planck Collaboration [41]. Compilation of CMB angular power spectrum measurements as of 2018. TOP: D_ℓ^{TT}, D_ℓ^{EE}, D_ℓ^{BB}, MIDDLE: D_ℓ^{TE}, BOTTOM: $D_\ell^{\phi\phi}$. Different colours correspond to different experiments, each retaining its original binning. The dashed line shows the best-fit Λ-CDM model to the Planck temperature, polarisation, and lensing data. See [41] for details and references.

The standard cosmological Λ-CDM model:

We summarise here the assumptions on which we relied in the two first chapters, which are the assumptions upon which is built the Λ-CDM model. Physics is the same in the whole observable Universe. The latter is well described by General Relativity, is homogeneous at the largest scales, expanding since early times, and is flat ($\Omega_K = 0$ assumed). It is filled by the fluids we have presented in Sec. 1.1, dark energy, dark matter, baryons, photons and non-interacting neutrinos. The latter are included in the particle-physics Standard Model. The scalar perturbations are Gaussian and adiabatic and have a power-law primordial power spectrum ($\alpha_\mathrm{s} = \beta_\mathrm{s} = 0$ assumed). There are no tensor perturbations ($r = 0$ assumed).

Only six parameters are required to describe such a Universe. Usually, one uses as parameters two density parameters, $\Omega_\mathrm{b}h^2$ and $\Omega_\mathrm{c}h^2$ (see Eqs. 1.21 and 1.9), the angular scale of the acoustic oscillations θ_*, τ, the optical depth on the light of sight due to the reionisation, and A_s and n_s the scalar amplitude and tilt (see Secs. 1.2.2.3 and 1.2.4.3).

In Fig. 2.6, the dashed line represents the best-fit Λ-CDM model to the data, showing an accurate agreement between data and theory. Additionally, the data leads to constraints on the cosmological parameters within Λ-CDM that are given by the first block of Tab. 2.1 [41]. The first column contains the results with Planck data alone, while the second one includes Baryon Acoustic Oscillations (BAO) [43, 44] measurements as seen in the distribution of galaxies, in this case measured by SDSS-III and BOSS Collaboration [45]. BAO data set is often combined with CMB because it allows the break of degeneracies between parameters, especially when opening extensions. One sees that apart from τ, which is known at 10% level, all main parameters are known with better precision than the per cent. The second block presents the resulting

Parameter	Planck alone	Planck + BAO
$\Omega_\mathrm{b}h^2$	0.02237 ± 0.00015	0.02242 ± 0.00014
$\Omega_\mathrm{c}h^2$	0.1200 ± 0.0012	0.11933 ± 0.00091
$100\theta_*$	1.04110 ± 0.00031	1.04119 ± 0.00029
τ	0.0544 ± 0.0073	0.0561 ± 0.0071
$\ln(10^{10}A_\mathrm{s})$	3.044 ± 0.014	3.047 ± 0.014
n_s	0.9649 ± 0.0042	0.9665 ± 0.0038
H_0	67.36 ± 0.54	67.66 ± 0.42
Ω_Λ	0.6847 ± 0.0073	0.6889 ± 0.0056
$\Omega_\mathrm{m}h^2$	0.1430 ± 0.0011	0.14240 ± 0.00087
z_re	7.67 ± 0.73	7.82 ± 0.71
Age[Gyr]	13.797 ± 0.023	13.787 ± 0.020
z_eq	3402 ± 26	3387 ± 21

Table 2.1: From [41]. Parameter 68% CL from Planck CMB temperature, polarisation, and lensing power spectra, and with the inclusion of BAO data. The first block contains the six Λ-CDM parameters, while the second one contains the derived quantities. More details can be found in the core of the text and in [41].

constraints on derived parameters: the Hubble parameter, the fraction of dark energy, the total fraction of matter, the reionisation redshift, the redshift at equality matter-radiation, and the

age of the Universe: 13.797 Giga-years. Within Λ-CDM, the latter is known with a precision below 0.2%.

Planck's results are internally consistent, with cross-checks between different fractions of data sets confirming the robustness of its findings [41]. Beyond Λ-CDM, Planck constrains extensions to the Λ-CDM model and demonstrates that opening most of those does not significantly shift the baseline parameter values [41], underscoring the robustness of Λ-CDM.

It also allows constraints to be placed on these extensions: the sum of the neutrino masses has been constrained to be less than 0.12 eV at 95% confidence limit [41], while the inferred number of light relics is 3.01 ± 0.35 at 95% confidence limit (the standard model prediction is 3.044 [46, 47]). We let the discussion of constraints on initial condition and on flatness to Sec. 2.4.4.

Let us mention that beyond its internal consistency, there exists tensions and anomalies with external data sets (see [48] for a review). The most important one is a $\sim 4-6\sigma$ tension between the Hubble parameter directly measured with late-time probes, typically supernovae of type IA calibrated with Cepheid [49], and the value given by the CMB within Λ-CDM. This tension could be explained by the presence of systematics in one or the other measurement, or could hint towards physics beyond Λ-CDM [50]. Another anomaly exists between the amplitude of perturbations inferred from CMB or from late-time measurements. We will not comment further on these hot topics, which are not the scope of our work.

New Planck data releases and reanalysis have been made since 2018, including the *Beyond-Planck* global Bayesian analysis of the Planck Low Frequency Instrument data [51], and the delivery of *Planck Data Release 4* with the *NPIPE* analysis [52] and associated constraints [53].

2.4.3. SPT, ACT and BICEP

The South Pole Telescope [54] (SPT, deployed in 2007) and the Atacama Cosmology Telescope [55] (ACT, deployed in 2007) have focused on measuring the CMB at small angular scales. SPTpol [56] (deployed in 2012) and ACTpol [57] (deployed in 2013) are the upgraded experiments that include sensitivity to polarisation. These observations complement those from Planck by providing high-resolution data that helps to refine the understanding of small-scale fluctuations. In Fig. 2.6, their data points from 2018 are represented in yellow and light blue for SPT and SPTpol and in orange and purple for ACT and ACTpol. Notice that these two experiments have been upgraded into SPT-3G [58] (deployment 2017) and Advanced ACTpol [59] (deployment 2016).

The BICEP (Background Imaging of Cosmic Extragalactic Polarization) experiment [60] is another ground-based telescope situated at the South Pole. Its observation strategy is to do a compact scan of a small patch of the sky in order to integrate over time and reach a high polarisation measurement sensitivity.

We will briefly mention again the future of these experiments in Sec. 4.5.

2.4.4. Current constraints on initial conditions

BICEP provides the best measurements of the non-primordial B-mode signal that is due to lensing (see Sec. 4.1.1). It places stringent constraints on the tensor-to-scalar ratio r, particularly when combined with Planck data, with a current 95% CL limit on r of $r < 0.032$ at $k_* = 0.05$ Mpc^{-1} [61]. The right panel of Fig. 2.7 illustrates this constraint in the r-n_s plane [61]. Moreover, it allows us to put constraints on the tensor tilt n_t [62, 63].

The left panel of Fig. 2.7illustrates the free-form reconstruction of the primordial scalar

power spectrum using the Planck and the BICEP/Keck-15 data. One observes the remarkable agreement of the scalar power spectrum with a simple power law. The pivot scale value k_* at which the primordial parameters are defined, see Sec. 1.2.2.3, is most often 0.05 Mpc^{-1}, which is in roughly in the middle of the CMB observable range (in order for the Taylor expansion to be as constrained as possible[2]).

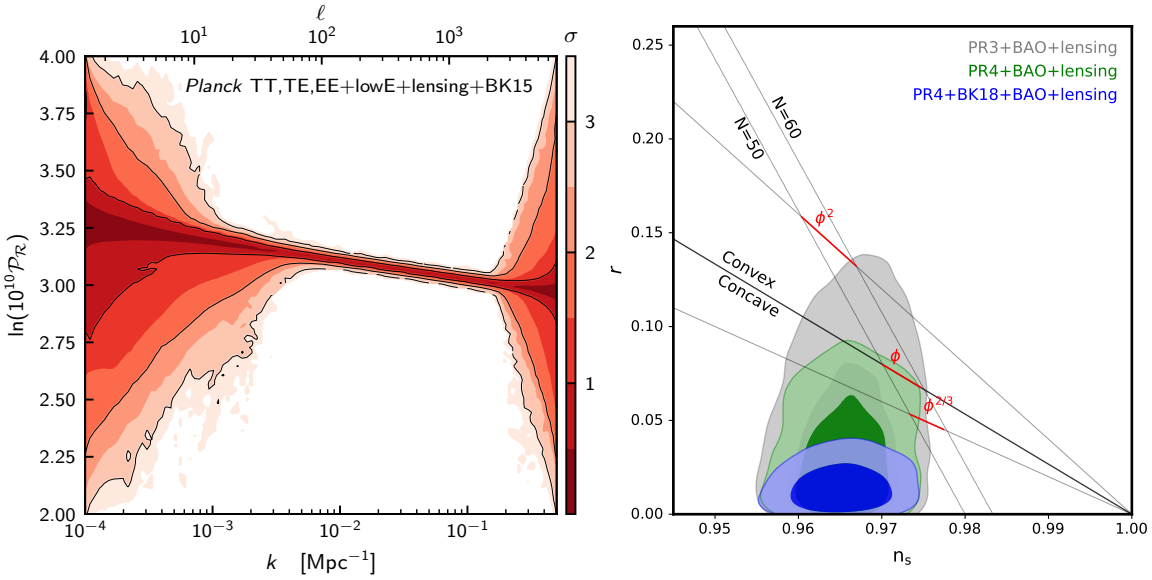

Figure 2.7: LEFT: From [64]. Free-form Bayesian reconstruction of the primordial scalar power spectrum using Planck TT,TE,EE+lowE+lensing+BK15. The reconstruction was made by fitting the spectrum values at certain wavelength knots, the number of which is marginalised. See [64] for details on the analysis. RIGHT: From [61]. The blue region represents the current best constraint on the n_s, r plane, combining BICEP/Keck 2018 [65] and Planck PR4 data TT,TE,EE+lowlEB+lensing+BAO [66]. We let for later on the explanation of the other feature on this figure.

Hence, the Planck data performed an unprecedented measurement of the scalar amplitude ($\ln(10^{10}A_s) = 3.047 \pm 0.014$) and tilt ($n_S = 0.967 \pm 0.004$), ruling out the Harrison-Zel'dovich flat spectrum by almost 9σ [64]. It also severely constrained the scalar running ($\alpha_S = -0.0042 \pm 0.0067$). Additionally, it observationally supported the assumptions on the perturbations we have worked with in these two first chapters: as far as constrained by Planck, the Universe is flat ($\Omega_K = 0.0007 \pm 0.0019$) and the perturbations are Gaussian and adiabatic ($f_{\rm NL} = -0.9 \pm 5.1$ and $\alpha_{-1} = 0.00013 \pm 0.00037$, which are two observables respectively quantifying amount of non-Gaussianities and amount of (non-adiabatic) isocurvature modes).

Conclusion

Building upon Chapter. 1, we have overview in Sec. 2.1 the various sources of anisotropies in the CMB: scalar Sachs-Wolfe, Doppler and integrated Sachs-Wolfe effects, linear polarisation

[2]For the same motivation, the tensor parameters are sometimes defined at a lower pivot scale of $k_* = 0.002$ Mpc^{-1}. It won't be the default choice in this work; we will instead stick with $k_* = 0.05$ Mpc^{-1}.

from the recombination and the reionisation scattering, as well as tensor-originated anisotropies. Then, we have defined in Sec. 2.2 the power spectra that are required to describe the anisotropy angular statistics, and summarised what standard cosmology predicts for these objects. In Sec. 2.3, we have presented the basic tools to analyse CMB pixelised maps. Finally, in Sec. 2.4, we have provided an overview of the past CMB experiments and the current status for Λ-CDM, a set of six parameters that describes very well the available data sets.

Despite this success, some observed peculiarities remain conceptually unexplained. We have seen that super-horizon modes have to exist, as hinted by the Sachs-Wolfe effect and confirmed by the detection of negative TE at large scales. With our description so far, the presence of such large-scale perturbation cannot be explained by causal effects. Moreover, the geometry of the Universe has been measured to be very close to flat. Finally, we have seen that Planck observes specific properties for the initial conditions, which would require a theoretical explanation. Answering these issues will be the focus of our next Chapter, through the introduction of *cosmological inflation*, a speculative high-energy process prior to the radiation era whose simplest realisations explain all our observations.

CHAPTER 3

Inflation as the cherry on the top

In this chapter, we first discuss the three main conceptual issues of Λ-CDM that led cosmologists to introduce an accelerated expansion era prior to the radiation era, *inflation*. By definition, accelerated expansion means that $\ddot{a} > 0$, which implies, according to the Raychaudhuri equation Eq. 1.23, that

$$w = P/\rho < -1/3. \tag{3.1}$$

We have seen that such an equation of state is already required for describing the late-time dynamic of the Universe. We will introduce the simplest mechanism that can explain inflation, *slow-roll*, which leads to an equation of state $w = -1$. Such a mechanism is also able to naturally predict initial conditions as inferred by modern cosmological observations, which is today the main motivation for the paradigm.

Contents

3.1 A solution to conceptual issues of the cosmological model . . 50
3.1.1 The Horizon Problem or how to correlate super-horizon scales . 50
3.1.2 The Flatness Problem 53
3.1.3 The Monopole problem 54
3.2 Single-field slow-roll inflation 54
3.2.1 Scalar-field inflation 55
3.2.2 Slow-roll approximation 55
3.3 Slow-roll inflationary perturbations 57
3.3.1 Slow-roll primordial power spectra 57
3.3.2 Success of slow-roll inflation 59
3.3.3 Motivation for measuring r 59
3.3.4 Constraints on the slow-roll parameters 60

3.1. A solution to conceptual issues of the cosmological model

We have introduced the standard model of cosmology, which explains the observations from various cosmological experiments, particularly those conducted by the Planck satellite. However, even before these observations, some conceptual issues emerged. To address these concerns, the concept of an accelerated expansion of the early Universe prior to the radiation era was introduced around 1980 in [67–69]. In this section, we go through the conceptual issues in turn and explain how inflation solves them.

3.1.1. The Horizon Problem or how to correlate super-horizon scales

The *horizon problem* was noticed as early as 1956 [70, 71].

3.1.1.1 The Horizon Problem

In a Universe described by an FLRW metric, the photon line element is $\mathrm{d}s^2 = 0$. Integrating the FLRW metric, we get a proper distance between the emission and detection of a photon:

$$\int_0^t \frac{\mathrm{d}t}{a(t)} = \int_0^r \frac{\mathrm{d}r}{\sqrt{1-kr^2}}. \tag{3.2}$$

We define *the particle horizon*, which has a length dimension, as

$$d_{\mathrm{hor}}(t) \equiv a(t) \int_0^t \frac{\mathrm{d}t'}{a(t')}. \tag{3.3}$$

Because the speed of light is finite, this distance is the maximum one a particle emitted at a time t in the direction of the observer can travel. During a given era, if $a(t) \propto t^p$,

$$d_{\mathrm{hor}}(t) = \frac{a(t)}{1-p}\left[t^{1-p} - t_i^{1-p}\right]. \tag{3.4}$$

We also introduce the *comoving Hubble radius* defined as $\frac{1}{aH}$, the typical size of the Universe. Then, the *comoving particle horizon*, defined as $d_{\mathrm{hor}}(t)/a(t)$ is proportional to the comoving Hubble radius:

$$\frac{1}{1-p}\left[t^{1-p} - t_i^{1-p}\right] \propto \frac{1}{aH}. \tag{3.5}$$

If two regions are separated by a comoving distance larger than the comoving particle horizon, it means that they are causally disconnected, *i.e.* they have never been in contact in the past [72, 73].

Let's substitute the time variable with the scale factor. Recalling that for a fluid parameter $w \neq -1$, $p = \frac{2}{3(1+w)}$, the comoving Hubble horizon of a Universe dominated by a fluid with an equation of state w between t_1 and t_2 can be written as

$$\ln \frac{1}{aH}\bigg|_{t_2} = \frac{1+3w}{2} \ln\left[\frac{a(t_2)}{a(t_1)}\right] + \ln \frac{1}{aH}\bigg|_{t_1}, \tag{3.6}$$

the formula that turns out to remain true in the special case $w = -1$. From Eq. 3.6, one sees that the comoving Hubble radius increases with the scale factor as long as $w > -\frac{1}{3}$. In this case, perturbation modes keep entering the Universe, establishing causal contact for distant regions

for the first time. This is what happens during the radiation era, during which the equation of state is $\frac{1}{3}$.

Today, we observe today CMB photons with similar temperatures from all directions of the sky ($\langle \delta T/T \rangle \simeq 10^{-5}$). Apart from accepting a severe fine-tuning of the initial conditions, a causal contact of these various patches before CMB emission is required. More quantitatively, the horizon at the recombination corresponds to the scale at which the cosine term in the acoustic oscillations, see Eq. 1.122, begins to contribute to the anisotropies. In ℓ-space, this roughly corresponds to the end of the Sachs-Wolfe plateau, where the causal acoustic-oscillation physics appears. Hence, at angular scales lower than 1°, no causal mechanism after the entry of the mode in the Universe can explain the homogeneity of the sky, nor does it explain the features in the spectra (especially concerning the TE spectra [30, 31]). How to reconcile these two pictures?

3.1.1.2 Inflation as a solution

According to Eq. 3.6, a simple way to explain this causal contact is to invoke an era dominated by a fluid with $w < -\frac{1}{3}$, during which the comoving Hubble horizon decreases, prior to radiation era. In this case, modes in causal contact exit the Hubble radius at an early time and freeze at super-horizon scales before reentering the horizon during the radiation-dominated era. This necessary condition on the equation of state coincides with the inflationary necessary condition Eq. 3.1. Consequently, a sufficiently long period of inflation is enough to solve the horizon problem.

To quantify the duration of inflation, one defines the number of elapsed e-folds between t and t_0 as:

$$N - N_0 \equiv \ln(a/a_0) = \int_{t_0}^{t} H \mathrm{d}t. \tag{3.7}$$

For the exercise, we place ourselves in a particular setup for the next discussion: during a *de-Sitter ($w = -1$) inflation*, the Hubble radius's logarithm scales as the scale factor's logarithm. If we only assume a period of inflation during which $w = -1$ (slope -1) and a radiation era $w = 1/3$ (slope $+1$), the horizon problem is solved if the largest comoving scale we observe today is smaller than the size of the comoving Hubble radius at the beginning of the inflation. Hence, one can derive the minimum number of e-folds elapsed during the beginning and the end of inflation $(N_{\rm end} - N_{\rm in})_{\rm min}$ that solves the Horizon problem:

$$(N_{\rm end} - N_{\rm in})_{\rm min} \equiv \ln \frac{a_{\rm end}}{a_{\rm in}}\bigg|_{\rm min} = \ln \frac{a_0}{a_{\rm end}}, \tag{3.8}$$

since, with our assumption, the radiation era and the inflation era are associated with an opposite evolution of the comoving Hubble radius. Thanks to Eq. 1.2, the minimum e-fold number can be rewritten as:

$$(N_{\rm end} - N_{\rm in})_{\rm min} = \ln \frac{T_{\rm end}}{T_0} \simeq 66, \tag{3.9}$$

where $T_0 \simeq h\nu_{\rm peak} = 0.6 \times 10^{-3}$ eV and we assumed that $T_{\rm end} = 3 \times 10^{16}$ GeV. The (ad-hoc) choice for the temperature at the end of inflation is motivated by the Grand Unified Theory (GUT) [74] breaking scale, above which infrared forces and their couplings are unified. We rediscuss this particular numerical value in Sec. 10.2.3 that focuses on the Minimal Supersymmetric Standard

Model specific case. This derivation is very simple for the illustration but can be done in general with fewer assumptions see [72, 75].

The first step to refine the picture would be to account for the late-time matter-dominated and dark-energy era. Moreover, one should include the transition era between inflation and radiation epoch, called *reheating*, which plays the role of converting the inflationary energy into light particles.

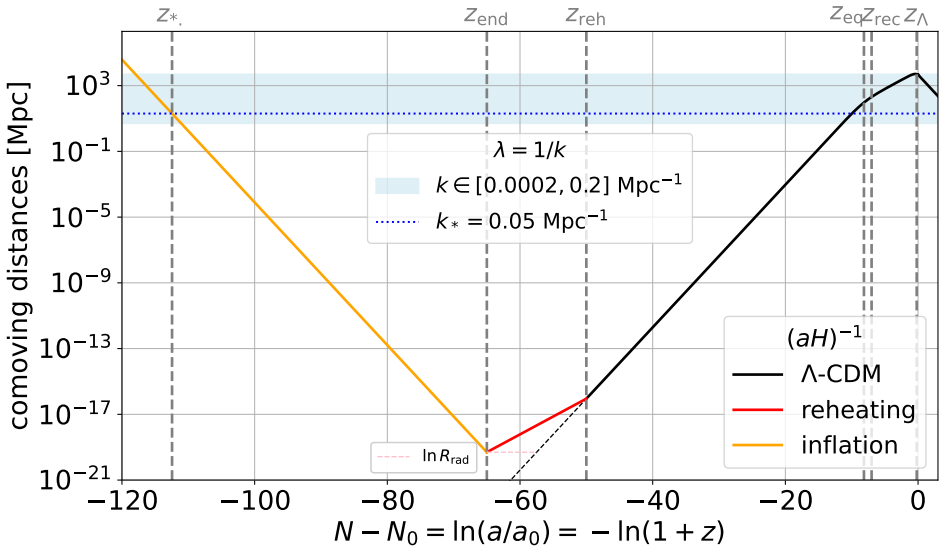

Figure 3.1: Evolution of the *comoving Hubble radius* including an inflation period, as compared to the comoving length of the characteristic wavelength of perturbations visible in the CMB. The chosen numerical Λ-CDM values are compatible with the current observational constraints and correspond to the fiducial we introduce later on in Tab. 4.1. To fix the ideas on the order of magnitude for the cosmic time for the numerical values of this particular figure, the time at the end of inflation is of order 10^{-37} s, while the one at reheating is of order 10^{-24} s. $\ln R_{\rm rad}$, the number of e-folds "saved" by the non-instantaneous reheating, will only be used in Part IV and introduced in Eq. 10.10.

We present a summary of the problem and its resolution, including these additional eras, in Fig. 3.1. In this figure, we show the comoving Hubble radius $1/(aH)$ compared to the typical comoving wavelengths of CMB perturbations. The wavenumber $k_* = 0.05$ Mpc^{-1} is represented by a blue dotted line, and the typical observable range of the CMB is shown in a light-blue contour. The slopes of $1/(aH)$ versus a during different eras can be derived from Eq. 3.6. For reference, we included a de Sitter inflation (slope = -1, in orange) and a matter-like reheating ($w = 0$ and slope = $1/2$, in red). In this case, the k_* mode exits the horizon at z_* around $N_* = 50$ e-folds before the end of inflation $z_{\rm end}$. In other words, it exceeds the horizon size, and the corresponding scales are no longer causally connected. This mode is *frozen*, it does not evolve during the superhorizon regime and it reenters the horizon slightly before the recombination ($z_{\rm rec}$).

We have shown that inflation provides a solution to the horizon problem, explaining the correlations between large-scale fluctuations that WMAP and Planck have precisely measured in temperature and polarisation. Alongside inflation, the existence of super-horizon fluctuations

and a solution to the horizon problem are also provided by the *pre-Big-Bang scenario* [76] and the *Ekpyrotic scenario* [77].

3.1.2. The Flatness Problem

Another important motivation for inflation is its role in solving the *flatness problem*, a role first proposed by Dicke and Peebles in 1986.

3.1.2.1 The problem

Let us start from equation 1.22. Dividing it by H^2, we get

$$\begin{aligned} 1 &= \Omega_{\rm tot} + \Omega_{\mathcal{K}} \quad \text{with } \Omega_{\mathcal{K}} \equiv \frac{\mathcal{K}}{a^2 H^2}, \\ &= \Omega_{\rm tot} + \frac{1}{1 + \frac{\Omega_{\rm tot}}{\Omega_{\mathcal{K}}}}. \end{aligned} \tag{3.10}$$

In the early Universe during the radiation-dominated era, $w = 1/3$. From Eq. 1.25 one gets $\Omega_{\rm tot} \simeq \Omega_{\rm rad} = \Omega_{\rm rad}^{(0)} \left(\frac{a_0}{a}\right)^4$, and following from its definition, $\Omega_{\mathcal{K}} = \Omega_{\mathcal{K}}^{(0)} \left(\frac{a_0}{a}\right)^2$. Using Eq. 3.10, the geometrical curvature density in the early Universe can be rewritten as

$$\begin{aligned} \Omega_{\mathcal{K}} &= \frac{1}{1 + \frac{\Omega_{\rm rad}^{(0)}}{\Omega_{\mathcal{K}}^{(0)}} \left(\frac{a_0}{a}\right)^2} \\ &\simeq \frac{\Omega_{\mathcal{K}}^{(0)}}{\Omega_{\rm rad}^{(0)}} \left(\frac{1}{1+z}\right)^2. \end{aligned} \tag{3.11}$$

Considering that the Planck constraints give $\Omega_{\rm rad}^{(0)} \simeq 10^{-4}$ and $\Omega_{\mathcal{K}}^{(0)} < 10^{-3}$ [41], one obtains the following value for the curvature at $T_{\rm GUT} = 3 \times 10^{16}$ GeV (*i.e.* $z_{\rm GUT} = 5 \times 10^{28} = \frac{3\times 10^{16}\ \text{GeV}}{0.6\times 10^{-3}\ \text{eV}}$):

$$\Omega_{\mathcal{K}} < 4 \times 10^{-57}. \tag{3.12}$$

Hence, to have such a flat Universe today, we need initial conditions that are extremely precise for the curvature of the Universe. This is a *fine-tuning* issue called the *flatness problem*.

3.1.2.2 Inflation as a solution

Once again, a period of *inflation* can solve it. Consider a fluid X with an equation of state $w_X < -1/3$, which dominates the Universe energy budget. Then, replacing $\Omega_{\rm rad}$ in Eq. 3.11 by Ω_X in 3.11, one can link the curvature at end of inflation $\Omega_{\mathcal{K}}^{\rm end}$ with the one at beginning of inflation $\Omega_{\mathcal{K}}^{\rm in}$:

$$\Omega_{\mathcal{K}}^{\rm end} =\simeq \frac{\Omega_{\mathcal{K}}^{\rm in}}{\Omega_X^{\rm in}} (1+z)^{-1-3w_X}. \tag{3.13}$$

To avoid fine-tuning for initial conditions, one can impose $\Omega_X^{\rm in} \sim \Omega_K^{\rm in}$. We can deduce from the last equation the minimum number of e-folds required to avoid fine-tuning. For instance, in the particular case in which $w_X = -1$, and where the end of inflation is at GUT,

$$(N_{\rm end} - N_{\rm in})_{\rm min} \equiv \ln \frac{a_{\rm end}}{a_{\rm in}} = -\frac{1}{2} \ln \Omega_{\mathcal{K}}^{\rm end} \simeq 65. \tag{3.14}$$

3.1.3. The Monopole problem

In most Grand Unified Theories, magnetic monopoles are expected to be abundantly produced [78–80] and should be dominating the current Universe energy budget [81, 82] while the magnetic monopoles today are typically constrained to be fewer than $\mathcal{O}(10^{-30})$ per nucleon.

3.1.3.1 The problem

Assuming conservatively that one Magnetic Monopole (MM) has been created in each Hubble volume at the GUT scale, one can compute the MM density today within the standard model of cosmology: there should be today one MM per volume $V_{\rm MM}^{(0)}$ given by

$$V_{\rm MM}^{(0)} = \left(\frac{a_0}{a_{\rm GUT}}\right)^3 \frac{4\pi}{3} d_{\rm hor}^3(t_{\rm GUT}), \tag{3.15}$$

where, $d_{\rm hor}(t_{\rm GUT})$ can be computed within a given cosmology from Eq. 3.2. It typically leads to $V_{\rm MM}^{(0)} = 1$ m^3 taking again $z_{\rm GUT} = 10^{28}$. Hence, if each monopole has for mass $M_{\rm GUT} = 3 \times 10^{16}$, $\Omega_{\rm MM}^{(0)}$ is expected to be of order

$$\Omega_{\rm MM}^{(0)} = \frac{1}{\rho_{\rm crit}^{(0)}} \frac{M_{\rm GUT}}{V_{\rm MM}^{(0)}} \simeq 10^{15}. \tag{3.16}$$

Hence, in the standard model, monopoles should dominate the energy content by 15 orders of magnitude today. One can also estimate the predicted number of monopole per nucleons, $\frac{\Omega_{\rm MM}^{(0)}}{\Omega_{\rm m}^{(0)}} \frac{m_{\rm nuc}}{M_{\rm GUT}} \simeq \mathcal{O}(1)$, to be compared to the experimental limit of 10^{-30} monopole per nucleon.

3.1.3.2 Inflation as a solution

A phase of inflation has the effect of diluting the monopole abundance. One can show [83] that in the case of an inflation phase of duration N e-folds characterised by an equation of state w, the typical volume in which we expect one monopole today is increased by

$$V_{\rm MM}^{(0)} \rightarrow V_{\rm MM}^{(0)} \times \exp\{(1-3w)N/2\}. \tag{3.17}$$

This formula and the experimental measurement of the matter fraction and constraint on the number of monopoles per nucleon can be used to put lower limits on the total number of e-folds. In the particular case $w = -1$, [83] obtains that $N > 61$ to satisfy these limits.

3.2. Single-field slow-roll inflation

We have seen that a phase of accelerated expansion in the very early Universe solves three conceptual issues of Λ-CDM. Today, the results from the Planck satellite reinforce this primordial *inflation* hypothesis [67–69, 84–88]. Measurements of the amplitude and spectral index of the primordial scalar power spectrum, and constraints on its running, on the amount of non-Gaussianities, on the amplitude of isocurvature modes and on the tensor-to-scalar ratio, favour single-field slow-roll inflationary models [64, 89]. Such phenomenological models are often inspired by high-energy constructions [90–92] in which one or more scalar fields acquire

a flat-enough potential. We present the single-field slow-roll inflationary framework in the two next sections, an introduction that is partly extracted from a work we have published [93], complemented by other references.

3.2.1. Scalar-field inflation

The simplest field compatible with space-time symmetries is a homogeneous scalar field ϕ, referred to as the *inflaton*. Let us first investigate what happens at background level, with a metric $g_{\mu\nu}$ given by Eq. 1.3. When minimally coupled to gravity, the scalar field action reads

$$S = -\int \mathrm{d}^4x\sqrt{-g}\left[\frac{1}{2}g^{\mu\nu}\partial_\mu\phi\partial_\nu\phi + V(\phi)\right], \tag{3.18}$$

where V is the potential on which the inflaton evolves. One can derive the stress-energy tensor associated with the field from Eq. 1.17:

$$T_{\mu\nu} = \partial_\mu\phi\partial_\nu\phi + g_{\mu\nu}\left[-\frac{1}{2}g^{\rho\sigma}\partial_\rho\phi\partial_\sigma\phi + V(\phi)\right], \tag{3.19}$$

which can be rewritten in the perfect-fluid generic form Eq. 1.17, allowing to identify the pressure and density associated with ϕ to:

$$\rho = \frac{\dot{\phi}^2}{2} + V \tag{3.20}$$

$$p = \frac{\dot{\phi}^2}{2} - V. \tag{3.21}$$

The so-called Klein-Gordon equation, which dictates the dynamic of the field in its potential, follows from injecting these density and pressure expressions into the continuity equation Eq. 1.24

$$\ddot{\phi} + 3H\dot{\phi} + V_\phi = 0, \tag{3.22}$$

where $V(\phi)$ is the potential energy stored in ϕ (and $V_\phi \equiv \mathrm{d}V/\mathrm{d}\phi$). Moreover, according to Friedmann's equation Eq. 1.19, the Hubble parameter is related to the energy density of the Universe,

$$3M_{\mathrm{Pl}}^2H^2 = V + \frac{\dot{\phi}^2}{2}, \tag{3.23}$$

where M_{Pl} is the reduced Planck mass. Inflation occurs when $\ddot{a} > 0$, which, from inserting the pressure and density expressions in Eq. 1.20, is equivalent to

$$V > \dot{\phi}^2. \tag{3.24}$$

This implies that the potential function $V(\phi)$ must be sufficiently flat for the field to roll it down sufficiently slowly, such that its kinetic energy does not exceed half its potential energy.

3.2.2. Slow-roll approximation

The Klein-Gordon equation (3.22) being second order, the dynamical phase-space $(\phi, \dot{\phi})$ has dimension two. When inflation takes place, there exists a dynamical attractor along which the acceleration term $\ddot{\phi}$ becomes negligible in Eq. 3.22, and the friction term $3H\dot{\phi}$ compensates the potential gradient V_ϕ. Such an attractor is called *slow-roll* [91] since for flat-enough potentials it

is also such that the condition of Eq. 3.24 is saturated, *i.e.* $\dot{\phi}^2 \ll V$. This attractive behaviour is an important conceptual argument for the naturalness of slow-roll. Additionally, [94] showed that the single-field setup act as an attractor for the inflationary trajectory, even in high-energy models with several degrees-of-freedom.

Hereafter, we will use the so-called *Hubble-flow parameters* to describe slow-roll inflation. The first of them is defined as $\varepsilon_0 \equiv H_{\rm in}/H$, which is constant in a de-Sitter Universe. The next-order parameters ε_n quantify the small deviations from de-Sitter and slow roll corresponds to the regime where they are small. Those are iteratively defined via

$$\varepsilon_{n+1} = \mathrm{d}\ln|\varepsilon_n|/\mathrm{d}N\,, \tag{3.25}$$

where N is the number of e-folds defined in Eq. 3.7. For instance, the first slow-roll parameter is given by

$$\varepsilon_1 = -\dot{H}/H^2 = 1 - a\ddot{a}/\dot{a}^2, \tag{3.26}$$

so inflation ($\ddot{a} > 0$) corresponds to $\varepsilon_1 < 1$. Using Eq. 3.22 and Eq. 3.23, one also has[1]

$$\varepsilon_1 = 3\dot{\phi}^2/(2V + \dot{\phi}^2) \tag{3.28}$$

and one recovers the condition of Eq. 3.24, with $\varepsilon_1 \ll 1$ corresponding to $\dot{\phi}^2 \ll V$. Injecting Eq. 3.26 into Eqs. 3.20 and 3.21, one obtains the equation of state during slow-roll

$$w \overset{\text{SRLO}}{\simeq} -1 + \frac{\varepsilon_1}{3}, \tag{3.29}$$

where $\overset{\text{SRLO}}{\simeq}$ indicates that we work at leading order in the slow-roll approximation. The second slow-roll parameter can be expressed as

$$\varepsilon_2 = 2\varepsilon_1 - 2V_\phi/(H\dot{\phi}) - 6 \tag{3.30}$$

so $|\varepsilon_2| \ll 1$ corresponds to $\dot{\phi} \simeq -V_\phi/(3H)$, *i.e.* when the acceleration term in the Klein-Gordon equation is subdominant. In this regime, $\dot{\phi}$ becomes a function of ϕ only,

$$\dot{\phi} \overset{\text{SRLO}}{\simeq} -M_{\rm Pl}\frac{V_\phi}{\sqrt{3V}}. \tag{3.31}$$

From the slow-roll Friedmann equation, $H^2 \overset{\text{SRLO}}{\simeq} V/(3M_{\rm Pl}^2)$, the number of e-folds elapsed between a given time (associated with a field value ϕ) and the end of inflation (associated with a field value $\phi_{\rm end}$) can be written as

$$\Delta N = \int_{\phi_{\rm end}}^{\phi} \frac{\mathrm{d}\phi}{\sqrt{2\varepsilon_1}} \overset{\text{SRLO}}{\simeq} \int_{\phi_{\rm end}}^{\phi} \frac{V(\phi)}{V_\phi(\phi)}\mathrm{d}\phi. \tag{3.32}$$

[1] combined with Eq. 3.23, this also gives an exact relation for the Hubble rate in slow-roll:

$$H^2 = \frac{V}{3M_{\rm Pl}^2}\left(1 - \frac{\varepsilon_1}{3}\right)^{-1}. \tag{3.27}$$

and approximate expressions for the slow-roll parameters only involving the potential function can be derived:

$$\varepsilon_1 \overset{\text{SRLO}}{\simeq} \frac{M_{\rm Pl}^2}{2}\left(\frac{V_\phi}{V}\right)^2, \tag{3.33}$$

$$\varepsilon_2 \overset{\text{SRLO}}{\simeq} 2M_{\rm Pl}^2\left[\left(\frac{V_\phi}{V}\right)^2 - \frac{V_{\phi\phi}}{V}\right], \tag{3.34}$$

$$\varepsilon_3 \overset{\text{SRLO}}{\simeq} \frac{2}{\varepsilon_2}M_{\rm Pl}^4\left[\frac{V_{\phi\phi\phi}V_\phi}{V^2} - 3\frac{V_{\phi\phi}}{V}\left(\frac{V_\phi}{V}\right)^2 + 2\left(\frac{V_\phi}{V}\right)^4\right], \tag{3.35}$$

where only the first three slow-roll parameters are given.

3.3. Slow-roll inflationary perturbations

Cosmological perturbations can be introduced on top of this homogeneous and isotropic expanding background. Only the scalar and tensor sectors propagate with equations of motion that are set by the background dynamics, which can be described by means of the slow-roll parameters. This is why, when the initial conditions for inflation are set in the quantum vacuum state, the Fourier modes of these perturbations only depend on the value of the slow-roll parameters at the time when they cross out the Hubble radius during inflation. At linear order in perturbation theory, their statistics is Gaussian, hence fully specified by the power spectra $\mathcal{P}_\mathcal{R}$ and $\mathcal{P}_h$ that we have introduced in Eq. 1.94 and Eq. 1.100.

3.3.1. Slow-roll primordial power spectra

For describing scalar perturbations, instead of doing the derivations with the curvature perturbation $\mathcal{R}$, one often define the Mukhanov variable

$$v \equiv a\sqrt{2\varepsilon_1}\mathcal{R}. \tag{3.36}$$

v is also a gauge-invariant quantity that satisfies an equation of motion simpler than $\mathcal{R}$, behaving as an oscillator with a time-dependent frequency:

$$v'' + \left[k^2 - \frac{\left(a\sqrt{\varepsilon_1}\right)''}{a\sqrt{\varepsilon_1}}\right]v = 0. \tag{3.37}$$

To integrate Eq. 3.37, one requires initial conditions for v. Classically, there is not natural way of chosing those. Instead, a quantum origin for the perturbations is natural: fundamentally, quantum fluctuations produce continuously perturbations, even in a vacuum state. The so-called *Bunch-Davies vacuum*, which describes the zero-point fluctuations of the theory leads to well-defined initial conditions [95, 96]:

$$\lim_{k\eta\to-\infty} v_k = \frac{1}{\sqrt{2k}}e^{-ik\eta}. \tag{3.38}$$

Solving Eq. 3.37 requires simultaneously solving the background evolution given by the Friedmann and Klein-Gordon equations, Eqs. 3.23 and 3.22. This can be done exactly numerically or

analytically at a chosen order in slow-roll expansion. Then, one replaces Eq. 3.36 in the definition of $\mathcal{P}_\mathcal{R}$, Eq. 1.94, which yields:.

$$\mathcal{P}_\mathcal{R}(k) = \frac{k^3}{4\pi^2 M_{\rm Pl}^2}\left|\frac{v(k)}{a\sqrt{\varepsilon_1}}\right|^2 . \tag{3.39}$$

For tensor perturbations, the equation of motion of the transverse amplitude of the gravitational wave $\mu(k)$ reads

$$\mu'' + \left(k^2 - \frac{a''}{a}\right)\mu = 0, \tag{3.40}$$

which can be solved with the same initial condition as for scalars. Then, the tensor power spectrum is written as

$$\mathcal{P}_h(k) = \frac{2k^3}{\pi^2}\left|\frac{\mu(k)}{a}\right|^2 . \tag{3.41}$$

Finally, at leading order in slow roll, the amplitudes of $\mathcal{P}_\mathcal{R}$ and $\mathcal{P}_h$ as introduced in Eqs. 1.94 and 1.100 can be written as [14, 15]

$$A_{\rm s} \equiv \mathcal{P}_\mathcal{R}|_{k_*} \overset{\text{SRLO}}{\simeq} \frac{V_*}{24\pi^2 M_{\rm Pl}^4 \varepsilon_{1*}} , \tag{3.42}$$

$$r \equiv \left.\frac{\mathcal{P}_h}{\mathcal{P}_\mathcal{R}}\right|_{k_*} \overset{\text{SRLO}}{\simeq} 16\varepsilon_{1*} , \tag{3.43}$$

where k_* is the CMB pivot scale that we chose to be $k_* = 0.05$ Mpc^{-1}. Stars indicate that the quantities are evaluated at the time of k_* Hubble crossing. The spectral tilts give

$$n_{\rm s} \equiv 1 + \left.\frac{{\rm d}\ln\mathcal{P}_\mathcal{R}}{{\rm d}\ln k}\right|_{k_*} \overset{\text{SRLO}}{\simeq} 1 - 2\varepsilon_{1*} - \varepsilon_{2*} \qquad \text{and} \qquad n_{\rm t} \equiv \left.\frac{{\rm d}\ln\mathcal{P}_h}{{\rm d}\ln k}\right|_{k_*} \overset{\text{SRLO}}{\simeq} -2\varepsilon_{1*}. \tag{3.44}$$

In the specific single-field slow-roll setup, a *consistency relation* appears between the tensor amplitude and tilt, $n_{\rm t} \overset{\text{SRLO}}{\simeq} -r/8$. The runnings give

$$\alpha_{\rm s} \equiv \left.\frac{{\rm d}^2\ln\mathcal{P}_\mathcal{R}}{({\rm d}\ln k)^2}\right|_{k_*} \overset{\text{SRLO}}{\simeq} -2\varepsilon_{1*}\varepsilon_{2*} - \varepsilon_{2*}\varepsilon_{3*} \qquad \text{and} \qquad \alpha_{\rm t} \equiv \left.\frac{{\rm d}^2\ln\mathcal{P}_h}{({\rm d}\ln k)^2}\right|_{k_*} \overset{\text{SRLO}}{\simeq} -2\varepsilon_{1*}\varepsilon_{2*} . \tag{3.45}$$

yielding a second consistency relation: $\alpha_{\rm t} \overset{\text{SRLO}}{\simeq} \frac{r}{8}\left[(n_{\rm s}-1) + \frac{r}{8}\right]$.

One can define several potential regimes. $\varepsilon_2 < 2\varepsilon_1$ yields a *convex* potential $V_{\phi\phi} < 0$. The main contribution to $n_{\rm s}$ comes from ε_1. $2\varepsilon_1 < \varepsilon_2 < 4\varepsilon_1$ also is related to a *convex* potential $V_{\phi\phi} < 0$ but with the main contribution to $n_{\rm s}$ coming from ε_2. $4\varepsilon_1 < \varepsilon_2$ yields a *concave* potential $V_{\phi\phi} > 0$ with a main contribution from ε_2 to $n_{\rm s}$. Let us also mention that expecting ε to be of the same order and consequently $r \sim (1 - n_{\rm s})$ is a too optimistic and naive thought. Tries to bound r from below based on V-shape naturalness have been made, showing that *a priori* low values of r cannot be excluded solely based on such arguments [97, 98].

Moreover, one can predict that the orders of magnitude of non-Gaussianities of any type are small. For instance, for describing the three-point correlations, one uses the bispectrum $B_\zeta(k_1, k_2, k_3)$. To quantify the amount of non-Gaussiantity as compared to the Gaussianity, one

uses a ratio between bispectrum and spectrum, called the $f_{\rm NL}$ parameter[2]. A particular example is the specific "*squeezed*" limit $k \equiv k_1 = k_2 \gg k_3$, in which single-field slow-roll predicts a third consistency relation [99]:

$$f_{\rm NL} \overset{\text{SRLO}}{\simeq} \frac{5}{12}\left(1 - n_s\right). \tag{3.47}$$

3.3.2. Success of slow-roll inflation

From these expressions follows the current *most solid motivation for inflation and, in particular, single-field slow-roll inflation*, which adds up to the super horizon fluctuation and flatness inflationary explanation. Because of the slow-roll parameter hierarchy, single-field slow-roll clearly predicts:

1) *Adiabatic initial conditions*: a single and homogeneous field ϕ is only able to produce perturbations that are described by a unique degree of freedom.

2) Perturbation almost completely *Gaussian*: $f_{\rm NL} \sim \mathcal{O}(\varepsilon)$.

3) Perturbations dominated by *scalar fluctuations*: $r \sim \mathcal{O}(\varepsilon)$.

4) Scalar perturbations with a power spectrum almost *scale invariant*[3]: $|1 - n_{\rm s}| \sim \mathcal{O}(\varepsilon)$.

5) Scalar perturbations with a power spectrum almost a *power law*: $\alpha_{\rm s} \sim \mathcal{O}(\varepsilon^2)$.

3.3.3. Motivation for measuring r

Now that $A_{\rm s}$ and $n_{\rm s}$ are well measured within Λ-CDM, to a precision better than one per cent, a huge success would be to detect the tensor modes and their amplitude, r. This is the main observational target for the next-generation CMB experiments. We quote two of the main motivations of the tensor-to-scalar ratio hunt:

- Through Eqs. 3.42 and 3.43, one sees that the value of r can be directly related to the energy scale of inflation:

$$\rho_* = V_* = 3M_{\rm Pl}^2 H_*^2 \overset{\text{SRLO}}{\simeq} \frac{3}{2}\pi^2 M_{\rm Pl}^4 A_{\rm s} r. \tag{3.48}$$

  Moreover, combining Eqs. 3.23 and 3.28, one can derive an approached expression for the *field excursion* over which slow-roll is supported,

$$\frac{\Delta\phi}{M_{\rm Pl}} \overset{\text{SRLO}}{\simeq} \sqrt{r/8}\Delta N, \tag{3.49}$$

  where ΔN is the total number of inflationary e-folds. Hence, within single-field slow-roll and since we know $A_{\rm s}$, r directly provides two precious clues about the physics governing the inflationary epoch: its energy scale and the typical field excursion.

[2]The scalar $f_{\rm NL}$ is defined as

$$\frac{6}{5} f_{\rm NL} \equiv \frac{B_\zeta\left(k_1, k_2, k_3\right)}{P_\zeta\left(k_1\right) P_\zeta\left(k_2\right) + P_\zeta\left(k_2\right) P_\zeta\left(k_3\right) + P_\zeta\left(k_3\right) P_\zeta\left(k_1\right)}. \tag{3.46}$$

[3]in most of cases, it is also slightly *red-tilted*, since ε_1, by definition, is positive. ε_2 is predicted to be positive for most but not every potential [92].

- A second motivation is to understand the role of quantumness in the generation of primordial perturbations. Most often, one uses quantum mechanics to define the initial conditions and explain naturally the presence of perturbation (see Eq. 3.38). Then, these initial perturbations undergo a *quantum-to-classical* transition on super-horizon scales, which is reviewed in [100]. An observational confirmation of the quantumness of scalar and tensor perturbations would be a major breakthrough for fundamental physics.

3.3.4. Constraints on the slow-roll parameters

We have already introduced in Sec. 2.4 the current observational constraints on the primordial power spectra ($\ln\left(10^{10}A_{\rm s}\right) = 3.047 \pm 0.014$, $r < 0.032$ at 95% CL, $n_{\rm s} = 0.967 \pm 0.004$, $\alpha_{\rm s} = -0.0042 \pm 0.0067$) which we have also illustrated in Fig. 2.7. Within the single-field slow-roll inflation framework, these can be translated into constraints on the slow-roll parameters directly, thanks to slow-roll equations Eqs. 3.42, 3.43, 3.44 and 3.45.

This has been, for instance, done in [101] within a Bayesian framework and assuming slow-roll prior $|\varepsilon_{i*}| < 0.2$. The data are from Planck 2018 [102], the BK21 BB [65], SPT TE and EE [103] and BAOs [104].

The one-dimensional marginalised posterior distributions for the non-primordial cosmological parameters and the third-order Hubble-flow parameters $A_{\rm s}$, ε_{1*}, ε_{2*}, ε_{3*} and ε_{4*} are represented on Fig. 3.2.

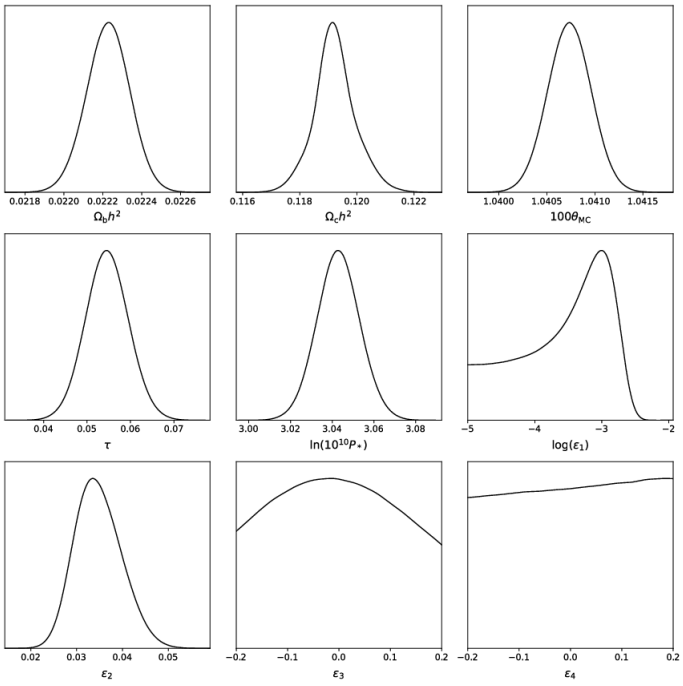

Figure 3.2: From [101]. One-dimensional marginalised posterior distributions for the non-primordial cosmological parameters and the third-order Hubble-flow parameters $A_{\rm s}$ (here denoted P_*), ε_{1*}, ε_{2*}, ε_{3*} and ε_{4*}. The used data set are Planck 2018 [102], the BK21 BB [65], SPT TE and EE [103] and BAOs [104].

Already, with constraints in this slow-roll parameter space, potential-independent conclusions

could be drawn. However, the most interesting results arise when propagating these constraints to the potential parameters [101]. In order to do so, the first step is to derive observational predictions from a given inflationary potential V. This can be done through solving the background dynamical equations Eqs. 3.22 and 3.23, computing the field value of the inflaton at the time of Hubble crossing of the pivot scale, evaluating Eqs. 3.33, 3.34 and 3.35 at that location, and insert the result into Eqs. 3.42 and 3.45. We defer this discussion for Part IV of the thesis.

Conclusion

In this chapter, we have shown how inflation solves three main conceptual issues with the standard model of cosmology, explaining the presence of correlation on superhorizon scales, the observed flatness of space-time today, and the absence of heavy GUT relics like magnetic monopoles. Inflation also provides a simple mechanism to explain the origin of perturbations: unavoidable quantum fluctuations in the inflationary vacuum are stretched to superhorizon scales and transition to the classical regime before reentering the horizon during the radiation era.

In the particular single-field slow-roll scenario, which is a simple and natural realisation of inflation, predictions of the initial scalar and tensor perturbations can be done. Namely, slow-roll perturbations are scalar-dominated, Gaussian and adiabatic, according to a quasi-scale-invariant power-law spectrum. These predictions are all in perfect agreement with the data

Detecting r, improving our constraints on the Hubble slow-roll parameters, and constraining inflation and its phenomenology are some of the main motivations for future CMB experiments. This ambitious science program comes with challenges. The signals to detect are faint and elusive, requiring unprecedented sensitivity, exquisite knowledge of the microwave sky and control of instrumental effects, as well as rigorous data-analysis strategies. We will depict this generically in the last introductory chapter.

CHAPTER 4

Challenges and target for future CMB experiments

This section aims to introduce the challenges and observational targets for next-generation CMB experiments. Our microwave sky is not limited to the CMB, whose main effects have been overviewed in Chapter 2 in temperature and polarisation. Therefore, we begin by discussing the foreground effects that overlay this picture, particularly the lensing of the CMB and various astrophysical emissions. Next, we review the instrumental effects that induce noise and systematic errors. We then introduce a tool for estimating the model parameter covariance matrix and forecasting future experiments. These initial sections set the stage for presenting the future prospects for CMB observation aimed at addressing various new scientific targets. Finally, we will introduce the main instruments planned for deployment in the coming years and their respective observational strategies.

Contents

4.1 Effects on the late line-of-sight 64
4.1.1 Lensing 64
4.1.2 Foregrounds 65
4.1.3 Atmosphere 72
4.2 Instrumental effects 72
4.2.1 Statistics 72
4.2.2 Beam size 73
4.2.3 Systematics 73
4.3 Estimating or forecasting cosmological constraints 74
4.3.1 Likelihood function 74
4.3.2 Fisher analysis 75
4.3.3 Actual data-analysis 77
4.4 Current and future observational targets 78

4.4.1 Small scales 78
4.4.2 Large scales 79
4.5 Next-generation CMB experiments 83
4.5.1 Ground-based experiments 83
4.5.2 Space-based experiments 83
4.5.3 Other perspectives in cosmology 84

4.1. Effects on the late line-of-sight

Cosmologists aim to extract cosmological parameters from the data. However, several effects intervene between the emission from the last scattering surface and the writing of a cosmology paper. These effects can be categorised into three classes: the first affects photons during their journey to Earth, the second is associated with instrumental effects, and the third relates to data analysis.

We have previously discussed three well-modelled late-time effects that contribute to altering the CMB power spectra during the photons' travel from the last scattering surface to the detector: the stretching of the wavelength, the Integrated Sachs-Wolfe (ISW) effect, and the reionisation rescattering. Additionally, there are two more challenging astrophysical effects that impact the maps: the lensing of the CMB and the microwave emission from various foregrounds.

4.1.1. Lensing

As photons from the CMB traverse the late-time Universe, they undergo deflection due to the gravitational lensing exerted by the large cosmic structures. This lensing effect distorts the primordial CMB spots around foreground masses, mainly around $z \simeq 1$.

The lensing of the CMB photons around Large Scale Structure (LSS) mass can be described by a remapping of the CMB photons. At linear order and in the flat-sky approximation, one can show that photons at a position of the sky $\boldsymbol{\theta}_{\rm true} = (\theta^1_{\rm true}, \theta^2_{\rm true})$ are deflected to $\boldsymbol{\theta}_{\rm obs} = (\theta^1_{\rm obs}, \theta^2_{\rm obs})$ by [105, 106]

$$\Delta\boldsymbol{\theta} = \frac{\partial}{\partial\boldsymbol{\theta}}\phi^L(\boldsymbol{\theta}), \tag{4.1}$$

where $\phi^L(\boldsymbol{\theta})$ is the lensing potential that can be expressed as an integral along the line of sight of all the deflections occurring between ourselves and the last scattering surface, hence involving the gravitational potential ψ of an LSS at a comoving distance $\chi_{\rm LSS}$, times a geometric factor that accounts for where the deflection happens:

$$\phi^L(\boldsymbol{\theta}) = -2\int_0^{\chi_{\rm LSS}} {\rm d}\chi_1 \psi\left[\boldsymbol{x}(\boldsymbol{\theta}, \chi_1), \eta_0 - \chi_1\right)\left(\frac{\chi_{\rm LSS} - \chi_1}{\chi_1\chi_{\rm LSS}}\right). \tag{4.2}$$

At linear order in $\Delta\boldsymbol{\theta}$, $\Delta T = T(\theta_{\rm obs}) - T(\theta_{\rm true})$ and $\Delta(Q \pm iU) = (Q \pm iU)(\theta_{\rm obs}) - (Q \pm iU)(\theta_{\rm true})$ write as

$$\Delta T = \Delta\boldsymbol{\theta}\frac{\partial T}{\partial\boldsymbol{\theta}}, \tag{4.3}$$

$$\Delta(Q \pm IU) = \Delta\boldsymbol{\theta}\frac{\partial(Q \pm iU)}{\partial\boldsymbol{\theta}}. \tag{4.4}$$

Consequently, this deflection modifies the observed statistics of the maps at the power spectrum level and introduces non-gaussian features in the maps. The power spectra expressions can be found, for instance, in [106]. They involve $C_\ell^{\phi\phi}$, the angular power spectrum of the gravitational lensing itself. The latter can be linked to the matter power spectrum today, P_ψ, which is another important cosmological observable that can be modelled within $\Lambda-$CDM

$$C_\ell^{\phi\phi} = 4\int_0^{\chi_{\rm LSS}} d\chi_1 \left(\frac{\chi_{\rm LSS}-\chi_1}{\chi_1^2\chi_{\rm LSS}}\right)^2 P_\psi\left(k=\frac{2\ell+1}{2\chi_1}, z(\chi_1)\right). \tag{4.5}$$

We show in Fig. 4.1 the effect of the lensing on the various CMB spectra. For all spectra, large-scale lenses blur the peaks of the observed power spectrum, while small-scale lenses induce leakage from large-scale towards smaller scales. Furthermore, a consequent increase in B-mode power is sourced by the lensing of E modes. Below $\ell = 300$, this effect can be described as an extra $\mathcal{O}(5\ \mu{\rm K}-{\rm arcmin})$ white noise component. Its exact amplitude and shape depend on the assumed Λ-CDM and on its extensions (in particular in the neutrino sector). The primordial tensor modes being already constrained to be faint, they lead to a BB contribution that is very subdominant at medium and small scales with respect to the lensing contribution. Hence, the lensing cosmic variance dominates the BB variance, representing the first significant challenge for primordial tensor discovery.

To tackle this lensing BB challenge, a technique called *CMB delensing* consists of subtracting the lensing at the map level in order to decrease the non-primordial signal and the associated variance. It consists in estimating the deflection angle $\Delta\theta$, and correcting the $Q \pm iU$ maps for it. To do so, a good knowledge of ϕ^L is required, which itself requires a matter tracer to estimate ψ, see Eq. 4.2. This tracer can be measurements of the matter power spectrum today or CMB gravitational lensing maps measurements.

From a CMB experiment, gravitational lensing maps themselves can be measured. They are a remarkable tracer of mass in the Universe at low redshift. In Fig. 4.2, we show the ACT DR6 CMB lensing mass map [107], which is the best CMB lensing measurement done to date. Then $C_\ell^{\phi\phi}$ can be estimated and compared to the model given by Eq. 4.5, yielding promising (and late time) cosmological information. See the lower panel of Fig. 2.6 for the 2018 picture of lensing, including data from ACTpol [108] and Planck [109], which has been improved since by advance ACTpol [107]. This extraction of the lensing maps, power spectra and cosmological information from CMB data is not easy task, and an active field of research focuses on providing optimal estimators of the observables (see *e.g.* [110, 111]).

Beyond the power spectra, to catch the non-Gaussian features due to the lensing, one typically assumes that $\boldsymbol{\Delta\theta}$ is small (Born approximation). The lensing effect can hence be simulated from generating mass haloes that seed a Websky lensing convergence map [112]. A more detailed approach is the ray-tracing technique, see *eg.* [113].

4.1.2. Foregrounds

Several astrophysical foregrounds emitting in the microwave domain are superimposed on the (almost) black-body CMB signal. They are characterised by their Spectral Energy Density (SED), their locality, and their statistical properties through the angular power spectrum.

The SED can be given in three different units. The first is the *surface brightness per solid angle*, $I_\nu(\hat{n})$, which quantifies the intensity of radiation coming from a given line of sight $\hat{n}$ as a function of wavelength. This measure represents the amount of energy per surface area, per solid

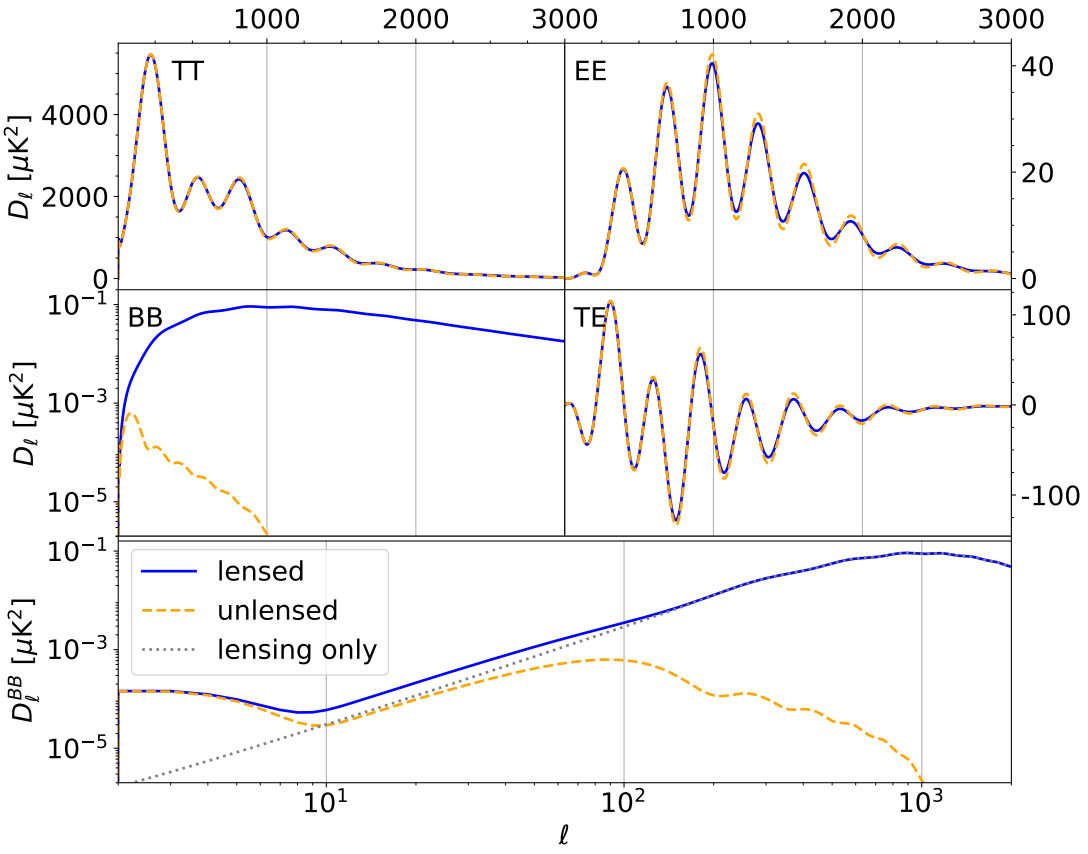

Figure 4.1: *TT*, *EE*, *BB*, *TE* Λ-CDM theoretical D_ℓ power spectra computed including (plain blue curve) or not (dashed orange curve) the lensing effect, assuming Tab. 4.1 and $r = 10^{-2}$. The lower panel illustrates the *BB* spectrum in logarithmic x-scale.

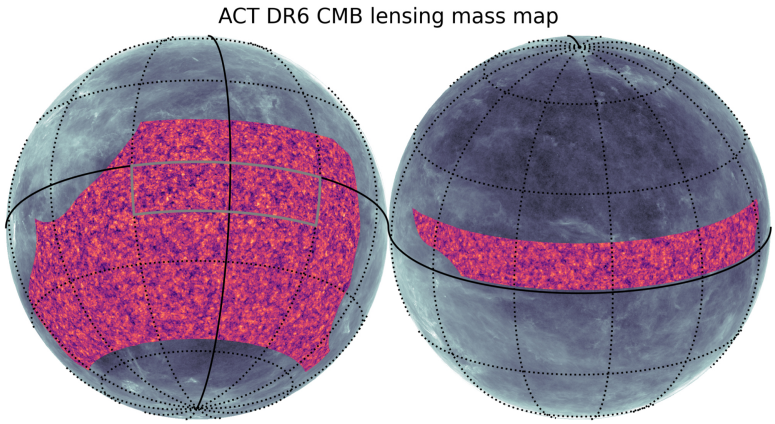

Figure 4.2: ACT lensing map.

angle, and per frequency, and it is the most natural and common way to measure the intensity of an astrophysical source, with units of MJy sr^{-1}.

In the CMB context, a more practical measure of intensity (which is also the one we have used so far) is the *thermodynamic temperature* T, obtained by identifying I_ν with the Planck law Eq. 1.38. Since the CMB is nearly a perfect black body, its thermodynamic temperature is frequency-independent along a given line of sight, hence $T(\nu, \hat{n}) = T_0(\hat{n})$. In this work, unless otherwise specified, the sky maps and spectra describing the variation of temperature with respect to the mean are given as a thermodynamic temperature,

$$\Delta T(\hat{n}) = T_0(\hat{n}) - T_0, \tag{4.6}$$

where $T_0 \equiv \langle T_0(\hat{n}) \rangle = 2.7255$ K. The unit of thermodynamic temperature is denoted as K hereafter unless the context involves standard temperatures, in which case we specify the unit notation as $\mathrm{K_{CMB}}$. One can compute the relation to convert a surface brightness into CMB unit:

$$\Delta T = I_\nu \frac{T_0^2}{2\nu^4} \frac{\left(e^{\nu/T_0} - 1\right)^2}{e^{\nu/T_0}}. \tag{4.7}$$

The last measure of radiation intensity is the *Rayleigh-Jeans (RJ) temperature*, T_{RJ}, defined by identifying I_ν to the Rayleigh-Jeans law, the low-frequency branch of the Planck law,

$$T_{\mathrm{RJ}} = \frac{I_\nu}{2\nu^2}. \tag{4.8}$$

The diffuse Galactic foregrounds contribute significantly to the total observed signal [51].

Galactic dust emission

Dust grains in the interstellar medium are constantly heated by the interstellar radiation field to a temperature between 10 and 30 K, depending on the dust grain and incident radiation properties [51, 114]. They thermally re-emit photons, whose distribution peaks between 1000 and 3000 GHz and is often modelled as the product between the blackbody spectrum and a power law:

$$I_\nu^{\mathrm{d}} \propto A_{\mathrm{d}} \nu^{\beta_{\mathrm{d}}} B_\nu \left(T_{\mathrm{d}}\right). \tag{4.9}$$

where A_{d} is the dust amplitude at a reference frequency, β_{d} is the dust power-law index and T_{d} is the dust blackbody temperature. The resulting root-mean-square RJ temperature frequency spectrum, for numerical values consistent with intensity observations, is represented in red on the upper side of Fig. 4.3 (we will describe this figure iteratively throughout the next paragraphs).

The localisation of the emission is illustrated in the left panel of Fig. 4.4, which represents the 353 GHz RJ total amplitude map as observed by Planck. Moreover, the dust grains' rotation axes, which correspond to their short axes, tend to align with the local magnetic field [115]. Consequently, the re-emitted radiation is significantly polarised, up to 25% [116, 117]. The resulting polarised dust emission follows a modified blackbody distribution with parameters *a-priori* independent of the total intensity ones [118]. A sketch of the root-mean-square distribution is represented on the lower panel of Fig. 4.3, while the corresponding 353 GHz RJ polarised amplitude map from Planck is displayed on the right panel of Fig. 4.4. Galactic dust emission is very localised in the Galactic plane.

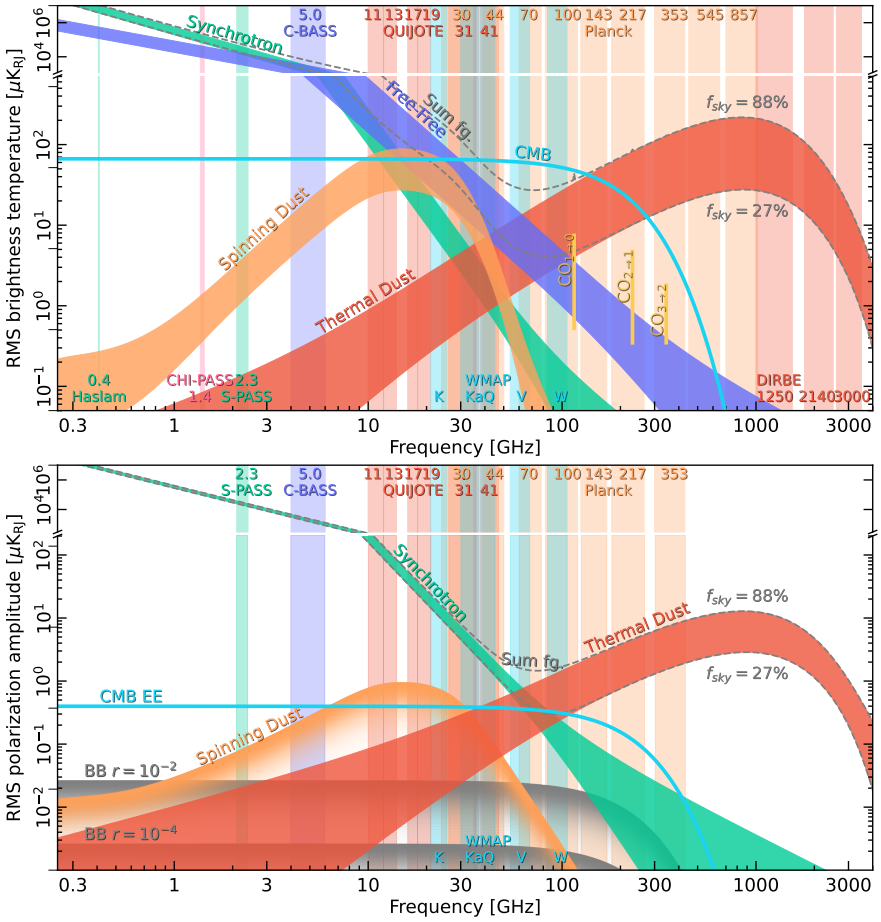

Figure 4.3: Root-mean-square TOP: temperature, BOTTOM: polarisation amplitude, in Rayleigh-Jeans temperature units, as a function of the frequency for the various sky components, obtained by fitting Planck data with BEYONDPLANCK [51] models. Vertical bands indicate the frequency ranges of various experiments. The effect of masking the sky on the spectra amplitudes is illustrated by the widths of the curves. See [51] for more details.

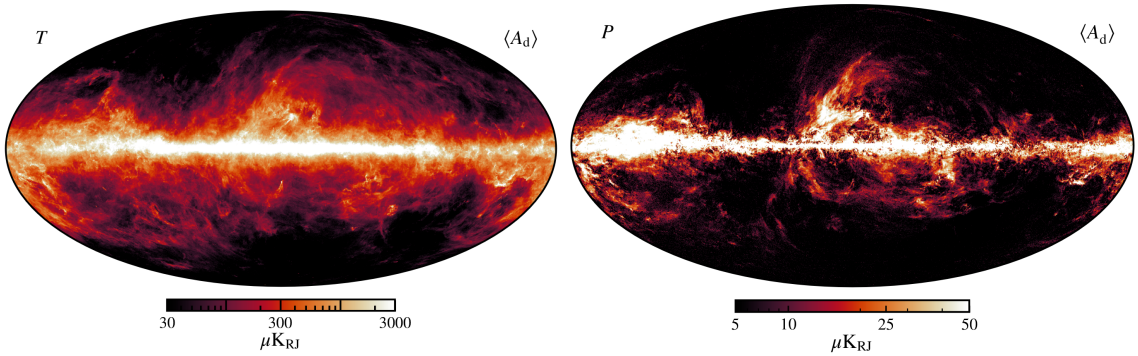

Figure 4.4: Dust amplitude maps from Planck as obtained by the COMMANDER component separation. LEFT: total intensity map at 545 GHz, from [51], RIGHT: total polarisation map at 353 GHz, from [118].

Galactic synchrotron emission

Synchrotron emission is generated by relativistic electrons spiralling around magnetic field lines in the interstellar medium [51, 114]. The emitted radiation is characterised by a power-law spectrum, which can be expressed as [119]:

$$I_\nu^{\mathrm{s}} \propto A_{\mathrm{s}} \nu^{-\beta_{\mathrm{s}}}, \tag{4.10}$$

where A_{s} is the synchrotron amplitude at a reference frequency and β_{s} is the synchrotron spectral index, typically ranging from 2.5 to 3.0. The resulting root-mean-square RJ temperature distribution for synchrotron emission, using values consistent with intensity observations, is depicted in green on the upper panel of Fig. 4.3. The spatial distribution of this emission is illustrated in the left panel of Fig. 4.5, showing the 30 GHz RJ total amplitude map as observed by Planck.

Furthermore, synchrotron emission is highly polarised due to the alignment of the electron orbits with the magnetic field lines. The degree of polarisation can be as high as 70%, depending on the magnetic field's regularity and the electrons' energy distribution. The polarisation follows the same power-law distribution but with parameters that can differ from the total intensity. The root-mean-square distribution of the polarised synchrotron emission is illustrated on the lower panel of Fig. 4.3, while the corresponding 30 GHz RJ polarised amplitude map from Planck is shown at the right of Fig. 4.5. Synchrotron emission is more diffuse compared to galactic dust emission but is still more prominent along the Galactic plane and regions with strong magnetic fields.

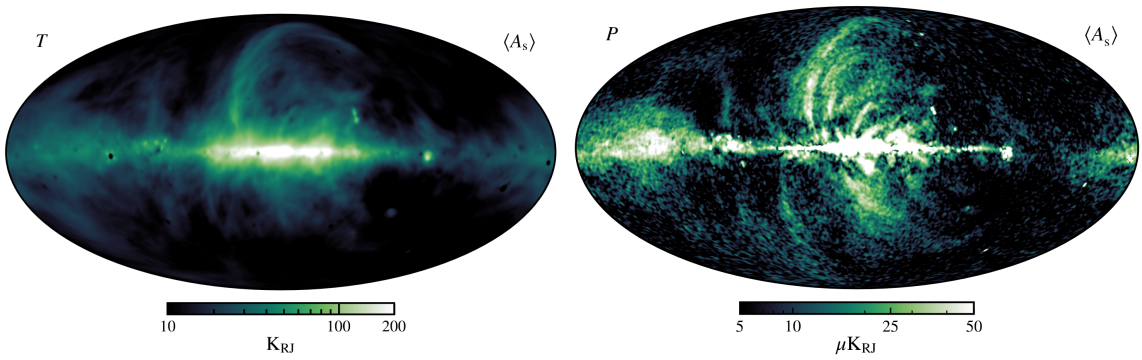

Figure 4.5: Synchrotron amplitude maps from Planck as obtained by the COMMANDER component separation. LEFT: total intensity map at 30 GHz, from [51], RIGHT: total polarisation map at 30 GHz, from [118].

Free-free emission

Free-free or bremsstrahlung emission is generated by the acceleration of free electrons in the vicinity of ions in ionised regions of the interstellar medium, such as HII regions [51, 114]. This process results in the emission of photons with a well-known spectrum shape [120, 121], which can be approximated as [51]

$$I_\nu^{\mathrm{ff}} = A_{\mathrm{ff}} \nu^{-2}, \tag{4.11}$$

where A_{ff} is the free-free amplitude at a reference frequency; see the blue curve in Fig. 4.3. The associated Planck intensity map is represented in the right panel of Fig. 4.6. Unlike synchrotron

and dust emissions, free-free emission is mostly unpolarised [122]. The intensity of free-free emission, which primarily contributes to the CMB foreground at frequencies below 100 GHz, can be modelled using a power-law spectrum.

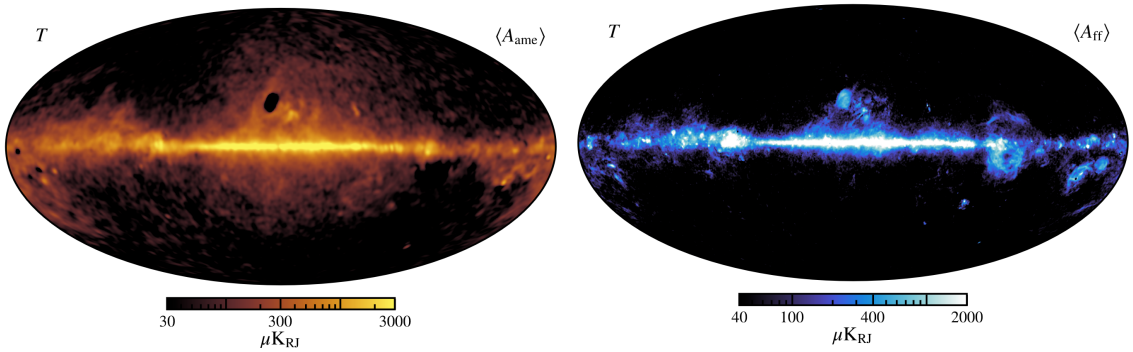

Figure 4.6: Total intensity amplitude maps from Planck as obtained by the COMMANDER component separation. LEFT: AME at 22 GHz, from [51], RIGHT: Free-free emission at 40 GHz, from [51].

Anomalous Microwave Emission

Anomalous Microwave Emission (AME) is usually believed to originate from the rapid rotation of small, interstellar dust grains, also known as spinning dust grains [51, 123–127]. These grains emit in the microwave range, with a spectrum that peaks between 20 and 60 GHz; see orange curve in Fig. 4.3. The associated Planck intensity map is represented in the left panel of Fig. 4.6. AME constitutes a CMB intensity foreground that can be distinguished by its distinct frequency dependence and lack of polarisation [128].

CO lines emissions

Carbon monoxide (CO) rotational transition lines are another significant foreground in CMB measurements. These spectral lines arise from the emission of CO molecules in molecular clouds and are found at specific frequencies, such as 115 GHz and 230 GHz; see yellow lines in Fig. 4.3. The associated molecular cloud map emitting intensity from CO lines is represented in the right panel of Fig. 4.7. CO emission in intensity is generally modelled with a delta function and knowledge of the molecular cloud distribution in the Milky Way, and its emission is only weakly polarised [129].

Point source emissions

Point source emissions in the microwave range are primarily from extragalactic sources, such as active galactic nuclei, quasars, and radio galaxies [114]. These sources emit across a broad range of frequencies and are characterised by their compact nature, appearing as distinct points in CMB maps. Point source emissions can significantly contaminate the CMB signal, especially at higher frequencies. They are typically subtracted from CMB data by using specific source catalogues and frequency-dependent templates.

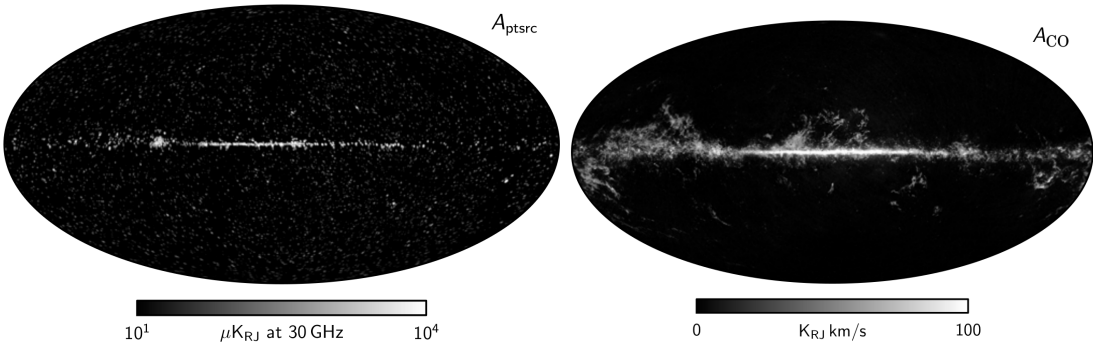

Figure 4.7: Total intensity amplitude maps from Planck as obtained by the COMMANDER component separation. LEFT: Point sources at 30 GHz, from [114] RIGHT: CO lines at 100 GHz integrated over all velocities emission from [114].

Cosmic Infrared Background emission

Cosmic Infrared Background (CIB) emission originates from the cumulative infrared light emitted by distant and unresolved galaxies, detected for the first time by COBE [130]. Even though the CIB contributes to the microwave foregrounds, particularly at higher frequencies above 100 GHz, it is less intense than the other foregrounds at large scales. The CIB emission is diffuse and unpolarised.

Sunyaev-Zel'dovich effect

The Sunyaev-Zel'dovich (SZ) [131] effect consists of two components: thermal SZ (tSZ) and kinetic SZ (kSZ). The tSZ effect is caused by the scattering of CMB photons off hot electrons in galaxy clusters, resulting in a distortion of the CMB spectrum that increases the intensity at high frequencies and decreases it at low frequencies. The kSZ effect arises from the motion of galaxy clusters relative to the CMB, causing a Doppler shift in the scattered photons. Both effects are crucial for studying the properties of galaxy clusters. They are known to be weakly polarised, probably non-coherently across the sky.

Mitigating the foreground emissions

An efficient strategy to limit the impact of foregrounds in CMB maps and angular power spectra involves masking the Galactic regions where foregrounds are localised. Doing so increases the CMB-to-foreground ratio because the CMB is diffuse over the entire sky. The fraction of the sky kept in the analysis is hereafter denoted as $f_{\rm sky}$. The impact of sky masking on the foreground level is illustrated in Fig. 4.3 by the width of the curves.

An approximate formula to generalise the covariance and variance expressions Eqs. 2.49 and 2.51 to partial-sky power-spectrum estimate reads

$$\sigma^2 \to \sigma^2 / f_{\rm sky} \tag{4.12}$$

$$\Xi \to \Xi / f_{\rm sky}. \tag{4.13}$$

The additional complexity induced in the analysis due to cutting the sky will be discussed in Sec. 8.

A complementary approach to cleaning CMB maps and gaining information on foregrounds involves applying component separation techniques that take advantage of the multiple observed frequencies. We will introduce and investigate such foreground cleaning in Sec. 7.

4.1.3. Atmosphere

Atmospheric effects play a significant role in ground-based CMB experiments [60, 132]. The atmosphere introduces additional noise and systematic errors due to its emission and absorption properties, which vary with weather conditions and frequency. Water vapour and oxygen are the primary sources of atmospheric emission, creating a variable and frequency-dependent foreground. Fluctuations in atmospheric conditions can introduce time-varying signals, known as atmospheric noise, which can correlate with the CMB signal and complicate data analysis. Mitigating atmospheric effects involves conducting observations from high-altitude, dry sites, such as the Atacama Desert or the South Pole, using atmospheric monitoring data, optimised scanning strategies and dedicated data processing techniques.

4.2. Instrumental effects

Now that we have reviewed the various astrophysical effects along the line of sight, we also need to introduce the different instrumental effects inherent to any CMB experiment: statistical noise, finite resolution of the detectors, and various potential systematic effects.

4.2.1. Statistics

In order to control instrumental and astrophysical systematics, modern CMB experiments observe the microwave sky at several frequencies, each coming with its own sensitivity to CMB. The sensitivity σ_ν, typically in μK·rad or μK·arcmin, is directly related to the number of detectors $N_{\rm det}$ in the considered frequency channel as well as the observation time $t_{\rm obs}$ and the sky fraction $f_{\rm sky}$:

$$\sigma_\nu \propto \sqrt{\frac{f_{\rm sky}}{N_{\rm det} t_{\rm obs}}}. \tag{4.14}$$

In CMB unit and for a fixed design, the frequencies near the CMB blackbody peak will come with a higher signal-to-noise ratio. Of course, at the end of the day, this number depends mainly on the instrumental design. We further discuss how the sensitivities are obtained in the particular LiteBIRD case in Sec. 6.3.

For an experiment that is perfectly sensitive to polarisation, *i.e.* we can access the polarisation state of each detected photon, the polarisation sensitivity is worsened by a factor $\sqrt{2}$ with respect to the temperature sensitivity.

$$\sigma_P = \sigma_E = \sigma_B = \sqrt{2}\sigma_T. \tag{4.15}$$

This is due to the polarised power that is shared between the two states of polarisation. Assuming no noise correlation between the frequency channels, the total combined sensitivity of the experiment is given by:

$$\sigma_{X,\rm comb} = \left(\sum_\nu \sigma_{X,\nu}^{-2} \right)^{-1/2}, \tag{4.16}$$

where $X = T$ or P. In angular power spectrum space, if the data are uncorrelated, this noise level translates into a white noise, which results a flat N_ℓ spectrum with respect to ℓ. The corresponding amplitude is σ_P^2 for EE or BB, σ_T^2 for TT, and $\sigma_P\sigma_T$ for TE. If instead, the data are correlated, they can lead to a non-flat noise power spectrum, we introduce an example in Sec. 4.2.3.

The covariance (and the variance) of the full-sky $\hat{C}_\ell^{XY}$ estimator can be obtained with the following substitution in Eq. 2.51 (and Eq. 2.49):

$$C_\ell \to C_\ell + N_\ell. \tag{4.17}$$

Moreover, in the case of auto power-spectrum, the estimate will also be biased by N_ℓ. Instead, when doing cross-spectra between independent data splits, the power spectrum will remain unbiased on average (see *eg.* [133]).

For what regards the pixels distribution, adding noise implies doing the substitution $\mathbf{M} \to \mathbf{M} + \mathbf{N}$ in Eq. 2.34, where $\mathbf{N}$ is the noise covariance matrix of shape $(3n, 3n)$. Equivalently, the noise can be specified at the power spectrum level in Eqs. 2.42 thanks to the substitution Eq. 4.17. In the case of a simple white noise specified by its dispersion σ in μK·rad, $\mathbf{N}$ reduces to a diagonal matrix with a constant $\frac{12N_{\rm side}^2}{4\pi}\sigma^2$ [μK^2] on the diagonal.

4.2.2. Beam size

Each detector exhibits a specific angular beam response to signal on the sky that determines the angular resolution of the resulting map. The beam effect is a decrease in the signal-to-noise for small angular separations, which one can quantify as a beam window function W_ℓ that erases the signal at the highest multipoles. Assuming each frequency ν is associated with a symmetric Gaussian beam response parameterised by a full-width half maximum angle $\theta_{\rm FWHM,\nu}$ (as well as a sensitivity $\sigma_{X,\nu}$), the white noise spectrum is not flat anymore, and instead follows

$$N_\ell^{XY,\ \rm inst} = \left[\sum_\nu \frac{1}{\sigma_{X,\nu}\sigma_{Y,\nu}} \exp\left(-\ell(\ell+1)\frac{\theta^2_{\rm FWHM,\nu}}{8\log 2}\right)\right]^{-1}. \tag{4.18}$$

Modifying the pixel distribution Eq. 2.34 to account for the beam can be done through applying a first window function $W_\ell^{{\rm beam},XY}$ in multipole space to Eq. 2.42. On top of this unavoidable instrumental beam, to implement successfully Eq. 2.42, an additional beam window function has to be included, which damps smoothly the scales smaller than the characteristic size of the pixelisation ($W_\ell^{{\rm smooth},XY}$ goes smoothly from 1 to 0 at $\ell = 3N_{\rm side} - 1$). This is a precaution to avoid map aliasing caused by a too-sharp multipole cut. Lastly, one should include the pixel window function $W_\ell^{\rm pix}$ that we have introduced in Sec. 2.3.3.

To account for these various window functions, one performs the transformation $C_\ell^{XY} \to C_\ell^{XY}\left(W_\ell^{XY}\right)^2$ in Eqs. 2.42, where $W_\ell^{XY} \equiv W_\ell^{{\rm beam},XY} \times W_\ell^{{\rm smooth},XY} \times W_\ell^{\rm pix}$ is the total window function. We will introduce in Sec. 8.3.1 the different window functions we will work with in this thesis.

4.2.3. Systematics

CMB measurement experiments are highly sensitive and require meticulous handling of various systematic effects that can introduce errors in the data. We will rediscuss some of them later

on in the LiteBIRD context as correlated $(1/f)$ noise, unmodelled-beam systematics, unmodelled-bandpass systematic, or misorientation of the detectors.

On top of these instrumental systematics, there usually is a contribution to the systematic budget due to foregrounds and to component separation, which also induces a degradation of the statistical noise (both of which we further discuss in Sec. 7).

Treating these effects requires accurate calibration and characterisation of the instrument before a mitigation that can be done at different stages of the data analysis. Without entering into this, to simplistically account for the induced residual systematic and statistical after component separation in a simple ℓ-space model, one can include a degradation factor Δ and a low-ℓ residual tail in the noise curve Eq. 4.18 (which can be parameterised by a power law of amplitude $A_{\rm fg}$ and spectral index $n_{\rm fg}$):

$$N_\ell^{XY,\rm tot} = \Delta^{XY} N_\ell^{XY,\ \rm inst} + A_{\rm syst}^{XY} \ell^{n_{\rm syst}^{XY}} \tag{4.19}$$

In the last two sections, we have introduced the various sources of astrophysical and instrumental contamination and noise. Eq. 4.19 applied to Eq. 4.17 and 2.49 is a rough model that summarises the degradation of the variance with respect to noiseless formula Eq. 2.49.

4.3. Estimating or forecasting cosmological constraints

Before introducing the next-generation CMB experiments and their respective observational strategy in Sec. 4.5, we want to identify in Sec. 4.4 the various regions in the CMB spectra that need to be targeted to address the main cosmological science goals. To do this, we need a tool that can predict the expected uncertainty on a given parameter, accounting for cosmic variance and any N_ℓ. We will use the *Fisher formalism*, which derives this order of magnitude from the curvature around the maximum of the *likelihood function*, that we introduce now. We will also use this section to outline a typical data-analysis procedure, which we will elaborate on in Part. III.

4.3.1. Likelihood function

First, we introduce the so-called *likelihood function*. Let X be a random variable following an absolutely continuous probability distribution with density function $f_\theta(x)$ (a function of x that depends on a parameter θ). Then the *likelihood function* is

$$\mathcal{L}(\theta \mid x) = f_\theta(x), \tag{4.20}$$

a function of θ, given the outcome $X = x$.

The *maximum likelihood estimate* of θ is given by

$$\hat{\theta} = \arg\max_\theta \mathcal{L}\left(\theta \mid x\right), \tag{4.21}$$

which is usually consistent, meaning that as the number of observations n goes to infinity, the estimator $\hat{\theta}$ converges in probability to its true value, which we call hereafter θ_0. Said otherwise, $\langle\hat{\theta}\rangle = \theta_0$ where $\langle . \rangle$ is the ensemble average.

For this reason, the maximum likelihood estimate is widely used in data analysis. We will apply it to several contexts in Part III. A word of caution is required: for limited statistics, the estimator may remain biased. This should be estimated through simulations. This remark

is especially relevant for CMB data analysis in which our sample is intrinsically limited in size because of cosmic variance. We shall rediscuss its consequence in Sec. 9.3.

4.3.2. Fisher analysis

The *Fisher information* [134] quantifies how much information an observable random variable X contains about an unknown parameter θ of the probability density function $f_\theta(x)$ that models X. This object is used to calculate the covariance matrix of the parameter estimate associated with maximum likelihood estimates. Hence, the Fisher analysis of the likelihood function has proven to be a very useful tool for forecasting future experiments with low computational cost. It is widely used in cosmological and CMB contexts (see for *e.g.* [135–138]). We have implemented this tool, and its results will be used for illustrative purposes throughout this manuscript. In this section, we introduce it for CMB application. Given some data vector $\mathbf{D}$ of shape n that plays the role of the outcome x, to estimate the expected uncertainty in the estimation $\{\hat{\theta}\}$ of true parameters $\{\theta_0\}$, one focuses on the shape of the likelihood function $\{\theta\} \to \mathcal{L}(\{\theta\} \mid \mathbf{D})$ around its maximum. This can be done by expanding the log-likelihood in a Taylor series around its maximum $\{\hat{\theta}\}$:

$$\ln \mathcal{L} = \ln \mathcal{L}\big|_{\{\hat{\theta}\}} - \frac{1}{2}(\theta_i - \hat{\theta}_i)\mathcal{H}^{ij}\big|_{\{\hat{\theta}\}}(\theta_j - \hat{\theta}_j) + ... \tag{4.22}$$

where the Hessian matrix is defined as

$$\mathcal{H}_{ij} \equiv -\frac{\partial^2 \ln \mathcal{L}}{\partial\theta_i \partial\theta_j}. \tag{4.23}$$

The Fisher matrix is defined as the variance of the *score*, which itself is the derivative of $\ln \mathcal{L}$ with respect to $\{\theta\}$:

$$F_{ij} \equiv \left\langle \frac{\partial \ln \mathcal{L}}{\partial\theta_i} \frac{\partial \ln \mathcal{L}}{\partial\theta_j}\bigg|_{\{\hat{\theta}\}} \right\rangle, \tag{4.24}$$

form that can be rewritten as[1]:

$$F_{ij} = \left\langle \mathcal{H}_{ij}|_{\{\hat{\theta}\}} \right\rangle = -\left\langle \frac{\partial^2 \ln \mathcal{L}}{\partial\theta_i \partial\theta_j}\bigg|_{\{\hat{\theta}\}} \right\rangle. \tag{4.25}$$

If some data sets are independent, their combined likelihood will be the product of the likelihoods of each data set, and subsequently, their combined Fisher matrix will be the sum of the individual Fisher matrices. If there happens to be some redundancies in the data sets, their contributions to the Fisher matrix should be accounted for only one time, so one should correct the naive sum with the Fisher matrix of the redundant data.

The Fisher matrix describes the curvature of the likelihood function around its maximum in the parameter space. If the higher order terms in the Taylor expansion Eq. 4.22 are negligible (asymptotic limit), the inverse Fisher matrix coincides with the covariance matrix of the maximum-likelihood estimated parameters. More generally, the Fisher matrix provides a lower bound on the covariance matrix elements of any unbiased estimator of θ,

$$\sigma_{ij}^2 \geq \left(F^{-1}\right)_{ij}, \tag{4.26}$$

[1] use that for $x \to u(x)$, $\frac{\partial^2 \ln u}{\partial x^2} = \frac{\partial^2 u}{u \partial x^2} - \left(\frac{\partial u}{u \partial x}\right)^2$

called *Cramér-Rao bound* [139, 140] and which follows from Cauchy-Schwarz inequality. Hereafter, we denote the Fisher uncertainty as $\sigma_{i,\mathrm{stat}}^{\mathrm{Fisher}} \equiv \left(F^{-1}\right)_{ii}$.

One can elaborate on Eq. 4.25 in the particular case of Gaussianly distributed data $\mathbf{D}$ of shape n with mean $\boldsymbol{\mu} = \langle \mathbf{D} \rangle$ and covariance matrix of shape (n, n) (in this discussion bold symbols represent vector or matrices over data indices):

$$\boldsymbol{\Sigma} \equiv \left\langle (\mathbf{D} - \boldsymbol{\mu})(\mathbf{D} - \boldsymbol{\mu})^{\mathrm{T}} \right\rangle, \tag{4.27}$$

where the T exponent denotes the transpose. Then, the likelihood takes the following multivariate Gaussian shape:

$$\mathcal{L} = \frac{1}{(2\pi)^{n/2} \ |\mathrm{det}\boldsymbol{\Sigma}|^{1/2}} \exp\left[-\frac{1}{2}(\mathbf{D} - \boldsymbol{\mu})^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} (\mathbf{D} - \boldsymbol{\mu})\right]. \tag{4.28}$$

In this specific case, the Fisher matrix can be rewritten as a Trace over the data indices:

$$F_{ij} = \frac{1}{2}\mathrm{Tr}\left[\boldsymbol{\Sigma}^{-1}\frac{\partial \boldsymbol{\Sigma}}{\partial \theta_i}\boldsymbol{\Sigma}^{-1}\frac{\partial \boldsymbol{\Sigma}}{\partial \theta_j}\right] + \frac{\partial \boldsymbol{\mu}^{\mathrm{T}}}{\partial \theta_i}\boldsymbol{\Sigma}^{-1}\frac{\partial \boldsymbol{\mu}}{\partial \theta_j}. \tag{4.29}$$

In this work, we will reuse this equation in two contexts: power spectrum estimation from pixel data (see Sec. 8.1.4) and cosmological parameters inference from the harmonic coefficients. Let's focus on the latter context for now and take as data the vector $\mathbf{A}_{\ell m}$ we have defined in Eq. 2.27 and Σ is given by $\mathbf{C}_\ell$ introduced in Eq. 2.33. Because of rotational invariance, one can simply transform the sum over m for each ℓ into $(2\ell + 1)$ factor, and Eq. 4.29 becomes

$$F_{ij} = \sum_\ell \frac{2\ell + 1}{2}\mathrm{Tr}\left(\mathbf{C}_\ell^{-1}\ \frac{\partial \mathbf{C}_\ell}{\partial \theta_i}\mathbf{C}_\ell^{-1}\ \frac{\partial \mathbf{C}_\ell}{\partial \theta_j}\right). \tag{4.30}$$

After a lengthy but straightforward matricial derivation, this equation can be equivalently rewritten in a more elegant form:

$$F_{ij} = \sum_\ell \mathrm{Tr}\left(\frac{\partial \tilde{\mathbf{C}}_\ell^T}{\partial \theta_i}\ \boldsymbol{\Xi}_\ell^{-1}\ \frac{\partial \tilde{\mathbf{C}}_\ell}{\partial \theta_j}\right), \tag{4.31}$$

where $\tilde{\mathbf{C}}_\ell$ is the vector we have previously defined in Eq. 2.50, and Ξ, the minimal covariance matrix that will come with any C_ℓ estimate due to the cosmic variance has been defined in Eq. 2.51.

The Eq. 4.31 form of the Fisher matrix can be simply interpreted as a usual *propagation of uncertainty* formula: the Fisher matrix converts information from C_ℓ to θ_i space. Each coefficient in the trace quantifies the contribution from one multipole of one spectrum to the total information on the parameter. Since the information we are discussing is the curvature in the N dimensional parameter space, it is simply given by a (N, N)-shaped matrix. Hence, by construction, the 1-σ contours one will be able to forecast by inverting the Fisher matrix will be optimistic ellipses, which will not capture non-Gaussian features in parameter space.

To include the statistical uncertainty coming from the instrumental statistical and systematic sources introduced in Sec. 4.2, one performs the substitution Eq. 4.17 in Eq. 4.31 and Eq. 2.51, where N_ℓ is as defined by Eq. 4.19. Consequently, the Fisher elements decrease and the forecasted σ_{ij}'s increase.

Furthermore, one can approximate the effect of partial but still significant coverage of the sky by rescaling the number of degrees of freedom, see Eq. 4.13. Hence, the forecasted σ_{ij}'s scale as $1/\sqrt{f_{\rm sky}}$. However, this rescaling is only approximate: when cutting the sky, $\mathbf{A}_{\ell m}$ is no longer distributed according to a multivariate Gaussian, and the derivation described in this section is no longer to be applied to the new problem.

Finally, there exists a simple prescription that has already been tested in a cosmological context to propagate a (small) systematic bias in ℓ space $\tilde{\mathbf{\Delta}}_\ell$ (defined similarly to Eq. 2.50) to a bias estimate for cosmological parameters $\sigma^{\rm Fisher}_{i,\rm syst}$, based on the Fisher matrix formalism [141–144]:

$$\sigma^{\rm Fisher}_{i,\rm syst} = \left(F^{-1}\right)_{ij} \sum_\ell {\rm Tr}\left(\frac{\partial \tilde{\mathbf{C}}_\ell^T}{\partial\theta_j}\, \mathbf{\Xi}_\ell^{-1}\, \tilde{\mathbf{\Delta}}_\ell\right) \tag{4.32}$$

We will use this formula alongside the $\sigma^{\rm Fisher}_{i,\rm stat}$ given by the Cramér-Rao bound in Sec. 7.4.3.

4.3.3. Actual data-analysis

Before continuing, it is important to emphasise that Fisher analysis is limited, as it only provides a covariance matrix. One adopts, in real cases, more sophisticated tools to extract the cosmological information from the likelihood function.

A significant part of the real data-analysis flowchart is dedicated to preprocessing the data set D to prepare the estimation of the parameters. Typically, the CMB data-analysis flowchart consists of transforming time-ordered data to maps, cleaning the foregrounds and possibly the lensing before building a power spectrum that is often the preferred input of (non-gaussian) likelihoods. The data-analysis flowchart in the particular context of LiteBIRD is introduced in Sec. 5.3 before we entirely describe a particular pipeline in Part III.

Then, two main prescriptions coexist for cosmological parameters inference from the obtained data:

- Frequentist perspective: To infer a set of parameters $\{\theta_i\}$, we maximise the likelihood function to obtain an estimate $\{\hat{\theta}_i\}$. Then, one typically works from the profile likelihoods for a subset of parameters, typically one specific θ_j, by sampling it and maximising over the remaining $\{\theta_i\}_{i\neq j}$. Finally, one can build confidence intervals from this profile.

- Bayesian perspective: To investigate the distribution of a set of parameters $\{\theta_i\}$ given the data D and a model $\mathcal{M}$, we need to estimate the posterior probability $\mathcal{P}(\{\theta_i\}) = f_{D,\mathcal{M}}(\{\theta_i\})$. The Bayes' theorem relates this posterior probability to the likelihood $\mathcal{L}(\{\theta_i\}) = f_{\{\theta_i\},\mathcal{M}}(D)$, prior $\Pi(\{\theta_i\}) = f_{\mathcal{M}}(\{\theta_i\})$ and to the Bayesian evidence $\mathcal{E} \equiv f_{\mathcal{M}}(D)$:

$$\mathcal{P}(\{\theta_i\}) = \frac{\mathcal{L}(\{\theta_i\})\Pi(\{\theta_i\})}{\mathcal{E}}. \tag{4.33}$$

  In the inference problem, we are interested in estimating $\{\theta_i\}$ for a fixed model, and the evidence normalisation can be forgotten. Then, a Bayesian analysis is interested in the posterior distribution, which is typically obtained by sampling it with Monte-Carlo Markov chains. The best-fit values and confidence intervals are then built from the parameter's marginal posterior distributions:

$$\mathcal{P}(\theta_j) = \int_{\Pi(\{\theta_i\})} {\rm d}\{\theta_i\}_{i\neq j} \mathcal{P}(\{\theta_i\}). \tag{4.34}$$

4.4. Current and future observational targets

In this section, we introduce non-exhaustively various science targets for future CMB experiments. To illustrate these, we use the Fisher formalism introduced in Sec. 4.3.2. This Sec. 4.4 aims at justifying the strategies adopted by future CMB experiments that we will introduce in Sec. 4.5. Hence, our aim is not to use Fisher formalism to forecast future experiments for now. For this reason, we will not use a N_ℓ specific to a particular experiment. Instead will mostly remain focused on the mock case in which cosmic variance dominates the error budget.

Throughout this thesis, we chose by default to work with given fiducial values for Λ-CDM that are summarised in Table 4.1. In particular, these values serve in the upcoming discussion as fiducial values for the computation of the Fisher matrices.

$\Omega_b h^2$	$\Omega_c h^2$	H_0	n_s	A_s	τ	r	n_t	Ω_k	$n_\mathrm{s,run}$	N_eff	m_ν
0.02237	0.12	67.36	0.965	2.099×10^{-9}	0.0544	0	0	0	0	3.046	0.06

Table 4.1: Fiducial cosmological parameters used throughout this work (including for producing the LiteBIRD simulations we will use in Part III).

4.4.1. Small scales

Small-scale measurements of temperature and polarisation will enable us to tackle many science goals, for instance:

1) refine our knowledge of Λ-CDM parameters globally,
2) leverage our constraints on the primordial scalar power spectrum,
3) refine our constraints on light relics and neutrino masses.

4.4.1.1 Refining our knowledge of Λ-CDM through small scales

The first target for future CMB experiments is to refine our constraints on the six Λ-CDM parameters (see Sec. 2.4). Planck temperature statistics measurement is cosmic variance limited for what concerns the larger scales ($\ell \lesssim 1500$), so we expect this improvement to primarily come from polarisation measurements at large and medium scales, which we will focus on in this thesis, but also from small-scale measurements.

An interesting quantity for evaluating the amount of information contained in a data subset is the Fisher *Figure of Merit* (FoM) (see *eg.* [145]). It can be defined as the characteristic size of the allowed fraction of parameter space relative to a reference. The characteristic size can be obtained from the n-th root of the allowed parameter-space volume, where n is the number of parameters. In our case, the FoM is normalised such that it is 100 at $\ell = 20000$, so

$$\mathrm{FoM} = 100 \times \left(\frac{V_\ell}{V_{\ell=20000}}\right)^{1/n} = 100 \times \left(\frac{\sqrt{|F_\ell^{-1}|}}{\sqrt{|F_{\ell=20000}^{-1}|}}\right)^{1/n} . \tag{4.35}$$

This quantity should be interpreted as the average factor that will be gained on the uncertainty of the parameters.

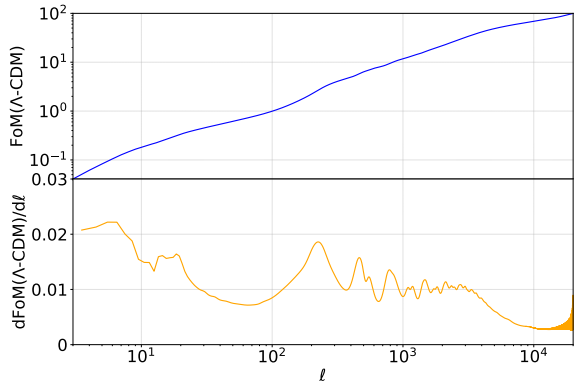

Figure 4.8: Λ-CDM figure of merit with respect to $\ell = 20000$. This quantity is more detailed in the core of the text.

For illustration, we represent in Fig. 4.8 the information that is still to be learnt from small scales up to $\ell = 20000$, including the scalar TT, EE, TE and BB ($r = 0$ and we only assume cosmic variance). The upper panel represents the FoM with respect to ℓ, and the lower panel is its derivative, which represents the Λ-CDM information that is carried in each multipole relative to the others.

This figure has an intuitive interpretation: dominated by cosmic variance, we will approximately gain one FoM decade per maximum-multipole decade. For instance, going from multipole 1500 to 15000, we gain a bit less than a factor of ten on the average parameter uncertainty.

4.4.1.2 Scalar initial conditions through small scales

Let us complement the previous discussion by mentioning the potential impact on our knowledge of the scalar primordial spectrum. $A_{\rm s}$, $n_{\rm s}$ but also its running will benefit from extending the CMB window in the high k region, see the Planck reconstructed primordial power spectrum Fig. 2.7 and the approximate correspondence between ℓ and k. In the particular slow-roll case, this would allow one to reduce the inflationary parameter space; see Sec. 3.3.4 and constrain the underlying models; see Part. IV.

4.4.1.3 Light relics and neutrino masses through small scales and lensing

Another important focus of future CMB experiments is to place tighter constraints on light relics, such as additional neutrino species [146], and neutrino masses [147]. Small-scale measurements of the CMB power spectrum and of CMB lensing are particularly sensitive to the presence of light relics through the $N_{\rm eff}$ parameter and the sum of neutrino masses $\sum m_\nu$.

4.4.2. Large scales

Several science goals can be tackled on large scales. In this section, we develop two of the most important ones:

1) improving our constraints on the reionisation parameter τ and history, which also allows degeneracy breaking with $A_{\rm s}$ and $\sum m_\nu$,

2) improving our constraints on the primordial BB spectrum and on r.

4.4.2.1 The reionisation history through large scale polarisation

As we have mentioned in Sec. 2.2.4, apart from low-ℓ polarisation, the effects of $A_{\rm s}$ an τ are extremely similar: both contribute to an overall amplitude. Consequently, the two quantities are degenerate along the $A_{\rm s}e^{-2\tau}$ direction and are hence positively correlated. This combination represents the overall amplitude that is measured by CMB experiments. This degeneracy is broken by adding the information from low-ℓ polarisation: a bump specific to the reionisation appears for multipoles lower than $\sim$ 20 in EE, TE, and BB; see Fig. 2.3. EE is the most informative because the effect can be much larger than cosmic variance, as we will illustrate soon. The Planck measurements of these scales allowed for an unprecedented measurement of the EE bump, sufficiently breaking the degeneracy to allow for a 10% measurement of τ [53]. However, the degeneracy in Planck data is still significant [148], and the uncertainty can still be improved by at least a factor 3 just from a better measurement of the reionisation bump.

For illustration, let us consider a noiseless full-sky measurement of temperature between $\ell = 2$ and $\ell = 2500$. The left panel of Fig. 4.9 illustrates the Fisher estimate of the correlation between the 6 Λ-CDM parameters when adding to this T measurement a $2 \leq \ell \leq 20$ E-mode measurement, which comes with a white noise whose level $\sigma^{{\rm low}EE}$ is represented in the x-axis (in μK·arcmin). One observes that several parameters are severely degenerated in temperature alone (*i.e.* when $\sigma^{{\rm low}EE}$ is high), especially $\Omega_{\rm c}$ and H_0 which are almost entirely anti-correlated, while τ and $A_{\rm s}$ are almost entirely correlated. In the other limit, where $\sigma^{{\rm low}EE}$ becomes negligible with respect to cosmic variance, one sees that τ and $A_{\rm s}$ lose a bit less than 10% of their correlation. On the right panel, we illustrate the errors for each parameter marginalised over the other ones, normalised by the error in TT only. This demonstrates that the impact of this degeneracy breaking is massive: more than a factor of five is gained on both $A_{\rm s}$ and τ uncertainties.

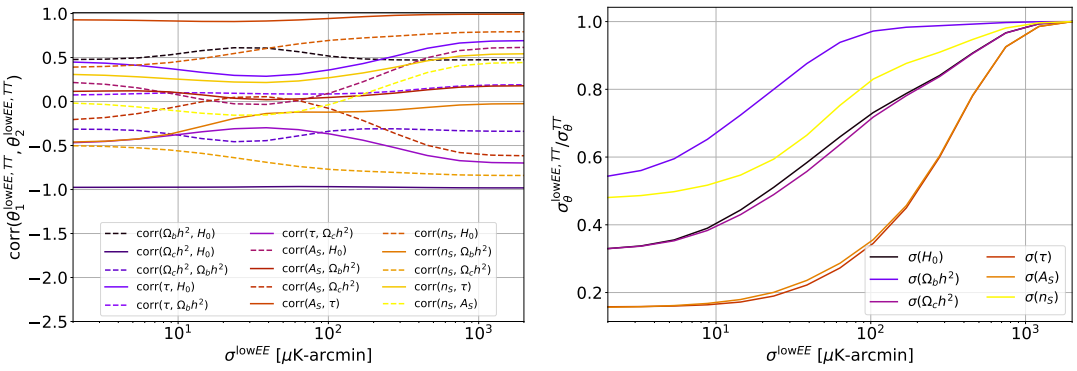

Figure 4.9: Fisher estimate of the correlation between the 6 Λ-CDM parameters (left panel) and their errors (right panel) when adding to this T measurement a $2 \leq \ell \leq 20$ E-mode measurement, which comes with a white noise whose level $\sigma^{{\rm low}EE}$ is represented in the x-axis (in μK·arcmin). The errors on the parameters are marginalised over Λ-CDM and normalised to the TT-only uncertainties.

We have identified how τ interacts with the other parameters and, in particular, with $A_{\rm s}$. We have seen that $A_{\rm s}$ is degenerated with τ even after including the information from the reionisation

bump. The reason for that is the extremely limited number of modes in the reionisation bump with respect to the rest of the spectra: only a few small ℓ that come with large cosmic variance. It is technically impossible for this effect to give more information on τ than the information the remaining of the spectra give on the overall amplitude $A_{\rm s}e^{-2\tau}$.

We illustrate in the upper-left panel of Fig. 4.10 the *EE* power spectrum for $\tau \in [0.9, 1.1] \times \tau_{\rm fid}$ and $A_{\rm s} = A_{\rm s,fid}$ where $\tau_{\rm fid}$, $A_{\rm s,fid}$ and the rest of Λ-CDM are given in Table 4.1. At other parameters fixed, the reionisation bump is clearly distinct, with a maximum multipole of $\ell = 9$ to $\ell = 22$ for τ going from 0.02 to 0.09. One also sees the effect of τ at higher multipoles as the overall amplitude.

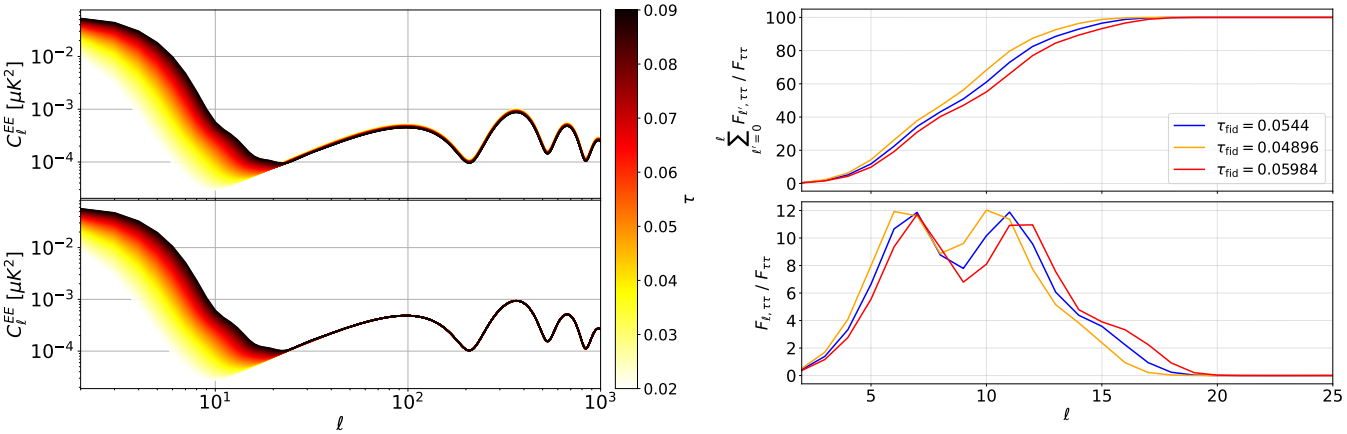

Figure 4.10: UPPER-LEFT: *EE* power spectrum for $\tau \in [0.9, 1.1] \times \tau_{\rm fid}$ and $A_{\rm s} = A_{\rm s,fid}$ and the rest of Λ-CDM are given in Table 4.1. LOWER-LEFT: Same but also fixing $A_{\rm s} = A_{\rm s,fid} \exp\{-2(\tau_{\rm fid} - \tau)\}$ instead. LOWER-RIGHT: Fisher information on r in percents of the total information in the 300 first multipoles as a function of ℓ, for $\tau = 0.9 \times \tau_{\rm fid}$ (orange), $\tau = \tau_{\rm fid}$ (blue), $\tau = 0.9 \times \tau_{\rm fid}$ (red). UPPER-RIGHT: Cumulative of the upper-right-panel curves.

In the following, we will never fix $A_{\rm s}$ while letting τ free: we have seen that the residual correlation between the two parameters is too high. Instead, we prefer to fix $A_{\rm s}e^{-2\tau}$ directly to its fiducial value[2]. It is provided with a <1% precision by the high-ℓ spectra alongside the other Λ-CDM parameters, and it is not more correlated to τ in the reionisation bump than the other Λ-CDM parameters. Moreover, with this new parameterisation, when varying the model, nothing varies but the reionisation bump. The lower-left panel of Fig. 4.10 shows the resulting spectra, which we will use in Part. III.

Finally, with this particular parameterisation, we can derive the Fisher information multipole per multipole with Eq. 4.31 (and to which the FoM introduced previously reduces). We illustrate it in the lower-right panel of Fig. 4.10 for three different τ fiducial values and normalised to the total Fisher information contained in the reionisation bump. The upper panel represents the same quantity in cumulative. It confirms that all the information in contained between $\ell = 2$ and $\ell \simeq 19$, 50% below $\ell = 2$ and $\ell = 7$ and the remaining fraction $\ell = 8$ and $\ell \simeq 19$.

The focus of Part. III will be to derive a τ estimate from LiteBIRD simulated frequency maps, based on one of these fiducials, $\tau = 0.0544$, including noise and foregrounds. We will concentrate

[2]Our reason is data-driven, but one should remember that the two parameters originate from two very different physics and hence should be inferred separately in the end

on the constraints gained on τ from measuring the reionisation bump. It's important to notice that these constraints will subsequently improve the current constraints on $A_{\rm s}$ indirectly, by nearly the same order of magnitude, since $A_{\rm s}$ and τ exhibit a nearly 100% degeneracy when all information is considered. Additionally, the sum of the neutrino masses, $\sum m_\nu$, measured by the small scales and by the lensing, is also degenerate with τ and will thus benefit from the reionisation τ measurement.

4.4.2.2 Tensor initial conditions through B-modes

Finally, as we mentioned in Sec. 3.3.3, the most important target for next-generation CMB experiments is r. Even though we can constrain r from temperature and E-modes, see Sec. 2.3, the constraints are driven by BB. As we have illustrated in Fig. 4.1, the BB signal, hence its variance, is also contributed to by lensing B-modes, which is one of the (several) reasons that make a r detection particularly challenging.

The left panel of Fig. 4.11 represents the C_ℓ^{BB} spectrum for r values between 10^{-5} and 10^{-1}. In this particular case, we fix $n_{\rm t}$ with the slow-roll consistency relation Eq. 3.44. As we mentioned in Sec. 4.1.1, one notices that the lensing C_ℓ, like a white noise, is quasi-flat with respect to ℓ. For the particular fiducial models illustrated here, the tensor part of the spectrum is visible up to $\ell \simeq 200$, with two characteristic features: the *reionisation bump*, up to $\ell \simeq 12$, followed by the *recombination bump* from $\ell \simeq 12$.

In order to visualise at which scales the constraint on r comes from, we represent on the right panel the Fisher information multipole-per-multipole on the lower panel while the cumulative information is illustrated on the upper panel, for a noiseless 50% of the sky experiment[3]. Both y axes are expressed as a percentage of the total Fisher information contained in the fiducial BB spectrum under scrutiny. We represent eight of these for four different values of r and two delensing setups; we denote A_L, the fraction of lensing removed in map-space, which allows one to cancel the associate spectrum and its variance by the same fraction. We consider a full-lensed and a 95%-delensing cases (*resp.* $A_L = 1$ and $A_L = 0.05$).

One sees that in the case in which $r = 0$, depicted by the superimposed green dashed and plain curves, the maximum amount of information comes from the lowest multipoles, $\ell = 2, 3, 4$. Instead, for the high values of r, most information comes from the smallest scales. The delensing favours the impact of the recombination bump. The second focus of Part. III will consist of deriving a r estimate from LiteBIRD simulated frequency maps based on one of these fiducials, $r = 0$ and $A_L = 1$, including noise and foregrounds.

We have mentioned that on top of r, one can put constraints on the tensor tilt [62, 63], with the potential to exclude slow-roll if Eq. 3.44 is broken. However, we show in App. A that confirming this relation is out of the technological reach of future CMB experiments, requiring unrealistic combinations of sensitivity, resolution and delensing level.

Of course, behind these constraints on initial conditions, the motivation is to understand better inflation as well as reheating, excluding models and restricting theory parameter spaces. This will be the topic of our Part. IV.

Our list of observational CMB targets is not exhaustive, omitting for instance to discuss the search for cosmic birefringence [149, 150] or for non-Gaussianities [151].

[3] of course, these results will change a lot when including noise and foregrounds.

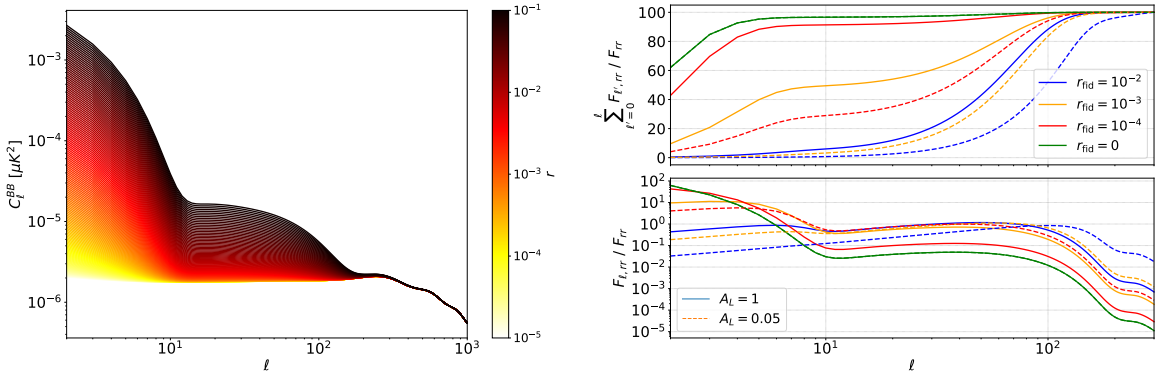

Figure 4.11: LEFT: *BB* power spectrum for fixed Λ-CDM given in Table 4.1 + $r \in [10^{-5}, 10^{-1}]$ and $n_T = -r/8$. BOTTOM-RIGHT: Fisher information on r in percents of the total information in the 300 first multipoles as a function of ℓ, for $r = 10^{-2}$ (blue), $r = 10^{-3}$ (orange), $r = 10^{-4}$ (red), for $r = 0$ (green), on 50% of the sky, when delensing 95% of the B maps (dashed) or not (plain). UPPER-RIGHT: Cumulative of the upper-right-panel curves.

4.5. Next-generation CMB experiments

As a last discussion of this introductory part, we introduce the various experiments to come in the next decade, especially focusing on CMB ones. Today, two main strategies coexist for measuring CMB anisotropies: ground-based experiments and space-mission telescopes.

4.5.1. Ground-based experiments

Future ground-based experiments are designed to maximise sensitivity to the CMB polarisation. They include BICEP, SPT, and the Simons Observatory (SO), each employing proper observational strategies.

The BICEP Array, which is located at the South Pole, focuses on deep observations of a small sky patch. Future iterations of BICEP aim to reach a sensitivity of $r < 0.01$ within the next few years [60]. A delensing program jointly with small-scale measurements from SPT in the same patch of the sky will be used.

The Advanced Simons Observatory [152], based in the Atacama Desert in Chile, aims to provide complete measurements of the CMB temperature and polarisation across various angular scales. SO employs both large aperture telescopes (LATs) for wide-field surveys and small aperture telescopes (SATs) for deep observations of smaller regions. The observatory is currently beginning to take data and expects to reach a sensitivity of $r \simeq 0.0012$. SO's detector number (60 000) and technology are expected to significantly improve over current experiments, yielding a constraint on r competitive with space missions.

4.5.2. Space-based experiments

Space CMB experiments offer several key advantages over ground-based experiments [153].

Firstly, space-based observations allow access to all frequencies, which is not possible from the ground due to interference from water and oxygen lines. This limitation on the ground prevents the possibility of having a leverage arm on the foregrounds. Secondly, space experiments have enhanced sensitivity compared to those on the ground because they are free from atmo-

spheric loading. This difference in sensitivity increases with frequency; hence, a significantly more powerful leverage arm on dust is possible from space. As a rule of thumb, one space-based detector can be as effective as around 100 ground-based detectors of similar quality. Thirdly, space-based measurements avoid the substantial brightness fluctuations caused by atmospheric emissions. Finally, satellite experiments can observe the entire sky, allowing measurement down to the lowest multipole $\ell = 2$. Optimal scanning strategies in space allow full and uniform scanning of the sky while maintaining a substantial angular separation between the telescope's boresight and the Earth and the Sun, minimising the induced systematic errors.

The choice of space has been previously made on several occasions. Three satellites have been launched and delivered CMB data in the last decades, COBE, WMAP, and Planck that we have already introduced in Sec. 2.4.

The LiteBIRD mission [153] is the proposed spacecraft mission that is currently in the most advanced phase, with a launch expected around 2032 Japanese Fiscal Year. The study of its Instrument Model and the preparation of its data analysis occupies a central place in this thesis.

4.5.3. Other perspectives in cosmology

Once again, we have not been exhaustive. Another important CMB perspective that we did not mention involves refining the COBE-FIRAS measurements of the spectral shape of the CMB emission, which might present distortions that could contain rich information about cosmology [154]. It is, for instance, the goal of a balloon project called BISOU [155].

Of course, the CMB is not the only cosmological probe. Several large-scale survey experiments aiming to describe the galaxy distribution statistics are expected to start taking data in the next decades, *e.g.*, the Vera C. Rubin Observatory [156], Euclid [157], and the WFIRST / Nancy Grace Roman Space Telescope [158, 159]. The BAO peak in the matter correlation will be measured with increasing sensitivity at various redshifts, *e.g.*, by DESI [160]. Radio telescopes like the Square Kilometre Array (SKA) will map the distribution of neutral hydrogen in the universe, opening an entirely new view of cosmology. Finally, it will be crucial to cross-correlate the data from all these upcoming datasets to extract all the information from them [161].

Conclusion

In this chapter, we have detailed the various challenges that any forthcoming CMB experiments aiming to measure polarisation will have to address. The targeted signals are very subtle, requiring highly sensitive experiments. One of the main targets for future CMB experiments is r, which must be constrained from scales where the BB lensing significantly affects the data. Furthermore, the presence of intense polarised foregrounds around the characteristic CMB frequency $\nu_{\rm peak}$ necessitates multifrequency measurements of the sky, masking of the most contaminated regions, and cleaning of the maps. An exquisite knowledge of instruments is needed to control and mitigate systematic effects. Finally, optimised data-analysis strategies are required in order to derive robust constraints from the complex data.

We have also seen that two main strategies, ground-based experiments and space missions, address these challenges differently, both with ambitious science goals (targeting an r uncertainty on the order of 10^{-3}). The focus of the two coming parts of this thesis will be on investigating the space-mission strategy, through the preparation of the LiteBIRD spacecraft. In Part. II, we

will introduce the LiteBIRD Instrument Model and its implementation. This model is essential for maintaining, tracking, and updating our understanding of the instrument, and it will be used as a basis for simulations and data analysis. In Part. III, we will develop a specific data-analysis pipeline designed to extract information from the large-scale E- and B-modes.

Part II

The LiteBIRD mission and its Instrument Model

In the introductory Part. I, we have detailed the standard model of cosmology, cosmological inflation, the CMB and its measurement challenges, and its current and future status. This Part. II is dedicated to introducing a spacecraft experiment to be launched in 2032 Japanese fiscal year, LiteBIRD, whose main goal is to refine our constraints on the tensor-to-scalar ratio. In Chapter. 5, we will provide an overview of the mission and of the experiment. In Chapter. 6, we will focus on the LiteBIRD Instrument Model (IMo) and will detail our choices in its implementation for the collaboration.

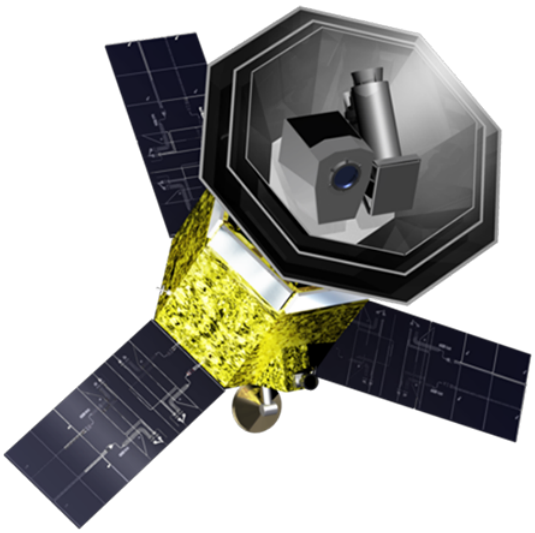

Figure II.1: From *ISAS/JAXA*. Artistic view of the LiteBIRD spacecraft.

CHAPTER 5

LiteBIRD overview

LiteBIRD, the *Lite (Light) satellite for the study of B-mode polarisation and Inflation from cosmic background Radiation Detection*, is a space mission that aims at mapping the cosmic microwave background polarisation over the entire sky for three years from the Sun-Earth Lagrangian point L2. The targeted sensitivity on the tensor-to-scalar ratio (δr=0.001 assuming r=0) poses a significant challenge: the statistical errors being pushed down, and the error budget becomes dominated by systematics. To make measurements at the required level, nano-Kelvin precision rejection of foregrounds and systematic artefacts is required. This challenge will be first addressed by a multi-frequency observation of the sky, by deploying three telescopes operating across 22 frequency bands spanning from 34 to 448 GHz. In addition, mitigating as many instrumental systematic effects as possible and a meticulous understanding of the residual ones will be required. This chapter provides an overview of the LiteBIRD requirements and instrumental design, widely based on [153].

Contents

5.1 Requirements 91

5.2 Characteristics and design 92

5.2.1 Spacecraft and Service Module 93

5.2.2 Payload Module and telescopes 94

5.2.3 Detection chain 96

5.3 Data-analysis processes and analyses 98

5.1. Requirements

The *full success criterion* of LiteBIRD is to achieve $\delta r < 0.001$ for a fiducial model with $r = 0$, where δr is the total error on the tensor-to-scalar ratio, including every source of uncertainties.

The LiteBIRD requirements are organised into five levels. The first one, Lv1, contains the top-level quantitative science requirements directly connected to the full success of the mission. They consist of two requirements [153].

> Lv1.01: *The mission shall measure r with a total uncertainty of $\delta r < 1 \times 10^{-3}$. This value shall include contributions from instrumental statistical noise fluctuations, instrumental systematics, residual foregrounds, lensing B-modes, and observer bias, and shall not rely on future external data sets.*

For more context on the various sources of uncertainty quoted here, one may refer to Secs. 4.1 and 4.2. The second requirement, made to cover the case where r turns out to be large, precise the required polarisation angular power spectrum measurement capability:

> Lv1.02: *The mission shall obtain full-sky CMB linear polarisation maps for achieving $> 5\sigma$ significance using $2 \leq \ell \leq 10$ and $11 \leq \ell \leq 200$ separately, assuming $r = 0.01$. We adopt a fiducial optical depth of $\tau = 0.05$ for this calculation.*

This requirement indicates that the sensitivity on the BB spectrum measurement assuming $r = 0.01$ should come independently from reionisation and recombination bumps, *cf.* Fig. 4.11.

The second level, Lv2, contains the measurement requirements needed to achieve Lv1 given the program-level constraints. There are eleven of them. The error budget is defined so $\sigma_{r,\mathrm{stat}} = \sigma_{r,\mathrm{syst}} < 0.6 \times 10^{-3}$ (Lv2.01, Lv2.02, Lv2.06), and the error allocation of the systematic error is detailed by (Lv2.07). The observer bias (Lv2.10) should be much smaller than $\sigma_{r,\mathrm{syst}}$. The angular resolution at the lowest frequency should have an FWHM smaller than 80 arcmin to reconstruct the signal up to $\ell = 200$ (Lv2.04). At the same time, the scanning strategy should allow a full scan of the sky to obtain the B modes to the lowest multipole of $\ell = 2$ (Lv2.03). The duration of the normal observation phase will be 3 years (Lv2.08) and the spacecraft will be in a Lissajous orbit around the Sun-Earth L2 point (Lv2.09). The required calibration measurements and noise-covariance knowledge are respectively developed in (Lv2.05) and (Lv2.11).

These two first levels contain mission and measurement requirements independently of any instrumental design choice. The lower-level system requirements are derived from them based on instrumental trade-off studies. Lv3 contains the top-level instrument requirements needed to achieve Lv2, while Lv4 and Lv5, respectively, contain the lower-level instrument and unit requirements. We do not detail the requirements composing these three lower levels, but the LiteBIRD instrumental design we will overview in the next two sections is based on them and has been built to satisfy the top-level requirements.

5.2. Characteristics and design

This section provides an overview of the LiteBIRD baseline design at the spacecraft and service module stages and at the payload and telescope stages.

5.2.1. Spacecraft and Service Module

LiteBIRD will be launched on an H3, Japan's new flagship rocket, from the Tanegashima Space Center to the Lagrange Point L2. It should reach its destination within approximately 100 days before being inserted on a Lissajous orbit around L2. After cooling, test, and calibration observations, the entire sky will be surveyed for three years.

The determination of LiteBIRD's spacecraft structure stems directly from the mission requirements. An axisymmetric shape is chosen to facilitate spinning and minimise the moment of inertia. The *Payload Module* (PLM), housing the telescopes, is positioned at the spacecraft's top, while the solar panels are placed at its bottom, perpendicular to the spin axis. On the satellite's underside, opposite the mission instruments, the high-gain antenna is positioned to face Earth, avoiding interference with the telescopes. Fig. 5.1 illustrates the spacecraft's basic structure. This whole structure slowly spins (0.05 rotation per minute) with a precession angle optimised for an observation of the entire sky (see Sec. 6.2.1).

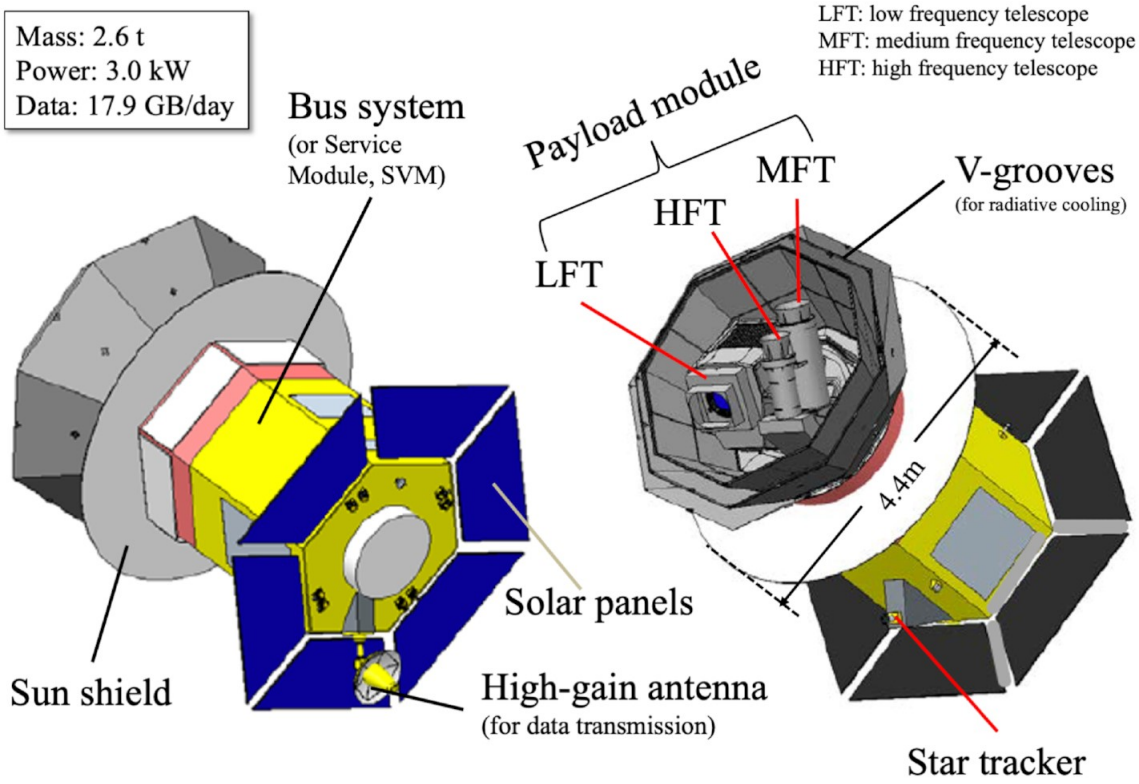

Figure 5.1: From [153]. Conceptual design of the LiteBIRD spacecraft. The Payload Module (PLM) houses the Low-Frequency Telescope (LFT), the Mid-Frequency Telescope (MFT), and the High-Frequency Telescope (HFT)

Several important subsystems are included in the *Service Module* (SVM). Momentum wheels and reaction control systems allow the control of the attitude and orbit of the spacecraft. This is required because the satellite's spinning and its orbit around L2 are slightly unstable. To determine the spacecraft's attitude at each time, the attitude orbit-control system uses the star tracker and inertial reference units. A thermal control system keeps the temperature of the onboard components in the required range thanks to the radiators placed on the upper parts of

the side panels of the spacecraft. The SVM will communicate with the Japanese station GREAT using the X-band, thanks to the high-gain antenna that will transfer 17.9 GB of scientific data daily at a 10 MBps downlink rate. The data handling and the power systems are also part of the SVM. The solar array panels will collect the power.

5.2.2. Payload Module and telescopes

The LiteBIRD payload module includes the global cooling chain from 300 K to 4.8 K and room-temperature elements, such as drivers and warm readout electronics of the detectors. Moreover, the PLM will contain the telescopes and their associated rotating half-wave plate.

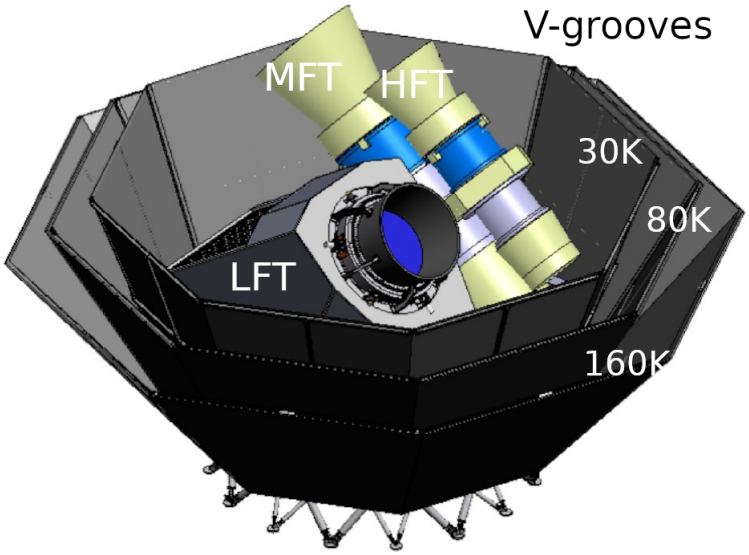

Figure 5.2: From [153]. Design of the LiteBIRD Payload Module (PLM), which contains the Low-Frequency Telescope (LFT), the Mid-Frequency Telescope (MFT), and the High-Frequency Telescope (HFT).

To achieve the science requirements, the choice was made to adopt as the first optical element of the telescopes a *Half-Wave Plate* (HWP) continuously rotating at a frequency $f_{\rm HWP}$ [162–165]. This allows one to distinguish between the instrumental polarised signal and the sky signal, which is modulated at $4f_{\rm HWP}$. The presence of the continuously rotating HWP effectively suppresses the polarised $1/f$ noise, which would otherwise be the dominant noise source at low multipoles (see Sec. III for more details). The *Polarisation Modulator Units* (PMUs) will continuously rotate at a few Hz around a stable temperature below 20 K, using a magnetic levitating mechanism with a superconducting bearing.

The current design involves three telescopes; a reflective one at low frequency, the *LFT* (34–161 GHz) [166], and two refractive ones at medium and high frequencies [167, 168], the *MFT* (89–225 GHz) and *HFT* (166–448 GHz). Notice the *frequency band* overlap, allowing LiteBIRD to control the systematics better. The MFT and HFT telescopes share the same mechanical structure and point in the opposite direction compared to the LFT but cover the same circle over the sky when spinning. The focal planes of the three telescopes have a large field of view. They are populated with *multichroic polarised transition-edge sensor* (TES) detectors (one to three bands per pixel). LiteBIRD will observe with 1030 multichroic pixels, containing a

total of 4508 TES's (1080 in LFT, 2074 in MFT and 1354 in HFT, distribution further detailed in Table 6.2), the signal of which is transmitted through a readout electronic system. Two detector technologies are used: lenslet-coupled detectors for the LFT and MFT and horn-coupled detectors for the HFT.

While the detectors are at 0.1 K, the telescopes themselves are entirely cooled down to 4.8 K, to ensure the stability of the noise and of the pointings. To achieve a cooling to 30 K, the telescopes are surrounded by a *sun shield* and *V-grooves* that prevent the heat from sun radiation. Then, the 4.8 K are achieved thanks to the mechanical cryo-cooler. Finally, a Joule-Thompson cooler decreases the temperature to 1.75 K, followed by four successive adiabatic demagnetisation refrigerators: two in series, bringing the temperature to 0.35 K and two in parallel, allowing the maintenance of a stable temperature of 0.1 K at the focal plane level.

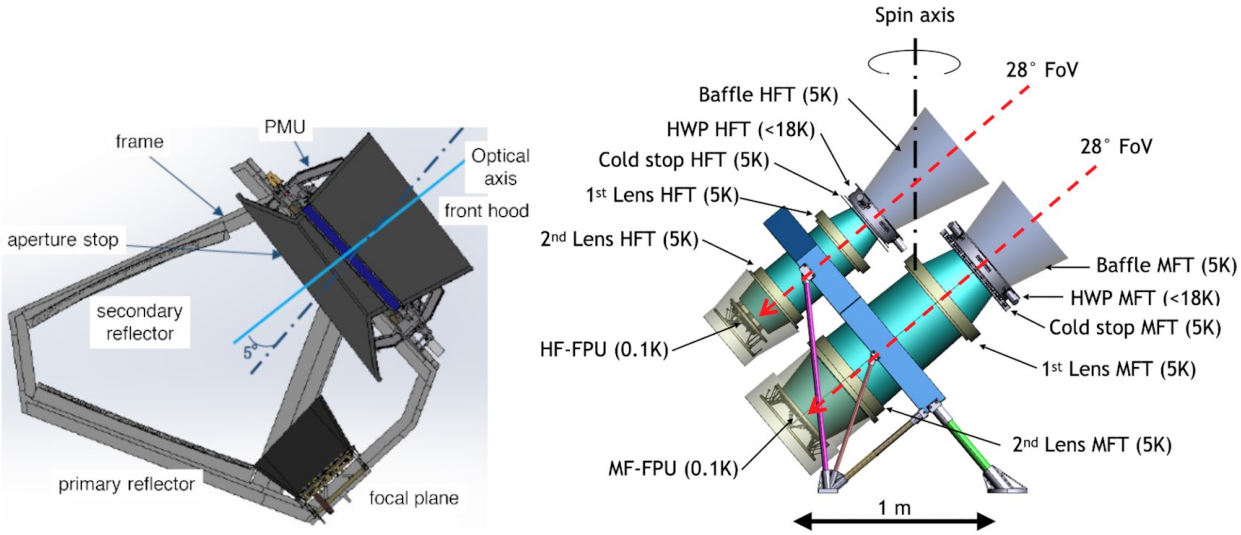

Figure 5.3: From [153]. Left: LFT structure. Right: MHFT structure. More commented in the text.

LFT is a 200 kg telescope consisting of nine broad frequency bands with angular resolution ranging from 24 to 71 arcminutes, spanning from 34 to 161 GHz, covering the spectral domains of CMB and synchrotron radiation emission. Its optical design follows a crossed-Dragone configuration, with two mirrors and an aluminum antenna. The PMU is mounted in front of an aperture stop with a diameter of 400 mm and made in a millimetre absorber on an aluminium plate. This PMU consists of an achromatic half-wave plate made of a multi-layer sapphire stack. The HWP, two mirrors and additional absorbers are supported by a 4.8 K frame structure alongside the focal plane. The field of view is wide (18° × 9°), so each pixel's optical properties vary a lot. The ray diagram of LFT is displayed on the left-hand side of Fig. 5.4. The photons enter along the optical axis from the upper-right aperture stop. The two mirrors successively reflect them before they are collected by the lenslets and hit the detectors placed on the focal plane.

MFT and HFT share a common mount and are very similar in terms of design. Hereafter, we shortcut "MFT and HFT" by "MHFT" when discussing the ensemble. These telescopes consist of five bands each, with angular resolution ranging from 28 to 38 arcminutes for MFT and 18 to 29 arcminutes for HFT, spanning from 77 to 254 GHz for MFT and 136 to 495 GHz for HFT. This range covers the spectral domains where mostly CMB and Galactic dust

radiation emission dominate. They are fully refractive telescopes, allowing compactness and lightness while observing a wide frequency range. Each telescope comes with a front baffle, PMU and aperture stop. The latter has a diameter of 300 mm for MFT and 200 mm for HFT (see Fig. 5.4). The PMUs consist of mesh transmissive HWPs. This technology emulates the behaviour of birefringent materials while being lighter than the LFT sapphire HWP and adapted to the MHFT bandwidth. Unlike the reflective LFT design, the telescopes use two polypropylene lenses to focalise light rays. As for LFT, absorbers will also prevent or mitigate spurious in-band reflections across the optics tubes. Filters are also included in the design to ensure out-of-band rejection of the radiation.

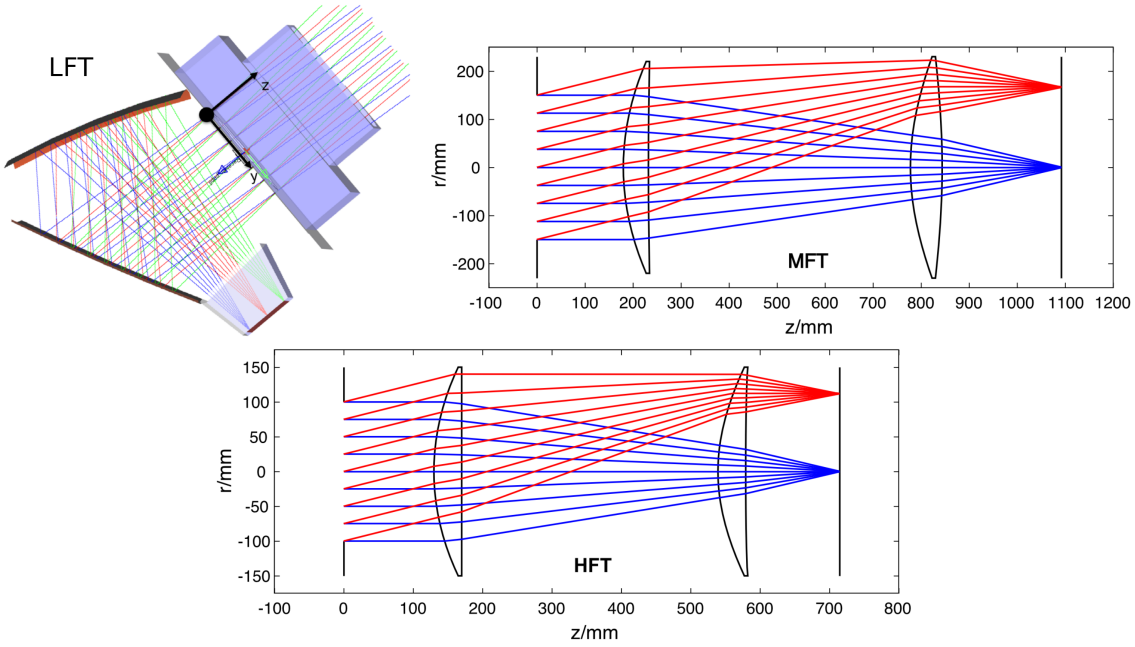

Figure 5.4: From [153]. Ray tracing diagram obtained from optical simulations. Upper-left: LFT, Upper-right: MFT, Bottom: HFT. For the three telescopes, the on-axis incident light hit the center of the focal plane, while the off-axis incident light hit the edge of the focal plane. The telescope apertures are located at $z = 0$.

5.2.3. Detection chain

In this section, we introduce the detection chain of the three telescopes. Each includes a *focal-plane unit* (FPU). A structure supports the refrigerators handling the 0.1 K cooling and provides the interface to the hotter stages. The second constituent is a module filled with *lenslet-* or *horn-coupled* TES bolometers arrays. These are fabricated on silicon *wafers* coupled to a multiplexed readout system. Due to the difference in design between LFT and MHFT, the LFT FPU is rectangular, while MHFT ones are hexagonal arrays. For illustration, the focal plane units of the various instruments are represented in Fig. 5.5.

A *pixel* is a collection of bolometers positioned at the same location on the focal plane. Depending on the LiteBIRD wafer being considered, pixels can be monochroic, dichroic, or trichroic, containing between 2 and 6 bolometers accordingly. Each pixel contains two orthogonal bolometer antennas for every observed frequency, enabling independent polarisation measurements. The

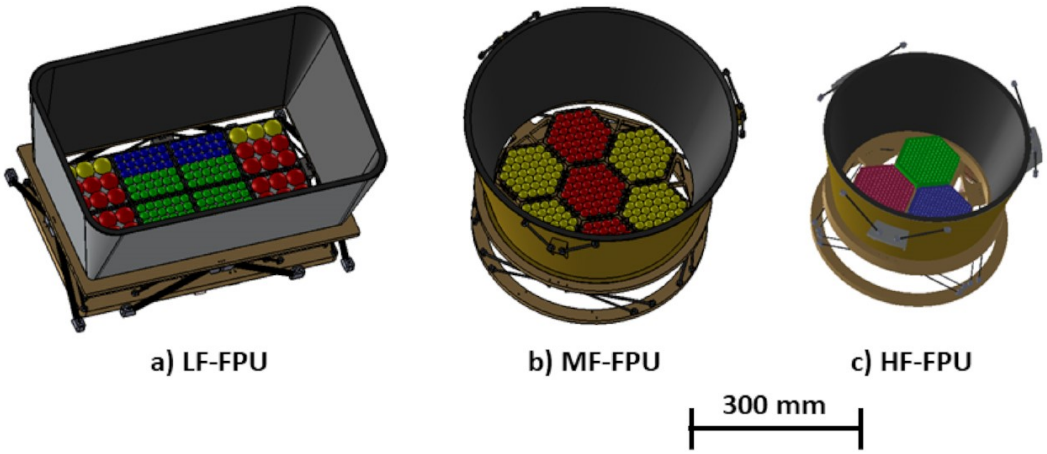

Figure 5.5: From [153]. Focal plane units of the three LiteBIRD telescopes. The colours of the detectors represented here will be discussed when we will detail further the LiteBIRD instrument model in Chapter 6.

diameter of a pixel is directly linked to the size of the antenna and, thus, to the observed frequency. LFT pixels are larger than MFT pixels, which are, in turn, larger than HFT pixels. Details regarding the number of pixels in each frequency channel, their diameter, bolometer content, *etc.*, will be provided in Secs. 6.3 and Sec. 6.4 (and most of this information will be summarised in Table 6.2 and Fig. 6.10).

Transition edge sensors (TES) are widely used bolometers in the context of CMB experiments. They operate with a thin film of superconducting material (picture in the top-left edge of Fig. 5.6), cooled close to its critical temperature, at which the resistance drops steeply to zero. This transition is illustrated in Fig. 5.6, representing the resistance versus the temperature around critical temperature. When electromagnetic radiation is absorbed, it causes a slight increase in temperature ΔT, pushing the TES above its critical temperature and inducing a measurable change of resistance ΔR.

The latter is typically read out using a *Superconducting Quantum Interference Device* (SQUID) current amplifier. It is a very sensitive magnetometer based on superconducting loops containing Josephson junctions. It can detect the faint magnetic field generated by an inductance connected to the TES, as illustrated on the right-hand side electric circuit in Fig. 5.6. This way, the SQUID amplifies and processes the signal from the TES, providing a highly sensitive measurement of the absorbed energy.

Instead of each bolometer being read out individually, the LiteBIRD readout electronics design uses *digital frequency-domain multiplexing*: signals from multiple detectors are simultaneously distributed over a shared data path, reducing the number of required connections or channels and enabling more efficient data acquisition. A *DfMux system* is adopted: each TES is placed in series with an *inductor-capacitor* bandpass filter, which induces a bias tone on top of the signal that allows the detectors to be operated independently. Hence, the *cold-readout* circuit is more complex than the one of Fig. 5.6 while being entirely contained in the 0.1-K stage. A warm electronics stage includes signal-processing, digitising and controlling the SQUID.

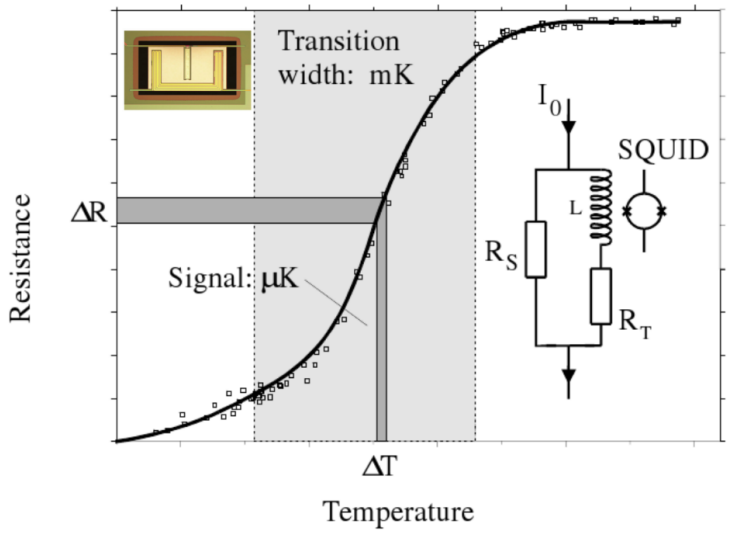

Figure 5.6: Adapted from [169] and [153]. Transition edge sensor.

5.3. Data-analysis processes and analyses

LiteBIRD will conduct continuous scans of the celestial sphere with 4508 detectors measuring both temperature and linear polarisation while keeping track of the telescope's orientation at each time. These time-ordered data (TOD) will be heavy objects comprising CMB, foregrounds, noise, and systematic effects.

The CMB data analysis strategy consists of reducing the dimensionality of the dataset. Typically, to analyse the data, one first projects pre-processed TODs ($>10^{10}$ parameters) onto frequency maps ($\sim 10^7$ parameters) using the pointing knowledge at each time and mitigating TOD systematics. Then, these are combined to construct a CMB map ($\sim 10^6$ parameters). Finally, we compare the statistics of this map ($\sim 10^3$ parameters) to the cosmological predictions from theoretical models to estimate their (~ 10) parameters by maximising the likelihood function of these parameters.

The mitigation of various systematic effects, from foregrounds or the instrument, can be performed at various stages of the analysis: in the TOD domain, map space, spectra domain, or by including them in the likelihood. In any case, this mitigation requires meticulous knowledge of the instrument and realistic modelling of the systematics.

This strategy is more precisely described in Fig. 5.7. From calibration, instrument knowledge, and raw TODs, one can build a mission model and the associated Instrument Model (IMo). Combined with a sky model, this object allows the production of data simulations that are useful for tuning the subsequent data analysis chain at its various steps. A first cleaning is performed in the time domain to eliminate the time-dependent systematic effects before projecting the maps, thanks to the detector pointing knowledge at each time. The component separation is usually done at the map level, where the CMB signal is recovered, as well as the foreground components by fitting models of their frequency and angular scale dependency. One obtains cleaned CMB maps that can be used to estimate the CMB power spectrum, which contains most of the cosmological information (in fact, all, in the case of Λ-CDM). Lastly, one can estimate the cosmological parameters from the obtained spectra.

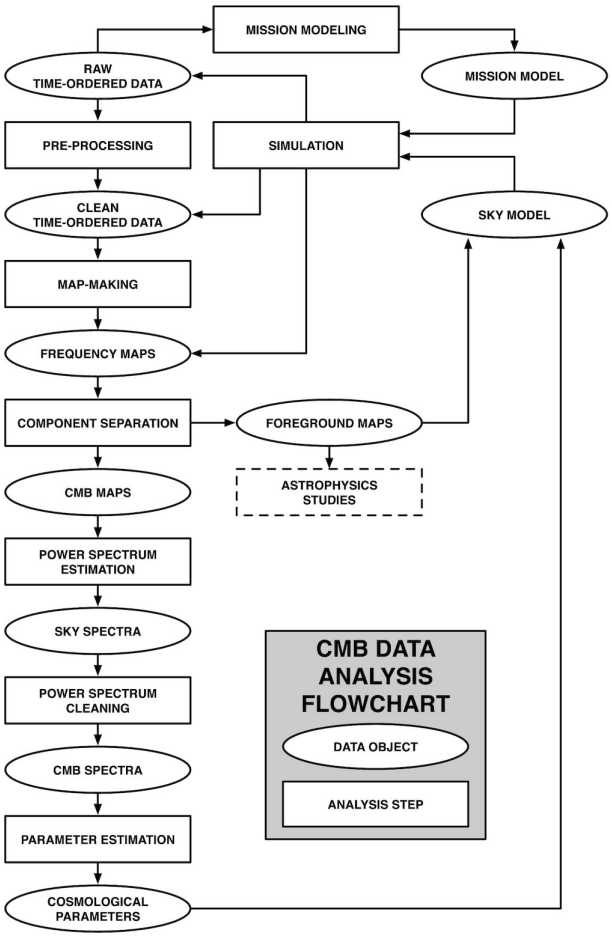

Figure 5.7: From [153]. LiteBIRD data-analysis flowchart summary.

The object of Part III will be to dig into the details of a data-analysis strategy for LiteBIRD E and B modes, especially component separation, power spectrum estimation, and likelihood approximation for optimally estimating the (minimum-variance and minimum-bias) cosmological parameters. This data analysis requires precise instrument knowledge to produce realistic simulations and mitigate the instrumental systematics through the analysis.

CHAPTER 6

LiteBIRD Instrument Model

In this chapter, we introduce the *LiteBIRD Instrument Model* (IMo), for which I have collaborated. The IMo is a quantitative description of the entire LiteBIRD instrument and is used to store and track changes in the instrumental parameters and their knowledge. It is used as input for forecasts, simulations and data-analysis pipelines. This object is stored in a database accessible to the collaboration. To begin with, we introduce the IMo as defined by the collaboration and emphasise the role of the IMo team. Then, we go through most of the quantities in the IMo, their identifiers and localisation in the database, their description and, if public and generic enough, their value in the current design.

Contents

6.1 IMo definition and organisation 102
6.1.1 Definition and specification 102
6.1.2 Software infrastructure: database 102
6.1.3 Software infrastructure: data-access layer and web interface . . 104
6.1.4 IMo team organisation and personal contributions to the IMo . 104
6.2 Low-level instrumental quantities 105
6.2.1 Observation and payload characterisation 105
6.2.2 Telescope characterisation 108
6.2.3 Bolometers characterisation 111
6.3 High-level derived quantities 112
6.3.1 Bandpasses 112
6.3.2 Sensitivities 112
6.3.3 White and $1/f$ noises 118
6.4 Detectors characteristics and mapping 120

6.4.1 Beams 120
6.4.2 Detectors information 121
6.4.3 Detector pointings and quaternions 123

6.1. IMo definition and organisation

The IMo team is in charge of building the IMo to fulfil the definition and specification fixed by the collaboration [170].

6.1.1. Definition and specification

The key specifications for the IMo are:

- The IMo should describe the entire LiteBIRD experiment that will ultimately include measured instrumental parameters with their error bars. Today, the values stored in the IMo come from requirements, outputs of prediction tools and real measurements.
- There are two kinds of parameters. On the first hand, the low-level ones (*e.g.*: lenses dimensions, properties of various optical elements...). They can be used to feed dedicated instrumental simulations and modelling. On the other hand, the high-level ones (*e.g.*: NEP predictions, beams simulations...). They can be used to feed the LiteBIRD simulations and data analysis codes.
- The IMo should be documented: what are the quantities, where they come from, the associated assumptions, how to use them...
- The IMo is a collaborative tool that should be used to optimise the design and check the impact on the LiteBIRD physics outputs. All the information shared within the collaboration about instrumental parameters has to appear in the IMo. The IMo should be accessible to everybody in the collaboration.

The IMo team is in charge of developing a tool satisfying this definition. In practice, the IMo team regularly releases new versions of the IMo to get iteratively closer to the actual instrument by including novel models and quantities. When releasing a new version, a first process of verification is performed. Ultimately, it should be followed by a validation step, during which the LiteBIRD Joint Study Groups and simulation team compare the new version with the older one, perform cross-checks to ensure the consistency and validity of the new version and provide feedback to the IMo team. When this process is done, the IMo is finally qualified, announced and widely released to the collaboration. The current IMo version, called version 2.1, was finalised in October 2023 and has been officially qualified and released in March 2024. It is the focus of the end of this chapter. This release also came with a software infrastructure that we detail below.

The software to organise the IMo data consists of several codes that interact closely with each other. They serve to fill or to build the *database* or to make the data accessible through a *data access layer* and a *web interface*.

6.1.2. Software infrastructure: database

A *database* (DB), in which objects are stored and organised in a tree architecture detailed later, is populated thanks to `IMo_LiteBIRD` [171], where the interface with the data files is done.

These come with various extensions and can include technical value tables, bandpasses files, beam files, *etc*, which we introduce later on in this section. `IMo_LiteBIRD` builds the DB thanks to a library called `instrument_db` [172], the development of which is mainly driven to answer LiteBIRD requirements.

The architecture tree, which includes as main levels (we list a few examples of associated quantities for each of them):

- `Payload Module` (*e.g.:* scanning strategy, data rate...)
- `Instrument` (*e.g.:* HWP, lens, mirror properties, mechanical dimensions...)
- `Frequency channel` (*e.g.:* polarisation sensitivity, design bandwidth, design beam FWHM...)
- `Detectors` (*e.g.:* angle of polarisation, diameter, pointing, beam accurate simulation...)

We provide one instance of IMo tree architecture implemented in IMo v2.1 for one particular level: the `frequency-channel` level is displayed in Fig. 6.1. The same figures for the `observation`, `MHFT`, and `detector` levels are provided in App. B. We detail all the various quantities appearing in these Figures in Secs. 6.2, 6.3 and 6.4.

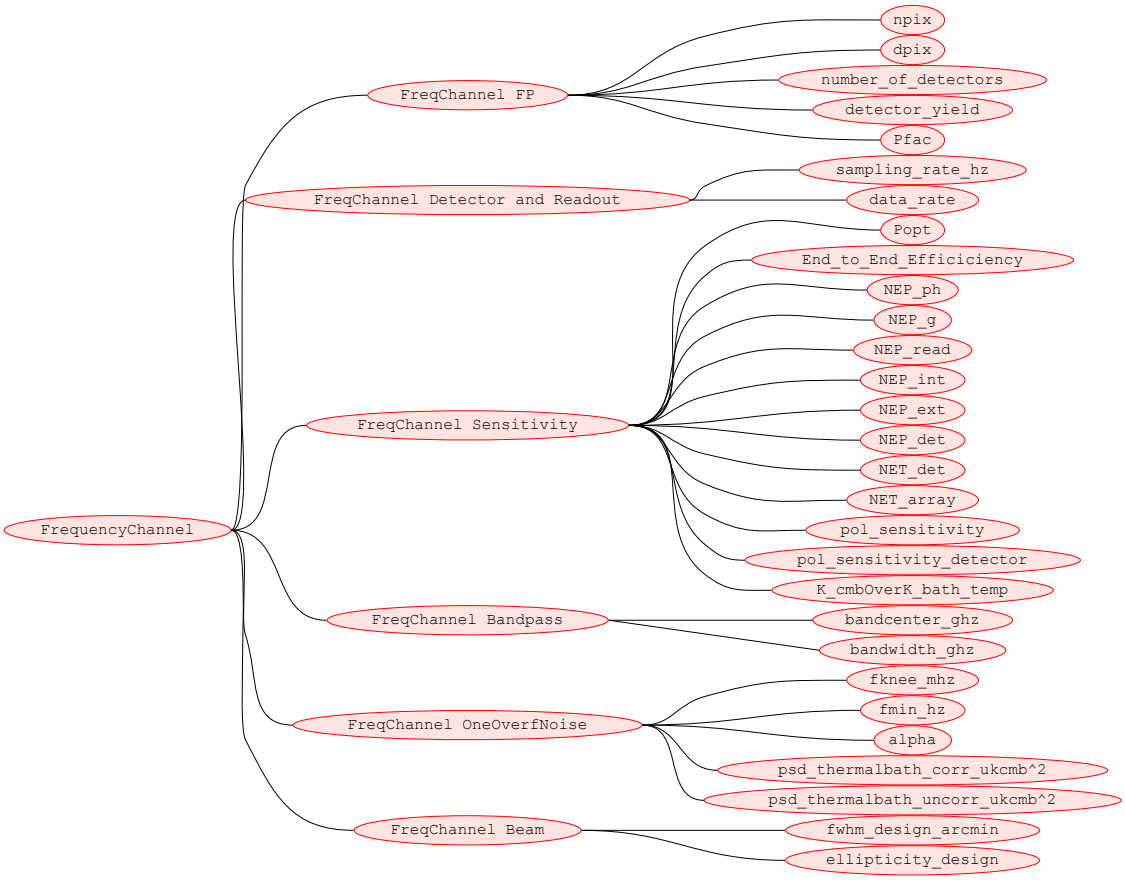

Figure 6.1: IMo tree architecture for frequency-channel level (there are 22 instances of this tree architecture, one per frequency channel).

6.1.3. Software infrastructure: data-access layer and web interface

The *data-access layer* (DAL) is used for accessing, writing, adding, and ensuring the consistency of the data stored in the DB. Two tools can be used as DAL: a parent library of `instrument_db` called `libinsdb` [173], or a Python class directly included inside `IMo_LiteBIRD`. The former calls for the latter, which includes access to lower-level functionalities. Simulation, performance, or data-analysis codes can be plugged into either of these tools. For example, a LiteBIRD simulation tool, `litebird_sim` [174], uses `libinsdb`, while a LiteBIRD performance code, `IMo_Perf` [175] is directly built upon `IMo_LiteBIRD`.

Additionally, a web interface enables any LiteBIRD collaborator to interact with and visualise the DB and download data. This interface is hosted by the *Space Science Data Center* (SSDC), a website of the *Italian Space Agency* (ASI). The various IMo software components introduced thus far and their interconnections are illustrated in Fig. 6.2.

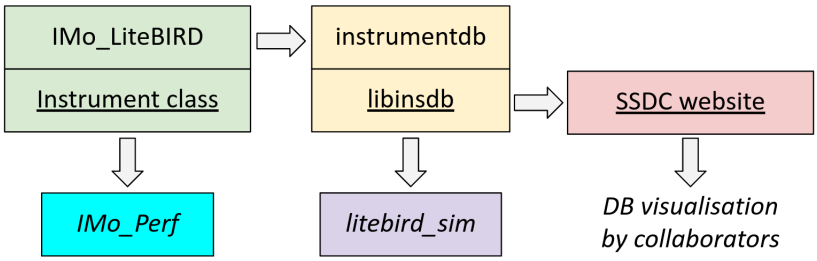

Figure 6.2: Pieces of the IMo software and their interconnections. `IMo_LiteBIRD` is the central piece that manages the implementation of the quantities. The DAL role is fulfilled by its Instrument class and by `libinsdb`, a module which allows accessing the database built by the `instrumentdb` software. To visualise or download part of the database, one can instead pass through the SSDC website.

6.1.4. IMo team organisation and personal contributions to the IMo

The IMo-team work consists of four main tasks that work closely together, each focusing on a different task required to answer the specification and handle a specific part of the software:

1) Data and needs collection for future IMo versions.
2) Development and exploitation of the software for DB filling & IMo access.
3) Scientific IMo data validation.
4) Development and deployment of the software for DB generic tools and web interface

During my PhD, I served as the coordinator and primary investigator of task 2. The requirements for new data in the IMo are collected by the coordinators of task 1, who request them from the instrumental experts. I am responsible for developing `LiteBIRD_IMo`, integrating these new IMo inputs into the tree architecture, ensuring the internal consistency of the model, and subsequently improving and updating the data-access layer. Task 3 is responsible for validating the new version, which includes conducting cross-checks with previous versions. Finally, 4 is tasked with developing `instrument_db` and `libinsdb` and deploying the database on the SSDC website.

On top of this task, I am in charge of deriving some high-level parameters that will be detailed later on. I have been involved in the release of three IMo versions:

- IMo v1.3 (development: June 2022 - release: January 2023).
  Personal contribution: First implementation by the IMo team of all the parameters inside `instrument_db`. Production of quaternions and Euler angles from updated pointings (coming from optical simulations) and new polarisation pattern. Definition of a new naming scheme for the detectors.

- IMo v2.0 (development: June 2023 - release: June 2023).
  Personal contribution: Implementation of the squid mapping and of the correlated-1/f-noise model.

- IMo v2.1 (development: October 2023 - release: March 2024).
  Personal contribution: Participation in redefining the IMo tree architecture. Input of many new PLM, instruments, frequency channels and detectors low-level quantities used in sensitivity computation. Implementation and renaming of the files containing realistic beams and bandpasses.

I have also served as a point of contact between the IMo and the simulation teams, as I am also a member of the latter. In this role, I provide information on the IMo to the simulation team and feedback from them to the IMo team. The LiteBIRD simulations are described at the beginning of Part III.

Now that we have defined the broad definition and goals of the IMo, as well as the organisation of the IMo team, the following sections will provide a more detailed description of the low-level and high-level quantities.

6.2. Low-level instrumental quantities

The low-level quantities are fundamental parameters inherent to the mission's or instrument's design, not derived from other values. They include the adopted scanning strategy, the mechanical design of the payload module, the system and sub-system characteristics of the three telescopes, and the technical specifications of the bolometers.

6.2.1. Observation and payload characterisation

The first class of IMo quantities pertains to the payload module and to LiteBIRD's observation strategy and the associated pointing of the instruments. The latter is crucial for determining the direction of observation for each detector at any given time. The current strategy is outlined in Fig. 6.3, illustrating the spacecraft's orientation relative to the anti-sun axis, as well as defining the boresight axis and spin axis in terms of the IMo quantities and introduced angles below:

- `spin_sun_angle_deg`: Angle between the spin axis and the anti-sun axis (α in Fig. 6.3) [deg].
  Baseline value: 45°
- `spin_boresight_angle_deg`: Angle between the spin axis and the boresight angle (β in Fig. 6.3) [deg].
  Baseline value: 50°.
- `precession_period_min`: Precession Period [minutes].
  Baseline value: 192.348 minutes.
- `spin_rate_rpm`: Payload Spin Rate [rpm].
  Baseline value: 0.05 rotations per minute.

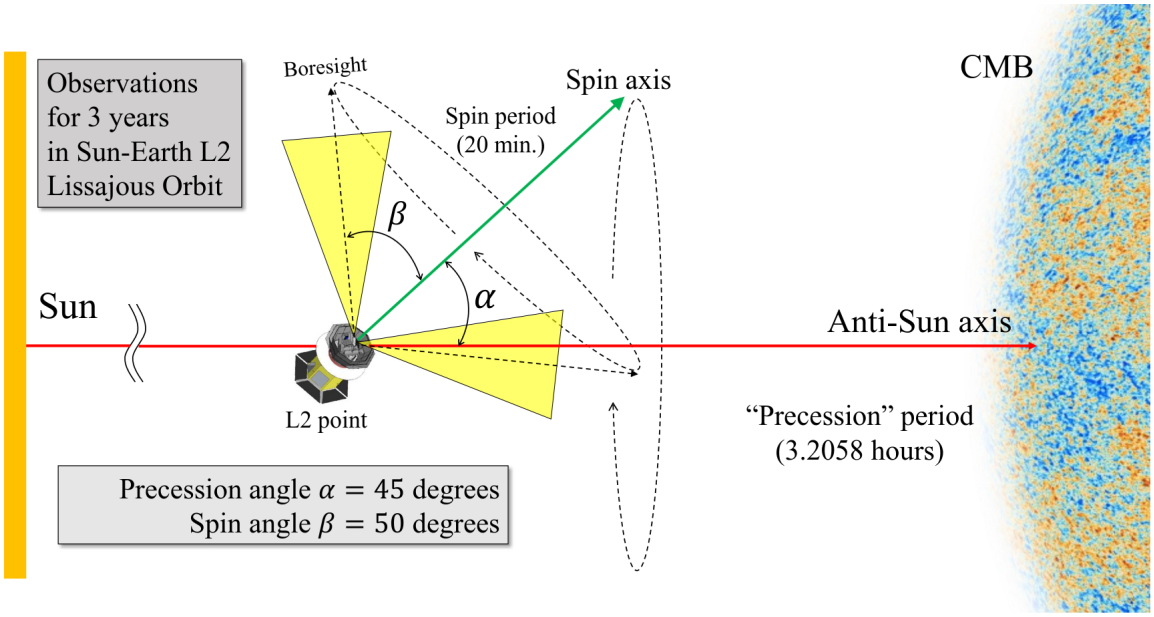

Figure 6.3: From [153]. Overview of the LiteBIRD scanning strategy and definition of the quantities. The boresight of a telescope is its line-of-sight axis.

There are two additional quantities, specified at the telescope levels, that define the three instruments' boresights alongside with `spin_boresight_angle_rad`:

- `boresight_rotangle_rad`, Angle between the focal plane with respect to the boresight direction itself [rad].
  Baseline value: 0° for all instruments.
- `spin_rotangle_rad`: Additional angle by which the boresight is rotated around the spin axis [rad].
  Baseline values: 180° for MHFT, 0° for LFT.

These baseline values have been optimised to achieve hit maps that are as uniform as possible across the entire sky. For example, the baseline value for spin rate is set at the minimum necessary to maintain this uniformity. Further details justifying these baseline choices are provided in [176]. The baseline hit map in Galactic coordinates for the three-year LiteBIRD survey is illustrated

in Fig. 6.4, from [177]. The precession period is optimised to avoid any Moiré pattern, and each pixel is at least scanned 200 times.

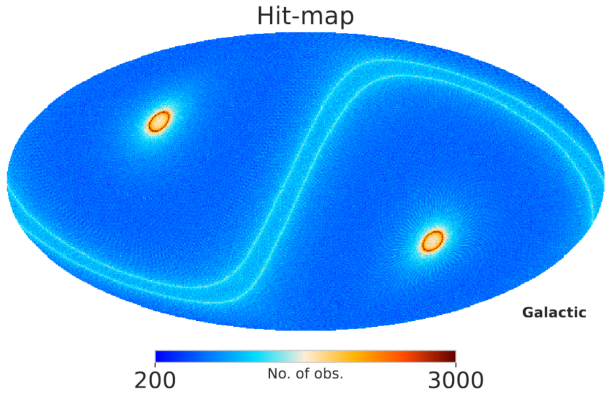

Figure 6.4: From [177]. Hit map in Galactic coordinates of the 3-year LiteBIRD mission in the design scanning strategy. This map is simulated for a single boresight detector, assuming a sampling rate of 19 Hz and $N_{\rm side} = 128$.

A second category of quantities includes general parameters related to the effective observation time.

- `mission_duration_year`: Duration of the Mission $t_{\rm obs}$ [years].
  Baseline value: 3 years.
- `observation_duty_cycle`: Duty Cycle Efficiency, margin due to the download of the data dead time, $\eta_{\rm duty}$ [unitless].
  Baseline value: 0.85.
- `cosmic_ray_loss`: Efficiency Loss due to Cosmic Rays, $\eta_{\rm CR}$ [unitless].
  Baseline value: 0.95.
- `margin`: Margin in efficiency inputs for sensitivity calculations, $\eta_{\rm margin}$ [unitless].
  Baseline value: 0.95.

The margin $\eta_{\rm duty}$ is a simplistic way of including a dead time, by only accounting for it as a sensitivity degradation. One can derive the observation or total time efficiency as $\eta_{\rm tot} = \eta_{\rm duty} * \eta_{\rm CR} * \eta_{\rm margin}$ and an effective observation time $\eta_{\rm tot} * t_{\rm obs}$, which is two years and four months in the baseline design. These quantities are, for instance, important in relating the total sensitivity of a given frequency channel, as we will show later on.

The crucial quantities regarding the detector readout are the sampling rate, chosen to be twice the Nyquist frequency of the HWP modulation, and the data rate:

- `sampling_rate_hz`: Readout Sampling Rate [Hz].
  Baseline value: 19.1 Hz,
- `data_rate`: Data Rate [GB/day].
  Baseline value: 17.9 GB day^{-1}.

6.2.2. Telescope characterisation

This section summarises the various quantities attached to the three LiteBIRD telescopes. These are organised as *system* and *sub-system* quantities depending on whether they are linked to the former or to one of the latter.

6.2.2.1 System temperature stages

First, as we introduced above, there are four temperature stages in LiteBIRD. The design temperatures are linked to the name of the stage and are common to all three instrument systems:

- `T_5K`: 5K stage [K]
- `T_2K`: 2K stage [K]
- `T_20K`: 20K stage [K]
- `T_01K`: 0.1K stage bolometer operation [K]

6.2.2.2 System optical parameters

In the instruments, there can be several optical effects that affect the light travelling, and that need to be modelled.

Spillover refers to the phenomenon where light extends beyond its intended boundaries. It can result in unwanted effects such as ghost images and reflections.

Ghost images are faint, secondary images of an object that appear alongside the primary image. They are typically caused by light reflecting off surfaces within the optical system, such as lenses or mirrors, and reaching the image sensor or observer. Spillover of light can contribute to the formation of these ghost images by allowing stray light to interact with optical components in unintended ways.

Reflections, on the other hand, occur when light bounces off a surface and changes direction. In an optical system, reflections can occur at various interfaces between different materials, such as air and glass. When light spills over beyond its intended path, it can increase the likelihood of reflections by encountering additional surfaces within the system. These reflections can then contribute to unwanted glare or artefacts in the final image.

In the IMo, the optical parameters specifying spillover, reflection and the derived optical efficiency, which are computed from the optical model, are stored at the level of the instrument system. For LFT:

- `det_eff`: Detector efficiency including lenslets and horns η_{det} [unitless] per frequency band.
- `spill_20K`: Spillover at 20K HWP mount [unitless] per frequency band
- `spill_5Kstop`: Spillover at the 5K stop [unitless]
- `spill_5Kenve`: Spillover at the 5K envelop [unitless]
- `spill_2Khood`: Spillover at the Hood [unitless]

While for MHFT:

- `det_eff`: Detector efficiency including lenslets and horns [unitless] noted $\eta_{\rm det}$.
- `spill_apt`: Aperture spillover [unitless] per frequency band.
- `spill_5K`: 5K spillover [unitless] per frequency band.
- `spill_2K`: 2K spillover [unitless] per frequency band.
- `sky_eff`: Sky efficiency [unitless] per frequency band.
- `ref1_5K`: 1st order reflection terminated to 5K [unitless] per frequency band.
- `ref1_2K`: 1st order reflection terminated to 2K [unitless] per frequency band.
- `ref1_FP`: 1st order reflection terminated to FP [unitless] per frequency band.
- `ref2_sky`: 2nd order reflection goes to sky [unitless] per frequency band.
- `ref2_5K`: 2nd order reflection terminated to 5K [unitless] per frequency band.
- `ref2_apt`: 2nd order reflection terminated to stop [unitless] per frequency band.

These parameters are derived from the instrument's optical model. In LiteBIRD, optics experts use GRASP, a software to analyse light path in an optical system [178]. The GRASP input files and the HWP and lens coefficient datasheet are also stored in the IMo to keep track of the model used to compute the optical parameters, HWP and lens parameters:

- `optical_design`: optical model file.
- `lenslet_or_feed_model`: files used for the feedhorns or lenslet model.
- `front_baffle_model`: Front Baffle file.
- `comment`: additionnal comment for the overall description.
- `HWP_data_sheet`: HWP coefficient datasheet calculations.
- `Lens_data_sheet`: lens coefficient datasheet calculations.

The optical design file itself has been built from the technical specifications of the system and the sub-systems.

6.2.2.3 System mechanical and thermic parameters

The mechanical parameters including the geometry related to the different sub-systems are put at the System level of MFT and HFT:

- `H_FP_L2`: Distance between the focal plane and L2 [m].
- `H_HWP_L1`: Distance between the HWP and L1 [m].
- `H_FP_HWP`: Distance between the focal plane and the HWP [m].
- `H_L1_L2`: Distance between L1 and L2 [m].
- `H_baf`: Baffle height [m].
- `FoV_diameter`: Diameter of the field of view [deg].

Concerning MHFT, the two quantities linked to the front baffle are its temperature and aperture angle:

- `Temp_FrontBaffle`: FrontBaffle Temperature [K].
- `theta_baf`: Baffle half angle [deg].

Lastly, the MHFT aperture parameters are:

- `Diameter_Aperture`: Diameter of the HFP aperture [m].
- `Temp_Aperture`: Aperture Temperature [K].
- `beam_waist_factor`: Beam waist factor [unitless].
- `Fnumber`: F number [unitless].

To derive the optical parameters, one also needs the sub-system parameters.

6.2.2.4 Sub-system parameters

In the IMo, we documented three LFT sub-systems: its half-wave-plate, its filter and its mirror, and four MHFT ones: its half-wave-plate, its filter, its two lenses and its focal plane.

The LFT and MHFT half-wave plate parameters are stored as:

- `T_HWP`: HWP temperature [K].
- `D_HWP`: HWP diameter [m].
- `emiss_hwp`: Emissivity [unitless] per frequency band
- `ref_hwp`: Reflectance [unitless] per frequency band
- `eff_hwp`: Transmission efficiency [unitless] per frequency band (LFT only)
- `pol_hwp`: polarisation efficiency [unitless] per frequency band
- `hwp_rpm`: HWP revolution rate [rpm].
  Baseline value: LFT: 46 rpm, 39 (MFT), 61 (HFT).
- `t_hwp`: Thickness [m] from CDR (LFT only).
- `n_hwp`: Refractive index (LFT only) [unitless].
- `tan_hwp`: Loss tangent (LFT only) [unitless].
- `pol_dilu`: HWP rotation dilution factor (MHFT only) [unitless].

The second current-design LFT sub-system is its 2K filter, which is specified by:

- `Temp_Filter`: Filter Temperature [K].
- `t_fil`: Thickness [m].
- `n_fil`: Refractive index [unitless].
- `tan_fil`: Loss tangent [unitless].
- `ref_fil`: Reflectance [unitless].
- `Fil_oob_reject`: Out-of-band rejection (0 is no filter) [dB].

The MHFT filter is currently assumed to be identical to the LFT one, so MHFT shares these sub-system quantities. The last LFT sub-system is its mirror characterised by:

- `rho`: Resistivity [Ohm.m].
- `rms`: RMS of the surface roughness [m].

In MHFT, there will be two lenses, L1 and L2. The numerical values of their main characteristics, which appear on the IMo, vary according to the considered lens. Their identifiers are

the following with $i \in 1, 2$:

- `D_L`i: Li diameter [m].
- `t_L`i: thickness of the lens [m].
- `Temp_L`i: Li Temperature [K].
- `emiss_L`i: Emissivity of first lens [unitless] per frequency band
- `ref_L`i: Reflectance of first lens [unitless].
- `material`: material of the lens.
- `n`: Refractive index [unitless].

Finally, the focal plane characteristics are its diameter and its temperature:

- `D_FP`: Focal plane diameter [m].
- `Temp_FocalPlane`: Focal Plane Temperature [K].

6.2.3. Bolometers characterisation

This pool of low-level parameters is related to the detector and readout electronic characteristics. It has the same value for every detector in the baseline design, though it is susceptible to vary according to the detector in the future.

- `squid_input_inductor`: SQUID input inductor [Henry].
- `shunt_resistor`: shunt resistor [Ohm].
- `temperature_focal_plane`: focal plane temperature [K].
- `tes_normal_resistance`: TES normal resistance [Ohm].
- `tes_log_sensitivity_alpha`: ES alpha = dlogR/dlogT.
- `tes_leg_thermal_carrier_exponent`: for phonons n = 4 (for electrons n = 2).
- `tes_normal_time_constant`: thermal time-constant in normal state [seconds].
- `optical_loading_power`: optical loading power [Watts].
- `tes_saturation_power`: tes saturation power [Watts].
- `tes_transition_temperature`: tes transition temperature [Kelvin].
- `tes_leg_thermal_conductivity`: tes thermal conductivity [Watts/Kelvin].
- `tes_heat_capacity`: tes heat capacity [Joule/Kelvin].
- `bias_current_amplitude`: bias current [Ampere].
- `ac_frequency`: AC bias driving frequency [Hz].

- `mux_frequency`: LC filter frequency [Hz].
- `mux_lc_inductor`: Mux LC filter inductor [Henry].
- `mux_lc_capacitor`: Mux LC filter capacitor [Farad].
- `bath_temperature`: bath temperature [K].
- `thermal_carrier_n`: thermal_carrier input for NEP_g calculations: assuming phonons, n = 3 [unitless].

6.3. High-level derived quantities

To derive the expected noise and sensitivity levels of each detector and frequency channel, many of the low-level quantities mentioned in the previous subsection are required. Each optical element contributes to the total optical power of the instrument through factors such as reflectance, emissivity, and efficiency. Sensitivity-related quantities, among others, are often referred to as "high-level" quantities because they are directly to be used as inputs in the simulation and data-analysis codes. They often depend on several "lower-level" ones. In this sense, bandpasses and beams, based on models rather than solely design specifications, are also high-level quantities. For certain high-level characteristics of the instrument, the IMo offers multiple parameterisations that can be employed for varying levels of refinement in simulations and analyses.

6.3.1. Bandpasses

There are various ways to parameterise the bandpasses, with the simplest being a top-hat function characterised by an associated band center and bandwidth. Today, these parameters are specified at the frequency-channel level of the IMo.

- `bandcenter_ghz`: Center Frequency of the band [GHz]
- `bandwidth_ghz`: Bandwidth given a TopHat description of the bandpass [GHz]

The top-hat description is depicted by rectangular boxes in the left panel of Fig. 6.5. The values of the centers and widths for each channel are crucial for simulations, forecasts, and data-analysis codes, and they are summarised in Table 6.1.

In order to study more realistic bandpass descriptions, the IMo also provides random bandpasses realisation for each detector, generated from a Chebyshev model. The right panel of Fig. 6.5 illustrates the spectral transmission of these realisations for a few detector examples spanning the LiteBIRD frequency range . In this particular representation, we show the transmission whose maximum is fixed to one for any detector, bandpass function we call hereafter $B(\nu)$. The top-hat bandpass function simply reduces to 1 if $\nu_1 < \nu < \nu_2$, else 0. One can tell from this figure that the Chebyshev bandpasses significantly deviate from the top-hat ones. These bandpasses will soon be useful in producing more realistic simulations.

- `bandpass_chebyshev_ghz`: Bandpass per detector (Chebyshev function); structure: index, frequency [GHz], bandpass - sampling: 64 points

Bandpasses are important to be well-controlled through calibration measurements and filtering of the out-of-band power: a bad knowledge of them leads to a significant systematic error.

6.3.2. Sensitivities

In this section, we introduce the LiteBIRD statistical noise budget and how the resulting sensitivities and levels of white and 1/f noise have been derived and are stored in the IMo.

Some quantities stored at the frequency-channel level and useful for deriving sensitivities are the count and diameter of the detectors in a given frequency channel. These are

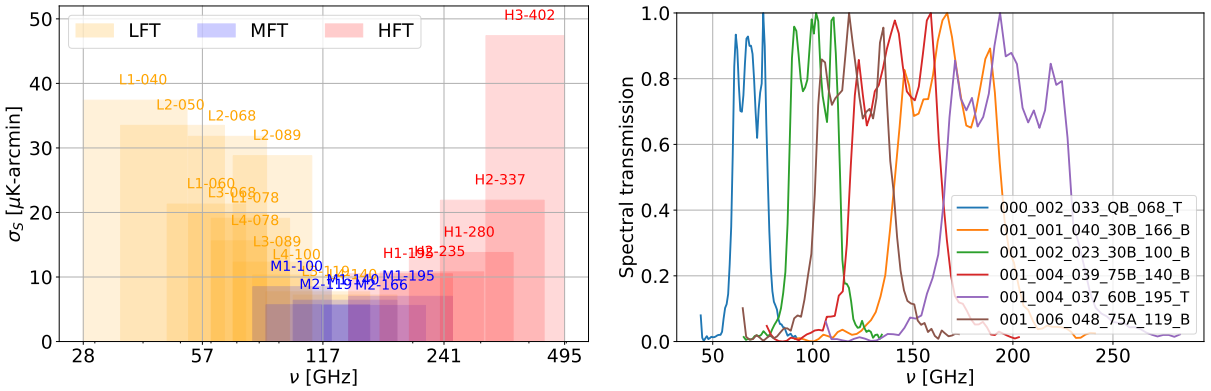

Figure 6.5: LEFT: Polarisation sensitivities per frequency channel as a function of the frequency assuming top-hat bandpasses for the 22 LiteBIRD frequency channels. The numerical values used for the plot can be found in Table 6.1. RIGHT: IMo realisation of Chebyshev bandpasses for a few detectors belonging to various frequency channels. The y-axis represents the spectral unitless spectral transmission whose maximum is fixed to one for each bandpass.

- `npix`: Number of pixels
- `dpix`: Pixel diameter [mm]
- `number_of_detectors`: Number of detectors

These three pieces of information will be represented for the different instruments in Fig. 6.10.

Let us introduce the ingredients required to derive the LiteBIRD sensitivities. The conventions and notations adopted hereafter mostly come from [153, 179] . The total operating power of the bolometer $P_{\rm oper}$ is the sum of its optical power $P_{\rm opt}$ and its electrical power $P_{\rm elec}$, all in Watts. The parameter $P_{\rm fac}$ gives the ratio between the total power and the optical power. In the baseline design, the optical power amounts to 40% of the operating power (*i.e.* $P_{\rm fac} = 2.5$).

- `Popt`: Optical power per detector for a given frequency channel (assuming CMB at 2.7K only) [pW]
- `Pfac`: Factor for bolometer operation power Pfac = Poper (operating power of the bolometer = Popt + Pelec) / Popt [unitless]

Assuming a fiducial sky model (for the baseline, a CMB at 2.7K), $P_{\rm opt}$ is estimated as an integral over the band:

$$P_{\rm opt} = \int_0^\infty {\rm d}\nu \; B(\nu) p_{\rm opt}(\nu) \tag{6.1}$$

where $p_{\rm opt}$ is the optical power per frequency. From now on, we will assume a top-hat bandpass, so

$$P_{\rm opt} = \int_{\nu_1}^{\nu_2} {\rm d}\nu \; p_{\rm opt}(\nu). \tag{6.2}$$

This choice is based on the fact that the sensitivities we will later provide are currently derived assuming a top-hat bandpass. Of course, the substitution $\int_{\nu_1}^{\nu_2} {\rm d}\nu \to \int_0^\infty {\rm d}\nu \; B(\nu)$ allows to account for a general bandpass shape $B(\nu)$. $p_{\rm opt}$ can be expressed as a sum over all optical

Frequency Channel	β_0 [GHz]	$\Delta\beta$ [GHz]	# of dets	$\mathrm{NET}_{\mathrm{det}}$ [$\mu\mathrm{K}/\sqrt{\mathrm{s}}$]	$\mathrm{NET}_{\mathrm{arr}}$ [$\mu\mathrm{K}/\sqrt{\mathrm{s}}$]	$\sigma_{S,\mathrm{det}}$ [$\mu\mathrm{K}\cdot\mathrm{arcmin}$]	σ_S [$\mu\mathrm{K}\cdot\mathrm{arcmin}$]	θ_{FWHM} [arcmin]
L1-040	40	12.0	48	114.6	18.5	259.3	37.4	70.5
L2-050	50	15.0	24	72.5	16.5	163.9	33.5	58.5
L1-060	60	13.8	48	65.3	10.5	147.7	21.3	51.1
L3-068	68	15.6	144	105.6	9.8	238.9	19.9	41.6
L2-068	68	15.6	24	68.8	15.7	155.6	31.8	47.1
L4-078	78	17.9	144	82.5	7.7	186.6	15.6	36.9
L1-078	78	17.9	48	58.6	9.5	132.6	19.1	43.8
L3-089	89	20.5	144	65.2	6.1	147.4	12.3	33.0
L2-089	89	20.5	24	62.3	14.2	141.0	28.8	41.5
L4-100	100	23.0	144	54.9	5.1	124.1	10.3	30.2
L3-119	119	35.7	144	40.8	3.8	92.2	7.7	26.3
L4-140	140	42.0	144	38.4	3.6	86.9	7.2	23.7
M1-100	100	23.0	366	71.7	4.2	162.2	8.5	37.8
M2-119	119	35.7	488	55.7	2.8	125.9	5.7	33.6
M1-140	140	42.0	366	54.0	3.2	122.1	6.4	30.8
M2-166	166	49.8	488	54.4	2.8	123.0	5.6	28.9
M1-195	195	58.5	366	59.6	3.5	134.8	7.0	28.0
H1-195	195	58.5	254	74.0	5.2	167.3	10.5	28.6
H2-235	235	70.5	254	76.1	5.3	172.0	10.8	24.7
H1-280	280	84.0	254	97.3	6.8	220.0	13.8	22.5
H2-337	337	101.1	254	154.6	10.8	349.8	21.9	20.9
H3-402	402	92.5	338	385.7	23.5	872.3	47.4	17.9

Table 6.1: LiteBIRD frequency channels with bandpass center (β_0, ID: `bandcenter_ghz`) and width ($\Delta\beta$, ID: `bandwidth_ghz`), number of detectors (ID: `number_of_detectors`), NET per detector and in the frequency channel ($\mathrm{NET}_{\mathrm{det}}$ and $\mathrm{NET}_{\mathrm{arr}}$, ID's: `NET_det` and `NET_array`), polarisation sensitivity per detector and in the frequency channel ($\sigma_{S,\mathrm{det}}$ and σ_S, ID's: `pol_sensitivity_detector` and `pol_sensitivity`), and beam FWHM size (θ_{FWHM}, ID: `fwhm_design_arcmin`).

elements from the sky to the detector,

$$p_{\mathrm{opt}}(\nu) = \sum_{i=0}^{N_{\mathrm{elem}}} P_i\left(T_i, T_{r;i}, \nu\right), \tag{6.3}$$

where P_i is the optical power of the optical element $i \in [0, N_{\mathrm{elem}}]$. In the case of MHFT, these elements from the sky to the detector are the CMB, the HWP, the aperture stop, the objective lens L1, the field lens L2, the 2K filter, the lenslet or horn and the detector. In the case of LFT: the CMB, the HWP, the aperture stop, the forebaffle, the primary and secondary mirror, the telescope shield, the HWP mount, the 2K filter, the lenslet and the detector. At the considered frequency ν, each component is assumed to radiate as a black body at temperature T_i and reflect part of its power toward an element at the black body temperature $T_{r;i}$. Noting S the Planck spectral density for a diffraction-limited single-moded polarimeter, P_i can be expressed as

$$P_i\left(T_i, T_{r;i}, \nu\right) = \prod_{j=i+1}^{N_{\mathrm{elem}}} \eta_j(\nu) \times \left[\varepsilon_i(\nu) S\left(T_i, \nu\right) + r_i(\nu) S\left(T_{r;i}, \nu\right)\right], \tag{6.4}$$

where the prefactor represents the cumulative transmission efficiency η_j of all components detector-side of element i, ε_i is the dielectric emissivity and r_i the reflectivity of element i. η, ε, and r vary between 0 and 1 according to the element considered. The computation of these factors

(which involve almost all the IMo quantities of the previous sections) are not detailed here but can be found in [179].

The optical power allows the computation of several Noise Equivalent Power (NEP) sources that constitute the total LiteBIRD statistical noise budget, adding in quadrature.

The photon shot noise is due to the random nature of photon arrival times: there is an inherent uncertainty in the number of photons detected within a given time interval. The total photon noise consists of the shot noise added to the wave noise related to the photon bosonic nature. The associated NEP, $\mathrm{NEP}_{\mathrm{ph}}$, can be written as:

$$\mathrm{NEP}_{\mathrm{ph}} = \sqrt{\int_{\nu_1}^{\nu_2} \mathrm{d}\nu \ [2p_{\mathrm{opt}}(\nu)h\nu + 2p_{\mathrm{opt}}(\nu)^2]} \tag{6.5}$$

- `NEP_ph`: Photon noise contribution per detector for a given frequency channel [aW/$\sqrt{\mathrm{Hz}}$]

The thermal carrier noise is caused by the random thermal motion of electrons within the detector material. According to the Johnson-Nyquist theorem, any resistor at a non-zero temperature will produce thermal noise. In the optical chain, it is due to thermal fluctuations between the absorbing element and the bath to which it is connected. $\mathrm{NEP}_{\mathrm{th}}$ reads

$$\mathrm{NEP}_{\mathrm{th}} = \sqrt{4k_B P_{\mathrm{opt}} P_{\mathrm{fac}} T_{\mathrm{bath}} \frac{(n+1)^2}{2n+3} \frac{\left(\frac{T_{\mathrm{crit}}}{T_{\mathrm{bath}}}\right)^{2n+3} - 1}{\left[\left(\frac{T_{\mathrm{crit}}}{T_{\mathrm{bath}}}\right)^{n+1} - 1\right]^2}}, \tag{6.6}$$

where P_{fac} is the already defined `Pfac`, where $n = 3$ (low-level `thermal_carrier_n`) for phonons, the critical temperature $T_{\mathrm{crit}} = 0.171$ K is the transition temperature of the LiteBIRD TES (read Sec. 6.4 and see 5.6) and the bath temperature of the detector is $T_{\mathrm{bath}} = 0.1$ K.

- `NEP_g`: Thermal carrier noise contribution per detector for a given frequency channel [aW/$\sqrt{\mathrm{Hz}}$].

These two NEPs, photon and thermal, are fundamental physical noise. On top of them is a noise related to the readout chain, $\mathrm{NEP}_{\mathrm{read}}$. In the baseline design, it amounts to

$$\mathrm{NEP}_{\mathrm{read}} = 3.5\sqrt{2P_{\mathrm{opt}}}. \tag{6.7}$$

It appears in the IMo as

- `NEP_read`: Readout noise contribution per detector for a given frequency channel [aW/$\sqrt{\mathrm{Hz}}$]

The readout noise NEP^2 should amount for less than 21% of the fundamental NEP. The quadrature sum of the two fundamental noise sources plus the readout noise is called the internal noise,

$$\mathrm{NEP}_{\mathrm{int}} = \sqrt{\mathrm{NEP}_{\mathrm{ph}}^2 + \mathrm{NEP}_{\mathrm{th}}^2 + \mathrm{NEP}_{\mathrm{read}}^2} \tag{6.8}$$

by opposition to the external noise, $\mathrm{NEP}_{\mathrm{ext}}$. The latter can be due to microphonic noise due to vibrations of the focal plane, $\mathrm{NEP}_{\mathrm{vib}}$, thermal fluctuations in bath temperature, $\mathrm{NEP}_{\mathrm{TF}}$, cosmic rays causing bolometer heating $\mathrm{NEP}_{\mathrm{CR}}$, magnetic flux fluctuations across the bolometer, $\mathrm{NEP}_{\mathrm{mag}}$, electromagnetic inference within the readout system, $\mathrm{NEP}_{\mathrm{emi}}$, or environmental unknowns. Its total should not exceed 57% of the internal noise. Conservatively, the design value implemented in the IMo comes from saturating this requirement:

$$\mathrm{NEP}_{\mathrm{ext}} = \sqrt{0.32\ \mathrm{NEP}_{\mathrm{int}}{}^2} \simeq 0.57\ \mathrm{NEP}_{\mathrm{int}}. \tag{6.9}$$

$\mathrm{NEP}_{\mathrm{ext}}$ and $\mathrm{NEP}_{\mathrm{int}}$ appear in the IMo as:

- `NEP_int`: Internal noise contribution per detector for a given frequency channel [aW/$\sqrt{\mathrm{Hz}}$]
- `NEP_ext`: External noise contribution (vibration, thermal fluctuations, CR, magnetic, EMI..) per detector for a given frequency channel [aW/$\sqrt{\mathrm{Hz}}$]

For a given detector, the total NEP is the quadrature sum of all NEP's:

$$\mathrm{NEP}_{\mathrm{det}} = \sqrt{\mathrm{NEP}_{\mathrm{int}}^2 + \mathrm{NEP}_{\mathrm{ext}}^2}, \tag{6.10}$$

and appears in the IMo as:

- `NEP_det`: Total detector noise per detector for a given frequency channel [aW/$\sqrt{\mathrm{Hz}}$] care: the detector yield is not taken into account here

We summarise in Fig. 6.6 the various statistical noise sources we have introduced so far.

The Noise Equivalent Temperature (NET) is defined from the Noise Equivalent Power (NEP) by

$$\mathrm{NET} = \frac{\mathrm{NEP}}{\sqrt{2}(\mathrm{d}P/\mathrm{d}T_{\mathrm{CMB}})}, \tag{6.11}$$

where the factor $\sqrt{2}$ is needed because the NEP is defined in terms of output bandwidth while the NET is defined in terms of integration time, and where $\mathrm{d}P/\mathrm{d}T_{\mathrm{CMB}}$ is the conversion factor from power to CMB thermodynamic temperature,

$$\mathrm{d}P/\mathrm{d}T_{\mathrm{CMB}} = \int_{\nu_1}^{\nu_2} \left\{ \frac{\eta(\nu)}{k_B} \left[\frac{h\nu}{T_{\mathrm{CMB}} \left(e^{h\nu/k_B T_{\mathrm{CMB}}} - 1 \right)} \right]^2 e^{h\nu/k_B T_{\mathrm{CMB}}} \right\} \mathrm{d}\nu, \tag{6.12}$$

where η, for a given instrument, is the overall end-to-end optical efficiency, which is the product of the efficiencies η_j of each optical element j, which we have introduced in Eq. 6.4.:

- `End_to_End_Efficiciency`: Overall (end-to-end) optical efficiency for a frequency channel [unitless]

The detector noise equivalent temperature per detector, $\mathrm{NET}_{\mathrm{det}}$, in $\mu\mathrm{K}\cdot\sqrt{s}$, quantifies the noise in each detector and is key, for instance, to simulate independent timelines for each detector. It appears in the IMo as

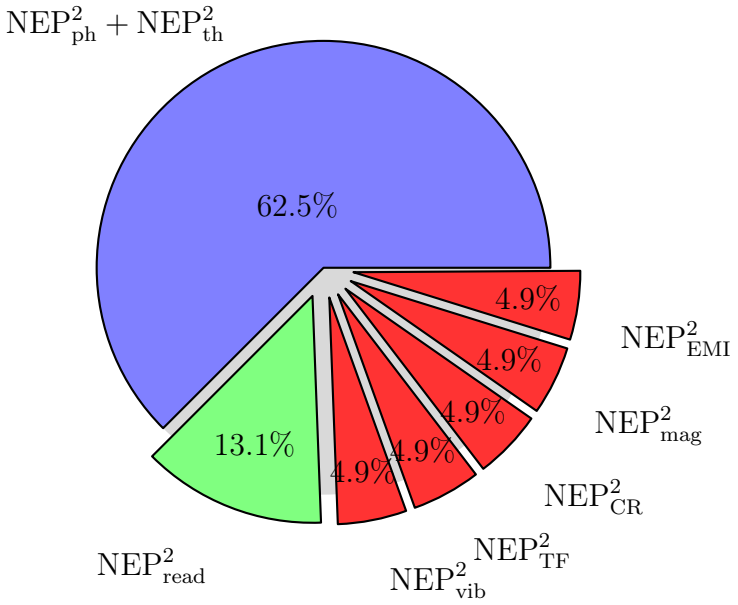

Figure 6.6: Statistical noise budget in terms of NEP^2. The fundamental noise in blue is composed of photon and thermal noise in a proportion that depends on the considered frequency channel. With the readout noise, whose maximum budget is represented in green (this proportion may also vary with the frequency), it represents 75.6% of the total NEP^2 budget. The remaining portion is attributed to various external noise sources, depicted in the pie chart as contributing equally to the budget.

- `NET_det`: Total noise equivalent temperature per detector for a given frequency channel [$\mu K \cdot \sqrt{s}$], care: the detector yield is not taken into account here

Assuming that there is no correlated noise between detectors, the NET per frequency channel NET_{arr} by dividing the detector noise by the square root of the number of detectors operating in the considered channel. Furthermore, we assume in the IMo that a fraction η_{yield} of the total number of pixels by channel is off. This degradation factor appears as

- `detector_yield`: Detector Yield, noted here η_{yield} [unitless]

and NET_{arr} is linked to NET_{det} through

$$NET_{arr} = \frac{NET_{det}}{\sqrt{N_{det}\ \eta_{yield}}}, \tag{6.13}$$

and appears in the IMo as

- `NET_array`: Total noise equivalent temperature for the frequency channel, it considers the detector yield [$\mu K \cdot \sqrt{s}$]

Another key parameter that allows directly generating one map-based noise simulation per frequency channel is the total sensitivity to polarisation of a given frequency channel in $\mu K \cdot arcmin$.

It can be obtained from $\mathrm{NET}_{\mathrm{det}}$ including a similar scaling as the Eq. 4.14 one:

$$\sigma_{\mathrm{S}}\ [\mu\mathrm{K}\cdot\mathrm{arcmin}] = \sqrt{\frac{4\pi f_{\mathrm{sky}}\times 2}{\eta_{\mathrm{tot}} t_{\mathrm{obs}}\times \eta_{\mathrm{yield}} N_{\mathrm{det}}}}\left(\frac{10800}{\pi}\right)\times \mathrm{NET}_{\mathrm{det}}\ [\mu\mathrm{K}\cdot\sqrt{\mathrm{s}}], \tag{6.14}$$

where η_{tot} is the total time efficiency introduced in Sec. 6.2.1, η_{yield} is the detector yield, N_{det} is the number of detectors in the channel considered, f_{sky} is the expected sky-fraction, and an additional factor accounts for the radian to arcmin conversion. The factor 2 in the numerator is due to the need for two detectors to measure CMB polarisation. It appears in the IMo, alongside the sensitivity per detector, which is the frequency-channel sensitivity multiplied by $\sqrt{N_{\mathrm{det}}}$ (hence both include the detector yield).

- `pol_sensitivity`: Sensitivity per frequency channel, it accounts for the detector yield and margins, BUS duty cycle and CR loss [μK·arcmin], also noted σ_S.
- `pol_sensitivity_detector`: Sensitivity to the polarisation of one detector taking into account yield/margin/BUS duty cycle and CR loss [μK·arcmin], also noted $\sigma_{S,\mathrm{det}}$.

Given the importance of these numbers for the subsequent simulation, forecast and data-analysis codes, we summarise for each of the channels their baseline design values in the last four columns of Table 6.1 and represent on the y-axis of the left panel of Fig. 6.5 the frequency-channel polarisation sensitivity. Assuming no foreground degradation and no correlated noise between frequency channels, the total sensitivity of LiteBIRD (all frequencies combined) can be computed from Eq. 4.16, yielding 2.16 $\mu\mathrm{K}\cdot\mathrm{arcmin}$ for LiteBIRD. We will see in Sec. 7.4.4 that, in the presence of foregrounds, this value is significantly degraded by the data analysis.

The coupling coefficients `K_cmbOverK_bath_temp`, which are the values in W/K that convert the thermal fluctuation of the 100 mK stage into a power fluctuation on the detectors, which we denote hereafter $C_{100\mathrm{mK}}$ are specified:

- `K_cmbOverK_bath_temp`: $C_{100\mathrm{mK}}$, the coupling coefficients in K_cmb/K_bath_temp as described in LB_thermal_fluctuations_note.pdf v1

6.3.3. White and $1/f$ noises

The noise levels we have introduced so far can directly be used to model the **white noise** amplitude in the detectors. On top of it, many of the noise sources discussed so far present low-frequency variations and generate **1/f noise**. Sub-systems, detectors and electronics can source this type of noise, as well as thermal fluctuations, cosmic ray interferences, vibrations...

A simple $1/f$-noise model requires introducing three additional parameters: the total knee frequency, f_{knee}, which is the frequency at which the white noise and the $1/f$ noise are at the same noise level (NET), the minimum frequency above which there is time-correlated noise, f_{min}, and α the total spectral index of the spectrum. Then, the power spectral density of the $1/f$ noise can be written as

$$\mathrm{PSD} = \mathrm{NET}^2\left(\frac{f_{\mathrm{knee}}}{f}\right)^{\alpha}\ \text{if } f > f_{\mathrm{min}},\ 0 \text{ else.} \tag{6.15}$$

In the database, these new parameters are identified as:

- `fknee_mhz`: Knee Frequency for 1/f noise description [MHz]
- `fmin_hz`: Min Frequency for 1/f noise description [MHz]
- `alpha`: Alpha for 1/f noise description

However, there is typically a correlation of the $1/f$ noise between detectors. For instance, the readout system partially correlates the noise across all detectors connected to it. Another example is correlated noise due to cosmic rays heating a given wafer that couples thermally the detectors. As a result, more sophisticated models than the parameters introduced so far are required. Dedicated models for the two source examples provided above have been defined in [180] and incorporated into LiteBIRD IMo v2.0.

The gain fluctuations in the digital-to-analog converter (DAC) induce a **readout 1/f noise**, with a power spectral index of $\alpha = 1$ and with a knee frequency $f_{\rm knee} = 40$ mHz for a single detector. Then, for a detector in a channel with a given $\mathrm{NET}_{\rm det}$, the induced per-detector $1/f$ component is given by the following power spectral density (PSD):

$$\mathrm{PSD}_{\rm DAC1/f} = 2\ \mathrm{NET}_{\rm det}^2 \frac{40 \times 10^{-3}}{f}. \tag{6.16}$$

For example taking the $\mathrm{NET}_{\rm det}$ in Table 6.1 for M1-140 leads to a readout $1/f$ noise spectrum of 233 $\mu\mathrm{K}^2 \left(\frac{1}{f}\right)$. Because related to the readout, this component is common for detectors that share a given SQUID.

Cosmic rays continuously heat the detector wafers, leading to fluctuations of the wafer temperature, and hence to a **cosmic-rays 1/f noise** component. In Planck, these fluctuations were consistent with $\alpha = 1$ [181]. The requirement on the focal plane temperature stability is 0.3 $\mu\mathrm{K}/\sqrt{\mathrm{Hz}}$ at 2.4 Hz (no CMB unit). Under the assumption that it is barely achieved, one can derive an amplitude for this noise, involving the optical coupling coefficients at 100 mK $C_{100\mathrm{mK}}$ (`K_cmbOverK_bath_temp`) that convert a 100 mK-stage thermal fluctuation into a detector power fluctuation:

$$\mathrm{PSD}_{\rm thermalbath} = (0.3\ C_{100\mathrm{mK}})^2\ \frac{2.4}{f}. \tag{6.17}$$

For example, for M1-140, $C_{100\mathrm{mK}} = 53\ \frac{\mu\mathrm{K_{CMB}}}{\mu\mathrm{K}}$ so the cosmic-ray 1/f amplitude is of about 611 $\mu\mathrm{K}^2_{\rm CMB} \left(\frac{1}{f}\right)$ (we precise the CMB unit here to avoid ambiguity). Taking a typical 60% correlation of this noise across the heated wafer leads to, for M1-140, an amplitude of correlated 1/f of 367 $\mu\mathrm{K}^2_{\rm CMB} \left(\frac{1}{f}\right)$ and 244 $\mu\mathrm{K}^2_{\rm CMB} \left(\frac{1}{f}\right)$ is uncorrelated.

These two PSD amplitudes, which specify a readout $1/f$ 100% correlated over a SQUID, and a cosmic ray $1/f$ 60% correlated over a wafer, appear in the IMo under the following identifiers:

- `psd_dac_ukcmb2`: readout 1/f due to gain fluctuations within the DAC
- `psd_thermalbath_corr_ukcmb2`: correlated 1/f due to cosmic ray heating at the detector wafer
- `psd_thermalbath_uncorr_ukcmb2`: uncorrelated 1/f due to cosmic ray heating at the detector wafer

`psd_thermalbath_corr_ukcmb2` and `psd_thermalbath_uncorr_ukcmb2` are stored at the frequency level channel, while `psd_dac_ukcmb2` is kept at the detector level since the noise level of a given detector could a priori vary within a given channel.

6.4. Detectors characteristics and mapping

6.4.1. Beams

Simple representation

As for bandpasses, we can use a simple representation of the beams and provide the design values of the full-width half maximum (FWHM) and ellipticity of the beam. These design values are common to all detectors of a given frequency channel, and we provide these at the frequency channel level.

- `fwhm_design_arcmin`: Design beam $\theta_{\rm FWHM}$ [arcmin] cf right columns of Table 6.1.
- `ellipticity_design`: Ellipticity - same as Frequency Channel for now - set to zero

The LiteBIRD beams FWHM and polarisation sensitivities can be combined together, assuming a Gaussian beam through Eq. 4.18 to provide a first estimate of the frequency-channel white-noise spectra. We delay their illustration to Fig. III.3.

GRASP simulation

One can numerically compute the expected beam detector per detector using GRASP from the optical models. Different levels of sophistication can be adopted in this computation, including or not the V-Grooves. Currently, MHFT beams have been computed for a subset of detectors, typically located at the corners and centers of the wafers: 119 out of 2074 MFT beams, 15 out of 1354 HFT beams, including the effects of V-Grooves, and 270 out of 1080 LFT beams (owing to the focal plane geometry, it is possible to recover the full focal plane). These are included in the IMo database as beam files. They contain both the real and imaginary parts of the electric field component perpendicular to the plane of observation (E_{co}) and the electric field component parallel to the plane of observation (E_{cx}), spanning θ, the angle between the main detector boresight and the line of sight, and ϕ, the angle of the cut around the boresight.

- `beam_file`: link to detector beam file without Vgrooves (if it exists)
- `beam_file_vgroove`: link to detector beam file with Vgrooves (if it exists)

We illustrate on the left panel of Fig. 6.7 a few detector beam cuts for $\phi = 0°$, the y-axis representing the amplitude in dB and the x-axis representing the θ angle. Complex structures appear in the far side lobes ($|\theta| \gtrsim 4°$) due to reflections and diffractions from the detector's neighbourhood. They can lead to high-power pick-up at a large angle that must be accurately estimated by calibration on the ground or in flight. The lack of knowledge of beam far side lobes is dominant among systematic effects. The beam knowledge to be achieved by calibration to satisfy the LiteBIRD requirements is derived in [182].

The right panel illustrates a full-2π beam of an HFT detector at 280 GHz located at the center of the W0 wafer. One can distinctly see the effect of the orientation of the HFT boresight with respect to the spacecraft: a well-localised shadow between $\phi = 30°$ and $\phi = 160°$ in the far-side lobes, which is due to V-grooves and the fact that the boresight of the focal planes dos not coincide with the boresight of the spacecraft (angle given by `spin_boresight_angle_deg`, `boresight_rotangle_rad`, `spin_rotangle_rad`) as illustrated in Fig. 5.2.

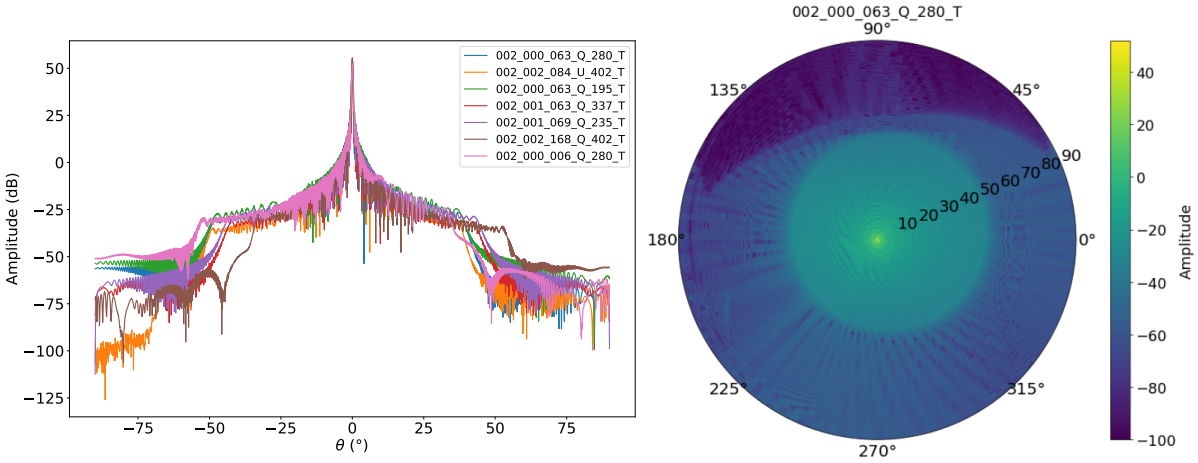

Figure 6.7: LEFT: few HFT GRASP beams cut in dB for fixed $\phi = 0°$. RIGHT: Full-2π representation of an HFT beam at 280 GHz, the detector being situated at the center of the W0 wafer.

6.4.2. Detectors information

In the IMo at the detector level, some specifications vary for each detector in the baseline design, such as the detector's name, bolometer type, wafer, SQUID and pixel IDs, and pointing-related quantities.

As emphasised in Sec. 6.4, there is a certain amount of wafers (*i.e.* detectors table) per instrument, in which there are detectors belonging to various frequency channels. The wafer ID mapping for the different instruments is represented in Fig. 6.10. In the IMo:

- `wafer`: Corresponding Wafer ID.

The detector ID mapping for the different instruments is represented in Fig. 6.10. In the IMo, it is stored as

- `pixel`: Corresponding Pixel ID.
- `pixtype`: Corresponding Pixel Type.

while the pixel type identifies the pixel technology. The LiteBIRD pixel types are quoted in Table 6.2: the trichroic ones are LF1-4 and MF1, the dichroic ones MF2, HF1 and HF2, and the monochroic one HF3.

We recall the NET without detector yield at the detector level:

- `NET_detector_ukrts`: NET detector : be careful the detector yield is NOT taken into account $[\mu\mathrm{K}\cdot\sqrt{\mathrm{s}}]$

The SQUID ID inside a given wafer indicates how the detectors of the wafer are mapped to a common SQUID (see Sec. 5.2.3). The SQUID ID mapping for the MFT is represented on the bottom-right of Fig. 6.10. As explained in Sec. 6.3.3, this information is crucial for generating correlated readout noise. In the IMo, the SQUID IDs are stored as:

Pixel type	Pixels count	Bolometer count	Frequency [GHz]	Bandwidth Δ_ν/ν	Pixel size [mm]
LF1	24	48	40/60/78	0.30/0.23/0.23	32
LF2	12	24	50/68/89	0.30/0.23/0.23	32
LF3	72	144	68/89/119	0.23/0.23/0.30	16
LF4	72	144	78/100/140	0.23/0.23/0.30	16
MF1	183	366	100/140/195	0.23/0.30/0.30	11.6
MF2	244	488	119/166	0.30/0.30	11.6
HF1	127	254	195/280	0.30/0.30	6.6
HF2	127	254	235/337	0.30/0.30	6.6
HF3	169	338	402	0.23	5.7

Table 6.2: The pixel in LiteBIRD: their ID, their amount, their frequency and their bandwidths

- `squid`: Corresponding SQUID ID.

The polarisation angle of the detector should be specified by three pieces of information about the detector: its handedness, its antenna type and its pixel orientation.

The detector handedness contains the information of the sign of the polarisation angle ψ and has to do with how the detector is set up on the wafer. We denote the handedness as $\mathcal{S}$: $\mathcal{S}_\mathrm{A} = 1$ and $\mathcal{S}_\mathrm{B} = -1$. Fig. 6.8 illustrates the effect of the handedness on the total polarisation angle. We represent in Fig. 6.10 the handedness pattern across the focal plane. In the IMo:

- `handedness`: detector handedness.

As explained in Sec. 6.4, for each pixel detecting photons at a certain frequency, there are two perpendicular bolometer antennas sensitive to the light of perpendicular polarisation. Fig. 6.8 illustrates the resulting polarisation angle from two distinct antenna types; the contribution to the polarisation angle from the antenna type can be viewed as $\psi_T = 0°$ while $\psi_B = 90°$. In the IMo this information is stored as

- `antenna_type`: detector antenna type.

A given pixel comes with a global orientation of the bolometers. The simplest prescription is with Q and U detectors (corresponding to orientation angles $\psi_Q = 0°$ and $\psi_U = 90°$) that are distributed uniformly inside the wafers. For MFT, the wafers are made with sub-wafers rotated by 30° angles, so, in the case of MFT detectors, this quantity can be an arbitrary angle in degree between 0° and 90°. We represent in Fig. 6.10 the angle pattern resulting from the polarisation orientation and antenna type. In the IMo:

- `polarisation_orientation`: polarisation angle around the boresight axis [deg].

From the polarisation information coming from `polarisation_orientation`, `antenna_type`, `handedness` one can derive the total polarisation angle ψ around the boresight of a given detector $\psi = (\psi_\mathrm{QU} + \psi_\mathrm{TB}) \times \mathcal{S}_\mathrm{AB}$. Fig. 6.8 summarises this prescription.

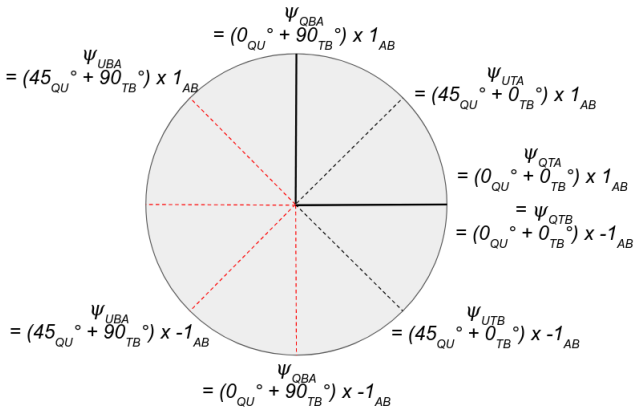

Figure 6.8: Prescription to derive the total polarisation angle from the orientation, antenna type and handedness of a bolometer.

For MFT detectors, orientation angles can differ from simple QU angles, and the polarisation angle computation simply generalises. The naming scheme of the LiteBIRD detectors synthesises the information we have introduced so far. For example, the ID of the detector 001_002_010_75B_140_B informs on the following characteristics:

- 001: It is on the instrument of ID 001, MFT (000: LFT, 001: MFT, 002: HFT),
- 002: ... in the wafer of ID 001,
- 010: ... at the pixel of ID 010,
- 010: ... at the pixel 010,
- 75B: ... with an orientation of 75° ("Q" or "U" would mean 0° or 45°) and an antenna type "B",
- 140: ... detecting at 140 GHz,
- B: ... and it is an "B" handedness.

For the 001_002_010_75B_140_B bolometer, the polarisation angle is $\psi = (+1) * (75 + 0)°$.

6.4.3. Detector pointings and quaternions

Each detector is oriented towards a specific region of the observed sky, a direction determined by the optical properties of the system. Generally, with knowledge of its optical design parameters and the position of a pixel on the focal plane, one can determine the corresponding direction using tools like GRASP, which simulates detector pointings based on the detector position and the shared telescope boresight axis. The pointing of a given bolometer refers to the direction it observes, expressed in any reference frame and represented in some coordinate system. In the IMo, we use three different coordinate systems to represent the pointings: Cartesian, Euler, and quaternion representations, the latter of which I have been responsible for producing.

Cartesian representation:

Any point on the unit sphere that a bolometer observes can be represented by its three Cartesian coordinates (u, v, w), subject to the normalisation constraint $w^2 = 1 - u^2 - v^2$ (whereby w is often not explicitly mentioned). The pointings provided by GRASP simulations are based

on the (u, v) direction conventions according to the following relations:

$$\vec{u}_{\text{pointing}} = -\vec{x}_{\text{focalplane}} \tag{6.18}$$

$$\vec{v}_{\text{pointing}} = \vec{y}_{\text{focalplane}} \tag{6.19}$$

$$\vec{w}_{\text{pointing}} = \vec{z}_{\text{focalplane}} \tag{6.20}$$

Fig. 6.9 schematises an optical system deflecting a light ray. Especially, the two coordinate bases are represented: $(\hat{x}, \hat{y}, \hat{z})$ associated with the physical positions of the detectors and $(\hat{u}, \hat{v}, \hat{w})$ associated with the pointings towards the sky. The cartesian paintings are stored in the IMo as:

- `pointing_u_v`: Vector of (u, v) on the sky describing the pointing of each detector

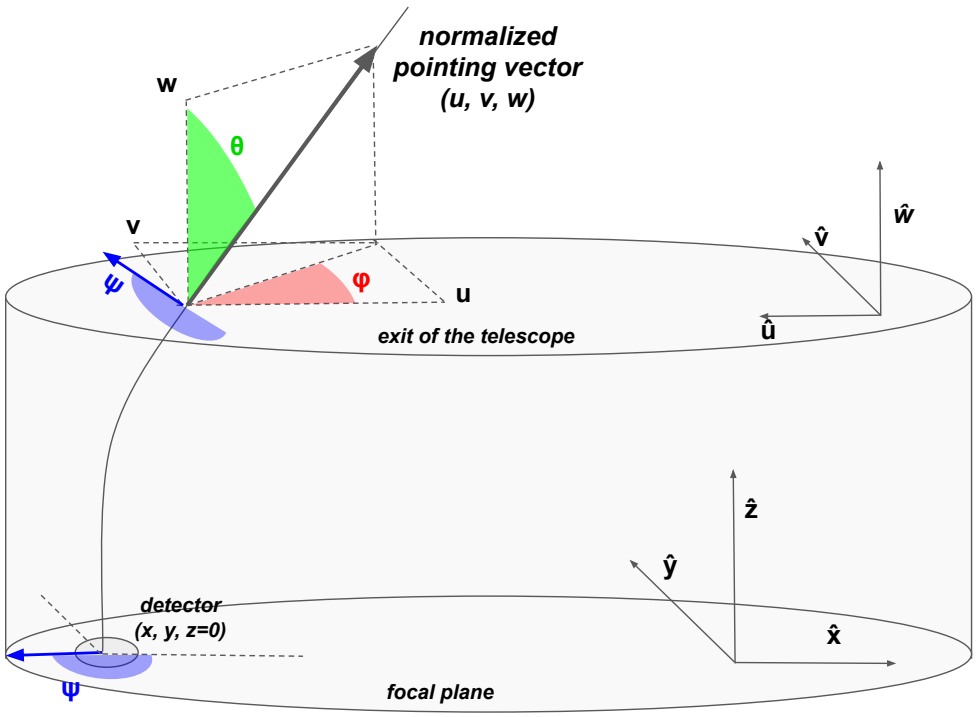

Figure 6.9: Schematic illustration of a specific detector pointing and its light ray in order to define the mechanical Cartesian coordinates (x, y, z), and the Euler (θ, ϕ, ψ) and Cartesian (u, v, w) pointing representation (beware of the difference between $\hat{x}$ and $\hat{u}$). For the purpose of the illustration, this drawing is not at scale and the θ angle is exaggerated.

We show the pointings computed from GRASP simulations in Cartesian representation in Fig. 6.10.

Euler representation:

The second pointing representation makes use of Euler angles $(\theta,\ \phi,\ \psi)$, with $\theta \in [0, 2\pi]$ and $\phi \in [0, \pi]$ containing the pointing information u, v, and the polarisation information in ψ. The conversion from (u, v) to (θ, ϕ) is given by $\phi = \arctan(v/u)$ and $\theta = \arcsin\sqrt{u^2 + v^2}$. Conversely, the Cartesian representation can be recovered with the relations $u = \sin(\theta)\cos(\phi)$

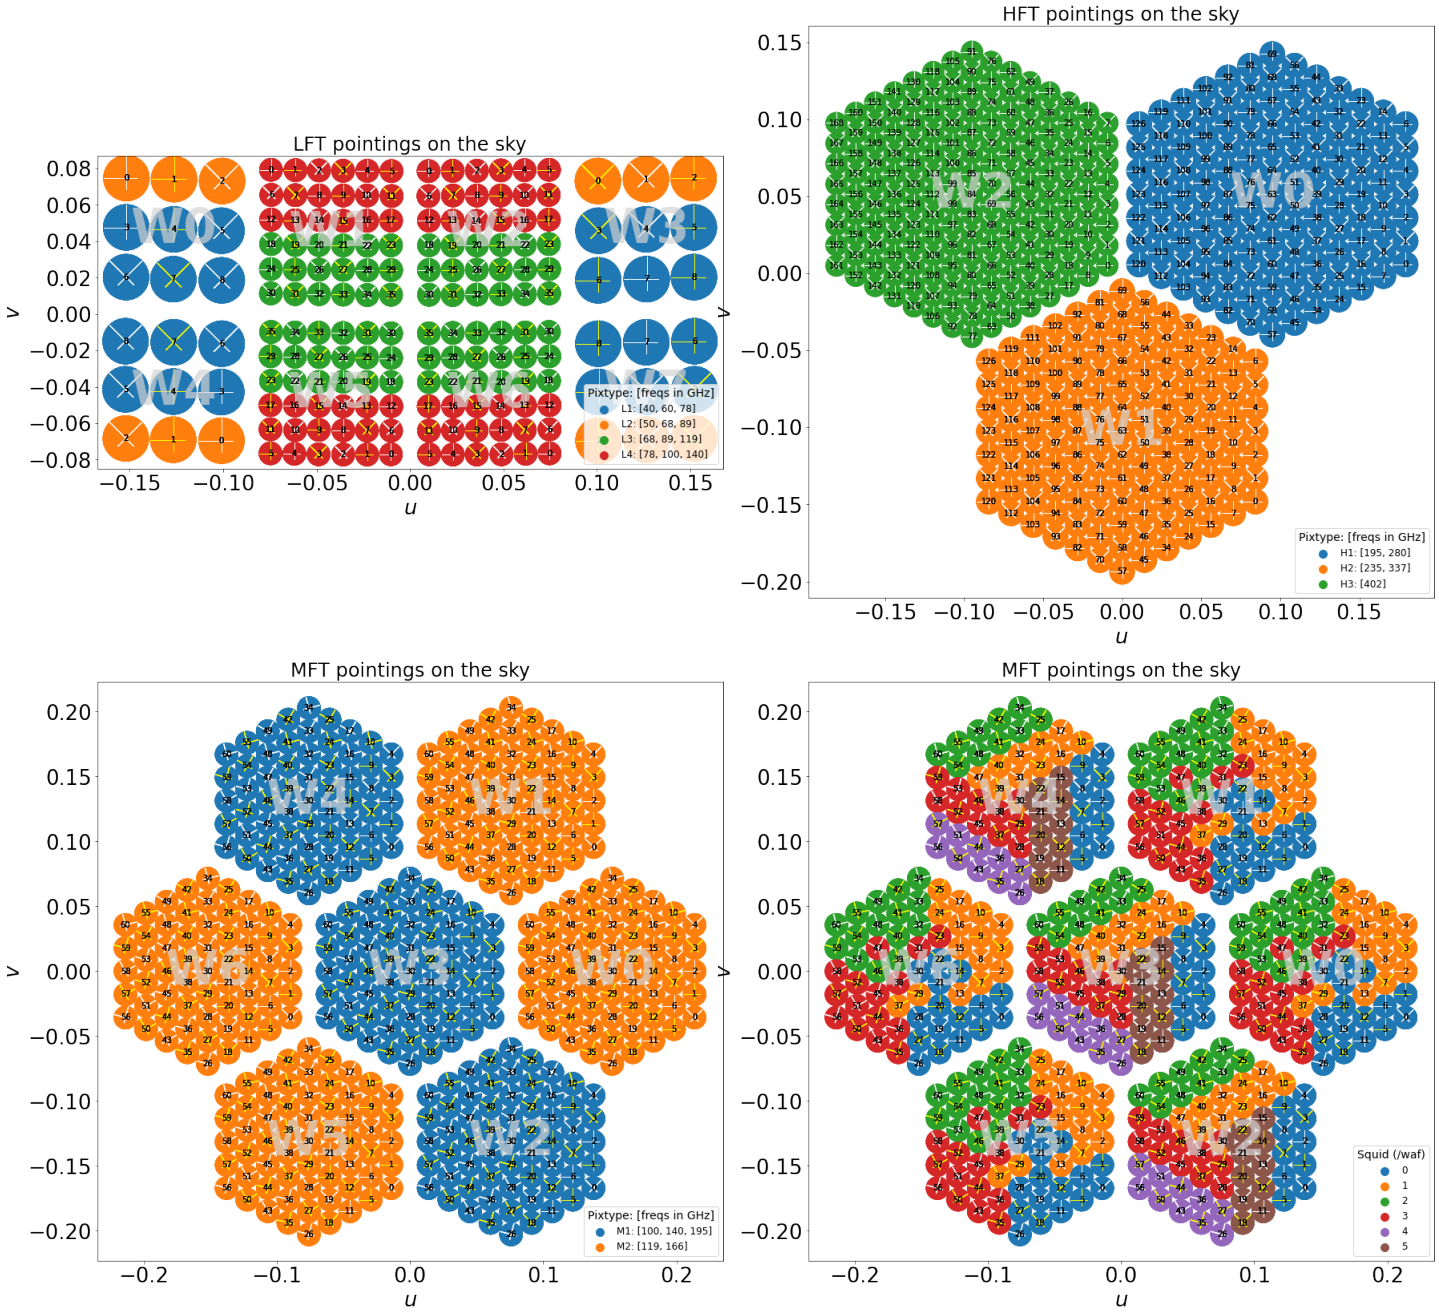

Figure 6.10: LiteBIRD pointings from GRASP simulation for each focal plane. The wafer and detector IDs are also displayed. Upper left: LFT, Upper right: HFT, Bottom left: MFT. On these three, the colours represent the frequency channels specified in the legend. Bottom right: MFT, with the colour representing the SQUID mapping. The crosses show the antenna angle pattern, in white when a detector is 'A' handed and in yellow otherwise. The circle markers' size scales with the pixels' mechanical diameter (and not with the beam size).

and $v = \sin(\theta)\sin(\phi)$. These Cartesian and Euler representations are summarised in Fig. 6.9. The Euler pointings appear in the IMo with the identifiers:

- `pointing_theta_phi_psi`: (theta,phi,psi) on the sky describing the pointing of each detector [deg]

Quaternion representation:

Lastly, a bolometer pointing and its polarisation can be represented by a unit quaternion $q = (q_0, q_1, q_2, q_3)$, with $q_i \in [-1, 1]$ and $\sum_i q_i^2 = 1$. The hypercomplex representation of a quaternion is $\mathbf{q} = q_0 + q_1\mathbf{i} + q_2\mathbf{j} + q_3\mathbf{k}$ where q_0 is called the real part and $q_1\mathbf{i} + q_2\mathbf{j} + q_3\mathbf{k}$ the imaginary or vectorial part.

A linear combination between two quaternions q and p can be written as

$$q \pm \lambda p = (q_0 \pm \lambda p_0, q_1 \pm \lambda p_1, q_2 \pm \lambda p_2, q_3 \pm \lambda p_3) \tag{6.21}$$

with λ a scalar, and the mutiplication between them is defined using the Hamilton product:

$$qp \equiv \begin{bmatrix} q_0p_0 - q_1p_1 - q_2p_2 - q_3p_3 \\ q_0p_1 + q_1p_0 + q_2p_3 - q_3p_2 \\ q_0p_2 + q_2p_0 + q_3p_1 - q_1p_3 \\ q_0p_3 + q_3p_0 + q_1p_2 - q_2p_1 \end{bmatrix} \tag{6.22}$$

Futhermore, one can define the inverse of the unit quaternion such that $q^{-1}q = 1$ as:

$$\bar{q} \equiv (q_0, -q_1, -q_2, -q_3) \tag{6.23}$$

One can go from the Euler representation of a pointing or a rotation (θ, ϕ) to its quaternion representation q with

$$q = (\sin(\phi)\sin(\theta/2), -\cos(\phi)\sin(\theta/2), 0, \cos(\theta/2)). \tag{6.24}$$

When we are solely interested in the pointing direction in a quaternion, it encodes only two pieces of information for its three degrees of freedom. Hence, one can impose a vanishing real part. In this case, there is a simple one-to-one relation between the cartesian representation (u, v) of the pointing and the vectorial part of the quaternion representation q,

$$\texttt{quat2vec} : q \to (u, v) = (q_1, q_2), \tag{6.25}$$

$$\texttt{vec2quat} : (u, v) \to q = (0, u, v, \sqrt{1 - u^2 - v^2}). \tag{6.26}$$

Rotations and changes of referential can also be handled within the quaternion formalism. One can change the referential of a pointing quaternion q_{R1} from $R1$ to $R2$ thanks to a quaternion encoding the rotation from $R1$ to $R2$, $q_{R1\to R2}$ and the following formula:

$$q_{R2} = q_{R1\to R2}\ q_{R1}\ q^{-1}_{R1\to R2} \tag{6.27}$$

This quaternion formalism is equivalent to the rotation matrix formalism. For instance, the previous rule in terms of rotation matrices becomes

$$q_{R2} = \texttt{rot2quat}\left(\texttt{rot2quat}(q_{R1\to R2}) * \texttt{rot2quat}(q_{R1}) * \texttt{rot2quat}(q_{R1\to R2})^{-1}\right), \tag{6.28}$$

where the operation `quat2rot` that allows the recovery of a quaternion q from its matrix representation R, is given by

$$\texttt{quat2rot} : q = \begin{bmatrix} q_0 \\ q_1 \\ q_2 \\ q_3 \end{bmatrix} \rightarrow R = \begin{bmatrix} q_0^2 + q_1^2 - q_2^2 - q_3^2, & 2q_1q_2 - 2q_0q_3, & 2q_1q_3 - 2q_0q_2 \\ 2q_1q_2 + 2q_0q_3, & q_0^2 - q_1^2 + q_2^2 - q_3^2, & 2q_2q_3 - 2q_0q_1 \\ 2q_1q_3 - 2q_0q_2, & 2q_2q_3 + 2q_0q_1, & q_0^2 - q_1^2 - q_2^2 + q_3^2 \end{bmatrix}, \tag{6.29}$$

and `rot2quat`, that allows the recovery of a quaternion q from its matrix representation R, is given by

$$\texttt{rot2quat} : R = \begin{bmatrix} R_{xx}, R_{xy}, R_{xz} \\ R_{yx}, R_{yy}, R_{yz} \\ R_{zx}, R_{zy}, R_{zz} \end{bmatrix} \rightarrow q = \begin{bmatrix} \frac{1}{2}\sqrt{1 + \mathrm{Tr}\ R} \\ \frac{1}{2\sqrt{1+\mathrm{Tr}\ R}}(R_{zy} - R_{yz}) \\ \frac{1}{2\sqrt{1+\mathrm{Tr}\ R}}(R_{xz} - R_{zx}) \\ \frac{1}{2\sqrt{1+\mathrm{Tr}\ R}}(R_{yx} - R_{xy}) \end{bmatrix}. \tag{6.30}$$

A useful quaternion property is that the composition of two rotations given by the quaternions q and p is the rotation given by their quaternion product qp. In matrix representation:

$$\texttt{quat2rot}(q) * \texttt{quat2rot}(p) = \texttt{quat2rot}(qp) \tag{6.31}$$

To build the IMo quaternions, we first use Eq. 6.24 to include the pointing information and get a q_{pos}. To include the polarisation information in a total quaternion q_{tot}, we construct the quaternion q_ψ as $q_\psi = (0, 0, \sin(\psi/2), \cos(\psi/2))$, and uses this composition rule Eq. 6.31 and

$$q_{\mathrm{IMo}} = q_{\mathrm{pos}} q_\psi = \texttt{rot2quat}(\texttt{quat2rot}(q_{\mathrm{pos}}) * \texttt{quat2rot}(q_\psi)). \tag{6.32}$$

q_{IMo}, which contains the information of the pointing and polarisation angle of the detectors, is included in the IMo database as

- `quaternions`: Vector of quaternions describing the pointing and polarisation orientation of each detector

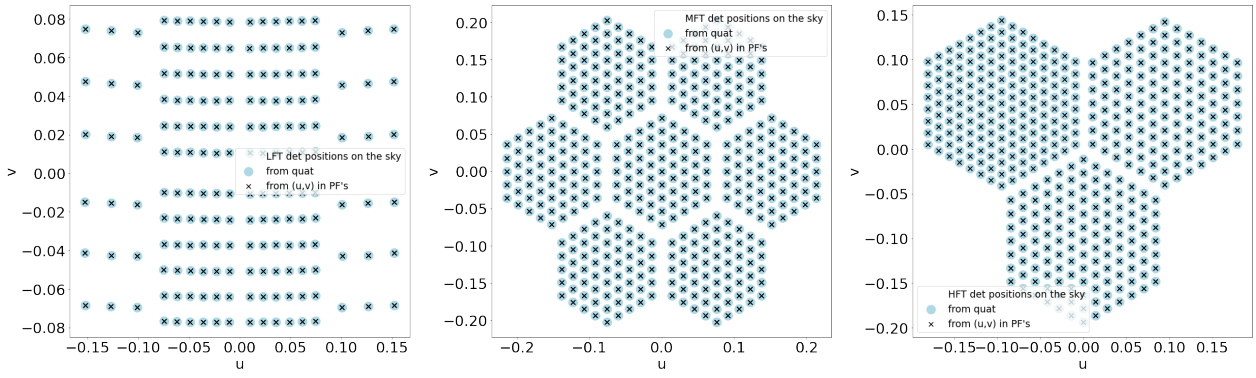

Figure 6.11: Cross-check of the validity of the IMo quaternion pointings

In Fig. 6.11, we verify that putting the IMo quaternions into Eq. 6.25 (black crosses) allows the recovery of the cartesian pointings (u, v) provided in the GRASP output files (blue points).

Finally, Eqs. 6.27 and 6.31 justify the extensive use of quaternion representation for LiteBIRD simulations: the changes of reference of the pointing vector become easy to perform. For example, to know the pointing and polarisation direction of a given detector on a given telescope in the ecliptic plane reference frame, one has to perform consecutively the following changes of referential:

1) From the detector reference frame, where the main beam of the radiation pattern is oriented along the z-axis, to the focal plane reference frame, rotation that can be described with $q_{\mathrm{FP,det}} \equiv q_{\mathrm{IMo}}$. This is the quaternion inside the IMo and the way to obtain it is detailed above.
2) From the focal plane reference frame, where the z-axis is aligned with the instrument boresight, to the spacecraft reference frame, a rotation that can be described with $q_{\mathrm{PLM,FP}}$. This quaternion itself results from the composition of various rotations introduced in the IMo at the System level of the instruments:
    a) A rotation of the focal plane with respect to the boresight direction by `boresight_rotangle_rad`,
    b) A rotation from the boresight direction to the spin axis by `spin_boresight_angle_rad`,
    c) An additional rotation boresight around the spin axis by `spin_rotangle_rad`.
3) From the spacecraft reference frame, where the z-axis is aligned with the spin axis, to the ecliptic reference frame, which is a fixed reference frame in relation to the plane of the Earth's orbit around the Sun, a rotation that can be described with $q_{\mathrm{ecl,PLM}}$. This quaternion is directly linked to the scanning strategy and depends on the time. It can itself be viewed as a quaternion composition encoding:
    a) the rotation of the spacecraft around the spin axis by an angle of $2 * \pi *$`spin_rate_rpm`$/60 * t$,
    b) the inclination of the spin axis with respect to the Ecliptic plane by an angle of $2 * \pi /$`precession_period_min`$/60 * t$,
    c) the rotation of the spin axis because of the precessional motion by an angle of $\pi/2 -$`spin_sun_angle_deg`$/180 * \pi$,
    d) the yearly spacecraft revolution around the Sun.

Hence, the quaternion associated with the main beam of the radiation pattern in the detector frame, $q^{\mathrm{det}}_{\mathrm{det}} \equiv$ `vec2quat`$((u = 0, v = 0, w = 1)) = (0, 0, 0, 1)$, transforms into the ecliptic frame like:

$$q^{\mathrm{det}}_{\mathrm{ecl}} = q_{\mathrm{ecl,PLM}}\ q_{\mathrm{PLM,FP}}\ q_{\mathrm{FP,det}}\ q^{\mathrm{det}}_{\mathrm{det}}\ q^{-1}_{\mathrm{FP,det}}\ q^{-1}_{\mathrm{PLM,FP}}\ q^{-1}_{\mathrm{ecl,PLM}}. \tag{6.33}$$

$q^{\mathrm{det}}_{\mathrm{ecl}}$ contains the information of the pointing and of the polarisation angle in the ecliptic plane and hence can be directly used to simulate time-ordered data consistent with the scanning strategy.

Finally, in the LiteBIRD case, one includes the effect of the HWP [183] by transforming the polarisation angle:

$$\psi \rightarrow \psi + 2\omega_{\mathrm{HWP}} t, \tag{6.34}$$

where ω_{HWP} is given by multiplying the IMo `hwp_rpm` by $2\pi/60$. This can be performed before the quaternion transformation or after having re-obtained the polarisation angle from the total quaternion.

LiteBIRD Instrument Model: Conclusion

In this section, we introduced the LiteBIRD mission and experiment. LiteBIRD is a space mission aimed at mapping the full-sky cosmic microwave background polarisation. Its goal is to achieve a total uncertainty on the tensor-to-scalar ratio, $\delta r < 1 \times 10^{-3}$, which will significantly enhance our understanding of inflation and fundamental physics, as emphasised in Sec. 3.3.3. This challenging requirement necessitates a thorough understanding of the instrument, motivating the creation and maintenance of a database to store all instrumental quantities and measurements, known as the Instrument Model or IMo for short.

We have provided a complete overview of the quantities currently stored in the IMo, especially those we have produced or implemented. The IMo is a key tool for the scientific collaboration, serving as input for simulation software, performance codes, and data analysis pipelines. After explaining how LiteBIRD's time-ordered data and frequency-map simulations are generated, the next section will focus on developing a full-analysis pipeline to extract τ and r from these simulations.

Part III

LiteBIRD data analysis preparation

This section is dedicated to describing a full analysis pipeline that I have developed to forecast LiteBIRD's ability to infer τ and r using simulations and to compare the results of different analysis assumptions.

Several options exist for extracting cosmological parameters θ_k from TOD_t^ν, the time-ordered data for each frequency ν. These options are summarised in Fig. III.1.

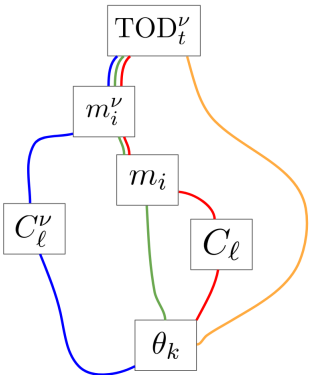

Figure III.1: Various paths from time-ordered data to cosmological parameters inference. The methodology we will focus on in this work is illustrated by the red path and has the most intermediary layers: the time-ordered data are projected to frequency maps that are cleaned before estimating the power spectrum and the associated cosmological parameters.

The most common method that is the focus of this work is illustrated by the red path. It involves first mitigating most of the systematics in the TOD and in the mapmaking process to obtain frequency maps m_i^ν, where i indexes the pixel for each observed frequency. The subsequent steps are the focus of the next three chapters in the particular LiteBIRD context:

1) First, we will present the results of **foreground cleaning** method performed at the frequency maps level (Chapter 7).
2) Then, we will use the resulting cleaned map m_i to estimate the CMB **power spectrum** C_ℓ, which contains most of the cosmological information (Chapter 8).
3) From this reduced-dimensionality dataset, we will discuss the **likelihood** of the model parameters, τ or r, before finally estimating them with their errorbars (Chapter 9).

In real cases, making the green path (which involves skipping the power spectrum estimation and directly estimating θ_k from m_i) efficient is tough: it requires an exquisite knowledge of the pixel-pixel covariance matrix. However, this path can be followed for analysing simulations whose noise properties are known, allowing for a very clean parameter estimation that can be used as a reference to test cleanly the other paths. Therefore, we will also discuss this approach. The two remaining paths represented in Fig. III.1, the blue one, which consists of working with a multi-frequency likelihood, and the orange one, which aims to conduct an end-to-end analysis from TODs to cosmology, avoiding all intermediate steps, are not the focus of this work.

The work presented hereafter is based on simulations. Simulations will also be a key object for the actual data analysis; for instance, estimating covariance matrices of the power spectra relies on these. Therefore, we begin this part with a preliminary chapter that discusses how to simulate LiteBIRD TODs and maps and introduces the input of the work.

Disclaimer:

The work presented in this part (the four next chapters) is partly based on research undertaken for the LiteBIRD project, includes unpublished material and has not been validated by the collaboration.

Several figures and results are based on inputs from the Simulation Team and have benefited from interactions within the E-modes team.

Specifically, Figs. III.4, 7.1 - 7.13, 8.9 - 8.17, 9.1, and 9.16 - 9.24, along with the associated results, are affected.

Therefore, these figures and results should be considered preliminary.
(And they will be identified by a "***Preliminary!***" mention in the caption).

Preliminary: simulating a microwave sky as seen by detectors

This preliminary chapter is dedicated to introducing the LiteBIRD simulations that will serve as input to the work.

We have already discussed the various components of the sky that contribute to the detector signal in Secs. 2.1, 4.1.1, 4.1.2, 4.2: the sky signal consists of a lensed polarised CMB contaminated by foreground emissions (mainly dust and synchrotron in polarisation). It is then affected by instrumental effects, including detector noise and systematic effects.

To simulate this, both an instrument and a sky model are required. The LiteBIRD simulations used as input for this part have been produced with the `litebird_sim` software developed by the LiteBIRD Simulation Team [174, 184]. They are based on the IMo that we have introduced in Chapter 6 and on a specific sky model that we detail hereafter. First, the fiducial cosmological parameters of the theoretical input CMB spectrum are given in Tab. 4.1. We chose to work with a $r = 0$ fiducial model since the full-success requirement of the mission, see Sec. 5.1, is defined assuming that there is no primordial B-modes.

One can derive the noiseless **signal** matrix from the resulting theoretical spectrum thanks to Eq. 2.35. To generate a realisation of CMB pixels consistent with the input spectrum, one simply does a random draw of the probability density function Eq. 2.34. To include the non-Gaussian effect of the lensing, the LiteBIRD sky model is based on a Websky convergence map, which allows to include the lensing deflection pixel per pixel (see Sec. 4.1.1 and [112, 185, 186]). The **extragalactic effects** like CIB (Sec. 4.1.2) and SZ (Sec. 4.1.2) trace the distribution of masses in the Websky map and are, therefore, correlated with the convergence map. Simulated radio sources based on catalogues are also included.

Since radio sources, CIB and SZ depend on the frequency, one map per frequency is generated. Then, they are coadded to a sky model of the Galactic emissions. These are produced by PySM 3 [187], a software that generates full-sky maps of **Galactic foregrounds** in intensity and polarisation. The foreground to be considered are thermal dust, synchrotron, AME, and free-free emissions, which we have introduced in Sec. 4.1.2. The PySM models are consistent with the various observations, especially Planck maps, and present varying complexity. Three combinations of PySM models will be discussed hereafter, which we quickly introduce here. For more details, one can refer to [188]. The first model, which is the one that has been used in [153]

(hence we call it "baseline"), includes a simple model of dust and synchrotron components:

baseline - `d1, s1`:

- `d1`: Thermal dust is modelled as a single-component modified black body with spectral variations [188, 189].
- `s1`: Synchrotron emission modelled as a power law with a spatially varying spectral index [40, 188, 190–192].

The second model includes higher complexity dust and synchrotron, alongside polarised CO and unpolarised AME and free-free:

medium complexity - `d10, s5, f1, a1, co3`:

- `d10`: Updated modified black body model with spectral variations [188].
- `s5`: Updated power law model with spectral variations [188].
- `co3`: Includes polarised CO lines [188].
- `a1`: AME emission modelled as a sum of two spinning dust populations. The AME emission is not polarised [188].
- `f1`: Free-free emission modeled using Commander Planck15 analytic model (unpolarised) [188].

The third model includes more complex dust and synchrotron as well as a polarised AME component.

high complexity - `d12, s7, f1, a2, co3`:

- `d12`: 3D polarised dust model with 6 dust clouds along the line of sight. Each cloud has a different spectral index and dust temperature templates [188].
- `s7`: Synchrotron emission modelled as a power law with a curved index [188].
- `a2`: AME has 2% polarisation fraction [188].

On the left panel of Fig. III.2 we represent the three sky models introduced above, in polarisation $P = \sqrt{Q^2 + U^2}$, in the frequency band and resolution of M1-100 (at 100 GHz) with increasing complexity from top to bottom. Differences are hard to distinguish and interpret at map-level, so we represent instead the associated beam-corrected TT, EE, and BB spectra on the upper right panel. One observes qualitatively that, in all cases, the trend of the foregrounds is a power law in multipole space. We represent the relative difference in percentage in the lower panel with respect to baseline foregrounds. One observes that for EE and BB, the medium- and small-scale behaviour deviates significantly from the baseline case: around $\ell = 100$ for M1-100, the high-complexity foregrounds are respectively two times and three times the baseline ones in BB and EE, while the medium-complexity ones are almost five times the baseline. The same figures are shown in the case of the L2-050 and H3-402 frequency channels in App. C. The question of which, among the medium or high complexity foregrounds, will be most impactful for cosmological parameter extraction is not trivial at this stage. One aim of this part will be to answer this question.

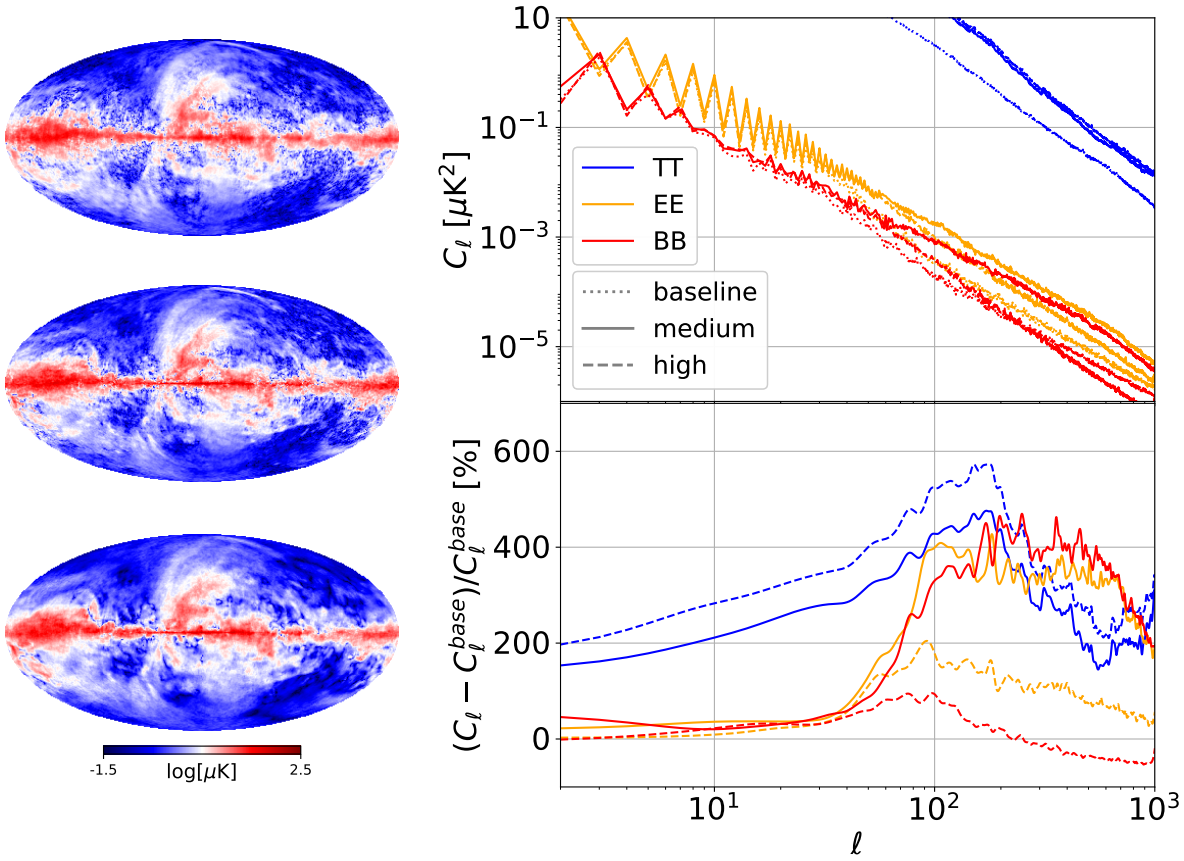

Figure III.2: LEFT: PySM template at $N_{\rm side} = 512$ of foreground $P = \sqrt{Q^2 + U^2}$ maps (no noise, no CMB) in the frequency band and at a resolution associated with the M1-100 frequency channel. From top to bottom, the foreground models increase in complexity. The colormap is in logarithmic scale. Top: baseline - `d1, s1`, middle: medium complexity - `d10, s5, f1, a1, co3`, bottom: high complexity - `d12, s7, f1, a2, co3`.
RIGHT: The associated beam-corrected M1-100 spectra are represented in the upper panel: *TT* (blue), *EE* (orange), *BB* (red), while the relative error to the baseline template in per cent is represented on the lower panel (on which we applied a Gaussian smoothing over ℓ with $\sigma = 4$ for illustration purpose).

On top of the sky map, one needs to compute the effect of the **motion of the spacecraft**. The latter is the composition of the motion of the spacecraft around L2, of the L2 point in the Ecliptic plane, of the motion of the Solar System around the Galactic Centre and of the motion of the Milky Way. The resulting velocity $\vec{\beta}$ induces a Doppler shift. The observed temperature $\tilde{T}$ coming from a line of sight direction $\hat{n}$ is linked to the actual temperature T by

$$T(\vec{\beta}, \hat{n}) = \frac{T_0}{\gamma(1 - \vec{\beta} \cdot \hat{n})}. \tag{III.1}$$

where $\gamma = (1 - \vec{\beta}\vec{\beta})^2$. One can apply this formula for a thermodynamic temperature shift in order to include the Doppler effect in the microwave-sky maps. The effect of the motion of the spacecraft with respect to the sun can be modelled and corrected, while the motion of the Sun and of the Galaxy results in a **dipole** in the measured signal that cannot be corrected for

and which is entirely degenerated with any astrophysical or cosmological dipole. Hereafter, the simulations we use have been generated with a dipole that is corrected for in a preprocessing step that we don't detail.

For the run of simulations used in this part, each of the sky components has been simulated assuming a top-hat **bandpass** for each LiteBIRD frequency channel, and the obtained maps are 4π-convoluted[1] by a Gaussian **beam**, with the numerical values given in the IMo (see Table 6.1).

At this stage, for each LiteBIRD detector, noiseless time-ordered data are produced by scanning the simulated sky during the LiteBIRD three years of survey. This is done thanks to the total quaternion given in Eq. 6.33, which describes the IMo detector's pointings and polarisation in ecliptic coordinates (hence including the effect of the LiteBIRD scanning strategy).

One includes **white noise** into the TOD simulations, using $\mathrm{NET}_{\mathrm{det}}$ as given by the IMo, which is degraded by a $1/\sqrt{\eta_{\mathrm{tot}}\eta_{\mathrm{yield}}}$ factor, see Sec. 6.3. We represent in Fig. III.3 the white noise spectra for each frequency channel deconvoluted for the beam. We also display the *EE* and *BB* fiducial spectra (fiducial values in Table 4.1). Each channel is associated with a transition multipole between signal-dominated and noise-dominated scales. If one takes the *EE* signal as a reference, this multipole is $\ell \simeq 120$ for L1-040, the lowest-resolution LiteBIRD channel, and $\ell \simeq 720$ for H3-402, the highest-resolution LiteBIRD channel.

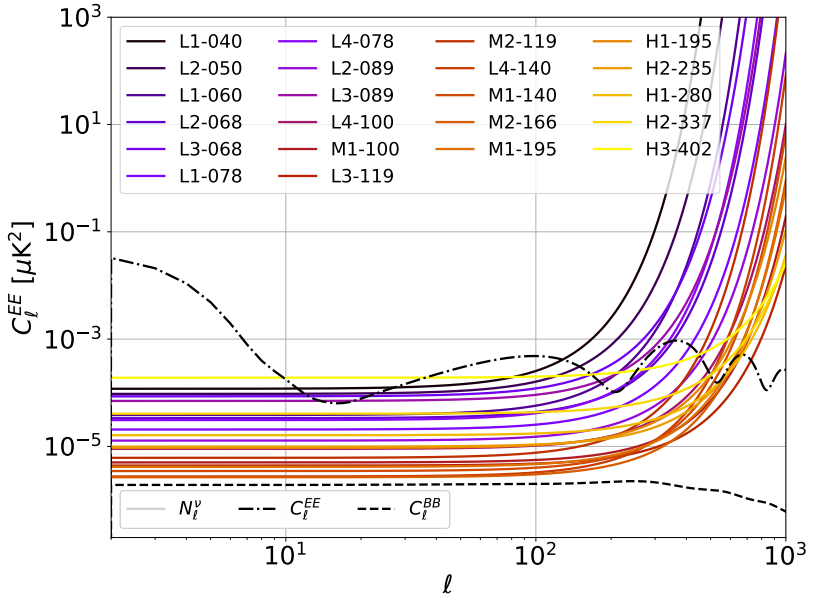

Figure III.3: Beam-corrected LiteBIRD frequency-channel white-noise polarisation spectra as compared with *EE* and *BB* fiducial spectra.

Finally, one adds a $1/f$ **noise** component at the TOD level using in our case a f_{knee} of 30 mHz. The $1/f$ noise is assumed to be uncorrelated between detectors for this run of simulations. We end up with TOD's, for each frequency, which contain foreground, white and $1/f$ noise. Other instrumental effects that can lead to **systematic uncertainties** are not included in these TODs, and current collaboration efforts are ongoing to include some of the effects introduced

[1]The convolution of two functions f and g defined over the surface of the unit sphere S^2 is given by $(f * g)(\mathbf{r}) = \int_{S^2} f(\mathbf{r}')\, g(\mathbf{r}\cdot\mathbf{r}')\, d\Omega(\mathbf{r}')$

in Sec. 4.2 (including full beam and bandpass shape detector per detector, cross-talk between detectors, *etc*).

From these simulated TOD's, the last step consists of building maps. We present here the standard way to perform *mapmaking*, through the use of a *binner*, which projects into a specific pixel the average of all data corresponding to the pixel, and of a *destriper*, which removes the effect of correlated instrumental noise from the TODs.

The mapmaking problem [193–196] consists of finding the three Stokes maps given a data timeline and the pointings information of the detectors. This problem can be formulated as finding the pixel vector $\mathbf{m}$, of shape $3N_{\rm pix}$, solution of the following linear problem:

$$y_t = \mathbf{P}_t^T \mathbf{m} + n'_t, \tag{III.2}$$

where $\mathbf{P}$ is the *pointing matrix* of shape $(N_t,\, 3N_{\rm pix})$, which relates times to pixels for each of the three Stokes parameters and n'_t is the noise timeline. For each time and pixel, the three remaining pointing matrix IQU elements are $(1, \cos(2\psi_t), \sin(2\psi_t))$ where ψ_t is the vector containing all time-ordered detector's polarisation angles including HWP modulation. Hence, for each time, the observed time-ordered data y_t can be decomposed in the I, Q, U basis as

$$y_t = I_t + Q_t \cos(2\psi_t) + U_t \sin(2\psi_t)\,. \tag{III.3}$$

Moreover, the noise n'_t can be decomposed in correlated and uncorrelated parts:

$$n'_t = F_t^{t'} a_{t'} + n_t, \tag{III.4}$$

where F correlates noise over time. The total noise covariance, which includes white and correlated noises, reads as

$$C_{{\rm total},tt'} = \left\langle n'_t n'_{t'} \right\rangle = F_t^{\tilde{t}} C_{{\rm a},\tilde{t}\tilde{t}'} F_{\tilde{t}'}^{t'} + C_{{\rm n},tt'}, \tag{III.5}$$

with $C_{{\rm a},tt'} = \langle a_t a_{t'} \rangle$ and $C_{{\rm n},tt'} = \langle n_t n_{t'} \rangle$.

Up to a constant, the Gaussian likelihood of the pixels can be written as

$$-2\ln\mathcal{L}(\mathbf{m}) \propto (y^t - F_{\tilde{t}}^t a^{\tilde{t}} - \mathbf{P}^{t,T}\mathbf{m}) \left(C_{\rm n}^{-1}\right)_t^{t'} (y_{t'} - F_{t'}^{\tilde{t}'} a_{\tilde{t}'} - \mathbf{P}_{t'}^T \mathbf{m}) + a^t \left(C_{\rm a}^{-1}\right)_t^{t'} a_{t'}. \tag{III.6}$$

The mathematically optimal (with minimum variance) solution of Eq. 7.1 is given by the *maximum-likelihood estimator*

$$\mathbf{m} = \left[\mathbf{P}^t \left(C_{\rm n}^{-1}\right)_t^{t'} \mathbf{P}_{t'}^T\right]^{-1} \mathbf{P}^t \left(C_{\rm n}^{-1}\right)_t^{t'} (y_{t'} - F_{t'}^{\tilde{t}'} a_{\tilde{t}'}), \tag{III.7}$$

which can be referred to as *binning equation*. a can be derived from requiring that the derivative of Eq. III.6 with respect to a vanishes, leading to an expression called *destriping equation*:

$$\left[F_{t'}^t \left(C_{\rm n}^{-1}\right)_t^{t'} \mathbf{Z} F_{t'}^t + \left(C_{\rm a}^{-1}\right)_t^{t'}\right] a_t = F_t^{t'} \left(C_{\rm n}^{-1}\right)_t^{t'} \mathbf{Z} y_t, \tag{III.8}$$

where $\mathbf{Z}$ is written as

$$\mathbf{Z} = \mathbf{I} - \mathbf{P}^{t'} \left[\mathbf{P}_{\tilde{t}}^T \left(C_{\rm n}^{-1}\right)_{\tilde{t}'}^{\tilde{t}} \mathbf{P}^{\tilde{t}'}\right]^{-1} \mathbf{P}_t^T \left(C_{\rm n}^{-1}\right)_{t'}^t. \tag{III.9}$$

Note that in practice, one is free to use a binning pointing matrix ($\mathbf{P}$ in Eq. III.7) at a different resolution than the one of the destriper ($\mathbf{P}$ in Eq. III.9) [196].

This destriping step allows us to handle most of the stripes, *ie.* time-correlated (and low-frequency) effects, which could remain after TODs' preprocessing. The role of the rotating HWP now becomes explicit: it is to modulate the polarised signal, shifting it to higher frequencies in the TOD's, while the unmodulated time-correlated instrumental noise remains at low frequencies. This separation eases the destriper task, which enable the mitigation of any polarised $1/f$ noise from the signal.

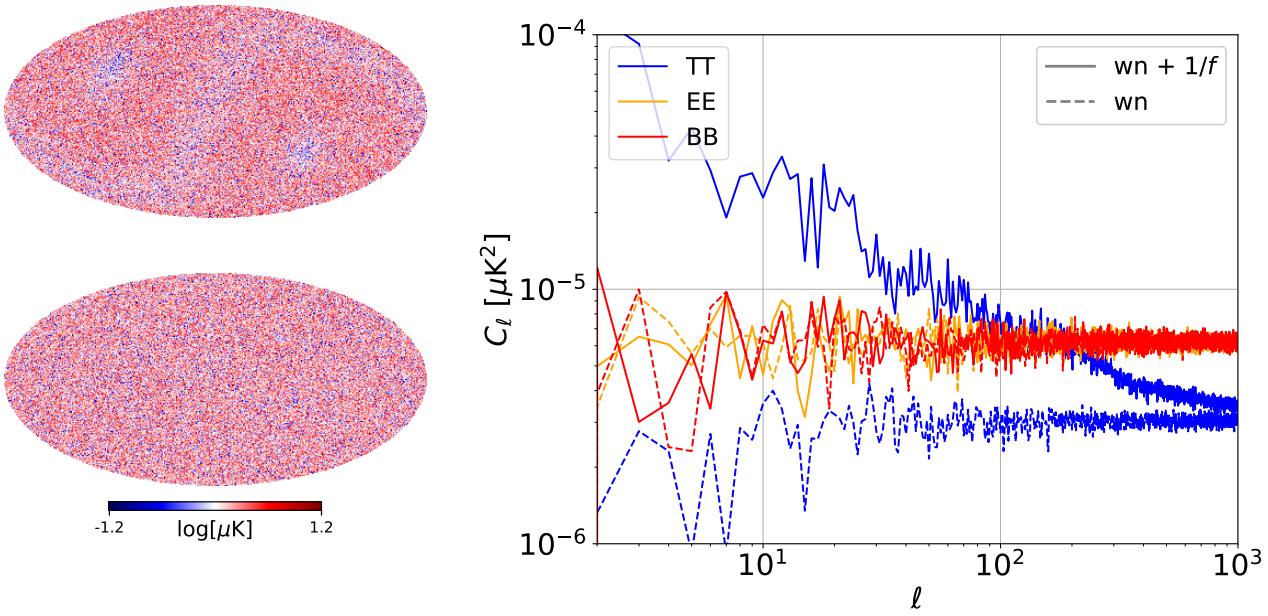

Figure III.4: ***Preliminary!*** LEFT: Particular realisation of $P = \sqrt{Q^2 + U^2}$ noise map for the M1-100 frequency channel with $1/f$ (top) or without (bot), at $N_{\rm side} = 512$. The white noise + $1/f$ power spectrum has been produced with a TOD simulation, destriping mapmaking, and full-sky power spectrum estimation steps, while the white-noise-only map is a map-based simulation. RIGHT: For the same two realisations, noise power spectrum for TT (blue), EE (orange), and BB (red), including $1/f$ noise (plain lines) or not (dashed lines), as detected by a full mission M1-100. For this figure, the spectra are not corrected for the beam window function.

On the left panel of Fig. III.4, we illustrate M1-100 noise polarisation ($P = \sqrt{Q^2 + U^2}$) maps obtained in two distinct ways: the top panel represents a TOD-based realisation of white noise + $1/f$ noise that includes the effect of the non-uniform scanning strategy, while the lower panel one is a simple map-based realisation of white noise. The white noise level is the same in the two maps. One observes patterns in the TOD-based realisation that are due to the non-exactly uniform scanning-strategy hits. The figure can be compared to the hit map, see Fig. 6.4: in a given sky region, the fewer hits there are, the higher the noise is.

On the right panel, we represent the TT (blue), EE (orange), and BB (red) power spectra corresponding to these two realisations. One notices the expected factor 2 between the flat spectra levels due to white noise in TT and the polarisation spectra. This factor is due to the splitting of the polarised power between the two antennas, see Eq. 4.15. Finally, one can observe the effect of $1/f$ noise by comparing dashed with plain lines. In temperature, it appears as a significant $1/\ell$ TT noise spectrum. The destriping has reduced but not cancelled this low-ℓ

temperature tail. Instead, as expected, the destriping induces a full suppression of the *1/f* noise in polarisation, thanks to the modulation of the signal by the rotating HWP.

In this part, we will use maps at resolution $N_{\rm side} = 512$, and will explore four simulation sets to see the impact of the simulation type on the data-analysis techniques we will investigate:

Four sets of simulations:

- "wn + baseline": map-based simulations that contain white noise and baseline foregrounds (used previously in [166]),
- "wn1f + baseline": TOD-based simulations that contain scanning strategy effect, white noise and residual $1/f$ noise after destriping, and baseline foregrounds,
- "wn1f + medium": TOD-based simulations that contain scanning strategy effect, white noise and residual $1/f$ noise after destriping, and medium-complexity foregrounds,
- "wn1f + high": TOD-based simulations that contain scanning strategy effect, white noise and residual $1/f$ noise after destriping, and high-complexity foregrounds.

Each of them will consist of 100 simulations for "wn + baseline" and 200 for the three other ones. We will also use the noise and foreground simulated maps independently for assessment purposes (even though they will not be available inputs in a real scenario).

CHAPTER 7

Component separation and foreground cleaning

This chapter is dedicated to assessing the performance of a particular foreground-cleaning method in the context of LiteBIRD called *Harmonic Internal Linear Combination* (HILC). First, in Sec. 7.1, we will provide an overview of the various classes of component separation methods. Then, we will introduce step-by-step the equations that we have used for our implementation of HILC in Sec. 7.2. Then, we will discuss the particular setup we have worked with and present the first results in Sec. 7.3. In Sec. 7.4, we will explain how to derive optimal HILC configurations that maximise the r and τ information contained in the EE and BB recovered spectra. Finally, we will conclude on the performance of HILC for cleaning the various foreground complexities introduced in the preliminary chapter of Part. III.

Contents

7.1 Methods overview 144

7.2 Harmonic Internal Linear Combination 146

7.3 HILC for LiteBIRD: setup and products 149

7.3.1 Inputs 149

7.3.2 Binned and smoothed weights 150

7.3.3 Output maps 151

7.3.4 Output power spectra 151

7.3.5 Two classes of mask 153

7.4 Optimisation and results 153

7.4.1 Masking input maps 153

7.4.2 Masking output maps 154

7.4.3 Optimising binning and smoothing 156

7.4.4 More complex foregrounds and noise 160

7.1. Methods overview

In this section, we introduce the basics of the component separation problem.

Let us start from some data object $\mathbf{d}$ of shape (n_p, n_ν) coming from an experiment measuring the sky at n_ν frequencies with n_p data points. Let us assume that the total signal comes from the contribution of n_c components whose signal we want to estimate. The data can be modelled as a linear combination of the various component signals:

$$d_{p\nu} = A^c_{p\nu} s_{pc} + n_{p\nu}, \tag{7.1}$$

where s is the components object of shape (n_p, n_c), n is the noise object of shape (n_p, n_ν), and A is the mixing matrix between the components, of shape (n_p, n_ν, n_c). What we want to estimate is s; the game is then to invert Eq. 7.1.

The various methodologies for separating components or removing foregrounds lie in distinct choices regarding which components to extract and the form of the mixing matrix. The first (and most interesting to us) microwave component to consider is the CMB, which is the component with the best-known frequency spectrum (see Sec. 2.4). In thermodynamic temperature units (see Sec. 4.1.2), the mixing matrix CMB column is exactly filled with ones.

Parametric methods

If one wants to also reconstruct the other components contributing to the microwave sky intensity and polarisation, which we introduced in Sec. 4.1.2, one can assume a model for the frequency dependency of the components and parameterise the mixing matrix as $A = A(\{\beta_i\}_{i\in[1,n_\beta]})$, where $\{\beta_i\}$ are the n_β parameters one chooses to parameterise the mixing matrix with. The minimal condition for inverting Eq. 7.1 and estimating the best parameters from the data is $n_\beta < n_\nu$. The class of methods aiming at doing this are called *parametric*.

Let us stress that we have not specified the nature of the data so far. There exist several distinct choices: for instance, the data might be map pixels (for several Stokes parameters or not), harmonic coefficients $a_{\ell m}$ or needlet coefficients. *A priori*, they contain the same amount of information. In practice for foreground cleaning and component separation methods, these choices have implications for the foreground cleaning methods and will lead to different cleaned outputs. Since these data types are all Gaussianly distributed, very similarly to the mapmaking problem introduced in the previous section, one can write the likelihood for s as:

$$-2\ln\mathcal{L}(s, \{\beta_i\}) \propto (d^{p\nu} - A^{p\nu}_c s^{pc}) \left(N^{-1}\right)^{p'\nu'}_{p\nu} (d_{p'\nu'} - A^{c'}_{p'\nu'} s_{p'c'}) \tag{7.2}$$

where N is the noise matrix of the multi-frequency data, organised in blocks and of total shape $(n_p \times n_\nu, n_p \times n_\nu)$. The mathematically optimal (with minimum variance) solution of Eq. 7.1 is given by the maximum-likelihood estimator [197–199], which allows to recover β_i from:

$$\frac{\partial \ln\mathcal{L}(s, \beta_i)}{\partial\beta_i} = -2\frac{\partial A^{p\nu}_c}{\partial\beta_i} s^{pc} \left(N^{-1}\right)^{p'\nu'}_{p\nu} (d_{p'\nu'} - A^{c'}_{p'\nu'} s_{p'c'}) = 0, \tag{7.3}$$

while $\hat{s}$ is given by

$$\hat{s}^{c'}_p = \left\{ \left[A_{\tilde{p}'\tilde{\nu}'} \left(N^{-1}\right)^{\tilde{p}'\tilde{\nu}'}_{\tilde{p}\tilde{\nu}} A^{\tilde{p}\tilde{\nu},T} \right]^{-1} \right\}^{c'}_{c} A^c_{p'\nu'} \left(N^{-1}\right)^{p'\nu'}_{p\nu} d^\nu_{p'} \equiv w^{c'p'}_{p\nu} d^\nu_{p'}. \tag{7.4}$$

where w, the so-called *weights*, are defined as the coefficients linking d to s for each frequency. The maximum-likelihood weights that can be read in Eqs. 7.3 and 7.4 noticeably involve the

noise matrix N. Hence, in real-case contexts, this method is delicate to apply while staying agnostic because it requires knowing N (which has a lot more elements - $n_p^2 \times n_\nu^2$ - than the number of pixels in the maps) and being able to invert it. Moreover, prior foreground knowledge is required to model the mixing matrix.

Non-parametric methods

A second way of using Eq. 7.1 without assuming the model for other components than CMB involves solving for the elements of A directly rather than the SED parameters, see for instance [200]. However, the knowledge of N is still required.

Internal Linear Combination methods

Another linear method called the *minimal variance Internal Linear Combination (ILC)* approaches the problem of estimating weights with a different philosophy, allowing for a workaround to the issue of the unknown N. One can directly look for an estimate $\hat{s}$ of the CMB signal as a *internal linear combination* of the multi-frequency data vector $d_{p\nu}$,

$$\hat{s}_p = w_p^\nu d_{p\nu}. \tag{7.5}$$

Typically, this data vector will be

$$d_{p\nu} = \mathrm{cmb}_p + \mathrm{fg}_{p\nu} + n_{p\nu} \tag{7.6}$$

where "cmb" is the contribution from the CMB, "fg" from the foregrounds, and "n" from the noise. In the absence of noise and foreground, the CMB sky should exactly match its estimate, $\mathrm{cmb}_p \equiv \hat{s}_p = w_p^\nu \mathbb{1}_\nu \mathrm{cmb}_p$. Hence, to avoid losing or increasing the CMB signal in the foreground cleaning, the weights have to sum up to one $\forall p,\ w_p^\nu \mathbb{1}_\nu = 1$.

The *minimal variance ILC* estimator, first introduced in the CMB context in [201–203] for WMAP, consists in finding the weights at each data point such that the variance of $\hat{s}_p$ is minimal, while preserving the previous constraint:

$$\begin{cases} \frac{\partial \langle \hat{s}^2 \rangle}{\partial w_p^\nu} = 0, \\ w_p^\nu \mathbb{1}_\nu = 1. \end{cases} \tag{7.7}$$

The $\hat{s}$ variance can be linked to the covariance of d for a given data point[1]:

$$\left\langle \hat{s}^2 \right\rangle_p = w_p^\nu \left\langle d_{p\nu} d_p^{\nu'} \right\rangle w_{p\nu'} \tag{7.8}$$

$$= w_p^\nu C_{p\nu}^{\nu'} w_{p\nu'}, \tag{7.9}$$

where C are the n_p covariance matrices of the multi-frequency data, of shape (ν, ν'). Then, the equations to solve jointly are:

$$\begin{cases} \frac{\partial}{\partial w_p^\nu} \left(w_p^\nu C_{p\nu}^{\nu'} w_{p\nu'} \right) = 0, \\ w_p^\nu \mathbb{1}_\nu = 1. \end{cases} \tag{7.10}$$

This constrained minimisation problem for w can be rewritten in an equivalent way as

$$\frac{\partial}{\partial w_p^\nu} \mathcal{L} = \frac{\partial}{\partial \lambda} \mathcal{L} = 0 \tag{7.11}$$

[1] We assumed the correlation between distinct data points p to vanish. Correlation between distinct data points might become relevant to consider when 1) working in harmonic space with polarisation and willing to consider correlation between T, E, B at a given scale, 2) cutting the sky in harmonic space, 3) working at high resolution in pixel space.

where λ is a Lagrange multiplier that defines the Lagrangian $\mathcal{L}$ as

$$\mathcal{L} = w_p^\nu C_{p\nu}^{\nu'} w_{p\nu'} + \lambda \left(1 - w_p^\nu \mathbb{1}_\nu\right). \tag{7.12}$$

Eqs. 7.10 are solved by

$$\lambda = \frac{2}{\mathbb{1}^{\nu'} \left(C_p^{-1}\right)_{\nu'}^{\nu} \mathbb{1}_\nu}, \tag{7.13}$$

$$w_p^\nu = \frac{\lambda \mathbb{1}^{\nu'} \left(C_p^{-1}\right)_{\nu'}^{\nu}}{2}. \tag{7.14}$$

Eventually, the weights reads

$$w_p^\nu = \frac{\mathbb{1}^{\nu'} \left(C_p^{-1}\right)_{\nu'}^{\nu}}{\mathbb{1}^{\nu'} \left(C_p^{-1}\right)_{\nu'}^{\nu} \mathbb{1}_\nu}, \tag{7.15}$$

which we can reinsert in Eq. 7.9 to get the total variance at a given data point:

$$\hat{\sigma}^2_{\rm tot} \equiv \left\langle \hat{s}^2 \right\rangle_p = \frac{1}{\mathbb{1}^{\nu'} \left(C_p^{-1}\right)_{\nu'}^{\nu} \mathbb{1}_\nu}, \tag{7.16}$$

which can *per se* be used as a criterion to compare methods or strategies of implementation.

In practice, and as we have already done earlier, the ensemble average in Eq. 7.9 has to be estimated as averages over subsets $\mathcal{D}_p$ of $N_{\mathcal{D}_p}$ data:

$$C_{p\in\mathcal{D}_p,\nu}^{\nu'} = \frac{1}{N_{\mathcal{D}_p}} \sum_{p'\in\mathcal{D}_p} d_{p'\nu} d^{p'\nu'}. \tag{7.17}$$

For any p, $N_{\mathcal{D}_p}$ needs to be large enough for the covariance estimate to be unbiased, invertible and induce no extra variance. Ensuring that $C_{p\in\mathcal{D}_p,\nu}^{\nu'}$ is invertible (hence minimally that $N_{\mathcal{D}_p} > n_\nu$) is often not enough to get a low $\hat{\sigma}^2_{\rm tot}$. One needs to ensure that the estimation of the matrix itself is not too noisy in order to ensure that $\hat{\sigma}^2_{\rm tot}$ remains reliably low.

As for the non-parametric methods, the minimum variance ILC method provides weights in a totally blind way, assuming nothing else than the fact that the CMB is a blackbody. No further assumption about CMB, foregrounds or noise is required. Furthermore, it does not require a knowledge of the noise covariance matrix, which makes the method particularly agnostic. However, [204] showed that it can be sensitive to calibration errors. ILC methods have been previously explored and developed in pixel space [202, 205], needlet space [206, 207], and harmonic space [201, 208].

7.2. Harmonic Internal Linear Combination

Harmonic ILC is an application of the minimum-variance ILC method described above, using harmonic coefficients as data, first introduced in the CMB context in [201]. This section is dedicated to introducing the important notations and equations for this particular method, which I have implemented, improved and optimised in the LiteBIRD context. Starting from some observed or simulated maps $\mathbf{m}^\nu$ at various resolutions θ^ν, one estimates $a_{\ell m}^\nu$ from Eq. 2.16 that should be corrected by the (beam and pixel) window functions. In the context of polarisation

foreground cleaning, one can work either with $a^{E,\nu}_{\ell m}$ or $a^{B,\nu}_{\ell m}$ without the need to treat them both simultaneously since they are *a priori* uncorrelated in a non-birefringent and full-sky setup. The HILC estimate of the CMB $a_{\ell m}$'s is given

$$\hat{a}_{\ell m} = \sum_{\nu} w_{\ell\nu} a^{\nu}_{\ell m}. \tag{7.18}$$

where we have assumed that the weights are common to all m at a given ℓ. To ensure the quality of the required frequency-frequency covariance matrix, we need enough m per estimation. To ensure this, we chose a list of bins over the multipoles, b, then we empirically estimate the covariance over the entire bin, according to Eq. 7.17:

$$C^{\nu\nu'}_{\ell\in[b_L,b_{L+1}[} = \frac{1}{\tilde{N}_m} \sum_{\ell=b_L}^{b_{L+1}-1} \sum_{m=-\ell}^{\ell} a^{\nu}_{\ell m} a^{*\nu'}_{\ell m}, \tag{7.19}$$

where $\tilde{N}_m$ is the number of $a_{\ell m}$ per bin, $\sum_{\ell=b_L}^{b_{L+1}-1}(2\ell+1)$. Typically, to build b, one wants to impose an approximate number of modes per bin, N_m. Then the bin starting at b_L will end at b_{L+1} the smallest multipole satisfying $\sum_{\ell=b_L}^{b_{L+1}}(2\ell+1) \geq N_m$. This is the choice made in this work.

Then, one can build the binned HILC weights according to Eq. 7.15:

$$\tilde{w}^{\nu}_{\ell} = \frac{\mathbb{1}^{\nu'} \left(C^{-1}_{\ell}\right)^{\nu}_{\nu'}}{\mathbb{1}^{\nu'} \left(C^{-1}_{\ell}\right)^{\nu}_{\nu'} \mathbb{1}_{\nu}}. \tag{7.20}$$

$\tilde{w}^{\nu}_{\ell}$ are discontinuous over the multipoles, which is not physical and propagates to discontinuities in the estimated power spectrum. This is the first and main motivation to apply a low-pass filter using a Gaussian window function over ℓ to the weights. This is a feature newly introduced by this work that we investigate in Sec. 7.3. Formally, this consists of defining the smoothed weights w^{ν}_{ℓ} as

$$w^{\nu}_{\ell} = \sum_{\tilde{\ell}\in\mathbb{Z}} \frac{1}{\sqrt{2\pi}\sigma_{w_\ell}} \exp\left(-\frac{\tilde{\ell}^2}{2\sigma^2_{w_\ell}}\right) \tilde{w}^{\nu}_{\ell-\tilde{\ell}}, \tag{7.21}$$

with σ_{w_ℓ} the smoothing size in multipole space, and $\tilde{\ell}$ are integers. The boundary conditions for this sum are defined such that $w_{\ell<2} = w_2$ and $w_{\ell>\ell_{\max}} = w_{\ell_{\max}}$. The property $\sum w = 1$ remains true after smoothing:

$$\sum_{\nu} w^{\nu}_{\ell} = \sum_{\tilde{\ell}\in\mathbb{Z}} \frac{1}{\sqrt{2\pi}\sigma_{w_\ell}} \exp\left(-\frac{\tilde{\ell}^2}{2\sigma^2_{w_\ell}}\right) \sum_{\nu} \tilde{w}^{\nu}_{\ell-\tilde{\ell}} \tag{7.22}$$

$$= \sum_{\tilde{\ell}\in\mathbb{Z}} \frac{1}{\sqrt{2\pi}\sigma_{w_\ell}} \exp\left(-\frac{\tilde{\ell}^2}{2\sigma^2_{w_\ell}}\right) \tag{7.23}$$

$$= 1. \tag{7.24}$$

Simultaneous optimisation of $\tilde{N}_m$ and σ_{w_ℓ} will be done in Sec. 7.4.3 in order to get the best results from HILC. The final products of HILC are either the estimated $\hat{a}_{\ell m}$, the cleaned map obtained applying Eq. 2.54 to $\hat{a}_{\ell m}$, or the total variance of the estimator, which is itself the

estimator of the power spectrum, $\hat{C}_\ell$. Let us develop Eq. 7.16 in the harmonic context. Eq. 7.6 becomes $a^\nu_{\ell m} = \mathbb{1}_\nu \mathrm{cmb}_{\ell m} + \mathrm{fg}^\nu_{\ell m} + \mathrm{n}^\nu_{\ell m}$, whose variance reads:

$$\hat{C}_\ell = \left\langle |\hat{a}_{\ell m}|^2 \right\rangle \tag{7.25}$$

$$= \left\langle |w_{\ell\nu} a^\nu_{\ell m}|^2 \right\rangle \tag{7.26}$$

$$= \left\langle |\mathrm{cmb}_{\ell m} + w_{\ell\nu}\mathrm{fg}^\nu_{\ell m} + w_{\ell\nu}\mathrm{n}^\nu_{\ell m}|^2 \right\rangle . \tag{7.27}$$

Developing the modulus square leads to

$$\begin{aligned}\hat{C}_\ell = \Big\langle\ & |\mathrm{cmb}_{\ell m}|^2 + \mathrm{cmb}_{\ell m} w_{\ell\nu} \mathrm{fg}^{\nu\,*}_{\ell m} + \mathrm{cmb}_{\ell m} w_{\ell\nu} \mathrm{n}^{\nu\,*}_{\ell m} \\ & + w_{\ell\nu}\mathrm{fg}^\nu_{\ell m}\mathrm{cmb}^*_{\ell m} + |w_{\ell\nu}\mathrm{fg}^\nu_{\ell m}|^2 + w_{\ell\nu}\mathrm{fg}^\nu_{\ell m} w_{\ell\nu'} \mathrm{n}^{\nu'\,*}_{\ell m} \\ & + w_{\ell\nu}\mathrm{n}^\nu_{\ell m}\mathrm{cmb}^*_{\ell m} + w_{\ell\nu}\mathrm{n}^\nu_{\ell m} w_{\ell\nu'} \mathrm{fg}^{\nu'\,*}_{\ell m} + |w_{\ell\nu}\mathrm{n}^\nu_{\ell m}|^2 \Big\rangle \end{aligned} \tag{7.28}$$

$$= C_\ell^{\mathrm{cmb}} + C_\ell^{\mathrm{resid}}, \tag{7.29}$$

where the residual C_ℓ^{resid} is the difference between the output spectrum after HILC and the input CMB spectrum and consists of a positive bias in the power-spectrum estimate. We call it the total residual. Assuming that the CMB, foregrounds and noise are independent leads to splitting the total residual into two terms,

$$C_\ell^{\mathrm{resid}} = w_{\ell\nu} \left\langle \mathrm{fg}^\nu_{\ell m} \mathrm{fg}^{*\nu'}_{\ell m} \right\rangle w_{\ell\nu'} + w_{\ell\nu} \left\langle \mathrm{n}^\nu_{\ell m} \mathrm{n}^{*\nu'}_{\ell m} \right\rangle w_{\ell\nu'} \tag{7.30}$$

$$\equiv w_{\ell\nu} C_\ell^{\mathrm{fg},\nu\nu'} w_{\ell\nu'} + w_{\ell\nu} C_\ell^{\mathrm{n},\nu\nu'} w_{\ell\nu'}. \tag{7.31}$$

where the first term stands for the foreground residual, C_ℓ^{fg}, and the second term for the noise residual, C_ℓ^{n}. Under the same assumption and similarly to Eq. 7.18, one can define the $a_{\ell m}$ *residuals* by linearly combining $\mathrm{n}^\nu_{\ell m}$ or $\mathrm{fg}^\nu_{\ell m}$ to the weights. The *map residuals* are obtained by applying Eq. 2.54 to the $a_{\ell m}$ residuals.

Both noise and foreground residuals contribute to the variance of the spectrum estimate. Since the HILC weights sum up to one, the resulting noise residual is bounded between the power spectra associated with the most sensitive and the least sensitive frequency channels (*i.e.*, cases where the weight associated with the best/worst channel is one while the others vanish). Consequently, the noise residual is always larger than the noise spectrum one would obtain by combining all frequency channels, see Eq. 4.16. Additionally, the cosmic variance of the foreground residual increases the total variance of the estimated power spectrum. The degradation in the variance due to the foreground cleaning with respect to the quadrature sum of the (foreground-free) sensitivity is referred to as the *degradation factor* [138], that we denote later on by Δ:

$$\Delta \equiv \left(\frac{\sigma(\hat{s})}{\sigma_{\mathrm{comb}}} \right)^2 \geq 1, \tag{7.32}$$

where σ_{comb} is defined in Eq. 4.16.

Moreover, contrary to the noise residual bias, which will cancel out when working with cross-spectra, the foreground residual will consist in a net positive bias[2]. When working on simulations

[2]there is also a subdominant part of the foregrounds residual which leads to statistical noise since the foreground residual implicitly depends on the noise through the weights.

with known inputs, these two terms can easily be estimated and are useful to assess the quality of the method (of course, the same does not hold in real life).

We have stated that HILC degrades the noise and is a biased method. However, let us notice a simple fact that clearly demonstrates the significant power of HILC. The fundamental philosophy of a foreground cleaning method is that the estimated cosmology should be as close as possible, in terms of bias and variance, to the cosmology that would have been estimated in the absence of foregrounds and noise. Since cosmology mainly resides in the power spectrum, this requirement is equivalent to stating that the power spectrum estimate should be as close as possible to the true power spectrum without foreground and noise. Hence, the residuals in ℓ space should be minimal. In Eq. 7.29, $C_\ell^{\rm cmb}$ is independent of the weights. Thus, the weights that minimise the residual spectrum $C_\ell^{\rm resid}$ are the ones that minimise the total variance of the $a_{\ell m}$, which are the HILC weights. In summary, in a linear setup (relating d to $s_{\rm cmb}$ through weights) and when restricting ourselves to assuming nothing beyond the CMB blackbody spectrum, the best blind linear method for recovering cosmology should be HILC by construction.

That being said, the devil is in the details. Two degrees of freedom remain in the method, N_m and σ_{w_ℓ}, which should be tuned to each specific experiment and scientific target in order to decrease the HILC bias and degradation factor as low as possible, based on some criterion that handles the bias-variance trade-off. In the next section, we explore the optimisation of this method and produce results for simulations in the context of the LiteBIRD large-scale polarisation recovery.

7.3. HILC for LiteBIRD: setup and products

The aim of this section is to apply and optimise the HILC method to the LiteBIRD simulations introduced in the preliminary chapter of Part. III. First, we introduce the simulations we use before providing illustrations of products obtained with HILC, weights, maps, and power spectra. Even if not strictly speaking directly related to foreground cleaning, we then present the different masks needed for the following discussions.

7.3.1. Inputs

As described in the preliminary chapter of Part. III, our input frequency maps contain CMB, foregrounds and noise with a resolution of $N_{\rm side} = 512$, for the three Stokes parameters and the twenty-two LiteBIRD frequency channels. These maps come with a Gaussian beam in pixel space whose angular size is given in Tab. 6.1 and are affected by the pixel window effect corresponding to their $N_{\rm side}$ described in Sec. 2.3.3. We compute the full-sky $a_{\ell m}^E$ and $a_{\ell m}^B$ with the usual Eq. 2.52 (the full-sky assumption will be relaxed and discussed in Sec. 7.4.2). Then, we correct the pixel beam by dividing $a_{\ell m}$ by the corresponding window function. Hence, our weights are estimated at infinite resolution. In a real-data scenario, the picture will be more complex because the window functions cannot be perfectly known. This is beyond the scope of this work.

In the three next sections, we present the HILC products for specific N_m and σ_{w_ℓ}, namely $N_m = 1000$ and $\sigma_{w_\ell} = 30$, for illustration. We let the actual optimisation of these degrees of freedom to Sec. 7.4.3. As a first step, we present the results for the 100 "wn + baseline" simulations (`d1s1`). The HILC products are the estimated weights, maps or power spectra.

7.3.2. Binned and smoothed weights

On the left panel of Fig. 7.1, we illustrate the E-mode binned weights (faint dashed lines) averaged over the simulations for each LiteBIRD frequency channel, superimposed with the associated smoothed weights (plain lines). We illustrate the B-mode smoothed weights on the right panel, whose contour represents the region where 68% of the simulations lie.

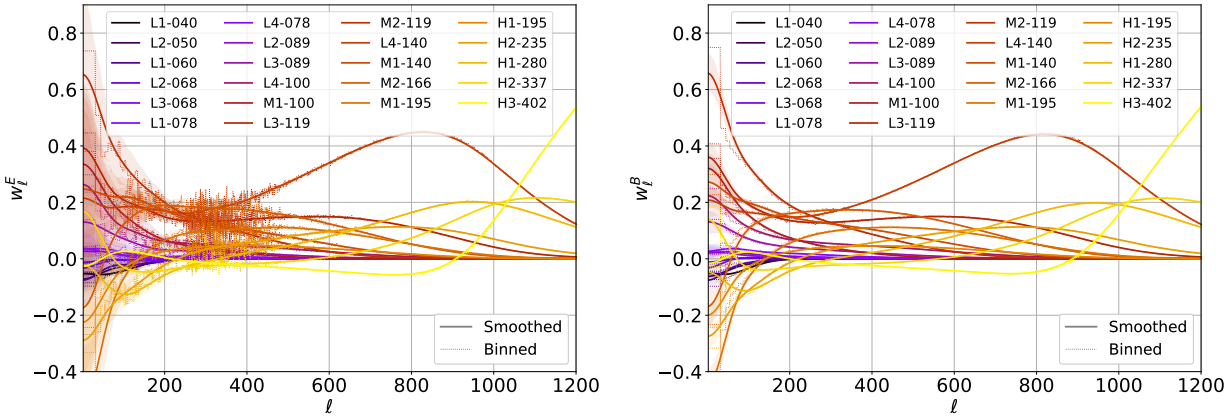

Figure 7.1: ***Preliminary!*** Frequency-per-frequency weights averaged over 100 simulations, each estimated from full-sky CMB + d1s1 + white noise maps. LEFT: for E-modes. RIGHT: for B-modes. The binned weights are represented by the dashed lines ($N_m = 1000$ and $\sigma_{w_\ell} = 0$), while the smoothed weights are in plain lines ($N_m = 1000$ and $\sigma_{w_\ell} = 30$). We also superimpose with their 68% confidence limit estimated over a hundred simulations (hence, it includes cosmic variance).

As previously emphasised, frequency-per-frequency, the 22 curves sum up to one at each scale. The dispersion of the weights is lower at small scales where the signal cosmic variance is lower. For the same reason, the weights are more stable over the simulations for B-modes, given that the expected ratio foregrounds over CMB is higher for B-modes than for E-modes.

The smoothing of the weights that I have introduced allows the erasing of unphysical discontinuities in ℓ space. From the right panel of Fig. 7.1, one can already anticipate that the arbitrary chosen (N_m, σ_{w_ℓ}) is slightly suboptimal at largest scales. For instance, the smoothed curve corresponding to the channel dominating at low multipoles (the one such that $w^B_{\ell=2} \simeq 0.65$), M2-119, clearly deviates from its binned weights, due to a too severe σ_{w_ℓ}.

We also see that HILC operates in two distinct regimes. Generally, the frequency-per-frequency $a_{\ell m}$ are foreground-dominated at the largest scales and noise-dominated at the smallest scales.

- At the largest scales, the weights corresponding to M2-119 are the largest, which can be explained because it is the channel where the CMB signal represents the largest part of the budget over foregrounds and noise. The other positive curves mainly correspond to channels where CMB is competitive with noise and foregrounds. Conversely, the lowest weights at large scales correspond to M1-195 and H2-235 frequency channels, which are low-noise channels that are dominated by dust and hence subtracted in the HILC.
- The highest-frequency channels are associated with the smallest beam FWHM (see Tab. 6.1), so at small scales, the noise dominates first in the low-frequency channels (see Fig. III.3).

Eventually, at the highest multipoles, H3-402 ends up being entirely selected by HILC over the other channels (and we are in the case where the noise residual completely follows the H3-402 noise power spectrum).

7.3.3. Output maps

From these weights, one can build the associated cleaned, foreground residual and noise residual maps with Eqs. 7.18 and 2.54. In Fig. 7.2, we show each of the E- and B-mode maps for a particular simulation. We also display the P maps, where $P = \sqrt{Q^2 + U^2}$. All three are smoothed with a 40 arcmin Gaussian beam. The E-modes residuals are subdominant with respect to the cleaned CMB. However, this is not the case for B-modes, and the recovered map is strongly contaminated by both noise and foregrounds. We see that the foreground residuals after HILC cleaning are very localised around the galactic plane.

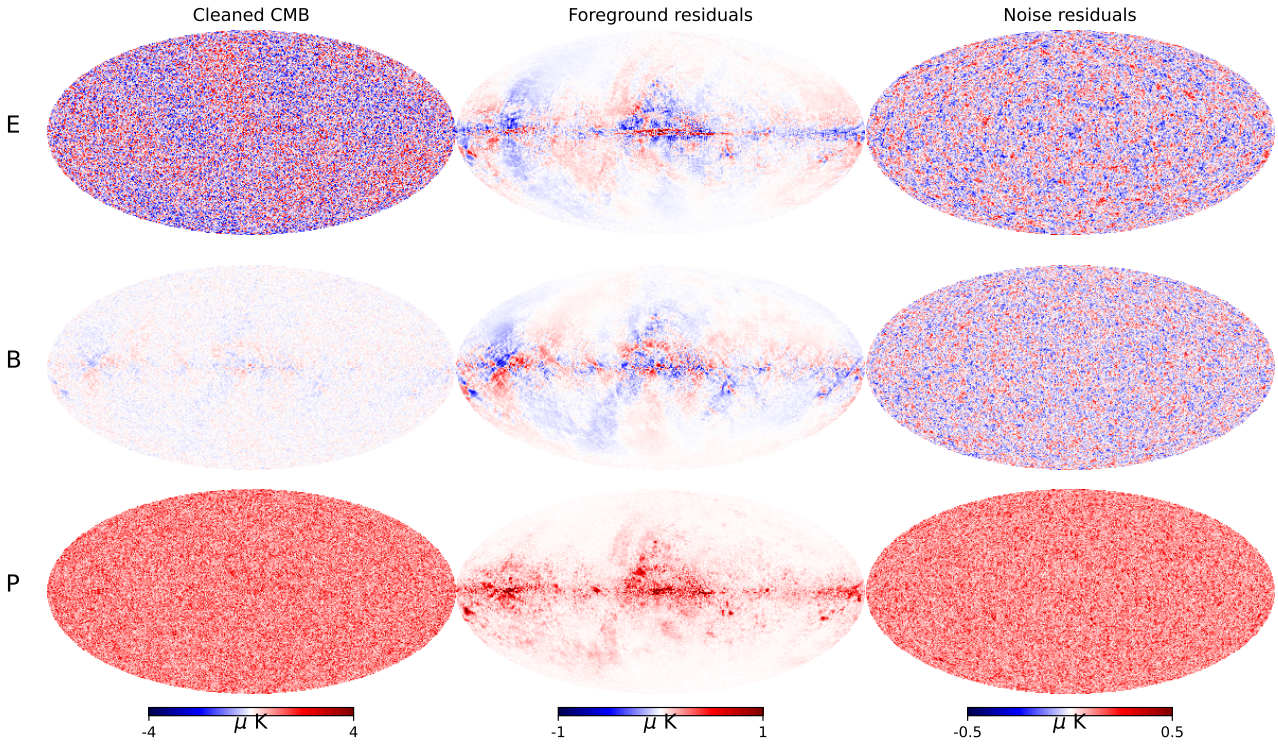

Figure 7.2: ***Preliminary!*** TOP: E-modes MIDDLE: B-modes and BOTTOM: total P maps after HILC for a particular simulation. LEFT: Cleaned CMB, MIDDLE: foreground residuals, RIGHT: noise residuals. The maps are produced with a 40 arcminutes beam.

7.3.4. Output power spectra

We represent in Fig. 7.3 the associated full-sky cleaned CMB, foreground and noise auto power spectrum residuals, which can be estimated from the weights and $a^{\nu}_{\ell m}$ with Eq. 2.47. The plain lines represent the power spectra obtained with the smoothed weights, while the spectra obtained with the binned weights are represented in dashed lines. The smoothing of the weights leads to two benefits. It allows us to recover continuous spectra, and to reduce the bin size without increasing weight variance. Consequently, it allows us to gain on the residuals. The investigation and exploitation of this gain will be the topic of Sec. 7.4.3.

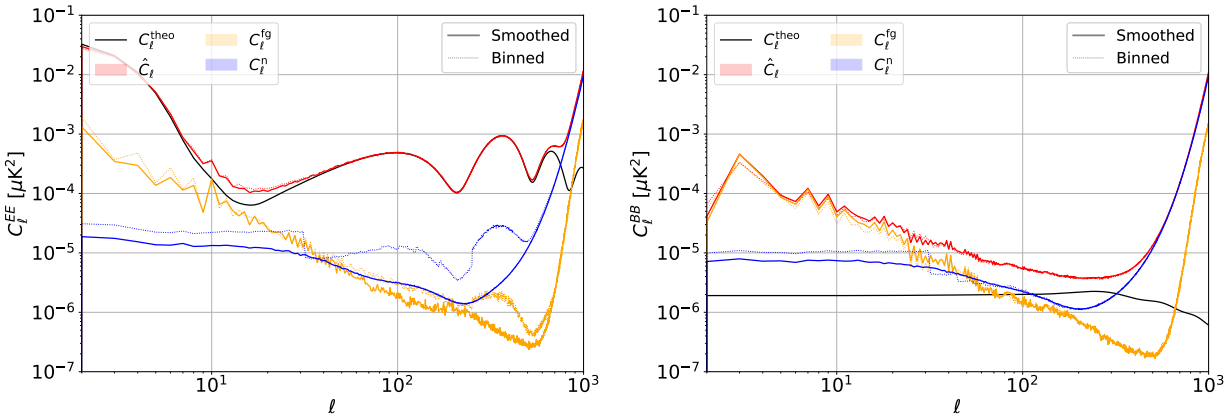

Figure 7.3: ***Preliminary!*** Cleaned CMB (red), foreground (orange), and noise (blue) power spectrum residuals estimated full-sky. LEFT: *EE*, RIGHT: *BB* ($r = 0$). The plain lines represent the power spectra obtained from the smoothed weights, on which we superimpose the contour in which lies 68% of the simulations, while the spectra obtained from the binned weights are represented in dashed lines.

Notice that the red curve includes the noise bias since it represents an auto power spectrum. As mentioned in Sec. 7.2, this will be eventually cancelled out by working with cross-power spectra. Hence, one should view the orange plus blue curves as a total N_ℓ contributing to the variance of the recovered estimate alongside the CMB signal variance itself, while the orange curve consists of a net bias for the estimated signal.

Fig. 7.3 allows us to derive first conclusions about the performance of HILC on full-sky for this particular configuration.

- For *EE* and in the case of `d1s1` foregrounds, the two full-sky residual power spectra are always lower than the CMB signal. Hence, in terms of variance, one expects that the recovered *EE* signal will be close to cosmic-variance limited on the whole multipole range, except for the region between multipoles of approximately 6 to 40, which remains affected by the foreground residuals. The other consequence is a non-negligible bias in the same region of multipoles that we will quantify in the next sections.

- In the right panel, one observes that the foreground residuals are above the CMB *BB* signal for $\ell < 80$. Hence, in this setup, an almost unbiased full-sky measurement of the lensed *BB* signal is only possible to achieve for multipoles higher than approximately 100. However, this conclusion will be improved by masking the galactic region. This measurement comes with a high associated variance, the noise residual being only slightly lower than the signal for $80 < \ell < 310$, and exploding above the signal for $\ell > 80$.

In order to reject the regions mostly contaminated by foregrounds, we might want to mask the sky at two distinct steps of the analysis. 1) Mask the frequency maps to estimate weights on a partial sky. 2) Mask the output map after component separation to estimate the power spectrum. First, let us introduce two categories of masks that can be used for either of these purposes.

7.3.5. Two classes of mask

This section aims to introduce the masks we will use for assessing the performance of HILC. To build the first class of masks, we consider the input Q and U maps at 402 GHz for each realisation, in which the dust dominates. Indeed, it is the polarised foreground with the highest signal to noise at high frequency, and which traces well the galactic structures, see Sec. 4.1.2. We build the total polarisation amplitude map $P = \sqrt{Q^2 + U^2}$, which we smooth by a θ_1 Gaussian beam. We chose a target sky fraction $f_{\rm sky}$ that will remain after the application of the mask (see Sec. 4.1.2) and set the $f_{\rm sky}$ per cent brightest pixels to 0 and the other ones to 1. Finally, we apodise the obtained binary mask by θ_2. This mask is *realistic* in the sense that it is built upon inputs that will be available in a real-case scenario. An instance of this mask, with $\theta_1 = \theta_2 = 2^\circ$ and built from a given simulation, is shown on the left panel Fig. 7.4.

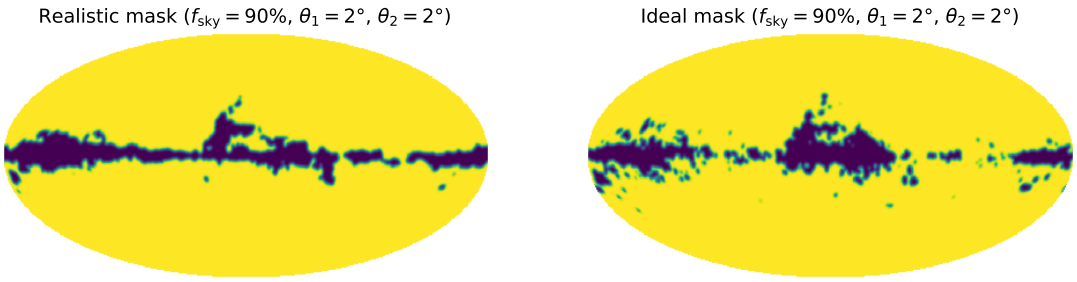

Figure 7.4: ***Preliminary!*** For $\theta_1 = \theta_2 = 2^\circ$, $f_{\rm sky} = 90\%$, and a given simulation: LEFT: realistic mask based on 402 GHz P input, RIGHT: ideal mask based on foreground residual P output.

The second class of masks, which we call *ideal*, is designed to mask the sky after component separation and is optimistic and *a priori* non-realistic. It is constructed from the total polarisation amplitude map of the foreground residuals P, which is only available in a controlled setup. Similar to the first class of masks, we smooth P by a θ_1 Gaussian beam before creating a binary mask from it, which is then apodised by θ_2. An example of this mask, built from the same simulation, is shown on the right-hand side of Fig. 7.4. When not precised otherwise, we will take hereafter $\theta_1 = \theta_2 = 2^\circ$, a conservative choice that is further commented in Sec. 7.4.2.

7.4. Optimisation and results

Now that we have presented our working setup and the HILC products we can derive, we discuss how to we have optimised the method in order to propose a HILC configuration that minimises the output variance and bias in spectrum space, allowing for optimal estimation of τ and r. Then, we assess the performance of HILC on foregrounds of varying complexity.

7.4.1. Masking input maps

First, we discuss the impact of masking the input frequency maps before the weights' estimation.

In Fig. 7.5, we compare the HILC residuals when various sky fraction configuration are applied to the input frequency maps (prior weight processing), for the baseline setup already discussed in Sec. 7.3. Then, the foreground (*resp.* noise) residuals are shown in plain (*resp.* dashed) lines

and are estimated on full-sky HILC maps. The foreground residuals increase dramatically when we estimate the weights on cut-sky maps. Instead, one should maximise the fraction of the map kept in order to get as much information as possible on the foreground cleaning, increasing the accuracy of the weights estimation and reducing the resulting residuals in the HILC output. This conclusion also justifies our legitimate use of the simple covariance estimator Eq. 7.19, which is optimal and unbiased as long as we remain full-sky. Henceforth, we will use full-sky frequency maps to determine the weights. However, it is important to notice that working with full-sky frequency maps will be impractical with real data due to the high point-source concentration on the Galactic plane (not taken into account in our input simulations), necessitating specific masking.

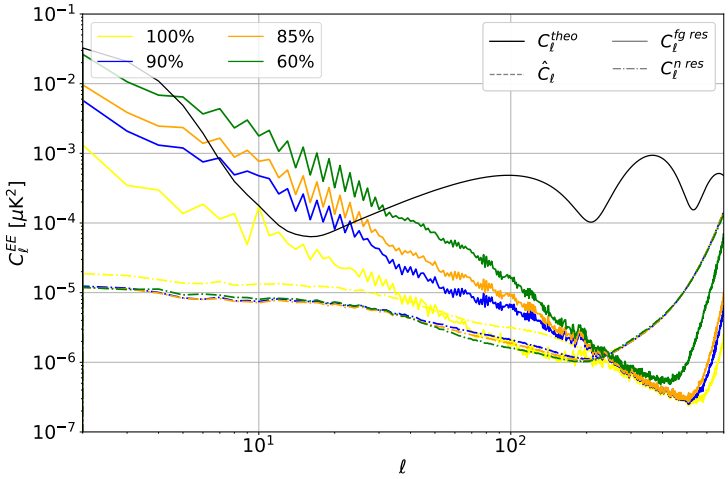

Figure 7.5: ***Preliminary!*** Foregrounds (plain) and noise (dashed) residuals power spectrum estimated on full-sky maps, themselves obtained from weights that have been estimated excluding the most polarised regions of the sky in input frequency maps. We also represent the theoretical CMB *EE* spectrum in plain black line.

7.4.2. Masking output maps

Instead of masking frequency maps, we illustrate in this section the impact of masking cleaned output maps to avoid highly foreground-contaminated regions.

Fig. 7.2 demonstrated that the foreground residuals are highly localised around the Galactic plane. Although masking the foreground-dominated region increases the variance of the spectra and the correlations between the modes by decreasing the sky fraction, it also reduces the bias and variance associated with the foreground residuals. We use the specific HILC setup introduced earlier (without masking the input frequency maps) to demonstrate this. We derive a set of weights and produce full-sky Q and U maps, which are then masked with the ideal mask before estimating the anafast power spectrum (which is rescaled by $f_{\rm sky}$ to approximately account for the loss of modes). The *EE* and *BB* residuals are presented in Fig. 7.6, where the averaged foreground residual is shown in solid lines and the averaged noise residual in dashed-dotted lines. The two curves are shown for several sky fractions between 50% and 100%.

The dashed curve represents the HILC-recovered CMB signal, which should not be compared directly to the input when $f_{\rm sky} < 100\%$. The anafast estimation of this curve mixes the *E*- and

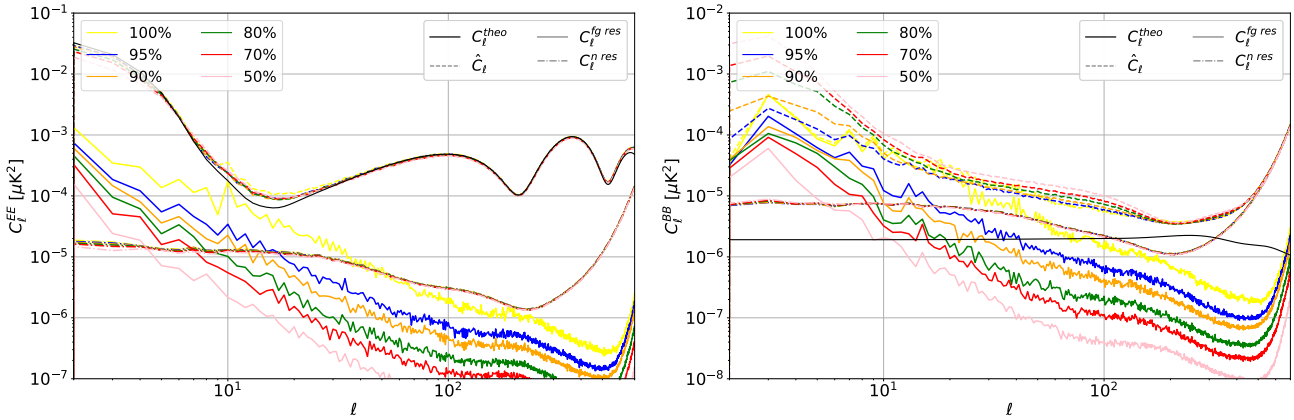

Figure 7.6: ***Preliminary!*** Cleaned CMB (dashed), foreground (plain) and noise (dashed-dotted) residuals power spectra for several sky fractions LEFT: *EE* RIGHT: *BB*.

B-modes, leading to biased spectra. It is only used here for a fast relative comparison of the residuals. Notably, Fig. 7.6 shows a loss of *EE* power and a gain of *BB* power at small scales. Instead, we will address the reconstruction of the signal power spectrum on cut-sky only in Chapter 8.

Nonetheless, even on cut-sky, one can roughly estimate the quality of HILC by comparing the amplitude of the two residual curves (the noise residual in dashed-doted and the foreground residual in plain line) with respect to the input using this simple spectrum estimation.

One observes that the noise residual remains at the same level, while the foreground residual significantly decreases as the sky fraction is reduced. Especially, one sees that masking 5% of the sky allows one to go from the yellow to the blue foreground residual curve, hence subsequently decreasing the bias. A first assessment of the implied impact on cosmology extraction will be done in Sec. 7.4.3, before entering in the details in Chapters 8 and 9.

We also compare on the left-hand side Fig. 7.7 the effect of masking Q and U or E and B for the residuals with the ideal or realistic masks. For both mask classes, we observe that the optimal strategy to reject as much foreground residual as possible is to mask Q and U rather than E and B (the blue curve is below the red one and the yellow one below the orange one). The foregrounds leak across the entire E and B sky, whereas they are more localised and hence easier to mask in Q and U. Consequently, the power spectrum estimation will have to correct for the E to B leakage that occurs when working with masked Q and U maps, a topic we will detail in Chapter 8. Moreover, as expected, the ideal mask is better at rejecting foreground residuals than the realistic mask.

Finally, we illustrate on the right-hand side of Fig. 7.7 the impact of θ_1 and θ_2 on the residuals for the ideal mask. The more we smooth the mask, the higher the foreground residuals. Moreover, we emphasise that this dependence on the smoothing angles only affects the most optimistic and ideal masks, and *a fortiori* the ones with small features. Indeed, such an effect is not visible at all for our realistic mask. For this reason, and since smoothing is mandatory in a realistic power spectrum estimation scenario, we stick to the conservative choice of 2° for both θ_1 and θ_2 in this work. With better masks that get closer to the ideal one (that would present more resolved features), reducing the smoothing angle would be a good asset to further reduce

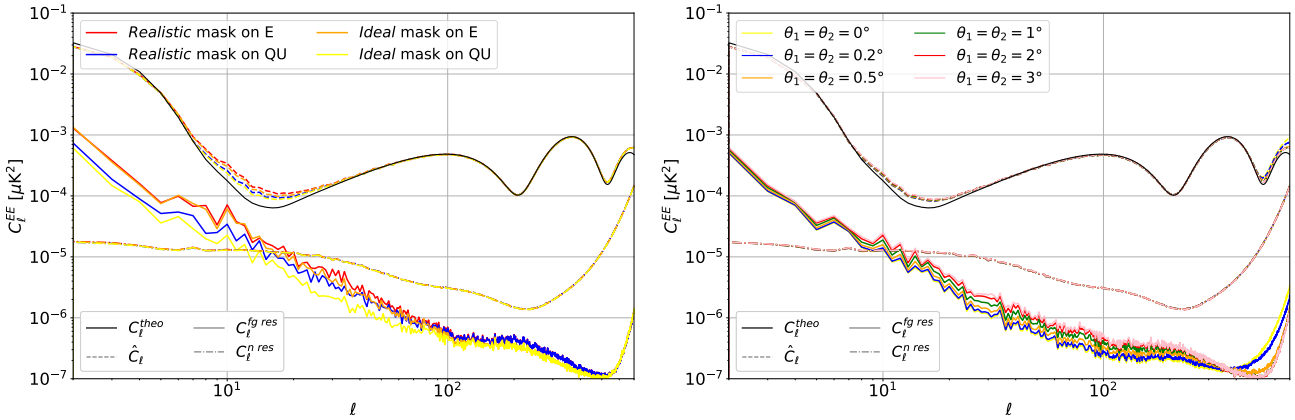

Figure 7.7: ***Preliminary!*** Cleaned *EE* CMB (dashed), foreground (plain) and noise (dashed-dotted) residuals power spectra for various masking strategies and $f_{\rm sky} = 90\%$. LEFT: masking Q & U or E with ideal or realistic masks, RIGHT: impact of the free smoothing angles in the mask definition.

the residuals.

Given its significant impact on the residual levels, we will leave the sky fraction used to estimate the power spectra free. Otherwise, we choose to use the optimistic ideal mask coupled with the pessimistic 2° smoothing, which we will apply to the output Q and U maps.

7.4.3. Optimising binning and smoothing

This discussion is dedicated to optimising the HILC degrees of freedom, N_m and σ_{w_ℓ}, while keeping the sky fraction free for masking output maps. To do so, we look for a criterion in power-spectrum space that traces the quality of the estimation of the parameter of interest without going (for now) into the proper power spectrum estimation and maximum-likelihood estimator of the parameter in question.

We base this discussion on the Fisher formalism, Eq. 4.30. It does not encapsulate all the complexity of a proper inference procedure; we will test proper C_ℓ likelihoods in Chapter 9. This approach is, however, sufficient for the exercise that follows.

7.4.3.1 *E*-modes HILC optimisation for τ

First, we aim at reconstructing τ from the *E*- modes reionisation bump. For the exercise, we fix all the other parameters to their fiducial values (*cf.* Table 4.1): these are known below the per cent level from other observables than large-scale *E*- modes. As detailed and explained in Sec. 4.4.2.1, we fix $A_{\rm s}\exp\{-2\tau\}$ instead of $A_{\rm s}$. The region where we should optimise the residuals as possible is the one where the information is maximal, *i.e.* $5 < \ell < 18$ according to Fig. 4.10. With Fisher, this multipole weighting is directly handled by the derivative of C_ℓ with respect to the parameter. From Eq. 4.30, one can reduce the Fisher $\tau\tau$ matrix element to

$$F_{\tau\tau} = \sum_{\ell=2}^{26} \frac{(2\ell+1) f_{\rm sky}}{2\left(C_\ell^{EE,\rm CMB} + C_\ell^{EE,\rm fg} + C_\ell^{EE,\rm n}\right)^2} \left(\frac{\partial C_\ell^{EE,\rm CMB}}{\partial \tau}\right)^2 , \tag{7.33}$$

where the notations have been already defined earlier. Only the reionisation bump contributes to this summation, so we bound it to $2 \leq \ell \leq 26$.

The statistical error is simply given by

$$\sigma^{\rm Fisher}_{{\rm stat},\tau} = F^{-1/2}_{\tau\tau}, \tag{7.34}$$

which follows from the definition of the Fisher matrix, while a systematic error criterion valid for small residual systematics [144], Eq. 4.32, which we have introduced in Sec. 4.3.2, is given by

$$\sigma^{\rm Fisher}_{{\rm syst},\tau} = F^{-1}_{\tau\tau} \sum_{\ell=2}^{26} \frac{(2\ell+1) f_{\rm sky}}{2\left(C_\ell^{EE,{\rm CMB}} + C_\ell^{EE,{\rm fg}} + C_\ell^{EE,{\rm n}}\right)^2} \, C_\ell^{EE,{\rm fg}} \, \frac{\partial C_\ell^{EE,{\rm CMB}}}{\partial \tau}, \tag{7.35}$$

Notice that since for EE, $C_\ell^{\rm fg}$ is a positive quantity, if we want to estimate the bias on the parameter i, $\sigma^{\rm Fisher}_{i,{\rm syst}}$ has the sign of $\frac{\partial C_\ell^{\rm CMB}}{\partial i}$ (positive for τ in reionisation bump). Hence, in all generality, τ is biased up by component separation methods.

To obtain the cosmic-variance estimate of $\sigma^{\rm Fisher}_{{\rm syst},\tau}$, one can apply this formula with $C_\ell^{EE,{\rm fg}} = C_\ell^{EE,{\rm n}} = 0$. This leads to $\sigma^{{\rm Fisher},\ {\rm CV}}_{{\rm syst},\tau} = 2 \times 10^{-3}$, which will serve of reference later on.

To reduce systematic and statistical error to a single criterion, one can sum the two, $\delta^{\rm Fisher} \equiv \sigma^{\rm Fisher}_{\rm stat} + \sigma^{\rm Fisher}_{\rm syst}$. This criterion indicates the total uncertainty one will have after having measured the parameter of interest, but it should not be overinterpreted since the two types of errors are of a distinct nature (one is statistical, and the other is a positive bias). Another criterion that we will discuss consists in ensuring that the systematic error is lower than the statistical one. Keeping the entire sky for the spectrum estimation does not allow us to fulfil this criterion. Hence, we also explore cut-sky configurations.

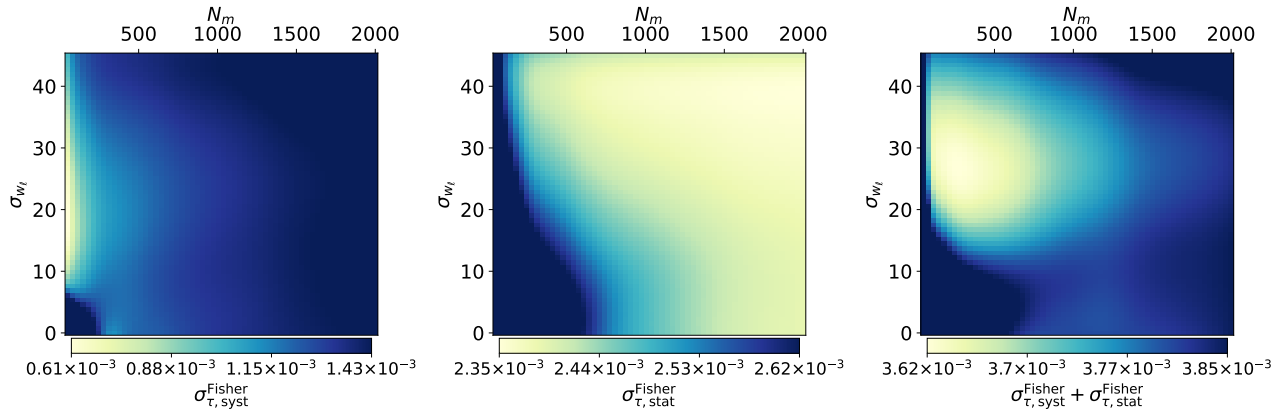

Figure 7.8: ***Preliminary!*** Rough estimation of the statistical and systematic error on τ, $\sigma^{\rm Fisher}_{\tau,{\rm stat}}$ and $\sigma^{\rm Fisher}_{\tau,{\rm syst}}$, from the residuals in the two dimensional space $(N_m,\, \sigma_{w_\ell})$, for 85% of the sky.

Fig. 7.8 illustrates the Fisher systematic, statistical and sum of both for the τ reconstruction, in a $(N_m,\, \sigma_{w_\ell})$ plan for 85% of the sky. To obtain it, we run HILC with 130 distinct configurations that span this 2D space. Then, we do a cubic interpolation between the points of the original grid. One sees that the optimal region for $\sigma^{\rm Fisher}_{\rm syst}$ does not coincide with the optimal region for $\sigma^{\rm Fisher}_{\rm stat}$. Given our limited-extent grid and this particular $f_{\rm sky}$, the optimal configuration for

minimising the systematic error is $(N_m, \sigma_{w_\ell}) = (50, 18)$, giving a systematic error of 6×10^{-4} and a statistical one of 36×10^{-4}. The optimal configuration for minimising the statistical error is $(N_m, \sigma_{w_\ell}) = (2000, 39)$, giving a systematic error of 24×10^{-4} and a statistical one of 15×10^{-4}. This justifies the need for a criterion combining both like their sum. The latter indicates an optimal configuration at $(N_m, \sigma_{w_\ell}) = (281, 27)$, yielding a systematic error of 25×10^{-4} and a statistical one of 11×10^{-4}.

One observes in Fig. 7.8 that for 85% of the sky, small bin configurations are interesting for τ. They especially allow a significant decrease in the bias on the parameter estimate. To understand this behaviour, we represent in Fig. 7.9 the full-sky noise residuals in dashed lines and the full-sky foreground residuals in plain lines. The colours correspond to different N_m, with σ_{w_ℓ} being fixed to 20. It illustrates the point of small bin configurations: the smaller the bin size is, the smaller the foreground residuals and the resulting systematic error on τ. Small bins allow for a more precise foreground cleaning but at the price of the noise residual. However, the latter only becomes impactful for the statistical error on τ when it gets larger than the foreground residual.

This leads small-bin configurations to be favoured at large sky fraction, where the foreground residual significantly exceed the noise residual. Consequently, we also expect an evolution of the best configuration with $f_{\rm sky}$: the lower the sky fraction, the closer the noise residual is from the foreground residual, and the less the bins need to be small. We verify this hypothesis with the blue path of Fig. 7.9, which illustrates how the best E-modes-HILC configuration evolves in the $(N_m,\ \sigma_{w_\ell})$ space. At 50%, the best configuration is located at $(975, 38)$, with a systematic uncertainty 33×10^{-4} and a statistical one of 3×10^{-4}, and it evolves towards $(50, 21)$ for the full-sky configuration, which corresponds to a systematic uncertainty of 32×10^{-4} and a statistical uncertainty of 35×10^{-4}. The conclusion is that the best HILC configuration for maximising the τ information strongly depends on the sky fraction. The final choice of configuration will be made in Sec. 7.4.3.3.

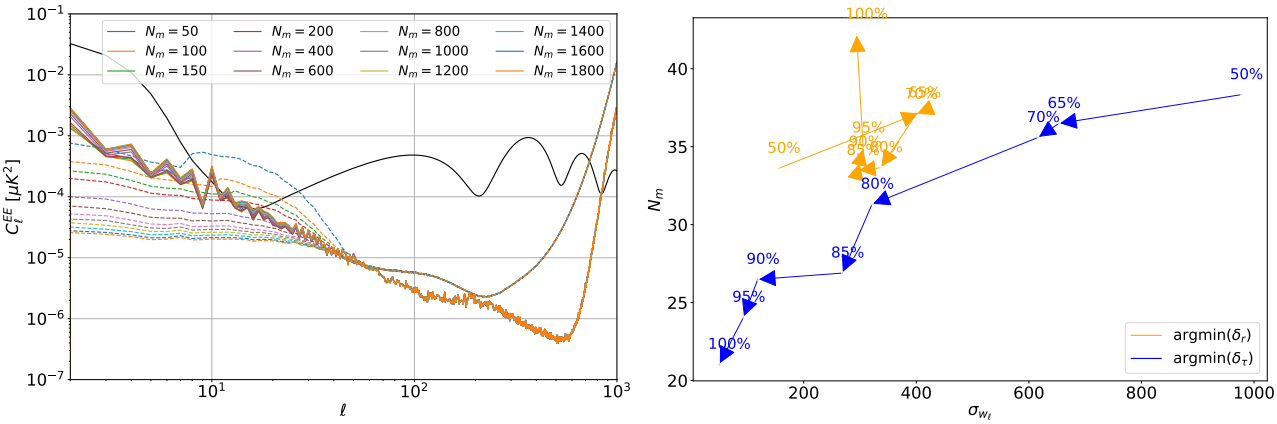

Figure 7.9: ***Preliminary!*** LEFT: Impact of N_m (the bin size) on the level of EE noise residuals in full-sky (dashed lines) and on the level of EE foreground residuals in full-sky (plain lines). σ_{w_ℓ} is fixed to 20. In black, we represent the EE fiducial power spectrum. RIGHT: Evolution with $f_{\rm sky}$ of the configuration that minimises the total uncertainty δ in the (N_m, σ_{w_ℓ}) plan.

7.4.3.2 B-modes HILC optimisation for r

In the same way, we estimate the Fisher statistical and systematic error on r from the B-modes reionisation and recombination bumps including the HILC residuals. Once again, we fix all the Λ-CDM parameters. According to Fig. 4.11, the region where we should optimise as much as possible the residuals is the one where this information curve is maximum: approximately 2 to 200 depending on the fiducial r value. Eqs. 7.33, 7.34 and 7.35 can also been applied to this new problem by replacing $C_\ell^{EE} \to C_\ell^{BB}$ and $\tau \to r$ and with a wider ℓ range that we take to be $2 \leq \ell \leq 200$.

Fig. 7.10 illustrates the Fisher systematic, statistical and sum of both for the r reconstruction, in a (N_m, σ_{w_ℓ}) plan for 50% of the sky. The computation of these arrays is done similarly to the τ-from-E-modes' ones. Given the limited extent of our grid and this particular $f_{\rm sky} = 50\%$, the optimal configuration for minimising the systematic error is $(N_m, \sigma_{w_\ell}) = (50, 8)$, giving a systematic error equal to the statistical one, 24×10^{-4}. The optimal configuration for minimising the statistical error instead is $(N_m, \sigma_{w_\ell}) = (678, 38)$, giving a systematic error of 32×10^{-4} and a statistical one of 17×10^{-4}. Once again, we use the total uncertainty to find an optimal configuration at $(N_m, \sigma_{w_\ell}) = (149, 34)$, yielding a systematic error of 28×10^{-4} and a statistical one of 20×10^{-4}. In our particular setup, this HILC preliminary uncertainty estimate is almost five times the LiteBIRD requirement, which has been set using parametric component-separation methods. In the next sections, we will investigate whether this result is confirmed after power spectrum estimation and likelihood analysis.

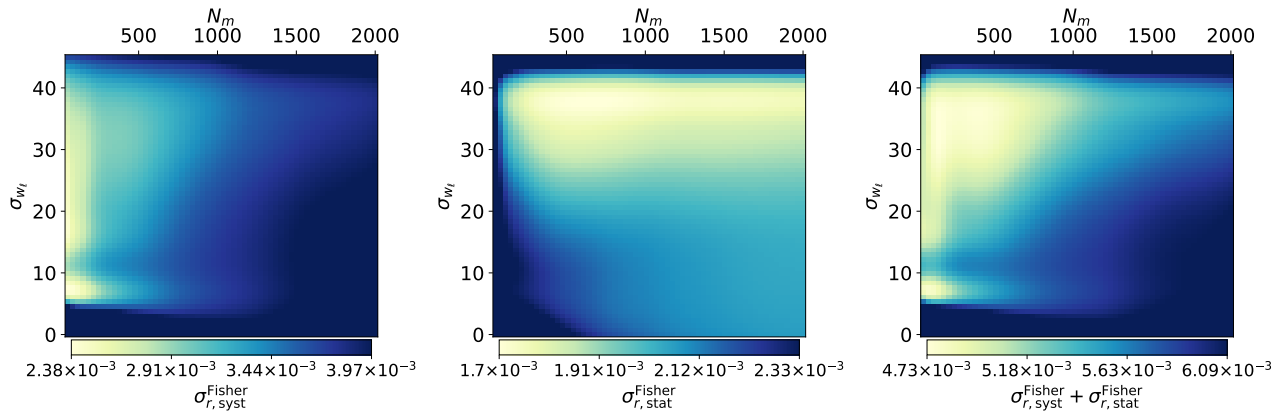

Figure 7.10: ***Preliminary!*** Rough estimation of the statistical and systematic error on r, $\sigma^{\rm Fisher}_{r,\rm stat}$ and $\sigma^{\rm Fisher}_{r,\rm syst}$ and the sum of the two, in the two dimensional space (N_m, σ_{w_ℓ}), for 50% of the sky. The fiducial spectrum that is also the input of the simulations assumes $r = 0$.

We illustrate with the orange path of Fig. 1.26 the evolution of the best B-HILC configuration for r in the (N_m, σ_{w_ℓ}) space. It starts from the optimal configuration at 50% that we already discussed, and ends at (N_m, σ_{w_ℓ}) = (282, 42) that yields a systematic uncertainty of 281×10^{-4} and a statistical one of 17×10^{-4}. One observes that this configuration stays more localised than for E- modes, and does not leave the region delimited by $200 \leq \sigma_{w_\ell} \leq 400$ and $32 \leq N_m \leq 42$.

7.4.3.3 Configuration choice

We could select the minimum total error configuration to fix a N_m and a σ_{w_ℓ} for the following. For r, it would lead us to choosing $N_m = 149$, $\sigma_{w_\ell} = 34$ at 50% (giving $\delta_r^{\text{Fisher}} = 47 \times 10^{-4}$), and for τ $N_m = 645$, $\sigma_{w_\ell} = 37$ at 65%, yielding $\delta_\tau^{\text{Fisher}} = 31 \times 10^{-4}$ (5×10^{-4} of systematics and 26×10^{-4} of statistics). If one wants to let free f_{sky}, at this stage, we could also select optimal HILC configurations at each f_{sky}, and keep going with that.

Instead, for this work, for simplicity, we chose to only select one configuration to pursue the analysis with through the next parts. In Fig. 7.11, we illustrate the evolution of the total uncertainties for the various configurations, normalised by the best configuration total uncertainty. For each sky fraction, the best configuration curve is at 1. In blue, we emphasise the configuration we used as a reference so far. We see that there exist better choices both for r or for τ. Once again, for simplicity reasons and even though it is not strictly needed, we prefer a configuration that can accommodate both τ and r. Some good options appear in Fig. 7.11, and we choose to work with the yellow configuration, $N_m = 400$ and $\sigma_{w_\ell} = 35$, which is optimal for τ around 80% and close to optimal for r for all sky fractions.

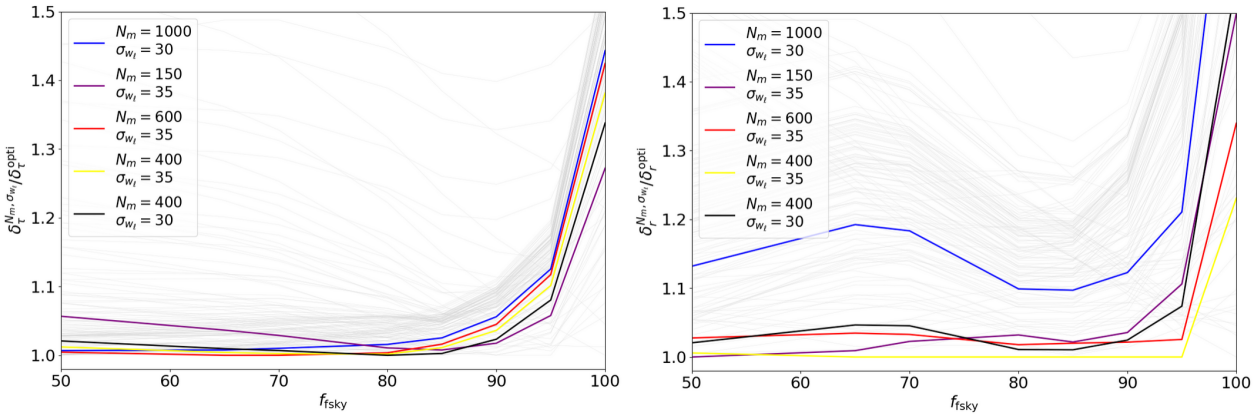

Figure 7.11: ***Preliminary!*** Total uncertainty on LEFT: τ RIGHT: r for various f_{sky}, normalised by the total uncertainty of the best configuration.

In this section and for the rest of the part, we mainly focus on the reconstruction of τ and r from *EE* and *BB* spectra with this particular configuration, $N_m = 400$ and $\sigma_{w_\ell} = 35$. We have also applied HILC to the temperature maps, which yields excellent results. All six HILC power spectra are ready to be analysed for the whole Λ-CDM + r, a task that is left for a future work.

7.4.4. More complex foregrounds and noise

To understand what happens when changing the noise or foreground type in the input simulations, we use the simulations described in the preliminary chapter of Part. III, "wn1f + baseline", "wn1f + medium", "wn1f + high", which include TOD-based noise simulation and more complex foregrounds sky model. We compare the results with the "wn + baseline" case studied so far. Fig. 7.12 shows the residuals power spectra when varying the foreground model and the noise model, for *EE* ($f_{\text{sky}} = 90\%$) and *BB* ($f_{\text{sky}} = 60\%$). The green curves illustrate the residuals in the "wn + baseline" case, while the orange, blue and red curves show the residuals for "wn1f + baseline", "wn1f + medium" and "wn1f + high" (see the preliminary chapter of Part. III).

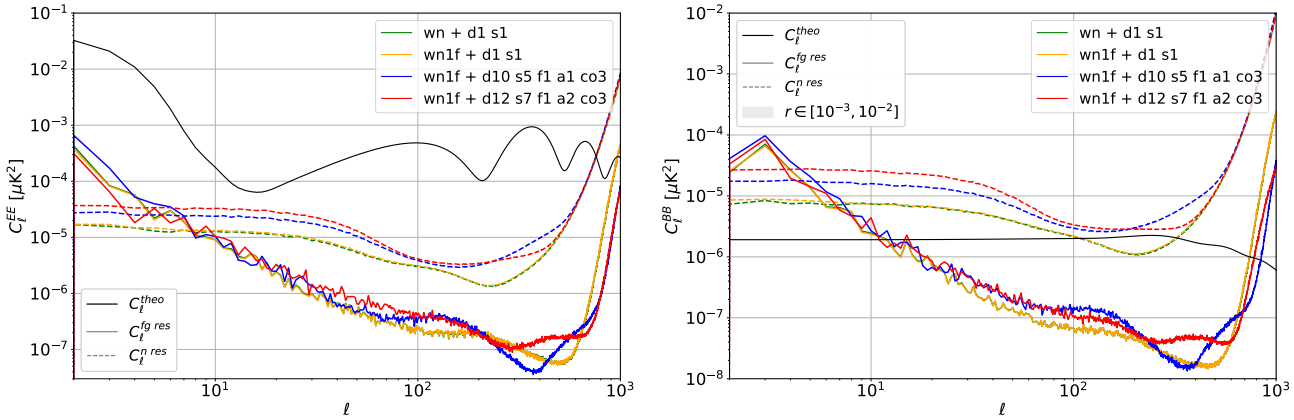

Figure 7.12: ***Preliminary!*** Comparison of the foregrounds (plain lines) and the noise (dashed lines) residuals for several foreground models and two noise approximations. LEFT: *EE* residuals for $f_{\rm sky} = 90\%$. RIGHT: *BB* residuals for $f_{\rm sky} = 60\%$. Even though our fiducial is $r = 0$, we display for reference the $10^{-3} < r < 10^{-2}$ theoretical primordial *BB* spectrum in lightgrey.

The dashed lines in Fig. 7.12 inform on the level of statistical uncertainty for each foregrounds complexity. These curves are almost independent of the sky fraction and of the multipole between $\ell = 2$ and $\ell = 30$. We quote this noise level averaged between these two multipoles in the two first lines of Tab. 7.1, in μK·arcmin, accompanied by the corresponding degradation factor for *EE* and *BB* defined by Eq. 7.32, which we compute here as the ratio between the square root of the noise residual averaged over $2 \leq \ell \leq 30$ and the LiteBIRD combined polarisation sensitivity of 2.16 μK·arcmin. The foreground and noise curves also translate into systematic and statistical errors for the estimation of the parameters, whose Fisher estimate is also given in Tab. 7.1 for indication and relative comparisons.

The similarity between the orange (wn + d1s1) and the green (wn + $1/f$ + d1s1) set of curves in Fig. 7.12 and between the first and second columns of Tab. 7.1 indicates that including $1/f$ noise to the standard white noise does not have a strong impact on foreground cleaning for *E* and *B*. This is an expected conclusion that we anticipated in Sec. 5.2.2 and in Fig. III.4: the rotating LiteBIRD HWP mitigates the $1/f$ noise in polarisation. Of course, this translates into comparable systematic and statistical errors for the estimation of the parameters, whose Fisher estimates are given in Tab. 7.1. Then, there is an approximate progression in the power spectrum residual amplitude and in the associated Fisher uncertainties, which follows the foreground complexity. The noise residual is especially affected and is the main driver of the increase in $\sigma_{\rm stat}^{\rm Fisher}$ that one can observe both for τ and r. Above $\ell = 200$, the noise residual ends up being the largest in the medium complexity case both for τ and r. The medium complexity case also has the strongest foreground residual for $\ell < 10$ both for τ and r. Fig. III.2 already hinted that medium and high complexity foregrounds are dominant at different multipoles, and we verify that this also leads to residuals exhibiting the same kind of trends.

	wn + baseline	wn1f + baseline	wn1f + medium	wn1f + high
N_ℓ^{EE} (Δ^{EE})	11.47 (5.31)	11.7 (5.42)	15.72 (7.28)	18.21 (8.43)
N_ℓ^{BB} (Δ^{BB})	8.65 (4.01)	8.8 (4.08)	12.69 (5.87)	15.99 (7.40)
$10^3 \times \sigma^{\rm Fisher}_{\tau,\rm syst}$	0.91	0.89	0.89	0.98
$10^3 \times \sigma^{\rm Fisher}_{\tau,\rm stat}$	2.43	2.43	2.54	2.63
$10^3 \times \sigma^{\rm Fisher}_{r,\rm syst}$	4.68	4.75	7.28	6.41
$10^3 \times \sigma^{\rm Fisher}_{r,\rm stat}$	1.80	1.85	2.65	2.98

Table 7.1: ***Preliminary!*** TOP: Large scale noise level in μK·arcmin (average over $2 \leq \ell < 30$) and the corresponding noise degradation factor in parentheses, BOTTOM: Fisher estimation of the systematic and statistical uncertainties on τ and r. The reference setup is $N_m = 1000$, $\sigma_{w_\ell} = 30$, with sky fraction of $f_{\rm sky} = 90\%$ for τ and $f_{\rm sky} = 60\%$ for r.

Conclusion

To conclude this foreground-cleaning section, HILC is a simple method that allows us to minimise the residuals in the reconstructed map. It presents the advantage of requiring no prior knowledge of foreground properties. It allows us to produce cleaned maps and power spectra in the context of LiteBIRD.

Keeping the entire sky, we summarise on the top-left panel of Fig. 7.13 how the *EE* noise and foreground levels in each of the 22 LiteBIRD frequency channels are reduced by HILC, resulting in an *EE* spectrum that is nearly unbiased and cosmic-variance dominated between multipoles 2 and 700. This is qualitatively shown by both residuals being lower than the signal on the whole multipole range[3]. The critical region where the residuals are high enough to potentially imply extra variance or bias on τ lies between $\ell \simeq 8$ and $\ell \simeq 30$. The input simulation for this figure includes white noise and d1s1 foregrounds.

Instead, the bottom-left panel displays the same figure in the case of wn1f + high complexity input simulations. In this configuration, the foreground residual slightly overcomes the signal in the critical region. Though, the situation is not dramatically worsened: with this complexity, only masking the 10% most foreground-contaminated regions in the maps already allows to get a Fisher estimate $\sigma^{\rm Fisher}_{\tau,\rm tot} = 3.6 \times 10^{-3}$, see Table 7.1, which is only 80% above the reionisation-bump cosmic variance of 2×10^{-3}.

The equivalent figures for *BB* are displayed on the right panels of Fig. 7.13, which illustrate that in the full sky, the *BB* spectrum for $r = 0$ is dominated by the foreground residuals at large scales and has a large variance for any multipole. On $f_{\rm sky} = 60\%$ and for the highest foreground complexity, we have shown that the Fisher estimate $\sigma^{\rm Fisher}_{r,\rm tot}$ is of order 10^{-2}. This result is preliminary and will be updated in Sec. 9.

To obtain a definitive answer on the possibility of inferring the cosmological parameters τ and r, as well as to prepare the LiteBIRD data-analysis strategy, we estimate in the next section more rigorously the power spectra for masked sky.

[3] $C_\ell > N_\ell$ in the variance formula Eq. 2.49 which is the definition of cosmic-variance dominated

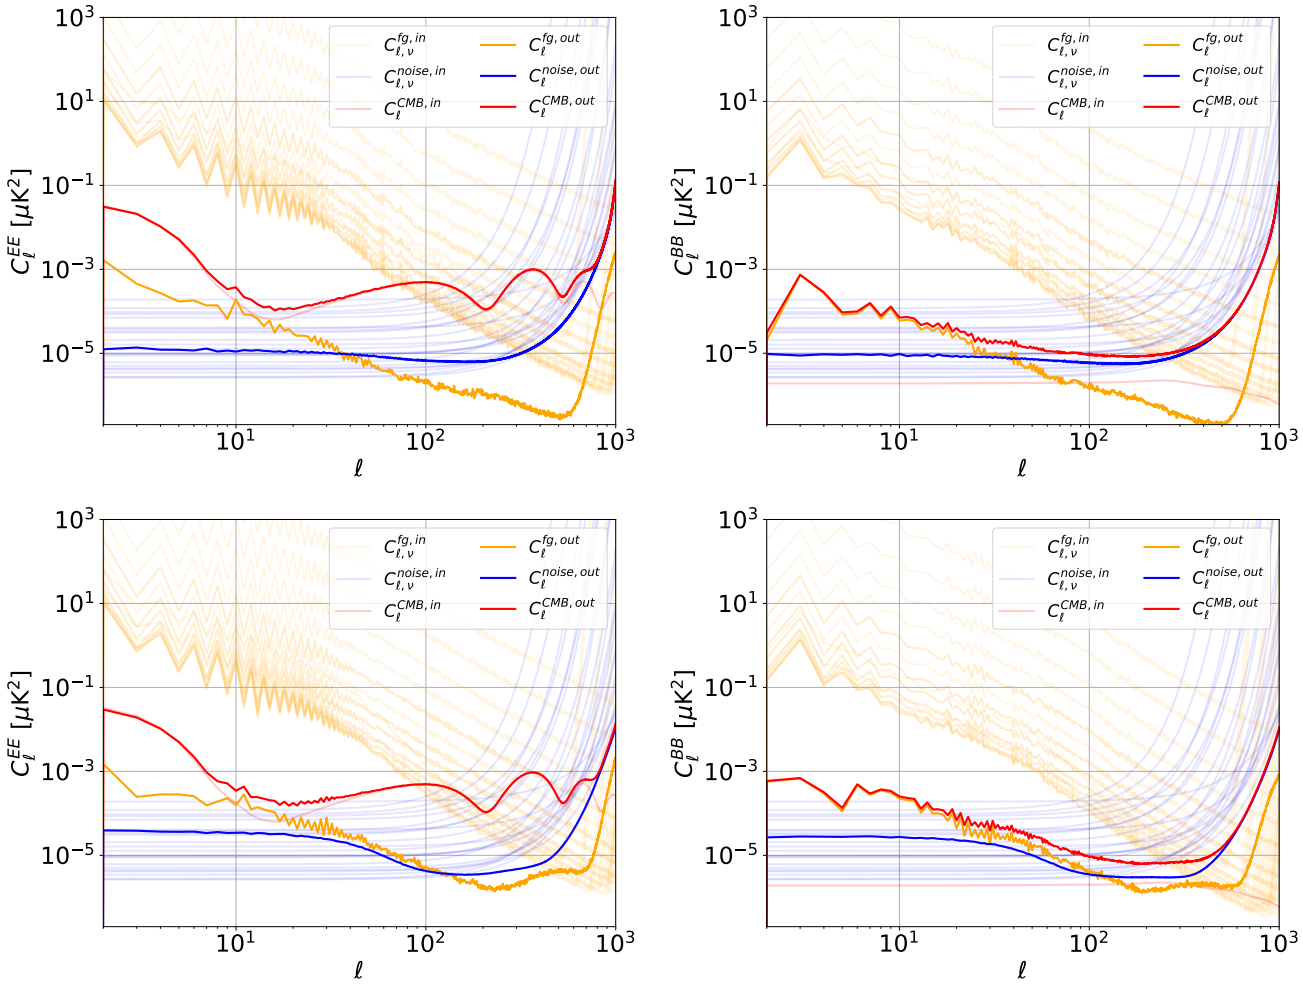

Figure 7.13: ***Preliminary!*** Summary figure for TOP-LEFT: *EE* TOP-RIGHT: *BB*. The clear curves represent the power spectra estimated from the full-sky HILC map, itself derived from the 22 input d1s1 maps (each corresponding to a LiteBIRD frequency channel). The spectra are coloured by component type: Galactic foregrounds in orange, noise in blue, CMB signal in red. The opaque curves represent the power spectra after HILC cleaning for the three components (associated with the same colours as before). BOTTOM: same with high-complexity foregrounds.

CHAPTER 8

Power-spectrum estimation

Even with an experiment scanning the entire sky as LiteBIRD, we have seen that masking the sky is mandatory to eliminate extra-galactic astrophysics emission and prevent excessive Galactic foreground contamination. After the HILC foreground cleaning, we have already anticipated in Sec. 7.4.3 that as much as 50% of the sky could be required for the challenging reconstruction of the tensor-to-scalar ratio. Estimating spectra on cut-sky is an important problem of CMB analysis that we will address in this chapter, first introducing (non-exhaustively) various methods and then testing them in the LiteBIRD context.

Contents

8.1 Methods overview 166
8.1.1 Full-sky standard estimator 166
8.1.2 Pseudo power spectrum 167
8.1.3 Maximum-Likelihood estimator 168
8.1.4 Quasi-Maximum-Likelihood Quadratic Estimator (QML) 169
8.2 Comparison of cross-spectra in a white noise case 171
8.2.1 High signal-to-noise 171
8.2.2 Low signal-to-noise 173
8.2.3 Hybrid estimator 174
8.2.4 ℓ transition and varying $f_{\rm sky}$ 176
8.3 Application to LiteBIRD and results 177
8.3.1 QML for LiteBIRD 177
8.3.2 Pseudo power spectrum 181
8.3.3 Comparison 181
8.3.4 Hybrid estimator 183

8.1. Methods overview

In this section, we overview various power spectrum estimation methods that we could apply in the context of LiteBIRD. Especially, we discuss the properties satisfied by these estimators, whether they are *unbiased* (*i.e.* give in average the *true* power spectrum), *optimal* (*i.e.* variance is minimal), or *lossless* (conserve all the information of the data). Of course, a lossless estimator is optimal, but not all optimal estimators are lossless.

8.1.1. Full-sky standard estimator

The first power spectrum estimator is the one introduced in Eq. 2.52 and Eq. 2.47, and implemented on the pixelised sky by the function `anafast` we introduced in Sec. 2.3.3. This estimator is *unbiased*, *optimal* and *lossless* on full-sky, up to $\ell \simeq 2N_{\rm side} - 1$, and is the one we used in the previous section. On full-sky, the blue curve of Fig. 8.1 displays the EE residuals over 20000 noiseless $N_{\rm side} = 32$ simulations with respect to the input and illustrates that the bias induced by the method is statistically consistent with zero for $\ell < \ell_{max} = 2N_{\rm side}$. Note that one can push this $\ell_{\rm max}$ value, beyond which integration errors appear, by iterating Eq. 2.52 [209].

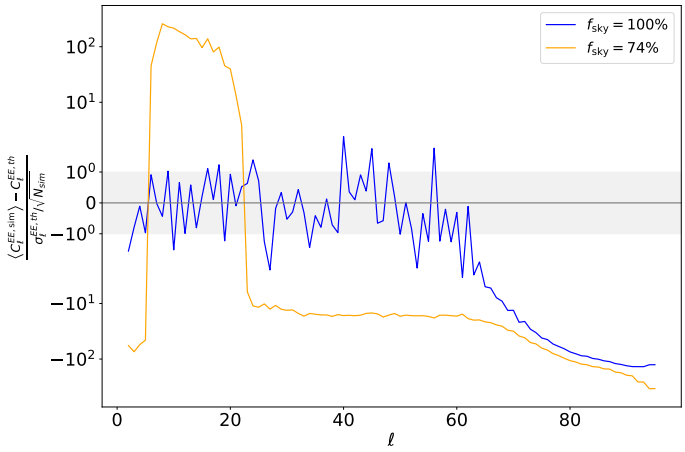

Figure 8.1: Residuals of the `anafast` estimator from two sets of $N_{\rm sim} = 20000$ simulated $N_{\rm side} = 32$ maps, one full-sky (blue), one 74% cut-sky (orange). These are normalised by the square root of the ratio cosmic variance over $\sqrt{N_{\rm sim}}$. In the full-sky case, this quantity scatters fairly around zero so the estimator is unbiased and optimal up to $\ell \simeq 63 = 2N_{\rm side}$.

When cutting the sky, this estimator induces E- and B-modes mixing, *i.e.* some modes are misinterpreted due to the effect of the mask. Effectively, since the E- modes are dominant over the B-modes, this mixing acts as a leakage of modes from E to B. Consequently, both EE and BB spectra end up being *biased*. In Fig. 8.1, the orange curve illustrates that the standard estimator is biased for $f_{\rm sky} = 74\%$.

The aim of this section is to discuss other power spectrum estimators that are unbiased even when the sky is partly masked. As we shall see, by knowing the pixel-pixel covariance matrix, some of them allow us to get close to optimality, but an unavoidable effect of masking is that power spectra become correlated over multipoles.

8.1.2. Pseudo power spectrum

The pseudo power spectrum aims to treat the effect of the mask directly in Eq. 2.52, in the same spirit of the pixel window function presented in Sec. 2.3.3. It consists of the minimal change in the standard full-sky estimator $\tilde{C}_\ell$ (obtained by applying Eqs. 2.52 and 2.47 to the masked map) that allows the correction of the E-to-B leakage effect. For each pixel, the sky signal f becomes:

$$f(p) \to W(p) f(p), \tag{8.1}$$

where W is the mask, a function over the pixel taking values between 0 and 1. Contrary to the beam effects that can be simply corrected by a window function, the effect of W in harmonic space can be described as a *mixing (or mode-coupling) matrix* $M_{\ell\ell'}$. More precisely, one obtains that the average of the measured spectrum $\tilde{C}_\ell^X$ $(X = TT, TE, EE, BB)$ is related to the true power spectrum C_ℓ^X by [210–214]:

$$\left\langle \begin{pmatrix} \tilde{C}_\ell^{\,TT} \\ \tilde{C}_\ell^{\,TE} \\ \tilde{C}_\ell^{\,EE} \\ \tilde{C}_\ell^{\,BB} \end{pmatrix} \right\rangle = M_\ell^{\ell'} \begin{pmatrix} C_{\ell'}^{TT} \\ C_{\ell'}^{TE} \\ C_{\ell'}^{EE} \\ C_{\ell'}^{BB} \end{pmatrix}, \tag{8.2}$$

where the mixing matrix reads

$$M_{\ell\ell'} \equiv \begin{pmatrix} M_{\ell\ell'}^{TT} & 0 & 0 & 0 \\ 0 & M_{\ell\ell'}^{TE} & 0 & 0 \\ 0 & 0 & M_{\ell\ell'}^{EE} & M_{\ell\ell'}^{EB} \\ 0 & 0 & M_{\ell\ell'}^{BE} & M_{\ell\ell'}^{BB} \end{pmatrix}, \tag{8.3}$$

whose matrix elements are given by

$$M_{\ell_1\ell_2}^{TT} = \frac{(2\ell_2+1)}{4\pi} \sum_{\ell_3} (2\ell_3+1)\, \tilde{W}_{\ell_3} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}^2, \tag{8.4}$$

$$M_{\ell_1\ell_2}^{TE} = \frac{(2\ell_2+1)}{8\pi} \sum_{\ell_3} (2\ell_3+1)\, \tilde{W}_{\ell_3} \left(1+(-1)^L\right) \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ -2 & 2 & 0 \end{pmatrix}, \tag{8.5}$$

$$M_{\ell_1\ell_2}^{EE} = M_{\ell_1\ell_2}^{BB} = \frac{(2\ell_2+1)}{16\pi} \sum_{\ell_3} (2\ell_3+1)\, \tilde{W}_{\ell_3} \left(1+(-1)^L\right)^2 \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ -2 & 2 & 0 \end{pmatrix}^2, \tag{8.6}$$

$$M_{\ell_1\ell_2}^{EB} = M_{\ell_1\ell_2}^{BE} = \frac{(2\ell_2+1)}{16\pi} \sum_{\ell_3} (2\ell_3+1)\, \tilde{W}_{\ell_3} \left(1-(-1)^L\right)^2 \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ -2 & 2 & 0 \end{pmatrix}^2, \tag{8.7}$$

where $\begin{pmatrix} \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \end{pmatrix}$ denotes $3j$-Wigner symbols, where $L = \ell_1 + \ell_2 + \ell_3$, and $\tilde{W}_\ell$ is the power spectrum of the mask, obtained by applying Eqs. 2.52 and 2.47 to $W(p)$.

Hence, one can define an *unbiased* pseudo power-spectrum estimator $\hat{C}_\ell^X$ by deconvolving the measured power spectrum $\tilde{C}_\ell^X$:

$$\begin{pmatrix} \hat{C}_\ell^{TT} \\ \hat{C}_\ell^{TE} \\ \hat{C}_\ell^{EE} \\ \hat{C}_\ell^{BB} \end{pmatrix} = (M^{-1})_\ell^{\ell'} \begin{pmatrix} \tilde{C}_{\ell'}^{TT} \\ \tilde{C}_{\ell'}^{TE} \\ \tilde{C}_{\ell'}^{EE} \\ \tilde{C}_{\ell'}^{BB} \end{pmatrix}. \tag{8.8}$$

This requires inverting the coupling matrix, which is singular in general. The issue can be solved by binning the multipoles before performing the inversion. The required binning depends, in general, on the situation encountered.

The pseudo-spectrum estimator method can be generalised to cross-spectrum estimation [214], using $\tilde{W}_\ell^{AB} = \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell} w_{\ell m}^A w_{\ell m}^{B*}$ instead of $\tilde{W}_\ell$, where $w_{\ell m}^Z$ are the mask harmonic coefficients applied to map Z. Eq. 8.8 is then straightforwardly generalised. When the measured spectra come with beam and noise, Eq. 8.2 is modified as $\left\langle \tilde{C}_\ell \right\rangle = M_\ell^{\ell'} W_{\ell'}^2 C_{\ell'} + N_\ell$ (see Sec. 4.2), effect that can be easily corrected for in the estimator Eq. 8.8.

This method also provides a way to express analytically the covariance matrix between pseudo-cross-power spectra $\hat{C}_\ell^{AB}$ and $\hat{C}_\ell^{CD}$, which is defined as

$$\Xi_{\ell\ell'}^{AB,CD} = \left\langle \left(\hat{C}_\ell^{AB} - \left\langle \hat{C}_\ell^{AB} \right\rangle\right) \left(\hat{C}_{\ell'}^{CD} - \left\langle \hat{C}_{\ell'}^{CD} \right\rangle\right)^* \right\rangle. \tag{8.9}$$

[211] and [214] showed that for high multipoles and large sky coverage that is common to both maps, this matrix can be approximated as

$$\Xi_{\ell\ell'}^{AB,CD} \simeq \frac{1}{\tilde{\nu}_{\ell'}} \left[\hat{C}_\ell^{AC}\hat{C}_{\ell'}^{BD} + \hat{C}_\ell^{AD}\hat{C}_{\ell'}^{BC}\right], \tag{8.10}$$

where $\tilde{\nu}_\ell$ is the effective number of degrees of freedom of the C_ℓ distribution assumed to be a χ^2, see [212]. One can simplistically account for the loss of modes when masking the sky in the following way:

$$\tilde{\nu}_\ell = (2\ell+1)\Delta_\ell \frac{w_2^2}{w_4}, \tag{8.11}$$

where w_i is the i-th moment of the mask:

$$w_i = \frac{1}{4\pi}\int_{4\pi} d\Omega W^i(\Omega). \tag{8.12}$$

This approximate result for the effective number of degrees of freedom becomes wrong when the sky fraction is too low.

To reduce the estimator error bars, the mask should be optimised in order to make ν_ℓ maximum and the variance and correlations over multipoles minimum. The pseudo-power-spectrum estimate is close to optimal for small scales, and white-noise dominated regimes [211]. However, for the largest scales, it is no longer the case: part of the information is lost in the inversion of the mixing matrix, and other estimators perform better at minimising the variance.

In the following, we will use the `XPol` implementation of the pseudo-cross-power spectrum estimator, including polarisation [215] based on [214].

8.1.3. Maximum-Likelihood estimator

The maximum likelihood power-spectrum estimator consists of estimating the power-spectrum by directly maximising Eq. 2.34 with respect to C_ℓ [216, 217]. In the asymptotic limit, this estimator is *optimal*, *unbiased* and *lossless*. However, this estimate depends on the pixel vector $\mathbf{m}$ in a nonlinear way. Implementing it in practice requires sampling numerically the likelihood over C_ℓ, which is time-consuming. Moreover, the resulting $\hat{C}_\ell$ probability density functions are analytically unknown [218], which is a consequent drawback for the subsequent parameter estimation. Instead, a simpler method approaching the ML results as closely as possible exists.

8.1.4. Quasi-Maximum-Likelihood Quadratic Estimator (QML)

We search for a simpler and well-behaved power-spectrum estimator [218]. A key criterion to determine if an estimator is lossless is whether it saturates the Cramer-Rao inequality (see Eq. 4.26), *i.e.* its covariance matrix is the inverse of the Fisher matrix of the pixel-based likelihood Eq. 2.34 [218]. Since this likelihood is simply a Gaussian, the Fisher matrix in this context directly follows from Eq. 4.29 (with the expectation values of the pixels vanishing):

$$F^C_{\ell\ell'} = \frac{1}{2}\,\mathrm{tr}\left(\mathbf{M}^{-1}\frac{\partial \mathbf{M}}{\partial C_\ell}\mathbf{M}^{-1}\frac{\partial \mathbf{M}}{\partial C'_\ell}\right) = \frac{1}{2}\,\mathrm{tr}\left(\mathbf{M}^{-1}\mathbf{P}_\ell\mathbf{M}^{-1}\mathbf{P}_{\ell'}\right), \tag{8.13}$$

where $\mathbf{M} = \mathbf{S} + \mathbf{N}$ is the pixel-pixel covariance matrix introduced in Sec. 2.3.1, and the $\mathbf{P}_\ell$ matrices are defined as

$$\mathbf{P}_\ell \equiv \frac{\partial \mathbf{M}}{\partial C_\ell}, \tag{8.14}$$

in order for the signal covariance matrix to simply write

$$\mathbf{S} = \sum_\ell \mathbf{P}_\ell C_\ell. \tag{8.15}$$

Notice that defining these $\mathbf{P}$ matrices is done in order to simplify the notations hereafter, but that they are not a new objects, and we can directly compute them from Eq. 2.42. Moreover, Eq. 8.15 emphasises a linear relation between map space (to the square) and spectrum space (regardless of the presence or absence of masked pixels). Hence, it is legitimate to expect a solution for C_ℓ that is linear in $m_i m_j$ where m_i are the element of the pixel vector $\mathbf{m}$. Moreover, *a priori*, such an estimator could well be lossless: the data set formed with n^2 data $m_i m_j$ contains as much information as $\mathbf{m}$ that contains n data m_i [218].

This leads to the search for an estimator for the power spectrum from partial-sky maps $\mathbf{m}$ of shape $(3n)$ with the following quadratic form:

$$\hat{C}_\ell = \mathbf{m}^T\mathbf{E}_\ell\mathbf{m} - b_\ell, \tag{8.16}$$

where $\mathbf{E}_\ell$ are matrices of shape $(3n \times 3n)$ and b a vector of shape ℓ. *A priori*, any choice of $\mathbf{E}$ and b define a valid *Quadratic Estimator (QE)*. Let us show that unbiasedness and optimality are respectively ensured by:

$$b_\ell = \mathrm{tr}(\mathbf{E}_\ell\mathbf{N}), \tag{8.17}$$

$$\mathbf{E}_\ell = \frac{1}{2}(F^{-1})^{\ell'}_\ell\mathbf{M}^{-1}\mathbf{P}_{\ell'}\mathbf{M}^{-1}. \tag{8.18}$$

To do so, we insert Eqs. 8.17 and 8.18 into Eq. 8.16 to get our estimator, for which we search for the expectation value and covariance. For a quadratic estimator, these are textbook relations.

Firstly, the expectation value gives

$$\begin{aligned}
\mathrm{E}(\hat{C}_\ell) &= \mathrm{tr}\,(\mathbf{E}_\ell\mathbf{M}) - \mathrm{tr}\,(\mathbf{E}_\ell\mathbf{N}) \\
&= \mathrm{tr}\,[\mathbf{E}_\ell(\mathbf{M}-\mathbf{N})] \\
&= \mathrm{tr}\left[\frac{1}{2}\left(F^{-1}\right)_\ell^{\ell'}\mathbf{M}^{-1}\mathbf{P}_{\ell'}\mathbf{M}^{-1}\left(\mathbf{P}^{\ell''}C_{\ell''}\right)\right] \\
&= \left(F^{-1}\right)_\ell^{\ell'}\frac{1}{2}\,\mathrm{tr}\left[\mathbf{M}^{-1}\mathbf{P}_{\ell'}\mathbf{M}^{-1}\mathbf{P}^{\ell''}\right]C_{\ell''} \\
&= \left(F^{-1}\right)_\ell^{\ell'}F_{\ell'}^{\ell''}C_{\ell''} \\
&= \delta_\ell^{\ell''}C_{\ell''} \\
&= C_\ell,
\end{aligned} \tag{8.19}$$

which shows the unbiasedness of the estimator. Secondly, the covariance gives

$$\begin{aligned}
\mathrm{cov}(\hat{C}_\ell,\hat{C}_{\ell'}) &= 2\,\mathrm{tr}\left(\mathbf{E}_\ell\mathbf{M}^{-1}\mathbf{E}_{\ell'}\mathbf{M}^{-1}\right) \\
&= 2\,\mathrm{tr}\left(\frac{1}{2}\left(F^{-1}\right)_\ell^{\tilde{\ell}}\mathbf{M}^{-1}\mathbf{P}_{\tilde{\ell}}\mathbf{M}^{-1}\mathbf{M}^{-1}\frac{1}{2}\left(F^{-1}\right)_{\ell'}^{\tilde{\ell}'}\mathbf{M}^{-1}\mathbf{P}_{\tilde{\ell}'}\mathbf{M}^{-1}\mathbf{M}^{-1}\right) \\
&= \frac{1}{2}\left(F^{-1}\right)_\ell^{\tilde{\ell}}\mathrm{tr}\left[\mathbf{M}^{-1}\mathbf{P}_{\tilde{\ell}}\mathbf{M}^{-1}\mathbf{P}_{\tilde{\ell}'}\right]\left(F^{-1}\right)_{\ell'}^{\tilde{\ell}'} \\
&= \left(F^{-1}\right)_\ell^{\tilde{\ell}}F_{\tilde{\ell}\tilde{\ell}'}\left(F^{-1}\right)_{\ell'}^{\tilde{\ell}'}. \\
&= \delta_{\ell\tilde{\ell}'}\left(F^{-1}\right)_{\ell'}^{\tilde{\ell}'}. \\
&= \left(F^{-1}\right)_{\ell\ell'}.
\end{aligned} \tag{8.20}$$

This teaches us that this particular QE saturates the Cramer-Rao inequality. Consequently, it is lossless and also optimal.

Furthermore, the statistics followed by a quadratic form is well-studied in the literature [219]. On top of its expectation value and its covariance, we know the whole probability density function of a quadratic form: the *generalised* χ^2. One notices that this time, it is the Fisher matrix that fulfils the role of the mode mixing matrix; see Eqs. 8.16 and 8.18. Hence, the QE power spectra are also correlated over multipoles.

However, an apparent major drawback of the estimator is that all the conclusions we have derived about its statistical properties assume that we know independently the three covariance matrices $\mathbf{M}$, $\mathbf{N}$ and $\mathbf{S}$.

Concerning $\mathbf{S}$, a fiducial power spectrum is required in order to compute Eq. 2.42. The ML estimator reduces to iterating the QE over the prior power spectrum until it equals the measured one. Conversely, if the initial guess for the power spectrum is close to the true one, then the resulting QE estimate will be close to the ML solution. Nowadays, we have a good idea of the actual CMB power spectrum, which has been estimated on large sky fractions with high signal-to-noise on a wide range of scales. Thus, the prior power spectrum does not limit the optimality of the estimator, which is close to the ML one. It justifies why this particular QE is often referred to as *Quasi-Maximum-Likelihood estimator (QML)*, a naming that we adopt hereafter.

In practice, one has to verify on simulations that an error on the pixel-pixel covariance matrix does not increase dramatically the errors on the estimate. Furthermore, not knowing $\mathbf{N}$ prevents a good estimate of the uncertainty, forcing one to work with realistic simulations.

The last worrying consequence of knowing imperfectly $\mathbf{N}$ when estimating auto power spectra is a poor estimation of the noise bias, which, in consequence, cannot be perfectly cancelled in Eq. 8.16. The obvious solution for this issue is to improve our model for the noise covariance matrix by running realistic simulations and mastering the instrument. A simpler turnaround is to estimate cross spectra between distinct data instead of auto spectra.

The QML estimator can be generalised for cross spectra between two maps A and B [220] and becomes:

$$\hat{C}_\ell^{AB} = \mathbf{m}^{A^T} \mathbf{E}_\ell^{AB} \mathbf{m}^B - b_\ell^{AB}, \tag{8.21}$$

where $b_\ell^{AB} = \mathrm{tr}\left(\mathbf{E}_\ell \mathbf{N}^{AB}\right) = 0$ and

$$\mathbf{E}_\ell^{AB} = \frac{1}{2}(F^{-1})_\ell^{\ell'} (\mathbf{M}^{AA})^{-1} \mathbf{P}_{\ell'} (\mathbf{M}^{BB})^{-1}. \tag{8.22}$$

This allows us to estimate the cross-power spectra between various Stokes parameters or data splits. One shows that in an ideal case and similarly to the auto QML estimator, the cross QML estimator is (very close to) optimal [220]. Moreover, contrary to the auto case, under the condition that the noise realisation in A is independent of the one in B, the cross QML estimator is unbiased even with poor knowledge of the noise covariance matrix.

We use our own code to derive auto-power spectra, which computes efficiently $\mathbf{S}$ thanks to `numpy` broadcasting and has been cross-checked against the publicly available code `xQML` [221]. The cross-spectra computation has directly been performed with `xQML`.

8.2. Comparison of cross-spectra in a white noise case

In this section, we compare the performance of the cross pseudo-Cl estimator to the cross QML estimator for *EE* and *BB* spectra reconstruction. We work with a specific setup for pseudo-Cl estimation, multipole per multipole (one bin = one ℓ). We draw conclusions for two specific noise regimes, a high signal-to-noise one with $\sigma = 0.01$ μK·arcmin, and a low signal-to-noise one of 10 μK·arcmin, closer to the actual LiteBIRD setup post foreground cleaning.

The input maps for QML are at $N_{\rm side} = 16$ and the maps for pseudo-power spectrum estimation at $N_{\rm side} = 512$ resolution. We focus on a mildly cut sky of 72% for this discussion. The associated latitudinal masks are obtained by cutting 16° above and below the Galactic plane; the QML one is binary at $N_{\rm side} = 16$ and the pseudo-power spectrum one is apodised by a 6° Gaussian beam at $N_{\rm side} = 512$. The cross-QML spectra are computed between $\ell = 2$ and $\ell = 30$, while the cross-pseudo ones are computed up to $\ell = 1000$.

8.2.1. High signal-to-noise

Fig. 8.2 illustrates the performance of the cross QML and the cross pseudo spectrum estimation methods on $N_{\rm sim} = 1000$ high signal-to-noise ($\sigma = 0.01$ μK·arcmin) simulations from an input power spectrum $C_\ell^{\rm in}$, with a $f_{\rm sky} = 72\%$ mask.

The left-hand side shows the normalised *EE* **residuals** defined from the estimate $\hat{C}_\ell$ and the input spectrum $C_\ell^{\rm in}$:

$$R_\ell \equiv \frac{\langle \hat{C}_\ell \rangle - C_\ell^{\rm in}}{\sqrt{\sigma^2(\hat{C}_\ell)/N_{\rm sim}}}, \tag{8.23}$$

and the $\pm 1\sigma$ band is highlighted in lightgrey. For this particular set of realisations, the mean and standard deviation of the residuals estimated over 30 multipoles for cross QML give respectively -0.28 and 1.18, while the same quantities for cross pseudo-Cl give respectively -0.01 and 0.99. These residuals are statistically in perfect agreement with 0 in both cases (given the number of multipoles over which we infer this statistics), illustrating that both methods yield unbiased results, an expected result since the noise bias cancels out for cross-spectra of independent noise maps.

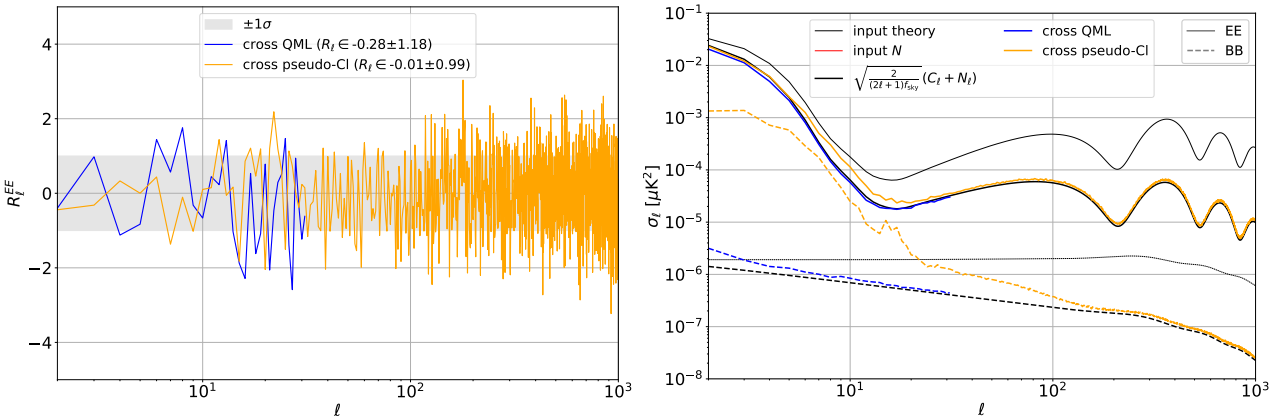

Figure 8.2: LEFT: Empirical normalised residuals, see Eq. 8.23, for the two *EE* cross power spectrum estimation methods (in blue cross QML, in orange, cross pseudo-Cl) as compared with the $\pm 1\sigma$ region. For this particular set of realisations, we display in the legend the averages and standard deviations of the residuals for both methods, estimated over the ℓ-range of the figure. RIGHT: standard deviation for the same methods (associated with the same colours), for *EE* (plain curves) and *BB* (dashed curves), as compared with the formula Eq. 2.49, which is an approximation to the theoretically optimal standard deviation. The empirical curves are estimated over 1000 simulations, which have been generated from the fiducial parameters of Tab. 4.1, with no foregrounds and high signal-to-noise regime (white noise with $\sigma = 0.01$ μK·arcmin), and an $f_{\rm sky} = 72\%$ mask. The noise curve (labelled in red) is below the y-axis range, at $8.5 \times 10^{-12} \mu\text{K}^2\cdot\text{rad}^2$.

In the same setup, the right panel illustrates the spectrum **standard deviation** that results from each of the two methods (same colours as the left panel), as compared with the input signal (black thin lines) and the cosmic-sample variance formula given by Eq. 2.49 (black thick lines). It shows that the power spectra estimated via the QML method come with smaller error bars than pseudo-Cl's. This is because the QML method is more optimal than the pseudo-Cl one, which is expected at the largest scales and in cases where the pixel covariance matrix is well-known. In our setup, at the largest scales, the variance on cross QML is more than 2 orders of magnitude below the one for cross pseudo spectrum for *BB* and also yields improved error bars for *EE*, almost a factor 2 around $\ell = 10$. The variance on cross QML almost coincides with the theoretical variance formula rescaled by $f_{\rm sky}^{-1/2}$.

On top of the variance, the **correlation** of the estimate over multipoles will impact the parameter estimation. We illustrate in Fig. 8.5 the correlation matrix for *EE* and *BB* for $2 \leq \ell \leq 31$ in the setup of 0.01 μK·arcmin noise at $f_{\rm sky} = 72\%$. The left panel illustrates the cross-pseudo power spectrum situation, which at low ℓ induces significant correlation patterns over

BB multipoles. The cross-correlations $EE \times BB$ are not as strong, and EE multipole cross-correlations are almost inexistent. For cross-QML illustrated on the right panel, there is almost no visible correlation between multipoles and modes. In this particular setup, the figure hints at a better constraining power with cross-QML spectra than with cross-pseudo spectra.

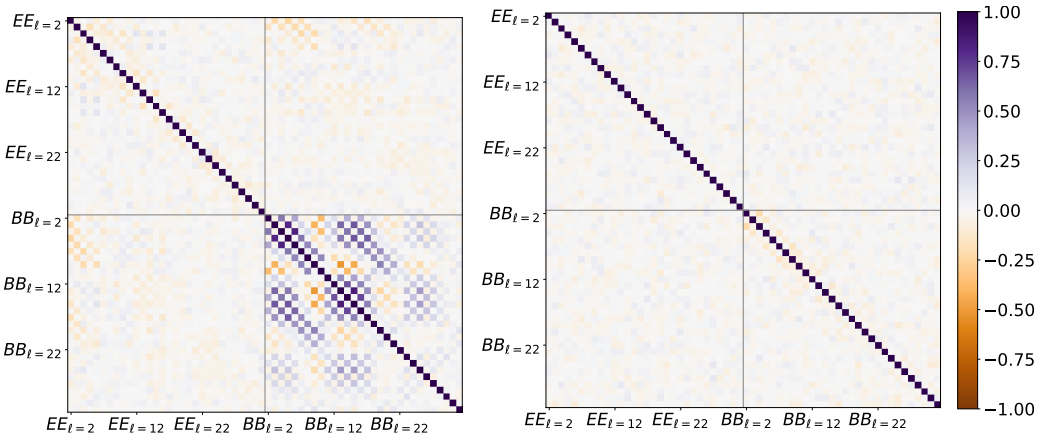

Figure 8.3: Correlation matrices between EE and BB multipoles at $f_{\rm sky} = 72\%$, with LEFT: cross-pseudo spectra RIGHT: cross-QML. The white noise is negligible with respect to the signal ($\sigma = 0.01$ μK·arcmin).

8.2.2. Low signal-to-noise

We show in Fig. 8.4 the case where the white noise level is comparable in magnitude to the signal: 10 μK·arcmin. This represents well the level we expect with LiteBIRD HILC, as emphasised by the two first rows of Table 7.1. The red flat curve illustrates this level as compared with the signal. The setup is otherwise the same as the one of Fig. 8.2.

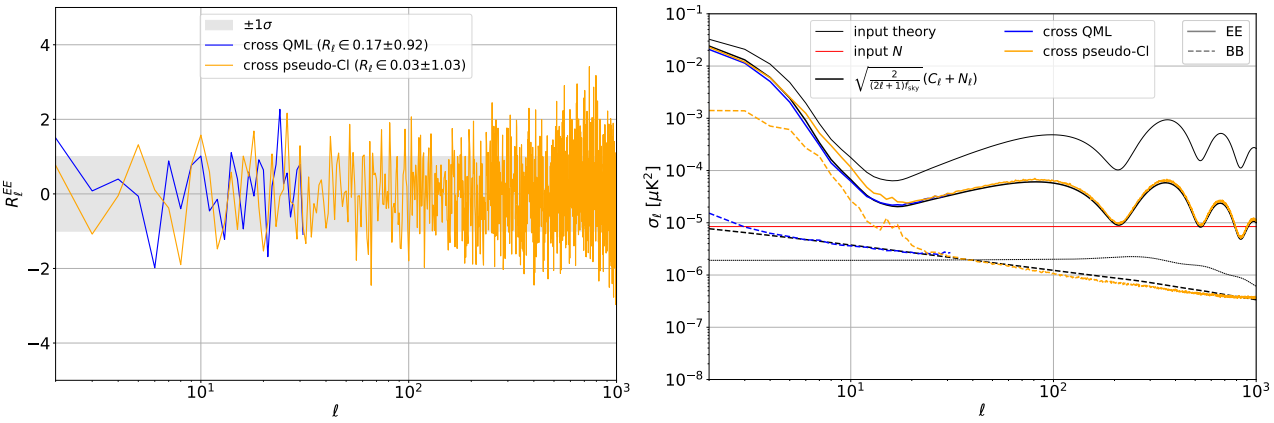

Figure 8.4: Same two figures as Fig. 8.2 for a white noise of 10 μK·arcmin. This time, the noise curve (red) is well visible and dominates the BB lensing signal by almost one order of magnitude.

The conclusion regarding the bias, illustrated on the right panel by the **residuals**, is similar to the high signal-to-noise one: both methods are unbiased. Regarding the estimate **uncertainty**, the conclusion is once again similar to the high signal-to-noise case: in our setup, the cross

QML method is more efficient than the cross pseudo one at yielding power spectra with low uncertainty.

We also display in Fig. 8.3 the **correlation matrices** for cross pseudo (left panel) and for cross QML (right panel). This figure is similar to Fig. 8.5 but with a 10 μK·arcmin noise. Comparing the left panels of Fig. 8.3 and Fig. 8.5, one sees that adding noise slightly decreases the correlations between the *BB* multipoles. Otherwise, the two noise levels give really similar results. The cross QML spectra remain mostly uncorrelated over multipoles. Hence, the conclusion with higher noise is similar to the one with low noise: in our setup, we expect to have better constraints from cross-QML spectra than from cross-pseudo spectra.

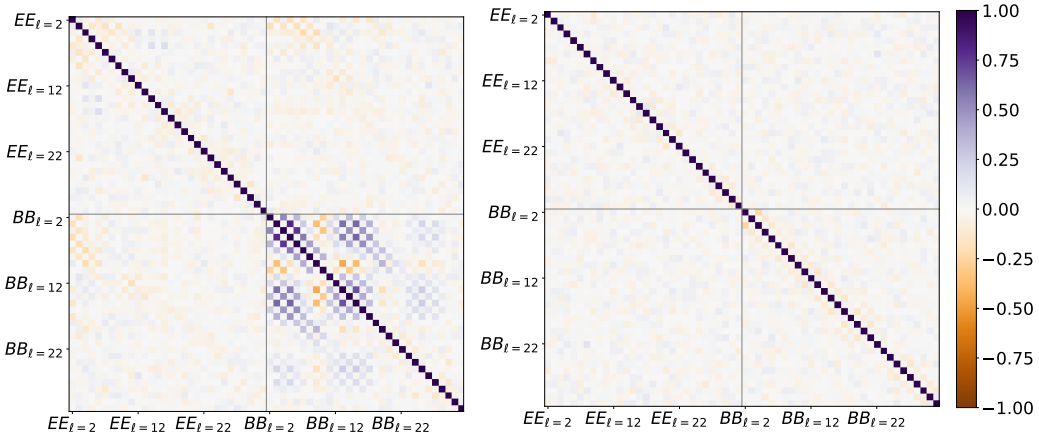

Figure 8.5: Correlation matrices between *EE* and *BB* multipoles at $f_{\rm sky} = 72\%$, with LEFT: cross-pseudo spectra RIGHT: cross-QML. The white noise is negligible with respect to the signal ($\sigma = 10$ μK·arcmin).

We conclude that, in optimal setups, one should use QML estimates of the power spectrum, the difference in variance and correlation being significant at the largest scales.

8.2.3. Hybrid estimator

For computational reasons, the QML method can realistically be applied only to a limited number of pixels and multipoles. Indeed, it requires manipulating **E**, an object of shape $(n\times n\times \ell)$. Consequently, the standard implementation typically requires $N^2_{\rm stokes} \times N^2_{\rm pix} \times N_{\rm spec} \times N_\ell$ Bits of RAM. In our use cases, we limit ourselves to $N_{\rm side} = 16$ for estimating polarisation spectra up to $\ell = 32$ (which is enough for reconstructing the reionisation bumps in most cases, see Figs. 4.10 and 4.11) [1]. The unbinned pseudo spectra method is close to being an optimal estimator for high ℓ (>150 for *BB*). In order to get better results at low multipoles, one usually bins Eq. 8.3 before inverting it. The pseudo-C_ℓ method is computationally inexpensive and does not require knowledge of the pixel-pixel covariance matrix (hence needs fewer simulations).

Depending on the problems addressed, both QML and pseudo-Cl have been used in past works. When the target is information that lies in the largest scales, QML is to be preferred (see *eg.* [223]), while pseudo spectra are good enough in other cases (see *eg.* [58]). Analyses that aim

[1] A recent work, [222], proposes approximations that allow pushing the computation up to $N_{\rm side} = 64$ without departing too much from optimality

at reconstructing a wide-scale range mix the two methods [102]. One usually uses simulations to estimate the covariance matrix of the resulting hybrid spectrum[2].

We built such a hybrid estimator based on the two methods we have introduced and detailed, with a simple transition from cross QML to cross-pseudo-power spectrum at $\ell_{\text{transi}} = 32$. The resulting averaged power spectra with their estimated variance (EE in blue and BB in orange) are illustrated in the left panels of Fig. 8.6, while the associated estimated correlation matrix is represented in the right panels. The upper panels correspond to the high signal-to-noise setup, and the bottom panel corresponds to the low signal-to-noise one. The rest of the setup is the same as previously.

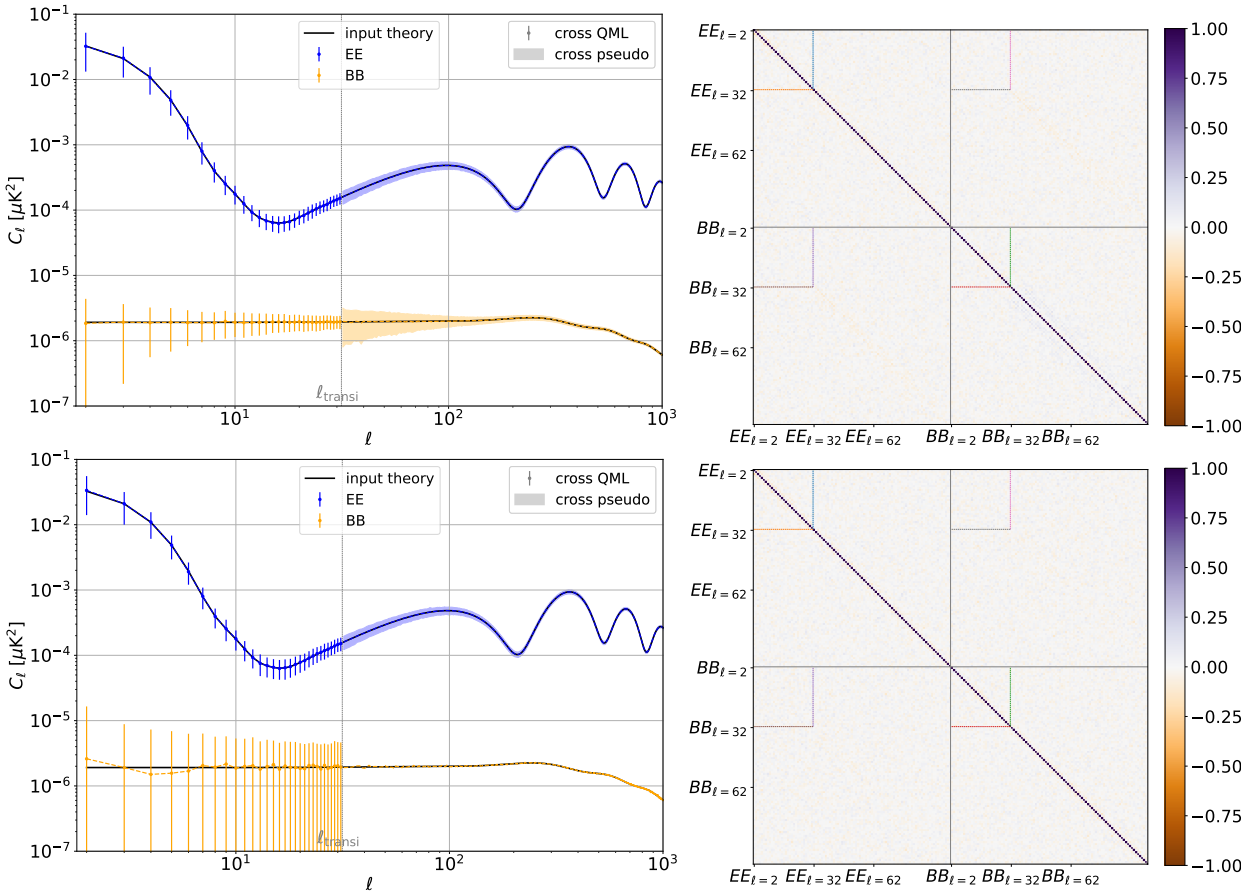

Figure 8.6: Hybrid EE and BB estimator which consists in cross-QML spectrum for $\ell < 32$ and cross-pseudo power spectrum for $\ell \geq 32$ LEFT: Estimated signal and its associated standard deviation, blue: EE, orange BB RIGHT: correlation matrix over spectra and multipoles TOP: $\sigma = 0.01$ μK·arcmin. BOTTOM: $\sigma = 10$ μK·arcmin.

On the left panels of Fig. 8.6, the QML part of the hybrid estimator is depicted by 68% asymmetric error bars, while the pseudo spectrum part is represented with 68% asymmetric coloured region. The resulting hybrid estimator is unbiased and quasi-optimal at all scales,

[2] [210, 211] present a similar hybrid estimator and give the approximate and analytical expression of the covariance matrix.

apart in the high signal-to-noise case around $\ell_{\rm transi}$ in BB. To solve this issue, one could choose a higher $\ell_{\rm transi}$.

We show the correlation between the 100 first multipoles for EE and BB on the right panels. We delimited the top corner QML×QML sub-block with a dotted square in each of the four main blocks. These figures illustrate that all EE and BB multipoles are essentially uncorrelated two-by-two. By avoiding the use of pseudo power spectra at low ℓ, we have circumvented the problems of correlations induced by this estimator. Furthermore, the abrupt transition between the two estimators doesn't introduce additional off-diagonal terms in the correlation matrix.

8.2.4. ℓ transition and varying $f_{\rm sky}$

Finally, we apply our hybrid estimator to a wider range of sky fractions: $f_{\rm sky} = 31.25$, 52.05, 72.87, 100% where the masks investigated are still latitudinal. We keep focusing on white noise cases $\sigma = 0.01, 10$ μK·arcmin. Moreover, we tune $\ell_{\rm transi}$ for each setup under scrutiny in order to make the variance of the hybrid spectrum estimate as low as possible. More precisely, we automatically chose $\ell_{\rm transi}$ as the first multipole such that QML gives a higher variance than pseudo-Cl. For computational reasons, we still restrict to $\ell_{\rm transi} \leq 47$.

We show the standard deviation of the resulting hybrid BB spectra in Fig. 8.7. We represent the various sky fractions with various colours, the low signal-to-noise regimes in dashed lines, and the high signal-to-noise regimes in plain lines. One sees that for $f_{\rm sky} = 31.25$, 52.05, 72.87%, $\ell_{\rm transi}$ is between $\ell = 31$ and $\ell = 47$. It means that one benefits from applying QML at large scales, especially at the lowest sky fractions. Instead, in the full-sky case and with our configuration, QML does not yield smaller error bars than pseudo-Cl, even at the largest scales. On full-sky, the cross pseudo-Cl estimate is already optimal without the need for cross QML.

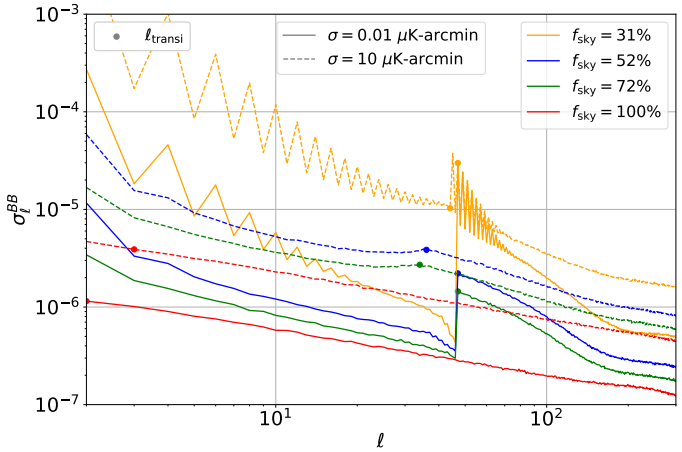

Figure 8.7: Variance of the hybrid BB estimator for various sky fractions and the two levels of white-noise; 0.01 and 10 μK·arcmin. The transition multipole between QML and pseudo-Cl estimator, $\ell_{\rm transi}$, is represented for each setup as a dot on the hybrid spectrum. It is optimised to minimise the variance of the resulting hybrid estimator.

The associated hybrid BB spectra will serve us as input in Sec. 9.3.2 (because we need medium scales to constrain r). One also observes some oscillating patterns in the variance of the spectra estimated on low sky fraction, especially in the $f_{\rm sky} = 31\%$ case, which might be due to the particular choice of latitudinal masks.

8.3. Application to LiteBIRD and results

In order to analyse the LiteBIRD simulations and prepare a robust extraction of the cosmological parameters from realistic data, we present in this section how we applied the estimators to the simulations cleaned by HILC.

8.3.1. QML for LiteBIRD

We have seen that in order to apply QML, we need to work at low resolution. Hence, the first step for processing the LiteBIRD (HILC-cleaned) simulations involves downgrading the maps from $N_{\rm side} = 512$ to $N_{\rm side} = 16$, without losing information or adding artificial features at large scales.

To avoid the aliasing effect, which occurs when the sampling rate of a signal is below the Nyquist rate, we need to eliminate the small scales before increasing the size of the pixels. To achieve this, we introduce a beam window function that vanishes at $3N_{\rm side}$ while remaining smooth:

$$\begin{aligned} W_\ell^{S,\ N_{\rm side}} &= \frac{1}{2}\left[1+\cos\left(\pi\ \frac{\ell-2}{3N_{\rm side}-1}\right)\right]\ \text{if}\ \ell \leq 3N_{\rm side}, \\ &= 0\ \text{else}. \end{aligned} \tag{8.24}$$

We summarise in Fig. 8.8 the various useful multipole window functions in this part:

- in red the beam given by Eq. 8.24,
- in pink the pixel window function at $N_{\rm side} = 16$, $W_\ell^{p,16}$,
- in orange the beam window function of 30 arcmin adopted for the HILC maps, $W_\ell^{B,\text{output HILC}}$,
- in blue the pixel window function of the native maps $W_\ell^{p,512}$.
- For comparison, we show the lowest and highest LiteBIRD-design beam functions that respectively correspond to the L1-040 ($W_\ell^{B,\rm L1-040}$) and H3-402 ($W_\ell^{B,\rm H3-402}$) frequency channels.

We transform the output HILC maps to harmonic space, correct for the native pixel window function $W_\ell^{p,512}$ and apply $W_\ell^{S,\ 16}$. Then, we transform back to map space to obtain a $N_{\rm side} = 16$ map, which will still carry the intrinsic HILC beam $W_\ell^{B,\text{output HILC}}$, the beam due to the new pixelisation $W_\ell^{p,16}$, and the artificial beam we introduced in the downgrading $W_\ell^{S,\ 16}$. All these beams will be corrected after the QML power spectrum estimation. The resulting map is displayed on the middle panel of Fig. 8.9, while the original map from which it has been derived is displayed on the left panel.

The other required ingredient for applying QML is the mask. To remain as realistic as possible, these are based on the realistic mask configuration introduced in Sec. 7.4.2, without apodisation ($\theta_2 = 0$) and downgrading to $N_{\rm side} = 16$.

After HILC and downgrading, the noise matrix is *a priori* unknown, even for the simulations with known inputs. A solution to obtain it would require doing more simulations before empirically estimating it. Typically, for reconstructing such a matrix, which should be invertible, the strict minimum required amount of simulations is the rank of the matrix, $N_{\rm stokes} \times N_{\rm pix}$ (and usually much more are needed to have a correct estimate). Being limited by our $\sim \mathcal{O}(100)$ simulations, we adopt another strategy.

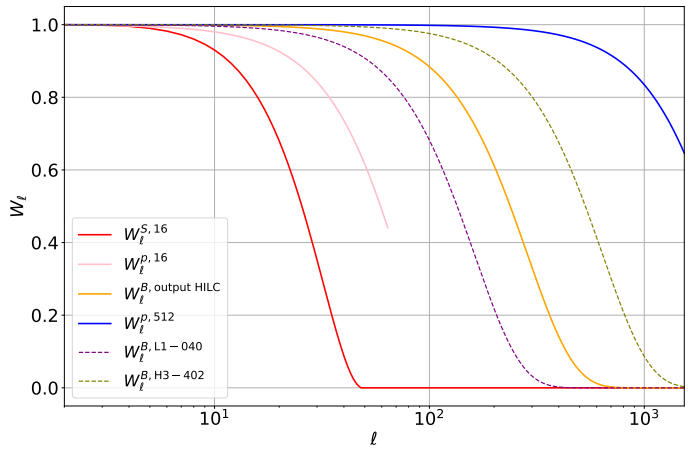

Figure 8.8: Different beams introduced in this work. In red the downgrade beam, in pink the pixel window function at $N_{\rm side} = 16$, in orange the beam window function of 30 arcmin adopted for the HILC maps, in blue the pixel window function for $N_{\rm side} = 16$. The purple (*resp.* green) dashed line illustrates the window function of the widest (*resp.* tightest) LiteBIRD beam, which corresponds to the L1-040 (*resp.* H3-402) frequency channel.

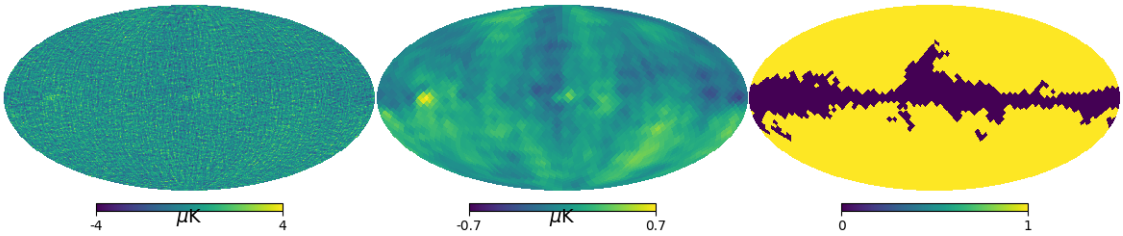

Figure 8.9: ***Preliminary!*** LEFT: particular simulation of a Q HILC map at $N_{\rm side} = 512$ (wn1f + d1s1), MIDDLE same map downgraded to $N_{\rm side} = 16$ (explanations in the text), RIGHT realistic binary mask we use to estimate the QML spectra.

We want to reconstruct the *EE* and *BB* spectra, so we are considering the Q and U Stokes parameters. We start from the singular covariance matrix $\hat{\mathbf{M}}$, estimated on the available simulations and displayed on the right-hand side of Fig. 8.10. One notices its noisiness. We build a simplistic covariance matrix $\tilde{\mathbf{M}}$ by setting the off-diagonal terms to 0. Then, we assume a fiducial power spectrum that matches the fiducial cosmological model Tab. 4.1, allowing to derive $\mathbf{S}$ thanks to Eq. 2.42, which is displayed on the left-hand side of Fig. 8.10. We can derive the corresponding $\tilde{\mathbf{N}} = \tilde{\mathbf{M}} - \mathbf{S}$ that is exact on the diagonal but that is simplified outside the diagonal.

These $\tilde{\mathbf{N}}$, $\tilde{\mathbf{M}}$, $\mathbf{S}$ are all invertible by construction, and they can be adopted in the QML. However, the QML estimate will be based on a diagonal total covariance matrix, which will lead to two consequences: it will deviate from optimality and be biased by the noise bias (which will be poorly estimated).

We show on the left panel of Fig. 8.11 the biased auto QML power spectra (orange) for *EE* and their residuals (blue) at low multipoles for various sky fractions and for baseline complexity foregrounds + wn + 1f. The left panel of Fig. 8.12 displays a similar representation for *BB*. The

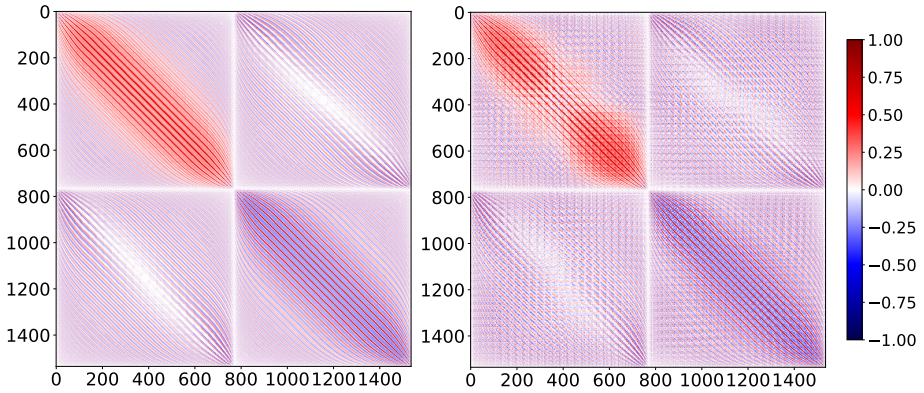

Figure 8.10: ***Preliminary!*** LEFT: Analytical signal covariance matrix $\mathbf{S}$ from the simulation-input C_ℓ, RIGHT: Empirical total covariance matrix $\hat{\mathbf{M}}$ estimated on the d1s1 + wn + $1/f$ HILC-cleaned simulations, inaccurate and not invertible due to lack of statistics. Both are obtained for a full sky at $N_{\rm side} = 16$.

figure allows to quantify the order of magnitude of the noise bias in the LiteBIRD case. This residual would be even more present when the foregrounds get more complex, as we previously discussed based Fig. 7.12. The issue seems less critical in *EE* than in *BB* if one forgets that the bias stands in a critical region for estimating τ.

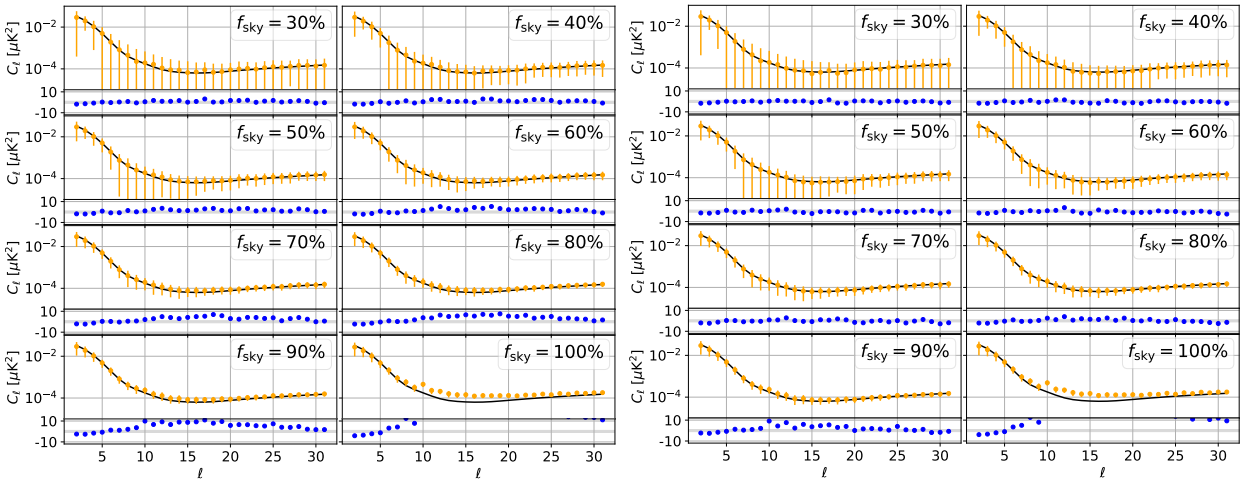

Figure 8.11: ***Preliminary!*** LEFT: *EE* auto QML power spectra (orange) and associated residuals (blue) for various sky fractions,RIGHT: cross QML spectra. They are based on the baseline complexity + wn + $1/f$ simulations, as compared with the input signal. The residuals in per cent are normalised by $\sqrt{N_{\rm sim}} = \sqrt{200}$ and compared to the ± 1 grey contour that represents the statistical fluctuation range.

We implement two approaches in order to address this issue by cancelling the noise bias:

1) Cross-correlating split maps allows for the estimation of unbiased cross-power spectra (at the price of increasing the variance).
2) Estimating the QML power spectrum of the noise residual maps allows us to estimate

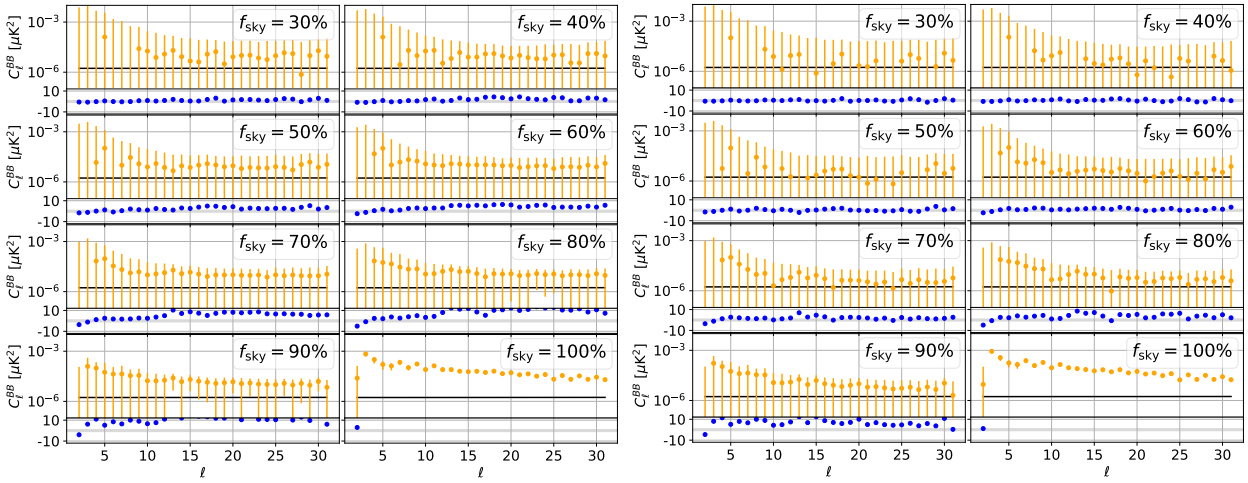

Figure 8.12: ***Preliminary!*** For various sky fractions LEFT: *BB* auto QML power spectra (orange) and associated residuals (blue), RIGHT: *BB* cross QML power spectra (orange) and associated residuals (blue). Based on the baseline complexity + wn + $1/f$ simulations, as compared with the input signal. The residuals in per cent are normalised by $\sqrt{N_{\rm sim}} = \sqrt{200}$ and compared to the ± 1 grey contour that represents the statistical fluctuation range.

the bias that can hence be subtracted. Of course, it can only be realised for a study on simulations and not with real data.

We focus on the first option for the discussions to come, which is the strategy closest to the actual data analysis. First, to mimic real data splits (across mission time or detectors), we multiply each input noise frequency map by $\sqrt{2}$ to build mock split data from the available simulations. Then, we add each CMB simulation with two independent noise splits before computing the associated HILC weights to produce two splits of HILC-cleaned maps. CMB simulation per CMB simulation, we estimate the cross-power spectra between the two split maps.

The resulting cross **spectra** and their **residuals** are displayed on the right-hand sides of Fig. 8.11 for *EE* and Fig. 8.12 for *BB*, with the same conventions than the left-hand sides to ease the comparison between auto and cross outputs. One notices that on cross spectra, most of the bias is cancelled between $f_{\rm sky} = 40\%$ and 90% for *EE*, and between $f_{\rm sky} = 30\%$ and 60% for the $\ell \geq 20$ for *BB*. These are the regions where the noise bias dominates the systematic budget (see Fig. 7.6). With this correction, the systematic residual in the spectrum essentially only contains the contribution from the foreground residuals.

We have seen that the pixel-pixel Fisher matrix Eq. 8.13 gives the covariance matrix of the power spectra. However, in this real case, we cannot use it: we have estimated the matrix $\tilde{\mathbf{M}}$ over 200 simulations, and we have noticed that this matrix is not close enough to the true $\mathbf{M}$ to give a reliable estimate of the Fisher matrix (an expected remark since 200 simulations for estimating a (60, 60)-shaped $\mathbf{M}$ is quite a low amount). Instead, we estimate empirically the covariance matrix from the obtained spectra. We represent the $f_{\rm sky} = 30\%$ (left) and $f_{\rm sky} = 70\%$ (right) **correlation matrices** in Fig. 8.13, for cross-QML *EE* and *BB* spectra. One observes that the off-diagonal blocks corresponding to the correlation between *EE* and *BB* spectra are consistent with zero, demonstrating that on real simulations also, QML performs well in unmixing *B*-modes from *E*-modes on cut-sky. Multipoles are only mildly coupled at $f_{\rm sky}$= 70%, while correlation

patterns appear for lower sky fractions, with negative terms on even diagonals. This effect has already been pin-pointed by previous works (see *eg.* [210]). We have verified that the auto-spectra matrix is extremely similar to the cross-spectra ones.

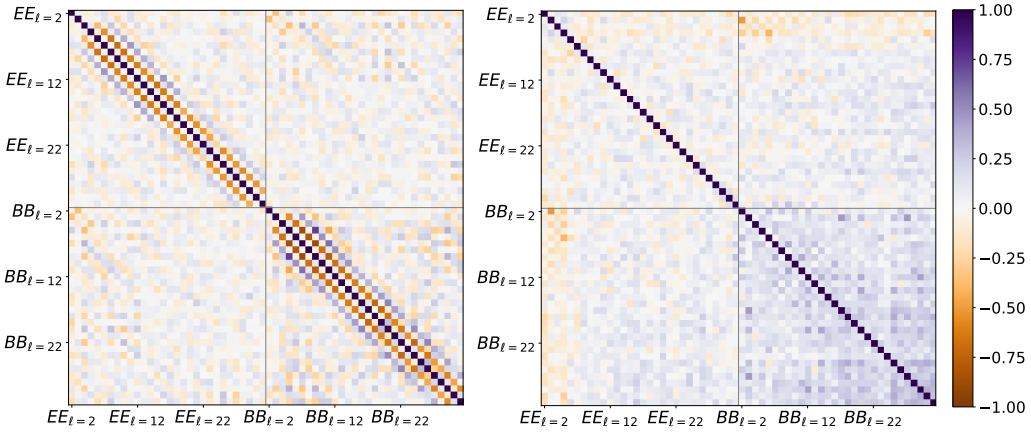

Figure 8.13: ***Preliminary!*** Correlation matrices of cross-QML LiteBIRD *EE* and *BB* spectra LEFT: $f_{\rm sky} = 30\%$, RIGHT: $f_{\rm sky} = 70\%$.

8.3.2. Pseudo power spectrum

We apply the pseudo-power spectrum estimation to the HILC-cleaned maps. The *EE* and *BB* **spectra** obtained from an *ideal* mask with $\theta_1 = \theta_2 = 2^\circ$ are shown in Fig. 8.14. For this particular figure, in order to represent the noise residual, we adopt the second option for correcting the output auto spectrum by directly subtracting the noise residuals at the power spectrum level[3]. The cross-pseudo spectrum yields very similar recovered spectra. The left panel Fig. 8.14 indicates that, as anticipated, the recovered spectrum is extremely close to the input for *EE* at all scales at $f_{\rm sky} = 90\%$ since the residuals are lower than the signal. In the next paragraph, we illustrate the residuals and compare them to QML ones. From the right panel, one observes that the only significant bias that remains at $f_{\rm sky} = 50\%$ in *BB* bias is at $\ell < 15$ and at $\ell > 750$. We illustrate by a grey contour the region in which a $10^{-3} < r < 10^{-2}$ signal would stand. Qualitatively, one sees that the low-ℓ bias could mimic a primordial B signal, challenging a robust detection of the latter. This also explains the high $\sigma^{\rm Fisher}_{r,{\rm syst}}$ we saw in the latter chapter, which we will also rediscuss in the next chapter.

Additionally, we represent in Fig. 8.15 the cross-pseudo spectrum correlation matrices for $f_{\rm sky} = 30\%$ (left) and $f_{\rm sky} = 70\%$ (right). As for QML, correlation patterns appear at $f_{\rm sky} = 30\%$ with negative terms on even diagonals. At $f_{\rm sky} = 70\%$ instead, the multipoles are close to be uncorrelated.

8.3.3. Comparison

Let us compare the performance of the cross-QML and the cross-pseudo spectrum methods for LiteBIRD. The right panels of Fig. 8.16 illustrate a comparison between the pseudo-power spectrum and QML *EE* spectrum **residuals**. The upper panel focuses on the $f_{\rm sky} = 70\%$ setup,

[3]More precisely, considering Eq. 7.29, we define the red curve such that $\hat{C}_\ell \to \hat{C}_\ell - C^{\rm n}_\ell - 2C^{\rm n\times fg}_\ell$ to ensure no cross terms due to noise remain in the corrected spectrum.

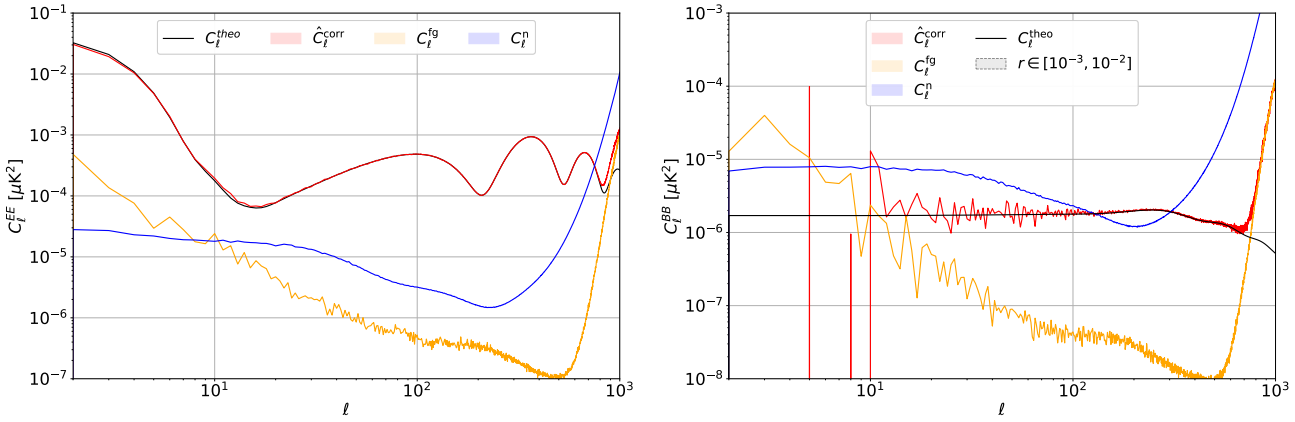

Figure 8.14: ***Preliminary!*** Pseudo power spectrum estimation averaged over 200 wn1f + baseline simulations, cleaned with HILC ($N_m = 400$, $\sigma_{w_\ell} = 35$), with an ideal mask. The theoretical spectrum is displayed in black curves, the noise-corrected output spectrum with the associated 68% region in red (see the core of the text), and the noise and the foreground residuals, respectively represented in orange and blue. LEFT: *EE* with $f_{\rm sky} = 90\%$ RIGHT: *BB* with $f_{\rm sky} = 50\%$. The $r \in [10^{-3}, 10^{-2}]$ theoretical region is displayed as a grey contour.

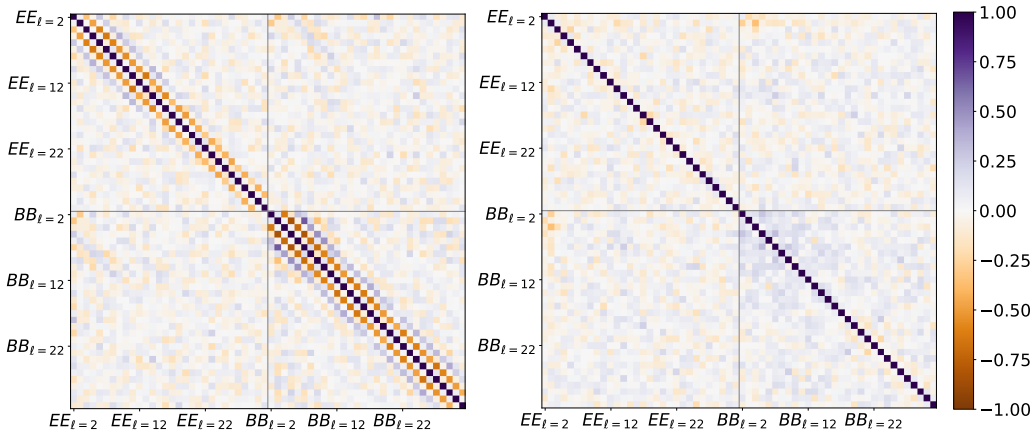

Figure 8.15: ***Preliminary!*** Correlation matrices of cross-pseudo LiteBIRD *EE* and *BB* spectra LEFT: $f_{\rm sky} = 30\%$, RIGHT: $f_{\rm sky} = 70\%$.

while the lower one represents the $f_{\rm sky} = 30\%$ setup. On top of the effect of the previously-discussed foreground residuals, one observes a slight decrease in the residuals around $\ell = 2$ and 3 for both methods, which is not statistically significant given the number of simulations. With more simulations, we expect to distinguish the feature origin, statistical flux in the simulations, or regular loss of power at the largest scales in the simulations.

The right panels of Fig. 8.16 illustrate a comparison between the pseudo-power spectrum and QML spectrum **uncertainties**. Contrary to the simple white-noise case (see Fig. 8.2), QML doesn't provide an estimate with lower variance than the pseudo C_ℓ method anymore neither at $f_{\rm sky} = 30\%$ or 70%. Possible reasons for this observed deviation from optimality are that additional steps and assumptions have been required to apply QML to the LiteBIRD simulations, especially concerning the pixel-pixel covariance matrix, see Sec. 8.3.1.

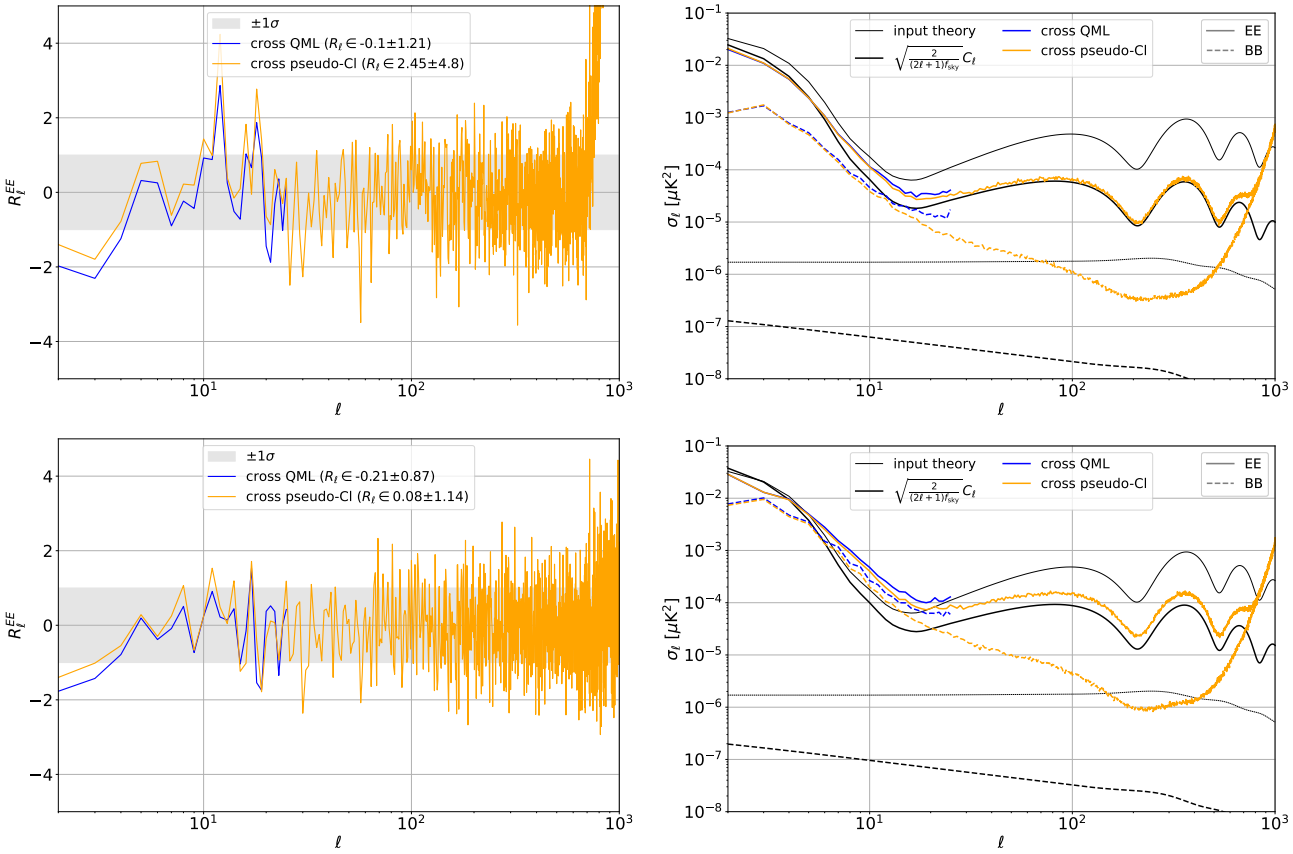

Figure 8.16: ***Preliminary!*** The two panels contain similar figures as Fig. 8.2 representing the biases and variances in the *EE* and *BB* estimates for QML and PCl, with 200 wn + $1/f$ + d1s1 HILC simulations. The upper panels focus on the $f_{\rm sky} = 70\%$ setup, while the lower ones illustrate the $f_{\rm sky} = 30\%$ setup.

Regarding the correlations induced by both methods, one can compare Figs. 8.13 and 8.15. In the LiteBIRD realistic case, the multipoles are no longer more correlated for cross-QML than for cross-pseudo-spectrum estimates. Once again, this has to originate from the additional steps and assumptions needed to apply QML to LiteBIRD simulations.

8.3.4. Hybrid estimator

Finally, we illustrate the hybrid estimator we can build from cross-QML and cross-pseudo spectrum applied to the LiteBIRD HILC simulations. We illustrate respectively the $f_{\rm sky} = 70\%$ and $f_{\rm sky} = 30\%$ on the upper and lower panels of Fig. 8.17. In this case, we choose $\ell_{\rm transi} = 25$. The left panel illustrates that the recovered *EE* cross spectra are almost bias-free while a significant bias and a high variance characterise the *BB* large scales.

As wcan be anticipated from Sec. 8.3.3, such a hybrid estimator is not optimal given our implementation of QML for LiteBIRD simulations. The cross pseudo spectra give better estimates in terms of variance and correlation at all multipoles. Consequently, optimising $\ell_{\rm transi}$ at any $f_{\rm sky}$, as done in Sec. 8.2.4, yields $\ell_{\rm transi} = 2$, *i.e.* the hybrid spectrum coincides with the pseudo-Cl estimator.

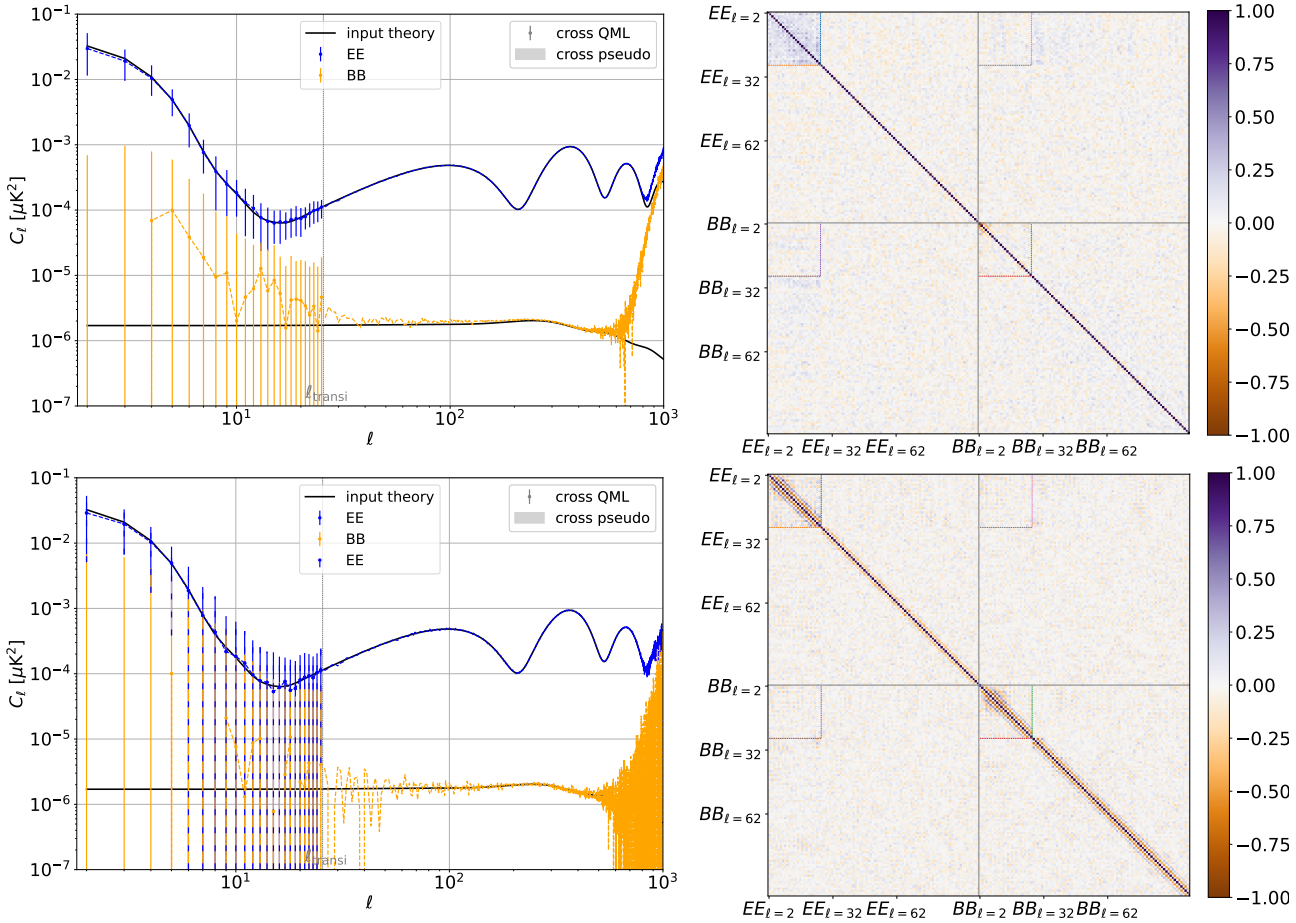

Figure 8.17: ***Preliminary!*** Same figures as Fig. 8.6 representing the hybrid spectra, with 200 wn + $1/f$ + d1s1 HILC simulations, $f_{\rm sky}$=70%. The upper panels illustrate the $f_{\rm sky}$ = 70% setup, while the ones illustrate the $f_{\rm sky}$ = 30% setup.

Conclusion

In this section, we have presented how we have implemented the QML and pseudo-Cl methods to estimate power spectra on cut-sky. We have applied these two methods on simple white noise simulations and on LiteBIRD simulations. We have shown that QML yields better estimates than pseudo-Cl for spectra on white-noise cut-sky, justifying why we build a hybrid estimator. However, on LiteBIRD simulations, the situation is less clear and QML also suffer a significant departure from optimality, which is probably due to the preprocessing of the QML inputs (especially the downgrading and the simplification of the pixel covariance matrix). This preprocessing is to be improved in the future.

In this section, we have illustrated the baseline + wn + $1/f$ scenario. We have verified that the main conclusions regarding power spectrum estimation methods for higher complexity foregrounds remain unchanged. Conclusions about the residual levels for different foreground complexities have already been discussed in Sec. 7.4.4. The spectra associated with more complex foreground setups have also been produced and will be used later on.

In the next chapter, consistent with the conclusions of this section, we will use only QML

and hybrid spectra for the white noise cases and adopt simple cross-pseudo spectra for LiteBIRD applications. Moving forward, our power spectrum estimation remains an area of investigation that could lead to improvements. The figures and results in this chapter and the upcoming one are based on mock masks, which require further investigation. No apodisation or multipole-bin optimisations have been made for the pseudo power spectrum method. We also aim to improve the performance of our QML implementation on the most realistic simulations despite their low available amount.

CHAPTER 9

Likelihood investigation and parameter estimation

This chapter investigates various methods for estimating the τ and r parameters from large-scale EE and BB power spectra. First, we do a quick preliminary Fisher status on the expected order of magnitude we expect for r and τ uncertainties. Then, we assess the ability of various C_ℓ-based likelihood functions to provide optimal estimators for the parameters under scrutiny and their uncertainties in white-noise cases. This optimality is tested by comparing the results of each likelihood function to the exact pixel-based likelihood. We consider various scenarios for the input simulations and the assumed fiducial models (auto or cross-spectra, low or high noise, empirical or analytical covariance matrix, correct or incorrect fiducial model).

Then, we apply the C_ℓ-based likelihood functions to LiteBIRD simulations with corrected multipole covariance matrices. Finally, we conclude on the status of this LiteBIRD investigation, providing conclusions on the expected error bars on τ and r obtained by our full-analysis pipeline.

Contents

9.1 Preliminary τ and r uncertainty estimate 188

9.2 Likelihoods approximations overview 189

9.2.1 Pixel-space 189

9.2.2 Full-sky spectrum-space 190

9.2.3 Cut-sky spectrum-space 192

9.3 Application to the white-noise case and results 199

9.3.1 τ estimation from white-noise case 200

9.3.2 r estimation from white-noise case 208

9.4 Application to LiteBIRD and results 210

9.4.1 Covariance matrix correction 211

9.4.2 τ estimation from LiteBIRD simulations 213

9.4.3 r estimation from LiteBIRD simulations 218

9.1. Preliminary τ and r uncertainty estimate

Before performing a proper maximum-likelihood estimation of the τ and r parameters from the spectra derived in this section, we use the pseudo power spectra average residuals to re-derive Fisher estimates of the systematic and statistical uncertainties on the two parameters, as initially explored in Sec. 7.4.3.

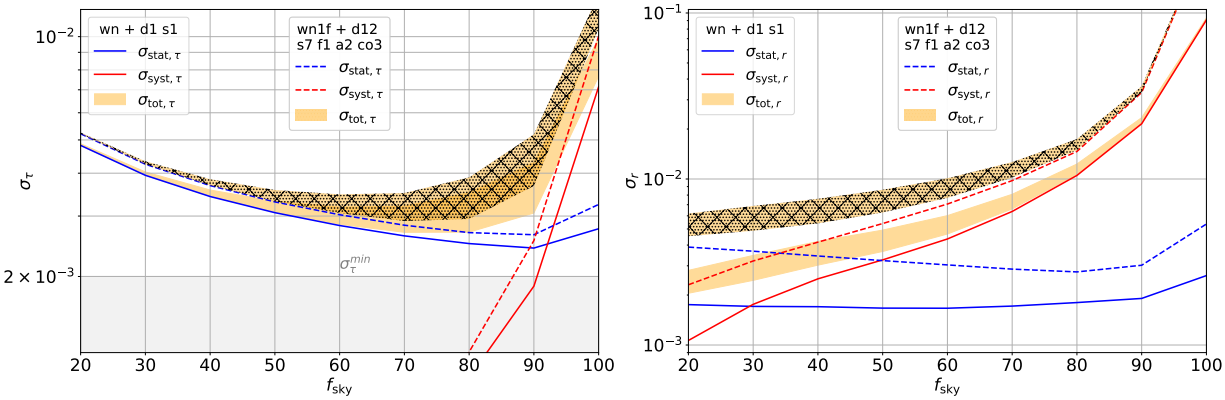

Figure 9.1: ***Preliminary!*** Fisher estimate of the $\sigma^{\mathrm{Fisher}}_{\mathrm{stat}}$ (blue) and $\sigma^{\mathrm{Fisher}}_{\mathrm{syst}}$ (red) for τ (left panel) and r (right panel), obtained from the (unbiased) EE and BB pseudo power spectrum estimation of the residuals averaged over 200 wn1f + d1 s1 simulations, cleaned with HILC ($N_m = 400$, $\sigma_{w_\ell} = 35$), with an ideal mask. The total uncertainty $\sigma^{\mathrm{Fisher}}_{\mathrm{tot}}$ will lie between the root sum square and the sum ($\equiv \delta$) of the two residuals, represented by the orange contour. The corresponding uncertainty and the σ_{tot} area with high complexity foregrounds (wn1f + d12 s7 f1 a2 co3) are represented with dashed lines and a hashed contour. The best σ_τ accessible in the reionisation bump due to cosmic variance, which is exactly 2×10^{-3} for our fiducial model, is also shown.

We show in Fig. 9.1 shows these quantities with respect to the sky fraction. An orange contour represents the total uncertainty, which combines both errors. Its lower bound is the root sum square of the two errors, treating them as uncorrelated statistical errors. The upper bound represents the sum of the two sources. A grey line shows the best σ_τ that can be extracted from the reionisation bump due to cosmic variance (assuming no noise or foregrounds), which is exactly 2×10^{-3} for our fiducial model. The corresponding quantity for r, *i.e.*, the best σ_r achievable with our fiducial ($r = 0$) in the setup without delensing of BB, noise, or foregrounds, is $\sigma_r^{\mathrm{min}} = 1.1 \times 10^{-4}$.

Given our fiducial cosmology and HILC setup, we can draw preliminary Fisher-based conclusions from this analysis:

- τ from EE:
  - The systematic-dominated regime transitions to a statistical-dominated regime around f_{sky}=90%,
  - The f_{sky} configuration minimising the total uncertainty lies between 70% and 80%,
  - At this f_{sky}, τ will be inferred at the level of 3×10^{-3}, a budget dominated by statistical uncertainty and not systematic one[1]. This amounts to 1.5 times the best achievable measurement from the EE reionisation bump due to cosmic variance.

[1]This is a level a bit worse than what is mentioned in [145, 153], in which a cosmic variance limited

- r from BB (assuming $r = 0$):

  - The systematic-dominated regime transitions to a statistical-dominated regime around $f_{\rm sky}$=30%,
  - Consequently, the smaller the sky, the smaller the uncertainty, down to 20%. At the price of losing in agnosticism, this would require a proper modelling of the foreground residual in the likelihood function. This would allow the removal of the red curve in Fig. 9.1, changing the situation. This will be a next goal beyond the current work.
  - Around $f_{\rm sky}$=50%, depending on the foreground model, the total uncertainty ranges from 3×10^{-3} and 8×10^{-3}, 3 to 8 times above the LiteBIRD requirements[2].

It is important to notice that this result is based on a Fisher analysis, which does not render the full shape of the likelihood function in parameter space and is not exact on a cut-sky. Instead, we aim to build, maximise, and fully analyse a reliable likelihood function in spectrum space to setup the final piece of our pipeline and to get a definitive assessment of τ and r.

9.2. Likelihoods approximations overview

In this section, we will introduce various likelihood functions in pixel-space and C_ℓ-space across different configurations (single/multi-field, full/cut-sky, noise/no noise). This overview is not meant to be exhaustive. In some configurations, one can derive analytical likelihoods, while in others, especially when working with spectra on a cut sky, only approximate likelihoods are possible to derive. As usually done, we will provide likelihood expressions as $-2 \log \mathcal{L}$ up to a constant term by analogy with the Gaussian case. Since we are not testing non-Gaussianities, we assume that all cosmological information is contained in the angular power spectrum and that $\mathcal{L}(\theta|D) = \mathcal{L}(C_\ell(\theta)|D)$.

9.2.1. Pixel-space

The likelihood of argument C_ℓ given the outcome $\mathbf{m}$, which derives from the probability density function Eq. 2.34, has already been discussed in Sec. 8.1.3 in the context of power spectrum estimation. Based on the same Eq. 2.34, the so-called *pixel-based log-likelihood* $\mathcal{L}_{\rm PB}$, of argument $\{\theta_i\}$ upon which C_ℓ depends, reads:

$$-2 \log \mathcal{L}_{\rm PB} \left(\{\theta_i\} \mid \mathbf{m} \right) = \log \left| \mathbf{M} \left(\{\theta_i\} \right) \right| + \mathbf{m}^T \mathbf{M}^{-1} \left(\{\theta_i\} \right) \mathbf{m}, \tag{9.1}$$

where $\mathbf{M}$ is the pixel covariance matrix defined in Sec. 2.3.1, and where the constant term has been omitted. The theory is present in both terms. We represent in Fig. 9.2 the two (-2 log)-terms (blue and orange) as well as their sum (red). For this illustration, the data consists of one QU maps realisation generated at resolution $N_{\rm side} = 8$ from the fiducial cosmology defined in

measurement of τ from reionisation bump is forecasted. One should remember that our methodology is minimal and intentionally simple at this stage, with no assumptions about the foregrounds. Improvements can be made by parametrising the mixing matrix and marginalising over the foreground residuals in the likelihood function.

[2]This result is not inherently concerning for meeting the LiteBIRD requirements for the same reason as for τ. The enhancements mentioned in *footnote 1* also hold in the r case, although they are beyond the scope of the present work, they should help to achieve the r-requirement as emphasised by [153].

Table 4.1 and a white noise of 20 μK·arcmin. The parameter of the likelihood is τ. Furthermore, we normalise the three curves to 0 when -2 log $\mathcal{L}$ is minimal. For the implementation, we use the identity $\log|\mathbf{M}| = \mathrm{Tr}[\log(\mathbf{M})]$.

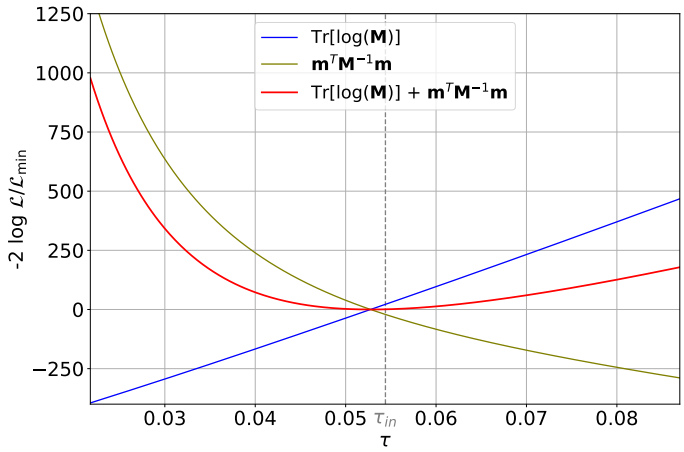

Figure 9.2: τ pixel-based likelihood and its two constituting subterms for one QU realisation with fiducial cosmology shown in Table 4.1 and white noise of 20 μK·arcmin. The three curves are normalised to 0 when -2 log $\mathcal{L}$ is minimal.

The presence of two terms with opposite trends with respect to the parameter is a general feature, as the covariance matrix appears in one term while its inverse appears in the other. In principle, this likelihood is exact, with no approximations made. However, the computation of the two terms must be very accurate. Hence, this likelihood becomes practically unusable in real-life scenarios, where neither the noise covariance matrix nor the foregrounds are perfectly known. Generally, and in this work, we primarily employ it in reference cases, where input data is simulated and thus known, to evaluate the performance of the spectrum-based likelihoods. Subsequently, the latter type of likelihood is favoured for real data because they are much more robust to unmodeled noise or foregrounds than the pixel-based one described here.

9.2.2. Full-sky spectrum-space

Let us assume that instead of pixels, the outcome is a measured power spectrum, $\hat{C}_\ell$, estimated from full-sky maps, to begin with, and we want to build a likelihood for the theoretical $C_\ell(\{\theta_i\})$. In full-sky the probability density function $f_{C_\ell}(\hat{C}_\ell)$ can be analytically written down for each independent multipoles, and the likelihood function simply follows:

$$\mathcal{L}(\{\theta_i\} \mid \hat{C}_\ell) = \prod_\ell f_{C_\ell(\{\theta_i\})}(\hat{C}_\ell). \tag{9.2}$$

The derivations presented below assume noiseless spectra but would remain the same when including a noise spectrum N_ℓ. In this case, for an auto-spectrum C_ℓ^{XX}, the substitution to perform for $\hat{C}_\ell^{XX} \to \hat{C}_\ell^{XX} + N_\ell^{XX}$.

Let's say we measure a single field, such as temperature. The associated harmonics coefficients are statistically independent and randomly distributed and follow a Gaussian distribution with vanishing average and covariance given by Eq. 2.17. In practice, by ergodicity hypothesis, one can

estimate this ensemble average from the $2\ell+1$ available coefficients per multipole, see Eq. 2.47, an estimator that is unbiased and optimal in full-sky.

Let's derive the distribution of this estimator (see *eg.* [34, 224, 225]). By definition, the sum of the squares of ν independent (complex) standard normal random variables follows the χ^2 distribution with k degrees of freedom. In particular, $\frac{a_{\ell m}}{\sqrt{C_\ell}} \sim \mathcal{CN}(0,1)$, where $\mathcal{CN}$ is the complex normal distribution, so $\hat{Y}_\ell \equiv \sum_m \left|\frac{a_{\ell m}}{\sqrt{C_\ell}}\right|^2$ follows a χ^2 distribution:

$$f_{C_\ell}(\hat{Y}_\ell) = \chi^2(\hat{Y}_\ell;\nu_\ell) = \frac{\hat{Y}_\ell^{\nu_\ell/2-1}}{\Gamma(\nu_\ell/2)2^{\nu_\ell/2}} \exp\left(-\frac{\hat{Y}_\ell}{2}\right), \tag{9.3}$$

where $\nu_\ell = 2\ell+1$, see Eq. 2.48. Additionally, it follows from Eq. 2.47 that $\hat{Y}_\ell = \nu_\ell \frac{\hat{C}_\ell}{C_\ell}$, hence,

$$f_{C_\ell}(\hat{C}_\ell) = \frac{{\nu_\ell}^{\nu_\ell/2-1}}{\Gamma(\nu_\ell/2)2^{\nu_\ell/2}} \frac{\hat{C}_\ell^{\nu_\ell/2-1}}{C_\ell^{\nu_\ell/2}} \exp\left(-\frac{\nu_\ell}{2}\frac{\hat{C}_\ell}{C_\ell}\right). \tag{9.4}$$

This is a Γ distribution of probability density $f_\Gamma(x;k,\theta) = \frac{1}{\Gamma(k)\theta^k} x^{k-1} e^{-x/\theta}$ with shape $k = \nu_\ell/2$ and scale $\theta = 2C_\ell/\nu_\ell$. One can use the textbook properties of $f_\Gamma(x;k,\theta)$ to derive the statistics of the probability density function of $\hat{C}_\ell$: its expectation value is

$$\left\langle \hat{C}_\ell \right\rangle = k\theta = C_\ell, \tag{9.5}$$

its variance, as already emphasised by Eq. 2.49, is

$$\sigma^2(\hat{C}_\ell) = k\theta^2 = \frac{2}{\nu_\ell} C_\ell^2, \tag{9.6}$$

its skewness is

$$\mathrm{skew}(\hat{C}_\ell) = \frac{2}{\sqrt{k}} = \sqrt{\frac{8}{\nu_\ell}}, \tag{9.7}$$

and its mode is

$$\mathrm{mode}(\hat{C}_\ell) = (k-1)\theta = C_\ell\left(1-\frac{2}{\nu_\ell}\right). \tag{9.8}$$

When ν_ℓ increases, the χ^2 distribution tends to a Gaussian distribution. In our context, $f_{\hat{C}_\ell}(C_\ell) \underset{\ell\to\infty}{\to} \mathcal{N}(\hat{C}_\ell; C_\ell, \frac{2}{\nu_\ell}C_\ell^2)$. Consequently, the higher moments of the distribution vanish, like the skewness, and the mode coincides with the expectation value and the theoretical C_ℓ.

In the more complex case of multiple correlated fields, typically the temperature T and two polarisation modes E and B, each element of the data vector $\mathbf{A}_{\ell m}$ given in Eq. 2.27 follows a multivariate Gaussian distribution $\mathbf{A}_{\ell m} \sim \mathcal{N}_3(0, \mathbf{C}_\ell)$ where $\mathbf{C}_\ell$ is given in Eq. 2.33. A straightforward generalisation of Eq. 2.47 for the estimation of the full-sky covariance matrix $\mathbf{C}_\ell$ reads

$$\hat{\mathbf{C}}_\ell = \frac{1}{2l+1}\sum_m \mathbf{A}_{\ell m}\mathbf{A}_{\ell m}^\dagger = \begin{pmatrix} \hat{C}_\ell^{\mathrm{TT}} & \hat{C}_\ell^{\mathrm{TE}} & \hat{C}_\ell^{\mathrm{TB}} \\ \hat{C}_\ell^{\mathrm{TE}} & \hat{C}_\ell^{\mathrm{EE}} & \hat{C}_\ell^{\mathrm{EB}} \\ \hat{C}_\ell^{\mathrm{TB}} & \hat{C}_\ell^{\mathrm{EB}} & \hat{C}_\ell^{\mathrm{BB}} \end{pmatrix}, \tag{9.9}$$

which, for each ℓ, follows the generalisation of the Gamma distribution for several dimensions, the Wishart distribution: $W_p(\mathbf{V}_\ell, n)$ with $p=3$, $\mathbf{V}_\ell = \frac{\mathbf{C}_\ell}{\nu_\ell}$ and $n = \nu_\ell$. Its probability density

function is

$$f_{\mathbf{C}_\ell}\left(\hat{\mathbf{C}}_\ell\right) = \frac{|\mathbf{C}_\ell|^{(\nu_\ell-p-1)/2}\exp\left[-\operatorname{Tr}\left(\nu_\ell\hat{\mathbf{C}}_\ell^{-1}\mathbf{C}_\ell/2\right)\right]}{2^{p\nu_\ell/2}\left|\hat{\mathbf{C}}_\ell/\nu_\ell\right|^{\nu_\ell/2}\Gamma_p(\nu_\ell/2)}, \tag{9.10}$$

where Γ_p is the multivariate Gamma function,

$$\Gamma_p(\nu_\ell/2) = \pi^{p(p-1)/4}\prod_{i=1}^{p}\Gamma[(\nu_\ell+1-i)/2]. \tag{9.11}$$

One recovers Eq. 9.4 from Eq. 9.10 when a single field is considered. In particular, the distribution of $\hat{C}_\ell^{BB}$, which is uncorrelated to the temperature and E-modes, reads $f_{\mathbf{C}_\ell}(\hat{C}_\ell^{BB}) = f_{C_\ell^{BB}}(\hat{C}_\ell^{BB})$. The marginal distribution of $\hat{C}_\ell^{TT}$ and $\hat{C}_\ell^{EE}$ also reduces to a Gamma distribution, and more complicated marginal distributions for the cross-spectra can also be derived analytically. The generalisation of the single-field expectation value expression Eq. 9.5 reads

$$\left\langle\hat{\mathbf{C}}_\ell\right\rangle = n\mathbf{V}_\ell = \mathbf{C}_\ell, \tag{9.12}$$

while the generalisation of the variance Eq. 9.6 expression for the vector $\tilde{\hat{\mathbf{C}}}_\ell$ (defined similarly as in Eq. 2.50 substituting C_ℓ^{XY} by $\hat{C}_\ell^{XY}$) gives $\hat{\mathbf{\Xi}}_\ell$ (defined similarly as Eq. 2.51 substituting C_ℓ^{XY} by $\hat{C}_\ell^{XY}$). Similarly to the single-field case, the distribution Eq. 9.10 tends to become a multivariate Gaussian at high multipoles.

We drop the normalisation factor of the likelihood Eq. 9.2 (which is independent of the model) and fix it so the value of the likelihood is one when the data corresponds to the model. It leads to writing the log-likelihood for independent ℓ corresponding to Eq. 9.10, which we hereafter call $\mathcal{L}_{\text{Wishart}}$, as [224, 225]

$$-2\ln\mathcal{L}_{\text{Wishart}}\left(\{\theta_i\}\mid\left\{\hat{\boldsymbol{C}}_\ell\right\}\right) = \sum_\ell(2\ell+1)\left\{\operatorname{Tr}\left[\hat{\boldsymbol{C}}_\ell\boldsymbol{C}_\ell(\{\theta_i\})^{-1}\right] - \ln\left|\hat{\boldsymbol{C}}_\ell\boldsymbol{C}_\ell(\{\theta_i\})^{-1}\right| - n\right\}. \tag{9.13}$$

9.2.3. Cut-sky spectrum-space

When working with maps that have been partially masked, pixel-based methods still provide the exact likelihood. However, once again, this approach is almost always inapplicable to real cases because it is highly sensitive to the noise covariance matrix knowledge.

Furthermore, Eq. 9.13 is not exact anymore on cut-sky for two reasons. First, the power spectrum at any multipole does not trivially follow a Gamma distribution (like Eq. 9.4) anymore (instead, QML spectra follow the more complex generalised Gamma distribution). Second, correlations between ℓ, which, as discussed in Chapter 8, arise from any power-spectrum estimation from a partial sky, are not included in Eq. 9.13.

A standard way to introduce correlations between the data in the likelihood is to work with a quadratic form, $\ln\mathcal{L}_{\text{quad}}$, in which the covariance matrix appears:

$$-2\ln\mathcal{L}_{\text{quad}}\left(\{\theta_i\}\mid\left\{\hat{\boldsymbol{C}}_\ell\right\}\right) = \mathbf{X}_\ell(\hat{\boldsymbol{C}}_\ell,\{\theta_i\})\left[\mathbf{\Xi}(\{\theta_i\})^{-1}\right]^{\ell\ell'}\mathbf{X}_{\ell'}(\hat{\boldsymbol{C}}_\ell,\{\theta_i\}). \tag{9.14}$$

where $\mathbf{X}_\ell$ is a function of $\hat{\boldsymbol{C}}_\ell$ and $\boldsymbol{C}_\ell(\{\theta_i\})$ to be specified. *A priori*, $\mathbf{\Xi}_{\ell\ell'} = \left\langle\mathbf{X}_\ell\mathbf{X}_{\ell'}^T\right\rangle$ depends on the cosmology, and an additional term arises because of the normalising determinant, which

is often ignored. Instead, one frequently uses a fiducial covariance matrix computed assuming a fixed cosmology. Hereafter, when not specified otherwise, we will follow this approximation and will further comment on it in Sec. 9.4.1. Moreover, since we will focus on single-field applications (*EE* and *BB* spectra independently), we simplify the notations by considering only one field. (One can refer to the notations in [225] to generalise the equations given hereafter.) With these simplifications, the quadratic-form log-likelihood reads

$$-2\ln\mathcal{L}_{\rm quad}\left(\{\theta_i\}\mid\left\{\hat{C}_\ell\right\}\right)=X_\ell(\hat{C}_\ell,\{\theta_i\})\left[\Xi^{-1}\right]^{\ell\ell'}X_{\ell'}(\hat{C}_\ell,\{\theta_i\}). \tag{9.15}$$

Let's introduce the different choices of X_ℓ that lead to the likelihood functions we will apply in this work: the Gaussian likelihood, the Hamimeche-Lewis likelihood (HL), and the offset-Hamimeche-Lewis likelihood (oHL).

9.2.3.1 Gaussian likelihood

A first simple choice of X_ℓ consists in:

$$X_\ell(\hat{C}_\ell,\{\theta_i\})=\hat{C}_\ell-C_\ell(\{\theta_i\}) \tag{9.16}$$

This choice leads to a likelihood commonly called *Gaussian likelihood* $\mathcal{L}_{\rm Gauss}$. In this case, the matrix $\Xi_{\ell\ell'}$ is usually simply the covariance matrix of a $C_\ell\left(\{\theta_i\}^{\rm fid}\right)$. This approximation is effective for high multipoles where the likelihood becomes more Gaussian, but it is a poor approximation for extracting cosmological information from the largest scales, as we will verify later. This approximation has been used for large-ℓ analyses in several works, *eg.* [102, 108, 226, 227].

9.2.3.2 Hamimèche-Lewis likelihood

An idea introduced in [225] consists of rewriting the exact log-likelihood given in Eq. 9.13 as a quadratic form. It leads to the following choice of X_ℓ in the single-field case:

$$X_\ell(\hat{C}_\ell,\{\theta_i\})=C_\ell\left(\{\theta_i\}^{\rm fid}\right)g\left[\frac{\hat{C}_\ell}{C_\ell(\{\theta_i\})}\right], \tag{9.17}$$

with

$$g(x)\equiv{\rm sign}(x-1)\sqrt{2(x-\ln(x)-1)}. \tag{9.18}$$

The quadratic form Eq. 9.15 applied to these X_ℓ is called *Hamimeche-Lewis likelihood* (HL), $\mathcal{L}_{\rm HL}$. Applied to a full-sky (diagonal) fiducial covariance matrix that simply reads $\Xi_{\ell\ell'}=\frac{2C_\ell\left(\{\theta_i\}^{\rm fid}\right)^2}{2\ell+1}\delta_{\ell\ell'}$ in the auto-single-field case, this likelihood leads to

$$-2\ln\mathcal{L}_{\rm HL}\left(\{\theta_i\}\mid\left\{\hat{C}_\ell\right\}\right)=C_\ell\left(\{\theta_i\}^{\rm fid}\right)g\left[\frac{\hat{C}_\ell}{C_\ell(\{\theta_i\})}\right]\times\left[\Xi^{-1}\right]^{\ell\ell'}\times C_{\ell'}\left(\{\theta_i\}^{\rm fid}\right)g\left[\frac{\hat{C}_{\ell'}}{C_{\ell'}(\{\theta_i\})}\right] \tag{9.19}$$

$$=\sum_\ell\frac{2\ell+1}{2C_\ell\left(\{\theta_i\}^{\rm fid}\right)^2}\times C_\ell\left(\{\theta_i\}^{\rm fid}\right)^2\times2\left(\frac{\hat{C}_\ell}{C_\ell}-\ln\left|\frac{\hat{C}_\ell}{C_\ell}\right|-1\right) \tag{9.20}$$

$$=-2\ln\mathcal{L}_{\rm Wishart}. \tag{9.21}$$

In fact, by construction, the HL likelihood applied to the entire sky matches perfectly the exact likelihood.

Moreover, Eq. 9.15 is now able to account for non-diagonal terms in the covariance matrix, resulting in an approximation on cut-sky that gets closer and closer to the exact likelihood as the sky fraction increases. This can be generalised to multi-field scenarios, as demonstrated in [225], resulting in a likelihood widely applied to CMB analyses [102, 108, 226, 227]. It is a very good approximation for high signal-to-noise regimes, and it is computationally equivalent to the Gaussian likelihood. For small masks, by construction, it provides better results than the Gaussian approximation for both small and large scales.

In Fig. 9.3, we represent the distributions of X_ℓ^{EE} in the Gaussian-likelihood case (blue), given by Eq. 9.16 and in the HL case (orange), given by Eq. 9.17. These distributions are estimated from 50 000 noiseless full-sky auto-power-spectrum simulations. A misinterpretation of the HL transform is to view it as a transformation of $\hat{C}_\ell$ distributions into Gaussian distributions centered on 0, to be used in a Gaussian likelihood. However, the orange distributions in Fig. 9.3 are not Gaussian distributions, nor are they centered on 0. This point is further emphasised and explained in App. D in the full-sky case.

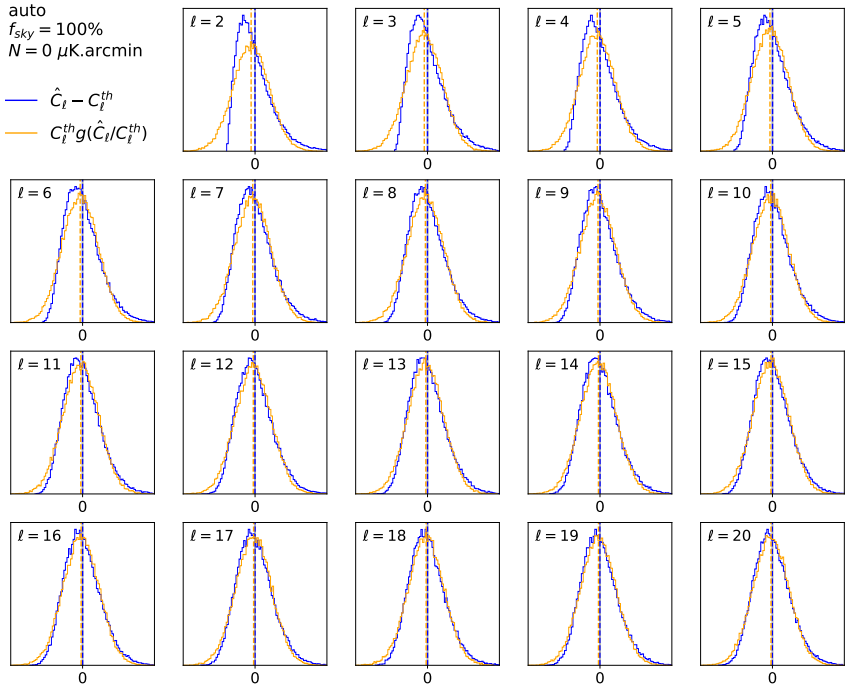

Figure 9.3: Distribution of 50 000 simulated X_ℓ^{EE} in the Gaussian-likelihood case, given by Eq. 9.16 (blue) and in the HL case, given by Eq. 9.17 (orange). The set-up is noiseless auto power spectra estimated on full-sky, from our usual fiducial Tab. 4.1. We superimpose the means of both distributions with corresponding colours.

A limit of this likelihood happens for cross spectra in low signal-to-noise regime where the dispersion of $\frac{\hat{C}_\ell}{C_\ell}$ around one increases and leads to invalid Eq. 9.18. A second scenario that can also lead to negative auto $\hat{C}_\ell$ is when the noise is over-corrected, which can happen in low signal-to-noise regime by statistical fluctuations. A way of solving this issue consists in replacing

Eq. 9.18 by [228]

$$g(x) \equiv \mathrm{sign}(x)\,\mathrm{sign}(|x|-1)\sqrt{2(|x|-\ln|x|-1)}, \tag{9.22}$$

which still leads to the exact likelihood in full-sky.

However, the trick of Eq. 9.22 alone is imperfect and can lead to a strongly biased likelihood if too many $\hat{C}_\ell$ are negative, as we will verify in the concrete cases studied in Sec. 9.3 and Sec. 9.4. Instead, for low signal-to-noise regime, a complementary solution has to be adopted.

9.2.3.3 Offset Hamimèche-Lewis likelihood

A second way of solving the issue of negative $\hat{C}_\ell$ in the Hamimèche-Lewis likelihood is to include offsets in X_ℓ [228]:

$$X_\ell(\hat{C}_\ell, \{\theta_i\}) = \left[C_\ell\left(\{\theta_i\}^{\mathrm{fid}}\right) + O_\ell\right] g\left[\frac{\hat{C}_\ell + O_\ell}{C_\ell(\{\theta_i\}) + O_\ell}\right]. \tag{9.23}$$

Previously in the literature, offsets have been typically introduced the context of cross-spectrum-based Hamimèche-Lewis likelihood, [61, 66, 228, 229], derived from simulations in order to mimic the noise bias present in auto-spectra.

The motivation behind our choice of offset is similar. However, we derive a new formula starting from a slightly different requirement. We want the Hamimèche-Lewis transform to be applied to a $\hat{C}_\ell + O_\ell$ whose distribution looks as close as possible to the noisy full-sky auto-spectra one. Let us assume that we have a map generated from a noisy model $C_\ell + N_\ell$. We have seen that the resulting auto spectrum estimated on full-sky, $\hat{C}_\ell$, follows a Γ distribution, with expectation value, variance and skewness given in Eqs. 9.5, 9.5 and 9.7, where we have to perform $C_\ell \to C_\ell + N_\ell$.

Instead, if one cuts the same map and estimates the noisy cross-spectrum $\tilde{C}_\ell$ from it, its distribution is no longer biased by N_ℓ, neither it is trivially a Γ distribution. However, one can still assume that for some effective degree of freedom $\tilde{\nu}_\ell$, the distribution approximately matches a Γ distribution[3]. In this scenario, the variance and skewness of the distribution are also given by Eqs. 9.5 and 9.7, substituting ν_ℓ by an effective number of degrees of freedom $\tilde{\nu}_\ell$, and C_ℓ by $C_\ell + N_\ell$.

Hence, to match the $\tilde{C}_\ell$ distribution with a Γ distribution with degree of freedom $\tilde{\nu}_\ell$, we only have to offset $\tilde{C}_\ell$ by N_ℓ. If one knows from simulations the expectation value, variance and skewness of $\hat{C}_\ell$, from the considerations we have given, the offset can easily be deduced:

$$O_\ell = N_\ell = 2\frac{\sigma(\tilde{C}_\ell)}{\mathrm{skew}(\tilde{C}_\ell)} - \left\langle \tilde{C}_\ell \right\rangle, \tag{9.24}$$

alongside with the effective number of degrees of freedom,

$$\tilde{\nu}_\ell = \frac{8}{\mathrm{skew}(\tilde{C}_\ell)^2}. \tag{9.25}$$

[3]To verify this hypothesis, we have fitted such noisy-cross-spectrum distributions by a Γ, see Fig. 9.4 that we comment in the upcoming Sec. 9.2.3.4.

Hereafter, we call *offset-Hamimècle-Lewis likelihood* (oHL) the likelihood approximation obtained by applying the particular offsets (Eq. 9.24 estimated over simulations) to the Hamimèche-Lewis X_ℓ (Eq. 9.15) themselves applied to the quadratic form (Eq. 9.23).

All three quadratic forms we have introduced so far are approximate on the cut sky. We will compare their performances in the next sections, Secs. 9.3 and 9.4. To be possibly even closer to the exact likelihood, a solution is to fully commit to building our approximation using some knowledge on the shape of the ℓ-per-ℓ probability density function $f_{C_\ell}(\hat{C}_\ell)$.

9.2.3.4 Copula likelihood

In this spirit, we present here how we adapt the so-called *copula likelihood* to the cosmological-parameter inference problem from cut-sky low-ℓ polarisation spectra. The copula likelihood has been introduced to correct the correlated Gaussian likelihood by including information on the shape of the independent likelihoods data-point-per-data-point [230]. This information can *a priori* be obtained by any means, for instance, by fitting the shape of the distribution of data simulations. The likelihood has already been applied to astrophysical and cosmological problems [231–234], as well as to CMB low-ℓ temperature [235]. In the CMB context, denoting $\mu_\ell \equiv \left\langle \hat{C}_\ell \right\rangle$ and $\sigma_\ell^2 \equiv \sigma^2(\hat{C}_\ell)$, the likelihood reads

$$-2\ln\mathcal{L}_{\rm cop}\left(\{\theta_i\} \mid \{\hat{C}_\ell\}\right) = \sum_{\ell\ell'}(q_\ell-\mu_\ell)\left[\Xi^{-1}\right]_{\ell\ell'}(q_{\ell'}-\mu_{\ell'}) - \sum_\ell \frac{(q_\ell-\mu_\ell)^2}{\sigma_\ell^2} - 2\sum_\ell \ln f_{C_\ell(\{\theta_i\})}\left(\hat{C}_\ell\right) \tag{9.26}$$

where $f_{C_\ell(\{\theta_i\})}\left(\hat{C}_\ell\right)$ is the *probability density function* of $\hat{C}_\ell$ given the theory, Ξ is the covariance matrix to be estimated over a set of fiducial simulations $\hat{C}_\ell$, and q_ℓ is defined as

$$q_\ell = \sigma_\ell \Psi_1^{-1}\left[F_{C_\ell(\{\theta_i\})}\left(\hat{C}_\ell\right)\right] + \mu_\ell, \tag{9.27}$$

where Ψ_1 is the *cumulative standard normal density function*, whose inverse is the so-called *percent point function*, and $F_{C_\ell(\{\theta_i\})}$ is the *cumulative density function* of $\hat{C}_\ell$ given the theory.

Notice that if the data are distributed according to a Gaussian distribution, only the first term remains, and $\mathcal{L}_{\rm cop}$ reduces to the correlated Gaussian likelihood $\mathcal{L}_{\rm gauss}$. In the case where there is no correlation between the data, the two first terms cancel out, and we recover the definition of the likelihood function for independent data. Especially, in the full-sky CMB context, $\mathcal{L}_{\rm cop}$ reduces to $\mathcal{L}_{\rm Wishart}$, and hence is exact.

We aim to assess the performance of this likelihood on cut-sky in our specific context alongside the quadratic forms. In practice, we estimate $f_{C_\ell(\{\theta_i\})}$ by fitting the $\hat{C}_\ell$ distributions, themselves estimated from maps generated by an input theoretical C_ℓ. The fit is done by a shifted/scaled gamma distribution:

$$f_{\tilde{\Gamma}}(x;k,\mathrm{loc},\mathrm{scale}) \equiv \frac{1}{\mathrm{scale}}\Gamma\left(\frac{x-\mathrm{loc}}{\mathrm{scale}};k\right) \tag{9.28}$$

with an additional constraint being that $\left\langle \hat{C}_\ell \right\rangle = C_\ell$ (*i.e.* the power spectrum estimation methods are in principle unbiased). This leads to a 2-parameter distribution that can be viewed as an offset-Γ distribution with an offset given by $-\mathrm{loc}$, an effective number of degree of freedom of $2k$, and where the scale is fixed by the unbiasedness constraint[4].

[4]Instead of numerically fitting the distributions, similarly to what we have done for determining the

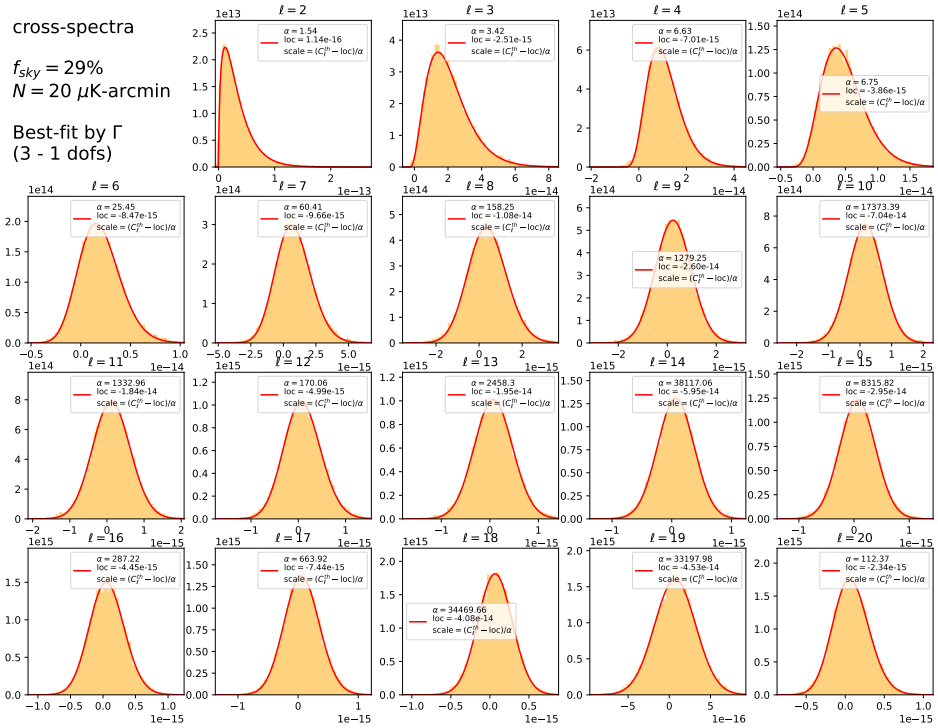

Figure 9.4: Distribution of 20 000 $\hat{C}_\ell^{EE}$ (in K^2) cross spectra from $f_{\rm sky} = 29\%$ maps containing a 20 μK·arcmin white noise. The distributions are superimposed with the best fit by a two-parameter offset-Γ distribution (here α stands for $k = \tilde{\nu}_\ell/2$).

These fits' results are presented in Fig. 9.4. It represents cross $\hat{C}_\ell^{EE}$ obtained from 2×20 000 simulated $f_{\rm sky} = 29\%$ maps containing a $\sigma = 20$ μK·arcmin white noise. We superimpose in red a Γ distribution whose parameters are fitted on the distributions. One observes that the fits yield remarkably good results. The setup illustrated in this figure depicts the worst-case scenario of low signal-to-noise in low sky fraction, and the other setups yield results that are at least as good.

We also present the fitted effective number of degrees of freedom on the left panel of Fig. 9.5, as compared to the number of m per ℓ, which is $2\ell + 1$. In the figure, the three sets of points correspond to three setups: QML cross spectra on $f_{\rm sky} = 29\%$ and $\sigma = 20$ μK·arcmin (blue), QML cross spectra on $f_{\rm sky} = 62\%$ and $\sigma = 20$ μK·arcmin (orange), and noiseless full-sky (red). One notices the deviation from the expected formula with the increasing complexity of the simulation. The effective number of freedom increases with the fraction of removed sky, *i.e.* the

Hamimèche-Lewis offsets, we can directly determine the Γ parameters from the only knowledge of the distribution moments. Textbook relations lead to:

$$k_\ell = 4/\mathrm{skew}(\hat{C}_\ell)^2 \tag{9.29}$$

$$\mathrm{scale}_\ell = \sigma(\hat{C}_\ell)|\mathrm{skew}(\hat{C}_\ell)|/2 \tag{9.30}$$

$$\mathrm{loc}_\ell = \left\langle \hat{C}_\ell \right\rangle - 2\sigma(\hat{C}_\ell)/|\mathrm{skew}(\hat{C}_\ell)| \tag{9.31}$$

However, we will stick hereafter to using the fitted parameter because we find that they allow a slightly better fit to the actual distribution.

distributions are more symmetric at low $f_{\rm sky}$. This observation can be explained by the fact that at low $f_{\rm sky}$, multipoles mix in the power spectrum estimation (see Eqs. 8.16 and 8.18 for QML), facilitating the reach of the central limit theorem.

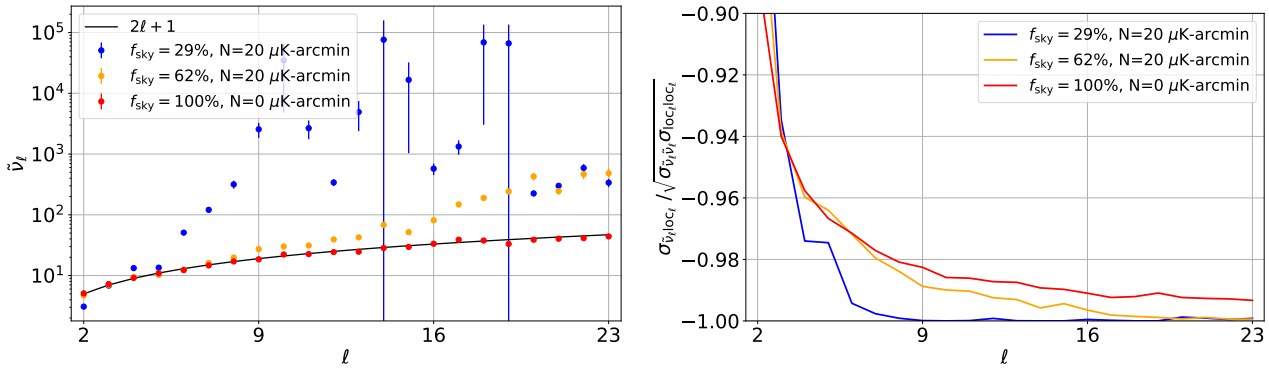

Figure 9.5: LEFT: Effective degrees of freedom fitted on the $\hat{C}_\ell$ distributions ℓ per ℓ and their uncertainty, for a noise of 20 μK·arcmin and $f_{\rm sky}$=29% (blue), for a noise of 20 μK·arcmin and $f_{\rm sky}$=62% (orange), and for no noise and $f_{\rm sky}$=100% (red), as compared to the standard full-sky theoretical formula Eq. 2.49 (black line). RIGHT: Correlation between the fitted effective degrees of freedom and the fitted loc, with the same colours corresponding to the same scenarios as on the left panel.

A limitation of the Γ fit is a degeneracy between the fitted parameters that becomes important at low $f_{\rm sky}$. We illustrate on the right panel of the same figure the correlation between the estimated loc and $\tilde{\nu}_\ell$. This significant degeneracy between the Γ parameter explains the huge errorbars at low $f_{\rm sky}$ on the left panel of Fig. 9.5. For now, we neglect this effect, which, we expect, can be solved by adopting an offset lognormal distribution that is also able to correctly fit the distributions. That being said, this is not a big issue for the copula application: the only actual requirement is a good fit of the $\hat{C}_\ell$ probability functions (which is achieved, see Fig. 9.5).

We need to specify their dependence on the model to use these fitted parameters as inputs for the copula likelihood. In the full-sky setup, combining Eqs.9.29, 9.30, 9.31 with Eqs. 9.5, 9.6, 9.7 yields the conclusion that

$$k_\ell \propto \tilde{\nu}_\ell \propto 1, \tag{9.32}$$
$$\mathrm{scale}_\ell \propto C_\ell, \tag{9.33}$$
$$\mathrm{loc}_\ell \propto C_\ell. \tag{9.34}$$

This remains a good approximation in the more complex setup (cut-sky, cross-spectra, noise): for illustration, in a $\sigma = 20$ μK·arcmin and $f_{\rm sky} = 29\%$ case, we show in Fig. 9.6 in red-dotted lines the fit done on $\hat{C}_\ell$ distributions from $\tau = 0.0544$, as compared with the $\hat{C}_\ell^{EE}$ distribution at $\tau = 1.1 \times 0.0544$. See Tab. 4.1 and Fig. 4.10 for more details on the low-ℓ *EE* fiducial. Without surprise, one observes that the two do not match. Instead, the blue curve represents the modified Γ distribution whose parameters are the scale, loc and k fitted on the $\tau = 0.0544$ distributions for which the scale and loc are multiplied by $C_\ell(\tau = (1 \pm 0.1) \times 0.0544)/C_\ell(\tau = 0.0544)$. This rescaling allows the recovery of a good match with the distribution without the need for an

actual fitting. We have verified that this remains true when going even further than 10% from the baseline τ, up to 60%.

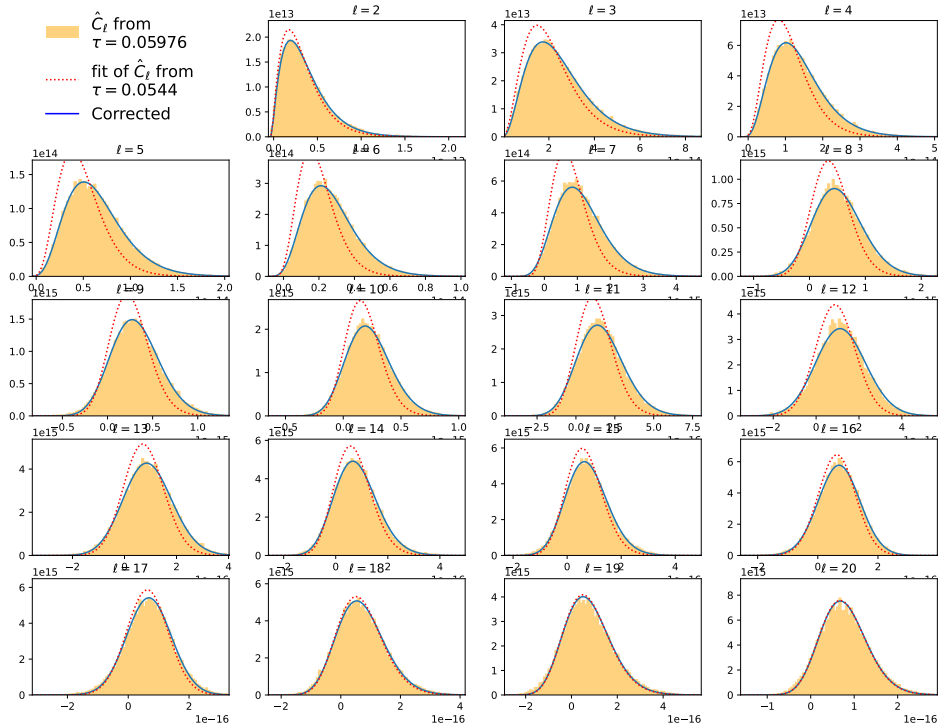

Figure 9.6: For $\sigma = 20$ μK·arcmin and $f_{\rm sky} = 29\%$: the red-dotted lines illustrate the fit done on $\hat{C}_\ell$ distributions from $\tau = 0.0544$, as compared with the $\hat{C}_\ell^{EE}$ distribution from $\tau = 1.1 \times 0.0544$ (right). The blue curve represents the modified Γ distribution whose parameters are the scale, loc and k fitted on the $\tau = 0.0544$ distributions and the scale and loc are rescaled by $C_\ell(\tau = (1 \pm 0.1) \times 0.0544)/C_\ell(\tau = 0.0544)$.

Hence, we only need to perform and use one fiducial fit per analysis setup (times the number of ℓ's), and no refitting is required for exploring the copula likelihood. Of course, the theory appears in μ_ℓ, but also implicitly in $f_{C_\ell(\{\theta_i\})}$ and $F_{C_\ell(\{\theta_i\})}$ that both depend on the rescaled Γ parameters. To remain similar to the setup of the three quadratic forms introduced before, we do not include the dependence of the model in the covariance nor in σ_ℓ.

In this section, we have introduced several likelihoods approximations. Hereafter, our goal will be to assess their performance for recovering τ and r from EE and BB power-spectra estimating on cut-sky, based on their maximum-likelihood estimate.

9.3. Application to the white-noise case and results

Before applying the various likelihood approximations to the LiteBIRD simulations introduced and processed in Sections III, 7.3, and 8.3, we first test them in a controlled white noise case. The practical reason for this is to enable us to evaluate the exact pixel-based likelihood and use its result as a reference case, facilitating the other methods' assessment, especially their performance severely masking sky, their robustness to a poor fiducial choice or to a low signal-to-

noise. This discussion will also lead us to introduce a criterion that doesn't require a reference, the pull, and which will be especially useful in the LiteBIRD context.

9.3.1. τ estimation from white-noise case

First, we focus on the reconstruction of the τ information from the *EE* reionisation bump. To do so, we only need the lowest multipoles of our hybrid power-spectrum estimator, QML.

9.3.1.1 Different setups

For this discussion, we generate 20 000 $N_{\rm side} = 8$ map simulations and associated QML spectra in each combination of the following setups:

Input theory: Seven input theory configurations

- *base*: $\Omega_b h^2$, $\Omega_c h^2$, H_0, $n_{\rm s}$, $A_{\rm s}$, τ, r given in Table 4.1.
- *low*: $\tilde{\Omega}_b h^2$, $\tilde{\Omega}_c h^2$, $\tilde{H}_0$, $\tilde{n}_{\rm s}$, $\tilde{A}_{\rm s}$ = 0.99 $\Omega_b h^2$, 0.99 $\Omega_c h^2$, 0.99 H_0, 0.99 $n_{\rm s}$, 0.99 $A_{\rm s}$
- *high*: $\tilde{\Omega}_b h^2$, $\tilde{\Omega}_c h^2$, $\tilde{H}_0$, $\tilde{n}_{\rm s}$, $\tilde{A}_{\rm s}$ = 1.01 $\Omega_b h^2$, 1.01 $\Omega_c h^2$, 1.01 H_0, 1.01 $n_{\rm s}$, 1.01 $A_{\rm s}$
- *lowtau*: $\tilde{\tau} = 0.9\ \tau$ and $\tilde{A}_{\rm s} = A_{\rm s} \exp\{-2\tau\} \exp\{2\tilde{\tau}\}$
- *hightau*: $\tilde{\tau} = 1.1\ \tau$ and $\tilde{A}_{\rm s} = A_{\rm s} \exp\{-2\tau\} \exp\{2\tilde{\tau}\}$
- *verylowtau*: $\tilde{\tau} = 0.4\ \tau$ and $\tilde{A}_{\rm s} = A_{\rm s} \exp\{-2\tau\} \exp\{2\tilde{\tau}\}$
- *veryhightau*: $\tilde{\tau} = 1.6\ \tau$ and $\tilde{A}_{\rm s} = A_{\rm s} \exp\{-2\tau\} \exp\{2\tilde{\tau}\}$

These will allow robustness checks over the choice of the fiducial model for estimating the covariance matrix and for fitting the distributions. The choice of the *low* and *high* models is driven by the current per cent uncertainty we have on the various Λ-CDM parameters, while the choice of *lowtau* and *hightau* models is driven by the current 10% uncertainty we have on τ. *verylowtau* and *veryhightau* aim to extend this study in more problematic cases.

Mask: Eight galactic mask (at $N_{\rm side} = 8$) configurations:

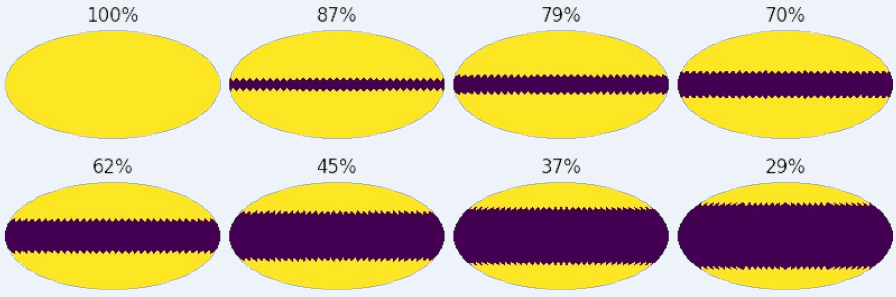

We chose to work with galactic masks as a reference case.

Noise: Two white noise levels:

- *low SNR*: $\sigma = 20\ \mu$K·arcmin
- *high SNR*: $\sigma = 0.01\ \mu$K·arcmin

The *high SNR* noise simulations aim at testing the likelihood in an ideal case, while the *low SNR* regime is representative of the worse case scenarios of noise we will encounter post component separation (see Tab. 7.1).

cross or auto:

- *cross*
- *auto*

We will explore both configurations, QML *auto* and *cross* spectra. Notice that the adopted multipole range is $2 \leq \ell \leq 31$ for τ inference and $2 \leq \ell \leq 200$ for r inference.

covariance: Two setups for the covariance computation:

- *empirical*: empirical from the simulations
- *analytical*: analytical `xQML` derivation

In the known white-noise scenario, we can compare the results obtained with an *empirical* covariance matrix with respect to those obtained from an analytical covariance matrix *analytical*.

9.3.1.2 Results for the baseline setup

We first compare the likelihoods for various sky fractions in the baseline case: *base* input theory, *high SNR* regime, with *auto* QML spectra and *empirical* covariance matrix.

We illustrate in Fig. 9.7 the shape of -2 log-likelihood for three realisations of this setup (one per line), and for four different $f_{\rm sky}$, from left to right 29%, 45%, 70%, 100%. This figure illustrates the global shape of the likelihood branches before going into quantitative conclusions. For visual purposes, we offset the curves by 50 and put a logarithmic y-scale. The black curve represents the reference (exact) $\mathcal{L}_{\rm PB}$, $\mathcal{L}_{\rm cop}$ in dashed-red, $\mathcal{L}_{\rm oHL}$ in dashed-dotted-blue, $\mathcal{L}_{\rm HL}$ in dotted-green, $\mathcal{L}_{\rm gauss}$ in orange. The input τ is represented by a vertical grey line.

The exact likelihood has a steep branch at low τ values. The Gaussian likelihood presents an opposite behaviour for all $f_{\rm sky}$, with a sharp branch at high τ. It means that already qualitatively, $\tau = 0$ is excluded with smaller significance if one adopts the Gaussian approximation. Said otherwise, the Gaussian likelihood will lead to a less-stringent 1-σ limit for low-τ and a more stringent one for high-τ, an effect that we quantify in the next paragraph. As expected, the other approximations converge to the exact likelihood on full-sky and only deviate significantly for more than 40%-masked sky. To be more quantitative on the impact of these differences for the maximum-likelihood estimate of τ as well as the estimate of its $\pm 1\sigma$ interval, we estimate for the 20 000 simulations the *maximum of the likelihood* $\hat{\tau}$ as well as $\sigma_\tau^\pm$ defined as

$$\sigma_\tau^- = \hat{\tau} - \tau|_{-2\ln\mathcal{L}/\mathcal{L}_{\rm min}=1 \text{ and } \tau<\hat{\tau}}, \tag{9.35}$$

$$\sigma_\tau^+ = -\hat{\tau} + \tau|_{-2\ln\mathcal{L}/\mathcal{L}_{\rm min}=1 \text{ and } \tau>\hat{\tau}}. \tag{9.36}$$

One should remember in the following that all results we will show are based on this particular way of estimating $\hat{\tau}$ and $\sigma_\tau^\pm$.

We summarise in Fig. 9.8 the averages and standard deviations with respect to $f_{\rm sky}$ of $\hat{\tau}$ (upper panel), σ_τ^- (middle) and σ_τ^+ (bottom). One observes that, as long as $f_{\rm sky} \geq 62\%$, all likelihood

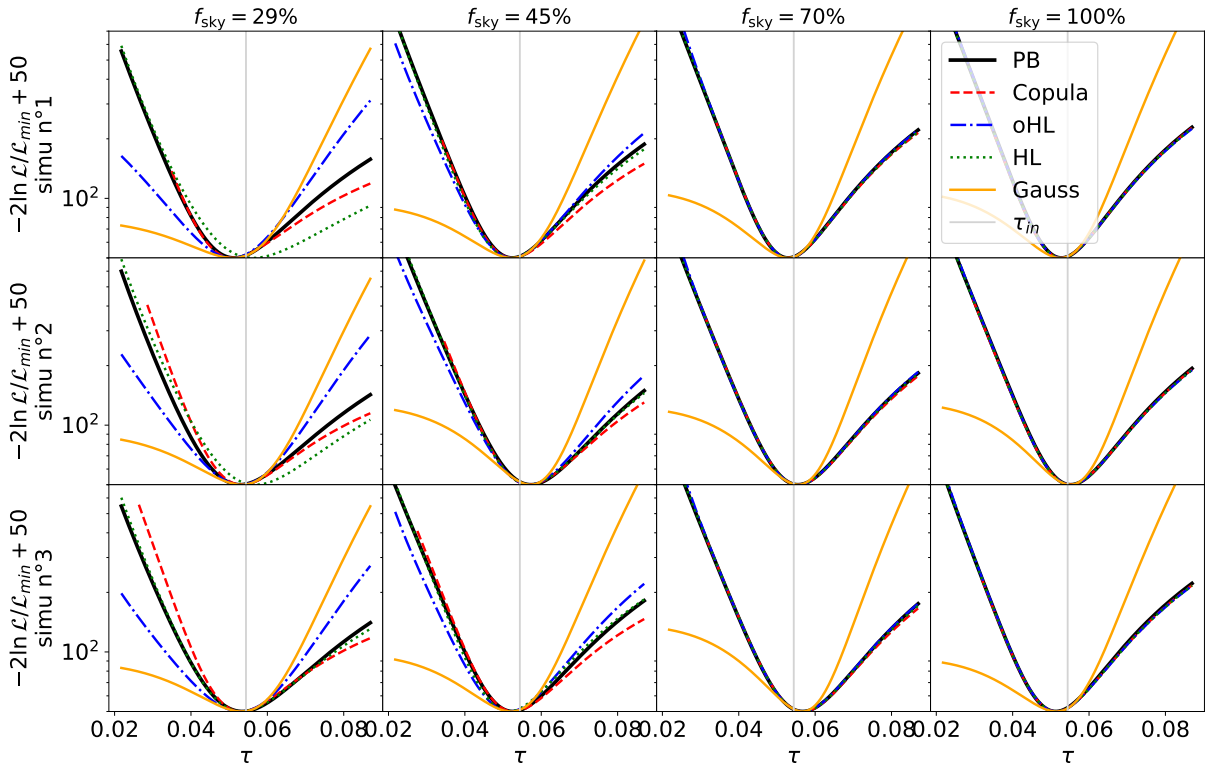

Figure 9.7: Shape of -2 log-likelihood for three different realisations (from top to bottom) and for four different $f_{\rm sky}$, 29%, 45%, 70%, 100% (from left to right). The black curve represents the reference (exact) $\mathcal{L}_{\rm PB}$, $\mathcal{L}_{\rm cop}$ in dashed-red, $\mathcal{L}_{\rm oHL}$ in dashed-dotted-blue, $\mathcal{L}_{\rm HL}$ in dotted-green, $\mathcal{L}_{\rm gauss}$ in orange. The input τ is represented by a vertical grey line. The minimum of the log-likelihood is fixed at 50. The setup is *base* input theory, *low SNR* regime, with *auto* QML spectra and *empirical* covariance matrix described in Sec. 9.3.1.1

approximations give in average similar results for the best τ value and for the $\sigma_\tau^\pm$ estimate, apart the Gaussian approximation which, as anticipated, is the only one yielding $\sigma_\tau^- > \sigma_\tau^+$. We conclude that the estimation of the error on τ from this approximation is significantly miss-estimated (overestimated for σ_τ^- and underestimated for σ_τ^+). For low $f_{\rm sky}$, the HL likelihood becomes significantly biased on average towards high values, while the oHL likelihood becomes slightly biased towards low values (by $\sim -\sigma_\tau/3$ at $f_{\rm sky} = 29\%$). Finally, the copula likelihood remains approximately unbiased on average but comes with error bars slightly off with respect to the optimal ones provided by the pixel-based likelihood. The latter effect also happens in the case of oHL.

An interesting result one obtains from this figure is that $\sigma_\tau^\pm$ doesn't scale as the usual scaling $f_{\rm sky}^{-1/2}$, which is the one we adopt for correcting the Fisher estimate in cut-sky, see Eq. 4.13. Especially, while one would expect to degrade the error on τ by $1 - 0.7^{-1/2} = 20\%$ by cutting 30% of the sky, we find that this error is almost not degraded going from $f_{\rm sky} = 100\%$ to $f_{\rm sky} = 70\%$.

Another noticeable feature is a slight but statistically significant bias towards low τ values, see upper panel, which is common to all likelihoods in full-sky, including to the exact one. In the latter case, this bias is of order 1.1×10^{-4}, *i.e.* 5.6% of the uncertainty due to cosmic

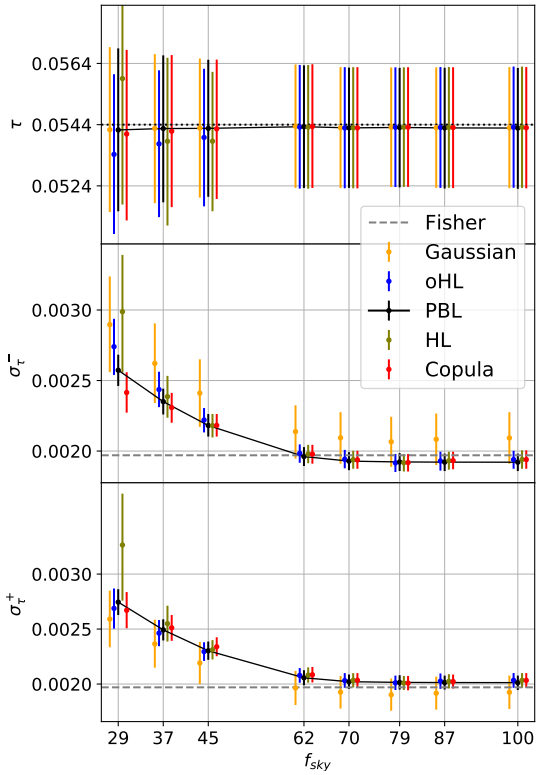

Figure 9.8: Estimate of $\hat{\tau}$ (upper panel), σ_τ^- (middle panel), σ_τ^+ (bottom panel) (see Eqs. 9.35, 9.36), from $\mathcal{L}_{\rm gauss}$ (orange), $\mathcal{L}_{\rm oHL}$ (blue), $\mathcal{L}_{\rm PB}$ (black), $\mathcal{L}_{\rm HL}$ (green), $\mathcal{L}_{\rm cop}$ (red) for varying $f_{\rm sky}$. The setup is *base* input theory, *low* regime, with *auto* QML spectra and *empirical* covariance matrix described in Sec. 9.3.1.1. The errorbars displayed here represent the central region in which lie 68% of the realisations (caution, *a priori* $\sigma(\hat{\tau}) \neq \sigma_\tau^\pm$). We also display the input τ by a dotted black horizontal line on the upper panel, and the Fisher estimate of the noiseless full-sky symmetric error by a dashed grey line on the other panels.

variance. Hence, this low bias does not raise worries for practical inference. However, it remains intriguing. We propose an explanation: the significant skewness of the likelihood, itself caused by the non-asymptotic regime of large-scale CMB data analysis (*i.e.* a limited number of data). In Sec. 4.3, we have mentioned that the consistency of the maximum-likelihood estimator does not guarantee the unbiasedness of an estimate in the non-asymptotic regime. This effect has been extensively studied in specific contexts (*e.g.* for a gamma likelihood [236]), though it is not discussed a lot in the CMB context in which many analyses are Bayesian, sampling and marginalising the likelihood instead of maximising it. Moreover, the τ estimation is particularly problematic since the information is contained in the largest scales where only a few modes bring information. However, even in this case, we find that the non-asymptotic effect is negligible with respect to cosmic variance, suggesting that it can be reasonably ignored.

In Fig. 9.9, to ease the comparison with respect to the pixel-based likelihood estimates, we represent the distributions of the relative errors induced by each method with respect to the exact $\mathcal{L}_{\rm PB}$ one, simulation per simulation. The colour correspondences are the same as the ones used for the previous plots. For $f_{\rm sky} \geq 62\%$, one observes that all likelihoods yield without bias the correct estimate of τ within less than 1% of the actual τ value. This small residual difference

is probably due to a slight loss of information in the power spectrum estimation. Similarly, apart from the Gaussian approximations, all likelihoods yield without bias the correct estimate of $\sigma_\tau^\pm$ within 1%. Instead, for low sky fractions, every τ and $\sigma_\tau^\pm$ estimates significantly lose in reliability.

A summary in terms of bias and dispersion of the Gaussian, HL, oHL and copula estimators with respect to the optimal one (expressed in per cent of the latter) is presented below, for $f_{\rm sky}$=29%. We recall that the cosmic variance on τ imposes a precision significantly smaller than $2\times10^{-3}/0.544 \simeq 3.6\%$ on the τ estimate.

- Gaussian: The Gaussian approximation yields an unbiased estimate of $\hat{\tau}$ at all $f_{\rm sky}$ (within 1% of the τ value). It is the only approximation that does so, which is an impressive property for this apparently naive approximation. However, the Gaussian approximation fails to yield the correct estimation of the uncertainties, its shape deviating too much from the exact one when moving away from the maximum.
- HL: The HL τ-estimates are biased in average (+3% of the τ value) with a distribution significantly dispersed and a long tail, leading for some realisations up to overestimated $\hat{\tau}$ by 15% of the actual τ values. The estimated lower and upper uncertainties are also overestimated on average, by 16% and 19% respectively and with a significant dispersion.
- oHL: The oHL τ-estimates are slightly under-estimated in average (-1.5% of the τ value) with a 1.5% dispersion. The estimated lower uncertainty is over-estimated by 6.6%.
- Copula: The copula τ-estimates are unbiased in average but still present a 2.2% dispersion. Contrary to oHL, this time, the lower uncertainty is under-estimated on average by 6.1%.

To assess the quality of an estimator at being unbiased and providing a correct coverage of the inferred parameters from simulations, and without the need for a reference optimal estimator, a useful criterion is the pulls distribution [237]. For an estimate $\hat{\theta}$ of a true θ_0, associated with asymmetric errorbars $\sigma_\theta^\pm$ (see Eqs. 9.35 and 9.36), the pull g_θ is defined as

$$g_\theta \equiv \frac{\theta_0-\hat{\theta}}{\sigma_\tau^+} \text{ if } \hat{\theta}\leq\theta_0,$$
$$g_\theta \equiv \frac{\theta_0-\hat{\theta}}{\sigma_\tau^-} \text{ if } \hat{\theta}>\theta_0. \tag{9.37}$$

Assuming that we know the coverage of the error interval, *ie.* the fraction of realisations between σ^- and σ^- that is typically 68.27%, a number that we take for reference, then even for skewed likelihood, the fraction of pulls between -1 and +1 should equal this number[5] In the asymptotic limit where the likelihood function reaches Gaussianity, $\sigma_\theta^- = \sigma_\theta^+$ and Eq. 9.37 is simply a standardisation and normalisation, *i.e.* $g_\theta \sim \mathcal{N}(0,1)$. For not-too-skewed likelihood, the pull distribution should be close to a Gaussian distribution, or minimally, 68.27% of the

[5] Proof:

$$\mathrm{P}(-1<g_\theta<1) = \mathrm{P}(-1<g_\theta<0)+\mathrm{P}(0\leq g_\theta<1) \tag{9.38}$$
$$= \mathrm{P}(-\sigma_\theta^-<\theta_0-\hat{\theta}<0)+\mathrm{P}(0\leq\theta_0-\hat{\theta}<\sigma_\theta^+) \tag{9.39}$$
$$= \mathrm{P}(-\sigma_\theta^-<\theta_0-\hat{\theta}<\sigma_\theta^+). \tag{9.40}$$

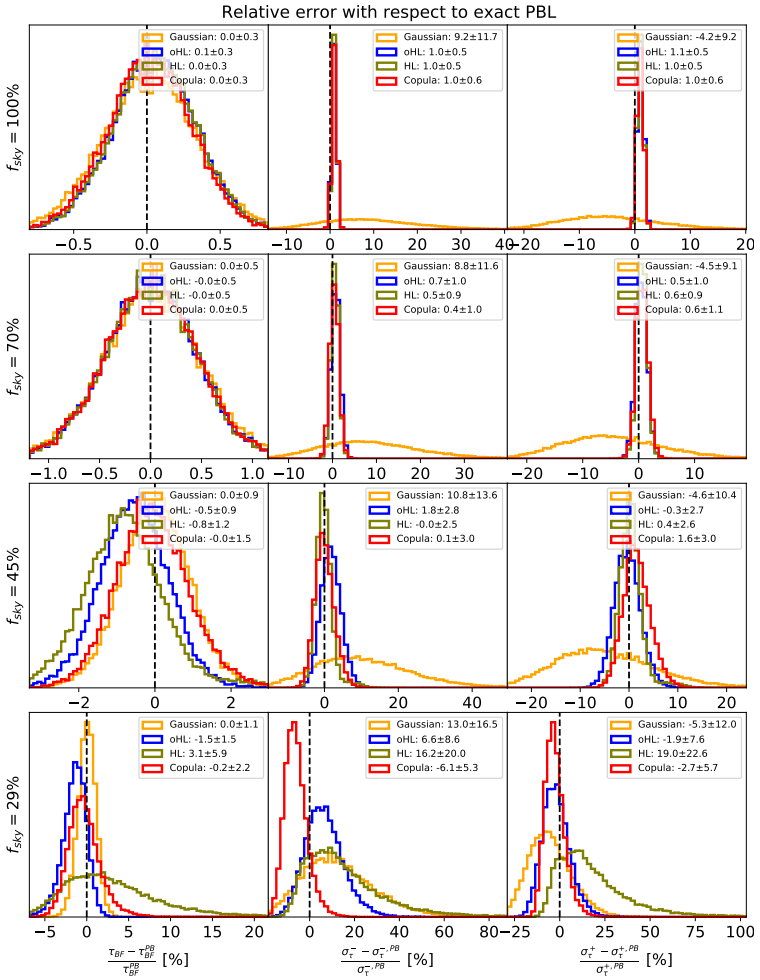

Figure 9.9: Distribution of the relative error of the estimate for $\hat{\tau}$ (upper panel), σ_τ^- (middle panel), σ_τ^+ (bottom panel) from $\mathcal{L}_{\rm gauss}$ (orange), $\mathcal{L}_{\rm oHL}$ (blue), $\mathcal{L}_{\rm HL}$ (green), $\mathcal{L}_{\rm cop}$ (red) with respect to the exact pixel-based estimate, in %. The setup is *base* input theory, *low SNR* regime, with *auto* QML spectra and *empirical* covariance matrix described in Sec. 9.3.1.1

simulations should be between -1 and 1. Moreover, the distribution should not have long tails. Hence, representing the pull distribution is a good habit for verifying the good coverage of an estimator.

We represent in Fig. 9.10 the pull distributions for the various likelihoods on four selected sky fractions, 100%, 70%, 45% and 29%. The best-fit Gaussian parameters, mean and standard deviation, are displayed in the legend for each of the likelihoods considered.

As long as the sky fraction is high, the pull distributions are consistent with a standard normal distribution, with a slight bias with the same origin as the one previously discussed. Especially we verify the good coverage of the simulations by checking that around 68.27% of the simulations present a pull between -1 and 1. Instead, at low $f_{\rm sky}$, say $f_{\rm sky} = 29\%$, the pull distributions for all likelihoods (except the exact one) depart significantly from a standard normal distribution. This hints towards a wrong estimate of the error bars from the likelihood under scrutiny. For instance, the HL distribution being significantly shifted towards low values of the

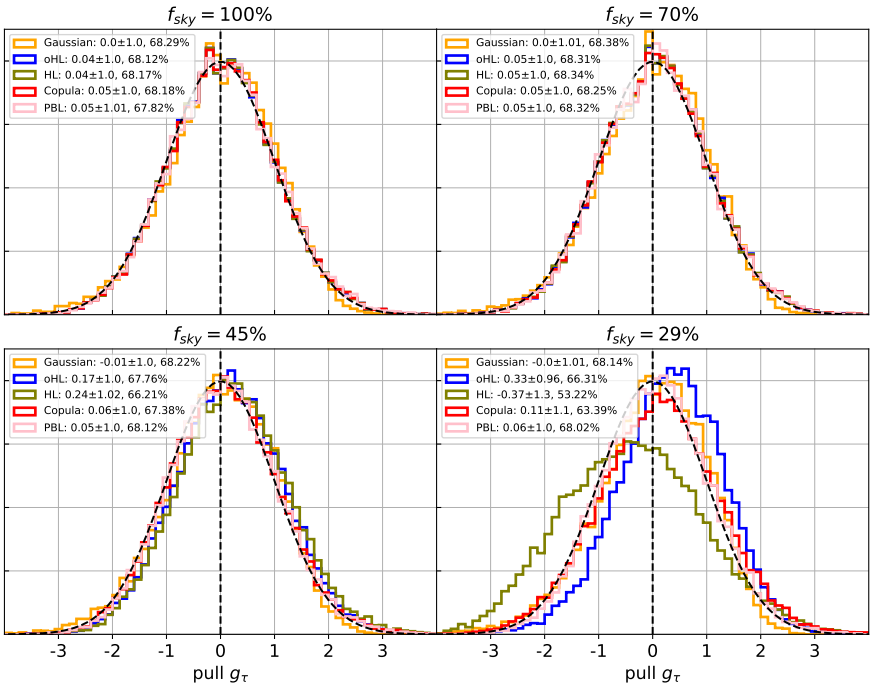

Figure 9.10: Normalised τ-pull distributions for $f_{\rm sky}$ =100, 70, 45, 29% (from top-left to bottom-right), superimposed to the standard normal distribution in black dashed line. g_τ has been estimated from the following likelihoods: $\mathcal{L}_{\rm PB}$ (pink), $\mathcal{L}_{\rm gauss}$ (orange), $\mathcal{L}_{\rm oHL}$ (blue), $\mathcal{L}_{\rm HL}$ (green), $\mathcal{L}_{\rm cop}$ (red). We also display in the legend Gaussian distribution parameters, mean and standard deviation, fitted on the distributions, as well as the fraction of realisations giving a pull between -1 and 1. The setup of the 20 000 simulations is a *base* input theory, *high SNR* regime, with *auto* QML spectra and *empirical* covariance matrix described in Sec. 9.3.1.1.

pull indicates that the τ value is over-estimated while an abnormal dispersion of the distribution due to an underestimated uncertainty. These translate into a wrong coverage with only 53% of the pulls lying in [-1, 1].

Concerning oHL and the copula, our conclusions are essentially similar to the one we already presented during the last discussion based on the comparison to the pixel-based likelihood. This demonstrates the strength of pull distribution tests: we can assess the quality of an estimator, without the need for an optimal reference, solely based on simulations.

9.3.1.3 Effect of noise and cross-spectra

To check the effect of including noise in the simulations and using cross spectra as input of the likelihoods, we redo the analysis in the *low SNR* and *cross* setup. The pull distributions are displayed in Fig. 9.11. They are extremely similar to the ones we have discussed previously in the noiseless auto case, demonstrating that the likelihoods are not significantly affected by the presence of noise and by the use of cross-spectra.

This observation justifies that we focus in the next LiteBIRD sections on cross-spectra-based estimation of τ.

We have also verified that all tests performed in this section yield similar results when taking

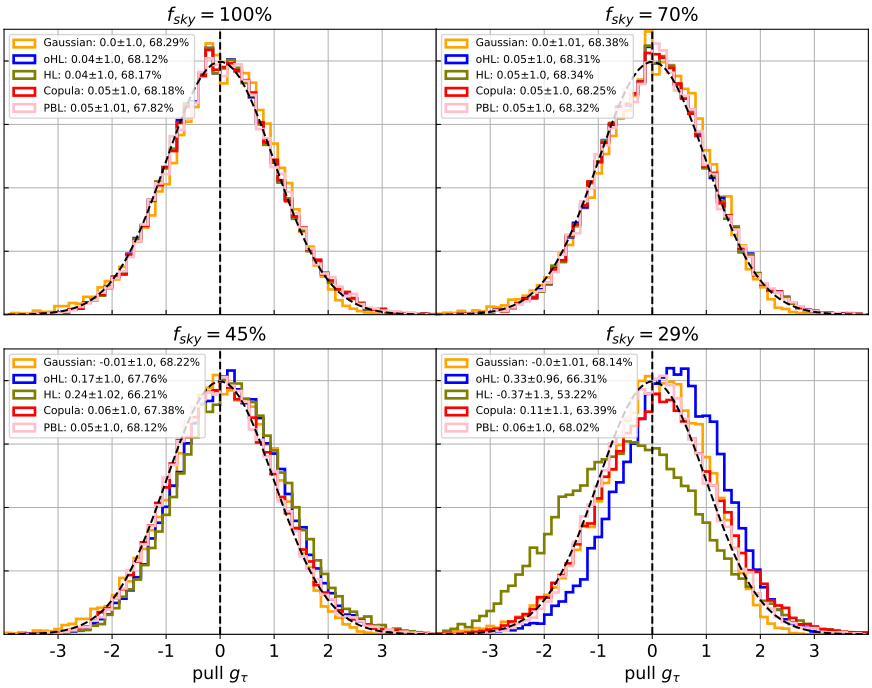

Figure 9.11: Refer to the caption of Fig. 9.10. This figure is similar but with simulations in the *low SNR* regime, and where the power-spectrum based likelihood have been built from with *cross* QML spectra (see Sec. 9.3.1.1).

the *analytical* covariance matrix instead of the empirical one.

9.3.1.4 Robustness to fiducial

All correlated-C_ℓ-based likelihoods depend on the fiducial model through $\Xi_{\ell\ell'}(\theta_i^{fid})$. Moreover, $\mathcal{L}_{\rm cop}$ rely on the C_ℓ dependence hypothesis Eq. 9.32 for the three parameters fitted on fiducial $\hat{C}_\ell$ distributions. These approximations could well result in a wrong estimate of τ when the fiducial does not match the truth. To quantify this effect, we use the sets of simulations for the other fiducial models introduced in Sec. 9.3.1.1. Fig. 9.12 represents the τ pull distributions for the various likelihoods in a full-sky setup. On the left panel, the input of the simulation, which is the *base* model, matches with the fiducial model assumed to derive the multipole covariance matrix and on which the copula parameters are based. Instead, on the middle and right panels, the input of the simulations, which still is the *base* model, mismatches with the used fiducial models, *low* (middle - all Λ-CDM parameters but τ are 1% lower than the input) or *lowtau* (right - τ is 10% lower than the input).

The middle panel of Fig. 9.12 indicates that using the *low* fiducial model doesn't affect the pull distributions. It demonstrates that a mismatch on Λ-CDM by 1% on all parameters does not affect the performance of the likelihood. Such a mismatch is conservative given the current Planck constraints, which are actually already considerably better than the per cent level (see Sec. 2.4).

Instead, the right panel of Fig. 9.12 illustrates that a mismatch of 10% on τ leads to a severe error on the errorbar estimation for all likelihoods studied so far. This effect is mainly driven by

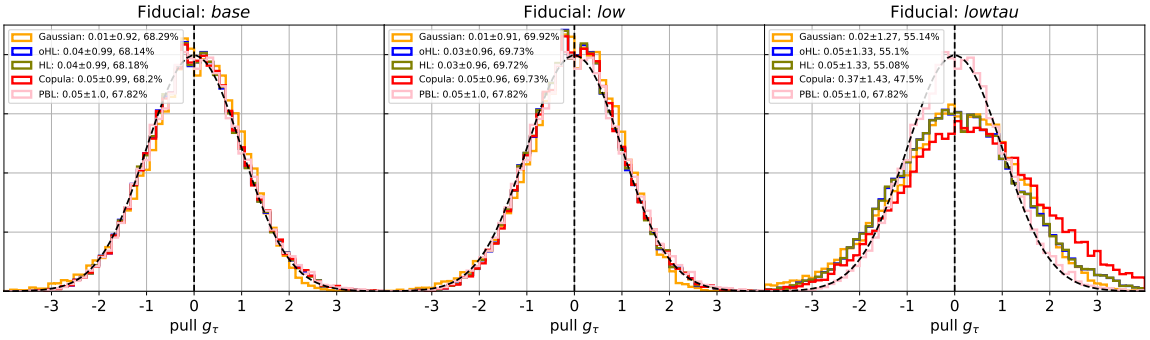

Figure 9.12: Pulls distributions for full-sky for *base*, *lowtau* and *low* fiducials (from left to right). $\hat{\tau}$ has been estimated from the following likelihoods: $\mathcal{L}_{\rm PB}$ (pink), $\mathcal{L}_{\rm gauss}$ (orange), $\mathcal{L}_{\rm oHL}$ (blue), $\mathcal{L}_{\rm HL}$ (green), $\mathcal{L}_{\rm cop}$ (red). We also display in the legend Gaussian parameters, mean and standard deviation, fitted on the distributions, as well as the fraction of realisations giving a pull between -1 and 1. The setup otherwise is *base* input theory, *high SNR* regime, with *cross* QML spectra and *empirical* covariance matrix described in Sec. 9.3.1.1

fixing the covariance matrix to a fiducial value. In the specific case of *lowtau*, the fiducial assumes a too-low value of τ, which implies an under-estimated EE spectrum and associated covariance matrix, which leads to an under-estimated τ uncertainty and to a too dispersed pull distribution. The possibility of such an under-estimated uncertainty for an estimated cosmological is a worrying prospect.

To address this issue, one typically includes the dependence of the covariance matrix in the model. The application of this idea requires further work: for now, our tries have been unfruitful and led to severely modified likelihood shapes that significantly deviate from the exact likelihood. This investigation will be let for a future work.

9.3.2. r estimation from white-noise case

In this section, let us focus on the r inference from white-noise BB simulated spectra.

9.3.2.1 Setup

In this discussion, we compare the oHL and the Gaussian likelihood performances for estimating r from 3000 hybrid BB spectra in low and high signal-to-noise regimes as introduced in Sec. 8.2.4. The hybrid estimator is built from cross pseudo-Cl and cross QML estimators with a transition multipole of $\ell_{\rm transi}$ optimised to each each situation encountered.

We use as a model the BB spectra illustrated in Fig. 9.13 in order to allow exploration of the negative value of r, even though those are unphysical r being the ratio between positive amplitudes[6]. Our motivation for doing that is to fully represent the likelihood without shrinkage, which noticeably allows us to verify the average unbiasedness of the maximum-likelihood estimator of r. In this case, we don't want to assume *a priori* the slow-roll relation between $n_{\rm t}$ and r. Instead, we fix $n_{\rm t}$ to zero.

[6]In practice, we extrapolate the definition of a primordial C_ℓ^{BB} for negative r by defining $\tilde{C}_\ell^{BB}(r) = \mathrm{sign}(r) C_\ell^{BB}(|r|)$.

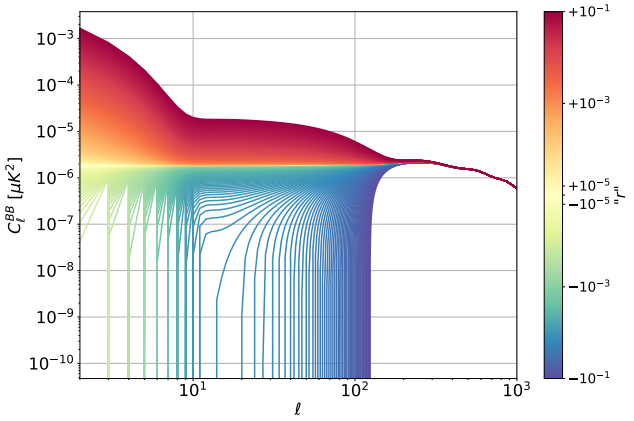

Figure 9.13: *BB* power spectrum for fixed Λ-CDM given in Table 4.1 + $r \in [-10^{-1}, 10^{-1}]$ with a symmetric logarithmic scale represented on the colorbar and $n_t = 0$. These are used as models for building the r-likelihoods.

9.3.2.2 Results

We represent on the left panel of Fig. 9.14 the statistics of the $\hat{r}$ (top), σ_r^- (middle), σ_r^+ (bottom) estimates for the oHL and Gaussian likelihoods, respectively represented by blue and orange errorbas, as a function of the sky fraction. As for τ, simulation-per-simulation, $\hat{r}$ maximises the likelihood while σ_r^- and σ_r^+ are defined similarly to Eqs. 9.35 and 9.36. The errorbars for each quantity give the 68% interval estimated over the 1000 simulations. We did the study for $f_{\rm sky} = 31.25, 52.05$ and 72.87% and $\sigma = 0.01$ and 10 μK·arcmin (represented respectively by plain and dashed bars).

As for τ, the Gaussian likelihoods yield $\hat{r}$ estimates that are, on average, well centered on the input value, $r = 0$. Instead, oHL follows a decreasing trend with decreasing $f_{\rm sky}$: in the high signal-to-noise regime, the bias is 4×10^{-5}, almost one order of magnitude below the uncertainty.

Contrary to our results on τ, for r, the $\sigma_r^\pm$ estimates from oHL and Gaussian likelihoods yield a result that is very similar, at the order of the per cent (instead of ten per cents for τ). Moreover, for the two likelihoods, one can also compare σ_r^+ with σ_r^-. One finds that the upper and lower uncertainties take extremely close values. This indicates that the shapes of the r likelihoods are much more symmetric than the τ ones. The explanation is that more multipoles contribute to the information r than for τ, multipoles which are higher and are associated with less skewed $\hat{C}_\ell$ probability density functions.

What this demonstrates is that for r, when a significant part of the information comes from the recombination bump at $\ell > 25$, the choice of oHL over the Gaussian approximation is not mandatory anymore. It even becomes a poor choice at the lowest $f_{\rm sky}$ where some C_ℓ's become negative and where a slight bias appears for oHL, while the Gaussian likelihood remains unbiased.

This conclusion is drawn from our particular setup, in which we include the information from the recombination bump in the fit, and could change if one focuses only on the reionisation bump. Also notice that when delensing more and more information comes from the recombination bump relative to the reionisation bump; see Fig. 4.11. In this situation, we expect the conclusion of this part to be strengthened and the Gaussian approximation to be good enough.

Finally, we represent in Fig. 9.15 the r-pull distributions. The first row illustrates the high

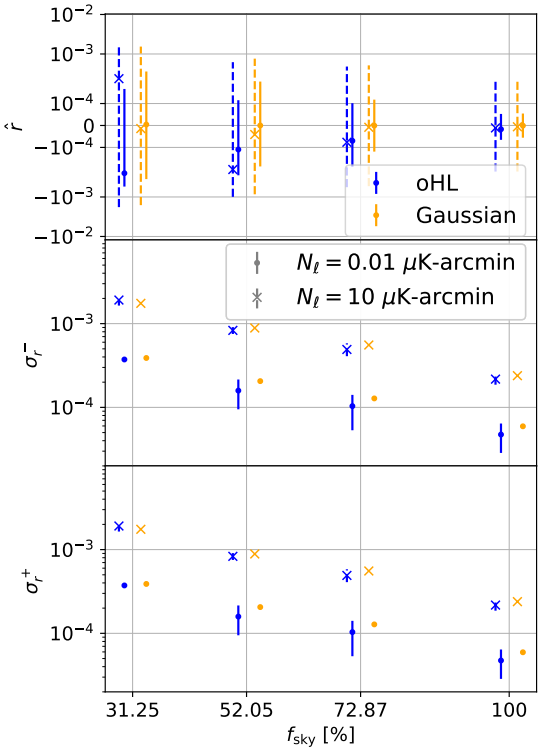

Figure 9.14: Estimate of $\hat{r}$ (upper panel), σ_τ^- (middle panel), σ_r^+ (bottom panel) (see Eqs. 9.35, 9.36), from $\mathcal{L}_{\rm gauss}$ (orange), and $\mathcal{L}_{\rm oHL}$ (blue) for varying $f_{\rm sky}$. The setup is the fiducial theory of Tab 4.1 (input $r = 0$) with 1000 hybrid cross spectra built in Sec. 8.6, and *empirical* covariance matrix estimated over the simulations. The simulations include a white noise of $\sigma = 0.01$ μK·arcmin illustrated by the plain errorbars, and $\sigma = 10$ μK·arcmin illustrated by the dashed errorbars. The latter represent the central region in which lie 68% of the realisations.

signal-to-noise regime, while the low signal-to-noise regime is represented on the lower panels. The sky fractions vary from the left to the right from 52% to 100%.

The possible next step for this investigation of r-likelihoods will be to make a comparison to the exact pixel-based likelihood, which is more challenging to implement than for τ because more pixels will be involved ($N_{\rm side} = 64$ at least). An important test would be to assume a poor fiducial model, as we did for τ, to test the likelihood's robustness. Then, we could vary the fraction of lensing in the simulations, which we expect to change the conclusions of the work. We also would like to adapt the copula likelihood to the r-inference problem (which is more challenging that for EE since most of the ℓ-per-ℓ probability density function are negative).

9.4. Application to LiteBIRD and results

In this section, we apply the maximum likelihood estimators, which we have previously introduced and tested, to the LiteBIRD simulations. We will focus on the simulations with wn+1/f, baseline and high-complexity foregrounds, applying the very conservative realistic mask with $\theta_1 = \theta_2 = 2°$. We start by introducing a correction to the quadratic likelihoods that is relevant due to the low number of simulations. Then, we re-assess the τ and r estimators in the

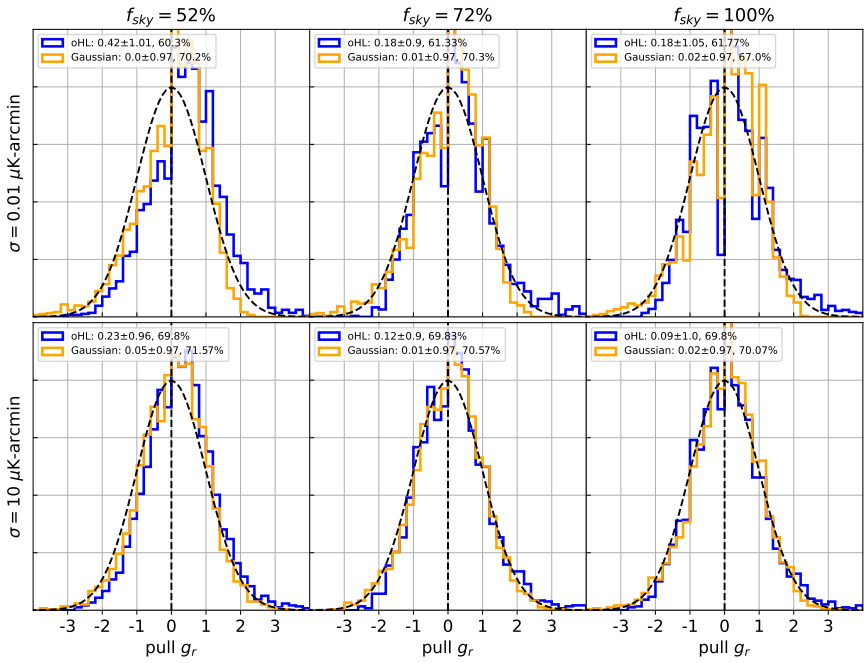

Figure 9.15: Normalised r-pull distributions for $f_{\rm sky}$ =52.05, 72.87, 100% (from left to right), and $\sigma = 0.01$ μK·arcmin (upper panel) or $\sigma = 10$ μK·arcmin (lower panel) superimposed to the standard normal distribution in black dashed line. We also display in the legend Gaussian distribution parameters, mean and standard deviation, fitted on the distributions, as well as the fraction of realisations giving a pull between -1 and 1. the setup consists of 1000 hybrid spectra simulations and the model is given in Fig. 9.13.

LiteBIRD setup, this time without the optimal pixel-based likelihood estimator, the noise being unknown. Finally, we conclude this LiteBIRD data-analysis part.

9.4.1. Covariance matrix correction

In the LiteBIRD simulations setup, we don't know the noise properties due to the previous steps of the analysis, and we no longer have access to the pixel-pixel noise covariance matrix. Consequently, the Fisher matrix of the power spectrum is not analytically computable, and we don't have a simple analytical (ℓ, ℓ) covariance matrix.

Therefore, we wish to use an empirical covariance matrix instead. However, we now only have 200 simulations (as opposed to the 20 000 used in the last section), from which an empirical covariance can only be poorly estimated, leading to a biased estimate of its inverse. This bias affects the parameter uncertainty estimates when used in likelihoods. The effect is discussed in [238], which proposes a simple formula for a corrected inverse covariance matrix, $\tilde{\Xi}^{-1}$. Given the estimated inverse covariance matrix Ξ, the number of simulations N and the number of data points p, the correction reads

$$\tilde{\Xi}^{-1} \equiv \alpha \Xi^{-1} \tag{9.41}$$

where the so-called *Hartlap factor* α is defined as

$$\alpha = \frac{N - p - 2}{N - 1}. \tag{9.42}$$

However, simply applying this formula in the quadratic form Eq. 9.15 relies on the assumption that the likelihood shape is not itself altered by the wrong inverse covariance estimate, which is *a priori* not the case. A more rigorous way of including the effect of the estimated covariance matrix is to marginalise over the true covariance matrix, conditioned on its estimated value, see [239]. This leads to rewriting any likelihood based on Eq. 9.15 as

$$\mathcal{L}\left(\{\theta_i\} \mid \left\{\hat{C}_\ell\right\}\right) = \frac{\bar{c}_p |\Xi|^{-1/2}}{\left[1 + \frac{X_\ell(\hat{C}_\ell,\{\theta_i\})[\Xi^{-1}]^{\ell\ell'} X_{\ell'}(\hat{C}_\ell,\{\theta_i\})}{N-1}\right]^{N/2}}. \tag{9.43}$$

with

$$\bar{c}_p = \frac{\Gamma\left(\frac{N}{2}\right)}{[\pi(N-1)]^{p/2}\Gamma\left(\frac{N-p}{2}\right)}. \tag{9.44}$$

Dropping the dependence of the covariance matrix in the model as well as the normalisation factors, the quadratic form $-2\ln\mathcal{L}_{\rm quad}$ is to be replaced by

$$-2\ln\mathcal{L}_{\rm quad}\left(\{\theta_i\} \mid \left\{\hat{C}_\ell\right\}\right) \to N \ln\left[1 + \frac{X_\ell(\hat{C}_\ell,\{\theta_i\})\left[\Xi^{-1}\right]^{\ell\ell'} X_{\ell'}(\hat{C}_\ell,\{\theta_i\})}{N-1}\right]. \tag{9.45}$$

The effect of this so-called *Sellentin-Heavens* prescription on the oHL τ likelihood for a given simulation is superimposed to the Hartlap correction in Fig. 9.16. For reference, we display the likelihood without correction in dashed blue lines. Our setup for estimating the covariance matrix in multipole space consists of 100 cross-QML *EE* LiteBIRD spectra between $\ell = 2$ and $\ell = 31$, which have been estimated from 30% (left panel) and 100% (right panel) maps, themselves being downgraded HILC split outputs that we have previously already introduced and used in Sec. 8.3.1. One of these realisations provides the data associated with the illustrated likelihood. We also display the estimated 1σ confidence level computed according to Eqs. 9.35 and 9.36.

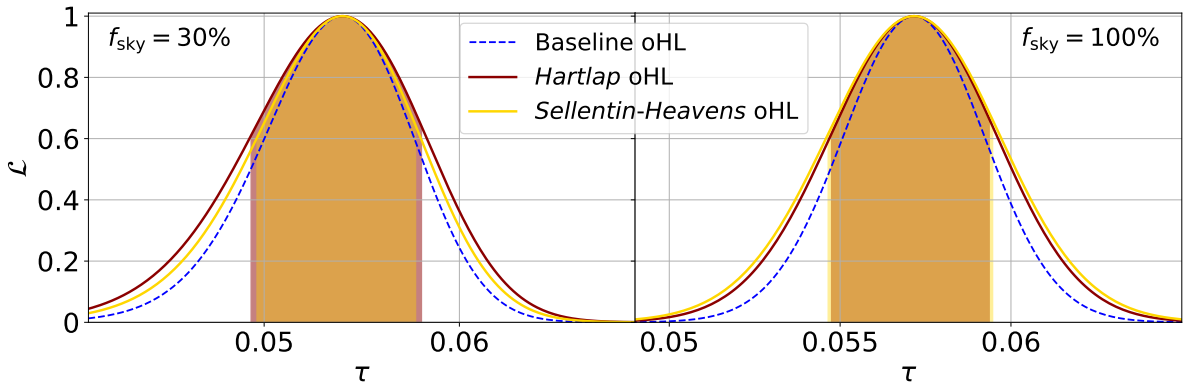

Figure 9.16: ***Preliminary!*** Effect of correcting the covariance matrix in the oHL likelihood by the Hartlap factor (red curve) or by the Sellentin-Heavens prescription (gold curve), where the empirical covariance is estimated on 100 cross-spectra estimated from the previous analysis steps, HILC and QML from wn + d1s1 input maps. This is shown for full-sky (left) and $f_{\rm sky} = 30\%$ (right). The fiducial model is given in Table 4.1.

Of course, all likelihoods share a common maximum since none of the corrections affect it. One notices that even for a matrix derived from 100 simulations, both corrections significantly

affect the likelihood shape and the reconstructed error bar on the parameter. More precisely, without correction, the error bars are underestimated. For this particular realisation

- for $f_{\rm sky}$=30%, the Hartlap-corrected oHL uncertainty is 18.4% greater than the non-corrected one, while this number is 9.5% for the Sellentin-Heavens prescription.
- In the full-sky setup, the Hartlap-corrected uncertainty is 18.6% greater than the non-corrected one, while this number is 23.4% for the Sellentin-Heavens prescription.

Later on, we use the Sellentin-Heavens prescription for the quadratic forms (Gaussian, HL and oHL approximations) and the Hartlap factor correction for the copula likelihood to correct this non-negligible effect.

9.4.2. τ estimation from LiteBIRD simulations

This section focuses on the τ inference in the LiteBIRD context. We have seen in Chapter 8 that while our implementation of the pseudo-power spectrum estimation leads to greater errorbars than QML for simplest white noise setups, the reverse happens for realistic LiteBIRD (and downgraded and HILC-cleaned) simulations. For this reason, and contrary to what we have done in Sec. 9.3, we use here cross pseudo-power spectra instead of the QML ones. The analysis could be replicated in the future with more optimal QML spectra. For τ inference, as in Sec. 9.3.1, we use a multipole range $2 \leq \ell \leq 31$.

In Fig. 9.17, we illustrate the branches of the τ likelihoods from cross pseudo-power spectra obtained from split HILC-cleaned maps, themselves estimated from frequency maps that contain white noise, $1/f$ noise and high-complexity foregrounds (see preliminary chapter of Part. III). Contrary to our similar representation in the white-noise case, we could not compute the pixel-based likelihood since we no longer have access to an accurate pixel covariance matrix. Furthermore, contrary to the high-SNR regime of Fig. 9.7, our simulations are noisier, so we chose to represent the log-likelihood in logarithmic scale with an offset of 5 units to visualise the branches.

Similarly to Fig. 9.7, the Gaussian likelihood presents its characteristic branches with a moderate slope for $\tau < \hat{\tau}$. However, the branch behaviour of the other likelihoods for $f_{\rm sky} \geq 70\%$ is less similar than it was in the white-noise-only scenario. The reason for that is the presence of the foreground residual that is unmodelled in the theoretical spectra, as well as the small number of simulations. The latter significantly impacts the performances of the copula likelihood especially, because it relies on a fit over 200 simulations, explaining why it deviates at large sky fractions from the other likelihoods.

In the same setup, in Fig. 9.18, we summarise how the estimations of $\hat{\tau}$ (upper panel), σ_τ^- (middle panel), σ_τ^+ (bottom panel) vary with the sky fraction. The left panel illustrates the results for high-complexity foreground inputs, while the baseline complexity results are shown in the right panel. Various likelihoods are represented: $\mathcal{L}_{\rm gauss}$ (orange), $\mathcal{L}_{\rm oHL}$ (blue), $\mathcal{L}_{\rm HL}$ (green), $\mathcal{L}_{\rm cop}$ (red).

First, from the upper panel, one sees that the trend for the estimated $\hat{\tau}$ in this realistic case is similar to the white-noise case:

1) As for the white noise case, a slight bias on the τ estimate remains at all sky fractions for all likelihood approximations, even though the spectra were not biased. This bias is especially visible in the right panel with baseline complexity, where all likelihoods apart from HL yield an underestimated $\hat{\tau}$ at $f_{\rm sky} \leq 70$. Despite being a poor approximation of

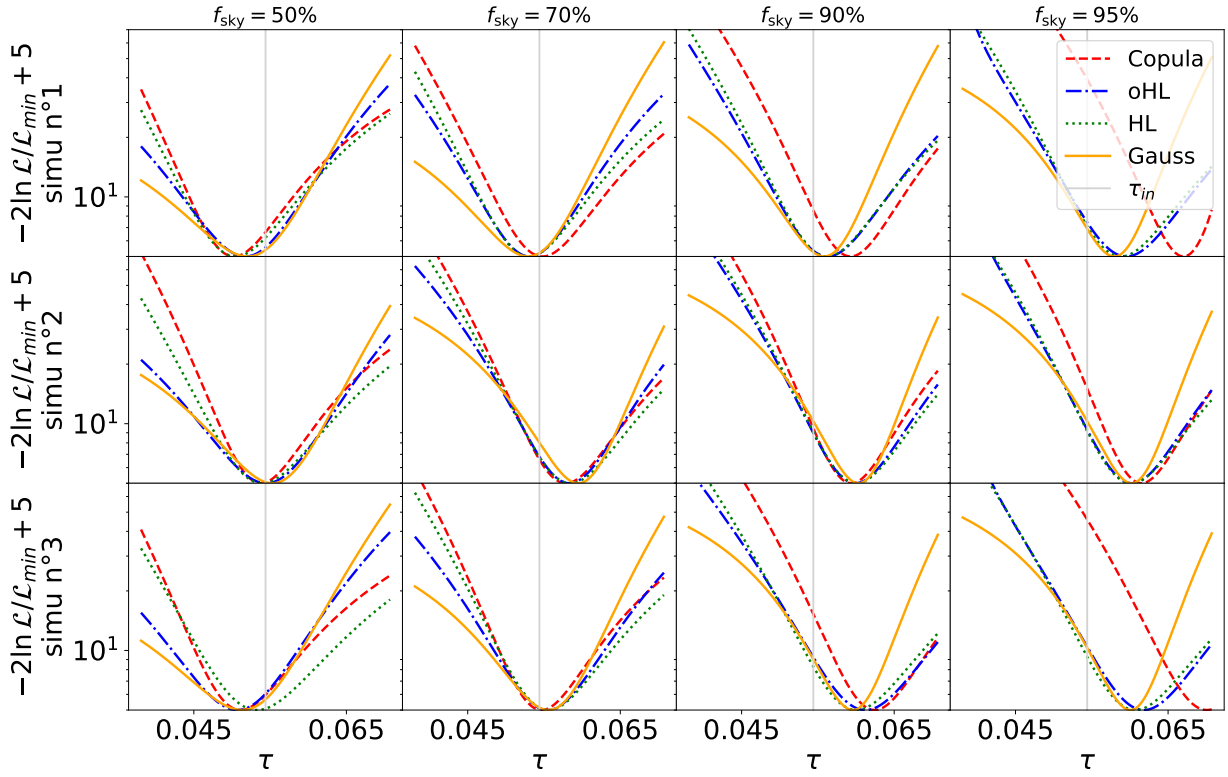

Figure 9.17: ***Preliminary!*** Shape of -2 log-τ-likelihood for three different LiteBIRD wn1f + high-complexity-foreground realisations (from top to bottom) and for four different sky fractions, 50%, 70%, 90%, 95% (from left to right). $\mathcal{L}_{\rm cop}$ is represented in dashed-red, $\mathcal{L}_{\rm oHL}$ in dashed-dotted-blue, $\mathcal{L}_{\rm HL}$ in dotted-green, $\mathcal{L}_{\rm gauss}$ in orange. The input τ is represented by a vertical grey line. For illustration purposes, the minimum log-likelihood is fixed to 5. the setup consists of cross pseudo-power spectra obtained from split HILC-cleaned maps. The covariance matrix is estimated from 200 simulations, and the likelihoods are corrected by Sellentin-Heavens prescription (for Gaussian, HL, oHL) or by the Hartlap factor (for copula). The pixel-based likelihood is not applicable to these unknown-covariance case.

the likelihood far from its maximum (and thus yielding incorrect uncertainty estimates), we have seen in Sec. 9.3 that the Gaussian likelihood provides a reliable estimate of $\hat{\tau}$ in average. From this estimate, we infer a bias on the order of 6×10^{-4} (1% of the τ value). This bias is stable across different $f_{\rm sky}$ values and likelihoods and remains robust when changing the ℓ range under scrutiny.

2) On top of this offset, as soon as $f_{\rm sky} \leq 60\%$, the HL likelihood slightly overestimates τ while the oHL likelihood slightly underestimates it, similarly to the white noise case.
3) The foreground cleaning significantly impacts the estimate for all likelihoods for $f_{\rm sky} \geq 80\%$. The resulting systematic error can be estimated by taking the Gaussian likelihood estimate at $f_{\rm sky}= 40\%$ as a reference. For the high-complexity foregrounds, this systematic error at 90% is of 1.7×10^{-3} while it is 1.3×10^{-3} for the baseline foregrounds.

To summarise, our τ estimation involves three interplaying biases that come from three different sources: 1) the non-asymptotic maximum-likelihood estimator, 2) poor likelihood approximations

at low $f_{\rm sky}$, and 3) foreground-biased power spectra at high $f_{\rm sky}$.

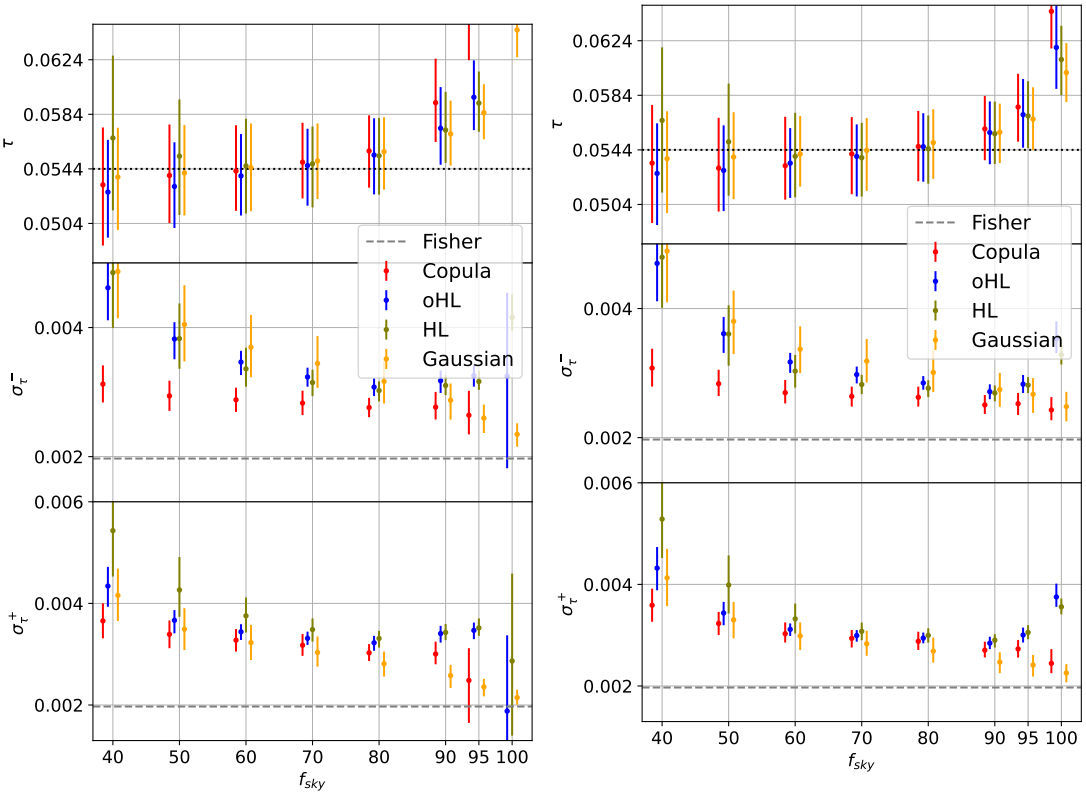

Figure 9.18: ***Preliminary!*** Estimate of $\hat{\tau}$ (upper panel), σ_τ^- (middle panel), σ_τ^+ (bottom panel), from $\mathcal{L}_{\rm gauss}$ (orange), $\mathcal{L}_{\rm oHL}$ (blue), $\mathcal{L}_{\rm HL}$ (green), $\mathcal{L}_{\rm cop}$ (red) for varying $f_{\rm sky}$. The setup consists of cross pseudo-power spectra obtained from split HILC-cleaned high-complexity maps, on the left, and baseline-complexity maps on the right. The covariance matrix is estimated from 200 simulations, and the likelihoods are corrected by the Sellentin-Heavens prescription (for Gaussian, HL, oHL) or by the Hartlap factor (for copula). The errorbars displayed here represent the central region in which lie 68% of the realisations. We also display the input τ by a dotted black horizontal line on the upper panel, and the Fisher estimate of the noiseless full-sky symmetric error by a dashed grey line on the other panels.

Second, one observes that $\sigma_\tau^\pm$ in this realistic setup is increased with respect to the high signal-to-noise case highlighted in Fig. 9.8. As anticipated in Sec. 9.1, around $f_{\rm sky}$= 80%, the uncertainty on τ is of order 3×10^{-3}. This uncertainty can be estimated simulation per simulation, an estimation that is reliable at less than 5×10^{-4} for HL and oHL and less than 10^{-3} for the Gaussian and copula likelihoods. At low $f_{\rm sky}$, the estimation of $\sigma_\tau^\pm$ is far more uncertain than in the white-noise case:

1) As usual, the Gaussian likelihood leads to a high σ_τ^-. Because of the considerations done in the white-noise setup, we can eliminate this approximation for uncertainty estimation.
2) Among the four approximations, the copula likelihood is the one that presents the steepest branch at low τ, which leads to a low estimated σ_τ^- that disagrees with the oHL and HL likelihoods.

We have verified that the remarks concerning the likelihood approximations also hold for the baseline and white-noise complexities.

Before concluding, we emphasise that the higher complexity of this setup is corroborated by the pull distributions, which are represented in Fig. 9.19 (with the same plotting conventions as Fig. 9.10 ones).

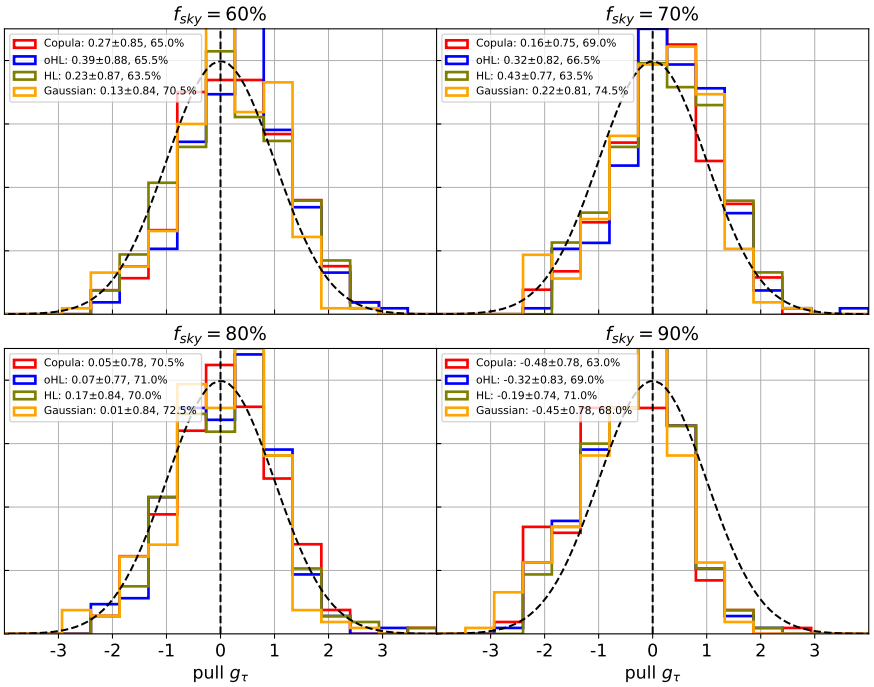

Figure 9.19: ***Preliminary!*** Normalised τ-pull distributions for $f_{\rm sky} = 60,\ 70,\ 80,\ 90\%$ (from top-left to bottom-right), superimposed to the standard normal distribution in black dashed line. $\hat{\tau}$ has been estimated from the following likelihoods: $\mathcal{L}_{\rm gauss}$ (orange), $\mathcal{L}_{\rm oHL}$ (blue), $\mathcal{L}_{\rm HL}$ (green), $\mathcal{L}_{\rm cop}$ (red). We also display in the legend Gaussian distribution parameters, mean and standard deviation, fitted on the distributions, as well as the fraction of realisations giving a pull between -1 and 1. The setup consists of cross pseudo-power spectra obtained from split HILC-cleaned maps.

1) One notices the different biases commented earlier on these plots in the mean of the fitted normal distribution. For instance, the positive bias on $\hat{\tau}$ due to foreground residuals at 90% of the sky leads as expected to negative pull averages (-0.19, -0.32, -0.45 and -0.48 for the HL, oHL, Gaussian and copula likelihoods respectively).
2) The standard deviations of the fitted normal distributions are all significantly lower than one, hinting at underestimated uncertainties. The reason for this could be simple: our 200 simulations all contain the same foreground template. This means that the cosmic sample variance of the foregrounds is not being emulated. Consequently, the dispersion of spectra over simulations is smaller than the actual uncertainty on the spectra. The same holds for $\hat{\tau}$, which are less dispersed than the actual uncertainty on the parameter simulation per simulation. The direct impact on the pull is a standard deviation of smaller than 1.

3) The number of simulations lying between -1 and 1 is more difficult to interpret than in our 20000 simulations' case. Instead, with 200 simulations, the expected 1σ statistical fluctuation around 68.27% is given by a 3.29% standard error. Hence, all per cent deviations seeable in Fig. 9.19 can be interpreted as statistical fluctuations.

These pull distributions are similar in the cases of setups other than the high-complexity one.

We have applied various power-spectrum-based τ likelihoods to our LiteBIRD setup. The main conclusions concerning the likelihood-based τ estimates are

1) We have shown that despite deviating at low $f_{\rm sky}$, the maximum-likelihood $\hat{\tau}$ estimations are quite reliable for any likelihoods at large $f_{\rm sky}$, with an excellent agreement of all four approximations at $f_{\rm sky} = 90\%$ to 95%. For lower sky fractions, the HL and oHL likelihoods lose in reliability for $\hat{\tau}$ estimation.
2) We have demonstrated how to estimate the τ from the likelihood. The result of this estimation is directly linked to the shape of the likelihood branches distancing from the maximum. On cut-sky, we have seen that the oHL and the copula approximations are the best at providing estimates that are close to optimal, even at low-sky fractions, while HL is mainly reliable at high $f_{\rm sky}$. For τ inference at high $f_{\rm sky}$, the safest choice in the realistic setup with a low number of simulations is to use the oHL likelihood uncertainty estimates.
3) We have identified a slight offset in the maximum-likelihood estimates common to all sky fractions and to all likelihoods, which we believe is due to the non-asymptotic regime of the τ inference from CMB. The bias is also present in the simplest test cases, while no bias in the spectra is noticed. Even if it does not represent a big issue in practice, we aim to confirm, investigate and explain this effect in future work.

With this tool, we can also conclude on the performance of the full data-analysis pipeline for τ estimation from simulated frequency maps (namely HILC-cleaning into cross pseudo-power-spectrum estimation into maximum spectra-based-likelihood τ estimation). We summarise this result in Fig. 9.20, which illustrates similar curves as Fig. 9.1. The difference lies in the manner in which the uncertainties have been estimated, which is now more realistic and agnostic.

$\sigma_{{\rm stat},\tau}$ is estimated as the mean of σ^- and σ^+, each of them estimated with oHL and averaged over the 200 simulations. Instead, $\tilde{\sigma}_{{\rm stat},\tau}$ is estimated as the standard deviation of the simulation-per-simulation oHL $\hat{\tau}$. We chose not to include in $\sigma_{{\rm syst},\tau}$ the systematic error due to the estimator, which is still to be investigated: we estimate $\sigma_{{\rm syst},\tau}$ as the difference between the average of the Gaussian $\hat{\tau}$ and a reference, which is the average of the Gaussian $\hat{\tau}$ for 30% of the sky (where foregrounds are mostly masked). Remark 3 of the previous discussion justifies the choice of these likelihoods for these contexts. As in Fig. 9.1, the region between the root sum square of the systematic and statistical uncertainties and their sum is delimited as orange contours, which is dashed for the high-complexity case.

Notice that for this discussion, we allow ourselves to use two different likelihoods to estimate the τ and its uncertainty, which is not the preferred solution with real data.

As we have mentioned when discussing the pull distributions, $\sigma_{{\rm stat},\tau} > \tilde{\sigma}_{{\rm stat},\tau}$: the foreground uncertainty in spectrum space is not accounted for because we only have one foreground template. This is even more visible by considering the decreasing $\tilde{\sigma}_{{\rm stat},\tau}$ trend with $f_{\rm sky}$ at high $f_{\rm sky}$. There, the presence of more and more foreground residuals does not increase the cosmic variance as it should. Hence, the statistical error should not be estimated from the statistics of the τ estimate.

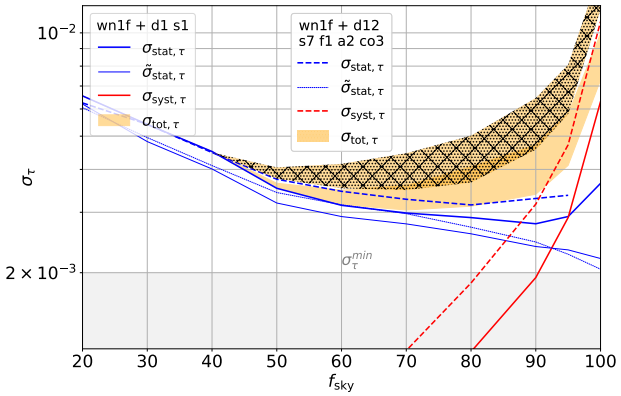

Figure 9.20: ***Preliminary!*** Maximum likelihood estimate of the $\sigma_{\mathrm{stat}} \equiv \sigma(\hat{\tau}^{\mathrm{oHL}})$ (blue) and $\sigma_{\mathrm{syst}}\langle\hat{\tau}\rangle^{\mathrm{Gauss}} - \langle\hat{\tau}\rangle^{\mathrm{Gauss}}(f_{\mathrm{sky}} = 30\%)$ (red) obtained from the cross EE pseudo power spectrum estimation of the residuals averaged over 200 wn1f + d1 s1 simulations, cleaned with HILC ($N_m = 400$, $\sigma_{w_\ell} = 35$), with the realistic mask. The total uncertainty σ_{tot} will lie between the root sum square and the sum ($\equiv \delta$) between the two residuals, an area represented by the orange contour. We represent the corresponding two uncertainties and the σ_{tot} area with high complexity foregrounds (wn1f + d12 s7 f1 a2 co3) with dashed lines and a hashed contour. We also represent the best σ_τ accessible in the reionisation bump due to cosmic variance, which is exactly 2×10^{-3} for our fiducial.

One observes that the overall behaviour of the curves is similar to the one we obtained with Fisher analysis of the residuals, with a foreground-driven systematic error dominating the statistical one around 85-90% of the sky, the statistical error slowly increasing when going toward lower sky fractions.

9.4.3. r estimation from LiteBIRD simulations

In this section, we apply the oHL and the Gaussian likelihoods to cross-pseudo spectra estimated on HILC-cleaned LiteBIRD simulations, corrected by the Sellentin-Heavens prescription. No delensing is assumed. We use as a model the $C_\ell^{BB}(r)$ as illustrated in Fig. 9.13, in order to allow exploration of the negative value of r. In this case, we limit our ℓ range to 190, beyond which no more information is brought to the fit.

In Fig. 9.21, we illustrate the r likelihoods from cross pseudo-power spectra obtained from split HILC-cleaned maps, themselves estimated from frequency maps that contain white noise, $1/f$ noise and high-complexity foregrounds. We represent -2 log-likelihood, offset by 0.5 units for visualisation reasons. As compared to τ, the r-likelihoods are, on average, slightly more symmetric with respect to $\hat{r}$.

Similarly to what we did in Sec. 9.3.2, we estimate $\hat{r}$, σ_r^- and σ_r^+ from these likelihoods. The results as a function of the sky fraction are displayed in Fig. 9.22, the high-complexity foregrounds' results are represented on the left panel while the baseline are illustrated in the right panel.

Let us comment on the $\hat{r}$ estimate first. Similarly to what we seen in the white noise case on hybrid spectra, oHL and Gaussian likelihood follow the same trend in terms of $\hat{r}$ estimate. However, oHL always leads to slightly underestimated r with respect to the $r = 0$ input. Apart from the lowest f_{sky} point of oHL, which is in average negative at $\langle\hat{r}\rangle = -9 \times 10^{-3}$, the other

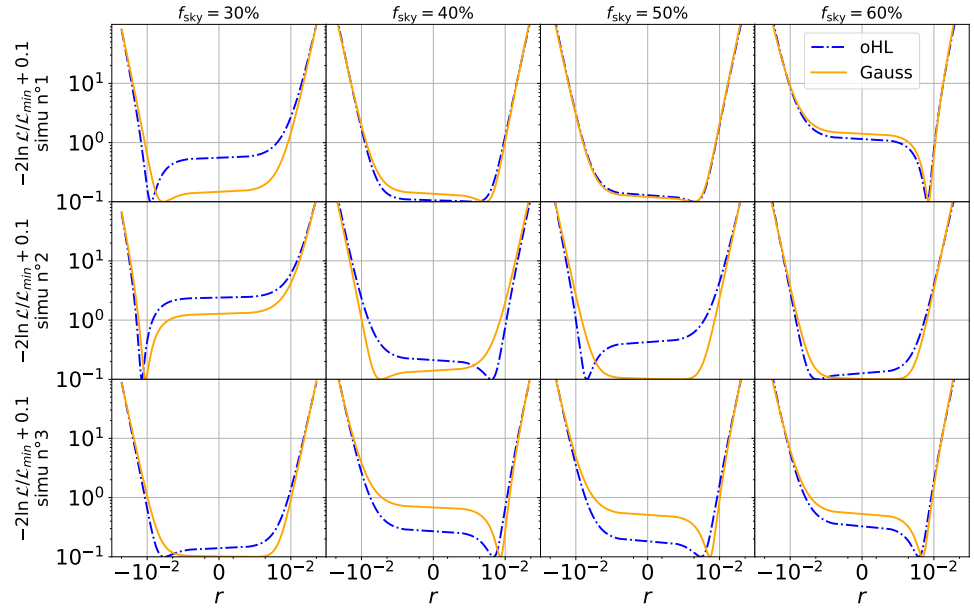

Figure 9.21: ***Preliminary!*** Shape of -2 log-r-likelihood for three different LiteBIRD wn1f + high-complexity-foreground realisations (from top to bottom) and for four different sky fractions, 50%, 60%, 70%, 80% (from left to right). $\mathcal{L}_{\rm oHL}$ is represented in dashed-dotted-blue, $\mathcal{L}_{\rm gauss}$ in orange. The fiducial model is given in Table 4.1 with $r = 0$. For illustration purposes, the minimum log-likelihood is fixed to 0.5. The setup consists of cross pseudo-power spectra obtained from split high-complexity HILC-cleaned maps. The covariance matrix is estimated from 200 simulations, and the likelihoods are corrected by the Sellentin-Heavens prescription.

averages lie in the positive regions.

From the Fisher analysis, we have concluded that $\sigma^{\rm Fisher}_{r,{\rm stat}}$ is quasi flat with $f_{\rm sky}$, slightly decreasing for high-complexity foregrounds, and slightly increasing for baseline. Here, instead, the uncertainty estimate decreases with $f_{\rm sky}$. We find again that the likelihood is nearly symmetric since $\sigma_r^- \simeq \sigma_r^+$. Contrary to what we have seen for τ, Gaussian and oHL approximations yield similar uncertainty estimates.

In Fig. 9.23, we represent the high-complexity-foregrounds r-pull distributions for $f_{\rm sky}$ between 30% and 60%. One directly sees that they are too sharp with respect to expected. Instead of a unit standard deviation of the distributions, their fitted Gaussian σ is between 0.5 and 0.6 for all $f_{\rm sky}$ and both likelihoods. This indicates that $\hat{r}$ are less dispersed with respect to the average estimated uncertainty simulation-per-simulation. As previously mentioned, we explain this by the fact that the power spectrum simulations are less dispersed than they should be because of our use of a unique foreground template common to every simulation. Pulls cannot be interpreted further in this case.

Finally, we reproduce the right panel of Fig. 9.24 from our updated $\hat{r}$ and $\sigma_r^\pm$ estimates. Similarly to what we have done for τ, $\sigma_{{\rm stat},r}$ is estimated as the mean of σ^- and σ^+, each of them estimated with oHL and averaged over the 200 simulations. Instead, $\tilde{\sigma}_{{\rm stat},r}$ is estimated as the standard deviation of the simulation-per-simulation oHL $\hat{r}$. We estimate $\sigma_{{\rm syst},r}$ as the difference between the average of the Gaussian $\hat{r}$ and a reference, which is the average of the Gaussian $\hat{r}$ for 30% of the sky (where foregrounds are mostly masked). Our reason for doing this is

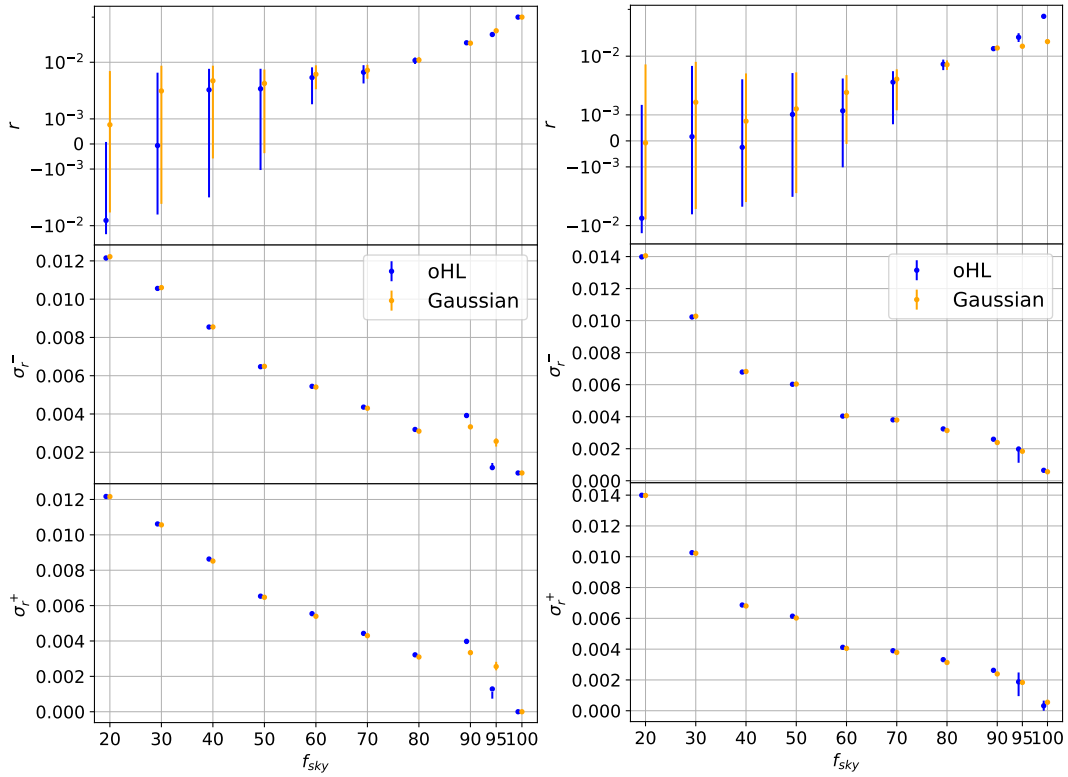

Figure 9.22: ***Preliminary!*** Estimate of $\hat{r}$ (upper panel), σ_r^- (middle panel), σ_r^+ (bottom panel), from $\mathcal{L}_{\rm gauss}$ (orange) and $\mathcal{L}_{\rm oHL}$ (blue) for varying $f_{\rm sky}$. The setup consists of cross pseudo-power spectra obtained from split HILC-cleaned maps, with realistic masks 400 35. Left: high complexity right: baseline. The covariance matrix is estimated from 200 simulations, and the likelihoods are corrected by the Sellentin-Heavens prescription. The errorbars displayed here represent the central region in which lie 68% of the realisations.

to avoid including in the systematic error budget the unexplained maximum-likelihood-estimator bias we have previously discussed. As in Fig. 9.1, the region between the root sum square of the systematic and statistical uncertainties and their sum is delimited as orange contours, which is dashed for the high-complexity case.

One observes that averaged estimates yield noisy curves. This representation would benefit from more than 200 simulations. One can still observe that the behaviour of the statistical error curve significantly differs from what we obtained with Fisher propagation of the residuals; see Fig. 9.1: for both foreground complexities, the blue curves significantly increase when going to low $f_{\rm sky}$. The reason for that, once again, is the fact that foreground sample variance does not appear in this curve estimate.

Conclusion

In this chapter, we introduced likelihood approximations for extracting large-scale polarisation information (τ and r) from masked maps and associated power spectra.

Regarding τ recovery, we have shown that despite deviating at low $f_{\rm sky}$, the maximum-

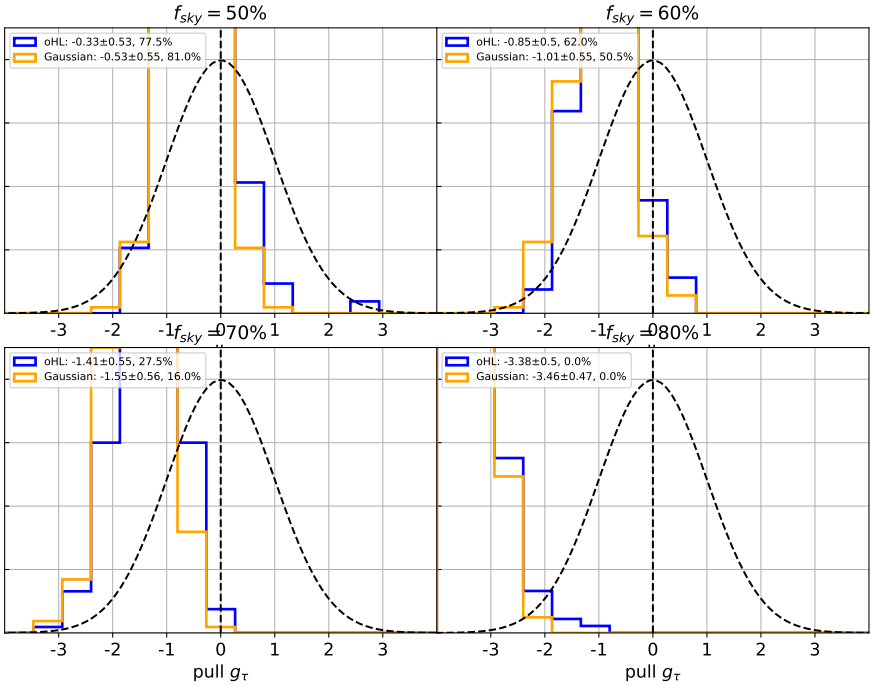

Figure 9.23: ***Preliminary!*** Normalised pulls distributions for $f_{\rm sky}$ =30, 40, 50, 60% (from top-left to bottom-right), superimposed to the standard normal distribution in black dashed line. $\hat{\tau}$ has been estimated from the following likelihoods: $\mathcal{L}_{\rm gauss}$ (orange), $\mathcal{L}_{\rm oHL}$ (blue). We also display in the legend Gaussian distribution parameters, mean and standard deviation, fitted on the distributions, as well as the fraction of realisations giving a pull between -1 and 1. The setup consists of cross pseudo-power spectra obtained from split HILC-cleaned maps.

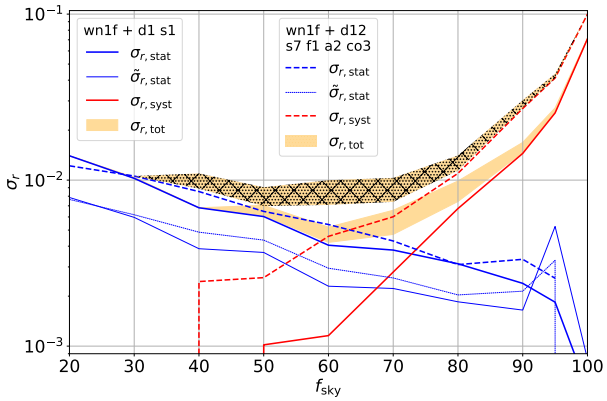

Figure 9.24: ***Preliminary!*** Maximum likelihood estimate of the $\sigma_{\rm stat} \equiv \sigma(\hat{r}^{\rm oHL})$ (blue) and $\sigma_{\rm syst}\langle\hat{r}\rangle^{\rm Gauss} - \langle\hat{r}\rangle^{\rm Gauss}(f_{\rm sky} = 30\%)$ (red) obtained from the cross BB pseudo power spectrum estimation of the residuals averaged over 200 wn1f + d1 s1 simulations, cleaned with HILC ($N_m = 400$, $\sigma_{w_\ell} = 35$), with the realistic mask. The total uncertainty $\sigma_{\rm tot}$ will lie between the root sum square and the sum ($\equiv \delta$) between the two residuals, an area represented by the orange contour. We represent the corresponding two uncertainties and the $\sigma_{\rm tot}$ area with high complexity foregrounds (wn1f + d12 s7 f1 a2 co3) with dashed lines and a hashed contour.

likelihood $\hat{\tau}$ estimates are quite reliable for large $f_{\rm sky}$, with excellent agreement among all four approximations at $f_{\rm sky} = 90\%$ to 95%. For lower sky fractions, the HL and oHL likelihoods lose reliability in $\hat{\tau}$ estimation. The statistical uncertainty on the parameter is best recovered from HL and oHL for $f_{\rm sky} \geq 70\%$, and is minimal around $f_{\rm sky} = 90\%$. However, the variance due to foregrounds is not included in this number since foreground variance is not included in our simulations. For all levels of foreground complexity, we find that the best τ estimate occurs around $f_{\rm sky} \simeq 70\%$, yielding a total uncertainty of approximately $\sigma_\tau^{\rm tot} = 3 \times 10^{-3}$, which is roughly in agreement with our Fisher estimate.

We have mentioned that using both HL and oHL for very low $f_{\rm sky}$ should be done with caution. The approximations break down, the estimators deviate from optimality, and significant bias can appear in the parameter estimate. This issue is less concerning for τ estimation on a slightly cut sky, unlike r recovery that has to be performed on a more severely cut sky. Moreover, with the LiteBIRD noise and residual post-HILC, the r information arises largely from the recombination bump, where the use of the Gaussian likelihood is satisfactory for both estimating $\hat{r}$ and its uncertainties. The r-likelihoods are nearly symmetric, and both oHL and Gaussian likelihoods give similar σ_r^+ and σ_r^-. The conclusion of this chapter regarding the r uncertainty value should not be taken too seriously: the variance due to foregrounds is not included, and we have seen through preliminary Fisher studies that this plays a crucial role (see Fig. 9.1). Moreover, without marginalisation over the foreground residual, we observe a significant residual systematic in the $\hat{r}$ estimate of the order 3 to 5×10^{-3}, depending on the foreground complexity at $f_{\rm sky} = 60\%$. The next step in this investigation will be to perform such marginalisation.

Data-analysis for LiteBIRD: conclusion

To conclude, we have developed and investigated a comprehensive pipeline to analyze LiteBIRD simulations, from frequency maps to τ and r. This pipeline comprises three main components, which we have detailed in Chapters 7, 8 and 9.

The **foreground cleaning** with HILC is a blind method that enables CMB recovery without requiring prior foreground knowledge. We have optimised it for the context of LiteBIRD, and we have introduced a binning/smoothing procedure in ℓ space that minimises residuals while preserving their physical continuity. Although this thesis does not include this comparison, we evaluated the performance of our method against a parametric method and an ILC method in needlet space within the context of LiteBIRD E-modes. This comparison, to be published by the LiteBIRD collaboration, will show that ILC performs similarly in harmonic and needlet space across all levels of foreground complexity and all multipoles. The parametric method, while outperforming the ILC methods at large scales, necessitates modelling the mixing matrix, thus losing in agnosticism. Therefore, HILC, demonstrated to be robust against increasing foreground complexity without requiring manipulation, will serve as a powerful method to cross-check results from more complex methods in a straightforward manner. The output maps it produces are agnostic, reliable, and provide reasonably good results.

The second part of our pipeline consists of **power spectrum** estimation, which we have performed using the (cross) QML and (cross) pseudo-spectrum methods. In the case of white noise, we have shown that the optimal solution at large scales would be to apply a QML method, which theoretically allows for low variance and low correlations. The challenge in a realistic context is the need for a well-controlled covariance matrix in pixel space. Currently, our implementation of QML on 200 LiteBIRD simulations does not outperform the pseudo-power spectrum method. We have also prepared for the possibility of using a hybrid power spectrum estimator with a flexible transition multipole. In white noise cases, this hybrid estimator allows for a quasi-optimal reconstruction of the spectrum from $\ell = 2$ to $\ell = 1000$. With upcoming improvements to our QML method implementation, we still hope to use this hybrid estimator for future LiteBIRD applications.

Finally, the last part of our pipeline involves **estimating τ, r** and their respective uncertainties. For τ, we have compared various C_ℓ-based likelihood approximations to a reference, the

pixel-based likelihood. In particular, we have presented the copula likelihood and applied it for the first time to infer τ from low-ℓ EE power spectra, yielding promising results.We also estimated r using two specific likelihoods. The maximum-likelihood estimation of these parameters in the LiteBIRD context still requires consolidation and would especially benefit from marginalisation over the foreground residuals. Preliminary results indicate that our current foreground removal approach is too simplistic to achieve a full cosmic-variance measurement of τ and an r uncertainty compatible with LiteBIRD requirements [153].

Our pipeline is intended to be one of several analysis pipelines for LiteBIRD. It is still under development, and we have suggested several improvements throughout this thesis. We aim to keep it as agnostic as possible, believing its value will lie in this characteristic. Blind and robust methods will be needed to cross-check results with pipelines that yield better results but require more prior knowledge of the instrument and the microwave sky. Given the potential for significant discoveries, LiteBIRD results must be extremely well understood to ensure solid and reliable claims.

Beyond constraining r and τ, polarisation maps will allow us to characterise reionisation beyond τ, test the Gaussianity of perturbations, constrain primordial magnetic fields, investigate parity violation and cosmic birefringence, address current anomalies, and tackle Galactic science.

In Part IV, we focus on a scientific topic that will benefit greatly from LiteBIRD results: inflation phenomenology. While awaiting LiteBIRD's data, we will focus on the available Planck dataset, which provides the best constraints to date on the primordial parameters. However, this work will undoubtedly need to be updated after the LiteBIRD flight.

Part IV

Slow-roll inflation phenomenology & MSSM inflation

The goal of this part is to demonstrate how one can combine cosmological and particle-physics constraints to constrain unified models able to explain inflation, dark-matter and particle-physics constraints. It will contain material based and partly extracted from [93], and additional material and references.

We first introduce the required equations for deriving cosmological predictions from slow-roll models. We illustrate these for two specific slow-roll phenomenologies. Very few of the proposed slow-roll models come with a complete embedding within particle physics model, which would allow a coherent description of physics from the high-energy scales of inflation down to the lower energy scales probed in particle physics experiments. The second chapter of this part focuses on a well-defined theoretical framework, used as a test case to study all aspects of a fully specified embedding. To introduce it, we recapitulate the motivations for extending beyond the particle-physics *Standard Model* (SM) before introducing one of its extensions, the *Minimal Supersymmetric Standard Model* (MSSM). Finally, we explain how inflation can be realised in an extended MSSM framework, how to derive inflationary predictions, and how to perform rigorous combined fits of MSSM with particle physics and cosmological constraints.

Figure IV.1: Adaptation from artistic views of LiteBIRD (credit: *ISAS/JAXA*), Planck (credit: *TICRA*), Xenon1T (credit: *CERN*), and LHC (credit: *CERN*). The final goal of this part is to perform a combined fit of the Minimal Supersymmetric Standard Model based on the constraints of these various experiments.

CHAPTER 10

Slow-roll predictions and BSM particle-physics

The first section of this chapter, Sec. 10.1, aims to complement our inflation and slow-roll introduction by explaining how, in practice, one can derive cosmological predictions from a given slow-roll model. High-energy physics constructions sometimes inspire the existing models, but the precise link is often not established. One needs to work within a fully specified theoretical theory to do that. In a second section, Sec. 10.2, we will overview the reasons that motivate going beyond the standard model of particle physics and its minimal supersymmetric extension, the Minimal Supersymmetric Standard Model. We will introduce notions and definitions useful for Chapter 11.

Contents

10.1 Predicting the primordial power spectra in slow-roll approximation 230

10.1.1 Validity of the slow-roll trajectory and initial conditions 230

10.1.2 Reheating consistent predictions of slow-roll inflation 232

10.1.3 Predicting cosmological observables 233

10.2 Going beyond the particle standard model 234

10.2.1 The Standard Model of particle physics and its open questions 235

10.2.2 Supersymmetry as an extension to the Standard Model 235

10.2.3 The Minimal Supersymmetric Standard Model 238

10.1. Predicting the primordial power spectra in slow-roll approximation

We begin this section illustrating in a concrete case the slow-roll equations introduced in Sec. 3. Then, we explain how one can predict the Hubble slow-roll parameters given the potential and duration of the reheating era. This will clarify how each inflationary model predicts the primordial power spectra observed by CMB experiments such as LiteBIRD.

10.1.1. Validity of the slow-roll trajectory and initial conditions

Before introducing the reheating era, let us illustrate how the slow-roll assumption translates for the simple *Large-Field Inflation* (LFI) potential [92, 240–243] defined as

$$V(\phi) = M^4 \left(\frac{\phi}{M_{\rm Pl}}\right)^p , \tag{10.1}$$

where M and p are the two parameters of the potential (respectively of mass dimension and dimensionless). This potential can be motivated by various theoretical models, for instance, supergravity [244, 245]. At slow-roll leading order, one can compute their first Hubble slow-roll parameters from Eqs. 3.33 and 3.34 [92]:

$$\varepsilon_1 \overset{\text{SRLO}}{\simeq} \frac{p^2}{2\phi^2/M_{\rm Pl}^2} \text{ and } \varepsilon_2 \overset{\text{SRLO}}{\simeq} \frac{2p}{\phi^2/M_{\rm Pl}^2}. \tag{10.2}$$

As long as $\varepsilon_1 < 1$, *i.e.*, in slow-roll, $-pM_{\rm Pl}/\sqrt{2} < \phi < pM_{\rm Pl}/\sqrt{2}$, the slow-roll approximation is valid and all trajectories are attracted to the slow-roll trajectory given by $(\phi,\ \dot\phi) = (\phi,\ -M_{\rm Pl}\, p\, \phi^{p/2-1}/\sqrt{3})$, see Eq. 3.31.

On the lower-left panel of Fig. 10.1 and in the particular case in which $p = 2$ and $m_\phi \equiv \sqrt{2}M^2/M_{\rm Pl} = 10^{-6}M_{\rm Pl}$, we represent this slow-roll trajectory with a dotted-red line. Additionally, we show a few examples of numerical exact trajectories obtained from directly integrating the Friedmann equation, assuming various initial field values (depicted on the top-right panel) and a zero initial velocity. One observes the attractor behaviour of the slow-roll trajectory, which is quickly joined independently on the initial condition. As long as $\varepsilon_1 < 1$, all trajectories stay in the accelerating region, represented by the blue region. When ε_1 becomes greater than 1 (in our case when $|\phi| < \sqrt{2}$), depicted by red crosses in Fig. 10.1, the slow-roll regime is violated and inflation stops.

In the right panel, we quantify what is implied for the background evolution of the Universe for a particular initial field value, $\phi_i = 15M_{\rm Pl}$, and initial velocity, $\dot\phi_i$, chosen so that the initial condition is already slow-rolling. The upper panel represents the Hubble rate as a function of the time. The two curves correspond to exact (dotted blue) or slow-roll (orange line) calculation, *i.e.* including or not the kinetic term in the Friedmann equation (both numerically obtained). One observes that as long as $\varepsilon_1 < 1$, H is quasi-constant in both cases. In this regime, the associated number of e-folds (that is obtained from integrating the curves on the upper panel) follows a trend that can be analytically obtained in slow-roll approximation from Eq. 3.32, namely $\Delta N = \frac{1}{2p}(\phi^2/M_{\rm Pl}^2 - \phi_{\rm end}^2/M_{\rm Pl}^2)$ [92]. Black-dotted points represent it. One could also represent the energy density as a function of the scale factor, obtaining a quasi-flat behaviour as long as slow roll is realised.

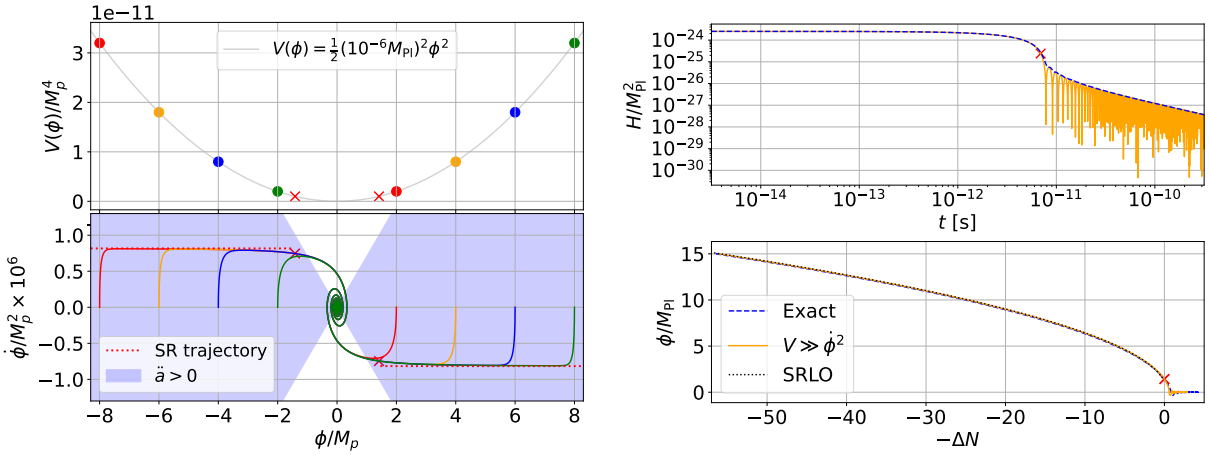

Figure 10.1: LEFT: On the upper panel, we illustrate in lightgrey the potential $V(\phi) = \frac{1}{2}m_\phi^2\phi^2$ where $m_\phi = 10^{-6}M_{\rm Pl}$ for eight different initial conditions ϕ_i. On the lower panel, we illustrate the $(\dot\phi,\ \phi)$ phase space of the field. The colourful trajectories are obtained from numerically integrating Eq. 3.23 with the initial conditions ϕ_i illustrated above and $\dot\phi_i = 0$. The red dotted line corresponds to the slow-roll trajectory Eq. 3.31. The blue region displays the region where the expansion is accelerated, *i.e.* when Eq. 3.24 is satisfied. The end of inflation, $\phi_{\rm end}$ is defined as the moment when the trajectory exits for the first time the blue region.
RIGHT: For the same potential, and with an initial condition $\phi_i = 15M_{\rm Pl}$ at slow-roll velocity, we represent on the upper panel the value of the Hubble expansion rate as a function of time. The dashed blue line is obtained from applying the Friedmann equation to the numerical trajectory, as for the orange line, but for which we have neglected $\dot\phi^2/2$ with respect to V. The lower panel displays the field value as a function of the number of e-folds, which is obtained with Eq. 3.7 applied the numerical (blue) and the slow-roll Hubble expansion rate (orange). We represent the opposite of ΔN to keep left-to-right time reading. On these two panels, the red cross illustrates the end of inflation.

When the inflationary trajectory leaves the slow-roll trajectory (*i.e.* exceed the red-cross point), inflation stops, and the energy contained in the Hubble rate is dissipated while ϕ coherently oscillates in the bottom of the potential [246].

This phase is correctly described by the simple Friedmann equation (*i.e.* by the blue curve of Fig. 10.1 in the case of ϕ^2 inflation) only if the inflaton ϕ have no coupling to other species. During this phase, the overall fluid parameter $\langle w\rangle$ is no longer minus one. Averaging the Friedmann equation Eq. 3.23 over time in which one neglects the friction term with respect to the acceleration term, one gets:

$$\left\langle \dot\phi^2 \right\rangle \simeq \langle \phi V_\phi(\phi)\rangle\,. \tag{10.3}$$

Inserted in Eqs. 3.21 and 3.20, it yields [246]:

$$\langle w\rangle = \frac{\langle \phi V_\phi(\phi)\rangle - 2\langle V(\phi)\rangle}{\langle \phi V_\phi(\phi)\rangle + 2\langle V(\phi)\rangle} = \frac{p-2}{p+2}, \tag{10.4}$$

assuming $V(\phi) \sim \phi^p$ in the oscillating region around the potential minimum. For the simplest potentials, which are well described by $p = 2$, this phase gives $\langle w\rangle = 0$, *i.e.* it behaves as a matter-dominated era.

However, to explain the transition from inflation to radiation-dominated Universe, one needs to transfer the vacuum energy ρ_ϕ into light species of density ρ_r with some inflaton decay rate Γ_ϕ. The time of equality $\rho_\phi = \rho_r$ is called *reheating*, and the era between the end of inflation and reheating is called *reheating phase*. Hereafter, we call $w_{\rm reh}$ the average equation of state during reheating, which simplest approximation is $\bar{w}_{\rm reh} = \langle w \rangle$, and $\rho_{\rm reh}$ is the total density energy at reheating. In the case of these two interacting species, the two background-level Boltzmann equations and the Friedmann equations lead to the following system [247]

$$\dot{\rho}_\phi + 3H\rho_\phi = -\Gamma_\phi(1 + \langle \bar{w}_{\rm reh} \rangle)\rho_\phi \tag{10.5}$$

$$\dot{\rho}_r + 4H\rho_r = \Gamma_\phi(1 + \langle \bar{w}_{\rm reh} \rangle)\rho_\phi \tag{10.6}$$

$$3M_{\rm Pl}^2 H^2 = \rho_\phi + \rho_r, \tag{10.7}$$

which, in general, can be easily solved.

Now that we have introduced reheating, let us explain how we treat it in the context of slow-roll prediction. The non-trivial remaining step to derive observables from the theories is the calculation of the field value at Hubble crossing, ϕ_*, which appears in the slow-roll expressions for the primordial observables (see Eqs. 3.42, 3.43, 3.44, 3.45).

10.1.2. Reheating consistent predictions of slow-roll inflation

This is the topic of this section, partly extracted from [93], which is itself based on [92, 248, 249]. We have seen that on ϕ^2 potential, any trajectory is attracted towards the slow-roll trajectory. In the general case, slow roll is a dynamical attractor, so in single-field models, it singles out one phase-space trajectory. Therefore, if one knows the end point of inflation $\phi_{\rm end}$ as well as the number of e-folds ΔN_* elapsed between the time the pivot scale k_* crosses out the Hubble radius and the end of inflation, the value of ϕ_* can be derived from

$$\Delta N_* = \int_{\phi_{\rm end}}^{\phi_*} \frac{{\rm d}\phi}{\sqrt{2\varepsilon_1}} \simeq \int_{\phi_{\rm end}}^{\phi_*} \frac{V(\phi)}{V_\phi(\phi)} {\rm d}\phi \,. \tag{10.8}$$

Here, the first part of the formula was obtained from Eq. 3.7 jointly with the expression of ε_1 in terms of ϕ and $\dot{\phi}$ given by Eq. 3.25, and the second part of the formula follows from Eq. 3.33. The end point $\phi_{\rm end}$ can be obtained from solving $\varepsilon_1 = 1$, using slow-roll formula (3.33) for ε_1.

Fig. 3.1 illustrates that a geometric-closure relationship can be established, allowing us to determine ϕ_*. The number of e-folds ΔN_* can be obtained from its definition: the CMB pivot scale exited the horizon when $\phi = \phi_*$. ΔN_* depends on the Hubble scale during inflation and on the expansion history from the end of inflation until now. More precisely, it is given by [249–251]

$$\Delta N_* = \ln R_{\rm rad} - \ln\left(\frac{k_*}{a_0 \tilde{\rho}_\gamma^{1/4}}\right) - \frac{1}{4}\ln\left[\frac{9V_{\rm end}}{\varepsilon_{1*}(3 - \varepsilon_{1{\rm end}})V_*}\right] + \frac{1}{4}\ln\left(8\pi^2 A_{\rm s}\right), \tag{10.9}$$

where "end" denotes quantities computed at the end of inflation. This equation is implicit for ΔN_*. In this expression, $\tilde{\rho}_\gamma$ is the energy density of radiation today rescaled by the change in the number of relativistic degrees of freedom between the completion of reheating and today, and $\ln R_{\rm rad}$ is the so-called "reheating parameter". It quantifies the additional number of e-folds induced by a reheating with respect to an instantaneous transition to a radiation-dominated era. This definition is illustrated in Fig. 3.1. Formally, this quantity can be expressed as

$$\ln R_{\rm rad} = \frac{1 - 3\bar{w}_{\rm reh}}{12(1 + \bar{w}_{\rm reh})} \ln\left(\frac{\rho_{\rm reh}}{\rho_{\rm end}}\right), \tag{10.10}$$

where $\bar{w}_{\rm reh}$ is the mean equation-of-state parameter during reheating and $\rho_{\rm reh}$ is the Universe's energy density at the onset of the radiation era. If reheating is instantaneous ($\rho_{\rm reh} = \rho_{\rm end}$), or if its equation of state is the one of radiation ($\bar{w}_{\rm reh} = 1/3$), then $\ln R_{\rm rad} = 0$. Note that it is common practice to approximate the third term in Eq. 10.9 by its slow-roll limit; the impact of this approximation is further discussed in Sec. 11.4.1.2.

In principle, when we know the decay rate from inflaton to light particles, Γ_ϕ, $\ln R_{\rm rad}$ could be computed. In most cases, we don't want to assume a particular decay rate, and we simply parameterise reheating by $\ln R_{\rm rad}$. However, this parameter is not totally arbitrary. One should ensure that the Universe is totally reheated at Big-Bang Nucleosynthesis (BBN) and that reheating happens after inflation ends, leading to bounds on $\rho_{\rm reh}$

$$\rho_{\rm BBN} < \rho_{\rm reh} < \rho_{\rm end}, \tag{10.11}$$

$$-\frac{1}{3} < w_{\rm reh} < 1. \tag{10.12}$$

where the bounds on w ensure that $\ddot{a} < 0$, *i.e.* inflation is ended, while keeping the physical upper bound. Finally, one can recover the energy density at reheating by inverting Eq. 10.10.

10.1.3. Predicting cosmological observables

Eventually, we have everything to predict cosmological observables, given a potential shape. We can compute the field value of the inflaton at the time of Hubble crossing of the pivot scale, evaluate Eqs. 3.33, 3.34 and 3.35 at that location, and insert the result into Eqs. 3.42, 3.43, 3.44 and 3.45. Because $A_{\rm s}$ is well measured, as done in [92], when we can write $V(\phi) = M^4 v(\phi)$, M is no longer a degree of freedom since it is given by

$$\left(\frac{M}{M_{\rm Pl}}\right)^4 = 24\pi^2 \frac{\varepsilon_{1*}}{v\left(\phi_*\right)} P_*, \tag{10.13}$$

which directly derives from Eq. 3.42. Hence, we decrease the number of parameters by one.

For instance, we illustrate for LFI the (p, $\log_{10}\rho_{\rm reh}$) grid of ($n_{\rm s}$, r) predictions. These predictions are obtained by ensuring that the bounds Eqs. 10.11 and 10.12 are satisfied. For illustration purposes, we have superimposed Fisher ellipses for mock Planck data (which corresponds to Planck constraints on r) and mock LiteBIRD data (with expected LiteBIRD noise and $\sigma_r = 10^{-3}$ which corresponds to the requirement on δr the total uncertainty). The fiducial Λ-CDM + r parameters are given in Table 4.1 ($r_{\rm fid} = 0$). The 1σ and 2σ contours are displayed. One observes that Planck mock data already severely constrains most of the scenarios. This is discussed with actual Planck data in [64].

We also provide a second historically important example, the *Higgs-Starobinsky Inflationary* (SHI) potential:

$$V(\phi) = M^4 \left(1 - \exp^{-\sqrt{\frac{2}{3}}\frac{\phi}{M_P}}\right)^2. \tag{10.14}$$

Once again, this potential can originate from different phenomenologies. The first most standard one is modified gravity [67, 252, 253]. In this case, one adds in the Einstein-Hilbert action a quadratic term in the Ricci scalar, whose strength is determined by a mass parameter μ. The Ricci scalar treated as the inflation field evolves in Eq. 10.14, where As essentially fixes $\mu = 2\sqrt{2}M^2/M_{\rm Pl}$. For the points displayed in Fig. 10.2, this yields a $\mu \in [2.85, 3.75] \times 10^{-5} M_{\rm Pl}$.

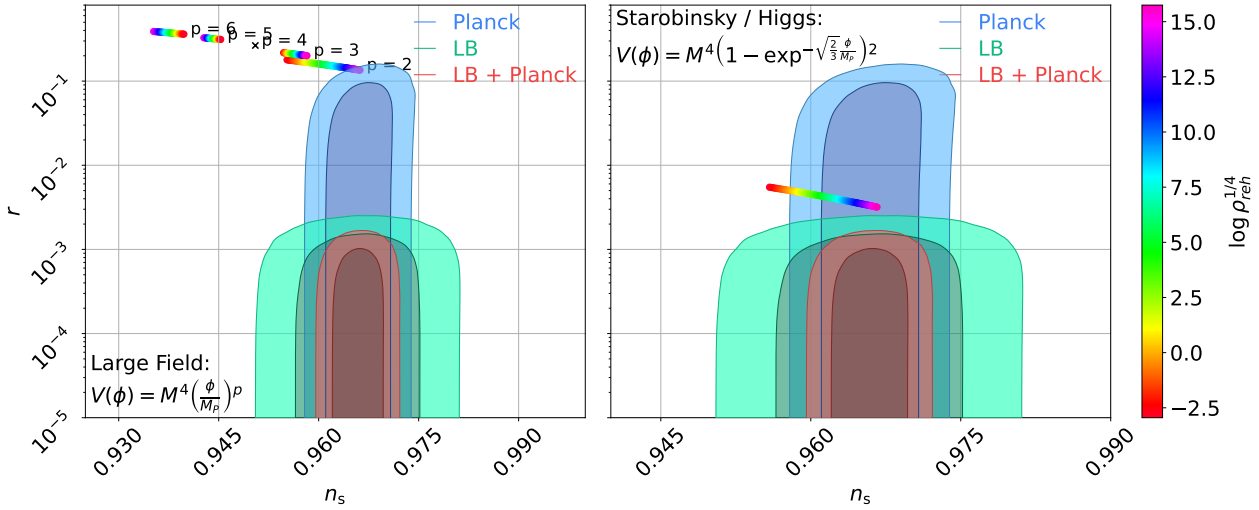

Figure 10.2: LEFT: LFI RIGHT: SHI (p, $\log_{10}\rho_{\rm reh}$) grid of ($n_{\rm s}$, r) predictions. $\rho_{\rm reh}$ is in GeV^4. Fisher ellipses for mock Planck data and mock LiteBIRD data are superimposed, with fiducial Λ-CDM + r parameters given in Table 4.1 (and $r_{\rm fid} = 0$).

The second most standard phenomenology is taking the Higgs boson, non-minimally coupled to gravity with a coupling constant ξ, as the inflaton [254–256]. This scalar field evolves in Eq. 10.14, where $\mu = 2\sqrt{2}M^2/M_{\rm Pl}$. The ratio $\xi/\sqrt{\lambda} = \frac{M_{\rm Pl}^2}{2M^2}$, where $\lambda \simeq 0.13$ is the Standard Model Higgs self-interacting coupling constant, is fixed by $A_{\rm s}$. For the points displayed in Fig. 10.2, this yields a $\xi/\sqrt{\lambda} \in [3.76, 4.96] \times 10^4$ (this high value of the coupling constant is commented in [257]).

This potential and most of the associated reheating scenarios are currently allowed by Planck measurement of $A_{\rm s}$ and $n_{\rm s}$, as well as existing constraints on r. However, the SHI predictions are well inside the sensitivity of the LiteBIRD satellite, which should be able to exclude the model completely or put constraints on the associated reheating scenarios.

One can refer to [92] for a review of many more potentials. The next step in the analysis performed, for instance, in [101], is to provide an overview of the constraints on this wide landscape of slow-roll potentials and define a criterion for comparing them. In [101], they use the Bayesian evidence and provide a complete Bayesian ranking of 286 slow-roll models.

Most of these models are physically motivated, but they rarely come with a complete embedding in a theoretically well-defined background. The drawback of this could be ignoring potential issues of realistic slow-roll models. The aim of this third part is to address this problem.

10.2. Going beyond the particle standard model

In this section, we first present the motivations for going beyond the particle-physics standard model. Then, we introduce the Minimal Supersymmetric Standard Model, which is one of the most standard and minimal supersymmetric extension to the Standard Model. Among other motivations, it can provide a stable and neutral Weakly Interacting Massive Particle (WIMP), which is a dark matter candidate.

10.2.1. The Standard Model of particle physics and its open questions

The Standard Model (SM) of particle physics is a gauge quantum field theory with the internal symmetries of $\mathrm{SU}(3)_C \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$. It describes the fundamental particles: leptons, quarks, gauge bosons, and the Higgs boson (see [258] for an introduction to Standard Model and [259] for a state of the art on Standard Model constraints in 2024). The Standard Model has been, so far, very successful in making experimental predictions.

However, various interrogations remain:

- The SM initially assumed neutrinos were massless, which conflicts with the observed phenomenon of neutrino oscillations [260], indicating they have mass [259].
- The characteristic scale of the SM, the Higgs boson mass, is orders of magnitude below the Planck scale [261]. This vast difference is called the *Hierarchy Problem*.
- The SM does not explain why the strong force does not break CP symmetry, a problem known as the *strong CP problem* [262].

Viewed from the cosmology perspective, other unexplained phenomena are worth to mention:

- The Standard Model does not include dark matter candidates, hence it cannot explain the $\Omega_\mathrm{c} h^2 = 0.11933 \pm 0.00091$ measured by the Planck satellite + BAO (see Tab. 2.1).
- The Standard Model vacuum energy prediction is orders of magnitude larger than $\Omega_\Lambda = 0.6889 \pm 0.0056$ measured by the Planck satellite + BAO (see Tab. 2.1), a problem usually referred to as the *cosmological constant problem*.
- The observed baryon asymmetry in the universe (the imbalance between matter and antimatter) is not adequately explained by the SM.
- The inclusion of gravity with the SM remains an open problem, as the SM does not incorporate a quantum theory of gravity.
- The SM does not provide an inflaton candidate (Higgs inflation that we have discussed in Sec. 10.1.3 requires a non-minimal coupling to gravity).

10.2.2. Supersymmetry as an extension to the Standard Model

In this perspective, the extensions of the Standard Model of particle physics based on supersymmetry (SUSY) (see *e.g.* [263, 264] for reviews), such as the Minimal or Next-to-Minimal Supersymmetric Standard Model (MSSM, NMSSM) provide an appealing theoretical framework that:

- naturally provides a *Weakly Interacting Massive Particle* (WIMP) that can explain the fraction of dark matter observed in the Universe today $\Omega_\mathrm{c} h^2$,
- allows one to describe the inflation era (as they naturally include flat directions [265–268] that could support inflation [269–280]),
- while explaining the physical processes that can be measured at the LHC [281].

Such a theory can predict multiple observables, both for cosmology and particle-physics (HEP), which can then be compared to measurements to more accurately assess the favoured/disfavoured area in parameter space within a coherent description of our universe.

Let us first briefly overview supersymmetry and the MSSM. We mainly focus on introducing the notations, while a complete introduction to MSSM can be found in [264].

10.2.2.1 Supersymmetry: definition

Supersymmetry (SUSY) is a speculative additional spacetime symmetry upon the Poincaré one, which relates fermions and bosons. On top of the translation and Lorentz generators, P^μ and $M^{\mu\nu}$, supersymmetry[1] introduces the discrete supersymmetric transformation Q such that

$$Q|\text{boson}\rangle = |\text{fermion}\rangle \text{ and } Q|\text{fermion}\rangle = |\text{boson}\rangle, \tag{10.15}$$

which satisfies the following algebra:

$$\left\{Q, Q^\dagger\right\} = P^\mu, \tag{10.16}$$

$$\{Q, Q\} = \left\{Q^\dagger, Q^\dagger\right\} = 0, \tag{10.17}$$

$$[P^\mu, Q] = \left[P^\mu, Q^\dagger\right] = 0. \tag{10.18}$$

where ($\{.,.\}$) $[.,.]$ denotes the (anti-)commutator of two operators.

A *supermultiplet* is the irreducible representation of this algebra, which consists of a pair of *superpatners*. According to the spin-statistics theorem, one can show that the number of baryons n_B equals the number of fermions n_F within a supermultiplet. Any supermultiplet can be reduced to the three following categories:

1) the *chiral supermultiplet*: 1 Weyl fermion (*ie.* a massless spin-1/2 particle with $n_\text{F} = 2$) and 2 real scalars (each with $n_\text{B} = 1$),
2) the *gauge supermultiplet*: 1 massless spin-1 boson ($n_\text{B} = 2$) and 1 Weyl fermion ($n_\text{F} = 2$) called *gaugino* because the superpartner of a gauge boson,
3) the *gravity supermultiplet*: 1 massless spin-2 *graviton* ($n_\text{B} = 2$) + 1 spin-3/2 *gravitino* (F = 2).

10.2.2.2 A supersymmetric extension of the SM

SUSY extends the SM, by describing each known particle as part of a supermultiplet, conserving the standard model local symmetries under any $\text{SU}(3)_C \times \text{SU}(2)_L \times \text{U}(1)_Y$ transformation. Hence, it introduces as many new supersymmetric particles as known SM ones. The standard notation consists of denoting with a $\tilde{.}$ the superpartner of a given SM particle. Then, the field content is defined as follows [264]:

- Since the SM quarks and leptons do not transform under the gauge group in the same way as their right-handed parts, they must be part of chiral supermultiplets with ***scalar*** bosonic superpartners, which are called ***squarks*** and ***sleptons***. The left-handed chiral supermultiplets are $\text{SU}(2)_L$-doublets Q or L, while the right-handed supermultiplets are $\text{SU}(2)_L$-singlets. The squark/quark supermultiplets are listed in the first block of Tab. 10.1, while the slepton/lepton supermultiplets are displayed in the second block. The conserved charges under Gauge transformations are displayed in the last column.

[1] To be able to incorporate a description of gravity in the theory, one usually promotes supersymmetry to a local symmetry.

Names		spin 0	spin 1/2	$SU(3)_C, SU(2)_L, U(1)_Y$
squarks, quarks	Q	$\left(\widetilde{u}_L\ \widetilde{d}_L\right)$	$(u_L\ d_L)$	$\left(\mathbf{3}, \mathbf{2}, \frac{1}{6}\right)$
(3 families)	$\bar{u}$	$\widetilde{u}_R^*$	$u_R^\dagger$	$\left(\overline{\mathbf{3}}, \mathbf{1}, -\frac{2}{3}\right)$
	$\bar{d}$	$\widetilde{d}_R^*$	$d_R^\dagger$	$\left(\overline{\mathbf{3}}, \mathbf{1}, \frac{1}{3}\right)$
sleptons ,& leptons	L	$(\widetilde{\nu}\ \widetilde{e}_L)$	$(\nu\ e_L)$	$\left(\mathbf{1}, \mathbf{2}, -\frac{1}{2}\right)$
(3 families)	$\bar{e}$	$\widetilde{e}_R^*$	$e_R^\dagger$	$(\mathbf{1}, \mathbf{1}, 1)$
Higgs & Higgsinos	H_u	$(H_u^+\ H_u^0)$	$\left(\widetilde{H}_u^+\ \widetilde{H}_u^0\right)$	$\left(\mathbf{1}, \mathbf{2}, +\frac{1}{2}\right)$
	H_d	$\left(H_d^0\ H_d^-\right)$	$\left(\widetilde{H}_d^0\ \widetilde{H}_d^-\right)$	$\left(\mathbf{1}, \mathbf{2}, -\frac{1}{2}\right)$

Table 10.1: Chiral supermultiplets in the Minimal Supersymmetric Standard Model. From [264]

- For electroweak consistency reasons, the Higgs sector requires two chiral supermultiplets. H_u and H_d are the two $\mathrm{SU}(2)_L$-doublet complex scalar fields corresponding to (H_u^+, H_u^0), (H_d^0, H_d^-). They are associated with Weyl fermions called *Higgsinos*. These are displayed in the third block of Tab. 10.1.
- The SM spin-1 bosons, g, $W^\pm$, W^0, B^0 appear in gauge supermultiplets and come with gaugino partners (respectively called *gluino*, *wino*, *bino*). These are displayed on Tab. 10.2.

Names	spin 1/2	spin 1	$SU(3)_C, SU(2)_L, U(1)_Y$
gluino, gluon	$\widetilde{g}$	g	$(\mathbf{8}, \mathbf{1}, 0)$
winos, W bosons	$\widetilde{W}^\pm\widetilde{W}^0$	$W^\pm W^0$	$(\mathbf{1}, \mathbf{3}, 0)$
bino, B boson	$\widetilde{B}^0$	B^0	$(\mathbf{1}, \mathbf{1}, 0)$

Table 10.2: Gauge supermultiplets in the Minimal Supersymmetric Standard Model. From [264]

The superpartners have the same gauge couplings as the SM ones. However, they do not have the same mass (otherwise, they would have been detected by now). Hence, supersymmetry has to be *broken*, typically *softly*, meaning that the breaking has a positive mass dimension (hence a typical mass M_{SUSY}) and that the complete Lagrangian of the theory can be decomposed as $\mathcal{L} = \mathcal{L}_{\mathrm{SUSY}} + \mathcal{L}_{\mathrm{soft}}$. The mass splitting (which is of order M_{SUSY}) and the mixing between the SUSY states is determined by $\mathcal{L}_{\mathrm{soft}}$.

10.2.2.3 Generic SUSY Lagrangian and scalar potential

Let us first focus on $\mathcal{L}_{\mathrm{SUSY}}$, whose detailed generic derivation is done in [264]. It involves every supermultiplets; hence, in particular, complex scalar fields φ_i, left-handed Weyl fermions ψ_i, as well as bosonic field strengths $F^a_{\mu\nu}$ appear. This Lagrangian is often described in terms of *superpotential* defined as

$$W \equiv \frac{1}{2} M^{ij} \varphi_i \varphi_j + \frac{1}{6} y^{ijk} \varphi_i \varphi_j \varphi_k, \tag{10.19}$$

where M^{ij} is a symmetric mass matrix for the fermion fields, and y^{ijk} is a Yukawa coupling of a scalar ϕ_k and two fermions $\psi_i \psi_j$ that must be symmetric under interchange of i, j, k. From this

Lagrangian, one can derive the *scalar potential* that consists of two terms

$$V \equiv V_{\rm F} + V_{\rm D}, \tag{10.20}$$

respectively called *F-terms* and *D-terms* which can be derived from

$$V_{\rm F} = \sum_i \left|F^i\right|^2, \tag{10.21}$$

$$V_{\rm D} = \frac{1}{2}\sum_a g_a^2 \vec{D}^a \cdot \vec{D}^a, \tag{10.22}$$

where the *complex auxiliary field* F^i can be derived from the superpotential

$$F^i \equiv \frac{\partial W}{\partial \varphi_i}, \tag{10.23}$$

and the real *bosonic auxiliary field* D^a from

$$\vec{D}^a = \sum_j \varphi_j^* \vec{T}^a \varphi_j, \tag{10.24}$$

where the i index runs over the three gauge groups, with the corresponding gauge coupling g_i, and generators $\vec{T}^i$. The *flat directions* $\varphi_{\rm flat}$ of a SUSY scalar potential are defined such that they are both *D-flat* and *F-flat*, *i.e.*:

$$V(\varphi_{\rm flat}) = V_{\rm F}(\varphi_{\rm flat}) = V_{\rm D}(\varphi_{\rm flat}) = 0. \tag{10.25}$$

We have mentioned the second contribution to the Lagrangian of the theory, which consists in soft-SUSY-breaking terms, $\mathcal{L}_{\rm soft}$. A generic form for these terms can be found in [264]. These also contribute to the scalar potential through F and D terms, similarly to the SUSY part. A flat direction in regards to $\mathcal{L}_{\rm SUSY}$ will typically be *lifted* by the effect of the soft-SUSY-breaking (*i.e.* V becomes $\neq 0$). This lifting effect will be key to understanding how one embeds inflation in MSSM in Sec. 11. This soft-SUSY-breaking is the first lifting mechanism we will consider, the second one being the contribution of effective nonrenormalisable terms to the potential. Such beyond-MSSM terms are expected at very high energy, where the theory becomes sensitive to ultraviolet completion effects. We will see what kind of term this effect induces for the scalar potential in Chapter 11.

10.2.3. The Minimal Supersymmetric Standard Model

Today, we know experimentally that the proton decay time is $> 3.6 \times 10^{29}$ years [259]. More generally, lepton number L and baryon number B are severely constrained from being conserved in all observed processes. To account for this, the MSSM is built upon the assumption of the foundational conservation of the *matter parity* [282, 283], which prevents B or L violating terms from appearing in the renormalisable superpotential. With this conservation law, the B and L violations are severely suppressed.

10.2.3.1 *R*-parity and MSSM

Additionally, to ensure that particles in the same multiplet come with different parity numbers, one often works with the *R-parity* [284], defined as:

$$P_R \equiv (-1)^{3(\mathrm{B-L})+2s}. \tag{10.26}$$

where s is the spin of the particle considered. This parity definition is useful: on top of enforcing B and L conservation, SM particles come with even R-parity, $P_R = 1$, while SUSY particles come with odd R-parity, $P_R = -1$. The phenomenological consequence of this parity conservation is that no mixing between SUSY and SM particles can happen and that the *Lightest Supersymmetric Particle* (LSP) is stable since it cannot decay in lighter SM states.

Without further assumption, the MSSM superpotential that conserves R-parity can be written as

$$W_{\mathrm{MSSM}} = \bar{u}\mathbf{y_u}QH_u - \bar{d}\mathbf{y_d}QH_d - \bar{e}\mathbf{y_e}LH_d + \mu H_u H_d. \tag{10.27}$$

Additionally, the most general soft-SUSY breaking Lagrangian that conserves R-parity leads to considering 105 parameters [285], which results in a significant arbitrariness. Hence, on top of conserving R-parity, MSSM is usually defined with additional assumptions phenomenologically well-justified, flavour mixing and CP-violating processes being experimentally severely restricted. In the following, we only consider a subset of the unconstrained parameters, as is frequently done in various studies (see *e.g.* [286]). They are listed in Table 10.3. We will refer to this *phenomenological MSSM*, thus defined, as *pMSSM*. We assume all the parameters to be real-

Origin	Parameters
Higgs-sector	$m^2_{H_u}$, $m^2_{H_d}$, $\tan\beta$, $\mathrm{sgn}(\mu)$
Gaugino masses	M_1, M_2, M_3
Slepton masses	$m_{\tilde{l}_L^{12}}$, $m_{\tilde{e}_R^{12}}$, $m_{\tilde{\tau}_L}$, $m_{\tilde{\tau}_R}$
Squark masses	$m_{\tilde{q}_L^{12}}$, $m_{\tilde{u}_R^{12}}$, $m_{\tilde{d}_R^{12}}$, $m_{\tilde{q}_L^3}$, $m_{\tilde{t}_R}$, $m_{\tilde{b}_R}$
A-terms	$A_{u_{12}}$, $A_{d_{12}}$, $A_{l_{12}}$, A_t, A_b, A_τ

Table 10.3: Parameters of the phenomenological MSSM

valued, which implies, among other things, no extra sources of explicit *CP*-violation from the SUSY extension of the Standard Model. As the $\mathrm{K}_0 - \overline{\mathrm{K}}_0$ mixing limits the mass splitting between the first and second squark generations [286, 287], we assume that the squark masses of the first two generations are equal. The same applies to the slepton masses and the A-terms. The additional hypotheses with which we will work in Chapter 11 will be explicited in Sec. 11.2.1. Under these approximations, many terms of W_{MSSM} can be ignored.

The *electroweak symmetry breaking* ($\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y \to \mathrm{U}(1)_{\mathrm{EM}}$), observed experimentally is described within the MSSM, allowing the prediction of the Higgs mass m_h. The scale of ElectroWeak Symmetry Breaking (EWSB) is given by [288, 289]:

$$M_{\mathrm{EWSB}} = \sqrt{m_{\tilde{t}_1} m_{\tilde{t}_2}}\ . \tag{10.28}$$

The mass eigenstates of the four Higgs neutral degrees of freedom are called *neutralinos* denoted χ_i^0 with $i = 1, 2, 3, 4$ (by order of mass). Instead, charged ones are the two positive and negative

charginos. The associated soft masses can be computed from the MSSM parameters' values, as well as the sleptons, squarks and gluinos mass matrices. Gluinos do not mix with other states as they belong to a color octet.

Similar to the SM parameters, the MSSM parameters require renormalisation. Their scale dependence is governed by their is governed by their *renormalisation group equations* (RGEs), resulting in an interdependence among many of the MSSM parameters. We denote β_X the derivative of a parameter X with respect to the logarithm of the scale Q:

$$\beta_X \equiv \frac{\partial X}{\partial \ln Q}. \tag{10.29}$$

We will introduce RGEs relevant to our work in Sec. 11.1.2. Let's mention that the MSSM RGEs of the gauge couplings (g_1, g_2, g_3) play in favour of MSSM over SM: these quantities' running meet at high energies, typically around $\sim 3 \times 10^{16}$ GeV in MSSM as opposed to in SM. This energy scale (which we have already mentioned in Chapter 3) is called *Grand Unified Theory scale*, at which we expect that the local symmetry under some simple group, as SU(5) or SO(10), broke into the standard gauge group $\mathrm{SU}(3)_C \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$.

10.2.3.2 Soft-SUSY breaking

Boundary conditions for the RGEs can be defined at a high scale (typically the soft-SUSY-breaking scale M_{SUSY}). This requires a description of the soft-SUSY-breaking mechanism. Several ones have been proposed. For instance, the *O'Raifeartaigh SUSY breaking* consists of lifting a flat direction with quantum corrections to break SUSY. The associated *"hidden" soft-SUSY-breaking sector* is mostly independent of the *"visible" MSSM sector*, so the transition is soft. The only interaction between the two sectors is the one responsible for generating the MSSM soft terms through a mediator called the *goldstino*.

Two main mediation scenarios have been proposed, *gauge-* or *high-scale-mediated SUSY breaking*. The *minimal supergravity model* (mSUGRA) is a high-scale-mediated SUSY breaking scenario that transmits the explicit SUSY breaking in the hidden sector as a soft breaking in the visible sector through *gravitational interactions*[2]. mSUGRA restricts the MSSM parameter space by unifying all scalar masses, all trilinear couplings and all Gaugino masses to common values at M_{SUSY}. These are respectively denoted m_0, A_0, and $m_{1/2}$. The Polonyi model proposes a further simplification [291], in which:

$$m_0^2 = m_{3/2}^2, \tag{10.30}$$

$$A_0 = (3 - \sqrt{3}) m_{3/2}, \tag{10.31}$$

$$m_{1/2} = \mathcal{O}\left(m_{3/2}\right) \tag{10.32}$$

at M_{SUSY}. Even though we will not entirely rely on the severe assumptions of the Polonyi-model mSUGRA in this work, we will bound the RGEs for the trilinear coupling with Eq. 10.31. As a further simplification inspired by mSUGRA, we will also set in this work

$$m_{\tilde{q}_L^{12}} = m_{\tilde{u}_R^{12}} = m_{\tilde{d}_R^{12}}\,. \tag{10.33}$$

[2]In this case, the SUSY breaking scale yields a mass to the gravitino $m_{3/2}$ that is not *a priori* heavier than the rest of the MSSM spectrum. We will mention the resulting *gravitino problem* [284, 290] in Sec. 11.5.2.2.

10.2.3.3 Dark matter as an LSP WIMP

We have seen in Sec. 2.4 that the cold-dark matter measured by Planck today amounts to $0.12/h^2 \simeq 25\%$ of the total energy content of the Universe. A popular candidate for dark matter is a *Weakly Interacting Massive Particle* (WIMP), say of mass m_X. As soon as the temperature of the Universe drops below m_X, the WIMP decouples from the primordial plasma and *freezes out*. Then, Eq. 10.5 allows to predict that [292]:

$$\Omega_X h^2 \sim 3 \times 10^{-27}\ \mathrm{cm^3 s^{-1}} \times \frac{m_X^2}{\alpha^2} \sim 3 \times (100\ \mathrm{TeV})^{-2} \times \frac{m_X^2}{\alpha^2} \tag{10.34}$$

where α is the weak coupling constant of X. If the WIMP considered has a mass $m_X = 1$ TeV (*resp.* 100 GeV), one needs a weak coupling constant $\alpha \sim 10^{-4}$ (*resp.* 10^{-2}) to have a heavy relics' density compatible with the Planck measurement.

Among the several WIMP dark matter candidates in the MSSM, which are neutral and stable LSP, we will focus on the case of the lightest neutralino $\tilde{\chi}_1^0$. To predict its freeze-out temperature and the resulting relic density today, one needs to know its annihilation channel towards a lighter final state before decoupling. For instance, neutralinos can lead to the right relics' density by annihilating into [293]:

- Two W bosons via a t-channel: $\tilde{\chi}_1^0\tilde{\chi}_1^0 \to W^+W^-$. This scenario is called *Higgsino-like dark matter* [294] since, in this case, $\tilde{\chi}_1^0$ is mainly a Higgsino.
- One Higgs scalar boson (which usually decays itself into two standard model fermions): $\tilde{\chi}_1^0\tilde{\chi}_1^0 \to h$. This scenario is called *h-funnel* [295]
- One A pseudo-scalar boson (which usually decays itself into two standard model fermions): $\tilde{\chi}_1^0\tilde{\chi}_1^0 \to A$. This scenario is called *A-funnel* [295].

One can select the MSSM spectrum configurations that provide such candidates and fit for the right amount of relic density while keeping the interaction rate compatible with the direct dark-matter research programs. This will be further discussed in Sec. 11.5, in which we introduce the relevant constraints and the tools implementing the relics' density predictions.

Conclusion

This chapter served to introduce the following in two ways. First, we have discussed the role of the initial conditions in the field phase space, and explained how to make predictions for the primordial cosmological observables including the effect of the reheating era. These notions will be key for Chapter 11. Second, we have provided an overview of the theoretical background in which we will embed inflation, the MSSM. The next chapter aims to present how such an embedding can be realised and demonstrate how, as a result, the MSSM parameter space can be constrained by cosmological observations.

CHAPTER 11

MSSM-inflation: Towards a coherent description of particle physics and cosmology

In this chapter, the test case we consider is an inflationary scenario associated with two sets of flat directions, dubbed LLe and udd, naturally encompassed in the MSSM. Most of the chapter is an extraction from [93]. The potential of the inflaton candidates can then be generated only when these flat directions are lifted by the combined effects of soft-SUSY-breaking masses already present in the MSSM, and effective non-renormalisable operators that should be added to the model [270]. We will thus refer to this scenario as the effective MSSM (eMSSM) to avoid confusion when discussing the particle-physics features of the strict MSSM. The analysis in [271, 272, 275, 278] or the Generalized MSSM Inflation Model (GMSSM) analysis of [92] have been, for example, motivated by such configurations. Since the inflaton candidates are gauge invariant, the parameters of the potential depend on the energy scale at which they are evaluated through the MSSM Renormalisation Group Equations. This well-known aspect has already been studied in this specific eMSSM case to relate the inflaton mass at the scale of inflation to the one probed at the LHC [273–275]. In this part, we go one step further and study how the RGEs at one-loop level affect all parameters in the inflationary potential, and impact the way we connect cosmology and particle-physics constraints.[1]

In a first section, we revisit the eMSSM inflationary potential. In a second section, we set the stage of the analysis with the description of the phenomenological MSSM we are using, together with the observational constraints, the tools and the methodology. In a third section, we address the conditions required on the inflationary potential parameters to fulfill slow-roll inflation. In particular, we discuss the initial conditions for inflation and the required level of fine-tuning of the parameters, comparing the tree-level and one-loop inflationary potentials. In a fourth section, we identify the region of the parameter space where inflation takes place and yields to predicted values of the amplitude, $A_{\rm s}$, and tilt, $n_{\rm s}$, of the primordial spectrum that are compatible with CMB experiments. We discuss the LHC phenomenology in this region, and also

[1] The inclusion of RGEs corrections in other inflationary potentials has already been performed for instance in [296–305].

study the impact of neglecting the RGEs (as done in previous works [92, 275]) in this analysis. Finally, we illustrate the impact of the use of the one-loop potential on several eMSSM points compatible with HEP observations, the cold-dark-matter energy density and the inflationary observables.

Contents

11.1 Slow-rolling the MSSM 245
11.1.1 eMSSM flat directions and the inflationary potential 245
11.1.2 Renormalisation-group-improved potential 249
11.1.3 Slow roll in the eMSSM 251
11.2 Analysis framework 252
11.2.1 pMSSM benchmark points 252
11.2.2 Observational constraints 253
11.2.3 Tools and methodology 254
11.3 Slow-roll conditions, initial conditions, fine-tuning 256
11.3.1 Shape of the potential 256
11.3.2 Field phase-space 257
11.3.3 Fine-tuning of the inflationary potential parameters 259
11.4 Inflationary constraints in the parameter space 262
11.4.1 Results for the potential at tree level 263
11.4.2 Results for the one-loop RGEs potential 264
11.4.3 Error-budget summary 268
11.5 Combining HEP and cosmological constraints 269
11.5.1 Combined fit 269
11.5.2 Additional cross-checks 270

11.1. Slow-rolling the MSSM

In our work, inflation is realised within an effective MSSM (eMSSM), in which the LLe and udd flat directions are lifted by the combined effect of soft-SUSY-breaking masses already present in the MSSM, together with the addition of effective non-renormalisable operators. The goal of this section is to introduce the relevant parameters and derive the inflationary potential. We will treat the renormalisation-group-improved inflationary potential differently from previous studies, reaching different conclusions concerning the amount of fine-tuning in the model (cf. Sec. 11.3.3).

11.1.1. eMSSM flat directions and the inflationary potential

The scalar potential of the R-parity conserving MSSM has a large number of flat directions in the supersymmetric limit that can be lifted by soft-SUSY-breaking terms and/or by higher-dimensional supersymmetry-preserving operators [268]. This offers, in principle, various possibilities for implementing inflationary scenarios, with the inflaton being a combination of scalar fields within the MSSM. The model we study in the present part extends the MSSM by specific non-renormalisable superpotential terms that lift the so-called LLe or udd flat directions in order to trigger an inflationary phase [270].

We will first go through the main ingredients of the model. This will allow us to recall the underlying physical assumptions and possible uncertainties, introduce the notations, and clarify discrepancies between the different parameter normalisations performed in the literature.

The notations LLe and udd stand for the $\mathrm{SU(3)_c \times SU(2)_L \times U(1)_Y}$ gauge-invariant operators, $L_i \cdot L_j e_k \equiv \varepsilon_{\alpha\beta}\, L_i^\alpha L_j^\beta\, e_k$ and $(u_i \times d_j)\cdot d_k \equiv \varepsilon_{abc}\, u_i^a d_j^b\, d_k^c$, that characterise the corresponding flat directions [265–268] for specific lepton or quark generation indices i, j, k. L denotes an $\mathrm{SU(2)_L}$ doublet scalar field and e, u, d, $\mathrm{SU(2)_L}$ singlet scalar fields; a, b, c are color indices and α, β, $\mathrm{SU(2)_L}$ flavor indices. Because the ε symbols are antisymmetric with respect to all their indices, the relevant cases are obviously $i \neq j$ for LLe and $j \neq k$ for udd. The associated flat directions of the potential correspond to scalar field components satisfying the following configurations for a fixed choice of generation, flavour and color indices:

$$L_i^1 = L_j^2 = e_k \equiv \ell(x),\ (i \neq j) \tag{11.1}$$

$$u_i^a = d_j^b = d_k^c \equiv \mathrm{II}(x),\ (j \neq k, a \neq b \neq c) \tag{11.2}$$

where $\ell(x)$ and $\mathrm{II}(x)$ denote arbitrary complex-valued scalar fields, and all other scalar field components are set to zero.

In the absence of renormalisable R-parity violating terms in the superpotential, these flat directions are lifted by dimension-6 operators of the form,

$$\begin{aligned} W_6^{(L_i L_j e_k)} =& \lambda_{LLe} \frac{(L_i \cdot L_j e_k)^2}{M_{\mathrm{Pl}}^3}, \\ W_6^{(u_i d_j d_k)} =& \lambda_{udd} \frac{(u_i \times d_j \cdot d_k)^2}{M_{\mathrm{Pl}}^3}. \end{aligned} \tag{11.3}$$

As explained in Sec. 10.2.2.3, these contributions to the superpotential (where we denote the superfields by the same letters as their scalar components) can be viewed as effective operators originating from an UltraViolet (UV) completion of the MSSM after integrating out the corresponding heavy fields. The dimensionless couplings (that we take to be real-valued) λ_{LLe} and

λ_{udd} are thus expected to be typically of order one. For definiteness, we choose the corresponding mass scale to be the Planck scale. However, depending on the content of the UV completion, similar operators with lower mass scales, such as the GUT scale $M_{\rm GUT}$, can also arise. We will come back to the consequences of such a change in Eq. 11.3 in Sec. 11.5. Since we are interested in single-field inflation, we will be considering one flat direction at a time, *i.e.* one given choice of (i, j, k) for the slepton or squark generation content. (For the latter case, we also avoid extra color factor enhancement of λ_{udd} by assuming for simplicity a fixed choice for the color indices $a \neq b \neq c$)[2].

We will use the common notation $\varphi(x)$ to denote either $\ell(x)$ or $\Pi(x)$. Since in Eqs. 11.1 and 11.2 there are always three distinct canonically normalised complex-valued fields involved, the three corresponding kinetic terms in the MSSM immediately lead to the normalisation

$$\ell(x), \Pi(x) = \frac{1}{\sqrt{3}}\varphi(x) \tag{11.4}$$

in order for the complex-valued $\varphi(x)$ to have a canonical kinetic term. Taking into account the (soft) mass terms of the three fields leads also to the canonically normalised mass for $\varphi(x)$:

$$m_\varphi^2 = \frac{m_{\tilde{u}_R^i}^2 + m_{\tilde{d}_R^j}^2 + m_{\tilde{d}_R^k}^2}{3} \quad (udd), \tag{11.5}$$

$$m_\varphi^2 = \frac{m_{\tilde{l}_L^i}^2 + m_{\tilde{l}_L^j}^2 + m_{\tilde{e}_R^k}^2}{3} \quad (LLe). \tag{11.6}$$

The normalisations in Eqs. 11.4, 11.5 and 11.6 are in agreement with the literature [271, 273–275]. However, we found differences with respect to the same literature in other parts related to the normalisation of the couplings in W_6 as well as in the soft-SUSY-breaking parameter that modify the inflationary potential. To ease the discussion, we will adopt three notations to distinguish among the various occurrences of W_6. We will generically denote by $\widetilde{\lambda_6}$ the couplings appearing in Eq. 11.3,

$$\widetilde{\lambda_6} = \lambda_{LLe} \text{ or } \lambda_{udd}, \tag{11.7}$$

and by $\widetilde{\widetilde{\lambda_6}}$ the coupling appearing in W_6 along a given flat direction, defined as:

$$W_6 = \frac{\widetilde{\widetilde{\lambda_6}}}{6}\frac{\varphi^6}{M_{\rm Pl}^3}. \tag{11.8}$$

Note that the 1/6 normalisation is the one adopted in the literature [270, 271, 273–275]. From Eqs. 11.1, 11.2, 11.3 and 11.4, one obviously finds the reduction $\widetilde{\widetilde{\lambda_6}} = (2/9)\widetilde{\lambda_6}$, which is inconsistent with the identification $\widetilde{\lambda_6} = \widetilde{\widetilde{\lambda_6}} \equiv \lambda_6$ performed in the literature, where λ_6 denotes the actual coupling entering the inflaton potential. We will see that λ_6 comes with a prefactor with respect to $\widetilde{\lambda_6}$ and $\widetilde{\widetilde{\lambda_6}}$. Indeed, the proper normalisation that leads to a canonical kinetic term for the real-valued inflaton field $\phi(x)$ is

$$\varphi(x) = \frac{1}{\sqrt{2}}\phi(x)e^{i\theta(x)}, \text{ with } \theta(x) \in [0, \pi]\,. \tag{11.9}$$

[2]This does not prevent other similar W_6 operators from being generated by the UV physics (note that, beyond Eq. 11.3, operators of the form $(L_i \cdot L_j e_k)(L_l \cdot L_m e_n)$ or $(u_i \times d_j \cdot d_k)(u_l \times d_m \cdot d_n)$ or $(L_i \cdot L_j e_k)(u_l \times d_m \cdot d_n)$ can also arise). Rather, it corresponds to setting initial field conditions near the considered flat directions at the onset of inflation.

Note that $\phi(x)$ is defined by $\phi(x) \propto \pm|\varphi(x)|$ which is why $\theta(x)$ takes values between 0 and π only. From Eqs. (11.3), (11.4), (11.9) and the supersymmetric contribution $\left|\frac{\partial W_6}{\partial \varphi}\right|^2$ to the potential, one finds that the coupling λ_6 appearing in the last term on the right-hand side of Eq. 11.13 below, is given by

$$\lambda_6 = \frac{\widetilde{\widetilde{\lambda_6}}}{4\sqrt{2}} = \frac{\widetilde{\lambda_6}}{18\sqrt{2}} \lesssim \mathcal{O}\left(\frac{1}{18\sqrt{2}}\right). \tag{11.10}$$

To summarise: on the one hand, λ_6 differs by a factor $4\sqrt{2}$ from the normalisation found in the literature [270, 271, 273–275] where $\widetilde{\widetilde{\lambda_6}}$ was identified with λ_6. On the other hand, it also differs by a factor $18\sqrt{2}$ from $\widetilde{\lambda_6}$ which satisfies $\widetilde{\lambda_6} \lesssim \mathcal{O}(1)$ and which was also identified with λ_6 in the literature. This latter difference should be considered when assessing the consistency of the magnitude of λ_6. The former difference could have been just an unphysical redefinition of the coupling, were it not for the presence of the extra term in the potential generated by SUSY breaking that does not scale similarly to the supersymmetric term, as we now discuss.

For simplicity, we take here, as in [271, 273–275], the example of mSUGRA mediation of SUSY-breaking that we have introduced in Sec. 10.2.3.2, and in particular the soft-SUSY-breaking terms corresponding to W_6. In mSUGRA (minimal Kähler potential, see *e.g.* [263, 306] for reviews), the soft-SUSY-breaking part of the potential at the SUSY breaking scale $M_{\rm SUSY}$ has the universal form

$$V_{\rm soft} = \left|m_{3/2}\right|^2 \sum_i |\varphi_i|^2 + \left\{ m_{3/2} \left[\sum_i \varphi_i \frac{\partial W}{\partial \varphi_i} + (a-3)W \right] + h.c. \right\} \tag{11.11}$$

where the φ_i , W are the complex-valued scalar superfields and the superpotential of the MSSM visible sector (both introduced in Sec. 10.2.2), $m_{3/2}$ the gravitino mass introduced in Sec. 10.2.3, and where a a parameter related to the model-dependent mechanism of SUSY breaking. The two latter parameters are fixed by the vacuum expectation values of scalar fields and superpotential in the hidden sector and are in general complex-valued. The first term in Eq. 11.11 corresponds to the soft-SUSY-breaking masses of all the scalar fields in the observable sector taking the universal value $|m_{3/2}|$ at the scale $M_{\rm SUSY}$. As for the second term, it is easy to see that any monomial W_n contributing to W and containing a product of n φ_i's, satisfies the identity $\sum_i \varphi_i \frac{\partial W_n}{\partial \varphi_i} = nW_n$. The overall contribution of a given W_n is thus $\widetilde{A}_n W_n + h.c.$ where $\widetilde{A}_n = m_{3/2}(a+n-3)$ defines the corresponding soft-SUSY-breaking scalar coupling, the A-term, which is in general complex-valued. In the eMSSM, the relevant terms in W are W_3 (comprising the full MSSM renormalisable superpotential) and W_6 given by Eq. 11.1 and 11.2. The expression for $\widetilde{A}_n$ noted above then immediately leads to a relation between the trilinear soft-SUSY-breaking coupling $\widetilde{A}_3$ and the bi-trilinear coupling $\widetilde{A}_6$ at the SUSY breaking scale:

$$\widetilde{A}_3 = a m_{3/2}, \ \widetilde{A}_6 = (3+a) m_{3/2}, \ \text{whence } \widetilde{A}_6 = \frac{3+a}{a} \widetilde{A}_3. \tag{11.12}$$

$\widetilde{A}_3$ stands for the universal value of all the soft-SUSY-breaking trilinear scalar couplings of the MSSM at the $M_{\rm SUSY}$ scale. Eq. 11.12 is a key relation. It links, at the SUSY-breaking scale, $\widetilde{A}_6$ to a universal value of the soft trilinear couplings in the squark/slepton sectors of the MSSM. It will therefore allow us, together with Eqs. 11.5, 11.6, to relate the particle-physics features of the MSSM to the inflation features of the eMSSM. To obtain the full potential, $V_{\rm soft}$ should be

added to the supersymmetric F-term, see Eq. 10.21, and D-term contributions, see Eq. 10.22. Along the *LLe* or *udd* flat directions, $V_{\rm D}$ vanishes, as well as the W_3 contribution to $V_{\rm F}$ and to $V_{\rm soft}$. Adding the F-term contribution of Eq. 11.8 to its contribution to $V_{\rm soft}$, one finally finds the potential,

$$V(\phi) = \frac{1}{2}m_\phi^2\phi(x)^2 + \sqrt{2}\left|\widetilde{A}_6\right|\cos\left[6\theta(x)+\theta_6\right]\frac{\lambda_6\phi(x)^6}{6M_{\rm Pl}^3} + \lambda_6^2\frac{\phi(x)^{10}}{M_{\rm Pl}^6}, \tag{11.13}$$

where we wrote $\widetilde{A}_6 = |\widetilde{A}_6|e^{i\theta_6}$. To reach this form of the potential, we used Eq. 11.9 which leads consistently to $m_\phi^2 = m_\varphi^2$, but requires reabsorbing some normalisation factors in the redefinition of the coupling, $\widetilde{\lambda}_6 = 4\sqrt{2}\lambda_6$, as was already anticipated in Eq. 11.10. We also took into account a factor 2 coming from the Hermitian conjugacy in Eq. 11.11. The outcome differs from the literature by the relative factor $\sqrt{2}$ in the A-term and a factor $18\sqrt{2}$ between λ_6 and the effective coupling in W_6.

Back to Eq. 11.12, we note that, in practice, the universal value $\widetilde{A}_3$ of the trilinear couplings in the various squark/slepton sectors is lost at lower scales, as these couplings run differently with the RGEs. Since the MSSM spectrum is mostly sensitive to the third-generation quark/squark sector, we will take $\widetilde{A}_3$ to be $A_{\rm t}$, the trilinear coupling involving the top squark at $M_{\rm SUSY}$. Furthermore, we will follow [271, 273–275] by making use of the knowledge of the a-parameter in the simple Polonyi model for the hidden sector, see Eq. 10.31 from which one sees that $a = 3-\sqrt{3}$ [263, 291]. Hence, from Eq. 11.12, one obtains the relation

$$\widetilde{A}_6(M_{\rm SUSY}) = \frac{6-\sqrt{3}}{3-\sqrt{3}}A_{\rm t}(M_{\rm SUSY}) \tag{11.14}$$

between $\widetilde{A}_6$ and $A_{\rm t}$, where we now indicate explicitly the scale dependence.[3] Note, however, that Eq. 11.14 is found to be the reverse of the relation given in [271, 273–275]. Eq. 11.14 will be assumed in the rest of the analysis, where we will also equate for simplicity the two scales $M_{\rm SUSY}$ and $M_{\rm GUT}$. Since in our analysis all parameters, including $A_{\rm t}$, are taken real-valued, cf. Sec. 11.2.1, Eq. 11.14 implies real-valued positive or negative $\widetilde{A}_6$, and thus $\theta_6 = 0$ or π. The actual values of θ_6 are, however, irrelevant due to the preferred alignment of the $\theta(x)$ field during inflation, as we now explain.

Let us note that, since φ is a complex field, Eq. 11.13 is *a priori* a two-field model, for ϕ and θ. The angular direction θ is however heavy and can therefore be integrated out. Indeed, in the θ direction, the potential is minimal when $\cos(6\theta+\theta_6) = -1$. Around this minimum, when θ is properly normalised so to have a canonical kinetic term, its mass is of order $HM_{\rm Pl}/\phi$, hence it is much larger than H as long as the inflaton takes on sub-Planckian field values, which is always the case hereafter. This implies that θ decays to the configuration $\cos(6\theta+\theta_6) = -1$ in a small fraction of an e-fold, and that its excitations are sufficiently suppressed to yield small amounts of isocurvature modes and of non-Gaussianities (note that, in principle, one would have to check that all other orthogonal directions are stabilised too [307]). We can therefore consider the one-field potential:

[3] In subsequent sections where the framework of our analysis is described, we depart from the strict assumption of high-scale universality of soft-SUSY-breaking parameters. We will however stick to Eq. 11.14 as an illustrative example of possible correlations.

$$V_{\rm tree}(\phi) \equiv V(\phi) = \frac{1}{2} m_\phi^2 \phi^2 - \sqrt{2} A_6 \frac{\lambda_6 \phi^6}{6 M_{\rm Pl}^3} + \lambda_6^2 \frac{\phi^{10}}{M_{\rm Pl}^6}, \tag{11.15}$$

where, from now on, we denote $|\widetilde{A}_6|$ by A_6 (> 0) in accordance with the common notation used in the literature.[4]

Eq. 11.15, together with the definition range of the parameters,

$$m_\phi^2 > 0, \tag{11.16}$$

$$A_6 > 0, \tag{11.17}$$

$$0 < \lambda_6 \lesssim \mathcal{O}\left(\frac{1}{18\sqrt{2}}\right), \tag{11.18}$$

will be our starting point for the analysis of inflation.

11.1.2. Renormalisation-group-improved potential

Beyond the tree-level approximation, the *effective* potential has a more involved dependence on ϕ than that given by Eq. 11.15. It is important to carefully examine this aspect as it can lead to sizeable modifications of the inflationary predictions of the model. As well known, a powerful approach to capture classes of loop contributions to the effective potential is the requirement of Renormalisation Group Invariance of the full effective potential [308–311]. This allows a resummation to all orders in perturbation theory of powers of logarithms appearing in the loops. Obtaining the effective potential including the leading-order logarithms boils down to replacing all parameters in the tree-level potential by their one-loop running counterparts. In our case, the inflationary potential becomes

$$V_{\rm RGE}(\phi) \equiv V(\phi) = \frac{1}{2} m_\phi^2(\phi) \phi^2 - \sqrt{2} A_6(\phi) \frac{\lambda_6(\phi) \phi^6}{6 M_{\rm Pl}^3} + \lambda_6(\phi)^2 \frac{\phi^{10}}{M_{\rm Pl}^6}, \tag{11.19}$$

where the running quantities $m_\phi(\phi)$, $A_6(\phi)$ and $\lambda_6(\phi)$ are governed by the RGEs given below. This has two potential impacts. First, although this introduces essentially a mild (logarithmic) dependence, it can lead to substantial modifications of inflationary predictions, which is reminiscent of the so-called η problem [312, 313]. Second, this allows one to relate the physics at the scale of inflation to the one observed by HEP experiments. We will discuss both aspects in the rest of the chapter.

The operators LLe and udd labelling the considered flat directions are R-parity violating. One can thus use the general results of [314] to extract the one-loop Renormalisation Group Equations governing $A_6(\phi)$ and $\lambda_6(\phi)$, including multiplicative factors for the non-renormalisable operators [315], and neglecting contributions suppressed by $M_{\rm Pl}$.

In the following, we use the conventions of `SuSpect3` [288, 289] and the notations summarised in Sec. 11.2.1. In particular, our sign convention for the gaugino masses is that of [316], opposite to the one adopted in [271, 273–275]. We denote the energy scale by Q. Since, as stated

[4] A word of caution: This notation should not lead to believe that a negative sign is not allowed for $A_{\rm t}$ in view of Eq. 11.14; indeed the latter reads now $A_6(M_{\rm SUSY}) = \frac{6-\sqrt{3}}{3-\sqrt{3}} |A_{\rm t}(M_{\rm SUSY})|$; $A_{\rm t}$ can have either sign, which is phenomenologically important when identifying parameter-space regions compatible with the observed Higgs mass, but its sign does not affect the inflationary potential.

previously, Eq. 11.14 is defined at the GUT scale, we use the RGEs with the GUT scale as boundary and adopt the SU(5)-GUT normalisation $g_1 = \sqrt{5/3} g_Y$ (where g_Y denotes the SM hypercharge gauge coupling). Within the above-mentioned assumptions, the RGEs of the potential parameters read:

For $L^i L^j e^k$:

$$Q\frac{\mathrm{d}m_\phi^2}{\mathrm{d}Q} = -\frac{1}{6\pi^2}\Big(\frac{9}{10}M_1^2 g_1^2 + \frac{3}{2}M_2^2 g_2^2 + \mathrm{Y}_{m_\phi}^{L^i L^j e^k}\Big), \tag{11.20}$$

$$Q\frac{\mathrm{d}\widetilde{A}_6}{\mathrm{d}Q} = \frac{1}{2\pi^2}\left(\frac{9}{10}M_1 g_1^2 + \frac{3}{2}M_2 g_2^2 + \mathrm{Y}_{A_6}^{L^i L^j e^k}\right), \tag{11.21}$$

$$Q\frac{\mathrm{d}\lambda_6}{\mathrm{d}Q} = -\frac{\lambda_6}{4\pi^2}\Big(\frac{9}{10}g_1^2 + \frac{3}{2}g_2^2 + \mathrm{Y}_{\lambda_6}^{L^i L^j e^k}\Big), \tag{11.22}$$

For $u^i d^j d^k$:

$$Q\frac{\mathrm{d}m_\phi^2}{\mathrm{d}Q} = -\frac{1}{6\pi^2}\Big(\frac{2}{5}M_1^2 g_1^2 + 4M_3^2 g_3^2 + \mathrm{Y}_{m_\phi}^{u^i d^j d^k}\Big), \tag{11.23}$$

$$Q\frac{\mathrm{d}\widetilde{A}_6}{\mathrm{d}Q} = \frac{1}{2\pi^2}\left(\frac{2}{5}M_1 g_1^2 + 4M_3 g_3^2 + \mathrm{Y}_{A_6}^{u^i d^j d^k}\right), \tag{11.24}$$

$$Q\frac{\mathrm{d}\lambda_6}{\mathrm{d}Q} = -\frac{\lambda_6}{4\pi^2}\Big(\frac{2}{5}g_1^2 + 4g_3^2 + \mathrm{Y}_{\lambda_6}^{u^i d^j d^k}\Big), \tag{11.25}$$

with $A_6(Q) = |\widetilde{A}_6(Q)|$. It is worth noting that the $U(1)$ gauge contributions that depend on a universal combination of *all* soft-SUSY-breaking squared scalar masses, present in the β-function of the RGEs of each of these masses individually (see *e.g.* [316]), cancel out exactly in Eqs. 11.20 and 11.23 as a consequence of Eqs. 11.6 and 11.5. It follows that the running of m_ϕ^2 is generally given by Eqs. 11.20 and 11.23 (even without a universality assumption of the soft-SUSY-breaking scalar masses at some scale, in which case the aforementioned combination would have vanished at all scales). Moreover, the one-loop runnings of the gauge couplings g_i and gaugino masses M_i for $i = 1, 2, 3$ are obtained from the following RGEs (see *e.g.* [317]):

$$Q\frac{\mathrm{d}\ln M_i}{\mathrm{d}Q} = Q\frac{\mathrm{d}\ln g_i^2}{\mathrm{d}Q} = \frac{b_i}{8\pi^2}g_i^2 \quad \text{with } b_1 = \frac{33}{5},\ b_2 = 1 \text{ and } b_3 = -3. \tag{11.26}$$

The Yukawa terms contributing to Eqs. 11.20, 11.21, 11.22, 11.23, 11.24 and 11.25 depend on the inflaton type. They are mainly functions of the renormalisable Yukawa and trilinear soft-SUSY-breaking couplings, the dependence of Y_{A_6} and Y_{λ_6} on λ_6 and A_6 themselves being negligible due to Planck-mass suppression. They are given for completeness in App. A of [93]. Eq. 11.20, 11.21, 11.22, 11.23, 11.24 and 11.25 should thus be coupled to the RGEs governing the runnings of the Yukawa and trilinear soft-SUSY-breaking couplings on top of Eq. 11.26 and, in general, cannot be solved analytically. Numerical solutions for the full set of RGEs can be obtained from `SuSpect3` [288]. Analytical solutions become available when the Yukawa terms can be neglected, which is the case for the LLe and udd flat directions involving first and/or second lepton and quark generations. For simplicity, this is the case studied in this work. One can find the RGEs' solutions within this approximation [271] in App. A of [93].[5] When the

[5]Note that the energy scale Q should be replaced by ϕ when the runnings are used in V_{RGE}. Indeed

radiative corrections are included, one needs to ensure that Eq. 11.16 is verified at all scales of the theory, and that Eq. 11.18 is validated at high scale (cf. Sec. 11.5.2.3).

Although we rely in the sequel on Eq. 11.19 to evaluate the effects of the loop corrections on inflation, it is useful to keep in mind some possible caveats. In principle, in the monomials appearing in Eq. 11.19, one should replace the ϕ field itself by its running counterpart taking into account the corresponding anomalous dimension. Furthermore, a constant can be in general added to the effective potential. This was considered by [318] in the context of the present model. However, when the potential is RGEs improved in the presence of such a constant, the latter does not remain constant and induces an extra field dependence [310, 311]. Last but not least, the non-renormalisable operators in Eq. 11.15, when inserted in loops, induce in the one-loop Coleman-Weinberg potential other effective operators of the form ϕ^n, with $n = 4, 8,$ ... that were absent at tree-level. Some of these operators would cancel out in the large ϕ limit, due to supersymmetry. Others will have to be renormalised away through counterterms at some boundary scale, say $M_{\rm GUT}$, but will be regenerated at lower scales.

11.1.3. Slow roll in the eMSSM

With a fully defined potential like $V_{\rm tree}$ or $V_{\rm RGE}$, we can further discuss slow-roll conditions and the calculation of ϕ_* in order to determine the predicted values of the inflation observables. To ensure that slow-roll inflation can occur, the potential needs to exhibit an inflection point at some field value ϕ_0 defined by

$$V_{\phi\phi}(\phi_0) = 0 \quad \text{and} \quad \nu \equiv V_\phi(\phi_0). \tag{11.27}$$

The slow-roll parameters (see Eqs. (3.33), (3.34) and (3.35)) should be small enough around $\phi = \phi_0$. In particular, the condition $\varepsilon_1 < 1$ reads

$$|\nu| < \frac{\sqrt{2}}{M_{\rm Pl}} |V(\phi_0)|. \tag{11.28}$$

We will refer to ϕ_0 as "quasi-flat inflection point" whenever it satisfies Eqs. 11.27 and 11.28. For $V_{\rm tree}$, this condition delineates a region of size $|\phi - \phi_0|/M_{\rm Pl} \sim (\phi_0/M_{\rm Pl})^{3/2}$, while a more stringent requirement comes from the condition $|\varepsilon_2| < 1$, which holds in the range

$$|\phi - \phi_0| \simeq \frac{\phi_0^3}{60 M_{\rm Pl}^2} \ . \tag{11.29}$$

Note that a similar condition is obtained from $|\varepsilon_3| < 1$, and we have numerically verified that this relation still holds for $V_{\rm RGE}$. In the range defined by Eq. 11.29, Eq. 10.9 can be used and the inflection point is very close to the value of the field at Hubble crossing, ϕ_*.

In order to determine ϕ_*, one needs to solve Eq. 10.9, of which two important ingredients are $\ln R_{\rm rad}$ and $\tilde{\rho}_\gamma$. The former describes how the reheating occurs, and the latter is a function of the variation in the number of relativistic degrees of freedom between the reheating epoch and today, which depends on the field content of the theory. An interesting feature of eMSSM inflation is that, in principle, both are entirely determined by the MSSM spectrum.

the same RGEs play a double role: They allow one to improve the effective potential and to improve the energy dependence of physical scattering processes. In the latter case Q stands for the typical energy of the process.

Based on the reheating continuity relations Eqs. 10.5, 10.6, 10.7, one can relate the reheating to the slow-roll potential V and some coupling of the inflaton to radiation species. Hence, $\ln R_{\rm rad}$ can *a priori* be derived from Γ_ϕ. We illustrate in App. E the exact numerical solution for $\ln R_{\rm rad}$ assuming a particular eMSSM slow-potential and an arbitrary Γ_ϕ of 10^7. The calculation of Γ_ϕ could also be *a priori* performed, but they are however complex and go beyond the scope of the present work. The analyses of [319, 320] suggest that $0 < \ln R_{\rm rad} \ll 1$ (*i.e.* quasi-instantaneous reheating) both for udd and LLe, so, hereafter, we will assume $\ln R_{\rm rad} = 0$ unless specified otherwise.

The effective number of relativistic degrees of freedom at reheating is 427/4 in the Standard Model of particle physics if reheating occurs above ~ 100 GeV, while it is expected to be 915/4 in the MSSM [321]. Instead, in the following, we do not include such details and we assume that the ratio of the relativistic degrees of freedom between the reheating epoch and today is equal to one. It only leads to differences in the value of ΔN_* inferred from Eq. 10.9 of the order of 0.4 e-folds. This error can be neglected compared to the other sources of uncertainties discussed below.

The impact of assuming different values of ΔN_* (and different $\ln R_{\rm rad}$) is further discussed in Secs. 11.4.1 and 11.4.3.

11.2. Analysis framework

In this part, we introduce the phenomenological MSSM (see Sec. 10.2.3) benchmark points that we will use, the observational constraints and the tools and methodology used to perform the analysis.

11.2.1. pMSSM benchmark points

The pMSSM parameters are listed in Table 10.3. We additionally rely on the assumption of Eq. 10.33. In the following, gauge couplings are unified at the GUT scale ($M_{\rm GUT}$). Unless otherwise indicated, it is set to 3×10^{16} GeV in the following. In Sec. 11.5, it is directly calculated for a given MSSM spectrum within `SuSpect3` following [322]. The EWSB scale is computed thanks to Eq. 10.28. In what follows, we will assume that the Lightest Supersymmetric Particle is the neutralino χ_1^0 and it acts as the dark-matter candidate.

As shown in Sec. 11.1.2, the runnings of the parameters of $V_{\rm RGE}$ depend on the inflaton type, but also on the values of the gauge couplings, the masses of the gauginos, the Yukawa couplings and the trilinear soft-SUSY-breaking scalar couplings (the latter two will be neglected in this work). For this reason, we define in this section two examples that will be used for illustration in the following sections. They are hereafter called "benchmark points", and their characteristics at the GUT scale are given in Table 11.1. The chosen values are representative of the MSSM

	$g_1^{\rm GUT}$	$g_2^{\rm GUT}$	$g_3^{\rm GUT}$	$M_1^{\rm GUT}$ (GeV)	$M_2^{\rm GUT}$ (GeV)	$M_3^{\rm GUT}$ (GeV)
BP1	0.70	0.69	0.68	136	1143	899
BP2	0.70	0.69	0.68	898	1790	883

Table 11.1: Benchmark points. The gauge couplings and gauginos masses are given at $M_{\rm GUT}$.

points discussed in Sec. 11.5: BP1 (*resp.* BP2) corresponds to h_1 (*resp.* A_1), run to GUT scale. We only specify the quantities entering the potential parameters RGEs at this stage. Note that to remain as general as possible we clearly depart here from a universality assumption for the gaugino soft-SUSY-breaking masses.

11.2.2. Observational constraints

11.2.2.1 Cosmological observables

The main constraints from cosmological observations that are used in this part are the cold-dark-matter energy density ($\Omega_{\rm cdm}h^2$), the amplitude ($A_{\rm s}$) and spectral index ($n_{\rm s}$) of the primordial scalar perturbations. For $\Omega_{\rm cdm}h^2$, we are using [323]

$$\Omega_{\rm cdm}h^2 = 0.1187 \pm 0.0017\,. \tag{11.30}$$

On top of the experimental uncertainty quoted above, we consider a theoretical uncertainty of 0.012 associated with the prediction of $\Omega_{\rm cdm}h^2$ in the MSSM [324]. For $A_{\rm s}$ and $n_{\rm s}$, we are using the measurements inferred from the combination of *Planck* temperature, polarisation, lensing, and BAO data assuming a Λ-CDM model [148]. The corresponding measurements of the parameters inferred for $k_* = 0.05$ Mpc^{-1} are given in Table 11.2. In the following, when we refer to the measurements instead of the parameters, we make use of the notations $\overline{n_{\rm s}} \pm \sigma_{n_{\rm s}}$, and $\overline{A_{\rm s}} \pm \sigma_{A_{\rm s}}$.

Parameter	Value and error
$\ln(10^{10}A_{\rm s})$	3.047 ± 0.014
$n_{\rm s}$	0.9665 ± 0.0038
$\alpha_{\rm s}$	-0.0042 ± 0.0067
r	< 0.032

Table 11.2: Measurements of the amplitude ($A_{\rm s}$) and spectral index ($n_{\rm s}$) of the primordial scalar spectrum [64] used in this analysis. For reference, we also give the running of the scalar index ($\alpha_{\rm s}$) [64] and 95% C.L. upper limit on the tensor-to-scalar ratio r [61].

As shown in [92], the tensor-only contribution to the tensor-to-scalar ratio, r, for this potential is beyond the reach of present and future experimental constraints (and below the threshold of secondary gravitational waves induced by scalar fluctuations through gravitational non-linearities (see *e.g.* [325]). We therefore consider a (not r-extended) Λ-CDM model for the $n_{\rm s}$ and $A_{\rm s}$ constraints and we do not make use of the current experimental constraint on the tensor-to-scalar ratio. We will discuss the consistency between the $\alpha_{\rm s}$ predictions and the measured value in Sec. 11.5.

Another experimental constraint is the one on $\rho_\gamma = \Omega_\gamma \rho_{\rm cri}$ with Ω_γ the density parameter of radiation today and $\rho_{\rm cri}$ the critical density which leads to $\rho_\gamma = (4.645 \pm 0.004) \times 10^{-34}$ [326]. This observable enters in Eq. 10.9. The propagation of the experimental error on this quantity leads to a negligible contribution to the error on ΔN_* of $(1/4)\ln(1 + 0.004/4.645) = 2 \times 10^{-4}$ (similar to $\ln R_{\rm rad}$ and the number of relativistic degrees of freedom discussed in the previous section, see Eq. 10.9). The value of ρ_γ together with the measured values of $A_{\rm s}$ and $n_{\rm s}$ are hereafter called "inflationary observables".

11.2.2.2 Particle-physics observables

On the particle-physics side, the values of the main measurements considered in this analysis are summarised in Table 11.3 with their statistical and systematic errors. Where relevant, the theoretical errors associated with the supersymmetric predictions are also indicated (with the subscript "th").

Measurement	Value and error
m_h [GeV]	$125.10 \pm 0.14 \pm 3.00_{\rm th}$
$\mathrm{BR}(B_S \to \mu^+\mu^-)$	$(30 \pm 4 \pm 2_{\rm th}) \times 10^{-10}$
$\mathrm{BR}(b \to s\gamma)$	$(33.2 \pm 1.5) \times 10^{-5}$
Δa_μ	$(26.1 \pm 7.9 \pm 2_{\rm th}) \times 10^{-10}$
$m_{\rm t}$ [GeV]	172.76 ± 0.30

Table 11.3: Main particle-physics measurements used in the analysis [326]. The last number is the theoretical uncertainty on the supersymmetric prediction, except for $\mathrm{BR}(b \to s\gamma)$ and $m_{\rm t}$ for which no such error is considered.

To complete the Higgs-sector constraints, on top of the mass quoted in Table 11.3, we include the Higgs couplings, which are taken from [326–335]. Also, we use the Large Electron-Positron collider limit on the mass of the first generation chargino: $m(\tilde{\chi}_1^+) > 103.5\,\mathrm{GeV}$ [336]. Finally, on the DM searches' side, we also consider the limits on the direct detection rate, which are provided by the XENON1T experiment, with a rate above 0.4 keV to be less than 1 event/(tonne $\times$ day $\times$ $\mathrm{keV_{ee}}$) [337].

The measurements of the W and Z masses, the Higgs boson width together with the forward-backward asymmetries, the left-right asymmetries, the effective weak mixing angle, and the hadronic corrections to the QED coupling are also included.

As far as the MSSM parameter space is concerned (without considering the need to embed inflation), the modeling of the parameter space comes essentially from the intertwined constraints of the Higgs mass and the cold-dark-matter energy density [324, 338].

11.2.3. Tools and methodology

Several tools have been used to perform the analysis described in this part:

- `ASPIC`[6][92]: even though not directly interfaced to our framework, we have adapted `ASPIC` to derive reheating-consistent observable predictions.
- `SuSpect3` [288]: it calculates the MSSM physical masses and couplings, taking into account the dominant radiative corrections, the requirement of EWSB, and the running of the eMSSM parameters through their RGEs, relating the low-energy physics to several high-energy model assumptions. In particular, on top of the Z-mass and the EWSB scales, it allows one to define boundary conditions for the relevant running parameters at up to three different physical scales that can be chosen to be the GUT, the SUSY-breaking and the inflection point scales. The version of `SuSpect3` used in this analysis is the `3.1.1`.
- `SFitter` [339]: it allows one to confront the experimental data with predictions determined from the spectrum calculated by `SuSpect3`. The statistical analysis is performed on the

[6] http://cp3.irmp.ucl.ac.be/~ringeval/aspic.html

basis of a global χ^2 calculated from the individual $\Delta\chi^2$'s of the measurements versus predictions for the observables described in the previous section. We then make use of `MINUIT` [340], included in `SFitter`, to infer the values of the underlying parameters. This is done in particular through interfacing with `Micromegas` [341] for the predictions of the dark-matter energy density and the rare B-decays branching ratios. `SusyPope` [342] and `HDecay` [343] are used to calculate the predictions of the Z-pole observables and the Higgs couplings, respectively. In `SFitter`, the statistical errors on the measurements are treated as Gaussian or Poisson where appropriate. The systematic errors are correlated if originating from the same source. Theoretical uncertainties are treated using flat errors.

The methodology we adopt to infer the parameters of an inflationary potential given some cosmological observations is inspired by `ASPIC`. Starting from an MSSM potential, $V_{\rm tree}$ (Eq. 11.15) or $V_{\rm RGE}$ (Eq. 11.19), we have identified the area of the parameter space that satisfies the conditions of a quasi-flat inflection point (defined in Sec. 11.1.3), then we solve numerically Eq. 3.42 and Eq. 3.44, setting $A_{\rm s}$ and $n_{\rm s}$ to their measured values at $\phi = \phi_*$ which, itself, is determined through Eqs. 10.8 and 10.9 (see blue box below for a concrete example of this last step).

For more concreteness, in the figure below, we illustrate how the cosmological observables are obtained from a tree-level potential with some parameters that yield $\phi_0 = 3 \times 10^{14}$ GeV precisely. In orange, we show the "late" ΔN calculated from the k_* definition; see Eq. 10.9. In blue, we show the "slow-roll" ΔN, derived from Eq. 3.32. ϕ_0 is represented by a vertical black dotted line. The intersection of the two curves allows us to numerically compute ϕ_*, the field value at pivot-scale Hubble-crossing, represented by a vertical grey dashed line. One observes that ϕ_* is relatively very close to ϕ_0. However, this small shift is responsible for the observable prediction. Indeed, with ϕ_*, one can compute V_*, ε_{1*} (middle panel) and ε_{2*} (bottom panel) which allows us to deduce $A_{\rm s}$, $n_{\rm s}$, r (*etc...*).

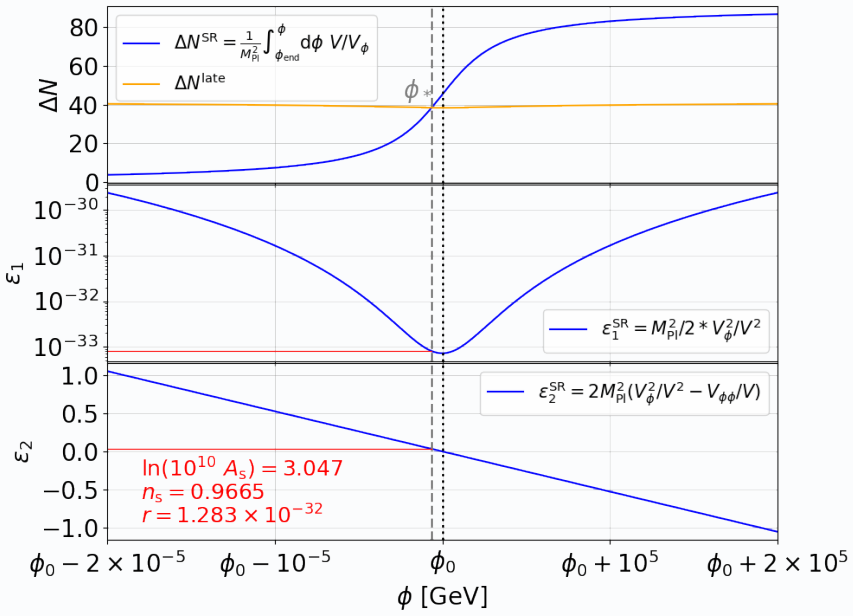

In practice, since the potential has three main free parameters (m_ϕ, A_6, λ_6), one needs to fix a degree of freedom, eg: A_6 (*resp.* ϕ_0), on top of the two constraints from n_s and A_s, in order to be able to determine the two other ones. For reasons that will be further detailed in Sec. 11.3.3, the two degrees of freedom we choose to tune are ν and ϕ_0 (*resp.* A_6). To be able to do that, we need the expressions for the potential and its derivatives taking ϕ_0 and ν directly as inputs: their full expressions are given in Sec. 11.3.3. Choosing to fix A_6 or ϕ_0 obviously leads to the same results; this choice is made according to the situation.

Moreover, in the specific case of the one-loop V_RGE, one also needs to specify the gaugino masses and gauge couplings (cf. Eq. 11.20 to Eq. 11.25). We make use of `SuSpect3` for the calculations of the RGEs for the inflationary parameters, as well as for the calculation of the physical spectra given a point in the MSSM parameter space. The correspondence between the squark or slepton soft masses and the inflaton mass (Eq. 10.33) as well as the relation between $A_\mathrm{t}$and A_6 (Eq. 11.14) are added at this stage. Finally, `SFitter` is used to perform the χ^2 calculation and its minimisation by comparing the HEP predictions and actual measurements (cf. Table 11.3).

11.3. Slow-roll conditions, initial conditions, fine-tuning

In this section, after an illustration of the shape of the potential, we address two essential questions: on the one hand, the field phase-space initial conditions required for inflation to take place, and on the other hand, the conditions the inflationary potential parameters must satisfy to yield predictions in agreement with observations. This will lead us to an assessment of the level of fine-tuning involved in this model.

11.3.1. Shape of the potential

To begin with, we discuss and illustrate the shapes of the V_tree and V_RGE potentials. For V_RGE, we take the example of a *LLe* inflaton for BP1 (Table 11.1). The parameter sets (p_tree and p_{LLe}^BP1) have been respectively determined for V_tree and V_RGE following the methodology described in Sec. 11.2.3, so that both potentials exhibit a common inflection point at $\phi_0 \simeq 1.2 \times 10^{15}$ GeV and match the inflationary observables (Table 11.2). They are shown in blue and red in Fig. 11.1. We do not propagate here the experimental errors on the observables and assume an instantaneous reheating for the purpose of illustration.

Around this inflection point, the slow-roll approximated first Hubble-flow parameter, as given by Eq. 3.33, is plotted in the upper-right panel of the same figure. Note that this approximation applies only when $\varepsilon_1 \ll 1$ and that, in practice, ε_1 can never be larger than 3 — thus values of ε_1 of order one or larger in Fig. 11.1 simply signal a break down of the slow-roll approximation. Slow roll takes place when $\varepsilon_1 < 1$ (orange region). A zoom of the potential is also given in the lower-right panel in the same field range (the red curve is above the Y-axis range of the plot).

One sees that the inflection points are very close to flat for $V_\mathrm{tree}\left(p_\mathrm{tree}\right)$ and $V_\mathrm{RGE}\left(p_{LLe}^\mathrm{BP1}\right)$, with ε_1 diving well below one in the slow-roll region (whose narrowness has been quantified in Eq. 11.29).

Applying the parameters obtained with V_tree (i.e. p_tree) to V_RGE leads to a tilted inflection point shown in green in Fig. 11.1. In this case, $V_\mathrm{RGE}\left(p_\mathrm{tree}\right)$ does not satisfy the slow-roll conditions: the corresponding values of the parameters are thus ruled out.

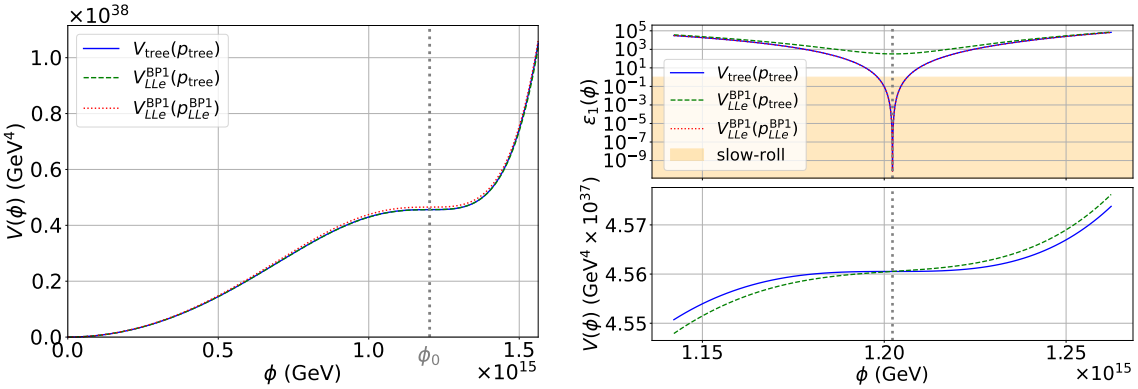

Figure 11.1: LEFT PANEL: Example of the shape of the inflationary potentials as a function of the field value for $\phi_0 = 1.2 \times 10^{15}$ GeV. $V_{\rm tree}$ is shown in blue and LLe BP1 $V_{\rm RGE}$ in green (dashed line) for the parameters determined to match the inflationary observables assuming $V_{\rm tree}$. The red-dotted line illustrates the LLe BP1 $V_{\rm RGE}$ for the parameters determined by taking into account the one-loop RGEs' corrections. RIGHT PANEL: The associated ε_1 function (defined by Eq. 3.33, more details in the text) is represented on the upper-right panel (the slow-roll region, where $\varepsilon_1 < 1$, is identified in orange). Zooms of the potentials around ϕ_0 are given on the lower-right panel.

This illustrates that the parameters determined for $V_{\rm tree}$ cannot be simply applied to a one-loop corrected $V_{\rm RGE}$ potential, suggesting that the one-loop correction of the inflationary potential is a feature that cannot be ignored (contrary to what is often done in the literature). Instead, one needs to re-determine the new set of parameters specifically for $V_{\rm RGE}$. Proceeding this way, we recover a shape similar to the $V_{\rm tree}$ one, with the same order of magnitude for ε_1 at the inflection point as shown with $V_{\rm RGE}\left(p_{LLe}^{\rm BP1}\right)$ in red.

11.3.2. Field phase-space

As mentioned in Sec. 3.2.2, slow roll is a dynamical attractor. This implies that, in a given inflationary potential, a successful phase of inflation takes place starting from a large set of initial conditions $(\phi, \dot{\phi})$, all attracted towards the same slow-roll solution [344, 345]. However, it is also known that the size of the basin of attraction depends on the shape of the inflationary potential (for instance, it is larger for plateau and large-field models than for hilltop potentials [345]). In order to determine to which extent eMSSM inflation is robust under changing the field initial conditions, let us thus study its field phase-space structure.

Both $V_{\rm RGE}$ and $V_{\rm tree}$ have a slow-roll region whose (narrow) extent is given by Eq. 11.29. We therefore expect the conclusions of this section to be identical when one uses either potential. Hence, we consider for explicitness $V_{\rm tree}$, where we arbitrarily set $\phi_0 = 0.395 M_{\rm Pl}$, and apply the methodology defined in Sec. 11.2.3 to determine the potential parameters such that the predicted values of $n_{\rm s}$ and $A_{\rm s}$ match the measurements given in Table 11.2. This leads to: $m_\phi = 4.29 \times 10^9$ GeV, $A_6 = 5.42 \times 10^{10}$ GeV and $\lambda_6 = 2.29 \times 10^{-8}$.

In the left panel of Fig. 11.2, the phase-space trajectory obtained by numerically integrating Eq. 3.22 and Eq. 3.23 is shown for three different initial conditions. In particular, the light-green one is starting near the inflection point, in the slow-roll region. In this figure, the blue area indicates the region where inflation takes place (*i.e.* where Eq. 3.24 is satisfied). The red-dotted

lines correspond to the slow-roll trajectory (3.31), where inflation proceeds at small velocity close to the inflection point, *i.e.* around ($\phi = \phi_0,\ \dot{\phi} \simeq 0$). If initial conditions are set close enough to the slow-roll trajectory, it acts as an attractor, as expected. This is the case for the light-green trajectory.

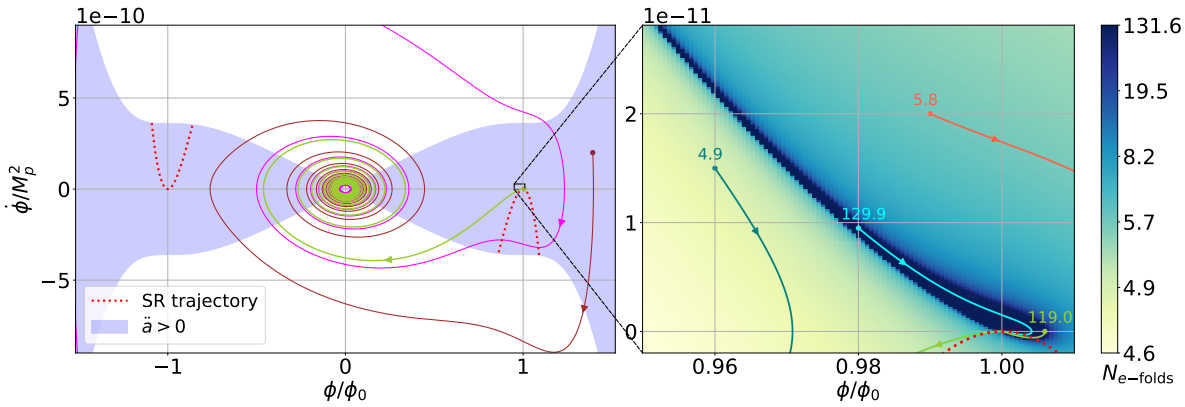

Figure 11.2: Phase-space diagram of the eMSSM inflationary potential V_{tree}, with $m_\phi = 4.29 \times 10^9$ GeV, $A_6 = 5.42 \times 10^{10}$ GeV and $\lambda_6 = 2.29 \times 10^{-8}$ (such that the model is compatible with cosmological observations). LEFT PANEL: A few trajectories (brown, magenta and light-green curves) are compared to the slow-roll attractor Eq. 3.31 (red-dotted lines). The blue shaded region indicates where inflation takes place, *i.e.* where the condition of Eq. 3.24 is satisfied. RIGHT PANEL: Zoom-in on the low-velocity inflection-point region (black square on the left). The color encodes the total number of inflationary e-folds generated by the trajectory starting at the given point of the phase space. A few examples of initial conditions (the light-green trajectory is the same on both panels) with their associated e-folds are also displayed.

Otherwise, the inflection point is overshot (this is the case for the magenta and brown trajectories), and the field quickly oscillates around the minimum of its potential at $\phi = 0$. Let us note that, during this oscillating phase, the system repeatedly crosses the inflating region, but it does so across very short periods of time and when averaging over several oscillations it does not inflate (but rather behaves as pressure-less matter for a quadratic minimum, see Eq. 10.4, [246], and Fig. 5.1 in App.E). Therefore, the relevant phase of inflation to consider is the one taking place before the oscillations. Unless initial conditions are chosen close to the slow-roll attractor, that phase is very short-lived, as illustrated in the right panel of Fig. 11.2. In this panel, some trajectories and their generated total number of e-folds $N_{e\text{-folds}}$ (given by Eq. 3.7, evaluated between $t = 0$ and the time of the end of inflation) are represented. One can see that only the ones being attracted by the slow-roll trajectory generate enough e-folds (*i.e.* originates in the dark-blue region). From Eq. 10.9, inflation must at least generate $N_{e\text{-folds}} \simeq 50$, and one can see that this requires to fine-tune the initial conditions close to the slow-roll attractor. The basin of attraction is therefore very narrow in this model, which constitutes a first level of fine-tuning.

Let us note that, at large-field values, $\phi \gg \phi_0$, the potential is of the "large-field type", $V_{\text{tree}} \propto \phi^{10}$, hence the slow-roll attractor is very powerful in that region [345] (*i.e.* its basin of attraction is very large). If initial conditions are set in that region, it uniquely determines the trajectory along which the inflection point is approached at lower field values. This trajectory is the magenta one in the left panel of Fig. 11.2, which starts at $\phi = 100 M_{\text{Pl}}$, oscillates a few

times around the potential minimum, before entering the range of the plot and overshooting the slow-roll attractor around the inflection point.

Therefore, another mechanism must be invoked to set the initial conditions close to slow roll at the inflection point, possibly involving additional dynamical fields [346]. Hereafter, we will assume that such a mechanism takes place and we will restrict the analysis to the slow-roll attractor.

Finally, let us point out that the reason why the slow-roll attractors at large-field values and around the inflection point are disconnected is because, when ϕ_0 is sub-Planckian, the slow-roll conditions are violated between these two regions. This implies that the slow-roll attractor is broken, and we found that the problem becomes worse when decreasing ϕ_0. Here, for illustrative convenience, we have set ϕ_0 to a mildly sub-Planckian value, but as will be made explicit below, ϕ_0 is usually expected to be much lower. This implies that, in practice, the fine-tuning problem of the field initial conditions is even worse than what can be seen in Fig. 11.2. To this respect, eMSSM inflation (and inflection-point models in general) behaves like small-field hilltop models [345].

11.3.3. Fine-tuning of the inflationary potential parameters

A peculiarity of the eMSSM potential that has been pointed out, in particular in [92, 270, 271, 274, 275], is that there is a high level of fine-tuning of the parameters when one imposes the condition of a quasi-flat inflection point. We discuss this point in the case of $V_{\rm tree}$ and $V_{\rm RGE}$ in this section.

11.3.3.1 At tree level

To ensure that the potential function remains monotonic with a quasi-flat inflection point for $V_{\rm tree}$, the parameters have to fulfill the following very restrictive relation between the bi-trilinear coupling and the combination of soft-SUSY-breaking scalar masses appearing in m_ϕ [92, 270, 271, 273–275]:

$$0 < 1 - \frac{A_6^2}{20{m_\phi}^2} \ll 1 \ , \tag{11.31}$$

where the additional factor $\sqrt{2}$, cf. Eq. 11.15, is included, as compared to the literature.

For later use, we define[7]

$$\alpha \equiv 1 - \frac{A_6^2}{20{m_\phi}^2}. \tag{11.32}$$

We recall here the origin of the fine-tuned requirement, Eq. 11.31, at tree-level. One can determine the value of ϕ_0 that satisfies the inflection-point requirement, the first equation in Eq. 11.27. Indeed, $V_{{\rm tree},\phi\phi}$ being quadratic in ϕ^4 irrespective of the magnitude of α, there is a two-branch solution for ϕ_0^4,

$$\phi_{0\pm}^4 = \frac{m_\phi M_{\rm Pl}^3}{9\lambda_6\sqrt{10}}\left(5\sqrt{1-\alpha} \pm \sqrt{16-25\alpha}\right). \tag{11.33}$$

When $\alpha < \frac{16}{25}$, the two branches are *a priori* acceptable since $\phi_{0\pm}^4$ remain real-valued and positive. One also finds that $V_{{\rm tree},\phi}(\phi_{0+})$ vanishes for $\alpha = 0$ and $V_{{\rm tree},\phi}(\phi_{0-})$ vanishes asymptotically for α

[7] Beware the different conventions between [274, 275] and [92], and the one adopted in this work in the definition of α. The relations between the different conventions are as follows: $\alpha_{[275]} = \sqrt{\alpha}/2$ and $\alpha_{[92]} = 1 - \alpha$.

very large and negative, so that the second equation in Eq. 11.27 can be satisfied with vanishingly small ν. However the "-" branch cannot satisfy Eq. 11.28 and will not be further considered. As for the "+" branch, ε_1, cf. Eq. 3.33, is found to scale as $\lambda_6^{\frac{1}{2}} M_{\rm Pl}^{\frac{1}{2}} m_\phi^{-\frac{1}{2}}$ and vanishes for $\alpha = 0$, whence the required high degree of fine-tuning on α around zero in order to keep $\varepsilon_1 < 1$.

It is thus justified to rely on the lowest order expansion in α for the field value and the potential function and its first derivative at the inflection point, $\phi_0 = \phi_{0+}$, that can be derived analytically from Eq. 11.33:

$$\phi_0^4 = \frac{m_\phi M_{\rm Pl}^3}{\lambda_6 \sqrt{10}} + \mathcal{O}(\alpha)\,, \tag{11.34}$$

$$V_{\rm tree}(\phi_0) = \frac{4}{15} m_\phi^2 \phi_0^2 + \mathcal{O}(\alpha)\,, \tag{11.35}$$

and

$$V_{{\rm tree},\phi}(\phi_0) = m_\phi^2 \phi_0 \alpha + \mathcal{O}\left(\alpha^2\right). \tag{11.36}$$

As clear from the above equation, relaxing the monotonicity assumption boils down to allowing for negative α (since $V(\phi)$ increases monotonically for $0 < \phi \ll \phi_0$ and $\phi \gg \phi_0$). Such potentials predict values for $n_{\rm s}$ that are always incompatible with measurements (see [92]), and they will not be further investigated here.

Restricting to $\alpha > 0$, Eq. 3.33, Eq. 11.35 and Eq. 11.36 imply a tight relation between α and the first slow-roll parameter at ϕ_0:

$$\alpha = \frac{4\sqrt{2}}{15} \frac{\phi_0}{M_{\rm Pl}} \sqrt{\varepsilon_1(\phi_0)} + \mathcal{O}\left[\frac{\phi_0^2}{M_{\rm Pl}^2} \varepsilon_1(\phi_0)\right]. \tag{11.37}$$

The required level of fine-tuning of the parameter α is illustrated by this equation as slow-roll conditions impose $\varepsilon_1 \ll 1$. More precisely, $\varepsilon_1 < 1$ implies that $\alpha < 4 \times 10^{-3}$ for $\phi_0 = M_{\rm GUT}$, and this required level of precision increases when ϕ_0 decreases. It gets even more stringent when one requires that cosmological observables are correctly reproduced (cf. Table 11.2). This is further illustrated in Sec. 11.3.3.2.

In practice, the fine-tuning of α requires a high level of (quadratic) precision in the numerical determination of the parameter space if we solve Eq. 10.9 directly for $V_{\rm tree}$ (see [92]). However, this would lead to computational accuracy mismatch when interfacing with lower (double-)precision codes such as `SuSpect3` and `SFitter`. To circumvent this problem, and also set the stage for the generalisation to the one-loop effective potential in the next subsection, we give the exact solution of Eq. 11.27:

$$m_\phi^2 = \frac{1}{40}\left(45\frac{\nu}{\phi_0} + A_6^2 + A_6\sqrt{A_6^2 - 10\frac{\nu}{\phi_0}}\right), \tag{11.38}$$

$$\lambda_6 = \frac{M_{\rm Pl}^3}{20\sqrt{2}\phi_0^4}\left(A_6 + \sqrt{A_6^2 - 10\frac{\nu}{\phi_0}}\right). \tag{11.39}$$

This form is different, though equivalent, to the one discussed above. To obtain these equations we used the fact that both $V_{{\rm tree},\phi\phi}(\phi)$ and $\frac{1}{\phi}V_{{\rm tree},\phi}(\phi)$ are quadratic in $\lambda_6\phi^4/M_{\rm Pl}^3$ and retained the "+" branch solution as discussed previously. Using this form of the solution one can reexpress $V_\phi(\phi)$ and $V_{\phi\phi}(\phi)$ by expanding $\lambda_6\phi^4$ around $\lambda_6\phi_0^4$ so that large cancellations are already

effected analytically, thus bypassing the need for numerical quadratic precision. Here, we give the corresponding exact expressions:

$$V_{\mathrm{tree},\phi}(\phi) = (1+\Delta_4)^{\frac{1}{4}} \left\{ \nu + \Delta_4 \lambda_6 \frac{\phi_0^5}{M_{\mathrm{Pl}}^3} \left[10(2+\Delta_4)\lambda_6 \frac{\phi_0^4}{M_{\mathrm{Pl}}^3} - \sqrt{2} A_6 \right] \right\}, \tag{11.40}$$

$$V_{\mathrm{tree},\phi\phi}(\phi) = 5\Delta_4 \lambda_6 \frac{\phi_0^4}{M_{\mathrm{Pl}}^3} \left[18(2+\Delta_4)\lambda_6 \frac{\phi_0^4}{M_{\mathrm{Pl}}^3} - \sqrt{2} A_6 \right], \tag{11.41}$$

where $\Delta_4 \equiv \phi^4/\phi_0^4 - 1$ and ν (Eq. 11.27) can be taken as input. These expressions are at the core of the methodology described in Sec. 11.2.3: given a value for A_6, one tunes ϕ_0 and ν such that the predicted values for n_{s} and A_{s} match the observations, then one gets λ_6 and m_ϕ thanks to Eq. 11.39 and Eq. 11.38.

11.3.3.2 At one-loop level

We now consider the one-loop potential, V_{RGE}. To estimate the level of fine-tuning in this case, we first note that the successive derivatives of $V_{\mathrm{RGE}}(\phi)$ will have the same polynomial dependences on ϕ as those obtained when differentiating $V_{\mathrm{tree}}(\phi)$. This is a direct consequence of the form of the RGEs. The first and second derivatives of the potential (11.19) read

$$V_{\mathrm{RGE},\phi}(\phi) = \phi \times \left[m_\phi^2(\phi) + \frac{1}{2}\beta_m(\phi) - \xi_1(\phi)\frac{\lambda_6(\phi)\phi^4}{6M_{\mathrm{Pl}}^3} + \xi_2(\phi)\frac{\lambda_6(\phi)^2\phi^8}{6M_{\mathrm{Pl}}^6} \right] \tag{11.42}$$

$$V_{\mathrm{RGE},\phi\phi}(\phi) = m_\phi^2(\phi) + \frac{1}{2}\mathcal{B}_1(\phi) + \xi_3(\phi)\frac{\lambda_6(\phi)\phi^4}{6M_{\mathrm{Pl}}^3} - \xi_4(\phi)\frac{\lambda_6(\phi)^2\phi^8}{6M_{\mathrm{Pl}}^6}, \tag{11.43}$$

where $\mathcal{B}_1$ and the ξ_i's depend on the running of A_6, the gauge couplings and the gauginos masses, as well as the Yukawa and trilinear soft-SUSY-breaking scalar couplings. Their explicit expressions do not require the knowledge of the solutions of the RGEs. They are given in App. B of [93] in the approximation of negligible Yukawa terms.

At the inflection point introduced in Eq. 11.27 and taking into account the ϕ dependence in Eqs. (11.42) and (11.43), one finds:

$$m_\phi^2(\phi_0) = \frac{\mathcal{A}^2(\phi_0,\nu)}{20}, \tag{11.44}$$

$$\lambda_6(\phi_0) = 3\frac{M_{\mathrm{Pl}}^3}{\phi_0^4} \frac{\mathcal{B}_1(\phi_0)\xi_2(\phi_0) - \mathcal{B}_2(\phi_0)\xi_4(\phi_0) + 2m_\phi^2(\phi_0)\left[\xi_2(\phi_0)+\xi_4(\phi_0)\right]}{\xi_1(\phi_0)\xi_4(\phi_0) - \xi_2(\phi_0)\xi_3(\phi_0)}, \tag{11.45}$$

where $\mathcal{B}_2(\phi)$ is given in App. B of [93] for the LLe case, and $\mathcal{A}^2(\phi,\nu)$ is given in App. C of [93].

In view of Eq. 11.44, a straightforward generalisation at one-loop of the tree-level α parameter, see Eq. 11.32, is given by:

$$\alpha^{(\mathrm{loop})} \equiv 1 - \frac{\mathcal{A}^2(\phi_0,\nu=0)}{20 m_\phi^2(\phi_0)}. \tag{11.46}$$

Taking Eq. 11.44 into account, one can remark that $\alpha^{(\mathrm{loop})}$ vanishes in the flat-inflection limit, $\nu = 0$, and deviates from 0 for the quasi-flat-inflection, cosmology-consistent configurations. In

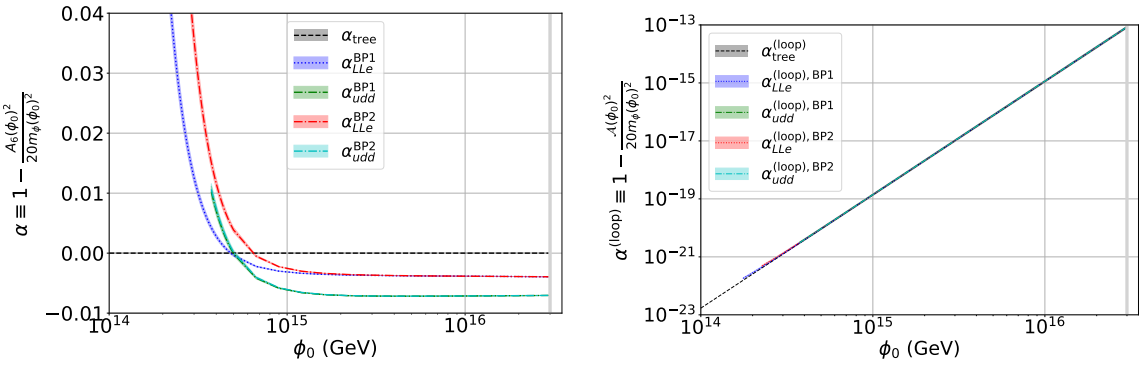

Figure 11.3: α (left panel) and $\alpha^{(\rm loop)}$ (right panel) as a function of ϕ_0, such that $V_{\rm tree}$ and $V_{\rm RGE}$ lead to predictions on $n_{\rm s}$ and $A_{\rm s}$ in agreement with measurements in the case of the two inflaton types (LLe and udd) and of the two benchmark points (BP1 or BP2). On the left panel, the green and cyan udd contours are almost superimposed and the black contour close to zero corresponds to $\alpha_{\rm tree}$. The vertical grey line indicates the GUT scale.

the limit of negligible one-loop corrections, Eqs. (11.44) and (11.45) give back the exact tree-level relations (Eq. 11.38 and Eq. 11.39) and $\alpha^{(\rm loop)} = \alpha$.

To illustrate the level of fine-tuning, for $V_{\rm tree}$ and $V_{\rm RGE}$, one can compute and compare α and $\alpha^{(\rm loop)}$ at ϕ_0 once we have applied the methodology described in Sec. 11.2.3. This is done in Fig. 11.3, where the tree level case is shown in black, and the LLe and udd cases for the two benchmark points are illustrated in colors[8]. The right panel of this figure provides evidence that the level of fine-tuning needed in $V_{\rm RGE}$ is of the same order of magnitude as the one in $V_{\rm tree}$ (contrary to what is suggested for instance in [274]). However, it cannot be estimated by simply replacing the tree-level quantities by their running counterparts in the definition of α (Eq. 11.32), as illustrated on the left panel. One needs to compute $\alpha^{(\rm loop)}$ instead.

Similarly to the $V_{\rm tree}$ case, we can circumvent the issue of the quadratic precision calculations required to solve Eqs. (10.8) and (10.9) for $V_{\rm RGE}$ directly by expanding the first and second derivatives of the potential around ϕ_0. In contrast to the tree-level case where exact forms can be obtained, Eqs. 11.40 and 11.41, here the non-polynomial dependence in ϕ is handled analytically through a Taylor expansion around $\phi = \phi_0$, where only the knowledge of the β-functions (and not the RGEs solutions) is needed. We have performed this calculation up to the tenth order in $\delta\phi = \phi - \phi_0$. This is absolutely required when one needs to make use of the double-precision calculations of `SuSpect3` to estimate the radiative corrections. We make use of this trick in Sec. 11.5.

11.4. Inflationary constraints in the parameter space

The aim of this section is to study the inflationary potential parameter space that satisfies the inflationary constraints. We first address this question using $V_{\rm tree}$ (cf. Sec. 11.4.1). We then compare the results to those obtained with $V_{\rm RGE}$ (Sec. 11.4.2).

[8] In the right panel, the different cases do not overlap exactly at low ϕ_0: this is linked to the fact that the inflaton mass needs to satisfy Eq. 11.16. This is further discussed in the next section.

11.4.1. Results for the potential at tree level

In this section, we consider the case of the tree-level inflationary potential, $V_{\rm tree}$.

11.4.1.1 Parameter-space constraints

Following the methodology defined in Sec. 11.2.3, one can determine m_ϕ, A_6 and λ_6 such that inflation takes place at ϕ_0, given the current measurements of the inflationary observables. The result depends on the assumptions made on reheating through Eq. 10.9. Fig. 11.4 shows the parameter-space projections as a function of ϕ_0, in log-scale. In red, we illustrate an instantaneous reheating, and in blue, we assume $\ln R_{\rm rad} = -10$ for illustration (see next section).

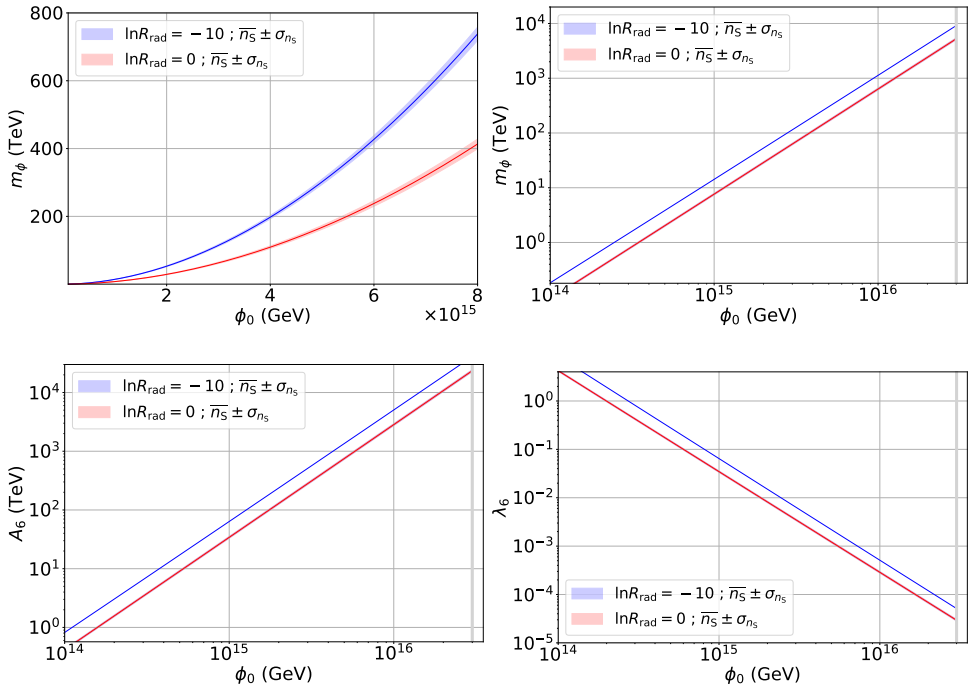

Figure 11.4: Tree-level inflationary potential parameters compatible with the inflationary observables as a function of ϕ_0: m_ϕ on the upper-left panel (low ϕ_0 in linear scale) and upper-right panel (wider range in log-scale), A_6 on the lower-left, and λ_6 the lower-right (log-scale). Two assumptions on $\ln R_{\rm rad}$ are illustrated: -10 in blue and 0 (instantaneous reheating) in red.

The figure of the upper-left panel is a zoom on the low-ϕ_0 region (in linear scale). At tree-level, the factor $\sqrt{2}$ in front of the ϕ^6 term that was introduced in Eq. 11.19 with respect to [275] may appear equivalent to a simple redefinition of the A_6 parameters. But this is not the case when one wants to relate inflation and HEP constraints and make use of Eq. 11.14 which connects A_6 to $A_{\rm t}$ at the GUT scale (this is further discussed in Sec. 11.5).

The width of the contour corresponds to the one-σ error on $n_{\rm s}$. The relative error due to the $A_{\rm s}$ uncertainty (cf. Table 11.3) is roughly of the order of seven per mille. It is therefore not propagated. The $\sigma_{n_{\rm s}}$ error translates into a σ_{m_ϕ}/m_ϕ of the order of 3% for a given ϕ_0. This compares to the approximately 15 % previously obtained in [275]. This difference is mainly due

to the reduction of the error bar on $n_{\rm s}$ between WMAP and *Planck* (the improvement factor is approximately 3.5 between the two measurements).

11.4.1.2 Number of e-folds and reheating

In order to obtain the contours shown in Fig. 11.4, one has to solve Eq. 10.9, which gives the number of e-folds ΔN_* between the time the pivot scale k_* crosses out the Hubble radius and the end of inflation (at next-to-leading order in slow-roll approximation). The larger ϕ_0, the larger ΔN_*: typically at the GUT scale for an instantaneous reheating, $\Delta N_* \simeq 45$, and for $\phi_0 \simeq 1\times10^{14}$ GeV, ΔN_* reaches 37.5. For slow-rolling potentials with $\phi_* \simeq \phi_0$, the ΔN_* expression (10.9) is often reduced to its slow-roll leading-order approximation[9],

$$\Delta N_* \simeq -\ln\left(\frac{k_*}{a_0\tilde{\rho}_\gamma^{1/4}}\right) - \frac{1}{4}\ln(9) + \frac{1}{4}\ln\frac{V_{\rm tree}(\phi_0)}{M_{\rm Pl}^4} \tag{11.47}$$

$$\simeq 61.2 + \frac{1}{4}\ln\frac{V_{\rm tree}(\phi_0)}{M_{\rm Pl}^4}\ , \tag{11.48}$$

for $k_* = 0.05$ Mpc^{-1}, ρ_γ being given in Sec. 11.2.2.1. The difference between both estimates for the number of e-folds is of the order of 0.1, showing that Eq. 11.47 is a very good approximation for this potential. We nevertheless keep solving Eqs. (10.8) and (10.9) explicitly in the following.

So far, we have not propagated any error on the prediction of ΔN_*. As already discussed in Secs. 11.2.2.1 and 11.1.3, such an error can be sourced by various effects: a deviation from an instantaneous reheating (*i.e.* $\ln R_{\rm rad} \neq 0$), a corrective term to account for the MSSM relativistic number of degrees of freedom at reheating, the experimental uncertainty of the ρ_γ measurement or the use of the slow-roll approximation. In Fig. 11.4, in blue, we illustrate the case where the contributions from these different terms end up shifting ΔN_* by -10 e-folds. It is a very pessimistic example since one expects the various contributing errors to be at most of order 1. Such a value of ΔN_* could, for example, arise from an extremely long reheating scenario with $\ln R_{\rm rad} = -10$. Such a change is equivalent, for the tree potential, to a shift of the $n_{\rm s}$ value from 0.9665 to 1.0353 (while keeping the instantaneous reheating assumption). Conversely, the current constraint on $n_{\rm s}$ would propagate into an error on $\ln R_{\rm rad}$ of 0.6 if all other parameters were to be fixed.

11.4.2. Results for the one-loop RGEs potential

In the two next sections, we compare the parameter space constraints assuming a tree-level potential to the ones obtained when taking into account the one-loop RGEs in the potential.

11.4.2.1 Comparisons for given ϕ_0

In this section, we determine the parameters such that, for both potentials $V_{\rm tree}$ and $V_{\rm RGE}$, inflation takes place at the same value for ϕ_0 and we first compare their values at this scale. We refer to each set of values as: $p^{V_{\rm tree}|\phi_0}$ (*resp.* $p^{V_{\rm RGE}|\phi_0}$) where p can be m_ϕ, A_6, λ_6. Note that $m_\phi^{V_{\rm tree}|\phi_0}$ corresponds to m_ϕ of the previous section (the same applies for the other parameters). Fig. 11.4 gives the absolute scale of these parameters at tree level.

[9]Beware the difference with [275] linked to the k_* values at which the inflationary observables are estimated.

To proceed, we introduce, for each parameter p, the following notations:

$$\Delta_i^{\mathrm{BP}j}[p] = p^{\mathrm{V_{tree}}|\phi_0}(Q=\phi_0, n_\mathrm{s}=\overline{n_\mathrm{s}}) - p^{\mathrm{V_{RGE}}|\phi_0}(Q=\phi_0, n_\mathrm{s}=\overline{n_\mathrm{s}}), \tag{11.49}$$

$$\sigma_{n_\mathrm{s},i}^{\mathrm{BP}j}[p] = \frac{1}{2}\left|p^{\mathrm{V_{RGE}}|\phi_0}(Q=\phi_0, n_\mathrm{s}=\overline{n_\mathrm{s}}+\sigma_{n_\mathrm{s}}) - p^{\mathrm{V_{RGE}}|\phi_0}(Q=\phi_0, n_\mathrm{s}=\overline{n_\mathrm{s}}-\sigma_{n_\mathrm{s}})\right|, \tag{11.50}$$

where i indicates the inflaton type (udd or LLe) and the subscript j refers to the benchmark point (as defined in Table 11.1). The Δ's give the biases that are related to the use of V_tree instead of the one-loop V_RGE on the values of the parameter p. The σ's indicate the statistical errors on the parameter p given the current 1σ error on n_s. In Fig. 11.5, these quantities are represented as a function of ϕ_0 for m_ϕ (upper-left panel), A_6 (upper-right panel) and λ_6 (bottom panel) for all benchmark points.

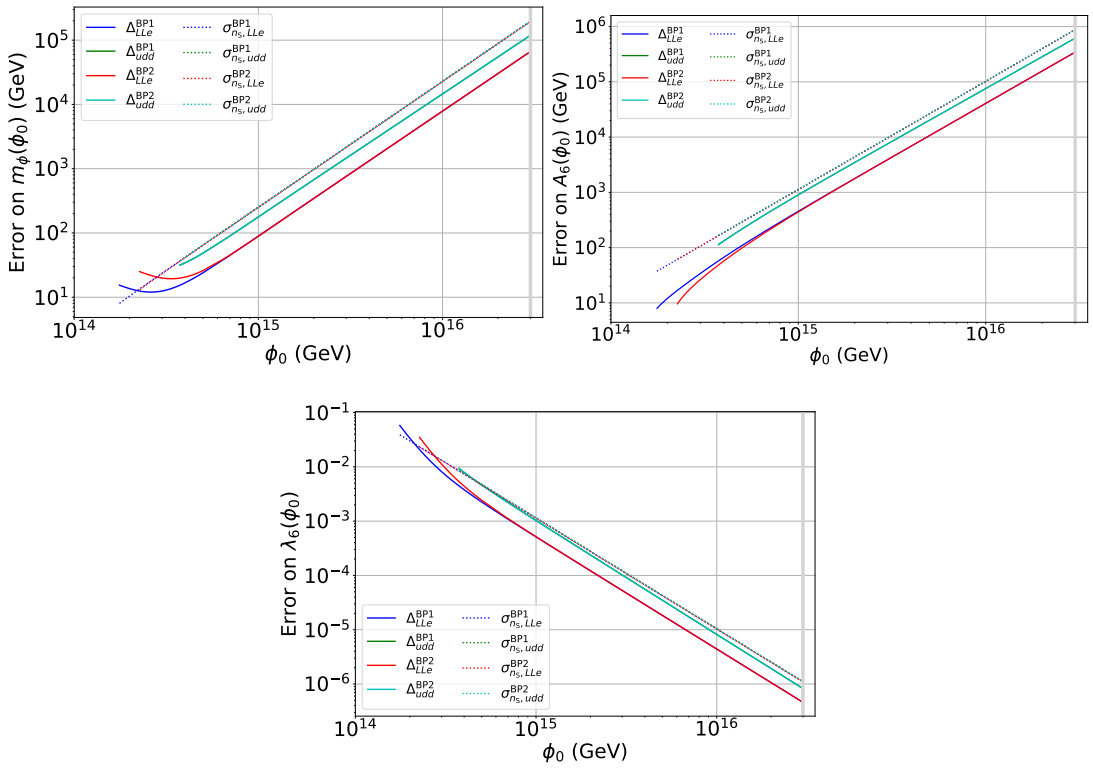

Figure 11.5: Systematic biases - Δ's as defined by Eq. 11.49 in solid lines - and statistical uncertainties - σ's as defined by Eq. 11.50 in dotted lines - on the determination of the inflationary potential parameters as a function of ϕ_0: in the upper-left for m_ϕ, in the upper-right for A_6, and in the lower-left for λ_6. BP1 (*resp.* BP2) is illustrated in dark blue (*resp.* red) for an LLe inflaton, and in green (*resp.* light blue) for udd.

The dotted lines in Fig. 11.5 show that at any inflation scale, for all parameters, and for both inflaton types and benchmark points, we recover the 3% relative statistical error that we have quantified for V_tree (cf Sec. 11.4.1).

In the high-ϕ_0 region, as illustrated by the comparison of the solid lines, the parameters are systematically shifted towards larger values when using V_RGE instead of V_tree. In these examples

where the Yukawa's terms can be neglected, this bias depends mainly on the inflaton type (the red and blue curves are perfectly superimposed). It is almost insensitive to the values of the gaugino masses and the gauge couplings assumed at the GUT scale. It is of the order of 253 GeV at $\phi_0 \simeq 1.2 \times 10^{15}$ GeV for $m_\phi(\phi_0)$ (upper-left panel) and $\simeq 1.3$ TeV for $A_6(\phi_0)$ (upper-right panel) for *udd*. This compares to the precisions on the determination of the parameters that are linked to the current uncertainty on $n_{\rm s}$ which are of the order of 360 GeV and 1.6 TeV respectively. The bias is not negligible: it is roughly 70% of the current statistical error due to $\sigma_{n_{\rm s}}$.

The low-ϕ_0 area exhibits two distinctive behaviors. On the one hand, in the RGEs case, there are unphysical regions where the inflaton masses become tachyonic (*i.e.* Eq. 11.16 is not satisfied). This implies that there exist lower bounds in the parameter space that depend on the benchmark point and on the inflaton type. This is an important difference to the tree-level potential for which the domain of definition of the parameters is unrestricted. In our examples, ϕ_0 cannot be lower than 2×10^{14} GeV to 5×10^{14} GeV depending on the cases. On the other hand, one can see a change of slope in the Δ's curves at low-ϕ_0. It is the same feature one could already see in Fig. 11.3. In this area, the systematic bias gets even larger than the statistical error and dominates the error budget. These two combined effects show how important it is to properly take into account the runnings of the parameters in the potential.

To assess the effect of these biases on the particle-physics phenomenology, the RGEs (Eq. 11.20 to Eq. 11.25) are used to calculate the values of these parameters at the EWSB scale taken to be 2 TeV as an illustration. For $V_{\rm tree}$, by definition, the parameters do not depend on ϕ in the potential, the RGEs are therefore taken into account only at this stage, as done in [275].

The induced difference between the parameters using $V_{\rm tree}$ and $V_{\rm RGE}$ is almost the same at 2 TeV as at ϕ_0. To give orders of magnitude, for the *LLe* BP1, while fixing $\phi_0 = 1.2 \times 10^{15}$ GeV, if one assumes $V_{\rm tree}$ and runs the parameters from ϕ_0 to 2 TeV, one obtains:

$$m_\phi^{\rm V_{tree}|\phi_0}(Q = 2{\rm TeV}) = 10892\ ^{358}_{-349}\ {\rm GeV}, \tag{11.51}$$

$$A_6^{\rm V_{tree}|\phi_0}(Q = 2{\rm TeV}) = 47681\ ^{1605}_{-1565}\ {\rm GeV}, \tag{11.52}$$

$$\lambda_6^{\rm V_{tree}|\phi_0}(Q = 2{\rm TeV}) = 0.045 \pm 0.001\ . \tag{11.53}$$

In addition, while using $V_{\rm RGE}$ and running the parameters to 2 TeV, one gets:

$$m_\phi^{\rm V_{RGE}|\phi_0}(Q = 2{\rm TeV}) = 11020\ ^{363}_{-354}\ {\rm GeV}, \tag{11.54}$$

$$A_6^{\rm V_{RGE}|\phi_0}(Q = 2{\rm TeV}) = 48336\ ^{1627}_{-1586}\ {\rm GeV}, \tag{11.55}$$

$$\lambda_6^{\rm V_{RGE}|\phi_0}(Q = 2{\rm TeV}) = 0.045 \pm 0.001. \tag{11.56}$$

In such an example, the predicted value of the inflaton running mass at the EWSB scale assuming $V_{\rm tree}$ instead of $V_{\rm RGE}$, is under-estimated by 128 GeV (more than one third of the statistical error bar linked to the propagation of the current value of $\sigma_{n_{\rm s}}$). For A_6, this systematic shift is of the order of 655 GeV, about 40% of the statistical error, hence it is not negligible.

The precise values of the induced systematic effect does depend, eventually, on the inflaton type, the gauginos masses and the gauge couplings. Still, as illustrated in the present example, using the simplified tree version of the inflationary potential may induce non-negligible bias in the end-results when one wants to combine constraints from cosmological and HEP observables:

a shift of more than 100 GeV on sparticle masses could lead to erroneous conclusions in the determination of the favored/disfavored area of the eMSSM parameter space if one wants to sample it extensively and couple it to particle-physics observables.

11.4.2.2 Comparisons for given $A_6(M_{\rm GUT})$

To illustrate the favored areas of the parameter space, one can also choose to compare their values assuming the same value of A_6 at the GUT scale, a particularly interesting quantity since it relates to $A_{\rm t}$ through the Polonyi relation, and, thus, to the particle-physics phenomenology (see next section). This implies that the comparison of the values of the $V_{\rm tree}$ and $V_{\rm RGE}$ parameters is performed at different values of ϕ_0.

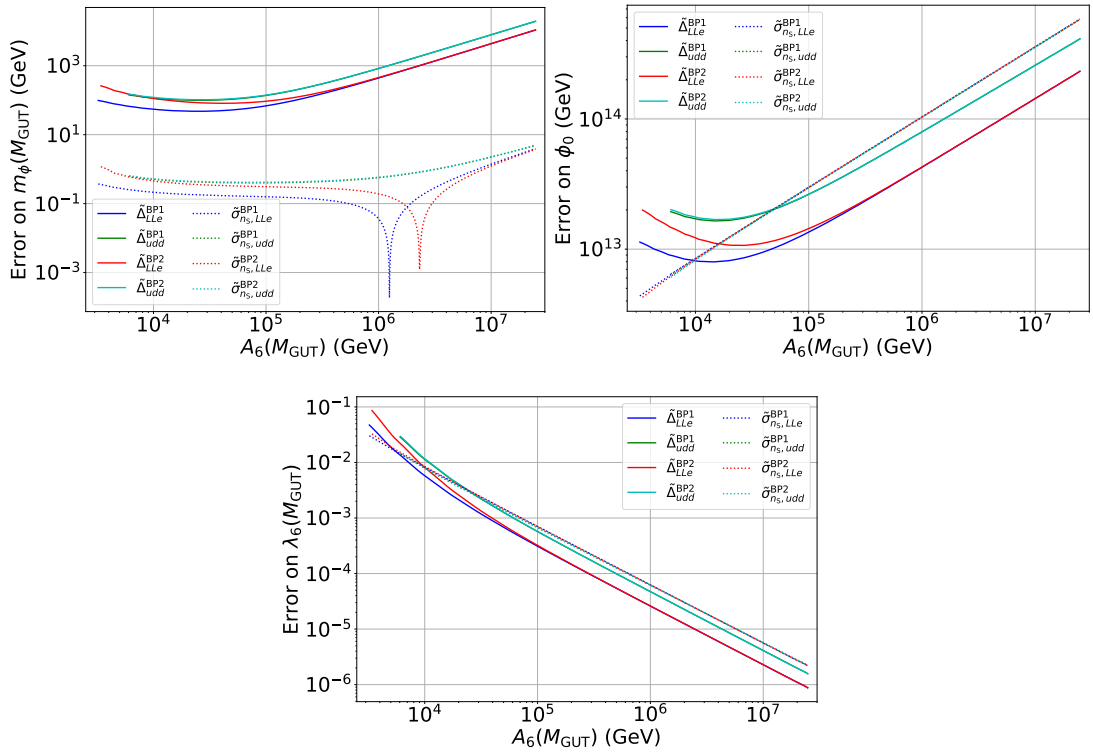

Figure 11.6: Systematic biases - $\tilde{\Delta}$'s as defined by Eq. 11.57 in solid lines - and statistical uncertainties - $\tilde{\sigma}$'s as defined by Eq. 11.58 in dotted lines - on the determination of the inflationary potential parameters as a function of $A_6(M_{\rm GUT})$: on the upper-left panel for $m_\phi(M_{\rm GUT})$, on the upper-right panel for ϕ_0, and on the lower panel for $\lambda_6(M_{\rm GUT})$. BP1 (*resp.* BP2) is illustrated in dark blue (*resp.* red) for an *LLe* inflaton, and in green (*resp.* light blue) for *udd*.

Similarly to what was done in Sec. 11.4.2, one defines:

$$\tilde{\Delta}_i^{{\rm BP}j}[p] = p^{{\rm V_{tree}}|_{A_6(M_{\rm GUT})}}(Q = M_{\rm GUT}, n_{\rm s} = \overline{n_{\rm s}}) - p^{{\rm V_{RGE}}|_{A_6(M_{\rm GUT})}}(Q = M_{\rm GUT}, n_{\rm s} = \overline{n_{\rm s}}), \quad (11.57)$$

$$\tilde{\sigma}_{n_{\rm s},i}^{{\rm BP}j}[p] = \frac{1}{2}\left| p^{{\rm V_{RGE}}|_{A_6(M_{\rm GUT})}}(Q = M_{\rm GUT}, n_{\rm s} = \overline{n_{\rm s}} + \sigma_{n_{\rm s}}) - p^{{\rm V_{RGE}}|_{A_6(M_{\rm GUT})}}(Q = M_{\rm GUT}, n_{\rm s} = \overline{n_{\rm s}} - \sigma_{n_{\rm s}}) \right|. \quad (11.58)$$

These errors are shown in Fig. 11.6 for m_ϕ, λ_6 and ϕ_0. When proceeding this way, because of the fine-tuning relation the propagation of the error on n_s on the m_ϕ axis is almost negligible (of the order of a few 0.1 GeV). The systematic bias when one compares the results using $V_{\rm RGE}$ versus $V_{\rm tree}$ for m_ϕ at the GUT scale is 250 times larger than this statistical error. This implies that tree-level and one-loop predicted inflaton masses are not compatible with each other. Such a high precision in the prediction of the inflaton mass comes with a worse determination of the ϕ_0 values.

For example, for $A_6(M_{\rm GUT}) = 49562$ GeV, one gets for $V_{\rm RGE}$:

$$m_\phi^{V_{\rm RGE}|A_6(M_{\rm GUT})}(Q = M_{\rm GUT}, n_s = \overline{n_s}) = 10999.4 \pm 0.2 \text{ GeV}, \tag{11.59}$$

$$\lambda_6^{V_{\rm RGE}|A_6(M_{\rm GUT})}(Q = M_{\rm GUT}, n_s = \overline{n_s}) = 0.022 \pm 0.001, \tag{11.60}$$

$$\phi_0^{V_{\rm RGE}|A_6(M_{\rm GUT})}(Q = M_{\rm GUT}, n_s = \overline{n_s}) = (1.20 \pm 0.02) \times 10^{15} \text{ GeV}, \tag{11.61}$$

which compare, for $V_{\rm tree}$, with:

$$m_\phi^{V_{\rm tree}|A_6(M_{\rm GUT})}(Q = M_{\rm GUT}, n_s = \overline{n_s}) = 11049.03 \pm 0.01 \text{ GeV}, \tag{11.62}$$

$$\lambda_6^{V_{\rm tree}|A_6(M_{\rm GUT})}(Q = M_{\rm GUT}, n_s = \overline{n_s}) = 0.021 \pm 0.001, \tag{11.63}$$

$$\phi_0^{V_{\rm tree}|A_6(M_{\rm GUT})}(Q = M_{\rm GUT}, n_s = \overline{n_s}) = (1.21 \pm 0.02) \times 10^{15} \text{ GeV}. \tag{11.64}$$

11.4.3. Error-budget summary

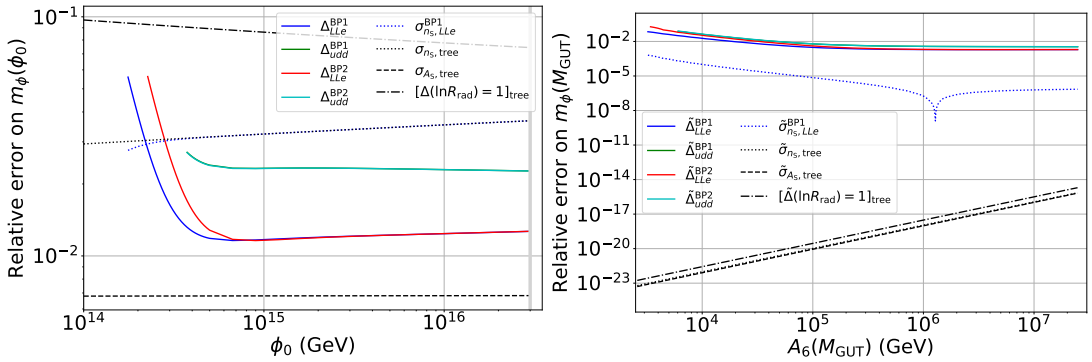

Figure 11.7: Summary of the different contributions to the relative errors on the predicted inflaton mass for different benchmark points and inflaton types. LEFT PANEL: for $Q = \phi_0$ as a function of ϕ_0. RIGHT PANEL: for $Q = M_{\rm GUT}$ as a function of $A_6(Q = M_{\rm GUT})$

.

A summary plot of the different contributions to the error budget on the determination of the inflaton mass is given in Fig. 11.7, which shows the relative errors on the determination of m_ϕ for $V_{\rm tree}$ and $V_{\rm RGE}$: at $Q = \phi_0$ assuming inflation occurs at ϕ_0 (left panel) and at $Q = M_{\rm GUT}$ assuming given values of $A_6(Q = M_{\rm GUT})$ (right panel). The solid lines show the systematic bias induced when adopting a tree-level approximation instead of taking into account the one-loop corrections in the potential for the four configurations already discussed in the previous sections.

The dotted lines represent the propagation of the error on n_s in the LLe BP1 case (in blue) and in the tree case (in black). Finally, the black-dash-dotted line represents the propagation of the A_s error, and the black-dotted line corresponds to the propagation of a theoretical error of one unit on $\ln R_\mathrm{rad}$.

This shows that the one-loop corrections to the inflationary potentials cannot be neglected. The tree-level fine-tuning that implies a tight relation between m_ϕ and A_6 at ϕ_0 is broken by the radiative corrections. It leads to a systematic underestimation of the parameters when one uses V_tree instead of V_RGE at high ϕ_0. In addition, because of the RGEs' runnings of the parameters, their definition domain is reduced, implying a minimal energy scale below which the parameters are not physical. Such a behavior is completely absent from the tree-level treatment of the potential. Moreover, for the same $A_6(Q = M_\mathrm{GUT})$ the predicted inflaton masses using V_tree or V_RGE are not compatible with each other anymore, the systematic error dominating the statistical error by several orders of magnitude. Finally, we also show the contribution of a theoretical error of $\ln R_\mathrm{rad}$ of one unit, which dominates the error budget if one wants to fix ϕ_0 in the analysis.

We have demonstrated that the measurements of the inflationary observables are, to date, accurate enough for this analysis to be sensitive to radiative corrections on the parameters of the inflationary potential. RGEs cannot be ignored anymore.

In addition, the theoretical error on the reheating phase is, today, a limiting factor in the constraints one can put on the eMSSM parameter space.

11.5. Combining HEP and cosmological constraints

In this section, we combine cosmological and HEP measurements using V_RGE and discuss some MSSM points in the context of the current data. We assume an instantaneous reheating (see Sec. 11.1.3), and we do not propagate any error on the prediction of ΔN_*. The goal of this section is to pin-point some specific examples of a combined fit but not to perform an exhaustive scan of the parameter space.

11.5.1. Combined fit

11.5.1.1 Method

We base our study on the best-fit MSSM points that have been identified in [324], which were found to match particle-physics observations and $\Omega_\mathrm{cdm}h^2$ measurements. We have updated the experimental constraints (cf. Table 11.3 and Table 11.2) and refitted the MSSM parameters for the corresponding predictions of the observables to be in agreement with the current measurements, using `SFitter`.

For each of these points, we proceed as follows. We set A_t at the GUT scale such that the Higgs sector is almost unchanged (while keeping a low χ^2). This allows us to fix $A_6(M_\mathrm{GUT})$ through Eq. 11.14. We are therefore in the case discussed in Sec. 11.4.2.2. We then determine the corresponding values of m_ϕ, λ_6 and ϕ_0 imposed by A_s and n_s. To do this, we bypass the quad-precision computation as explained at the end of Sec. 11.3.3.2, to compute the one-loop corrections with the double-precision code `Suspect3` and feed the `ASPIC`-like part of the code. Finally, we tune the eMSSM parameters for the soft-SUSY-breaking masses to match the inflaton mass (according to Eq. 11.5 or Eq. 11.6): either the mass of the second generation squarks (for a udd inflaton), or the mass of the sleptons (for LLe), given the relations defined by Eq. 11.5

and Eq. 11.6. To keep the Yukawa's terms low enough to be neglected, we only focus on the first two generations for the inflaton candidates: $u_1 d_1 d_2$ and $L_1 L_2 e_1$. Finally, we ensure the full consistency of the procedure in a global χ^2 minimisation using `SFitter`. In this way, the obtained corresponding MSSM spectra are compatible with all current observational constraints described in Sec. 11.2.2.

11.5.1.2 Results for different dark-matter annihilation channels

We have considered the MSSM spectra that correspond to three dark-matter annihilation channels: a Higgsino[10] channel, for example $\tilde{\chi}_1^0\tilde{\chi}_1^0 \to W^+W^-$ via a t-channel (hereafter points H_1 and H_2), a h-funnel channel $\tilde{\chi}_1^0\tilde{\chi}_1^0 \to h$ (points h_1 and h_2), and a A-funnel channel $\tilde{\chi}_1^0\tilde{\chi}_1^0 \to A$ (points A_1, A_2 and A_3).

The corresponding spectra are summarised in Table 11.4 for different inflaton hypotheses (udd and LLe). For each point, the first block provides the MSSM fundamental parameters at the EWSB scale. The second one gives the determined mass spectrum (the slepton physical masses are of the same order as the soft masses), together with the inflationary parameters, m_ϕ, A_6 and ϕ_0 determined with $V_{\rm RGE}$, whose values are given for $Q = \phi_0$. A χ^2 value is also given, it refers to the difference between predictions and all HEP observables described in Sec. 11.2.2.2 as well as $\Omega_{\rm cdm}h^2$. The contributions to the χ^2 are evenly distributed among the different observational constraints taken into account.

This table shows that we are able to combine all current observations to study, within the coherent framework of the eMSSM, the underlying parameter space and relate inflation to LHC physics at the level of the one-loop RGEs corrections on the inflationary potential. For the different DM annihilation channels we have considered, we have found eMSSM points that are compatible with the Higgs mass and the cold-dark-matter energy density but also with $n_{\rm s}$ and $A_{\rm s}$ (together with all the HEP observables detailed in Sec. 11.2.2). Apart from A_3, all the points have a similar value of $A_{\rm t}$, for this reason the inflaton masses are of the same order of magnitude ($\simeq 1 \times 10^4$ GeV), so are the values for ϕ_0 ($\simeq 1 \times 10^{15}$ GeV).

Using $V_{\rm tree}$, we would have obtained an inflaton mass (hence squark masses) roughly 30 GeV below the one obtained with $V_{\rm RGE}$ for the A_1 case. This difference is significantly larger than the statistical error linked to $\sigma_{n_{\rm s}}$, which is lower than one GeV. Even though these 30 GeV will highly depend on the inflaton type and on the values of the gauge couplings and gaugino masses, it is far from being negligible when considering HEP data. This further reinforces the fact that the measurements of inflationary observables are, to date, sufficiently accurate for inflationary potential analyses to be sensitive to radiative corrections on the parameters.

11.5.2. Additional cross-checks

11.5.2.1 LHC direct searches

As the gluino and the neutralino masses are, for some points, close to the limit of the current LHC searches, we have performed additional *a posteriori* cross-checks of these points, adding the constraints from direct searches for supersymmetric particles at the LHC.

For this, we use `SmodelS` [347–352] with cross-sections calculated using `Pythia6` [353] and `NLL-fast` [354–360] comparing to a database of Run 2 analyses [361]. It should be noted that

[10] The channel is called as such because the neutralino is mainly Higgsino in this example.

ID	H_1	H_2	h_1	h_2	A_1	A_2	A_3
DM channel	Higgsino		h-funnel		A-funnel		
inflaton	$u_1d_1d_2$	$L_1L_2e_1$	$u_1d_1d_2$	$L_1L_2e_1$	$u_1d_1d_2$	$L_1L_2e_1$	$L_1L_2e_1$
EWSB scale	2556.9	2556.8	2713.1	2713.0	4008.9	4008.9	4009.7
$\tan\beta$	29		26.6		18.3		24.7
sgn(μ)	+		+		-		
$M_{\rm GUT}$	1.237×10^{16}		3×10^{16}		1.295×10^{16}		
M_1	1571		61		400		362
M_2	2917		967		1515		1662
M_3	1931		1934		1898		1098
$M_{\tilde{\mu}_L}$	2864	12108	4210	9892	2564	10567	4773
$M_{\tilde{\mu}_R}$	2785	12108	4221	9892	3976	10400	4723
$M_{\tilde{\tau}_L}$	3342		4068		2039		
$M_{\tilde{\tau}_R}$	2064		3931		2801		
$M_{\tilde{q}_{2L}}$	11959	9090	9860	7370	10452	7299	2984
$M_{\tilde{q}_{3L}}$	2176		3155		3950		
$M_{\tilde{t}_R}$	3003		2330		4066		
$M_{\tilde{b}_R}$	3474		1952		2422		
A_τ	-3499		-2564		-3010		
$A_{\rm t}$	3180	3223	2985	3020	3450	3493	1190
A_b	148		143		187		
$m(\tilde{\chi}_1^0)$	1108	1108	60.0	60.1	397	398	357
$m(\tilde{\chi}_2^0)$	-1113	-1112	497	496	766	766	760
$m(\tilde{\chi}_3^0)$	1578	1587	-504	-503	-769	- 770	-763
$m(\tilde{\chi}_4^0)$	2992	3005	1041	1041	1589	1592	1710
$m(\tilde{\chi}_1^+)$	1111	1111	496	495	765	766	760
$m(\tilde{\chi}_2^+)$	2992	3005	1041	1041	1589	1592	1711
$m(\tilde{g})$	2371	2333	2351	2310	2382	2332	1389
$m(\tilde{q}_L)$	12072	9173	9956	7442	10557	7376	3036
$m(\tilde{q}_R)$	12056	9162	9939	7429	10540	7364	3024
$m(\tilde{b}_1)$	2243	2244	2031	2031	2499	2499	2453
$m(\tilde{b}_2)$	3512	3512	3177	3177	3974	3974	3971
$m(\tilde{t}_1)$	2244	2244	2347	2347	3946	3945	3972
$m(\tilde{t}_2)$	2992	2992	3188	3188	4074	4075	4059
$m(\tilde{\tau}_1^-)$	2066	2066	3931	3931	2055	2055	2055
m_h	125.3	125.2	125.3	125.2	125.4	125.3	122.2
m_A	784.9	783.9	3625.9	3625.8	782.2	784.4	757.2
$m_{3/2}$	12596	12557	10352	10383	11005	11010	4886
ϕ_0	1.25×10^{15}	1.26×10^{15}	1.13×10^{15}	1.14×10^{15}	1.16×10^{15}	1.17×10^{15}	7.68×10^{14}
$m_\phi(\phi_0)$	11982	11973	9847	9892	10459	10487	4661
$A_6(\phi_0)$	53757	53593	44181	44312	46970	46990	20855
$\lambda_6(\phi_0)$	0.0224	0.0218	0.0278	0.0269	0.0260	0.0252	0.0609
$\chi^2_{\rm HEP}$ (d.o.f. = 78)	50.9	51.7	46.1	46.5	47.5	47.0	49.2
LHC searches			✗	✗			✗

Table 11.4: Fitted benchmark MSSM spectra and their status regarding the observables. The parameters are given at $M_{\rm EWSB}$. All the masses and energy scales are given in GeV. The ✗ symbol highlights the points excluded by direct LHC searches (see 11.5.2.1).

the phenomenology of the points illustrated in Table 11.4 is far from the "simplified models" topologies used to express exclusions in LHC experiments. In particular, gluinos masses are close to the excluded regions for simplified models but in these points they rather tend to decay via sbottom or in long cascade decays. For most of the points, a definite conclusion would require

a complete event simulation and analysis. Nevertheless, the h-funnel points (h_1 and h_2) with their low-mass charginos and neutralinos are excluded with a small margin by searches for di-leptons and missing transverse momentum (MET). The A-funnel A_3 point with a low-mass gluino has a large enough cross section to be excluded despite the complex decay chain by multiple analyses for topologies like jets and MET, and leptons jets and MET.

These additional cross-checks on the h_1, h_2 and A_3 points are interesting examples of how particle-physics measurements can disfavor values of the inflationary potential parameters. It also shows how important it is to perform dedicated and detailed analysis of the particle-physics inputs from observations if one wants to perform a full analysis of the eMSSM (as was done for instance in [362] for the extensive study of the pMSSM).

11.5.2.2 Cosmology and LSP

We proceed here to two other cross-checks: one on the parameters related to cosmology, and the other on the nature of the LSP.

As discussed in Sec. 11.2.2.1, the predicted value of the tensor-to-scalar ratio r is too low to be meaningful, since it is expected to be lower than the threshold of secondary gravitational waves induced by scalar fluctuations through gravitational non-linearities. This remains true for $V_{\rm RGE}$, implying that r cannot be used in the global fit. We have also checked the predicted values of $\alpha_{\rm s}$. For the eMSSM points discussed in this section, we obtain $\alpha_{\rm s} \simeq -4.4 \times 10^{-3}$. This is perfectly consistent with current measurements (quoted in Table 11.2) and reinforces the fact that the eMSSM points are robust to the current observational constraint. As a side comment and to give orders of magnitude, the energy density at the pivot scale can also be derived thanks to the following relation: $\rho_* = 3M_{\rm Pl}^2 H_*^2 \overset{\text{SRLO}}{\simeq} 24\pi^2 M_{\rm Pl}^4 A_{\rm s} \varepsilon_{1*}$. We obtain $\rho_*^{1/4} \simeq 2 \times 10^9$ GeV (corresponding to $H_* \simeq 1$ GeV) for the points of Table 11.4.

The second a posteriori cross-check one can perform is linked to the nature of the LSP. We have made the assumption for our analysis that the LSP is the lightest neutralino, which is ensured by the global fit for all sparticles except for the gravitino, which is not part of the MSSM. Given the mSUGRA relation (Eq. 11.12), one can deduce $m_{3/2}$ from A_6: the corresponding values are given in Table 11.4, where we show that the gravitino mass is always larger than $m(\tilde{\chi}_1^0)$, ensuring the consistency of the analysis. Sufficiently fast gravitino decays will thus evade a potential problem of thermal gravitino overabundance (a detailed study taking into account the reheating and decay temperatures is however beyond the scope of the present part, see *e.g.* [363, 364]).

11.5.2.3 Consistency of the theory

Finally we have checked two theoretical assumptions: the bounds on the parameters and the choice of the UV scale.

In order to avoid tachyonic masses for the inflaton, we have checked that Eq. 11.16 is satisfied at all scales. One also needs to check that $\widetilde{\lambda_6}$ remains of order 1 between $M_{\rm Pl}$ and ϕ_0. Since the RGEs enhance the values of λ_6 at low scales, one only needs to check that $\widetilde{\lambda_6}(Q = \phi_0) = 18\sqrt{2}\lambda_6(Q = \phi_0) \lesssim 1$. As shown by the values of $\lambda_6(Q = \phi_0)$ of Table 11.4, this is verified for our points.

The last check concerns the choice of the UV completion scale in the definition of λ_6, cf.

Eqs. 11.3, 11.7, 11.10. It has to be noted that the $M_{\rm Pl}$ dependence in the potential only appears through $\lambda_6/M_{\rm Pl}^3$. The choice of a different UV scale (say $M_{\rm GUT}$), as mentioned in Sec. 11.1.1, would be equivalent to a rescaling of λ_6 by a factor $(M_{\rm GUT}/M_{\rm Pl})^3$. In addition, such a rescaling is valid at any energy scale, due to the fact that Eqs. 11.22, 11.25 are invariant under a re-scaling of λ_6. Assuming $M_{\rm UV} = M_{\rm GUT}$ instead of $M_{\rm Pl}$ preserves the consistency of the theory, since λ_6 (required to be $\lesssim \mathcal{O}\big(1/(18\sqrt{2})\big)$) now takes an additional 10^{-6} factor.

Since new physics is expected both at $M_{\rm Pl}$ and at $M_{\rm GUT}$, one could instead consider natural to have contributions from these two scales to the superpotential W_6 (Eqs. 11.3 and 11.7). In that case, one would need to apply the following transformation:

$$\frac{\widetilde{\lambda_6}}{M_{\rm Pl}^3} \to \frac{\widetilde{\lambda_{6,M_{\rm Pl}}}}{M_{\rm Pl}^3} + \frac{\widetilde{\lambda_{6,M_{\rm GUT}}}}{M_{\rm GUT}^3} \ , \tag{11.65}$$

where $\widetilde{\lambda_{6,M_{\rm Pl}}}$ and $\widetilde{\lambda_{6,M_{\rm GUT}}}$ are expected to be of order 1. Thus, unless $\widetilde{\lambda_{6,M_{\rm GUT}}}$ is unnaturally suppressed, the dominant contribution is the one from $M_{\rm GUT}$. One can then read off the resulting inflationary configuration from our study by replacing therein formally $\lambda_6(\phi)$ by $\big(M_{\rm Pl}^3/M_{\rm GUT}^3\big)\,\lambda_{6,M_{\rm GUT}}(\phi)$. Again, unless $\lambda_{6,M_{\rm GUT}}(\phi)$ is unnaturally suppressed, these significantly increased values of $\lambda_6(\phi)$ would come with significantly smaller values of ϕ_0, A_6 and m_ϕ, leading to too light sfermions and Higgs mass that are excluded by HEP constraints. Our conclusion is strengthened when $M_{\rm UV}$, the scale at which new physics arise, is smaller than $M_{\rm GUT}$, but still large enough to justify the use of effective operators. Hence, for the *udd* and *LLe* inflatons studied in this work, our results strongly suggest that potential new physics effects at the GUT (or any other sufficiently heavy UV) scale have to be (surprisingly) suppressed with respect to those originating from the Planck scale.

Conclusion

In this work, we have shown how a consistent analysis of cosmological and particle-physics constraints can be performed in the context of the eMSSM model.

We have studied the field phase-space structure and showed that eMSSM inflation behaves as a small-field model, with a narrow basin of attraction.

We have identified the region of parameter space that can support inflation when the one-loop RGEs corrections are included and compared it with the one obtained at tree level. We have developed a new way to estimate the level of fine-tuning, including one-loop RGE corrections. We have demonstrated that this level remains the same as the tree-level one. We have proposed a solution to overcome the resulting accuracy requirement. Furthermore, we have shown that the parameters of the one-loop inflationary potential are bounded at a low scale in order to avoid tachyonic inflaton masses.

We then detailed the area of the parameter space compatible with the $A_{\rm s}$ and $n_{\rm s}$ measurements when neglecting or not the one-loop corrections in the expression of the inflationary potential. We have shown that the small changes in the potential due to the RGEs induce a significant modification in the prediction for the model. While this shift depends on the masses of the gauginos and on the gauge couplings at the GUT scale, we have given examples for which the induced bias can be, for example, for the inflaton mass, almost comparable to the statistical error linked to the $n_{\rm s}$ measurements for a fixed value of ϕ_0. We also demonstrated that this effect

is even more important when one fixes A_6 at the GUT scale. This shows that, with the current constraints on n_s, one cannot neglect the one-loop corrections in the inflationary potential.

We have also compared these shifts to the ones induced by the uncertainty on the reheating duration and have found that the value predicted for n_s is very sensitive to the reheating details. More precisely, changing the reheating temperature by one order of magnitude is enough to shift the spectral index by more than its measurement error. On the one hand, this means that CMB measurements weakly constrain the MSSM parameters due to this large degeneracy with the reheating sector. On the other hand, this also implies that if those parameters were measured in particle-physics experiments, the CMB data would already be accurate enough to deliver a precise measurement of the reheating temperature. This contrasts with other single-field models of inflation where the reheating is still poorly constrained, even when the inflationary potential does not contain additional parameters [365, 366].

Finally, we have found points in the MSSM that are compatible with current measurements (in particular, the Higgs mass, the cold-dark-matter energy density, and the inflationary observables). We have shown how conclusions about their compatibility with inflationary constraints can be affected by the way we take into account RGEs at the inflation scale. In particular, beyond the Higgs funnel example, we have highlighted an A-funnel point compatible with most HEP and cosmological observations yet excluded by the LHC beyond Standard Model searches, which opens the door to constraining inflation using HEP measurements. We have also given other examples of Higgsinos and A-funnel points that are at the limit of the current LHC constraints, for which the SUSY phenomenology shows many cascade decays. For all these reasons, new insights into the eMSSM inflationary potential are expected in the coming years. However, a detailed implementation of the full exclusions for all analyses would be required to exploit the LHC data fully.

Inflation phenomenology: Conclusion

In this phenomenological part of the thesis, we have started discussing the slow-roll approximation and how it can yield cosmological predictions, a discussion that we have illustrated with standard slow-roll examples. Then, we have introduced the particle-physics SM and the MSSM.

We have then introduced an extended version of MSSM, eMSSM, in which some flat directions can undergo a slow-roll regime. We have investigated it as a test case to combine all HEP and cosmological observations. In order to proceed, we had to make choices to specify the theoretical framework. For example, we relied on the simple correlation between A_6 and A_t obtained in mSUGRA assuming a Polonyi hidden sector and setting the SUSY breaking energy scale to the GUT scale, although we carried out the analysis in the more general phenomenological MSSM. Finally, we did not consider additional terms in the inflationary potentials when we derived $V_{\rm RGE}$ from $V_{\rm tree}$, such as anomalous dimension running effects in the inflaton field, or possible induced runnings if considering an additive constant to the potential, or non-RGEs loop-induced operators in the effective potential. These assumptions would require further investigation in future work.

Needless to say that a detection of the tensor-to-scalar ratio (for instance by LiteBIRD [153]), or of primordial non-Gaussianities, would question this model, see App. F for more details on this. However, we have demonstrated that this work is very timely given the fact that, to date, the measurements of n_s and A_s are already sufficiently accurate for inflationary potential analyses to be sensitive to radiative corrections on the parameters (and their accuracy will be significantly reduced in a close future [138]). This analysis opens the door towards full-scan studies of the parameter space combining all observational constraints, but it also paves the way for future studies of theories that are able to describe the physical processes at the low-energy scale of the LHC up to the inflation scale in a well-defined theoretical framework.

Conclusions and perspectives

Measuring with precision the CMB anisotropies in temperature and in polarisation over the last few decades has opened an exciting window into our understanding of cosmological inflation. Quantum fluctuations in the vacuum during this speculative phase of explosive expansion in the very young Universe may have acted as the initial seeds for the structure of our perturbed Universe. Understanding this era would shed light on the very high-energy physics processes that triggered it. The next key objective in cosmologists' agenda is the first indirect detection of inflationary-induced gravitational waves in the CMB polarisation. These waves are expected to leave characteristic parity-odd patterns at angular scales larger than approximately 1°, with an amplitude directly proportional to the energy scale of inflation.

A space-based experiment is particularly suited for this challenging detection as it allows for full-sky and multifrequency measurement of the microwave sky. This approach enables access to the largest angular scales while distinguishing foregrounds from the primordial signal, which do not share the same statistical characteristics over the sky and frequencies. This strategy is adopted by the LiteBIRD mission.

After an introductory Part I, in which we reviewed the standard model of cosmology, the slow-roll paradigm for inflation, and the status and forthcoming challenges for CMB experiments, we focused on the LiteBIRD mission and instrument in Part II.

Since the scientific goals of this mission impose extremely stringent requirements on sensitivity, a meticulous understanding of the instrument is essential to address all identified systematic effects. In this context, we presented the Instrument Model (IMo), which gathers all relevant instrumental quantities. This model is *key for the collaboration*, serving as input for simulation tools, performance codes, and data-analysis pipelines. The IMo is now in a *stable version*, with most important design quantities already included. We provided an overview of its current content, which may serve as documentation for the collaboration.

The IMo will evolve throughout future mission phases according to the collaboration's needs. In the short term, we plan to incorporate more systematic models to enable setting up more realistic simulations.

In Part III, we described a comprehensive pipeline designed to extract cosmological information from large-scale polarisation maps.

It involves producing a foreground-cleaned CMB map from the sky measured in each frequency channel, estimating the associated power spectrum, and estimating τ and r, the reionisation depth and the tensor-to-scalar ratio, respectively, along with their uncertainties. We demonstrated that our pipeline enables a fully agnostic analysis, *i.e.*, with *no need for prior foreground knowledge*, while *yielding reasonably good results*. It can be used as a cross-check analysis against more sophisticated methods that include prior knowledge on foreground models.

Further developments may include improving the power spectrum estimation at the largest scales. It should also include improving the likelihood function, possibly by modelling the foreground residuals with an effective shape.

Additionally, my future project includes incorporating information from experiments observing lower frequency bands (analysing C-BASS [367] data & combining it to SPASS [368], Quijote [369], Planck, SO, & mock LiteBIRD data), which will help characterise the foregrounds. This will be key to excluding any Galactic-induced false detection of the tensor-to-scalar ratio.

Finally, in Part IV, we explored how cosmological inflation can be explained by a high-energy theory that also describes particle physics.

This framework is an extended version of the Minimal Supersymmetric Standard Model (MSSM), which we have shown can robustly predict the cosmological observables related to primordial fluctuations while including the expected quantum corrections. For the first time, we have demonstrated that their inclusion is necessary to exclude parts of the parameter space with cosmological constraints. We have *performed a global fit that includes all current cosmological and particle-physics measurements*. This work illustrates the kind of combined analysis that can be performed within a given high-energy theory and that will be required in our quest to understand the most fundamental physics.

The possible next steps for this work are numerous and exciting, including a full exploration of the parameter space and computation of the reheating duration. Moreover, we identified conceptual issues that question the model's naturalness. A motivating follow-up for supersymmetric inflation would be to identify other inflaton candidates that do not suffer from the same fine-tuning issues and build a realistic model based on the robust methodology described in this work.

Appendices

A- Detectability of the slow-roll tensor tilt

In Sec. 3.3.1 we have discussed a very neat prediction of single-field slow-roll inflation for the tensor spectra index: $n_{\rm t} = -r/8$. The current best constraint on this parameter, including BK 2015 and 2018, Planck PR3 and PR4, and LIGO-Virgo-KAGRA, yield $-1.37 < n_{\rm t} < 0.42$ at 95% confidence level [62, 63]. If the relation is measured to be violated, slow-roll is ruled out. If the consistency relation is verified, it would be crucial evidence for single-field slow-roll. The question we want to answer in this appendix is whether such verification of the consistency relation will be possible with forthcoming experiments.

Let us assume that the single-field slow-roll relation is fulfilled, *i.e.* $n_{\rm t}^{\rm fid} = -r^{\rm fid}/8$, and place ourselves in an optimistic setup in which we have a noiseless experiment dominated by full-sky cosmic variance. We will derive $\sigma^{\rm Fisher}_{n_{\rm t},\,{\rm stat}}$ marginalised over r from a BB-based Fisher matrix. We furthermore assume conservatively that the true r value is at the current upper bound: $r^{\rm fid} = 0.032$ [61].

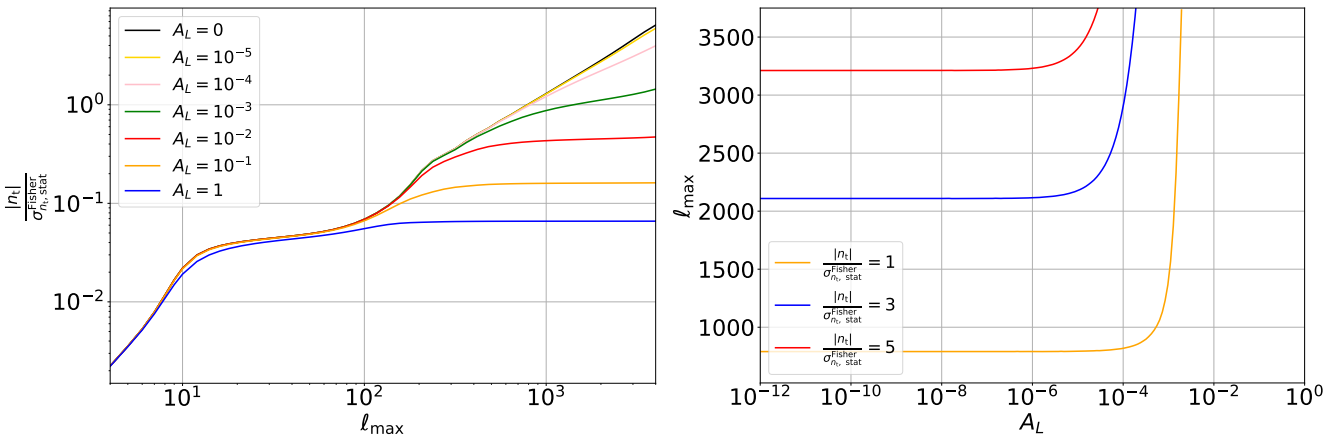

Figure 1.1: In these figures, we fix $r^{\rm fid} = 0.032$ and $n_{\rm t}^{\rm fid} = -0.032/8$. LEFT: For each fraction of removed lensing, A_L, we represent the signal-to-noise on $n_{\rm t}$ with respect to the maximum probed multipole $\ell_{\rm max}$. RIGHT: $\ell_{\rm max}$ needed to reach a 1σ (red), 3σ (orange) or 5σ (blue) detection of the slow-roll relation as a function of the delensing fraction.

In Fig. 1.1, for each fraction of removed lensing, A_L, we represent the signal to noise on $n_{\rm t}$ with respect to the maximum probed multipole $\ell_{\rm max}$. One observes that to reach a 3σ "detection" for $n_{\rm t}$ (which would confirm the slow-roll consistency relation), one can, for instance, remove 99.99% of the lensing while going at $\ell_{\rm max} = 2500$. Hence $n_{\rm t}^{\rm slow-roll}$ is extremely challenging to detect. The right panel provides another view illustrating a similar conclusion. This time, we show which $\ell_{\rm max}$ is needed to reach a 1σ (red), 3σ (orange) or 5σ (blue) detection of the slow-roll relation, as a function of the delensing fraction. We conclude that even for a 1σ "detection", a 99.9% delensing combined to a $\ell_{\rm max}$ of ~ 2000 is required.

If $r_{\rm fid}$ is actually lower than its current upper bound, or if one considers noise in the analysis, these conclusions are even worse. Hence, it seems out of reach from the next generation of CMB experiments to validate the slow-roll relation (though they might be able to exclude it). In Part. IV, we focus on extracting r solely from LiteBIRD BB power spectra.

B- IMo tree architecture

This appendix provides the tree architecture of various LiteBIRD IMo levels. The *observation* level is illustrated in Fig. 2.1, the *LFT* level in Fig. 2.2, the *MFT* and *HFT* levels in Fig. 2.3 and the *detector* level in Fig. 2.4. See Fig. 6.1 for an illustration of the frequency-channel-level architecture.

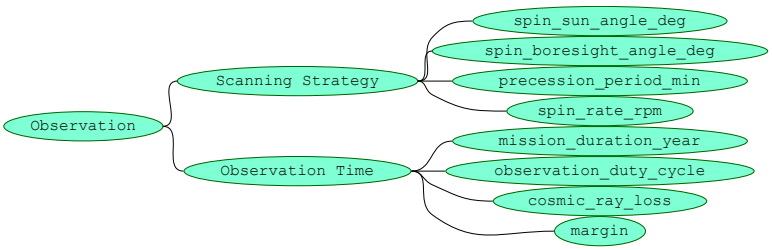

Figure 2.1: IMo tree architecture for the *observation* level.

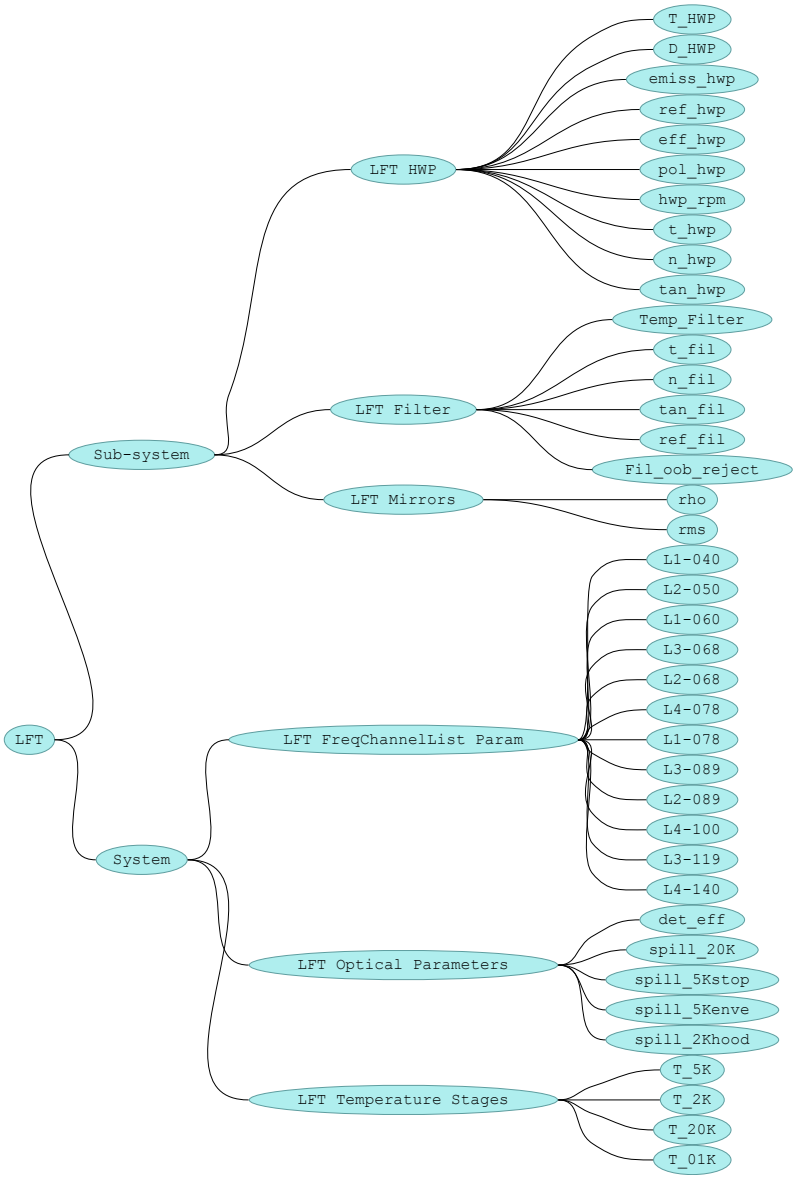

Figure 2.2: IMo tree architecture for *LFT* level.

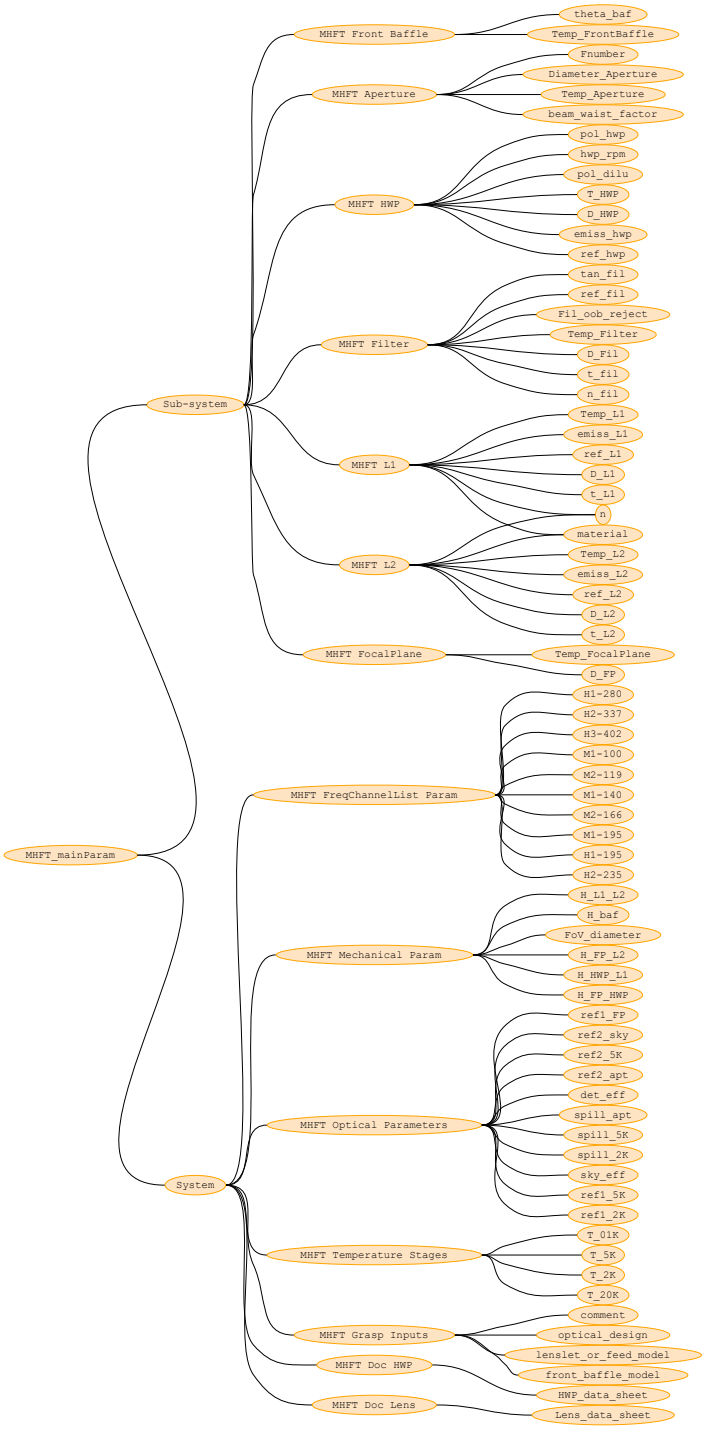

Figure 2.3: IMo tree architecture for *MFT* and *HFT* levels.

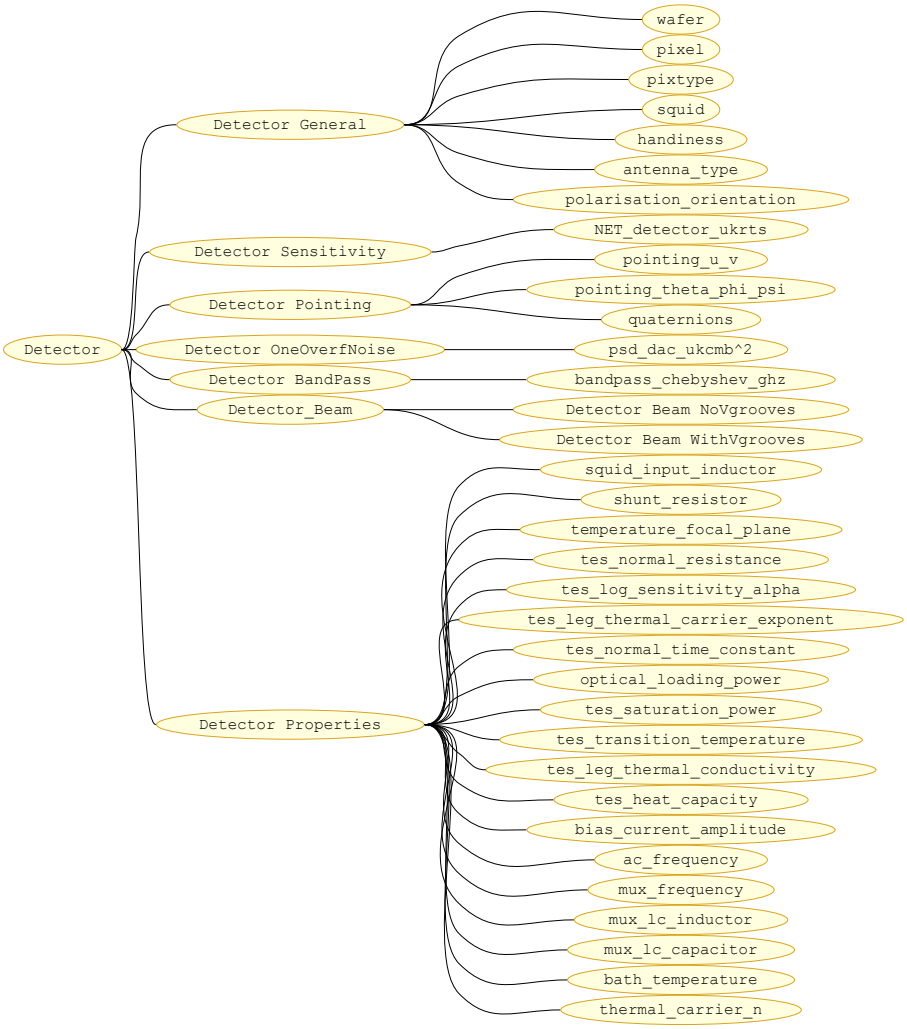

Figure 2.4: IMo tree architecture for the *detector* level.

C- Foreground complexity in L2-050 and H3-402

Figs. 3.1 and 3.2 are similar to Fig. III.2 respectively for the L2-050 and H3-402 frequency channels. The former corresponds to synchrotron-dominated frequencies, while the latter corresponds to dust-dominated frequencies (see Fig. 4.3).

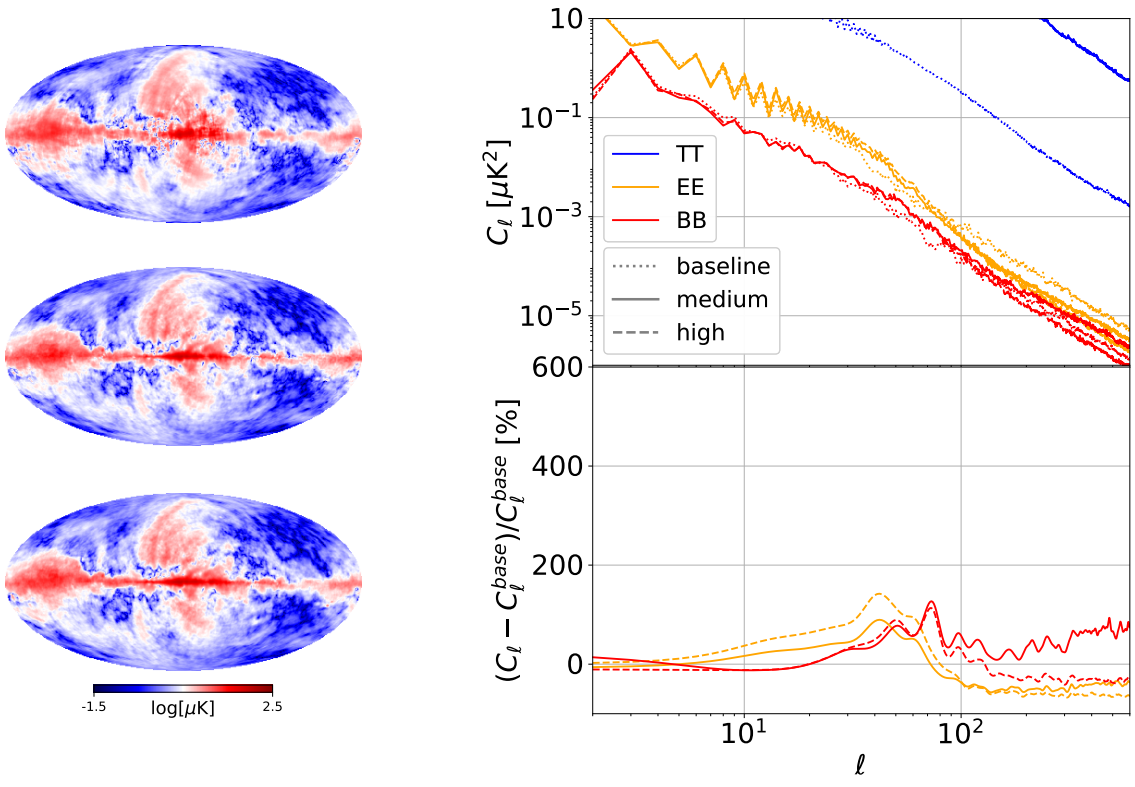

Figure 3.1: Refer to caption of Fig. III.2. This figure is similar but obtained for the L2-050 frequency channel observing at a central frequency of 50 GHz, a synchrotron-dominated regime.

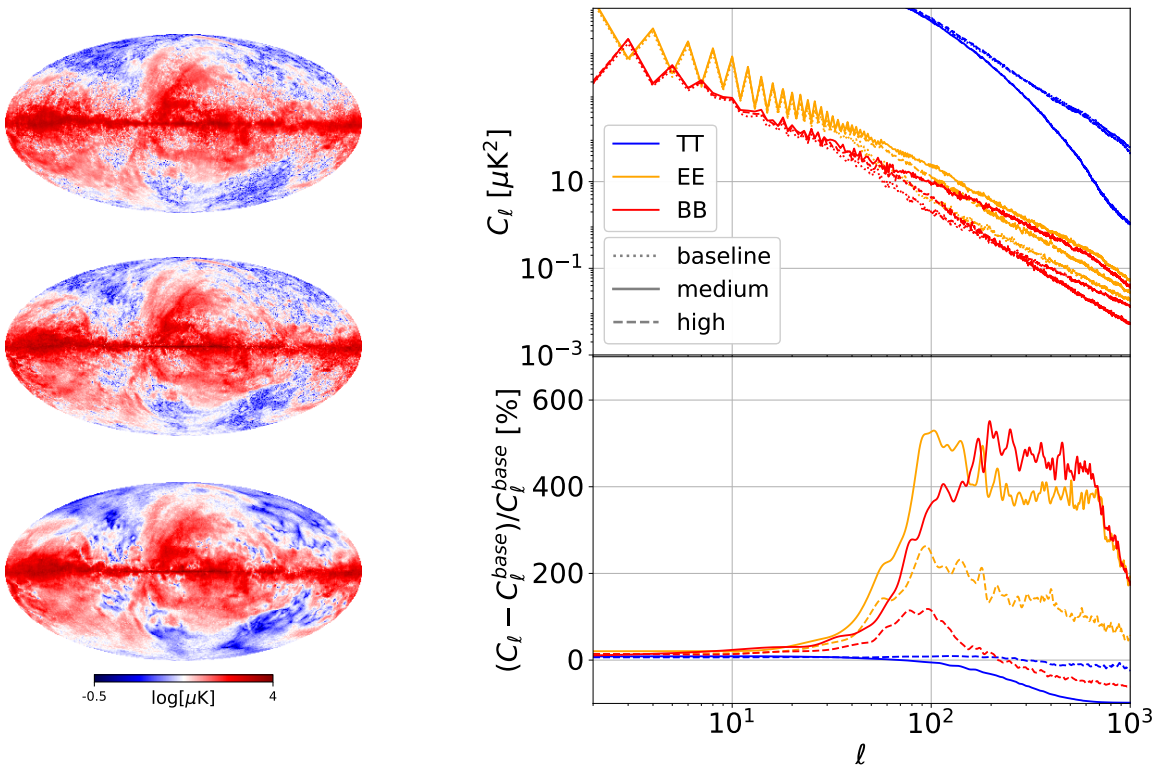

Figure 3.2: Refer to caption of Fig. III.2. This figure is similar but obtained for the H3-402 frequency channel observing at a central frequency of 402 GHz, a dust-dominated regime.

D- Comments on the HL transform

This appendix aims at showing in which sense the HL transform Eq. 9.17 does not simply consist in transforming $f_{C_\ell}(\hat{C}_\ell)$ given in Eq. 9.4 into a multivariate normal distribution $f_{C_\ell}(X_\ell) = \mathcal{N}(0,\ \Xi^{\rm fid}_{\ell\ell'})$ (contrary to what one could expect knowing that the quadratic form built from X_ℓ with covariance $\Xi^{fid}_{\ell\ell'}$ yields the exact full-sky likelihood).

For the explanation, let us place in the auto-power-spectrum full-sky setup. Let $\tau = \tau^{\rm fid} = \tau^{\rm th}$ (the associated power spectrum is denoted C_ℓ), and X_ℓ the HL transform of some noisy $\hat{C}_\ell$. From the basic property of C_ℓ distrib, $\varepsilon \equiv \frac{\hat{C}_\ell - C_\ell}{C_\ell}$ follows $\mathcal{N}(0, \frac{2}{2\ell+1})$ and we will assume that ℓ is large enough in order to expand in terms of ε. Then:

$$X_\ell = C_\ell\ g\left(\frac{\hat{C}_\ell}{C_\ell}\right) \tag{4.1}$$

$$= C_\ell\ g\left(\frac{C_\ell + \varepsilon C_\ell}{C_\ell}\right) \tag{4.2}$$

$$\xrightarrow[\varepsilon = \frac{\delta\hat{C}_\ell}{C_\ell} \ll 1]{} C_\ell\ \sqrt{2\left(1 + \varepsilon - \varepsilon + \frac{\varepsilon^2}{2} - \frac{\varepsilon^3}{3} - 1\right)} \tag{4.3}$$

$$\xrightarrow[\varepsilon \ll 1]{} C_\ell\ \sqrt{\varepsilon^2 - \frac{2}{3}\varepsilon^3} \tag{4.4}$$

$$\xrightarrow[\varepsilon \ll 1]{} C_\ell\ \varepsilon\left(1 - \frac{1}{3}\varepsilon\right). \tag{4.5}$$

Since $\langle\varepsilon\rangle = \left\langle\frac{\delta\hat{C}_\ell}{C_\ell}\right\rangle = 0$ and $\langle\varepsilon^2\rangle = \frac{2}{2\ell+1}$, one obtains that

$$\langle X_\ell\rangle \to \frac{-{\rm var}(\hat{C}_\ell)}{3C_\ell} = -\frac{2}{2\ell+1}\frac{C_\ell}{3}, \tag{4.6}$$

$${\rm var}(X_\ell) \to {\rm var}(\hat{C}_\ell) = \frac{2}{2\ell+1}C_\ell^2. \tag{4.7}$$

Equivalently, one can express the ratio between the mean and the standard deviation of X_ℓ as

$$\frac{\langle X_\ell \rangle}{\sigma(X_\ell)} \to -\frac{\sigma(\hat{C}_\ell)}{3C_\ell} \tag{4.8}$$

$$= -\frac{1}{3}\sqrt{\frac{2}{2\ell+1}} \tag{4.9}$$

We represent the formula Eq. 4.9 as a green dashed curve in Fig. 4.1, and compare it to empirical $\frac{\langle X_\ell \rangle}{\sigma(X_\ell)}$ estimated over 50,000 simulated full-sky spectra, in orange. For reference, we also display $\frac{\langle \hat{C}_\ell - C_\ell \rangle}{\sigma(\hat{C}_\ell)}$ in blue, as compared to the $\pm\frac{1}{\sqrt{5\times10^4}}$ gray band which represents the expected statistical flux. As expected, the empirical blue curve lies in the $\pm\frac{1}{\sqrt{5\times10^4}}$ region, meanwhile the empirical orange curve agrees with Eq. 4.9, even on the lowest multipoles where ε could become of order $\sqrt{2/5} \not\ll 1$.

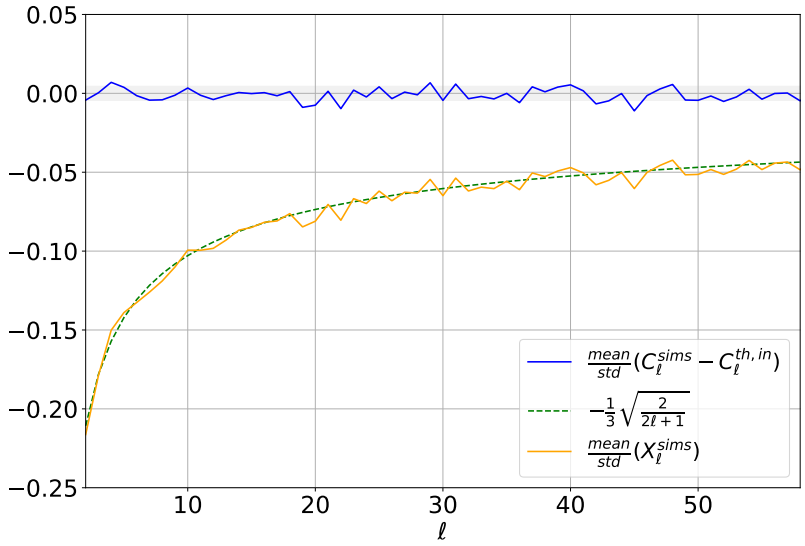

Figure 4.1: The green dashed curve represents Eq. 4.9, while $\frac{\langle X_\ell \rangle}{\sigma(X_\ell)}$ estimated over 50,000 simulated full-sky spectra is displayed in orange. For reference, we also show $\frac{\langle \hat{C}_\ell - C_\ell \rangle}{\sigma(\hat{C}_\ell)}$ in blue, as compared to the $\pm\frac{1}{\sqrt{5\times10^4}}$ gray band which represents the expected statistical flux.

E- Illustration of the $\ln R_{\rm rad}$ calculation from Γ_ϕ

We illustrate in Fig. 5.1, the exact numerical solution of the system Eqs. 10.5, 10.6, 10.7 in a particular tree-level eMSSM slow-rolling setup[1] and an arbitrary decay rate of the inflaton into light species: $\Gamma_\phi = 10^7$ GeV. We chose the latter small enough for the reheating to be not too efficient so it can be visualised in the figure.

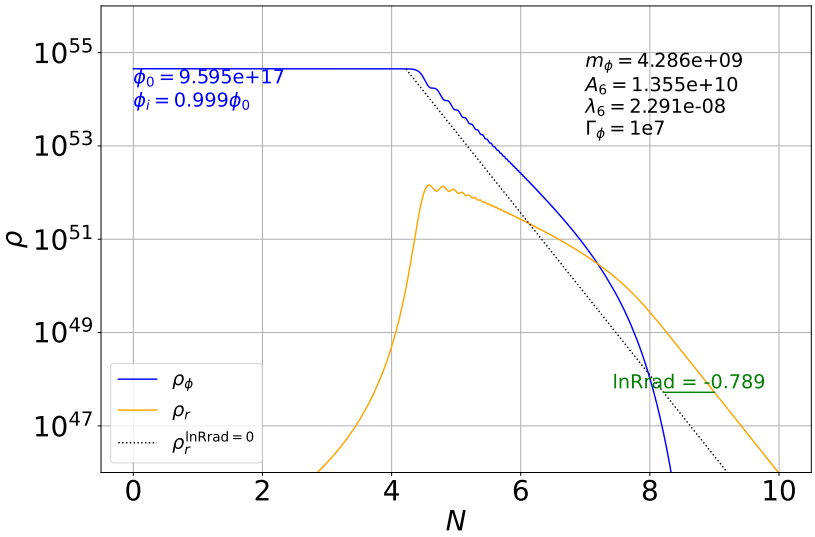

Figure 5.1: Solution of Eqs. 10.5, 10.6, 10.7. $m_\phi = 4.286 \times 10^9$ GeV, $A_6 = 1.355 \times 10^{10}$ GeV and $\lambda_6 = 2.291 \times 10^{-8}$

The crossing between the radiation density (orange) and the inflaton density (blue) determines the reheating time. The radiation density decreases as N^{-4}, while the inflaton density in the oscillating phase decreases on average as N^{-3}. The reason is that the bottom of the eMSSM potential is a ϕ^2 potential, which implies a matter-like era during reheating, see Eq. 10.4. $\ln R_{\rm rad}$

[1] which is able to predict $A_{\rm s}$ and $n_{\rm s}$: $\phi_0 = 9.595 \times 10^{17}$ GeV and $\phi_i = 0.999\phi_0$, corresponding to $m_\phi = 4.286 \times 10^9$ GeV, $A_6 = 1.355 \times 10^{10}$ GeV, $\lambda_6 = 2.291 \times 10^{-8}$

is the amount of e-folds removed by reheating with respect to an instantaneous transition. In this case, the predicted $\ln R_{\rm rad}$ is -0.789. In reality, one expects the coupling Γ_ϕ to be of order m_ϕ, so $\ln R_{\rm rad}$ is expected to be orders of magnitude below this mock numerical value [320].

F- Other eMSSM cosmological observables

In the framework of eMSSM, thanks to the usual relations, one can compute the predicted tensor-to-scalar ratio r and the running of the scalar index α_s for each allowed potential. To illustrate this discussion, we assume an instantaneous reheating, $\ln R_{\rm rad} = 0$, and stay at tree level.

The parameters' posterior distributions extracted from the MCMC chains of the Planck PR4 data release are shown in Fig. 6.1 together with the 68% and 95% CL contours. The superimposed dots on the (n_s, α_s) (left) and (n_s, r) (right) 2D distribution show the predictions of the MSSM potential (Eq. 11.15) that match the n_s and A_s, see Table 11.2, assuming an instantaneous reheating and no running of the parameters during the inflation period. The colour scale traces the field's value at the inflection point in GeV, which can be equivalently related to m_ϕ through Eq. 11.34. The low predicted values for the scalar-to-tensor ratio for this model justify why we have not used the upper limit on r in the core of the work.

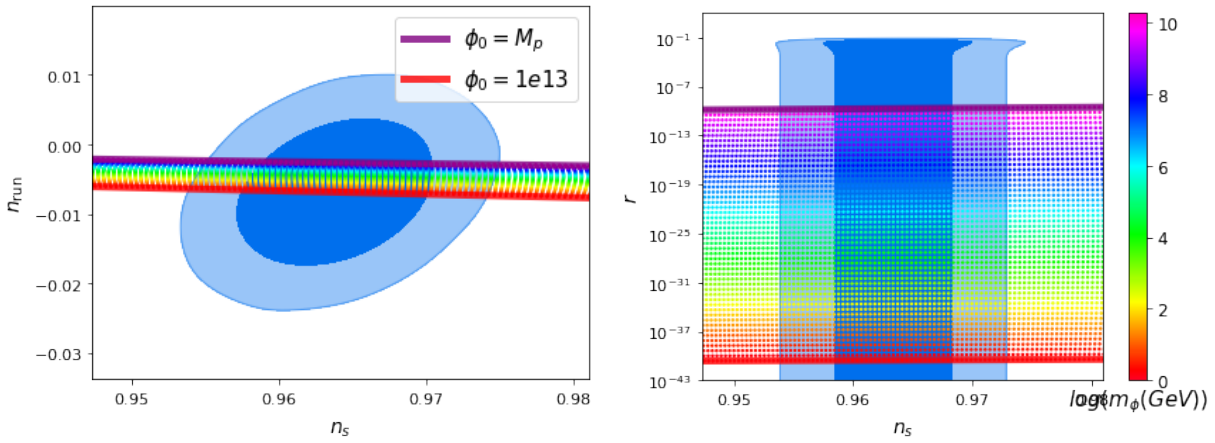

Figure 6.1: Predicted cosmological observables for potentials selected to produce A_s as currently measured by Planck, for a grid of ϕ_0 and α.

First, the right-hand-side figure illustrates that the tensor-to-scalar ratios predicted with MSSM-inflation are very small compared to other models (*e.g.*, Large Field Inflation, Starobinsky-Higgs inflation, Natural Inflation...) and far out of range from the current experiments. Concerning the High Energy Physics side, the model will be more and more constrained by future experiments as the lower m_ϕ will become invalidated, but Fig. 6.1 makes clear that even in a far

future, the parameter space will never be constrained simultaneously by both r-measurements and HEP experiments.

However, the same remark does not hold for the predicted α_s values. The current best standard deviation on this quantity is of order $\sigma(\alpha_s) \simeq 3 \times 10^{-3}$ and should be improved to $\sigma(\alpha_s) \simeq 1.7 \times 10^{-3}$ in the next decade with the near launch of SO, CMB-S4, LiteBIRD, and large-scale-structure surveys [138]. It could eventually become a constraint for eMSSM inflation in the next decades. Also, within this model, the predicted running always remains negative if one imposes A_s and n_s[1]. This implies that the predicted primordial power spectrum will be damped at higher wavelengths. Hence, the model predicts an absence of primordial black holes and induced gravitational waves. Eventually, in our single field framework, making use of Eq. 3.47 [99], the primordial non-gaussianities amplitude f_{NL} is predicted to be $\mathcal{O}(0.01)$, well below the current constraints.

If in the next years, one observes r, a bigger running than predicted, induced gravitational waves, primordial black holes, or primordial non-gaussianities, one should abandon the simplest version of the eMSSM.

The information in the CMB power spectra that can be directly constrained by the CMB (r, n_s or α_s, ...) is limited to a narrow range of k. Given a potential, the scalar power spectrum can be predicted much further than this narrow range. The comparison of the underlying spectra for the various cases studied in the previous sections is shown in Fig. 6.2 (the right-hand side is a zoom in the yellow region). All of them satisfy the constraint on A_s, n_s and α_s around $k = 0.05$ Mpc^{-1}: a quasi power-law slightly red-tilted. This information could become relevant thanks to new probes such as CMB spectral distortions or induced gravitational waves background.

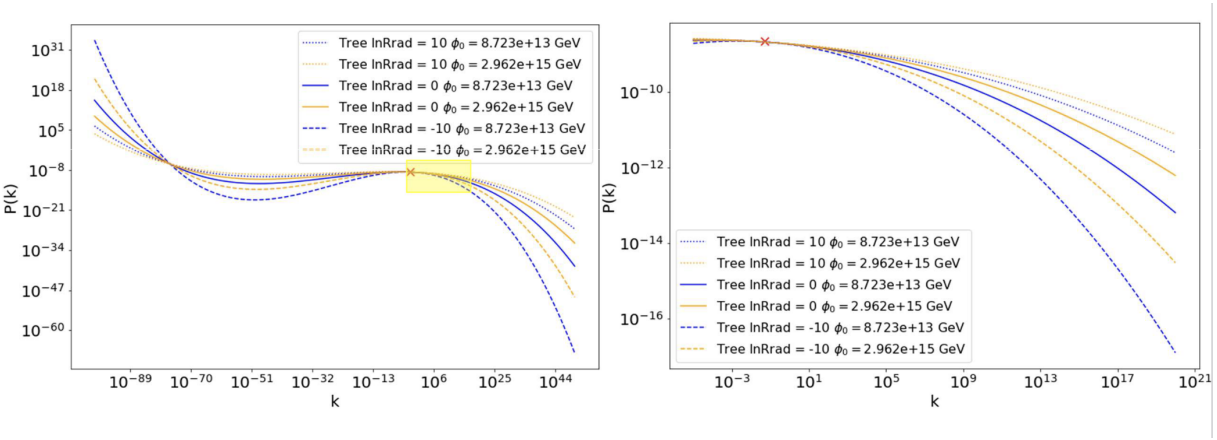

Figure 6.2: Full primordial power spectrum computed at third order in slow-roll on a wide wavelength range for various tree-level eMSSM-inflation configurations. k is given in Mpc^{-1} and the typical CMB observable window lies between $k = 0.0002$ and $k = 0.2$ Mpc^{-1}.

[1]It is also the case for most models currently allowed by A_s and n_s measurements, see [370].

Synthèse en français

La mesure des anisotropies du *Fond Diffus Cosmologique* (FDC) en *température* et en *polarisation* au cours des dernières décennies a ouvert une fenêtre fascinante sur notre compréhension de l'*inflation cosmologique*. Des fluctuations dans le vide quantique de cette phase d'expansion explosive de l'Univers naissant pourraient être à l'origine de la structure actuelle de notre Univers. Comprendre cette époque permettrait d'explorer la *physique à très hautes énergies* qui y est associée. À cet égard, la prochaine étape clé dans l'agenda des cosmologistes est la première détection indirecte d'un fond stochastique d'ondes gravitationnelles généré par l'inflation. Ce frémissement continu de l'espace-temps dans le jeune Univers aurait dû laisser des motifs caractéristiques dans les cartes de polarisation, à des échelles angulaires supérieures à environ 1°. Ce discret signal primordial devrait avoir une amplitude directement proportionnelle à l'échelle d'énergie de l'inflation. Le volet de gauche de la Figure FR.1 représente la carte de polarisation actuellement mesurée par le satellite Planck, pour l'instant dépourvue de tels motifs.

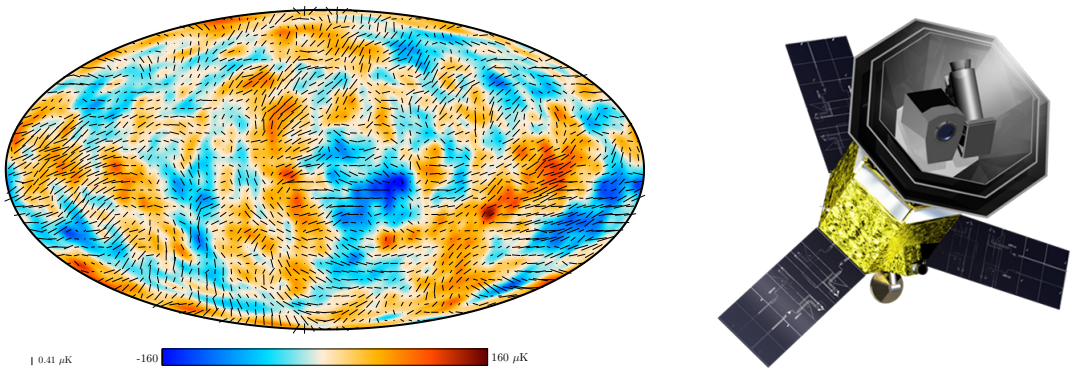

Figure FR.1: À GAUCHE : Carte du ciel délivré par le satellite Planck [41] illustrant le champ de polarisation du FDC, représenté par des segments de différentes longueurs superposés à la carte de température, en Kelvin en couleur (par rapport à la température moyenne).
À DROITE : Crédit : *ISAS/JAXA*. Vue d'artiste du satellite LiteBIRD.

Une expérience spatiale serait particulièrement adaptée à cette détection complexe, car elle permettrait de mesurer l'ensemble du ciel micro-onde sur plusieurs fréquences. Une telle approche offrirait l'accès aux plus grandes échelles angulaires et permettrait de distinguer les avant-plans galactiques du signal primordial, ces derniers n'ayant pas les mêmes caractéristiques statistiques sur le ciel ni sur les fréquences. Cette stratégie est adoptée par la JAXA et la collaboration scientifique internationale *LiteBIRD*, qui prévoient le lancement d'un satellite en 2032, représenté dans le volet de droite de la Figure FR.1.

Après une **Partie I** introductive, dans laquelle nous passons en revue le *modèle standard de la cosmologie*, le paradigme de l'inflation à *roulement lent* ainsi que les défis actuels et à venir pour les expériences du FDC, nous nous concentrons, dans la **Partie II**, sur l'introduction et l'étude de l'instrument LiteBIRD.

Les ambitieux objectifs scientifiques de cette mission exigent une sensibilité exceptionnelle de l'expérience, et une compréhension minutieuse de l'instrument est essentielle pour traiter tout potentiel effet systématique. Dans ce contexte, nous avons présenté le *Modèle de l'Instrument* (MdI) de LiteBIRD, qui rassemble toutes les quantités instrumentales dans un outil unique. Ce

modèle est central pour la collaboration car il sert d'entrée pour les codes de simulation, de performance et d'analyse de données. Le MdI est désormais dans une version stable et inclut la plupart des principales quantités instrumentales. Cette thèse décrit l'implémentation de cette première version et fourni un aperçu de son contenu actuel, servant de documentation pour la collaboration.

À titre d'illustration, la Figure FR.2 montre les pointages d'un des trois télescopes de LiteBIRD, appelé *LFT*. La dérivation de la représentation en quaternions de ces pointages, ainsi que leur implémentation au sein du MdI, constitue un exemple des contributions qui s'inscrivent au sein du premier volet de cette thèse. Le MdI est pensé pour évoluer tout au long des futures phases de la mission, en fonction des besoins de la collaboration. À court terme, nous prévoyons d'incorporer davantage de modèles pour les effets instrumentaux afin de mettre en place des simulations plus réalistes.

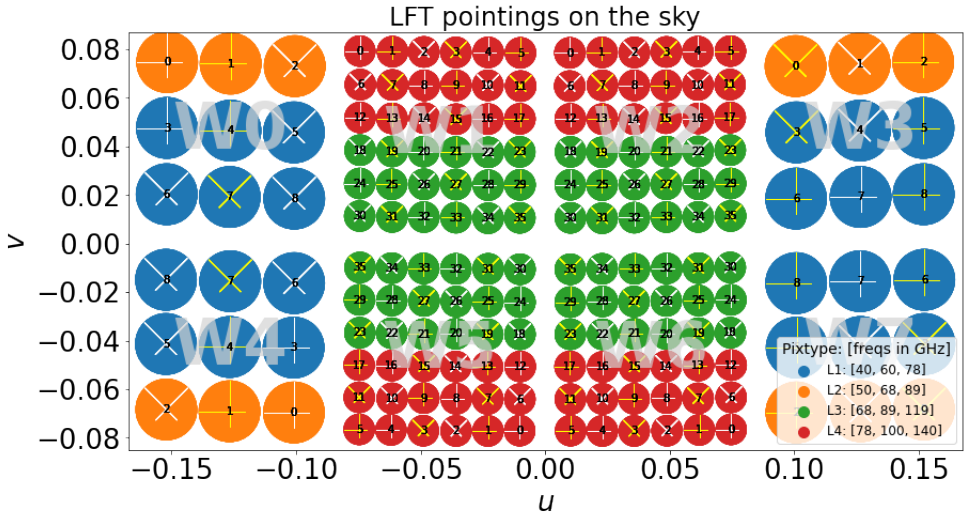

Figure FR.2: Pointages des différents détecteurs au sein du télescope LFT en coordonnées cartésiennes. L'orientation des antennes bolométriques est représentée par des segments blancs ou jaunes, avec les identifiants du panneau de détecteur, du pixel, du type de pixel et de la fréquence des détecteurs également spécifiés. Plus de détails sont disponibles dans la Section 6.4.

Dans la **Partie III**, nous avons décrit un pipeline complet conçu pour extraire l'information cosmologique à partir de cartes de polarisation à grande échelle.

Ce pipeline consiste dans la production d'une carte du CMB nettoyée des avant-plans à partir des observations dans chaque canal de fréquence, dans l'estimation du spectre en puissance associé, ainsi que de la détermination de τ et r, respectivement la profondeur optique de réionisation et le rapport tenseurs/scalaires, avec leurs incertitudes. La première de ces étapes est illustrée par la Figure FR.3, où nous comparons le niveau du signal E-mode aux niveaux de bruit et de résidus des avant-plans avant et après nettoyage. Nous avons appliqué ce pipeline à des simulations des futures données de LiteBIRD et démontré qu'il permet une analyse entièrement *agnostique*, c'est-à-dire, sans nécessiter de connaissances préalables sur les avant-plans, tout en produisant des résultats *robustes*. Une telle analyse servira de *référence aveugle* par rapport à d'autres méthodes reposant sur une connaissance préalable des avant-plans.

Les développements futurs pourront inclure une localisation dans l'espace des pixels de la méthode de suppression des avant-plans, l'amélioration de l'estimation du spectre en puissance aux plus grandes échelles, ainsi qu'une amélioration de la fonction de vraisemblance par la modélisation des résidus des avant-plans. Mon projet à venir comprend également l'incorporation d'informations provenant d'expériences observant à des fréquences plus basses (analyse des données de C-BASS [367] et leur combinaison avec SPASS [368], Quijote [369], Planck, SO et des données simulées de LiteBIRD), ce qui devrait aider à caractériser les avant-plans. Cela sera essentiel pour exclure toute détection erronée de r induite par la Galaxie.

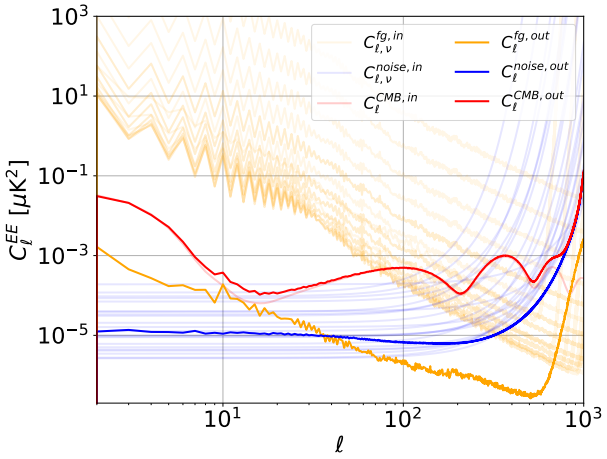

Figure FR.3: Les courbes en couleurs pastel représentent les spectres en puissance des 22 cartes d'entrée (chacune correspondant à un canal de fréquence de LiteBIRD), colorées en fonction du type de composant : les avant-plans galactiques en orange, le bruit en bleu, le signal CMB en rouge. Les courbes opaques représentent les spectres en puissance après le nettoyage pour les trois composants (associés aux mêmes couleurs qu'auparavant).

Enfin, dans la **Partie IV**, nous avons exploré comment l'inflation peut être expliquée par une théorie des hautes énergies également capable de prédire les observables de la physique des particules.

Le modèle en question consiste en une version du *Modèle Standard Supersymétrique Minimal* (MSSM). Dans ce cadre donné, nous avons démontré la possibilité de prédire de manière robuste les observables cosmologiques, tout en incluant les *corrections quantiques* inévitables. Pour la première fois, nous avons démontré que leur inclusion est nécessaire afin de dériver des contraintes valides sur l'espace des paramètres du MSSM. Nous avons *réalisé un ajustement global incluant toutes les mesures cosmologiques et de physique des particules actuelles*. Ce travail illustre le type d'analyse combinée qui peut être réalisée dans une théorie des hautes énergies donnée, et qui sera nécessaire dans notre quête pour lier la physique la plus fondamentale aux données du FDC.

Les prochaines étapes possibles pour ce travail sont nombreuses et enthousiasmantes, par exemple à travers une exploration complète de l'espace des paramètres ou encore en prédisant la rapidité de la production de radiation à la fin de l'inflation. De plus, nous avons identifié des problèmes conceptuels qui remettent en question la naturalité du modèle. Une suite motivante

pour l'inflation supersymétrique serait d'identifier d'autres candidats inflatons qui ne souffrent pas des mêmes problèmes et de construire un modèle réaliste basé sur la méthodologie robuste décrite dans ce travail.

En résumé, cette thèse vise à préparer l'avenir des expériences sur le FDC et leur possibilité de contraindre des modèles inflationnaires physiques. Avec une solide compréhension des instruments et du ciel micro-onde, combinée à des techniques rigoureuses d'analyse de données et à des prédictions robustes de modèles théoriques, les prochaines expériences sur le FDC ont le potentiel de nous apprendre beaucoup sur l'inflation cosmique et la physique des hautes énergies.

Bibliography

[1] *Esa - the history of the universe*, https://www.esa.int/ESA_Multimedia/Images/2015/02/The_history_of_the_Universe.

[2] D. J. Fixsen, E. S. Cheng, J. M. Gales, J. C. Mather, R. A. Shafer, and E. L. Wright, "The Cosmic Microwave Background spectrum from the full COBE FIRAS data set," *Astrophys. J.*, vol. 473, p. 576, 1996. DOI: 10.1086/178173. arXiv: astro-ph/9605054.

[3] A. D. Dolgov, "Neutrinos in cosmology," *Phys. Rept.*, vol. 370, pp. 333–535, 2002. DOI: 10.1016/S0370-1573(02)00139-4. arXiv: hep-ph/0202122.

[4] V. F. Mukhanov and G. V. Chibisov, "Quantum Fluctuations and a Nonsingular Universe," *JETP Lett.*, vol. 33, pp. 532–535, 1981.

[5] S. Dodelson, *Modern Cosmology*. Amsterdam: Academic Press, 2003, ISBN: 978-0-12-219141-1.

[6] J. M. Bardeen, "Gauge Invariant Cosmological Perturbations," *Phys. Rev. D*, vol. 22, pp. 1882–1905, 1980. DOI: 10.1103/PhysRevD.22.1882.

[7] V. F. Mukhanov, H. A. Feldman, and R. H. Brandenberger, "Theory of cosmological perturbations. Part 1. Classical perturbations. Part 2. Quantum theory of perturbations. Part 3. Extensions," *Phys. Rept.*, vol. 215, pp. 203–333, 1992. DOI: 10.1016/0370-1573(92)90044-Z.

[8] J. Martin and D. J. Schwarz, "The Influence of cosmological transitions on the evolution of density perturbations," *Phys. Rev. D*, vol. 57, pp. 3302–3316, 1998. DOI: 10.1103/PhysRevD.57.3302. arXiv: gr-qc/9704049.

[9] C.-P. Ma and E. Bertschinger, "Cosmological perturbation theory in the synchronous and conformal Newtonian gauges," *Astrophys. J.*, vol. 455, pp. 7–25, 1995. DOI: 10.1086/176550. arXiv: astro-ph/9506072.

[10] E. Lifshitz, "Republication of: On the gravitational stability of the expanding universe," *J. Phys. (USSR)*, vol. 10, no. 2, p. 116, 1946. DOI: 10.1007/s10714-016-2165-8.

[11] S. F. Dodelson S., *Modern Cosmology, 2nd edition*. Academic Press, 2021. DOI: 10.1016/C2017-0-01943-2.

[12] J. R. Bond and G. Efstathiou, "The statistics of cosmic background radiation fluctuations," *Mon. Not. Roy. Astron. Soc.*, vol. 226, pp. 655–687, 1987.

[13] D. Langlois, "Isocurvature cosmological perturbations and the CMB," *Comptes Rendus Physique*, vol. 4, pp. 953–959, 2003. DOI: 10.1016/j.crhy.2003.09.004.

[14] E. D. Stewart and D. H. Lyth, "A More accurate analytic calculation of the spectrum of cosmological perturbations produced during inflation," *Phys. Lett. B*, vol. 302, pp. 171–175, 1993. DOI: 10.1016/0370-2693(93)90379-V. arXiv: gr-qc/9302019.

[15] J.-O. Gong and E. D. Stewart, "The Density perturbation power spectrum to second order corrections in the slow roll expansion," *Phys. Lett. B*, vol. 510, pp. 1–9, 2001. DOI: 10.1016/S0370-2693(01)00616-5. arXiv: astro-ph/0101225.

[16] A. Kosowsky, "Cosmic microwave background polarization," *Annals Phys.*, vol. 246, pp. 49–85, 1996. DOI: 10.1006/aphy.1996.0020. arXiv: astro-ph/9501045.

[17] R. Crittenden, R. L. Davis, and P. J. Steinhardt, "Polarization of the microwave background due to primordial gravitational waves," *Astrophys. J. Lett.*, vol. 417, pp. L13–L16, 1993. DOI: 10.1086/187082. arXiv: astro-ph/9306027.

[18] M. Kamionkowski and E. D. Kovetz, "The Quest for B Modes from Inflationary Gravitational Waves," *Ann. Rev. Astron. Astrophys.*, vol. 54, pp. 227–269, 2016. DOI: 10.1146/annurev-astro-081915-023433. arXiv: 1510.06042 [astro-ph.CO].

[19] H. Kodama and M. Sasaki, "Cosmological Perturbation Theory," *Prog. Theor. Phys. Suppl.*, vol. 78, pp. 1–166, 1984. DOI: 10.1143/PTPS.78.1.

[20] W. Hu and N. Sugiyama, "Small scale cosmological perturbations: An Analytic approach," *Astrophys. J.*, vol. 471, pp. 542–570, 1996. DOI: 10.1086/177989. arXiv: astro-ph/9510117.

[21] J. Silk, "Cosmic black body radiation and galaxy formation," *Astrophys. J.*, vol. 151, pp. 459–471, 1968. DOI: 10.1086/149449.

[22] R. Barkana and A. Loeb, "In the beginning: The First sources of light and the reionization of the Universe," *Phys. Rept.*, vol. 349, pp. 125–238, 2001. DOI: 10.1016/S0370-1573(01)00019-9. arXiv: astro-ph/0010468.

[23] N. Y. Gnedin, "Cosmological reionization by stellar sources," *Astrophys. J.*, vol. 535, pp. 530–554, 2000. DOI: 10.1086/308876. arXiv: astro-ph/9909383.

[24] N. Y. Gnedin and J. P. Ostriker, "Reionization of the universe and the early production of metals," *Astrophys. J.*, vol. 486, p. 581, 1997. DOI: 10.1086/304548. arXiv: astro-ph/9612127.

[25] M. Zaldarriaga and D. D. Harari, "Analytic approach to the polarization of the cosmic microwave background in flat and open universes," *Phys. Rev. D*, vol. 52, pp. 3276–3287, 1995. DOI: 10.1103/PhysRevD.52.3276. arXiv: astro-ph/9504085.

[26] A. Lewis, "Cosmological parameters from WMAP 5-year temperature maps," *Phys. Rev. D*, vol. 78, p. 023 002, 2008. DOI: 10.1103/PhysRevD.78.023002. arXiv: 0804.3865 [astro-ph].

[27] M. Kaplinghat, M. Chu, Z. Haiman, G. Holder, L. Knox, and C. Skordis, "Probing the reionization history of the universe using the cosmic microwave background polarization," *Astrophys. J.*, vol. 583, pp. 24–32, 2003. DOI: 10.1086/344927. arXiv: astro-ph/0207591.

[28] M. Zaldarriaga and U. Seljak, "An all sky analysis of polarization in the microwave background," *Phys. Rev. D*, vol. 55, pp. 1830–1840, 1997. DOI: 10.1103/PhysRevD.55.1830. arXiv: astro-ph/9609170.

[29] D. Baumann *et al.*, "CMBPol Mission Concept Study: Probing Inflation with CMB Polarization," *AIP Conf. Proc.*, vol. 1141, no. 1, S. Dodelson *et al.*, Eds., pp. 10–120, 2009. DOI: 10.1063/1.3160885. arXiv: 0811.3919 [astro-ph].

[30] D. N. Spergel and M. Zaldarriaga, "CMB polarization as a direct test of inflation," *Phys. Rev. Lett.*, vol. 79, pp. 2180–2183, 1997. DOI: 10.1103/PhysRevLett.79.2180. arXiv: astro-ph/9705182.

[31] H. V. Peiris *et al.*, "First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Implications for inflation," *Astrophys. J. Suppl.*, vol. 148, pp. 213–231, 2003. DOI: 10.1086/377228. arXiv: astro-ph/0302225.

[32] M. Tegmark and A. de Oliveira-Costa, "How to measure CMB polarization power spectra without losing information," *Phys. Rev. D*, vol. 64, p. 063 001, 2001. DOI: 10.1103/PhysRevD.64.063001. arXiv: astro-ph/0012120.

[33] T. Louis, X. Garrido, A. Soussana, M. Tristram, S. Henrot-Versillé, and S. Vanneste, "Consistency of CMB experiments beyond cosmic variance," *Phys. Rev. D*, vol. 100, no. 2, p. 023 518, 2019. DOI: 10.1103/PhysRevD.100.023518. arXiv: 1905.06864 [astro-ph.CO].

[34] L. Knox, "Cosmic microwave background anisotropy observing strategy assessment," *Astrophys. J.*, vol. 480, p. 72, 1997. DOI: 10.1086/303959. arXiv: astro-ph/9606066.

[35] K. M. Gorski, B. D. Wandelt, F. K. Hansen, E. Hivon, and A. J. Banday, "The healpix primer," May 1999. arXiv: astro-ph/9905275.

[36] K. M. Górski *et al.*, "HEALPix - A Framework for high resolution discretization, and fast analysis of data distributed on the sphere," *Astrophys. J.*, vol. 622, pp. 759–771, 2005. DOI: 10.1086/427976. arXiv: astro-ph/0409513.

[37] *Healpy*, https://github.com/healpy/healpy/tree/main.

[38] A. Benoit *et al.*, "Cosmological constraints from Archeops," *Astron. Astrophys.*, vol. 399, pp. L25–L30, 2003. DOI: 10.1051/0004-6361:20021722. arXiv: astro-ph/0210306.

[39] D. N. Spergel *et al.*, "First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Determination of cosmological parameters," *Astrophys. J. Suppl.*, vol. 148, pp. 175–194, 2003. DOI: 10.1086/377226. arXiv: astro-ph/0302209.

[40] C. L. Bennett *et al.*, "Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results," *Astrophys. J. Suppl.*, vol. 208, p. 20, 2013. DOI: 10.1088/0067-0049/208/2/20. arXiv: 1212.5225 [astro-ph.CO].

[41] N. Aghanim *et al.*, "Planck 2018 results. I. Overview and the cosmological legacy of Planck," *Astron. Astrophys.*, vol. 641, A1, 2020. DOI: 10.1051/0004-6361/201833880. arXiv: 1807.06205 [astro-ph.CO].

[42] N. Aghanim *et al.*, "Planck intermediate results. XLIX. Parity-violation constraints from polarization data," *Astron. Astrophys.*, vol. 596, A110, 2016. DOI: 10.1051/0004-6361/201629018. arXiv: 1605.08633 [astro-ph.CO].

[43] R. A. Sunyaev and Y. B. Zeldovich, "Small scale fluctuations of relic radiation," *Astrophys. Space Sci.*, vol. 7, pp. 3–19, 1970.

[44] P. J. E. Peebles and J. T. Yu, "Primeval adiabatic perturbation in an expanding universe," *Astrophys. J.*, vol. 162, pp. 815–836, 1970. DOI: 10.1086/150713.

[45] K. S. Dawson *et al.*, "The Baryon Oscillation Spectroscopic Survey of SDSS-III," *Astron. J.*, vol. 145, p. 10, 2013. DOI: 10.1088/0004-6256/145/1/10. arXiv: 1208.0022 [astro-ph.CO].

[46] G. Mangano, G. Miele, S. Pastor, T. Pinto, O. Pisanti, and P. D. Serpico, "Relic neutrino decoupling including flavor oscillations," *Nucl. Phys. B*, vol. 729, pp. 221–234, 2005. DOI: 10.1016/j.nuclphysb.2005.09.041. arXiv: hep-ph/0506164.

[47] J. J. Bennett *et al.*, "Towards a precision calculation of $N_{\rm eff}$ in the Standard Model II: Neutrino decoupling in the presence of flavour oscillations and finite-temperature QED," *JCAP*, vol. 04, p. 073, 2021. DOI: 10.1088/1475-7516/2021/04/073. arXiv: 2012.02726 [hep-ph].

[48] E. Abdalla *et al.*, "Cosmology intertwined: A review of the particle physics, astrophysics, and cosmology associated with the cosmological tensions and anomalies," *JHEAp*, vol. 34, pp. 49–211, 2022. DOI: 10.1016/j.jheap.2022.04.002. arXiv: 2203.06142 [astro-ph.CO].

[49] A. G. Riess and L. Breuval, "The Local Value of H0," *IAU Symp.*, vol. 376, pp. 15–29, 2022. DOI: 10.1017/S1743921323003034. arXiv: 2308.10954 [astro-ph.CO].

[50] N. Schöneberg, G. Franco Abellán, A. Pérez Sánchez, S. J. Witte, V. Poulin, and J. Lesgourgues, "The H0 Olympics: A fair ranking of proposed models," *Phys. Rept.*, vol. 984, pp. 1–55, 2022. DOI: 10.1016/j.physrep.2022.07.001. arXiv: 2107.10291 [astro-ph.CO].

[51] K. J. Andersen *et al.*, "BeyondPlanck I. Global Bayesian analysis of the Planck Low Frequency Instrument data," in *BeyondPlanck Release Conference*, Nov. 2020. DOI: 10.1051/0004-6361/202244953. arXiv: 2011.05609 [astro-ph.CO].

[52] Y. Akrami *et al.*, "*Planck* intermediate results. LVII. Joint Planck LFI and HFI data processing," *Astron. Astrophys.*, vol. 643, A42, 2020. DOI: 10.1051/0004-6361/202038073. arXiv: 2007.04997 [astro-ph.CO].

[53] M. Tristram *et al.*, "Cosmological parameters derived from the final Planck data release (PR4)," *Astron. Astrophys.*, vol. 682, A37, 2024. DOI: 10.1051/0004-6361/202348015. arXiv: 2309.10034 [astro-ph.CO].

[54] J. E. Ruhl *et al.*, "The South Pole Telescope," *Proc. SPIE Int. Soc. Opt. Eng.*, vol. 5498, p. 11, 2004. DOI: 10.1117/12.552473. arXiv: astro-ph/0411122.

[55] A. Kosowsky, "The Atacama Cosmology Telescope," *New Astron. Rev.*, vol. 47, pp. 939–943, 2003. DOI: 10.1016/j.newar.2003.09.003. arXiv: astro-ph/0402234.

[56] J. E. Austermann *et al.*, "SPTpol: an instrument for CMB polarization measurements with the South Pole Telescope," *Proc. SPIE Int. Soc. Opt. Eng.*, vol. 8452, 84521E, 2012. DOI: 10.1117/12.927286. arXiv: 1210.4970 [astro-ph.IM].

[57] R. J. Thornton *et al.*, "The Atacama Cosmology Telescope: The polarization-sensitive ACTPol instrument," *Astrophys. J. Suppl.*, vol. 227, no. 2, p. 21, 2016. DOI: 10.3847/1538-4365/227/2/21. arXiv: 1605.06569 [astro-ph.IM].

[58] L. Balkenhol *et al.*, "Measurement of the CMB temperature power spectrum and constraints on cosmology from the SPT-3G 2018 TT, TE, and EE dataset," *Phys. Rev. D*, vol. 108, no. 2, p. 023 510, 2023. DOI: 10.1103/PhysRevD.108.023510. arXiv: 2212.05642 [astro-ph.CO].

[59] S. W. Henderson *et al.*, "Advanced ACTPol Cryogenic Detector Arrays and Readout," *J. Low Temp. Phys.*, vol. 184, no. 3-4, P. Camus, A. Juillard, and A. Monfardini, Eds., pp. 772–779, 2016. DOI: 10.1007/s10909-016-1575-z. arXiv: 1510.02809 [astro-ph.IM].

[60] P. A. R. Ade *et al.*, "Bicep/KeckXV: The Bicep3 Cosmic Microwave Background Polarimeter and the First Three-year Data Set," *Astrophys. J.*, vol. 927, no. 1, p. 77, 2022. DOI: 10.3847/1538-4357/ac4886. arXiv: 2110.00482 [astro-ph.IM].

[61] M. Tristram *et al.*, "Improved limits on the tensor-to-scalar ratio using BICEP and Planck data," *Phys. Rev. D*, vol. 105, no. 8, p. 083 524, 2022. DOI: 10.1103/PhysRevD.105.083524. arXiv: 2112.07961 [astro-ph.CO].

[62] G. Galloni, N. Bartolo, S. Matarrese, M. Migliaccio, A. Ricciardone, and N. Vittorio, "Updated constraints on amplitude and tilt of the tensor primordial spectrum," *JCAP*, vol. 04, p. 062, 2023. DOI: 10.1088/1475-7516/2023/04/062. arXiv: 2208.00188 [astro-ph.CO].

[63] G. Galloni, S. Henrot-Versillé, and M. Tristram, "Robust constraints on tensor perturbations from cosmological data: a comparative analysis from Bayesian and frequentist perspectives," May 2024. arXiv: 2405.04455 [astro-ph.CO].

[64] Y. Akrami *et al.*, "Planck 2018 results. X. Constraints on inflation," *Astron. Astrophys.*, vol. 641, A10, 2020. DOI: 10.1051/0004-6361/201833887. arXiv: 1807.06211 [astro-ph.CO].

[65] P. A. R. Ade *et al.*, "Improved Constraints on Primordial Gravitational Waves using Planck, WMAP, and BICEP/Keck Observations through the 2018 Observing Season," *Phys. Rev. Lett.*, vol. 127, no. 15, p. 151 301, 2021. DOI: 10.1103/PhysRevLett.127.151301. arXiv: 2110.00483 [astro-ph.CO].

[66] M. Tristram *et al.*, "Planck constraints on the tensor-to-scalar ratio," *Astron. Astrophys.*, vol. 647, A128, 2021. DOI: 10.1051/0004-6361/202039585. arXiv: 2010.01139 [astro-ph.CO].

[67] A. A. Starobinsky, "A New Type of Isotropic Cosmological Models Without Singularity," *Phys. Lett. B*, vol. 91, I. M. Khalatnikov and V. P. Mineev, Eds., pp. 99–102, 1980. DOI: 10.1016/0370-2693(80)90670-X.

[68] K. Sato, "First Order Phase Transition of a Vacuum and Expansion of the Universe," *Mon. Not. Roy. Astron. Soc.*, vol. 195, pp. 467–479, 1981.

[69] A. H. Guth, "The Inflationary Universe: A Possible Solution to the Horizon and Flatness Problems," *Phys. Rev. D*, vol. 23, L.-Z. Fang and R. Ruffini, Eds., pp. 347–356, 1981. DOI: 10.1103/PhysRevD.23.347.

[70] W. Rindler, "Visual Horizons in World Models," *Mon. Not. Roy. Astron. Soc.*, vol. 116, no. 6, pp. 662–677, 1956. DOI: 10.1093/mnras/116.6.662.

[71] C. W. Misner, "The Isotropy of the universe," *Astrophys. J.*, vol. 151, pp. 431–457, 1968. DOI: 10.1086/149448.

[72] J. Martin, "Inflation and precision cosmology," *Braz. J. Phys.*, vol. 34, L. A. Ferreira, Ed., pp. 1307–1321, 2004. DOI: 10.1590/S0103-97332004000700005. arXiv: astro-ph/0312492.

[73] G. F. R. Ellis and W. R. Stoeger S. J., "Horizons in Inflationary Universes," *Class. Quant. Grav.*, vol. 5, pp. 207–220, 1988. DOI: 10.1088/0264-9381/5/1/024.

[74] H. Georgi and S. L. Glashow, "Unity of All Elementary Particle Forces," *Phys. Rev. Lett.*, vol. 32, pp. 438–441, 1974. DOI: 10.1103/PhysRevLett.32.438.

[75] V. Vennin, "Cosmological Inflation: Theoretical Aspects and Observational Constraints," Ph.D. dissertation, Paris, Inst. Astrophys., 2014.

[76] M. Gasperini and G. Veneziano, "Pre - big bang in string cosmology," *Astropart. Phys.*, vol. 1, pp. 317–339, 1993. DOI: 10.1016/0927-6505(93)90017-8. arXiv: hep-th/9211021.

[77] J. Khoury, B. A. Ovrut, P. J. Steinhardt, and N. Turok, "Density perturbations in the ekpyrotic scenario," *Phys. Rev. D*, vol. 66, p. 046 005, 2002. DOI: 10.1103/PhysRevD.66.046005. arXiv: hep-th/0109050.

[78] G. 't Hooft, "Magnetic Monopoles in Unified Gauge Theories," *Nucl. Phys. B*, vol. 79, J. C. Taylor, Ed., pp. 276–284, 1974. DOI: 10.1016/0550-3213(74)90486-6.

[79] A. H. Guth and S. H. H. Tye, "Phase Transitions and Magnetic Monopole Production in the Very Early Universe," *Phys. Rev. Lett.*, vol. 44, p. 631, 1980, [Erratum: Phys.Rev.Lett. 44, 963 (1980)]. DOI: 10.1103/PhysRevLett.44.631.

[80] M. B. Einhorn, D. L. Stein, and D. Toussaint, "Are Grand Unified Theories Compatible with Standard Cosmology?" *Phys. Rev. D*, vol. 21, p. 3295, 1980. DOI: 10.1103/PhysRevD.21.3295.

[81] Y. B. Zeldovich and M. Y. Khlopov, "On the Concentration of Relic Magnetic Monopoles in the Universe," *Phys. Lett. B*, vol. 79, pp. 239–241, 1978. DOI: 10.1016/0370-2693(78)90232-0.

[82] J. Preskill, "Cosmological Production of Superheavy Magnetic Monopoles," *Phys. Rev. Lett.*, vol. 43, p. 1365, 1979. DOI: 10.1103/PhysRevLett.43.1365.

[83] V. Vennin, "Cosmological inflation : theoretical aspects and observational constraints," Theses, Université Pierre et Marie Curie - Paris VI, Sep. 2014.

[84] A. D. Linde, "A New Inflationary Universe Scenario: A Possible Solution of the Horizon, Flatness, Homogeneity, Isotropy and Primordial Monopole Problems," *Phys. Lett. B*, vol. 108, L.-Z. Fang and R. Ruffini, Eds., pp. 389–393, 1982. DOI: 10.1016/0370-2693(82)91219-9.

[85] A. D. Linde, "Chaotic Inflation," *Phys. Lett. B*, vol. 129, pp. 177–181, 1983. DOI: 10.1016/0370-2693(83)90837-7.

[86] A. Albrecht and P. J. Steinhardt, "Cosmology for Grand Unified Theories with Radiatively Induced Symmetry Breaking," *Phys. Rev. Lett.*, vol. 48, L.-Z. Fang and R. Ruffini, Eds., pp. 1220–1223, 1982. DOI: 10.1103/PhysRevLett.48.1220.

[87] A. D. Linde, *Particle physics and inflationary cosmology*. 1990, vol. 5, pp. 1–362.

[88] D. H. Lyth and A. Riotto, "Particle physics models of inflation and the cosmological density perturbation," *Phys. Rept.*, vol. 314, pp. 1–146, 1999. DOI: 10.1016/S0370-1573(98)00128-8. arXiv: hep-ph/9807278.

[89] J. Martin, C. Ringeval, R. Trotta, and V. Vennin, "The Best Inflationary Models After Planck," *JCAP*, vol. 03, p. 039, 2014. DOI: 10.1088/1475-7516/2014/03/039. arXiv: 1312.3529 [astro-ph.CO].

[90] B. Ratra and P. J. E. Peebles, "Cosmological Consequences of a Rolling Homogeneous Scalar Field," *Phys. Rev. D*, vol. 37, p. 3406, 1988. DOI: 10.1103/PhysRevD.37.3406.

[91] A. R. Liddle, P. Parsons, and J. D. Barrow, "Formalizing the slow roll approximation in inflation," *Phys. Rev. D*, vol. 50, pp. 7222–7232, 1994. DOI: 10.1103/PhysRevD.50.7222. arXiv: astro-ph/9408015.

[92] J. Martin, C. Ringeval, and V. Vennin, "Encyclopædia Inflationaris," *Phys. Dark Univ.*, vol. 5-6, pp. 75–235, 2014. DOI: 10.1016/j.dark.2014.01.003. arXiv: 1303.3787 [astro-ph.CO].

[93] G. Weymann-Despres, S. Henrot-Versillé, G. Moultaka, V. Vennin, L. Duflot, and R. von Eckardstein, "MSSM inflation revisited: Toward a coherent description of high-energy physics and cosmology," *Phys. Rev. D*, vol. 108, no. 2, p. 023 511, 2023. DOI: 10.1103/PhysRevD.108.023511. arXiv: 2304.04534 [hep-ph].

[94] K. Tokeshi and V. Vennin, "Why Does Inflation Look Single Field to Us?" *Phys. Rev. Lett.*, vol. 132, no. 25, p. 251 001, 2024. DOI: 10.1103/PhysRevLett.132.251001. arXiv: 2310.16649 [astro-ph.CO].

[95] T. S. Bunch and P. C. W. Davies, "Quantum Field Theory in de Sitter Space: Renormalization by Point Splitting," *Proc. Roy. Soc. Lond. A*, vol. 360, pp. 117–134, 1978. DOI: 10.1098/rspa.1978.0060.

[96] N. D. Birrell and P. C. W. Davies, *Quantum Fields in Curved Space* (Cambridge Monographs on Mathematical Physics). Cambridge, UK: Cambridge Univ. Press, Feb. 1984, ISBN: 978-0-521-27858-4, 978-0-521-27858-4. DOI: 10.1017/CBO9780511622632.

[97] S. Bird, H. V. Peiris, and R. Easther, "Fine-tuning criteria for inflation and the search for primordial gravitational waves," *Phys. Rev. D*, vol. 78, p. 083 518, 2008. DOI: 10.1103/PhysRevD.78.083518. arXiv: 0807.3745 [astro-ph].

[98] S. Hotchkiss, G. German, G. G. Ross, and S. Sarkar, "Fine tuning and the ratio of tensor to scalar density fluctuations from cosmological inflation," *JCAP*, vol. 10, p. 015, 2008. DOI: 10.1088/1475-7516/2008/10/015. arXiv: 0804.2634 [astro-ph].

[99] J. M. Maldacena, "Non-Gaussian features of primordial fluctuations in single field inflationary models," *JHEP*, vol. 05, p. 013, 2003. DOI: 10.1088/1126-6708/2003/05/013. arXiv: astro-ph/0210603.

[100] T. Colas, "Open Effective Field Theories for primordial cosmology : dissipation, decoherence and late-time resummation of cosmological inhomogeneities," Ph.D. dissertation, Institut d'astrophysique spatiale & AstroParticule et Cosmologie, 2023.

[101] J. Martin, C. Ringeval, and V. Vennin, "Cosmic Inflation at the Crossroads," Apr. 2024. arXiv: 2404.10647 [astro-ph.CO].

[102] N. Aghanim *et al.*, "Planck 2018 results. V. CMB power spectra and likelihoods," *Astron. Astrophys.*, vol. 641, A5, 2020. DOI: 10.1051/0004-6361/201936386. arXiv: 1907.12875 [astro-ph.CO].

[103] D. Dutcher *et al.*, "Measurements of the E-mode polarization and temperature-E-mode correlation of the CMB from SPT-3G 2018 data," *Phys. Rev. D*, vol. 104, no. 2, p. 022 003, 2021. DOI: 10.1103/PhysRevD.104.022003. arXiv: 2101.01684 [astro-ph.CO].

[104] K. S. Dawson *et al.*, "The SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Overview and Early Data," *Astron. J.*, vol. 151, p. 44, 2016. DOI: 10.3847/0004-6256/151/2/44. arXiv: 1508.04473 [astro-ph.CO].

[105] A. Lewis and A. Challinor, "Weak gravitational lensing of the CMB," *Phys. Rept.*, vol. 429, pp. 1–65, 2006. DOI: 10.1016/j.physrep.2006.03.002. arXiv: astro-ph/0601594.

[106] D. Hanson, A. Challinor, and A. Lewis, "Weak lensing of the CMB," *Gen. Rel. Grav.*, vol. 42, pp. 2197–2218, 2010. DOI: 10.1007/s10714-010-1036-y. arXiv: 0911.0612 [astro-ph.CO].

[107] M. S. Madhavacheril *et al.*, "The Atacama Cosmology Telescope: DR6 Gravitational Lensing Map and Cosmological Parameters," *Astrophys. J.*, vol. 962, no. 2, p. 113, 2024. DOI: 10.3847/1538-4357/acff5f. arXiv: 2304.05203 [astro-ph.CO].

[108] T. Louis *et al.*, "The Atacama Cosmology Telescope: Two-Season ACTPol Spectra and Parameters," *JCAP*, vol. 06, p. 031, 2017. DOI: 10.1088/1475-7516/2017/06/031. arXiv: 1610.02360 [astro-ph.CO].

[109] N. Aghanim *et al.*, "Planck 2018 results. VIII. Gravitational lensing," *Astron. Astrophys.*, vol. 641, A8, 2020. DOI: 10.1051/0004-6361/201833886. arXiv: 1807.06210 [astro-ph.CO].

[110] J. Carron and A. Lewis, "Maximum a posteriori CMB lensing reconstruction," *Phys. Rev. D*, vol. 96, no. 6, p. 063 510, 2017. DOI: 10.1103/PhysRevD.96.063510. arXiv: 1704.08230 [astro-ph.CO].

[111] L. Legrand and J. Carron, "Robust and efficient CMB lensing power spectrum from polarization surveys," *Phys. Rev. D*, vol. 108, no. 10, p. 103 516, 2023. DOI: 10.1103/PhysRevD.108.103516. arXiv: 2304.02584 [astro-ph.CO].

[112] G. Stein, M. A. Alvarez, J. R. Bond, A. van Engelen, and N. Battaglia, "The Websky Extragalactic CMB Simulations," *JCAP*, vol. 10, p. 012, 2020. DOI: 10.1088/1475-7516/2020/10/012. arXiv: 2001.08787 [astro-ph.CO].

[113] R. Takahashi *et al.*, "Full-sky Gravitational Lensing Simulation for Large-area Galaxy Surveys and Cosmic Microwave Background Experiments," *Astrophys. J.*, vol. 850, no. 1, p. 24, 2017. DOI: 10.3847/1538-4357/aa943d. arXiv: 1706.01472 [astro-ph.CO].

[114] Y. Akrami *et al.*, "Planck 2018 results. IV. Diffuse component separation," *Astron. Astrophys.*, vol. 641, A4, 2020. DOI: 10.1051/0004-6361/201833881. arXiv: 1807.06208 [astro-ph.CO].

[115] B. G. Andersson, "Interstellar Grain Alignment - Observational Status," Aug. 2012. arXiv: 1208.4393 [astro-ph.GA].

[116] P. A. R. Ade *et al.*, "Planck intermediate results. XXII. Frequency dependence of thermal emission from Galactic dust in intensity and polarization," *Astron. Astrophys.*, vol. 576, A107, 2015. DOI: 10.1051/0004-6361/201424088. arXiv: 1405.0874 [astro-ph.GA].

[117] N. Aghanim *et al.*, "Planck 2018 results. XII. Galactic astrophysics using polarized dust emission," *Astron. Astrophys.*, vol. 641, A12, 2020. DOI: 10.1051/0004-6361/201833885. arXiv: 1807.06212 [astro-ph.GA].

[118] T. L. Svalheim *et al.*, "BEYONDPLANCK - XIV. Polarized foreground emission between 30 and 70 GHz," *Astron. Astrophys.*, vol. 675, A14, 2023. DOI: 10.1051/0004-6361/202243160. arXiv: 2011.08503 [astro-ph.CO].

[119] A. Kogut, "Synchrotron Spectral Curvature from 22 MHz to 23 GHz," *Astrophys. J.*, vol. 753, p. 110, 2012. DOI: 10.1088/0004-637X/753/2/110. arXiv: 1205.4041 [astro-ph.GA].

[120] C. Dickinson, R. D. Davies, and R. J. Davis, "Towards a free-free template for CMB foregrounds," *Mon. Not. Roy. Astron. Soc.*, vol. 341, p. 369, 2003. DOI: 10.1046/j.1365-8711.2003.06439.x. arXiv: astro-ph/0302024.

[121] B. T. Draine, *Physics of the Interstellar and Intergalactic Medium*. 2011.

[122] B. Keating, P. Timbie, A. Polnarev, and J. Steinberger, "Large angular scale polarization of the cosmic microwave background and the feasibility of its detection," *Astrophys. J.*, vol. 495, p. 580, 1998. DOI: 10.1086/305312. arXiv: astro-ph/9710087.

[123] E. M. Leitch, A. C. S. Readhead, T. J. Pearson, and S. T. Myers, "An Anomalous component of galactic emission," *Astrophys. J. Lett.*, vol. 486, p. L23, 1997. DOI: 10.1086/310823. arXiv: astro-ph/9705241.

[124] A. Kogut, "Anomalous microwave emission," *ASP Conf. Ser.*, vol. 181, p. 91, 1999. arXiv: astro-ph/9902307.

[125] B. T. Draine and A. Lazarian, "Diffuse galactic emission from spinning dust grains," *Astrophys. J. Lett.*, vol. 494, pp. L19–L22, 1998. DOI: 10.1086/311167. arXiv: astro-ph/9710152.

[126] Y. Ali-Haimoud, "Spinning dust radiation: a review of the theory," *Adv. Astron.*, vol. 2013, p. 462 697, 2013. DOI: 10.1155/2013/462697. arXiv: 1211.2748 [astro-ph.GA].

[127] A. Macià Escatllar and S. T. Bromley, "Assessing the viability of silicate nanoclusters as carriers of the anomalous microwave emission: a quantum mechanical study," **??***jnlA&A*, vol. 634, A77, A77, Feb. 2020. DOI: 10.1051/0004-6361/201936419.

[128] T. Hoang, A. Lazarian, and P. G. Martin, "Constraint on the Polarization of Electric Dipole Emission from Spinning Dust," *Astrophys. J.*, vol. 779, p. 152, 2013. DOI: 10.1088/0004-637X/779/2/152. arXiv: 1305.0276 [astro-ph.GA].

[129] J. S. Greaves, W. S. Holland, P. Friberg, and W. R. F. Dent, "Polarized co emission from molecular clouds," *Astrophys. J. Lett.*, vol. 512, p. L139, 1999. DOI: 10.1086/311888. arXiv: astro-ph/9812428.

[130] J. L. Puget *et al.*, "Tentative detection of a cosmic far - infrared background with COBE," *Astron. Astrophys.*, vol. 308, p. L5, 1996.

[131] M. Birkinshaw, "The Sunyaev-Zel'dovich effect," *Phys. Rept.*, vol. 310, pp. 97–195, 1999. DOI: 10.1016/S0370-1573(98)00080-5. arXiv: astro-ph/9808050.

[132] T. W. Morris *et al.*, "The Atacama Cosmology Telescope: Modeling bulk atmospheric motion," *Phys. Rev. D*, vol. 105, no. 4, p. 042 004, 2022. DOI: 10.1103/PhysRevD.105.042004. arXiv: 2111.01319 [astro-ph.IM].

[133] G. Polenta *et al.*, "Unbiased estimation of angular power spectrum," *JCAP*, vol. 11, p. 001, 2005. DOI: 10.1088/1475-7516/2005/11/001. arXiv: astro-ph/0402428.

[134] R. A. Fisher, *Statistical methods for research workers*, O. Edinburgh and B. (1925), Eds., ISBN: ISBN 0-05-002170-2.

[135] A. Taylor, A. Heavens, B. Ballinger, and M. Tegmark, "Cosmological parameter estimation from the CMB," in *Particle Physics and the Early Universe Conference: In Celebration of the Centenary of the Discovery of the Electron: Conference on Structure of the Microwave Background*, Apr. 1997. arXiv: astro-ph/9707265.

[136] L. Verde, "Statistical methods in cosmology," *Lect. Notes Phys.*, vol. 800, pp. 147–177, 2010. DOI: 10.1007/978-3-642-10598-2_4. arXiv: 0911.3105 [astro-ph.CO].

[137] K. M. Smith *et al.*, "CMBPol Mission Concept Study: Gravitational Lensing," *AIP Conf. Proc.*, vol. 1141, no. 1, S. Dodelson *et al.*, Eds., p. 121, 2009. DOI: 10.1063/1.3160886. arXiv: 0811.3916 [astro-ph].

[138] J. Errard, S. M. Feeney, H. V. Peiris, and A. H. Jaffe, "Robust forecasts on fundamental physics from the foreground-obscured, gravitationally-lensed CMB polarization," *JCAP*, vol. 03, p. 052, 2016. DOI: 10.1088/1475-7516/2016/03/052. arXiv: 1509.06770 [astro-ph.CO].

[139] H. Cramer, *Mathematical Methods of Statistics*, P. U. Press, Ed. DOI: http://dx.doi.org/https://doi.org/10.1515/9781400883868.

[140] C. Rao, *Information and the accuracy attainable in the estimation of statistical parameters*, B. M. 81, Ed.

[141] D. Huterer and M. Takada, "Calibrating the nonlinear matter power spectrum: Requirements for future weak lensing surveys," *Astropart. Phys.*, vol. 23, pp. 369–376, 2005. DOI: 10.1016/j.astropartphys.2005.02.006. arXiv: astro-ph/0412142.

[142] D. Huterer, M. Takada, G. Bernstein, and B. Jain, "Systematic errors in future weak lensing surveys: Requirements and prospects for self-calibration," *Mon. Not. Roy. Astron. Soc.*, vol. 366, pp. 101–114, 2006. DOI: 10.1111/j.1365-2966.2005.09782.x. arXiv: astro-ph/0506030.

[143] A. Amara and A. Refregier, "Systematic Bias in Cosmic Shear: Beyond the Fisher Matrix," *Mon. Not. Roy. Astron. Soc.*, vol. 391, pp. 228–236, 2008. DOI: 10.1111/j.1365-2966.2008.13880.x. arXiv: 0710.5171 [astro-ph].

[144] J. L. Bernal, N. Bellomo, A. Raccanelli, and L. Verde, "Beware of commonly used approximations. Part II. Estimating systematic biases in the best-fit parameters," *JCAP*, vol. 10, p. 017, 2020. DOI: 10.1088/1475-7516/2020/10/017. arXiv: 2005.09666 [astro-ph.CO].

[145] E. Di Valentino *et al.*, "Exploring cosmic origins with CORE: Cosmological parameters," *JCAP*, vol. 04, p. 017, 2018. DOI: 10.1088/1475-7516/2018/04/017. arXiv: 1612.00021 [astro-ph.CO].

[146] Z. Hou, R. Keisler, L. Knox, M. Millea, and C. Reichardt, "How Massless Neutrinos Affect the Cosmic Microwave Background Damping Tail," *Phys. Rev. D*, vol. 87, p. 083 008, 2013. DOI: 10.1103/PhysRevD.87.083008. arXiv: 1104.2333 [astro-ph.CO].

[147] J. Lesgourgues, L. Perotto, S. Pastor, and M. Piat, "Probing neutrino masses with cmb lensing extraction," *Phys. Rev. D*, vol. 73, p. 045 021, 2006. DOI: 10.1103/PhysRevD.73.045021. arXiv: astro-ph/0511735.

[148] N. Aghanim *et al.*, "Planck 2018 results. VI. Cosmological parameters," *Astron. Astrophys.*, vol. 641, A6, 2020, [Erratum: Astron.Astrophys. 652, C4 (2021)]. DOI: 10.1051/0004-6361/201833910. arXiv: 1807.06209 [astro-ph.CO].

[149] Y. Minami and E. Komatsu, "New Extraction of the Cosmic Birefringence from the Planck 2018 Polarization Data," *Phys. Rev. Lett.*, vol. 125, no. 22, p. 221 301, 2020. DOI: 10.1103/PhysRevLett.125.221301. arXiv: 2011.11254 [astro-ph.CO].

[150] P. Diego-Palazuelos *et al.*, "Cosmic Birefringence from Planck Public Release 4," in *56th Rencontres de Moriond on Cosmology*, Mar. 2022. arXiv: 2203.04830 [astro-ph.CO].

[151] N. Bartolo, S. Matarrese, and A. Riotto, "Non-Gaussianity in the Cosmic Microwave Background Anisotropies at Recombination in the Squeezed limit," *JCAP*, vol. 02, p. 017, 2012. DOI: 10.1088/1475-7516/2012/02/017. arXiv: 1109.2043 [astro-ph.CO].

[152] N. Galitzki *et al.*, "The Simons Observatory: Instrument Overview," *Proc. SPIE Int. Soc. Opt. Eng.*, vol. 10708, G. Z. Angeli and P. Dierickx, Eds., p. 1 070 804, 2018. DOI: 10.1117/12.2312985. arXiv: 1808.04493 [astro-ph.IM].

[153] E. Allys *et al.*, "Probing Cosmic Inflation with the LiteBIRD Cosmic Microwave Background Polarization Survey," Feb. 2022. DOI: 10.1093/ptep/ptac150. arXiv: 2202.02773 [astro-ph.IM].

[154] J. Chluba and R. A. Sunyaev, "The evolution of CMB spectral distortions in the early Universe," *Mon. Not. Roy. Astron. Soc.*, vol. 419, pp. 1294–1314, 2012. DOI: 10.1111/j.1365-2966.2011.19786.x. arXiv: 1109.6552 [astro-ph.CO].

[155] B. Maffei *et al.*, "BISOU: a balloon project for spectral observations of the early universe," *Proc. SPIE Int. Soc. Opt. Eng.*, vol. 12190, 121900A, 2022. DOI: 10.1117/12.2630136.

[156] P. A. Abell *et al.*, "LSST Science Book, Version 2.0," Dec. 2009. arXiv: 0912.0201 [astro-ph.IM].

[157] R. Laureijs *et al.*, "Euclid Definition Study Report," Oct. 2011. arXiv: 1110.3193 [astro-ph.CO].

[158] R. Akeson *et al.*, "The Wide Field Infrared Survey Telescope: 100 Hubbles for the 2020s," Feb. 2019. arXiv: 1902.05569 [astro-ph.IM].

[159] J. J. Han *et al.*, "NANCY: Next-generation All-sky Near-infrared Community surveY," Jun. 2023. arXiv: 2306.11784 [astro-ph.IM].

[160] A. G. Adame *et al.*, "DESI 2024 VI: Cosmological Constraints from the Measurements of Baryon Acoustic Oscillations," Apr. 2024. arXiv: 2404.03002 [astro-ph.CO].

[161] A. Weltman *et al.*, "Fundamental physics with the Square Kilometre Array," *Publ. Astron. Soc. Austral.*, vol. 37, e002, 2020. DOI: 10.1017/pasa.2019.42. arXiv: 1810.02680 [astro-ph.CO].

[162] T. D. Hoang *et al.*, "Testbed preparation of a small prototype polarization modulator for LiteBIRD Low-Frequency Telescope," *Proc. SPIE Int. Soc. Opt. Eng.*, vol. 12190, p. 1118, 2022. DOI: 10.1117/12.2630653. arXiv: 2208.03673 [astro-ph.IM].

[163] Y. Sakurai *et al.*, "Breadboard model of the polarization modulator unit based on a continuously rotating half-wave plate for the low-frequency telescope of the LiteBIRD space mission," in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy X*, J. Zmuidzinas and J.-R. Gao, Eds., ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 11453, Dec. 2020, 114534E, 114534E. DOI: 10.1117/12.2560289.

[164] T. Matsumura *et al.*, "Prototype design and evaluation of the nine-layer achromatic half-wave plate for the LiteBIRD low frequency telescope," *Proc. SPIE Int. Soc. Opt. Eng.*, vol. 10708, p. 1 070 847, 2018. DOI: 10.1117/12.2312431.

[165] F. Columbro *et al.*, "Polarization Modulator Unit Harness Thermal Design for the Mid- and High-Frequency Telescopes of the LiteBIRD Space Mission," *Proc. SPIE Int. Soc. Opt. Eng.*, vol. 11443, 114436Z, 2020. DOI: 10.1007/s10909-022-02919-1. arXiv: 2101.05188 [astro-ph.IM].

[166] Y. Sekimoto *et al.*, "Concept Design of Low Frequency Telescope for CMB B-mode Polarization satellite LiteBIRD," *Proc. SPIE Int. Soc. Opt. Eng.*, vol. 11453, p. 1 145 310, 2020. DOI: 10.1117/12.2561841. arXiv: 2101.06342 [astro-ph.IM].

[167] L. Lamagna *et al.*, "Optical modeling for the LiteBIRD Medium and High Frequency Telescope," *Proc. SPIE Int. Soc. Opt. Eng.*, vol. 12190, 121901R, 2022. DOI: 10.1117/12.2629271.

[168] L. Montier *et al.*, "Overview of the Medium and High Frequency Telescopes of the LiteBIRD satellite mission," *Proc. SPIE Int. Soc. Opt. Eng.*, vol. 11443, 114432G, 2020. DOI: 10.1117/12.2562243. arXiv: 2102.00809 [astro-ph.IM].

[169] J. Höhne *et al.*, "High-resolution x-ray spectrometry using iridium-gold phase transition thermometers," *X-Ray Spectrometry*, vol. 28, pp. 396–398, 1999. [Online]. Available: https://api.semanticscholar.org/CorpusID:96985150.

[170] S. Henrot-Versillé, M. Tomasi, H. Ishino, T. Matsumura, D. Poletti, and G. Puglisi, *Litebird imo: Definition specification document for a software implementation*, LiteBIRD internal note n°58, 2020.

[171] *Imo_litebird*, https://github.com/litebird/IMo_LiteBIRD.

[172] *Instrumentdb*, https://github.com/ziotom78/instrumentdb.

[173] *Libinsdb*, https://github.com/ziotom78/libinsdb.

[174] *Litebird_sim*, https://github.com/litebird/litebird_sim.

[175] *Imo_perf*, https://github.com/litebird/IMo_perf.

[176] D. T. Hoang *et al.*, "Bandpass mismatch error for satellite CMB experiments I: estimating the spurious signal," *JCAP*, vol. 12, p. 015, 2017. DOI: 10.1088/1475-7516/2017/12/015. arXiv: 1706.09486 [astro-ph.CO].

[177] Y. Takase *et al.*, "Multi-dimensional optimisation of the scanning strategy for the LiteBIRD space mission," Aug. 2024. arXiv: 2408.03040 [astro-ph.IM].

[178] *Grasp10*, https://www.ticra.com/software/grasp/.

[179] T. Hasebe *et al.*, *Litebird sensitivity calculation version 28.0*, LiteBIRD internal report, 2019.

[180] T. de Haan, *Litebird note 81: 1/f model v0.1*, LiteBIRD internal note n°81, 2023.

[181] H. T. Ihle *et al.*, "BEYONDPLANCK - VI. Noise characterization and modeling," *Astron. Astrophys.*, vol. 675, A6, 2023. DOI: 10.1051/0004-6361/202243619. arXiv: 2011.06650 [astro-ph.CO].

[182] C. Leloup *et al.*, "Impact of beam far side-lobe knowledge in the presence of foregrounds for LiteBIRD," Dec. 2023. arXiv: 2312.09001 [astro-ph.CO].

[183] M. Tominaga, M. Tsujimoto, S. L. Stever, T. Ghigna, H. Ishino, and K. Ebisawa, "Simulation of the cosmic ray effects for the LiteBIRD satellite observing the CMB B-mode polarization," in *SPIE Astronomical Telescopes + Instrumentation 2020*, Jan. 2021. DOI: 10.1117/12.2576127. arXiv: 2101.08410 [astro-ph.IM].

[184] *Litebird_sim's documentation*, https://litebird-sim.readthedocs.io/en/latest/.

[185] M. Reinecke, S. Belkner, and J. Carron, "Improved cosmic microwave background (de-)lensing using general spherical harmonic transforms," *Astron. Astrophys.*, vol. 678, A165, 2023. DOI: 10.1051/0004-6361/202346717. arXiv: 2304.10431 [astro-ph.CO].

[186] *Lenspyx*, https://github.com/carronj/lenspyx.

[187] B. Thorne, J. Dunkley, D. Alonso, and S. Naess, "The Python Sky Model: software for simulating the Galactic microwave sky," *Mon. Not. Roy. Astron. Soc.*, vol. 469, no. 3, pp. 2821–2833, 2017. DOI: 10.1093/mnras/stx949. arXiv: 1608.02841 [astro-ph.CO].

[188] *Summary of pysm 3 models*, https://pysm3.readthedocs.io/en/latest/models.html#sunyaevzeldovich-emission.

[189] R. Adam *et al.*, "Planck 2015 results. X. Diffuse component separation: Foreground maps," *Astron. Astrophys.*, vol. 594, A10, 2016. DOI: 10.1051/0004-6361/201525967. arXiv: 1502.01588 [astro-ph.CO].

[190] A. Kogut *et al.*, "Three-year wilkinson microwave anisotropy probe (wmap)* observations: Foreground polarization," *The Astrophysical Journal*, vol. 665, no. 1, p. 355, 2007. DOI: 10.1086/519754. [Online]. Available: https://dx.doi.org/10.1086/519754.

[191] J. Delabrouille *et al.*, "The pre-launch Planck Sky Model: a model of sky emission at submillimetre to centimetre wavelengths," *Astron. Astrophys.*, vol. 553, A96, 2013. DOI: 10.1051/0004-6361/201220019. arXiv: 1207.3675 [astro-ph.CO].

[192] M. A. Miville-Deschênes *et al.*, "Separation of anomalous and synchrotron emissions using WMAP polarization data," **??***jnlA&A*, vol. 490, no. 3, pp. 1093–1102, Nov. 2008. DOI: 10.1051/0004-6361:200809484. arXiv: 0802.3345 [astro-ph].

[193] E. Keihanen, H. Kurki-Suonio, T. Poutanen, D. Maino, and C. Burigana, "A maximum likelihood approach to the destriping technique," *Astron. Astrophys.*, vol. 428, pp. 287–298, 2004. DOI: 10.1051/0004-6361:200400060. arXiv: astro-ph/0304411.

[194] D. Maino, C. Burigana, K. M. Gorski, N. Mandolesi, and M. Bersanelli, "Removing 1/f noise stripes in cosmic microwave background anisotropy observations," *Astron. Astrophys.*, vol. 387, p. 356, 2002. DOI: 10.1051/0004-6361:20020242. arXiv: astro-ph/0202271.

[195] C. Burigana *et al.*, "A Preliminary study on destriping techniques of Planck/LFI measurements versus observational strategy," Nov. 1997. arXiv: astro-ph/9906360.

[196] E. Keihanen, R. Keskitalo, H. Kurki-Suonio, T. Poutanen, and A. S. Sirvio, "Making CMB temperature and polarization maps with Madam," *Astron. Astrophys.*, vol. 510, A57, 2010. DOI: 10.1051/0004-6361/200912813. arXiv: 0907.0367 [astro-ph.CO].

[197] W. N. Brandt, C. R. Lawrence, A. C. S. Readhead, J. N. Pakianathan, and T. M. Fiola, "Separation of Foreground Radiation from Cosmic Microwave Background Anisotropy Using Multifrequency Measurements,", vol. 424, p. 1, Mar. 1994. DOI: 10.1086/173867.

[198] H. K. Eriksen *et al.*, "CMB component separation by parameter estimation," *Astrophys. J.*, vol. 641, pp. 665–682, 2006. DOI: 10.1086/500499. arXiv: astro-ph/0508268.

[199] R. Stompor, S. Leach, F. Stivoli, and C. Baccigalupi, "Maximum likelihood algorithm for parametric component separation in cosmic microwave background experiments," *Monthly Notices of the Royal Astronomical Society*, vol. 392, no. 1, pp. 216–232, Dec. 2008, ISSN: 0035-8711. DOI: 10.1111/j.1365-2966.2008.14023.x. eprint: https://academic.oup.com/mnras/article-pdf/392/1/216/3707574/mnras0392-0216.pdf. [Online]. Available: https://doi.org/10.1111/j.1365-2966.2008.14023.x.

[200] J. F. Taylor, M. A. J. Ashdown, and M. P. Hobson, "Spectal Matching Independent Component Analysis, a method of blind component separation for observations of the Cosmic Microwave Backgound polarization," in *41st Rencontres de Moriond: Workshop on Cosmology: Contents and Structures of the Universe*, Hanoi: The Gioi, 2006, pp. 290–292.

[201] M. Tegmark, A. de Oliveira-Costa, and A. Hamilton, "A high resolution foreground cleaned CMB map from WMAP," *Phys. Rev. D*, vol. 68, p. 123 523, 2003. DOI: 10.1103/PhysRevD.68.123523. arXiv: astro-ph/0302496.

[202] C. Bennett *et al.*, "First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Foreground emission," *Astrophys. J. Suppl.*, vol. 148, p. 97, 2003. DOI: 10.1086/377252. arXiv: astro-ph/0302208.

[203] H. K. Eriksen, A. J. Banday, K. M. Gorski, and P. B. Lilje, "Foreground removal by an internal linear combination method: Limitations and implications," *Astrophys. J.*, vol. 612, pp. 633–646, 2004. DOI: 10.1086/422807. arXiv: astro-ph/0403098.

[204] J. Dick, M. Remazeilles, and J. Delabrouille, "Impact of calibration errors on CMB component separation using FastICA and ILC," *Mon. Not. Roy. Astron. Soc.*, vol. 401, p. 1602, 2010. DOI: 10.1111/j.1365-2966.2009.15798.x. arXiv: 0907.3105 [astro-ph.CO].

[205] M. Remazeilles, A. Rotti, and J. Chluba, "Peeling off foregrounds with the constrained moment ILC method to unveil primordial CMB B-modes," *Mon. Not. Roy. Astron. Soc.*, vol. 503, no. 2, pp. 2478–2498, 2021. DOI: 10.1093/mnras/stab648. arXiv: 2006.08628 [astro-ph.CO].

[206] J. Delabrouille, J. F. Cardoso, M. L. Jeune, M. Betoule, G. Fay, and F. Guilloux, "A full sky, low foreground, high resolution CMB map from WMAP," *Astron. Astrophys.*, vol. 493, p. 835, 2009. DOI: 10.1051/0004-6361:200810514. arXiv: 0807.0773 [astro-ph].

[207] A. Carones *et al.*, "Multiclustering needlet ILC for CMB B-mode component separation," *Mon. Not. Roy. Astron. Soc.*, vol. 525, no. 2, pp. 3117–3135, 2023. DOI: 10.1093/mnras/stad2423. arXiv: 2212.04456 [astro-ph.CO].

[208] J. Kim, P. Naselsky, and P. R. Christensen, "CMB map derived from the WMAP data through harmonic internal linear combination," *Phys. Rev. D*, vol. 77, p. 103 002, 2008. DOI: 10.1103/PhysRevD.77.103002. arXiv: 0803.1394 [astro-ph].

[209] *Anafast*, https://healpix.sourceforge.io/html/fac_anafast.htm.

[210] G. Efstathiou, "Hybrid estimation of cmb polarization power spectra," *Mon. Not. Roy. Astron. Soc.*, vol. 370, pp. 343–362, 2006. DOI: 10.1111/j.1365-2966.2006.10486.x. arXiv: astro-ph/0601107.

[211] G. Efstathiou, "Myths and truths concerning estimation of power spectra," *Mon. Not. Roy. Astron. Soc.*, vol. 349, p. 603, 2004. DOI: 10.1111/j.1365-2966.2004.07530.x. arXiv: astro-ph/0307515.

[212] E. Hivon, K. M. Gorski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen, "Master of the cosmic microwave background anisotropy power spectrum: a fast method for statistical analysis of large and complex cosmic microwave background data sets," *Astrophys. J.*, vol. 567, p. 2, 2002. DOI: 10.1086/338126. arXiv: astro-ph/0105302.

[213] A. Kogut *et al.*, "Wilkinson Microwave Anisotropy Probe (WMAP) first year observations: TE polarization," *Astrophys. J. Suppl.*, vol. 148, p. 161, 2003. DOI: 10.1086/377219. arXiv: astro-ph/0302213.

[214] M. Tristram, J. F. Macias-Perez, C. Renault, and D. Santos, "Xspect, estimation of the angular power spectrum by computing cross power spectra," *Mon. Not. Roy. Astron. Soc.*, vol. 358, p. 833, 2005. DOI: 10.1111/j.1365-2966.2005.08760.x. arXiv: astro-ph/0405575.

[215] M. Tristram, *Xpol*, https://gitlab.in2p3.fr/tristram/Xpol.

[216] E. F. Bunn and M. J. White, "The Four year COBE normalization and large scale structure," *Astrophys. J.*, vol. 480, pp. 6–21, 1997. DOI: 10.1086/303955. arXiv: astro-ph/9607060.

[217] G. Hinshaw *et al.*, "Band power spectra in the COBE DMR 4-year anisotropy maps," *Astrophys. J. Lett.*, vol. 464, pp. L17–L20, 1996. DOI: 10.1086/310074. arXiv: astro-ph/9601058.

[218] M. Tegmark, "How to measure CMB power spectra without losing information," *Phys. Rev. D*, vol. 55, pp. 5895–5907, 1997. DOI: 10.1103/PhysRevD.55.5895. arXiv: astro-ph/9611174.

[219] J. P. BfflOF, "Computing the distribution of quadratic forms in normal variables," 2005. [Online]. Available: https://api.semanticscholar.org/CorpusID:202545615.

[220] S. Vanneste, S. Henrot-Versillé, T. Louis, and M. Tristram, "Quadratic estimator for CMB cross-correlation," *Phys. Rev. D*, vol. 98, no. 10, p. 103 526, 2018. DOI: 10.1103/PhysRevD.98.103526. arXiv: 1807.02484 [astro-ph.CO].

[221] *Xqml*, https://gitlab.in2p3.fr/xQML/xQML.

[222] J. D. Bilbao-Ahedo, R. B. Barreiro, P. Vielva, E. Martínez-González, and D. Herranz, "ECLIPSE: a fast Quadratic Maximum Likelihood estimator for CMB intensity and polarization power spectra," *JCAP*, vol. 07, p. 034, 2021. DOI: 10.1088/1475-7516/2021/07/034. arXiv: 2104.08528 [astro-ph.CO].

[223] N. Aghanim *et al.*, "Planck intermediate results. XLVI. Reduction of large-scale systematic effects in HFI polarization maps and estimation of the reionization optical depth," *Astron. Astrophys.*, vol. 596, A107, 2016. DOI: 10.1051/0004-6361/201628890. arXiv: 1605.02985 [astro-ph.CO].

[224] W. J. Percival and M. L. Brown, "Likelihood methods for the combined analysis of CMB temperature and polarisation power spectra," *Mon. Not. Roy. Astron. Soc.*, vol. 372, pp. 1104–1116, 2006. DOI: 10.1111/j.1365-2966.2006.10910.x. arXiv: astro-ph/0604547.

[225] S. Hamimeche and A. Lewis, "Likelihood Analysis of CMB Temperature and Polarization Power Spectra," *Phys. Rev. D*, vol. 77, p. 103 013, 2008. DOI: 10.1103/PhysRevD.77.103013. arXiv: 0801.0554 [astro-ph].

[226] J. W. Henning *et al.*, "Measurements of the Temperature and E-Mode Polarization of the CMB from 500 Square Degrees of SPTpol Data," *Astrophys. J.*, vol. 852, no. 2, p. 97, 2018. DOI: 10.3847/1538-4357/aa9ff4. arXiv: 1707.09353 [astro-ph.CO].

[227] N. Aghanim *et al.*, "Planck 2015 results. XI. CMB power spectra, likelihoods, and robustness of parameters," *Astron. Astrophys.*, vol. 594, A11, 2016. DOI: 10.1051/0004-6361/201526926. arXiv: 1507.02704 [astro-ph.CO].

[228] A. Mangilli, S. Plaszczynski, and M. Tristram, "Large-scale cosmic microwave background temperature and polarization cross-spectra likelihoods," *Mon. Not. Roy. Astron. Soc.*, vol. 453, no. 3, pp. 3174–3189, 2015. DOI: 10.1093/mnras/stv1733. arXiv: 1503.01347 [astro-ph.CO].

[229] R. Adam *et al.*, "Planck intermediate results. XLVII. Planck constraints on reionization history," *Astron. Astrophys.*, vol. 596, A108, 2016. DOI: 10.1051/0004-6361/201628897. arXiv: 1605.03507 [astro-ph.CO].

[230] M. Sklar, "Fonctions de répartition à N dimensions et leurs marges," *Annales de l'ISUP*, vol. VIII, no. 3, pp. 229–231, 1959. [Online]. Available: https://hal.science/hal-04094463.

[231] M. Sato, K. Ichiki, and T. T. Takeuchi, "Copula Cosmology: Constructing a Likelihood Function," *Phys. Rev. D*, vol. 83, p. 023 501, 2011. DOI: 10.1103/PhysRevD.83.023501. arXiv: 1011.4997 [astro-ph.CO].

[232] R. J. Scherrer, A. A. Berlind, Q. Mao, and C. K. McBride, "From Finance to Cosmology: The Copula of Large-Scale Structure," *Astrophys. J. Lett.*, vol. 708, pp. L9–L13, 2010. DOI: 10.1088/2041-8205/708/1/L9. arXiv: 0909.5187 [astro-ph.CO].

[233] T. T. Takeuchi, "Constructing a bivariate distribution function with given marginals and correlation: application to the galaxy luminosity function," *Mon. Not. Roy. Astron. Soc.*, vol. 406, p. 1830, 2010. DOI: 10.1111/j.1365-2966.2010.16778.x. arXiv: 1004.0439 [astro-ph.CO].

[234] I.-G. Jiang, L.-C. Yeh, Y.-C. Chang, and W.-L. Hung, "Construction of Coupled Period-Mass Functions in Extrasolar Planets through the Nonparametric Approach," *Astron. J.*, vol. 137, pp. 329–336, 2009. DOI: 10.1088/0004-6256/137/1/329. arXiv: 0909.0824 [astro-ph.IM].

[235] K. Benabed, J. F. Cardoso, S. Prunet, and E. Hivon, "TEASING: a fast and accurate approximation for the low multipole likelihood of the Cosmic Microwave Background temperature," *Mon. Not. Roy. Astron. Soc.*, vol. 400, p. 219, 2009. DOI: 10.1111/j.1365-2966.2009.15202.x. arXiv: 0901.4537 [astro-ph.CO].

[236] D. Giles and H. Feng, "Bias of the maximum likelihood estimators of the two-parameter gamma distribution revisited," *Department of Economics, University of Victoria, Econometrics Working Papers*, Jan. 2009.

[237] *Everything you always wanted to know about pulls*, CDF/ANAL/PUBLIC/5776, version Version 3.00, 2008.

[238] J. Hartlap, P. Simon, and P. Schneider, "Why your model parameter confidences might be too optimistic: Unbiased estimation of the inverse covariance matrix," *Astron. Astrophys.*, vol. 464, p. 399, 2007. DOI: 10.1051/0004-6361:20066170. arXiv: astro-ph/0608064.

[239] E. Sellentin and A. F. Heavens, "Parameter inference with estimated covariance matrices," *Mon. Not. Roy. Astron. Soc.*, vol. 456, no. 1, pp. L132–L136, 2016. DOI: 10.1093/mnrasl/slv190. arXiv: 1511.05969 [astro-ph.CO].

[240] G. Lazarides and Q. Shafi, "A Predictive inflationary scenario without the gauge singlet," *Phys. Lett. B*, vol. 308, pp. 17–22, 1993. DOI: 10.1016/0370-2693(93)90595-9. arXiv: hep-ph/9304247.

[241] M. S. Madsen and P. Coles, "CHAOTIC INFLATION," *Nucl. Phys. B*, vol. 298, pp. 701–725, 1988. DOI: 10.1016/0550-3213(88)90004-1.

[242] A. D. Linde, "CHAOTIC INFLATING UNIVERSE," *JETP Lett.*, vol. 38, pp. 176–179, 1983.

[243] L. Kofman, A. D. Linde, and A. A. Starobinsky, "Reheating after inflation," *Phys. Rev. Lett.*, vol. 73, pp. 3195–3198, 1994. DOI: 10.1103/PhysRevLett.73.3195. arXiv: hep-th/9405187.

[244] K. Nakayama and F. Takahashi, "Running Kinetic Inflation," *JCAP*, vol. 11, p. 009, 2010. DOI: 10.1088/1475-7516/2010/11/009. arXiv: 1008.2956 [hep-ph].

[245] F. Takahashi, "Linear Inflation from Running Kinetic Term in Supergravity," *Phys. Lett. B*, vol. 693, pp. 140–143, 2010. DOI: 10.1016/j.physletb.2010.08.029. arXiv: 1006.2801 [hep-ph].

[246] M. S. Turner, "Coherent Scalar Field Oscillations in an Expanding Universe," *Phys. Rev. D*, vol. 28, p. 1243, 1983. DOI: 10.1103/PhysRevD.28.1243.

[247] M. A. G. Garcia, K. Kaneta, Y. Mambrini, and K. A. Olive, "Inflaton Oscillations and Post-Inflationary Reheating," *JCAP*, vol. 04, p. 012, 2021. DOI: 10.1088/1475-7516/2021/04/012. arXiv: 2012.10756 [hep-ph].

[248] C. Ringeval, T. Suyama, and J. Yokoyama, "Magneto-reheating constraints from curvature perturbations," *JCAP*, vol. 09, p. 020, 2013. DOI: 10.1088/1475-7516/2013/09/020. arXiv: 1302.6013 [astro-ph.CO].

[249] J. Martin and C. Ringeval, "First CMB Constraints on the Inflationary Reheating Temperature," *Phys. Rev. D*, vol. 82, p. 023 511, 2010. DOI: 10.1103/PhysRevD.82.023511. arXiv: 1004.5525 [astro-ph.CO].

[250] J. Martin and C. Ringeval, "Inflation after WMAP3: Confronting the Slow-Roll and Exact Power Spectra to CMB Data," *JCAP*, vol. 08, p. 009, 2006. DOI: 10.1088/1475-7516/2006/08/009. arXiv: astro-ph/0605367.

[251] R. Easther and H. V. Peiris, "Bayesian Analysis of Inflation II: Model Selection and Constraints on Reheating," *Phys. Rev. D*, vol. 85, p. 103 533, 2012. DOI: 10.1103/PhysRevD.85.103533. arXiv: 1112.0326 [astro-ph.CO].

[252] P. C. W. Davies, S. A. Fulling, S. M. Christensen, and T. S. Bunch, "Energy Momentum Tensor of a Massless Scalar Quantum Field in a Robertson-Walker Universe," *Annals Phys.*, vol. 109, pp. 108–142, 1977. DOI: 10.1016/0003-4916(77)90167-1.

[253] T. S. Bunch and P. C. W. Davies, "Covariant Point Splitting Regularization for a Scalar Quantum Field in a Robertson-Walker Universe with Spatial Curvature," *Proc. Roy. Soc. Lond. A*, vol. 357, pp. 381–394, 1977. DOI: 10.1098/rspa.1977.0174.

[254] F. Bezrukov and M. Shaposhnikov, "Standard Model Higgs boson mass from inflation: Two loop analysis," *JHEP*, vol. 07, p. 089, 2009. DOI: 10.1088/1126-6708/2009/07/089. arXiv: 0904.1537 [hep-ph].

[255] F. L. Bezrukov, A. Magnin, and M. Shaposhnikov, "Standard Model Higgs boson mass from inflation," *Phys. Lett. B*, vol. 675, pp. 88–92, 2009. DOI: 10.1016/j.physletb.2009.03.035. arXiv: 0812.4950 [hep-ph].

[256] F. L. Bezrukov and M. Shaposhnikov, "The Standard Model Higgs boson as the inflaton," *Phys. Lett. B*, vol. 659, pp. 703–706, 2008. DOI: 10.1016/j.physletb.2007.11.072. arXiv: 0710.3755 [hep-th].

[257] J. L. F. Barbon and J. R. Espinosa, "On the Naturalness of Higgs Inflation," *Phys. Rev. D*, vol. 79, p. 081 302, 2009. DOI: 10.1103/PhysRevD.79.081302. arXiv: 0903.0355 [hep-ph].

[258] W. N. Cottingham and D. A. Greenwood, *An Introduction to the Standard Model of Particle Physics*. Cambridge University Press, Jul. 2023, ISBN: 978-1-009-40168-5, 978-1-009-40172-2, 978-1-009-40170-8, 978-0-511-27136-6, 978-0-521-85249-4. DOI: 10.1017/9781009401685.

[259] S. Navas *et al.*, "Review of Particle Physics," *to be published in Phys. Rev. D 110, 030001 (2024)*, vol. 2024, 2024.

[260] B. Pontecorvo, "Mesonium and anti-mesonium," *Sov. Phys. JETP*, vol. 6, p. 429, 1957.

[261] N. Arkani-Hamed, S. Dimopoulos, and G. R. Dvali, "The Hierarchy problem and new dimensions at a millimeter," *Phys. Lett. B*, vol. 429, pp. 263–272, 1998. DOI: 10.1016/S0370-2693(98)00466-3. arXiv: hep-ph/9803315.

[262] R. D. Peccei and H. R. Quinn, "CP Conservation in the Presence of Instantons," *Phys. Rev. Lett.*, vol. 38, pp. 1440–1443, 1977. DOI: 10.1103/PhysRevLett.38.1440.

[263] H. P. Nilles, "Supersymmetry, Supergravity and Particle Physics," *Phys. Rept.*, vol. 110, pp. 1–162, 1984. DOI: 10.1016/0370-1573(84)90008-5.

[264] S. P. Martin, "A Supersymmetry primer," *Adv. Ser. Direct. High Energy Phys.*, vol. 18, G. L. Kane, Ed., pp. 1–98, 1998. DOI: 10.1142/9789812839657_0001. arXiv: hep-ph/9709356.

[265] F. Buccella, J. P. Derendinger, S. Ferrara, and C. A. Savoy, "Patterns of Symmetry Breaking in Supersymmetric Gauge Theories," *Phys. Lett. B*, vol. 115, pp. 375–379, 1982. DOI: 10.1016/0370-2693(82)90521-4.

[266] I. Affleck, M. Dine, and N. Seiberg, "Dynamical Supersymmetry Breaking in Four-Dimensions and Its Phenomenological Implications," *Nucl. Phys. B*, vol. 256, pp. 557–599, 1985. DOI: 10.1016/0550-3213(85)90408-0.

[267] M. A. Luty and W. Taylor, "Varieties of vacua in classical supersymmetric gauge theories," *Phys. Rev. D*, vol. 53, pp. 3399–3405, 1996. DOI: 10.1103/PhysRevD.53.3399. arXiv: hep-th/9506098.

[268] T. Gherghetta, C. F. Kolda, and S. P. Martin, "Flat directions in the scalar potential of the supersymmetric standard model," *Nucl. Phys. B*, vol. 468, pp. 37–58, 1996. DOI: 10.1016/0550-3213(96)00095-8. arXiv: hep-ph/9510370.

[269] K. Enqvist and A. Mazumdar, "Cosmological consequences of MSSM flat directions," *Phys. Rept.*, vol. 380, pp. 99–234, 2003. DOI: 10.1016/S0370-1573(03)00119-4. arXiv: hep-ph/0209244.

[270] R. Allahverdi, K. Enqvist, J. Garcia-Bellido, and A. Mazumdar, "Gauge invariant MSSM inflaton," *Phys. Rev. Lett.*, vol. 97, p. 191 304, 2006. DOI: 10.1103/PhysRevLett.97.191304. arXiv: hep-ph/0605035.

[271] R. Allahverdi, K. Enqvist, J. Garcia-Bellido, A. Jokinen, and A. Mazumdar, "MSSM flat direction inflation: Slow roll, stability, fine-tuning and reheating," *JCAP*, vol. 06, p. 019, 2007. DOI: 10.1088/1475-7516/2007/06/019. arXiv: hep-ph/0610134.

[272] J. C. Bueno Sanchez, K. Dimopoulos, and D. H. Lyth, "A-term inflation and the MSSM," *JCAP*, vol. 01, p. 015, 2007. DOI: 10.1088/1475-7516/2007/01/015. arXiv: hep-ph/0608299.

[273] R. Allahverdi, B. Dutta, and A. Mazumdar, "Probing the parameter space for an MSSM inflation and the neutralino dark matter," *Phys. Rev. D*, vol. 75, p. 075 018, 2007. DOI: 10.1103/PhysRevD.75.075018. arXiv: hep-ph/0702112.

[274] R. Allahverdi, B. Dutta, and Y. Santoso, "MSSM inflation, dark matter, and the LHC," *Phys. Rev. D*, vol. 82, p. 035 012, 2010. DOI: 10.1103/PhysRevD.82.035012. arXiv: 1004.2741 [hep-ph].

[275] C. Boehm, J. Da Silva, A. Mazumdar, and E. Pukartas, "Probing the Supersymmetric Inflaton and Dark Matter link via the CMB, LHC and XENON1T experiments," *Phys. Rev. D*, vol. 87, no. 2, p. 023 529, 2013. DOI: 10.1103/PhysRevD.87.023529. arXiv: 1205.2815 [hep-ph].

[276] S. Choudhury and S. Pal, "Fourth level MSSM inflation from new flat directions," *JCAP*, vol. 04, p. 018, 2012. DOI: 10.1088/1475-7516/2012/04/018. arXiv: 1111.3441 [hep-ph].

[277] S. Choudhury, A. Mazumdar, and S. Pal, "Low & High scale MSSM inflation, gravitational waves and constraints from Planck," *JCAP*, vol. 07, p. 041, 2013. DOI: 10.1088/1475-7516/2013/07/041. arXiv: 1305.6398 [hep-ph].

[278] S. Choudhury, A. Mazumdar, and E. Pukartas, "Constraining $\mathcal{N} = 1$ supergravity inflationary framework with non-minimal Kähler operators," *JHEP*, vol. 04, p. 077, 2014. DOI: 10.1007/JHEP04(2014)077. arXiv: 1402.1227 [hep-th].

[279] M. N. Dubinin, E. Y. Petrova, E. O. Pozdeeva, and S. Y. Vernov, "MSSM inflation and cosmological attractors," *Int. J. Geom. Meth. Mod. Phys.*, vol. 15, p. 1 840 001, 2018. DOI: 10.1142/S0219887818400017. arXiv: 1712.03072 [hep-ph].

[280] M. N. Dubinin, E. Y. Petrova, E. O. Pozdeeva, M. V. Sumin, and S. Y. Vernov, "MSSM-inspired multifield inflation," *JHEP*, vol. 12, p. 036, 2017. DOI: 10.1007/JHEP12(2017)036. arXiv: 1705.09624 [hep-ph].

[281] R. Allahverdi, S. Downes, and B. Dutta, "Constructing Flat Inflationary Potentials in Supersymmetry," *Phys. Rev. D*, vol. 84, p. 101 301, 2011. DOI: 10.1103/PhysRevD.84.101301. arXiv: 1106.5004 [hep-th].

[282] S. Dimopoulos and H. Georgi, "Softly Broken Supersymmetry and SU(5)," *Nucl. Phys. B*, vol. 193, pp. 150–162, 1981. DOI: 10.1016/0550-3213(81)90522-8.

[283] S. Weinberg, "Supersymmetry at Ordinary Energies. 1. Masses and Conservation Laws," *Phys. Rev. D*, vol. 26, p. 287, 1982. DOI: 10.1103/PhysRevD.26.287.

[284] G. R. Farrar and P. Fayet, "Phenomenology of the Production, Decay, and Detection of New Hadronic States Associated with Supersymmetry," *Phys. Lett. B*, vol. 76, pp. 575–579, 1978. DOI: 10.1016/0370-2693(78)90858-4.

[285] S. Dimopoulos and D. W. Sutter, "The Supersymmetric flavor problem," *Nucl. Phys. B*, vol. 452, pp. 496–512, 1995. DOI: 10.1016/0550-3213(95)00421-N. arXiv: hep-ph/9504415.

[286] A. Djouadi *et al.*, "The Minimal supersymmetric standard model: Group summary report," in *GDR (Groupement De Recherche) - Supersymetrie*, Dec. 1998. arXiv: hep-ph/9901246.

[287] M. Ciuchini *et al.*, "Delta M(K) and epsilon(K) in SUSY at the next-to-leading order," *JHEP*, vol. 10, p. 008, 1998. DOI: 10.1088/1126-6708/1998/10/008. arXiv: hep-ph/9808328.

[288] J.-L. Kneur, G. Moultaka, M. Ughetto, D. Zerwas, and A. Djouadi, "SuSpect3: A C++ Code for the Supersymmetric and Higgs Particle Spectrum of the MSSM," Nov. 2022. arXiv: 2211.16956 [hep-ph].

[289] A. Djouadi, J.-L. Kneur, and G. Moultaka, "SuSpect: A Fortran code for the supersymmetric and Higgs particle spectrum in the MSSM," *Comput. Phys. Commun.*, vol. 176, pp. 426–455, 2007. DOI: 10.1016/j.cpc.2006.11.009. arXiv: hep-ph/0211331.

[290] H. Pagels and J. R. Primack, "Supersymmetry, Cosmology and New TeV Physics," *Phys. Rev. Lett.*, vol. 48, p. 223, 1982. DOI: 10.1103/PhysRevLett.48.223.

[291] J. Polonyi, *Generalization of the Massive Scalar Multiplet Coupling to the Supergravity*, 1977.

[292] G. Jungman, M. Kamionkowski, and K. Griest, "Supersymmetric dark matter," *Phys. Rept.*, vol. 267, pp. 195–373, 1996. DOI: 10.1016/0370-1573(95)00058-5. arXiv: hep-ph/9506380.

[293] T. Cohen and J. G. Wacker, "Here be Dragons: The Unexplored Continents of the CMSSM," *JHEP*, vol. 09, p. 061, 2013. DOI: 10.1007/JHEP09(2013)061. arXiv: 1305.2914 [hep-ph].

[294] S. Mizuta and M. Yamaguchi, "Coannihilation effects and relic abundance of Higgsino dominant LSP(s)," *Phys. Lett. B*, vol. 298, pp. 120–126, 1993. DOI: 10.1016/0370-2693(93)91717-2. arXiv: hep-ph/9208251.

[295] K. Griest and D. Seckel, "Three exceptions in the calculation of relic abundances," *Phys. Rev. D*, vol. 43, pp. 3191–3203, 1991. DOI: 10.1103/PhysRevD.43.3191.

[296] J. D. Barrow and P. Parsons, "Inflationary models with logarithmic potentials," *Phys. Rev. D*, vol. 52, pp. 5576–5587, 1995. DOI: 10.1103/PhysRevD.52.5576. arXiv: astro-ph/9506049.

[297] V. N. Senoguz and Q. Shafi, "Chaotic inflation, radiative corrections and precision cosmology," *Phys. Lett. B*, vol. 668, pp. 6–10, 2008. DOI: 10.1016/j.physletb.2008.08.017. arXiv: 0806.2798 [hep-ph].

[298] A. O. Barvinsky, A. Y. Kamenshchik, and A. A. Starobinsky, "Inflation scenario via the Standard Model Higgs boson and LHC," *JCAP*, vol. 11, p. 021, 2008. DOI: 10.1088/1475-7516/2008/11/021. arXiv: 0809.2104 [hep-ph].

[299] A. De Simone, M. P. Hertzberg, and F. Wilczek, "Running Inflation in the Standard Model," *Phys. Lett. B*, vol. 678, pp. 1–8, 2009. DOI: 10.1016/j.physletb.2009.05.054. arXiv: 0812.4946 [hep-ph].

[300] A. O. Barvinsky, A. Y. Kamenshchik, C. Kiefer, A. A. Starobinsky, and C. F. Steinwachs, "Higgs boson, renormalization group, and naturalness in cosmology," *Eur. Phys. J. C*, vol. 72, p. 2219, 2012. DOI: 10.1140/epjc/s10052-012-2219-3. arXiv: 0910.1041 [hep-ph].

[301] F. Bezrukov, A. Magnin, M. Shaposhnikov, and S. Sibiryakov, "Higgs inflation: consistency and generalisations," *JHEP*, vol. 01, p. 016, 2011. DOI: 10.1007/JHEP01(2011)016. arXiv: 1008.5157 [hep-ph].

[302] C. F. Steinwachs and A. Y. Kamenshchik, "Non-minimal Higgs Inflation and Frame Dependence in Cosmology," *AIP Conf. Proc.*, vol. 1514, no. 1, M. P. Dabrowski, A. Balcerzak, and T. Denkiewicz, Eds., pp. 161–164, 2013. DOI: 10.1063/1.4791748. arXiv: 1301.5543 [gr-qc].

[303] E. Elizalde, S. D. Odintsov, E. O. Pozdeeva, and S. Y. Vernov, "Renormalization-group improved inflationary scalar electrodynamics and SU(5) scenarios confronted with Planck 2013 and BICEP2 results," *Phys. Rev. D*, vol. 90, no. 8, p. 084001, 2014. DOI: 10.1103/PhysRevD.90.084001. arXiv: 1408.1285 [hep-th].

[304] D. P. George, S. Mooij, and M. Postma, "Quantum corrections in Higgs inflation: the Standard Model case," *JCAP*, vol. 04, p. 006, 2016. DOI: 10.1088/1475-7516/2016/04/006. arXiv: 1508.04660 [hep-th].

[305] V.-M. Enckell, K. Enqvist, S. Rasanen, and E. Tomberg, "Higgs inflation at the hilltop," *JCAP*, vol. 06, p. 005, 2018. DOI: 10.1088/1475-7516/2018/06/005. arXiv: 1802.09299 [astro-ph.CO].

[306] A. Brignole, L. E. Ibanez, and C. Munoz, "Soft supersymmetry breaking terms from supergravity and superstring models," *Adv. Ser. Direct. High Energy Phys.*, vol. 21, pp. 244–268, 2010. DOI: 10.1142/9789814307505_0004.

[307] K. Enqvist, D. G. Figueroa, and G. Rigopoulos, "Fluctuations along supersymmetric flat directions during Inflation," *JCAP*, vol. 01, p. 053, 2012. DOI: 10.1088/1475-7516/2012/01/053. arXiv: 1109.3024 [astro-ph.CO].

[308] S. Coleman and E. Weinberg, "Radiative corrections as the origin of spontaneous symmetry breaking," *Physical Review D*, vol. 7, no. 6, p. 1888, 1973.

[309] B. M. Kastening, "Renormalization group improvement of the effective potential in massive phi**4 theory," *Phys. Lett. B*, vol. 283, pp. 287–292, 1992. DOI: 10.1016/0370-2693(92)90021-U.

[310] M. Bando, T. Kugo, N. Maekawa, and H. Nakano, "Improving the effective potential," *Phys. Lett. B*, vol. 301, pp. 83–89, 1993. DOI: 10.1016/0370-2693(93)90725-W. arXiv: hep-ph/9210228.

[311] C. Ford, D. Jones, P. Stephenson, and M. Einhorn, "The Effective potential and the renormalization group," *Nucl. Phys. B*, vol. 395, pp. 17–34, 1993. DOI: 10.1016/0550-3213(93)90206-5. arXiv: hep-lat/9210033.

[312] E. J. Copeland, A. R. Liddle, D. H. Lyth, E. D. Stewart, and D. Wands, "False vacuum inflation with Einstein gravity," *Phys. Rev. D*, vol. 49, pp. 6410–6433, 1994. DOI: 10.1103/PhysRevD.49.6410. arXiv: astro-ph/9401011.

[313] D. Baumann, "Theoretical Advanced Study Institute in Elementary Particle Physics: Physics of the Large and the Small," 2011, pp. 523–686. DOI: 10.1142/9789814327183_0010. arXiv: 0907.5424 [hep-th].

[314] B. C. Allanach, A. Dedes, and H. K. Dreiner, "R parity violating minimal supergravity model," *Phys. Rev. D*, vol. 69, p. 115 002, 2004, [Erratum: Phys.Rev.D 72, 079902 (2005)]. DOI: 10.1103/PhysRevD.69.115002. arXiv: hep-ph/0309196.

[315] S. Antusch and M. Ratz, "Supergraph techniques and two loop beta functions for renormalizable and nonrenormalizable operators," *JHEP*, vol. 07, p. 059, 2002. DOI: 10.1088/1126-6708/2002/07/059. arXiv: hep-ph/0203027.

[316] D. J. Castano, E. J. Piard, and P. Ramond, "Renormalization group study of the Standard Model and its extensions. 2. The Minimal supersymmetric Standard Model," *Phys. Rev. D*, vol. 49, pp. 4882–4901, 1994. DOI: 10.1103/PhysRevD.49.4882. arXiv: hep-ph/9308335.

[317] S. P. Martin and M. T. Vaughn, "Two loop renormalization group equations for soft supersymmetry breaking couplings," *Phys. Rev. D*, vol. 50, p. 2282, 1994, [Erratum: Phys.Rev.D 78, 039903 (2008)]. DOI: 10.1103/PhysRevD.50.2282. arXiv: hep-ph/9311340.

[318] K. Enqvist, A. Mazumdar, and P. Stephens, "Inflection point inflation within supersymmetry," *JCAP*, vol. 06, p. 020, 2010. DOI: 10.1088/1475-7516/2010/06/020. arXiv: 1004.3724 [hep-ph].

[319] R. Allahverdi, A. Ferrantelli, J. Garcia-Bellido, and A. Mazumdar, "Non-perturbative production of matter and rapid thermalization after MSSM inflation," *Phys. Rev. D*, vol. 83, p. 123 507, 2011. DOI: 10.1103/PhysRevD.83.123507. arXiv: 1103.2123 [hep-ph].

[320] A. Ferrantelli, "Reheating, thermalization and non-thermal gravitino production in MSSM inflation," *Eur. Phys. J. C*, vol. 77, no. 10, p. 716, 2017. DOI: 10.1140/epjc/s10052-017-5269-8. arXiv: 1702.01051 [hep-ph].

[321] L. Husdal, "On Effective Degrees of Freedom in the Early Universe," *Galaxies*, vol. 4, no. 4, p. 78, 2016. DOI: 10.3390/galaxies4040078. arXiv: 1609.04979 [astro-ph.CO].

[322] C. Adam, J.-L. Kneur, R. Lafaye, T. Plehn, M. Rauch, and D. Zerwas, "Measuring Unification," *Eur. Phys. J. C*, vol. 71, p. 1520, 2011. DOI: 10.1140/epjc/s10052-010-1520-2. arXiv: 1007.2190 [hep-ph].

[323] P. A. R. Ade *et al.*, "Planck 2013 results. XVI. Cosmological parameters," *Astron. Astrophys.*, vol. 571, A16, 2014. DOI: 10.1051/0004-6361/201321591. arXiv: 1303.5076 [astro-ph.CO].

[324] S. Henrot-Versillé *et al.*, "Constraining Supersymmetry using the relic density and the Higgs boson," *Phys. Rev. D*, vol. 89, no. 5, p. 055 017, 2014. DOI: 10.1103/PhysRevD.89.055017. arXiv: 1309.6958 [hep-ph].

[325] G. Domènech, "Scalar Induced Gravitational Waves Review," *Universe*, vol. 7, no. 11, p. 398, 2021. DOI: 10.3390/universe7110398. arXiv: 2109.01398 [gr-qc].

[326] R. L. Workman *et al.*, "Review of Particle Physics," *PTEP*, vol. 2022, p. 083C01, 2022. DOI: 10.1093/ptep/ptac097.

[327] G. Aad *et al.*, "Search for the Standard Model Higgs boson in the decay channel $H \to ZZ^{(*)} \to 4\ell$ with 4.8 fb^{-1} of pp collision data at $\sqrt{s} = 7$ TeV with ATLAS," *Phys. Lett. B*, vol. 710, pp. 383–402, 2012. DOI: 10.1016/j.physletb.2012.03.005. arXiv: 1202.1415 [hep-ex].

[328] G. Aad *et al.*, "Search for the Standard Model Higgs boson in the diphoton decay channel with 4.9 fb^{-1} of pp collisions at $\sqrt{s} = 7$ TeV with ATLAS," *Phys. Rev. Lett.*, vol. 108, p. 111 803, 2012. DOI: 10.1103/PhysRevLett.108.111803. arXiv: 1202.1414 [hep-ex].

[329] G. Aad *et al.*, "Search for the Standard Model Higgs boson in the $H \to$ WW(*) $\to \ell\nu\ell\nu$ decay mode with 4.7 /fb of ATLAS data at $\sqrt{s} = 7$ TeV," *Phys. Lett. B*, vol. 716, pp. 62–81, 2012. DOI: 10.1016/j.physletb.2012.08.010. arXiv: 1206.0756 [hep-ex].

[330] G. Aad *et al.*, "Search for the Standard Model Higgs boson in the H to $\tau^+\tau^-$ decay mode in $\sqrt{s} = 7$ TeV pp collisions with ATLAS," *JHEP*, vol. 09, p. 070, 2012. DOI: 10.1007/JHEP09(2012)070. arXiv: 1206.5971 [hep-ex].

[331] G. Aad *et al.*, "Combined search for the Standard Model Higgs boson in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector," *Phys. Rev. D*, vol. 86, p. 032 003, 2012. DOI: 10.1103/PhysRevD.86.032003. arXiv: 1207.0319 [hep-ex].

[332] G. Aad *et al.*, "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC," *Phys. Lett. B*, vol. 716, pp. 1–29, 2012. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 [hep-ex].

[333] S. Chatrchyan *et al.*, "Search for the standard model Higgs boson in the decay channel H to ZZ to 4 leptons in pp collisions at $\sqrt{s} = 7$ TeV," *Phys. Rev. Lett.*, vol. 108, p. 111 804, 2012. DOI: 10.1103/PhysRevLett.108.111804. arXiv: 1202.1997 [hep-ex].

[334] S. Chatrchyan *et al.*, "Search for the standard model Higgs boson decaying to W^+W in the fully leptonic final state in pp collisions at $\sqrt{s} = 7$ TeV," *Phys. Lett. B*, vol. 710, pp. 91–113, 2012. DOI: 10.1016/j.physletb.2012.02.076. arXiv: 1202.1489 [hep-ex].

[335] S. Chatrchyan *et al.*, "Search for neutral Higgs bosons decaying to tau pairs in pp collisions at $\sqrt{s} = 7$ TeV," *Phys. Lett. B*, vol. 713, pp. 68–90, 2012. DOI: 10.1016/j.physletb.2012.05.028. arXiv: 1202.4083 [hep-ex].

[336] S. Schael *et al.*, "Precision electroweak measurements on the Z resonance," *Phys. Rept.*, vol. 427, pp. 257–454, 2006. DOI: 10.1016/j.physrep.2005.12.006. arXiv: hep-ex/0509008.

[337] E. Aprile *et al.*, "Light Dark Matter Search with Ionization Signals in XENON1T," *Phys. Rev. Lett.*, vol. 123, no. 25, p. 251 801, 2019. DOI: 10.1103/PhysRevLett.123.251801. arXiv: 1907.11485 [hep-ex].

[338] J. Ellis and K. A. Olive, "Revisiting the Higgs Mass and Dark Matter in the CMSSM," *Eur. Phys. J. C*, vol. 72, p. 2005, 2012. DOI: 10.1140/epjc/s10052-012-2005-2. arXiv: 1202.3262 [hep-ph].

[339] R. Lafaye, T. Plehn, and D. Zerwas, "SFITTER: SUSY parameter analysis at LHC and LC," Apr. 2004. arXiv: hep-ph/0404282.

[340] F. James and M. Roos, "Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations," *Comput. Phys. Commun.*, vol. 10, pp. 343–367, 1975. DOI: 10.1016/0010-4655(75)90039-9.

[341] G. Belanger, F. Boudjema, A. Pukhov, and A. Semenov, "micrOMEGAs_3: A program for calculating dark matter observables," *Comput. Phys. Commun.*, vol. 185, pp. 960–985, 2014. DOI: 10.1016/j.cpc.2013.10.016. arXiv: 1305.0237 [hep-ph].

[342] S. Heinemeyer, W. Hollik, A. Weber, and G. Weiglein, "Z pole observables in the mssm," *Journal of High Energy Physics*, vol. 2008, no. 04, p. 039, 2008. DOI: 10.1088/1126-6708/2008/04/039.

[343] A. Djouadi, J. Kalinowski, M. Muehlleitner, and M. Spira, "HDECAY: Twenty$_{++}$ years after," *Comput. Phys. Commun.*, vol. 238, pp. 214–231, 2019. DOI: 10.1016/j.cpc.2018.12.010. arXiv: 1801.09506 [hep-ph].

[344] G. N. Remmen and S. M. Carroll, "Attractor Solutions in Scalar-Field Cosmology," *Phys. Rev. D*, vol. 88, p. 083 518, 2013. DOI: 10.1103/PhysRevD.88.083518. arXiv: 1309.2611 [gr-qc].

[345] D. Chowdhury, J. Martin, C. Ringeval, and V. Vennin, "Assessing the scientific status of inflation after Planck," *Phys. Rev. D*, vol. 100, no. 8, p. 083 537, 2019. DOI: 10.1103/PhysRevD.100.083537. arXiv: 1902.03951 [astro-ph.CO].

[346] R. Allahverdi, B. Dutta, and A. Mazumdar, "Attraction towards an inflection point inflation," *Phys. Rev. D*, vol. 78, p. 063 507, 2008. DOI: 10.1103/PhysRevD.78.063507. arXiv: 0806.4557 [hep-ph].

[347] S. Kraml *et al.*, "SModelS: a tool for interpreting simplified-model results from the LHC and its application to supersymmetry," *Eur. Phys. J. C*, vol. 74, p. 2868, 2014. DOI: 10.1140/epjc/s10052-014-2868-5. arXiv: 1312.4175 [hep-ph].

[348] F. Ambrogi *et al.*, "SModelS v1.1 user manual: Improving simplified model constraints with efficiency maps," *Comput. Phys. Commun.*, vol. 227, pp. 72–98, 2018. DOI: 10.1016/j.cpc.2018.02.007. arXiv: 1701.06586 [hep-ph].

[349] F. Ambrogi *et al.*, "SModelS v1.2: long-lived particles, combination of signal regions, and other novelties," *Comput. Phys. Commun.*, vol. 251, p. 106 848, 2020. DOI: 10.1016/j.cpc.2019.07.013. arXiv: 1811.10624 [hep-ph].

[350] G. Alguero, S. Kraml, and W. Waltenberger, "A SModelS interface for pyhf likelihoods," *Comput. Phys. Commun.*, vol. 264, p. 107 909, 2021. DOI: 10.1016/j.cpc.2021.107909. arXiv: 2009.01809 [hep-ph].

[351] G. Alguero *et al.*, "Constraining new physics with SModelS version 2," *JHEP*, vol. 08, p. 068, 2022. DOI: 10.1007/JHEP08(2022)068. arXiv: 2112.00769 [hep-ph].

[352] A. Buckley, "PySLHA: a Pythonic interface to SUSY Les Houches Accord data," *Eur. Phys. J. C*, vol. 75, no. 10, p. 467, 2015. DOI: 10.1140/epjc/s10052-015-3638-8. arXiv: 1305.4194 [hep-ph].

[353] T. Sjostrand, S. Mrenna, and P. Z. Skands, "PYTHIA 6.4 Physics and Manual," *JHEP*, vol. 05, p. 026, 2006. DOI: 10.1088/1126-6708/2006/05/026. arXiv: hep-ph/0603175.

[354] W. Beenakker, R. Hopker, M. Spira, and P. M. Zerwas, "Squark and gluino production at hadron colliders," *Nucl. Phys. B*, vol. 492, pp. 51–103, 1997. DOI: 10.1016/S0550-3213(97)80027-2. arXiv: hep-ph/9610490.

[355] W. Beenakker, M. Kramer, T. Plehn, M. Spira, and P. M. Zerwas, "Stop production at hadron colliders," *Nucl. Phys. B*, vol. 515, pp. 3–14, 1998. DOI: 10.1016/S0550-3213(98)00014-5. arXiv: hep-ph/9710451.

[356] A. Kulesza and L. Motyka, "Threshold resummation for squark-antisquark and gluino-pair production at the LHC," *Phys. Rev. Lett.*, vol. 102, p. 111 802, 2009. DOI: 10.1103/PhysRevLett.102.111802. arXiv: 0807.2405 [hep-ph].

[357] A. Kulesza and L. Motyka, "Soft gluon resummation for the production of gluino-gluino and squark-antisquark pairs at the LHC," *Phys. Rev. D*, vol. 80, p. 095 004, 2009. DOI: 10.1103/PhysRevD.80.095004. arXiv: 0905.4749 [hep-ph].

[358] W. Beenakker, S. Brensing, M. Kramer, A. Kulesza, E. Laenen, and I. Niessen, "Soft-gluon resummation for squark and gluino hadroproduction," *JHEP*, vol. 12, p. 041, 2009. DOI: 10.1088/1126-6708/2009/12/041. arXiv: 0909.4418 [hep-ph].

[359] W. Beenakker, S. Brensing, M. Kramer, A. Kulesza, E. Laenen, and I. Niessen, "Supersymmetric top and bottom squark production at hadron colliders," *JHEP*, vol. 08, p. 098, 2010. DOI: 10.1007/JHEP08(2010)098. arXiv: 1006.4771 [hep-ph].

[360] W. Beenakker *et al.*, "Squark and Gluino Hadroproduction," *Int. J. Mod. Phys. A*, vol. 26, pp. 2637–2664, 2011. DOI: 10.1142/S0217751X11053560. arXiv: 1105.1110 [hep-ph].

[361] J. Dutta, S. Kraml, A. Lessa, and W. Waltenberger, "SModelS extension with the CMS supersymmetry search results from Run 2," *LHEP*, vol. 1, no. 1, pp. 5–12, 2018. DOI: 10.31526/LHEP.1.2018.02. arXiv: 1803.02204 [hep-ph].

[362] G. Aad *et al.*, "Summary of the ATLAS experiment's sensitivity to supersymmetry after LHC Run 1 — interpreted in the phenomenological MSSM," *JHEP*, vol. 10, p. 134, 2015. DOI: 10.1007/JHEP10(2015)134. arXiv: 1508.06608 [hep-ex].

[363] M. Kawasaki, K. Kohri, and T. Moroi, "Big-Bang nucleosynthesis and hadronic decay of long-lived massive particles," *Phys. Rev. D*, vol. 71, p. 083 502, 2005. DOI: 10.1103/PhysRevD.71.083502. arXiv: astro-ph/0408426.

[364] M. Kawasaki, K. Kohri, T. Moroi, and Y. Takaesu, "Revisiting Big-Bang Nucleosynthesis Constraints on Long-Lived Decaying Particles," *Phys. Rev. D*, vol. 97, no. 2, p. 023 502, 2018. DOI: 10.1103/PhysRevD.97.023502. arXiv: 1709.01211 [hep-ph].

[365] J. Martin, C. Ringeval, and V. Vennin, "Information Gain on Reheating: the One Bit Milestone," *Phys. Rev. D*, vol. 93, no. 10, p. 103 532, 2016. DOI: 10.1103/PhysRevD.93.103532. arXiv: 1603.02606 [astro-ph.CO].

[366] J. Martin, C. Ringeval, and V. Vennin, "Observing Inflationary Reheating," *Phys. Rev. Lett.*, vol. 114, no. 8, p. 081 303, 2015. DOI: 10.1103/PhysRevLett.114.081303. arXiv: 1410.7958 [astro-ph.CO].

[367] M. E. Jones *et al.*, "The C-Band All-Sky Survey (C-BASS): Design and capabilities," *Mon. Not. Roy. Astron. Soc.*, vol. 480, no. 3, pp. 3224–3242, 2018. DOI: 10.1093/mnras/sty1956. arXiv: 1805.04490 [astro-ph.IM].

[368] N. Krachmalnicoff *et al.*, "S–PASS view of polarized Galactic synchrotron at 2.3 GHz as a contaminant to CMB observations," *Astron. Astrophys.*, vol. 618, A166, 2018. DOI: 10.1051/0004-6361/201832768. arXiv: 1802.01145 [astro-ph.GA].

[369] R. Genova-Santos *et al.*, "The QUIJOTE experiment: project overview and first results," in *11th Scientific Meeting of the Spanish Astronomical Society*, Apr. 2015. arXiv: 1504.03514 [astro-ph.CO].

[370] J. Martin, C. Ringeval, and V. Vennin, "Vanilla Inflation Predicts Negative Running," Apr. 2024. arXiv: 2404.15089 [astro-ph.CO].

université PARIS-SACLAY

ÉCOLE DOCTORALE
Particules, hadrons, énergie et noyau : instrumentation, imagerie, cosmos et simulation (PHENIICS)

Title: Inflation: phenomenological study and preparation of the LiteBIRD space mission
Keywords: Cosmology, Inflation, LiteBIRD, Phenomenology, Cosmic microwave background, Supersymmetry

Abstract: This thesis is devoted to the study of cosmological inflation, a phase of accelerated expansion in the early universe that remains speculative to this day. The central observable for this study is the cosmic microwave background (CMB), the oldest light still visible today, whose statistical study enables cosmological inference.

We first approach the study from an experimental perspective, focusing on the preparation of the LiteBIRD satellite. Set to launch in the middle of the next decade, LiteBIRD will measure the large-scale polarisation of the CMB with unprecedented precision, allowing for stringent constraints on the presence of primordial gravitational waves generated during inflation. To achieve the required sensitivity and minimise systematic effects, we must ensure precise control of both the instrument and data analysis. As part of this effort, we have implemented the instrument model in a dedicated database, along with the tools necessary to produce key instrumental parameters. This includes generating quaternions that encode each detector's pointing and orientation information, as well as implementing beam models, bandpasses, the noise model, and the specification of the readout system.

Furthermore, we have developed a complete pipeline for analysing the polarisation maps that LiteBIRD will deliver. We have tested this pipeline on realistic simulations of the instrument with various levels of complexity. The analysis pipeline consists of three stages. The first stage involves component separation to remove foreground contamination from the maps. We optimise an agnostic method that does not rely on prior knowledge of the foreground properties. The second stage focuses on estimating power spectra from the cleaned and masked maps. To this end, we have implemented and tested various unbiased and quasi-optimal methods. Finally, we assess the performance of different likelihood functions to infer cosmological parameters. In addition to constraining primordial gravitational waves, this analysis will enhance our understanding of the epoch of reionisation, which is due to the intense radiation from the first generation of stars.

In the third section of the thesis, we focus on a phenomenological study of inflation, particularly on a model of inflation situated within a particle physics framework: the minimal supersymmetric model. In collaboration with cosmologists, theorists, and particle physicists, we demonstrate that the existing data from the Planck satellite are already precise enough that systematic errors in the model's predictions dominate the error budget in an inference context. These theoretical systematics arise from the non-inclusion of radiative corrections and an incomplete understanding of the end of inflation. We have included the necessary corrections and identified points in parameter space that satisfy both the observational constraints of particle physics (such as the Higgs mass and direct SUSY searches at the LHC) and cosmology (including the dark matter fraction in the universe and the properties of scalar perturbations as observed by Planck). Our work demonstrates the feasibility of unifying particle physics and cosmology descriptions within a single self-consistent model, paving the way for a comprehensive exploration of the inflationary MSSM or other high-energy physics models.