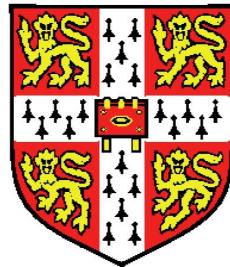


The Non-Gaussian Universe

Large-scale structure non-Gaussianities with modal methods,

and

joint analysis of CMB temperature and
lensing-reconstruction power spectra



Marcel M. Schmittfull

Department of Applied Mathematics and Theoretical Physics

St John's College, University of Cambridge

A thesis submitted for the degree of

Doctor of Philosophy

August 2013

Abstract

The Λ CDM model with an initial inflation phase has been confirmed with ever higher precision by recent cosmological observations. To constrain the underlying mechanism leading to inflation and to add to our understanding of dark energy, neutrino masses and spatial curvature, large-scale structure surveys and CMB lensing reconstructions are becoming increasingly important. This thesis presents the implementation of novel fast and general methods to study the non-Gaussian imprints of inflation on large-scale structures. Additionally, it provides an analysis of the question of whether the information from CMB lensing reconstruction using the CMB temperature 4-point function is independent from the lensing information contained already in the CMB power spectrum.

Based on a separable modal expansion of the (appropriately weighted) bispectrum and trispectrum, we present an implementation for generating general non-Gaussian N -body initial conditions. The method works for arbitrary bispectra and a special class of trispectra and is several orders of magnitude faster than previous approaches. Using a similar modal expansion, we present an implementation of a bispectrum estimator for 3D large-scale structures, which is again orders-of-magnitude faster than the usually adopted approach. We apply this framework to track the dark matter bispectrum in a large suite of Gaussian and non-Gaussian N -body simulations. We describe the effect of primordial non-Gaussianity with a simple time-shift model and generate new fitting formulae for gravitational and primordial contributions to the dark matter bispectrum.

Gravitational lensing affects both the 2- and 4-point functions of the lensed CMB temperature anisotropies (the latter corresponding to the lensing reconstruction power spectrum). These two observables are often treated as independent, although this is impossible since they are derived from the same CMB map. We analyse their covariance, finding two dominant contributions: (i) correlations between the disconnected noise bias in the trispectrum measurement and sample variance in the temperature power spectrum; and (ii) sample variance of the lenses themselves. We show that they can be ignored (or easily mitigated) for a Planck-like experiment, justifying treating the two observable spectra as independent. We also test simple likelihood approximations for the deflection power spectrum, finding that a Gaussian with a parameter-independent covariance performs well.

Supervisors and Examiners

First Supervisor: Prof. Paul Shellard

Second Supervisor: Dr. Anthony Challinor

External Examiner: Prof. Benjamin Wandelt

Internal Examiner: Dr. Daniel Baumann

Acknowledgements

[Adopted for online version of this thesis]

I am grateful that I had the opportunity to meet and interact with many fantastic people who have helped me in many ways to write this thesis and develop the research projects. In particular I would like to thank Paul Shellard and Anthony Challinor who were fantastic supervisors. I also thank James Fergusson, Duncan Hanson, Antony Lewis, Donough Regan, Daniel Baumann, Hiro Funakoshi, Uros Seljak, Helge Gruetjen, Simon Su, Camille Bonvin, Xingang Chen, Alex Hall, Eugene Lim, Sebastien Renaux-Petel, Jeremy Sakstein, Hadi Godazgar, Andrei Lazanu, Sean Lip, Andrey Kaliazin, Gabi Koren, Peter Krahmer, Herbert Bauer, Gudrun Nowak, Irene Kaindl, Iris Schreiner, Arnold Ruemmele, Hugo Reinhardt, Franz Pedit, Willy Kley, David Wharam, Matthias Bartelmann, Arthur Hebecker, Raquel Ribeiro. I am also grateful for support from DAAD, DAMTP Cambridge, Siegfried and Marlene Weiss, St John's College Cambridge, STFC and Studiens-tiftung des deutschen Volkes. Finally, very special thanks and warm hugs to Terry Farrelly, Jeng Rongchai, my parents Martin and Ingrid, my sisters Janine and Kristina, Jens and Simone Walther, and my wife Anett.

Contents

Contents	vii
1 Introduction	1
2 Background: Cosmological Perturbations	7
2.1 Primordial perturbations	8
2.1.1 Inflation	8
2.1.2 Single-field slow-roll inflation	9
2.1.3 Horizon crossings	10
2.1.4 Quantisation of a light scalar field during slow-roll inflation .	10
2.1.5 From quantum to classical fluctuations	14
2.1.6 Primordial power spectrum	14
2.1.7 Primordial non-Gaussianity	15
2.1.7.1 Order of magnitude estimate for single-field slow roll inflation	16
2.1.7.2 Primordial bispectrum and trispectrum definitions	18
2.1.7.3 Primordial bispectrum shapes	21
2.1.7.4 Estimating non-Gaussianity	23
2.2 CMB perturbations	27
2.2.1 CMB anisotropies from metric perturbations	28
2.2.2 Scattering of photons and electrons	31
2.2.3 Boltzmann hierarchy	32
2.2.4 CMB power spectrum	35
2.3 CMB lensing	37
2.3.1 Basics	38
2.3.1.1 CMB lensing potential	38
2.3.1.2 Lensing power spectrum	39
2.3.1.3 Lensing-temperature correlation from ISW-effect .	40

2.3.1.4	Lensed CMB temperature	40
2.3.1.5	Effect of lensing on the CMB temperature power spectrum	42
2.3.1.6	CMB lensing reconstruction	45
2.3.2	Observational status	48
2.3.3	Cosmological information probed by CMB lensing	49
2.3.3.1	Geometric degeneracy of the primary CMB	49
2.3.3.2	Breaking the geometric degeneracy: Evidence for dark energy and flatness from the CMB alone	51
2.3.3.3	Neutrino mass, optical depth and LSS bias	52
2.3.4	Cosmological parameter estimation with CMB lensing	53
2.4	Large-scale dark matter perturbations	56
2.4.1	Evolution of perturbations during radiation and matter domination	56
2.4.2	Transfer function	57
2.4.3	Growth function	58
2.4.4	Higher-order perturbation theory	59
2.4.5	Gravitational bispectrum	63
2.4.6	Zel'dovich approximation and N -body initial conditions	65
2.4.7	N -body simulations	66
2.4.8	Large-scale structure surveys	68
3	Non-Gaussian N-body initial conditions for arbitrary primordial bispectra	71
3.1	Introduction	72
3.2	Algorithm	73
3.2.1	Bispectrum	73
3.2.2	Trispectrum	76
3.3	Algorithm validation	77
3.3.1	Bispectrum estimation	78
3.4	Results	80
3.5	Conclusions	81

4	Fast estimation of gravitational and primordial bispectra in large-scale structures	85
4.1	Introduction	86
4.2	The distribution of matter	90
4.2.1	Power spectrum and transfer functions	90
4.2.2	Matter bispectrum	91
4.2.3	Tree-level gravitational matter bispectrum	92
4.2.4	Gravitational matter bispectrum beyond tree level	95
4.2.4.1	Loop corrections	95
4.2.4.2	Halo models	96
4.2.4.3	Constant ‘halo’ model	97
4.2.4.4	Alternative phenomenological fit	99
4.3	Primordial non-Gaussianity	100
4.3.1	Primordial contribution to the matter bispectrum	100
4.3.2	Non-Gaussianity as a halo model time-shift	102
4.4	Bispectrum estimation methodology	105
4.4.1	f_{NL} estimator	105
4.4.2	Shape and size comparisons	106
4.4.3	Separable mode expansion for fast f_{NL} estimation	107
4.4.4	Fast modal bispectrum estimator	109
4.4.5	Fast modal bispectrum correlations	111
4.4.6	Cumulative measures of non-Gaussianity	112
4.4.7	Towards experimental setups	112
4.4.8	Comparison to other bispectrum estimators	113
4.4.9	Bispectrum visualisation	114
4.5	Simulation setup, initial conditions and validation	116
4.5.1	N -body simulations setup	116
4.5.2	Regular grid vs glass initial conditions	119
4.5.3	Validation and convergence tests	120
4.5.4	Error bars	122
4.5.5	Discussion in light of previous analyses	124
4.6	Gravitational bispectrum results	125
4.6.1	Gravitational collapse and bispectrum evolution	125

4.6.1.1	Comparison with leading order PT	127
4.6.2	Fitting formulae for Gaussian simulations	129
4.6.2.1	Separable polynomial expansion	129
4.6.2.2	Time-shift model fit	134
4.6.2.3	Alternative phenomenological fit	138
4.6.2.4	Halo model	139
4.7	Primordial non-Gaussian bispectrum results	140
4.7.1	Primordial bispectrum measurements	140
4.7.1.1	Local shape	140
4.7.1.2	Other shapes	140
4.7.1.3	Loop corrections for primordial non-Gaussianity .	144
4.7.1.4	Linearity in input f_{NL}	144
4.7.2	Fitting formulae for non-Gaussian simulations	146
4.7.2.1	Separable polynomial fits	146
4.7.2.2	Time-shift model fits	146
4.7.2.3	Redshift at which 1-halo contribution becomes im- portant	151
4.8	Summary and conclusions	151
5	On the joint analysis of CMB temperature and lensing-reconstruction power spectra	155
5.1	Full-sky CMB lensing reconstruction	156
5.2	Simulations	162
5.3	Auto-correlations of power spectra	163
5.3.1	Lensed temperature	166
5.3.2	Lensing reconstruction	168
5.4	Temperature-lensing cross-correlation	170
5.4.1	Noise contribution	171
5.4.1.1	Perturbative derivation: Disconnected part of lensed temperature 4- and 6-point functions	171
5.4.1.2	Magnitude and structure of the correlation matrix	172
5.4.1.3	Comparison with simulations	175

5.4.1.4	Mitigating the noise contribution with the empirical $\hat{N}^{(0)}$ bias correction	176
5.4.2	Matter cosmic variance contribution	177
5.4.2.1	Warm-up: Power covariance of input lensing potential and lensed temperature	177
5.4.2.2	Power covariance of reconstructed lenses and lensed temperature from the connected 6-point function .	178
5.4.2.3	Magnitude and structure of the covariance matrix .	179
5.4.2.4	Comparison with simulations	180
5.4.2.5	Mitigating the matter cosmic variance contribution	183
5.4.3	Towards a complete model for power covariances	184
5.4.4	Impact of correlations on parameter estimation	185
5.4.4.1	Lensing amplitude estimates	185
5.4.4.2	Cosmological parameters	191
5.5	Towards a lensing likelihood	192
5.5.1	Lensing amplitude from lensing reconstruction	192
5.5.2	Two-parameter likelihood tests with lensing amplitude and lensing tilt	195
5.6	Conclusions	199
5.A	Why are the lensing amplitude cross-correlations so small?	202
5.A.1	Cosmic variance of the CMB	202
5.A.2	Cosmic variance of the lenses	205
5.B	Optimal trispectrum estimation and $N^{(0)}$ subtraction	208
5.C	$\hat{N}^{(0)}$ subtraction as efficient mitigation of chance alignments	210
5.D	Temperature-lensing power covariance from the CMB trispectrum .	214
5.E	Temperature-lensing power covariance from the connected CMB 6-point function	217
5.F	Effect of $C^{T\phi}$	219
6	Summary and outlook	223
References		227

Introduction

Non-Gaussian Relics from the Big Bang

Physical cosmology is now in exciting times because a simple model continues to successfully describe the increasingly high precision observational data even though this model is based on a big puzzle: The universe had to go through an *inflationary* phase of rapid accelerated expansion after the big bang, but the underlying physical mechanism is unknown. By confronting theoretical predictions with observations, cosmologists aim to reduce the number of viable inflation theories or even pin down the correct model, which will ultimately test ultra-high-energy physics beyond reach of any earth-bound accelerator.

During the inflationary expansion of the universe, tiny quantum fluctuations were “stretched” to become seeds for the formation of galaxies. Therefore information about inflation is imprinted in the distribution of galaxies, in particular in deviations from Gaussian statistics. This *primordial non-Gaussianity* (PNG) is directly related to interactions between fields during inflation. Galaxy catalogues will be the primary observational window for PNG over the next decade. However, there is no complete model against which observations can be compared, and optimal statistical methods to extract PNG have not been developed. A substantial part of this thesis aims to make progress with both of these problems.

We present the implementation of numerical methods to study the formation of cosmological structures for physical models of the inflationary expansion in Chapters 3 and 4. Such a numerical approach is crucial because non-linear physics relevant for upcoming observations cannot be modelled analytically. The method presented to create general non-Gaussian initial conditions is the fastest and most general currently available and it provides the only way to simulate arbitrary physical inflation models for realistic simulation sizes. We also develop a method to extract efficiently PNG from such simulations by measuring the full 3-point

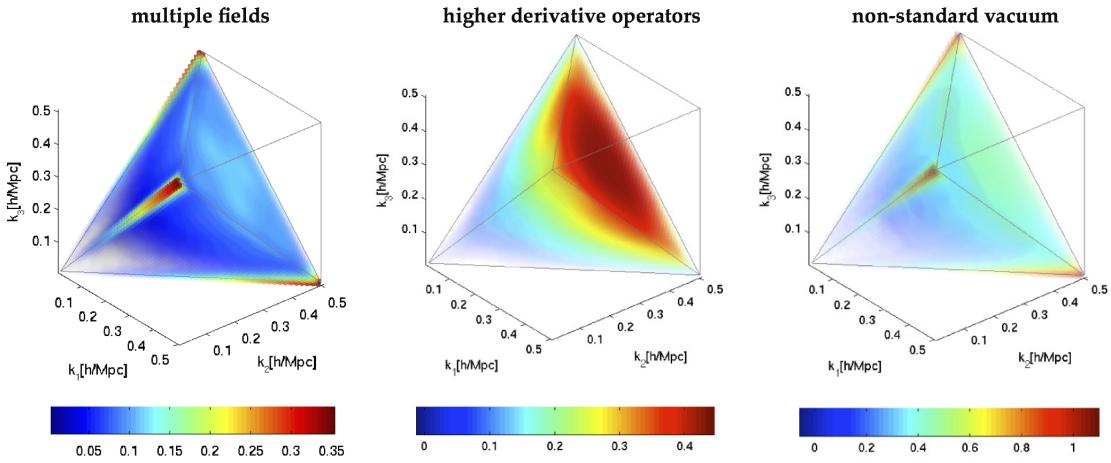


Figure 1.1: Deviations of the dark matter distribution from Gaussian statistics for three different classes of inflation models as measured for the first time in our simulations. The pictures show the excess of the (Fourier-transformed) 3-point correlation function over the simplest model of inflation as a function of wavenumbers (inverse length scales) k_1 , k_2 and k_3 . See Chapter 4 for details.

correlation function. While it is known that this has better sensitivity to PNG than the typically studied 2-point correlation function, the method presented here is for the first time sufficiently fast to study the complete 3-point function in high-resolution simulations and measure it in real data. The 3-point correlation function of the dark matter distribution is shown in Figure 1.1 for various inflation models. A phenomenological time-shift model is developed in Chapter 4 to describe the simulation results. This is used to create fitting formulae for the dark matter bispectrum in presence of PNG which should be useful for constraining PNG with galaxy surveys.

Gravitational Lensing of the Cosmic Microwave Background

The second part of this thesis concerns an emerging probe which is adding to our understanding of the geometry of the universe, neutrino masses and the late-time expansion of the universe (dark energy): Lensing of the cosmic microwave background (CMB). Photons from the big bang have been traveling essentially freely for more than 13 billion years to reach Earth today in the form of the CMB. Along their way these photons are deflected by the inhomogeneous distribution of

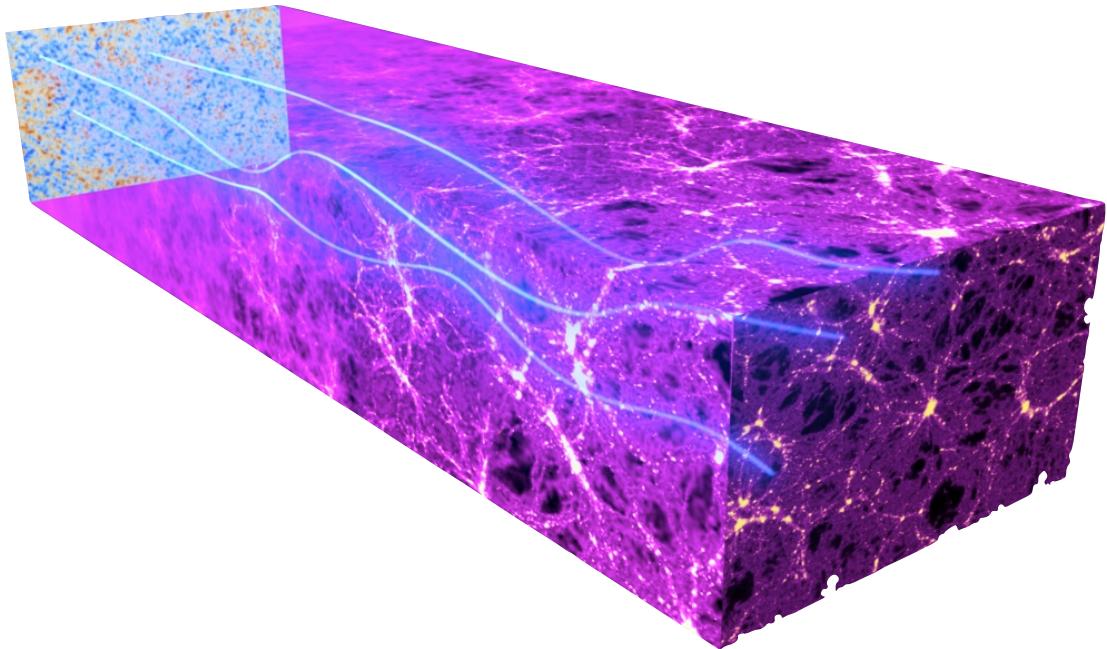


Figure 1.2: Illustration of gravitational lensing of CMB photons by large-scale structure (adapted from [1]).

matter, which acts as a collection of gravitational lenses (see Fig. 1.2). By reconstructing these lenses from the observed CMB we can obtain crucial information on e.g. dark energy or the geometry of the universe. Several experiments will provide high quality CMB data in the next few years, e.g. full-mission Planck data including polarization, ACT/ACTPol, Polarbear and SPT/SPTpol.

The contribution of this thesis, presented in Chapter 5, is a thorough analysis of how the reconstructed lensing information can be combined with the primary CMB data to perform reliably a joint parameter estimation. Such joint analyses are important to break degeneracies that limit the information that can be extracted from the CMB fluctuations laid down at recombination. The joint analysis is non-trivial because the part of the lensing information that is present in the primary CMB power spectrum as well as in the lensing reconstruction is potentially double-counted. We quantify the temperature-lensing cross-correlation analytically, finding two physical contributions and confirming the results with simulated lensed CMB maps. This cross-correlation has not been considered be-

fore and it could have turned out to be anywhere between zero and unity. We also use simulations to test approximations for the likelihood of the lensing reconstruction, which is non-trivial because the reconstruction is quadratic in the (nearly-Gaussian) CMB. As a result of this study, the Planck results in [2, 3], which combine lensing reconstruction with primary CMB data, use the lensing likelihood approximation proposed here and the result obtained here for the temperature-lensing cross-correlation, i.e. that this can be ignored and the lensing information from the lensing reconstruction is independent from that obtained from the primary CMB power spectrum. These results will also be relevant for obtaining cosmological information from future CMB lensing reconstructions.

Short overview of this thesis

This thesis is structured as outlined below. Every chapter begins with a short summary which contains a more detailed overview of its content.

- Chapter 2 provides background on cosmological perturbations. We start with primordial fluctuations seeded by inflation and discuss the induced primordial power spectrum and primordial non-Gaussianity, and how to estimate the latter from data. We then review the mechanism by which the primordial perturbations are transferred to CMB perturbations and discuss the effects of gravitational lensing on the CMB. We end the chapter by discussing large-scale dark matter perturbations seeded by inflationary fluctuations.
- Chapter 3 is about the efficient generation of general non-Gaussian N -body initial conditions using an expansion of the bispectrum kernel in terms of separable basis functions. A non-trivial consistency check of the implementation is presented by testing the non-Gaussianity of the initial conditions. This Chapter is based on [4].
- Chapter 4 starts with an overview of known analytical results for large-scale structure clustering and some phenomenological extensions. We then discuss efficient bispectrum estimation using a separable mode expansion and present the first implementation of this estimator in the 3D large-scale structure context. The implementation is validated and applied to N -body simulations

with and without primordial non-Gaussianity. We discuss results obtained for the gravitational bispectrum for Gaussian initial conditions and the excess bispectra measured for several types of non-Gaussian initial conditions. The bispectrum measurements are described with a time-shift model which is used to create novel fitting formulae for the dark matter bispectrum. This chapter is based on [5].

- In Chapter 5 we discuss how to use CMB lensing reconstruction for cosmological parameter analyses. After reviewing known results about the auto-correlations of the power spectra of the CMB temperature and the lensing reconstruction, we analyse the cross-correlation between these power spectra analytically using perturbative techniques. The results are tested against simulations and we assess the impact on parameter estimation. We then test approximations for the likelihood function of the lensing reconstructions and conclude by proposing a joint likelihood for combined CMB temperature and lensing reconstruction analyses. A series of appendices provide intuitive plausibility arguments for the size of the cross-correlations obtained by the perturbative calculations, as well as filling in several technical details. This chapter is based on [6].
- In Chapter 6 we provide a short summary of the thesis with the main conclusions.

Background: Cosmological Perturbations

Summary

In this chapter we provide some background on the generation and properties of cosmological perturbations seeded by inflation. We start with primordial quantum fluctuations and calculate the induced primordial power spectrum. After establishing an order of magnitude estimate indicating the smallness of primordial non-Gaussianity (PNG) for simple single-field inflation models, we define the bispectrum and trispectrum and briefly comment on several types of non-Gaussianities that can be produced at observable levels in specific inflation models. We give a brief discussion of the issue of estimating such non-Gaussianity from a realisation of the field. We then discuss the transfer of primordial perturbations to CMB temperature fluctuations using the Boltzmann equation. This is followed by a short review of weak gravitational lensing of the CMB. We discuss the impact of lensing on the CMB power spectrum and the anisotropy which is introduced in the CMB by fixed lenses. The latter is then used to discuss lensing reconstruction from the CMB temperature. We continue with the current observational status of CMB lensing and discuss the cosmological information that can in principle be extracted from CMB lensing. We comment on complications involved in parameter estimation with CMB lensing which will be considered in detail in Chapter 5. We end the chapter with a brief overview of perturbations of the large-scale dark matter distribution, discussing the evolution of perturbations during radiation and matter domination and the resulting transfer and growth function. We give a brief introduction to perturbation theory to derive the leading-order bispectrum from gravitational non-linearities, and comment on the Zel'dovich approximation and N -body simulations. The impact of PNG on large-scale structures will be discussed in later chapters. We use $\eta_{\mu\nu} = \text{diag}(1, -1, -1, -1)$ throughout this chapter. Large parts of this chapter are based on material from [7, 8, 9, 10, 11, 12, 13].

2.1 Primordial perturbations

2.1.1 Inflation

Up to tiny $\mathcal{O}(10^{-4})$ fluctuations, the observed CMB temperature is uniform across the sky. This implies that these photons must be emitted from regions that must have been in causal contact at some time in the past. However, two particles can only communicate with each other if their distance on a grid comoving with the cosmic fluid is smaller than the comoving Hubble scale¹ \mathcal{H}^{-1} . If we assume that gravity always decelerates the expansion of the universe, \mathcal{H}^{-1} increases in time over the whole history of the universe. Thus regions separated by a comoving distance larger than \mathcal{H}^{-1} at photon decoupling – corresponding to regions separated by more than a few degrees on the sky today (e.g. [14]) – must have been causally disconnected all the time before photon decoupling (because \mathcal{H}^{-1} was even smaller then). But if so many separate regions of the sky have never been in causal contact, how can photons from all these regions have the same temperature to one part in 10^4 ?

Alan Guth [15] solved this *horizon problem* by giving up the assumption that the expansion of the universe is always decelerating. Instead he suggested an epoch of accelerated expansion, *inflation*, during which the comoving Hubble scale \mathcal{H}^{-1} decreases. Then regions separated by a comoving distance larger than \mathcal{H}^{-1} at photon decoupling can have been in causal contact during inflation, when \mathcal{H}^{-1} was larger (see also Fig. 2.1 below). Notably, inflation also solves the flatness and magnetic monopole problems.

An additional striking feature of inflation is the prediction that all structures we see today were originally seeded by quantum fluctuations of the inflation-driving inflaton field. During inflation the quantum fluctuations are stretched to cosmologically relevant scales, leaving initial perturbations in the energy density and metric behind. After inflation these perturbations are amplified by gravitational

¹The distance light can travel from the big bang until some time $t(a)$ is given by the comoving horizon $\eta = \int_0^t \frac{dt}{a} = \int_0^a \frac{d\ln a}{aH(a)}$ (see e.g. [11]). Thus, the comoving Hubble scale $\mathcal{H}^{-1} = (aH)^{-1}$ is the contribution to the comoving horizon per logarithmic expansion interval $d\ln a$. Therefore, if two regions are separated by a distance larger than \mathcal{H}^{-1} [at the time $t(a)$] they cannot currently communicate.

instability, finally building galaxy clusters, galaxies, stars, etc. In the last decades the idea of inflation has been tested observationally with ever increasing precision, and so far it is consistent with all observations.

However, we do not know what actually caused inflation and several different inflationary mechanisms have been suggested, many of them being consistent with all observations (see e.g. [16, 17, 18, 19, 20]). One of the most powerful observables to distinguish these models are statistical properties of density perturbations, e.g. in the CMB or in large-scale structures. The standard model of single-field slow-roll inflation produces nearly Gaussian statistics, but other models can produce significant deviations from Gaussian statistics, i.e. *non-Gaussianity*.

In the following, over-dots denote derivatives with respect to coordinate time t , whereas primes ' denote derivatives with respect to conformal time η [the only exception being $V'(\varphi) \equiv \partial_\varphi V(\varphi)$].

2.1.2 Single-field slow-roll inflation

During inflation the expansion of the universe is accelerated, $\ddot{a} > 0$, where the scale factor a describes the expansion of the universe. The second Friedmann equation,

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3P), \quad (2.1)$$

which follows from the trace of the Einstein equations in the FLRW metric for a fluid with energy density ρ and pressure P , implies that we need $\rho + 3P < 0$ for inflation. This can be achieved with a single scalar field φ , the inflaton, if and only if the inflaton potential $V(\varphi)$ satisfies $V(\varphi) > \dot{\varphi}^2$.¹ Although not necessary for accelerated expansion (see e.g. [13]) we assume here that the *slow-roll approximation* applies. This is defined by neglecting $\ddot{\varphi}$ and $\dot{\varphi}^2$ in the first Friedmann equation and in the equation of motion for φ , which is equivalent to requiring

$$\epsilon \ll 1, \quad |\eta_s| \ll 1, \quad \text{and} \quad |\ddot{\varphi}| \ll |3H\dot{\varphi}|, \quad (2.2)$$

¹The stress-energy tensor of a homogeneous scalar field φ implies $\rho_\varphi = \dot{\varphi}^2/2 + V$ and $P_\varphi = \dot{\varphi}^2/2 - V$, so $\rho_\varphi + 3P_\varphi < 0 \Leftrightarrow \dot{\varphi}^2 < V$.

where the slow-roll parameters are defined by

$$\epsilon \equiv \frac{M_{\text{Pl}}^2}{2} \left(\frac{V'}{V} \right)^2 \quad \text{and} \quad \eta_s \equiv M_{\text{Pl}}^2 \frac{V''}{V}, \quad (2.3)$$

where $V' \equiv \partial_\varphi V$ and $M_{\text{Pl}} = \sqrt{\hbar c/(8\pi G)}$ is the reduced Planck mass. Then the inflaton field is slowly rolling towards the minimum of its potential V , keeping its kinetic energy below its potential energy (implying an accelerated expansion) for a sufficiently long time to solve the horizon problem. The slow-roll conditions imply that the effective inflaton mass is negligible compared to the Hubble rate, $m_\varphi^2 = |V''| \ll H^2$ (e.g. [7, 16]).

2.1.3 Horizon crossings

The properties of the quantum fluctuations during inflation are transferred to the spatial statistics of late-time observables such as the CMB or galaxies. To describe this mapping from early to late times note that all cosmologically relevant comoving scales $\lambda = k^{-1}$ start from inside the comoving Hubble radius $\mathcal{H}^{-1} = (aH)^{-1}$ and exit this horizon at some time during inflation, when $k = aH$; see Fig. 2.1. The subsequent super-horizon evolution is, luckily, trivial to follow because the comoving curvature perturbation \mathcal{R} stays constant. Finally, the modes re-enter the horizon at late times. The physics governing the sub-horizon evolution around recombination can be described by transfer functions. Before discussing these transfer functions we will characterise the fluctuations at horizon exit during inflation.

2.1.4 Quantisation of a light scalar field during slow-roll inflation

In the next few subsections we will follow [7, 9, 16, 21] to outline the quantisation of the inflaton field φ in slow-roll inflation. The discussion applies not only to the inflaton field but to any light scalar field minimally coupled to gravity.

comoving scales

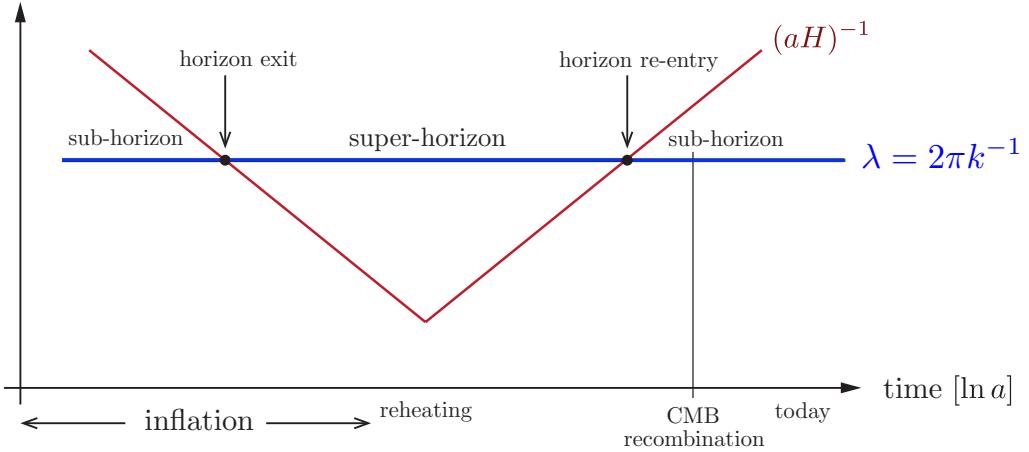


Figure 2.1: Evolution of the comoving Hubble radius $(aH)^{-1}$ compared to a constant comoving scale $\lambda = 2\pi k^{-1}$ (adapted from [8]).

Minimally coupling the inflaton field φ to gravity gives the action

$$S = \int d\eta d^3x \sqrt{-g} \left(\frac{1}{2} \nabla_\mu \varphi \nabla^\mu \varphi - V(\varphi) \right), \quad (2.4)$$

where $V(\varphi)$ is the potential of φ . In slow-roll inflation, metric perturbations can be neglected if we work in the spatially-flat gauge, corresponding to coordinates in which spatial hypersurfaces are flat [7, 8]. Thus $g_{\mu\nu} = a^2 \eta_{\mu\nu}$ and $g^{\mu\nu} = a^{-2} \eta^{\mu\nu}$, where $\eta_{\mu\nu} = \text{diag}(1, -1, -1, -1)$. Then the determinant of the metric tensor is $g \equiv \det g_{\mu\nu} = \det \text{diag}(a^2, -a^2, -a^2, -a^2) = -a^8$, so the action becomes

$$S = \int d\eta d^3x \left(\frac{a^2}{2} (\varphi'^2 - (\nabla \varphi)^2) - a^4 V(\varphi) \right). \quad (2.5)$$

It is convenient to split the field into an unperturbed homogeneous background value $\bar{\varphi}(\eta)$ and a small inhomogeneous fluctuation $\delta\varphi(\eta, \mathbf{x})$. Varying the action with respect to $\bar{\varphi}$ and neglecting all perturbations, we get the unperturbed equations of motion (note that the spatial $\nabla \bar{\varphi} = 0$)

$$\bar{\varphi}'' + 2\mathcal{H}\bar{\varphi}' + a^2 \partial_\varphi V(\bar{\varphi}) = 0. \quad (2.6)$$

Let us rescale $\delta\varphi$ by the scale factor, $u \equiv a\delta\varphi$. Our aim is to quantise this rescaled field fluctuation $u(\eta, \mathbf{x})$. We first need the classical equations of motion for u . Since u is small, we can analyse the equations of motion at different orders in u separately. At zeroth order in u we recover the equations of motion for the unperturbed field $\bar{\varphi}$. The $\mathcal{O}(u)$ part $S^{(1)}$ of the action vanishes because of the background equations of motion. To get the equations of motion at first order in u we need to vary the $\mathcal{O}(u^2)$ part $S^{(2)}$ of the action. We can neglect the inflaton mass $m_\varphi^2 = |\partial_\varphi^2 V|$ in slow-roll inflation to get¹

$$S^{(2)} = \frac{1}{2} \int d\eta d^3x \left(u'^2 + \frac{a''}{a} u^2 - (\nabla u)^2 \right). \quad (2.7)$$

The equations of motion linear in u follow from varying $S^{(2)}$ with respect to u , which gives in Fourier space

$$u''_{\mathbf{k}} + \left(k^2 - \frac{a''}{a} \right) u_{\mathbf{k}} = 0. \quad (2.8)$$

Now we can quantise the theory by promoting the classical fields u and $\pi = \partial\mathcal{L}/\partial u' = u'$ to operators \hat{u} and $\hat{\pi}$ with bosonic equal-time commutation relations

$$[\hat{u}(\eta, \mathbf{x}), \hat{\pi}(\eta, \mathbf{y})] = i\delta_D^{(3)}(\mathbf{x} - \mathbf{y}), \quad [\hat{u}(\eta, \mathbf{x}), \hat{u}(\eta, \mathbf{y})] = 0, \quad [\hat{\pi}(\eta, \mathbf{x}), \hat{\pi}(\eta, \mathbf{y})] = 0. \quad (2.9)$$

Furthermore we require that in the Heisenberg picture the time-dependent operators \hat{u} and $\hat{\pi}$ satisfy the analogues of the classical equations of motion, i.e. $\hat{\pi}_{\mathbf{k}} = \hat{u}'_{\mathbf{k}}$ and (2.8) if we replace $u_{\mathbf{k}}$ by $\hat{u}_{\mathbf{k}}$. Then the general Hermitian solution for $\hat{u}_{\mathbf{k}}$ is given by

$$\hat{u}(\eta, \mathbf{x}) = \int \frac{d^3k}{(2\pi)^3} \hat{u}_{\mathbf{k}}(\eta) e^{i\mathbf{k}\cdot\mathbf{x}}, \quad \hat{u}_{\mathbf{k}}(\eta) = \hat{a}_{\mathbf{k}} u_k(\eta) + \hat{a}_{-\mathbf{k}}^\dagger u_k^*(\eta), \quad (2.10)$$

where the mode functions $u_k(\eta)$ are solutions of the classical equation of motion (2.8) and $\hat{a}_{\mathbf{k}}^\dagger$ and $\hat{a}_{\mathbf{k}}$ are creation and annihilation operators, which satisfy commu-

¹Use partial integration for the term $\propto uu'$ to get a term $\propto u^2$ and Taylor expand V around $\bar{\varphi}$ (see e.g. [7]).

2.1 Primordial perturbations

tation relations that are determined by (2.9):¹

$$[\hat{a}_{\mathbf{k}}, \hat{a}_{\mathbf{k}'}^\dagger] = (2\pi)^3 \delta_D^{(3)}(\mathbf{k} - \mathbf{k}'), \quad [\hat{a}_{\mathbf{k}}, \hat{a}_{\mathbf{k}'}] = 0, \quad [\hat{a}_{\mathbf{k}}^\dagger, \hat{a}_{\mathbf{k}'}^\dagger] = 0, \quad (2.11)$$

To solve the classical equation of motion (2.8), let us now consider the time span from a few Hubble times before horizon exit until a few Hubble times after horizon exit. We define² the conformal time by

$$\eta = -\frac{1}{H_k a}, \quad (2.12)$$

where H_k is the Hubble parameter when the fixed comoving scale k^{-1} exits the horizon, i.e. $k = a_k H_k$. With this definition, $a''/a = (\eta^{-1})''\eta = 2/\eta^2$, so that (2.8) becomes

$$u_{\mathbf{k}}'' + \left(k^2 - \frac{2}{\eta^2} \right) u_{\mathbf{k}} = 0. \quad (2.13)$$

The solution with correct initial condition is given by³ [7, 9, 16]

$$u_k(\eta) = C e^{-ik\eta} \left(1 - \frac{i}{k\eta} \right), \quad (2.14)$$

where

$$C \equiv \frac{1}{\sqrt{2k}} \quad (2.15)$$

is a normalisation constant.

¹The normalisation of the first equation is chosen conveniently by rescaling the mode functions such that their (time-independent) Wronskian is $u_k(\eta)u_k^*(\eta) - u_k'(\eta)u_k^*(\eta) = i/(2\pi)^3$ (e.g. [7, 8, 9]).

²It is straightforward to check that this definition is consistent with $dt = ad\eta$ a few Hubble times before and after horizon exit when $H \approx H_k$: $d\eta = \frac{da}{H_k a^2} = \frac{\dot{a}dt}{H_k a^2} = \frac{H}{H_k} \frac{dt}{a}$. Note that $\eta < 0$ and η approaches zero with time, i.e. $|\eta|$ decreases with time. Also note that $\eta = -1/k$ at horizon exit.

³Well before horizon exit, $|\eta| \gg k^{-1}$, Eq. (2.8) becomes $u_{\mathbf{k}}'' + k^2 u_{\mathbf{k}} = 0$ with solution $u_k(\eta) \propto e^{-ik\eta}$. By neglecting the second solution $e^{+ik\eta}$ we select the ‘minimal excitation state’ as the vacuum, minimising the expectation value of the Hamiltonian (see e.g. [9]).

2.1.5 From quantum to classical fluctuations

The solution (2.14) behaves rather differently before and after horizon exit (see e.g. [13]). Well before horizon exit, $|\eta| \gg k^{-1}$, it can be approximated by $u_k(\eta) \approx Ce^{-ik\eta}$, so that the Heisenberg-picture field operator $\hat{u}_{\mathbf{k}}$ from (2.10) has a non-trivial time-dependence

$$\hat{u}_{\mathbf{k}}(\eta) = C \left(\hat{a}_{\mathbf{k}} e^{-ik\eta} + \hat{a}_{-\mathbf{k}}^\dagger e^{+ik\eta} \right). \quad (2.16)$$

The sub-horizon fluctuations are therefore quantum objects in the sense that performing a measurement at some time does not determine the outcome of later measurements because the state has evolved away from the eigenstate of the first measurement. In contrast, well after horizon exit, $|\eta| \ll k^{-1}$, we can approximate (2.14) by the purely imaginary solution $u_k(\eta) \approx -iC/(k\eta)$ to get a field operator with trivial time-dependence

$$\hat{u}_{\mathbf{k}}(\eta) = -\frac{iC}{k\eta} \left(\hat{a}_{\mathbf{k}} - \hat{a}_{-\mathbf{k}}^\dagger \right). \quad (2.17)$$

The operator $\hat{u}_{\mathbf{k}}(\eta)$ is only rescaled with time but does not involve different combinations of $\hat{a}_{\mathbf{k}}$ and $\hat{a}_{-\mathbf{k}}^\dagger$ at different times. Therefore the super-horizon fluctuation is a classical object in the sense that after a measurement the state stays in the same eigenstate [13].

2.1.6 Primordial power spectrum

A few e-folds after horizon exit the variance of the inflaton fluctuations, or their power spectrum $P_{\delta\varphi}(k)$, can be computed as

$$(2\pi)^3 P_{\delta\varphi}(k) \delta_D(\mathbf{k} - \mathbf{k}') \equiv \langle 0 | \hat{\delta\varphi}_{\mathbf{k}} \hat{\delta\varphi}_{\mathbf{k}'}^\dagger | 0 \rangle \stackrel{(2.17)}{=} \frac{\langle 0 | \hat{a}_{\mathbf{k}} \hat{a}_{\mathbf{k}'}^\dagger | 0 \rangle}{2\eta^2 a^2 k^3} = (2\pi)^3 \frac{H_k^2}{2k^3} \delta_D(\mathbf{k} - \mathbf{k}'), \quad (2.18)$$

where we assumed no interactions and used $u = a\delta\varphi$ and the super-horizon solution (2.17). In the last step we used (2.12) and $[\hat{a}_{\mathbf{k}}, \hat{a}_{\mathbf{k}'}^\dagger] = (2\pi)^3 \delta_D(\mathbf{k} - \mathbf{k}')$. From (2.18)

2.1 Primordial perturbations

we get for the dimensionless power spectrum $\mathcal{P} \equiv k^3 P / (2\pi^2)$

$$\mathcal{P}_{\delta\varphi}(k) = \left(\frac{H_k}{2\pi} \right)^2. \quad (2.19)$$

This is almost independent of scale k because the Hubble constant is only slowly varying during inflation. In the spatially-flat gauge the comoving curvature perturbation \mathcal{R} , which is conserved after horizon exit, is related to $\delta\varphi$ (interpreted as fluctuations on hypersurfaces of zero intrinsic curvature) via [7]

$$\mathcal{R} = -\mathcal{H} \frac{\delta\varphi}{\partial_\eta \bar{\varphi}}. \quad (2.20)$$

The corresponding power spectrum of \mathcal{R} at horizon exit is therefore

$$\mathcal{P}_{\mathcal{R}}(k) = \left(\frac{\mathcal{H}}{\partial_\eta \bar{\varphi}} \right)^2 \mathcal{P}_{\delta\varphi}(k) = \left(\frac{H^2}{2\pi \partial_t \bar{\varphi}} \right)_{k=aH}^2. \quad (2.21)$$

The small scale-dependence is described by the spectral index n_s

$$n_s - 1 \equiv \frac{d \ln \mathcal{P}_{\mathcal{R}}(k)}{d \ln k}, \quad (2.22)$$

i.e. $n_s = 1$ corresponds to a perfectly scale-invariant power spectrum. Under the slow-roll approximation one can show (see e.g. [7, 8, 9])

$$n_s - 1 = 2\eta_s - 6\epsilon, \quad (2.23)$$

i.e. the spectral index probes the slow-roll parameters (2.3) and therefore constrains the shape of the inflaton potential. Recently, Planck measured $n_s = 0.9603 \pm 0.0073$ (68% error) [2] which rules out perfect scale-invariance by more than 5σ .

2.1.7 Primordial non-Gaussianity

To constrain further the physics leading to inflation we consider higher n -point correlation functions beyond the power spectrum. We will first give a rough order-of-magnitude estimate for the amplitude of non-Gaussianity expected in single-field slow-roll inflation. We then characterise non-Gaussianity in more detail by

studying the momentum-dependence for some inflation models which can generate observable levels of non-Gaussianity.

2.1.7.1 Order of magnitude estimate for single-field slow roll inflation

If the 3-point quantum correlation function (of an interacting theory) is non-zero, the field perturbations $\delta\varphi$ are non-Gaussian. To estimate the order of magnitude of the quantum correlation functions for $n \geq 3$ we use the path integral [22]

$$\langle \Omega | \delta\varphi_{\mathbf{k}_1} \cdots \delta\varphi_{\mathbf{k}_n} | \Omega \rangle = \frac{\int \mathcal{D}[\delta\varphi_{\mathbf{k}}] \delta\varphi_{\mathbf{k}_1} \cdots \delta\varphi_{\mathbf{k}_n} \exp(i \int_{\gamma} \mathcal{L}(\delta\varphi_{\mathbf{k}}))}{\int \mathcal{D}[\delta\varphi_{\mathbf{k}}] \exp(i \int_{\gamma} \mathcal{L}(\delta\varphi_{\mathbf{k}}))}, \quad (2.24)$$

where the contour γ of the integration over the Lagrangian \mathcal{L} is such that it selects the vacuum $|\Omega\rangle$ of the interacting theory by an appropriate $i\epsilon$ -prescription [22] [23]. One can evaluate the path integral by rewriting it in terms of a generating functional with sources J_k and then replacing the field perturbations in the numerator of (2.24) by functional derivatives ∂_J with respect to the sources. Together with corresponding prefactors (including time and spatial derivatives) these functional derivatives can be pulled out of the path integral, leaving a Gaussian integral which can be solved. Acting with the functional derivatives ∂_J on this solution will generate terms corresponding to Feynman diagrams. In the following we will suppress all prefactors of $\delta\varphi$ (including time and spatial derivatives), which are pulled out of the path integral together with the functional derivatives ∂_J , because we only want to estimate the order of magnitude of the correlation functions, keeping only slow-roll parameters and $\delta\varphi$.

Our goal is to calculate the non-Gaussianity of the comoving curvature perturbation \mathcal{R} that results from non-Gaussianity of the field perturbation $\delta\varphi$. Assuming slow-roll, $H = \sqrt{V/(3M_{\text{Pl}}^2)}$ and $\partial_t \bar{\varphi} = -V'/3H$ imply that the conversion (2.20) becomes

$$\mathcal{R} = \frac{\delta\varphi}{\sqrt{2\epsilon} M_{\text{Pl}}} \sim \frac{\delta\varphi}{\sqrt{\epsilon}}. \quad (2.25)$$

The appropriate generalisation beyond linear order is obtained by considering ζ , which agrees with \mathcal{R} at linear order, is conserved on super-horizon scales (for adiabatic perturbations) and is defined non-linearly in terms of the spatial part of

2.1 Primordial perturbations

the metric by (see e.g. [24] and references therein)

$$g_{ij} = a^2 e^{2\zeta} \delta_{ij}. \quad (2.26)$$

We get a path integral formula for correlation functions $\langle \zeta_{\mathbf{k}_1} \cdots \zeta_{\mathbf{k}_n} \rangle$ if we replace all $\delta\varphi$ on the right hand side of (2.24) by $\sqrt{2\epsilon}M_{\text{Pl}}\zeta$.

If we only keep terms with two ζ in the Lagrangian $\mathcal{L}(\zeta)$, corresponding to (2.7), the exponent in the path integral (2.24) includes only these terms with two ζ and the n -point correlation function of ζ becomes (suppressing prefactors and spatial and time derivatives, and using the notation $\langle \cdot \rangle_2$ to indicate that only terms with two ζ in the action are considered)

$$\langle \zeta_{\mathbf{k}_1} \cdots \zeta_{\mathbf{k}_n} \rangle_2 \sim \int \mathcal{D}[\zeta_{\mathbf{k}}] \zeta_{\mathbf{k}_1} \cdots \zeta_{\mathbf{k}_n} e^{i \int_{\gamma} \zeta_{\mathbf{k}}^2}, \quad (2.27)$$

which corresponds to a Gaussian probability functional if we regard ζ as a classical random variable, i.e. the (quantum and ensemble) correlation function is zero for odd n and can be expressed by $\langle \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} \rangle_2$ for even n . From (2.25) and noting that $\langle \delta\varphi_{\mathbf{k}_1} \delta\varphi_{\mathbf{k}_2} \rangle_2 \sim \epsilon^0$ from (2.19), we find

$$\langle \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} \rangle_2 \sim \epsilon^{-1}. \quad (2.28)$$

We can get non-Gaussianity if higher-order terms with three or more $\zeta \propto \delta\varphi$ in the Lagrangian are considered, corresponding to interaction terms of the inflaton. They can arise from expanding the potential V to third order or from the $\sqrt{-g}$ term in the action. It turns out that the part of the action with three ζ is of order ϵ^2 [22, 23], i.e. $\mathcal{L}_3 \sim \epsilon^2 \zeta^3$, where we again suppress derivatives and factors that are not slow-roll parameters. The corresponding contribution to the 2-point function is

$$\begin{aligned} \langle \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} \rangle_3 &\sim \int \mathcal{D}[\zeta_{\mathbf{k}}] \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} e^{i \int_{\gamma} \mathcal{L}_2} \left(1 + i \int_{\gamma} \mathcal{L}_3 + \frac{i^2}{2} \left(\int_{\gamma} \mathcal{L}_3 \right)^2 + \dots \right) \\ &\sim \langle \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} \rangle_2 + \frac{i^2}{2} \epsilon^4 \int \mathcal{D}[\zeta_{\mathbf{k}}] \int_{\gamma} \int_{\gamma} \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} \zeta_{\mathbf{k}}^3 \zeta_{\mathbf{k}'}^3 e^{i \int_{\gamma} \mathcal{L}_2} \\ &\sim \epsilon^{-1} + o(\epsilon^0), \end{aligned} \quad (2.29)$$

where we used that integrals over an odd number of ζ 's vanish and that the integral involving eight ζ 's can be written as a linear combination of terms proportional to $(\langle \zeta^2 \rangle_2)^4 \sim \epsilon^{-4}$ because the exponential is a Gaussian.¹ Thus, at leading order in slow-roll parameters, the 2-point function with \mathcal{L}_3 is the same as if we only consider \mathcal{L}_2 . However, the 3-point function becomes non-zero if we use \mathcal{L}_3 instead of \mathcal{L}_2 ,

$$\begin{aligned} \langle \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} \zeta_{\mathbf{k}_3} \rangle_3 &\sim i\epsilon^2 \int \mathcal{D}[\zeta_{\mathbf{k}}] \int_{\gamma} \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} \zeta_{\mathbf{k}_3} \zeta_{\mathbf{k}}^3 e^{i \int_{\gamma} \mathcal{L}_2} + \dots \\ &\sim \epsilon^{-1} + o(\epsilon^0), \end{aligned} \quad (2.30)$$

where we used that the integral can be expressed as a linear combination of terms of the form $(\langle \zeta^2 \rangle_2)^3 \sim \epsilon^{-3}$.

We can quantify the level of non-Gaussianity by the ratio of the 3-point function (2.30) and the square of the 2-point function from (2.28) and (2.29) [25],

$$f_{\text{NL}}^{\zeta} \sim \frac{\langle \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} \zeta_{\mathbf{k}_3} \rangle_3}{(\langle \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} \rangle_2 + \langle \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} \rangle_3)^2} \sim \frac{\epsilon^{-1}}{(\epsilon^{-1})^2} \sim \epsilon. \quad (2.31)$$

Thus, in single-field slow-roll inflation (with standard Bunch-Davies vacuum), the non-Gaussianity parameter f_{NL}^{ζ} due to non-Gaussian field perturbations $\delta\varphi$ is generically suppressed by the slow-roll parameter $\epsilon \leq \mathcal{O}(10^{-2})$. Unfortunately, this level of non-Gaussianity is too small to be observable in the near future. Note that this result has been derived much more rigorously in [22, 23].

2.1.7.2 Primordial bispectrum and trispectrum definitions

Despite the slow-roll suppressed non-Gaussianity for single-field slow-roll inflation there are many alternative classes of inflation models which can generate observable levels of non-Gaussianity, e.g. if they involve additional fields, non-standard kinetic terms, or a modified initial vacuum state. These mechanisms produce different types of non-Gaussianity characterised by the momentum-dependence, or *shape*, of the non-Gaussian 3- and 4-point functions. It is convenient to work

¹The o -notation in Eq. (2.29) means that the term can be at most of order ϵ^0 , which takes account of the possibility that loop contributions can in principle cancel each other.

2.1 Primordial perturbations

with the Bardeen potential Φ which is nominally related to \mathcal{R} via $\Phi = -3\mathcal{R}/5$ on super-horizon scales during matter domination and does not change under gauge transformations (see e.g. [7, 26]; note that different signs are used in the literature).

We will describe $\Phi(\mathbf{x})$ as a classical random variable at each point \mathbf{x} , with zero mean. The probability of a field configuration $\Phi(\mathbf{x})$ is given by a probability distribution functional $P[\Phi(\mathbf{x})]$. Our goal is to estimate the form of this probability distribution, especially its deviations from a Gaussian distribution. Since the background cosmology is assumed to be homogeneous and isotropic, we assume that P is also homogeneous and isotropic, i.e. we assume that the probability of a field configuration does not change if we translate or rotate the field. In real space this means that if we fix Φ at position \mathbf{x}_1 then the probability to have some value $\Phi(\mathbf{x}_2)$ at \mathbf{x}_2 may only depend on the distance $|\mathbf{x}_1 - \mathbf{x}_2|$ between the two points. Therefore the 2-point correlation function of $\Phi(\mathbf{x})$ is of the form $\langle \Phi(\mathbf{x}_1)\Phi(\mathbf{x}_2) \rangle = \xi(|\mathbf{x}_1 - \mathbf{x}_2|)$. In Fourier space the 2-point function is therefore of the form

$$\langle \Phi(\mathbf{k}_1)\Phi(\mathbf{k}_2) \rangle = (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2) P_\Phi(k_1), \quad (2.32)$$

which follows by writing the $\Phi(\mathbf{k}_i)$ as Fourier transforms of $\Phi(\mathbf{x}_i)$. The quantity $P_\Phi(k)$ is the power spectrum of Φ . Since $\Phi(\mathbf{x})$ is real we have $\Phi(\mathbf{k}) = \Phi^*(-\mathbf{k})$ and therefore (2.32) implies that Fourier modes with different wavevectors \mathbf{k} are uncorrelated, i.e. if we measure $\Phi(\mathbf{k}_1)$ this tells us nothing about the likely value of $\Phi(\mathbf{k}_2)$ for $\mathbf{k}_2 \neq \pm \mathbf{k}_1$.

Similarly, statistical homogeneity and isotropy imply that the 3-point function is of the form

$$\langle \Phi(\mathbf{k}_1)\Phi(\mathbf{k}_2)\Phi(\mathbf{k}_3) \rangle = (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) f_{\text{NL}} B_\Phi(k_1, k_2, k_3), \quad (2.33)$$

where we introduced the *bispectrum* B_Φ and the non-linearity parameter f_{NL} , which is conventionally defined such that $B_\Phi(k, k, k)$ equals $P_\Phi^2(k)$ times a number of order unity. The bispectrum vanishes if Φ is Gaussian. If the bispectrum deviates from zero then Φ must be non-Gaussian and modes with different wavevectors are correlated, i.e. if we measure two modes $\Phi(\mathbf{k}_1)$ and $\Phi(\mathbf{k}_2)$, we can predict the most likely value of $\Phi(\mathbf{k}_3)$ using B_Φ [27]. Therefore the bispectrum represents the

lowest-order correlation function that distinguishes non-Gaussian from Gaussian perturbations. Note that the Dirac delta function in (2.33) imposes a triangle condition on the three wavevectors \mathbf{k}_i , i.e. the bispectrum is only defined for those triples (k_1, k_2, k_3) for which there is a triangle with lengths k_1, k_2, k_3 . Since the longest edge may not be longer than the sum of the other two edges, we need

$$\begin{cases} k_1 \leq k_2 + k_3 & \text{for } k_1 \geq k_2, k_3, \text{ or} \\ k_2 \leq k_1 + k_3 & \text{for } k_2 \geq k_1, k_3, \text{ or} \\ k_3 \leq k_1 + k_2 & \text{for } k_3 \geq k_1, k_2. \end{cases} \quad (2.34)$$

If we additionally demand $k \leq k_{\max}$, the allowed domain for (k_1, k_2, k_3) forms a ‘tetrapyd’ \mathcal{V}_B (see [28] and Fig. 4.1 below). The functional dependence of the bispectrum on the triangle configurations, i.e. the *bispectrum shape*, is sensitive to the specific inflation model under consideration. Different classes of inflation models can be distinguished observationally if they generate different bispectrum shapes with sufficient amplitude.

Extending the above to the 4-point function, we define the *trispectrum* T_Φ and its amplitude G_{NL} by

$$\langle \Phi(\mathbf{k}_1)\Phi(\mathbf{k}_2)\Phi(\mathbf{k}_3)\Phi(\mathbf{k}_4) \rangle_c = (2\pi)^3 \delta_D \left(\sum_{i=1}^4 \mathbf{k}_i \right) G_{\text{NL}} T_\Phi(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4), \quad (2.35)$$

where the Dirac delta enforces that the four wavevectors form a closed quadrilateral, and $\langle \cdot \rangle_c$ denotes the part of the correlation function that corresponds to connected Feynman diagrams and is given by the full correlation function minus the part corresponding to disconnected Feynman diagrams,

$$\langle \Phi(\mathbf{k}_1) \cdots \Phi(\mathbf{k}_4) \rangle_c \equiv \langle \Phi(\mathbf{k}_1) \cdots \Phi(\mathbf{k}_4) \rangle - (\langle \Phi(\mathbf{k}_1)\Phi(\mathbf{k}_2) \rangle \langle \Phi(\mathbf{k}_3)\Phi(\mathbf{k}_4) \rangle + 2 \text{ perms}). \quad (2.36)$$

If perturbations are Gaussian, we only have the disconnected part, whereas the perturbations must be non-Gaussian for the connected part and the trispectrum to deviate from zero.

2.1.7.3 Primordial bispectrum shapes

We provide a short overview over generic non-Gaussianity shapes that can be generated by inflation models.

Local shape The fiducial f_{NL} model of primordial non-Gaussianity is the local model [29], which is described in this way because it can be generated simply by squaring a Gaussian field Φ_G in real space,

$$\Phi(\mathbf{x}) = \Phi_G(\mathbf{x}) + f_{\text{NL}}[\Phi_G^2(\mathbf{x}) - \langle \Phi_G^2 \rangle]. \quad (2.37)$$

where subtracting the ensemble average ensures that $\Phi(\mathbf{x})$ has zero mean. The resulting bispectrum takes the form

$$B_{\Phi}^{\text{loc}}(k_1, k_2, k_3) = 2 [P_{\Phi}(k_1)P_{\Phi}(k_2) + 2 \text{ perms}], \quad (2.38)$$

which peaks for squeezed triangle configurations, where one wavenumber is much smaller than the other two (see Fig. 4.3 below). Multiple-field inflation models are one potential source of this shape (for a review see [17]). If a bispectrum signal in the squeezed limit is detected, then this will rule out all single-field models of inflation [23, 30, 31, 32].

Equilateral shape Higher derivative operators in the inflationary action, arising e.g. in DBI inflation [33] and in effective field theory approaches [34, 35, 36, 37], produce a shape that can be approximated by the separable equilateral template [17, 38, 39]

$$B_{\Phi}^{\text{eq}}(k_1, k_2, k_3) = 6 \left[- (P_{\Phi}(k_1)P_{\Phi}(k_2) + 2 \text{ perms}) - 2(P_{\Phi}(k_1)P_{\Phi}(k_2)P_{\Phi}(k_3))^{2/3} \right. \\ \left. + (P_{\Phi}^{1/3}(k_1)P_{\Phi}^{2/3}(k_2)P_{\Phi}(k_3) + 5 \text{ perms}) \right], \quad (2.39)$$

which peaks in the equilateral limit, $k_1 = k_2 = k_3$. For equilateral triangles with $k_1 = k_2 = k_3$ summing up the corresponding three plane waves, $\sum_j \Phi(\mathbf{k}_j) e^{i\mathbf{k}_j \cdot \mathbf{x}}$, gives filamentary overdensities, i.e. overdense cylinders along the direction perpendicular to the plane of the triangle of $(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)$, surrounded by underdensities (as

motivated in [27]). A primordial equilateral shape from higher derivatives, therefore, must be distinguished carefully from the stronger equilateral and (highly correlated) constant contribution produced at late times by nonlinear gravitational collapse and the emergence of filamentary and point-like structures (see e.g. [40, 41]). We will confirm this expectation by measuring the bispectrum in N -body simulations in Chapter 4.

Orthogonal shape Another shape that can arise from single-field inflation [42] or from multifield DBI Galileon inflation [43, 44] is the orthogonal shape, which is roughly orthogonal (see Sec. 2.1.7.4 below) to the equilateral and the local shape and peaks, with opposite sign, for both equilateral and flattened/elongated triangle configurations (e.g. $k_3 = k_1 + k_2$).

It can be approximated by the separable template [42]

$$B_{\Phi}^{\text{orth}}(k_1, k_2, k_3) = 6 \left[3(P_{\Phi}^{1/3}(k_1)P_{\Phi}^{2/3}(k_2)P_{\Phi}(k_3) + 5 \text{ perms}) - \frac{3}{2}B_{\Phi}^{\text{loc}}(k_1, k_2, k_3) - 8(P_{\Phi}(k_1)P_{\Phi}(k_2)P_{\Phi}(k_3))^{2/3} \right]. \quad (2.40)$$

We note that all the shapes above (2.38)–(2.40) are written explicitly as the sum of separable functions of the wavenumbers k_1, k_2, k_3 .

Flattened shape A non-Bunch-Davies vacuum leads to shapes which peak for flattened/elongated (e.g. $k_1 + k_2 = k_3$) and folded (e.g. $k_1 = 2k_2 = 2k_3$) triangle configurations¹ [17, 46, 47, 48]. This flattened shape depends on terms such as $1/(k_1 + k_2 - k_3)$ [46]. Therefore it provides an important example of a bispectrum which is inherently non-separable, so it is computationally difficult to generate initial conditions for this shape and to extract its amplitude from data. We will use an expansion in separable modes to solve both problems [4, 28, 49, 50]. There is a divergence for elongated triangles, $k_3 = k_1 + k_2$, which must be removed with a physically motivated cutoff. For definiteness we use the same shape and cutoff as was used for the CMB in [49, 51], setting $B = 0$ for $k_1 + k_2 - k_3 < 0.03(k_1 + k_2 + k_3)$

¹Sometimes other names are used for these triangles; we follow [45]. Note that shapes peaking for folded configurations can also arise without invoking a non-Bunch-Davies vacuum if curvature-related operators are included in the effective theory of single-field inflation [35].

and permutations, and then smoothing on $k_1 + k_2 + k_3 = \text{const}$ slices with a Gaussian filter with FWHM of $0.03/(k_1 + k_2 + k_3)$.

Templates vs shapes It is worth noting that the equilateral and orthogonal templates are separable approximations of inflationary shapes, with good agreement where the primordial bispectrum signal peaks. However late time observables may be sensitive to suppressed triangle configurations for which the template and physical shapes may differ significantly (e.g. scale-dependent halo bias mainly depends on the squeezed limit of the bispectrum [39, 52]). The methods described below enable us to simulate and estimate physical non-separable bispectra without using such separable templates. However we do not expect significant differences for the dark matter bispectrum, because its peaks are well approximated by the separable templates. For this reason and for easier comparison with other simulations, we use the equilateral and orthogonal templates instead of the physical shapes. Future work, especially on the power spectrum and bispectrum of halos, will instead focus on the physical shapes.

2.1.7.4 Estimating non-Gaussianity

As a first step towards confronting the theory predictions outlined above with observations, we discuss here how non-Gaussianity of the probability distribution of a generic 3D field $\delta(\mathbf{x})$ can be estimated from a realisation of that field. We leave the discussion of the transfer of primordial fluctuations to late-time observables for later, noting that the estimation formalism is essentially the same if linear transfer functions are included; in Chapters 3 and 4, $\delta(\mathbf{x})$ will denote the fractional perturbation of the late time dark matter density, while the CMB lensing reconstruction discussed in Section 2.3.1.6 and Chapter 5 can be understood as an optimal trispectrum estimator of CMB temperature fluctuations and the 3D $\delta(\mathbf{x})$ is replaced by $\Delta T(\hat{\mathbf{n}})$, or spherical harmonic coefficients a_{lm} on the 2D sphere.

To learn about non-Gaussianity of the 3D field $\delta(\mathbf{x})$, we derive an optimal estimator by reformulating the derivation in [53], where the 2D CMB case was considered. As a theoretical model we assume that $\delta(\mathbf{x})$ is drawn from a probability distribution function (pdf) which is fully specified by the covariance $C_{ij} \equiv \langle \delta_{\mathbf{k}_i} \delta_{\mathbf{k}_j}^* \rangle$ and the bispectrum $f_{\text{NL}} B(k_1, k_2, k_3)$, which we assume to have some fiducial shape

$B(k_1, k_2, k_3)$ but unknown amplitude f_{NL} [i.e. the pdf is given by (2.41) below; see e.g. [50, 54] for the generalisation to include the trispectrum]. What value of f_{NL} describes the data $\delta(\mathbf{x})$ best? This question can be answered by maximising the (posterior) probability $\Pr(f_{\text{NL}}|\delta)$ with respect to f_{NL} . If we assume a flat prior $\Pr(f_{\text{NL}}) = \text{const.}$, this is equivalent to maximising the likelihood $\mathcal{L}(\delta|f_{\text{NL}})$ with respect to f_{NL} . While the posterior probability is hard to write down, the likelihood can be obtained from a multi-variate Edgeworth expansion around a Gaussian pdf in the limit of weak non-Gaussianity (see e.g. [53, 55] and references therein):

$$\begin{aligned} \mathcal{L}(\delta|f_{\text{NL}}) &\propto \\ &\left[1 - \int \frac{d^3\mathbf{k}_1 \cdots d^3\mathbf{k}_3}{(2\pi)^9} \frac{1}{6} \langle \delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \delta_{\mathbf{k}_3} \rangle \frac{\partial}{\partial \delta_{\mathbf{k}_1}} \frac{\partial}{\partial \delta_{\mathbf{k}_2}} \frac{\partial}{\partial \delta_{\mathbf{k}_3}} \right] \frac{1}{\sqrt{\det C}} \prod_{ij} e^{-\frac{1}{2} \delta_{\mathbf{k}_i}^* (C^{-1})_{ij} \delta_{\mathbf{k}_j}}, \end{aligned} \quad (2.41)$$

where $C_{ij} \equiv \langle \delta_{\mathbf{k}_i} \delta_{\mathbf{k}_j}^* \rangle$. We will assume statistical homogeneity and isotropy so that $C_{ij} = \delta_{ij} C_{ii}$. However, in real-world applications inhomogeneous noise will lead to non-zero off-diagonal entries, which will lead to an additional term in the final estimator that is linear in the data δ [50, 56]. Calculating the functional derivatives in (2.41) yields

$$\begin{aligned} \mathcal{L}(\delta|f_{\text{NL}}) &\propto \left[1 - \int_{\mathbf{k}_1 \mathbf{k}_2 \mathbf{k}_3} \frac{f_{\text{NL}} (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B(k_1, k_2, k_3)}{6 C_{11} C_{22} C_{33}} \right. \\ &\times \left. (\delta_{\mathbf{k}_1}^* \delta_{\mathbf{k}_2}^* \delta_{\mathbf{k}_3}^* - \{\delta_D(\mathbf{k}_1 + \mathbf{k}_2) C_{11} \delta_{\mathbf{k}_3}^* + 2 \text{ perms}\}) + \mathcal{O}(f_{\text{NL}}^2) \right] \frac{\prod_i e^{-\frac{1}{2} |\delta_{\mathbf{k}_i}|^2 C_{ii}^{-1}}}{\sqrt{\prod_j C_{jj}}}, \end{aligned} \quad (2.42)$$

where we replaced the three point function by the bispectrum B using (2.33) and $\int_{\mathbf{k}_1 \mathbf{k}_2 \mathbf{k}_3}$ stands for the integral over \mathbf{k}_1 , \mathbf{k}_2 and \mathbf{k}_3 as in (2.41). The terms in the curly brackets, which are linear in the data δ , vanish because the two Dirac delta functions force one of the wavevectors to be zero and perturbations at $\mathbf{k} = 0$ vanish because the perturbations have zero mean in real space by definition. These terms do not vanish if C_{ij} is non-diagonal, and neglecting these linear terms in that

2.1 Primordial perturbations

case would lead to an unbiased but sub-optimal estimator [50, 56]. If we take the logarithm of (2.42) and use $\ln(1 + h) \approx h + \mathcal{O}(h^2)$ for $h \ll 1$ to expand in f_{NL} we get (up to additional constants, and up to $\mathcal{O}(f_{\text{NL}}^2)$ corrections)

$$-\ln \mathcal{L}(\delta|f_{\text{NL}}) \approx \left[\sum_i \frac{1}{2} C_{ii}^{-1} |\delta_{\mathbf{k}_i}|^2 + \int_{\mathbf{k}_1 \mathbf{k}_2 \mathbf{k}_3} \frac{f_{\text{NL}} (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B(k_1, k_2, k_3) \delta_{\mathbf{k}_1}^* \delta_{\mathbf{k}_2}^* \delta_{\mathbf{k}_3}^*}{6 C_{11} C_{22} C_{33}} \right]. \quad (2.43)$$

To maximise this with respect to f_{NL} , note that

$$-\frac{d \ln \mathcal{L}(\delta|f_{\text{NL}})}{d f_{\text{NL}}} = \int_{\mathbf{k}_1 \mathbf{k}_2 \mathbf{k}_3} \frac{(2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B(k_1, k_2, k_3) \delta_{\mathbf{k}_1}^* \delta_{\mathbf{k}_2}^* \delta_{\mathbf{k}_3}^*}{6 C_{11} C_{22} C_{33}} + \mathcal{O}(f_{\text{NL}}). \quad (2.44)$$

Setting this to zero gives the best-fit parameter \hat{f}_{NL} . This shows that we actually need the $\mathcal{O}(f_{\text{NL}})$ term here, i.e. we should have expanded $\mathcal{L}(\delta|f_{\text{NL}})$ up to $\mathcal{O}(f_{\text{NL}}^2)$ from the beginning. However it was shown in [53] that this term can be replaced by its Gaussian expectation value to a good approximation (this approximation is commonly used for Gaussian maximum likelihood estimators [57]). Therefore this term does not depend on the data δ and only leads to a normalisation factor for the f_{NL} estimator, so that

$$\hat{f}_{\text{NL}} \propto \int_{\mathbf{k}_1 \mathbf{k}_2 \mathbf{k}_3} \frac{(2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B(k_1, k_2, k_3) \delta_{\mathbf{k}_1}^* \delta_{\mathbf{k}_2}^* \delta_{\mathbf{k}_3}^*}{6 C_{11} C_{22} C_{33}}. \quad (2.45)$$

Using $C_{ii} = \langle \delta_{\mathbf{k}_i}^* \delta_{\mathbf{k}_i} \rangle = (2\pi)^3 \delta_D(\mathbf{0}) P_\delta(k_i)$ and introducing a normalisation factor $N_{f_{\text{NL}}}$ this estimator becomes [50]

$$\hat{f}_{\text{NL}}^B = \frac{1}{N_{f_{\text{NL}}}} \int \frac{d^3 \mathbf{k}_1 \cdots d^3 \mathbf{k}_3}{(2\pi)^9} \frac{(2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B(k_1, k_2, k_3) \delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \delta_{\mathbf{k}_3}}{P_\delta(k_1) P_\delta(k_2) P_\delta(k_3)}, \quad (2.46)$$

where we also used $\delta_{\mathbf{k}}^* = \delta_{-\mathbf{k}}$ and wrote out the integral over the \mathbf{k}_i . Given a theoretical bispectrum shape B and some data δ , Eq. (2.46) is an optimal estimator for the amplitude f_{NL}^B of the theoretical bispectrum shape B , in the limit of weak non-Gaussianity and under the assumption of statistical homogeneity and isotropy.

For clarity we wrote the amplitude of the bispectrum shape B as $f_{\text{NL}} = f_{\text{NL}}^B$.

The bispectrum B in the estimator (2.46) is the bispectrum of the theoretical model that we want to test, e.g. a local or equilateral shape bispectrum. To distinguish different models of non-Gaussianity let us assume that the correct model has a bispectrum given by B_{correct} , i.e. $\langle \delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \delta_{\mathbf{k}_3} \rangle \equiv (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B_{\text{correct}}(k_1, k_2, k_3)$, if we average over an ensemble of universes or independent patches of our universe. This correct bispectrum B_{correct} may of course be different from the theoretical bispectrum B whose amplitude we want to estimate in (2.46). The average of the estimator for f_{NL} is therefore given by

$$\langle \hat{f}_{\text{NL}}^B \rangle = \frac{1}{N_{f_{\text{NL}}}} \frac{V}{\pi} \int_{\mathcal{V}_B} dk_1 dk_2 dk_3 \frac{k_1 k_2 k_3 B_{\text{correct}}(k_1, k_2, k_3) B(k_1, k_2, k_3)}{P_\delta(k_1) P_\delta(k_2) P_\delta(k_3)}, \quad (2.47)$$

where \mathcal{V}_B is the tetrahedral domain allowed by the triangle condition for the wavenumbers k_i , $V = L^3 = (2\pi)^3 \delta_D(\mathbf{0})$ and we integrated over the angular parts of the \mathbf{k}_i [50]. Demanding $\langle \hat{f}_{\text{NL}}^B \rangle = 1$ for $B = B_{\text{correct}}$ fixes the normalisation to

$$N_{f_{\text{NL}}} = \frac{V}{\pi} \int_{\mathcal{V}_B} dk_1 dk_2 dk_3 \frac{k_1 k_2 k_3 [B(k_1, k_2, k_3)]^2}{P_\delta(k_1) P_\delta(k_2) P_\delta(k_3)}. \quad (2.48)$$

Equation (2.47) motivates the definition of an inner product between two different bispectra by [58], [50]

$$\langle B_i, B_j \rangle_{\text{est}} \equiv \frac{V}{\pi} \int_{\mathcal{V}_B} dk_1 dk_2 dk_3 \frac{k_1 k_2 k_3 B_i(k_1, k_2, k_3) B_j(k_1, k_2, k_3)}{P_\delta(k_1) P_\delta(k_2) P_\delta(k_3)}, \quad (2.49)$$

which can be normalised to a number between -1 and 1 by defining the shape correlation or shape cosine

$$\mathcal{C}(B_i, B_j) \equiv \frac{\langle B_i, B_j \rangle_{\text{est}}}{\sqrt{\langle B_i, B_i \rangle_{\text{est}} \langle B_j, B_j \rangle_{\text{est}}}}. \quad (2.50)$$

If the shape correlation \mathcal{C} of a theoretical bispectrum B and the correct bispectrum B_{correct} is equal to one, then (2.46) is an optimal estimator for the full amplitude f_{NL}^B . If instead $\mathcal{C}(B, B_{\text{correct}}) = 0$, we will measure $\hat{f}_{\text{NL}}^B = 0$ (in average). Therefore if two theoretical bispectra B_1 and B_2 have a small inner product, the optimal estimator for $f_{\text{NL}}^{B_1}$ will perform badly in detecting $f_{\text{NL}}^{B_2}$ and vice versa [58].

A similar optimal estimator can be derived for the trispectrum (see e.g. [50, 54]). The terms involving $B\delta^3$ and $P\delta$ (in the presence of inhomogeneous noise) for the bispectrum estimator turn into terms of the form $T\delta^4$, $P\delta^2$ and P^2 for the trispectrum estimator. In the context of CMB lensing reconstruction, the power spectrum of the reconstructed lensing potential corresponds to the $T\delta^4$ term, while $P\delta^2$ corresponds to a data-dependent bias correction (the empirical $\hat{N}^{(0)}$ in Chapter 5) and P^2 corresponds to a data-independent ($N^{(0)}$) bias correction. Notably, just as the linear $P\delta$ term was initially overlooked for bispectrum estimation, the $P\delta^2$ term was only recently included in the context of lensing reconstruction in [59] (for other reasons). However, current CMB bispectrum estimations and CMB lensing reconstructions use these correction terms to obtain optimal constraints (e.g. [3, 20]). Note that we will comment on the dominant Gaussian contribution to the variance of the estimated bispectrum in Chapter 4 below (see Eqs. (4.35) and (4.57)).

A bispectrum amplitude can be estimated very efficiently with (2.46) if the bispectrum shape is product-separable, i.e. if it can be written as a sum of terms of the form $f_1(k_1)f_2(k_2)f_3(k_3)$ for some 1D functions f_i . This was first appreciated for CMB applications by Komatsu, Spergel and Wandelt in [60]. Recently, this KSW-estimator has been applied in the context of the ISW-lensing bispectrum [61] and in the context of constraining specific shapes of primordial non-Gaussianity from WMAP (see [62] for the final WMAP9 results) and from Planck [20].

2.2 CMB perturbations

In this section we relate the primordial curvature fluctuations to CMB temperature perturbations with the Boltzmann equation, closely following the excellent pedagogical expositions in [7, 8, 10, 11, 12]. Throughout Section 2.2 we adopt the sign convention of Ref. [10] for the potentials (and for the metric, which we keep as $\eta_{\mu\nu} = \text{diag}(1, -1, -1, -1)$ throughout this chapter).

2.2.1 CMB anisotropies from metric perturbations

Primordial fluctuations imply metric perturbations which perturb the energy of observed CMB photons, giving rise to temperature anisotropies. If we add scalar perturbations to the homogeneous and isotropic FLRW metric and work in the conformal Newtonian gauge, the linearly perturbed metric in coordinates $x^\mu = (\eta, \mathbf{x})$ is given by

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu = a^2(\eta) [(1 + 2\Psi(\mathbf{x}, \eta)) d\eta^2 - (1 - 2\Phi(\mathbf{x}, \eta)) d\mathbf{x}^2], \quad (2.51)$$

where the position-dependent potentials Ψ and Φ break the homogeneity and isotropy of the background, which only depends on the conformal time η (recall $d\eta = dt/a$). The scale factor $a(\eta)$ describes the expansion of the background. If there is no anisotropic stress, then $\Psi = \Phi$ and the geodesic motion in a space-time with metric (2.51) is equivalent to Newtonian gravity with the gravitational potential given by Φ .

The Boltzmann equation involves the photon distribution function $f(x^\mu, p^\mu)$, which depends on time, position and 4-momentum p^μ . We parameterise the photon momentum by the energy E measured by an observer who does not move relative to our spatial coordinates \mathbf{x} . Let us first assume an unperturbed universe ($\Psi, \Phi = 0$). Then the comoving observer has the 4-velocity $u_\mu = \sqrt{g_{00}} \delta_\mu^0$ and measures the photon energy to be $E = p^\mu u_\mu = ap^0$, so that $p^0 = E/a$. If we specify the photon propagation direction by the 3-vector $\hat{\mathbf{p}}$ (normalised such that $\delta_{ij}\hat{p}^i\hat{p}^j = 1$), then $p^2 = 0$ for the photon implies $p^\mu = \frac{E}{a}(1, \hat{\mathbf{p}})$.

If we turn to a perturbed universe ($\Psi, \Phi \neq 0$) and work at first order in perturbations, imposing $E = p^\mu u_\mu = a(1 + \Psi)p^0$ gives $p^0 = \frac{E}{a}(1 - \Psi)$. Writing $p^i = C\hat{p}^i$, $p^2 = 0$ requires

$$p^\mu = \frac{\varepsilon}{a^2}(1 - \Psi, (1 + \Phi)\hat{\mathbf{p}}). \quad (2.52)$$

Here we introduced the comoving photon energy $\varepsilon \equiv aE$ which is conserved in absence of perturbations.¹

In absence of collisions the photon distribution function is conserved along

¹Heuristically, the wavelength is stretched, $E \propto 1/\lambda \propto 1/a$. More formally, $E \propto 1/a$ follows from the time-component of the geodesic equation in an unperturbed universe, or from the zeroth-order Boltzmann equation.

the path of the photons in phase space. The path is determined by the geodesic equation

$$\frac{dp^\mu}{d\lambda} + \Gamma_{\nu\rho}^\mu p^\nu p^\rho = 0, \quad (2.53)$$

where we introduced the affine parameter λ such that $p^\mu = dx^\mu/d\lambda$. The Christoffel symbol is defined by

$$\Gamma_{\nu\rho}^\mu \equiv \frac{1}{2} g^{\mu\lambda} (\partial_\nu g_{\rho\lambda} + \partial_\rho g_{\lambda\nu} - \partial_\lambda g_{\nu\rho}) \quad (2.54)$$

and can be computed from the metric (2.51) (e.g. [7, 8]). Combining $p^0 = dx^0/d\lambda$ with (2.52) relates λ to conformal time η so that we can replace λ in favor of η in the geodesic equation. At first order, after some algebra we find (e.g. [7, 10])

$$\frac{1}{\varepsilon} \frac{d\varepsilon}{d\eta} = -\frac{d\Psi}{d\eta} + \partial_\eta(\Phi + \Psi), \quad (2.55)$$

$$\frac{d\hat{p}^i}{d\eta} = -(\delta^{ij} - \hat{p}^i \hat{p}^j) \partial_j(\Phi + \Psi), \quad (2.56)$$

where total derivatives are along the photon path, $d/d\eta = \partial_\eta + \hat{p}^i \partial_i$. According to (2.55) the evolution of the comoving energy is given by changes of Ψ along the photon path and explicit time variations of the metric perturbations Φ and Ψ . Note that the background expansion changes the *non*-comoving photon energy E but not the comoving energy ε . Equation (2.56) describes gravitational lensing of the photons: The change of the propagation direction is given by the gradient of the potentials in the direction perpendicular to the photon path.

The photon distribution function f is given by a perturbed Bose-Einstein distribution which we write as

$$f(\eta, \mathbf{x}, \varepsilon, \hat{\mathbf{p}}) = \left[\exp \left(\frac{\varepsilon}{a(\eta)T(\eta)(1 + \Theta(\eta, \mathbf{x}, \hat{\mathbf{p}}))} \right) - 1 \right]^{-1}. \quad (2.57)$$

In absence of metric perturbations Ψ, Φ , the temperature perturbation Θ is zero and the zero-order distribution function is a function $f^{(0)}(\varepsilon)$ of ε alone (note that $\varepsilon/(aT)$ does not depend explicitly on time). The first-order correction for $\Psi, \Phi \neq 0$

can be written as¹

$$f^{(1)}(\eta, \mathbf{x}, \hat{\mathbf{p}}) = \frac{df}{d\Theta} \Theta = -\frac{df^{(0)}}{d\varepsilon} \varepsilon \Theta. \quad (2.58)$$

At first order, the collisionless Boltzmann equation is therefore

$$0 = \frac{df}{d\eta} = \frac{df^{(0)}}{d\varepsilon} \frac{d\varepsilon}{d\eta} + \frac{df^{(1)}}{d\eta} = \frac{df^{(0)}}{d\varepsilon} \frac{d\varepsilon}{d\eta} - \frac{df^{(0)}}{d\varepsilon} \varepsilon \frac{d\Theta}{d\eta}. \quad (2.59)$$

We did not write down zero order terms and we neglected second-order corrections to the derivative of $f^{(1)}$. The total derivative of Θ in the last term is

$$\frac{d\Theta}{d\eta} = \left(\partial_\eta + \hat{p}^i \partial_i + \frac{d\hat{p}^i}{d\eta} \partial_{\hat{p}^i} \right) \Theta. \quad (2.60)$$

The last term in the brackets can be neglected because the change of propagation direction along the photon path is already first order in the perturbations; see (2.56). Then, the evolution of the temperature perturbation along the photon path due to metric perturbations is at first order given by (e.g. [10])

$$\frac{d\Theta}{d\eta} = (\partial_\eta + \hat{p}^i \partial_i) \Theta = -\frac{d\Psi}{d\eta} + \partial_\eta(\Phi + \Psi), \quad (2.61)$$

where we used (2.60) in the first step and (2.59) and (2.55) in the second step. The observed temperature perturbation can be obtained by integrating (2.61) from recombination until today along the line of sight. The first term on the r.h.s. of (2.61) means that, relative to an unperturbed universe, a photon is redshifted if the potential Ψ is deeper at recombination than today. Therefore overdensities at recombination are converted to cold spots in the CMB (the photon has to climb out of a deeper potential well). The last two terms on the r.h.s. of (2.61) come from explicit time variations of the metric perturbations, which vanish on all scales during matter domination. However, they are relevant at recombination when the universe is still not fully matter dominated, leading to the early integrated Sachs-Wolfe (ISW) effect. The late ISW effect is due to dark energy domination at late times which leads to a decay of the potentials. If the potentials become shallower during the passage of a photon, $\partial_\eta(\Phi + \Psi) > 0$, the photon will gain more energy

¹For the last step, calculate the zero order part of $df/d\Theta$ and compare with $df^{(0)}/d\varepsilon$.

when it enters the potential well than the amount of energy it loses when it escapes from the potential well again, leading to a net blueshift corresponding to a hot spot in the CMB.

2.2.2 Scattering of photons and electrons

We have so far neglected any interactions between the photons and charged particles. However, before recombination, photons Thomson-scatter efficiently off non-relativistic free electrons, because the corresponding scattering rate is high compared to the Hubble expansion rate. The corresponding change of the photon distribution function is described by the collision terms (e.g. [10])

$$\left(\frac{1}{a} \frac{d\Theta(\eta, \mathbf{x}, \hat{\mathbf{p}})}{d\eta} \right)_{\text{collision}} \approx \Gamma_T(-\Theta + \Theta_0 + \hat{\mathbf{p}} \cdot \mathbf{v}_b), \quad (2.62)$$

where $\Gamma_T = \sigma_T n_e$ is the Thomson scattering rate, which is the product of the Thomson-scattering cross-section σ_T and the electron number density n_e . The quantity \mathbf{v}_b in the last term is the peculiar bulk velocity of the electrons. For simplicity we replaced the $1+\cos^2\theta$ dependence on the scattering angle θ by its angular average $4/3$, and we introduced the monopole of the temperature perturbation,

$$\Theta_0 \equiv \int \frac{d^2 \hat{\mathbf{p}}}{4\pi} \Theta(\eta, \mathbf{x}, \hat{\mathbf{p}}). \quad (2.63)$$

The first term in (2.62) describes how photons scatter from the phase space element around $(\eta, \mathbf{x}, \epsilon, \hat{\mathbf{p}})$ to some other phase space element (out of the beam, reducing the distribution function). The second term describes how photons with arbitrary propagation direction $\hat{\mathbf{p}}$ scatter such that they align with $\hat{\mathbf{p}}$ (into the beam, increasing the distribution function). The last term in (2.62) is a Doppler term which accounts for the energy change of photons when they scatter off moving electrons.

Adding the changes of the temperature perturbation due to gravity (2.61) and scattering (2.62) gives

$$\frac{d\Theta}{d\eta} \equiv \partial_\eta \Theta + \hat{p}^i \partial_i \Theta = -\frac{d\Psi}{d\eta} + \partial_\eta (\Phi + \Psi) + a\sigma_T n_e (-\Theta + \Theta_0 + \hat{\mathbf{p}} \cdot \mathbf{v}_b). \quad (2.64)$$

It is convenient to express the Thomson prefactor by $a\sigma_T n_e = -\partial_\eta \tau(\eta, \eta_R) \equiv -\tau'$, where the prime represents the partial derivative ∂_η with respect to conformal time η from now on, and

$$\tau(\eta, \eta_R) \equiv \int_\eta^{\eta_R} d\tilde{\eta} a n_e \sigma_T \quad (2.65)$$

is the optical depth as η goes back in time. Keeping the reception time η_R fixed, τ increases monotonically from zero at $\eta = \eta_R$ to large values before decoupling. If we switch to Fourier space,

$$\Theta(\eta, \mathbf{x}, \hat{\mathbf{p}}) = \int \frac{d^3 \mathbf{k}}{(2\pi)^3} e^{i\mathbf{k}\cdot\mathbf{x}} \Theta(\eta, k, \mu) \quad (2.66)$$

(noting that the Fourier transform of (2.64) depends on \mathbf{k} and $\hat{\mathbf{p}}$ only through $k \equiv |\mathbf{k}|$ and $\mu \equiv \hat{\mathbf{k}} \cdot \hat{\mathbf{p}}$ for vorticity-free \mathbf{v}_b), we get for a single \mathbf{k} mode

$$\Theta' + ik\mu\Theta - \tau'\Theta = -ik\mu\Psi + \Phi' - \tau'\Theta_0 - i\tau'\mu v_b, \quad (2.67)$$

where $\Psi = \Psi(\eta, k)$ and similarly for Φ , Θ_0 and v_b . Note that we used $d/d\eta \rightarrow \partial_\eta + ik\mu$, and $\hat{\mathbf{p}} \cdot \mathbf{v}_b \rightarrow \hat{\mathbf{p}} \cdot i\hat{\mathbf{k}}v_b = i\mu v_b$.

2.2.3 Boltzmann hierarchy

To simplify (2.67) further, note that the dependence of the r.h.s. on the photon propagation direction μ is very simple. We therefore expand $\Theta(\eta, k, \mu)$ in Legendre polynomials $P_l(\mu)$ (which have a simple dependence on μ and are orthogonal on the domain $\mu \in [-1, 1]$) as

$$\Theta(\eta, k, \mu) = \sum_l (2l+1)(-i)^l \Theta_l(\eta, k) P_l(\mu), \quad (2.68)$$

where the expansion coefficients Θ_l are

$$\Theta_l(\eta, k) = \frac{i^l}{2} \int_{-1}^1 d\mu \Theta(\eta, k, \mu) P_l(\mu). \quad (2.69)$$

2.2 CMB perturbations

If we use this expansion in (2.67), replace the term involving $\mu\Theta$ using the recurrence relation

$$\mu P_l(\mu) = \frac{l+1}{2l+1} P_{l+1}(\mu) + \frac{l}{2l+1} P_{l-1}(\mu), \quad (2.70)$$

multiply both sides of (2.67) by $P_{l'}$ and integrate over μ , we get the Boltzmann hierarchy (e.g. [12])

$$\Theta'_l - k \frac{l}{2l+1} \Theta_{l-1} + k \frac{l+1}{2l+1} \Theta_{l+1} - \tau' \Theta_l = \delta_{l0}(\Phi' - \tau' \Theta_0) + \frac{1}{3} \delta_{l1}(k\Psi + \tau' v_b), \quad (2.71)$$

where we used $P_0 = 1$ and $P_1 = \mu$. This is a system of coupled ordinary differential equations for the temperature perturbation coefficients $\Theta_l(\eta, k)$, where any coefficient is directly coupled to its two neighbouring coefficients. To close the system we need to include the evolution equations for dark matter, baryons and metric perturbations and specify the initial conditions. In practice, since the source term on the right of (2.71) only contributes for $l = 0$ and $l = 1$ (neglecting the $l = 2$ contribution from the angular dependence of Thomson scattering), we can solve for all quantities on the right of (2.71) by truncating the hierarchy at $l \sim 10$, because higher multipoles have only a small backreaction on the $l = 0, 1$ coefficients through a chain of coupled neighbouring coefficients. Since we then also know all quantities on the r.h.s. of (2.67), we can directly integrate (2.67) along the line of sight to obtain the full $\Theta(\eta, k, \mu)$, corresponding to Θ_l for arbitrary high l , without the need to solve (2.71) up to the high $l \sim \mathcal{O}(10^3)$ required for current CMB experiments.

Indeed, once the r.h.s. of (2.71) and therefore the r.h.s. of (2.67) are known, the evolution equation (2.67) can be solved by the line of sight integral¹

$$\begin{aligned} \Theta(\eta_R) &= e^{-ik\mu(\eta_R - \eta_I)} e^{-\tau(\eta_I, \eta_R)} \Theta(\eta_I) \\ &+ e^{-ik\mu\eta_R} \int_{\eta_I}^{\eta_R} d\eta [-ik\mu\Psi + \Phi' - \tau' \Theta_0 - i\tau' \mu v_b] e^{-\tau(\eta, \eta_R)} e^{ik\mu\eta}, \end{aligned} \quad (2.72)$$

where η_I is some early initial time and η_R is the time when the photon is received by the observer. The term with the initial condition $\Theta(\eta_I)$ in (2.72) can be neglected

¹To confirm that (2.72) is really a solution of (2.67) compute $\partial\Theta(\eta_R)/\partial\eta_R$ and note that the left-most exponential factors in (2.72) generate the terms on the l.h.s. of (2.67), while the η -dependent parts of the exponentials in the integrand of (2.72) are due to the r.h.s. of (2.67).

by choosing η_I to be a very early time before decoupling so that $\tau \rightarrow \infty$. To obtain the Legendre coefficients of the solution, we express $\mu e^{ik\mu\eta}$ in (2.72) as $-(i/k)\partial_\eta e^{ik\mu\eta}$ and integrate by parts to get (up to boundary terms which only contribute to the monopole and dipole [12, 63])

$$\Theta(\eta_R) = \int_{\eta_I}^{\eta_R} d\eta e^{-ik\mu(\eta_R-\eta)} S(\eta, k), \quad (2.73)$$

where the sources S are

$$S(\eta, k) = (\Psi' + \Phi')e^{-\tau(\eta, \eta_R)} + \left(\Psi + \Theta_0 - \frac{v'_b}{k} \right) \partial_\eta e^{-\tau} - \frac{v_b}{k} \partial_\eta^2 e^{-\tau}. \quad (2.74)$$

Using the plane-wave expansion

$$e^{-ik\mu(\eta_R-\eta)} = \sum_l i^l (2l+1) j_l(k(\eta - \eta_R)) P_l(\mu) \quad (2.75)$$

we get for the Legendre coefficients

$$\Theta_l(\eta_R, k) = \int_{\eta_I}^{\eta_R} d\eta S(\eta, k) j_l(k(\eta - \eta_R)). \quad (2.76)$$

Note that the exponentials in the source function (2.74) imply that the ISW-term $\Psi' + \Phi'$ gets contributions along the whole line of sight from decoupling until today (because, neglecting reionisation, $e^{-\tau} \approx 1$ for $z \lesssim z_{\text{dec}}$ and $e^{-\tau} \approx 0$ for $z \gtrsim z_{\text{dec}}$), while the other terms are only relevant around decoupling (because $\partial_\eta e^{-\tau}$ and $\partial_\eta^2 e^{-\tau}$ are strongly peaked at decoupling; see e.g. [12]). The line-of-sight approach described above has been implemented in CMBfast [63], CAMB [64] and CLASS [65].

Note that in order to solve the Boltzmann hierarchy up to multipole $l \sim 10$ we need to calculate the sources in (2.74). We also need to take into account that $4\Theta_0 = \delta_\gamma$ and $\Theta_1 \propto v_\gamma$, where v_γ is the photon bulk velocity. These perturbations are part of the photon energy momentum tensor which is related to metric perturbations via the Einstein equations. If we assume that until recombination photons and baryons are tightly coupled via Thomson scattering such that $v_b = v_\gamma$, energy-momentum conservation of the composite photon-baryon fluid in combina-

tion with the continuity equation for the photons implies for adiabatic fluctuations that the evolution of the photon density is given by a damped harmonic oscillator equation sourced by metric perturbations (see e.g. [7, 8, 10]):

$$\delta_\gamma'' + \frac{\mathcal{H}R}{1+R}\delta_\gamma' + c_s^2 k^2 \delta_\gamma = 4\Phi'' + \frac{4\mathcal{H}R}{1+R}\Phi' - \frac{4}{3}k^2\Psi, \quad (2.77)$$

where $R = 3\bar{\rho}_b/(4\bar{\rho}_\gamma)$, $\mathcal{H} = aH$ and the sound speed is given by $c_s^2 = 1/[3(1+R)]$. This equation describes acoustic oscillations of the baryon-photon fluid, whose pressure is due to the photons while its inertia comes from the baryon mass.

For adiabatic initial conditions from inflation all Fourier modes \mathbf{k} with the same wavenumber k oscillate in phase, reaching their maxima and minima at the same time [66]. At the time of recombination, when the oscillations end and photons start free-streaming, all Fourier modes $\delta_\gamma(\mathbf{k})$ with k such that they have reached their first extremum at the time of recombination lead to the first acoustic peak of the CMB power spectrum. Modes $\delta_\gamma(\mathbf{k})$ with somewhat larger k such that they are zero at recombination lead to the first acoustic trough, etc. The separation of the acoustic peaks and troughs is related to the sound horizon, which is given by the distance $\int d\eta c_s$ that the sound wave travels from early times until recombination. On small scales, photon diffusion washes out structures and leads to the exponential diffusion damping of the small-scale CMB power spectrum. Despite this small-scale suppression Planck was able to measure the first seven acoustic peaks [2] and small-scale experiments like ACT [67] and SPT [68] push to even smaller angular scales.

2.2.4 CMB power spectrum

The line-of-sight integration approach outlined above provides the solution of the photon temperature perturbation Θ for a given set of initial conditions from inflation. In general, this solution can be written in terms of transfer functions Δ_l as (e.g. [69])

$$\Theta_l(\mathbf{k}) = 4\Delta_l(k)\Phi_{\text{init}}(\mathbf{k}), \quad (2.78)$$

where e.g. Δ_0 provides the transfer function for δ_γ while Δ_1 represents the transfer function for v_γ . From the expansion coefficients Θ_l we get the full temperature

anisotropies $\Delta T/\bar{T} = \Theta$ by using (2.68),

$$\frac{\Delta T(\hat{\mathbf{n}}, \mathbf{k})}{\bar{T}} = \sum_l (2l+1)(-i)^l \Delta_l(k) \Phi_{\text{init}}(\mathbf{k}) \frac{4\pi}{2l+1} \sum_m Y_{lm}(\hat{\mathbf{n}}) Y_{lm}^*(\hat{\mathbf{k}}), \quad (2.79)$$

where we used

$$P_l(\mu) = \frac{4\pi}{2l+1} \sum_m Y_{lm}(\hat{\mathbf{n}}) Y_{lm}^*(\hat{\mathbf{k}}). \quad (2.80)$$

Equation (2.79) is an expansion in spherical harmonics $Y_{lm}(\hat{\mathbf{n}})$ with expansion coefficients given by

$$a_{lm}(\mathbf{x}) = (-i)^l 4\pi \int \frac{d^3k}{(2\pi)^3} e^{i\mathbf{k}\cdot\mathbf{x}} \Delta_l(k) \Phi_{\text{init}}(\mathbf{k}) Y_{lm}^*(\hat{\mathbf{k}}), \quad (2.81)$$

where we transformed back to real space using (2.66), noting that \mathbf{x} is the position of the observer receiving CMB photons from different directions on the sky.

We cannot predict the realisation $\Phi(\mathbf{k})$ of the primordial fluctuations but only their statistical properties. If we assume Gaussian statistically-homogeneous and isotropic perturbations with power spectrum

$$\langle \Phi_{\text{init}}(\mathbf{k}) \Phi_{\text{init}}^*(\mathbf{k}') \rangle = (2\pi)^3 \delta_D(\mathbf{k} - \mathbf{k}') P_\Phi(k), \quad (2.82)$$

the a_{lm} are also Gaussian random variables with covariance given by¹

$$\begin{aligned} \langle a_{lm}(\mathbf{x}) a_{l'm'}^*(\mathbf{x}) \rangle &= (-i)^{l-l'} (4\pi)^2 \int \frac{d^3k}{(2\pi)^6} \Delta_l^2(k) (2\pi)^3 P_\Phi(k) Y_{lm}^*(\hat{\mathbf{k}}) Y_{l'm'}^*(\hat{\mathbf{k}}) \\ &= \delta_{ll'} \delta_{mm'} \underbrace{\frac{2}{\pi} \int dk k^2 \Delta_l^2(k) P_\Phi(k)}_{\equiv C_l^{TT}}. \end{aligned} \quad (2.83)$$

In the second line we integrated over directions of \mathbf{k} using the orthonormality of the spherical harmonics. According to (2.83) the power spectrum C_l^{TT} of the temperature anisotropies is given by the product of the primordial power spectrum and the square of a transfer function, integrated over wavenumbers k . The diagonality of the covariance of the a_{lm} reflects the fact that statistical homogeneity and

¹In the first line the product of the exponentials from (2.81) is unity because of (2.82).

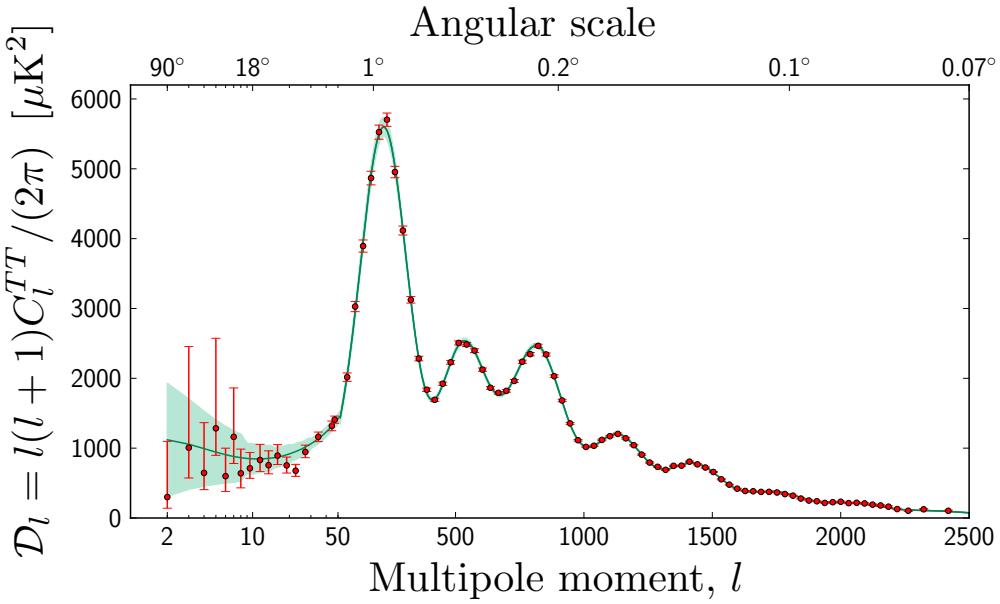


Figure 2.2: CMB temperature power spectrum measured by Planck (red points), including cosmic variance error (shaded green) and best-fit Λ CDM model prediction (solid green). The plot is from [70].

isotropy of the primordial fluctuations are preserved by the subsequent evolution.

Figure 2.2 shows the temperature power spectrum for the best-fit Λ CDM model obtained from the recent Planck data [2, 70]. The model agrees very well with the data (plotted as red points).

2.3 CMB lensing

We have so far neglected the gravitational deflection of CMB photons by intervening large-scale structures on their way from the last-scattering surface to the observer, because the RMS deflection of CMB photons is only ~ 2.5 arcmin, which is a small effect. However, these deflections are coherent over scales of several degrees and the effect of lensing on the CMB power spectrum is large enough that it must be included for high-resolution CMB experiments like Planck to obtain accurate cosmological parameter estimates. Indeed, lensing of the CMB can be exploited as a precious probe of the inhomogeneous distribution of dark matter along the line of sight which is otherwise hard to observe. Since most of the lens-

ing effect is caused by dark matter structures around redshift $z \sim 2$, CMB lensing can be used to break degeneracies that affect the primary CMB to improve constraints on spatial curvature, neutrino masses, dark energy and modified gravity (see e.g. [71, 72, 73, 74, 75, 76, 77, 78], and Section 2.3.3 below). For example, recent lensing reconstructions provide evidence for dark energy from the CMB alone. Measurements of lensing can also be used to study the bias relation between the dark matter distribution and directly observable tracers of large-scale structures like galaxies or quasars (see Section 2.3.3.3 below).

2.3.1 Basics

2.3.1.1 CMB lensing potential

The lensed CMB temperature anisotropies $\tilde{T}(\hat{\mathbf{n}})$, defined as perturbations to the mean temperature, can be described by a remapping of the unlensed CMB anisotropies $T(\hat{\mathbf{n}})$ by the deflection angle $\alpha(\hat{\mathbf{n}})$:

$$\tilde{T}(\hat{\mathbf{n}}) = T(\hat{\mathbf{n}} + \alpha(\hat{\mathbf{n}})). \quad (2.84)$$

The deflection angle can be written as the transverse gradient of the *lensing potential* $\phi(\hat{\mathbf{n}})$, which is given by an integral over the matter distribution along the line of sight (under the Born approximation that the integral can be taken along the unperturbed photon path; see e.g. [79, 80]):

$$\alpha(\hat{\mathbf{n}}) = \nabla\phi(\hat{\mathbf{n}}), \quad (2.85)$$

$$\phi(\hat{\mathbf{n}}) = -2 \int_0^{\chi_*} d\chi \frac{f_K(\chi_* - \chi)}{f_K(\chi)f_K(\chi_*)} \Psi_W(\chi\hat{\mathbf{n}}; \eta_R - \chi). \quad (2.86)$$

Here χ denotes comoving distances, χ_* is the comoving distance to last scattering and η_R is the conformal time at reception today (i.e. these photons last scattered at conformal time $\eta_R - \chi_*$). We describe the matter along the line of sight by the ‘Weyl potential’ $\Psi_W \equiv (\Psi + \Phi)/2$, where Ψ and Φ are the scalar perturbations about the FLRW metric in the conformal Newtonian gauge as defined in (2.51). Depending on the spatial curvature K of the universe, the angular diameter distance is given

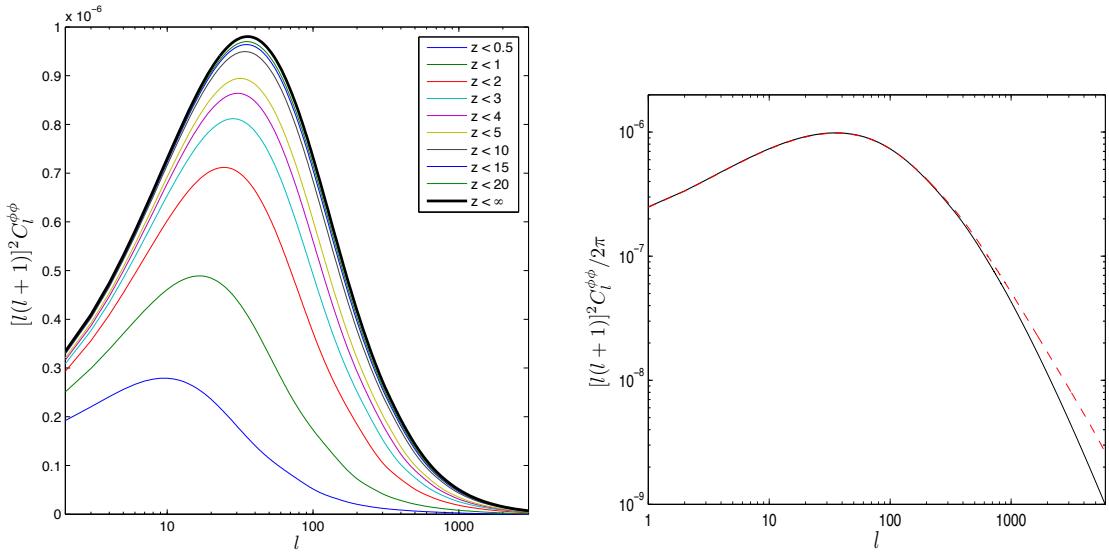


Figure 2.3: *Left:* Cumulative contributions from different redshifts to the theoretical lensing power spectrum (2.89) in a concordance Λ CDM model. *Right:* Linear deflection power spectrum (solid) and non-linear corrections from HALOFIT [81] (dashed). Both plots are taken from [79].

by

$$f_K(\chi) = \begin{cases} \frac{1}{\sqrt{K}} \sin(\sqrt{K}\chi) & \text{if } K > 0 \quad (\text{closed}) \\ \chi & \text{if } K = 0 \quad (\text{flat}) \\ \frac{1}{\sqrt{|K|}} \sinh(\sqrt{|K|}\chi) & \text{if } K < 0 \quad (\text{open}) . \end{cases} \quad (2.87)$$

We see that lensing explicitly depends on the geometry of the universe and can be used to constrain the spatial curvature (note that the evolution of Ψ_W is also affected). We will come back to this in Section 2.3.3.

2.3.1.2 Lensing power spectrum

Assuming a flat universe from now on, $K = 0$, and linear transfer of primordial perturbations \mathcal{R} to Ψ ,

$$\Psi(\mathbf{k}, \eta) = T_\Psi(k, \eta) \mathcal{R}(\mathbf{k}), \quad (2.88)$$

with linear transfer function T_Ψ , the lensing power spectrum is given by [79]

$$C_l^{\phi\phi} = 16\pi \int \frac{dk}{k} \mathcal{P}_{\mathcal{R}}(k) \left[\int_0^{\chi_*} d\chi T_{\Psi}(k, \eta_R - \chi) j_l(k\chi) \left(\frac{\chi_* - \chi}{\chi_* \chi} \right) \right]^2. \quad (2.89)$$

Figure 2.3 shows that the lensing kernel is most sensitive to structures around redshift $z \sim 2$. Therefore, CMB lensing can be used to measure the growth of structures around this epoch and constrain late-time parameters such as e.g. the dark energy density parameter today, Ω_Λ , or the sum of the neutrino masses, $\sum m_\nu$. The cosmological information probed by CMB lensing will be discussed in more detail in Section 2.3.3. Figure 2.3 also shows that non-linear corrections to the lensing power spectrum are small up to l of a few hundred, but should be included on smaller scales (which is done in CAMB).

It is convenient to plot $[l(l+1)]^2 C_l^{\phi\phi}/(2\pi)$ on a logarithmic l -scale because this shows the contribution of different scales to the RMS deflection angle (see e.g. [57, 80]):

$$\langle |\alpha|^2 \rangle \approx \int d \ln l \frac{[l(l+1)]^2 C_l^{\phi\phi}}{2\pi}. \quad (2.90)$$

2.3.1.3 Lensing-temperature correlation from ISW-effect

The Weyl potential Ψ_W , which describes the deflection of CMB photons along the line of sight, evolves with time. This time-evolution influences the CMB temperature anisotropies through the ISW-effect (see (2.61)). This leads to a correlation $C^{T\phi}$ of the temperature anisotropies and the lensing potential. Since the late ISW-effect is important for the large-scale temperature fluctuations but contributes only little to intermediate- and small-scale fluctuations, $C^{T\phi}$ is very small on all but the largest scales, see Fig. 2.4. Since lensing mostly affects intermediate- and small-scale temperature multipoles, this correlation is generally negligible for the purposes of this thesis. We neglect it in the rest of this thesis except where stated otherwise.

2.3.1.4 Lensed CMB temperature

CMB photons are remapped according to (2.84). To get some intuition for the effects of lensing on the CMB temperature anisotropies we first work in the flat-sky approximation, which approximates the two-dimensional sphere of the observed

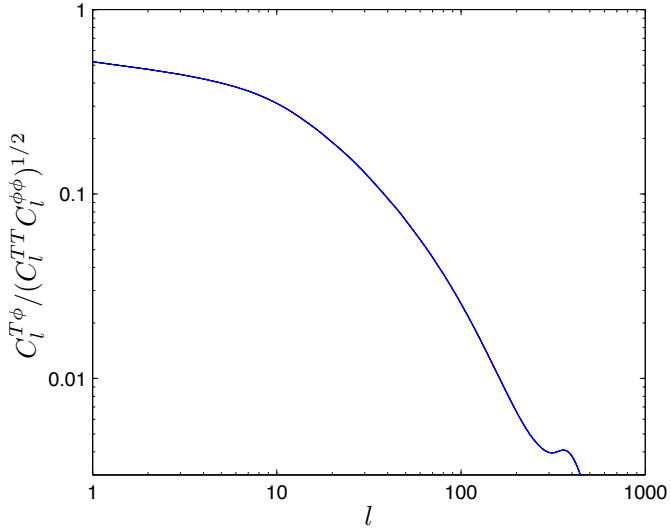


Figure 2.4: Typical theoretical correlation of temperature fluctuations and lensing potential due to the ISW-effect (from [79]).

sky as a flat two-dimensional plane. This is a good approximation on small scales. A generalisation to the full sky will be used in Chapter 5. Under the flat-sky approximation we can consider the two-dimensional Fourier transform of the temperature field $T(\mathbf{x})$ on the plane

$$T(\mathbf{x}) = \int \frac{d^2\mathbf{l}}{2\pi} T(\mathbf{l}) e^{i\mathbf{l}\cdot\mathbf{x}}, \quad T(\mathbf{l}) = \int \frac{d^2\mathbf{x}}{2\pi} T(\mathbf{x}) e^{-i\mathbf{l}\cdot\mathbf{x}}. \quad (2.91)$$

To see how the CMB is changed by lensing we Taylor expand the lensed temperature $\tilde{T}(\mathbf{x})$ around $\phi(\mathbf{x}) = 0$

$$\tilde{T}(\mathbf{x}) = T(\mathbf{x} + \nabla\phi(\mathbf{x})) = T(\mathbf{x}) + \nabla T(\mathbf{x}) \cdot \nabla\phi(\mathbf{x}) + \mathcal{O}(\phi^2). \quad (2.92)$$

In harmonic space this becomes

$$\tilde{T}(\mathbf{l}) = T(\mathbf{l}) - \underbrace{\int \frac{d^2\mathbf{l}'}{2\pi} \mathbf{l}' T(\mathbf{l}') \cdot (\mathbf{l} - \mathbf{l}') \phi(\mathbf{l} - \mathbf{l}')}_{\delta T(\mathbf{l})} + \mathcal{O}(\phi^2), \quad (2.93)$$

i.e. the lensed temperature at multipole \mathbf{l} depends on a range of unlensed temperature multipoles \mathbf{l}' due to lensing. This lensing-induced mode coupling is most

effective for CMB modes separated by $\delta l = |\mathbf{l} - \mathbf{l}'| \sim 50$ because the lensing power spectrum peaks at this scale. The mode coupling implies that fixed lenses ϕ introduce an anisotropy in the CMB so that *for fixed lenses* the off-diagonal CMB covariance $\langle \tilde{T}(\mathbf{l})\tilde{T}^*(\mathbf{l}') \rangle_{\text{CMB}}$ (averaged only over the unlensed CMB) depends on the lensing potential ϕ . This allows us to reconstruct the *realisation* of ϕ from the observed CMB temperature \tilde{T} (see Section 2.3.1.6). Note that the off-diagonal CMB covariance would vanish if we took the ensemble average over the unlensed CMB and the lenses, because both are statistically-isotropic and homogeneous random fields.

2.3.1.5 Effect of lensing on the CMB temperature power spectrum

The leading-order correction to the temperature power spectrum due to lensing is of order ϕ^2 , because $\langle \phi \rangle = 0$ if we average over the CMB and the lenses. If $\delta^n T$ denotes the change of the temperature at n -th order in the lensing potential, these corrections are of the form $\langle \delta T \delta T^* \rangle$, $\langle T \delta^2 T^* \rangle$ and $\langle \delta^2 T T^* \rangle$. Explicitly, from (2.93) we have

$$\langle \delta T(\mathbf{l}) \delta T^*(\mathbf{l}') \rangle = \delta(\mathbf{l} - \mathbf{l}') \int \frac{d^2 \mathbf{L}}{(2\pi)^2} C_L^{TT} C_{|\mathbf{l}-\mathbf{L}|}^{\phi\phi} [\mathbf{L} \cdot (\mathbf{l} - \mathbf{L})]^2. \quad (2.94)$$

From [79, 82]

$$\delta^2 T(\mathbf{l}) = -\frac{1}{2} \int \frac{d^2 \mathbf{l}_1}{2\pi} \frac{d^2 \mathbf{l}_2}{2\pi} [\mathbf{l}_1 \cdot (\mathbf{l}_1 + \mathbf{l}_2 + \mathbf{l})] [\mathbf{l}_1 \cdot \mathbf{l}_2] T(\mathbf{l}_1) \phi(\mathbf{l}_2) \phi^*(\mathbf{l}_1 + \mathbf{l}_2 - \mathbf{l}) \quad (2.95)$$

we get

$$\langle \tilde{T}(\mathbf{l}) \delta^2 \tilde{T}^*(\mathbf{l}') \rangle = -\frac{1}{2} \delta(\mathbf{l} - \mathbf{l}') C_l^{TT} \int \frac{d^2 \mathbf{L}}{(2\pi)^2} (\mathbf{l} \cdot \mathbf{L})^2 C_L^{\phi\phi}. \quad (2.96)$$

Noting that the angular part of the last integral is just π and adding the $\langle \delta^2 T T^* \rangle$ term gives

$$C_l^{\tilde{T}\tilde{T}} = C_l^{TT} + \int \frac{d^2 \mathbf{L}}{(2\pi)^2} C_L^{TT} C_{|\mathbf{l}-\mathbf{L}|}^{\phi\phi} [\mathbf{L} \cdot (\mathbf{l} - \mathbf{L})]^2 - R^\phi l^2 C_l^{TT}, \quad (2.97)$$

where the orientation of \mathbf{l} can be picked arbitrarily because we integrate over \mathbf{L} . We defined half of the total RMS deflection [79]

$$R^\phi \equiv \frac{1}{2} \langle |\nabla\phi|^2 \rangle = \frac{1}{4\pi} \int \frac{dL}{L} L^4 C_L^{\phi\phi} \sim 3 \times 10^{-7}. \quad (2.98)$$

The integral in the middle term in (2.97) is a convolution of the unlensed temperature power spectrum with the lensing power spectrum which smoothes the acoustic peaks and troughs over scales $\delta l \sim 50$. This can be seen in Fig. 2.5 which shows the changes of the temperature power due to lensing (on the full sky). Note that the last term in (2.97) cancels the overall scaling of the middle term in the acoustic range but it does not cancel the smoothing effect because it does not involve any convolution [57].

On very small scales, $l \gtrsim 2000$, the unlensed CMB is exponentially suppressed because of diffusion damping. Therefore, to a good approximation, only the C_L^{TT} modes with relatively low $L \ll l$ will contribute to the lensing changes in (2.97), simplifying the lensed temperature power spectrum to [57, 79, 82, 83]

$$C_l^{\tilde{T}\tilde{T}} \approx C_l^{\phi\phi} \int \frac{d^2\mathbf{L}}{(2\pi)^2} [\mathbf{L} \cdot \mathbf{l}]^2 C_L^{TT} = l^2 C_l^{\phi\phi} \int \frac{dL}{L} \frac{(L)^4 C_L^{TT}}{4\pi} = l^2 C_l^{\phi\phi} R^T, \quad (2.99)$$

where R^T is half of the RMS gradient of the unlensed CMB (defined analogously to R^ϕ in (2.98)). The right-hand side of (2.99) is positive and leads to an increase of CMB power on small scales $l \gtrsim 2000$ as visible in Fig. 2.5. This means that the lensing convolution transfers CMB power from large (non-exponentially suppressed) scales to small (otherwise exponentially suppressed) scales.

In summary, lensing smoothes the acoustic peaks and troughs of the temperature power spectrum and transfers power from large to small scales in the damping tail.

It should be noted that the expansion in the deflection angle used in this section provides useful intuition for the effect of lensing on the CMB power spectrum but it is in fact wrong by up to 5% at temperature multipoles $2000 \lesssim l \lesssim 5000$, where a non-perturbative calculation provides more accurate results [79, 84, 85, 86].

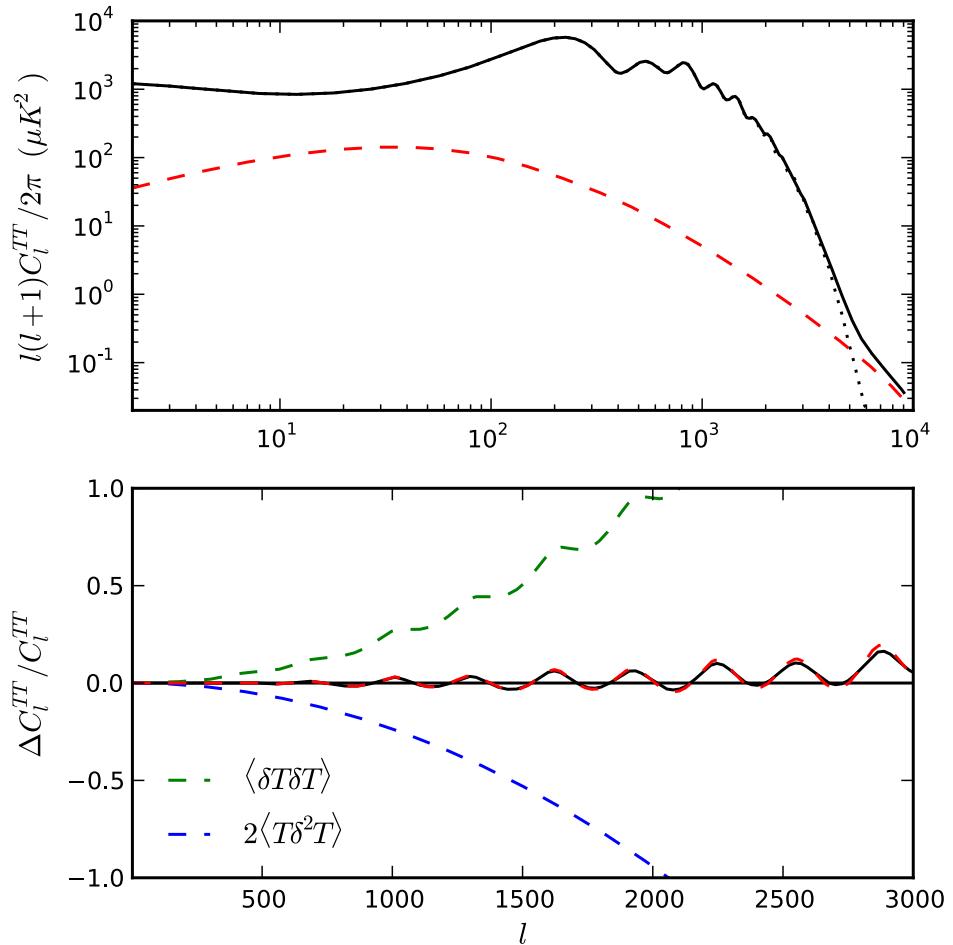


Figure 2.5: *Top:* Lensed (solid) and unlensed (dotted) temperature power spectrum. The red dashed curve shows the small-scale approximation of Eq. (2.99). *Bottom:* Fractional change of the temperature power spectrum due to lensing (solid black). The first order result (2.97) (red dashed) and its two constituent terms (blue and green dashed) are also shown. Both plots are from [57].

2.3.1.6 CMB lensing reconstruction

Although we can only observe the lensed CMB directly and do not know the unlensed CMB, it is possible to reconstruct the lensing potential $\phi(\hat{\mathbf{n}})$ from the lensed CMB by exploiting the fact that lensing changes the (well-known) statistical properties of the temperature anisotropies, introducing a local anisotropy of the CMB for a fixed realisation of lenses. With the reconstructed lensing potential we can, in principle, de-lens the lensed CMB using the inverse mapping of (2.84) to get the unlensed CMB anisotropies. In this section we review the reconstruction method, working in the flat-sky approximation for simplicity (see Section 5 for the generalisation to the full sky).

Let us consider a single fixed lensing mode $\phi(\mathbf{L})$ which we aim to reconstruct from the lensed temperature \tilde{T} . If we average over the unlensed CMB and over all lensing modes $\phi(\mathbf{L}')$ except $\phi(\mathbf{L})$, we find [87]

$$\langle \tilde{T}(\mathbf{l})\tilde{T}^*(\mathbf{l}') \rangle_{\mathbf{L}' \neq \mathbf{L}} = \delta_D(\mathbf{l} - \mathbf{l}')C_l^{TT} - \frac{1}{2\pi}\delta_D(\mathbf{l} - \mathbf{l}' - \mathbf{L})\mathbf{L}\phi(\mathbf{L}) \cdot \left[-\mathbf{l}C_l^{\tilde{T}\tilde{T}} + \mathbf{l}'C_{l'}^{\tilde{T}\tilde{T}} \right]. \quad (2.100)$$

This is correct at first order in $\phi(\mathbf{L})$ and at all orders in $\phi(\mathbf{L}')$ for $\mathbf{L}' \neq \mathbf{L}$.¹ As noticed already above, the lensing mode $\phi(\mathbf{L})$ couples CMB modes separated by $\mathbf{l} - \mathbf{l}' = \mathbf{L}$. Note that (2.100) contains *lensed* power spectra in the square brackets, accounting for non-perturbative corrections to the earlier leading-order results in [88, 89] which involved unlensed spectra.

Here and in the following \tilde{T} multipoles denote the observed, beam-deconvolved noisy CMB temperature with power spectrum

$$C_{l,\text{expt}}^{\tilde{T}\tilde{T}} = C_l^{\tilde{T}\tilde{T}} + \left(\frac{\sigma_N}{T_{\text{CMB}}} \right)^2 e^{l(l+1)\sigma_{\text{FWHM}}^2/(8\ln 2)}, \quad (2.101)$$

where $C_l^{\tilde{T}\tilde{T}}$ is the theoretical lensed power spectrum and σ_N and σ_{FWHM} characterise the detector noise level and the full width at half-maximum (FWHM) of the beam. The form of (2.101) is obtained by first convolving the CMB power spectrum with the beam by multiplying with $\exp[-l(l+1)\sigma_{\text{FWHM}}^2/(8\ln 2)]$, then

¹More correctly the approximation $\tilde{C}_l^{T\nabla T} \approx C_l^{\tilde{T}\tilde{T}}$ was used, where $\tilde{C}_l^{T\nabla T}$ is defined in [87]. This holds to about one percent [87].

adding white noise with power spectrum $(\sigma_N/T_{\text{CMB}})^2$ to the result, and finally deconvolving the beam by multiplying everything with $\exp[l(l+1)\sigma_{\text{FWHM}}^2/(8\ln 2)]$.

Noting that the realisation of the lensing potential $\phi(\mathbf{L})$ enters the covariance (2.100) of CMB modes \mathbf{l} and $\mathbf{l}' = \mathbf{l} - \mathbf{L}$, we aim to reconstruct $\phi(\mathbf{L})$ by summing over all products $\tilde{T}(\mathbf{l})\tilde{T}^*(\mathbf{l} - \mathbf{L})$ with arbitrary weight $\tilde{g}(\mathbf{l}, \mathbf{L})$ and normalisation $A(\mathbf{L})$ [89]:

$$\hat{\phi}(\mathbf{L}) = A(\mathbf{L}) \int \frac{d^2\mathbf{l}}{2\pi} \tilde{T}(\mathbf{l})\tilde{T}^*(\mathbf{l} - \mathbf{L})\tilde{g}(\mathbf{l}, \mathbf{L}). \quad (2.102)$$

The normalisation $A(\mathbf{L})$ can be chosen such that the estimator is unbiased, i.e. $\langle \partial\hat{\phi}(\mathbf{L})/\partial\phi(\mathbf{L}') \rangle = \delta_{\mathbf{L},\mathbf{L}'}$. Taking the average first over the unlensed CMB and all lensing modes $\mathbf{L}'' \neq \mathbf{L}'$ using (2.100), and then averaging over the remaining lensing mode \mathbf{L}' gives (for $\mathbf{L} \neq 0$)

$$\left\langle \frac{\partial\hat{\phi}(\mathbf{L})}{\partial\phi(\mathbf{L}')} \right\rangle = A(\mathbf{L})\delta_{\mathbf{L},\mathbf{L}'} \int \frac{d^2\mathbf{l}}{(2\pi)^2} \mathbf{L}' \cdot \left[\mathbf{l}C_l^{\tilde{T}\tilde{T}} + (\mathbf{L} - \mathbf{l})C_{|\mathbf{L}-\mathbf{l}|}^{\tilde{T}\tilde{T}} \right] \tilde{g}(\mathbf{l}, \mathbf{L}).$$

The normalisation is therefore

$$A(\mathbf{L}) = \left(\int \frac{d^2\mathbf{l}}{(2\pi)^2} \mathbf{L} \cdot \left[\mathbf{l}C_l^{\tilde{T}\tilde{T}} + (\mathbf{L} - \mathbf{l})C_{|\mathbf{L}-\mathbf{l}|}^{\tilde{T}\tilde{T}} \right] \tilde{g}(\mathbf{l}, \mathbf{L}) \right)^{-1}, \quad (2.103)$$

for any weight $\tilde{g}(\mathbf{l}, \mathbf{L})$. To optimise this weight we minimise the variance $\langle \hat{\phi}(\mathbf{L})\hat{\phi}^*(\mathbf{L}') \rangle$ of the estimator. From (2.102) we get

$$\begin{aligned} \langle \hat{\phi}(\mathbf{L})\hat{\phi}^*(\mathbf{L}') \rangle &= A(\mathbf{L})A(\mathbf{L}') \int \frac{d^2\mathbf{l}_1}{2\pi} \frac{d^2\mathbf{l}_2}{2\pi} \tilde{g}(\mathbf{l}_1, \mathbf{L})\tilde{g}(\mathbf{l}_2, \mathbf{L}') \\ &\quad \times \langle \tilde{T}(\mathbf{l}_1)\tilde{T}^*(\mathbf{l}_1 - \mathbf{L})\tilde{T}^*(\mathbf{l}_2)\tilde{T}(\mathbf{l}_2 - \mathbf{L}') \rangle. \end{aligned} \quad (2.104)$$

The leading-order contribution comes from the Gaussian part of the 4-point function, which is

$$\begin{aligned} \langle \tilde{T}(\mathbf{l}_1)\tilde{T}^*(\mathbf{l}_1 - \mathbf{L})\tilde{T}^*(\mathbf{l}_2)\tilde{T}(\mathbf{l}_2 - \mathbf{L}') \rangle &= \\ \delta_D(\mathbf{l}_1 - \mathbf{l}_2)\delta_D(-(\mathbf{l}_1 - \mathbf{L}) + (\mathbf{l}_2 - \mathbf{L}'))C_{l_1, \text{expt}}^{\tilde{T}\tilde{T}}C_{|\mathbf{L}-\mathbf{l}_1|, \text{expt}}^{\tilde{T}\tilde{T}} + 2 \text{ perms.} \end{aligned} \quad (2.105)$$

2.3 CMB lensing

For $\mathbf{L}, \mathbf{L}' \neq 0$ this leads to

$$\langle \hat{\phi}(\mathbf{L}) \hat{\phi}^*(\mathbf{L}') \rangle = 2A(\mathbf{L})^2 \delta_D(\mathbf{L} - \mathbf{L}') \int \frac{d^2 \mathbf{l}_1}{2\pi} C_{l_1, \text{expt}}^{\tilde{T}\tilde{T}} C_{|\mathbf{L}-\mathbf{l}_1|, \text{expt}}^{\tilde{T}\tilde{T}} \tilde{g}(\mathbf{l}_1, \mathbf{L})^2 \quad (2.106)$$

after integrating over \mathbf{l}_2 . Calculating the derivative with respect to the weight \tilde{g} (subject to the constraint (2.103)) and setting it to zero leads to

$$\tilde{g}(\mathbf{l}, \mathbf{L}) = \frac{\mathbf{L} \cdot \mathbf{l} C_l^{\tilde{T}\tilde{T}} + \mathbf{L} \cdot (\mathbf{L} - \mathbf{l}) C_{|\mathbf{L}-\mathbf{l}|}^{\tilde{T}\tilde{T}}}{C_{l, \text{expt}}^{\tilde{T}\tilde{T}} C_{|\mathbf{L}-\mathbf{l}|, \text{expt}}^{\tilde{T}\tilde{T}}}.$$

We may rescale \tilde{g} by a constant factor because such a rescaling is absorbed by $A(\mathbf{L})$, given by (2.103), when we evaluate the estimator (2.102). For later convenience we choose

$$\tilde{g}(\mathbf{l}, \mathbf{L}) = \frac{\mathbf{L} \cdot \mathbf{l} C_l^{\tilde{T}\tilde{T}} + \mathbf{L} \cdot (\mathbf{L} - \mathbf{l}) C_{|\mathbf{L}-\mathbf{l}|}^{\tilde{T}\tilde{T}}}{2C_{l, \text{expt}}^{\tilde{T}\tilde{T}} C_{|\mathbf{L}-\mathbf{l}|, \text{expt}}^{\tilde{T}\tilde{T}}}, \quad (2.107)$$

because this simplifies the variance of $\hat{\phi}$ to $\langle |\hat{\phi}(\mathbf{L})|^2 \rangle = \delta_D(\mathbf{0}) A(\mathbf{L})$ (at zeroth order in the lensing potential if we neglect the implicit dependence of the lensed power on the lensing potential). The delta function is $\delta_D(\mathbf{0}) = f_{\text{sky}}/\pi$.¹ Note that if we substitute the weight \tilde{g} from (2.107) back into the normalisation $A(\mathbf{L})$ given by (2.103) we see that $A(\mathbf{L})$ is actually a function of $|\mathbf{L}| = L$, because \mathbf{L} only appears as \mathbf{L}^2 or $\mathbf{l} \cdot \mathbf{L}$ and we integrate over \mathbf{l} . Therefore we will write $A(L)$ in the following.

The reconstructed lensing potential $\hat{\phi}(\mathbf{L})$ from (2.102) can be used to study the realisation of the large-scale structures by which the CMB photons are lensed, or to delense the observed CMB anisotropies. For comparison with theoretical expectations it is useful to estimate the power spectrum of the reconstructed lensing potential,

$$C_L^{\hat{\phi}\hat{\phi}} = [\delta_D(\mathbf{0})]^{-1} \frac{1}{n_L} \sum_{\mathbf{L}, |\mathbf{L}|=L} |\hat{\phi}(\mathbf{L})|^2, \quad (2.108)$$

where n_L is the number of vectors \mathbf{L} with length L and the inverse of the delta function comes from the definition of the power spectrum. The reconstructed lensing potential power spectrum has a signal contribution from the connected part of

¹From $(2\pi)^2 \delta_D(\mathbf{l}) = \int_{A_{\text{sky}}} d\mathbf{x} e^{i\mathbf{l} \cdot \mathbf{x}}$ we get $\delta_D(\mathbf{0}) = A_{\text{sky}}/(2\pi)^2 = f_{\text{sky}}/\pi$, where $f_{\text{sky}} = A_{\text{sky}}/(4\pi)$ is the fraction of the sky observed.

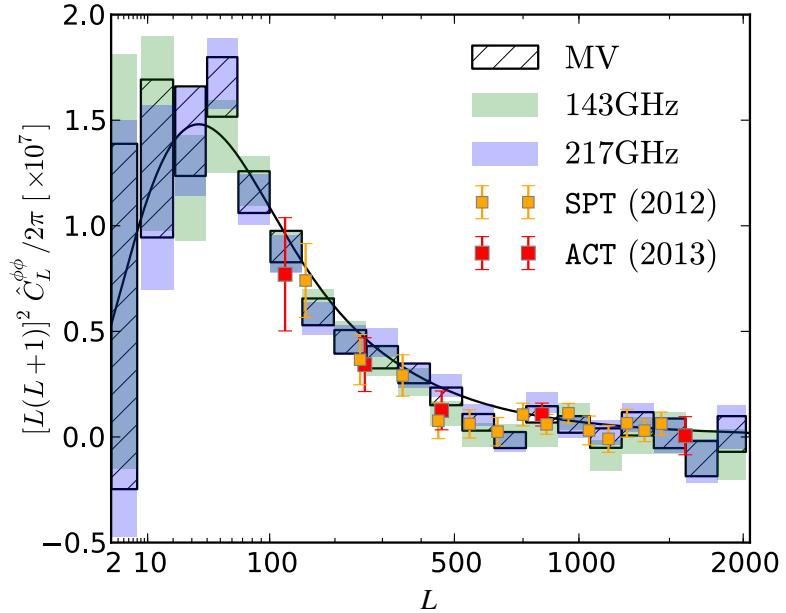


Figure 2.6: Theoretical lensing power spectrum $C_L^{\phi\phi}$ (black line) and reconstructions from Planck (143GHz, 217 GHz and minimum variance ‘MV’ combination; [3]), SPT [94] and ACT [93]. The plot is from [3].

the lensed CMB 4-point function. After subtracting the (typically large) Gaussian noise bias, this can be compared with the theoretical lensing power spectrum (from e.g. CAMB [64]) to test the theoretical prediction and to improve cosmological parameter estimates.

2.3.2 Observational status

Gravitational lensing of the CMB has recently made the transition from first detections to high quality signals with increasing importance for observational cosmology. CMB lensing was first detected by cross-correlating lensing reconstructions from WMAP data with other tracers of large-scale structure [90, 91]. First detections from the CMB alone were recently presented by ACT [92, 93], SPT [94] and Planck [3] with significances of 4.6, 6.3 and 25σ , respectively; see Fig. 2.6. Further improvements are expected from the full 2500 deg^2 SPT survey and the full-mission Planck data. This data also allows for polarization-based lensing reconstruction, which was performed for the first time very recently in [95] (using

SPT data). Less significant detections of lensing can be obtained from the effect of lensing on the CMB temperature power spectrum. Recent detections of this effect are at the 10σ level [2, 96].

2.3.3 Cosmological information probed by CMB lensing

2.3.3.1 Geometric degeneracy of the primary CMB

If the comoving size of the sound horizon at decoupling, λ_1 , and the redshift of decoupling, z_{dec} , are fixed (corresponding to fixed physical densities $\Omega_b h^2$ and $\Omega_c h^2$), the angular degree-scale separation under which the acoustic peaks appear in the observed temperature power spectrum depends on the curvature of the universe and on the expansion rate between decoupling and today, $H(z)$, which is in turn sensitive to Ω_m and Ω_Λ (note that $\Omega_r h^2$ is fixed by $T_0 = 2.73\text{K}$). Indeed, if we write out $H(z)$ explicitly in terms of the density parameters Ω_X , the multipole of the n -th acoustic peak is [12] (setting $a_0 = 1$ today)

$$l_n \approx \frac{\pi}{\theta_n} \approx \frac{2\pi D_A(z_{\text{dec}})}{\lambda_n} = \frac{n\pi f_K \left(\frac{1}{H_0} \int_0^{z_{\text{dec}}} \frac{dz}{[\Omega_r(z+1)^4 + \Omega_m(z+1)^3 + \Omega_\Lambda + \Omega_K(z+1)^2]^{1/2}} \right)}{\frac{1}{H_0} \int_{z_{\text{dec}}}^\infty \frac{c_s dz}{[\Omega_r(z+1)^4 + \Omega_m(z+1)^3 + \Omega_\Lambda + \Omega_K(z+1)^2]^{1/2}}}, \quad (2.109)$$

where θ_n is the angular scale of the n -th acoustic peak, $D_A(z)$ is the angular diameter distance to an event at redshift z and $\lambda_n = 2\pi/k_n$ is the comoving wavelength associated with the wavenumber k_n of the modes that lead to the n -th acoustic peak [$k_n r_s(\eta_{\text{rec}}) = k_n \int_0^{\eta_{\text{rec}}} c_s d\eta = n\pi$ for adiabatic fluctuations; $n = 1$ labels modes that underwent one compression after horizon entry and are maximal at decoupling]. Keeping the sound horizon in the denominator of (2.109) fixed, we see from (2.87) that the fixed physical scale λ_1 appears at lower multipole l_1 in a closed universe ($K > 0$) than in a flat universe ($K = 0$); but it appears at higher multipole l_1 if there is less dark energy and dark matter (and the universe expands slower today and is older). This is the *angular diameter distance degeneracy* or *geometric degeneracy*: The peak positions are determined by the three quantities $\Omega_b h^2$, $\Omega_m h^2$ and $D_A(z_{\text{dec}})$, but these involve four independent cosmological parameters, Ω_b , Ω_m , Ω_Λ and H_0 (any one of them can be replaced by $\Omega_K = 1 - \Omega_\Lambda - \Omega_m$). These four parameters can therefore not be determined simultaneously from the

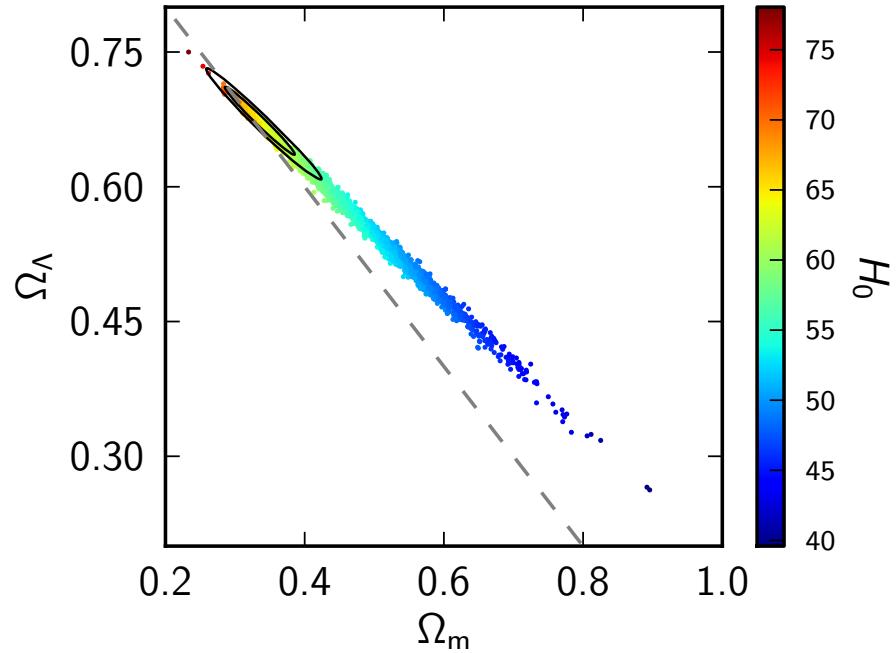


Figure 2.7: Illustration of the geometric degeneracy (from [3]). Colored points show Monte-Carlo samples for combined Planck temperature, WMAP polarization and ‘highL’ ground-based temperature data (*Planck+WP+highL*). The color indicates H_0 . The effect of lensing on the CMB power spectrum is taken into account (otherwise the degeneracy would extend to $\Omega_\Lambda = 0$). The black lines show 68% and 95% confidence contours if the Planck lensing reconstruction power spectrum (which probes the CMB trispectrum) is included using the lensing likelihood based on Chapter 5 of this thesis. The dashed line indicates flat models, $\Omega_K = 1 - \Omega_m - \Omega_\Lambda = 0$. The upper region has $\Omega_m + \Omega_\Lambda > 1$, i.e. $\Omega_K < 0$.

primary CMB. In particular, the primary CMB cannot constrain the curvature of the universe and it can only give marginal evidence for the existence of dark energy (if curved models are considered). This is illustrated in Fig. 2.7 where colored points correspond to parameter values in the Ω_m - Ω_Λ plane allowed by primary CMB data. The degeneracy follows a line which is slightly tilted with respect to the diagonal $\Omega_K = 0$ line. It corresponds to constant angular diameter distance $D_A(z_{\text{dec}})$ for fixed physical densities $\Omega_b h^2$ and $\Omega_m h^2$.

2.3.3.2 Breaking the geometric degeneracy: Evidence for dark energy and flatness from the CMB alone

The geometric degeneracy can be broken if the CMB is combined with an independent low-redshift observation of e.g. the Hubble constant, BAO or Type Ia supernova, because this probes a fourth independent parameter combination which then allows simultaneous measurements of Ω_b , Ω_m , Ω_Λ and Ω_K . Alternatively, without resorting to astrophysical non-CMB data, CMB lensing reconstruction can be used to probe the low-redshift universe and break the geometric degeneracy (see [3, 97, 98, 99] and Fig. 2.3). This is illustrated by the black contours in Fig. 2.7, which give tight constraints on Ω_m and Ω_Λ from the CMB alone. In particular there is clear evidence for dark energy (see [99] for the first demonstration of this using ACT) and for spatial flatness at the percent level from the CMB alone.

Planck constraints (combined with WMAP polarization and ground-based high-l data) for curvature in curved Λ CDM models with and without lensing reconstruction are [2]

$$100\Omega_K = -4.2^{+4.3}_{-4.8} \quad (95\%; Planck+WP+highL), \quad (2.110)$$

$$100\Omega_K = -1.0^{+1.8}_{-1.9} \quad (95\%; Planck+lensing+WP+highL). \quad (2.111)$$

Including lensing reconstruction reduces the errors by a factor of more than two. Note that the first constraint already includes the smoothing effect of lensing on the lensed CMB power spectrum. For *curved* Λ CDM models, Planck finds [3]

$$\Omega_\Lambda = 0.57^{+0.073}_{-0.055} \quad (68\%; Planck+WP+highL), \quad (2.112)$$

$$\Omega_\Lambda = 0.67^{+0.023}_{-0.023} \quad (68\%; Planck+lensing+WP+highL). \quad (2.113)$$

Including the lensing reconstruction improves the error again by a factor of more than two.

The likelihood used by the Planck collaboration to add the lensing reconstruction to the primary CMB data is based on the lensing likelihood advocated and validated in Chapter 5.

2.3.3.3 Neutrino mass, optical depth and LSS bias

CMB lensing can also be used to constrain additional cosmological parameters (see e.g. [100] for an overview). It can put an upper bound on the sum of neutrino masses $\sum m_\nu$ because increasing neutrino mass (while adjusting other parameters to stay consistent with observations) increases the expansion rate at $z \gtrsim 1$, which suppresses sub-horizon clustering probed by CMB lensing. The current combined constraint from the Planck temperature power spectrum, Planck lensing reconstruction, WMAP polarization and ground-based high-l data is [2]

$$\sum m_\nu < 0.85 \text{ eV} \quad (95\%; \textit{Planck+lensing+WP+highL}). \quad (2.114)$$

Future CMB lensing reconstructions should be able to improve this upper bound significantly. This will be very interesting because neutrino oscillations bound the mass-squared differences $m_i^2 - m_j^2$ from below, implying that at least two neutrino species must be massive and $\sum m_\nu \gtrsim 0.06 \text{ eV}$ [101].

During reionisation, luminous sources such as the first stars, galaxies or quasars formed and ionised the neutral hydrogen gas, liberating free electrons. CMB photons Thomson-scatter from these free electrons which leads to a $e^{-2\tau}$ damping of the CMB power spectrum on scales smaller than the Hubble radius at reionisation, where τ is the optical depth. The measured amplitude of small-scale CMB fluctuations is therefore given by $A_s e^{-2\tau}$ where A_s is the amplitude of the primordial scalar fluctuations. This leads to a near-perfect degeneracy of A_s and τ for the primary CMB, which is only partially broken by large-scales because they are affected by cosmic variance (WMAP polarization is included in the above Planck constraints to reduce this degeneracy; for Planck the correlation between A_s and τ is still 94% [70]). However, CMB lensing breaks this degeneracy because it is sensitive to the distribution of dark matter around $z \sim 2$ and is not affected by the τ -dependent damping at reionisation. Combining Planck's CMB temperature power spectrum and lensing reconstruction gives [2]

$$\tau = 0.089 \pm 0.032 \quad (68\%; \textit{Planck+lensing}). \quad (2.115)$$

Note that the WMAP-9 *polarization* data give a consistent constraint with about

twice the accuracy (and much higher accuracy is expected from Planck polarization data), but the constraint from CMB lensing is still useful because it has very different systematics as it relies on temperature data only.

Since CMB lensing depends directly on the dark matter distribution along the line of sight, with a broad kernel that peaks around $z \sim 2$, it can be used to study the bias relation between dark matter and luminous tracers of dark matter such as galaxies or quasars. Understanding the bias relation is crucial for obtaining robust cosmological information from any direct large-scale structure experiment surveying luminous matter. The cross-correlation between the lensing potential derived from WMAP and NVSS has been detected by [90, 91] at the 3σ level. More recently, Ref. [102] obtained a $4\text{-}5\sigma$ detection for the correlation between SPT lensing reconstruction and data from BCS, WISE and Spitzer/IRAC. By cross-correlating with the ACT CMB lensing reconstruction (obtaining a 3.8σ result), Ref. [103] constrained the bias parameter of $z \sim 1.4$ quasars extracted from SDSS DR8 data to be $b = 2.5 \pm 0.6$. Correlations of the Planck lensing reconstruction with large-scale structure data have been detected with roughly 20σ significance for NVSS, 10σ for SDSS LRGs and 7σ for both MaxBCG and WISE [3]. The correlation of Planck's lensing reconstruction with the Cosmic Infrared Background (CIB), emitted from dusty star-forming galaxies and measured with Planck's high-frequency channels, has been detected at 42σ significance [104]. This correlation has also been detected by cross-correlating a lensing reconstruction map derived from SPT data with Herschel/SPIRE maps, reaching significances from 6.7 to 8.8σ in different submillimeter-wavelength maps [105]. Combining recent CMB lensing reconstructions with near-future DES observations is forecasted to give percent-level constraints on the linear galaxy bias [106].

2.3.4 Cosmological parameter estimation with CMB lensing

While lensing has been measured from both the CMB 2- and 4-point function, before the work presented in this thesis it has not been known how independent these measurements are. If the two lensing measurements were perfectly correlated lensing reconstructions from the CMB 4-point function would not contain

any cosmological information which was not already contained in the CMB 2-point function. This would reduce the motivation for performing 4-point lensing reconstructions in the first place. To date, data analyses of lensing reconstructions have assumed the lensing measurements from the 2- and 4-point function to be perfectly *uncorrelated* [3, 94, 99]. Since both power spectra are derived from the same CMB temperature map, one might question the validity of this approach, raising the concern that lensing information is inadvertently being double counted. While early lensing forecasts [71, 72, 107] addressed this by using *unlensed* CMB power spectra, an optimal combination of the observed lensed CMB 2- and 4-point functions should model their cross-covariance.

Intuitively, we expect that two effects could lead to a relevant correlation of the 2- and 4-point lensing estimates. First, if due to cosmic variance some mode of the CMB fluctuations is high in our realisation of the universe, the Gaussian noise of the 4-point lens reconstruction will also be high because it depends on the CMB fluctuations. Cosmic variance of the CMB fluctuations will therefore imply a correlation between 2- and 4-point lensing measurements. We will show in Chapter 5 that this correlation comes from the fully disconnected part of the CMB 6-point function.

The second effect is due to cosmic variance of the dark matter distribution that lenses the CMB: If in our realisation of the universe the lensing potential is high at some scale, the 4-point lensing reconstruction will be high at the same scale and the lensed CMB 2-point function will be smoother, i.e. acoustic peaks will be low and troughs will be high. Therefore acoustic peaks of the CMB 2-point function will be anti-correlated with higher overall 4-point reconstruction, while troughs will be correlated. This effect follows from the connected part of the CMB 6-point function at fourth order in the lensing potential.

As we shall show in Chapter 5, the induced *correlations* are rather small (a few percent) for a Planck-like experiment for broad-band measures of the lensing power such as a lensing power spectrum amplitude. Essentially, this is because there are a limited number of modes of the lensing power spectrum that influence the acoustic part of the temperature power spectrum, and the correlation due to cosmic variance of these modes is diluted by the significant noise due to cosmic variance of the CMB and instrumental noise (i.e. the fact that lensing measurements from

the CMB 2- and 4-point function are not limited by cosmic variance of the lenses). Moreover, the first effect mentioned above produces rather small lensing amplitude correlations since CMB modes at different scales fluctuate independently, and most of the information on peak smearing in the CMB power spectrum comes from modes *near* the acoustic peaks and troughs, whereas the reconstruction of (large-scale) lenses is most effective at CMB scales where the CMB spectrum changes most rapidly, i.e. *between* acoustic peaks and troughs.

An additional effect comes from the late-time ISW effect [108], which accounts for the fact that the ISW effect is sourced by the same late-time metric perturbations as the ones that determine the lensing potential. However, this effect is only relevant on the very largest scales and turns out to be much smaller than the two cosmic variance effects.

We will quantify these arguments with a perturbative calculation of the CMB 6-point function and with simulations of lensed CMB maps in Chapter 5, where we will also assess the relevance of the correlations for the lensing likelihood. The size of the effects is consistent with intuitive plausibility arguments presented in Sec. 5.A.

Once the correlations between the CMB 2-point function and the 4-point lensing reconstruction are understood, we need to model the likelihood if we want to constrain both the CMB and the lenses from the observed CMB. Since both the unlensed CMB and the lensing potential are to a very good approximation Gaussian random fields, one can write down a formal expression for the likelihood of the lensed CMB [109]. However it is not practical to work with this exact likelihood directly given the large number of pixels observed by current CMB datasets. Instead, on small angular scales the data is typically compressed to empirical power spectra which are fast to compute and straightforward to model with an approximate Gaussian likelihood and a covariance which accounts for the noise characteristics of the experiment under consideration. To account for CMB lensing it is desirable to include the 4-point reconstruction as well as the empirical CMB power spectra in the likelihood. In Chapter 5, we will compare different approximate forms of the likelihood for the lensing reconstruction, finding that a Gaussian in the power spectrum of the reconstruction with a model-independent covariance performs well in constraining the amplitude and tilt of the fiducial lensing power spectrum. The

cross-correlation effects between CMB 2- and 4-point function mentioned above can then be incorporated in this likelihood in a straightforward manner.

2.4 Large-scale dark matter perturbations

An alternative observational probe of the primordial fluctuations is the distribution of galaxies whose positions are collected with galaxy surveys. We will briefly sketch how large-scale dark matter structures form from their primordial seeds generated during inflation, closely following [7, 8, 9, 11, 110].

2.4.1 Evolution of perturbations during radiation and matter domination

The evolution of the dark matter density and velocity perturbations δ_c and $k v_c$ is governed by the perturbed Einstein equations. Combining the 00 and ii Einstein equations yields for the evolution of the potential Φ (for $c_s^2 = \delta P / \delta \rho = \text{const}$ and $w = \bar{P} / \bar{\rho} = \text{const}$, neglecting anisotropic stress; see e.g. [8])

$$\ddot{\Phi} + 3(1 + c_s^2)\mathcal{H}\dot{\Phi} + c_s^2 k^2 \Phi = 0. \quad (2.116)$$

We can also exploit conservation of the energy momentum-tensor which gives for cold dark matter with negligible pressure and sound speed:

$$\dot{\delta}_c = -\theta_c + 3\dot{\Phi}, \quad \dot{\theta}_c = -\mathcal{H}\theta_c + k^2\Phi, \quad (2.117)$$

where θ_c is the velocity divergence (i.e. $\theta_c = \nabla \cdot \mathbf{v}_c$ in real space). Taking the time derivative of the first equation and substituting the second gives an evolution equation for the dark matter density for known potential Φ ,

$$\ddot{\delta}_c + \mathcal{H}\dot{\delta}_c = 3\ddot{\Phi} + 3\mathcal{H}\dot{\Phi} - k^2\Phi. \quad (2.118)$$

Solving the evolution equations numerically is straightforward. In some limits it is also possible to find approximate analytical solutions, which are summarised in Table 2.1. The evolution depends on the matter content of the universe, which

Perturbation scale	RD		MD	
	Φ	δ	Φ	δ
$k \gg k_{\text{eq}}$				
super-horizon	const.	const.	—	—
sub-horizon	$\sim a^{-2}$	$\propto \ln a$	const.	$\propto a$
$k \ll k_{\text{eq}}$				
super-horizon	const.	const.	const.	const.
sub-horizon	—	—	const.	$\propto a$

Table 2.1: Approximate analytical limits of the solutions for dark matter density contrast δ_c and potential Φ for small-scale modes with $k \gg k_{\text{eq}}$ which enter the horizon during RD, and for large-scale modes with $k \ll k_{\text{eq}}$ which enter the horizon during MD. The wavenumber k_{eq} corresponds to the mode which enters the horizon at matter radiation equality, $k_{\text{eq}} = a_{\text{eq}} H(a_{\text{eq}})$.

smoothly turns from radiation domination (RD) to matter domination (MD) at $z_{\text{eq}} \sim 3500$, and it is affected by the time at which the modes enter the horizon, which happens when $k = aH$. Both δ and Φ are constant during RD and MD as long as they are super-horizon. If their scales k^{-1} are small enough that they enter the horizon already during RD, radiation pressure prevents effective clustering and implies that the potential Φ oscillates with an amplitude that decays as a^{-2} , and the density only grows logarithmically, $\delta_c \propto \ln a$. During MD there is no pressure which allows the sub-horizon density to grow as $\delta_c \propto a$ and the potential Φ to stay constant (on all scales).

2.4.2 Transfer function

The evolution of the potential Φ from primordial times into MD can be expressed in terms of the transfer function $T_\Phi(k) = \Phi_{\text{late}}(\mathbf{k})/\Phi_{\text{init}}(\mathbf{k})$. For simplicity we neglect corrections to Table 2.1 which describe e.g. the transition from RD to MD (giving rise to a factor of 9/10 for Φ on large scales). In this simplified picture, the transfer function is unity, $T_\Phi(k) = 1$, on *large scales* $k \ll k_{\text{eq}}$ because Φ is always constant. On *small scales*, $k \gg k_{\text{eq}}$, the amplitude of Φ decays as $1/a^2$ during the sub-horizon RD stage. Since $aH \propto 1/a$ during RD, the mode

enters the horizon at $a_{\text{entry}} \propto 1/k$, i.e. earlier for smaller scales (see also Fig. 2.1). Therefore perturbations on smaller scales decay with $\Phi \propto a^{-2}$ for a longer period of time before they become constant at matter-radiation equality. This implies that the transfer function decreases with increasing k . Explicitly, on small scales, $k > k_{\text{eq}}$, we find from $\Phi_{\text{MD}} = \Phi_{\text{init}}(a_{\text{eq}}/a_{\text{entry}})^{-2} \propto \Phi_{\text{init}}k^{-2}$ that $T_\Phi(k) \propto 1/k^2$. Since the density power spectrum contains $T_\Phi^2(k)$ it scales with four powers of k less on small scales than on large scales. Indeed, if we use the sub-horizon Poisson equation during MD (see Eq. (4.2) in Chapter 4 for prefactors),

$$\delta_{c,\text{MD}}(\mathbf{k}) \propto k^2 \Phi_{\text{MD}}(\mathbf{k}), \quad (2.119)$$

we find that

$$P_{\delta,\text{MD}}(k) \propto k^4 P_{\Phi,\text{MD}}(k) \propto k^4 T^2(k) P_{\Phi,\text{prim}}(k) \propto \begin{cases} k, & k \ll k_{\text{eq}}, \\ k^{-3}, & k \gg k_{\text{eq}} \end{cases}, \quad (2.120)$$

where we assumed the primordial power spectrum to scale as k^{-3} , which is close to the observed scaling of $\sim k^{-3.04}$. The approximate scalings in (2.120) and the turn-over at $k_{\text{eq}} \sim 0.02 h/\text{Mpc}$ can be seen in Fig. 2.8, which also illustrates the agreement with observations.

2.4.3 Growth function

At late times, $z \lesssim 10$, all linear modes have entered the horizon and evolve identically independent of k . We describe their scale-independent growth with the growth function $D(a)$ which is proportional to a during MD. When dark energy starts to dominate, the growth is suppressed due to the exponential expansion of the universe. If the growth function is normalised such that $D(a) = a$ deep in MD, one can show that it is given by (see e.g. Eq. (7.77) in [11] or Eq. (2.10) in [112]; this follows from solving the first order parts of (2.127) and (2.129) below)

$$D(a) = \frac{5\Omega_m}{2} \frac{H(a)}{H_0} \int_0^a \frac{d\tilde{a}}{(\tilde{a}H(\tilde{a})/H_0)^3}. \quad (2.121)$$

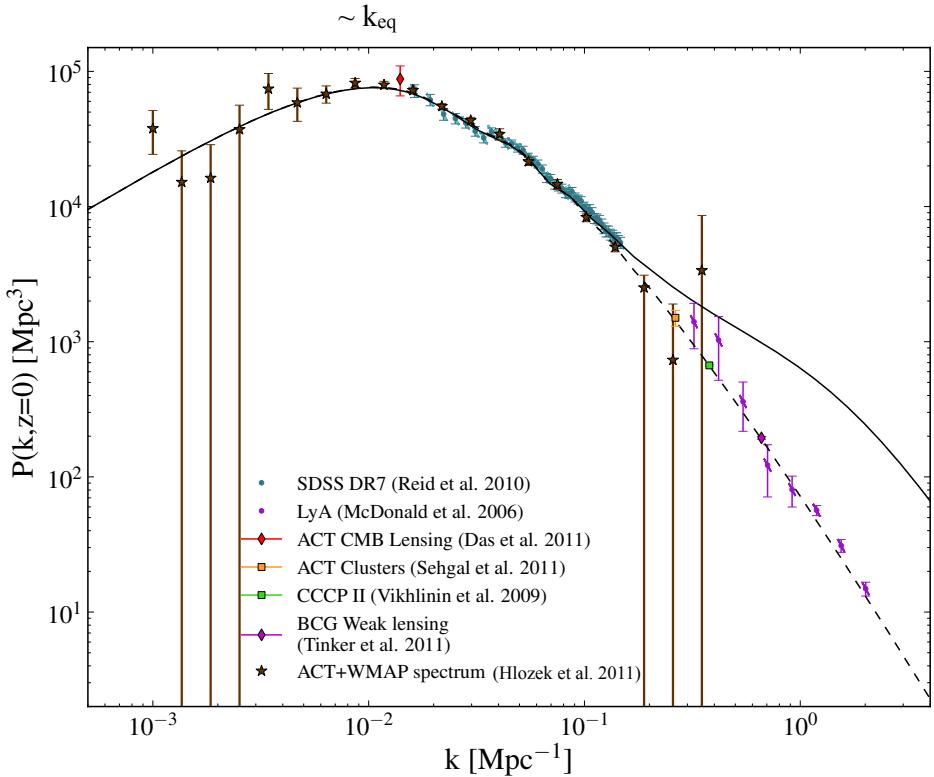


Figure 2.8: Reconstructed matter power spectrum today, adapted from Fig. 5 in [111]. The data points are from various large-scale structure and CMB experiments. The dashed line represents the linear power spectrum which follows approximately (2.120). The solid line contains non-linear corrections which are important on small scales.

The integral can be computed numerically using the Friedmann equation

$$\left(\frac{H(a)}{H_0}\right)^2 = \Omega_r a^{-4} + \Omega_m a^{-3} + \Omega_\Lambda + \Omega_k a^{-2}, \quad (2.122)$$

where Ω_X are the density parameters today (when $a = 1$). Note that we only consider the fastest growing mode here and in the following.

2.4.4 Higher-order perturbation theory

The late-time dark matter distribution has a non-trivial bispectrum even in the absence of primordial non-Gaussianity due to non-linearities in the equations of

motion. Deep in the matter-dominated universe on sub-horizon scales we can use the second-order Newtonian fluid equations (corresponding to the Einstein equations in conformal-Newtonian gauge); see e.g. [7, 8, 41, 110, 113]:

$$\partial_\eta \delta + \nabla \cdot [(1 + \delta) \mathbf{v}] = 0 \quad (\text{continuity}) \quad (2.123)$$

$$\partial_\eta \mathbf{v} + \mathcal{H} \mathbf{v} + (\mathbf{v} \cdot \nabla) \mathbf{v} = -\nabla \Phi \quad (\text{Euler}) \quad (2.124)$$

$$\nabla^2 \Phi = \frac{3}{2} \Omega_m(\eta) \mathcal{H}^2 \delta \quad (\text{Poisson}), \quad (2.125)$$

where \mathbf{v} is the velocity perturbation, assumed to be curl-free with $\theta \equiv \nabla \cdot \mathbf{v}$, and we used the time-dependent¹ density parameter $\Omega_m(\eta) = 8\pi G \bar{\rho}_{\text{DM}} / (3H^2)$ in the Poisson equation. The Fourier transform of $\nabla \cdot [(1 + \delta) \mathbf{v}]$ is

$$\theta(\mathbf{k}, \eta) + \int \frac{d^3 k_1 d^3 k_2}{(2\pi)^3} \delta_D(\mathbf{k} - \mathbf{k}_1 - \mathbf{k}_2) \left[\frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{k_1^2} \theta(\mathbf{k}_1, \eta) \delta(\mathbf{k}_2, \eta) + \theta(\mathbf{k}_1, \eta) \delta(\mathbf{k}_2, \eta) \right], \quad (2.126)$$

where we used $\mathbf{v}(\mathbf{k}_1, \eta) = i\mathbf{k}_1 \theta(\mathbf{k}_1, \eta) / k_1^2$. Therefore the Fourier transform of the continuity equation is given by (see e.g. [41, 110])

$$\partial_\eta \delta(\mathbf{k}, \eta) + \theta(\mathbf{k}, \eta) = - \int \frac{d^3 k_1 d^3 k_2}{(2\pi)^3} \delta_D(\mathbf{k} - \mathbf{k}_1 - \mathbf{k}_2) \alpha(\mathbf{k}, \mathbf{k}_1) \theta(\mathbf{k}_1, \eta) \delta(\mathbf{k}_2, \eta), \quad (2.127)$$

where we defined the mode-coupling function

$$\alpha(\mathbf{k}, \mathbf{k}_1) \equiv \frac{\mathbf{k} \cdot \mathbf{k}_1}{k_1^2}. \quad (2.128)$$

Similarly, the Fourier transform of the divergence of the Euler equation (2.124) is [41, 110]

$$\begin{aligned} \partial_\eta \theta(\mathbf{k}, \eta) + \mathcal{H} \theta(\mathbf{k}, \eta) + \frac{3}{2} \mathcal{H}^2 \Omega_m(\eta) \delta(\mathbf{k}, \eta) = \\ - \int \frac{d^3 k_1 d^3 k_2}{(2\pi)^3} \delta_D(\mathbf{k} - \mathbf{k}_1 - \mathbf{k}_2) \beta(\mathbf{k}, \mathbf{k}_1, \mathbf{k}_2) \theta(\mathbf{k}_1, \eta) \theta(\mathbf{k}_2, \eta), \end{aligned} \quad (2.129)$$

¹So far, Ω_m denoted the density parameter *today*. We will always write out the time-dependence explicitly in the rest of this Chapter if we mean the time-dependent quantity $\Omega_m(\eta)$.

where

$$\beta(\mathbf{k}, \mathbf{k}_1, \mathbf{k}_2) \equiv \frac{k^2(\mathbf{k}_1 \cdot \mathbf{k}_2)}{2k_1^2 k_2^2} \quad (2.130)$$

describes the third term on the left-hand side of (2.124), which is non-linear in \mathbf{v} , and we used (2.125).

To solve (2.127) and (2.129) we start with the first-order parts of the equations, denoting the solutions as δ_1 and θ_1 . The continuity equation (2.127) relates θ_1 to δ_1 ,

$$\theta_1(\mathbf{k}, \eta) = -\partial_\eta \delta_1(\mathbf{k}, \eta) = \frac{\delta_1(\mathbf{k}, \eta)}{D(\eta)} \frac{dD(\eta)}{d\eta} = -f(\eta)\mathcal{H}\delta_1(\mathbf{k}, \eta), \quad (2.131)$$

where we used that the time dependence of δ_1 is given by

$$\delta_1(\mathbf{k}, \eta) = D(\eta)\delta_1(\mathbf{k}), \quad (2.132)$$

and introduced the logarithmic growth rate $f(\eta) \equiv d \ln D / d \ln a$, which is unity during MD but decreases at late times when dark energy domination implies slower growth of the linear density contrast. Due to (2.131) the linear perturbations are completely determined by the linear density contrast δ_1 . To solve (2.127) and (2.129) at higher orders, including the non-linearities on the right-hand sides, we expand the n th-order solutions δ_n and θ_n in powers of the linear solution δ_1 as

$$\begin{aligned} \delta(\mathbf{k}, \eta) &= \sum_{n=1}^{\infty} \int \frac{d^3 q_1}{(2\pi)^3} \cdots \frac{d^3 q_n}{(2\pi)^3} (2\pi)^3 \delta_D \left(\mathbf{k} - \sum_{i=1}^n \mathbf{q}_i \right) \\ &\quad \times F_n^{(s)}(\mathbf{q}_1, \dots, \mathbf{q}_n, \eta) \delta_1(\mathbf{q}_1, \eta) \cdots \delta_1(\mathbf{q}_n, \eta), \end{aligned} \quad (2.133)$$

and

$$\begin{aligned} \theta(\mathbf{k}, \eta) &= -f\mathcal{H} \sum_{n=1}^{\infty} \int \frac{d^3 q_1}{(2\pi)^3} \cdots \frac{d^3 q_n}{(2\pi)^3} (2\pi)^3 \delta_D \left(\mathbf{k} - \sum_{i=1}^n \mathbf{q}_i \right) \\ &\quad \times G_n^{(s)}(\mathbf{q}_1, \dots, \mathbf{q}_n, \eta) \delta_1(\mathbf{q}_1, \eta) \cdots \delta_1(\mathbf{q}_n, \eta), \end{aligned} \quad (2.134)$$

where $F_n^{(s)}$ and $G_n^{(s)}$ are arbitrary kernels, symmetrised over their wavevector arguments. At first order we have $F_1^{(s)} = G_1^{(s)} = 1$. At higher order, substituting

the expansions (2.133) and (2.134) into (2.127) and (2.129) yields recursion relations relating the (unsymmetrised) kernels F_n, G_n to lower order kernels F_m, G_m ($m < n$; see [114, 115], or reviews in e.g. [41, 110, 112]):

$$F_n(\mathbf{q}_1, \dots, \mathbf{q}_n) = \sum_{m=1}^{n-1} \frac{G_m(\mathbf{q}_1, \dots, \mathbf{q}_m)}{(2n+3)(n-1)} \left[(2n+1)\alpha(\mathbf{k}, \mathbf{k}_1)F_{n-m}(\mathbf{q}_{m+1}, \dots, \mathbf{q}_n) + 2\beta(\mathbf{k}, \mathbf{k}_1, \mathbf{k}_2)G_{n-m}(\mathbf{q}_{m+1}, \dots, \mathbf{q}_n) \right], \quad (2.135)$$

$$G_n(\mathbf{q}_1, \dots, \mathbf{q}_n) = \sum_{m=1}^{n-1} \frac{G_m(\mathbf{q}_1, \dots, \mathbf{q}_m)}{(2n+3)(n-1)} \left[3\alpha(\mathbf{k}, \mathbf{k}_1)F_{n-m}(\mathbf{q}_{m+1}, \dots, \mathbf{q}_n) + 2n\beta(\mathbf{k}, \mathbf{k}_1, \mathbf{k}_2)G_{n-m}(\mathbf{q}_{m+1}, \dots, \mathbf{q}_n) \right], \quad (2.136)$$

where $\mathbf{k}_1 = \sum_{i=1}^m \mathbf{q}_i$, $\mathbf{k}_2 = \sum_{i=m+1}^n \mathbf{q}_i$, $\mathbf{k} = \mathbf{k}_1 + \mathbf{k}_2$, $F_1 = G_1 = 1$ and the symmetrised kernels are

$$F_n^{(s)}(\mathbf{q}_1, \dots, \mathbf{q}_n) = \frac{1}{n!} \sum_{\pi} F_n(\mathbf{q}_{\pi(1)}, \dots, \mathbf{q}_{\pi(n)}) \quad (2.137)$$

and similarly for G , where π denotes all permutations of $\{1, \dots, n\}$. Since non-trivial time-dependencies of the kernels F and G would lead to a coupled system of differential equations, the algebraic recursion relations (2.135) and (2.136) can only be obtained because the kernels F_n and G_n turn out to be time-independent for $\Omega_m = 1$, and for other cosmologies their time-dependence is generally negligibly small [41, 116]. We can therefore express the expansions (2.133) and (2.134) in terms of time-independent perturbations $\delta_n(\mathbf{k})$ and $\theta_n(\mathbf{k})$ which are of n -th order in δ_1 ,

$$\delta(\mathbf{k}, \eta) = \sum_{n=1}^{\infty} \delta_n(\mathbf{k}, \eta) = \sum_{n=1}^{\infty} D^n(\eta) \delta_n(\mathbf{k}), \quad (2.138)$$

$$\theta(\mathbf{k}, \eta) = \sum_{n=1}^{\infty} \theta_n(\mathbf{k}, \eta) = -f\mathcal{H} \sum_{n=1}^{\infty} D^n(\eta) \theta_n(\mathbf{k}). \quad (2.139)$$

The first non-trivial kernels can be obtained from (2.135) and (2.137) ($\mathbf{k} =$

$\mathbf{k}_1 + \mathbf{k}_2$; see e.g. [41]):

$$\begin{aligned} F_2^{(s)}(\mathbf{k}_1, \mathbf{k}_2) &= \frac{5}{14} \left[\alpha(\mathbf{k}, \mathbf{k}_1) + \alpha(\mathbf{k}, \mathbf{k}_2) \right] + \frac{2}{7} \beta(\mathbf{k}, \mathbf{k}_1, \mathbf{k}_2) \\ &= \frac{5}{7} + \frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{2} \left(\frac{1}{k_1^2} + \frac{1}{k_2^2} \right) + \frac{2}{7} \frac{(\mathbf{k}_1 \cdot \mathbf{k}_2)^2}{k_1^2 k_2^2}, \end{aligned} \quad (2.140)$$

and

$$G_2^{(s)}(\mathbf{k}_1, \mathbf{k}_2) = \frac{3}{7} + \frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{2} \left(\frac{1}{k_1^2} + \frac{1}{k_2^2} \right) + \frac{4}{7} \frac{(\mathbf{k}_1 \cdot \mathbf{k}_2)^2}{k_1^2 k_2^2}. \quad (2.141)$$

2.4.5 Gravitational bispectrum

Assuming δ_1 to be a Gaussian statistically-homogeneous and isotropic random field with linear power spectrum

$$\langle \delta_1(\mathbf{k}_1, \eta) \delta_1(\mathbf{k}_2, \eta) \rangle = (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2) P_\delta^L(k_1, \eta), \quad (2.142)$$

we find for the leading-order 3-point function, noting $\langle \delta_1^3 \rangle = 0$, that

$$\begin{aligned} \langle \delta(\mathbf{k}_1, \eta) \delta(\mathbf{k}_2, \eta) \delta(\mathbf{k}_3, \eta) \rangle &= \langle (\delta_2(\mathbf{k}_1, \eta) \delta_1(\mathbf{k}_2, \eta) \delta_1(\mathbf{k}_3, \eta)) + 2 \text{ perms} \rangle \\ &= (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B_\delta^{\text{grav}}(k_1, k_2, k_3; \eta), \end{aligned} \quad (2.143)$$

where

$$B_\delta^{\text{grav}}(k_1, k_2, k_3; \eta) = 2P_\delta^L(k_1, \eta) P_\delta^L(k_2, \eta) F_2^{(s)}(\mathbf{k}_1, \mathbf{k}_2) + 2 \text{ perms}. \quad (2.144)$$

This is the leading-order dark matter density bispectrum due to non-linear gravity in the absence of any primordial non-Gaussianity. The shape dependence on the triangle configuration is determined by the F_2 kernel which is directly related to the non-linear terms in the continuity and Euler equation. We will analyse the shape of this bispectrum in more detail in Chapter 4, where we also discuss some extensions beyond leading order and the impact of primordial non-Gaussianity.

Measuring the large-scale structure bispectrum, and correcting for that induced by gravity, can therefore be used to constrain primordial non-Gaussianity. Bispectrum measurements are expected to improve upon power spectrum measure-

ments because the scale-dependent halo-bias is mainly sensitive to the squeezed limit of the primordial bispectrum (which is tightly constrained by the CMB), but the bispectrum can in principle access all triangle configurations of the primordial bispectrum. On sufficiently small scales the halo bispectrum is therefore expected to contain more information than the power spectrum (see e.g. [117, 118] for local non-Gaussianity). Chapters 3 and 4 will discuss the effect of primordial non-Gaussianity on large-scale structure bispectra in more detail.

In the absence of primordial non-Gaussianity there are important late-time motivations for measuring the large-scale structure bispectrum in combination with the power spectrum, e.g. to constrain non-linear bias and break the degeneracy between linear bias and the power spectrum normalisation σ_8 (or Ω_m ; see e.g. [119, 120, 121, 122, 123]). To understand this, we can work with the naive local bias ansatz (e.g. [124]) relating the halo density δ_h to the dark matter density δ for Gaussian initial conditions by

$$\delta_h(\mathbf{x}) = b_1 \delta(\mathbf{x}) + \frac{1}{2} b_2 \delta(\mathbf{x})^2 + \dots, \quad (2.145)$$

where we truncated at second order in δ .¹ The correct bias relation is in fact more complicated and contains e.g. higher order terms, corrections from the tidal tensor [125] and potentially further non-local, stochastic and scale-dependent corrections (especially on small scales, but also on large scales in the presence of primordial non-Gaussianity, see e.g. [117, 126]), but we neglect them in the schematic discussion here. Assuming (2.145), the leading-order halo power spectrum and bispectrum are

$$P_h(k) \approx b_1^2 P_\delta(k), \quad (2.146)$$

$$B_h(k_1, k_2, k_3) \approx b_1^3 B_\delta(k_1, k_2, k_3) + b_1^2 b_2 (P_\delta(k_1) P_\delta(k_2) + 2 \text{ perms}). \quad (2.147)$$

If we only measure P_h , we cannot distinguish a rescaling in b_1 from a rescaling of the amplitude σ_8 of the dark matter power spectrum. However, this degeneracy can be broken when combining with B_h because of its distinctive dependence on

¹Corrections to Eqs. (2.146) and (2.147) below due to a third-order term bias $b_3 \delta^3 / 6$ are suppressed by a factor of $P_\delta \ll 1$ (if we assume that loop integrals are at most of order unity, which is a reasonable assumption on large scales).

triangle configurations (only the first term in (2.147) contains the F_2 kernel which vanishes in the squeezed limit). If one is only interested in testing analytic bias models against simulations instead of relating them to observations, one can also consider mixed bispectra like e.g. the halo-halo-dark-matter bispectrum $B_{hh\delta}$, see e.g. [125].

2.4.6 Zel'dovich approximation and N -body initial conditions

To represent an initial density perturbation δ (sampled from a pdf with the desired dark matter power spectrum, bispectrum and trispectrum) at redshift $z = z_i$ by particle positions and velocities we can use the Zel'dovich approximation [127] (see also [41, 128]), which states that the unperturbed comoving particle positions \mathbf{q}_j , which can be placed equally spaced on a grid, should be displaced by a displacement field $\Psi(\mathbf{q}, z)$ to perturbed positions

$$\mathbf{x}_j(z) = \mathbf{q}_j + \Psi(\mathbf{q}_j, z), \quad (2.148)$$

where the Fourier transform of Ψ is given by

$$\Psi(\mathbf{k}, z) = \frac{i\mathbf{k}}{k^2} \delta(\mathbf{k}, z). \quad (2.149)$$

The velocities \mathbf{v}_j of the particles can be obtained from (see e.g. [128])

$$\begin{aligned} \mathbf{v}_j(z) &= a\mathbf{x}'_j = a \frac{D'(z)}{D(z)} \Psi(\mathbf{x}_j, z) \\ &= \frac{a}{D(z)} \frac{\Omega_m H_0^2}{2a^2 H(a)} \left(5 - \frac{3D(z)}{a} - \frac{2\Omega_k D(z)}{\Omega_m} \right) \Psi(\mathbf{x}_j, z), \end{aligned} \quad (2.150)$$

where a is the scale factor at the initial redshift z and primes denote derivatives with respect to conformal time as usual. In going to the second line we expressed $D'(z)$ in terms of $D(z)$ using the definition of $D(z)$ in (2.121).

A derivation of the Zel'dovich approximation can be found in e.g. Chapter 2.6 of [41]. First, the divergence of the equation of motion of particle trajectories $\mathbf{x}(z)$ is simplified by using the Poisson equation and conservation of total mass

when displacing the particles. Then the resulting equation is solved at first order in perturbations by (2.149). Because this is an approximation, the density perturbation δ_{ZA} of the displaced particles is not exactly equal to the input density perturbation δ . Indeed it was shown [110, 129, 130] that the Zel'dovich approximation introduces a spurious bispectrum, i.e. even if the input density perturbation is Gaussian, the density perturbation of the displaced particles is slightly non-Gaussian. Since this will be potentially confusing when examining the effects of non-Gaussian input density perturbations due to primordial non-Gaussianity, we will use the second order Lagrangian perturbation theory (2LPT) generalisation of the Zel'dovich approximation as advertised in [130], because this should give no spurious bispectrum (but possibly a spurious trispectrum, which we neglect however).

In a practical implementation of the Zel'dovich approximation or its 2LPT generalisation, the displacement field $\Psi(\mathbf{q}, z)$ is calculated from the density perturbation $\delta(\mathbf{q}, z)$ at points \mathbf{q} on a grid. However the initially unperturbed particles may sit at other points \mathbf{q}_i inside the box.¹ Therefore the displacement field Ψ must be interpolated, e.g. using a cloud-in-cell scheme which averages over the field values at the eight nearest-neighbour grid points of \mathbf{q}_i , with weights given by the distances to the neighbours (see e.g. [133]). For higher accuracy one can also interpolate over the 27 nearest neighbours using a triangular cloud-in-cell scheme.

2.4.7 *N*-body simulations

To model the full non-linear clustering of dark matter we describe it as a self-gravitating fluid obeying the collisionless Boltzmann (or Vlasov) equation in co-moving coordinates (e.g. [41])

$$\frac{df}{d\eta} = \frac{\partial f}{\partial \eta} + \mathbf{v} \cdot \frac{\partial f}{\partial \mathbf{x}} - \frac{\partial \phi}{\partial \mathbf{x}} \cdot \frac{\partial f}{\partial \mathbf{v}} = 0, \quad (2.151)$$

¹This is the case if particles are not perturbed from a regular-spaced grid but from an initial glass distribution, which is obtained by distributing particles randomly in a box and then evolving them with the sign of the gravitational force flipped [131, 132].

where f is the mass density in phase space. The self-consistent potential ϕ solves the Poisson equation

$$\nabla^2 \phi(\mathbf{x}, \eta) = 4\pi G \int f(\mathbf{x}, \mathbf{v}, \eta) d\mathbf{v}, \quad (2.152)$$

where the last integral is the local mass density. Note that the continuity equation can be obtained from the zeroth moment of (2.151), and combining this with its first moment gives the Euler equation [41].

Solving the Boltzmann equation (2.151) is hard because it is a non-linear PDE in the seven variables $\mathbf{x}, \mathbf{v}, \eta$. Typically one employs the N -body method, which is a Monte-Carlo-like approach where $N_p^3 \gg 1$ “particles” sample the phase-space distribution and are displaced by the gravitational force at every time step.

Different numerical techniques can be used to implement this method (see e.g. [41, 134] and references therein). The most straightforward one is the direct summation particle-particle (PP) method where the force on a particle is computed by summing over all other particles. This is very CPU-intensive because it scales like $\mathcal{O}(N_p^6)$ for N_p^3 particles. A faster approach is the tree method, which stores particles in a tree to allow for a $\mathcal{O}(N_p^3 \log N_p)$ force-computation by truncating a hierarchical multipole-expansion of “clusters” containing many particles. The main drawback is the high memory requirement. A less memory demanding method is the Particle-Mesh (PM) method, where particles are converted to a mass density interpolated to a fixed 3D grid with N_g^3 grid points, and the Poisson equation is solved using a Fast Fourier Transform. The main drawback is that structures on scales of the grid separation or below cannot be resolved. Therefore, modern codes like Gadget-2 [132] often use a hybrid TreePM method, where short-range forces are computed with the tree method while the PM method is used for long-range forces on large scales.

While Gadget-2 also contains a smoothed particle hydrodynamics (SPH) module to incorporate collisional gas, this is not needed in large-scale cosmological studies like in this thesis, where it is sufficient to simulate dark matter only.

2.4.8 Large-scale structure surveys

To probe the dark matter distribution of the late-time universe, three-dimensional galaxy catalogues have been generated from a series of experiments such as the 2-degree Field Galaxy Redshift Survey (2dFGRS; around 250,000 galaxies with median $z = 0.11$; [135]), the 6-degree Field Galaxy Survey (6dFGS; [136]), the Sloan Digital Sky Survey I/II (SDSS-I, 2000-2005; SDSS-II, 2005-2008; 930,000 galaxies [137]) and the WiggleZ Dark Energy Survey (240,000 galaxies at $z = 0.2 - 1$; [138]), which are all completed. The Baryon Oscillation Spectroscopic Survey (BOSS) is part of the ongoing SDSS-III and will be completed in 2014, collecting about 1.5 million galaxies in a volume of about $6 \text{ Gpc}^3/h^3$, with mean redshift of $z \approx 0.5$ [139]. Various further improvements will be obtained by future planned or ongoing surveys like the Dark Energy Survey (DES; 300 million galaxies with photometric redshifts $z \lesssim 1$ in the next five years; [140, 141]), the Hobby-Eberly Telescope Dark Energy Experiment (HETDEX; 2013-2015; $z \approx 3$; [142]), the Extended Baryon Oscillation Spectroscopic Survey (e-BOSS; 2014-2018?; $z = 0.3 - 1.5$; [143]), the Prime Focus Spectrograph (PFS) on the Subaru Telescope (2017; few million galaxies at $z \approx 0.8 - 2.4$; [144]), the Medium Scale Dark Energy Spectroscopic Instrument (MS-DESI; 2018-2022?; 20 - 40 million redshifts at $z \approx 0.2 - 1.6$; based on BigBOSS and DESpec proposals; [145, 146, 147]), the EUCLID satellite mission (2020; 50 million spectroscopic redshifts at $z \approx 0.7 - 2$; [148, 149]), the space-based Wide Field Infrared Survey Telescope (WFIRST; 20 million galaxies at $z \approx 1 - 3$; 2023-2028?; [150]), and the ground-based Large Synoptic Survey Telescope (LSST; several billion galaxies with median photometric redshift $z \approx 1.2$; 2022-2032?; [151]). In the same decade, the Square Kilometer Array (SKA; [152]) will collect the 21cm-radiation of hydrogen (due to a hyperfine splitting of the ground state), which also traces the dark matter distribution.

Many of these surveys also map the high-redshift distribution of neutral hydrogen clouds in front of quasars by measuring the Lyman-alpha forest, which is formed from absorption lines in quasar spectra that occur when photons from quasars excite hydrogen atoms along the line of sight through the Lyman- α transition. (Since photons redshift on their way from quasars, the observed frequency of an absorption line corresponds to the location of the hydrogen along the line

2.4 Large-scale dark matter perturbations

of sight). The mapped hydrogen distribution is also a tracer of the dark matter distribution.

While the planned surveys mentioned above are often tailored towards constraining dark energy and measuring the BAO scale, the data delivered by all these surveys can also be used to constrain primordial non-Gaussianity with ever higher precision from the large scale structure clustering properties or the abundance of highly clustered objects. Measuring higher order correlations like the bispectrum from these surveys will not only be useful for testing the initial conditions but will also break degeneracies of late-time observables and nuisance parameters present at the 2-point level, e.g. by providing precious information on the bias relation between dark matter and galaxies and on the growth of structures due to specific induced triangle dependencies of the bispectrum.

Non-Gaussian N -body initial conditions for arbitrary primordial bispectra

Based on [4]

Summary

We present the implementation of an efficient framework for generating arbitrary non-Gaussian initial conditions for N -body simulations. The method is based on a separable modal expansion of the primordial bispectrum. This allows generating non-Gaussian initial conditions far more efficiently than described in the literature and far more generally as e.g. non-separable bispectra can also be studied for the first time. The modal expansion is also used to estimate the bispectra in the realisations of the initial conditions, which we use as a consistency check of the initial condition setup.

Overview

This chapter is organised as follows. After a short introduction in Sec. 3.1, we describe in Sec. 3.2 the algorithm for generating non-Gaussian initial condition using a separable mode expansion of the (appropriately weighted) bispectrum or trispectrum. In Sec. 3.3 we describe a non-trivial validation procedure by estimating the non-Gaussianity of the generated initial field. We present the results for a series of non-Gaussian shapes in Sec. 3.4 and conclude in Sec. 3.5.

3.1 Introduction

Non-Gaussianity of the primordial density perturbations provides an observational window to distinguish mechanisms that created the initial fluctuations (see e.g. [16, 17, 18, 42, 87, 153, 154, 155, 156, 157, 158]). So far most constraints on primordial non-Gaussianity (PNG) come from CMB observations, the most recent ones being from Planck [20]. Large-scale structures (LSS) provide an alternative probe of PNG. Since galaxy surveys provide three-dimensional datasets, LSS has in principle access to more modes than the two-dimensional CMB. Current and upcoming surveys like SDSS, DES and EUCLID can therefore provide tight estimates on PNG. However, the effect of PNG on LSS is difficult to model because of non-linear structure formation, which requires the use of N -body simulations on all but the largest scales, and the complicated relationship between dark matter structures and the observed tracers like galaxies, which requires detailed understanding of galaxy bias and redshift space distortions. We will only consider the first problem here; our aim is to study the late time dark matter clustering in presence of PNG using N -body simulations.

So far only the effect of the local, equilateral and orthogonal templates of PNG have been studied [39, 159]. Except for the local shape the generation of non-Gaussian initial conditions is computationally extremely expensive even if separable templates of the physical shapes are used. This prohibits non-local non-Gaussian N -body simulations with the grid sizes required for upcoming galaxy surveys. In addition, theoretical bispectra from inflation models are often non-separable. The separable equilateral and orthogonal templates are tailored for the

CMB (meaning that the measured f_{NL} of the templates is very close to the f_{NL} one would get for the correct physical shapes), but they should ultimately not be used for LSS, because LSS observables are sensitive to limits of the bispectrum which are not necessarily represented correctly by the CMB templates (e.g. halo bias strongly depends on the squeezed limit of the bispectrum). To be able to study structure formation for essentially any model of inflation without relying on separable templates, the goal in this chapter is to provide an efficient and general numerical framework to run N -body simulations for arbitrary types of PNG. The method presented here achieves this for arbitrary primordial bispectra.¹ The main idea is to expand the primordial bispectrum (or rather the associated kernel required for generating initial conditions) in separable basis functions. We generate initial conditions for every basis mode and sum up the contributions to represent the full bispectrum in the initial particle distribution. This approach is a natural LSS generalisation of the modal expansion technique pioneered for the CMB by Fergusson and Shellard [28, 49, 160]. While the theory has been presented in [50], we summarise the formalism here and present the implementation and numerical consistency checks of this framework.

3.2 Algorithm

3.2.1 Bispectrum

To generate a realisation of the primordial potential Φ if this has power spectrum $P(k)$ and arbitrary primordial bispectrum $B(k_1, k_2, k_3)$, we first draw Φ^G from a Gaussian with the desired power spectrum $P(k)$. To represent the bispectrum we compute

$$\Phi = \Phi^G + \frac{f_{NL}}{2} \Phi^B, \quad (3.1)$$

¹The method was also successfully applied to trispectra which do not depend on the diagonals of the quadrilateral [4]. Extending this to diagonal-dependent trispectra would be interesting to test models that depend non-trivially on diagonals, like e.g. some equilateral trispectra [54]. This is conceptually straightforward but computationally challenging [50].

where¹ [39, 50]

$$\begin{aligned}\Phi^B(\mathbf{k}) = & \int \frac{d^3\mathbf{k}' d^3\mathbf{k}''}{(2\pi)^6} (2\pi)^3 \delta_D(\mathbf{k} - \mathbf{k}' - \mathbf{k}'') \Phi^G(\mathbf{k}') \Phi^G(\mathbf{k}'') \\ & \times \frac{B(k, k', k'')}{P(k)P(k') + P(k)P(k'') + P(k')P(k'')}.\end{aligned}\quad (3.2)$$

For the local shape this reduces to $\Phi^G * \Phi^G$ in k -space, which corresponds to the squared Gaussian field in x -space. The symmetrised non-separable denominator in (3.2) is required to leave the power spectrum (to a very good approximation) unchanged [39]. Indeed, for small levels of non-Gaussianity we get, as desired,

$$\begin{aligned}\langle \Phi(\mathbf{k}_1)\Phi(\mathbf{k}_2) \rangle &= (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2) P(k_1), \\ \langle \Phi(\mathbf{k}_1)\Phi(\mathbf{k}_2)\Phi(\mathbf{k}_3) \rangle &= (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) f_{NL} B(k_1, k_2, k_3).\end{aligned}\quad (3.3)$$

The prescription (3.2) is a generalisation of the analogous expressions for the CMB [28, 154]. A computationally expensive brute-force evaluation of Eq. (3.2) was used to create N -body initial conditions for some separable shapes in [39, 159]. This approach is not efficient because the denominator of the kernel in (3.2) is not separable in the sense that it cannot be expressed in the form of a product $f_1(k_1)f_2(k_2)f_3(k_3)$ for some functions f_i . In analogy to the modal CMB techniques described in [28, 49, 160], we achieve efficiency by *expanding* the kernel in (3.2) in separable modes as

$$\frac{B(k, k', k'')}{P(k)P(k') + P(k)P(k'') + P(k')P(k'')} = \sum_{rst} \alpha_{rst}^Q q_r(k) q_s(k') q_t(k'').\quad (3.4)$$

The q_r are one-dimensional independent polynomials defined on the tetrapyd domain where the bispectrum is defined if the triangle condition due to statistical homogeneity is imposed. The specific form of the q_r can be arbitrary as long as they form a set of linearly independent functions. Here we choose simple mono-

¹The kernel in Eq. (3.2) is just one particular, non-unique choice that produces the correct power spectrum and bispectrum, but it does not necessarily generate the desired trispectrum and higher-order correlation functions beyond the trispectrum. Ref. [161] discusses attempts to improve the kernel to obtain correct primordial trispectra, which is an interesting field of future work.

3.2 Algorithm

mials $q_r(k) = k^r$ following [28]. For notational convenience we symmetrise over the indices r, s, t in (3.4), denoted by curly brackets, and label the different index combinations by n , so that

$$\sum_{rst} \alpha_{rst}^Q q_r(k) q_s(k') q_t(k'') = \sum_{n=\{rst\}} \alpha_n^Q q_{\{r}(k) q_s(k') q_{t\}}(k'') = \sum_n \alpha_n^Q Q_n(k, k', k''), \quad (3.5)$$

where we also introduced the 3D basis functions Q_n . The expansion coefficients α_n^Q are specific to a given inflation model which predicts a bispectrum B . While the expansion (3.4) is exact if a complete set of basis functions is used, truncating the expansion after $n_{\max} = \mathcal{O}(30 - 100)$ terms yields an accurate description for many inflation models. Expressing the Dirac delta function in terms of plane waves,

$$\delta_D(\mathbf{k}) = \frac{1}{(2\pi)^3} \int d^3 \mathbf{x} e^{i\mathbf{k}\cdot\mathbf{x}}, \quad (3.6)$$

and using the expansion (3.4) in (3.2) we find

$$\Phi^B(\mathbf{k}) = \sum_n \alpha_n^Q q_{\{r}(k) \int d^3 \mathbf{x} e^{i\mathbf{k}\cdot\mathbf{x}} M_s(\mathbf{x}) M_{t\}}(\mathbf{x}), \quad (3.7)$$

where $M_s(\mathbf{x})$ denotes the potential filtered by the basis functions,

$$M_s(\mathbf{x}) = \int \frac{d^3 \mathbf{k}}{(2\pi)^3} \Phi^G(\mathbf{k}) q_s(k) e^{-i\mathbf{k}\cdot\mathbf{x}}. \quad (3.8)$$

While a brute force evaluation of (3.2) requires $\mathcal{O}(N^6)$ operations (one \mathbf{k}' integral for every \mathbf{k}), where $N = \mathcal{O}(10^3)$ is the number of grid points per dimension, evaluating (3.7) requires only $\mathcal{O}(n_{\max} N^3)$ operations because it consists of a series of Fourier transforms which can be efficiently computed with Fast Fourier Transforms.

In practise the above expressions are evaluated in a box with side length L and N grid points per dimension. The allowed discrete wavenumbers are therefore

$$k = -\left(\frac{N}{2} - 1\right) \Delta k, \dots, \left(\frac{N}{2} - 1\right) \Delta k, \frac{N}{2} \Delta k, \quad (3.9)$$

where $\Delta k = 2\pi/L$. The minimal scale resolved by the grid is represented by the maximum wavenumber $k_{\max} = N\Delta k/2$. Care must be taken to account for unwanted realisations of the discretisation of the Dirac delta function when the wavevectors, \mathbf{k}_i , align.¹ This can be accounted for simply by restricting the range of the wavevectors to $|\mathbf{k}_i| < 2k_{\max}/3$ for the calculation of Φ^B . We will work with $|\mathbf{k}_i| \leq k_{\max}^{\text{eff}} \equiv k_{\max}/2$ in the following. This means that we cannot use the full grid that is available in memory. However this limitation is offset by the improvement in numerical speed that the modal method provides.

To obtain N -body initial conditions from the realisation of Φ generated with the above prescription we relate Φ to the initial density perturbation δ at the initial redshift of the simulation using the sub-horizon Poisson equation during matter domination and transfer function $T(k)$ (see Sec. 2.4.2 and Eq. (4.2) below for details). The Zel'dovich approximation [127] or a second order 2LPT generalisation thereof [128, 129] can then be used to obtain initial particle positions and velocities for the N -body simulation.

3.2.2 Trispectrum

The modal expansion method can also be applied to represent non-trivial trispectra in the initial conditions. To reduce the computational cost we only consider trispectra which do not depend on the diagonal of the quadrilateral, i.e.

$$\langle \Phi(\mathbf{k}_1)\Phi(\mathbf{k}_2)\Phi(\mathbf{k}_3)\Phi(\mathbf{k}_4) \rangle_c = (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3 + \mathbf{k}_4) G_{NL} T(k_1, k_2, k_3, k_4). \quad (3.10)$$

A field with this trispectrum and the desired power spectrum can be generated with

$$\Phi = \Phi^G + \frac{G_{NL}}{6} \Phi^T, \quad (3.11)$$

¹The discrete version of (3.6) is given by $\delta_D(\mathbf{k} - \mathbf{k}' - \mathbf{k}'') \propto \sum_{\mathbf{x}_i} e^{i(\mathbf{k} - \mathbf{k}' - \mathbf{k}'') \cdot \mathbf{x}_i}$, where the sum goes over all points \mathbf{x}_i on the grid in real space with spacing $\Delta x = L/N$. This representation of the discrete Dirac delta function only holds if $|\mathbf{k} - \mathbf{k}' - \mathbf{k}''| < 2k_{\max}$. Therefore we need $\mathbf{k}_i < 2k_{\max}/3$ for all \mathbf{k}_i appearing in the formulas in the main text. I am very grateful to Hiro Funakoshi for identifying this important issue. We only learned later that this was also found in e.g. [112].

3.3 Algorithm validation

where

$$\begin{aligned}\Phi^T(\mathbf{k}) &= \int \frac{d^3\mathbf{k}' d^3\mathbf{k}'' d^3\mathbf{k}'''}{(2\pi)^9} (2\pi)^3 \delta_D(\mathbf{k} - \mathbf{k}' - \mathbf{k}'' - \mathbf{k}''') \\ &\times \frac{T(k, k', k'', k''')}{P(k)P(k')P(k'')} \Phi^G(\mathbf{k}') \Phi^G(\mathbf{k}'') \Phi^G(\mathbf{k}''').\end{aligned}\quad (3.12)$$

Evaluation of this integral requires $\mathcal{O}(N^9)$ operations in general, which is computationally not feasible. However, efficiency is restored again by expanding the relevant kernel in (3.12) in separable modes,

$$\frac{T(k, k', k'', k''')}{P(k)P(k')P(k'')} = \sum_m \bar{\alpha}_m^Q \bar{q}_{\{r}(k) \bar{q}_s(k') \bar{q}_t(k'') \bar{q}_{u\}}(k''').\quad (3.13)$$

Here \bar{q}_r are independent one-dimensional polynomials defined on the trispectrum domain allowed by the quadrilateral condition on the wavenumbers. The label m represents the combination $\{rstu\}$ and the expansion coefficients $\bar{\alpha}_m^Q$ characterise the trispectrum under consideration. We denote the 4D basis functions as $\bar{Q}_n = \bar{q}_{\{r} \bar{q}_s \bar{q}_t \bar{q}_{u\}}$. Using (3.6) we find

$$\Phi^T(\mathbf{k}) = \sum_m \bar{\alpha}_m^Q \bar{q}_{\{r}(k) \int d^3\mathbf{x} e^{i\mathbf{k}\cdot\mathbf{x}} \bar{M}_s(\mathbf{x}) \bar{M}_t(\mathbf{x}) \bar{M}_{u\}}(\mathbf{x}),\quad (3.14)$$

where the potential is filtered with \bar{q}_s ,

$$\bar{M}_s(\mathbf{x}) = \int \frac{d^3\mathbf{k}}{(2\pi)^3} \Phi^G(\mathbf{k}) \bar{q}_s(k) e^{-i\mathbf{k}\cdot\mathbf{x}}.\quad (3.15)$$

Evaluating (3.14) requires only $\mathcal{O}(n_{\max} N^3)$ instead of $\mathcal{O}(N^9)$ operations. To avoid unwanted images of the discrete delta function when the wavevectors align we have to restrict the allowed wavenumbers to $|\mathbf{k}| < k_{\max}/2$.

3.3 Algorithm validation

To validate our implementation of the modal method we directly test if the generated field Φ has the desired bispectrum property by creating a set of realisations

of Φ and comparing the average estimated bispectrum with the theoretical input bispectrum. Note that the accuracy of the expansion of primordial bispectrum shapes has been tested in [28, 51, 162], reaching $\mathcal{O}(90 - 95\%)$ shape correlations for the bispectrum with $\mathcal{O}(30)$ modes and for the trispectrum with $\mathcal{O}(50)$ modes.

3.3.1 Bispectrum estimation

The amplitude of a given theoretical bispectrum shape B can be estimated with [50]

$$\mathcal{E} = \int \frac{\prod_{i=1}^3 d^3 k_i}{(2\pi)^6} \frac{\delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B(k_1, k_2, k_3)}{P(k_1)P(k_2)P(k_3)} [\phi_{\mathbf{k}_1} \phi_{\mathbf{k}_2} \phi_{\mathbf{k}_3} - 3 \langle \phi_{\mathbf{k}_1} \phi_{\mathbf{k}_2} \rangle \phi_{\mathbf{k}_3}]. \quad (3.16)$$

The ensemble average is

$$\langle \mathcal{E} \rangle = \frac{V}{\pi} \int_{\mathcal{V}_B} dk_1 dk_2 dk_3 \frac{k_1 k_2 k_3 B^2(k_1, k_2, k_3)}{P(k_1)P(k_2)P(k_3)}, \quad (3.17)$$

where \mathcal{V}_B denotes the bispectrum domain allowed by the triangle condition. The volume prefactor is $V = (2\pi)^3 \delta_D(\mathbf{0})$. The integrand in (3.16) contains the inverse-variance weighted potential ϕ projected on the theoretical bispectrum B , i.e. the estimator is just the cumulative signal to noise. Since the integral in (3.16) is numerically expensive (scaling as $\mathcal{O}(N^6)$) if the theoretical bispectrum is not separable, we expand the integrand in separable modes as

$$\frac{\sqrt{k_1 k_2 k_3} B(k_1, k_2, k_3)}{\sqrt{P(k_1)P(k_2)P(k_3)}} = \sum_n \alpha_n^Q q_{\{r}(k_1) q_s(k_2) q_{t\}}(k_3). \quad (3.18)$$

Note that this is a different expansion than in (3.4) where we expand the kernel required for initial conditions instead of the theoretical signal to noise relevant for bispectrum estimation. The two corresponding sets of expansion coefficients α_n^Q and $\alpha_n^{Q'}$ can be related with a conversion matrix which can be computed as a

3.4 Algorithm validation

one-off calculation. With the expansion (3.18) the estimator (3.16) becomes [50]

$$\mathcal{E} = \sum_n \alpha_n^{Q'} \underbrace{\int d^3\mathbf{x} [M_r(\mathbf{x}) M_s(\mathbf{x}) M_t(\mathbf{x}) - \langle M_{\{r}(\mathbf{x}) M_{s}(\mathbf{x}) \rangle M_{t\}}(\mathbf{x})]}_{\equiv \beta_n^{Q'}} \quad (3.19)$$

$$\langle \mathcal{E} \rangle = \sum_{nm} \alpha_n^{Q'} \alpha_m^{Q'} \gamma_{nm}, \quad (3.20)$$

where

$$\gamma_{nm} = \frac{V}{\pi} \int_{V_B} dk_1 dk_2 dk_3 Q_n(k_1, k_2, k_3) Q_m(k_1, k_2, k_3), \quad (3.21)$$

$$M_r(\mathbf{x}) = \int \frac{d^3\mathbf{k}}{(2\pi)^3} \frac{\phi_{\mathbf{k}} q_r(k)}{\sqrt{k P(k)}} e^{i\mathbf{k}\cdot\mathbf{x}}. \quad (3.22)$$

The $\beta_n^{Q'}$ coefficients defined by (3.19) depend on the data ϕ and the basis functions q_r , and satisfy

$$\langle \beta_n^{Q'} \rangle = \sum_m \alpha_m^{Q'} \gamma_{nm}. \quad (3.23)$$

The relationship between theory coefficients α and data coefficients β becomes more transparent if we transform Q_n to mode functions \mathcal{R}_n which are orthogonal with respect to the inner product $\langle fg \rangle = (V/\pi) \int_{V_B} dk_1 dk_2 dk_3 f(k_1, k_2, k_3) g(k_1, k_2, k_3)$ motivated by (3.21). Then, (3.23) yields the consistency relation

$$\langle \beta_n^{\mathcal{R}'} \rangle = \alpha_n^{\mathcal{R}'}, \quad (3.24)$$

i.e. in this basis the average data coefficients $\beta_n^{\mathcal{R}'}$ are just given by the expansion coefficients $\alpha_n^{\mathcal{R}'}$ of the theoretical bispectrum under consideration. Therefore, $\beta_n^{\mathcal{R}'}$ can be regarded as the estimated f_{NL} -amplitude of the n -th basis shape. The sum (3.19) sums up these individual amplitudes to obtain the estimated f_{NL} of the full theory bispectrum shape B under consideration.

3.4 Results

Our implementation of the modal method was used to generate non-Gaussian initial conditions of the local, equilateral, constant, orthogonal and flattened bispectrum shapes (see Sec. 2.1.7.3 and [51]¹). To generate initial conditions we first compute the expansion coefficients in (3.4) by integrating over the tetrapyd. We then compute the initial conditions using (3.7). To test their accuracy we estimate the bispectrum with (3.19) and check the consistency relation (3.24), where the theory expansion coefficients are obtained from (3.18). The mode expansions (3.4) and (3.18) were truncated after $n_{\max} = 30$ modes for all models except the non-separable flattened shape for which we used $n_{\max} = 80$ modes to improve accuracy. The term of the form $\langle MM \rangle M$ in (3.19) can be dropped when computing the β 's because we do not include any inhomogeneous noise or mask.

Generating the initial conditions is very efficient, taking about one hour on only 6 cores for one realisation of a 1024^3 grid. Since we are only presenting a proof of concept in this chapter, we will mostly work with smaller 256^3 grids.

Before validating the bispectrum of the initial conditions we note that the power spectrum of the initial conditions is almost unchanged when adding Φ^B to the Gaussian field. Indeed, Fig. 3.1 demonstrates that the non-Gaussian change of the power spectrum given by $f_{NL}^2 \langle \Phi^B \Phi^B \rangle / 4$ is negligible compared to the Gaussian power spectrum $P(k)$.

To validate the bispectrum in the initial conditions we compare in Fig. 3.2 the theoretical expansion coefficients $\alpha_n^{\mathcal{R}'}$ with the estimated $\beta_n^{\mathcal{R}}$ averaged over 5 realisations. The excellent agreement for both the local and equilateral bispectrum shape demonstrates the correctness of the implementation. A more quantitative validation is presented in Table 3.1, where the shape correlations between the theoretical bispectrum and the separable expansions (3.4) and (3.18) are shown to test the primordial expansions, while the last column shows the shape correlation of the expansion (3.18) of the theory bispectrum with the average measured bispectrum represented by the average β coefficients. Since this last column is at least 99.3% in all cases we see that the initial conditions represent the input decomposition

¹I implemented the modal expansion technique for generating LSS initial conditions for arbitrary bispectra in a hybrid MPI/OpenMP-parallelised C code. The code was later generalised to diagonal-independent trispectra by Donough Regan; see [4] for details.

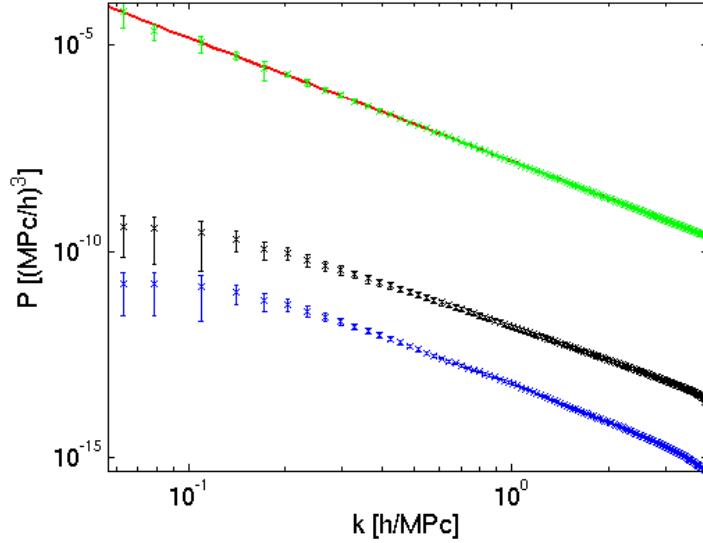


Figure 3.1: Plot of input power spectrum $P(k)$ for Φ^G (red) and measured power spectra of $f_{NL}\Phi^B/2$ for the local shape with $f_{NL} = 100$ (black), $f_{NL}\Phi^B/2$ for the equilateral shape with $f_{NL} = 400$ (blue) and $\Phi^G + f_{NL}\Phi^B/2$ for the two cases (green). Φ^B was calculated for 100 realisations of Φ^G using (3.7) on a 256^3 grid in a box with $L = 100$ Mpc/h.

of the primordial bispectrum very well. The correlation of the primordial shape with its expansions (3.4) and (3.18) is at least 90% in all cases. While this is already a reasonably good representation, further improvements can be obtained by using more basis functions (i.e. increasing n_{\max}) or by choosing a more suitable functional form for the basis functions instead of the simple monomials employed here.

Results for trispectrum initial conditions are presented in [4], where we have validated the modal trispectrum method for local, equilateral and constant trispectra.

3.5 Conclusions

We have presented the implementation of a method to create initial conditions for N -body simulations for arbitrary primordial bispectra, including non-separable shapes. The method is based on a separable modal expansion of the bispectrum

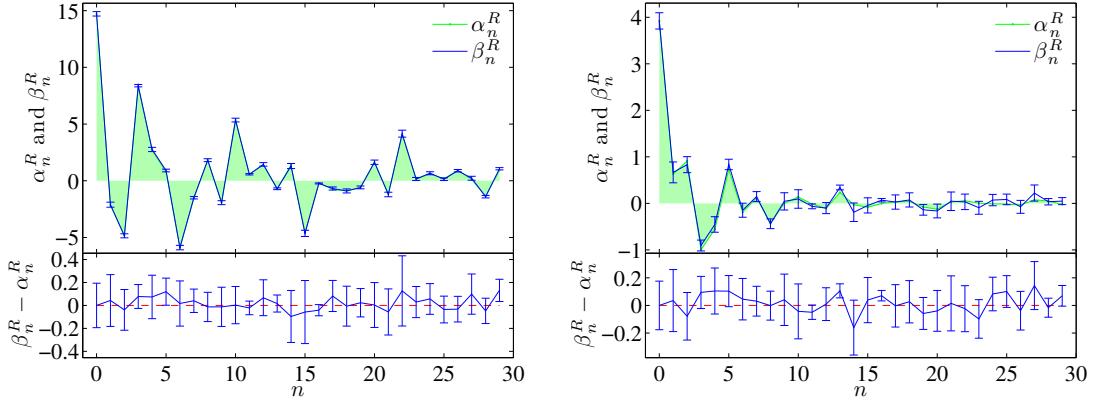


Figure 3.2: Theoretical expansion coefficients $\alpha_n^{R'}$ from (3.18) compared with the average estimated coefficients $\langle \beta_n^R \rangle$ from 5 realisations in a box with $N = 512$ and $L = 100\text{Mpc}/h$. *Left:* Local bispectrum shape with $f_{\text{NL}} = 100$, *right:* equilateral shape with $f_{\text{NL}} = 500$. Note that the left plot differs from the one shown in [4] because we use a different ordering of polynomials here and use different box and grid size. The error bars show 1σ errors. The lower panels show residual plots.

which reduces the generation of initial conditions to a series of efficient Fourier transforms. It requires $\mathcal{O}(n_{\text{max}}N^3)$ operations, where $N = \mathcal{O}(1000)$ is the number of grid points per dimension and $n_{\text{max}} = \mathcal{O}(30 - 100)$ is the number of basis functions which is kept in the expansion of the bispectrum. The computational cost is therefore reduced by a factor of $\mathcal{O}(10^7)$ compared to a brute force approach which requires $\mathcal{O}(N^6)$ operations and has been used in the literature so far. We have validated our implementation for the local, equilateral, orthogonal, constant and non-separable flattened bispectrum shapes with a non-trivial consistency check by comparing the bispectra of the initial conditions with the input bispectra.

We note that further work is required to represent arbitrary trispectra together with arbitrary bispectra in the initial conditions because the bispectrum term $f_{\text{NL}}\Phi^B/2$ in (3.1) introduces a spurious trispectrum which is non-trivial to mitigate (see [161] for first attempts in this direction for a few specific shapes). We could subtract this spurious trispectrum if the presented methods are generalised to trispectra which also depend on the diagonal of the quadrilateral, which is in principle straightforward though computationally challenging.

We hope that the work presented here proves useful for studying large-scale

Model	n_{\max}	f_{NL}	Shape vs Decomp (3.4)	Shape vs Decomp (3.18)	$\langle \beta_n^Q \rangle$ vs Decomp (3.18)
Local	30	100	100%	92.6%	99.5%
Equilateral	30	200	99.7%	99.7%	99.8%
Constant	30	200	99.9%	100%	99.3%
Orthogonal	30	200	98.7%	98.9%	99.7%
Flattened	80	200	91.8%	90.6%	99.1%

Table 3.1: Correlation between the primordial bispectrum shape and the separable modal expansions (3.4) and (3.18), as well as the correlation between the average of the realisations $\langle \beta_n^Q \rangle$ and the expansion (3.18).

structures with N -body simulations in presence of arbitrary primordial bispectra.

Acknowledgements

I am very grateful to my collaborators for this chapter: Donough Regan, James Fergusson and Paul Shellard. We thank Hiro Funakoshi for identifying the issue of images of the Dirac delta function. We are grateful to Andrey Kaliazin for his invaluable computational help. Simulations were performed on the COSMOS supercomputer (an SGI UV Altix) which is funded by STFC and DBIS.

Fast estimation of gravitational and primordial bispectra in large-scale structures

Based on [5]

Summary

We present the implementation of a fast estimator for the full dark matter bispectrum of a three-dimensional particle distribution relying on a separable modal expansion of the bispectrum. The computational cost of accurate bispectrum estimation is negligible relative to simulation evolution, so the isotropic bispectrum can be used as a standard diagnostic whenever the power spectrum is evaluated. As an application we measure the evolution of gravitational and primordial dark matter bispectra in N -body simulations with Gaussian and non-Gaussian initial conditions of the local, equilateral, orthogonal and flattened shape. The results are compared to theoretical models using a 3D visualisation, 3D shape correlations and the cumulative bispectrum signal-to-noise, all of which can be evaluated extremely quickly. Our measured bispectra are determined by $\mathcal{O}(50)$ coefficients, which can be used as fitting formulae in the nonlinear regime and for non-Gaussian initial conditions. In the nonlinear regime with $k < 2h \text{ Mpc}^{-1}$, we find an excellent correlation between the measured dark matter bispectrum and a simple model based on a ‘constant’ bispectrum plus a (nonlinear) tree-level gravitational bispectrum. In the same range for non-Gaussian simulations, we find the measured excess bispectrum (i.e. the difference with the measured gravitational bispectrum) has an excellent correlation with a phenomenological model consisting of the constant bispectrum plus a (nonlinear) tree-level primordial bispectrum. We demonstrate that the constant contribution to the non-Gaussian bispectrum can be understood as a time-shift of the constant mode in the gravitational bispectrum, which is motivated by the one-halo model. The final amplitude of this extra non-Gaussian constant contribution is directly related to the initial amplitude of the constant mode in the primordial bispectrum. We also comment on the effects of regular grid and glass initial conditions on the bispectrum.

Overview

The chapter is organised as follows. After a brief introduction in Section 4.1 we review the generation of non-Gaussianity due to nonlinear gravitational evolution in Section 4.2, together with the physically-motivated models that we will use to describe our results. In Section 4.3 we review primordial non-Gaussianity and discuss its impact on the matter bispectrum, proposing a new time-shift model. In Section 4.4 we discuss the bispectrum estimation methodology based on a separable expansion which allows extremely efficient estimation of the full N -body bispectrum, as well as the cumulative signal-to-noise and its correlation to theoretically-predicted bispectra. We discuss the simulation setup, impact of glass initial conditions, validations and convergence tests in Section 4.5. In Section 4.6 we present our results for the gravitational bispectrum and the simple fits thereof, while primordial bispectra in non-Gaussian simulations and their fits are discussed in Section 4.7. We summarise our conclusions in Section 4.8. Readers familiar with the modal methodology and interested primarily in the measured bispectra and testing of various fitting formulae may wish to start with Section 4.6 and follow the references to the earlier sections given there. Table 4.1 lists mathematical symbols used in this chapter.

4.1 Introduction

So far, CMB [51, 162, 163, 164] and large-scale structure (LSS; [165, 166]) observations have not found any evidence for primordial non-Gaussianity. However, a detection of primordial non-Gaussianity through a non-zero bispectrum or trispectrum would significantly improve our understanding of the physics of inflation. General methods that can distinguish different shapes of bispectra and trispectra are particularly important to discriminate between inflationary mechanisms. Fergusson and Shellard have developed such methods for the CMB using a separable expansion of the bispectrum or trispectrum to achieve computational feasibility [28, 49, 51, 54, 162].

While the most stringent constraints on primordial non-Gaussianity are currently given by the CMB Planck results [20], competitive or improved constraints are expected from the large-scale structure surveys listed in Sec. 2.4.8. Partially,

Symbol	Description	Definition
$\langle \cdot, \cdot \rangle$	S/N weighted scalar product between bispectra	(4.32)
$\langle \cdot, \cdot \rangle_{\text{noweight}}$	Unweighted scalar product between bispectra	(4.44)
$\ \cdot \ $	Bispectrum norm (cumulative S/N)	(4.34)
$\alpha_n^{Q,R}$	Theory bispectrum expansion coefficients	(4.38), (4.45), above (4.47)
β_n^R	Estimated bispectrum coefficients	(4.48)
α, β	Sides of $k_1 + k_2 + k_3 = \text{const}$ triangles	
$\gamma_{nm}, \lambda_{nm}$	Matrices related to Q_n and R_n bases	(4.41), above (4.47)
δ	Fractional dark matter density	
δ_D	Dirac delta function	
δ_n	n -th order density perturbation	(2.138)
Δz	Halo model time shift due to prim. non-Gaussianity	Section around (4.24)
ν	k -scaling of $B_{\delta, \text{const}}^{\text{grav}}$ for equilateral triangles	(4.12)
Φ	Primordial potential	above (4.2)
$B_{\{\delta, \Phi\}}$	Bispectrum of δ or Φ	(4.3), (4.21)
$B_{\delta}^{\text{fit}}, B_{\text{NG}}^{\text{fit}}$	Time shift model fits for bispectra in Gaussian and non-Gaussian simulations	(4.62), (4.67)
B_{δ}^{grav}	Perturbative gravitational bispectrum	(4.4)
$B_{\delta, \text{const}}^{\text{grav}}$	‘Constant’ bispectrum (approx. 1-halo bispectrum)	(4.12)
$B_{\delta, \text{NL}}^{\text{grav}}$	Phenomenological non-linear extension of B_{δ}^{grav}	(4.6)
B_{δ}^{opt}	Optimal linear combination of $B_{\delta, \text{NL}}^{\text{grav}}$ and $(k_1 + k_2 + k_3)^{\nu}$ with arbitrary weight $w(z)$	(4.63)
B_{δ}^{prim}	Perturbative density bispectrum from primordial B_{Φ}	(4.22)
$B_{\delta, \text{const}}^{\text{prim}}$	Time-shifted constant bispectrum due to prim. non-Gaussianity	(4.26)
$B_{\delta, \text{NL}}^{\text{prim}}$	Phenomenological non-linear extension of B_{δ}^{prim}	(4.23)
B_n^R	Orthonormal contributions to full bispectrum shape	(4.46), (4.51)
\hat{B}_{δ}	Measured density bispectrum	(4.51)
\hat{B}_{NG}	Measured primordial contribution to \hat{B}_{δ}	(4.52)
$\mathcal{C}(\cdot, \cdot)$	3D shape correlation (cosine)	(4.33)
$\mathcal{C}_{\alpha, \text{th}}$	Correlation of theory vs expansion	(4.53)
$\mathcal{C}_{\beta, \alpha}$	Correlation of simulations vs expansion	(4.54)
$\mathcal{C}_{\beta, \text{th}}$	Correlation of simulations vs theory	(4.55)
c_1, c_2	Amplitudes of $B_{\delta, \text{const}}^{\text{grav}}$ and $B_{\delta, \text{const}}^{\text{prim}}$	(4.12), (4.26)
$\bar{D}(z), D(z)$	Linear growth function, normalised to $\bar{D}(0) = 1$ and $D(z) = (1+z)^{-1}$ in matter dom.	below (4.2)
dV_k	Volume element $dk_1 dk_2 dk_3$	Below (4.30)

Chapter 4

$F_n^{(s)}$	Symmetrised density kernels	(2.133)
f_{NL}	Non-linear amplitude of B_Φ	(4.21)
$\hat{f}_{\text{NL}}^{\text{th}}$	Estimator of amplitude of theoretical bispectrum B_δ^{th}	(4.28)
\bar{F}_{NL}^B	Cumulative bispectrum S/N relative to linearly evolved local bispectrum	(4.36)
L	Side length of simulation box (in Mpc/h)	
$M(k, z)$	Poisson factor	(4.2)
M_r	Density perturbation filtered by 1D basis function q_r	(4.42)
N_{th}	Normalisation of $\hat{f}_{\text{NL}}^{\text{th}}$	(4.31)
N	Number of grid points per dimension	
N_{fit}	Overall normalisation factor for B_δ^{fit}	(4.64)
n_h, n_h^{prim}	Halo exponents describing time evolution of $B_{\delta, \text{const}}^{\text{grav}}$ and $B_{\delta, \text{const}}^{\text{prim}}$	(4.12), (4.26)
n_{\max}	Number of modal basis functions used	(4.38)
N_p	Number of particles per dimension	
N_r	Number of realisations simulated	
P_δ	Power spectrum of δ	(4.1)
P_δ^L	Linear power spectrum of δ_1	
q_r	1D basis functions (chosen to be monomials here)	
Q_n	3D basis functions obtained from q_r	(4.43)
R_n	Orthonormalised 3D basis functions (rotated Q_n 's)	above (4.47)
s_B	Symmetry factor	Below (4.34)
$T(k)$	Linear transfer function	below (4.2)
V	Volume factor	$V = L^3$
\mathcal{V}_B	Tetrahedral domain allowed by triangle condition	

Table 4.1: List of most important mathematical symbols used throughout this Chapter (alphabetically sorted).

this is expected because LSS data is three-dimensional and can therefore access many more modes than the 2D CMB. While the CMB 2-point statistics are basically unaffected by primordial non-Gaussianity, the power spectrum of biased tracers of the dark matter distribution receives corrections on large scales that can scale as k^{-2} if the primordial bispectrum has a significant squeezed limit [167]. Detection of this ‘transfer’ of the bispectrum to the power spectrum through a scale-dependent bias would rule out single field inflation. While this effect makes future LSS power spectrum constraints competitive with the CMB bispectrum,

some care needs to be taken to understand all systematics on very large scales in a robust manner. Including the LSS bispectrum is expected to improve the power spectrum constraints already on intermediate scales [117, 118], where systematics should be well controlled, but next-to-leading order contributions in the theoretical modeling may become relevant. Including smaller scales in the bispectrum analysis will improve constraints further. This is not limited by observations but only by the break-down of the theoretical modeling, i.e. any theoretical progress about understanding the impact of non-Gaussianity on small scales can yield major improvements for constraining primordial non-Gaussianity from LSS. Several groups have therefore studied the impact of primordial non-Gaussianity on matter, halo and galaxy power spectra and bispectra (see, e.g., [118, 168, 169, 170, 171, 172, 173, 174, 175, 176, 177] and references therein). Constraints on local non-Gaussianity from LSS have recently been obtained in [166, 178, 179].

Only few works have studied bispectrum shapes beyond the local template because of the difficulty to generate initial conditions for N -body simulations for non-local shapes. This changed when Wagner and Verde [39, 159], as well as Scoccimarro et al. [161] suggested methods to simulate some other specific shapes. However, the methods are not efficient for truly general non-separable shapes. As an alternative, [50] suggested to apply the separable expansion method which also works for non-separable shapes. This was successfully implemented and tested in [4], as described in Chapter 3 of this thesis.

In this chapter, we apply the initial conditions generator from Chapter 3 to run full N -body simulations with non-Gaussian initial conditions for several primordial bispectrum shapes. Exploiting again a separable expansion of the bispectrum we estimate and study the full resulting 3D dark matter bispectrum at a series of redshifts down to small scales. The estimator is not only orders of magnitude faster than previous N -body bispectrum estimators, but it also compresses the full 3D information of bispectrum shape to $\mathcal{O}(100)$ numbers, which greatly simplifies subsequent analyses based on 3D shape correlations and cumulative signal to noise estimates. We compare the N -body bispectra with tree level predictions and show the time and length scales at which they break down, both for Gaussian and non-Gaussian simulations. We describe the results in the highly non-linear

regime using a phenomenological time shift model according to which primordial non-Gaussianity leads to an earlier start of halo growth compared to a Gaussian universe. This motivates a simple form of the dark matter bispectrum, which we exploit to create fitting formulae valid down to small scales over a wide range of redshifts and for a series of different inflationary bispectrum shapes. We also study the impact of glass initial conditions in comparison to regular grid initial conditions. While the main purpose here is to demonstrate the efficiency and accuracy of the modal bispectrum estimation by applying it to the dark matter distribution in N -body simulations, we emphasise that the same method can in principle be applied to the 3D distribution of galaxies mapped by LSS surveys. Treating systematics, foregrounds and window functions with the modal method should in principle be similar to the CMB treatment used in the Planck analysis [20]. However, applying the modal method to real data will be left for future work.

4.2 The distribution of matter

4.2.1 Power spectrum and transfer functions

An important prediction of cosmological models is the probability of finding a certain configuration of the fractional overdensity $\delta(\mathbf{x})$ in the universe. The simplest possibility is a Gaussian pdf, which is determined by the two-point correlation function, or in Fourier space by the power spectrum P_δ ,

$$\langle \delta(\mathbf{k}_1) \delta(\mathbf{k}_2) \rangle = (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2) P_\delta(k_1). \quad (4.1)$$

Here we assumed statistical homogeneity, leading to the Dirac delta function, and statistical isotropy, which implies that the power spectrum only depends on the magnitude of the wavevector.

Perturbations from inflation are usually described by the comoving curvature perturbation \mathcal{R} , which is related to the primordial potential Φ by $\Phi = -3\mathcal{R}/5$, in linear perturbation theory during matter domination. The density perturbation δ

can be obtained using the linear Poisson equation

$$\delta(\mathbf{k}, z) = \frac{2}{3} \frac{k^2 T(k) D(z)}{\Omega_m H_0^2} \Phi(\mathbf{k}) \equiv M(k, z) \Phi(\mathbf{k}), \quad (4.2)$$

where $T(k)$ is the linear transfer function at low redshift normalised to $T(k) = 1$ on large scales and calculable with CAMB [64], and $D(z)$ is the linear growth function for $\Omega_{\text{rad}} = 0$ [180] normalised to $D(z) = 1/(1+z)$ during matter domination. Later we will also find it convenient to use the growth function normalised to unity today, $\bar{D}(z) = D(z)/D(0)$.

At late times and on small scales, the linear treatment breaks down and N -body simulations must be used to find the nonlinear transfer of primordial to late time perturbations. A widely used fit of the nonlinear power spectrum is HALOFIT [81], which is included in CAMB [64]. However, in our own bispectrum analysis we will generally use the actual nonlinear power spectrum measured from the N -body simulations.

4.2.2 Matter bispectrum

At late times, the model of a Gaussian pdf for the density perturbation is oversimplified because nonlinear gravitational collapse will lead to higher order correlations. To study deviations from Gaussianity, it is useful to consider the bispectrum B_δ (or three-point correlator transform), which is defined by

$$\langle \delta(\mathbf{k}_1) \delta(\mathbf{k}_2) \delta(\mathbf{k}_3) \rangle = (2\pi)^3 \delta_D(\sum_i \mathbf{k}_i) B_\delta(k_1, k_2, k_3), \quad (4.3)$$

assuming again statistical homogeneity and isotropy. The Dirac delta function imposes a triangle constraint on the bispectrum arguments k_1, k_2, k_3 . The space of triangle configurations is sometimes called a tetrapyd [28] and is shown in Fig. 4.1. Additionally to the triangle constraint, we impose a cut off at k_{max} , corresponding to the smallest scale under consideration. The bispectrum shape, i.e. its dependence on the wavenumbers k_i , gives important information about the physics which induced the bispectrum. Not only does it help to disentangle late time non-Gaussianity, e.g. induced by nonlinear gravitational collapse, from primordial inflationary non-Gaussianity, but it also opens the intriguing possibility of distin-

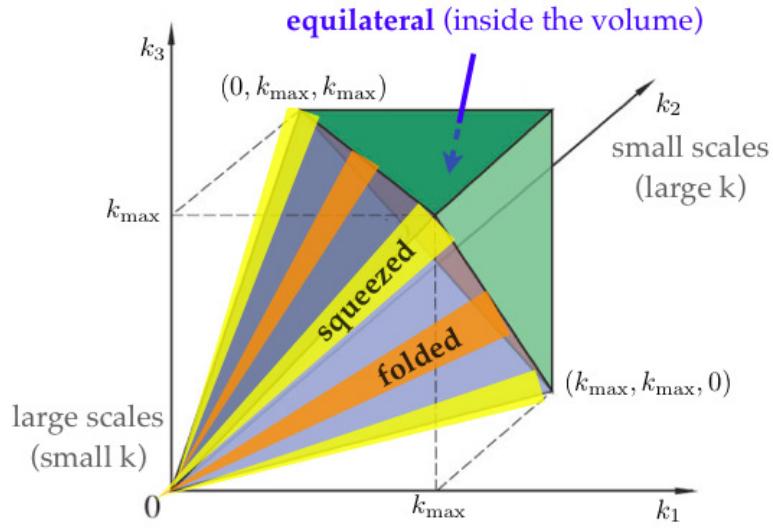


Figure 4.1: Space of triangles with sides k_1, k_2, k_3 , i.e. each point inside the tetrapyd volume corresponds to a triangle configuration. Squeezed, folded and equilateral configurations are highlighted.

guishing between different inflationary models. To see this explicitly we will review the perturbative treatment of such non-Gaussianities in the following sections.

4.2.3 Tree-level gravitational matter bispectrum

The equations of motion for a perfect pressureless fluid in a homogeneous and isotropic universe contain nonlinearities, which induce a non-vanishing bispectrum (see Section 2.4.4 or, e.g., [41, 110] for details). Specifically these are the $\nabla \cdot (\delta \mathbf{v})$ term in the continuity equation (2.123) and the $(\mathbf{v} \cdot \nabla)\mathbf{v}$ term in the Euler equation (2.124), where \mathbf{v} is the peculiar velocity, which describes deviations from the background Hubble flow. As reviewed in Sections 2.4.4 and 2.4.5, the induced tree level gravitational bispectrum is (see (2.144) and [40])

$$B_{\delta}^{\text{grav}}(k_1, k_2, k_3) = 2P_{\delta}^L(k_1)P_{\delta}^L(k_2)F_2^{(s)}(\mathbf{k}_1, \mathbf{k}_2) + 2 \text{ perms}, \quad (4.4)$$

where P_δ^L is the power spectrum of the linear perturbation δ_1 and

$$F_2^{(s)}(\mathbf{k}_1, \mathbf{k}_2) = \frac{5}{7} + \frac{1}{2} \frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{k_1 k_2} \left(\frac{k_1}{k_2} + \frac{k_2}{k_1} \right) + \frac{2}{7} \left(\frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{k_1 k_2} \right)^2. \quad (4.5)$$

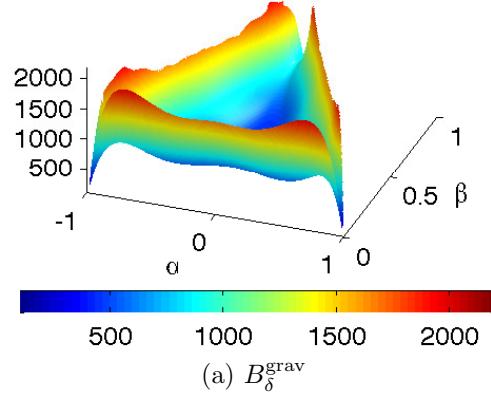
Here the first and second terms come, respectively, from $\delta \nabla \cdot \mathbf{v}$ and $\mathbf{v} \cdot \nabla \delta$ in the continuity equation, while the last term comes from $(\mathbf{v} \cdot \nabla) \mathbf{v}$ in Euler's equation [110].

To discuss the shape of this gravitational bispectrum let us consider two-dimensional slices through the tetrapyd shown in Fig. 4.1, with $k_1 + k_2 + k_3 = \text{const.}$ We denote one side of these two-dimensional triangular slices as α and parameterise the other direction with β .¹ A slice through the gravitational bispectrum (4.4) is shown in Fig. 4.2a (see [45] for more slices at slightly different length scales). The bispectrum is maximal at the edges of the plot, corresponding to flattened triangle configurations, where $k_1 + k_2 = k_3$ or permutations thereof, i.e. where the wavevectors \mathbf{k}_i are parallel or anti-parallel to each other. However there is a suppression in the corners of the plot, corresponding to squeezed triangle configurations with $k_1 \ll k_2 \approx k_3$ or permutations thereof. Non-flattened and equilateral triangle configurations in the centre of the plot are also suppressed. This tree level gravitational bispectrum is also illustrated in Fig. 4.3 as a function over the full tetrapyd.

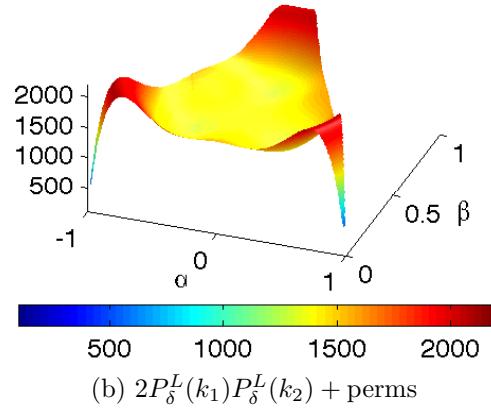
To understand the basic shape of the gravitational bispectrum (4.4) we plot in Fig. 4.2b the expression $2P_\delta^L(k_1)P_\delta^L(k_2) + 2$ perms, which corresponds to replacing $F_2^{(s)}$ in (4.4) by a constant. Comparing Fig. 4.2b with Fig. 4.2a shows that the configuration dependence of the kernel (4.5), which is induced by the terms containing scalar products, leads to an enhancement of flattened and particularly folded configurations, where two of the wavevectors \mathbf{k}_i equal each other. Non-flattened configurations are relatively suppressed. As we go along the edge of the plot in Fig. 4.2a, from folded ($k_1 = k_2 = k_3/2$ or permutations) to squeezed configurations ($k_1 \ll k_2 \approx k_3$ or permutations), the bispectrum shape reflects the shape of the power spectrum P_δ^L , which peaks at $k_{\text{eq}} \approx 0.02h/\text{Mpc}$ and then decreases with decreasing k_1 because $P_\delta^L(k_1) \propto k_1$ on large scales.

¹Details can be found in [49]. The slice parameters α, β should not be confused with the expansion coefficients $\alpha_n^{\{Q,R\}}$ and $\beta_n^{\{Q,R\}}$ to be defined later.

$$k_1 + k_2 + k_3 = 0.24 h/\text{Mpc}$$



$$k_1 + k_2 + k_3 = 0.24 h/\text{Mpc}$$



$$k_1 + k_2 + k_3 = 0.24 h/\text{Mpc}$$

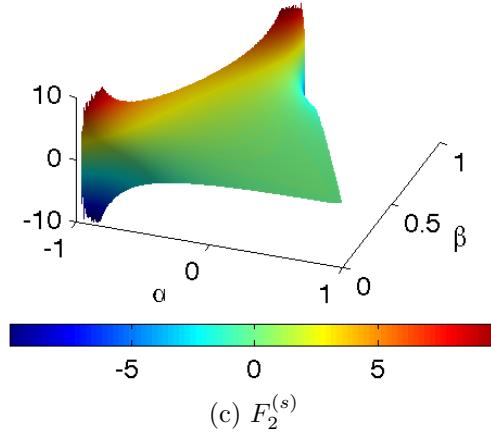


Figure 4.2: Two-dimensional slices of (a) the leading order gravitational bispectrum B_δ^{grav} from (4.4), (b) the expression $2P_\delta^L(k_1)P_\delta^L(k_2) + \text{perms}$ and (c) the kernel $F_2^{(s)}$ defined in (4.5) (with cut-offs in the squeezed limit where the kernel diverges). All slices are at fixed $k_1 + k_2 + k_3 = 0.24h/\text{Mpc}$ and we restrict the plot to $k_i \geq 0.0013h/\text{Mpc}$, corresponding to a $5000\text{Mpc}/h$ box. The linear power spectra are evaluated at $z = 30$.

Further discussion is required in the squeezed limit, where $\mathbf{k}_2 \approx -\mathbf{k}_3$. Let us consider the regime where $k_1 < k_{\text{eq}}$ and $k_2, k_3 > k_{\text{eq}}$. The term with $P_\delta^L(k_2)P_\delta^L(k_3)$ in (4.4) is small since the small scale power spectra decrease rapidly with increasing k_2, k_3 and $F_2^{(s)}(\mathbf{k}_2, \mathbf{k}_3)$ vanishes for $\mathbf{k}_2 = -\mathbf{k}_3$. The other two permutations in (4.4) depend on the angle between the large-scale wavevector \mathbf{k}_1 and the small-scale wavevectors $\mathbf{k}_2, \mathbf{k}_3$. First consider the case where the large scale is approximately perpendicular to the two small scales, i.e. $\mathbf{k}_1 \cdot \mathbf{k}_2 \approx -\mathbf{k}_1 \cdot \mathbf{k}_3 \approx 0$. Then $F_2^{(s)}(\mathbf{k}_1, \mathbf{k}_2) \approx F_2^{(s)}(\mathbf{k}_1, \mathbf{k}_3) \approx 5/7$ and $P_\delta^L(k_1)P_\delta^L(k_2)$ and $P_\delta^L(k_1)P_\delta^L(k_3)$ decrease as we approach more squeezed triangles, implying a suppressed bispectrum in the squeezed limit. In the other limit, the large-scale wavevector \mathbf{k}_1 is not orthogonal to the small scale wavevectors, and so is aligned with one or other of $\mathbf{k}_2, \mathbf{k}_3$. Then the squeezed limit $k_1 \rightarrow 0$ implies $F_2^{(s)}(\mathbf{k}_1, \mathbf{k}_2) \propto k_1^{-1} \rightarrow \infty$ and $F_2^{(s)}(\mathbf{k}_1, \mathbf{k}_3) \propto -k_1^{-1} \rightarrow -\infty$, which can be seen in Fig. 4.2c. However, in the sum

$$P_\delta^L(k_1)P_\delta^L(k_2)F_2^{(s)}(\mathbf{k}_1, \mathbf{k}_2) + P_\delta^L(k_1)P_\delta^L(k_3)F_2^{(s)}(\mathbf{k}_1, \mathbf{k}_3),$$

the two terms containing k_1^{-1} divergences in the kernel approximately cancel, because $k_2 \approx k_3$ and $\mathbf{k}_1 \cdot \mathbf{k}_2 \approx -\mathbf{k}_1 \cdot \mathbf{k}_3$. Also in the limit $k_1 \rightarrow 0$, the divergences of the kernels are regulated by the large-scale power spectrum because $k_1^{-1}P_\delta^L(k_1) = \text{const.}$ on very large scales. Fig. 4.2a shows that the divergences are indeed cancelled in the total bispectrum (4.4) and the squeezed limit is suppressed if the large-scale wavenumber satisfies $k_1 < k_{\text{eq}}$.

4.2.4 Gravitational matter bispectrum beyond tree level

4.2.4.1 Loop corrections

The tree level prediction for the gravitational matter bispectrum (4.4) is only a good approximation on large scales and can be improved by including so-called loop corrections, which were derived for Gaussian initial conditions in [110] and extended to include non-Gaussian initial conditions in [168]. Important loop corrections can be included simply by replacing the linear power spectrum P_δ^L by the nonlinear power spectrum P_δ in the tree level expression (4.4) for Gaussian initial

conditions, that is,

$$B_{\delta, \text{NL}}^{\text{grav}} \equiv 2P_\delta(k_1)P_\delta(k_2)F_2^{(s)}(\mathbf{k}_1, \mathbf{k}_2) + 2 \text{ perms}. \quad (4.6)$$

Although we expect this simple model for the gravitational bispectrum to fail on sufficiently small scales, we will later find the expression (4.6) to be a key approximation to the gravitational bispectrum in certain limits as we probe beyond the mildly nonlinear regime. Of course, this omits several loop corrections containing powers of $F_2^{(s)}$ or higher order kernels, but it is easy to evaluate with the nonlinear power spectrum from CAMB [64]. We leave a more quantitative discussion of further loop corrections for future work.

4.2.4.2 Halo models

At sufficiently small scales, well probed by our simulations, the perturbative treatment breaks down and simulations and phenomenological models must be used. In the strongly nonlinear regime the halo model can be used as a phenomenological model for the dark matter distribution (see [181] for a review). Recently in Ref. [182] it was demonstrated that combining 1-loop perturbation theory with the halo model describes the matter bispectrum in simulations at the $\mathcal{O}(10\%)$ level on nonlinear scales for equilateral and isosceles bispectrum slices. Similar results were obtained in Ref. [183], where local non-Gaussian initial conditions were also considered and compared with simulations for $k \leq 0.3h/\text{Mpc}$ at $z \leq 1$. In the mildly nonlinear regime, when the approximation that all dark matter particles are inside halos is no longer valid, the halo model becomes less accurate than in the strongly nonlinear regime. It should also be noted that important ingredients of the halo model like the halo density profile and halo mass function cannot be derived analytically but are obtained from fits to N -body simulations.

In detail, the halo model prediction for the dark matter bispectrum is given by three contributions corresponding to the three cases that the three points of the three point function lie in one, two or three halos (see e.g. [181, 182, 183]):

$$B_{\text{HM}} = B_{1\text{H}} + B_{2\text{H}} + B_{3\text{H}}, \quad (4.7)$$

where

$$B_{1H}(k_1, k_2, k_3) = \frac{1}{\bar{\rho}^3} \int dm n(m) \prod_{i=1}^3 \hat{\rho}(k_i, m), \quad (4.8)$$

$$B_{2H}(k_1, k_2, k_3) = \frac{1}{\bar{\rho}^3} \int dm_1 n(m_1) \hat{\rho}(k_1, m_1) \int dm_2 n(m_2) \hat{\rho}(k_2, m_2) \hat{\rho}(k_3, m_2) P_h(k_1, m_1, m_2) + \text{perms}, \quad (4.9)$$

$$B_{3H}(k_1, k_2, k_3) = \frac{1}{\bar{\rho}^3} \left[\prod_{i=1}^3 \int dm_i n(m_i) \hat{\rho}(k_i, m_i) \right] B_h(k_1, m_1; k_2, m_2; k_3, m_3). \quad (4.10)$$

Here $n(m)$ is the mass function and $\hat{\rho}(k, m)$ is the Fourier transform of the density profile of a halo of mass m , i.e. for spherically symmetric density profiles (like the commonly used NFW profile [184])

$$\hat{\rho}(k, m) = 4\pi \int dr r^2 \rho(r, m) \frac{\sin(kr)}{kr}. \quad (4.11)$$

P_h and B_h denote the halo power spectrum and bispectrum, which can be related to the tree level dark matter power spectrum and bispectrum on large scales using bias relations.

4.2.4.3 Constant ‘halo’ model

As we consider the more strongly nonlinear regime, $k \geq 1h/\text{Mpc}$ at $z = 0$, the halo model bispectrum is dominated by the 1-halo contribution B_{1H} . Neglecting the overall bispectrum amplitude, we find that in this regime the 1-halo shape for Gaussian initial conditions can be approximated by a simple constant bispectrum,

$$B_{\delta, \text{const}}^{\text{grav}}(k_1, k_2, k_3)|_z \equiv c_1 \bar{D}^{n_h}(z) (k_1 + k_2 + k_3)^\nu. \quad (4.12)$$

Here, ‘constant’ refers to constancy on $k_1 + k_2 + k_3 = \text{const}$ slices (see [49]). The expression (4.12) has an overall wavenumber scale-dependence with exponent ν and a time-dependence on the linear growth function $\bar{D}(z)$ with halo exponent n_h . The wavenumber scaling $\nu \approx -1.7$ is chosen such that it approximately reflects the scaling for equilateral triangle configurations in this regime as measured in our simulations (and which is approximately predicted by the halo model for

Gaussian initial conditions [182, 183]). The exponent n_h is similarly defined by the fast growth factor appropriate for the halo model for the scales under study $0.2h/\text{Mpc} \lesssim k \leq 2h/\text{Mpc}$, typically with $n_h \approx 6\text{--}8$. For $k_{\max} = 2h/\text{Mpc}$ and $z = 0$, this simple model achieves a shape correlation of 99.7% with the 1-halo contribution $B_{1\text{H}}$ and 99.3% with the full halo model bispectrum (4.7). While we found this approximation to be numerically accurate, analytic and physical motivations for the specific form (4.12) warrant further investigation (e.g. related to the facts that isolated point-like objects and shot noise induce constant bispectra).

In later sections, we will note that (4.12) provides an excellent approximation to the late-time bispectrum in the nonlinear regime when we use it in a simple fitting formula together with the modified tree-level gravitational bispectrum (4.6). We shall investigate the other halo contributions in more detail in the future.

As we shall see, on smaller nonlinear scales (large $k > 1h/\text{Mpc}$) the fast bispectrum growth begins to slow down by the present day $z = 0$. In this case, we should really replace the power law growth $D^{n_h}(z)$ for the constant mode using a more general growth factor $\mathcal{T}(\tilde{k}, z, z_i)$, where the ‘slice’ or ‘average’ wavenumber

$$\tilde{k} = \frac{1}{3}(k_1 + k_2 + k_3). \quad (4.13)$$

For future reference, it is convenient to use this growth rate in a more general integral form of the ‘constant’ model (4.12), that is,

$$\begin{aligned} B_{\delta, \text{const}}^{\text{grav}}(k_1, k_2, k_3)|_z &= \mathcal{T}(\tilde{k}, z, z_i) B_{\text{const}}^{\text{init}}(\tilde{k}, z_i) \\ &\equiv \int_{z_i}^z \mathcal{G}(\tilde{k}, z) dz B_{\text{const}}^{\text{init}}(\tilde{k}, z_i), \end{aligned} \quad (4.14)$$

where z_i is the redshift at which this rapid ‘halo’ growth takes hold (it has an implicit \tilde{k} dependence). The quantity $B_{\text{const}}^{\text{init}}(\tilde{k}, z_i)$ represents the initial condition at $z = z_i$ for the constant part of the bispectrum. Naively, we might take this to be $B_{\text{const}}^{\text{init}}(\tilde{k}, z_i) = B_{\delta}^{\text{grav}}(\tilde{k}, \tilde{k}, \tilde{k})|_{z=z_i}$, that is, the equilateral or constant part of the tree-level gravitational bispectrum (4.4) at $z = z_i$. However, to date, determining the amplitude of this initial constant bispectrum has relied on simulations. We will use this simple model to characterise both the time- and scale-dependence of the gravitational bispectrum, as well as primordial non-Gaussianity as we approach

the strongly nonlinear regime.

4.2.4.4 Alternative phenomenological fit

An alternative description of the gravitational matter bispectrum in the non-perturbative regime was proposed in Ref. [170], who constructed a fitting formula which interpolates between the perturbative prediction on large scales and a local-type bispectrum on small scales, which was suggested by early simulations. In their terminology, the small-scale bispectrum was denoted as a ‘constant reduced’ bispectrum, which implies that it takes the same form as the local shape $B_\delta \sim P_\delta(k_1)P_\delta(k_2) + \text{perms}$, in contrast to the constant shape (4.12) above. Recently Ref. [171] extended this fitting formula with updated simulations into mildly nonlinear scales $0.03h/\text{Mpc} \leq k \leq 0.4h/\text{Mpc}$ at $0 \leq z \leq 1.5$.

We review for the sake of completeness this phenomenological 9-parameter fit, for which we will test the regime of validity. The linear power spectrum in (4.4) is replaced by the nonlinear one P_δ , and the symmetrised kernel $F_2^{(s)}$ is replaced by

$$F_2^{(s)\text{eff}}(\mathbf{k}_1, \mathbf{k}_2) = \frac{5}{7}a(n_1, k_1)a(n_2, k_2) + \frac{1}{2}\frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{k_1 k_2} \left(\frac{k_1}{k_2} + \frac{k_2}{k_1} \right) b(n_1, k_1)b(n_2, k_2) + \frac{2}{7} \left(\frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{k_1 k_2} \right)^2 c(n_1, k_1)c(n_2, k_2), \quad (4.15)$$

where

$$a(n, k) = \frac{1 + \sigma_8^{a_6}(z)\sqrt{0.7Q_3(n)}(a_1 q)^{a_2+n}}{1 + (a_1 q)^{a_2+n}}, \quad (4.16)$$

$$b(n, k) = \frac{1 + 0.2a_3(n+3)(qa_7)^{n+3+a_8}}{1 + (qa_7)^{3.5+n+a_8}}, \quad (4.17)$$

$$c(n, k) = \frac{1 + 4.5a_4/[1.5 + (n+3)^4](qa_5)^{n+3+a_9}}{1 + (qa_5)^{3.5+n+a_9}}. \quad (4.18)$$

In these formulae, n represents the slope of the linear power spectrum at k , i.e. $n = d \ln P_\delta^L(k) / d \ln k$ (with an additional spline interpolation as described in [171]), $q = k/k_{\text{nl}}$ with k_{nl} defined by $k_{\text{nl}}^3 P_\delta^L(k_{\text{nl}}) / (2\pi^2) = 1$, and the function $Q_3(n)$ is

given by

$$Q_3(n) = \frac{4 - 2^n}{1 + 2^{n+1}}. \quad (4.19)$$

The best-fit values for the free parameters $a_1 - a_9$ were found by simulations [171]

$$\begin{aligned} a_1 &= 0.484, & a_2 &= 3.740, & a_3 &= -0.849, \\ a_4 &= 0.392, & a_5 &= 1.013, & a_6 &= -0.575, \\ a_7 &= 0.128, & a_8 &= -0.722, & a_9 &= -0.926. \end{aligned} \quad (4.20)$$

4.3 Primordial non-Gaussianity

Additional to the contribution B_δ^{grav} from nonlinear gravity, the matter bispectrum can have primordial contributions B_δ^{prim} from inflation or some other early universe model such as cosmic defects. While the simple model of single field slow roll inflation gives only a small primordial bispectrum, $f_{\text{NL}} \sim \mathcal{O}(10^{-2})$, other models can yield large non-Gaussianities with $f_{\text{NL}} > 1$ (see e.g. [16, 17, 18, 156, 185] for reviews). Such models can be distinguished if they induce different bispectrum shapes, i.e. different dependencies of the bispectra on the momenta k_1, k_2, k_3 as illustrated in Fig. 4.3. However, before reviewing primordial bispectrum shapes we describe how primordial non-Gaussianity changes the dark matter bispectrum.

4.3.1 Primordial contribution to the matter bispectrum

Let us assume that an inflationary model produces a primordial bispectrum B_Φ with nonlinear amplitude f_{NL} , i.e.¹

$$\langle \Phi(\mathbf{k}_1)\Phi(\mathbf{k}_2)\Phi(\mathbf{k}_3) \rangle = (2\pi)^3 \delta_D(\Sigma_i \mathbf{k}_i) f_{\text{NL}} B_\Phi(k_1, k_2, k_3). \quad (4.21)$$

¹Non-linear corrections to the mapping between \mathcal{R} and Φ , which are not taken into account here, induce an additional bispectrum of Φ , which is however smaller than the primordial contributions considered here [16].

4.3 Primordial non-Gaussianity

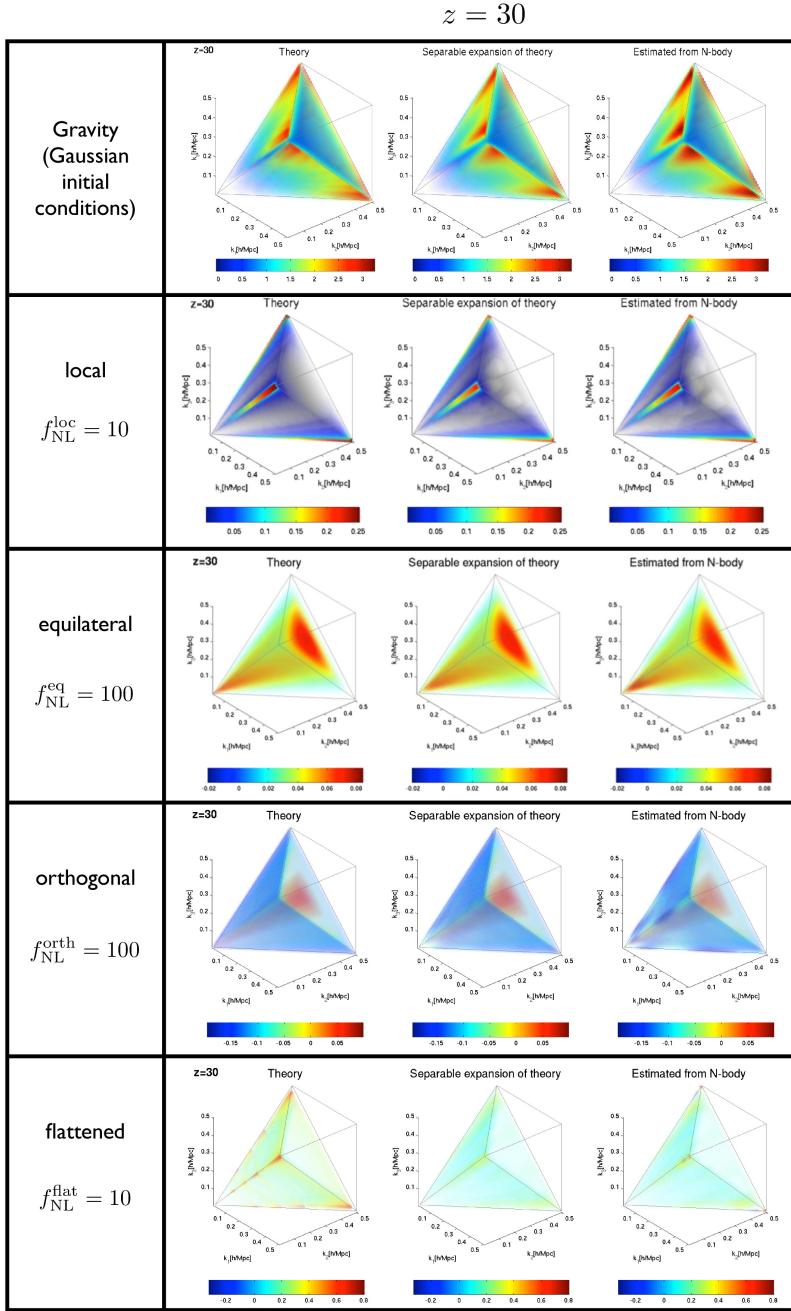


Figure 4.3: Illustration of separable expansion (4.38) for some theoretical tree level bispectra. Each black cell of 3 plots contains from left to right: Theoretical tree level prediction (4.4) or (4.22), separable expansion of the tree level theory (4.38) and estimated bispectrum from N -body simulations (4.51) (for simulations G512b, Loc10, Eq100, Orth100 and Flat10 from top to bottom; see Table 4.2 below). The plots show signal to noise weighted bispectra (as on the left hand side of (4.38)) on the tetrapyd domain in Fig. 4.1, evaluated at redshift $z = 30$ for 512^3 particles in a box with $L = 1600 \text{Mpc}/h$. The opaqueness of the points reflects the absolute value of the signal to noise weighted bispectrum (values close to 0 are completely transparent). Colors and transparency scales are consistent within each black cell but differ across different black cells. The plot axes are $k_1, k_2, k_3 \leq 0.5h/\text{Mpc}$. For better visibility we do not plot points with $k_1 + k_2 + k_3 > 2k_{\max}$ (corresponding to the green region in Fig. 4.1).

From the linear Poisson equation (4.2) we see that the primordial contribution to the matter bispectrum is given at leading order by

$$B_\delta^{\text{prim}}(k_1, k_2, k_3)|_z = M(k_1, z)M(k_2, z)M(k_3, z)f_{\text{NL}}B_\Phi(k_1, k_2, k_3), \quad (4.22)$$

which is valid on large scales.

A simple improvement to the tree level shape which incorporates some loop corrections can be obtained with the nonlinear power spectrum P_δ :

$$B_{\delta, \text{NL}}^{\text{prim}}(k_1, k_2, k_3) \equiv \sqrt{\frac{P_\delta(k_1)P_\delta(k_2)P_\delta(k_3)}{P_\Phi(k_1)P_\Phi(k_2)P_\Phi(k_3)}}B_\Phi(k_1, k_2, k_3), \quad (4.23)$$

where $P_\Phi(k_i)$ refers to the primordial power spectrum at some early time in matter domination. As will be demonstrated later in this chapter, this shape can be used to obtain simple fitting formulae for the primordial contribution to the matter bispectrum. A more systematic but also more cumbersome approach is to include loop corrections, which become important on small scales [168, 169]. We shall calculate more of these contributions in future work and test their correspondence to the N -body simulations.

We note that, while the time dependence of the tree level gravitational bispectrum (4.4) is given by $B_\delta^{\text{grav}} \propto D^4(z)$, the tree level primordial contribution (4.22) only grows like $B_\delta^{\text{prim}} \propto D^3(z)$, implying that it is easier to extract the primordial contribution to the dark matter bispectrum at early times. The simulations and fitting formulae discussed later will show that the gravitational bispectrum also grows faster than the primordial contribution in the strongly nonlinear regime (also by a factor of roughly $D(z)$).

4.3.2 Non-Gaussianity as a halo model time-shift

At sufficiently small scales the perturbative treatment breaks down and simulations and phenomenological models must be used. The phenomenological halo model prediction for local non-Gaussian initial conditions was computed in [183] by incorporating modified expressions for the halo model ingredients in presence of local non-Gaussianity. While first tests by [183] demonstrate that this approach works

well for some one-dimensional bispectrum slices in the mildly nonlinear regime, $k \leq 0.3h/\text{Mpc}$ and $z \leq 1$, comparisons to simulations in the strongly nonlinear regime have not been undertaken to date.

An alternative phenomenological model we propose here is to note that the primordial bispectrum can also contribute to the halo model bispectrum as an initial offset or time-shift. As we have seen the 1-halo model is highly correlated with the simple ‘constant’ bispectrum we described earlier (4.12). The key point is that this constant halo contribution grows much more rapidly than both the tree-level primordial bispectrum (4.22) or the tree-level gravitational bispectrum (4.4). Consequently, if the primordial bispectrum has a significant positive (or negative) constant component, then this can act as an initial condition for the halo bispectrum; the faster halo amplification growth will start earlier (or later) by a time or redshift offset Δz . The constant part of the primordial signal, then, will participate in the halo bispectrum growth and can be amplified much faster than expected from the tree-level result (4.22) or even with loop corrections (4.23). If this physical picture is correct, then the primordial constant contribution can be described perturbatively around the constant halo bispectrum (4.12) by expanding the growth factor

$$\bar{D}^{nh}(z + \Delta z) \approx \bar{D}^{nh}(z) + n_h \bar{D}^{nh-1}(z) \frac{d\bar{D}(z)}{dz} \Delta z. \quad (4.24)$$

We determine Δz by matching the projection of the total measured bispectrum \hat{B}_δ on the constant shape $B_{\delta,\text{const}}^{\text{grav}}$ in the Gaussian and non-Gaussian simulations:

$$(\hat{f}_{\text{NL}}^{\text{const}})_{\text{Gauss}}(z + \Delta z) = (\hat{f}_{\text{NL}}^{\text{const}})_{\text{NG}}(z). \quad (4.25)$$

From simulations, the time shift Δz and the corresponding shift of the growth function, $\Delta \bar{D} = \frac{d\bar{D}(z)}{dz} \Delta z$, are both found to vary over time at most by a factor of 2 for redshifts $10 \gtrsim z \gtrsim 1$ for local, equilateral and flattened initial conditions (to the extent to which this can be tested by linearly interpolating the limited number of output redshifts at which the bispectrum was measured). Provided $\Delta \bar{D}$ is time-independent this perturbation of the simple ‘constant’ model (4.12) means we

should be able to model the halo contribution from the primordial perturbation as

$$B_{\delta,\text{const}}^{\text{prim}}(k_1, k_2, k_3) \equiv c_2 \bar{D}^{n_h^{\text{prim}}}(z) (k_1 + k_2 + k_3)^\nu , \quad (4.26)$$

where we expect $n_h^{\text{prim}} = n_h - 1$ from (4.24) and again $\nu \approx -1.7$. The fitting parameter c_2 will be related to the correlation between the primordial shape¹ and the constant model at the time at which halos form for the length scale under consideration. Different non-Gaussian bispectrum models should show consistent behaviour in the nonlinear regime depending on the relative magnitude of their constant component. We shall define these quantities more precisely after introducing the bispectrum shape correlator in the next section, but (4.26) will be an important component in our later fitting formulae for primordial non-Gaussianity.

The simple power law in the linear growth function $\bar{D}(z)$ used to model the time dependence of the constant halo bispectrum can of course be extended to more general functions of time (4.14), whose time derivative would then enter in the expansion (4.24). The results presented later show that the overall normalisation of the simple fits can be improved by a more general modeling of the time dependence (reducing the overall growth rate in the strongly nonlinear regime). In this case, we can consider a time-shift for the more general ‘constant’ model (4.14)

$$\begin{aligned} B_{\delta,\text{const}}^{\text{prim}}(k_1, k_2, k_3)|_z &= B_{\delta,\text{const}}^{\text{grav}}(k_1, k_2, k_3)|_{z+\Delta z} - B_{\delta,\text{const}}^{\text{grav}}(k_1, k_2, k_3)|_z \\ &= \frac{d\mathcal{T}(\tilde{k}, z, z_i)}{dz} B_{\text{const}}^{\text{init}}(\tilde{k}, z_i) \Delta z + \dots , \end{aligned} \quad (4.27)$$

where we neglect higher order terms assuming them to be subdominant. We can determine the time-shift Δz by determining the magnitude of the constant part of the primordial bispectrum at $z = z_i$.

While more sophisticated models of the time evolution are left for future work, we note that the correlation with the measured bispectrum shape cannot be improved much, because simulations are so well described by a combination of the ‘constant’ shape (4.26) and the modified tree-level gravitational shape (4.23) used

¹We refer here to the excess bispectrum compared to Gaussian initial conditions as measured in N -body simulations at times when the constant contribution to the bispectrum is not negligible compared to the partially loop-corrected tree level contribution (4.23). Therefore c_2 is a fitting parameter to be determined from N -body simulations.

in the fitting formulae presented later.

The specific models of primordial non-Gaussianity studied in this chapter were introduced in Section 2.1.7.3.

4.4 Bispectrum estimation methodology

4.4.1 f_{NL} estimator

If our theoretical model is that the density perturbation δ has power spectrum $P_\delta(k)$ and bispectrum $f_{\text{NL}}^{\text{th}} B_\delta^{\text{th}}$, then the maximum likelihood estimator for the amplitude of this bispectrum in the limit of weak non-Gaussianity is given by (see [4, 50] and Section 2.1.7.4)¹

$$\hat{f}_{\text{NL}}^{\text{th}} = \frac{(2\pi)^3}{N_{\text{th}}} \int \frac{\prod_{i=1}^3 d^3 \mathbf{k}_i}{(2\pi)^6} \frac{\delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B_\delta^{\text{th}}(k_1, k_2, k_3) [\delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \delta_{\mathbf{k}_3} - 3 \langle \delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \rangle \delta_{\mathbf{k}_3}]}{P_\delta(k_1) P_\delta(k_2) P_\delta(k_3)} \quad (4.28)$$

where δ is the observed density perturbation. If this has a bispectrum B_δ^{obs} , i.e.

$$\langle \delta_{\mathbf{k}_1} \delta_{\mathbf{k}_2} \delta_{\mathbf{k}_3} \rangle = (2\pi)^3 \delta_D(\sum_i \mathbf{k}_i) B_\delta^{\text{obs}}(k_1, k_2, k_3), \quad (4.29)$$

then the expectation value of (4.28) is given by²

$$\langle \hat{f}_{\text{NL}}^{\text{th}} \rangle = \frac{1}{N_{\text{th}}} \frac{V}{\pi} \int_{\mathcal{V}_B} dV_k k_1 k_2 k_3 \frac{B_\delta^{\text{obs}}(k_1, k_2, k_3) B_\delta^{\text{th}}(k_1, k_2, k_3)}{P_\delta(k_1) P_\delta(k_2) P_\delta(k_3)}, \quad (4.30)$$

where $dV_k \equiv dk_1 dk_2 dk_3$, \mathcal{V}_B is the tetrahedral domain allowed by the triangle condition on the wavenumbers k_i , and V is a volume factor given by $V = (2\pi)^3 \delta_D(\mathbf{0}) =$

¹In the denominator we have assumed that $\langle \delta \delta \rangle$ is diagonal which is valid for a statistically homogeneous density perturbation. The linear term $\langle \delta \delta \rangle \delta$ in the numerator is written out for completeness but it is not used in our simulations because it vanishes for a statistically homogeneous field (because $\delta_{\mathbf{k}=0} = 0$).

²Performing the angular integrals shows for arbitrary $F(k_1, k_2, k_3)$

$$\int \frac{\prod_{i=1}^3 d^3 \mathbf{k}_i}{(2\pi)^9} (2\pi)^6 \delta_D^2(\sum_{j=1}^3 \mathbf{k}_j) F = \frac{V}{8\pi^4} \int_{\mathcal{V}_B} dk_1 dk_2 dk_3 k_1 k_2 k_3 F,$$

which corrects for a factor of $(2\pi)^3$ missing in [4, 50].

L^3 . Demanding $\langle \hat{f}_{\text{NL}}^{\text{th}} \rangle = 1$ for $B_{\delta}^{\text{th}} = B_{\delta}^{\text{obs}}$ fixes the normalisation such that

$$N_{\text{th}} = \frac{V}{\pi} \int_{\mathcal{V}_B} dV_k \frac{k_1 k_2 k_3 [B_{\delta}^{\text{th}}(k_1, k_2, k_3)]^2}{P_{\delta}(k_1) P_{\delta}(k_2) P_{\delta}(k_3)}. \quad (4.31)$$

4.4.2 Shape and size comparisons

Eq. (4.30) motivates the definition of a scalar product between two bispectrum shapes [50, 58] (see also (2.49)),

$$\langle B_i, B_j \rangle \equiv \frac{V}{\pi} \int_{\mathcal{V}_B} dV_k \frac{k_1 k_2 k_3 B_i(k_1, k_2, k_3) B_j(k_1, k_2, k_3)}{P_{\delta}(k_1) P_{\delta}(k_2) P_{\delta}(k_3)}, \quad (4.32)$$

which can be normalised to

$$\mathcal{C}(B_i, B_j) \equiv \frac{\langle B_i, B_j \rangle}{\sqrt{\langle B_i, B_i \rangle \langle B_j, B_j \rangle}}. \quad (4.33)$$

This shape correlation is a useful measure of the similarity of bispectrum shapes due to the reasons discussed after (2.50).

To measure the total integrated size of a bispectrum we define the squared norm as the cumulative signal to noise squared of the bispectrum [38, 51],

$$\|B\|^2 \equiv \frac{V}{(2\pi)^3} \int_{\mathcal{V}_B} dV_k \frac{s_B}{6} \frac{B^2(k_1, k_2, k_3)}{\text{var } B(k_1, k_2, k_3)} = \frac{\langle B, B \rangle}{6(2\pi)^3}, \quad (4.34)$$

which equals the quantity $(S/N)_B^2$ in [118]. The symmetry factor s_B is 6 if $k_1 = k_2 = k_3$, 2 if only two k_i equal each other and 1 if all k_i are different from each other. To obtain the expression on the right hand side of (4.34) we assumed that different Fourier modes are uncorrelated, $\langle \delta_{\mathbf{k}} \delta_{\mathbf{k}'} \rangle \propto \delta_D(\mathbf{k} + \mathbf{k}')$, and we only took the Gaussian contribution to the bispectrum noise into account, i.e.¹

$$\text{var } B(k_1, k_2, k_3) = (2\pi)^3 \frac{s_B P_{\delta}(k_1) P_{\delta}(k_2) P_{\delta}(k_3)}{8\pi^2 k_1 k_2 k_3}. \quad (4.35)$$

It is also convenient to normalise the total integrated bispectrum size with

¹We need an additional factor of $(2\pi)^3$ compared to [118] because $P_{\text{them}} = (2\pi)^3 P_{\text{us}}$ and $B_{\text{them}} = (2\pi)^3 B_{\text{us}}$.

respect to the linearly evolved local bispectrum defined in (2.38) and (4.22) (see [51] for a similar quantity in the CMB context):

$$\bar{F}_{\text{NL}}^B \equiv \frac{\|B\|}{\|B_{\delta}^{\text{prim,loc}}\|}. \quad (4.36)$$

We use the *linearly* evolved local bispectrum without loop corrections to obtain an expression which can be easily evaluated.

Using definitions (4.33) and (4.34), the expectation value (4.30) can be rewritten in the intuitive form

$$\langle \hat{f}_{\text{NL}}^{\text{th}} \rangle = \mathcal{C}(B_{\delta}^{\text{th}}, B_{\delta}^{\text{obs}}) \frac{\|B_{\delta}^{\text{obs}}\|}{\|B_{\delta}^{\text{th}}\|}. \quad (4.37)$$

Therefore the bispectrum amplitude f_{NL} can be interpreted as the projection of the observed bispectrum shape on the theoretical bispectrum shape, which is given by the cosine of the shapes times the ratio of their norms.

4.4.3 Separable mode expansion for fast f_{NL} estimation

Evaluating the estimator (4.28) is computationally very expensive because it requires $\mathcal{O}(N^6)$ operations for a grid with N points per dimension (typically $N = \mathcal{O}(10^3)$). For an efficient evaluation we expand the signal to noise weighted theoretical bispectrum in the separable form

$$\frac{\sqrt{k_1 k_2 k_3} B_{\delta}^{\text{th}}(k_1, k_2, k_3)}{\sqrt{P_{\delta}(k_1) P_{\delta}(k_2) P_{\delta}(k_3)}} = \sum_{n=0}^{n_{\text{max}}-1} \alpha_n^Q q_{\{r}(k_1) q_s(k_2) q_{t\}}(k_3), \quad (4.38)$$

where $\{rst\}$ denotes a symmetrisation over the indices, and a partial ordering has been introduced to enumerate the modes (for a fuller description see [28]). The expressions q_r could be any set of independent one-dimensional basis functions. For definiteness we choose q_r to be polynomials of order r defined in [28] and choose the slice ordering defined in Eq. (58) of [28]. We truncate the modal expansion (4.38) after $n_{\text{max}} = \mathcal{O}(50)$ modes, which provides an accurate representation for the shapes we are considering as illustrated in Fig. 4.3, which compares theoretical bispectra on the left with truncated expansions in the centre (and bispectra mea-

sured in N -body simulations on the right). More quantitative discussions of the error induced by the truncation of the expansions will be provided later in terms of shape correlations.

The separable expansion (4.38) allows us to write the estimator and its expectation value in the form

$$\hat{f}_{\text{NL}}^{\text{th}} = \frac{(2\pi)^3}{N_{\text{th}}} \sum_n \alpha_n^Q \int d^3\mathbf{x} [M_r(\mathbf{x}) M_s(\mathbf{x}) M_t(\mathbf{x}) - 3\langle M_{\{r}(\mathbf{x}) M_s(\mathbf{x}) \rangle M_{t\}}(\mathbf{x})], \quad (4.39)$$

$$\langle \hat{f}_{\text{NL}}^{\text{th}} \rangle = \frac{1}{N_{\text{th}}} \sum_{nm} \alpha_n^Q \alpha_m^Q \gamma_{nm}, \quad (4.40)$$

where

$$\gamma_{nm} = \frac{V}{\pi} \int_{V_B} dV_k Q_n(k_1, k_2, k_3) Q_m(k_1, k_2, k_3), \quad (4.41)$$

$$M_r(\mathbf{x}) = \int \frac{d^3\mathbf{k}}{(2\pi)^3} \frac{\delta_{\mathbf{k}} q_r(k)}{\sqrt{k} P_{\delta}(k)} e^{i\mathbf{k}\cdot\mathbf{x}} \quad (4.42)$$

and

$$Q_n(k, k', k'') \equiv q_{\{r}(k) q_s(k') q_{t\}}(k''). \quad (4.43)$$

This form of the estimator can be evaluated efficiently because it involves only Fourier transforms and one three-dimensional integral over position space, leading to $\mathcal{O}(n_{\max} N^3)$ operations instead of the $\mathcal{O}(N^6)$ of the brute force estimator (where $N = \mathcal{O}(10^3)$).

To motivate the choice of the prefactor in the expansion (4.38) let us transform to basis functions $R_n(k_1, k_2, k_3)$ that are orthonormal on the tetravert domain, $\langle R_n, R_m \rangle_{\text{noweight}} = \delta_{nm}$, with respect to the unweighted scalar product

$$\langle B_i, B_j \rangle_{\text{noweight}} \equiv \frac{V}{\pi} \int_{V_B} dV_k B_i(k_1, k_2, k_3) B_j(k_1, k_2, k_3). \quad (4.44)$$

4.4 Bispectrum estimation methodology

From $\sum_n \alpha_n^Q Q_n = \sum_n \alpha_n^R R_n$ and Eq. (4.38) we find that

$$B_\delta^{\text{th}}(k_1, k_2, k_3) = \sum_n \alpha_n^R B_n^R(k_1, k_2, k_3), \quad (4.45)$$

where the contributions

$$B_n^R(k_1, k_2, k_3) \equiv \sqrt{\frac{P_\delta(k_1)P_\delta(k_2)P_\delta(k_3)}{k_1 k_2 k_3}} R_n(k_1, k_2, k_3) \quad (4.46)$$

to the full bispectrum are orthonormal with respect to (4.32). Thus the expansion coefficients α_n^R measure the size of contributions to the bispectrum which are orthonormal with respect to the signal to noise weighted scalar product (4.32).

Note that when applying the bispectrum estimator in the non-perturbative regime, the power spectra appearing in e.g. (4.38), (4.42) and (4.46) correspond to full non-perturbative power spectra of the density field under consideration, because these are the 2-point statistics relevant for deriving the maximum-likelihood estimator (4.28). In our implementation we use (ensemble-averaged) power spectra estimated directly from our N -body simulations.

As a side remark, note that in contrast to the signal to noise weighted scalar product (4.32) induced by the estimator expectation value, the weight of the scalar product (4.44) is time-independent and does not depend on the range of scales under consideration. Therefore we can use the same set of orthonormal polynomials R_n at any time and for all length scales. If the expansion (4.38) does not converge as well as for the applications considered here, different (separable) scalar product weights can be used to orthonormalise the R_n 's.

4.4.4 Fast modal bispectrum estimator

Importantly the separable mode expansion allows us not only to measure $f_{\text{NL}}^{\text{th}}$ efficiently for any given theoretical bispectrum, but it also allows us to reconstruct the full bispectrum in a model-independent manner by measuring a set of independent f_{NL} 's and summing up the corresponding contributions to the bispectrum [4, 50]. To see this, note that the transformation $R_n = \sum_m \lambda_{nm} Q_m$ implies $\alpha_n^Q = \sum_p \lambda_{pn} \alpha_p^R$ and from $\langle R_n, R_m \rangle_{\text{noweight}} = \delta_{nm}$ we get $\gamma = \lambda^{-1}(\lambda^{-1})^T$. Plugging

this into (4.39) gives

$$\hat{f}_{\text{NL}}^{\text{th}} = \frac{1}{N_{\text{th}}} \sum_m \alpha_m^R \beta_m^R, \quad (4.47)$$

with

$$\beta_m^R \equiv (2\pi)^3 \sum_n \lambda_{mn} \int d^3\mathbf{x} [M_r M_s M_t - 3\langle M_{\{r} M_s \rangle M_{t\}}], \quad (4.48)$$

where the index n labels the combination r, s, t . If $B_{\delta}^{\text{th}} = \langle B_{\delta}^{\text{obs}} \rangle$ then (4.40) becomes

$$\langle \hat{f}_{\text{NL}}^{\text{th}} \rangle = \frac{1}{N_{\text{th}}} \sum_m \alpha_m^R \alpha_m^R. \quad (4.49)$$

Equations (4.47) and (4.49) imply

$$\langle \beta_n^R \rangle = \alpha_n^R. \quad (4.50)$$

Therefore we can estimate the full bispectrum with [4, 50]

$$\hat{B}_{\delta}(k_1, k_2, k_3) = \sum_{n=0}^{n_{\text{max}}-1} \beta_n^R B_n^R(k_1, k_2, k_3), \quad (4.51)$$

where B_n^R is given by (4.46). Intuitively the coefficients β_n^R measure the components of the orthonormal basis bispectra B_n^R in the data and \hat{B}_{δ} measures the projection of the bispectrum on the subspace of bispectra spanned by the B_n^R .

This approach is extremely efficient with the calculation of each β_n^R coefficient being equivalent to a single 3D integration over products of fast Fourier transforms. While theoretically a complete basis would require N^3 modes in practice far fewer modes ($\mathcal{O}(50)$) are necessary to reconstruct the bispectrum accurately.

The primordial contribution to the matter bispectrum will be extracted by measuring the difference of the bispectrum between non-Gaussian and Gaussian

initial conditions,

$$\hat{B}_{\text{NG}} = \sum_{n=0}^{n_{\max}-1} [\beta_n^R - (\beta_n^R)_{\text{Gauss}}] B_n^R. \quad (4.52)$$

As will be discussed in more detail later, the difference is computed mode by mode to reduce error bars.

The time-dependence of the expansion coefficients α_n^R and the β_n^R coefficients can be read off from (4.38) to be $(\alpha_n^R)_{\text{grav}} \propto D(z)$ for the gravitational bispectrum and $(\alpha_n^R)_{\text{NG}} \propto D(z)^0$ for primordial bispectra, assuming tree level perturbation theory where $P_\delta \propto D^2(z)$. In practice we must use the nonlinear power in (4.38) and (4.42) implying slightly different time dependences for α_n^R in the nonlinear regime (see e.g. the middle panel of Fig. 4.14a, which will be discussed in more detail later).

4.4.5 Fast modal bispectrum correlations

To analyse the reconstructed bispectrum and to control the accuracy of the separable method, the following shape correlations are computed using (4.33):

$$\mathcal{C}_{\alpha,\text{th}} \equiv \mathcal{C} \left(\sum_n \alpha_n^R B_n^R, B_\delta^{\text{th}} \right) = \sqrt{\frac{\sum_n (\alpha_n^R)^2}{\langle B_\delta^{\text{th}}, B_\delta^{\text{th}} \rangle}}, \quad (4.53)$$

$$\mathcal{C}_{\beta,\alpha} \equiv \mathcal{C} \left(\langle \hat{B}_\delta \rangle_{\text{sim}}, \sum_n \alpha_n^R B_n^R \right) = \frac{\sum_n \alpha_n^R \langle \beta_n^R \rangle_{\text{sim}}}{\sqrt{\sum_m (\alpha_m^R)^2 \sum_p \langle \beta_p^R \rangle_{\text{sim}}^2}}, \quad (4.54)$$

$$\mathcal{C}_{\beta,\text{th}} \equiv \mathcal{C} \left(\langle \hat{B}_\delta \rangle_{\text{sim}}, B_\delta^{\text{th}} \right) = \mathcal{C}_{\beta,\alpha} \mathcal{C}_{\alpha,\text{th}}. \quad (4.55)$$

Here $\langle \hat{B}_\delta \rangle_{\text{sim}}$ denotes the average of (4.51) over independent simulations and we truncate all sums appearing in the above expressions and in (4.51) after n_{\max} modes. The first shape correlation measures how well the separable expansion (4.38) of the theoretical bispectrum approximates the theoretical bispectrum. The second shape correlation quantifies how well the estimated bispectrum $\langle \hat{B}_\delta \rangle_{\text{sim}}$ agrees with the separable expansion of the theoretical bispectrum. The product of these two shape correlations gives the shape correlation between the reconstructed

bispectrum and the theoretical bispectrum.

4.4.6 Cumulative measures of non-Gaussianity

Additionally to these shape correlations we will also compare the projections (4.47) of the measured bispectrum on theory bispectra, as well as the cumulative signal to noise of the average reconstructed bispectrum, which can be expressed in terms of the measured average $\langle \beta_n^R \rangle_{\text{sim}}$ as¹

$$\|\langle \hat{B}_\delta \rangle_{\text{sim}}\| = \sqrt{\frac{\sum_n \langle \beta_n^R \rangle_{\text{sim}}^2}{6(2\pi)^3}}, \quad \bar{F}_{\text{NL}}^{\text{sim}} = \sqrt{\frac{\sum_n \langle \beta_n^R \rangle_{\text{sim}}^2}{N_{B_\delta^{\text{prim}, \text{loc}}}}}. \quad (4.56)$$

The norm was defined in (4.34) using the cumulative signal to noise squared of the bispectrum. As a consistency check we check if (4.56) also measures the cumulative signal to noise squared of the β_n^R coefficients. The theoretical covariance of β_n^R involves the 6-point function of the density perturbation. Taking only the Gaussian contribution into account and assuming that different Fourier modes are uncorrelated we find

$$\langle \beta_n^R \beta_m^R \rangle = 6(2\pi)^3 \delta_{nm}. \quad (4.57)$$

This predicts $\sigma_{\beta_n^R} = \sqrt{6(2\pi)^3}$ for Gaussian simulations and therefore confirms that (4.56) measures the cumulative signal to noise of the reconstructed bispectrum. We confirmed the prediction for $\sigma_{\beta_n^R}$ quantitatively in Gaussian simulations, but we have not assessed non-Gaussian corrections to the predicted noise (see e.g. [186]). Note that our plots show 1σ sample standard deviations obtained by running realisations with different random number seeds.

4.4.7 Towards experimental setups

For the bispectrum measurements presented below we will drop the second term in the square brackets in (4.48), i.e. we assume that different modes of the density perturbation are uncorrelated, $\langle \delta_{\mathbf{k}} \delta_{\mathbf{k}'} \rangle \propto \delta_D(\mathbf{k} + \mathbf{k}')$, which is true for a statistically

¹Alternatively using $\langle \sum_n (\beta_n^R)^2 \rangle_{\text{sim}}$ would add undesirable contributions from the variances of the β_n^R , which become relevant if the measurements are not signal-dominated.

homogeneous field. In experimental setups, which are not considered here, off-diagonal mode couplings from inhomogeneous noise can be incorporated in (4.48) with the linear term $\langle MM \rangle M$, without affecting the efficiency of the bispectrum estimator (4.51). Ref. [186] explores how modifications to the denominator of the estimator (4.28) in presence of off-diagonal mode couplings can be incorporated efficiently with modal expansions and an implementation for CMB experiments has been developed successfully (see also [187]).

4.4.8 Comparison to other bispectrum estimators

An alternative way to reconstruct the bispectrum is to estimate the three-point correlation function directly by averaging over orientations of triangles. This is computationally feasible if the analysis is restricted to the largest simulated scales or to 1D slices through the bispectrum (e.g. [161, 169, 188]). Alternatively, one can take a subset of all possible triangle orientations that yields the desired accuracy of the bispectrum, but is still noisier than estimators incorporating all triangle orientations (e.g. [171, 189]). Both approaches are typically combined with binning in k -space and reduce the available *data* (i.e. the number of used triangle configurations and/or orientations) to achieve efficiency.

In contrast to these methods, our estimator takes all possible triangle configurations and orientations down to the smallest simulated scales into account but is nevertheless extremely efficient. For example, with our method the extraction of the bispectrum of a 1024^3 grid takes about one hour on 6 cores, which is only a small fraction of the time required to run an N -body simulation of this size. This is achieved by reducing the allowed space of *theory* bispectra using a truncated separable basis. This leads to a further advantage of our estimator: It compresses the shape information of the bispectrum, which is a function of the whole three-dimensional tetrahedral domain allowed by the triangle condition, to n_{\max} numbers β_n^R , where n_{\max} is the number of basis functions used in the expansion (4.38). This simplifies further processing of the bispectrum, e.g. for comparisons with theoretical models, generation and testing of fitting formulae, or modeling bispectrum covariances in presence of experimental realism. The important issue is that the basis functions provide a convergent expansion for the model under investigation.

Once β_n^R is measured from the data, we can not only obtain the full bispectrum with (4.51), but we can also calculate full three-dimensional shape correlations, the cumulative signal to noise and the nonlinearity parameter $f_{\text{NL}}^{\text{th}}$ associated to any theoretical bispectrum B_{δ}^{th} without additional computational cost, using (4.54), (4.56) and (4.47), respectively. In presence of inhomogeneous noise we can in principle include off-diagonal covariance elements $\langle \delta_{\mathbf{k}} \delta_{\mathbf{k}'} \rangle$ in a straightforward way. Finally our approach allows us to estimate the trispectrum efficiently [50], which has been shown for a class of trispectrum shapes in [4] (see also [54, 162]), but we leave the application to N -body simulations for future work. Note that we provide a critical discussion of our bispectrum estimation methodology in light of previous analyses in Sec. 4.5.5 below, where we also mention some potential disadvantages.

4.4.9 Bispectrum visualisation

Often in the literature the bispectrum is visualised by plotting one- or two-dimensional slices through the tetrapyd shown in Fig. 4.1. E.g. the plots in [58] correspond to two-dimensional slices through the tetrapyd obtained by varying k_2 and k_3 at fixed k_1 and removing the $k_3 > k_2$ part of the slice by symmetry. Some plots in [49] and in this chapter show two-dimensional tetrapyd slices with $k_1 + k_2 + k_3 = \text{const.}$ While such slices are sufficient for scale-invariant primordial bispectra, late time bispectra typically have different triangle dependences at different overall scales $k_1 + k_2 + k_3$. This motivates plotting late time bispectra on the full three-dimensional tetrapyd instead of plotting particular slices through the tetrapyd.

Instead of the unweighted bispectrum the so-called reduced bispectrum

$$Q(k_1, k_2, k_3) \equiv \frac{B_{\delta}(k_1, k_2, k_3)}{P_{\delta}(k_1)P_{\delta}(k_2) + 2\text{perms}} \quad (4.58)$$

is often shown, because it reduces the dynamical range of the bispectrum and for the tree level gravitational contribution it is independent of time, overall scale and power spectrum normalisation and almost independent of cosmology [41]. Instead we will plot the signal to noise weighted bispectrum

$$\sqrt{\frac{k_1 k_2 k_3}{P_{\delta}(k_1)P_{\delta}(k_2)P_{\delta}(k_3)}} B_{\delta}(k_1, k_2, k_3) \quad (4.59)$$

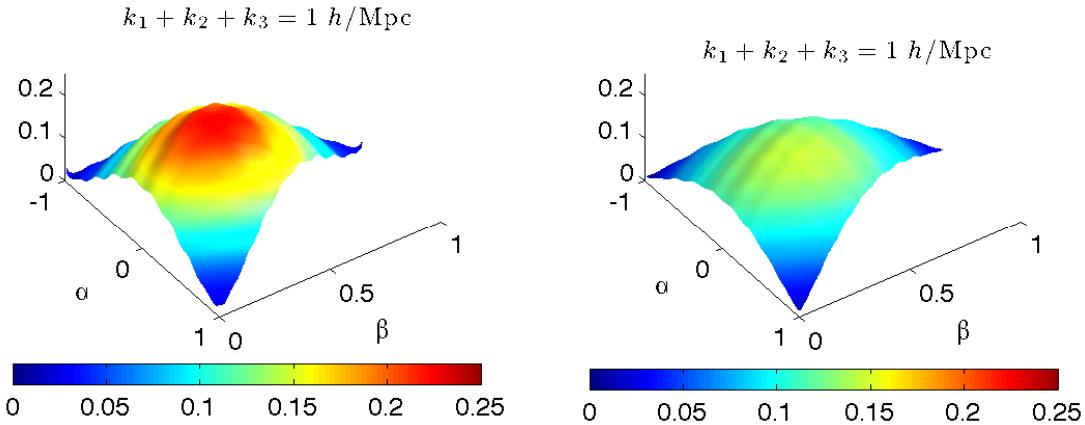


Figure 4.4: Bispectrum weights $[P_\delta(k_1)P_\delta(k_2) + 2 \text{ perms}]^{-1}$ (left) and $\sqrt{k_1k_2k_3/[P_\delta(k_1)P_\delta(k_2)P_\delta(k_3)]}$ (right) evaluated with CAMB [64] at redshift $z = 30$ on slices with $k_1 + k_2 + k_3 = 1 h/\text{Mpc}$.

because then the product of the shown functions gives the scalar product defined in (4.32) [58], visual similarity indicates a high shape correlation (4.33) and the dynamical range of the unweighted bispectrum is also greatly reduced, which is of advantage for plotting purposes. For linearly evolved primordial bispectra (4.22), the signal to noise weighted bispectrum is time-independent in the linear regime.

Qualitatively the two weights are quite similar in the regime relevant for most plots of this chapter, see Fig. 4.4. However they differ in the squeezed limit, because for decreasing k_1 , $P_\delta(k_1)P_\delta(k_2)$ turns over at $k_1 = k_{\text{eq}}$, whereas

$$\sqrt{\frac{P_\delta(k_1)P_\delta(k_2)P_\delta(k_3)}{k_1k_2k_3}} \quad (4.60)$$

turns constant where $P(k_1)/k_1$ turns constant, which is on somewhat larger scales than k_{eq} . While it is not straightforward to deduce the squeezed limit of an unweighted bispectrum from a plot of its weighted form, one can compare plots of different bispectra if they are weighted in the same way to deduce the relative behavior in the squeezed limit.

4.5 Simulation setup, initial conditions and validation

4.5.1 N -body simulations setup

We use the separable mode expansion method described in [4, 50] and Chapter 3 to generate realisations of the initial primordial potential $\Phi = -3\mathcal{R}/5$ during matter domination with the desired primordial power spectrum and bispectrum. For 512^3 particles this takes about 10 minutes per seed on one core and works for separable as well as non-separable bispectra, therefore being more general than other proposed methods [159, 161]. From Φ we calculate the linear density perturbation δ at the initial redshift of the simulation with the Poisson equation (4.2). We then use this initial density perturbation to displace the initial particles from an unperturbed distribution with the 2LPT method [129, 130], which also determines the initial particle velocities. Then the N -body code Gadget-3 [132, 190] simulates the time evolution until today and we use a cloud in cell scheme to calculate the density perturbation δ of the particle distribution on a grid at different redshifts. After deconvolving δ with the cloud in cell kernel we compute the power spectrum P_δ and the coefficients β_n^R from (4.48) using $n_{\max} = 50$ modes. Finally we reconstruct the full bispectrum with (4.51) and calculate its norm \hat{F}_{NL} (4.56) as well as its shape correlation $\mathcal{C}_{\beta,\alpha}$ (4.54) with theoretical bispectra and its nonlinear amplitude (4.47).

Table 4.2 lists the parameters of the N -body simulations that were performed in this work. All simulations assume a flat Λ CDM model with the WMAP-7 [163] parameters $\Omega_b h^2 = 0.0226$, $\Omega_c h^2 = 0.11$, $\Omega_\Lambda = 0.734$, $h = 0.71$, $\tau = 0.088$, $\Delta_{\mathcal{R}}^2(k_0) = 2.43 \times 10^{-9}$ and $n_s(k_0) = 0.963$, where $k_0 = 0.002 \text{Mpc}^{-1}$. Note that the primordial power spectrum is given by

$$P_\Phi(k) = \frac{9}{25} \frac{2\pi^2}{k^3} \Delta_{\mathcal{R}}^2(k_0) \left(\frac{k}{k_0} \right)^{n_s-1}. \quad (4.61)$$

Name	NG shape	f_{NL}	$L[\frac{\text{Mpc}}{h}]$	N_p	z_i	$L_s[\frac{\text{kpc}}{h}]$	N_r	glass
G512g	–	–	1600	512	49	156	3	yes
G512	–	–	1600	512	49	156	3	no
G_L^{512}	–	–	{400, 100}	512	49	{39, 9.8}	3	no
G768	–	–	2400	768	19	90	3	no
G1024	–	–	1875	1024	19	40	2	no
Loc10g	local	10	1600	512	49	156	3	yes
Loc10	local	10	1600	512	49	156	3	no
Loc10_L^{512}	local	10	{400, 100}	512	49	{39, 9.8}	3	no
Loc10^-	local	–10	1600	512	49	156	3	no
Loc20	local	20	1600	512	49	156	3	no
Loc50	local	50	1600	512	49	156	3	no
Eq100g	equil	100	1600	512	49	156	3	yes
Eq100	equil	100	1600	512	49	156	3	no
Eq100_L^{512}	equil	100	{400, 100}	512	49	{39, 9.8}	3	no
Eq100^-	equil	–100	1600	512	49	156	3	no
Orth100g	orth	100	1600	512	49	156	3	yes
Orth100	orth	100	1600	512	49	156	3	no
$\text{Orth100}_{400}^{512}$	orth	100	400	512	49	39	3	no
Orth100^-	orth	–100	1600	512	49	156	3	no
Flat10	flat	10	1600	512	49	156	3	no
$\text{Flat10}_{400}^{512}$	flat	100	400	512	49	39	3	no

Table 4.2: Parameters of N -body simulations: Non-linearity parameter f_{NL} , box size L , number of particles per dimension N_p , initial redshift of the simulations z_i , softening length L_s and number of realisations (i.e. random seeds) N_r for each parameter set. ‘glass’ indicates if the initial particles were displaced from a regular grid or from a glass configuration. Initial conditions for non-local non-Gaussian simulations were generated with the separable method described in [4, 50]. All simulations use 2LPT [129, 130] to get the initial particle distribution, which is then evolved with Gadget-3 [132, 190].

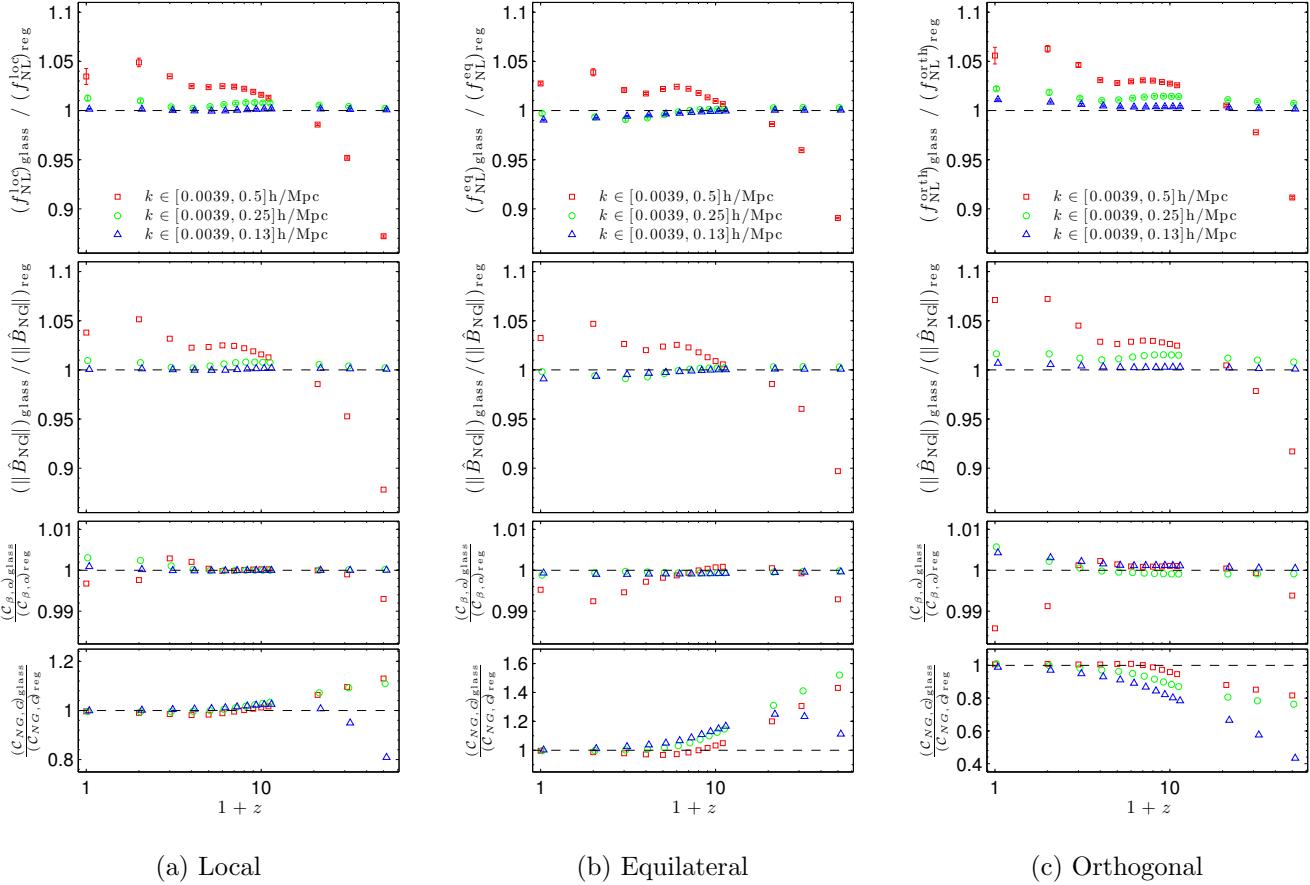


Figure 4.5: Impact of glass initial conditions on measured non-Gaussian bispectra in simulations Loc10, Eq100 and Orth100. We plot the ratio to simulations with regular grid initial conditions. The top panel contains error bars obtained by calculating the sample standard deviation of the ratio for each seed. We do not show error bars in the other panels because they depend on $\sum_n \langle \beta_n^R \rangle_{\text{sim}}^2$ and are therefore more difficult to estimate. However they should be similar to the errors in the top panel because of (4.37). The large deviations in the bottom panel are mainly due to the different gravitational bispectra for glass and regular grid initial conditions as shown in Fig. 4.14c.

4.5.2 Regular grid vs glass initial conditions

For the unperturbed particle distribution, from which initial particles are displaced using the 2LPT method, we either use a regular grid or a glass configuration, which is obtained by placing particles randomly in the box and then evolving them with the sign of gravity flipped in Gadget-3 [131, 132]. A disadvantage of glass initial conditions is that they require generating a glass and specifying when the glass configuration has converged (see e.g. [191] for a discussion). In contrast, regular grid initial conditions are easier to set up but break statistical isotropy, which may have implications for structure formation. We will compare the different initial condition setups first for Gaussian and then non-Gaussian initial conditions.

Measured bispectra for Gaussian N -body simulations will be shown in Figures 4.14 and 4.15 in Sec. 4.6. Simulations started from glass and regular grid initial conditions are shown in Fig. 4.14a and Fig. 4.14b, respectively. At early times regular grid initial conditions produce a bispectrum which is more than three times as large as the tree level prediction. This significant spurious bispectrum is not present for glass initial conditions and must therefore be caused by the breaking of isotropy induced by the regular grid. Fig. 4.14b shows that the spurious bispectrum decays with time as the regular grid structure is washed out by the growing gravitational perturbations. At late times, $z \leq 3$, the bispectra from regular grid and glass initial conditions differ by at most 10% in their cumulative signal to noise and by less than 0.5% in their shape correlation to the tree level prediction, see Fig. 4.14c.

We conclude that at the 10% accuracy level (for Gaussian simulations with $L = 1600\text{Mpc}/h$, 512^3 particles and $k_{\max} = 0.5h/\text{Mpc}$), both glass and regular grid initial conditions can be used to extract the gravitational bispectrum from Gaussian simulations as long as the bispectrum is measured at low redshifts when the regular grid is washed out. The difference between regular grid and glass initial conditions decreases when reducing the cutoff k_{\max} used for the bispectrum estimation because then structures of the size of the grid separation are smoothed out.

For non-Gaussian simulations, Fig. 4.5 shows how glass initial conditions change the non-Gaussian excess bispectrum \hat{B}_{NG} (4.52) compared to regular grid initial

conditions. At $z \leq 10$ the shape of \hat{B}_{NG} is consistent between glass and regular grid initial conditions at the 1.5% level for the simulations Loc10, Eq100 and Orth100. However, the total integrated bispectrum \bar{F}_{NL} differs by up to 7% for $z \leq 10$. At earlier times the deviations are somewhat larger. Compared to the full gravitational bispectrum for Gaussian initial conditions shown in Fig. 4.14a, the impact of glass initial conditions on \hat{B}_{NG} is quite small, possibly because the spurious bispectrum due to the regular grid is partly subtracted out when calculating $\hat{B}_{\text{NG}} = \hat{B} - \hat{B}_{\text{Gauss}}$.

Whether regular grid or glass initial conditions are representing the statistics of the non-Gaussian density perturbation more successfully is not unambiguous. Regular grid initial conditions introduce a spurious contribution to the full bispectrum \hat{B} at early times, which can be avoided using glass initial conditions. However glass initial conditions represent the primordial contribution \hat{B}_{NG} slightly less accurately than regular grid initial conditions at the starting redshift of the simulation. However these effects impact \hat{B}_{NG} only at the $\mathcal{O}(5\%)$ level at $z \leq 10$, so that both glass or regular grid initial conditions can be used at this level of precision. As for Gaussian simulations the two initial condition methods agree better when k_{max} is reduced (see Fig. 4.5) and grid discretization effects are minimised.

4.5.3 Validation and convergence tests

First we test the setup of the initial conditions and the N -body simulations by comparing the measured matter power spectrum with the power spectrum predicted by linear theory and by CAMB [64, 81] in Fig. 4.6. Next we perform a simple test of the bispectrum estimator by distributing particles randomly in a box and comparing the measured bispectrum with the pure shot noise bispectrum $B_{\delta}^{\text{shot}} = L^6/N_p^6$ [112], see Fig. 4.7. These two tests show that the basic setup of the simulations is correct and that the separable bispectrum estimator works well.

To check convergence of the N -body simulations and the bispectrum measurements we compare results for box sizes $L = \{1600, 400, 100\}\text{Mpc}/h$ with $N_p^3 = 512^3$ particles. For each simulation we measure the bispectrum from $k_{\text{min}} = 2\pi/L$ to $k_{\text{max}} = \{1, 0.5, 0.25\} \times (\frac{N_p}{4} \frac{2\pi}{L})$ by imposing a sharp k filter on the density perturbation $\delta_{\mathbf{k}}$.

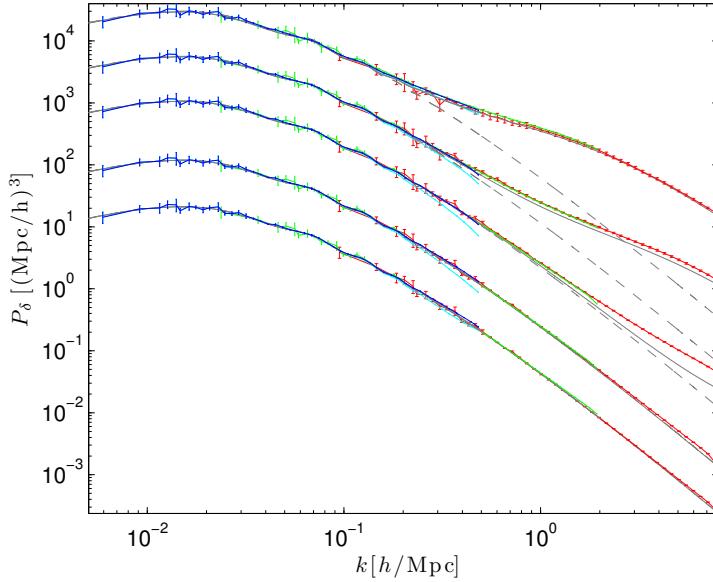


Figure 4.6: Measured matter power spectra in simulations G_{100}^{512} (red), G_{400}^{512} (green), $G512$ (blue) and $G512g$ (cyan) in comparison with power from linear perturbation theory (grey dashed) and CAMB (grey solid, [64, 81]). See Table 4.2 for parameters. Curves from bottom to top correspond to redshifts $z = 49, 20, 6, 2, 0$.

The measured late time bispectra are shown in Fig. 4.8 as functions of k_{\max} . Data points with the same k_{\max} measured in a low and high resolution simulation can differ in the plots because, for these data, k_{\min} differs by a factor of 4 and, therefore, the number of modes for the bispectrum estimation differs by a factor of $4^3 = 64$, which affects particularly the cumulative signal to noise $\|\hat{B}\|$. In contrast the quantity \bar{F}_{NL}^B is normalised with respect to the linearly evolved local shape over the given range of scales and therefore takes differences in the number of modes into account. Thus the agreement of \bar{F}_{NL}^B (and $\mathcal{C}_{\beta,\alpha}$) for the different simulations seen in Fig. 4.8 shows that the N -body simulations have converged and the bispectrum estimator gives consistent results.

It should be noted that the shown cumulative signal to noise takes only Gaussian noise into account, c.f. (4.35). This simplification is not expected to be valid in the strongly nonlinear regime so that the correct signal to noise may differ significantly. However to assess this issue in a robust manner we need to run more realisations of the N -body simulations or use larger boxes with more particles.

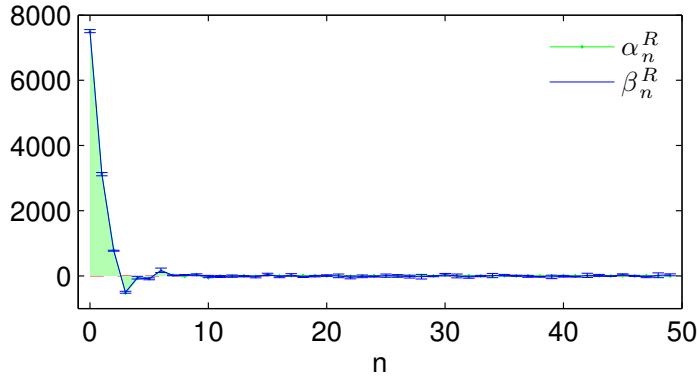


Figure 4.7: Shot noise validation test for the separable bispectrum estimator: Measured bispectrum coefficients β_n^R (4.48) (blue, averaged over 3 seeds) for 512^3 particles placed randomly in a $L = 1600 \text{Mpc}/h$ box, compared with the expansion coefficients α_n^R (4.38) (green) for the expected pure shot noise bispectrum. The shape correlation is $\mathcal{C}_{\beta,\alpha} = 0.9998$. For better visibility the region under the α_n^R curve is colored green.

4.5.4 Error bars

To obtain error bars for quantities based on the primordial contribution to the bispectrum, we use seed by seed subtracted coefficients $\langle \beta_m^R - (\beta_m^R)_{\text{Gauss}} \rangle_{\text{sim}}$ and calculate their sample standard deviations, which are then used for standard error propagation assuming uncorrelated β_m^R . The error bars therefore measure how much the primordial contribution to the matter bispectrum scatters for different seeds of the initial Gaussian field Φ_G in the simulations. Compared to calculating the difference between the average bispectra, $\langle \beta_m^R \rangle_{\text{sim}} - \langle (\beta_m^R)_{\text{Gauss}} \rangle_{\text{sim}}$, the seed by seed subtraction reduces the error bars, because the late time bispectra for Gaussian and non-Gaussian initial conditions are very correlated due to the presence of the large gravitational bispectrum in both cases.

In real observations we do not know the realisation of our universe with perfectly Gaussian initial conditions and therefore error bars will be larger if we measure $\beta_m^R - \langle (\beta_m^R)_{\text{Gauss}} \rangle_{\text{sim}}$. We do not discuss the interesting issue of observability of primordial non-Gaussianity here, because instead of dark matter bispectra this requires halo bispectra, which we leave for future work (but see [118] for the local shape).

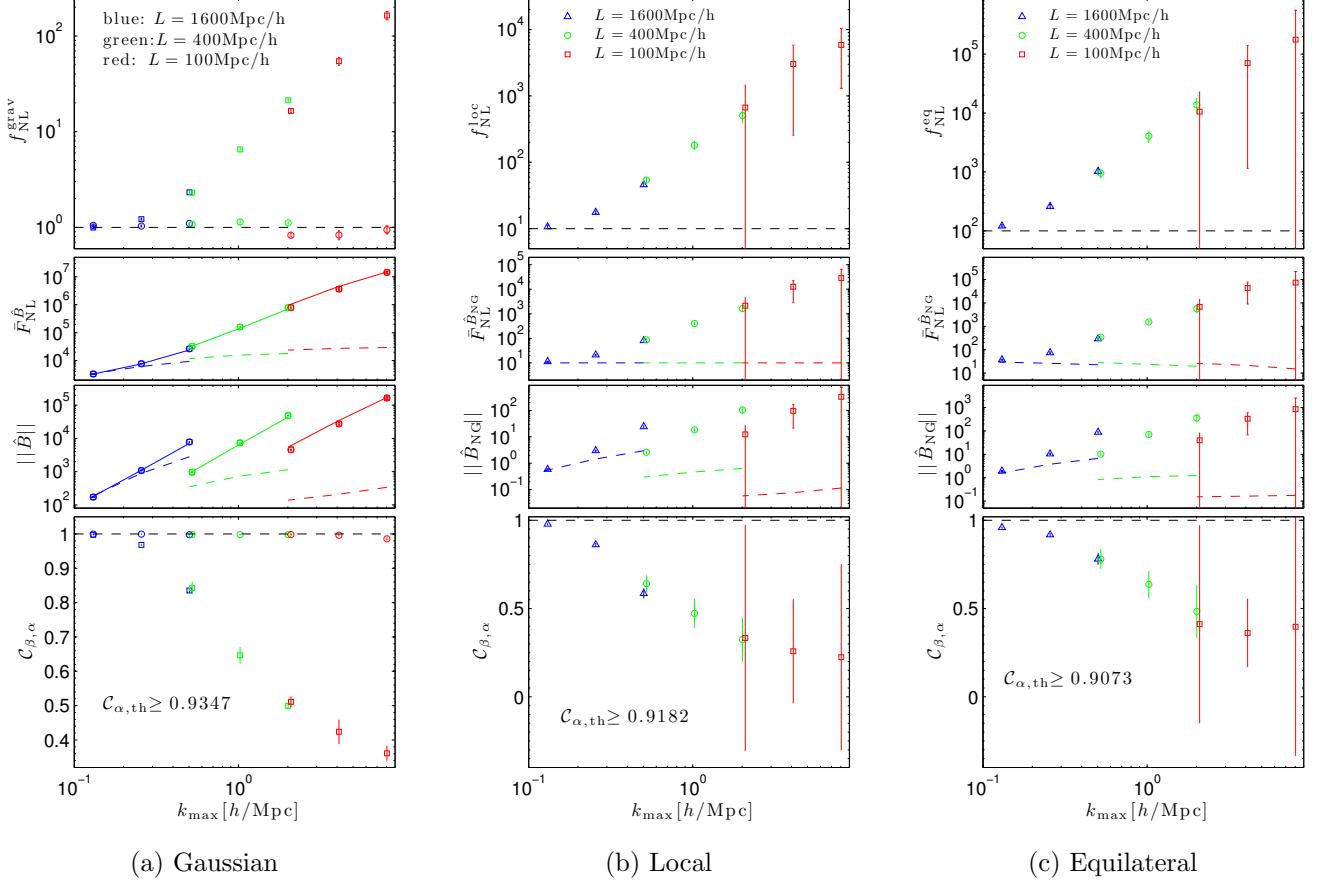


Figure 4.8: Convergence tests: Dependence of measured bispectra on k_{\max} at $z = 0$ for simulations with 512^3 particles and box sizes $L = \{1600, 400, 100\}\text{Mpc}/h$ ($\text{G512}, \text{G}_{\{400,100\}}^{512}, \text{Loc10}, \text{Loc10}_{\{400,100\}}, \text{Eq100}$ and $\text{Eq100}_{\{400,100\}}$). (a) *Top panel*: Projection of measured bispectrum on the tree level prediction (4.4) (squares) and on the fitting formula from [171] (circles). *Middle panels*: Cumulative signal to noise of measured bispectrum (squares/circles), tree level prediction (4.4) (dashed lines) and fitting formula from [171] (solid lines). *Bottom panel*: Shape correlation of measured bispectrum with tree level prediction (squares) and fitting formula from [171] (circles). (b-c) Measured non-Gaussian bispectra \hat{B}_{NG} (4.52) in simulations with local and equilateral initial conditions compared to the linearly evolved primordial bispectrum (4.22) (dashed lines). Note comments on the cumulative signal to noise on small scales in the main text.

4.5.5 Discussion in light of previous analyses

To our knowledge, we analysed for the first time the entire dark matter bispectrum in $N^3 = 512^3$ -grid N -body simulations in the strongly non-linear regime and provided a quantitative comparison of simulations and theory using full 3D shape correlations without restricting the comparison to 1D slices through the bispectrum or relying on binning in k -space, which unnecessarily reduces the number of grid points, i.e. the resolution of the bispectrum.¹ Such simplifications were required in previous analyses because of the high computational cost associated with bispectrum estimation. Overcoming the simplification of 1D slices is desirable because the fraction of bispectrum data points probed by a (manually chosen) 1D slice is only about one millionth of the amount of information in the full bispectrum (assuming no correlations); the fraction is roughly $\sim N/N^3 = 1/N^2 = \mathcal{O}(10^{-6})$ for $N = 512$ (corresponding to a line through the 3D tetrapyd plots). While comparing theory and simulations on such a small subset of the data points can indicate whether the theory is crudely wrong, it does in general not allow for quantitative conclusions about the full bispectrum to be made because it does not use all the statistical information available. Additionally, in light of observations and to improve upon the information already contained in the power spectrum, it is desirable to use and model the whole bispectrum data instead of restricting to small subsets of triangle configurations.

The main advantage of the modal bispectrum estimator is that it overcomes these restrictions by accessing all triangle configurations and orientations in an efficient way, being $\mathcal{O}(N^3/n_{\max}) = \mathcal{O}(10^7)$ times faster than a full brute force bispectrum estimation (for $N = 10^3$ and $n_{\max} = 10^2$). This speed-up allows for efficient comparisons of theoretical models, simulations and observations of the bispectrum without throwing away any bispectrum data. Additional advantages are the compression of the shape information to $\mathcal{O}(100)$ β_n^R coefficients and the applicability of the same method to estimate the trispectrum efficiently.

¹In fact, binning in k -space corresponds to the modal expansion estimator if the basis functions are chosen to be top-hat functions. While we have shown that the monomial basis used in this chapter represents gravitational and many primordial bispectra accurately in the sense that it achieves high shape correlations and therefore can be regarded as an almost loss-less compression of the full bispectrum information, this has not yet been tested for top-hat bases.

A potential disadvantage of the modal estimator method used here is that it requires a choice of a truncated set of basis functions, which reduces the dimensionality of the space of allowed bispectrum shapes (e.g. the monomial basis functions are not very sensitive to highly oscillatory bispectrum shapes; see e.g. [192]). However, it is straightforward to change the set of basis functions by including any type of functions expected to describe the problem at hand best (e.g. in [20] both Fourier and polynomial basis functions are used for the CMB bispectrum). The amount of information lost for a particular choice of basis functions can be quantified by fast-to-evaluate shape correlations. Although we obtained high shape correlations between theory and simulations for gravitational and primordial bispectra with the monomial basis used in this chapter, it would be interesting to extend or modify this basis (e.g. to improve the expansion in the squeezed limit which contributes only little to the 3D shape correlation but is crucial for halo bias).

4.6 Gravitational bispectrum results

4.6.1 Gravitational collapse and bispectrum evolution

The cosmic web simulated by N -body simulations has a complex filamentary structure, which is illustrated in Fig. 4.9a. The bispectrum of this dark matter distribution is shown in Fig. 4.9b. In the following we will discuss how the bispectrum can be used as a quantitative tool to characterise the structures formed by gravitational collapse and to study modifications due to primordial non-Gaussianity. We will demonstrate the unbiasedness of our bispectrum estimator by comparison with perturbation theory for large scales at early times.

We can develop a qualitative and visual understanding of the bispectrum by comparing it to the form of large-scale structures. Fig. 4.10 shows snapshots of the dark matter distribution at redshifts $z = 4, 2$ and 0 on the left and the corresponding measured bispectrum signal on the right. It is apparent that the shape of the bispectrum characterises the three-dimensional structures that have formed. The diffuse blob- and pancake-like structures at early times, $z = 4$, correspond to a flattened bispectrum, which peaks at the edges of the tetrapod and is predicted by leading order perturbation theory. At later times we find

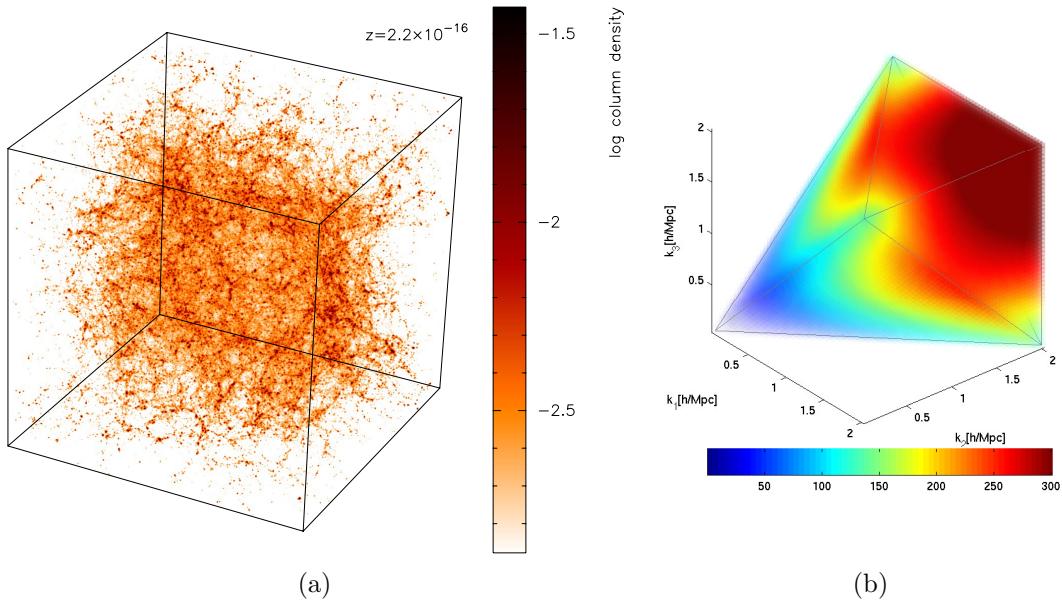


Figure 4.9: (a) Dark matter distribution in one of the G_{400}^{512} simulations of 512^3 particles in a box with $L = 400\text{Mpc}/h$ at redshift $z = 0$. (b) Measured (signal to noise weighted) bispectrum in the range $0.016h/\text{Mpc} \leq k \leq 2h/\text{Mpc}$, averaged over the simulation on the left and two additional seeds. The opaqueness of the plotted points reflects the absolute value of the bispectrum.

that filaments and clusters induce a bispectrum with a relatively enhanced signal for equilateral triangle configurations, as expected from Sections 4.2 and 4.3. As illustrated in Fig. 4.11, the bispectrum signal shows a similar transition for varying overall scale $k_1 + k_2 + k_3$ at fixed time $z = 0$, which indicates self-similar behaviour.

4.6.1.1 Comparison with leading order PT

On large scales and at early times we expect the gravitational bispectrum to agree with leading order perturbation theory (4.4). To test this, Fig. 4.12a shows the β_n^R coefficients (4.48) measured in the Gaussian N -body simulations G512 with $L = 1600\text{Mpc}/h$ in comparison with the expansion coefficients α_n^R (4.38) of the tree level gravity bispectrum (4.4). The corresponding theoretical and estimated bispectra (4.51) are plotted over the full tetrapyd in Fig. 4.3 for $z = 30$ and Fig. 4.13 for $z = 2$ and $z = 0$. At early times, $z = 30$, the bispectrum agrees with tree level perturbation theory (4.4), which predicts that the signal to noise weighted bispectrum is large for folded and elongated configurations but is suppressed in the squeezed limit. However at $z = 0$ for $k \gtrsim 0.25h/\text{Mpc}$ the N -body bispectrum has an enhanced amplitude for elongated configurations and an additional equilateral contribution not captured by tree level perturbation theory, which breaks down on these scales as expected.

To analyse the measured bispectrum more quantitatively and to illustrate its time and scale dependence we show in Figures 4.14 and 4.15 the measured amplitude $f_{\text{NL}}^{\text{grav}}$ of the gravitational bispectrum, its cumulative signal to noise $\|\hat{B}\|$ (4.34) and the cumulative signal to noise normalised to the local shape \hat{F}_{NL} (4.36), as well as the shape correlation \mathcal{C} (4.33) of the measured bispectra with the theoretical expectation. The meaning of these quantities is illustrated in Fig. 4.15c. We plot them as functions of redshift z and use different colors for different k ranges used for the bispectrum measurements. Note that $f_{\text{NL}}^{\text{grav}}$ is the amplitude of the gravitational bispectrum (4.4), i.e. tree level perturbation theory predicts $f_{\text{NL}}^{\text{grav}} = 1$ in this convention.

Fig. 4.14a and Fig. 4.14b show that for $0.0039h/\text{Mpc} \leq k \leq 0.5h/\text{Mpc}$ at $z = 0$ the tree level bispectrum under-predicts the total integrated bispectrum \bar{F}_{NL} by a factor of 2.8 and its shape correlation with the measured bispectrum drops to

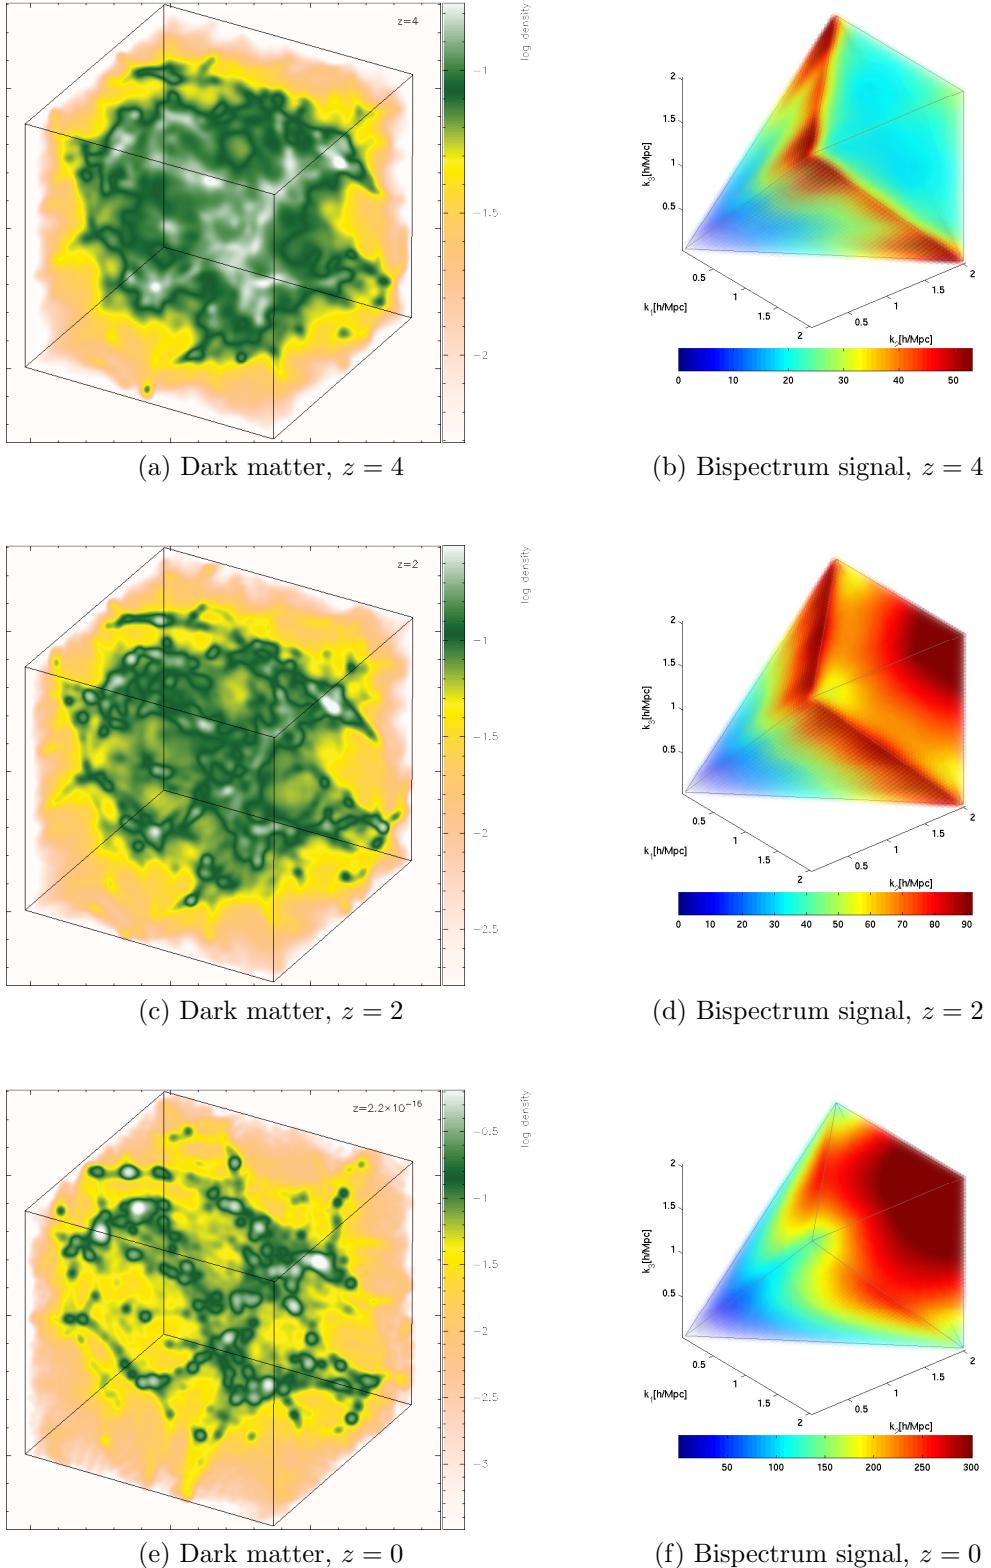


Figure 4.10: *Left:* Dark matter distribution in a $(40\text{Mpc}/h)^3$ subbox of one of the G_{400}^{512} simulations at redshifts $z = 4, 2$ and 0 , from top to bottom. *Right:* Measured (signal to noise weighted) bispectrum in the range $0.016h/\text{Mpc} \leq k \leq 2h/\text{Mpc}$, averaged over the simulation on the left and two additional seeds.

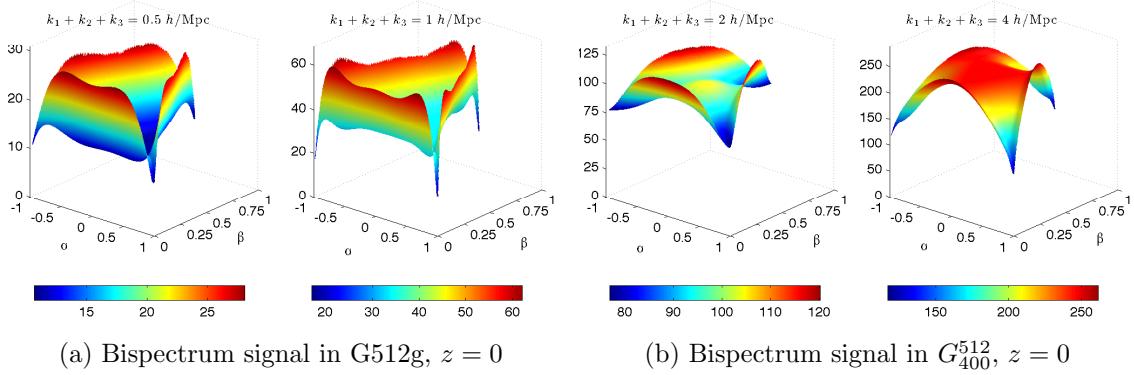


Figure 4.11: Bispectrum slices at fixed redshift $z = 0$ for different overall wavenumbers $k_{\text{sum}} = k_1 + k_2 + k_3$. For $k_{\text{sum}} \leq 1 \text{ h/Mpc}$ [panel (a)] the bispectrum signal dominates for flattened configurations, while an enhanced equilateral contribution clearly emerges for $k_{\text{sum}} = 2 \text{ h/Mpc}$ [left plot in panel (b)]. For $k_{\text{sum}} = 4 \text{ h/Mpc}$ [right plot in panel (b)] the signal is large for all triangle configurations away from the squeezed limit, that is, it is nearly constant. The slices show signal to noise weighted bispectra measured in Gaussian simulations (a) G_{512}^g and (b) G_{400}^{512} .

0.84, implying that the projection of the measured bispectrum on the tree level bispectrum is $f_{\text{NL}}^{\text{grav}} = 2.3$. At larger scales the tree level bispectrum describes the measured bispectra much more accurately until late times, as expected. Similar plots for Gaussian N -body simulations with 768^3 and 1024^3 particles are shown in Fig. 4.15a and Fig. 4.15b, demonstrating again the expected break down of tree level perturbation theory on small scales at late times.

4.6.2 Fitting formulae for Gaussian simulations

4.6.2.1 Separable polynomial expansion

The separable mode expansion allowed us to compress the measured N -body dark matter bispectra to $n_{\text{max}} = 50$ numbers β_n^R at each redshift. The first ten of these coefficients are listed in Table 4.3 and can be used as a polynomial fitting formula of the matter bispectrum by evaluating (4.51). The orthogonal contributions B_n^R appearing in (4.51) were defined in (4.46) and contain the nonlinear power spectrum P_δ as well as the orthogonal polynomials R_n , which can be obtained from

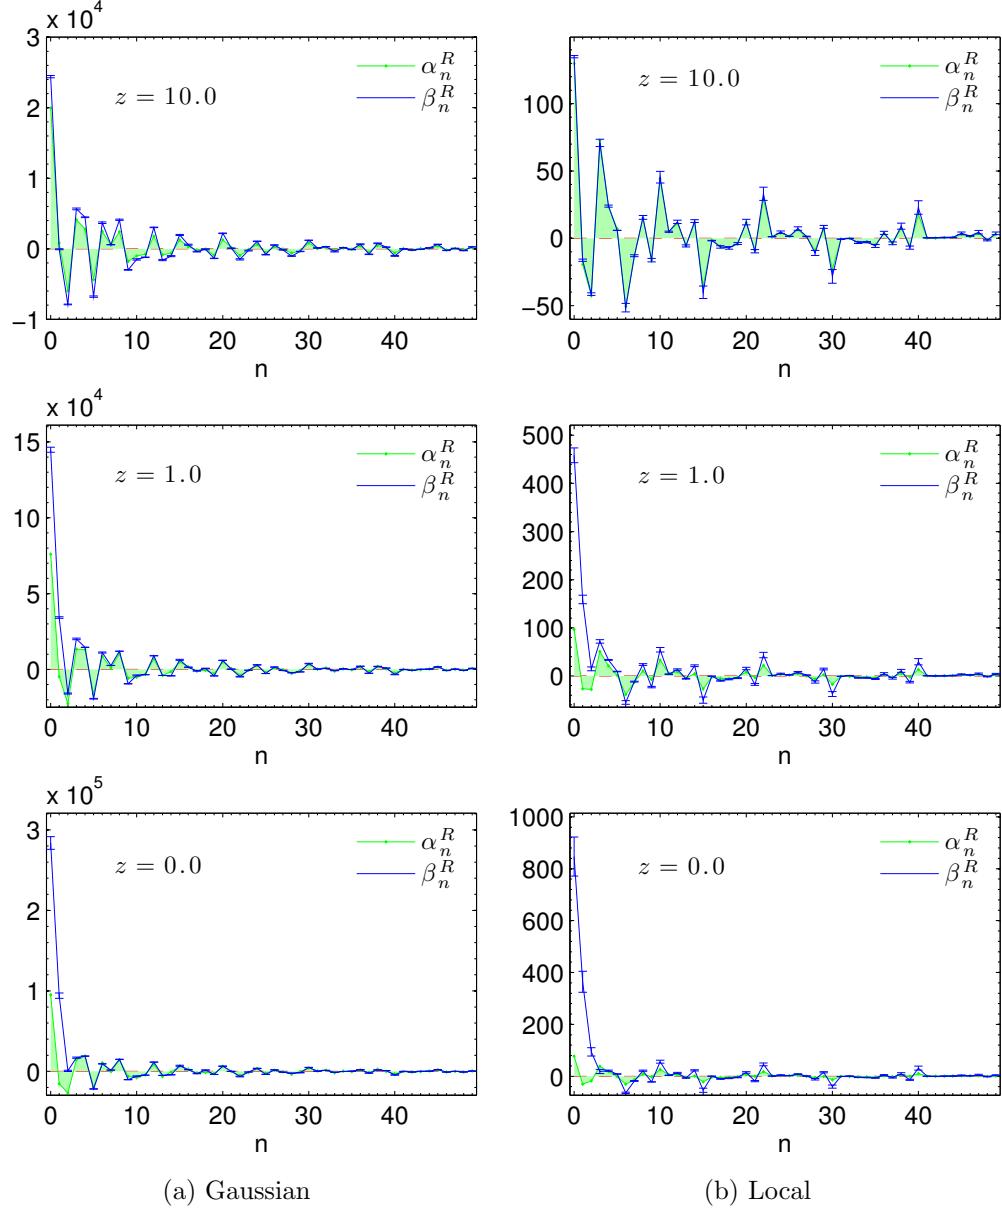


Figure 4.12: (a) β_n^R coefficients measured with (4.48) for the Gaussian simulations G512 compared to expansion coefficients α_n^R of the tree level gravitational bispectrum (4.4). (b) $\beta_n^R - (\beta_n^R)_{\text{Gauss}}$ measured in Loc10 simulations with $f_{\text{NL}}^{\text{loc}} = 10$ compared to α_n^R coefficients of the linearly evolved local shape (see (4.22) and (2.38)). β 's were calculated using all modes $\delta_{\mathbf{k}}$ with $0.0039h/\text{Mpc} \leq k \leq 0.5h/\text{Mpc}$. For better visibility the region under the α_n^R curves is colored green.

4.6 Gravitational bispectrum results

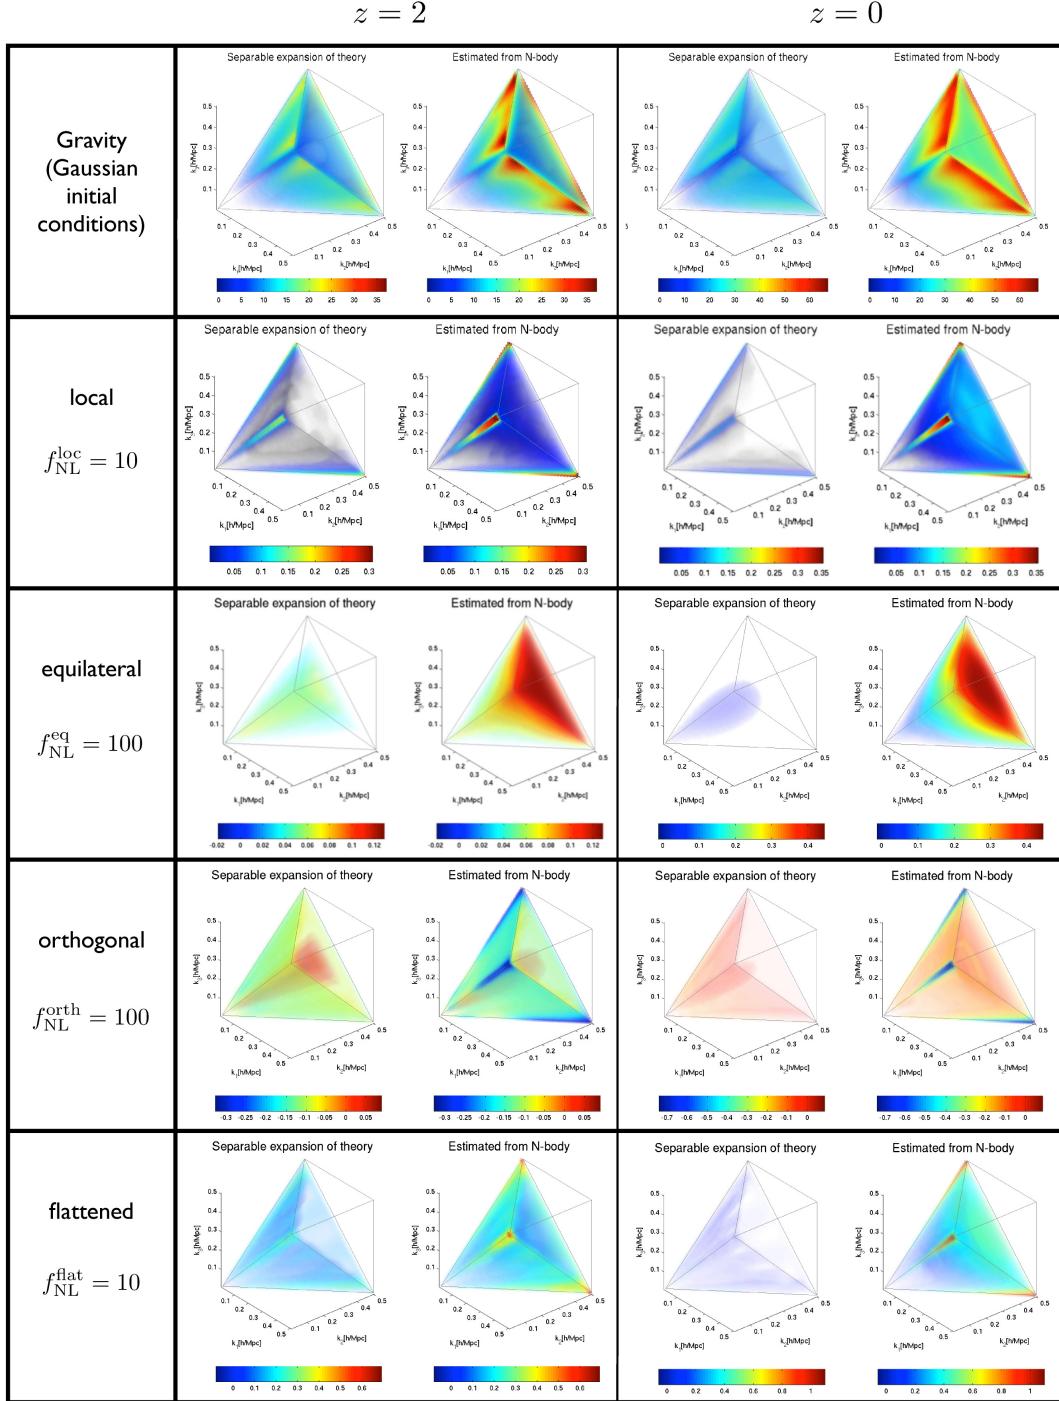


Figure 4.13: Comparison of measured N -body bispectra (4.51) (on the right in each black cell) with separable expansion of tree level theory (4.38) (on the left in each black cell) for simulations G512b, Loc10, Eq100, Orth100 and Flat10 at redshifts $z = 2$ and $z = 0$. The plot axes are $k_1, k_2, k_3 \leq 0.5 h/\text{Mpc}$ and we plot signal to noise weighted bispectra $\sqrt{k_1 k_2 k_3 / (P_\delta(k_1) P_\delta(k_2) P_\delta(k_3))} B_\delta$ (see Fig. 4.3 for comparisons at $z = 30$ and further descriptions of the plots).

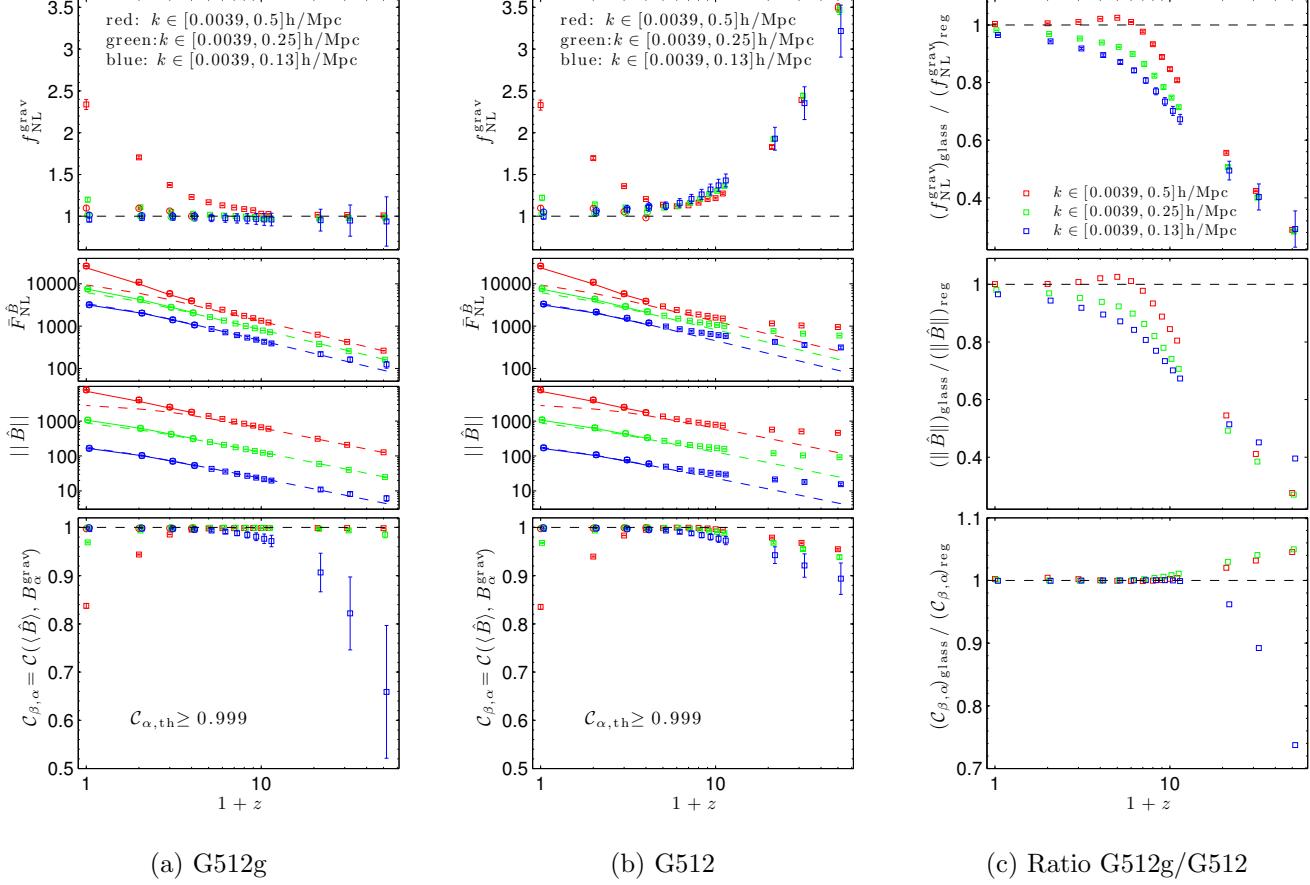


Figure 4.14: (a)-(c): Measured bispectra in *Gaussian N-body* simulations (see Table 4.2) as a function of redshift. *Upper panels*: Amplitude f_{NL}^{grav} of the tree level gravitational bispectrum (4.4) (squares), normalised to unity if the tree level prediction is correct, and amplitude of the fitting formula from [171] (circles). *Middle panels*: Total integrated bispectrum (4.56) of measured bispectrum (symbols), tree level prediction (dashed lines) and fitting formula by [171] (solid lines). *Lower panels*: Shape correlation (4.54) of reconstructed bispectrum with separable expansions of tree level theory (squares) and fitting formula from [171] (circles). Colors indicate different k ranges for the bispectrum estimation. The plotted quantities are visualised in Fig. 4.15c.

4.6 Gravitational bispectrum results

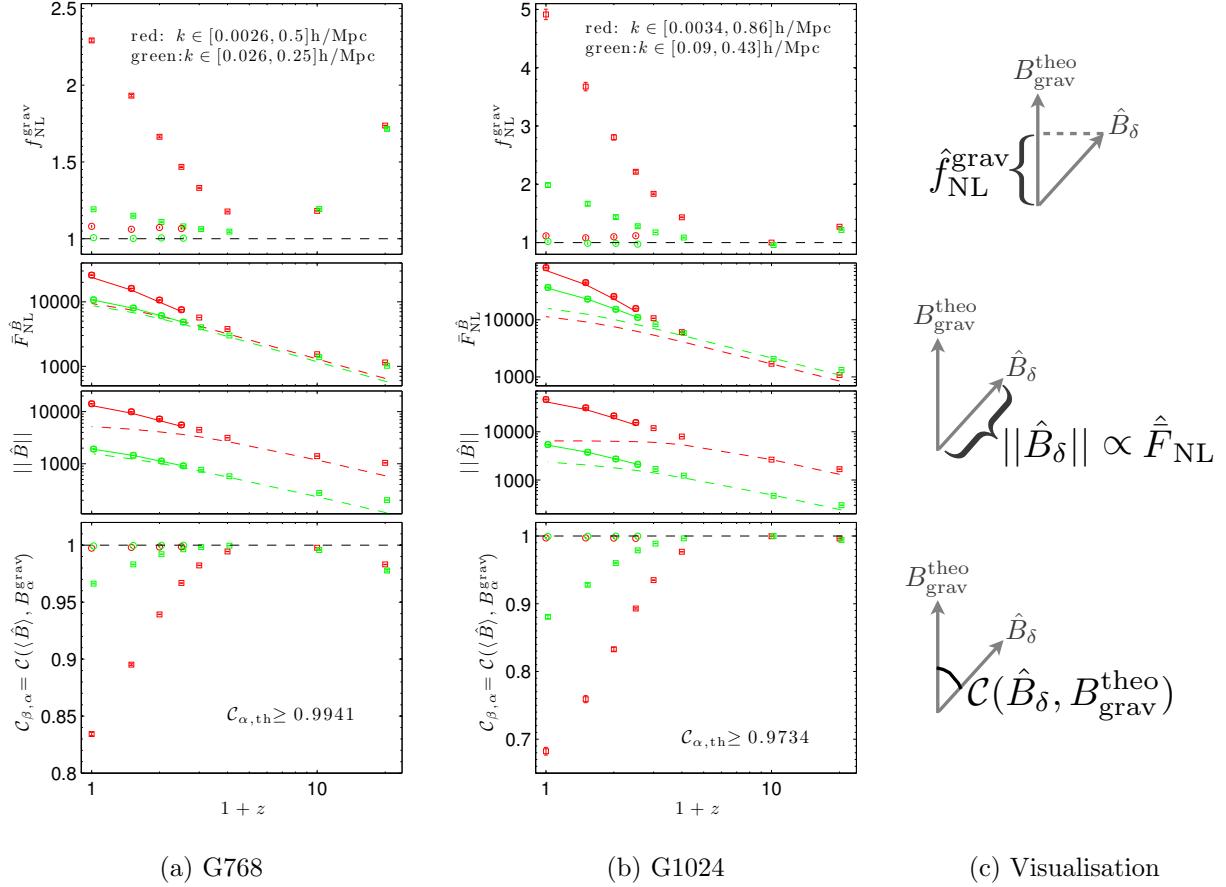


Figure 4.15: (a)-(b): Like Fig. 4.14 but for 768^3 and 1024^3 grids (see Table 4.2). (c): Visualisation of the quantities plotted in (a)-(b) and in Fig. 4.14.

Q_n , defined in (4.43), by the basis change¹ $R_n = \sum_m \lambda_{nm} Q_m$. The polynomials Q_n were sorted with the slice ordering defined in Eq. (58) of [28].

n	$\frac{(\beta_n^R)_{G512}}{1000}$	$\frac{(\beta_n^R)_{G1875}}{10000}$
0	283.8	167
1	94.16	63.06
2	0.6725	9.427
3	17.04	0.2799
4	18.89	5.774
5	-21.82	-7.911
6	9.395	1.329
7	1.624	-0.4791
8	15.04	5.825
9	-10.25	-3.548

Table 4.3: First 10 bispectrum expansion coefficients β_n^R from Gaussian N -body simulations specified in Table 4.2 for $z = 0$. The shape correlation between the full measured bispectrum (4.51) for 50 modes and the bispectrum corresponding to the shown first 10 modes is 99.7% and 99.9% for the middle and right column, respectively. We used $k_{\max} = 0.5h/\text{Mpc}$ for the middle column and $k_{\max} = 0.86h/\text{Mpc}$ for the G_{1875}^{1024} simulation in the column on the right.

4.6.2.2 Time-shift model fit

We find a remarkably accurate fit to the matter bispectrum for Gaussian initial conditions by combining the modified tree level gravity shape $B_{\delta,\text{NL}}^{\text{grav}}$, defined in (4.6), with the ‘constant’ model $B_{\delta,\text{const}}^{\text{grav}}$, defined in (4.12), which approximates the 1-halo shape, that is,

$$B_{\delta}^{\text{fit}}(k_1, k_2, k_3) \equiv B_{\delta,\text{NL}}^{\text{grav}} + B_{\delta,\text{const}}^{\text{grav}}, \quad (4.62)$$

with fitting parameters c_1 and n_h .

The combined shape (4.62) is dominated by the perturbative gravity bispectrum at early times and by the constant or (approximate) halo model prediction at late times. We will demonstrate that this combination can achieve a good fit

¹The transformation matrix λ can be obtained from the author of this thesis upon request.

4.6 Gravitational bispectrum results

Simulation	$L[\frac{\text{Mpc}}{h}]$	c_1	n_h	$\min_{z \leq 20} \mathcal{C}_{\beta,\alpha}$	$\mathcal{C}_{\beta,\alpha}(z=0)$
G512g	1600	4.1×10^6	7	99.8%	99.8%
G_{400}^{512}	400	1.0×10^7	8	99.8%	99.8%

Table 4.4: Fitting parameters c_1 and n_h for the fit (4.62) of the matter bispectrum for Gaussian initial conditions (simulations G512g and G_{400}^{512}). The two columns on the right show the minimum shape correlation with the measured (excess) bispectrum in N -body simulations, which was measured at redshifts $z = 49, 30, 20, 10, 9, 8, \dots, 0$, and the shape correlation at $z = 0$.

of the matter bispectrum at all redshifts $z \leq 20$, while both the perturbative and the halo model prediction individually break down at intermediate redshifts, when nonlinearities are important but not all dark matter particles can be treated as residing in halos.

The fitting parameters c_1 and n_h in (4.12) are obtained as follows. At each measured redshift the arbitrary weight $w(z)$ in

$$B_\delta^{\text{opt}} = B_{\delta,\text{NL}}^{\text{grav}} + w(z)(k_1 + k_2 + k_3)^\nu \quad (4.63)$$

is analytically determined such that the correlation with the measured bispectrum $\mathcal{C}(\hat{B}_\delta, B_\delta^{\text{opt}})$ is maximal (see green lines in Fig. 4.16). Then c_1 and n_h are chosen such that $c_1 \bar{D}^{n_h}(z)$ approximates the optimal weight $w(z)$ over all redshifts (see black dashed lines in Fig. 4.16). Table 4.4 lists the fitting parameters obtained by this procedure.

The shape correlation of the fitting formula (4.62) with the measured bispectrum is 99.8% or better at redshifts $z \leq 20$ for both $k_{\text{max}} = 0.5h/\text{Mpc}$ and $k_{\text{max}} = 2h/\text{Mpc}$ as shown in the upper panels of Fig. 4.17. The model (4.62), therefore, contains all the meaningful bispectrum shape information. All we require is the time dependence or growth rate of the bispectrum amplitude. As a first step, the fitting formula (4.63) can be normalised to the measured bispectrum size by multiplying it with the normalisation factor

$$N_{\text{fit}} \equiv \frac{\|\hat{B}\|}{\|B_\delta^{\text{fit}}\|}, \quad (4.64)$$

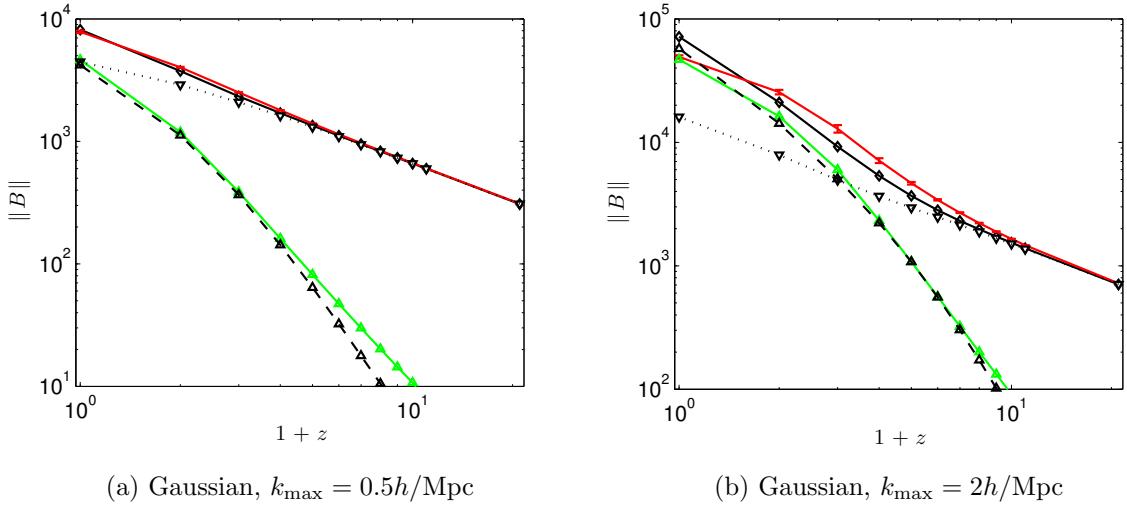


Figure 4.16: Motivation for using the growth function \bar{D} in the simple fitting formula (4.62). The arbitrary weight $w(z)$ in $B_{\delta}^{\text{opt}} = B_{\delta,\text{NL}}^{\text{grav}} + w(z)(k_1 + k_2 + k_3)^{\nu}$ is determined analytically such that $\mathcal{C}(\hat{B}_{\delta}, B_{\delta}^{\text{opt}})$ is maximal (for $\nu = -1.7$). We plot $\|B_{\delta,\text{NL}}^{\text{grav}}\|$ (black dotted), $\|w(z)(k_1 + k_2 + k_3)^{\nu}\|$ (green) and $B_{\delta,\text{const}}^{\text{grav}}$ (black dashed) as defined in (4.12) with fitting parameters given in Table 4.4, illustrating that $w(z) = c_1 \bar{D}^{n_h}(z)$ is a good approximation. The continuous black and red curves show $\|B_{\delta}^{\text{fit}}\|$ from (4.62) and the estimated bispectrum size $\|\hat{B}_{\delta}\|$, respectively. The overall normalisation can be adjusted with N_{fit} as explained in the main text [see Eq. (4.64)].

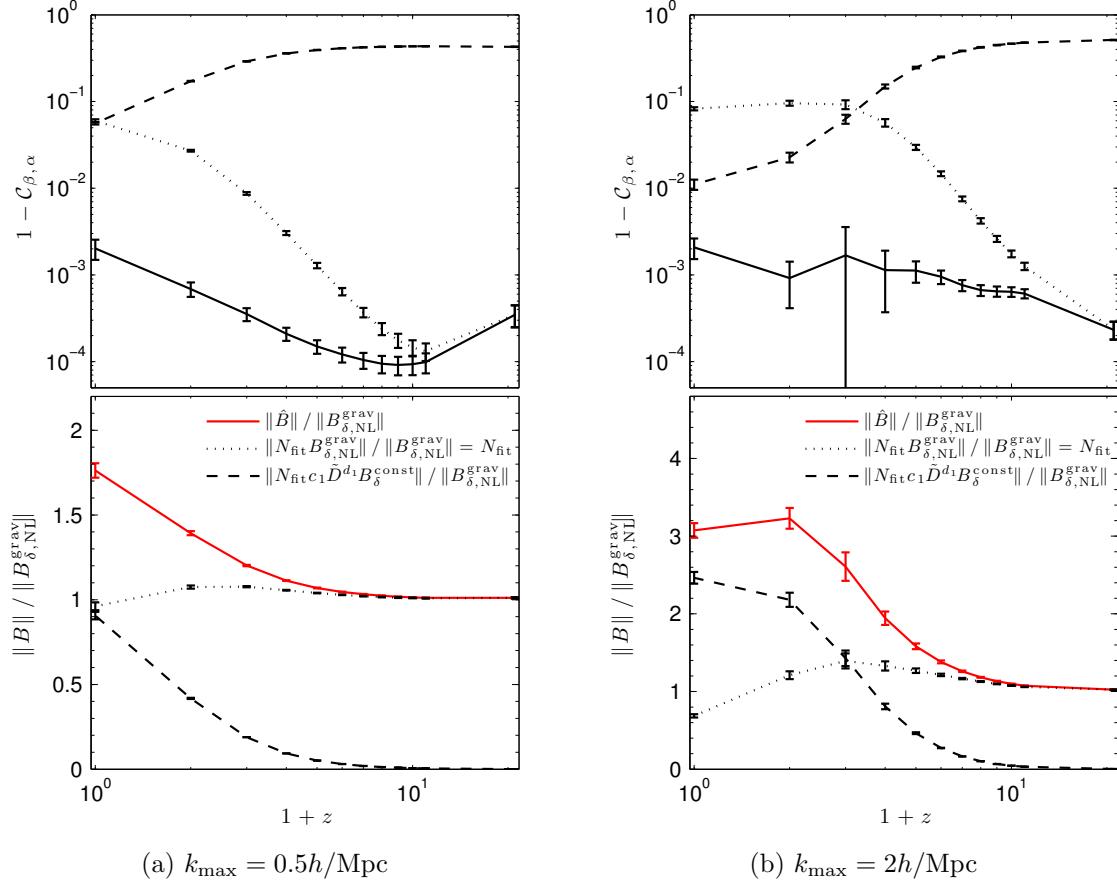


Figure 4.17: *Upper panels:* Shape correlation of estimated bispectrum in (a) G512g and (b) G₄₀₀⁵¹² simulations with the full fitting formula (4.62) (solid black), the shape $B_{\delta, \text{NL}}^{\text{grav}}$ as defined in (4.6) (dotted black) and the shape $B_{\delta, \text{const}}^{\text{grav}}$ as defined in (4.12) (dashed black). All correlations were calculated with 120 modes Q_n for (a) $k \in [0.004, 0.5]h/\text{Mpc}$ and (b) $k \in [0.016, 2]h/\text{Mpc}$. *Lower panels:* Integrated size of the measured bispectrum (red) and the contributions $N_{\text{fit}} B_{\delta, \text{NL}}^{\text{grav}}$ (dotted black) and $N_{\text{fit}} B_{\delta, \text{const}}^{\text{grav}}$ (dashed black). All these bispectrum sizes are divided by the size of $B_{\delta, \text{NL}}^{\text{grav}}$ for convenience. The normalisation factor N_{fit} is defined in (4.64) and equals the dotted black curve.

which is shown by the dotted line in the lower panels of Fig. 4.17. While it varies with redshift between 0.7 and 1.4 for $k_{\max} = 2h/\text{Mpc}$, it deviates by at most 8% from unity for $k_{\max} = 0.5h/\text{Mpc}$. The lower panels also show the measured integrated bispectrum size $\|\hat{B}\|$ and the two individual contributions to (4.62) when the normalisation factor N_{fit} is included. These quantities are divided by $\|B_{\delta,\text{NL}}^{\text{grav}}\|$ for convenience. At high redshifts the total bispectrum size is essentially given by the contribution from $B_{\delta,\text{NL}}^{\text{grav}}$, which equals the tree level prediction for the gravitational bispectrum in this regime. The contribution from $B_{\delta}^{\text{const}}$ dominates at $z \leq 2$ for $k_{\max} = 2h/\text{Mpc}$ when filamentary and spherical nonlinear structures are apparent. A similar transition can be seen at later times on larger scales in Fig. 4.17a, indicating self-similar behaviour.

It is worth noting that the high integrated correlation between the simple fit (4.62) and measurements does not imply that all triangle configurations agree perfectly and sub-percent level differences between shape correlations can in principle contain important information, e.g. about the observationally relevant squeezed limit which only makes a small contribution to the total tetrapyd integral over the signal-to-noise weighted dark matter bispectrum. However, if we observed the dark matter bispectrum directly, these shapes would be hard to distinguish because the shape correlation contains the signal-to-noise weighting. Modified shape correlation weights and additional basis functions have been used for better quantitative comparison of the squeezed limit of dark matter bispectra, but this is left for a future publication.

4.6.2.3 Alternative phenomenological fit

An alternative fitting formula with 9 fitting parameters and calibrated on larger scales was given in [171] and summarised in this chapter in Sec. 4.2.4.4. In the range of validity given by [171], $0.03h/\text{Mpc} \leq k \leq 0.4h/\text{Mpc}$ at $0 \leq z \leq 1.5$, we find good agreement with our N -body measurements, see green circles in Fig. 4.15a and Fig. 4.15b. Without having to run N -body simulations with higher resolutions, we extended the bispectrum measurement to $k_{\max} = 0.86h/\text{Mpc}$ with the fast separable estimator, see red symbols in Fig. 4.15a and Fig. 4.15b. In this extended regime the fitting formula of [171] still has a shape correlation of 99.5% or more

with the measured bispectrum at $z \leq 1.5$, but underestimates the cumulative signal to noise by up to 11%. Our measured β_n^R coefficients or our simple fit (4.62) can be used as alternative fitting formulae for the gravitational bispectrum valid to smaller scales, $k_{\max} \leq 2h/\text{Mpc}$, and for all redshifts $z \leq 20$.

4.6.2.4 Halo model

The halo model prediction for the Gaussian dark matter bispectrum yields a remarkably high shape correlation of more than 99.7% with the measured bispectrum at $z = 0$ for $k_{\max} = 2h/\text{Mpc}$. While the halo model bispectrum has been tested on some one-dimensional slices in [182] and on larger scales in [183], the result presented here demonstrates that at $z = 0$ the halo model shape is a good representation of the shape measured in N -body simulations over the full tetrapyd allowed by the triangle condition. At higher redshifts, when less dark matter resides in halos, the halo model prediction becomes worse and alternative phenomenological approaches like the ones discussed above yield higher shape correlations with the measured bispectrum. A more thorough examination of the halo model is left for future work.

4.7 Primordial non-Gaussian bispectrum results

4.7.1 Primordial bispectrum measurements

We have set up and evolved non-Gaussian initial conditions for local models with $f_{\text{NL}}^{\text{loc}} = -10, 10, 20, 50$ using (2.37), for equilateral models with $f_{\text{NL}}^{\text{eq}} = \pm 100$, for orthogonal models with $f_{\text{NL}}^{\text{orth}} = \pm 100$ and for the flattened model with $f_{\text{NL}}^{\text{flat}} = 10$ using separable expansions as in [4, 50]. Detailed parameters for the N -body simulations are given in Table 4.2.

4.7.1.1 Local shape

In Fig. 4.12b we show comparisons of the measured $\beta_n^R - (\beta_n^R)_{\text{Gauss}}$ coefficients with the expansion coefficients α_n^R of the linearly evolved primordial bispectrum (4.22) for $f_{\text{NL}}^{\text{loc}} = 10$, finding good agreement at early times, but deviations at late times which are scale-dependent. The most obvious feature as we consider increasingly nonlinear wavenumbers is the emergence of a large constant or equilateral signal. Indeed Fig. 4.13 shows this signal on small scales with the late time bispectrum also exhibiting an enhanced signal in the squeezed limit.

In Fig. 4.18a we plot the projection $f_{\text{NL}}^{\text{loc}}$ of the measured non-Gaussian bispectrum \hat{B}_{NG} on the linearly evolved primordial bispectrum (see (4.22) and (2.38)), their shape correlation $\mathcal{C}_{\beta,\alpha}$, cumulative signal to noise $\|\hat{B}_{\text{NG}}\|$ and total integrated bispectrum normalised to the local shape, $\bar{F}_{\text{NL}}^{\hat{B}_{\text{NG}}}$. We also show the shape correlation $\mathcal{C}_{\text{NG,G}}$ between \hat{B}_{NG} and the bispectrum measured in Gaussian simulations. For $z = 0$ and $k \in [0.0039, 0.5]h/\text{Mpc}$ the linearly evolved primordial bispectrum under-predicts the measured cumulative signal to noise $\|\hat{B}_{\text{NG}}\|$ by a factor of 8 and has a shape correlation with the measured bispectrum of about 0.6, implying that the projection $f_{\text{NL}}^{\text{loc}}$ is about 4.8 times the input value of $f_{\text{NL}}^{\text{loc}} = 10$. The green and blue symbols in Fig. 4.18a show that tree level perturbation theory improves if we consider larger scales, as expected.

4.7.1.2 Other shapes

We also ran simulations with equilateral, orthogonal and flattened non-Gaussian initial conditions using the separable expansion method [4, 50]. Bispectrum mea-

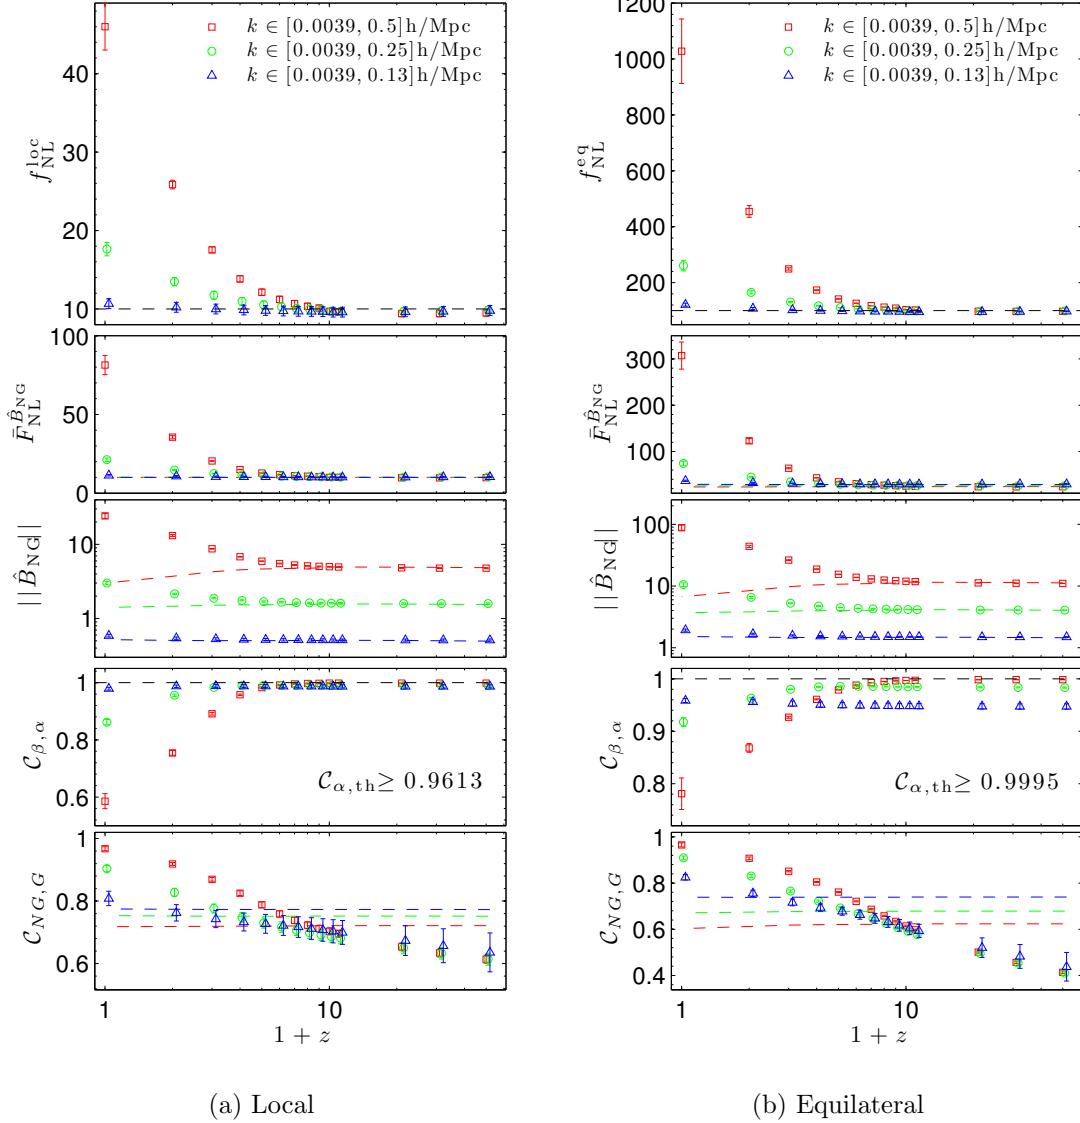


Figure 4.18: Dark matter bispectrum quantities as in Figures 4.14 and 4.15 but for the non-Gaussian simulations Loc10, Eq100 and Orth100 with local, equilateral and orthogonal initial conditions, respectively. The non-Gaussian bispectrum is computed with (4.52) and then compared to the tree level prediction (4.22) (dashed lines). The lowest panel shows the shape correlation of the measured non-Gaussian bispectrum with the measured bispectrum for Gaussian initial conditions.

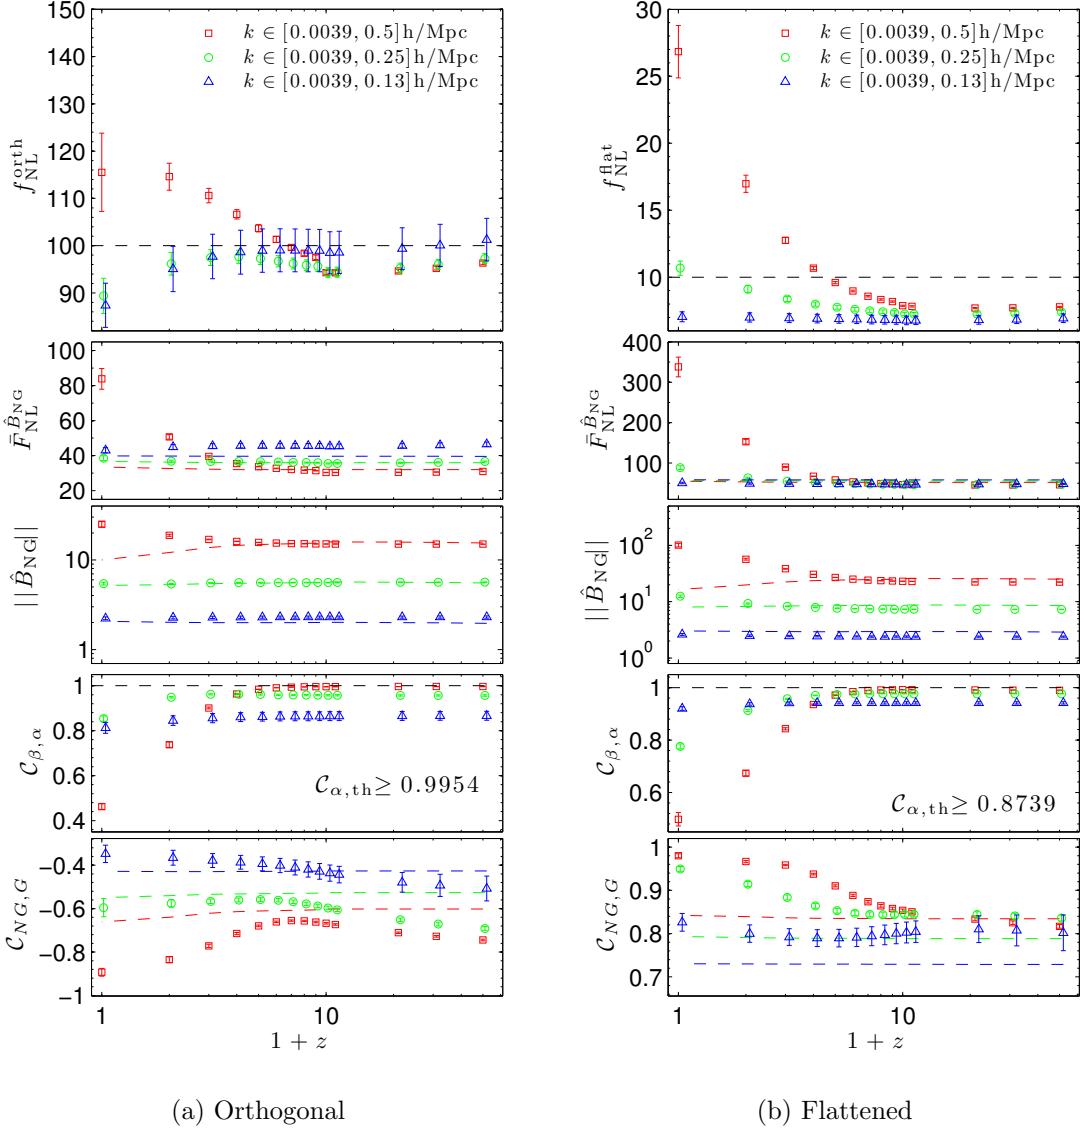


Figure 4.19: Like Fig. 4.18 but for simulations Orth100 and Flat10 with orthogonal and flattened non-Gaussian initial conditions.

surements are shown in Figures 4.3, 4.13, 4.18 and 4.19. As in the local case we find agreement with tree level perturbation theory at all times on large scales. However on small scales and at late times, Fig. 4.13 shows that the measured bispectrum for equilateral initial conditions has additional equilateral and flattened contributions compared to the linearly evolved input bispectrum (4.22). For initial conditions with the $f_{\text{NL}}^{\text{orth}} = 100$ orthogonal shape, which is positive for equilateral configurations and negative for flattened and squeezed configurations, we find that the late time small scale bispectrum turns slightly negative for equilateral configurations and is more negative in flattened and squeezed configurations than the linearly evolved input bispectrum. The bispectrum from flattened initial conditions shows additional contributions for squeezed, equilateral and flattened configurations compared to the tree level prediction. For orthogonal and flattened initial conditions the late time bispectra have a large signal in the squeezed limit, leading to possible confusion with local shape initial conditions (with negative $f_{\text{NL}}^{\text{loc}}$ for positive input $f_{\text{NL}}^{\text{orth}}$ and positive $f_{\text{NL}}^{\text{loc}}$ for positive input $f_{\text{NL}}^{\text{flat}}$). However scale dependent halo bias is sensitive to the scaling in the squeezed limit and can therefore help to disentangle such shapes (see e.g. [39, 167]).

It is worth noting that the shape correlation $\mathcal{C}_{\beta,\alpha}$ between measured and input bispectra drops as we reduce k_{\max} already at the initial time of the simulations, e.g. $\mathcal{C}_{\beta,\alpha} \approx 0.95$ for $k_{\max} = 0.13h/\text{Mpc}$ in the equilateral case. This indicates that the initial conditions generated with a separable mode expansion - which are generated at $k_{\max} = 0.5h/\text{Mpc}$ - are not a perfect fit to the shape when restricted to large scales. This is to be expected since the expansion is optimised to fit the total cumulative signal to noise up to $k_{\max} = 0.5h/\text{Mpc}$. There are however several possibilities of further improving the initial conditions on large scales. One could add basis functions which are localised on large scales or one could introduce optimised separable weighting functions in the mode expansions. This would be particularly important for examining halo bias which is very sensitive to the squeezed limit and therefore to large-scale modes. However for studying the global behaviour of the dark matter bispectrum we find it acceptable that the shape of the initial conditions is not perfect on very large scales and leave further improvements of the initial conditions for future work.

It is worth noting that the flattened shape initial conditions do not perfectly

represent the flattened template because the shape correlation between expansion and theory is only $\mathcal{C}_{\alpha,\text{th}} \geq 0.87$; see Figures 4.3 and 4.19b. However, the bispectrum estimation is very accurate which is reflected by the fact that $\mathcal{C}_{\beta,\alpha} \gtrsim 0.95$ on all scales at high redshifts (see fourth panel in Fig. 4.19b). The accuracy of the initial conditions can be improved by employing more monomial basis functions or by including basis functions which are better suited to cope with the divergence of the flattened template used here.

4.7.1.3 Loop corrections for primordial non-Gaussianity

As is apparent from Figures 4.13, 4.18 and 4.19, the tree level prediction for the non-Gaussian matter bispectrum breaks down for small scales and low redshift. The inclusion of loop corrections described in [168] is expected to improve the fit at such scales and redshifts. First results on the correlation between the observations and theoretical prediction $\mathcal{C}_{\beta,\alpha}$ indicate that 1-loop corrections can significantly improve the shape correlation in the nonlinear regime to $\mathcal{O}(0.8 - 0.9)$, while shape correlations of more than 0.9 at $k = 0.5h/\text{Mpc}$ and $z = 0$ may require higher order loop corrections. However we defer a detailed analysis of loop corrections and more phenomenological halo model approaches to future work.

4.7.1.4 Linearity in input f_{NL}

We test if $\mathcal{O}(f_{\text{NL}}^2)$ corrections are important by comparing simulations with $f_{\text{NL}}^{\text{loc}} = -10, 20$ and 50 to the $f_{\text{NL}}^{\text{loc}} = 10$ simulation in Fig. 4.20a (with box size $L = 1600\text{Mpc}/h$). The curves in all three panels would be exactly unity in case of perfect linearity in the input f_{NL} . Deviations from linearity are at most 1%.

In case of equilateral and orthogonal initial conditions we compare simulations with input $f_{\text{NL}} = \pm 100$. The equilateral shape gives similar results to the local shape, see Fig. 4.20b. For the orthogonal shape the shape correlation between theory and measurements deviates from linearity in the input f_{NL} by about 3% at $z = 0$ for $k_{\text{max}} = 0.5h/\text{Mpc}$ and the total integrated bispectrum differs by less than 5%. These two effects cancel each other approximately so that the measured projection $f_{\text{NL}}^{\text{orth}}$ only deviates by less than 1.5% from linearity.

We conclude that at the $\mathcal{O}(5\%)$ level the bispectra measured in our large-scale

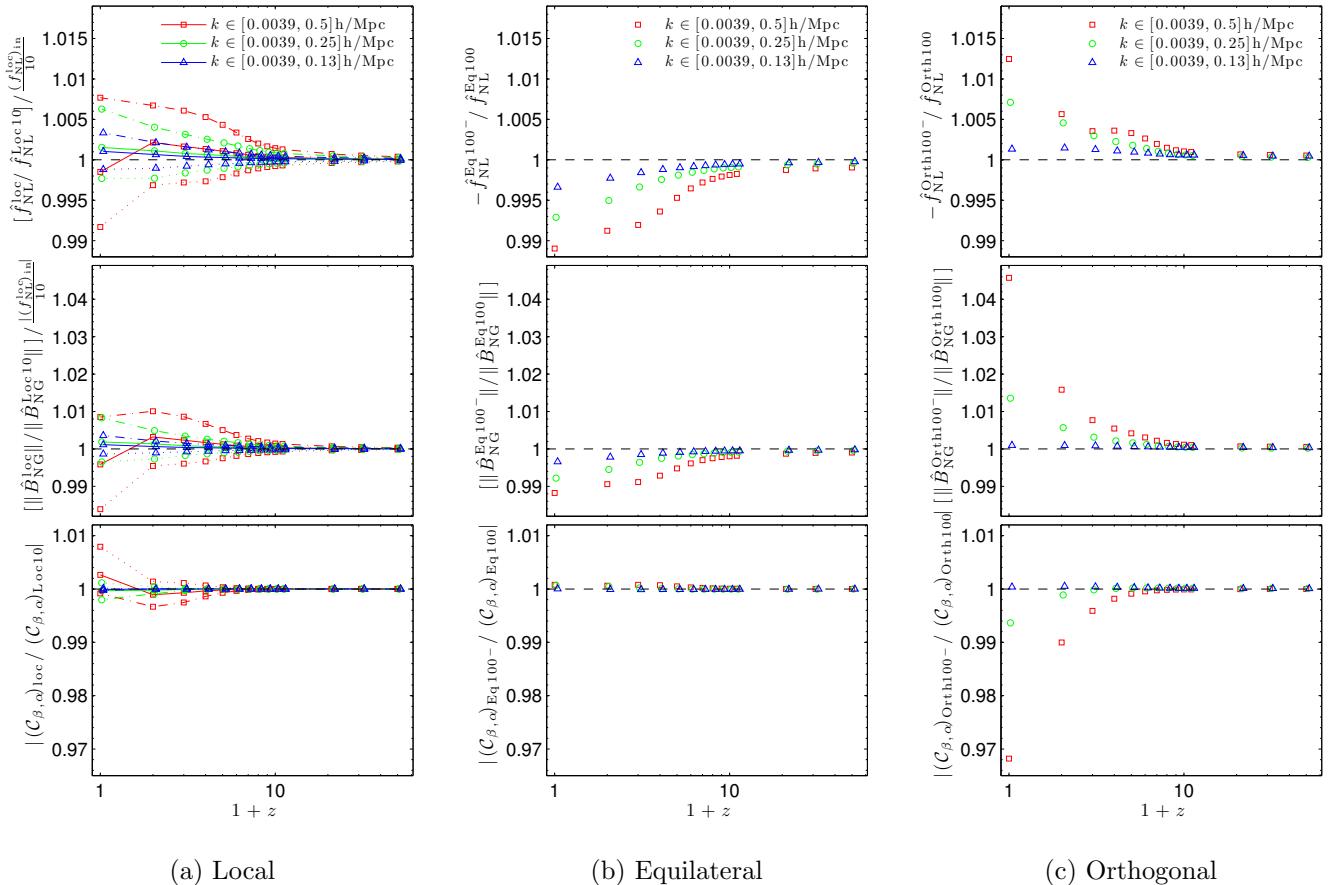


Figure 4.20: Test if the bispectrum contributions (4.52) from non-Gaussian initial conditions are linear in the input f_{NL} . (a) Ratios of Loc20 (solid line), Loc50 (dash-dot line) and Loc10 $^{-}$ (dotted line) to Loc10 simulation. The label 'loc' stands for Loc20, Loc50 and Loc10 $^{-}$. We scale the curves by $(f_{\text{NL}}^{\text{loc}})^{\text{in}}/10$ so that in case of perfect linearity in f_{NL} all curves were unity. (b) Ratio of Eq100 $^{-}$ to Eq100 simulation. (c) Ratio of Orth $^{-}$ to Orth100 simulation.

simulations can be scaled to other values $f_{\text{NL}}^{\text{new}}$ fulfilling $|f_{\text{NL}}^{\text{loc}}| \leq 50$, $|f_{\text{NL}}^{\text{Eq}}| \leq 100$ and $|f_{\text{NL}}^{\text{Orth}}| \leq 100$ using the linear scaling

$$B_{\text{NG}}(f_{\text{NL}}^{\text{new}}) = \frac{f_{\text{NL}}^{\text{new}}}{f_{\text{NL}}} \hat{B}_{\text{NG}}(f_{\text{NL}}). \quad (4.65)$$

4.7.2 Fitting formulae for non-Gaussian simulations

4.7.2.1 Separable polynomial fits

Matter bispectra for non-Gaussian initial conditions of the local, equilateral, orthogonal and flattened type are described by the β_n^R coefficients in Table 4.5. These polynomial fitting formulae can serve as a starting point for future work that relies on non-Gaussian dark matter bispectra in the nonlinear regime.

n	Loc10	Eq100	Orth100	Flat10
0	846.8	3065	-793.5	3544
1	364.3	1390	-240.4	1484
2	94.22	453	28.49	362.2
3	18.58	-183.5	-238.7	128.1
4	21.7	-47.71	-173.6	163.5
5	8.64	35.7	85.29	-43.72
6	-63.53	-6.603	154.1	-232.5
7	-17.96	-2.818	53.24	-77.86
8	21.4	-57.18	-186	171.9
9	-21.61	39.26	108.1	-109.2

Table 4.5: First 10 bispectrum expansion coefficients β_n^R from non-Gaussian N -body simulations specified in Table 4.2 for $z = 0$ and $k_{\text{max}} = 0.5h/\text{Mpc}$. The shape correlation between the full measured bispectrum (4.51) for 50 modes and the bispectrum corresponding to the shown first 10 modes is 99.2%, 99.95%, 95.8% and 99.7% for the columns from the left to the right.

4.7.2.2 Time-shift model fits

Simple fitting formulae for the primordial contribution to the matter bispectrum can be successfully obtained from the halo time-shift model described in Sec. 4.3.2. Before discussing this fit in detail, we perform a simple consistency check of the

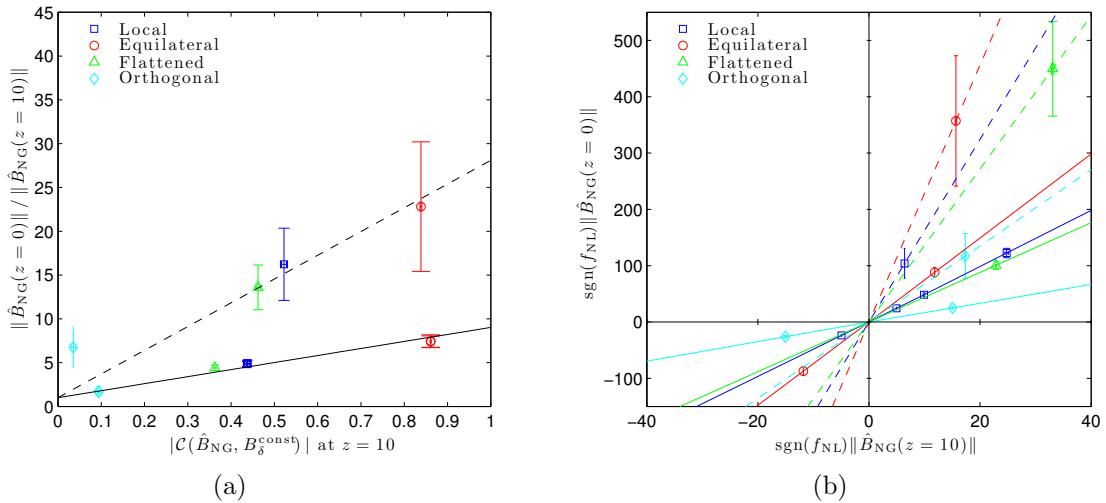


Figure 4.21: (a) Relative growth of bispectrum size between $z = 10$ and $z = 0$ as a function of the absolute value of the shape correlation between the measured non-Gaussian bispectrum \hat{B}_{NG} and the shape $B_{\delta}^{\text{const}}$ defined in (4.12) at redshift $z = 10$, testing relation (4.66). The plot contains all simulations shown in (b) but data points which only differ in the input f_{NL} are almost indistinguishable. Regression lines through $(0, 1)$ are shown for $k_{\text{max}} = 0.5h/\text{Mpc}$ (continuous) and $k_{\text{max}} = 2h/\text{Mpc}$ (dashed). (b) Bispectrum size at redshift $z = 0$ as a function of the bispectrum size at high redshift, $z = 10$, for different primordial shapes (see legend). Continuous lines correspond to $k_{\text{max}} = 0.5h/\text{Mpc}$ while dashed lines correspond to $k_{\text{max}} = 2h/\text{Mpc}$ simulations. Different points on one line show results for different input f_{NL} .

basic idea of the model. The relatively fast growth of the constant bispectrum implies that it constitutes the dominant contribution to the non-Gaussian bispectrum B_{NG} at sufficiently small scales and late time z_{late} . The amplitude of this constant bispectrum is related to the projection of the non-Gaussian bispectrum B_{NG} on the constant shape at the time z_{early} when halos start to form (and thereafter). Hence we expect that

$$\|B_{\text{NG}}(z_{\text{late}})\| \propto \mathcal{C}(B_{\text{NG}}(z_{\text{early}}), B_{\delta, \text{const}}^{\text{grav}}(z_{\text{early}})) \|B_{\text{NG}}(z_{\text{early}})\|. \quad (4.66)$$

This simple expectation is approximately seen in Fig. 4.21a for the local, equilateral and flattened shapes, confirming the basic idea of the time-shift model. The fact that the orthogonal shape deviates somewhat could be related to the change of sign of the orthogonal shape for different triangle configurations. Note that relation (4.66) and Fig. 4.21a are interesting results on their own because they show that the relative growth of the non-Gaussian bispectrum can be predicted from its correlation with the constant shape at early times. Fig. 4.21b illustrates the absolute values of the measured bispectrum sizes which were used to produce Fig. 4.21a.

In detail, the simple fitting formulae for the non-Gaussian bispectra are obtained by combining the partially loop-corrected tree level expression (4.23) with the constant shape (4.26) as

$$B_{\text{NG}}^{\text{fit}}(k_1, k_2, k_3) \equiv f_{\text{NL}} [B_{\delta, \text{NL}}^{\text{prim}} + B_{\delta, \text{const}}^{\text{prim}}]. \quad (4.67)$$

The fitting parameters c_2 and n_h^{prim} in (4.26) are listed in Table 4.6. Similar to the Gaussian case they were obtained by analytically determining the optimal weight $w(z)$ for the ‘constant’ $(k_1 + k_2 + k_3)^\nu$ contribution and approximating this with $c_2 \bar{D}(z)^{n_h^{\text{prim}}}$ (see green and black dashed lines in Fig. 4.22). As expected from Sec. 4.3.2 we find $n_h^{\text{prim}} = n_h - 1$ for local, equilateral and flattened initial conditions.

Table 4.6 also shows the shape correlation with the measured excess bispectrum (4.52). These shape correlations are remarkably good given the simplicity of (4.67), especially for local, equilateral and flattened initial conditions. The impact of

4.7 Primordial non-Gaussian bispectrum results

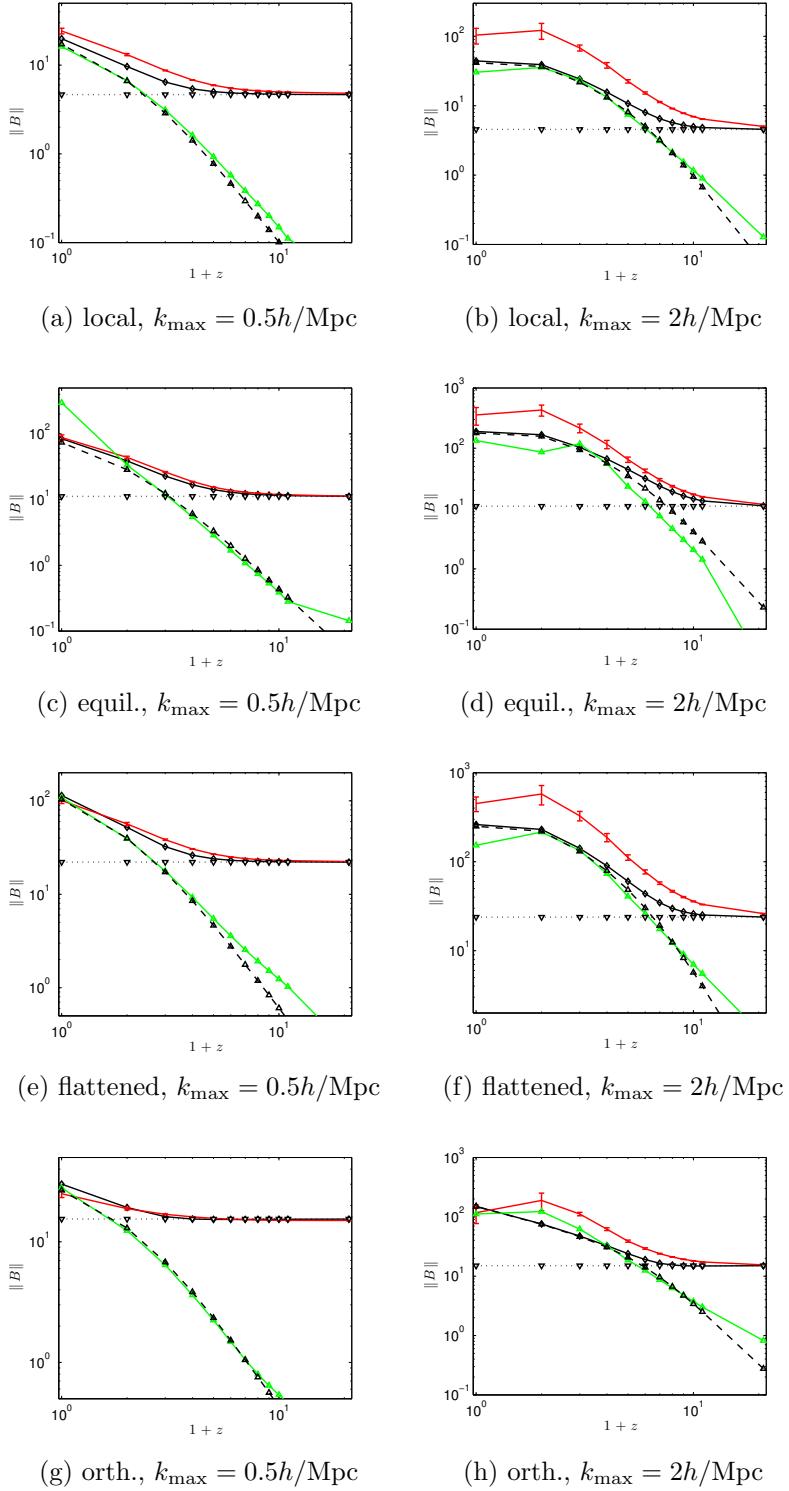


Figure 4.22: Contributions to simple non-Gaussian fits (4.67). The arbitrary weight $w(z)$ in $B_{\text{NG}}^{\text{opt}} = B_{\delta,\text{NL}}^{\text{prim}} + w(z)(k_1 + k_2 + k_3)^\nu$ is determined analytically such that $\mathcal{C}(\hat{B}_{\text{NG}}, B_{\text{NG}}^{\text{opt}})$ is maximal (for $\nu = -1.7$). We plot $\|B_{\delta,\text{NL}}^{\text{prim}}\|$ (black dotted), $\|w(z)(k_1 + k_2 + k_3)^\nu\|$ (green) and $B_{\delta,\text{const}}^{\text{prim}}$ (black dashed) as defined in (4.26) with fitting parameters given in Table 4.6. The continuous black and red curves show $\|B_{\text{NG}}^{\text{fit}}\|$ from (4.67) and the estimated primordial bispectrum size $\|\hat{B}_{\text{NG}}\|$, respectively.

Simulation	$L[\frac{\text{Mpc}}{\text{h}}]$	c_2	n_h^{prim}	$\min_{z \leq 20} \mathcal{C}_{\beta,\alpha}$	$\mathcal{C}_{\beta,\alpha(z=0)}$
Loc10	1600	2×10^3	6	99.7%	99.8%
Eq100	1600	8.6×10^2	6	97.9%	99.4%
Flat10	1600	1.2×10^4	6	98.8%	98.9%
Orth100	1600	-3.1×10^2	5.5	91.0%	91.0%
Loc10 ₄₀₀ ⁵¹²	400	$2 \times 10^3 dD/da$	7	98.2%	99.0%
Eq100 ₄₀₀ ⁵¹²	400	$8.6 \times 10^2 dD/da$	7	94.4%	97.9%
Flat10 ₄₀₀ ⁵¹²	400	$1.2 \times 10^4 dD/da$	7	97.7%	99.1%
Orth100 ₄₀₀ ⁵¹²	400	-2.6×10^2	6.5	97.3%	98.9%

Table 4.6: Fitting parameters c_2 and n_h^{prim} for the fit (4.67) of the primordial bispectrum (4.52). The two columns on the right show the minimum shape correlation with the measured (excess) bispectrum in N -body simulations, which was measured at redshifts $z = 49, 30, 20, 10, 9, 8, \dots, 0$, and the shape correlation at $z = 0$. For the equilateral case the minimum shape correlation can be improved to 99.4% if the term $4.6 \times 10^{-5} f_{\text{NL}} \bar{D}(z)^{0.5} [2P_\delta(k_1)P_\delta(k_2)F_2^{(s)}(\mathbf{k}_1, \mathbf{k}_2) + 2 \text{ perms}]$ is added to (4.67).

the orthogonal shape on the matter bispectrum seems to be somewhat harder to model, which is reflected in the fact that our simple model performs worse for this shape. This is expected because the orthogonal shape contains almost no constant component, the dominant evolution of which is the basis for the simple time-shift model. While the integrated bispectrum size \bar{F}_{NL} of the fitting formulae is consistent with the measured bispectrum size at high redshift and at $z = 0$, it underestimates the measured size at the 10 – 20% level at intermediate redshifts for $k_{\text{max}} = 0.5h/\text{Mpc}$. This could be improved by adding more shapes or by using a redshift-dependent normalisation factor in (4.67) similar to the N_{fit} factor in the Gaussian case. On smaller scales, $k_{\text{max}} = 2h/\text{Mpc}$, the non-Gaussian fitting formulae provide a less accurate overall fit, but we also list them in Table 4.6 for completeness and because the correlations at $z = 0$ are quite high.

It should be noted that simple fits like the ones presented here may be somewhat more important in the mildly nonlinear regime than in the strongly nonlinear regime because the halo model (extended to non-Gaussian initial conditions) should be able to describe the strongly nonlinear regime. We leave more accurate extensions of the simple fits for non-Gaussian initial conditions and more quan-

titative comparisons with loop corrections and halo model predictions for future work.

4.7.2.3 Redshift at which 1-halo contribution becomes important

We determine the redshift z_* at which the approximate 1-halo bispectrum B_δ^{const} in the simple fits of the bispectrum for non-Gaussian initial conditions becomes important by matching

$$\|\mathcal{C}(B_{\delta,\text{NL}}^{\text{prim}}(z_*), B_{\delta,\text{const}}^{\text{prim}}(z_*))B_{\delta,\text{NL}}^{\text{prim}}(z_*)\| = \|B_{\delta,\text{const}}^{\text{prim}}(z_*)\|. \quad (4.68)$$

For $k_{\max} = 0.5h/\text{Mpc}$ we find $z_* = 3.5, 3.3, 4.3$ for local, equilateral and flattened initial conditions, respectively. For $k_{\max} = 2h/\text{Mpc}$ we find instead $z_* = 8.5, 8$ and 9.6 for the same initial conditions. The orthogonal shape has almost no constant part. We therefore get somewhat different values of $z_* = 5.5$ and 11 for $k_{\max} = 0.5h/\text{Mpc}$ and $k_{\max} = 2h/\text{Mpc}$, respectively. This demonstrates that the time at which nonlinear structures contribute significantly to the perturbative prediction for the primordial bispectrum is earlier on smaller scales.

4.8 Summary and conclusions

We have presented an implementation of a bispectrum estimator for N -body simulations using a separable modal expansion of the bispectrum as described in [50]. While a brute force estimation of the full bispectrum is computationally expensive, requiring $\mathcal{O}(N^6)$ operations for N particles per dimension, we find that the gravitational and the most prominent primordial bispectra can be approximated by only $n_{\max} = \mathcal{O}(50)$ separable basis functions (for the range of scales relevant for N -body simulations). The bispectrum projection on the corresponding subspace of all possible bispectra is estimated with $\mathcal{O}(n_{\max}N^3)$ operations, which is faster than brute force estimation by a factor of $\mathcal{O}(N^3/n_{\max}) \sim \mathcal{O}(10^7)$ for typical simulations. Thus the computational cost for accurate 3D bispectrum estimation is almost negligible compared to the cost for running the N -body simulations (e.g. we can estimate the full bispectrum of a 1024^3 grid up to $k_{\max} = \frac{N}{4}\frac{2\pi}{L}$ in one hour on only 6 cores). This allows us to estimate the bispectrum as a standard simulation

diagnostic whenever the power spectrum is measured. Expressing the bispectrum using its n_{\max} separable components yields a radical compression of data which simplifies further analysis like comparisons between theory and simulations. The separable bispectrum estimator therefore provides a very useful additional statistic characterising the formation of structures in N -body simulations, with high sensitivity to different shapes of primordial non-Gaussianity corresponding to different models of inflation.

We have performed many N -body simulations with Gaussian initial conditions as well as non-Gaussian initial conditions of the local, equilateral, orthogonal and (non-separable) flattened shape, exploiting the separable bispectrum expansion for efficient generation of initial conditions as described for primordial fields in [4, 50]. On large scales the measured gravitational and primordial bispectra agree with leading order perturbation theory and with measurements for Gaussian initial conditions by [171] in the mildly nonlinear regime, demonstrating the unbiasedness of the estimator and the initial conditions. In the nonlinear regime, the gravitational bispectrum becomes dominated by a large ‘constant’ signal receiving elongated and equilateral contributions not captured by tree level perturbation theory. However, it remains suppressed in the squeezed limit, where primordial bispectrum signals can peak (see Fig. 4.13).

Our measured N -body bispectra for Gaussian and non-Gaussian simulations can be expressed by 50 components β_n^R , which we provide in Tables 4.3 and 4.5 for the key models. They can be used as fitting formulae for the gravitational and primordial bispectrum in the nonlinear regime.

Less accurate but simpler fitting formulae are obtained by modeling the bispectrum as a combination of partially loop-corrected perturbative bispectra and a simple ‘constant’ bispectrum, which is constant on slices $\sum k_i = \text{const.}$ and which is obtained as an approximation to the 1-halo bispectrum. While the former contribution dominates on large scales and early times, the latter constant contribution dominates in the nonlinear regime. Interpreting the effect of primordial non-Gaussianity on the constant bispectrum contribution as a time-shift with respect to Gaussian simulations allows us to model the time dependence of the constant bispectrum contribution for non-Gaussian initial conditions. For Gaussian initial conditions, the simple fits achieve a shape correlation of at least 99.8%

with the measured gravitational bispectrum for $z \leq 20$ and $k_{\max} = \{0.5, 2\}h/\text{Mpc}$. For local, equilateral and flattened non-Gaussian initial conditions the primordial bispectrum is fitted with a shape correlation of at least 97.9% at $z = 0$ and at least 94.4% for $z \leq 20$ (with correlation typically $\gtrsim 98\%$ for most shapes and redshifts, see Table 4.6). The impact of the orthogonal shape seems to be somewhat harder to model, because it does not have a constant component initially, but our simple fit still achieves correlations of at least 91% for $z \leq 20$.

Throughout this work we have visualised the measured bispectra in three-dimensional tetrapyd plots [28], which show the bispectrum shape and amplitude at different length scales and generalise commonly used plots of one- or two-dimensional slices through the tetrapyd volume. For a more quantitative analysis, particularly to test analytic predictions and fitting formulae, we have made extensive use of full three-dimensional shape correlations, the cumulative signal-to-noise of the bispectra and their projection f_{NL} on theoretical shapes. These quantities have been evaluated extremely efficiently using the bispectrum components obtained from the separable estimator.

We find that regular grid initial conditions produce an initial spurious bispectrum due to the anisotropy of the regular grid, which can be avoided by using glass initial conditions. However the difference between regular grid and glass initial conditions decreases with time as gravitational perturbations grow such that both initial conditions yield similar results at late times. Effects of order f_{NL}^2 were shown to affect the bispectrum measurements by less than 5% in our large-scale simulations.

Clearly, further work is required to study the effects of general primordial non-Gaussianity on observable quantities like the bispectrum of galaxies, particularly in the nonlinear regime. However the present work represents, we believe, an important step forward in the understanding of structure formation in the presence of primordial non-Gaussianity and the search for primordial non-Gaussianity in large-scale structures.

Acknowledgements

I am very grateful to my collaborators for this chapter: Donough Regan and Paul Shellard. We are especially grateful to James Fergusson for his work on the development of the separable bispectrum estimation formalism, for sharing his experience with the application to the CMB and for many informative and illuminating discussions regarding the application to large-scale structures. We thank Anthony Challinor for comments on the manuscript and useful discussions, and Daniel Baumann, Francis Bernardeau, Hiro Funakoshi, Helge Gruetjen, Martin Haehnelt, Antony Lewis and Emiliano Sefusatti for helpful discussions. We are also very grateful to Andrey Kaliazin for his invaluable computational help. We thank Volker Springel for providing us the Gadget-3 N -body code (see [132, 190] for older versions). We also thank Edwin Sirkо for his initial conditions code [128], parts of which we use for setting up initial conditions, Sebastian Pueblas and Roman Scoccimarro for their public 2LPT code for Gaussian initial conditions [130] and Antony Lewis and Anthony Challinor for CAMB [64]. We acknowledge the use of the FFTW3 library and SPLASH [193] for visualising particle distributions. Simulations were performed on the COSMOS supercomputer (an Altix 4700) which is funded by STFC, HEFCE and SGI.

On the joint analysis of CMB temperature and lensing-reconstruction power spectra

Based on [6]

Summary

Gravitational lensing provides significant cosmological information in modern CMB analyses. It is estimated from both the 2- and 4-point functions of the lensed CMB anisotropies (the latter corresponding to the lensing reconstruction power spectrum). These two observables are often treated as independent, although this is impossible since they are derived from the same CMB map. In this chapter, we analyse the covariance between the CMB temperature and lensing reconstruction power spectra. We find two dominant contributions coming from: (i) correlations between the disconnected noise bias in the trispectrum measurement and sample variance in the temperature power spectrum; and (ii) sample variance of the lenses themselves. The former is naturally removed when the dominant $N^{(0)}$ Gaussian bias in the reconstructed deflection spectrum is dealt with via a data-dependent correction, as advocated elsewhere for other reasons. The remaining lens-cosmic-variance contribution is easily modeled but can safely be ignored for a Planck-like experiment, justifying treating the two observable spectra as independent. We also test simple likelihood approximations for the deflection power spectrum, finding that a Gaussian with a parameter-independent covariance performs well.

Overview

A general introduction to CMB lensing was given in Sec. 2.3, including a more specific introduction to this chapter in Sec. 2.3.4. This chapter is organised as follows. We review CMB lensing reconstruction on the full sky in Sec. 5.1 and we describe our simulations of lensed CMB maps and the mechanics of our reconstructions in Sec. 5.2. Section 5.3 surveys known results for the auto-correlations of the lensed CMB temperature power spectrum and the reconstruction power spectrum. In Sec. 5.4 we present results for the cross-correlation of these power spectra and assess the importance of correlations for estimating cosmological parameters. We test likelihood approximations for the lensing reconstruction (in isolation) in Sec. 5.5 and conclude in Sec. 5.6. In Appendix 5.A we provide intuitive arguments for the magnitude of the temperature-lensing power correlation. A series of further appendices provide calculational details, and motivate some of the approaches taken in the main text. Table 5.1 summarises the key quantities, and their definitions and symbols, used in this chapter.

5.1 Full-sky CMB lensing reconstruction

We described lensing reconstruction from the CMB temperature anisotropies in Sec. 2.3. For enhanced clarity, there we worked under the flat-sky approximation. However, it is safer to work on the full sky for realistic studies of Planck-like experiments which include very large scales. We therefore generalise the flat-sky treatment from Sec. 2.3 to the full sky, which leads to a more involved notation but does not change the mechanics of the quadratic lensing reconstruction (see e.g. [59, 79, 88, 89]).

The lensed CMB temperature $\tilde{T}(\hat{\mathbf{n}})$ along direction $\hat{\mathbf{n}}$ is related to the unlensed temperature $T(\hat{\mathbf{n}})$ by the deflection field $\alpha(\hat{\mathbf{n}})$ ¹

$$\tilde{T}(\hat{\mathbf{n}}) = T[\hat{\mathbf{n}} + \alpha(\hat{\mathbf{n}})]. \quad (5.1)$$

¹The notation here is rather symbolic on the spherical sky: the point $\hat{\mathbf{n}} + \alpha(\hat{\mathbf{n}})$ is understood to be obtained from $\hat{\mathbf{n}}$ by displacing through a distance $|\alpha(\hat{\mathbf{n}})|$ along the geodesic that is tangent to α at $\hat{\mathbf{n}}$ [194].

5.1 Full-sky CMB lensing reconstruction

Symbol	Description	Definition
T	Unlensed CMB temperature	(5.1)
\tilde{T}	Lensed CMB temperature	(5.1)
ϕ_{in}	Input lensing potential field in simulations	
$\hat{\phi}$	Lensing reconstruction (quadratic in \tilde{T})	(5.9)
A_L, \tilde{g}	Normalisation and weights for $\hat{\phi}$	(5.10), (5.11)
C_l^{XX}	Fiducial theoretical power spectrum of $X = T, \tilde{T}$ or ϕ (without noise/beam)	
$C_{l,\text{expt}}^{\tilde{T}\tilde{T}}$	Lensed temperature power spectrum including beam-deconvolved noise	(2.101)
\hat{C}_l^{XX}	Empirical power spectrum of a realisation of $X = T, \tilde{T}, \phi_{\text{in}}$ or $\hat{\phi}$	(5.18)
$N_l^{(0)}$	Gaussian, fully disconnected noise bias of $\hat{C}_l^{\hat{\phi}\hat{\phi}}$	(5.16)
$N_l^{(1)}, N_l^{(2)}$	Biases of $\hat{C}_l^{\hat{\phi}\hat{\phi}}$ at $\mathcal{O}(C^{\phi\phi})$ and $\mathcal{O}[(C^{\phi\phi})^2]$	[59]
$\hat{N}_l^{(0)}$	Data-dependent, empirical bias subtraction term	(5.17), (5.31)
$\text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})$	Covariance of CMB temperature and lensing reconstruction power spectra	(5.32)
$\text{cov}(\dots)_{\text{disconn.}}$	Noise contribution (from Gaussian, fully disconnected CMB 6-point function)	(5.34), (5.35)
$\text{cov}(\dots)_{[(\text{non-})\text{primary}]}^{\text{conn.4pt.} X}$	Trispectrum contributions of type $X = A, B$ [from (non-)primary coupling]	Appendix 5.D
$\text{cov}(\dots)_{\langle \tilde{T}_1 \dots \tilde{T}_6 \rangle_{c,\text{dom.}}^{(4)}}$	Matter cosmic variance contribution [from $\mathcal{O}(\phi^4)$ connected CMB 6-point function]	(5.44), (5.106)
A	Overall lensing amplitude of a fiducial lensing power spectrum	$C_l^{\phi\phi} = AC_l^{\phi\phi} _{\text{fid}}$
\hat{A}	Estimator for A based on reconstruction power spectrum, i.e. CMB trispectrum	(5.50)
\hat{A}'	Estimator for A based on CMB power spectrum	(5.51)
n	Tilt of a fiducial lensing power spectrum $C_l^{\phi\phi} _{\text{fid}}$	(5.63)

Table 5.1: Summary of the main quantities used throughout this chapter (roughly ordered after first appearance).

The deflection angle can be written in the Born approximation as the angular gradient of the lensing potential, $\alpha(\hat{\mathbf{n}}) = \nabla\phi(\hat{\mathbf{n}})$, which is given by an integral along the (unperturbed) line of sight of the gravitational potentials (see Eq. (2.86)).

Analogously to Eqs. (2.92) and (2.93), on the full sky, the effect of lensing on the CMB temperature can be expressed perturbatively by Taylor expanding Eq. (5.1) in ϕ . Expanding in spherical harmonics, the multipoles of the lensed CMB, \tilde{T}_{lm} , are related to those of the unlensed CMB T_{lm} and the lensing potential ϕ_{lm} via [82]

$$\tilde{T}_{lm} = T_{lm} + \delta T_{lm} + \delta^2 T_{lm} + \dots, \quad (5.2)$$

where changes due to lensing $\delta^n T_{lm}$ are of order $\mathcal{O}(\phi^n)$ and linear in the unlensed temperature T :

$$\delta T_{lm} = \sum_{l_1, l_2} \phi_{l_1} T_{l_2} I_{ll_1 l_2}^{mm_1 m_2}, \quad (5.3)$$

$$\delta^2 T_{lm} = \sum_{l_1, l_2, l_3} \phi_{l_1} \phi_{l_2} T_{l_3} J_{ll_1 l_2 l_3}^{mm_1 m_2 m_3}, \quad (5.4)$$

where we have introduced the notation $\underline{l} \equiv (lm)$ and I denotes an angular integral over a product of (derivatives of) spherical harmonics given by [82]

$$I_{ll_1 l_2}^{mm_1 m_2} = (-1)^m \begin{pmatrix} l & l_1 & l_2 \\ -m & m_1 & m_2 \end{pmatrix} F_{l_1 l_2}. \quad (5.5)$$

Expressions for $J_{ll_1 l_2 l_3}^{mm_1 m_2 m_3}$ can be found in [59]. The geometrical factor $F_{l_1 L l_2}$, which is symmetric in the last two indices, is given by

$$\begin{aligned} F_{l_1 L l_2} &= [L(L+1) - l_1(l_1+1) + l_2(l_2+1)] \\ &\times \sqrt{\frac{(2l_1+1)(2L+1)(2l_2+1)}{16\pi}} \begin{pmatrix} l_1 & L & l_2 \\ 0 & 0 & 0 \end{pmatrix}, \end{aligned} \quad (5.6)$$

and describes the rotationally-invariant part of the coupling between the three multipoles. Note that the factor in the second line of Eq. (5.6) is symmetric in all three indices.

It is possible to reconstruct the lensing potential $\phi(\hat{\mathbf{n}})$ from the lensed CMB

5.1 Full-sky CMB lensing reconstruction

exploiting the fact that fixed lenses introduce correlations between temperature modes [195, 196]. Following the non-perturbative calculations in [87], we have

$$\left\langle \frac{\partial}{\partial \phi_{LM}} (\tilde{T}_{l_1 m_1} \tilde{T}_{l_2 m_2}) \right\rangle \approx (-1)^M \begin{pmatrix} L & l_1 & l_2 \\ -M & m_1 & m_2 \end{pmatrix} \tilde{f}_{l_1 L l_2}, \quad (5.7)$$

where $\tilde{f}_{l_1 L l_2}$ is symmetric in l_1 and l_2 and contains the *lensed* temperature power spectrum, $C_l^{\tilde{T}\tilde{T}}$:

$$\tilde{f}_{l_1 L l_2} = C_{l_2}^{\tilde{T}\tilde{T}} F_{l_1 L l_2} + C_{l_1}^{\tilde{T}\tilde{T}} F_{l_2 L l_1}. \quad (5.8)$$

The angle brackets in Eq. (5.7) denote the expectation value over ϕ and T (and noise) and we have neglected the T - ϕ correlation which is generally a good approximation for CMB lensing since we are usually interested in small-scale modes of the temperature where the correlation is small. More generally, we will neglect the T - ϕ correlation throughout this thesis, except where explicitly stated otherwise. Equation (5.7) motivates forming a quadratic estimator¹ for the lensing potential [89],

$$\hat{\phi}_{LM} = A_L \sum_{l_1 l_2} (-1)^M \begin{pmatrix} l_1 & l_2 & L \\ m_1 & m_2 & -M \end{pmatrix} \tilde{g}_{l_1 l_2}(L) \tilde{T}_{l_1} \tilde{T}_{l_2}, \quad (5.9)$$

which is the full-sky equivalent of (2.102). For any choice of weights $\tilde{g}_{l_1 l_2}(L)$, we determine the normalisation A_L by demanding that $\langle \partial \hat{\phi}_{LM} / \partial \phi_{L'M'} \rangle = \delta_{LL'} \delta_{MM'}$ which gives

$$A_L = (2L+1) \left(\sum_{l_1 l_2} \tilde{f}_{l_1 L l_2} \tilde{g}_{l_1 l_2}(L) \right)^{-1}. \quad (5.10)$$

One can determine optimal weights by minimising the variance of the estimator to find [analogously to Eq. (2.107)]

$$\tilde{g}_{l_1 l_2}(L) = \frac{\tilde{f}_{l_1 L l_2}}{2C_{l_1, \text{expt}}^{\tilde{T}\tilde{T}} C_{l_2, \text{expt}}^{\tilde{T}\tilde{T}}}, \quad (5.11)$$

¹Beyond the class of quadratic estimators, the lensing potential can potentially be estimated more optimally by maximising the full likelihood (see e.g. [109, 197, 198]). Since current data analyses employ the quadratic estimator we only consider this case throughout this thesis.

where the numerator contains the *lensed* temperature power $C_l^{\tilde{T}\tilde{T}}$, while the denominator involves the total power spectrum for the experiment $C_{l,\text{expt}}^{\tilde{T}\tilde{T}}$ including beam-deconvolved noise [see Eq. (2.101)].

The normalisation (5.10) and weight (5.11) differ slightly from the original works of Hu and collaborators [88, 89] since $\tilde{f}_{l_1 L l_2}$ involves the lensed instead of the unlensed CMB power spectrum. The form here is particularly suited to power spectrum estimation and incorporates the non-perturbative findings of [87]. The expectation value of the estimated reconstruction power spectrum

$$\hat{C}_L^{\hat{\phi}\hat{\phi}} \equiv \frac{1}{2L+1} \sum_{M=-L}^L \hat{\phi}_{LM} \hat{\phi}_{LM}^*, \quad (5.12)$$

involves the 4-point function of the lensed CMB. The connected part of the 4-point function can be written in terms of the fully-reduced trispectrum $\mathbb{T}_{l_3 l_4}^{l_1 l_2}(L)$ as [199]

$$\begin{aligned} & \langle \tilde{T}_{l_1 m_1} \tilde{T}_{l_2 m_2} \tilde{T}_{l_3 m_3} \tilde{T}_{l_4 m_4} \rangle_c \\ &= \frac{1}{2} \sum_{LM} (-1)^M \begin{pmatrix} l_1 & l_2 & L \\ m_1 & m_2 & M \end{pmatrix} \begin{pmatrix} l_3 & l_4 & L \\ m_3 & m_4 & -M \end{pmatrix} \mathbb{T}_{l_3 l_4}^{l_1 l_2}(L) + \text{perms}. \end{aligned} \quad (5.13)$$

By evaluating the trispectrum correct to $\mathcal{O}[(C_L^{\phi\phi})^2]$, Ref. [59] show that the dominant terms for lensing reconstruction can be approximated by

$$\mathbb{T}_{l_3 l_4}^{l_1 l_2}(L) \approx \frac{1}{4} C_L^{\phi\phi} \tilde{f}_{l_1 L l_2} \tilde{f}_{l_3 L l_4}, \quad (5.14)$$

which involves the lensed $C_l^{\tilde{T}\tilde{T}}$ (via $\tilde{f}_{l_1 L l_2}$ etc.). The normalisation of Eq. (5.10) therefore correctly normalises the power spectrum $\hat{C}_L^{\hat{\phi}\hat{\phi}}$, avoiding a $\mathcal{O}(\phi^4)$ bias of around -15% on large scales that results from using the unlensed spectra, i.e. $f_{l_1 L l_2}$, in the normalisation A_L (and weight \tilde{g}).

Taking the expectation value $\langle \hat{C}_L^{\hat{\phi}\hat{\phi}} \rangle$, and using Eq. (5.13), gives [59, 200]

$$\langle \hat{C}_L^{\hat{\phi}\hat{\phi}} \rangle = C_L^{\phi\phi} + N_L^{(0)} + N_L^{(1)} + \mathcal{O}((C^{\phi\phi})^3), \quad (5.15)$$

where $N_L^{(n)}$ is of order $(C^{\phi\phi})^n$ if we do not count appearances of $C^{\phi\phi}$ in the lensed

5.1 Full-sky CMB lensing reconstruction

temperature power spectrum. The disconnected part of the lensed temperature 4-point function leads to the Gaussian bias

$$N_L^{(0)} = \frac{2A_L^2}{2L+1} \sum_{l_1, l_2} \tilde{g}_{l_1 l_2}^2(L) C_{l_1, \text{expt}}^{\tilde{T}\tilde{T}} C_{l_2, \text{expt}}^{\tilde{T}\tilde{T}} = A_L, \quad (5.16)$$

where the last equality holds only if the weights $\tilde{g}_{l_1 l_2}(L)$ are given by Eq. (5.11). Since this bias is present even in the absence of lensing, it corresponds to the power spectrum generated by the statistical noise of the lensing reconstruction (see Appendix 5.C for a more detailed discussion of how this bias is generated). The $N^{(1)}$ bias is due to those permutations in Eq. (5.13) that mix multipoles between the primary $(l_1 m_1, l_2 m_2)$ and $(l_3 m_3, l_4 m_4)$ couplings, e.g. $1 \leftrightarrow 3$. It has been computed in [59, 200]; see also Fig. 5.1. An unbiased estimate for the lensing potential power spectrum can be obtained by subtracting $N^{(0)} + N^{(1)}$ from $\hat{C}^{\phi\phi}$. Here, for $N^{(1)}$ we subtract the bias evaluated in the fiducial cosmology used for our simulations. We also evaluate the normalisation A_L for this fiducial model.¹

The subtraction of the Gaussian $N^{(0)}$ bias can be done in various ways. The simplest is to subtract a fiducial model. For temperature reconstructions, accurate $N^{(0)}$ subtraction is critical because $N^{(0)}$ generally exceeds the signal power $C^{\phi\phi}$ on all scales (by around two orders of magnitude at multipoles $l \sim 1000$). For the idealised isotropic surveys considered here, this is perhaps not too problematic since only $C_{l, \text{expt}}^{\tilde{T}\tilde{T}}$ is required to calculate $N^{(0)}$ and this can be estimated from a smoothed version of the measured CMB power spectrum. However, in the presence of a mask and other survey anisotropies, the Gaussian bias must generally be subtracted via simulations which requires an accurate procedure for simulating maps including all relevant real-world effects such as noise inhomogeneities and correlations, beam asymmetries and unresolved foreground emission. A work-around to these issues is

¹Note that the normalisation A_L is dependent on the cosmological model through e.g. $f_{l_1 L l_2}$. However, the lensed temperature power spectrum, and hence A_L , generally varies much less across models consistent with the measured CMB power spectrum than $C_L^{\phi\phi}$ (and so $N^{(1)}$). In practical applications, if a fiducial model is assumed to normalise the 4-point function, the parameter-dependent $C_L^{\phi\phi}$ can easily be renormalised in the likelihood by the ratio of the fiducial A_L to that at the current location in parameter space. This is not possible for $N^{(1)}$ because it is very slow to evaluate. Therefore, some modeling uncertainty for $N^{(1)}$ should be included in the covariance of $\hat{C}^{\phi\phi}$.

to use alternative data-dependent forms of $N^{(0)}$ (see [201] and references therein) whose expectation value either equals $N^{(0)}$ or closely approximates it. Here, we use the form (for an isotropic survey) $\hat{C}_L^{\hat{\phi}\hat{\phi}} - (2\hat{N}_L^{(0)} - N_L^{(0)})$, advocated in [59, 201], where

$$\hat{N}_L^{(0)} = \frac{2A_L^2}{2L+1} \sum_{l_1,l_2} \tilde{g}_{l_1 l_2}^2(L) \hat{C}_{l_1,\text{expt}}^{\tilde{T}\tilde{T}} C_{l_2,\text{expt}}^{\tilde{T}\tilde{T}}, \quad (5.17)$$

i.e. we replace one occurrence of $C_{l,\text{expt}}^{\tilde{T}\tilde{T}}$ in $N^{(0)}$ with the empirical temperature power spectrum of our given sky. Note that the expectation value $\langle \hat{N}_L^{(0)} \rangle = N_L^{(0)}$. As well as reducing the impact of modelling errors, this form of $N^{(0)}$ subtraction has the benefit of greatly reducing correlations between the $\hat{C}_L^{\hat{\phi}\hat{\phi}}$ estimates that arise from the disconnected part of the CMB 8-point function; see Sec. 5.3.2 and [59]. As we shall see in Sec. 5.4.1, it also eliminates the Gaussian contribution to the covariance between $\hat{C}_L^{\hat{\phi}\hat{\phi}}$ and the measured temperature power spectrum. Further motivation for the $2\hat{N}_L^{(0)} - N_L^{(0)}$ construction comes from considering optimal trispectrum estimation for weakly non-Gaussian fields [54]; these arguments are discussed further in the context of CMB lensing in Appendices 5.B and 5.C.

Throughout this chapter we assume that the instrumental beam has been deconvolved from the \tilde{T}_{lm} . The total power spectrum of the experiment is then of the form (2.101). Following [59], we will use $\sigma_N = 27 \mu\text{K arcmin}$ and $\sigma_{\text{FWHM}} = 7 \text{arcmin}$, which is roughly appropriate for Planck.

5.2 Simulations

We use simulations to verify our analytic arguments. These are based on 1000 realisations of a flat ΛCDM cosmology with WMAP7+BAO+ H_0 parameters [202] $h = 0.704$, $\Omega_b h^2 = 0.0226$, $\Omega_c h^2 = 0.1123$, $\tau = 0.087$, $n_s = 0.963$, $\Delta_{\mathcal{R}}^2 = 2.441 \times 10^{-9}$ at $k_0 = 0.002 \text{ Mpc}^{-1}$ and three massless neutrino species. We start with realisations of the unlensed temperature and lensing potential up to $l_{\max}^{\text{in}} = 3000$ (including the $C^{T\phi}$ correlation on large scales), which are then lensed using LensPix [203]. The lensed temperature up to $l_{\max}^T = 2750$ is used to reconstruct the lensing potential up to $l_{\max}^\phi = 2650$ with the full-sky simulation setup of Hanson *et al.* [59]. The convolution in harmonic space in Eq. (5.9) is evaluated as a product in pixel

space, where spin-1 spherical harmonic transforms are taken using HEALPix [204]. As a slight modification to [59] we use the reconstruction normalisation A_L and weights \tilde{g} given in Eqs. (5.10) and (5.11) to avoid the $N^{(2)}$ bias.

Figures 5.1 and 5.2 confirm that the power spectrum of the lensing reconstruction agrees with the input lensing power spectrum if the $N^{(0)}$ and $N^{(1)}$ biases are taken into account. Similar plots in [59] contain the $N^{(2)}$ bias because their simulations used reconstruction weights and normalisation with unlensed instead of lensed temperature power spectra. The realisation-dependent $\hat{N}^{(0)}$ bias correction does not change the expectation value of the reconstruction power spectrum but reduces its covariance.

5.3 Auto-correlations of power spectra

We will argue later that parameter estimation with the lensing reconstruction should be based on empirical power spectra defined by

$$\hat{C}_l^{XX} \equiv \frac{1}{2l+1} \sum_{m=-l}^l (-1)^m \hat{X}_{lm} \hat{X}_{l,-m}, \quad (5.18)$$

where \hat{X}_{lm} are the multipole coefficients of the realisation $\hat{X}(\hat{\mathbf{n}})$ of a field on the sphere. Here, $X = \tilde{T}$ or ϕ . To construct a likelihood for the empirical power spectra we will model auto- and cross-correlations of the empirical power spectra of the observed lensed temperature and the reconstructed lensing potential. If X is a statistically-isotropic, Gaussian field, then the power covariance is diagonal with

$$\text{cov}(\hat{C}_l^{XX}, \hat{C}_{l'}^{XX}) = \delta_{ll'} \frac{2}{2l+1} \langle \hat{C}_l^{XX} \rangle^2. \quad (5.19)$$

In the following we will abbreviate the Gaussian variance with

$$\text{var}_G(C_l^{XX}) \equiv \frac{2}{2l+1} (C_l^{XX})^2. \quad (5.20)$$

We demonstrate in Appendix 5.F that for most applications we can neglect the effect of $C^{T\phi}$ on the covariance between the power spectra of the lensed temperature and the lens reconstruction. Because the Taylor expansion of Eq. (5.2) is linear

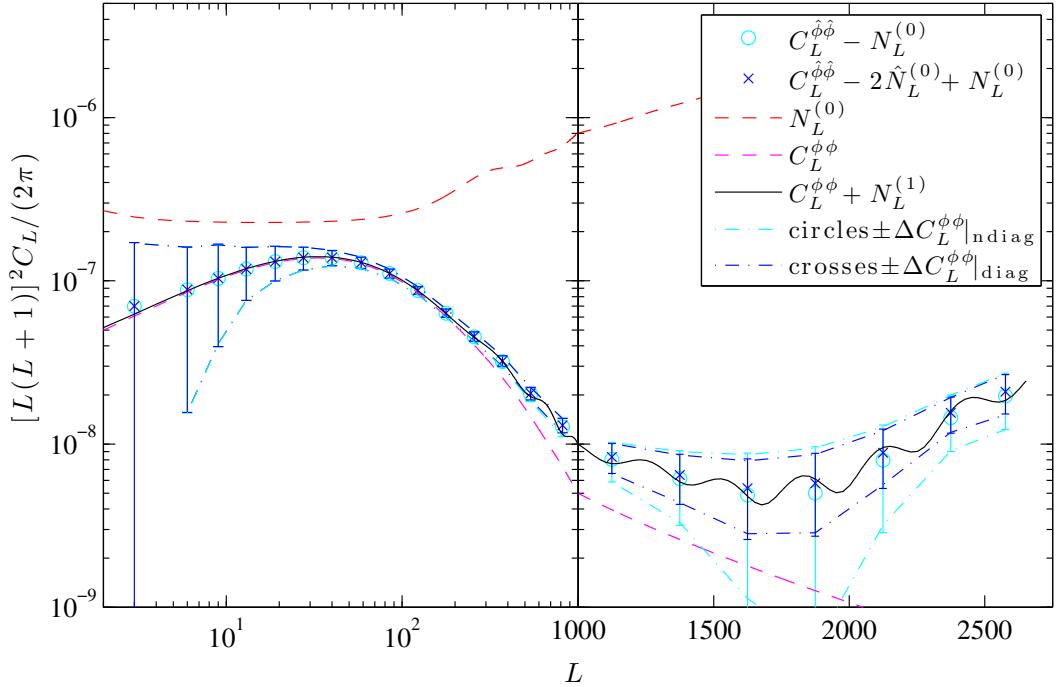


Figure 5.1: Reconstructed lensing potential power spectrum $\hat{C}_L^{\hat{\phi}\hat{\phi}}$ averaged over 1000 simulations, after subtracting the analytically-calculated $N^{(0)}$ bias (cyan circles). The magenta line is the input lensing potential power. Blue symbols include the empirical reconstruction bias correction of Eq. (5.17). The error bars show the estimated standard deviation of the binned power spectrum for a single realisation of the lensed CMB. Cyan error bars are only visible when they disagree from the blue error bars. Cyan dash-dotted lines show theoretical error bars obtained by binning Eqs. (5.27) and (5.28) below. Blue dash-dotted lines are theoretical error bars from Eq. (5.28) only. The biases $N^{(0)}$ (red) and $N^{(1)}$ (black, which also includes $C_L^{\phi\phi}$) are calculated analytically. Note that the left-hand panel uses a log scale for multipole L , whereas the right-hand panel uses a linear scale.

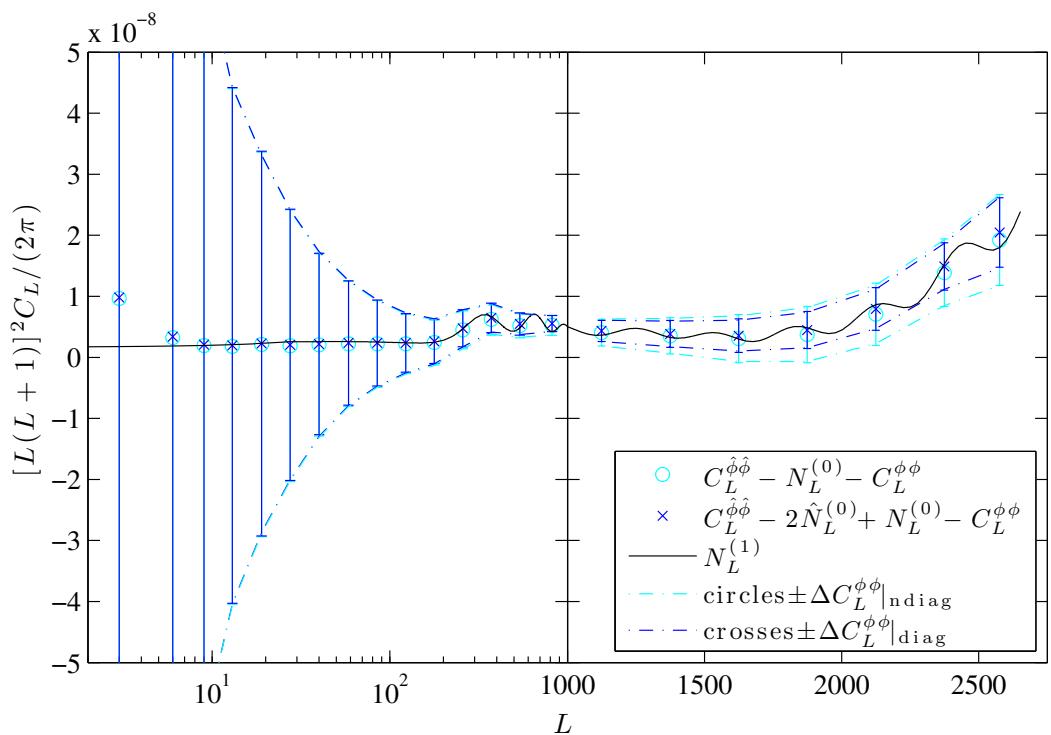


Figure 5.2: Same as Fig. 5.1 after subtracting the theoretical lensing power spectrum $C_L^{\phi\phi}$.

in the unlensed temperature all odd n -point functions of the lensed temperature vanish.

5.3.1 Lensed temperature

The auto-correlation of the lensed temperature power spectrum has been computed at first order in $C^{\phi\phi}$ in [98] under the flat-sky approximation and in [205] on the full sky. A contribution at second order in $C^{\phi\phi}$ was recently identified in [206]. The power covariance is given by

$$\begin{aligned} \text{cov}(\hat{C}_{l,\text{expt}}^{\tilde{T}\tilde{T}}, \hat{C}_{l',\text{expt}}^{\tilde{T}\tilde{T}}) &= \delta_{ll'} \text{var}_G(C_{l,\text{expt}}^{\tilde{T}\tilde{T}}) + \frac{1}{(2l+1)(2l'+1)} \sum_{m,m'} (-1)^{m+m'} \langle \tilde{T}_{lm} \tilde{T}_{l,-m} \tilde{T}_{l'm'} \tilde{T}_{l',-m'} \rangle_c, \end{aligned} \quad (5.21)$$

where $\langle \rangle_c$ denotes the connected part of the 4-point function, which is at $\mathcal{O}(C^{\phi\phi})$ [199]

$$\begin{aligned} &\langle \tilde{T}_{l_1} \tilde{T}_{l_2} \tilde{T}_{l_3} \tilde{T}_{l_4} \rangle_c \\ &= \frac{1}{2} C_{l_2}^{TT} C_{l_4}^{TT} \sum_{LM} (-1)^M C_L^{\phi\phi} \begin{pmatrix} l_1 & l_2 & L \\ m_1 & m_2 & -M \end{pmatrix} \begin{pmatrix} l_3 & l_4 & L \\ m_3 & m_4 & M \end{pmatrix} F_{l_1 L l_2} F_{l_3 L l_4} \\ &\quad + \text{all perms.} \end{aligned} \quad (5.22)$$

Here and in the following ‘all perms’ denotes permutations in all non-contracted multipole indices, i.e. permutations of 1, 2, 3, and 4 in Eq. (5.22). If we also

5.3 Auto-correlations of power spectra

include the contribution at second order in $C^{\phi\phi}$ from [206] we get [205, 206]¹

$$\begin{aligned} \text{cov}(\hat{C}_{l,\text{expt}}^{\tilde{T}\tilde{T}}, \hat{C}_{l',\text{expt}}^{\tilde{T}\tilde{T}}) &= \delta_{ll'} \text{var}_G(C_{l,\text{expt}}^{\tilde{T}\tilde{T}}) + \frac{2}{(2l+1)(2l'+1)} \sum_L C_L^{\phi\phi} f_{lLl'}^2 \\ &\quad + \sum_L \frac{\partial C_l^{\tilde{T}\tilde{T}}}{\partial C_L^{\phi\phi}} \frac{2}{2L+1} \left(C_L^{\phi\phi} \right)^2 \frac{\partial C_{l'}^{\tilde{T}\tilde{T}}}{\partial C_L^{\phi\phi}} + \mathcal{O}[(C^{\phi\phi})^3], \end{aligned} \quad (5.23)$$

where $f_{l_1 L l_2}$ is from the unlensed version of Eq. (5.8): $f_{l_1 L l_2} = C_{l_2}^{TT} F_{l_1 L l_2} + C_{l_1}^{TT} F_{l_2 L l_1}$. The third term on the right of Eq. (5.23) arises from cosmic variance of the lenses. Fluctuations at lens multipole L produce fluctuations in the empirical lensed temperature power spectrum over a range of multipoles. The fluctuations in the lens power, $\Delta C_L^{\phi\phi}$, propagate to the empirical temperature power spectrum approximately as $\partial C_l^{\tilde{T}\tilde{T}} / \partial C_L^{\phi\phi}$. The power derivative here can be calculated perturbatively by noting that at $\mathcal{O}(C^{\phi\phi})$ [82]

$$C_l^{\tilde{T}\tilde{T}} = C_l^{TT} [1 - l(l+1)R] + \sum_{l_1, l_2} C_{l_1}^{\phi\phi} C_{l_2}^{TT} \frac{F_{l_1 l_2}^2}{2l+1}, \quad (5.24)$$

where

$$R = \frac{1}{8\pi} \sum_l (2l+1) l(l+1) C_l^{\phi\phi} \quad (5.25)$$

is half the mean-squared deflection. Therefore

$$\frac{\partial C_l^{\tilde{T}\tilde{T}}}{\partial C_L^{\phi\phi}} = \sum_{l'} C_{l'}^{TT} \frac{F_{l' l}^2}{2l+1} - \frac{L(L+1)(2L+1)}{8\pi} l(l+1) C_l^{TT}. \quad (5.26)$$

While Ref. [206] included higher-order corrections to this expression by taking numerical derivatives of lensed power spectra computed non-perturbatively with the CAMB code [64, 84], these corrections are not expected to be important for our purposes here.

The off-diagonal contributions to the *correlation* between the empirical $\hat{C}_{l,\text{expt}}^{\tilde{T}\tilde{T}}$

¹The $\mathcal{O}(\phi^4)$ part of Eq. (5.23) was not derived in a rigorous perturbative analysis, which would imply additional corrections. For example, the unlensed $f_{lLl'}$ in the second term on the right could be replaced by its lensed counterparts, as in Eq. (5.14). We do not investigate such corrections to the temperature power auto-covariance further because corrections to the leading Gaussian term are negligible for all applications in this chapter.

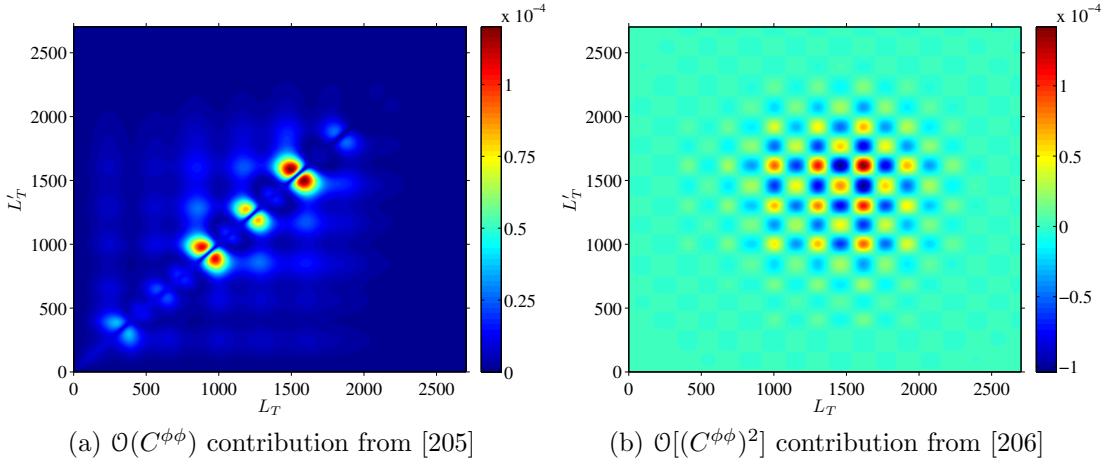


Figure 5.3: (a) Theoretical off-diagonal part of the correlation $\text{correl}(\hat{C}_{L_T}^{\tilde{T}\tilde{T}}, \hat{C}_{L'_T}^{\tilde{T}\tilde{T}})$ to first order in $C^{\phi\phi}$, given by the second term on the right of Eq. (5.23) (see [205]). The covariance is converted to a correlation by dividing by $[\text{var}_G(C_{L_T,\text{expt}}^{\tilde{T}\tilde{T}})\text{var}_G(C_{L'_T,\text{expt}}^{\tilde{T}\tilde{T}})]^{1/2}$. (b) $\mathcal{O}[(C^{\phi\phi})^2]$ contribution to the correlation identified in [206].

are shown in Fig. 5.3 (see [205] and [206] for similar plots). The checkerboard structure of the $\mathcal{O}[(C^{\phi\phi})^2]$ contribution arises because fluctuations in the lensing power produce changes in the lensed temperature spectra of opposite signs at the acoustic peaks and troughs. Both of the corrections in Eq. (5.23) give correlations that are at most of order 10^{-4} and are rather localised in the (l, l') plane. (Note that the correlations are suppressed on small scales where noise dominates the diagonal variance.) The impact of these non-diagonal contributions is found to be negligible for all calculations in this chapter, i.e. we can assume a Gaussian diagonal auto-covariance of the temperature power spectrum.

5.3.2 Lensing reconstruction

The auto-correlation of the lensing reconstruction power spectrum involves the 8-point function of the lensed temperature. Hanson *et al.* [59] found that the dominant off-diagonal contributions on the full sky are given by disconnected terms

5.3 Auto-correlations of power spectra

that contribute as (see [200] for a similar calculation on the flat sky)

$$\text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L'}^{\hat{\phi}\hat{\phi}})_{\text{non-diag}}^{\text{dom}} = \frac{32 A_L^2 A_{L'}^2}{(2L+1)(2L'+1)} \\ \times \sum_{l_1} \frac{1}{2l_1+1} (C_{l_1, \text{expt}}^{\tilde{T}\tilde{T}})^2 \left[\sum_{l_2} \tilde{g}_{l_1 l_2}^2(L) C_{l_2, \text{expt}}^{\tilde{T}\tilde{T}} \right] \left[\sum_{l_3} \tilde{g}_{l_1 l_3}^2(L') C_{l_3, \text{expt}}^{\tilde{T}\tilde{T}} \right]. \quad (5.27)$$

This dominates over more tightly-coupled terms that involve products of four weights \tilde{g} that do not factor in the form $\tilde{g}^2 \tilde{g}^2$ and therefore enforce a reduced summation volume (because the $3j$ -symbols appearing in the weights \tilde{g} impose more triangle conditions). The variance ($L = L'$) on the full sky is predominantly

$$\text{var}(\hat{C}_L^{\hat{\phi}\hat{\phi}}) = \frac{2}{2L+1} \langle C_L^{\hat{\phi}\hat{\phi}} \rangle^2, \quad (5.28)$$

with small corrections from Eq. (5.27) (for $L = L'$).

As shown in [59] the off-diagonal reconstruction power correlation can reach a level of 0.5% and is rather broad-band. If the reconstructed power is binned this can induce correlations of $\mathcal{O}(10\%)$ between different bins. Physically, these broad-band correlations arise because cosmic-variance fluctuations in the CMB at a given scale produce fluctuations in $\hat{C}^{\hat{\phi}\hat{\phi}}$ that are coherent over a broad range of scales due to the mode-coupling nature of lens reconstruction (small-scale CMB fluctuations are used to reconstruct large-scale lenses). To make this physical interpretation more explicit we note that the dominant non-diagonal covariance contribution of Eq. (5.27) can be written as

$$\text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L'}^{\hat{\phi}\hat{\phi}})_{\text{non-diag}}^{\text{dom}} = \sum_{l_1} \frac{\partial(2\hat{N}_L^{(0)})}{\partial \hat{C}_{l_1, \text{expt}}^{\tilde{T}\tilde{T}}} \frac{2}{2l_1+1} (C_{l_1, \text{expt}}^{\tilde{T}\tilde{T}})^2 \frac{\partial(2\hat{N}_{L'}^{(0)})}{\partial \hat{C}_{l_1, \text{expt}}^{\tilde{T}\tilde{T}}}, \quad (5.29)$$

where the (realisation-independent) derivative is given by

$$\frac{\partial(2\hat{N}_L^{(0)})}{\partial \hat{C}_{l_1, \text{expt}}^{\tilde{T}\tilde{T}}} = \frac{4A_L^2}{2L+1} \sum_{l_2} \tilde{g}_{l_1 l_2}^2(L) C_{l_2, \text{expt}}^{\tilde{T}\tilde{T}}, \quad (5.30)$$

which is non-zero even in the absence of lensing. Equation (5.30) describes the

change of the Gaussian reconstruction noise¹ resulting from fluctuations of the observed temperature realisation. In propagating these changes through to the covariance of $\hat{C}^{\hat{\phi}\hat{\phi}}$, one picks up the sample variance of the total lensed temperature power spectrum, $\text{var}_G(C_{l,\text{expt}}^{\tilde{T}\tilde{T}})$.

As noted in Sec. 5.1, using the realisation-dependent $\hat{N}^{(0)}$ bias correction of Eq. (5.17) significantly reduces the off-diagonal covariance of the reconstruction power spectrum. To help interpret the $\hat{N}^{(0)}$ correction, we write it in the form

$$\hat{C}_L^{\hat{\phi}\hat{\phi}} - 2\hat{N}_L^{(0)} + N_L^{(0)} = \hat{C}_L^{\hat{\phi}\hat{\phi}} - \sum_l \frac{\partial(2\hat{N}_L^{(0)})}{\partial \hat{C}_{l,\text{expt}}^{\tilde{T}\tilde{T}}} \hat{C}_{l,\text{expt}}^{\tilde{T}\tilde{T}} + N_L^{(0)}. \quad (5.31)$$

Therefore, for a given realisation of the lensed temperature, the empirical $\hat{N}^{(0)}$ bias correction partly removes the response of the Gaussian reconstruction noise to changes in the lensed temperature realisation. To see that this removes the non-diagonal power covariance of the lensing reconstruction caused by cosmic variance of the lensed temperature, note that both $\text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, 2\hat{N}_{L'}^{(0)})$ and $\text{cov}(2\hat{N}_L^{(0)}, 2\hat{N}_{L'}^{(0)})$ equal the right hand side of Eq. (5.29) at $\mathcal{O}(\phi^0)$. The empirical $\hat{N}^{(0)}$ correction leads to a small reduction in the *variance* of the binned reconstructed power spectrum, as shown in Figs. 5.1 and 5.2. Any residual covariance after empirical $\hat{N}^{(0)}$ subtraction is too small to be detected in our simulations.

5.4 Temperature-lensing cross-correlation

For a joint analysis involving the empirical power spectra of the lens reconstruction and the lensed temperature anisotropies, the likelihood should model their cross-correlation to avoid potential double-counting of lensing effects. In this section we calculate the cross-correlation. We recover the two main physical effects introduced in Sec. 2.3.4, i.e. a ‘‘noise contribution’’ from the cosmic variance of the lensed temperature affecting the noise in the reconstruction over a wide range of scales, and a ‘‘matter cosmic variance contribution’’ from cosmic variance of the lenses

¹When squaring the reconstruction $\hat{\phi}_{LM}$ to form the reconstruction power spectrum, we pick up not only the signal power $C_L^{\phi\phi}$ but also the noise of the reconstruction. Further details on the correspondence between noise terms in the reconstruction power and the $2\hat{N}^{(0)}$ expression are provided in Appendix 5.C.

altering the smoothing of the acoustic peaks in the temperature power spectrum. We will show in Sec. 5.4.1 that the noise contribution is due to the disconnected part of the lensed temperature 6-point function, while in Sec. 5.4.2 we show that the matter cosmic variance contribution is due to the connected part of the 6-point function. Corrections from the lensed temperature trispectrum generally have a sub-dominant effect on parameter estimation and are discussed in Appendix 5.D (see also Fig. 5.9 below).

5.4.1 Noise contribution

5.4.1.1 Perturbative derivation: Disconnected part of lensed temperature 4- and 6-point functions

Since the reconstructed lensing potential is quadratic in the lensed temperature the covariance $\text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})$ involves the 4- and 6-point functions of the lensed temperature. Using the definition of $\hat{\phi}$ in (5.9) we find

$$\begin{aligned} & \text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}}) \\ &= \frac{A_L^2}{(2L+1)(2L'+1)} \sum_{l_1, \dots, l_4, M, M'} (-1)^{M+M'} \begin{pmatrix} l_1 & l_2 & L \\ m_1 & m_2 & -M \end{pmatrix} \begin{pmatrix} l_3 & l_4 & L \\ m_3 & m_4 & M \end{pmatrix} \\ & \quad \times \tilde{g}_{l_1 l_2}(L) \tilde{g}_{l_3 l_4}(L) \left[\langle \tilde{T}_{l_1} \tilde{T}_{l_2} \tilde{T}_{l_3} \tilde{T}_{l_4} \tilde{T}_{L'M'} \tilde{T}_{L',-M'} \rangle - \langle \tilde{T}_{l_1} \tilde{T}_{l_2} \tilde{T}_{l_3} \tilde{T}_{l_4} \rangle \langle \tilde{T}_{L'M'} \tilde{T}_{L',-M'} \rangle \right]. \end{aligned} \quad (5.32)$$

Since all connected terms vanish in the absence of lensing we expect the noise contribution to come from the fully disconnected part. If we only keep disconnected terms, the second line of Eq. (5.32) can be replaced by

$$\begin{aligned} & \left[\langle \tilde{T}_{l_1} \tilde{T}_{l_2} \tilde{T}_{l_3} \tilde{T}_{l_4} \tilde{T}_{L'M'} \tilde{T}_{L',-M'} \rangle - \langle \tilde{T}_{l_1} \tilde{T}_{l_2} \tilde{T}_{l_3} \tilde{T}_{l_4} \rangle \langle \tilde{T}_{L'M'} \tilde{T}_{L',-M'} \rangle \right] \\ & \rightarrow 4 \left[\langle \tilde{T}_{l_1} \tilde{T}_{l_3} \rangle \langle \tilde{T}_{l_2} \tilde{T}_{L'M'} \rangle \langle \tilde{T}_{l_4} \tilde{T}_{L',-M'} \rangle + (M' \leftrightarrow -M') \right], \end{aligned} \quad (5.33)$$

where we exploited symmetry under relabelling ($l_1 \leftrightarrow l_2$) and/or ($l_3 \leftrightarrow l_4$). We also used that the contractions $\langle \tilde{T}_{l_1} \tilde{T}_{l_2} \rangle$ and $\langle \tilde{T}_{l_3} \tilde{T}_{l_4} \rangle$ do not contribute because

$\langle \hat{\phi}_{LM} \rangle = 0$. The fully disconnected part of Eq. (5.32) is therefore

$$\text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})_{\text{disconn.}} = \frac{8A_L^2}{(2L+1)(2L'+1)} \left(C_{L',\text{expt}}^{\tilde{T}\tilde{T}} \right)^2 \sum_{l_1} \tilde{g}_{l_1 L'}^2(L) C_{l_1,\text{expt}}^{\tilde{T}\tilde{T}}, \quad (5.34)$$

where the weight $\tilde{g}_{l_1 L'}(L)$ enforces $l_1 + L + L'$ to be even.

To interpret this result note that it can be expressed in terms of the derivative in Eq. (5.30) as

$$\text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})_{\text{disconn.}} = \frac{\partial(2\hat{N}_L^{(0)})}{\partial \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}}} \frac{2}{2L'+1} \left(C_{L',\text{expt}}^{\tilde{T}\tilde{T}} \right)^2. \quad (5.35)$$

This part of the covariance is therefore due to the response of the Gaussian reconstruction noise to changes in the observed temperature realisation and the resulting covariance with the observed temperature power. Based on this intuition, we anticipate that the covariance can be mitigated by the realisation-dependent $\hat{N}^{(0)}$ correction of the reconstruction power bias (see Sec. 5.4.1.4 for confirmation).

5.4.1.2 Magnitude and structure of the correlation matrix

In Fig. 5.4a we plot the power correlation resulting from the power covariance in Eq. (5.34) (denoting $L_\phi = L$ and $L_T = L'$ for convenience),

$$\text{correl}(\hat{C}_{L_\phi}^{\hat{\phi}\hat{\phi}}, \hat{C}_{L_T,\text{expt}}^{\tilde{T}\tilde{T}}) = \frac{\text{cov}(\hat{C}_{L_\phi}^{\hat{\phi}\hat{\phi}}, \hat{C}_{L_T,\text{expt}}^{\tilde{T}\tilde{T}})}{\sqrt{\text{var}_G(C_{L_\phi}^{\phi\phi} + N_{L_\phi}^{(0)} + N_{L_\phi}^{(1)}) \text{var}_G(C_{L_T,\text{expt}}^{\tilde{T}\tilde{T}})}}. \quad (5.36)$$

We plot the correlation of the unbinned spectra. Note that if the covariance is broad-band (i.e. roughly constant over the bin width) the *correlation* of (sufficiently finely) binned power spectra will increase roughly proportionally to the square root of the product of the two bin widths. The Gaussian variance of $C_{L_T,\text{expt}}^{\tilde{T}\tilde{T}}$ in the denominator of Eq. (5.36) contains the beam-deconvolved noisy temperature power spectrum (2.101), so that high temperature multipoles are suppressed.

The unbinned power correlation shown in Fig. 5.4a is mostly constrained to a cone-like region in the L_ϕ vs L_T plane, with the maximum correlation of 0.5%

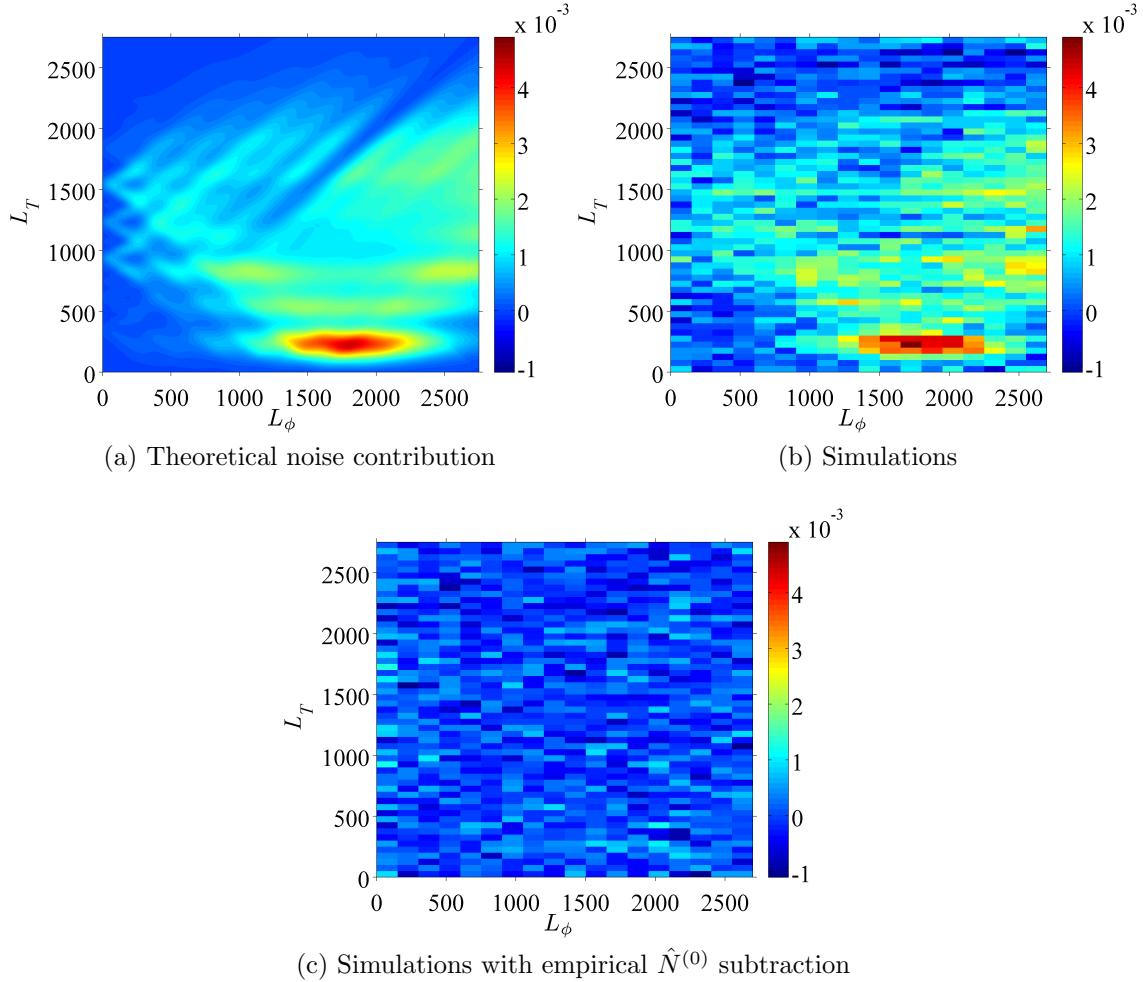


Figure 5.4: (a) Theoretical noise contribution from Eq. (5.34) to the correlation of unbinned power spectra of the lensed temperature and the reconstructed lensing potential, $\text{correl}(\hat{C}_{L_\phi}^{\hat{\phi}\hat{\phi}}, \hat{C}_{L_T, \text{expt}}^{\hat{T}\hat{T}})$. The acoustic peaks of the temperature power spectrum are visible in the vertical direction. (b) Estimate of the correlation of unbinned power spectra from 1000 simulations. (c) Same as (b) after subtracting the empirical $\hat{N}^{(0)}$ bias of Eq. (5.17) from $\hat{C}^{\hat{\phi}\hat{\phi}}$.

located at the first acoustic peak $L_T \sim 200$ and lensing reconstruction multipoles $1600 \lesssim L_\phi \lesssim 1900$. To understand the basic structure of the correlation in Fig. 5.4a we compute approximations to Eq. (5.34) in the limits $L_T \ll L_\phi$ and $L_\phi \ll L_T$.

For $L_T \ll L_\phi$, the weights in Eq. (5.34) restrict the summation from $l_1 = L_\phi - L_T$ to $l_1 = L_\phi + L_T$. If we Taylor expand in l_1 around L_ϕ we get

$$\begin{aligned} & \text{correl}(\hat{C}_{L_\phi}^{\hat{\phi}\hat{\phi}}, \hat{C}_{L_T, \text{expt}}^{\tilde{T}\tilde{T}})^{L_T \ll L_\phi}_{\text{disconn.}} \\ & \approx \frac{[L_\phi(L_\phi + 1)]^2 A_{L_\phi}}{8\pi} \sqrt{(2L_T + 1)(2L_\phi + 1)} \frac{L_T(L_T + 1) C_{L_T}^{\tilde{T}\tilde{T}}}{L_\phi(L_\phi + 1) C_{L_\phi, \text{expt}}^{\tilde{T}\tilde{T}}}. \end{aligned} \quad (5.37)$$

Recalling that $A_{L_\phi} = N_{L_\phi}^{(0)}$ for optimal weights, we see from Fig. 5.1 that the first term slightly increases with L_ϕ . The last term is maximised at the first acoustic peak $L_T \sim 200$ and at the reconstruction multipole $1600 \lesssim L_\phi \lesssim 1900$ where the observed temperature power $C_{L_\phi, \text{expt}}^{\tilde{T}\tilde{T}}$ is minimal (for the Planck-like noise and beam considered here). In this region Eq. (5.37) gives a correlation of around 0.4–0.5%, which agrees with Fig. 5.4a. Equation (5.37) also implies that lower noise in the temperature power spectrum would move the region where the correlation is maximal to higher reconstruction multipoles L_ϕ . The cone structure in Fig. 5.4a, with apex at $(L_\phi, L_T) \sim (1600–1900, 200)$ and edges $L_\phi - L_T \lesssim (1600–1900) \lesssim L_\phi + L_T$, encloses the region for which the sum over l_1 includes the maximum of $1/C_{l_1, \text{expt}}^{\tilde{T}\tilde{T}}$ [which is contained in the summand in Eq. (5.34)] around $l_1 = 1600–1900$. A similar argument can be applied for the cone patterns in the $L_\phi \lesssim 200$ region.

For high temperature and low reconstruction multipoles, $L_\phi \ll L_T$, we can Taylor expand Eq. (5.34) in l_1 around L_T (see [59] for a similar calculation):

$$\begin{aligned} & \text{correl}(\hat{C}_{L_\phi}^{\hat{\phi}\hat{\phi}}, \hat{C}_{L_T, \text{expt}}^{\tilde{T}\tilde{T}})^{L_\phi \ll L_T}_{\text{disconn.}} \approx \frac{[L_\phi(L_\phi + 1)]^2 A_{L_\phi}^2}{2\pi(A_{L_\phi} + C_{L_\phi}^{\phi\phi})} \frac{(C_{L_T}^{\tilde{T}\tilde{T}})^2}{(C_{L_T, \text{expt}}^{\tilde{T}\tilde{T}})^2} \sqrt{\frac{(2L_\phi + 1)(2L_T + 1)}{4}} \\ & \times \frac{1}{4} \left[\left(\frac{d \ln(L_T^2 C_{L_T}^{\tilde{T}\tilde{T}})}{d \ln L_T} \right)^2 + \frac{1}{2} \left(\frac{d \ln C_{L_T}^{\tilde{T}\tilde{T}}}{d \ln L_T} \right)^2 \right], \end{aligned} \quad (5.38)$$

where we have neglected $N_{L_\phi}^{(1)}$ and used $A_{L_\phi} = N_{L_\phi}^{(0)}$. The quadrature sum of

derivatives in the last term is maximal between acoustic peaks and troughs at $L_T \sim 350, 625, 925, 1225, 1550, 1850$ etc., which agrees with the temperature multipoles where the full correlation shown in Fig. 5.4a is maximal (for $L_\phi \lesssim 200$).¹ The maximum value of the second line of Eq. (5.38) is around 36 (for $L_T \lesssim 2700$). If we neglect $C_{L_\phi}^{\phi\phi}$ compared to A_{L_ϕ} , which is roughly acceptable for $L_\phi \lesssim 10$, then the first term in Eq. (5.38) is around 2×10^{-7} (see Fig. 5.1), i.e. for $L_\phi \lesssim 10$,

$$\text{correl}(\hat{C}_{L_\phi}^{\hat{\phi}\hat{\phi}}, \hat{C}_{L_T, \text{expt}}^{\tilde{T}\tilde{T}})^{L_\phi \ll L_T}_{\text{disconn.}} \lesssim 2 \times 10^{-7} \times \frac{1}{2} \sqrt{(2L_\phi + 1)(2L_T + 1)} \times 36. \quad (5.39)$$

For example, for $L_\phi = 10, L_T = 1500$ the bound is 9×10^{-4} which is consistent with the full result shown in Fig. 5.4a.

5.4.1.3 Comparison with simulations

Before assessing the relevance of the noise contribution to the covariance for parameter estimation we compare the analytic result in Eq. (5.34) with the full covariance estimated from our simulations. For the latter, we use

$$\widehat{\text{cov}}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L', \text{expt}}^{\tilde{T}\tilde{T}}) = \frac{1}{N_{\text{sims}} - 1} \sum_{s=1}^{N_{\text{sims}}} (\hat{C}_{L,s}^{\hat{\phi}\hat{\phi}} - \langle \hat{C}_L^{\hat{\phi}\hat{\phi}} \rangle_{\text{sims}})(\hat{C}_{L', \text{expt},s}^{\tilde{T}\tilde{T}} - \langle \hat{C}_{L', \text{expt}}^{\tilde{T}\tilde{T}} \rangle_{\text{sims}}), \quad (5.40)$$

where s labels different realisations and $\langle \cdot \rangle_{\text{sims}}$ denotes the average over N_{sims} realisations. To reduce the noise of the estimates from the finite number of simulations,

¹The correlation at high L_T is suppressed by the noise in $C_{L_T, \text{expt}}^{\tilde{T}\tilde{T}}$ which is used to normalise the covariance in Eq. (5.36).

we average the measured covariance over a range of L and L' values,¹

$$\begin{aligned} \widehat{\text{cov}}(\hat{C}_{\bar{L}_i}^{\hat{\phi}\hat{\phi}}, \hat{C}_{\bar{L}'_j, \text{expt}}^{\tilde{T}\tilde{T}}) &= \frac{1}{\Delta L_i \Delta L'_j} \frac{1}{[\bar{L}_i(\bar{L}_i + 1)]^2 [\bar{L}'_j(\bar{L}'_j + 1)]} \\ &\times \sum_{L=L_i}^{L_{i+1}-1} \sum_{L'=L'_j}^{L'_{j+1}-1} [L(L+1)]^2 L'(L'+1) \widehat{\text{cov}}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L'}^{\tilde{T}\tilde{T}}), \end{aligned} \quad (5.41)$$

where L_i and L'_j are bin boundaries for lensing and temperature powers, respectively, and \bar{L}_i and \bar{L}'_j denote the corresponding bin centres. The bin widths are $\Delta L_i = L_{i+1} - L_i$ and $\Delta L'_j = L'_{j+1} - L'_j$. We divide out $L(L+1)$ prefactors to average over relatively slowly varying quantities. In Fig. 5.4b we plot the estimate of the correlation that *unbinned* power spectra would have. This is obtained by dividing the covariance estimate of Eq. (5.41) by the theoretical Gaussian variance of unbinned power spectra as in Eq. (5.36) [evaluated at (\bar{L}_i, \bar{L}'_j)]. Within the random scatter from the finite N_{sims} , the estimated correlation agrees with the theoretical noise contribution of Eq. (5.34). We can also conclude that the noise contribution is the dominant part of the temperature-lensing power correlation.

5.4.1.4 Mitigating the noise contribution with the empirical $\hat{N}^{(0)}$ bias correction

In practice, it is desirable that the temperature and reconstruction power spectra are uncorrelated so that their respective likelihoods may be simply combined. Indeed, this is the assumption that has been made in all joint analyses to date [3, 94, 99].

Fortunately, the empirical $\hat{N}^{(0)}$ subtraction, which was originally proposed in [59] to eliminate the non-diagonal reconstruction power auto-covariance (5.27), also removes the noise contribution [Eq. (5.34)] to the temperature-lensing power cross-covariance. To see this note that if the empirical bias correction of Eq. (5.17)

¹Note that for broad-band covariances this binning procedure does not bias the covariance estimate, i.e., the binned covariance agrees with the unbinned covariance in the limit of averaging over infinitely many simulations. However, for a finite number of simulations, the binned covariance is less noisy than the covariance estimate at a single (L, L') pair. Note also that binning the estimated covariance is equivalent to estimating the covariance of the binned spectra.

is used, the cross-covariance changes by

$$\begin{aligned} -\text{cov}(2\hat{N}_L^{(0)}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}}) &= -\sum_l \frac{\partial(2\hat{N}_L^{(0)})}{\partial \hat{C}_{l,\text{expt}}^{\tilde{T}\tilde{T}}} \text{cov}(\hat{C}_{l,\text{expt}}^{\tilde{T}\tilde{T}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}}) \\ &= -\text{cov}(\hat{C}_L^{\phi\phi}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})_{\text{disconn.}} + \mathcal{O}(\phi^2). \end{aligned} \quad (5.42)$$

We establish a more general version of this result in [6], where we show that the generalisation of the $\hat{N}^{(0)}$ correction for anisotropic surveys removes the noise contribution to the covariance with any quadratic estimate (including e.g. cross-spectra) of the temperature power spectrum. We confirm the reduction in the cross-covariance with simulations in Fig. 5.4c. Corrections to Eq. (5.42) from the non-Gaussian terms in the covariance of the temperature power spectra [see Eq. (5.23)] at $\mathcal{O}(\phi^2)$ and $\mathcal{O}(\phi^4)$ reach at most 3×10^{-5} in the correlation, which is roughly one order of magnitude smaller than the matter cosmic variance contribution discussed below (see Fig. 5.5a). These corrections are too small to be visible in Fig. 5.4c and we neglect them in the following.

5.4.2 Matter cosmic variance contribution

5.4.2.1 Warm-up: Power covariance of input lensing potential and lensed temperature

We expect the cosmic variance of the lenses to induce a power-correlation of the lensed temperature with the lensing reconstruction since greater lensing power in a given realisation leads to additional smoothing of the empirical temperature power spectrum. As a warm-up, we calculate how the same effect gives rise to a covariance between the power spectrum of the temperature and the (unobservable) power spectrum of the lensing potential, as if we were able to measure ϕ directly with no noise (i.e. assuming that we knew the input lensing potential). This correlation can be extracted from simulations simply by measuring the correlation of the power spectra of the lensed temperature and the input lensing potential without performing any lensing reconstruction. To calculate the covariance perturbatively, note that for a fixed realisation of the input lensing potential ϕ_{in} the lensed temperature power spectrum obtained by averaging only over the unlensed

CMB is given by ($L' \neq 0$)

$$\begin{aligned} & \frac{1}{2L'+1} \sum_{M'} \langle |\tilde{T}_{L'M'}|^2 \rangle_{\text{CMB}} \\ &= C_{L'}^{TT} + \frac{1}{2L'+1} \sum_{l_1 l_2} \hat{C}_{l_1}^{\phi_{\text{in}} \phi_{\text{in}}} C_{l_2}^{TT} F_{L'l_1 l_2}^2 - \hat{R} L'(L'+1) C_{L'}^{TT} \\ &= C_{L'}^{TT} + \sum_{L''} \frac{\partial C_{L'}^{\tilde{T}\tilde{T}}}{\partial C_{L''}^{\phi\phi}} \hat{C}_{L''}^{\phi_{\text{in}} \phi_{\text{in}}}. \end{aligned} \quad (5.43)$$

Here $\hat{C}^{\phi_{\text{in}} \phi_{\text{in}}}$ denotes the empirical power spectrum of the input lensing potential realisation ϕ_{in} ¹ and \hat{R} is defined by replacing $C^{\phi\phi}$ in Eq. (5.25) by $\hat{C}^{\phi_{\text{in}} \phi_{\text{in}}}$. The first line of Eq. (5.43) can be derived from the expansion in Eq. (5.2) at second order in ϕ_{in} . The term linear in ϕ_{in} vanishes because it is proportional to the monopole of ϕ_{in} . In Eq. (5.43) we neglected noise in the temperature measurement since this does not contribute to the covariance that we aim to calculate. The derivative in the second line of Eq. (5.43) is given by Eq. (5.26). Neglecting the T - ϕ correlation, we then have

$$\begin{aligned} \text{cov}(\hat{C}_L^{\phi_{\text{in}} \phi_{\text{in}}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}}) &= \langle \hat{C}_L^{\phi_{\text{in}} \phi_{\text{in}}} \langle \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}} \rangle_{\text{CMB}} \rangle_{\text{LSS}} - \langle \hat{C}_L^{\phi_{\text{in}} \phi_{\text{in}}} \rangle \langle \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}} \rangle, \\ &= \frac{2}{2L+1} (C_L^{\phi\phi})^2 \frac{\partial C_{L'}^{\tilde{T}\tilde{T}}}{\partial C_L^{\phi\phi}}, \end{aligned} \quad (5.44)$$

where $\langle \cdot \rangle_{\text{LSS}}$ denotes an average over realisations of ϕ (i.e., over large-scale structure).

5.4.2.2 Power covariance of reconstructed lenses and lensed temperature from the connected 6-point function

The power covariance [Eq. (5.32)] of the *reconstructed* lensing potential and the lensed temperature receives contributions from the connected 4- and 6-point functions, as well as the Gaussian (disconnected) part. The leading-order trispectrum is linear in $C^{\phi\phi}$ and so cannot give rise to the expected matter cosmic variance

¹We denote this as ϕ_{in} , while the reconstructed lensing potential is $\hat{\phi}$. Recall that C denotes theoretical and \hat{C} empirical power spectra.

contribution calculated above. The trispectrum contributions are discussed in detail in Appendix 5.D, and are shown to have a sub-dominant effect on parameter estimation (see Fig. 5.9 below). We therefore focus here on the contribution from the connected 6-point function. Terms independent of ϕ do not contribute to the connected part and terms linear in ϕ vanish when averaging over large-scale structure. Since the ϕ^2 contribution can be shown to vanish [207] as well as ϕ^3 terms which vanish for Gaussian ϕ , the leading-order contribution is of order ϕ^4 . The $\mathcal{O}(\phi^4)$ contribution from the connected 6-point function to the temperature-lensing power covariance is calculated in detail in Appendix 5.E. The dominant contribution there [Eq. (5.106)] is exactly of the form calculated above [Eq. (5.44)] for the power covariance of the input lensing potential and the lensed temperature.

5.4.2.3 Magnitude and structure of the covariance matrix

The *correlation* of unbinned temperature and lensing power spectra due to (5.106) is at most $\sim 0.04\%$ (see Fig. 5.5a). Since higher lensing power lowers the peaks and increases the troughs of the temperature power, the correlation is negative at multipoles L_T where the temperature power peaks and positive where it has a trough. The correlation is small for $L_\phi \gtrsim 150$ because the acoustic peak smearing is mainly caused by large-scale lenses [98].

The structure of the correlation in Fig. 5.5a suggests that it can be approximated by a low-rank matrix. Indeed, if we perform a singular-value-decomposition of the covariance of $[L_\phi(L_\phi + 1)]^2 \hat{C}_{L_\phi}^{\hat{\phi}\hat{\phi}}/(2\pi)$ and $L_T(L_T + 1) \hat{C}_{L_T}^{\tilde{T}\tilde{T}}/(2\pi)$ induced by (5.106), the first singular value is 34 times larger than the second (see Fig. 5.6). Truncating after the first singular value, the rank-one approximation of the covariance is

$$\text{cov} \left(\frac{[L_\phi(L_\phi + 1)]^2 \hat{C}_{L_\phi}^{\hat{\phi}\hat{\phi}}}{2\pi}, \frac{L_T(L_T + 1) \hat{C}_{L_T}^{\tilde{T}\tilde{T}}}{2\pi} \right)_{(5.106)} \approx \lambda_1 u_{L_\phi} v_{L_T}, \quad (5.45)$$

where the maximum singular value is $\lambda_1 = 8.3 \times 10^{-8} \mu\text{K}^2$ and the associated left and right singular vectors u_{L_ϕ} and v_{L_T} are shown in Fig. 5.7. The approximation in (5.45) is so accurate that differences between the l.h.s. and r.h.s. are not visible

¹More accurately, we estimate the covariance and normalise it as in Eq. (5.36).

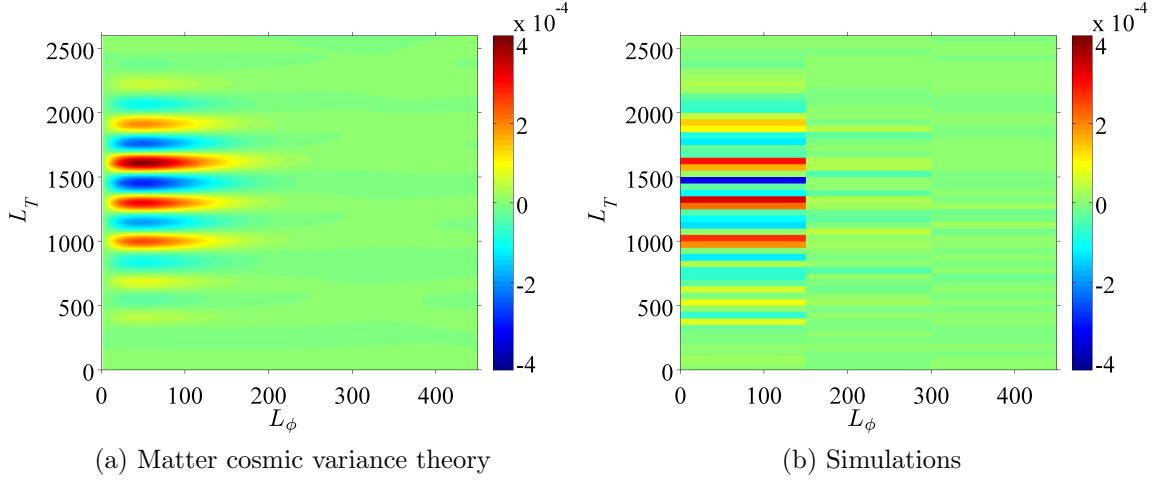


Figure 5.5: (a) Theoretical matter cosmic variance contribution (5.106) to the correlation of the unbinned power spectra of reconstructed lensing potential and lensed temperature. The covariance (5.106) is converted to a correlation using the same conversion factor as in (5.36). (b) Measured covariance $\hat{\text{cov}}(\hat{C}_{L_\phi}^{\phi_{\text{in}}\phi_{\text{in}}}, \hat{C}_{L_T}^{TT} - \hat{C}_{L_T}^{TT})$ of the input lensing potential power spectrum and the difference of noise-free lensed and unlensed temperature in 1000 simulations, normalised as in Eq. (5.36).

if both sides are plotted (see Fig. 5.8). Therefore, the power covariance (5.106) due to matter cosmic variance is caused by the fact that cosmic variance fluctuations of the lensing power spectrum that have the shape of the mode u_{L_ϕ} induce correlated changes of the temperature power spectrum that have the form of v_{L_T} . The accuracy of the rank-one approximation can be understood analytically by Taylor-expanding the temperature power spectrum appearing in the derivative (5.26) in the limit $L_\phi \ll L_T$ and noting that the result has a rank-one structure (see [6] for details). This result is consistent with [98] who found that the lensed CMB temperature power spectrum is mostly sensitive to a single mode of the lensing power spectrum $C^{\phi\phi}$.

5.4.2.4 Comparison with simulations

Before assessing the importance for parameter estimation of the matter cosmic variance contribution to the cross-covariance, we validate it against simulations. The measurements of the temperature-lensing power correlation in Fig. 5.4b are

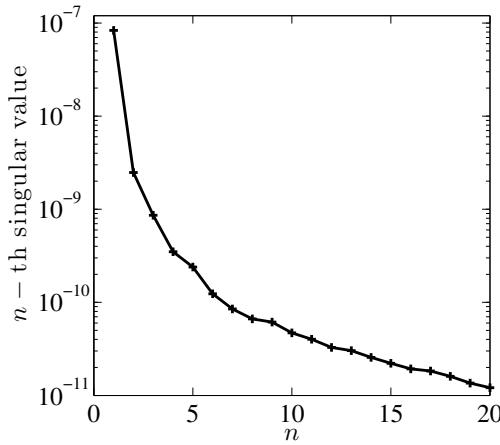


Figure 5.6: Twenty largest singular values of the covariance between $[L_\phi(L_\phi + 1)]^2 \hat{C}_{L_\phi}^{\phi\phi}/(2\pi)$ and $L_T(L_T + 1) \hat{C}_{L_T}^{TT}/(2\pi)$ induced by the matter cosmic variance contribution (5.106), for $l_{\max}^\phi = 500$ and $l_{\max}^T = 2500$.

too noisy to resolve clearly the matter cosmic variance contribution. However, we can test Eq. (5.44) by correlating the lensed temperature with the empirical power spectrum $\hat{C}^{\phi_{\text{in}}\phi_{\text{in}}}$ of the realisation of the input lensing potential without performing any reconstruction. To reduce the noise of the covariance estimate we work with the empirical power of the lensed temperature without including beam effects or noise, which do not affect the matter cosmic variance effect we are looking for. Since we have so far neglected T - ϕ correlations in all calculations, we try to eliminate correlations between the lensing potential and the unlensed temperature in the simulations by calculating $\widehat{\text{cov}}(\hat{C}_L^{\phi_{\text{in}}\phi_{\text{in}}}, \hat{C}_{L'}^{TT} - \hat{C}_{L'}^{TT})$. Here, $\hat{C}^{\phi_{\text{in}}\phi_{\text{in}}}$ is the empirical power of the input lensing potential and \hat{C}^{TT} is the empirical power of the unlensed temperature. Additionally, subtracting the unlensed from the lensed empirical temperature power reduces the noise of the covariance estimate because it eliminates the scatter due to cosmic variance of the unlensed temperature. Otherwise, our estimate of the covariance follows the procedure described in Sec. 5.4.1.3. As shown in Fig. 5.5b, these broad-band estimates are consistent with the theoretical expectation from Eq. (5.44).

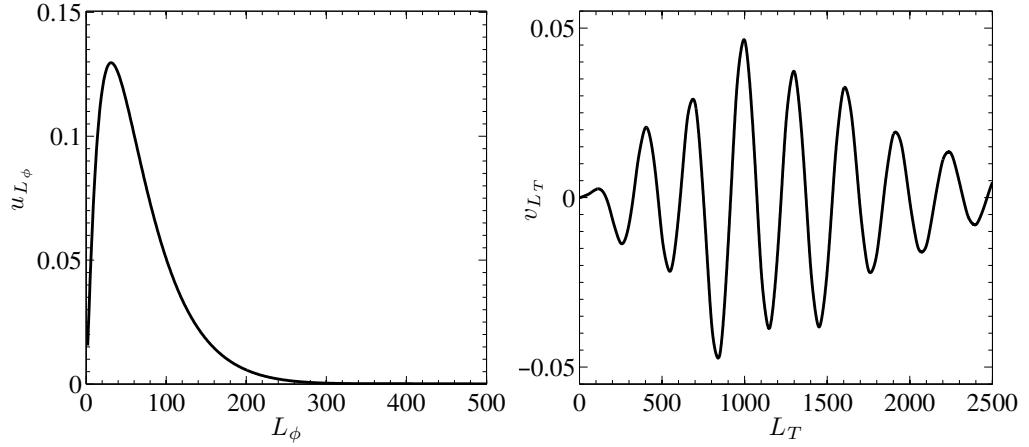


Figure 5.7: Left and right singular vectors associated with the largest singular value of the covariance between $[L_\phi(L_\phi + 1)]^2 \hat{C}_{L_\phi}^{\hat{\phi}\hat{\phi}}/(2\pi)$ and $L_T(L_T + 1) \hat{C}_{L_T}^{\tilde{T}\tilde{T}}/(2\pi)$ induced by Eq. (5.106), for $l_{\max}^\phi = 500$ and $l_{\max}^T = 2500$.

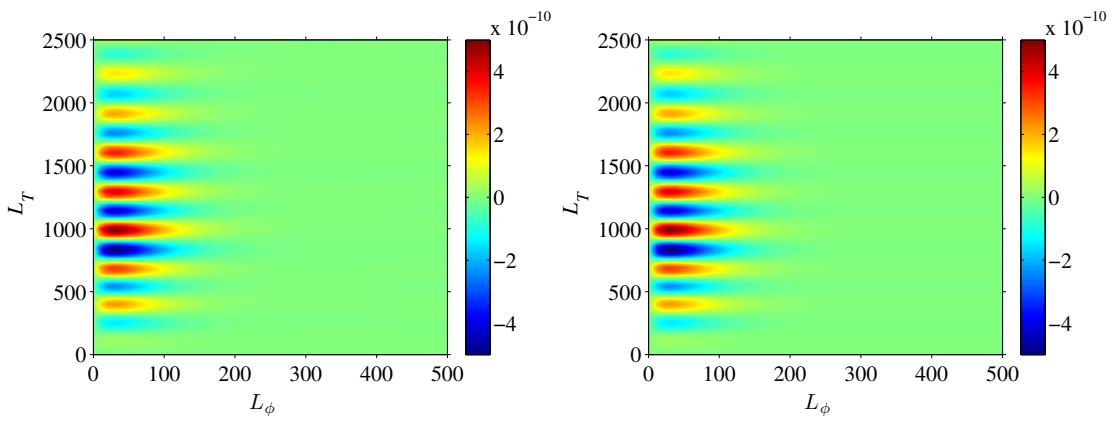


Figure 5.8: *Left:* The approximate contribution to the covariance between $[L_\phi(L_\phi + 1)]^2 \hat{C}_{L_\phi}^{\hat{\phi}\hat{\phi}}/(2\pi)$ and $L_T(L_T + 1) \hat{C}_{L_T}^{\tilde{T}\tilde{T}}/(2\pi)$ from cosmic variance of the lenses, derived from Eq. (5.106). *Right:* The rank-one approximation to the matrix on the left from retaining only the largest singular value.

5.4.2.5 Mitigating the matter cosmic variance contribution

As we shall show in Sec. 5.4.4, the impact on parameter errors of ignoring the covariance between the lensed temperature and reconstruction power due to cosmic variance of the lenses is small for an experiment like Planck. Indeed, this is why the covariance is not accounted for in the current Planck likelihood [3]. However, the covariance is simple to model using Eq. (5.106) and could easily be included in a joint analysis. Equivalently, the covariance could be diagonalised by appropriate modifications of the empirical temperature or lensing power spectra. A symmetric way to do this, mirroring the empirical $N^{(0)}$ correction that we advocate applying to the empirical lensing reconstruction power, is to modify the measured temperature power spectrum by

$$\hat{C}_{l,\text{expt}}^{\tilde{T}\tilde{T}} \rightarrow \hat{C}_{l,\text{expt}}^{\tilde{T}\tilde{T}} - \sum_{l'} A_{ll'} (\hat{C}_{l'}^{\hat{\phi}\hat{\phi}} - 2\hat{N}_{l'}^{(0)}) + \left\langle \sum_{l'} A_{ll'} (\hat{C}_{l'}^{\hat{\phi}\hat{\phi}} - 2\hat{N}_{l'}^{(0)}) \right\rangle, \quad (5.46)$$

where

$$A_{ll'} = \frac{\partial C_l^{\tilde{T}\tilde{T}}}{\partial C_{l'}^{\hat{\phi}\hat{\phi}}} \left(\frac{C_{l'}^{\phi\phi}}{\langle \hat{C}_{l'}^{\hat{\phi}\hat{\phi}} \rangle} \right)^2. \quad (5.47)$$

The expectation value is introduced in Eq. (5.46) to preserve the mean value of $\hat{C}_{l,\text{expt}}^{\tilde{T}\tilde{T}}$. The construct $(C_{l'}^{\phi\phi}/\langle \hat{C}_{l'}^{\hat{\phi}\hat{\phi}} \rangle)^2$ acts like a Wiener filter on the empirical spectrum $\hat{C}_{l'}^{\hat{\phi}\hat{\phi}} - 2\hat{N}_{l'}^{(0)}$. The effect is to ‘delens’ the lensed temperature power spectrum with the (CMB-averaged) lensing effect from any scales in the reconstructed potential power spectrum where the signal-to-noise on the reconstruction is high. In practice, high signal-to-noise lens reconstructions are never achieved with reconstructions based only on the temperature.

An alternative procedure is suggested by Eq. (5.45): project out from $\hat{C}_L^{\hat{\phi}\hat{\phi}}$ the dominant singular vector u_L . Although this is lossy, it does have the virtue that the $\tilde{T}\tilde{T}$ part of the likelihood is left unchanged. We discuss the effect of such a lossy projection on the errors of amplitude estimates in Sec. 5.4.4.1.

XY, ZW	UU	VV	Equation here	Reference
Any combination of $\tilde{T}\tilde{T}$, $\tilde{T}\tilde{E}$, $\tilde{E}\tilde{E}$	$\phi\phi$	--	(5.23)	[206]
$\tilde{B}\tilde{B}$, $\tilde{B}\tilde{B}$	$\phi\phi$	EE		[206]
$\hat{\phi}\hat{\phi}$, $\hat{\phi}\hat{\phi}$	--	$\tilde{T}\tilde{T}$	(5.28), (5.29)	[59, 200], this work
$\hat{\phi}\hat{\phi} - 2\hat{N}^{(0)} + N^{(0)}$, $\hat{\phi}\hat{\phi} - 2\hat{N}^{(0)} + N^{(0)}$	--	--	(5.28)	[59]
$\hat{\phi}\hat{\phi}$, $\tilde{T}\tilde{T}$	$\phi\phi$	$\tilde{T}\tilde{T}$	(5.106), (5.35)	this work
$\hat{\phi}\hat{\phi} - 2\hat{N}^{(0)} + N^{(0)}$, $\tilde{T}\tilde{T}$	$\phi\phi$	--	(5.106)	this work

Table 5.2: Covariance types UU and VV that are picked up by power covariances $\text{cov}(\hat{C}_L^{XY}, \hat{C}_{L'}^{ZW})$ according to Eq. (5.48). Lensed power spectra are denoted with tildes and for $\tilde{T}\tilde{T}$ the VV covariance should be understood to correspond to $C_{l,\text{expt}}^{\tilde{T}\tilde{T}}$ including noise. The notation $\hat{\phi}\hat{\phi} - 2\hat{N}^{(0)} + N^{(0)}$ stands for $\hat{C}_L^{XY} = \hat{C}_L^{\hat{\phi}\hat{\phi}} - 2\hat{N}_L^{(0)} + N_L^{(0)}$, and in this case the Gaussian variance in the first term of Eq. (5.48) involves $\langle \hat{C}_L^{\hat{\phi}\hat{\phi}} \rangle^2$, while the Gaussian covariance with $\hat{C}_{L'}^{\tilde{T}\tilde{T}}$ involves $(C_L^{T\phi})^2$. The symbol -- means that the corresponding contribution is cancelled or can be neglected. For the temperature-lensing covariances in the bottom two rows the leading trispectrum contribution of Eq. (5.97) should be added to Eq. (5.48).

5.4.3 Towards a complete model for power covariances

The power covariances of Eqs. (5.29), (5.35) and (5.106) can be regarded as a natural extension of the model for temperature and polarization power covariances found in [206]. We can summarise the covariances in the unified form

$$\begin{aligned} \text{cov}(\hat{C}_L^{XY}, \hat{C}_{L'}^{ZW}) &= \delta_{LL'} \text{cov}_G(\hat{C}_L^{XY}, \hat{C}_{L'}^{ZW}) + \sum_{ll'} \frac{\partial C_L^{XY}}{\partial C_l^{UU}} \text{cov}(\hat{C}_l^{UU}, \hat{C}_{l'}^{UU}) \frac{\partial C_{L'}^{ZW}}{\partial C_{l'}^{UU}} \\ &\quad + \sum_{ll'} \frac{\partial C_L^{XY}}{\partial C_l^{VV}} \text{cov}(\hat{C}_l^{VV}, \hat{C}_{l'}^{VV}) \frac{\partial C_{L'}^{ZW}}{\partial C_{l'}^{VV}}, \end{aligned} \quad (5.48)$$

where U and V are listed in Table 5.2 for different combinations of XY and ZW . In this context we make the identifications

$$\frac{\partial C_L^{\hat{\phi}\hat{\phi}}}{\partial C_l^{\phi\phi}} \equiv \delta_{ll}, \quad \frac{\partial C_L^{\hat{\phi}\hat{\phi}}}{\partial C_l^{\tilde{T}\tilde{T}}} \equiv \frac{\partial(2\hat{N}_L^{(0)})}{\partial \hat{C}_{l,\text{expt}}^{\tilde{T}\tilde{T}}}. \quad (5.49)$$

The first term on the right of Eq. (5.48) is the Gaussian covariance that would arise for Gaussian fields X , Y , Z and W .

The general formula of Eq. (5.48) does not include trispectrum contributions to

the temperature-lensing covariance, but the dominant correction has a simple form [Eq. (5.97)], which can be added straightforwardly to the covariance. Generally, terms involving $C^{T\phi}$ in Eq. (5.48) can be neglected. While we evaluated derivatives perturbatively in ϕ , non-perturbative corrections can be included from numerical derivatives of accurate lensed power spectra [64, 206]. However we do not expect these corrections to be significant here. All combinations listed in Table 5.2 have been verified with simulations in [59, 206] or in this work. Extending the covariance model to polarization-based lensing reconstructions would be interesting but is beyond the scope of this chapter.

5.4.4 Impact of correlations on parameter estimation

5.4.4.1 Lensing amplitude estimates

As a first step in assessing the impact of covariances between the temperature and lens-reconstruction power spectra on parameter estimation, we consider constraining an overall amplitude parameter A of a fiducial lensing power spectrum, $C_l^{\phi\phi} = AC_l^{\phi\phi}|_{\text{fid}}$ [75] with all other parameters fixed. The value of $A = 1$ corresponds to lensing at the level expected in the fiducial model, while $A = 0$ corresponds to no lensing.

The lensing amplitude can be estimated from the reconstructed lensing potential with

$$\hat{A} = \frac{\sum_{l,l'} C_l^{\phi\phi} \left(\text{cov}_{\hat{\phi}\hat{\phi}}^{-1} \right)_{ll'} \left(\hat{C}_{l'}^{\hat{\phi}\hat{\phi}} - N_{l'}^{(0)} - N_{l'}^{(1)} \right)}{\sum_{L,L'} C_L^{\phi\phi} \left(\text{cov}_{\hat{\phi}\hat{\phi}}^{-1} \right)_{LL'} C_{L'}^{\phi\phi}}, \quad (5.50)$$

where $\text{cov}_{\hat{\phi}\hat{\phi}}$ denotes the auto-covariance of the reconstructed lensing power given by Eqs. (5.27) and (5.28), evaluated for $A = 1$. Equation (5.50) is the maximum-likelihood estimator for the lensing amplitude if the likelihood is modeled as a multi-variate Gaussian in the empirical power spectrum of the lensing reconstruction. This form of the likelihood will be motivated later. Alternatively, the lensing amplitude can be extracted directly from the lensed temperature power spectrum

without invoking lensing reconstruction by

$$\hat{A}' = \frac{\sum_l (\hat{C}_{l,\text{expt}}^{\tilde{T}\tilde{T}} - C_{l,\text{expt}}^{TT}) (\text{cov}_{\tilde{T}\tilde{T},\text{expt}}^{-1})_{ll} (C_l^{\tilde{T}\tilde{T}} - C_l^{TT})}{\sum_{l'} (C_{l'}^{\tilde{T}\tilde{T}} - C_{l'}^{TT})^2 (\text{cov}_{\tilde{T}\tilde{T},\text{expt}}^{-1})_{l'l'}}, \quad (5.51)$$

where the auto-covariance of the temperature power is approximated by its leading-order diagonal piece [see Eq. (5.23)].

Since the reconstruction-based amplitude \hat{A} is linear in the empirical reconstruction power and the temperature-based amplitude \hat{A}' is linear in the empirical lensed temperature power, the covariance of \hat{A} and \hat{A}' involves the temperature-lensing power covariance that we computed earlier,

$$\begin{aligned} \text{cov}(\hat{A}, \hat{A}') \\ = \sigma_A^2 \sigma_{A'}^2 \sum_{l,l',l''} C_l^{\phi\phi} (\text{cov}_{\hat{\phi}\hat{\phi}}^{-1})_{ll'} \text{cov}(\hat{C}_{l'}^{\hat{\phi}\hat{\phi}}, \hat{C}_{l'',\text{expt}}^{\tilde{T}\tilde{T}}) (\text{cov}_{\tilde{T}\tilde{T},\text{expt}}^{-1})_{l'l''} (C_{l''}^{\tilde{T}\tilde{T}} - C_{l''}^{TT}), \end{aligned} \quad (5.52)$$

where the standard deviations are¹

$$\sigma_A = \left[\sum_{l,l'} C_l^{\phi\phi} (\text{cov}_{\hat{\phi}\hat{\phi}}^{-1})_{ll'} C_{l'}^{\phi\phi} \right]^{-1/2}, \quad (5.53)$$

$$\sigma_{A'} = \left[\sum_l (C_l^{\tilde{T}\tilde{T}} - C_l^{TT})^2 (\text{cov}_{\tilde{T}\tilde{T},\text{expt}}^{-1})_{ll} \right]^{-1/2}. \quad (5.54)$$

¹If no empirical $\hat{N}^{(0)}$ subtraction is used we evaluate σ_A with non-diagonal reconstruction power auto-covariance, which gives $\sigma_A \approx 2.7\%$ if $l_{\max}^\phi \geq 500$ for our noise and beam specifications. The estimated sample standard deviation of \hat{A} from simulations is larger by a factor of up to 1.07 compared to the theoretical expectation. If the $\hat{N}^{(0)}$ subtraction is used we evaluate σ_A with diagonal reconstruction power auto-covariance, which yields $\sigma_A \approx 2.5\text{--}2.6\%$ for $l_{\max}^\phi \geq 500$. The estimated sample standard deviation from simulations is larger by a factor of at most 1.05. The modest reduction in σ_A with empirical $\hat{N}^{(0)}$ subtraction is expected given the origin of this estimator as the approximate maximum-likelihood estimator for the trispectrum (see Appendix 5.B). For the lensing amplitude A' estimated from the temperature power spectrum, we find $\sigma_{A'} \approx 3.9\%$ for $l_{\max}^T = 2002$ and our Planck-like noise model. The estimated sample standard deviation from simulations is larger by a factor of 1.01.

5.4 Temperature-lensing cross-correlation

The covariance of \hat{A} and \hat{A}' can be measured in N_{sims} simulations with

$$\widehat{\text{cov}}(\hat{A}, \hat{A}') = \frac{1}{N_{\text{sims}} - 1} \sum_{s=1}^{N_{\text{sims}}} (\hat{A}_s - \langle \hat{A} \rangle_{\text{sims}})(\hat{A}'_s - \langle \hat{A}' \rangle_{\text{sims}}), \quad (5.55)$$

where s labels different realisations. We get the correlation by dividing by the theoretical standard deviations of Eqs. (5.53) and (5.54). Approximating the spectra as Gaussian variables, the variance of the estimated covariance is $\sigma_A^2 \sigma_{A'}^2 + [\text{cov}(\hat{A}, \hat{A}')]^2$, i.e. the variance of $(\hat{A} - 1)(\hat{A}' - 1)$, divided by $N_{\text{sims}} - 1$.¹ If we ignore the (small) correlation between \hat{A} and \hat{A}' , the theoretical standard error of the measured covariance is therefore $\sigma_A \sigma_{A'} / \sqrt{N_{\text{sims}} - 1}$, i.e. the theoretical error of the estimated correlation is roughly $1/\sqrt{N_{\text{sims}} - 1} \approx 3.2\%$ for 1000 simulations.

Figure 5.9 shows that the lensing amplitude correlation measured in our simulations agrees well with the theoretical correlation of Eq. (5.52) if all contributions to the temperature-lensing power covariance are taken into account. As one of the main results of this chapter we find that the correlation is at most 7% if the realisation-dependent $\hat{N}^{(0)}$ subtraction [Eq. (5.17)] is used. Without this $\hat{N}^{(0)}$ subtraction the correlation can reach up to 10% because the disconnected noise contribution of Eq. (5.34) is not cancelled. A plausibility argument for the relatively small level of amplitude correlations is presented in Appendix 5.A. Briefly, the disconnected noise contribution is small since the temperature modes that bring most information to the reconstruction are in between acoustic peaks and troughs, but the temperature modes that influence \hat{A}' most strongly are at the peaks and troughs. Since these disjoint modes vary independently, the amplitude correlation is suppressed. The matter cosmic variance contribution to the amplitude correlation is small since the errors in the measurements of \hat{A} and \hat{A}' are dominated by cosmic variance of the temperature, not the lenses.

Figure 5.9 also illustrates the relative importance of the individual covariance contributions derived above. The dominant effect comes from the matter cosmic variance contribution [Eq. (5.106)] which induces an amplitude correlation of around 4–5% for any $l_{\text{max}}^\phi \gtrsim 100$. The disconnected noise contribution [Eq. (5.34)]

¹As a product of two approximately normally distributed variables the random variable $(\hat{A}_s - \langle \hat{A} \rangle_{\text{sims}})(\hat{A}'_s - \langle \hat{A}' \rangle_{\text{sims}})$ is not normally distributed. However, the average over $N_{\text{sims}} = 1000$ realisations is approximately normally distributed due to the central limit theorem.

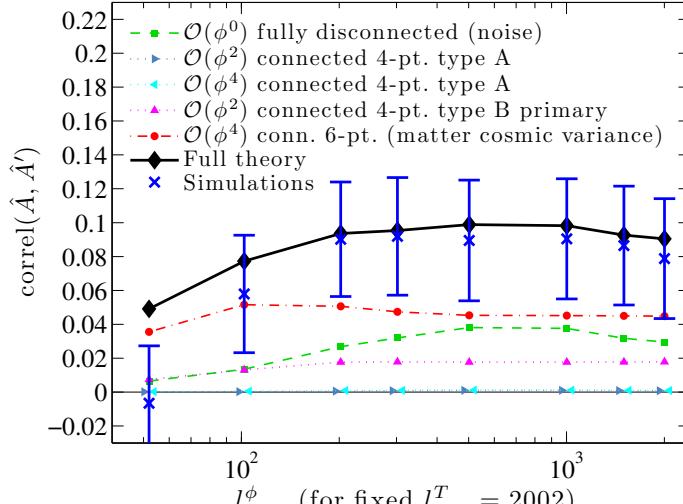
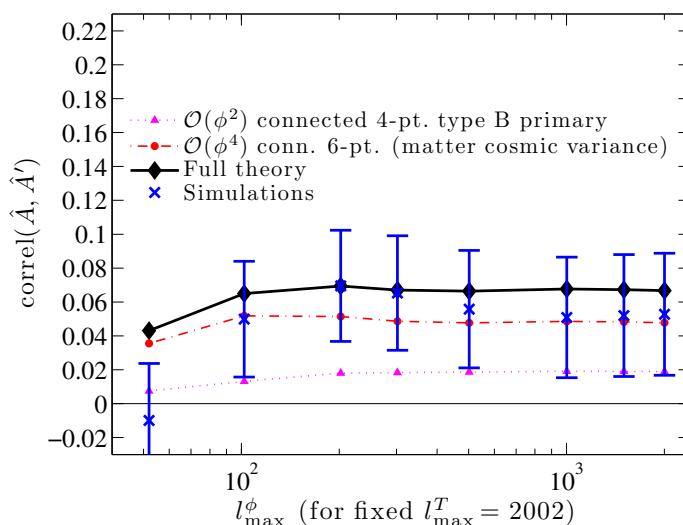

 (a) No $\hat{N}^{(0)}$ subtraction

 (b) With $\hat{N}^{(0)}$ subtraction

Figure 5.9: (a) Correlation $\widehat{\text{correl}}(\hat{A}, \hat{A}') = \widehat{\text{cov}}(\hat{A}, \hat{A}') / (\sigma_A \sigma_{A'})$ of the lensing amplitude estimates of Eqs. (5.50) and (5.51) measured in 1000 simulations (blue crosses), if no empirical $\hat{N}^{(0)}$ subtraction is applied. The theoretical correlation from Eq. (5.52) (black) includes the noise contribution of Eq. (5.34), the matter cosmic variance contribution of Eq. (5.106) and the dominant connected 4-point contributions from Eqs. (5.96) and (5.97). In the simulations the lensing potential is reconstructed from the lensed temperature power spectrum up to $l_{\max}^{\text{rec}} = 2750$. Then, \hat{A} is estimated from $\hat{C}_l^{\hat{\phi}\hat{\phi}}$ from $l_{\min}^\phi = 2$ up to the multipole l_{\max}^ϕ , which is varied along the horizontal axis. The amplitude \hat{A}' is estimated using $\hat{C}_{l,\text{expt}}^{\hat{T}\hat{T}}$ up to $l_{\max}^T = 2002$. Theoretical expressions are evaluated with the same cut-offs (non-diagonal covariance matrices are cut off before inverting them). Error bars show standard errors on the measured correlation. The errors are very correlated because they all involve overlapping low- l reconstruction modes. (b) Same as (a) but with empirical $\hat{N}^{(0)}$ subtraction. Since this removes some covariance contributions, the theoretical covariance is given by Eqs. (5.97) and (5.106) only.

implies a slightly smaller amplitude correlation of 3–4% for $l_{\max}^\phi \gtrsim 300$.¹ An additional contribution to the temperature-lensing power covariance comes from the lensed temperature trispectrum discussed in Appendix 5.D. The dominant term [Eq. (5.97)] gives rise to a 2% lensing amplitude correlation. The agreement between simulations and theory in Fig. 5.9 gives us confidence that the power correlations modeled here include all relevant contributions for amplitude measurements.

Instead of fixing $l_{\min}^\phi = 2$ and varying l_{\max}^ϕ it is worthwhile to consider the disjoint reconstruction bins $[l_{\min}^\phi, l_{\max}^\phi] = [40, 84], [85, 129], [130, 174], \dots, [355, 400]$ used for the Planck analyses in [2, 3]. If the realisation-dependent $\hat{N}^{(0)}$ is used, the theoretical correlation of the lensing amplitude estimated from one of these bins alone with the lensing amplitude estimated from the temperature power (for $l_{\max}^T = 2002$) is 5%, 3.6% and 2% for the first three reconstruction bins, and decreases further for the remaining higher- l bins. This is consistent with the correlations estimated from our 1000 simulations. In particular, this result shows that the lensing amplitude estimated from the temperature power spectrum and the reconstruction amplitudes used in the Planck lensing likelihood [3] are nearly uncorrelated, which justifies neglecting this correlation in the likelihood.

For experiments with superior noise and beam characteristics the matter cosmic variance contribution to the temperature-lensing power covariance does not change, but the lensing amplitude errors decrease. We therefore expect the corresponding amplitude correlation to increase. For example, for a full-sky experiment with SPT-like noise and beam specifications, $\sigma_N = 18 \mu\text{K}\text{arcmin}$ and $\sigma_{\text{FWHM}} = 1 \text{arcmin}$, the amplitude correlation from the matter cosmic variance contribution alone is around 10–11% for $l_{\max}^T = 2002$ and $l_{\max}^\phi = 500\text{--}1000$.

Combined lensing amplitude estimate

We have presented two estimators of the lensing amplitude: \hat{A} is linear in the reconstruction power and \hat{A}' is linear in the CMB power. These two estimates can

¹Although for the *power spectrum* cross-correlation the maximal noise contribution is about an order of magnitude larger than the maximal matter cosmic variance contribution, the latter can be more relevant for the correlation of amplitude estimates because of the phase argument given in the text.

be combined with inverse variance weighting,

$$\hat{A}_C = \frac{1}{\sigma_A^{-2} + \sigma_{A'}^{-2}} \left(\frac{\hat{A}}{\sigma_A^2} + \frac{\hat{A}'}{\sigma_{A'}^2} \right). \quad (5.56)$$

This combined estimator is the maximum-likelihood estimator for the lensing amplitude if \hat{A} and \hat{A}' are assumed to be uncorrelated. If there is a correlation between \hat{A} and \hat{A}' this does not change the expectation value of \hat{A}_C , but it does change its variance, which is then given by¹

$$\text{var}(\hat{A}_C) = \sigma_A \sigma_{A'} \left(\frac{\sigma_A}{\sigma_{A'}} + \frac{\sigma_{A'}}{\sigma_A} \right)^{-2} \left[\left(\frac{\sigma_A}{\sigma_{A'}} + \frac{\sigma_{A'}}{\sigma_A} \right) + 2 \text{correl}(\hat{A}, \hat{A}') \right]. \quad (5.57)$$

A correlation between \hat{A} and \hat{A}' therefore increases the 1σ error of the combined estimator (5.56) by a factor of

$$\frac{\sigma_{A_C}|_{\text{cov}(A,A') \neq 0}}{\sigma_{A_C}|_{\text{cov}(A,A')=0}} = \sqrt{1 + \frac{2(\sigma_A/\sigma_{A'})}{1 + (\sigma_A/\sigma_{A'})^2} \text{correl}(\hat{A}, \hat{A}')} \leq \sqrt{1 + \text{correl}(\hat{A}, \hat{A}')} \quad (5.58)$$

compared to the error if \hat{A} and \hat{A}' were uncorrelated. Since correlations between \hat{A} and \hat{A}' were found to be at most 7% if the empirical $\hat{N}^{(0)}$ subtraction is used, the 1σ error of the combined lensing estimate changes by at most 3.5% for Planck (5.5% for the full-sky SPT-like experiment mentioned above). Noting that this is the error on the error bar, the correlations between \hat{A} and \hat{A}' found above can be safely neglected when combining these two estimates of the lensing amplitude.

We briefly introduced a projection technique in Sec. 5.4.2.5 to remove the covariance between the reconstructed lensing power and the lensed temperature power spectrum due to the cosmic variance of the lenses. A simple way to perform the projection is to modify the covariance matrix $(\text{cov}_{\hat{\phi}\hat{\phi}})_{ll'}$ in Eq. (5.50) by adding $\lambda' u_l u_{l'}$, where u_l is the dominant singular vector in Eq. (5.45), and taking λ' to infinity. To the extent that the covariance between the lensing and temperature power spectra is really rank-one, this procedure removes the correlation between \hat{A} and \hat{A}' exactly. However, the variance of \hat{A} is increased by projection: it is

¹To first order in the correlation, this sampling variance of the combined \hat{A}_C is the same as the sampling variance of the optimal combined estimate that takes account of the correlations.

still given by Eq. (5.53) but with the modified $(\text{cov}_{\hat{\phi}\hat{\phi}})_{ll'}$. For $l_{\max}^\phi = 500$, we find that σ_A is increased from 0.025 to 0.042, i.e. a 70% increase. This falls to 60% for $l_{\max}^\phi = 1000$. The reason for the large increase is that u_l is rather similar in shape to the signal whose amplitude we are trying to reconstruct. Given the large increase in the error on \hat{A} , and that ignoring the effect of the covariance between the lensing reconstruction and temperature power spectra is relatively harmless, we do not advocate the use of projection to remove the correlations.

5.4.4.2 Cosmological parameters

We naively expect the impact of power correlations on cosmological parameters to be smaller than for the lensing amplitude, because the latter is directly related to the lensing potential on all scales and can therefore accumulate contributions from the full power covariances. We confirm this with a simple Fisher analysis. The covariance matrix for the joint data vector $\hat{\underline{C}} = (\hat{C}_{\text{expt}}^{\tilde{T}\tilde{T}}, \hat{C}^{\hat{\phi}\hat{\phi}} - 2\hat{N}^{(0)} + N^{(0)})$ is

$$\begin{aligned} \text{cov}_{LL',\text{joint}} &\equiv \text{cov}(\hat{C}_L, \hat{C}_{L'}) \\ &= \begin{pmatrix} \delta_{LL'} \text{var}_G(C_{L,\text{expt}}^{\tilde{T}\tilde{T}}) & \text{cov}(\hat{C}_{L,\text{expt}}^{\tilde{T}\tilde{T}}, \hat{C}_{L'}^{\hat{\phi}\hat{\phi}} - 2\hat{N}_{L'}^{(0)}) \\ \text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}} - 2\hat{N}_L^{(0)}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}}) & \delta_{LL'} \text{var}_G(\langle \hat{C}_L^{\hat{\phi}\hat{\phi}} \rangle) \end{pmatrix}. \end{aligned} \quad (5.59)$$

Fisher errors are obtained by taking the square root of the diagonal entries of the inverse of the Fisher matrix

$$F_{ij} = \sum_{LL'} \frac{\partial C_L}{\partial p_i} (\text{cov}_{\text{joint}}^{-1})_{LL'} \frac{\partial C_{L'}}{\partial p_j} \quad (5.60)$$

for cosmological parameters $\mathbf{p} = (\Omega_b h^2, \Omega_c h^2, h, \tau, A_s, n_s, \Omega_\nu h^2, \Omega_K)$ and theoretical power spectra $\underline{C} = (C^{\tilde{T}\tilde{T}}, C^{\phi\phi})$ (assuming cosmology-independent $N^{(0)}$ for simplicity). Including the off-diagonal temperature-lensing covariances of Eqs. (5.97) and (5.106) in Eq. (5.59) increases the Fisher errors for these parameters by at most 0.7% (0.5% if only Eq. (5.106) is used) compared to a completely diagonal covariance matrix.¹ The off-diagonal part of the joint covariance matrix can therefore

¹To obtain accurate derivatives for the Fisher matrix we assumed a fiducial cosmology with massive neutrinos, $\mathbf{p}_{\text{fid}} = (0.0226, 0.1123, 0.704, 0.087, 2.167 \times 10^{-9}, 0.963, 0.005, 0)$, which differs slightly from the cosmology used throughout the rest of the chapter. The Fisher errors were

fore safely be neglected for cosmological parameter estimation with a Planck-like experiment.

5.5 Towards a lensing likelihood

As argued in Sec. 2.3.4, dealing with the exact likelihood for the lensed CMB temperature is generally computationally prohibitive. For this reason, we have focussed on a form of data compression whereby the non-Gaussian lensed CMB is represented by its 2- and 4-point functions (the latter via the lensing reconstruction power spectrum). In computing the correlations between these spectra, we have implicitly been assuming that the likelihood takes the form of a multi-variate Gaussian in the spectra. In this section, we test the accuracy of this assumption in simple parameter-estimation exercises.

5.5.1 Lensing amplitude from lensing reconstruction

As a toy model, we first aim to constrain the lensing amplitude A from the lensing reconstruction alone. Considering an isotropic CMB survey, with Planck-like noise as described earlier, we consider two simple models for the likelihood, both of which depend only on the empirical power spectrum of the reconstruction. The first is the usual isotropic likelihood for a Gaussian field:

$$-2 \ln \mathcal{L}_1(\hat{\phi}|A) = \sum_l (2l+1) \left(\frac{\hat{C}_l^{\hat{\phi}\hat{\phi}}}{AC_l^{\phi\phi} + N_l} + \ln |AC_l^{\phi\phi} + N_l| \right) + (\text{const.}) . \quad (5.61)$$

This would be correct if $\hat{\phi}$ were a Gaussian field. However, since the reconstruction is manifestly non-Gaussian, we do not expect this likelihood to perform well. The

computed for $l_{\max}^T = 2002$ and $l_{\max}^\phi = 1002$.

second is Gaussian in the empirical power spectrum $\hat{C}^{\hat{\phi}\hat{\phi}}$:

$$\begin{aligned} & -2 \ln \mathcal{L}_2(\hat{C}^{\hat{\phi}\hat{\phi}} | A) \\ &= \sum_{l,l'} \left[\hat{C}_l^{\hat{\phi}\hat{\phi}} - (AC_l^{\phi\phi} + N_l) \right] (\text{cov}_{\hat{\phi}\hat{\phi}}^{-1})_{ll'} \left[\hat{C}_{l'}^{\hat{\phi}\hat{\phi}} - (AC_{l'}^{\phi\phi} + N_{l'}) \right] + (\text{const.}). \end{aligned} \quad (5.62)$$

The theoretical reconstruction power auto-covariance $\text{cov}_{\hat{\phi}\hat{\phi}}$, given by Eqs. (5.27) and (5.28), and the bias of the reconstructed lensing power, $N = N^{(0)} + N^{(1)}$, are evaluated for the fiducial amplitude $A = 1$. The empirical $\hat{N}^{(0)}$ subtraction is obtained by replacing N with $\hat{N} = 2\hat{N}^{(0)} - N^{(0)} + N^{(1)}$. For \mathcal{L}_1 , the maximum-likelihood estimate for A (given a realisation of the lensing reconstruction) is found numerically by direct evaluation of \mathcal{L}_1 for various A . The maximum likelihood estimator for A based on the second likelihood \mathcal{L}_2 is given by Eq. (5.50).

We compute estimates \hat{A} of the lensing amplitude for 1000 realisations of the lensed CMB. The sample mean of \hat{A} should be unity and the sample variance of \hat{A} , i.e. the scatter of the best-fit amplitude over different realisations, should agree with the typical width of the likelihood evaluated for a single realisation. Checking these two properties provides a non-trivial test of the likelihood \mathcal{L}_1 . In contrast, for \mathcal{L}_2 , rather than testing the accuracy of \mathcal{L}_2 , the sample mean and sample variance of \hat{A} just test our understanding of the mean and covariance of $\hat{C}^{\hat{\phi}\hat{\phi}}$.¹ This test is still useful to check for residual biases and the accuracy of our model for the reconstruction power covariance.

Figure 5.10 compares the likelihood evaluated for several individual realisations (coloured) with a Gaussian (black) with mean and standard deviation given by sample mean and standard deviation of \hat{A} averaged over all 1000 realisations (for $l_{\max}^\phi = 2650$ and our Planck-like noise model). Including the $N^{(1)}$ bias is important at high multipoles for both likelihoods: e.g. without it, \mathcal{L}_1 overestimates the lensing amplitude by 9%; see Fig. 5.10a. Including the $N^{(1)}$ bias in \mathcal{L}_1 yields the correct lensing amplitude $A = 1$ in the mean, but the scatter of \hat{A} over realisations is more

¹This is because the estimated lensing amplitude [Eq. (5.50)] is linear in $\hat{C}^{\hat{\phi}\hat{\phi}}$. For example, if the true likelihood depends on the third power of $\hat{C}^{\hat{\phi}\hat{\phi}}$, it would be possible that this only shows up in the skewness of \hat{A} . This issue will be addressed later by considering the tilt of the lensing power spectrum, which depends non-linearly on the reconstruction power.

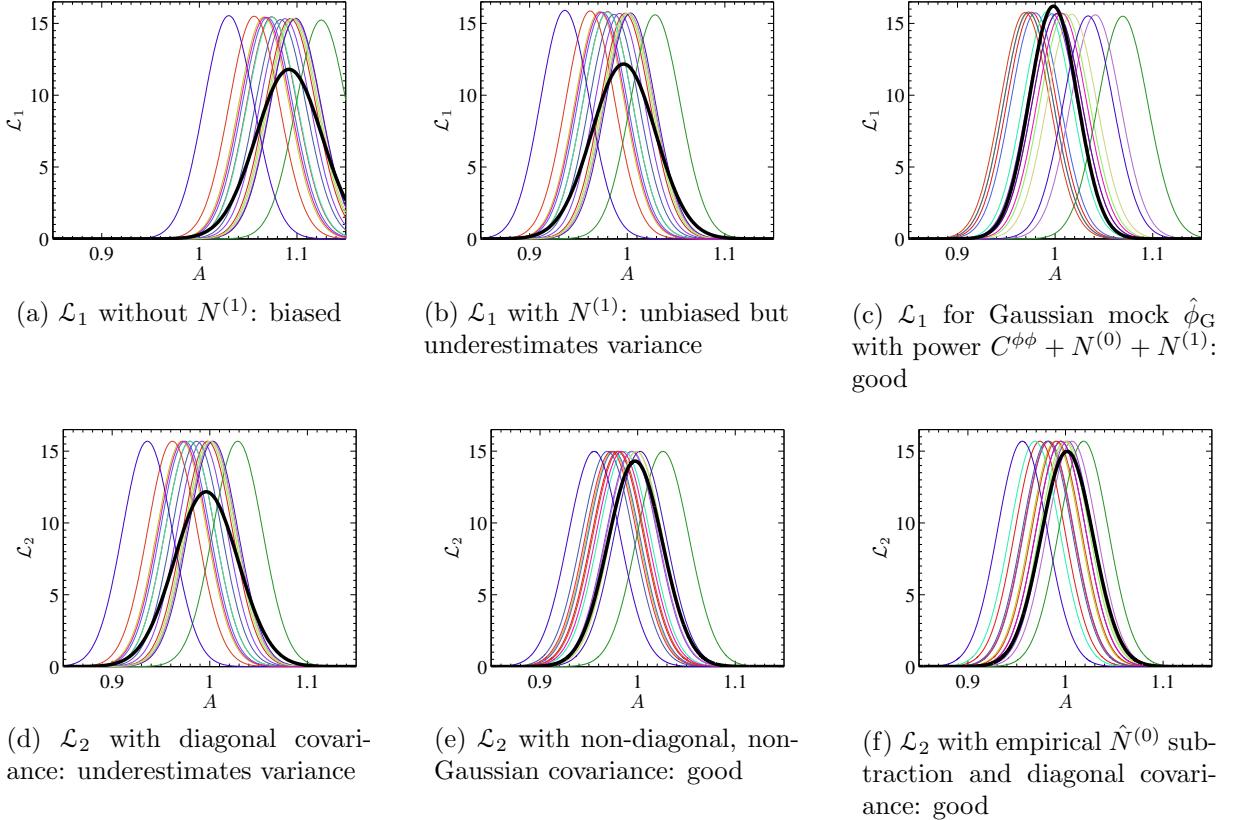


Figure 5.10: Likelihoods for the lensing amplitude for 15 different realisations (thin, coloured lines), using the lensing reconstruction up to $l_{\max}^\phi = 2650$, compared with the scatter in the best-fit amplitude obtained over 1000 simulations. Each likelihood peaks at some best-fit parameter \hat{A} , and the mean and scatter of these best-fitting amplitudes over all 1000 simulations correspond to the Gaussian curves (thick black). The upper panels show \mathcal{L}_1 from Eq. (5.61), while the lower panels show \mathcal{L}_2 given by Eq. (5.62). All curves are normalised such that their integral over A is unity.

than 30% larger than the typical width of \mathcal{L}_1 in a single realisation; see Fig. 5.10b. The likelihood \mathcal{L}_1 underpredicts the error of A because of the non-Gaussianity of $\hat{\phi}$. This is demonstrated in Fig. 5.10c, for which we replace the reconstructions with Gaussian simulations of a field with power spectrum $C_l^{\phi\phi} + N_l^{(0)} + N_l^{(1)}$. In this case, \mathcal{L}_1 should be exact and the scatter does indeed match the widths of individual realisations.

In contrast, \mathcal{L}_2 can partly model the non-Gaussianity of $\hat{\phi}$ through the non-diagonal reconstruction power auto-covariance. We compute \hat{A} based on \mathcal{L}_2 for different forms of the lensing covariance. Neglecting off-diagonal contributions to $\text{cov}_{\hat{\phi}\hat{\phi}}$ gives likelihood-based errors for A less than 80% of the scatter of \hat{A} across the simulations; see Fig. 5.10d. If we include non-Gaussian, off-diagonal contributions given by Eq. (5.27), we find that \mathcal{L}_2 predicts the scatter in A to better than 5%; see Fig. 5.10e. Similar results are achieved with the empirical $\hat{N}^{(0)}$ correction of Eq. (5.17) and diagonal (Gaussian) reconstruction power covariance; see Fig. 5.5.1.

5.5.2 Two-parameter likelihood tests with lensing amplitude and lensing tilt

To test the likelihood approximation \mathcal{L}_2 we use the lensing reconstruction additionally to constrain the lensing tilt n defined by

$$C_l^{\phi\phi} = A \left(\frac{l}{l_*} \right)^n C_l^{\phi\phi}|_{\text{fid}}. \quad (5.63)$$

The pivot multipole $l_* = 124$ is chosen such that the Fisher matrix associated with \mathcal{L}_2 is diagonal [for $l_{\max}^\phi = \mathcal{O}(10^3)$], implying that the parameters A and n are approximately uncorrelated. The likelihoods for nine realisations are compared with the scatter of the best-fit parameters over 1000 realisations in Fig. 5.11. If the non-diagonal lensing power covariance of Eq. (5.27) is included we find good agreement (without empirical $\hat{N}^{(0)}$ subtraction). Note that we have binned the reconstruction power in bins with boundaries at

$$l = 2, 13, 35, 75, 115, 155, 195, \dots \quad (\text{increasing by 40 above } l = 35). \quad (5.64)$$

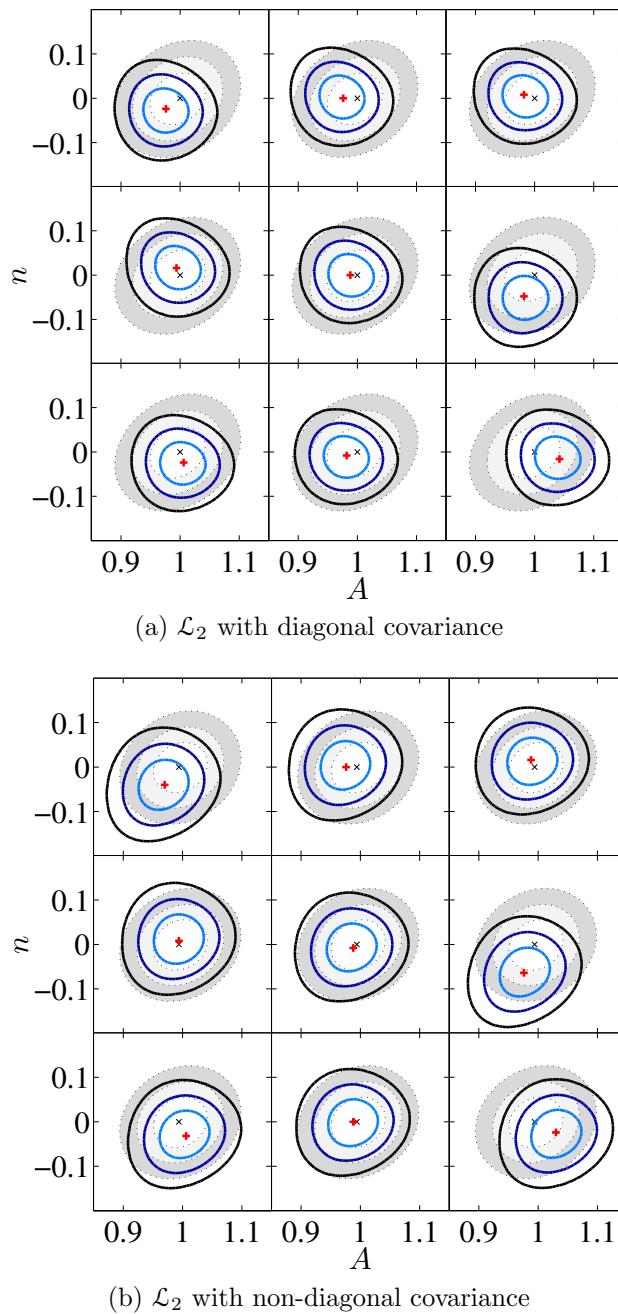


Figure 5.11: Likelihood test if lensing amplitude A (horizontal axis) and tilt n (vertical axis) are varied. *Thick lines:* Contours enclosing 68%, 95% and 99.7% of the probability for the likelihood \mathcal{L}_2 (5.62) evaluated for nine realisations. *Gray filled ellipses:* Contours of a Gaussian with central point and covariance matrix estimated from the scatter of the best-fit parameters over 1000 realisations. The lensing power auto-covariance $\text{cov}_{\hat{\phi}\hat{\phi}}$ is assumed to be diagonal in (a), while (b) also includes the non-diagonal contribution in Eq. (5.27). The reconstruction power is used up to $l_{\max}^\phi = 1002$ and is binned as described in the text.

Similar results for the unbinned case will be summarised in Fig. 5.12 below.

To quantify the agreement between the likelihoods for individual realisations and the scatter of their best-fit parameters, we compare the areas of the confidence contours shown in Fig. 5.11. We show in Fig. 5.12 the fractional deviation of the areas of the Gaussian, with sample mean and sample covariance matched to the scatter of the best-fit parameters over realisations, from the average area of the likelihoods for individual realisations (i.e. the fractional deviation of gray background areas from average areas enclosed by the solid lines in Fig. 5.11).

Neglecting the off-diagonal contribution to the lensing power covariance, which is largest at high reconstruction multipoles [59], gives narrow misshapen likelihoods that underestimate the scatter across simulations. This is particularly so for $l_{\max}^\phi \gtrsim 1000$ where the confidence areas disagree by around 40–65%. Binning does not help because it does not reduce the broad-band correlations of the reconstruction power. The agreement is better when the non-diagonal reconstruction power covariance is used (the disagreement of confidence areas is at most 14%). Alternatively, if the empirical $\hat{N}^{(0)}$ bias correction and the diagonal covariance is used, the confidence areas deviate by at most 22%. If we assume circular contours the fractional deviation of the contour radius is $\sqrt{1+d} - 1 < d/2$ if d is the fractional deviation of the contour areas. Taking this as the approximate fractional error of the marginalised error bars of A or n shows that the error on the error bars is smaller than 11% if the non-diagonal reconstruction covariance or empirical $\hat{N}^{(0)}$ subtraction are used in \mathcal{L}_2 . Therefore these two cases provide a reasonably accurate model for the lensing likelihood in this test. If diagonal reconstruction power covariance is assumed, and no empirical $\hat{N}^{(0)}$ subtraction performed, the error on the error bars can reach 30% even if binning is used. It is also worth noting that in this last case only the confidence areas *increase* with l_{\max}^ϕ , i.e. the analysis is clearly non-optimal.

Note that in the above, ideally, we should use a histogram of the best-fit parameters instead of fitting a Gaussian to their scatter. This would test the tails of the distribution better because it would include possible skewness etc. However, histograms from our 1000 simulations are too noisy to be useful for this purpose, giving results that scatter significantly with changes in histogram binning widths ΔA and Δn .

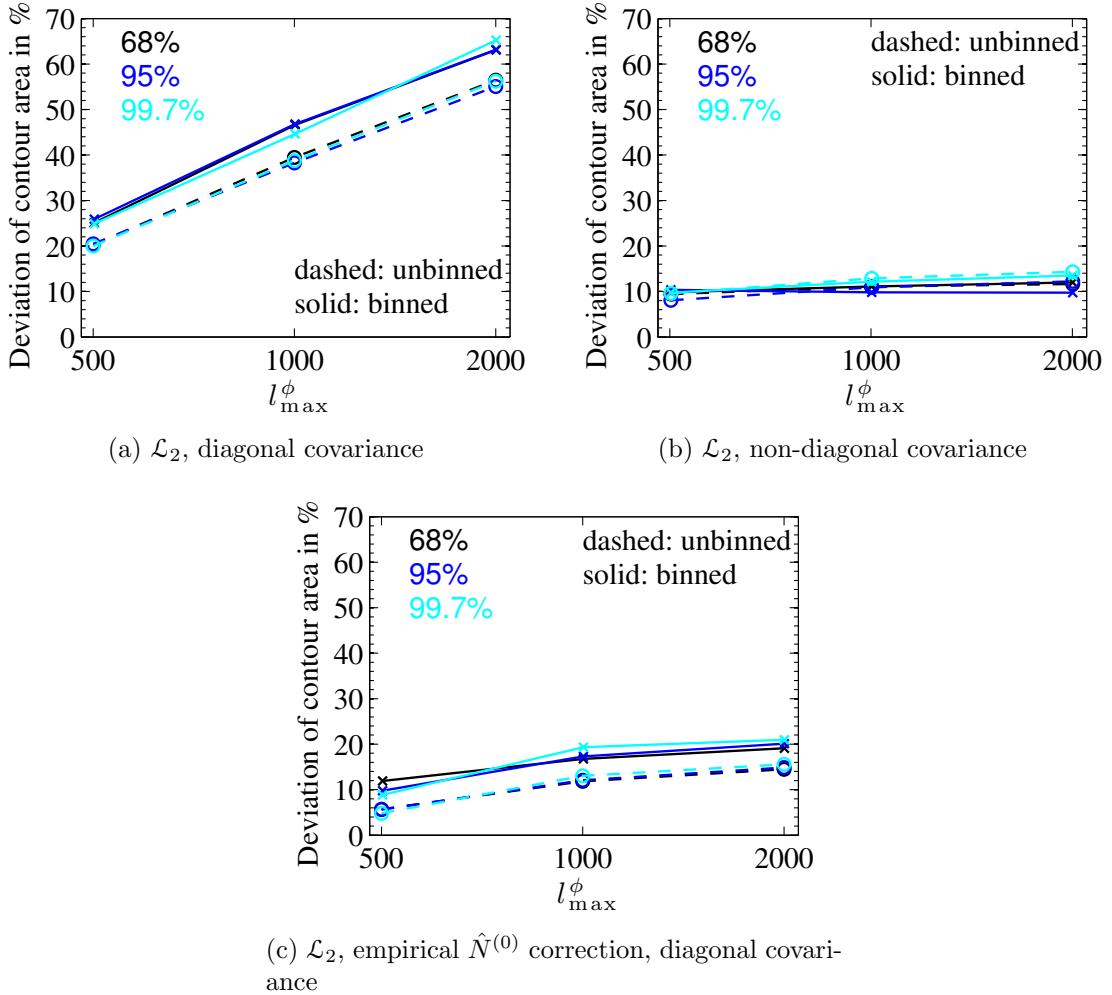


Figure 5.12: Quantitative comparison of areas in the A vs. n plane enclosed by contours of the \mathcal{L}_2 likelihood. We show the fractional deviation of the areas of Gaussians (with sample mean and sample covariance derived from the scatter of the best fit-parameters over 1000 realisations) from the average area in a single realisation. Positive values mean that the width of the likelihood typically underestimates the scatter of its peak across realisations. Results are shown for unbinned power spectra (dashed lines) and power spectra binned according to the scheme of Eq. (5.64) (solid lines); and for 68% (black), 95% (blue) and 99.7% (cyan) confidence levels. The maximum multipole of the reconstruction power is varied along the horizontal axis (only for three values which are connected by straight lines to guide the eye). The lensing power auto-covariance $\text{cov}_{\hat{\phi}\hat{\phi}}$ in \mathcal{L}_2 is assumed to be diagonal in (a), while (b) also includes the non-diagonal contribution of Eq. (5.27). Panel (c) is for empirical $\hat{N}^{(0)}$ bias correction and diagonal covariance. The contours in Fig. 5.11 correspond to the crosses at $l_{\max}^\phi = 1002$ in (a) and (b).

5.6 Conclusions

To include the CMB lensing reconstruction power spectrum in a joint likelihood analysis with the power spectrum of the temperature anisotropies requires knowledge of the cross-covariance of the two spectra. We computed this cross-covariance between the CMB 4-point and 2-point functions perturbatively, identifying two physical contributions. The disconnected part of the 6-point function of the lensed temperature leads to a noise contribution which can be interpreted as the response of the statistical noise in the lens reconstruction to fluctuations in the underlying CMB temperature field. The connected $\mathcal{O}(\phi^4)$ piece of the 6-point function gives rise to a second contribution attributable to the cosmic variance of the lenses, which causes the power spectrum of the lens reconstruction and the smoothing effect in the anisotropy power spectrum to covary. The temperature-lensing power covariance can therefore be written as

$$\begin{aligned} & \text{cov}(\hat{C}_L^{\phi\phi}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}}) \\ &= \frac{\partial(2\hat{N}_L^{(0)})}{\partial\hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}}} \frac{2}{2L'+1} (C_{L',\text{expt}}^{\tilde{T}\tilde{T}})^2 \left[1 + 2\frac{C_L^{\phi\phi}}{A_L} \right] + \frac{2}{2L+1} (C_L^{\phi\phi})^2 \frac{\partial C_{L'}^{\tilde{T}\tilde{T}}}{\partial C_L^{\phi\phi}}, \end{aligned} \quad (5.65)$$

where perturbative expressions for the derivatives are given in Eqs. (5.26) and (5.30). Both contributions were confirmed with simulations. The second term in the square brackets represents the leading correction from the connected 4-point function; see Eq. (5.97). The $C^{T\phi}$ correlation gives a diagonal contribution to the temperature-lensing power correlation, which is less than 5% for Planck-like specifications and falls rapidly with L . This generally has a negligible impact on parameter constraints derived jointly from the CMB 2- and 4-point functions.

We showed that correcting for the Gaussian $N^{(0)}$ bias in the reconstruction power with the data-dependent $\hat{N}^{(0)}$, advocated by [59] to remove auto-covariances of the lensing reconstruction power spectrum, also removes the noise contribution to the temperature-lensing power correlation and we provided an intuitive interpretation of this result.

For Planck-like specifications, estimates of the lensing amplitude A based on the lensing reconstruction or the peak smearing of the lensed temperature power

spectrum can be correlated at around the 10% level due to the power correlations. If the correlations are ignored, this gives a mis-estimate of the error on a joint amplitude estimate of only 5%, which should be negligible. The data-dependent $\hat{N}^{(0)}$ bias correction reduces the amplitude correlation further to 7% and the error of the error to 3.5%. Intuitively, we can understand the smallness of the correlation (found perturbatively and with simulations) by noting that: (i) covariance of the amplitude estimates due to cosmic variance of the lenses is limited by the small number of modes of $C^{\phi\phi}$ that influence the acoustic region of the temperature power spectrum, and is diluted significantly by CMB cosmic variance (and noise); and (ii) roughly disjoint scales in the CMB contribute to the amplitude determination from peak smearing and to the lens reconstruction limiting the correlation due to CMB cosmic variance. (See Appendix 5.A for further details of these arguments.) For a joint analysis of the power spectrum of a temperature-based CMB lensing reconstruction and the power spectrum of the temperature anisotropies themselves, the likelihoods for these two observables can therefore be simply combined for a Planck-like experiment (as was the case for the 2013 Planck analysis [3]).

Non-Gaussianity of the lensing reconstruction complicates the construction of a likelihood. We showed that the usual likelihood for isotropic Gaussian fields does not perform well for lens reconstruction in simple parameter tests, significantly underestimating the scatter seen in the best-fitting parameters across simulations. We obtained better results with simple likelihoods that are Gaussian in the measured spectra (with fiducial covariance matrix) provided that power spectrum covariances were properly modeled or data-dependent $N^{(0)}$ subtraction included. In two-parameter tests based on the amplitude and tilt of a fiducial lensing power spectrum, the widths of these Gaussian likelihoods reproduce the scatter in parameters across simulations at the 10% level.

With polarization-based reconstructions becoming feasible with current observations, it will be important to extend the analysis presented here to polarization (see [208] for work in this direction). While we expect that many of our results can be simply applied to reconstructions based on the temperature *and* polarization, the correlations are likely to be much more significant and particularly so for the most powerful *EB*-based reconstructions. We leave this to future work.

Acknowledgements

I am very grateful to my collaborators for this chapter: Anthony Challinor, Duncan Hanson and Antony Lewis. We would like to thank Uros Seljak for illuminating discussions, particularly related to Appendix 5.A.2. We also thank Helge Gruetjen, Simon Su and Oliver Zahn for helpful discussions. The numerical calculations for this chapter were performed on the COSMOS supercomputer, part of the DiRAC HPC Facility jointly funded by STFC and the Large Facilities Capital Fund of BIS. We are grateful to Andrey Kaliassin for computational support. We acknowledge use of LensPix [203] and HEALPix [204]. I thank the University of Zurich, Institut d’astrophysique de Paris and Berkeley Center for Cosmological Physics for hospitality and the opportunity to present this work.

Appendices of Chapter 5

5.A Why are the lensing amplitude cross-correlations so small?

The calculations in this chapter give a rigorous derivation of the cross-correlation of the power spectra of the lensing reconstruction and the temperature anisotropies. In this appendix we present simple physical arguments for why the correlation of the lensing amplitudes estimated from the reconstruction power and the anisotropy power are so small. We present these arguments first for the correlation due to CMB cosmic variance and then for the correlation due to cosmic variance of the lenses.

5.A.1 Cosmic variance of the CMB

Due to the smoothing effect of lensing, most of the constraint on the lensing amplitude \hat{A}' estimated from the CMB power spectrum comes from the CMB on scales of the acoustic peaks and troughs. This can be seen directly from the contribution $s_l^{\tilde{T}\tilde{T}}$ to the total signal-to-noise squared $[(S/N)^2]$ associated with the CMB cosmic variance at multipole l (see the blue curve in Fig. 5.13):

$$\left(\frac{S}{N}\right)_{A'}^2 = \frac{1}{\sigma_{A'}^2} = \sum_{l=l_{\min}^T}^{l_{\max}^T} \underbrace{\frac{(C_l^{\tilde{T}\tilde{T}} - C_l^{TT})^2}{\text{var}_G(C_{l,\text{expt}}^{\tilde{T}\tilde{T}})}}_{s_l^{\tilde{T}\tilde{T}}}. \quad (5.66)$$

In contrast, in the limit of very large-scale lenses, and as argued in more detail below, the reconstruction combines local convergence and shear measurements, for which scales in the CMB where the power spectrum changes rapidly are most informative. For large-scale lenses, the $(S/N)^2$ on the reconstruction-based amplitude estimate \hat{A} is thus expected to be dominated by CMB modes *between* acoustic peaks and troughs. Therefore, the lensing amplitude estimates \hat{A} and \hat{A}' are determined by rather disjoint CMB modes with independent CMB cosmic-variance fluctuations. We therefore expect the amplitude correlation due to CMB cosmic

5.A Why are the lensing amplitude cross-correlations so small?

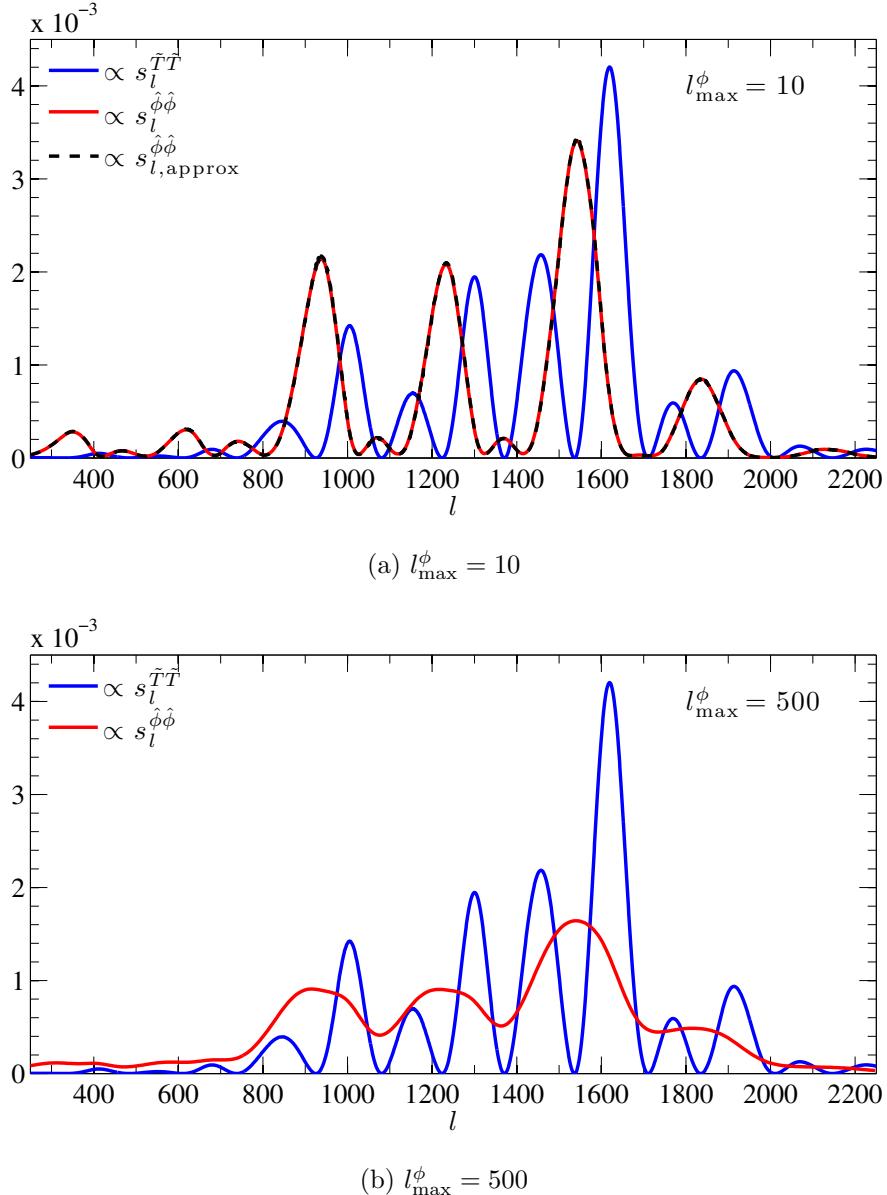


Figure 5.13: Contribution of different CMB multipoles l to the $(S/N)^2$ of the lensing amplitude for estimators based on the CMB power spectrum [blue; Eq. (5.66)] and on the lensing reconstruction [red; Eq. (5.67)] for $l_{\max}^{\phi} = 10$ (left) and $l_{\max}^{\phi} = 500$ (right). The approximation in Eq. (5.68) for large-scale lenses is shown in black dashed in the left plot. Since we are not interested in the total S/N but only in its distribution over different CMB scales, all curves are normalised such their integral over l is unity.

variance to be suppressed (in case this correlation is not mitigated by the empirical $\hat{N}^{(0)}$ subtraction anyway).

To make this point more quantitative, note that the reconstruction power $\hat{C}^{\phi\phi}$ is affected by CMB cosmic variance through the disconnected CMB 4-point contribution $N^{(0)}$. Keeping the estimator normalisation A_L and weights \tilde{g} fixed in Eq. (5.16), the contribution from the CMB at multipole l to the $(S/N)^2$ of the reconstruction-based amplitude estimate \hat{A} is monitored by

$$\begin{aligned} s_l^{\phi\phi} &= \left| \frac{\delta\sigma_A^{-2}}{\delta \ln C_{l,\text{expt}}^{\tilde{T}\tilde{T}}} \right|_{\text{fix } A_L, \tilde{g}} \\ &= \sum_{L=l_{\min}^\phi}^{l_{\max}^\phi} \left| \frac{-2(C_L^{\phi\phi})^2}{\frac{2}{2L+1}(N_L^{(0)} + C_L^{\phi\phi} + N_L^{(1)})^3} \frac{4A_L^2}{2L+1} \sum_{l_1} \tilde{g}_{l_1 l}^2(L) C_{l_1,\text{expt}}^{\tilde{T}\tilde{T}} \right| C_{l,\text{expt}}^{\tilde{T}\tilde{T}}, \end{aligned} \quad (5.67)$$

where we used σ_A from Eq. (5.53) and kept only the dominant diagonal part of the reconstruction power auto-covariance. As shown in Fig. 5.13, $s_l^{\phi\phi}$ (red) is out of phase compared to $s_l^{\tilde{T}\tilde{T}}$ (blue). To understand this structure, we separate the sums over L and l_1 in Eq. (5.67) by restricting ourselves to very large-scale lenses, $l_{\max}^\phi \ll l$, and using the large scale approximation for $N^{(0)}$ derived in [59] [see their Eq. (19)], to find, for $[l_{\max}^\phi \lesssim \mathcal{O}(10)]$,

$$s_l^{\phi\phi} \approx s_{l,\text{approx}}^{\phi\phi} = \frac{f(l_{\min}^\phi, l_{\max}^\phi)}{\text{var}_G(C_{l,\text{expt}}^{\tilde{T}\tilde{T}})} \left[\left(C_l^{\tilde{T}\tilde{T}} \frac{d \ln(l^2 C_l^{\tilde{T}\tilde{T}})}{d \ln l} \right)^2 + \frac{1}{2} \left(C_l^{\tilde{T}\tilde{T}} \frac{d \ln C_l^{\tilde{T}\tilde{T}}}{d \ln l} \right)^2 \right]. \quad (5.68)$$

Here, the prefactor f depends on the minimum and maximum reconstruction multipole but not on the CMB multipole l . The terms in square brackets have the form of the quadrature sum of the information in convergence and shear. Convergence changes locally the angular scale of the CMB anisotropies and so would contribute nothing to the $(S/N)^2$ for a scale-invariant spectrum, $l^2 C_l = \text{const.}$, while shear contributes nothing for a white-noise spectrum, $C_l = \text{const.}$ ¹ Thus,

¹The relation between large-scale lenses and the induced local convergence and shear is discussed in detail by [209], who also find agreement between the $(S/N)^2$ of a combined convergence and shear estimate with the large-scale limit of the $(S/N)^2$ of the trispectrum reconstruction.

5.A Why are the lensing amplitude cross-correlations so small?

for large-scale lenses, the $(S/N)^2$ for the reconstruction-based amplitude gets most contributions from CMB scales where the gradient of the CMB power spectrum is maximal, i.e. between acoustic peaks and troughs (see red and black curves in Fig. 5.13a).

In reality, temperature multipoles that are not precisely at peaks or troughs and not precisely in between them will affect both amplitude estimates, which implies a small amplitude correlation. Intermediate- and small-scale lenses can mix CMB modes over multipole ranges comparable to the acoustic peak separation so that they are affected by wider ranges of CMB multipoles than argued above (see red curve in Fig. 5.13b), which implies a somewhat larger amplitude correlation. However, since the CMB scales that are most important for the reconstruction still have negligible impact on the amplitude estimated from the temperature power, we expect the correlation of the amplitudes to stay rather small.

5.A.2 Cosmic variance of the lenses

We now consider the contribution of cosmic variance of the lenses to the covariance of the lensing amplitude estimates given in Eq. (5.52). It is instructive to consider a toy-model where the reconstruction “noise” power is proportional to $C_l^{\phi\phi}$, i.e. $N_l^{(0)} + N_l^{(1)} = \lambda C_l^{\phi\phi}$. Taking the limit $\lambda \rightarrow 0$ is equivalent to being able to observe ϕ directly with no measurement error, while $\lambda \rightarrow \infty$ corresponds to there being no information in the reconstruction. With $N_l^{(0)} + N_l^{(1)} = \lambda C_l^{\phi\phi}$, the weighting of the reconstruction power spectrum in \hat{A} is the same as for an ideal reconstruction (i.e. one with no $N^{(0)}$ and $N^{(1)}$ noise). Provided we then determine \hat{A} from all those ϕ modes that influence the temperature power spectrum, the contribution to the amplitude *covariance* from cosmic variance of the lenses simplifies significantly to give¹

$$\text{cov}(\hat{A}, \hat{A}') = \sigma_{A,\text{ideal}}^2. \quad (5.69)$$

Here, $\sigma_{A,\text{ideal}}^2 = [\sum_l (l + 1/2)]^{-1} \approx 2/(l_{\max}^\phi)^2$ is the variance of the reconstruction-based amplitude in the ideal limit $\lambda \rightarrow 0$ using all modes up to l_{\max}^ϕ . Since only

This correspondence has also been used to approximate the squeezed limit of the ISW-lensing bispectrum [87].

¹Note that $\sum_{l=1}^{l_{\max}^\phi} (\partial C_l^{\tilde{T}\tilde{T}} / \partial C_l^{\phi\phi}) C_l^{\phi\phi} = C_l^{\tilde{T}\tilde{T}} - C_l^{TT}$ for sufficiently large l_{\max}^ϕ .

a few (large-scale) lensing modes affect \hat{A}' , including more lensing modes in the reconstruction dilutes the covariation of \hat{A} and \hat{A}' over different realisations of the lenses, because there are increasingly more lensing modes in \hat{A} whose fluctuations do not enter \hat{A}' . The amplitude covariance falls inversely as the number of modes in the reconstruction since the weight in \hat{A} given to those (few) modes of $C_l^{\phi\phi}$ that influence the temperature power spectrum falls as the total number of modes. Note that the covariance is independent of the weighting of the measured temperature power spectrum in \hat{A}' , provided \hat{A}' is appropriately normalised, and it is also independent of additional contributions to the reconstruction noise (e.g. from CMB cosmic variance, for fixed l_{\max}^ϕ). The *variance* of \hat{A} does depend on the reconstruction noise level, with

$$\begin{aligned}\sigma_A^{-2} &\approx \sum_l \frac{2l+1}{2} \left(\frac{C_l^{\phi\phi}}{C_l^{\phi\phi} + N_l^{(0)} + N_l^{(1)}} \right)^2 \\ &= \frac{1}{(1+\lambda)^2} \sum_l \frac{2l+1}{2} = \frac{1}{(1+\lambda)^2} \sigma_{A,\text{ideal}}^{-2}.\end{aligned}\quad (5.70)$$

The result $\text{cov}(\hat{A}, \hat{A}') = \sigma_{A,\text{ideal}}^2$ for ideal weighting is necessary to ensure that the lensed CMB spectrum adds no further information on the lensing amplitude when combined with an ideal measurement of ϕ itself on all scales that are relevant for peak smearing of the temperature power spectrum. To see this, note that we can combine the amplitude estimates \hat{A} and \hat{A}' optimally into a single estimate $\hat{A}_{\text{opt.}}$, properly taking account of their correlation. If we do this, the inverse variance of the optimal estimate is given by contracting the inverse covariance matrix of the estimates:

$$\sigma_{A,\text{opt.}}^{-2} = \text{cov}^{-1}(\hat{A}, \hat{A}) + 2\text{cov}^{-1}(\hat{A}, \hat{A}') + \text{cov}^{-1}(\hat{A}', \hat{A}').\quad (5.71)$$

This evaluates to

$$\sigma_{A,\text{opt.}}^{-2} = \frac{\sigma_A^2 + \sigma_{A'}^2 - 2\sigma_{A,\text{ideal}}^2}{\sigma_A^2 \sigma_{A'}^2 - \sigma_{A,\text{ideal}}^4},\quad (5.72)$$

on using Eq. (5.69) for the covariance. In the ideal case, taking the limit $\lambda \rightarrow 0$, we have $\sigma_A = \sigma_{A,\text{ideal}}$ so that $\sigma_{A,\text{opt.}} = \sigma_{A,\text{ideal}}$. This as it must be – the observation of the peak smearing in the power spectrum adds no new information to that

5.A Why are the lensing amplitude cross-correlations so small?

obtained from the ideal measurement of ϕ . In the opposite limit, $\lambda \rightarrow \infty$, we have $\sigma_A \rightarrow \infty$ and $\sigma_{A,\text{opt.}} = \sigma_{A'}$ and all information is coming from the temperature power spectrum.

The *correlation* induced by matter cosmic variance,

$$\text{correl}(\hat{A}, \hat{A}') = \frac{\sigma_{A,\text{ideal}}^2}{\sigma_A \sigma_{A'}} = \frac{\sigma_A}{(1 + \lambda)^2 \sigma_{A'}}, \quad (5.73)$$

reaches its maximal value of $\sigma_{A,\text{ideal}}/\sigma_{A'}$ if the variance of the reconstruction power spectrum is only due to matter cosmic variance, $\lambda = (N_L^{(0)} + N_L^{(1)})/C_L^{\phi\phi} \rightarrow 0$; and it falls monotonically with increasing λ , tending to zero as $\lambda \rightarrow \infty$ (when CMB cosmic variance dominates the reconstruction uncertainty). This is expected since we assume matter and CMB fluctuations to be independent. More generally, σ_A is determined by the number of high S/N modes in the reconstruction, but $\sigma_{A'}$ depends not only on the number of CMB modes but also the fractional size of the power spectrum corrections from lensing relative to the total spectrum, $C_{l,\text{expt.}}^{\tilde{T}\tilde{T}}$. The result is that both factors $\sigma_{A,\text{ideal}}/\sigma_A$ and $\sigma_{A,\text{ideal}}/\sigma_{A'}$ in Eq. (5.73) are less than 1, diluting the amplitude correlation.

For our Planck-like parameters, the power spectrum corrections from lensing are only ever a few percent of the total spectrum and so cosmic variance of the CMB limits $\sigma_{A'} \approx 0.04$. Statistical noise in the lens reconstruction limits $\sigma_A \approx 0.025$. It is clear from Fig. 5.1 that a constant λ is not a good approximation for lens reconstruction, but we can crudely limit $\lambda \gtrsim 2$ in which case we expect the amplitude correlation to be less than $(0.025/0.04)/3^2 \sim 0.06$ which is close to the value plotted in Fig. 5.9.

To summarise, the correlation of the lensing amplitudes due to the cosmic variance of the lenses is generally small since there are a limited number of modes of $C^{\phi\phi}$ that influence the acoustic part of the temperature power spectrum (so the covariance for an ideal reconstruction scales inversely as the number of reconstruction modes), and the small covariance [less than $O(10^{-4})$] is diluted by cosmic variance of the CMB (and noise), which dominates the error on \hat{A}' and contributes significantly to the error on \hat{A} . We emphasise that these conclusions assume that the temperature power spectrum at multipoles $l \gtrsim 3000$, where the lensing-induced power from small-scale lenses dominates the unlensed power, does not influence

the amplitude estimate (i.e. the spectrum is limited by noise or foregrounds there).

5.B Optimal trispectrum estimation and $\hat{N}^{(0)}$ subtraction

We show here that the data-dependent $\hat{N}^{(0)}$ subtraction introduced in (5.17) to remove the Gaussian reconstruction bias follows naturally from optimal trispectrum estimation in the limit of weak non-Gaussianity. For simplicity we assume an idealised isotropic survey but the calculations can straightforwardly be generalised to include anisotropies from e.g. the mask by using a non-diagonal covariance $\langle \tilde{T}_{\underline{l}_1} \tilde{T}_{\underline{l}_2} \rangle$ (see [6], which also treats the more general case where the CMB temperature power is obtained by cross-correlating maps with uncorrelated noise, and shows moreover that the correlation cancellation properties are retained).

In the limit of weak non-Gaussianity, the Edgeworth expansion of the PDF for the zero-mean temperature fluctuation $\tilde{T}_{\underline{l}m}$ for zero bispectrum and non-zero trispectrum is (e.g. [54])

$$\Pr(\tilde{T}) = \frac{e^{-\tilde{T}_{\underline{l}}^2/(2C_{\underline{l}})}}{\sqrt{\det(2\pi C)}} \left\{ 1 + \frac{1}{24} \frac{\langle \tilde{T}_{\underline{l}_1} \tilde{T}_{\underline{l}_2} \tilde{T}_{\underline{l}_3} \tilde{T}_{\underline{l}_4} \rangle_c}{C_{\underline{l}_1, \text{expt}}^{\tilde{T}\tilde{T}} C_{\underline{l}_2, \text{expt}}^{\tilde{T}\tilde{T}} C_{\underline{l}_3, \text{expt}}^{\tilde{T}\tilde{T}} C_{\underline{l}_4, \text{expt}}^{\tilde{T}\tilde{T}}} \left[\tilde{T}_{\underline{l}_1} \tilde{T}_{\underline{l}_2} \tilde{T}_{\underline{l}_3} \tilde{T}_{\underline{l}_4} - (C_{\underline{l}_1 \underline{l}_2} \tilde{T}_{\underline{l}_3} \tilde{T}_{\underline{l}_4} + 5 \text{ perms}) + (C_{\underline{l}_1 \underline{l}_2} C_{\underline{l}_3 \underline{l}_4} + 2 \text{ perms}) \right] \right\}, \quad (5.74)$$

where summation over \underline{l} and \underline{l}_i is implicitly assumed and we defined

$$C_{\underline{l}_i \underline{l}_j} \equiv \langle \tilde{T}_{\underline{l}_i} \tilde{T}_{\underline{l}_j} \rangle = \delta_{\underline{l}_i \underline{l}_j} \delta_{m_i, -m_j} (-1)^{m_1} C_{\underline{l}_i, \text{expt}}^{\tilde{T}\tilde{T}}. \quad (5.75)$$

To obtain an estimator for the lensing power spectrum $C^{\phi\phi}$ we maximise (5.74) with respect to $C^{\phi\phi}$. A practical approximate solution is to take only the first step of a Newton-Raphson scheme starting from $C^{\phi\phi} = 0$, so that the estimator is proportional to $\partial \Pr(\tilde{T}) / \partial C_L^{\phi\phi}$ evaluated at $C^{\phi\phi} = 0$. Since we are only interested in the trispectrum reconstruction (and not in the effect of lensing on the CMB 2-point statistics), we consider only the impact of $C^{\phi\phi}$ on the connected 4-point function and neglect its effect on the covariance C . Then, the trispectrum-based

estimator for the lensing power involves the derivative of the connected CMB 4-point function (5.13) with respect to $C^{\phi\phi}$, which follows directly from (5.14). Since the data combination in the square brackets in (5.74) is symmetric under permutations, we only need to retain the primary coupling [i.e. the first term in (5.13)]. We therefore have (with the square brackets in the first line abbreviating the terms in square brackets in (5.74); $L > 0$)

$$\begin{aligned} \hat{C}_L^{\phi\phi} &\sim \frac{\partial \langle \tilde{T}_{l_1} \tilde{T}_{l_2} \tilde{T}_{l_3} \tilde{T}_{l_4} \rangle_c^{\text{primary}}}{\partial C_L^{\phi\phi}} [\dots] \\ &\propto \sum_{l_i, M} (-1)^M \begin{pmatrix} l_1 & l_2 & L \\ m_1 & m_2 & M \end{pmatrix} \begin{pmatrix} l_3 & l_4 & L \\ m_3 & m_4 & -M \end{pmatrix} \tilde{g}_{l_1 l_2}(L) \tilde{g}_{l_3 l_4}(L) \\ &\quad \times \left[\tilde{T}_{l_1} \tilde{T}_{l_2} \tilde{T}_{l_3} \tilde{T}_{l_4} \right. \\ &\quad \left. - (C_{l_1 l_3} \tilde{T}_{l_2} \tilde{T}_{l_4} + C_{l_1 l_4} \tilde{T}_{l_2} \tilde{T}_{l_3} + C_{l_2 l_3} \tilde{T}_{l_1} \tilde{T}_{l_4} + C_{l_2 l_4} \tilde{T}_{l_1} \tilde{T}_{l_3}) \right. \\ &\quad \left. + (C_{l_1 l_3} C_{l_2 l_4} + C_{l_1 l_4} C_{l_2 l_3}) \right] \end{aligned} \quad (5.76)$$

$$\propto \hat{C}_L^{\phi\hat{\phi}} - 2\hat{N}_L^{(0)} + N_L^{(0)}, \quad (5.77)$$

where $\hat{C}^{\phi\hat{\phi}}$ is the empirical power spectrum of the reconstruction $\hat{\phi}_{LM}$ in (5.9). Note that couplings where $l_1 = l_2$ and $m_1 = -m_2$, or $l_3 = l_4$ and $m_3 = -m_4$, do not contribute because the sum over m_1 or m_3 implies $L = 0$. From (5.77) we see that the data-dependent $\hat{N}^{(0)}$ bias subtraction indeed follows from optimal trispectrum estimation.

Conceptually, the $\hat{N}^{(0)}$ subtraction is similar to the linear term required for optimal bispectrum estimation [see Eq. (2.42)]; neither are necessary to obtain unbiased estimators but they both reduce the variance of the estimators, making them more optimal. While the linear term in the context of optimal bispectrum estimation is only required in presence of inhomogeneities (e.g. due to noise), the quadratic $\hat{N}^{(0)}$ term for optimal trispectrum estimation is required even in the case of an ideal isotropic survey.

5.C $\hat{N}^{(0)}$ subtraction as efficient mitigation of chance alignments

In this appendix we provide further details to understand why the empirical $\hat{N}^{(0)}$ bias subtraction cancels not only the non-diagonal reconstruction power autocovariance, but also the noise contribution to the temperature-lensing power cross-covariance. After identifying chance-aligned terms in the empirical reconstruction power spectrum in Eq. (5.79) below (see also [210]), we will show that avoiding these terms naturally cancels the noise contribution to the auto- and cross-covariance. We then show that empirical $\hat{N}^{(0)}$ subtraction achieves the same but in a faster way (due to the specific multipole couplings that are relevant for the covariances).

Given a CMB realisation \tilde{T}_{lm} , we split the *empirical* reconstruction power spectrum into two contributions,

$$\begin{aligned}\hat{C}_L^{\hat{\phi}\hat{\phi}} &= \frac{A_L^2}{2L+1} \sum_{l_1 \dots l_4 M} (-1)^M \begin{pmatrix} l_1 & l_2 & L \\ m_1 & m_2 & -M \end{pmatrix} \begin{pmatrix} l_3 & l_4 & L \\ m_3 & m_4 & M \end{pmatrix} \\ &\quad \times \tilde{g}_{l_1 l_2}(L) \tilde{g}_{l_3 l_4}(L) \tilde{T}_{l_1} \tilde{T}_{l_2} \tilde{T}_{l_3} \tilde{T}_{l_4} \\ &= \hat{C}_{L,\text{noise}}^{\hat{\phi}\hat{\phi}} + \hat{C}_{L,\text{rest}}^{\hat{\phi}\hat{\phi}},\end{aligned}\tag{5.78}$$

where the noise term contains the part of the sum over l_i where CMB multipoles are chance-aligned,¹

$$\begin{aligned}\hat{C}_{L,\text{noise}}^{\hat{\phi}\hat{\phi}} &= \frac{2A_L^2}{2L+1} \sum_{l_1 l_2 m_3 m_4 M} (-1)^M \begin{pmatrix} l_1 & l_2 & L \\ m_1 & m_2 & -M \end{pmatrix} \begin{pmatrix} l_1 & l_2 & L \\ m_3 & m_4 & M \end{pmatrix} \\ &\quad \times \tilde{g}_{l_1 l_2}^2(L) s_{l_1 l_2} \tilde{T}_{l_1 m_1} \tilde{T}_{l_2 m_2} \tilde{T}_{l_1 m_3} \tilde{T}_{l_2 m_4},\end{aligned}\tag{5.79}$$

taking only half of the terms for $l_1 = l_2$ by defining $s_{l_1 l_2} = 1 - \delta_{l_1 l_2}/2$. The

¹It may be possible to modify the specific splitting of Eq. (5.78) chosen here without significantly changing the conclusions below (e.g. by coupling m_3 and m_4 to m_1 and m_2). Our choice allows for a relatively simple analytical assessment of the mean and covariance properties, and is sufficiently accurate for the discussion here.

5.C $\hat{N}^{(0)}$ subtraction as efficient mitigation of chance alignments

expectation values are as desired, at sub-percent level accuracy¹,

$$\langle \hat{C}_{L,\text{noise}}^{\hat{\phi}\hat{\phi}} \rangle \approx N_L^{(0)}, \quad \langle \hat{C}_{L,\text{rest}}^{\hat{\phi}\hat{\phi}} \rangle \approx C_L^{\phi\phi} + N_L^{(1)}, \quad (5.85)$$

To assess the covariance properties of the noise-corrected reconstruction power spectrum $\hat{C}^{\hat{\phi}\hat{\phi}} - \hat{C}_{\text{noise}}^{\hat{\phi}\hat{\phi}}$, note that the $\mathcal{O}(\phi^0)$ noise contribution (5.34) to the temperature-lensing power cross-covariance is sourced by contractions

$$\tilde{T}_{l_1} \tilde{T}_{l_3} \tilde{T}_{l_2} \tilde{T}_{l_4} \quad | \quad \tilde{T}_{L'M'} \tilde{T}_{L',-M'}, \quad (5.86)$$

which imply $l_1 = l_3$ and $l_2 = l_4$. Since all terms with coinciding multipoles of this form are contained in the noise term (5.79), this covariance is cancelled if $\hat{C}^{\hat{\phi}\hat{\phi}} - \hat{C}_{\text{noise}}^{\hat{\phi}\hat{\phi}}$ is used instead of $\hat{C}^{\hat{\phi}\hat{\phi}}$. To see this explicitly note that

$$\text{cov}(\hat{C}_{L,\text{noise}}^{\hat{\phi}\hat{\phi}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})_{\text{disconn.}} = \text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})_{\text{disconn.}}. \quad (5.87)$$

¹The disconnected part is $\langle \hat{C}_{L,\text{noise}}^{\hat{\phi}\hat{\phi}} \rangle_{\text{disconn.}} = N_L^{(0)}$. Using Eq. (5.14) and simplifying the product of a 3j- and a 6j-symbol [211], we have

$$\begin{aligned} \langle \hat{C}_{L,\text{noise}}^{\hat{\phi}\hat{\phi}} \rangle_{\text{conn.}} &= \frac{2A_L^2}{(2L+1)^2} \sum_{l_1 l_2} s_{l_1 l_2} \tilde{g}_{l_1 l_2}^2(L) \tilde{f}_{l_1 L l_2}^2 C_L^{\phi\phi} + \frac{A_L^2}{16(2L+1)} \sum_{m' l_1 l_2 L'} \frac{(-1)^{m'} C_{L'}^{\phi\phi} s_{l_1 l_2}}{(C_{l_1,\text{expt}}^{\tilde{T}\tilde{T}} C_{l_2,\text{expt}}^{\tilde{T}\tilde{T}})^2} \\ &\times \left[(\tilde{f}_{l_1 L l_2}^{(u)})^2 \tilde{f}_{l_1 L' l_2}^{(u)} \tilde{f}_{l_1 L' l_2}^{\text{no3j}} \mathcal{G}_{0-m'm'}^{l_1 L l_2} \mathcal{G}_{-m'm'}^{l_1 L l_2} \mathcal{G}_{m'0-m'}^{l_1 L' l_2} \right. \\ &\quad \left. + \tilde{f}_{l_1 L l_2}^{\text{no3j}} \tilde{f}_{l_1 L l_2}^{(u)} \tilde{f}_{l_1 L' l_2}^{(u)} \tilde{f}_{l_1 L' l_2}^{(u)} \mathcal{G}_{-m'0m'}^{l_1 L l_2} \mathcal{G}_{0-m'm'}^{l_1 L' l_1} \mathcal{G}_{-m'm'0}^{l_2 L' l_2} \right], \end{aligned} \quad (5.80)$$

where we defined the Gaunt coefficients, $\tilde{f}^{(u)}$ with unsymmetric terms and \tilde{f}^{no3j} with no 3j-symbols by

$$\mathcal{G}_{m_1 m_2 m_3}^{l_1 l_2 l_3} \equiv \sqrt{\frac{(2l_1+1)(2l_2+1)(2l_3+1)}{4\pi}} \begin{pmatrix} l_1 & l_2 & l_3 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} l_1 & l_2 & l_3 \\ m_1 & m_2 & m_3 \end{pmatrix}, \quad (5.81)$$

$$\tilde{f}_{l_1 L l_2}^{(u)} \equiv [L(L+1) - l_1(l_1+1) + l_2(l_2+1)] C_{l_2}^{\tilde{T}\tilde{T}} + (l_1 \leftrightarrow l_2), \quad (5.82)$$

$$\tilde{f}_{l_1 L l_2}^{\text{no3j}} \equiv \sqrt{\frac{(2l_1+1)(2L+1)(2l_2+1)}{16\pi}} \tilde{f}_{l_1 L l_2}^{(u)}. \quad (5.83)$$

We find that (5.80) is smaller than $C_L^{\phi\phi} + N_L^{(1)}$ by a factor of more than 400 for any L , so that [with Eq. (5.15); $q_L \lesssim 0.002$]

$$\langle \hat{C}_{L,\text{noise}}^{\hat{\phi}\hat{\phi}} \rangle = N_L^{(0)} + q_L(C_L^{\phi\phi} + N_L^{(1)}), \quad \langle \hat{C}_{L,\text{rest}}^{\hat{\phi}\hat{\phi}} \rangle = (1-q_L)(C_L^{\phi\phi} + N_L^{(1)}). \quad (5.84)$$

The dominant non-diagonal contribution of Eq. (5.27) to the auto-covariance of the uncorrected reconstruction power $\hat{C}^{\hat{\phi}\hat{\phi}}$ is due to couplings [see Eq. (46c) in Ref. [59]]

$$\tilde{T}_{l_1} \tilde{T}_{l_3} \tilde{T}_{l_2} \tilde{T}_{l_4} \quad | \quad \tilde{T}_{l_7} \tilde{T}_{l_5} \tilde{T}_{l_6} \tilde{T}_{l_8}, \quad (5.88)$$

which implies $l_1 = l_3$ and $l_6 = l_8$. Writing out the contractions following [59] and summing over m_1 and M in Eq. (42) of [59] enforces $l_2 = l_4$, and therefore also $l_5 = l_7$. Since all terms of the reconstruction power (5.78) with these coinciding multipoles are in the noise term (5.79), the auto-covariance of $\hat{C}^{\hat{\phi}\hat{\phi}} - \hat{C}_{\text{noise}}^{\hat{\phi}\hat{\phi}}$ does not contain the non-diagonal contribution (5.27). In contrast, the covariance contributions that involve a product of two trispectra and lead to the dominant diagonal variance (5.28) [59] are still present, so that

$$\text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}} - \hat{C}_{L,\text{noise}}^{\hat{\phi}\hat{\phi}}, \hat{C}_{L'}^{\hat{\phi}\hat{\phi}} - \hat{C}_{L',\text{noise}}^{\hat{\phi}\hat{\phi}}) \approx \delta_{LL'} \frac{2}{2L+1} \langle C_L^{\hat{\phi}\hat{\phi}} \rangle^2. \quad (5.89)$$

Thus, the chance-aligned noise terms (5.79), which lead to the $N^{(0)}$ bias, are responsible for both the dominant non-diagonal reconstruction power auto-covariance (5.27) and the noise contribution (5.34) to the temperature-lensing power cross-covariance. It is therefore desirable to avoid these noise terms. In practice, a brute-force way to achieve this would be to subtract $\hat{C}_{\text{noise}}^{\hat{\phi}\hat{\phi}}$ from $\hat{C}^{\hat{\phi}\hat{\phi}}$ directly, or, equivalently, to restrict the summation over the l_i in (5.78) appropriately. A more efficient method is obtained from a “partial-averaging” procedure, which will turn out to be equivalent to the empirical $\hat{N}^{(0)}$ subtraction.¹ Guided by the fact that both covariances (5.27) and (5.34) are due to disconnected terms containing the contraction

$$\tilde{T}_{l_1} \tilde{T}_{l_3} \quad (5.90)$$

[see Eqs. (5.86) and (5.88)], we take the corresponding expectation value already at the level of evaluating (5.79), before computing covariances, by defining the

¹An alternative would be “phase randomisation” following [92]. Alternatively, one could split \tilde{T}_{lm} into in- and out-annuli [210], which however reduces the signal-to-noise [94]. Note that while the goal of Ref. [210] was to avoid the Gaussian $N^{(0)}$ noise bias for any CMB realisation, our goal is to simplify the reconstruction power auto- and cross-covariance.

5.D $\hat{N}^{(0)}$ subtraction as efficient mitigation of chance alignments

partial averaging operation, \mathcal{R} , by

$$\mathcal{R}[\tilde{T}_{\underline{l}_1}\tilde{T}_{\underline{l}_2}\tilde{T}_{\underline{l}_3}\tilde{T}_{\underline{l}_4}] \equiv \langle \tilde{T}_{\underline{l}_1}\tilde{T}_{\underline{l}_2} \rangle \tilde{T}_{\underline{l}_3}\tilde{T}_{\underline{l}_4} + 5 \text{ perms.}, \quad (5.91)$$

i.e. we average out two of the four modes while keeping the unaveraged realisation of the other two modes. This leaves contractions of the form (5.86) in the disconnected (4+2)-point function and contractions of the form (5.88) in the disconnected (4+4)-point function invariant in the sense that

$$\langle \mathcal{R}[\tilde{T}_{\underline{l}_1}\tilde{T}_{\underline{l}_2}\tilde{T}_{\underline{l}_3}\tilde{T}_{\underline{l}_4}]\tilde{T}_{L'M'}\tilde{T}_{L',-M'} \rangle_{\text{disconn.}}^{(5.86) \text{ terms}} = \langle \tilde{T}_{\underline{l}_1}\tilde{T}_{\underline{l}_2}\tilde{T}_{\underline{l}_3}\tilde{T}_{\underline{l}_4}\tilde{T}_{L'M'}\tilde{T}_{L',-M'} \rangle_{\text{disconn.}}^{(5.86) \text{ terms}} \quad (5.92)$$

and

$$\begin{aligned} \langle \mathcal{R}[\tilde{T}_{\underline{l}_1} \cdots \tilde{T}_{\underline{l}_4}]\mathcal{R}[\tilde{T}_{\underline{l}_5} \cdots \tilde{T}_{\underline{l}_8}] \rangle_{\text{disconn.}}^{(5.88) \text{ terms}} &= \langle \mathcal{R}[\tilde{T}_{\underline{l}_1} \cdots \tilde{T}_{\underline{l}_4}](\tilde{T}_{\underline{l}_5} \cdots \tilde{T}_{\underline{l}_8}) \rangle_{\text{disconn.}}^{(5.88) \text{ terms}} \\ &= \langle \tilde{T}_{\underline{l}_1} \cdots \tilde{T}_{\underline{l}_8} \rangle_{\text{disconn.}}^{(5.88) \text{ terms}}. \end{aligned} \quad (5.93)$$

Therefore both covariance contributions (5.27) and (5.34) are still eliminated if instead of subtracting the full noise term (5.79) from the reconstruction power we subtract the partial average of this noise term, i.e. if we consider (up to realisation-independent bias mitigation terms) $\hat{C}_L^{\hat{\phi}\hat{\phi}} - \mathcal{R}[\hat{C}_{L,\text{noise}}^{\hat{\phi}\hat{\phi}}]$. We find that this partially-averaged noise mitigation reduces to the empirical $\hat{N}^{(0)}$ subtraction,

$$\mathcal{R}[\hat{C}_{L,\text{noise}}^{\hat{\phi}\hat{\phi}}] = 2\hat{N}_L^{(0)}. \quad (5.94)$$

Thus, the empirical $\hat{N}^{(0)}$ subtraction can be interpreted as an efficient method to mitigate disconnected (auto- and cross-) covariance contributions generated by the chance-aligned noise terms (5.79) in the reconstruction power spectrum. In contrast to the noise terms (5.79), the empirical $\hat{N}^{(0)}$ defined in Eq. (5.17) can be evaluated very efficiently because the empirical temperature power spectrum is isolated in the sum.

5.D Temperature-lensing power covariance from the CMB trispectrum

We will show here that contributions from the lensed CMB trispectrum to the temperature-lensing power covariance have a sub-dominant effect on parameter estimation compared to the fully disconnected $\mathcal{O}(\phi^0)$ contribution and the $\mathcal{O}(\phi^4)$ contribution from the connected 6-point function. The connected 4-point function contributes to Eq. (5.32) with couplings of the form $\langle \tilde{T}_1 \tilde{T}_3 \rangle \langle \tilde{T}_2 \tilde{T}_4 \tilde{T}_{L'M'} \tilde{T}_{L',-M'} \rangle_c$ ('type A') and $\langle \tilde{T}_1 \tilde{T}_{L'M'} \rangle \langle \tilde{T}_2 \tilde{T}_3 \tilde{T}_4 \tilde{T}_{L',-M'} \rangle_c$ ('type B'). Other couplings either cancel in Eq. (5.32) or vanish because $\langle \hat{\phi}_{LM} \rangle = 0$. The contribution from type A can be expressed non-perturbatively in terms of the derivative in Eq. (5.30) and the connected 4-point contribution to the temperature power auto-covariance as

$$\text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})^{\text{conn.4pt.A}} = \sum_{l_2} \frac{\partial(2\hat{N}_L^{(0)})}{\partial \hat{C}_{l_2,\text{expt}}^{\tilde{T}\tilde{T}}} \text{cov}(\hat{C}_{l_2,\text{expt}}^{\tilde{T}\tilde{T}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})^{\text{conn.4pt.}}. \quad (5.95)$$

Perturbatively, up to $\mathcal{O}(\phi^4)$, we have

$$\begin{aligned} & \text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})^{\text{conn.4pt.A}} \\ &= \sum_{l_2} \frac{\partial(2\hat{N}_L^{(0)})}{\partial \hat{C}_{l_2,\text{expt}}^{\tilde{T}\tilde{T}}} \left[\text{cov}(\hat{C}_{l_2,\text{expt}}^{\tilde{T}\tilde{T}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})|_{\mathcal{O}(\phi^2)} + \text{cov}(\hat{C}_{l_2,\text{expt}}^{\tilde{T}\tilde{T}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})|_{\mathcal{O}(\phi^4)} \right], \end{aligned} \quad (5.96)$$

where the perturbative temperature covariances are given by Eq. (5.23). The covariance (5.96) can be interpreted as the correction to the noise contribution of Eq. (5.35) due to the non-diagonal $\mathcal{O}(\phi^2)$ and $\mathcal{O}(\phi^4)$ parts of the temperature power auto-covariance (5.23). The correlation corresponding to (5.96) is at most 5×10^{-5} , which is two orders of magnitude smaller than the dominant noise contribution. The induced correlation of the lensing amplitude estimates \hat{A} and \hat{A}' is less than 0.15% (see Fig. 5.9), i.e. the effect of (5.96) is negligible.

The type-A contribution to the covariance is removed by the empirical $\hat{N}^{(0)}$ correction introduced in Sec. 5.1. As for the disconnected contribution to the covariance (see Appendix 5.B), this is actually a more general result that applies

5.D Temperature-lensing power covariance from the CMB trispectrum

for anisotropic surveys and for an arbitrary quadratic estimate of the temperature power spectrum [6].

Terms of coupling type B have contributions from the primary trispectrum term, where sums over m_i simplify due to orthogonality relations of $3j$ -symbols,

$$\begin{aligned} \text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})^{\text{conn.4pt.B}}_{\text{primary}} &= 2 \frac{C_L^{\phi\phi}}{A_L} \frac{\partial(2\hat{N}_L^{(0)})}{\partial \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}}} \frac{2}{2L'+1} (C_{L',\text{expt}}^{\tilde{T}\tilde{T}})^2 \\ &= 2 \frac{C_L^{\phi\phi}}{A_L} \text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})_{\text{disconn.}}, \end{aligned} \quad (5.97)$$

and from non-primary trispectrum terms, where sums over m_i lead to a non-trivial $6j$ -symbol:

$$\begin{aligned} &\text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}})^{\text{conn.4pt.B}}_{\text{non-primary}} \\ &= \sum_{\substack{m' L'' \\ l_2 l_3 l_4}} \frac{A_L^2 (-1)^{m'} C_{L''}^{\phi\phi} \tilde{f}_{L'Ll_2}^{\text{no3j}} \tilde{f}_{l_3 Ll_4}^{(u)} \tilde{f}_{L'L''l_3}^{(u)} \tilde{f}_{l_2 L''l_4}^{(u)} G_{-m'0m'}^{l_3 Ll_4} G_{0-m'm'}^{L'L''l_3} G_{0m'-m'}^{l_2 L''l_4}}{4(2L+1)(2L'+1) C_{l_2,\text{expt}}^{\tilde{T}\tilde{T}} C_{l_3,\text{expt}}^{\tilde{T}\tilde{T}} C_{l_4,\text{expt}}^{\tilde{T}\tilde{T}}} + (l_3 \leftrightarrow l_4). \end{aligned} \quad (5.98)$$

Here we used Eq. (5.14) and expressed the product of a $6j$ -symbol with a $3j$ -symbol as a sum over a product of three $3j$ -symbols [211]. The \tilde{f} factors are defined in Appendix 5.C. Neither contributions in Eqs. (5.97) and (5.98) are cancelled by empirical $\hat{N}^{(0)}$ subtraction.

The correlation of unbinned power spectra from the primary term of Eq. (5.97) is shown in Fig. 5.14. It is at most 0.08%, which is almost a factor of two larger than the maximum of the matter cosmic variance contribution [Eq. (5.106)] shown in Fig. 5.5a. The structure is very similar to that of the disconnected (noise) contribution [Eq. (5.34)] shown in Fig. 5.4a, but with the additional signal-to-noise factor of $C_L^{\phi\phi}/A_L$ that falls off rapidly for $L > 200$. (Recall, $A_L = N_L^{(0)}$ for our choice of optimal weights.) Therefore the contribution of the primary type-B covariance to correlations between the lensing amplitudes \hat{A} and \hat{A}' is suppressed compared to that of the disconnected noise contribution, reaching at most 2% (see Fig. 5.9). It is also suppressed compared to the matter cosmic variance contribution because \hat{A}' gives most weight to CMB modes at the acoustic peaks and troughs

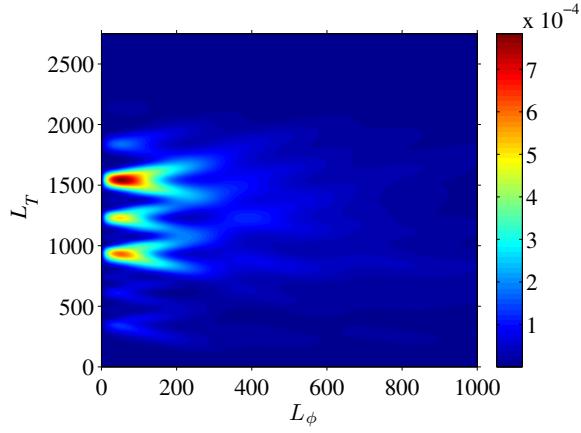


Figure 5.14: Theoretical contribution from the primary type-B connected 4-point function to the correlation of the unbinned power spectra of the lens reconstruction and the (lensed) temperature power spectrum. The covariance (5.97) is converted to a correlation using the same conversion factor as in Eq. (5.36). Correlations for $L_\phi > 1000$ are too small to be visible in this plot.

where the primary type-B power correlation is small.

While the measured correlation of \hat{A} and \hat{A}' is consistent with our theoretical expectations, the measurements are too noisy to test sub-dominant terms such as that in Eq. (5.97). In Fig. 5.4c, the lowest L_ϕ pixels show hints of the structure plotted in Fig. 5.14, but the measurements are, again, too noisy to provide a conclusive test of Eq. (5.97). The simulations do however imply that there cannot be significantly larger covariance contributions than the ones we model analytically.

Although the summation in the expression (5.98) for the non-primary type-B terms is restricted by triangle inequalities, its evaluation is still numerically challenging. We could only evaluate Eq. (5.98) for 30 (L, L') pairs at $L \lesssim 400$, which is relevant for the current Planck lensing likelihood [3]. At these points the power correlation from Eq. (5.98) is at most around 4×10^{-5} , which is an order of magnitude smaller than the maximum correlation from the primary term [Eq. (5.97)]. The structure of the correlation matrix seems similar to that of the primary term in the L_T direction, but seems to peak towards higher L_ϕ (but we cannot assess the structure reliably from the small number of evaluation points). Moreover, given that our simulations do not show any significant excess covariance we expect the non-primary term of Eq. (5.98) to be negligible.

5.E Temperature-lensing power covariance from the connected CMB 6-point function

The $\mathcal{O}(\phi^2)$ contribution to the connected part of the lensed temperature 6-point function vanishes [207]. Here we will compute the $\mathcal{O}(\phi^4)$ terms. There are five types of terms, which involve $\delta^4 T$, $\delta^3 T \delta T$, $\delta^2 T \delta^2 T$, $\delta^2 T \delta T \delta T$ or $\delta T \delta T \delta T \delta T$ with the appropriate number of factors of the unlensed CMB. Only the last two types of terms contribute to the connected 6-point function; the remaining terms contribute to the full six-point function but are cancelled when the connected 4-point and fully disconnected parts are removed.

Writing out only the two relevant types of terms, the full 6-point function at $\mathcal{O}(\phi^4)$ is ($\tilde{T}_i \equiv \tilde{T}_{l_i m_i}$)

$$\begin{aligned} \langle \tilde{T}_1 \tilde{T}_2 \tilde{T}_3 \tilde{T}_4 \tilde{T}_5 \tilde{T}_6 \rangle^{(4)} &= \frac{1}{12} [\langle \delta^2 T_1 \delta T_2 \delta T_3 T_4 T_5 T_6 \rangle + \text{all perms}] \\ &\quad + \frac{1}{48} [\langle \delta T_1 \delta T_2 \delta T_3 \delta T_4 T_5 T_6 \rangle + \text{all perms}]. \end{aligned} \quad (5.99)$$

From this we must subtract the following terms involving the connected 4-point function,

$$\begin{aligned} &\frac{1}{48} [\langle \tilde{T}_1 \tilde{T}_2 \tilde{T}_3 \tilde{T}_4 \rangle_c \langle \tilde{T}_5 \tilde{T}_6 \rangle + \text{all perms}] \\ &= \frac{1}{48} [6 \langle \delta T_1 \delta T_2 T_3 T_4 \rangle \langle \delta T_5 \delta T_6 \rangle + 12 \langle \delta T_1 \delta T_2 T_3 T_4 \rangle \langle \delta^2 T_5 T_6 \rangle \\ &\quad + 12 \langle \delta^2 T_1 \delta T_2 \delta T_3 T_4 \rangle \langle T_5 T_6 \rangle + \langle \delta T_1 \delta T_2 \delta T_3 \delta T_4 \rangle \langle T_5 T_6 \rangle \\ &\quad - (9 \langle \delta T_1 \delta T_2 \rangle + 24 \langle \delta^2 T_1 T_2 \rangle) \langle \delta T_3 \delta T_4 \rangle \langle T_5 T_6 \rangle + \text{all perms}] + \dots, \end{aligned} \quad (5.100)$$

and the fully disconnected part,

$$\begin{aligned} \langle \tilde{T}_1 \tilde{T}_2 \tilde{T}_3 \tilde{T}_4 \tilde{T}_5 \tilde{T}_6 \rangle_{\text{disc}}^{(4)} &= \frac{1}{2} [\langle \delta^2 T_1 \delta T_2 \rangle \langle \delta T_3 T_4 \rangle \langle T_5 T_6 \rangle + \text{all perms}] \\ &\quad + \frac{1}{4} [\langle \delta^2 T_1 T_2 \rangle \langle \delta T_3 \delta T_4 \rangle \langle T_5 T_6 \rangle + \text{all perms}] \\ &\quad + \frac{1}{2} [\langle \delta^2 T_1 T_2 \rangle \langle \delta T_3 T_4 \rangle \langle \delta T_5 T_6 \rangle + \text{all perms}] \\ &\quad + \frac{1}{16} [\langle \delta T_1 \delta T_2 \rangle \langle \delta T_3 \delta T_4 \rangle \langle T_5 T_6 \rangle + \text{all perms}] \\ &\quad + \frac{1}{4} [\langle \delta T_1 \delta T_2 \rangle \langle \delta T_3 T_4 \rangle \langle \delta T_5 T_6 \rangle + \text{all perms}] + \dots. \end{aligned} \quad (5.101)$$

Chapter 5

To evaluate these expressions, we use, for example,

$$\begin{aligned}\langle \delta T_1 \delta T_2 T_3 T_4 \rangle &= \langle \delta T_1 \delta T_2 \rangle \langle T_3 T_4 \rangle + \langle \langle \delta T_1 T_3 \rangle_{\text{CMB}} \langle \delta T_2 T_4 \rangle_{\text{CMB}} \rangle_{\text{LSS}} \\ &\quad + \langle \langle \delta T_1 T_4 \rangle_{\text{CMB}} \langle \delta T_2 T_3 \rangle_{\text{CMB}} \rangle_{\text{LSS}},\end{aligned}\tag{5.102}$$

where $\langle \rangle_{\text{CMB}}$ denotes averaging over unlensed CMB realisations and $\langle \rangle_{\text{LSS}}$ denotes averaging over realisations of the lensing potential ϕ . Rewriting the other 4- and 6-point functions in a similar way gives the following final expression for the connected 6-point function:

$$\begin{aligned}\langle \tilde{T}_1 \tilde{T}_2 \tilde{T}_3 \tilde{T}_4 \tilde{T}_5 \tilde{T}_6 \rangle_c^{(4)} &= \frac{1}{2} \text{cov}_{\text{LSS}} (\langle \delta^2 T_1 T_2 \rangle_{\text{CMB}}, \langle \delta T_3 T_4 \rangle_{\text{CMB}} \langle \delta T_5 T_6 \rangle_{\text{CMB}}) + \text{all perms} \\ &\quad + \frac{1}{4} \text{cov}_{\text{LSS}} (\langle \delta T_1 \delta T_2 \rangle_{\text{CMB}}, \langle \delta T_3 T_4 \rangle_{\text{CMB}} \langle \delta T_5 T_6 \rangle_{\text{CMB}}) + \text{all perms} \\ &= \frac{1}{16} \text{cov}_{\text{LSS}} \left(\langle \tilde{T}_1 \tilde{T}_2 \rangle_{\text{CMB}}^{(2)}, \langle \tilde{T}_3 \tilde{T}_4 \rangle_{\text{CMB}}^{(1)} \langle \tilde{T}_5 \tilde{T}_6 \rangle_{\text{CMB}}^{(1)} \right) + \text{all perms.}\end{aligned}\tag{5.103}$$

The result is written in terms of the covariance over different realisations of the lensing potential, $\text{cov}_{\text{LSS}}(X, Y) \equiv \langle XY \rangle_{\text{LSS}} - \langle X \rangle_{\text{LSS}} \langle Y \rangle_{\text{LSS}}$.

We expect only a few of the 720 permutations in Eq. (5.103) to be relevant for the lensing-temperature power covariance [Eq. (5.32)]. The weights \tilde{g} in Eq. (5.32) impose triangle conditions on l_1, l_2, L and l_3, l_4, L , which constrain the summation volume (especially at low L where the lensing power dominates). Terms that couple e.g. $\langle \tilde{T}_1 \tilde{T}_3 \rangle_{\text{CMB}}$ impose an additional triangle constraint on l_1, l_3 and a lensing multipole \tilde{L} , which reduces the summation volume further. This is not the case for couplings of the type $\langle \tilde{T}_1 \tilde{T}_2 \rangle_{\text{CMB}}$ and $\langle \tilde{T}_3 \tilde{T}_4 \rangle_{\text{CMB}}$. For this reason we expect the dominant terms to come from the couplings 12, 34 and $(L', M'; L', -M')$ in (5.32), i.e. terms which factor most under the weights [59]. The dominant contribution of the connected 6-point function at $\mathcal{O}(\phi^4)$ to the lensing-temperature power covariance is therefore expected to come from

$$\begin{aligned}(-1)^{M'} \langle \tilde{T}_1 \tilde{T}_2 \tilde{T}_3 \tilde{T}_4 \tilde{T}_{L'M'} \tilde{T}_{L', -M'} \rangle_{c, \text{dom}}^{(4)} &= \text{cov}_{\text{LSS}} \left(\langle \tilde{T}_{L'M'} \tilde{T}_{L'M'}^* \rangle_{\text{CMB}}^{(2)}, \langle \tilde{T}_1 \tilde{T}_2 \rangle_{\text{CMB}}^{(1)} \langle \tilde{T}_3 \tilde{T}_4 \rangle_{\text{CMB}}^{(1)} \right) \\ &\quad + 2 \text{cov}_{\text{LSS}} \left(\langle \tilde{T}_1 \tilde{T}_2 \rangle_{\text{CMB}}^{(2)}, \langle \tilde{T}_{L'M'} \tilde{T}_{L'M'}^* \rangle_{\text{CMB}}^{(1)} \langle \tilde{T}_3 \tilde{T}_4 \rangle_{\text{CMB}}^{(1)} \right).\end{aligned}\tag{5.104}$$

Here, we have used the result that permutations of the form $1 \leftrightarrow 2$, $3 \leftrightarrow 4$, $(12) \leftrightarrow (34)$ and $(L', M') \leftrightarrow (L', -M')$ all lead to the same contribution to Eq. (5.32).

The $\langle \tilde{T}_1 \tilde{T}_2 \rangle_{\text{CMB}}^{(1)}$ term on the right of Eq. (5.104) evaluates to

$$\langle \tilde{T}_1 \tilde{T}_2 \rangle_{\text{CMB}}^{(1)} = \sum_{L_1 M_1} (-1)^{M_1} \begin{pmatrix} l_1 & l_2 & L_1 \\ m_1 & m_2 & -M_1 \end{pmatrix} f_{l_1 L_1 l_2} \phi_{L_1 M_1}, \quad (5.105)$$

which combines with the weights, normalisation and one of the $3j$ symbols in Eq. (5.32) to give $\phi_{LM} + \mathcal{O}(\phi^2)$. Similarly, the $\langle \tilde{T}_3 \tilde{T}_4 \rangle_{\text{CMB}}^{(1)}$ term returns $\phi_{L, -M}$. Equation (5.105) also shows that the second term on the right of Eq. (5.104) does not contribute to the power covariance since $\langle \tilde{T}_{L'M'} \tilde{T}_{L'M'}^* \rangle_{\text{CMB}}^{(1)}$ returns the monopole of ϕ on summing over M' .

Putting these pieces together, and using Eq. (5.43), finally gives

$$\begin{aligned} \text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L'}^{\tilde{T}\tilde{T}})_{\langle \tilde{T}_1 \dots \tilde{T}_6 \rangle_{c,\text{dom.}}^{(4)}} &= \sum_{L''} \frac{\partial C_{L'}^{\tilde{T}\tilde{T}}}{\partial C_{L''}^{\phi\phi}} \text{cov}_{\text{LSS}}(\hat{C}_{L''}^{\phi\phi}, \hat{C}_L^{\phi\phi}) + \mathcal{O}(\phi^5) \\ &= \frac{2}{2L+1} (C_L^{\phi\phi})^2 \frac{\partial C_{L'}^{\tilde{T}\tilde{T}}}{\partial C_L^{\phi\phi}} + \mathcal{O}(\phi^5). \end{aligned} \quad (5.106)$$

5.F Effect of $C^{T\phi}$

In this appendix we discuss the contribution of the ISW-induced large-scale $C^{T\phi}$ correlation on the temperature-lensing power covariance. We first compute corrections due to the ISW-lensing bispectrum and then the ones due to corrections of the lensed temperature power spectrum.

The lensed temperature 6-point function in the temperature-lensing power covariance of Eq. (5.32) involves the following 3-point terms:

$$\langle \tilde{T}_{l_1} \tilde{T}_{l_2} \tilde{T}_{l_3} \tilde{T}_{l_4} \tilde{T}_{L'M'} \tilde{T}_{L'-M'} \rangle|_{3\text{-pt}, C^{T\phi}} = \langle \tilde{T}_{l_1} \tilde{T}_{l_2} \tilde{T}_{l_3} \rangle^{\text{ISW}} \langle \tilde{T}_{l_4} \tilde{T}_{L'M'} \tilde{T}_{L'-M'} \rangle^{\text{ISW}} + 9 \text{ perms}, \quad (5.107)$$

where the non-perturbative ISW-lensing bispectrum is approximately given by

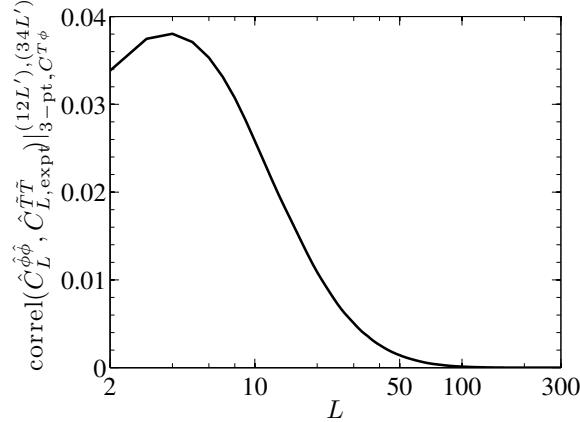


Figure 5.15: ISW-lensing bispectrum contribution of Eq. (5.109) to the diagonal unbinned temperature-lensing power correlation. The difference between the full expression and the approximation in Eq. (5.109) is too small to be visible in this plot.

[87, 212]

$$\langle \tilde{T}_{l_1} \tilde{T}_{l_2} \tilde{T}_{l_3} \rangle^{\text{ISW}} = \begin{pmatrix} l_1 & l_2 & l_3 \\ m_1 & m_2 & m_3 \end{pmatrix} B_{l_1 l_2 l_3}^{\text{ISW}} = \begin{pmatrix} l_1 & l_2 & l_3 \\ m_1 & m_2 & m_3 \end{pmatrix} C_{l_1}^{T\phi} \tilde{f}_{l_2 l_1 l_3} + 2 \text{ perms.} \quad (5.108)$$

While the contribution to Eq. (5.32) from the coupling $(123)(4L'L')$, which is written out explicitly in Eq. (5.107), vanishes on summing over M' , the couplings of type $(12L')(34L')$ yield the diagonal covariance contribution¹

$$\begin{aligned} \text{cov}(\hat{C}_L^{\hat{\phi}\hat{\phi}}, \hat{C}_{L',\text{expt}}^{\tilde{T}\tilde{T}}) |^{(12L')(34L')}_{3\text{-pt., } C^{T\phi}} &= \delta_{LL'} \frac{2A_L^2}{(2L+1)^3} \left[\sum_{l_1 l_2} \tilde{g}_{l_1 l_2}(L) B_{l_1 l_2 L}^{\text{ISW}} \right]^2 \\ &\approx \delta_{LL'} \frac{2}{2L+1} (C_L^{T\phi})^2. \end{aligned} \quad (5.109)$$

The approximation on the right is simply the covariance between the power spectra of the input lensing potential and the lensed temperature, an intuitive result

¹To obtain the approximate result in the last step we use Eq. (5.10) and neglect all terms depending on $C_{l_1}^{T\phi}$ or $C_{l_2}^{T\phi}$, because they enforce l_1 or l_2 to be small, which reduces the remaining summation volume due to the triangle condition on l_1 , l_2 and L . We find that using the approximation instead of the full expression leads a temperature-lensing correlation which is wrong by at most 3×10^{-6} for $L \leq 300$.

that we might have anticipated. For our Planck-like parameters, the corresponding power correlation is always less than 5% and rapidly decreases with increasing multipoles, being less than 0.1% for $L \geq 60$ (see Fig. 5.15). The induced correlation of the lensing amplitudes A and A' is less than 10^{-5} and therefore negligible. This small correlation arises because most of the information on the lensing amplitude from the temperature power spectrum comes from small scales. The power correlation due to the ISW-lensing effect is small on these scales, and, additionally, there is limited information in the lensing reconstruction on such scales as the signal-to-noise is very low there. We expect that couplings of the type $(13L')(24L')$ in Eq. (5.107), that do not factor under the weights, are further suppressed as they limit the summation volume.

The ISW-induced change in the lensed temperature power spectrum,

$$\Delta C_L^{\tilde{T}\tilde{T}} = C_L^{\tilde{T}\tilde{T}} - C_L^{\tilde{T}\tilde{T}}|_{C^{T\phi}=0} = \frac{1}{2L+1} \sum_{L_1, L_2} F_{LL_1L_2}^2 C_{L_1}^{T\phi} C_{L_2}^{T\phi}, \quad (5.110)$$

leads to a correction of the disconnected noise contribution [Eq. (5.34)] to the temperature-lensing power covariance which is also second order in $C^{T\phi}$. The correction has a similar structure to the noise contribution itself (shown in Fig. 5.4a), but it is around 10^6 times smaller and induces a lensing amplitude correlation of $\mathcal{O}(10^{-9})$ that is totally negligible.

Summary and outlook

Summary

Power spectra of CMB temperature anisotropies and of galaxy clustering have been crucial for pinning down today’s standard Λ CDM model. Current observations set the stage for analysing non-Gaussian statistics beyond the 2-point statistics like the bispectrum and trispectrum. This provides a window into the physics of the very early universe and allows us to disentangle effects that are degenerate at the level of 2-point statistics. In this thesis we have studied several non-Gaussian aspects of large-scale structure and the CMB.

In Chapter 3 we presented the implementation of a fast method to create general non-Gaussian initial conditions for N -body simulations to study structure formation for general models of inflation. The method works for arbitrary primordial bispectra, and a generic class of primordial trispectra. Using an expansion of the primordial bispectrum and trispectrum in a truncated set of separable basis functions, the method is several orders-of-magnitude faster than previous approaches. Even very large N -body simulations required for current and future analyses of galaxy surveys can therefore be started from general non-Gaussian initial conditions to study their impact on observables. This should prove useful in the near future for constraining primordial non-Gaussianity (PNG) beyond the typically considered linear-regime local template from large-scale structure.

In Chapter 4, we presented the implementation of a fast method to estimate the full bispectrum of a three-dimensional distribution of objects. Again, computational efficiency is achieved by a separable mode expansion of the bispectrum, which requires only $\mathcal{O}(100N^3)$ instead of $\mathcal{O}(N^6)$ operations, where $N = \mathcal{O}(10^3)$ is the number of grid points per dimension. Essentially, the idea is to measure around 100 independent projections f_{NL} of separable basis bispectra and then sum up these contributions to reconstruct the full bispectrum in the data. Instead of

the (potentially very) lossy data compression associated with restricting bispectrum analyses to a subset of triangle configurations, our method automatically takes all triangle configurations into account. This comes at the expense of a restricted space of theoretical bispectra because our method is only sensitive to bispectrum shapes that are covered by the truncated separable basis. However, we find that typically only a few percent of the cumulative signal-to-noise is lost, which should be acceptable given the computational speed-up by several orders-of-magnitude compared to brute-force estimators. An additional advantage of our method is that the information of the entire 3D functional form of the bispectrum is compressed to about 100 numbers, which should be useful for modeling covariances and comparing observations against theory and simulations. Indeed, the separable-mode framework has already proved successful in its application to CMB data, for which it was originally introduced [20, 28, 49].

We applied the separable 3D bispectrum estimator to track the dark matter bispectrum in a large suite of Gaussian and non-Gaussian N -body simulations (which were started from initial conditions obtained from the method implemented in Chapter 3). As expected, the perturbative description breaks down at late times when filaments and halos form, which affect the functional form of the dark matter bispectrum significantly. Using the cumulative signal-to-noise and full 3D shape correlations, we presented a quantitative analysis of the time evolution of the full dark matter bispectrum down to small scales. In a simple model to describe the effect of PNG in the non-linear regime we assume that (positive) PNG leads to slightly earlier gravitational collapse compared to a Gaussian universe due to enhanced initial peaks. Based on this ‘time-shift’ model we generated fitting formulae for dark matter bispectra at $0 \leq z \leq 20$ for $k \leq 2h/\text{Mpc}$ for local, equilateral, flattened and orthogonal PNG. Full 3D shape correlations with the measured N -body bispectra confirm the high quality of the fitting formulae, with typical shape correlations of 95% or more.

Chapter 5 dealt with another emerging cosmological probe beyond the Gaussian power spectrum – the trispectrum of the CMB temperature anisotropies induced by gravitational lensing along the line of sight. Lensing also affects the CMB power spectrum and it was not clear before our work to what extent the CMB trispectrum provides *independent* information about the lenses. To avoid inadvertent double-

counting of lensing information in a joint analysis of the CMB power spectrum and the lensing reconstruction power spectrum (probing the CMB trispectrum), the cross-covariance between these power spectra needs to be known. We computed this cross-covariance of the CMB 2- and 4-point functions perturbatively and verified the results with simulations. The dominant effects on parameter constraints are due to two physical contributions: (i) a noise contribution, which describes the covariation of lensing reconstruction noise with cosmic variance fluctuations of the CMB; and (ii) a matter cosmic variance contribution which describes the effect that if the lensing power spectrum fluctuates high at some scale due to cosmic variance, this implies a smoother lensed CMB power spectrum. For Planck-like specifications, broad-band measures of the lensing amplitude from the CMB and lensing reconstruction power spectra can be correlated at the 10% level due the power cross-covariance (see Sec. 5.A for an intuitive order-of-magnitude estimate). This leads to a 5% mis-estimate of the error of a joint lensing amplitude estimate, which is negligible. This is a positive result in the sense that lensing information in the CMB trispectrum is almost perfectly independent of that in the lensed CMB power spectrum, maximising the science return that can be obtained from lens reconstruction. As a consequence, the Planck 2013 analysis [3] neglected this cross-covariance in their joint likelihood. We also found that the lensing likelihood can be modeled accurately as a Gaussian in the reconstruction power spectrum with parameter-independent covariance, which should include non-Gaussian corrections if the Gaussian reconstruction bias is not removed in a realisation-dependent manner.

Outlook

Applying the separable mode expansion framework of Chapters 3 and 4 to real data would be interesting given the vast amount of high-quality large-scale structure data that will be collected in the coming decade. To achieve this goal we need to add experimental realism to the dark matter simulations studied in this thesis. First, the bispectrum of simulated halo catalogues should be characterised and compared with theoretical expectations (using bias models). As a next step we will include redshift space distortions and deal with survey geometries, fore-

grounds and systematics (e.g. star contamination). Once the separable bispectrum framework will work for realistic mock catalogues and the expected signal is well characterised for Gaussian and non-Gaussian initial conditions, we can measure the bispectrum of real datasets to constrain primordial non-Gaussianity and break (cosmological and systematic) parameter degeneracies present at the power spectrum level. Finally, the analysis can be extended to the trispectrum.

With polarization-based lensing reconstruction becoming feasible with current CMB data, it will be important to extend the CMB lensing analysis of Chapter 5 to polarization. In particular, B -modes generated by primordial gravitational waves must be cleaned from lensing contamination to obtain optimal constraints on the tensor-to-scalar ratio r . Additionally, lensing reconstructions based on polarization maps will significantly improve upon pure temperature-based reconstructions, which will improve constraints on e.g. neutrino masses, curvature and bias. Understanding and quantifying correlations between CMB power spectra and temperature- and polarization-based lensing reconstructions with simulations and analytical arguments will be crucial to perform cosmological analyses of these precious datasets without inadvertently double-counting lensing information.

References

- [1] ESA Science and Technology (2013), URL http://sci.esa.int/science-e-media/img/96/Planck_gravitational_lensing_CMB_orig.jpeg. 3
- [2] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, et al., ArXiv e-prints (2013), 1303.5076. 4, 15, 35, 37, 49, 51, 52, 189
- [3] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, et al., ArXiv e-prints (2013), 1303.5077. 4, 27, 48, 50, 51, 53, 54, 176, 183, 189, 200, 216, 225
- [4] D. M. Regan, M. M. Schmittfull, E. P. S. Shellard, and J. R. Fergusson, PRD **86**, 123524 (2012), 1108.3813. 4, 22, 71, 73, 80, 81, 82, 89, 105, 109, 110, 114, 116, 117, 140, 152
- [5] M. M. Schmittfull, D. M. Regan, and E. P. S. Shellard, PRD **88**, 063512 (2013), 1207.5678. 5, 85
- [6] M. M. Schmittfull, A. Challinor, D. Hanson, and A. Lewis, Phys. Rev. D **88**, 063012 (2013), 1308.0286. 5, 155, 177, 180, 208, 215
- [7] A. Challinor, *Large-scale structure formation (part iii lecture notes)* (2010), URL http://camd05.ast.cam.ac.uk/Cosmology/COSMOLOGY_files/Structure_formation.pdf. 7, 10, 11, 12, 13, 15, 19, 27, 29, 35, 56, 60
- [8] D. Baumann, *Cosmology (part iii lecture notes)* (2013). 7, 11, 13, 15, 27, 29, 35, 56, 60
- [9] D. Baumann, *The physics of inflation* (2013), URL <http://www.damtp.cam.ac.uk/user/db275/TEACHING/INFLATION/Lectures.pdf>. 7, 10, 13, 15, 56
- [10] A. Challinor and H. Peiris, in *American Institute of Physics Conference Series*, edited by M. Novello and S. Perez (2009), vol. 1132 of *American Institute of Physics Conference Series*, pp. 86–140, 0903.5158. 7, 27, 29, 30, 31, 35
- [11] S. Dodelson, *Modern Cosmology*, Academic Press (Academic Press, 2003), ISBN 9780122191411, URL <http://books.google.co.uk/books?id=3oPRxdXJexcC>. 7, 8, 27, 56, 58
- [12] R. Durrer, *The Cosmic Microwave Background* (Cambridge University Press, 2008), ISBN 9780521847049, URL <http://books.google.co.uk/books?id=WZkN1QDOD1sc>. 7, 27, 33, 34, 49
- [13] D. Lyth and A. Liddle, *The Primordial Density Perturbation: Cosmology, Inflation and the Origin of Structure* (Cambridge University Press, 2009), ISBN 9780521828499, URL http://books.google.co.uk/books?id=Mor_eh1Yc4QC. 7,

9, 14

-
- [14] A. Liddle, *An Introduction to Modern Cosmology* (Wiley, 2003), 2nd ed., ISBN 0470848359, URL <http://books.google.com/books?vid=ISBN0470848359>. 8
 - [15] A. H. Guth, Physical Review D **23**, 347 (1981), URL <http://dx.doi.org/10.1103/PhysRevD.23.347>. 8
 - [16] N. Bartolo, E. Komatsu, S. Matarrese, and A. Riotto, Physics Reports **402**, 103 (2004), arXiv:astro-ph/0406398. 9, 10, 13, 72, 100
 - [17] X. Chen, Advances in Astronomy **2010**, 638979 (2010), 1002.1416. 9, 21, 22, 72, 100
 - [18] M. Liguori, E. Sefusatti, J. R. Fergusson, and E. P. S. Shellard, Advances in Astronomy **2010**, 980523 (2010), 1001.4707. 9, 72, 100
 - [19] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, et al., ArXiv e-prints (2013), 1303.5082. 9
 - [20] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, et al., ArXiv e-prints (2013), 1303.5084. 9, 27, 72, 86, 90, 125, 224
 - [21] M. M. Schmittfull, Part III Essay, University of Cambridge (2010). 10
 - [22] D. Seery and J. E. Lidsey, JCAP **6**, 003 (2005), arXiv:astro-ph/0503692. 16, 17, 18
 - [23] J. Maldacena, JHEP **5**, 13 (2003), arXiv:astro-ph/0210603. 16, 17, 18, 21
 - [24] D. H. Lyth and Y. Rodríguez, PRD **71**, 123508 (2005), arXiv:astro-ph/0502578. 17
 - [25] D. Seery and J. E. Lidsey, JCAP **9**, 011 (2005), arXiv:astro-ph/0506056. 18
 - [26] D. H. Lyth, C. Ungarelli, and D. Wands, PRD **67**, 023503 (2003), arXiv:astro-ph/0208055. 19
 - [27] A. Lewis, JCAP **10**, 26 (2011), 1107.5431. 19, 22
 - [28] J. R. Fergusson, M. Liguori, and E. P. S. Shellard, PRD **82**, 023502 (2010), 0912.5516. 20, 22, 73, 74, 75, 78, 86, 91, 107, 134, 153, 224
 - [29] D. S. Salopek and J. R. Bond, PRD **42**, 3936 (1990). 21
 - [30] P. Creminelli and M. Zaldarriaga, JCAP **10**, 6 (2004), arXiv:astro-ph/0407059. 21
 - [31] C. Cheung, A. L. Fitzpatrick, J. Kaplan, and L. Senatore, JCAP **2**, 21 (2008), 0709.0295. 21
 - [32] V. Assassi, D. Baumann, and D. Green, JCAP **11**, 047 (2012), 1204.4207. 21
 - [33] M. Alishahiha, E. Silverstein, and D. Tong, PRD **70**, 123505 (2004), arXiv:hep-th/0404084. 21
 - [34] C. Cheung, A. L. Fitzpatrick, J. Kaplan, L. Senatore, and P. Creminelli, Journal of High Energy Physics **3**, 014 (2008), 0709.0293. 21
 - [35] N. Bartolo, M. Fasiello, S. Matarrese, and A. Riotto, JCAP **8**, 008 (2010), 1004.0893. 21, 22

- [36] D. Baumann and D. Green, *JCAP* **9**, 14 (2011), 1102.5343. 21
- [37] P. Creminelli, M. A. Luty, A. Nicolis, and L. Senatore, *Journal of High Energy Physics* **12**, 080 (2006), arXiv:hep-th/0606090. 21
- [38] P. Creminelli, A. Nicolis, L. Senatore, M. Tegmark, and M. Zaldarriaga, *JCAP* **5**, 4 (2006), arXiv:astro-ph/0509029. 21, 106
- [39] C. Wagner and L. Verde, *JCAP* **3**, 2 (2012), 1102.3229. 21, 23, 72, 74, 89, 143
- [40] J. N. Fry, *Astrophys. J.* **279**, 499 (1984), URL <http://dx.doi.org/10.1086/161913>. 22, 92
- [41] F. Bernardeau, S. Colombi, E. Gaztañaga, and R. Scoccimarro, *Physics Reports* **367**, 1 (2002), arXiv:astro-ph/0112551. 22, 60, 62, 63, 65, 66, 67, 92, 114
- [42] L. Senatore, K. M. Smith, and M. Zaldarriaga, *JCAP* **1**, 028 (2010), 0905.3746. 22, 72
- [43] S. Renaux-Petel, *Classical and Quantum Gravity* **28**, 249601 (2011), 1105.6366. 22
- [44] S. Renaux-Petel, S. Mizuno, and K. Koyama, *JCAP* **11**, 042 (2011), 1108.0305. 22
- [45] D. Jeong and E. Komatsu, *APJ* **703**, 1230 (2009), 0904.0497. 22, 93
- [46] X. Chen, M.-x. Huang, S. Kachru, and G. Shiu, *JCAP* **1**, 2 (2007), arXiv:hep-th/0605045. 22
- [47] R. Holman and A. J. Tolley, *JCAP* **5**, 1 (2008), 0710.1302. 22
- [48] P. D. Meerburg, J. P. van der Schaar, and P. Stefano Corasaniti, *JCAP* **5**, 18 (2009), 0901.4044. 22
- [49] J. R. Fergusson and E. P. S. Shellard, *PRD* **80**, 043510 (2009), 0812.3413. 22, 73, 74, 86, 93, 97, 114, 224
- [50] J. R. Fergusson, D. M. Regan, and E. P. S. Shellard, *PRD* **86**, 063511 (2012), 1008.1730. 22, 24, 25, 26, 27, 73, 74, 78, 79, 89, 105, 106, 109, 110, 114, 116, 117, 140, 151, 152
- [51] J. R. Fergusson, M. Liguori, and E. P. S. Shellard, *JCAP* **12**, 032 (2012), 1006.1642. 22, 78, 80, 86, 106, 107
- [52] F. Schmidt and M. Kamionkowski, *PRD* **82**, 103002 (2010), 1008.0638. 23
- [53] D. Babich, *PRD* **72**, 043003 (2005), arXiv:astro-ph/0503375. 23, 24, 25
- [54] D. M. Regan, E. P. S. Shellard, and J. R. Fergusson, *PRD* **82**, 023520 (2010), 1004.2915. 24, 27, 73, 86, 114, 162, 208
- [55] L. Amendola, *Mon. Not. Roy. Astron. Soc.* **283**, 983 (1996), URL http://adsabs.harvard.edu/cgi-bin/nph-bib_query?bibcode=1996MNRAS.283..983A. 24
- [56] P. Creminelli, A. Nicolis, L. Senatore, M. Tegmark, and M. Zaldarriaga, *Journal of Cosmology and Astroparticle Physics* **0605**, 004 (2006), astro-ph/0509029. 24, 25
- [57] D. Hanson, Ph.D. thesis, University of Cambridge (2010). 25, 40, 43, 44
- [58] D. Babich, P. Creminelli, and M. Zaldarriaga, *JCAP* **8**, 9 (2004), arXiv:astro-ph/0405356. 26, 106, 114, 115

-
- [59] D. Hanson, A. Challinor, G. Efstathiou, and P. Bielewicz, PRD **83**, 043005 (2011), 1008.4403. 27, 156, 157, 158, 160, 161, 162, 163, 168, 169, 174, 176, 184, 185, 197, 199, 204, 212, 218
 - [60] E. Komatsu, D. N. Spergel, and B. D. Wandelt, APJ **634**, 14 (2005), arXiv:astro-ph/0305189. 27
 - [61] A. Mangilli, B. Wandelt, F. Elsner, and M. Liguori, Astronomy and Astrophysics **555**, A82 (2013), 1303.1722. 27
 - [62] C. L. Bennett, D. Larson, J. L. Weiland, N. Jarosik, G. Hinshaw, N. Odegard, K. M. Smith, R. S. Hill, B. Gold, M. Halpern, et al., ArXiv e-prints (2012), 1212.5225. 27
 - [63] U. Seljak and M. Zaldarriaga, APJ **469**, 437 (1996), arXiv:astro-ph/9603033. 34
 - [64] A. Lewis, A. Challinor, and A. Lasenby, APJ **538**, 473 (2000), arXiv:astro-ph/9911177. 34, 48, 91, 96, 115, 120, 121, 154, 167, 185
 - [65] J. Lesgourgues, ArXiv e-prints (2011), 1104.2932. 34
 - [66] S. Dodelson, in *Neutrinos, Flavor Physics, and Precision Cosmology*, edited by J. F. Nieves and R. R. Volkas (2003), vol. 689 of *American Institute of Physics Conference Series*, pp. 184–196, arXiv:hep-ph/0309057. 35
 - [67] ACT collaboration, *Act atacama cosmology telescope* (2013), URL <http://www.princeton.edu/act/index.xml>. 35
 - [68] SPT collaboration, *South pole telescope* (2013), URL <http://pole.uchicago.edu/>. 35
 - [69] D. M. Regan, Ph.D. thesis, University of Cambridge (2011), 1112.5899. 35
 - [70] Planck collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, et al., ArXiv e-prints (2013), 1303.5075. 37, 52
 - [71] W. Hu, PRD **65**, 023003 (2002), arXiv:astro-ph/0108090. 38, 54
 - [72] M. Kaplinghat, L. Knox, and Y.-S. Song, Physical Review Letters **91**, 241301 (2003), arXiv:astro-ph/0303344. 38, 54
 - [73] L. Verde and D. N. Spergel, PRD **65**, 043007 (2002), arXiv:astro-ph/0108179. 38
 - [74] V. Acquaviva, C. Baccigalupi, and F. Perrotta, PRD **70**, 023515 (2004), arXiv:astro-ph/0403654. 38
 - [75] E. Calabrese, A. Slosar, A. Melchiorri, G. F. Smoot, and O. Zahn, PRD **77**, 123531 (2008), 0803.2309. 38, 185
 - [76] A. C. Hall and A. Challinor, Mon. Not. R. Astron. Soc. **425**, 1170 (2012), 1205.6172. 38
 - [77] L. Santos, P. Cabella, A. Balbi, and N. Vittorio, ArXiv e-prints (2013), 1307.2919. 38
 - [78] B. Hu, M. Liguori, N. Bartolo, and S. Matarrese, ArXiv e-prints (2013), 1307.5276. 38
 - [79] A. Lewis and A. Challinor, Physics Reports **429**, 1 (2006), arXiv:astro-

- ph/0601594. 38, 39, 41, 42, 43, 156
- [80] D. Hanson, A. Challinor, and A. Lewis, General Relativity and Gravitation **42**, 2197 (2010), 0911.0612. 38, 40
- [81] R. E. Smith, J. A. Peacock, A. Jenkins, S. D. M. White, C. S. Frenk, F. R. Pearce, P. A. Thomas, G. Efstathiou, and H. M. P. Couchman, Mon. Not. R. Astron. Soc. **341**, 1311 (2003), arXiv:astro-ph/0207664. 39, 91, 120, 121
- [82] W. Hu, PRD **62**, 043007 (2000), arXiv:astro-ph/0001303. 42, 43, 158, 167
- [83] M. Zaldarriaga, PRD **62**, 063510 (2000), arXiv:astro-ph/9910498. 43
- [84] A. Challinor and A. Lewis, PRD **71**, 103010 (2005), arXiv:astro-ph/0502425. 43, 167
- [85] U. Seljak, APJ **463**, 1 (1996), arXiv:astro-ph/9505109. 43
- [86] M. Zaldarriaga and U. Seljak, PRD **58**, 023003 (1998), arXiv:astro-ph/9803150. 43
- [87] A. Lewis, A. Challinor, and D. Hanson, JCAP **3**, 018 (2011), 1101.2234. 45, 72, 159, 160, 205, 220
- [88] W. Hu, APJL **557**, L79 (2001), arXiv:astro-ph/0105424. 45, 156, 160
- [89] T. Okamoto and W. Hu, PRD **67**, 083002 (2003), arXiv:astro-ph/0301031. 45, 46, 156, 159, 160
- [90] K. M. Smith, O. Zahn, and O. Doré, PRD **76**, 043510 (2007), 0705.3980. 48, 53
- [91] C. M. Hirata, S. Ho, N. Padmanabhan, U. Seljak, and N. A. Bahcall, PRD **78**, 043520 (2008), 0801.0644. 48, 53
- [92] S. Das, B. D. Sherwin, P. Aguirre, J. W. Appel, J. R. Bond, C. S. Carvalho, M. J. Devlin, J. Dunkley, R. Dünner, T. Essinger-Hileman, et al., Physical Review Letters **107**, 021301 (2011), 1103.2124. 48, 212
- [93] S. Das, T. Louis, M. R. Nolta, G. E. Addison, E. S. Battistelli, J. Bond, E. Calabrese, D. C. M. J. Devlin, S. Dicker, J. Dunkley, et al., ArXiv e-prints (2013), 1301.1037. 48
- [94] A. van Engelen, R. Keisler, O. Zahn, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, T. M. Crawford, et al., APJ **756**, 142 (2012), 1202.0546. 48, 54, 176, 212
- [95] D. Hanson, S. Hoover, A. Crites, P. A. R. Ade, K. A. Aird, J. E. Austermann, J. A. Beall, A. N. Bender, B. A. Benson, L. E. Bleem, et al., ArXiv e-prints (2013), 1307.5830. 48
- [96] K. T. Story, C. L. Reichardt, Z. Hou, R. Keisler, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. Cho, et al., ArXiv e-prints (2012), 1210.7231. 49
- [97] R. Stompor and G. Efstathiou, Mon. Not. R. Astron. Soc. **302**, 735 (1999), arXiv:astro-ph/9805294. 51
- [98] K. M. Smith, W. Hu, and M. Kaplinghat, PRD **74**, 123002 (2006), arXiv:astro-ph/0607315. 51, 166, 179, 180

-
- [99] B. D. Sherwin, J. Dunkley, S. Das, J. W. Appel, J. R. Bond, C. S. Carvalho, M. J. Devlin, R. Dünner, T. Essinger-Hileman, J. W. Fowler, et al., Physical Review Letters **107**, 021302 (2011), 1105.0419. 51, 54, 176
- [100] K. M. Smith, A. Cooray, S. Das, O. Doré, D. Hanson, C. Hirata, M. Kaplinghat, B. Keating, M. Loverde, N. Miller, et al., in *American Institute of Physics Conference Series*, edited by S. Dodelson, D. Baumann, A. Cooray, J. Dunkley, A. Fraisse, M. G. Jackson, A. Kogut, L. Krauss, M. Zaldarriaga, and K. Smith (2009), vol. 1141 of *American Institute of Physics Conference Series*, pp. 121–178, 0811.3916. 52
- [101] M. C. Gonzalez-Garcia, M. Maltoni, J. Salvado, and T. Schwetz, Journal of High Energy Physics **12**, 123 (2012), 1209.3023. 52
- [102] L. E. Bleem, A. van Engelen, G. P. Holder, K. A. Aird, R. Armstrong, M. L. N. Ashby, M. R. Becker, B. A. Benson, T. Biesiadzinski, M. Brodwin, et al., APJL **753**, L9 (2012), 1203.4808. 53
- [103] B. D. Sherwin, S. Das, A. Hajian, G. Addison, J. R. Bond, D. Crichton, M. J. Devlin, J. Dunkley, M. B. Gralla, M. Halpern, et al., PRD **86**, 083006 (2012), 1207.4543. 53
- [104] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, et al., ArXiv e-prints (2013), 1303.5078. 53
- [105] G. P. Holder, M. P. Viero, O. Zahn, K. A. Aird, B. A. Benson, S. Bhattacharya, L. E. Bleem, J. Bock, M. Brodwin, J. E. Carlstrom, et al., APJL **771**, L16 (2013), 1303.5048. 53
- [106] A. Vallinotto, ArXiv e-prints (2013), 1304.3474. 53
- [107] J. Lesgourgues, L. Perotto, S. Pastor, and M. Piat, PRD **73**, 045021 (2006), arXiv:astro-ph/0511735. 54
- [108] D. M. Goldberg and D. N. Spergel, PRD **59**, 103002 (1999), arXiv:astro-ph/9811251. 55
- [109] C. M. Hirata and U. Seljak, PRD **67**, 043001 (2003), arXiv:astro-ph/0209489. 55, 159
- [110] R. Scoccimarro, APJ **487**, 1 (1997), arXiv:astro-ph/9612207. 56, 60, 62, 66, 92, 93, 95
- [111] R. Hlozek, J. Dunkley, G. Addison, J. W. Appel, J. R. Bond, C. Sofia Carvalho, S. Das, M. J. Devlin, R. Dünner, T. Essinger-Hileman, et al., APJ **749**, 90 (2012), 1105.4887. 59
- [112] D. Jeong, dissertation (2010), URL http://www.pha.jhu.edu/~djeong/dissertation/djeong_diss.pdf. 58, 62, 76, 120
- [113] P. J. E. Peebles, *The large-scale structure of the universe* (1980). 60
- [114] M. H. Goroff, B. Grinstein, S. Rey, and M. B. Wise, Astrophys. J. **311**, 6 (1986), URL <http://dx.doi.org/10.1086/164749>. 62
- [115] B. Jain and E. Bertschinger, Astrophys. J. **431**, 495 (1994), astro-ph/9311070,

REFERENCES

- URL <http://dx.doi.org/10.1086/174502>. 62
- [116] R. Takahashi, Progress of Theoretical Physics **120**, 549 (2008), 0806.1437. 62
- [117] T. Baldauf, U. Seljak, and L. Senatore, JCAP **4**, 006 (2011), 1011.1513. 64, 89
- [118] E. Sefusatti, M. Crocce, and V. Desjacques, Mon. Not. R. Astron. Soc. **425**, 2903 (2012), 1111.6966. 64, 89, 106, 122
- [119] J. N. Fry, Phys. Rev. Lett. **73**, 215 (1994), URL <http://link.aps.org/doi/10.1103/PhysRevLett.73.215>. 64
- [120] S. Matarrese, L. Verde, and A. F. Heavens, Mon. Not. R. Astron. Soc. **290**, 651 (1997), arXiv:astro-ph/9706059. 64
- [121] L. Verde, A. F. Heavens, S. Matarrese, and L. Moscardini, Mon. Not. R. Astron. Soc. **300**, 747 (1998), arXiv:astro-ph/9806028. 64
- [122] R. Scoccimarro, H. M. P. Couchman, and J. A. Frieman, APJ **517**, 531 (1999), arXiv:astro-ph/9808305. 64
- [123] E. Sefusatti, M. Crocce, S. Pueblas, and R. Scoccimarro, PRD **74**, 023522 (2006), arXiv:astro-ph/0604505. 64
- [124] J. N. Fry and E. Gaztanaga, APJ **413**, 447 (1993), arXiv:astro-ph/9302009. 64
- [125] T. Baldauf, U. Seljak, V. Desjacques, and P. McDonald, PRD **86**, 083540 (2012), 1201.4827. 64, 65
- [126] D. Baumann, S. Ferraro, D. Green, and K. M. Smith, JCAP **5**, 001 (2013), 1209.2173. 64
- [127] Y. B. Zel'Dovich, A& A **5**, 84 (1970). 65, 76
- [128] E. Sirkо, APJ **634**, 728 (2005), arXiv:astro-ph/0503106. 65, 76, 154
- [129] R. Scoccimarro, Monthly Notices of the Royal Astronomical Society **299**, 1097 (1998), URL <http://dx.doi.org/10.1046/j.1365-8711.1998.01845.x>. 66, 76, 116, 117
- [130] M. Crocce, S. Pueblas, and R. Scoccimarro, Mon. Not. R. Astron. Soc. **373**, 369 (2006), arXiv:astro-ph/0606505. 66, 116, 117, 154
- [131] S. D. M. White, ArXiv Astrophysics e-prints (1994), arXiv:astro-ph/9410043. 66, 119
- [132] V. Springel, Mon. Not. R. Astron. Soc. **364**, 1105 (2005), arXiv:astro-ph/0505010. 66, 67, 116, 117, 119, 154
- [133] M. Bartelmann, ArXiv Astrophysics e-prints (2003), arXiv:astro-ph/0304162. 66
- [134] V. Springel, Ph.D. thesis, PhD Thesis, Ludwig-Maximilians-Universitaet Muenchen, 1999 (1999), URL <http://www.mpa-garching.mpg.de/~volker/publications/thesis.ps.gz>. 67
- [135] M. Colless, G. Dalton, S. Maddox, W. Sutherland, P. Norberg, S. Cole, J. Bland-Hawthorn, T. Bridges, R. Cannon, C. Collins, et al., Mon. Not. R. Astron. Soc. **328**, 1039 (2001), arXiv:astro-ph/0106498. 68
- [136] D. H. Jones, M. A. Read, W. Saunders, M. Colless, T. Jarrett, Q. A. Parker, A. P. Fairall, T. Mauch, E. M. Sadler, F. G. Watson, et al., Mon. Not. R. Astron. Soc.

399, 683 (2009), 0903.5451. 68

- [137] K. N. Abazajian, J. K. Adelman-McCarthy, M. A. Agüeros, S. S. Allam, C. Allende Prieto, D. An, K. S. J. Anderson, S. F. Anderson, J. Annis, N. A. Bahcall, et al., *Astrophys. J. Supp.* **182**, 543 (2009), 0812.0649. 68
- [138] M. J. Drinkwater, R. J. Jurek, C. Blake, D. Woods, K. A. Pimbblet, K. Glazebrook, R. Sharp, M. B. Pracy, S. Brough, M. Colless, et al., *Mon. Not. R. Astron. Soc.* **401**, 1429 (2010), 0911.4246. 68
- [139] D. Schlegel, M. White, and D. Eisenstein, in *astro2010: The Astronomy and Astrophysics Decadal Survey* (2009), vol. 2010 of *Astronomy*, p. 314, 0902.4680. 68
- [140] DES collaboration, *The dark energy survey* (2013), URL <https://www.darkenergysurvey.org/>. 68
- [141] DES collaboration, *The dark energy survey* (2013), URL <https://des.fnal.gov/>. 68
- [142] HETDEX collaboration, *Hetdex - hobby-eberly telescope dark energy experiment* (2013), URL <http://hetdex.org/>. 68
- [143] eBOSS collaboration, *Project eboss* (2013), URL http://lamwws.oamp.fr/cosmowiki/Project_eBoss. 68
- [144] R. Ellis, M. Takada, H. Aihara, N. Arimoto, K. Bundy, M. Chiba, J. Cohen, O. Dore, J. E. Greene, J. Gunn, et al., ArXiv e-prints (2012), 1206.0737. 68
- [145] BigBOSS collaboration, *Bigboss* (2013), URL <http://bigboss.lbl.gov/>. 68
- [146] F. Abdalla, J. Annis, D. Bacon, S. Bridle, F. Castander, M. Colless, D. DePoy, H. T. Diehl, M. Eriksen, B. Flaugher, et al., ArXiv e-prints (2012), 1209.2451. 68
- [147] M. Levi, *Ms-desi intro* (2013), URL <http://ms-desi2013.lbl.gov/talks/Levi.pptx>. 68
- [148] R. Laureijs, J. Amiaux, S. Arduini, J. Augu  res, J. Brinchmann, R. Cole, M. Cropper, C. Dabin, L. Duvet, A. Ealet, et al., ArXiv e-prints (2011), 1110.3193. 68
- [149] ESA, *Esa science and technology: Euclid* (2013), URL <http://sci.esa.int/euclid/>. 68
- [150] WFIRST collaboration, *Wfirst* (2013), URL <http://wfirst.gsfc.nasa.gov/>. 68
- [151] LSST collaboration, *The new sky, lsst* (2013), URL <http://www.lsst.org/lsst/>. 68
- [152] SKA collaboration, *The square kilometer array* (2013), URL <http://www.skatelescope.org/>. 68
- [153] P. Creminelli, L. Senatore, and M. Zaldarriaga, *Journal of Cosmology and Astroparticle Physics* **3**, 19 (2007), arXiv:astro-ph/0606001. 72
- [154] K. M. Smith and M. Zaldarriaga, *Mon. Not. R. Astron. Soc.* **417**, 2 (2011), arXiv:astro-ph/0612571. 72, 74
- [155] A. P. S. Yadav, E. Komatsu, and B. D. Wandelt, *APJ* **664**, 680 (2007), arXiv:astro-ph/0701921. 72

REFERENCES

- [156] E. Komatsu, Classical and Quantum Gravity **27**, 124010 (2010), 1003.6097. 72, 100
- [157] V. Desjacques and U. Seljak, Advances in Astronomy **2010**, 908640 (2010), 1006.4763. 72
- [158] L. Verde, Advances in Astronomy **2010** (2010), 1001.5217. 72
- [159] C. Wagner, L. Verde, and L. Boubekeur, JCAP **10**, 22 (2010), 1006.5793. 72, 74, 89, 116
- [160] J. R. Fergusson and E. P. S. Shellard, Phys. Rev. **D76**, 083523 (2007), astro-ph/0612713. 73, 74
- [161] R. Scoccimarro, L. Hui, M. Manera, and K. C. Chan, PRD **85**, 083002 (2012), 1108.5512. 74, 82, 89, 113, 116
- [162] J. R. Fergusson, D. M. Regan, and E. P. S. Shellard, ArXiv e-prints (2010), 1012.6039. 78, 86, 114
- [163] E. Komatsu, K. M. Smith, J. Dunkley, C. L. Bennett, B. Gold, G. Hinshaw, N. Jarosik, D. Larson, M. R. Nolta, L. Page, et al., Astrophys. J. Supp. **192**, 18 (2011), 1001.4538. 86, 116
- [164] J. Smidt, A. Amblard, C. T. Byrnes, A. Cooray, A. Heavens, and D. Munshi, PRD **81**, 123007 (2010), 1004.1409. 86
- [165] A. Slosar, C. Hirata, U. Seljak, S. Ho, and N. Padmanabhan, JCAP **8**, 31 (2008), 0805.3580. 86
- [166] J.-Q. Xia, C. Baccigalupi, S. Matarrese, L. Verde, and M. Viel, JCAP **8**, 33 (2011), 1104.5015. 86, 89
- [167] N. Dalal, O. Doré, D. Huterer, and A. Shirokov, PRD **77**, 123514 (2008), 0710.4560. 88, 143
- [168] E. Sefusatti, PRD **80**, 123002 (2009), 0905.0717. 89, 95, 102, 144
- [169] E. Sefusatti, M. Crocce, and V. Desjacques, Mon. Not. R. Astron. Soc. **406**, 1014 (2010), 1003.0007. 89, 102, 113
- [170] R. Scoccimarro and H. M. P. Couchman, Mon. Not. R. Astron. Soc. **325**, 1312 (2001), arXiv:astro-ph/0009427. 89, 99
- [171] H. Gil-Marin, C. Wagner, F. Fragkoudi, R. Jimenez, and L. Verde, JCAP **2**, 47 (2012), 1111.4477. 89, 99, 100, 113, 123, 132, 138, 152
- [172] K. C. Chan, R. Scoccimarro, and R. K. Sheth, PRD **85**, 083509 (2012), 1201.3614. 89
- [173] N. Hamaus, U. Seljak, and V. Desjacques, PRD **84**, 083509 (2011), 1104.2321. 89
- [174] M. Manera, R. K. Sheth, and R. Scoccimarro, Mon. Not. R. Astron. Soc. **402**, 589 (2010), 0906.1314. 89
- [175] V. Desjacques, D. Jeong, and F. Schmidt, PRD **84**, 063512 (2011), 1105.3628. 89
- [176] K. M. Smith and M. LoVerde, JCAP **11**, 9 (2011), 1010.0055. 89
- [177] V. Desjacques and U. Seljak, Classical and Quantum Gravity **27**, 124011 (2010), 1003.5020. 89

-
- [178] J.-Q. Xia, A. Bonaldi, C. Baccigalupi, G. De Zotti, S. Matarrese, L. Verde, and M. Viel, *JCAP* **8**, 13 (2010), 1007.1969. 89
 - [179] T. Giannantonio, A. J. Ross, W. J. Percival, R. Crittenden, D. Bacher, M. Kilbinger, R. Nichol, and J. Weller, ArXiv e-prints (2013), 1303.1349. 89
 - [180] K. M. Smith, S. Ferraro, and M. LoVerde, *JCAP* **3**, 32 (2012), 1106.0503. 91
 - [181] A. Cooray and R. Sheth, *Physics Reports* **372**, 1 (2002), arXiv:astro-ph/0206508. 96
 - [182] P. Valageas and T. Nishimichi, *Astronomy and Astrophysics* **532**, A4 (2011), 1102.0641. 96, 98, 139
 - [183] D. G. Figueroa, E. Sefusatti, A. Riotto, and F. Vernizzi, *JCAP* **8**, 036 (2012), 1205.2015. 96, 98, 102, 139
 - [184] J. F. Navarro, C. S. Frenk, and S. D. M. White, *APJ* **462**, 563 (1996), arXiv:astro-ph/9508025. 97
 - [185] A. P. S. Yadav and B. D. Wandelt, *Advances in Astronomy* **2010**, 565248 (2010), 1006.0275. 100
 - [186] J. R. Fergusson and E. P. S. Shellard, ArXiv e-prints (2011), 1105.2791. 112, 113
 - [187] H. F. Gruetjen and E. P. S. Shellard, ArXiv e-prints (2012), 1212.6945. 113
 - [188] R. Scoccimarro, S. Colombi, J. N. Fry, J. A. Frieman, E. Hivon, and A. Melott, *APJ* **496**, 586 (1998), arXiv:astro-ph/9704075. 113
 - [189] R. E. Smith, R. K. Sheth, and R. Scoccimarro, *PRD* **78**, 023523 (2008), 0712.0017. 113
 - [190] V. Springel, N. Yoshida, and S. D. M. White, *New Astronomy* **6**, 79 (2001), arXiv:astro-ph/0003162. 116, 117, 154
 - [191] S. H. Hansen, O. Agertz, M. Joyce, J. Stadel, B. Moore, and D. Potter, *APJ* **656**, 631 (2007), arXiv:astro-ph/0606148. 119
 - [192] T. Battefeld and J. Grieb, *JCAP* **12**, 003 (2011), 1110.1369. 125
 - [193] D. J. Price, *Publications of the Astronomical Society of Australia* **24**, 159 (2007), 0709.0832. 154
 - [194] A. Challinor and G. Chon, *PRD* **66**, 127301 (2002), arXiv:astro-ph/0301064. 156
 - [195] M. Zaldarriaga and U. Seljak, *PRD* **59**, 123507 (1999), arXiv:astro-ph/9810257. 159
 - [196] W. Hu and T. Okamoto, *APJ* **574**, 566 (2002), arXiv:astro-ph/0111606. 159
 - [197] C. M. Hirata and U. Seljak, *PRD* **68**, 083002 (2003), arXiv:astro-ph/0306354. 159
 - [198] E. Anderes, L. Knox, and A. van Engelen, *PRD* **83**, 043523 (2011), 1012.1833. 159
 - [199] W. Hu, *PRD* **64**, 083005 (2001), arXiv:astro-ph/0105117. 160, 166
 - [200] M. Kesden, A. Cooray, and M. Kamionkowski, *PRD* **67**, 123507 (2003), arXiv:astro-ph/0302536. 160, 161, 169, 184
 - [201] T. Namikawa, D. Hanson, and R. Takahashi, *Mon. Not. R. Astron. Soc.* **431**, 609 (2013), 1209.0091. 162
 - [202] NASA, *Lambda*, <http://lambda.gsfc.nasa.gov/> (2013), URL <http://lambda>.

REFERENCES

- gsfc.nasa.gov/. 162
- [203] A. Lewis, PRD **71**, 083008 (2005), arXiv:astro-ph/0502469. 162, 201
 - [204] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, APJ **622**, 759 (2005), arXiv:astro-ph/0409513. 163, 201
 - [205] C. Li, T. L. Smith, and A. Cooray, PRD **75**, 083501 (2007), arXiv:astro-ph/0607494. 166, 167, 168
 - [206] A. Benoit-Lévy, K. M. Smith, and W. Hu, PRD **86**, 123008 (2012), 1205.0474. 166, 167, 168, 184, 185
 - [207] M. Kesden, A. Cooray, and M. Kamionkowski, PRD **66**, 083007 (2002), arXiv:astro-ph/0208325. 179, 217
 - [208] O. Zahn et al., *in prep.* (2013). 200
 - [209] M. Bucher, C. S. Carvalho, K. Moodley, and M. Remazeilles, PRD **85**, 043016 (2012), 1004.3285. 204
 - [210] B. D. Sherwin and S. Das, ArXiv e-prints (2010), 1011.4510. 210, 212
 - [211] NIST, *NIST Digital Library of Mathematical Functions*, <http://dlmf.nist.gov/>, Release 1.0.5 of 2012-10-01, online companion to [213], URL <http://dlmf.nist.gov/>. 211, 215, 237
 - [212] D. Hanson, K. M. Smith, A. Challinor, and M. Liguori, PRD **80**, 083004 (2009), 0905.4732. 220
 - [213] F. W. J. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark, eds., *NIST Handbook of Mathematical Functions* (Cambridge University Press, New York, NY, 2010), print companion to [211]. 237