Inflation and Cosmological Parameter Estimation

Dissertation
zur Erlangung des Doktorgrades
des Departments Physik
der Universität Hamburg

vorgelegt von
Jan Hamann
aus Hamburg

Hamburg
2007

Gutachter der Dissertation:	Prof. Dr. W. Buchmüller
	Prof. Dr. J. Louis
Gutachter der Disputation:	Prof. Dr. W. Buchmüller
	Prof. Dr. G. Kramer
Datum der Disputation:	23. 04. 2007
Vorsitzender des Prüfungsausschusses:	Prof. Dr. J. Bartels
Vorsitzender des Promotionsausschusses:	Prof. Dr. G. Huber
Dekan der Fakultät MIN:	Prof. Dr. A. Frühwald

Abstract

In this work, we focus on two aspects of cosmological data analysis: inference of parameter values and the search for new effects in the inflationary sector.
Constraints on cosmological parameters are commonly derived under the assumption of a minimal model. We point out that this procedure systematically underestimates errors and possibly biases estimates, due to overly restrictive assumptions. In a more conservative approach, we analyse cosmological data using a more general eleven-parameter model. We find that regions of the parameter space that were previously thought ruled out are still compatible with the data; the bounds on individual parameters are relaxed by up to a factor of two, compared to the results for the minimal six-parameter model.
Moreover, we analyse a class of inflation models, in which the slow roll conditions are briefly violated, due to a step in the potential. We show that the presence of a step generically leads to an oscillating spectrum and perform a fit to CMB and galaxy clustering data. We do not find conclusive evidence for a step in the potential and derive strong bounds on quantities that parameterise the step.

Zusammenfassung

In dieser Arbeit behandeln wir zwei Aspekte der statistischen Analyse kosmologischer Daten: Die Genauigkeit der Bestimmung von Modellparametern sowie die Suche nach neuen physikalischen Effekten.
Die Werte kosmologischer Parameter werden im Allgemeinen unter der Annahme eines minimalen Modells bestimmt. Bei dieser Vorgehensweise werden allerdings generell die Fehler unterschätzt. Im Rahmen eines konservativeren Ansatzes analysieren wir kosmologische Daten in einem erweiterten Modell mit elf freien Parametern. Im Vergleich zu Analysen mit dem minimalen 6-Parameter-Modell vergrößert sich der erlaubte Bereich im Parameterraum deutlich. Für einzelne Parameter wächst die Unsicherheit um einen Faktor zwei.
Außerdem untersuchen wir eine Klasse von Inflationsmodellen, in der das Inflatonpotential eine Stufe aufweist. Dies führt zu einer Oszillation im Spektrum der primordialen Dichtefluktuationen. Ein Vergleich der Vorhersagen dieser Modelle mit Daten aus Messungen der Hintergrundstrahlung und der Verteilung von Galaxien ergibt keine zwingende Evidenz für eine Stufe im Potential. Wir finden starke Einschränkungen für den erlaubten Parameterbereich.

Contents

Introduction

In the past two decades, the field of theoretical cosmology has come a long way. Twenty years ago, our knowledge about the beginnings of the Universe was, for a lack of solid observational data, speculative at best. These days, the situation presents itself in a completely different light. High precision measurements of various cosmological probes, such as the anisotropies of the Cosmic Microwave Background (CMB) or the distribution of galaxies, have brought a wealth of information, enabling us to answer fundamental questions about the nature of the Universe.

The recent advances have led to the emergence of the so-called cosmological concordance model, whose predictions are consistent with all current data. But at the same time, new questions and challenges have arisen. It seems that only about five percent of the energy density of the Universe are made up of particles of the Standard Model of particle physics. Another quarter consists of the unknown *dark matter*, the remainder being taken up by the even more enigmatic *dark energy*. In the case of dark matter, there is well-founded hope that its mysteries will be unravelled at the upcoming Large Hadron Collider at CERN, scheduled to start operating this year. The detection and identification of a particle with the right properties to constitute the dark matter would be a spectacular confirmation of our cosmological picture. A solution to the dark energy conundrum, on the other hand, will most likely require a deeper understanding of quantum gravity.

The data also provide substantial evidence that the infant Universe underwent a period of *inflation*, an exponential growth of its scale factor driven by the potential energy of a scalar field, the *inflaton*. Initially postulated to explain the observed large scale homogeneity and isotropy of the Universe [1], it was soon realised that inflation also offers a mechanism for the generation of small density perturbations [2, 3]. During inflation, quantum fluctuations of the inflaton field are blown up to macroscopic scales and eventually "freeze in", as they are stretched beyond the horizon scale. The so created inhomogeneities eventually start growing under the influence of gravity and can be considered the seeds of structure formation.

Whereas there is little doubt that inflation has taken place in some form, it is not clear how it can be embedded in the framework of particle physics; the inflaton has not yet revealed its identity to us. Consequently, there is a large number of suggestions on the market, differing slightly in their predictions of the perturbation spectra. The current data put us in a position where we can begin to rule out candidates. This can give us important clues about the physics happening at energies only a few orders of magnitude below the Planck scale.

While most inflation models predict a smooth and featureless spectrum of primordial perturbations, there also exist slightly more exotic models that predict certain features. The discovery of evidence for such a feature in the data would have far-reaching consequences for inflationary model-building. In this work, we analyse a particular class of non-standard inflation models, where the scalar potential of the inflaton field has a step-like form. We show that the presence of a step generically leads to an oscillating spectrum and perform a fit to CMB and galaxy clustering data. Our results show that there is no conclusive evidence for a step in the potential and we derive strong bounds on quantities that parameterise the step.

Another important problem in the era of precision cosmology is that of parameter estimation. Most cosmological parameters of interest cannot be measured directly. Rather, their magnitudes and uncertainties have to be inferred from their effects on observable quantities. The results of the inference depend, among other things, on the underlying model one assumes for the analysis. In the literature, it is a common practice to quote the values derived under the assumption of a minimal model. We point out that this procedure systematically underestimates errors and possibly biases the inferred parameters, due to the restrictive assumptions made in that model.

Considering that estimates of, for instance, the dark matter density, are often used as input to constrain other theories, we advocate a more conservative approach. To this end, we analyse cosmological data using a more general eleven-parameter model, incorporating a number of well-motivated physical effects, e.g., a neutrino mass. We find that regions of the parameter space that were previously thought ruled out are still compatible with the data; the bounds on individual parameters are relaxed by up to a factor of two, compared to the results for the minimal six-parameter model.

This thesis is organised as follows. In Chapter I, the reader will be reminded of some of the basic concepts of Friedmann-Lemaître-Robertson-Walker cosmology and the cosmological concordance model. Chapter II will illuminate the rôle of inflation in the generation of initial perturbations, which determine the dynamics of structure formation during later stages of cosmic evolution. In Chapter III we will review the statistical concepts and methods that form the foundations of the analysis of cosmological data and pave the way for an application of these techniques in Chapters IV and V, where the main results of our work will be presented, largely based on our Refs. [4–6].

Throughout this work we will use natural units with $c = \hbar = 8\pi G = 1$.

I. The Cosmological Concordance Model

This introductory chapter serves to introduce the notation and the reader will be given a very brief overview of the cosmological concordance model, the current state of observations, and possible extensions of the basic model. In the following, we will adhere to the sign conventions of Ref. [7].

1. Basics

The quantitative description of cosmology rests on a fundamental postulate, known as the *cosmological principle*, which states that when averaged over large scales, our Universe is spatially homogeneous and isotropic. In other words, no single point or direction should be distinguished. The most general metric possessing this maximal spatial symmetry is the Friedmann-Lemaître-Robertson-Walker metric:

$$\mathrm{d}s^2 = g_{\mu\nu}\mathrm{d}x^\mu \mathrm{d}x^\nu = \mathrm{d}t^2 - a^2(t)\left[\frac{\mathrm{d}r^2}{1-\mathcal{K}r^2} + r^2\left(\mathrm{d}\theta^2 + \sin^2\theta\,\mathrm{d}\phi^2\right)\right], \tag{I.1}$$

where a is the scale factor and $\mathcal{K} \in \{-1, 0, 1\}$ corresponds to a spatially open, flat, or closed Universe, respectively. Defining the conformal time interval $\mathrm{d}\tau \equiv \mathrm{d}t/a$, the metric can also be written as

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

Einstein's equations

$$G_{\mu\nu} = T_{\mu\nu} \tag{I.3}$$

relate the Universe's geometry, encoded in the Einstein tensor $G_{\mu\nu}$, to its energy content represented by the energy-momentum tensor $T_{\mu\nu}$. Obviously, $T_{\mu\nu}$ must obey the same symmetries as $G_{\mu\nu}$ and is typically parameterised as an ideal fluid with pressure p and energy density ρ,

$$T_{\mu\nu} = -p\, g_{\mu\nu} + (\rho + p)\, u_\mu u_\nu, \tag{I.4}$$

with $u_\mu = (1, 0, 0, 0)$ the velocity of a comoving observer. With $T_{\mu\nu}$ of this form, Einstein's equations reduce to two independent equations, Friedmann's equation and the acceleration equation

$$H^2 \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{\rho}{3} - \frac{\mathcal{K}}{a^2}, \tag{I.5}$$

$$6\,\frac{\ddot{a}}{a} = -\left(\rho + 3p\right), \tag{I.6}$$

where a dot represents a derivative with respect to t, and H is the Hubble parameter. Its current value H_0 is often given in terms of the dimensionless quantity

$$h = H_0/(100\ \mathrm{km\,s^{-1}Mpc^{-1}}). \tag{I.7}$$

Defining a critical density $\rho_\mathrm{c} \equiv 3H^2$ and $\Omega_i \equiv \rho_i/\rho_\mathrm{c}$, Friedmann's equation can be rewritten as

$$1 - \sum_i \Omega_i = -\frac{\mathcal{K}}{a^2H^2} \equiv \Omega_\mathcal{K}, \tag{I.8}$$

where the summation index i runs over all components of the cosmic fluid. To our current knowledge, we need at least five different fluids to model our Universe halfway faithfully: photons (Ω_γ), neutrinos (Ω_ν), baryonic matter (Ω_b), dark matter (Ω_DM) and dark energy (Ω_DE), each of them equipped with an equation of state relating their pressure and energy density:

$$p_i = w_i(t)\,\rho_i. \tag{I.9}$$

In the case of radiation or relativistic matter, for example, we expect $w = \frac{1}{3}$. Non-relativistic matter has $w = 0$, and a cosmological constant corresponds to $w = -1$.

It is clear that the idealised homogeneous and isotropic description of the Universe cannot be wholly realistic. It does turn out, however, that a large number of phenomenological aspects of the Universe can be very well approximated by considering small perturbations $\delta g_{\mu\nu}$ of the metric (I.1),

$$\tilde{g}_{\mu\nu}(\boldsymbol{x}, t) = g_{\mu\nu}(t) + \delta g_{\mu\nu}(\boldsymbol{x}, t), \tag{I.10}$$

and, consequently, small perturbations $\delta T_{\mu\nu}$ of the homogeneous and isotropic energy-momentum tensor,

$$\tilde{T}_{\mu\nu}(\boldsymbol{x}, t) = T_{\mu\nu}(t) + \delta T_{\mu\nu}(\boldsymbol{x}, t). \tag{I.11}$$

As long as the perturbations are sufficiently small with respect to the background, linear perturbation theory will be an appropriate tool for the treatment of the inhomogeneities.

For a statistical description of a perturbation δf (where f could be, for instance, the energy density ρ_i of component i), we define the real-space *correlation function* ξ,

$$\xi_f(r) = \langle \delta f(\boldsymbol{x}_1)\, \delta f(\boldsymbol{x}_2) \rangle, \tag{I.12}$$

with $r = |\boldsymbol{x}_1 - \boldsymbol{x}_2|$ and $\langle\cdot\rangle$ denoting the ensemble average.[1] Its counterpart in momentum space, the *power spectrum* $\mathcal{P}$, is defined by

$$\mathcal{P}_f(k) = 4\pi \int \mathrm{d}r\; r^2 \xi_f(r)\, \frac{\sin kr}{r}. \tag{I.13}$$

In linear perturbation theory, modes of different wavenumbers will decouple, and their equations can be evolved separately. In general, the fluctuations will of course be

---

[1] The field δf is taken to be an ergodic, statistically homogeneous and isotropic random field, see, e.g., Refs. [8, 9] for a more complete account.

time-dependent. The time evolution can be expressed in terms of a *transfer function* $\Delta^2(\tau_1, \tau_2)$:

$$\mathcal{P}(k, \tau_1) = \Delta^2(\tau_1, \tau_2)\, \mathcal{P}(k, \tau_2). \tag{I.14}$$

An accurate calculation of the transfer function is fairly involved and a detailed discussion of this topic would be beyond the scope of this work. We refer the interested reader to, e.g., Refs. [7, 10–12]. Let us just briefly outline the basic idea here.

To calculate the transfer function, one needs to track the behaviour of the constituents of the energy-momentum tensor, i.e., the different cosmological fluids and the corresponding perturbations. Their properties can be described in terms of the *distribution function* $\mathfrak{f}(\tau, \boldsymbol{x}, \boldsymbol{p})$, which gives the number $\mathrm{d}N$ of particles per volume element of phase space $\mathrm{d}^3\boldsymbol{x}\, \mathrm{d}^3\boldsymbol{p}$, with $\boldsymbol{p}$ denoting the canonically conjugate momentum of $\boldsymbol{x}$. It is related to the energy-momentum tensor by

$$\tilde{T}_i^{\mu\nu} = \frac{g_i}{(2\pi)^3} \int \frac{\mathrm{d}^3p}{E_i}\, p^\mu p^\nu\, \mathfrak{f}(\tau, \boldsymbol{x}, \boldsymbol{p}), \tag{I.15}$$

where g_i is the number of spin states and $E_i = \sqrt{m_i^2 + p^2}$ is the energy of the particles making up fluid i. The dynamics of $\mathfrak{f}$ are governed by the *Boltzmann equation*, in an abstract form given by

$$\frac{\mathrm{d}\mathfrak{f}}{\mathrm{d}\tau} = C\,[\mathfrak{f}]\,. \tag{I.16}$$

The total derivative on the left hand side encodes the effects of gravity while the functional C on the right hand side describes the interactions of the fluids. If a fluid only interacts gravitationally, such as for instance the dark matter, then $C = 0$ and Equation (I.16) is called the *collisionless* Boltzmann equation. In a Universe filled with photons, neutrinos, baryonic matter, dark matter, and dark energy, the only relevant interaction after Big Bang Nucleosynthesis (BBN) is Compton scattering between the baryons and the photons (taking place before recombination and after reionisation, see below; this also ignores highly non-linear effects such as star formation). So the task at hand is to evolve a set of four (five, if one allows perturbations of the dark energy component) coupled Boltzmann equations, and construct the power spectrum from the distribution functions, which can be used to retrieve the transfer function.

There exist a number of publicly available numerical programmes dedicated to determining the transfer functions [13, 14], based on the `cmbfast` code by Seljak and Zaldarriaga [15].

The advantage of this approach is that the transfer function is independent of the initial spectrum of perturbations, as long as linear perturbation theory is valid. We only need to know the type of perturbations: generally, any initial state can be expressed in terms of a linear combination of an adiabatic mode and isocurvature modes. Let us define the entropy, or fluctuation in the number density n, between two components i and j as

$$S_{ij} = \delta(n_i/n_j), \tag{I.17}$$

and the curvature perturbation on comoving hypersurfaces $\mathcal{R}$ (more on this quantity in Chapter II). Adiabatic initial conditions correspond to pure curvature perturbations ($\mathcal{R} \neq 0$ and $S = 0$), while isocurvature modes correspond to perturbations in the entropy with no perturbation in the curvature ($\mathcal{R} = 0$ and $S \neq 0$).

In this work, we shall mostly be concerned with the form of the initial spectrum of perturbations, where "initial" refers to a time during the radiation dominated phase, after BBN. But what determines these initial conditions? In principle, they could be chosen arbitrarily, leading to equally arbitrary observable spectra today and hence a theory without predictivity, certainly not an attractive feature. We will see in Chapter II though, that we can get a quite definite prediction for the form of the initial power spectra, if we postulate a period of cosmic inflation to have taken place at an energy beyond the BBN scale of $\mathcal{O}(\mathrm{MeV})$. As it happens, this prediction appears to describe the observed data quite well. But before we get too far ahead, let us first present what kind of observational information we have about the state of our Universe.

2. Cosmological Data

The most powerful probes of the history of the Universe are those that track the perturbations of its energy components. Mapping these perturbations requires a messenger which, on the one hand, enjoys a more or less free propagation from its source to us and, on the other hand, is also detectable. The first requirement rules out charged particles, the second one precludes us from using neutrinos as messengers. This leaves only the photon to track the perturbations. Due to these limitations, only the perturbations of the photon background and the baryonic matter component are available for direct measurements. We can, however, indirectly infer information about, e.g., the dark matter perturbations via their gravitational interaction.

Furthermore, if we want to make reliable theoretical predictions for the spectra, it is essential that the perturbations are small enough, such that nonlinear and backreaction effects can be kept under control. This puts a lower limit to the wavelengths of the perturbations one can potentially probe at a given redshift.

In the following, we will briefly describe the most important observable quantities and give examples of the current state of data.

2.1. Matter Perturbations

If we require the baryonic matter to be in the linear (or quasi-linear) regime, it will essentially consist of a sparse gas of neutral hydrogen and helium atoms. One way of tracing it is through its absorption of photons if it lies between a photon source and the observer. Such a source could, for instance, be a distant quasar, in whose spectra one can detect absorption lines from the Lyman-α transition of neutral hydrogen in the interstellar medium. From these measurements of the "Lyman-α-forest", the power spectrum can be reconstructed [16].

Additionally, one could try to detect absorption lines in the spectra of CMB photons due to the 21 cm spin flip in neutral hydrogen, and possibly also emission at lower redshifts [17,18]. This method can potentially cover an extremely large range of scales, but the signal is very faint and may be hard to separate from foreground emission; its practical application remains a challenge for the future.

At the scales where baryonic matter forms luminous objects, such as stars or galaxies, nonlinear effect are dominant by far. However, galaxies are more likely to form in the potential wells of dark matter overdensities. They can therefore be used as a tracer of the dark matter perturbations. The galaxy power spectrum $\mathcal{P}_\mathrm{g}$ and the dark matter power spectrum $\mathcal{P}_\mathrm{DM}$ are assumed to be proportional to each other,

$$\mathcal{P}_\mathrm{g}(k) = b'^2\, \mathcal{P}_\mathrm{DM}(k), \tag{I.18}$$

and the bias factor b' is taken to be scale independent on scales in the linear regime.

Given a catalogue of galaxy positions and redshifts, one can try to reconstruct the correlation function and the power spectrum. While the theoretical correlation function and power spectrum contain exactly the same information (since they are related by a Fourier transform), the same is not true for the observational data sets, even though they may be derived from the same raw data. This difference is due to a loss of information in the individual reconstruction processes, involving e.g., a binning of the raw data.

Galaxy Power Spectrum

We plot the galaxy power spectrum from the luminous red galaxy (LRG) sample of the Sloan Digital Sky Survey (SDSS) [19] in the left panel of Figure I.1. At large scales the accuracy is limited by boundary effects of the survey geometry; at small scales $k/h \gtrsim 0.06$ Mpc^{-1}, the effects of nonlinear structure growth begin to set in. For reference, we show the theoretical prediction of the concordance model for the linear power spectrum and an empirical correction for the nonlinearities.

Galaxy Correlation Function

The SDSS LRG correlation function data [20] are shown in the right panel of Figure I.1. Unlike the power spectrum data, which are relatively uncorrelated, this data set is highly correlated, but, on the other hand, less susceptible to the effects of nonlinear structure growth. The peak at $s \simeq 110$ Mpc$/h$ corresponds to the acoustic oscillations of the plasma prior to decoupling and singles out an important physical scale, the size of the sound horizon at decoupling.

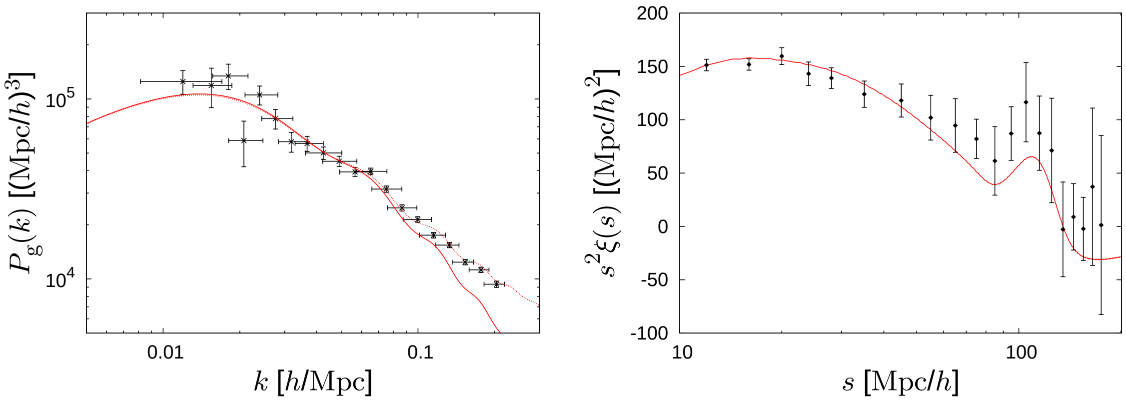

Figure I.1.: *Left*: Power spectrum of luminous red galaxies measured by SDSS [19]. The solid line is the power spectrum from linear perturbation theory, the dotted line includes corrections due to scale dependent bias and nonlinear growth of structure. *Right*: Real-space correlation function of luminous red galaxies vs. comoving separation s, measured by SDSS [20].

Weak Lensing

It is also possible to circumvent the galaxy-dark matter bias problem and directly estimate the underlying dark matter power spectrum from the distortion of galaxy shapes due to gravitational lensing effects of the dark matter [21]. The presently available data, however, cannot yet compete with that of conventional galaxy surveys [22].

2.2. Cosmic Microwave Background Anisotropies

After the recombination of protons and free electrons at a redshift of $z_{\rm rec} \simeq 1100$, the Universe became transparent to photons. Since then, the propagation of the CMB photons has been almost undisturbed. Only with the ignition of the first stars ("reionisation") around redshift $z_{\rm ri} \simeq 10$ and the repopulation of the Universe with ions and free electrons from stellar winds was there a chance for the photons to be scattered again. This issue is commonly parameterised in terms of the optical depth to reionisation τ (not to be confused with the conformal time): a fraction of $1 - e^{-\tau}$ of the CMB photons have scattered once since recombination. The remaining fraction of $e^{-\tau}$ has reached us directly from the surface of a sphere centred around us, the surface of last scattering.

While the information we can gather about, say, the galaxy power spectrum, is three-dimensional, the CMB can only give us two-dimensional information about perturbations. The definition of the power spectrum as given in Equation (I.13) will not be applicable here. Instead, since the perturbations lie on the surface of a sphere, one resorts to an expansion in terms of spherical harmonics (see Appendix A.1 for details). The angular power spectrum $\mathcal{C}_\ell$ is related to the three-dimensional power spectrum $\mathcal{P}(k)$ via

$$\mathcal{C}_\ell = \int \frac{\mathrm{d}k}{k}\, \mathcal{P}(k)\, \Delta_\ell^2(k). \qquad \text{(I.19)}$$

In the line-of-sight approach of Ref. [15], the transfer function $\Delta_\ell^2(k)$ can be written as an integral in conformal time τ over a geometrical part and a part that depends on the cosmology:

$$\Delta_\ell(k) = \int_0^{\tau_0} \mathrm{d}\tau \, j_\ell(k(\tau - \tau_0)) \, S(k, \tau), \tag{I.20}$$

where τ_0 corresponds to today, j_ℓ is a spherical Bessel function and $S(k, \tau)$ is a source function, calculated from the Boltzmann equations.

The anisotropies of the CMB manifest themselves in two ways: firstly, as temperature anisotropies $\Delta T/T$ sourced by inhomogeneities in the gravitational potential, density perturbations in the photon-baryon gas and Doppler redshift from fluctuations in the peculiar velocity of the baryons which emitted the photons. In addition to the temperature fluctuations, one finds linear polarisation induced by Thomson scattering in the inhomogeneous plasma at recombination and after reionisation. Linear polarisation corresponds to a two-dimensional vector field on the last scattering sphere, which can be split into two scalar quantities, a curl-free E-mode and a divergence-free B-mode, resembling the vector field properties of classical electrostatics.

Figure I.2 shows the Wilkinson Microwave Anisotropy Probes (WMAP) measurements of the temperature and E-polarisation angular power spectra, as well as the cross-correlation between the two [23–26]. The B-mode, being notoriously hard to separate from foregrounds and leaking of E-mode power due to incomplete sky coverage, has eluded detection so far and only upper limits exist.

2.3. Background Probes

Apart from measurements of the perturbations, there are numerous observations which bring additional information about the background model, i.e., the averaged quantities. These include for instance

- **Big Bang Nucleosynthesis**:
  By comparison of measurements of the primordial abundances of light elements with the predictions of BBN, one can find constraints on the baryon density: $0.017 < \Omega_b h^2 < 0.024$ at 95% confidence level [27].

- **Type Ia Supernovæ (SNIa)**:
  Due to a simple empirical relation between their absolute luminosity and width of the light curve [28], supernovæ of type Ia can be considered standard candles. Measurements of their luminosity distance allow a reconstruction of the expansion history of the Universe at redshifts up to $z \simeq 1$ [29, 30]. Hence, this data set is particularly sensitive to the presence of dark energy, which dominates the energy density in recent times.

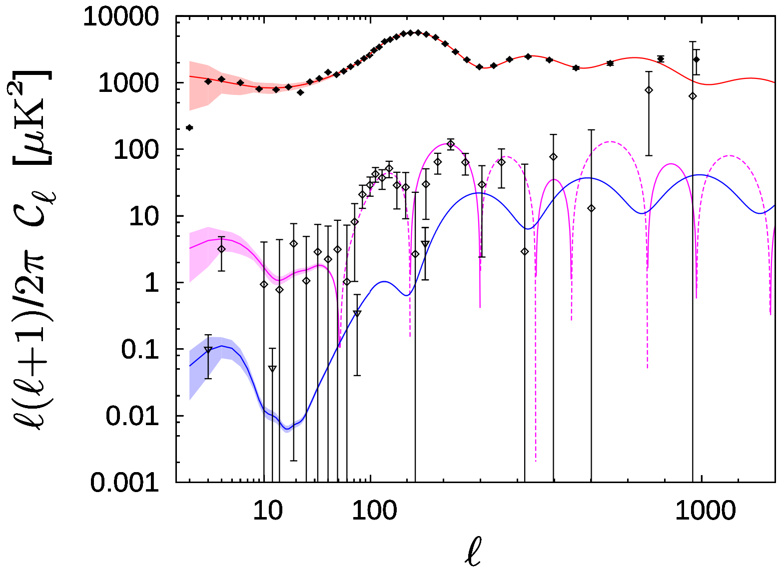

Figure I.2.: Three-year WMAP data of the angular power spectra of temperature (red) and E-polarisation (blue) anisotropies and their cross-correlation (purple), dashed lines corresponding to anticorrelation. The shaded bands indicate the uncertainties due to cosmic variance.

- **Measurements of the Hubble parameter**:
  The history of observational determinations of the Hubble parameter has been full of controversies [31], due to various sources of systematic errors. The most-cited recent value is $h = 0.72 \pm 0.08$ at 68% C.L. [32].

3. The Vanilla Model

Given the cosmological data presented in the previous section, one could ask, 'what is the model with the least number of free parameters that describes these data well?' This model is known under various names, ΛCDM-model, the concordance model or the "vanilla" model. It is based on the FLRW-cosmology of Section I.1 and the following additional assumptions:

- The Universe is flat, $\Omega_K = 0$, as predicted by inflation.
- It is filled with
  - the photons of the CMB, which has a temperature of $T = 2.726$ K,
  - a cosmic neutrino background (CνB), consisting of three species of *massless* neutrinos,
  - baryonic matter (this includes charged leptons), with a primordial Helium fraction of $Y_p = 0.24$ as predicted by a naïve estimate [33],

  - cold dark matter, and
  - dark energy in the form of a cosmological constant, i.e., $w_{\mathrm{DE}} = -1$.

- The initial perturbations are adiabatic, Gaussian and scalar and their spectrum is given by a power-law $\mathcal{P}(k) = A_{\mathrm{S}}(k/k_0)^{n_{\mathrm{S}}-1}$, with a spectral index $n_{\mathrm{S}} \approx 1$, and a normalisation A_{S} defined at a pivot scale k_0, as predicted by single field slow roll inflation (this point will be treated in more detail in the next chapter).

- Due to reionisation, a fraction of $1 - e^{-\tau}$ of the CMB photons have scattered since recombination.

- Its topology is trivial.

This model has six free parameters: the baryon and cold dark matter densities $\Omega_{\mathrm{b}}h^2$ and $\Omega_{\mathrm{c}}h^2$, the Hubble parameter H_0, the optical depth to reionisation τ, the normalisation of the primordial power spectrum A_{S} and its tilt n_{S}. Choosing the best-fit values of these parameters gives a chi-squared per degree of freedom (DOF) of $\chi^2_{\mathrm{eff}}/\mathrm{DOF} \simeq 1.02$. The 95% C.L. intervals of these parameters for a fit of the vanilla model to the WMAP data are listed in Table I.1. These numbers should only be taken as rough guidelines though; we will delve deeper into this issue in Chapters III and IV. Removing any of these parameters and fixing them to their "natural" value (e.g., a model with $\Omega_{\mathrm{c}} = 0$, or a scale independent spectrum with $n_{\mathrm{S}} = 1$) would significantly worsen the quality of the fit.

The point that has made the vanilla model so successful and earned it the epithet "concordance", is that it appears to fit all currently available cosmological data from various independent sources without major discrepancies. Some people have even called it the standard model of cosmology, but we should remember that it is only an empirical model and not a fundamental theory, like for instance the Standard Model of particle physics.

Parameter		95% C.L. region
Baryon density	$\Omega_{\mathrm{b}}h^2$	$0.0207 \rightarrow 0.0236$
Cold dark matter density	$\Omega_{\mathrm{c}}h^2$	$0.089 \rightarrow 0.121$
Hubble parameter	h	$0.67 \rightarrow 0.80$
Reionisation optical depth	τ	$0.04 \rightarrow 0.14$
Normalisation of scalar spectrum	$\ln[10^{10}A_{\mathrm{S}}]$	$2.89 \rightarrow 3.17$
Scalar spectral index	n_{S}	$0.927 \rightarrow 0.993$

Table I.1.: Parameters of the vanilla model and their 95% C.L. regions for a fit to the three year WMAP data.

4. Extensions of the Basic Model

There are dozens of more or less well-motivated extensions to this basic model. They all have in common that none of them explains the observed data so much better that it would have to be deemed absolutely necessary to introduce these extra effects. We cannot hope to give a complete list of alternatives here, but let us at least mention a few:

- **Massive neutrinos**:
  Cosmological data can possibly complement oscillation experiments in providing evidence for neutrino masses. Massive neutrinos suppress the growth of structure at small scales due to free streaming. A convenient parameterisation is in terms of $f_\nu = \Omega_\nu/\Omega_\mathrm{c}$, which is approximately related to the sum of neutrino masses by $\Omega_\nu h^2 \simeq \sum m_\nu/93$ eV.

- **Dynamical dark energy**:
  Dark energy appears to have the largest share of the cosmic energy budget, yet we know very little about it. Departures from a cosmological constant can be as simple as $w_\mathrm{DE} \neq -1$, but the equation of state parameter could also be time-dependent as, e.g., in quintessence models [34]. In these scenarios one would also have to take into account the effects of perturbations in the dark energy.

- **Non-standard initial perturbations**:
  From non-power-law spectra over tensor perturbations, isocurvature modes and non-Gaussianities to subdominant contributions from topological defects, there are many models which predict deviations from the vanilla setting. We shall grant aspects of this point a more in-depth coverage in Chapter V.

- **Non-zero spatial curvature**:
  This is one of the benchmark tests for the theory of inflation; if it could be unequivocally shown that the spatial curvature of the Universe is significantly different from zero, inflation would be in a bad shape.

- **Multiply connected Universe**:
  General Relativity is a local theory, it cannot make statements about global properties of the Universe. If the Universe had periodic boundary conditions on a scale smaller than the current Hubble scale, one might be able to find characteristic signatures in the CMB, see for instance Ref. [35].

- **Alternative theories of gravity**:
  General Relativity has been experimentally tested on scales between $\mathcal{O}(0.1\mathrm{mm})$ and roughly the radius of the solar system. At larger or smaller scales, there might be corrections to the standard picture and several alternatives have been suggested (cf., e.g., Refs. [36, 37]).

II. Inflation and Cosmological Perturbations

The aim of this chapter is to outline how the initial perturbations are generated from quantum fluctuations during inflation. We will see why they survive until a much later time to set the stage for their subsequent evolution during the radiation- and matter-dominated phases and show how to calculate their spectrum.

We employ the gauge invariant approach for the description of perturbations, following Refs. [38, 39]. For simplicity, we restrict ourselves to the single-field inflation scenario in which only one scalar field is of dynamical relevance during inflation.

We will derive a set of differential equations which allow us to calculate the initial perturbation spectra exactly, within the limits of linear perturbation theory. This system of equations does not generally have an analytic solution, but can easily be solved numerically.

One does not, however, always have to resort to the numerical brute force method. It is possible to find an approximate analytical solution, the slow roll approximation, which is applicable to most viable inflation models on the market these days. If a model happens to be endowed a very flat and smooth potential, the inflaton field will slowly roll towards the vacuum. The flatness and smoothness of the potential (or, alternatively, the slow variation of the Hubble parameter) can be quantified in terms of the slow roll parameters and the resulting perturbation spectra can be expanded in terms of these parameters [40].

In the end, we shall see that single field slow roll inflation predicts an almost scale-invariant power-law spectrum of scalar and tensor perturbations, thus justifying the parameterisation used in the concordance model.

Our journey begins with an action $\mathfrak{S}$: for single-field inflation, we will take a canonically normalised, minimally coupled scalar field ϕ with potential $V(\phi)$ in Einstein gravity:

$$\mathfrak{S} = \int \mathrm{d}^4x \, \sqrt{-g} \left[-\tfrac{1}{2}R + \tfrac{1}{2}\partial_\mu\phi\partial^\mu\phi - V(\phi) \right] . \tag{II.1}$$

The (unperturbed) background metric is taken to be the flat Friedmann-Lemaître-Robertson-Walker metric

$$\mathrm{d}s^2 = \mathrm{d}t^2 - a^2(t)\mathrm{d}\boldsymbol{x}^2 = a^2(\tau)\left[\mathrm{d}\tau^2 - \mathrm{d}\boldsymbol{x}^2\right] . \tag{II.2}$$

In the following, derivatives with respect to conformal time τ will be denoted with a prime. Consider now a perturbed metric and scalar field

$$\tilde{g}_{\mu\nu}(\boldsymbol{x}, t) = g_{\mu\nu}(t) + \delta g_{\mu\nu}(\boldsymbol{x}, t) \, , \tag{II.3}$$

$$\tilde{\phi}(\boldsymbol{x}, t) = \phi(t) + \delta\phi(\boldsymbol{x}, t). \tag{II.4}$$

The perturbed metric has ten degrees of freedom, which can be divided into three classes: scalar (4 DOF), vector (4 DOF) and tensor (2 DOF) perturbations, which do not couple to each other at the linear order and can thus be treated separately. In fact, due to the invariance of General Relativity under gauge (i.e., general coordinate) transformations which removes four degrees of freedom, there are only six independent DOF left, two for each type of perturbation.

In the inflationary scenario, vector perturbations will not be excited. Therefore we will focus on the other two, beginning with the scalar ones.

1. Scalar Perturbations

If we limit ourselves to scalar linear perturbations, the most general form of the perturbed metric is given by [38]

$$\mathrm{d}s^2 = a^2(\tau)\left\{(1+2A)\,\mathrm{d}\tau^2 - 2\,\partial_i B\,\mathrm{d}x^i\,\mathrm{d}\tau - \left[(1+2\psi)\delta_{ij} + 2\,\partial_i\partial_j E\right]\,\mathrm{d}x^i\mathrm{d}x^j\right\}, \tag{II.5}$$

where A, B, ψ and E are real numbers, representing the four scalar degrees of freedom before fixing a gauge. After choosing a suitable gauge, two degrees of freedom will remain, one of which can be eliminated by imposing the condition that the energy momentum tensor $T_{\mu\nu}$ be stress-free: $T^i_j = 0$, $(i \neq j)$. This is well motivated since the quantum fluctuations of a scalar field can be regarded as a statistically homogeneous and isotropic random field [41] which satisfies this requirement. The one remaining degree of freedom will eventually be related to the field perturbation by Einstein's equations.

We can now plug the perturbed quantities into the action. It is easy to see that the first variation of the perturbed action is zero, since it is proportional to the solutions of the background equations of motion. The first non-vanishing term is therefore the second variation:

$$\widetilde{\mathfrak{S}}^{\mathrm{S}} = \mathfrak{S}^{\mathrm{S}} + \delta^2\mathfrak{S}^{\mathrm{S}}. \tag{II.6}$$

It can be shown that the action for the perturbations can be expressed in terms of $z \equiv a\frac{\dot{\phi}}{H}$ and the Mukhanov variable u [42,43],

$$\delta^2\mathfrak{S}^{\mathrm{S}} = \int \mathrm{d}\tau\,\mathrm{d}^3\boldsymbol{x}\left[\tfrac{1}{2}\left((u')^2 - \partial_i u\,\partial^i u\right) + \tfrac{1}{2}\,\frac{z''}{z}\,u^2\right], \tag{II.7}$$

up to total derivatives. The Mukhanov variable u is a gauge invariant combination of field and metric perturbations and can be defined via $\mathcal{R}$, the curvature perturbation (i.e., the 3D Ricci scalar) on comoving hypersurfaces:

$$u \equiv -z\mathcal{R}. \tag{II.8}$$

In the longitudinal gauge ($E = B = 0$), the no-anisotropic-stress requirement reads $A = -\psi$, and u can be expressed as

$$u = -z\left(\psi - \frac{H}{\dot{\phi}}\,\delta\phi\right). \tag{II.9}$$

Equation (II.7) is the action of a scalar field with a time-dependent squared mass. If we now promote u to an operator and quantise it,

$$\hat{u}(\tau, \boldsymbol{x}) = \frac{1}{(2\pi)^{3/2}} \int \mathrm{d}^3k \left\{ \hat{\mathfrak{a}}_k \, u_k(\tau) \, e^{\mathrm{i}\boldsymbol{k}.\boldsymbol{x}} + \hat{\mathfrak{a}}_k^\dagger \, u_k^*(\tau) \, e^{-\mathrm{i}\boldsymbol{k}.\boldsymbol{x}} \right\}, \tag{II.10}$$

imposing the standard commutation relations for the creation and annihilation operators $\hat{\mathfrak{a}}^\dagger$ and $\hat{\mathfrak{a}}$, it follows that the Fourier components of momentum k are decoupled from other momenta and obey a simple equation,

$$u_k'' + \left(k^2 - \frac{z''}{z} \right) u_k = 0. \tag{II.11}$$

Finally, we can define the primordial power spectrum of curvature perturbations $\mathcal{P}_\mathcal{R}(k)$ via the two-point correlation function in k-space

$$\langle \mathcal{R}_{\boldsymbol{k}} \mathcal{R}^*_{\boldsymbol{k}'} \rangle = \frac{2\pi^2}{k^3} \mathcal{P}_\mathcal{R}(k) \, \delta^{(3)}(\boldsymbol{k} - \boldsymbol{k}'). \tag{II.12}$$

It is related to u_k and z by

$$\mathcal{P}_\mathcal{R}(k) = \frac{k^3}{2\pi^2} \left| \mathcal{R}_k \right|^2 = \frac{k^3}{2\pi^2} \left| \frac{u_k}{z} \right|^2. \tag{II.13}$$

Assuming Gaussianity and adiabaticity, this quantity contains all the necessary information for a complete statistical description of the fluctuations. Note that the perturbations generated during single field inflation are necessarily adiabatic, since there is only one fluid present.

1.1. Background Equations of Motion

In order to find a solution to Equation (II.11), one needs to know the behaviour of the term z''/z. Its evolution is determined by the dynamics of the Hubble parameter and the unperturbed inflaton field, governed by Friedmann's equation

$$H^2 = \tfrac{1}{3}(V + \tfrac{1}{2}\dot{\phi}^2), \tag{II.14}$$

and the Klein-Gordon equation for ϕ

$$\ddot{\phi} + 3H\dot{\phi} + \frac{\mathrm{d}V}{\mathrm{d}\phi} = 0. \tag{II.15}$$

For our purposes, it is convenient to introduce another time parameter, the number of e-foldings, defined by $N \equiv \ln a/a_0$. In terms of N, Equations (II.11), (II.14) and (II.15) read

$$\begin{aligned} & u_{k,NN} + \left(\frac{H_{,N}}{H} + 1\right) u_{k,N} \\ & \quad + \left[\frac{k^2}{e^{2(N-N_0)}H^2} \left(2 - 4\frac{H_{,N}}{H}\frac{\phi_{,NN}}{\phi_{,N}} - 2\left(\frac{H_{,N}}{H}\right)^2 - 5\frac{H_{,N}}{H} - \frac{1}{H^2}\frac{\mathrm{d}^2V}{\mathrm{d}\phi^2}\right)\right] u_k = 0, \end{aligned} \tag{II.16}$$

$$H_{,N} = -\tfrac{1}{2}H\phi_{,N}^2, \tag{II.17}$$

$$\phi_{,NN} + \left(\frac{H_{,N}}{H} + 3\right)\phi_{,N} + \frac{1}{H^2}\frac{\mathrm{d}V}{\mathrm{d}\phi} = 0, \tag{II.18}$$

where “$, N$” denotes a derivative with respect to N, and N_0 determines the normalisation of the scale factor: i.e., at $N = N_0$ we have $a = 1$. This coupled system of differential equations can easily be solved numerically, once a suitable set of initial conditions has been chosen.

1.2. Initial Conditions

Supposing that at a time N_{sr} the system has reached the inflationary attractor solution

$$\ddot{\phi} \ll 3H\dot{\phi}, \tag{II.19}$$

and is rolling slowly,

$$\dot{\phi}^2 \ll V(\phi), \tag{II.20}$$

the initial conditions for ϕ and H will be given by

$$\phi(N_{\mathrm{sr}}) = \phi_{\mathrm{sr}}\,, \tag{II.21}$$

$$\phi_{,N}(N_{\mathrm{sr}}) = -\frac{1}{V(\phi_{\mathrm{sr}})}\left.\frac{\mathrm{d}V}{\mathrm{d}\phi}\right|_{\phi_{\mathrm{sr}}}\,, \tag{II.22}$$

$$H(\phi_{\mathrm{sr}}) = \sqrt{\frac{V(\phi_{\mathrm{sr}})}{3}}. \tag{II.23}$$

In fact, the existence of the attractor is one of the properties that make inflation such an elegant mechanism: no matter what the initial state of the system is, after a few e-foldings, it will have reached the attractor and its dynamics will be described by Equations (II.21)-(II.23).

The initial conditions for u_k can be obtained by requiring the late time solution of (II.11) to match the solution of a field in the Bunch-Davies vacuum of de Sitter space [44], given by

$$u_k(\tau) = \frac{e^{-\mathrm{i}k\tau}}{\sqrt{2k}}\left(1 + \frac{\mathrm{i}}{k\tau}\right), \tag{II.24}$$

at early times, well before the observationally relevant scales leave the horizon. For $k \gg z''/z$ (or, equivalently, $k\tau \gg 1$) this can be approximated by the free field solution in flat space

$$u_k = \frac{1}{\sqrt{2k}}\, e^{-\mathrm{i}k\tau}. \tag{II.25}$$

Fixing the irrelevant phase, we obtain the initial conditions for a mode k

$$u_k(\tau_0) = \frac{1}{\sqrt{2k}}, \tag{II.26}$$

$$u'_k(\tau_0) = -\mathrm{i}\sqrt{\frac{k}{2}}, \tag{II.27}$$

at a time τ_0 satisfying $k \gg z''/z|_{\tau_0}$.

1.3. Slow Roll

Following Ref. [45], we define the nth, ($n \geq 1$), Hubble slow roll parameter by

$$ {}^{n}\beta_H \equiv \left\{ \prod_{i=1}^{n} \left[-\frac{\mathrm{d}\ln H^{(i)}}{\mathrm{d}\ln a} \right] \right\}^{1/n} = 2 \left(\frac{(H^{(1)})^{n-1} H^{(n+1)}}{H^n} \right)^{1/n}, \tag{II.28}$$

with a superscript "(i)" denoting the ith derivative with respect to ϕ. In addition to that, we define ${}^{0}\beta_H \equiv 2(H^{(1)}/H)^2$. The first three parameters of the Hubble slow roll hierarchy read

$$\epsilon_H \equiv {}^{0}\beta_H = 2\left(\frac{H^{(1)}(\phi)}{H(\phi)}\right)^2 = -\frac{\dot{H}}{H^2}, \tag{II.29}$$

$$\eta_H \equiv {}^{1}\beta_H = 2\,\frac{H^{(2)}(\phi)}{H(\phi)} = -\frac{\ddot{\phi}}{\dot{\phi}H}, \tag{II.30}$$

$$\xi_H^2 \equiv ({}^{2}\beta_H)^2 = 4\,\frac{H^{(1)}(\phi)H^{(3)}(\phi)}{H^2(\phi)} = \frac{\dddot{\phi}}{H^2\dot{\phi}} - \eta_H^2. \tag{II.31}$$

Using these definitions it can be shown that the mode equation (II.11) can be written as

$$u''_k + \left(k^2 - 2a^2H^2\left[1 + \epsilon_H - \tfrac{3}{2}\eta_H + \epsilon_H^2 - 2\epsilon_H\eta_H + \tfrac{1}{2}\eta_H^2 + \tfrac{1}{2}\xi_H^2\right]\right)u_k = 0. \tag{II.32}$$

Note that this expression is exact: it does not assume the slow roll parameters to be small.

From a model-building point of view, where one regards the Lagrangian (or the scalar potential) of the theory as the input quantity, the calculation of the Hubble slow roll parameters can be quite involved. In this sense it may be more convenient to work with the potential slow roll parameters instead, which use derivatives of the potential instead of derivatives of the Hubble parameter. The first three potential slow roll parameters are defined by

$$\epsilon \equiv \tfrac{1}{2}\left(\frac{V^{(1)}}{V}\right)^2, \tag{II.33}$$

$$\eta \equiv \frac{V^{(2)}}{V}, \tag{II.34}$$

$$\xi^2 \equiv \frac{V^{(1)}V^{(3)}}{V^2}. \tag{II.35}$$

If the attractor condition (Equation (II.19)) is satisfied, the two are related via [45]

$$\epsilon_H = \epsilon - \tfrac{4}{3}\epsilon^2 + \tfrac{2}{3}\epsilon\eta + \mathcal{O}_3, \tag{II.36}$$

$$\eta_H = \eta - \epsilon + \tfrac{8}{3}\epsilon^2 + \tfrac{1}{3}\eta^2 - \tfrac{8}{3}\epsilon\eta + \tfrac{1}{3}\xi^2 + \mathcal{O}_3, \tag{II.37}$$

$$\xi_H^2 = \xi^2 - 3\epsilon\eta + 3\epsilon^2 + \mathcal{O}_3, \tag{II.38}$$

up to corrections of third and higher orders, denoted by $\mathcal{O}_3$. Expressed in terms of the potential slow roll parameters, z''/z is given by

$$\frac{z''}{z} = 2a^2H^2\left[1 + \tfrac{5}{2}\epsilon - \tfrac{3}{2}\eta - \tfrac{1}{3}\epsilon^2 + \tfrac{1}{6}\epsilon\eta + \mathcal{O}_3\right]. \tag{II.39}$$

It is commonly assumed that the first two slow roll parameters vary slowly with time (i.e., $\xi^2_{(H)} \ll 1$). Then it follows that, if one wants to sustain inflation for long enough to solve the horizon and flatness problems, $\epsilon_{(H)}$ and $|\eta_{(H)}|$ will also have to be much smaller than unity. In this ("slow roll") limit, we have $z''/z \approx 2a^2H^2$, $\dot{H} \approx 0$ and $a \propto \exp[Ht]$.

Let us now turn back to Equation (II.11), which is basically the equation of an oscillator with a time dependent mass term, and discuss its solutions. The initial conditions imply that for wavenumbers with $k/a \gg H$, i.e., with wavelengths much smaller than the horizon, the solution is given by Equation (II.25) and u_k describes a circular motion in the complex plane. Due to the exponential growth of the scale factor, the physical wavelengths will be blown up and leave the horizon, eventually satisfying $k/a \ll H$. In this limit, the solution for u_k is given by

$$u_k \simeq C_1 z + C_2\, z \int \frac{\mathrm{d}\tau}{z^2}, \tag{II.40}$$

with integration constants C_1 and C_2. The second piece represents the decaying mode and can be neglected. Plugging $u_k \propto z$ into Equation (II.13), we see that the spectrum $\mathcal{P}_\mathcal{R}$ will converge to a constant value for super-Hubble modes, i.e., the perturbations "freeze in". We can also conclude that the fate of a perturbation with wavelength k is decided when $k/a \sim H$ and the spectrum will have its final shape imprinted on horizon exit.

During the reheating phase after the end of inflation, when the potential energy of the inflaton is converted to particles again, the perturbations in the inflaton field will be transferred to perturbations in the fluids making up the later Universe, such as photons, baryons or dark matter. It can in fact be shown that super-horizon modes remain constant not only during inflation, but also in a radiation or matter dominated Universe. So it is not until much later, when the modes reenter the horizon during radiation or matter domination, that they will exhibit dynamical behaviour again. In this way, inflation generates the initial conditions for the subsequent evolution of the perturbations, and the asymptotic value of $\mathcal{P}_\mathcal{R}(k)$ is the input quantity for Equation (I.19).

But what will the spectrum look like? Generically, it will not be exactly scale independent, with a scale dependence being induced by the variation of, e.g., the potential energy and the Hubble parameter as the inflaton field rolls down the potential. In the slow roll regime, however, the scale dependence is rather weak and one might want to approximate $\mathcal{P}_\mathcal{R}$ by a power-law (i.e., a linear approximation in log-space around a pivot scale k_0):

$$\mathcal{P}_\mathcal{R}(k) \simeq A_\mathrm{S} \left(\frac{k}{k_0}\right)^{n_\mathrm{S}-1}, \tag{II.41}$$

with the normalisation A_S given by

$$A_\mathrm{S} \simeq \frac{1}{24\pi^2} \frac{V}{\epsilon}\bigg|_{k_0=aH}, \tag{II.42}$$

and the spectral index

$$n_\mathrm{S} - 1 = \frac{\mathrm{d}\ln\mathcal{P}_\mathcal{R}}{\mathrm{d}\ln k}\bigg|_{k=k_0} \simeq -6\epsilon + 2\eta. \tag{II.43}$$

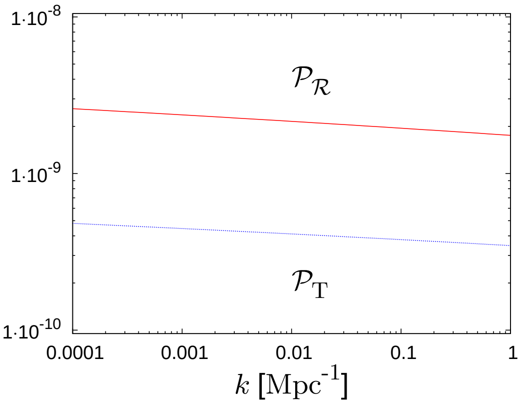

Figure II.1.: This plot shows the exact solution for the spectrum of curvature perturbations (solid red line) and tensor perturbations (dotted blue line) for inflation with a potential $V = \frac{1}{2}m^2\phi^2$, with $n_\mathrm{S} \simeq 0.96$ and $n_\mathrm{T} \simeq 0.98$.

As an example, we plot the spectrum of a chaotic inflation model with $V = \frac{1}{2}m^2\phi^2$ in Figure II.1. The validity of the power-law approximation can be estimated by looking at the derivative of n_S,

$$\alpha_\mathrm{S} \equiv \frac{\mathrm{d}n_\mathrm{S}}{\mathrm{d}\ln k} \simeq 16\epsilon\eta - 24\eta^2 - 2\xi^2. \tag{II.44}$$

For the vast majority of viable particle-physics inspired inflationary models, the variation of the slope over the observationally relevant scales is indeed negligible (see, e.g., Ref. [46] for an extensive overview), so unless one were to look for particular signatures of more exotic inflationary effects, this parameterisation would do just fine for performing data analysis.

2. Tensor Perturbations

So far, we have focused our attention exclusively on scalar perturbations of the metric. Let us now turn to tensor perturbations or gravity waves. Their treatment is somewhat easier, since we do not need to worry about gauge issues; tensor perturbations are gauge invariant by construction. Otherwise, we can proceed just like in the scalar case.

Tensor linear perturbations to the metric (II.2) can most generally be expressed as

$$\mathrm{d}s^2 = a^2(\tau)\left[\mathrm{d}\tau^2 - (\delta_{ij} + 2h_{ij})\,\mathrm{d}x^i\mathrm{d}x^j\right]. \tag{II.45}$$

The tensor h_{ij} is symmetric $h_{ij} = h_{ji}$, traceless $\delta^{ij}h_{ij} = 0$ and transverse $\partial^i h_{ij} = 0$ and hence represents the two degrees of freedom associated with the graviton. Since the energy-momentum tensor of the inflaton field perturbations is assumed to be stress-free, there will be no source term for gravity waves, so we can take the Einstein-Hilbert action for the unperturbed part:

$$\mathfrak{S}^\mathrm{T} = -\tfrac{1}{2}\int \mathrm{d}^4x\,\sqrt{-g}\,R. \tag{II.46}$$

Inserting the perturbed metric (II.45) yields the perturbed action

$$\mathfrak{S}^\mathrm{T}_\mathrm{pert} = \tfrac{1}{8}\int \mathrm{d}\tau\,\mathrm{d}^3x\;a^2\left[(h'_{ij})^2 - \partial_l h^{ij}\,\partial^l h_{ij}\right]. \tag{II.47}$$

Before we can quantise, we should make sure that we actually have a canonically normalised field in the action. This can be accomplished by defining $v_{ij} \equiv \frac{a}{2}\,h_{ij}$ and eliminating h_{ij} in the equations of motion. We can then deduce the action in terms of v_{ij}, which is equivalent, up to total derivatives, to the action (II.47):

$$\mathfrak{S}^\mathrm{T}_\mathrm{pert} = \int \mathrm{d}\tau\,\mathrm{d}^3x\;\left[\tfrac{1}{2}\left((v'_\lambda)^2 - \partial_l v^{ij}\,\partial^l v_{ij}\right) - \tfrac{1}{2}\,\frac{a''}{a}v_\lambda^2\right]. \tag{II.48}$$

This is indeed the action of a canonically normalised scalar field v_{ij} with a time-dependent mass term. We can now quantise v_{ij}:

$$\hat{v}_{ij}(\tau,\boldsymbol{x}) = \frac{1}{(2\pi)^{3/2}}\int \mathrm{d}^3k\left\{\hat{\mathfrak{a}}_k\,(v_k)_{ij}(\tau)\,e^{\mathrm{i}\boldsymbol{k}.\boldsymbol{x}} + \hat{\mathfrak{a}}^\dagger_k\,(v_k)^*_{ij}(\tau)\,e^{-\mathrm{i}\boldsymbol{k}.\boldsymbol{x}}\right\}. \tag{II.49}$$

One will typically express the two degrees of freedom of $(v_k)_{ij}$ in terms of two polarisations $\lambda \in \{+, \times\}$

$$(v_k)_{ij} = v_k^+ \, e_{ij}^+(\boldsymbol{k}) + v_k^\times \, e_{ij}^\times(\boldsymbol{k}), \tag{II.50}$$

with polarisation tensors e^+, $e^\times$ satisfying

$$\begin{gathered} e_{ij} = e_{ji}, \quad k^i e_{ij} = 0, \quad e_{ii} = 0, \\ e_{ij}(-\boldsymbol{k}, \lambda) = e_{ij}^*(\boldsymbol{k}, \lambda), \quad e_{ij}^*(\boldsymbol{k}, \lambda) \, e^{ij}(\boldsymbol{k}, \lambda') = \delta_{\lambda'}^{\lambda}. \end{gathered} \tag{II.51}$$

Again, we can use the two-point correlation function

$$\sum_\lambda \langle h_{k,\lambda} h_{k',\lambda}^* \rangle = \frac{2\pi^2}{k^3} \mathcal{P}_\mathrm{T}(k) \; \delta^{(3)}(\boldsymbol{k} - \boldsymbol{k}'), \tag{II.52}$$

to define the primordial power spectrum of gravitational waves, which is given by

$$\mathcal{P}_\mathrm{T}(k) = 2 \times \frac{k^3}{2\pi^2} \left| h_k \right|^2 = 8 \, \frac{k^3}{2\pi^2} \left| \frac{v_k}{a} \right|^2, \tag{II.53}$$

the factor of 2 arising from the two polarisations. The solution for v_k can be obtained from its equation of motion, which follows from the action (II.48)

$$v_k'' + \left(k^2 - \frac{a''}{a} \right) v_k = 0, \tag{II.54}$$

with initial conditions analogous to the ones for u_k, discussed in Section 1.2 above. This equation bears a close resemblance to its scalar counterpart (II.11), so it should not come as a surprise that the behaviour of the resulting spectrum will in general be very similar as well. This can be readily seen if we express the mass term a''/a in terms of the slow roll parameters:

$$\begin{aligned} \frac{a''}{a} &= 2a^2 H^2 \left[1 - \tfrac{1}{2}\epsilon_H \right] \\ &\simeq 2a^2 H^2 \left[1 - \tfrac{1}{2}\epsilon + \tfrac{2}{3}\epsilon^2 - \tfrac{1}{3}\epsilon\eta + \mathcal{O}_3 \right]. \end{aligned} \tag{II.55}$$

So in the slow roll regime, we have $a''/a \approx 2a^2H^2 \approx z''/z$, the tensor perturbations will freeze in at horizon exit, just like the scalar modes and we would also expect an almost scale independent spectrum

$$\mathcal{P}_\mathrm{T}(k) \simeq A_\mathrm{T} \left(\frac{k}{k_0} \right)^{n_\mathrm{T}}, \tag{II.56}$$

with the tensor spectral index given by [47]

$$n_\mathrm{T} \simeq -2\epsilon, \tag{II.57}$$

and normalisation

$$A_\mathrm{T} \simeq \frac{2}{3\pi^2} V \bigg|_{k_0 = aH}. \tag{II.58}$$

2.1. Tensor to Scalar Ratio and the Consistency Relation

Instead of the normalisation of the tensor spectrum, one often finds a derived quantity in the literature, the tensor to scalar ratio. Unfortunately, there are several different definitions floating around, which can be a source of confusion. From the model-builder's point of view, a natural definition would be

$$r = \left.\frac{\mathcal{P}_{\mathrm{T}}(k_0)}{\mathcal{P}_{\mathcal{R}}(k_0)}\right|_{k_0=aH} . \tag{II.59}$$

Note that, apart from the unlikely event that $n_{\mathrm{S}} = n_{\mathrm{T}}$ holds, r will be mildly scale dependent, so if one quotes a value for r, it should always be accompanied by the pivot scale k_0 at which it is evaluated. From Equations (II.42) and (II.58) it follows that

$$r \simeq 16\epsilon|_{k_0=aH} , \tag{II.60}$$

so the contribution of gravitational waves in slow roll inflation is expected to be subdominant.

In a more observer-friendly oriented approach one could take the ratio of the tensor and scalar contributions to the CMB temperature anisotropies at a given multipole ℓ (usually at large scales $\ell \lesssim 10$)

$$r_\ell = \frac{\mathcal{C}_\ell^{\mathrm{T}}}{\mathcal{C}_\ell^{\mathrm{S}}}. \tag{II.61}$$

This however also requires knowledge about the temperature anisotropy source function and is thus not independent of other cosmological parameters, particularly Ω_Λ, due to the integrated Sachs-Wolfe effect which affects only the scalar perturbations. Assuming scale invariant spectra and $\Omega_\Lambda = 0$, Ref. [12] gives $r_4 \simeq 13.8\epsilon$. This relation asymptotes to $r_\ell \simeq 12.4\epsilon$ for larger multipoles, but under more realistic conditions, i.e., $\Omega_\Lambda \neq 0$ or scale dependent spectra, these numbers should be taken with care.

Now if we take the third and higher derivatives of the potential to be zero, we are left with three free parameters (e.g., V, ϵ and η) on the one side, opposed to four (potentially) observable quantities (say, A_{S}, n_{S}, A_{T} and n_{T}). We can exploit this by eliminating one of these variables, for instance,

$$n_{\mathrm{T}} = -\frac{A_{\mathrm{T}}}{8A_{\mathrm{S}}}. \tag{II.62}$$

This equation is known as the slow roll single field *consistency relation.* An experimental verification of this relation, however, will be extremely challenging, at the least. At present, there is no observational evidence for inflationary gravitational waves (or any other kinds of gravitational waves, for that matter). The most promising avenue for their detection is to look for traces of tensor perturbations in the power spectrum of the CMB B-mode polarisation B-mode. The PLANCK satellite [48] will likely be able to make a positive detection if r is not much smaller than 0.1 [49]. Even under the most optimistic assumptions, though, tensor to scalar ratios of $r \lesssim 10^{-4}$ may be undetectable [50]. So if inflation took place at a low scale, we may never be able to check the consistency relation. Of course, even if we did detect a primordial gravitational wave signal, determining the tilt of its spectrum to sufficient accuracy to make a meaningful statement about whether the consistency relation holds, may not be possible either [51].

It is likely that only in the case of a gross violation will we be able to pass a final verdict on the consistency relation. We will return to the issue of how to decide whether a statement such as “the consistency relation holds” should be refuted or accepted, in the context of model selection in Section III.4.

What if it really did turn out that equation (II.62) is violated in Nature? We made several simplifying assumptions in its derivation, some of which may not be true. It might be that more than one field plays an important rôle during inflation, or perhaps our choice of initial conditions was not correct and the Universe was not in the Bunch-Davies vacuum when the relevant scales left the horizon [52]. Another possibility, which will be pursued in more detail in Chapter V, is that the higher derivatives of the potential do give a non-negligible contribution to the inflaton dynamics, possibly even violating the slow roll conditions, if only for an instant.

III. Data Analysis and Statistics in Cosmology

Having seen in the previous chapter how the inflaton potential determines the form of the primordial perturbation power spectra, we will now explore the connection between the theoretical predictions and actual data.

Let us ask more generally: Given a data set, how can we determine the free parameters of a theory and give reasonable error estimates? And how can we decide whether a theory is compatible with the data or should be considered ruled out? The answers to these questions involve a few subtleties and great care has to be taken to interpret the results correctly.

The purpose of this section is to familiarise the reader with the statistical concepts and techniques used for the analysis of cosmological data. We will begin with the discussion of an almost trivial example, the direct measurement of a single quantity, which serves to introduce the terminology and illustrates some of the conceptual differences between the two different approaches to statistics, the Bayesian and the Frequentist one. Which of these two methods should be used is a subject of fierce philosophical debate among statisticians.

Adopting the Bayesian standpoint, we will discuss how to infer the values of the free parameters of a given model from the data. In practice, this procedure is complicated by the fact that data sets are, in general, sensitive to variations of more than just one single parameter. We will see that instead one needs to consider probability densities in multi-dimensional parameter spaces. This is problematic in two ways: One needs to find ways to condense information about the function in the higher-dimensional space to lower-dimensional, “human-readable” output, a process known as marginalisation. On the other hand, one needs to gather this information in the first place, which turns out to be a computationally very demanding process.

Fortunately, we have a very elegant and flexible solution at hand to deal with these difficulties: Markov Chain Monte Carlo (MCMC) algorithms. The third section of this chapter will be devoted to a discussion of their practical application.

In the last section, we will briefly discuss how to compare models in the framework of Bayesian statistics.

1. Frequentist or Bayesian?

As a warm-up exercise, let us follow the example given in Ref. [53]. Consider an apparatus which measures the unknown mass of an elementary particle. We are not concerned about any technical details; let us just imagine the apparatus as a black box. A measurement consists of inserting the particle into the box and reading off the measured mass m from a display. If we want to get a meaningful result in the end, we need to make sure that the apparatus is carefully calibrated by measuring the mass of test particles whose true mass m_t is known beforehand. We will assume that our machine has been optimally calibrated, so we know the probability density of reading off a mass m, given that the true mass of the particle is m_t, this will be denoted by $\mathcal{L}(m|m_\mathrm{t})$. In our hypothetical machine, let this probability density function (pdf) be given by a Gaussian centred around m_t, with a width $\sigma_m \ll m_\mathrm{t}$:

$$\mathcal{L}(m|m_\mathrm{t}) = \frac{1}{\sqrt{2\pi\sigma_m^2}}\, e^{-(m-m_\mathrm{t})^2/2\sigma_m^2}. \tag{III.1}$$

We will call $\mathcal{L}$ the likelihood function.

After performing the measurement, we would like to give a point estimate $\hat{m}$ for the mass of the particle and a confidence interval (m_1, m_2), that quantifies our uncertainty. Let us see how this is done in the classical (Frequentist) picture.

1.1. The Classical View

In the above example, it would appear natural to take $\hat{m}$ as the value that maximises $\mathcal{L}(m|m_\mathrm{t})$, this is known as the maximum likelihood estimator.

For the definition of a classical confidence interval [54], remember the definition of Frequentist probability:

Definition 1 In a repeatable experiment with a sample space S of outcomes A_i, the probability of attaining outcome A_i is given by the limiting frequency

$$P(A_i) \equiv \lim_{n\to\infty} \frac{\#(A_i)}{n}, \tag{III.2}$$

where n is the number of times the experiment is repeated and $\#(A_i)$ the number of occurrences of outcome A_i.

These outcomes could, for instance, be taken to be hypotheses of the form: "The estimator $\hat{m}$ is smaller than the value m_1". For any given measurement, a statement of this kind is either true or false. We can now define a $100\gamma\%$-confidence interval:

Definition 2 The interval $[m_1, m_2]$ is called a $100\gamma\%$-confidence interval, if, out of an ensemble of identical experiments, a fraction of γ finds the hypothesis: "The estimator $\hat{m}$ satisfies $m_1 < \hat{m} < m_2$" to be true.

The above definition does not lead to a unique confidence interval. Two common choices for confidence intervals are

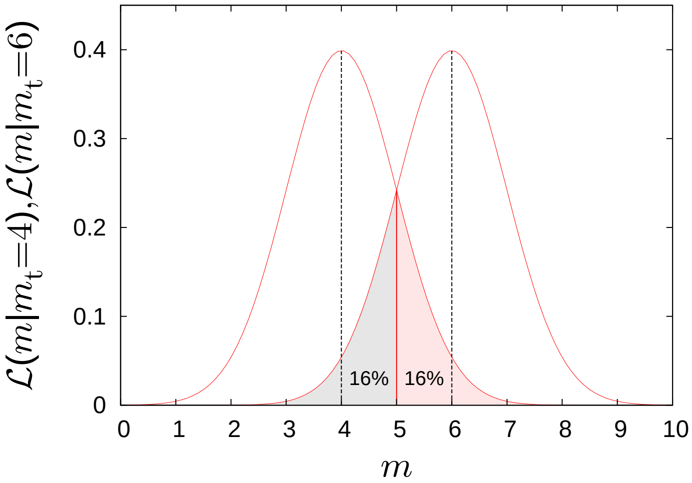

Figure III.1.: Frequentist construction of the central 68%-confidence interval, after a single measurement with $\sigma_m = 1$ yields $\hat{m} = 5$. If the true mass was $m_t = 4$, one would find a result $\hat{m} > 5$ in 16% of all repeated realisations of the experiment (left curve, red shaded area). If, on the other hand, the true mass was $m_t = 6$, then 16% of all experiments would find a result $\hat{m} < 5$ (right curve, grey shaded area). Hence, the central 68%-confidence interval is given by $(4 \to 6)$.

- Central confidence interval
  Here, a fraction of $(1-\gamma)/2$ of an ensemble of experiments would find $\hat{m} < m_1$ and a fraction of $(1-\gamma)/2$ would likewise find $\hat{m} > m_2$.

- Minimal confidence interval
  Here, m_1 and m_2 are chosen such that $m_2 - m_1$ is minimised.

For the likelihood function in our example the two happen to coincide, the central interval is also minimal. In Figure III.1 we illustrate the construction of the central 68%-confidence interval.

A commonly used approximate method of finding the 68%-confidence interval is sketched in Figure III.2. Based on the method of least squares, one can define an effective χ^2:

$$\chi^2_{\text{eff}} = -2\ln\left(\mathcal{L}/\mathcal{L}_{\max}\right). \tag{III.3}$$

In the special case of a Gaussian likelihood pdf this coincides with the classical definition of χ^2 [55]. Given a measurement resulting in an estimate $\hat{m}$, one could take as upper and lower limits of the 68%-confidence interval those values of m_t, for which $\Delta\chi^2_{\text{eff}} \equiv \chi^2_{\text{eff}}(m_t) - \chi^2_{\text{eff}}(\hat{m}) = 1$. Note the shift in interpretation here; while $\mathcal{L}$ is typically regarded as a function of m with m_t constant, in χ^2_{eff}, one keeps m constant and varies m_t.

We stress that this prescription does not in general lead to an exact confidence interval in the sense of Definition 2. Generally, this approach will only be accurate in the case of

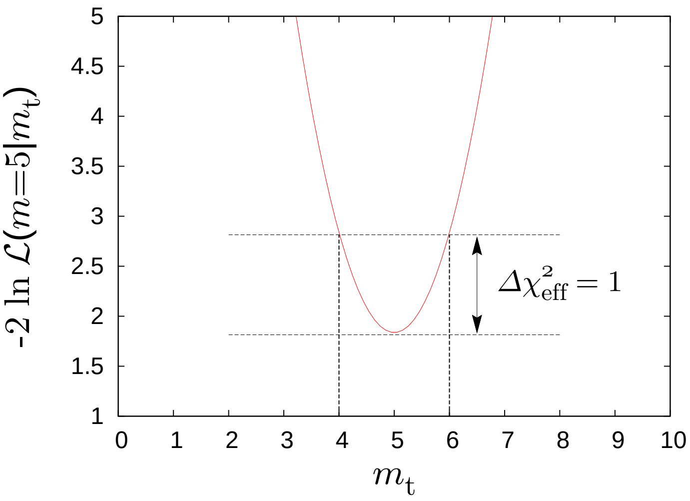

Figure III.2.: Approximate construction of the 68%-confidence interval with the $\Delta\chi^2 = 1$ method.

a Gaussian likelihood pdf. Given that it is often numerically easier to construct limits in this way and the fact that, by virtue of the central limit theorem, most likelihoods one would encounter in physical applications do resemble a Gaussian to a fair degree, this method is quite popular. Care should be taken, however, if one has to deal with more unusual likelihood functions.

In our example, though, the two thus derived intervals agree. Having found a lower limit of $\hat{m} - \sigma_m$ and an upper limit of $\hat{m} + \sigma_m$ one conventionally states that the measured mass is "$\hat{m} \pm \sigma_m$". Intuitively, one might be tempted to imagine this result as a pdf in the variable m_t, centred around $\hat{m}$, as depicted in Figure III.3. In the Frequentist definition of probability however, this is *not* permitted!

For a Frequentist, the true mass of a particle is not a random variable, either it falls in a given interval or it does not. What can be varied in this picture are the limits of the confidence interval, the true mass remains fixed. In fact, in the construction of the classical confidence intervals, the use of probability densities in the variable m_t was carefully avoided, for such a quantity would simply be ill-defined.[1]

This is where the Bayesian ansatz comes into play. It can be shown that there exists an axiomatic, mathematically consistent way to extend the definition of probabilities to accommodate *degrees of belief* [56]. This allows us to make probabilistic statements about propositions that do not involve random variables, such as, e.g., "Will it rain tomorrow?", or, applied to our problem, "What is the probability of the true mass of the particle being between m_1 and m_2 given that our apparatus displays a value of $\hat{m}$?".

To a Frequentist, a hypothesis can either be true or false. A Bayesian could say that it is probably true (or false), with a probability P quantifying the degree of belief. This

---

[1] While in χ^2, the true mass is taken to be variable, it is not interpreted as a pdf in m_t!

increased flexibility does not come for free though: as the “belief” part of the preceding sentence implies, a Bayesian statement relies on certain assumptions on the part of the statistician. The price we have to pay is a loss of objectivity, a point that has raised a lot of criticism.

We have seen that the interpretation of Frequentist confidence levels relies on the idea that an experiment can be repeated a large number of times under the same circumstances. Situations, in which these additional measurements could not even in principle be made, may call for a somewhat more pragmatic approach.

This would seem to explain why most particle physicists adhere to the Frequentist philosophy, while in cosmology, it is not so popular: We do not have an ensemble of Universes to observe, there is only one realisation!

1.2. What would the Reverend say?[2]

Let us turn back to our example. In the Bayesian picture, we can now make sense of a probability distribution $\mathcal{P}(m_\mathrm{t}|m)$ of the true mass m_t, given that we have measured a mass m.

Invoking an obvious relation for conditional probabilities known as *Bayes' theorem*

$$P(A|B) \cdot P(B) = P(B|A) \cdot P(A), \tag{III.4}$$

and exploiting the law of total probability

$$P(B) = \sum_i P(B|A_i)P(A_i), \tag{III.5}$$

we can determine $\mathcal{P}$ from the likelihood via

$$\mathcal{P}(m_\mathrm{t}|m) = \frac{\mathcal{L}(m|m_\mathrm{t})\pi(m_\mathrm{t})}{\mathcal{E}(m)} = \frac{\mathcal{L}(m|m_\mathrm{t})\,\pi(m_\mathrm{t})}{\int \mathrm{d}m_\mathrm{t}\;\mathcal{L}(m|m_\mathrm{t})\pi(m_\mathrm{t})}\,, \tag{III.6}$$

in plain text:

$$\text{Posterior} = \frac{\text{Likelihood} \cdot \text{Prior}}{\text{Evidence}}\,. \tag{III.7}$$

In order to define the posterior probability, we had to introduce a new quantity, the prior probability $\pi(m_\mathrm{t})$. It is at this point that the statistician's prejudice enters, $\pi(m_\mathrm{t})$ describes what we know about the parameter *prior* to taking data. Let us defer a more detailed discussion about the choice of priors to the next section and assume for now that there is no theoretical bias for particular values of m_t. Thus, one might want to assign

$$\pi(m_\mathrm{t}) = \begin{cases} (m_\mathrm{max} - m_\mathrm{min})^{-1} & \text{if } m_\mathrm{max} > m_\mathrm{t} > m_\mathrm{min} \\ 0 & \text{otherwise} \end{cases} \tag{III.8}$$

---

[2] Reverend Thomas Bayes, (1702-1761)

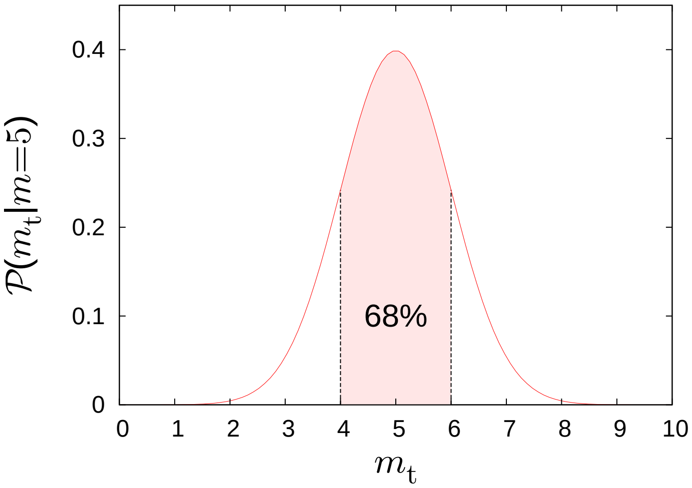

Figure III.3.: Bayesian 68%-confidence interval with a flat prior on m_t.

with the limits of the support of π chosen sufficiently large or small, respectively. This is known as a flat, or top hat prior on m_t.

Applied to our example, imposing a flat prior on m_t will lead to a Gaussian posterior pdf:

$$\mathcal{P}(m_\mathrm{t}|m) = \frac{1}{\sqrt{2\pi\sigma_m^2}}\, e^{-(m_\mathrm{t}-m)^2/2\sigma_m^2}\;. \tag{III.9}$$

Once we know the posterior pdf, there is a very natural way to define a Bayesian 100γ%-confidence interval: it is simply an interval I_γ that contains a fraction γ of the total area under the posterior pdf,

$$\int\limits_{I_\gamma} \mathrm{d}m_\mathrm{t}\, \mathcal{P}(m_\mathrm{t}|m) = \gamma. \tag{III.10}$$

Once again, this does not distinguish a unique interval. Like in the Frequentist case, we can for instance choose a central or a minimum length interval. We exemplify the construction of the central 68%-confidence interval in Figure III.3.

For the special case of a Gaussian likelihood pdf and the choice of a flat prior, the limits we derived with different methods agree with each other. In general, however, one cannot expect this to be the case. This will become clearer when we move to the discussion of multivariate likelihood pdfs in the next section.

2. Inference in Cosmology

In the above example, we assumed that there was no dependence of the result of our measurement on parameters other than the mass of the particle and that we could determine the mass directly. In cosmology there are dozens of parameters of interest, but the vast majority of them are not directly accessible in such a fashion. Notable exceptions to this rule include for instance the temperature of the CMB and the photon energy density Ω_γ.

Estimates for the values of other parameters will have to be inferred from the measurements. In the following we will outline how this can be accomplished. Let us rewrite Equation (III.6) in a more general way:

$$\mathcal{P}(\boldsymbol{\theta}|\mathcal{D},\mathcal{M}) = \frac{\mathcal{L}(\mathcal{D}|\boldsymbol{\theta},\mathcal{M}) \cdot \pi(\boldsymbol{\theta}|\mathcal{M})}{\mathcal{E}(\mathcal{D}|\mathcal{M})} = \frac{\mathcal{L}(\mathcal{D}|\boldsymbol{\theta},\mathcal{M}) \cdot \pi(\boldsymbol{\theta}|\mathcal{M})}{\int \mathrm{d}\boldsymbol{\theta}\ \mathcal{L}(\mathcal{D}|\boldsymbol{\theta},\mathcal{M}) \cdot \pi(\boldsymbol{\theta}|\mathcal{M})}\,, \tag{III.11}$$

where $\mathcal{D}$ represents a given data set and $\mathcal{M}$ is the theoretical model we use to describe the data, dependent on a set of free parameters $\boldsymbol{\theta}$. For instance, $\mathcal{D}$ could be a measurement of the temperature anisotropies of the CMB, $\mathcal{M}$ a set of rules to determine the theoretical predictions for the fluctuations (e.g., the concordance model) from the standard six cosmological parameters $\boldsymbol{\theta} = (\Omega_\mathrm{b}, \Omega_\mathrm{c}, H_0, \tau, A_\mathrm{S}, n_\mathrm{S})$.

At this point, a few comments about the constituents of Equation (III.11) are in order:

- **Evidence** $\mathcal{E}(\mathcal{D}|\mathcal{M})$
  By definition, the evidence is the probability to measure the data $\mathcal{D}$ given that the underlying model $\mathcal{M}$ is true. It is not dependent on the parameters, since they are integrated over. So for the purpose of parameter estimation, one can regard it as a mere normalisation constant of the posterior and its actual value is not relevant. It does, however, play an important part in determining the quality of $\mathcal{M}$ with respect to other models. We will return to this point in Section 4 below.

- **Prior** $\pi(\boldsymbol{\theta}|\mathcal{M})$
  The prior probability represents our personal belief in what the true value of the parameters $\boldsymbol{\theta}$ should be, before any data is taken (it is independent of $\mathcal{D}$!), under the condition that the model is correct. The choice of prior is a subjective process and one might object that scientific results should depend solely on facts, not on opinions. We saw that for the choice of a flat prior, Bayesian confidence intervals agree with their Frequentist counterparts.

  In practical applications, most people advocate a pragmatic approach of "letting the data decide"; as long as the prior is reasonably close to constant over the parameter ranges where the bulk of the volume of the posterior pdf lies, the resulting confidence intervals will not be terribly sensitive to the exact choice of prior. This practice is quite common, but not entirely without pitfalls. Unlike the Frequentist construction, Bayesian confidence intervals are not invariant under transformations of the parameters. Returning to the mass measurement example, if we found a 68%-confidence interval for m_t given by $[m_1, m_2]$ with a flat prior on m_t, the corresponding 68%-interval for a derived parameter which is related to m_t by a nonlinear function f, will in general *not* be given by $[f(m_1), f(m_2)]$. Imposing a flat prior on $m_{\rm t}$ gives a non-flat prior on $f(m_{\rm t})$. Frequentist confidence intervals, on the other hand, do satisfy this property.

What is to decide which is the fundamental parameter that should be assigned a flat prior? Should it be $m_{\rm t}$, or rather $m_{\rm t}^2$, or possibly even $\ln m_{\rm t}$?

This choice is not always obvious, see for instance Ref. [57] for an application where the parameterisation affects the results.

But of course, not in all situations would one want to choose a non-informative prior. In fact the subjectivity of Bayesian analysis could be considered a blessing as well as a curse. It allows us to incorporate external knowledge into the analysis, which is not connected to the experiment. This could range from theoretical input (e.g., "$m_{\rm t}$ is a mass and it better be positive, or else...") to information garnered from independent experiments in the form of limits on individual parameters.

- **Likelihood $\mathcal{L}(\mathcal{D}|\boldsymbol{\theta},\mathcal{M})$**
  The likelihood pdf quantifies the relative agreement of the data with the theoretical predictions made by the model $\mathcal{M}$ with fixed parameters $\boldsymbol{\theta}$. A more technical discussion of how to actually calculate $\mathcal{L}$ given a data set $\mathcal{D}$ shall be relegated to Appendix A.2.

- **Posterior $\mathcal{P}(\boldsymbol{\theta}|\mathcal{D},\mathcal{M})$**
  This probability density function combines the information from the data with our prior guesses about the values of the parameters. Naturally, if we had to quote a point estimate of the parameter vector $\boldsymbol{\theta}$, it would be the one that maximises $\mathcal{P}$. And of course, it is also the quantity used to derive limits on the parameters.

Given the posterior pdf in the D-dimensional parameter space, it is now straightforward to generalise the one-dimensional concept of a Bayesian confidence interval and define a $100\gamma\%$ confidence level hypersurface:

Definition 3 A closed (but not necessarily connected) hypersurface ∂A_γ is called $100\gamma\%$ confidence level hypersurface if

$$\int\limits_{A_\gamma} \mathrm{d}\boldsymbol{\theta}\, \mathcal{P}(\boldsymbol{\theta}|\mathcal{D},\mathcal{M}) = \gamma. \tag{III.12}$$

Unless explicitly stated otherwise, we will always assume that ∂A_γ is chosen such that the volume A_γ is minimal. This is equivalent to demanding that the posterior probability of any point within the confidence hypersurface is larger than that of any point outside of it.

Unfortunately, it may be a bit difficult to visualise this $D-1$-dimensional hypersurface if $D > 3$. This deficiency can be alleviated by mapping $\mathcal{P}$ to a lower-dimensional subspace, a process known as marginalisation. The interpretation of $\mathcal{P}$ as a probability density in the Bayesian picture suggests the following definition of a marginalised probability:

$$P_{\text{marge}}(\tilde{\boldsymbol{\theta}}) \propto \int \mathrm{d}\boldsymbol{\theta}\, P(\boldsymbol{\theta})\, \delta(h(\boldsymbol{\theta}) - \tilde{\boldsymbol{\theta}}), \tag{III.13}$$

where h is usually taken to be a projection operator on the subspace spanned by $\tilde{\boldsymbol{\theta}}$. This process favours regions in the marginalised directions of parameter space that contain a large volume of the probability density over those that may have a sharp maximum which only contains a relatively small volume. This can sometimes lead to somewhat counterintuitive results, like for instance in the example given in Figure III.4 below. As is generally the case with surjective mappings, one loses information about the original function in the process of marginalisation, here in particular about the shape of the distribution in the marginalised directions. Under certain conditions, it may therefore be enlightening to quote additional quantities, that may preserve better other aspects of the original function, such as for instance the projection of P onto the subspace spanned by the θ_i:

$$P_{\text{proj}}(\tilde{\boldsymbol{\theta}}) \propto \max \left[P(\boldsymbol{\theta}) \int \mathrm{d}\boldsymbol{\theta}\, \delta(h(\boldsymbol{\theta}) - \tilde{\boldsymbol{\theta}}) \right] . \tag{III.14}$$

In practice, P_{proj} can sometimes be hard to evaluate, therefore one frequently states a *mean probability* [13] defined by

$$P_{\text{mean}}(\tilde{\boldsymbol{\theta}}) \propto \frac{\int \mathrm{d}\boldsymbol{\theta}\, P^2(\boldsymbol{\theta})\, \delta(h(\boldsymbol{\theta}) - \tilde{\boldsymbol{\theta}})}{\int \mathrm{d}\boldsymbol{\theta}\, P(\boldsymbol{\theta})\, \delta(h(\boldsymbol{\theta}) - \tilde{\boldsymbol{\theta}})} \tag{III.15}$$

instead.

Let us conclude this discussion with a few points that should be kept in mind when interpreting Bayesian results.

- The Bayesian construction will always produce confidence levels, no matter how badly (or well) the data is described by the model. This is in contrast to, e.g., a p-value goodness-of-fit analysis. It may therefore sometimes be helpful to indicate also the χ^2_{eff}/d.o.f. of the best fit point to give a rough idea of the quality of the model. On the other hand, one could argue that it does not help to know that a certain model is unlikely to be true unless one can actually present a better one. This question of comparison of different models can be addressed in a strictly Bayesian manner and we will return to it in the next section.

- The most likely point in parameter space need not necessarily coincide with the best fit point. This can be a consequence of the choice of prior, or, if one uses marginalised posterior pdfs, due to the marginalisation procedure, see for instance Figure III.4.

- Similarly, if a parameter value θ lies outside a certain confidence interval, it does not inevitably mean that it would provide a bad fit to the data. Remember that the confidence levels are derived from marginalised posteriors. As mentioned above, the only thing that the marginalisation cares about is the volume contained in the marginalised directions, values of parameters that give a good fit on average are preferred over parameters that generally fit the data badly but may provide an excellent fit in a tiny corner of parameter space. In other words, the Bayesian analysis rewards predictivity and punishes fine-tuning.

- Looking at Equation (III.11), the observant reader will have noticed that every single term represents a conditional probability with the condition being that the model $\mathcal{M}$ used to make the theoretical predictions for observable quantities is true. Different underlying models, even if they share the same parameters, will generally not lead to the same confidence intervals. The issue of model dependent limits on the standard cosmological parameters will be the subject of discussion in Chapter IV.

- Finally, at the risk of pointing out the obvious, let us emphasise: constraints derived with Bayesian inference should be taken *cum grano salis*. In light of the points raised above and given the dependence on one's choice of model and priors, limits thus inferred are not set in stone and one should not be too quick to declare something "ruled out".

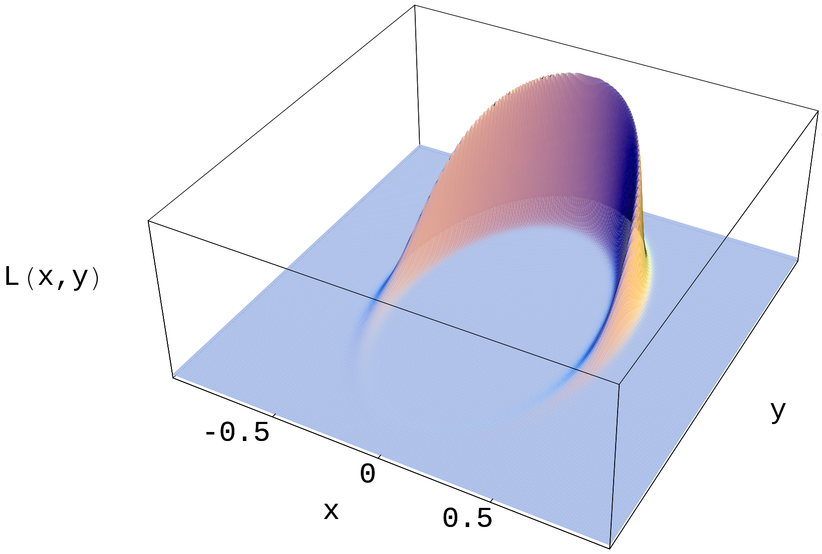

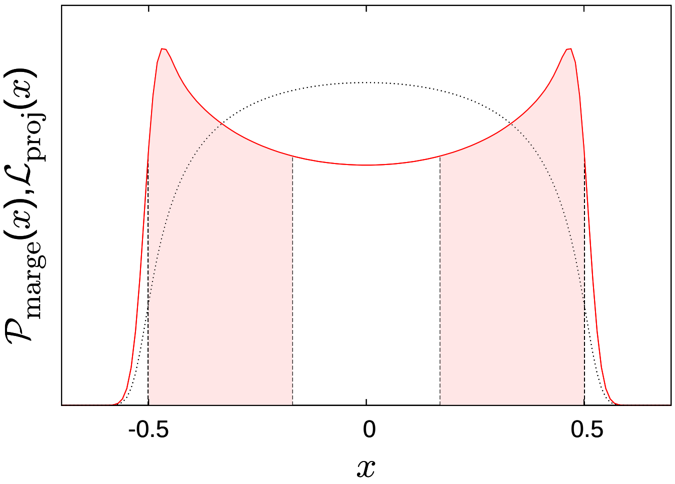

Figure III.4.: This admittedly somewhat unusual example illustrates some of the less intuitive possible consequences of marginalisation. Consider the bivariate likelihood pdf $\mathcal{L}(x, y)$ depicted in the upper diagram and assume flat priors on x and y. If we marginalise over the parameter y, the resulting posterior in x is given by the red solid line in the lower diagram. For comparison, we have also plotted the projection of $\mathcal{L}$ on the x-axis (dotted line). The marginalised posterior is found to be bimodal, the 68%-confidence interval (red shaded area) is not connected, and $x = 0$, the value that maximises the $2D$-likelihood pdf, lies outside this interval. Had we chosen to construct the confidence interval centred around $x = 0$ instead, then the maxima of the marginalised posterior, i.e., the most likely values of x would lie outside the 68%-confidence interval, a similarly paradoxical situation.

3. Markov Chain Monte Carlo Techniques

Having discussed the statistical underpinning of parameter estimation we now face the question of how to calculate the confidence intervals in practice. Of course, before we can think about this, we need to have a good idea about the shape of the posterior pdf. Naïvely, one might consider the following algorithm:

1. Pick a point $\boldsymbol{\theta}$ in parameter space.
2. Calculate the theoretical predictions $\boldsymbol{d}^{\text{th}}$ for the observed quantities.
3. Compare with the observed values $\boldsymbol{d}^{\text{obs}}$ and determine $-2\ln\mathcal{L}(\mathcal{D}|\boldsymbol{\theta})$ (cf. Appendix A.2).
4. Apply the prior to find $-2\ln\mathcal{P}(\boldsymbol{\theta})$, up to an additive constant.
5. Repeat steps 1.-4. on a lattice of points in parameter space.
6. Interpolate between the lattice sites and find the confidence intervals via marginalisation.

While in principle a valid approach, this prescription does not work so well for cosmological data analysis. The numerical calculation of the CMB angular power spectra (step 3.) takes $\mathcal{O}(1)$ second on current computers, so even generating a lattice with a modest resolution of a measly 10 points per parameter direction for the six parameters of the vanilla model would take several days. Considering that the total computation time increases exponentially with the number of dimensions of the parameter space, varying more than six parameters would hardly be feasible.

There is, fortunately, a cleverer way of reconstructing the posterior pdf. It is based on the following idea: $\mathcal{P}$ is a probability density function, so if one could draw a large number of uncorrelated random samples from it, the *density* of these samples would be proportional to $\mathcal{P}$. Ironically, this is a very Frequentist concept.

3.1. The Metropolis Algorithm

A widely used realisation of this idea is the Metropolis algorithm [58] for reconstructing $\mathcal{P}(\boldsymbol{\theta})$ which utilises a random walk ansatz:

1. Pick a starting point $\boldsymbol{\theta}_0$ with $\mathcal{P}(\boldsymbol{\theta}_0) \neq 0$.
2. "Propose the next step":
   Generate a vector $\boldsymbol{\theta}_p$, drawn randomly from a fixed probability density function $Q(\boldsymbol{\theta}_p)$, the *proposal density.*

3. Evaluate the posterior at the proposed point $\mathcal{P}(\boldsymbol{\theta}_i + \boldsymbol{\theta}_p)$, (up to a constant factor, see above).
4. "Decide whether the step should be taken or not":
    - "If the proposed point is more likely, take the step":
      If $\frac{\mathcal{P}(\boldsymbol{\theta}_i+\boldsymbol{\theta}_p)}{\mathcal{P}(\boldsymbol{\theta}_i)} \geq 1$, accept the proposal and set $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i + \boldsymbol{\theta}_p$.
    - "If the proposed point is less likely, roll a dice to decide":
      If $\frac{\mathcal{P}(\boldsymbol{\theta}_i+\boldsymbol{\theta}_p)}{\mathcal{P}(\boldsymbol{\theta}_i)} < 1$, generate a random number ρ from the interval $[0, 1]$.
        - If $\rho > \frac{\mathcal{P}(\boldsymbol{\theta}_i+\boldsymbol{\theta}_p)}{\mathcal{P}(\boldsymbol{\theta}_i)}$, accept the proposal and set $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i + \boldsymbol{\theta}_p$.
        - If $\rho < \frac{\mathcal{P}(\boldsymbol{\theta}_i+\boldsymbol{\theta}_p)}{\mathcal{P}(\boldsymbol{\theta}_i)}$, reject the proposal and set $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i$.
5. Return to step 2. with new starting point $\boldsymbol{\theta}_{i+1}$.

This algorithm generates a series $\{\boldsymbol{\theta}_i\}$, known as a Markov chain. It can be shown that the distribution of points in the chain converges to a static limit and that this limit is given by $\mathcal{P}$. A proof can be found, e.g., in Ref. [59]. Remarkably, the result is independent of the form of the proposal density Q. We have seen that the chain generated thus will converge to the sampling distribution in the limit of an infinite number of steps, but in practice one deals with finite chains. It is hence an important question how to estimate the degree of convergence. Furthermore, for the sake of efficiency, one wants to ensure that convergence is reached as quickly as possible. There exist a number of refinements of the basic algorithm that are motivated by this goal.

The beauty of the MCMC technique lies in the fact that the number of required evaluations of the likelihood grows only roughly linearly with the number of free parameters. Additionally, it is extremely easy to reconstruct the marginalised posteriors, all one has to do is to generate a histogram of the points in the chain, sorted by the parameter one wants to constrain.[3] In general, we can evaluate the expectation value $\langle X \rangle$ of any quantity X via

$$\langle X \rangle = \int \mathrm{d}\boldsymbol{\theta}\, X\, \mathcal{P}(\boldsymbol{\theta}) \simeq \frac{1}{N} \sum_i^N X_i, \tag{III.16}$$

where the chain runs from $i = 1$ to N and X_i is the value of X at the ith point of the chain. Note that the samples generated by the Metropolis algorithm are not statistically independent, consecutive points will be correlated depending on the form of Q and $\mathcal{P}$. As long as the chain has converged, however, this can be ignored for the purposes of inference.

The application of Monte Carlo methods for cosmological parameter estimation was first suggested in Refs. [60, 61] and there exist two independent publicly available packages employing MCMC to analyse cosmological data, `cosmomc` [13] and `AnalyzeThis`

---

[3] To be more precise, one can find the marginalised posteriors up to a constant normalisation factor, which is irrelevant for constructing the confidence intervals.

[62], both conveniently equipped with their respective Boltzmann codes for generating CMB and matter power spectra, `camb` [63] and `cmbeasy` [14].

3.2. Technical Issues

Let us mention some of the practical details one will encounter when applying MCMC methods.

- Proposal density:
  The choice of the proposal density crucially affects the efficiency of the algorithm. It is common to take Q as a multivariate Gaussian with widths depending on the shape of the posterior pdf. If the distribution is too wide, the average step size will be large and it is likely that the proposal will be in a region with considerably lower $\mathcal{P}$, leading to a low acceptance rate. If, on the other hand, the proposal distribution is too narrow, the acceptance rate will be high, but it will take the chain a long time to cover the entire volume of $\mathcal{P}$, which results in slow convergence. To optimise the process, one has to balance these two effects and find a safe middle ground. If the posterior is (approximately) Gaussian with widths σ_i, the authors of Ref. [64] suggest proposal widths of $\sim 2.4\sigma_i$. Of course, this requires one to have some idea of the shape of $\mathcal{P}$ in advance. It can therefore be helpful to run a short simulation with a test proposal density to estimate an efficient Q for the main run.

- Basis of parameter space:
  The performance of the algorithm also depends on the choice of basis in parameter space. If $\mathcal{P}$ has degeneracies between parameters, convergence will be delayed. Hence, one may want to choose a set of parameters that is as uncorrelated as possible.

- Burn in:
  If the starting point of the chain lies in a region of parameter space where the posterior density is very low, the initial points cannot be regarded as samples of the target distribution and have to be discarded. It will usually take the chain $\lesssim \mathcal{O}(100)$ steps to find a region near a local maximum of $\mathcal{P}$, lose memory of its starting point and start sampling from the equilibrium distribution. The period of equilibration is known as "burn in".

- Number of chains and choice of starting points:
  Several chains generated sampling from the same posterior pdf can be concatenated for the purposes of parameter inference, provided the burn in phase has been removed. Using more than one chain has several advantages: they can be run in parallel (if the computing power is available), which saves time. Secondly, in cases where one cannot be certain that the posterior is not multimodal (i.e., has several local maxima), using only one chain might be dangerous. If it gets stuck in a deep local maximum it could look like it has converged, when in reality it has

not. This can be avoided (or at least made very unlikely) by using several chains with starting points scattered widely over the available parameter space. Lastly, running multiple chains in parallel also allows us to easily monitor convergence during the run, as we will see in the next point.

- Convergence tests:
  It is important to have a measure to help one decide when the point density of the chain(s) is close enough to the sample distribution for one to safely infer parameter values. Gelman and Rubin [65] suggest tracking the following quantity to estimate convergence when using multiple chains:

  $$R(\theta_i) \equiv \frac{\text{Variance of chain means}}{\text{Mean of chain variances}} \,, \tag{III.17}$$

  convergence being likely to have occurred if $R-1 \ll 1$ for all parameters of interest. The `cosmomc` code allows one to automatically stop the algorithm once a certain minimum value of $R-1$ has been reached. One can also check for convergence using the Raftery and Lewis criterion [66]. However, these criteria should only be regarded as a guideline and there is no absolute guarantee that the point density resembles the sampling pdf accurately, particularly so if $\mathcal{P}$ is multimodal.

- Sampling of tails and peaks:
  By construction, the chains produced with the Metropolis algorithm will not spend much time scanning the regions of low posterior density, simply because proposals suggesting a move there are likely to be rejected. The number of points in the chain that fall into a given interval around $\ln\mathcal{P}$ is exponentially suppressed with $-\Delta\chi^2_{\mathcal{P}}/2 \equiv -\ln\mathcal{P}_{\text{max}} + \ln\mathcal{P}$, in the limit of large $-\Delta\chi^2_{\mathcal{P}}$. So we can only expect a small number of samples in the tail of the distribution, which makes it difficult to reliably estimate high-confidence intervals ($\geq 99\%$). In Figure III.5 we plot a histogram for an example chain sampling a 6-dimensional parameter space. We notice the exponential suppression at large $-\Delta\chi^2_{\mathcal{P}}/2$. In addition, we find that the region of highest posterior density, i.e., the peak of the distribution, is also relatively under-sampled. This can be attributed to a volume effect: in general, the volume of the peak region is fairly small, so it will be easily missed by the random walk, leading to a suppression proportional to some power of $-\Delta\chi^2_{\mathcal{P}}/2$, the exact value depending on the number of dimensions and the geometry of $\mathcal{P}$.

Before we proceed to the application of the techniques presented in this section, we will devote a few lines to the matter of model selection.

4. Bayesian Model Selection

Due to the model dependence of statistical data analysis, it is a natural question to ask whether a given model $\mathcal{M}_1$ provides an appropriate description of the data or whether

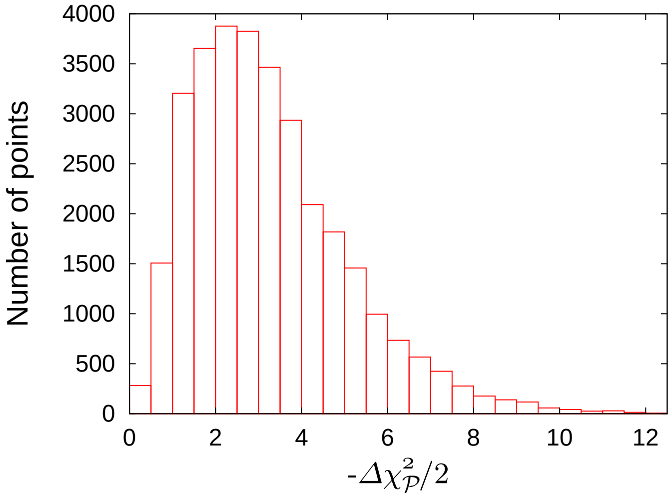

Figure III.5.: This figure shows a histogram of the number of MCMC points per interval of $\Delta\chi^2_{\mathcal{P}}$. The chain has roughly 30000 likelihood evaluations and was generated using the Metropolis algorithm, WMAP data and the vanilla model; $\Delta\chi^2_{\mathcal{P}} = 0$ corresponds to the maximum of $\mathcal{P}$. The chain spends most of the time sampling a region of intermediate values of the posterior; the peak and tails are covered by only a small number of samples.

perhaps a different model $\mathcal{M}_2$ will do a better job. This problem often occurs if one is looking for evidence of new physical effects in the data. Now if we take a given model and add extra free parameters which describe these effects, the quality of the fit will, in general, improve, i.e., the likelihood of the best-fit point in parameter space will be higher. But how can we decide whether the improvement is relevant enough to claim a detection of the effect?

Bayesian statistics provides the tools needed to deal with this issue in a straightforward way (cf., e.g., Refs. [67,68] and [69,70] for applications to cosmology). The "goodness" of a model $\mathcal{M}_i$ can be regarded as the conditional probability $\mathcal{P}_{\mathcal{M}_i}(\mathcal{M}_i|\mathcal{D})$. We can relate this to the evidence $\mathcal{E}$ of Equation (III.11) by applying Bayes' theorem again:

$$\mathcal{P}_{\mathcal{M}_i}(\mathcal{M}_i|\mathcal{D}) = \frac{\mathcal{E}(\mathcal{D}|\mathcal{M}_i)\,\pi(\mathcal{M}_i)}{p(\mathcal{D})}\ , \tag{III.18}$$

where $\pi(\mathcal{M}_i)$ is the model prior and $p(\mathcal{D})$ a normalisation constant. If we want to decide between two models, i.e., pick the model that is more likely to describe the data well, we should consider the ratio of the two posteriors $\mathcal{P}_{\mathcal{M}_i}$, known as the *Bayes factor*

$$B_{12} \equiv \frac{\mathcal{P}_{\mathcal{M}_1}}{\mathcal{P}_{\mathcal{M}_2}} = \frac{\mathcal{E}(\mathcal{D}|\mathcal{M}_1)\,\pi(\mathcal{M}_1)}{\mathcal{E}(\mathcal{D}|\mathcal{M}_2)\,\pi(\mathcal{M}_2)}\ , \tag{III.19}$$

which conveniently removes the normalisation $p(\mathcal{D})$. If there is no theoretical reason to prefer one model over the other, we can set $\pi(\mathcal{M}_1) = \pi(\mathcal{M}_2) = \frac{1}{2}$, and the Bayes factor reduces to the ratio of evidences. We can interpret the Bayes factor as betting odds in favour of $\mathcal{M}_1$, so the larger B_{12}, the more convinced we can be that $\mathcal{M}_1$ should be regarded as being superior to $\mathcal{M}_2$. There exists a somewhat arbitrary mapping from numbers to semantics [67], stating that $\ln B_{12} > 5$ represents "decisive" evidence in favour of $\mathcal{M}_1$, $\ln B_{12} > 1$ is "significant", while $\ln B_{12} < 1$ is "barely worth mentioning".

Note that the evidence and hence the Bayes factor is still dependent on the priors imposed on the parameters of the models. In particular, for flat priors there is a linear dependence on the width of the top hat function, i.e., the length of the interval over which the prior is non-zero. In parameter estimation, this quantity is not relevant for the inferred confidence intervals, as long as it is big enough to cover most of the volume of the likelihood pdf. Here however, it matters. This adds another degree of subjectivity to the procedure and should be kept in mind when interpreting the results.

Since the evidence is essentially a marginalisation of the prior weighted likelihood over the whole parameter space, the above comments on marginalisation also apply here: it punishes fine-tuning of parameters and prefers theories that make definite predictions (that happen to reasonably agree with the data). Thus, this method can be viewed as an implementation of *Ockham's razor*, sorting out theories that are baroquely embellished with needless parameters in favour of their simpler brethren, but at the same time ruling out models that are too simple to fit the data well.

Bayesian model selection has been applied to existing cosmological data, finding significant evidence for a non-zero tilt of the primordial power spectrum [71] and spatial flatness of the Universe as well as no evidence for isocurvature contributions to the initial perturbations [69, 72] or unusual topologies of the Universe [73]. A further area of application lies in the prediction of the model selection power of future experiments; this has been done, e.g., in Ref. [74] for dynamical dark energy and in Refs. [75, 76] for the spectral tilt and its running.

This framework would also enable one to assess the ability of future observations to test the inflationary consistency relation by comparing a model that has n_T fixed by Equation (II.62) with a model that treats n_T as a free parameter for a given fiducial data set.

The trouble with Bayesian model selection, however, is that it is very computationally intensive to evaluate the evidence, a multi-dimensional integral, to a required accuracy. Despite recent algorithmic advances [77, 78], one might also want to consider approximative methods for model comparison. Ref. [79] gives a nice overview on different methods derived from information theory. Let us list the two most widely used information criteria. They both attempt to strike a balance between the goodness of the fit and the number of parameters used in the model and are fairly simple to calculate:

- **Akaike Information Criterion (AIC)** [80]
  The AIC is defined by
  $$\mathrm{AIC} = -2\,\ln \mathcal{L}_{\mathrm{max}} + 2K, \tag{III.20}$$
  where $\mathcal{L}_{\mathrm{max}}$ is maximum likelihood attainable in the model, and K the number of free parameters.

- **Bayesian Information Criterion (BIC)** [81]
  The BIC is given by
  $$\mathrm{BIC} = -2\,\ln \mathcal{L}_{\mathrm{max}} + K \ln N, \tag{III.21}$$
  where N is the number of data points used in the fit.

The lower their value, the better the model, by the respective standards. The difference $\Delta\mathrm{AIC} \equiv \mathrm{AIC}(\mathcal{M}_1) - \mathrm{AIC}(\mathcal{M}_2)$, or ΔBIC, respectively, can be used for comparison, negative values preferring $\mathcal{M}_1$.

IV. Vanilla and Beyond

In the last chapter we have seen how the construction of confidence intervals depends on the choice of the underlying model. We have also outlined how to rate different models based on the available data. More or less rigorous application of the model selection formalism to current cosmological data, ranging from measurements of the CMB anisotropies by the WMAP satellite [23–25] and the large scale structure of the Universe by the Sloan Digital Sky Survey [19, 82], to the observation of distant type Ia supernovæ [83], seems to favour the six-parameter vanilla concordance model, in the sense that adding further free parameters will not significantly improve the quality of the fit.

The values and confidence intervals of these six parameters are commonly inferred in the vanilla setting [23, 84]. Similarly, if one wants to look for physical effects that go beyond the vanilla model, the bounds on the parameters which quantify these effects are usually derived from a fit to just the minimal "vanilla+1" extension [19, 85], motivated by a naïve application of Ockham's razor.

While the simplicity of this approach is certainly appealing, it is in fact quite restrictive. The resulting error estimates can by no means be considered conservative because they are derived under the condition that other physically well motivated effects are non-existent. For instance, the vanilla model assumes neutrinos to be massless. While the presently available cosmological data do not seem to require neutrinos to be massive, we know from the observation of neutrino flavour oscillations [86, 87], that massless neutrinos are ruled out at high confidence level. If one were to apply this information, in the form of a model prior, to a Bayesian model comparison analysis and compared the vanilla model with its vanilla+m_ν extension, one would actually have to conclude that vanilla is ruled out! This should not come as a surprise though. The concordance model is, after all, just an empirical attempt at explaining cosmological data and does not claim to be the ultimate theory of the physics of the early Universe.

That said, the vanilla(+1) method is certainly a valid way of probing cosmological data for new physical effects. Given that we cannot rule out physics beyond vanilla with absolute certainty, however, its unreflected usage can lead to significant underestimation of parameter errors, as well as bias in the parameter estimates, particularly if the new parameters have a similar effect on the spectra as existing ones, i.e., if they introduce new degeneracies.[1] One well-known example is the interplay between the dark energy

---

[1] It should be noted that, owing to the nature of Bayesian inference, the addition of extra parameters need not lead to larger errors in all cases, opposed to a Frequentist analysis. A nice counterexample is given in Ref. [88]. The extra parameters considered in the following happen to adhere to what intuition would suggest, though.

equation of state and the neutrino mass [89–91]. When the dark energy equation of state is allowed to vary, the neutrino mass bound is relaxed by almost a factor of three if only CMB and LSS power spectrum information is used. Conversely, by imposing a prior on the neutrino masses according to the Heidelberg–Moscow claims [92–94], the authors of Ref. [90] find that dark energy in the form of a cosmological constant is ruled out at more than 95% C.L. by CMB+LSS+SNIa data.

To put it another way, bounds derived using the vanilla model do not take into account the intrinsic uncertainty about whether the model one is using to fit the data is actually appropriate. In a way, these limits feign an accuracy that is only tenable under the restrictive assumptions of the model. This becomes of special importance if one wants to use these results to draw conclusions about other models, e.g., to rule out a particular model of inflation based on limits on the spectral index, or constraining the MSSM parameter space with inferred values of Ω_c.

We argue that when constraining or excluding specific theoretical models, one should in principle allow for uncertainties in all physically well-motivated parameters, even if they have *a priori* no direct link to the models concerned.

It is therefore important to perform a parameter estimation analysis in the framework of a much more general model, compare the results to the vanilla bounds and see how this would affect conclusions about the nature of the Universe.

1. The Model

In principle, there is of course an arbitrary number of extensions to the concordance model, of which maybe 20 or so extra parameters have been discussed in the literature. In some cases though, the precision of the present data is not yet sufficient to constrain them (e.g., the primordial helium fraction and the effective sound speed of dark energy). We cannot presume to present a completely general analysis here. In the present work we will therefore limit ourselves to a model with eleven free parameters, given by

$$\boldsymbol{\theta} = (\Omega_\mathrm{c}, \Omega_\mathrm{b}, H_0, \tau, A_\mathrm{S}, n_\mathrm{S}, f_\nu, \Omega_\mathcal{K}, w, r, \alpha_\mathrm{S}), \tag{IV.1}$$

and subsets thereof. The vanilla model is defined by $f_\nu = \Omega_\mathcal{K} = r = \alpha_\mathrm{S} = 0$, and $w = -1$. In addition, we marginalise over a nuisance parameter b' which describes the relative bias between the observed galaxy power spectrum $P_\mathrm{g}(k)$ and the underlying dark matter spectrum $P_\mathrm{DM}(k)$ via $P_\mathrm{DM}(k) = b'^2 P_\mathrm{g}(k)$. We impose flat priors on all parameters except the normalisation, which is customarily assigned a flat prior on $\ln A_\mathrm{S}$.

1.1. Initial Conditions

We assume the initial conditions for structure formation to be set by inflation. The spectra of adiabatic scalar and tensor perturbations will be parameterised by amplitudes

A_S and $A_\mathrm{T} = rA_\mathrm{S}$ and spectral indices n_S and n_T, respectively. As discussed at the end of Chapter II, current data does not allow us to constrain the tensor spectral index. We will therefore not treat it as a free parameter and instead take the single field consistency relation $n_\mathrm{T} \sim -r/8$ to hold. For the same reason it would not be sensible to include a running of the tensor spectral index. We do allow for a running of the scalar spectral index though, parameterised by α_S. We define α_S and r at the pivot scale $k = 0.002\ \mathrm{Mpc}^{-1}$, in concordance with most recent analyses (see also Ref. [95] for a discussion of the optimal choice of pivot scale). This parameterisation will cover all slow roll single field inflationary models that do not have other unusual features, such as, e.g., particle production.

In the following, we will consider three different parameter subsets, corresponding to different degrees of restrictions on the inflationary model:

A. All 11 parameters.

B. A 10-parameter set with $\Omega_\mathcal{K} = 0$. Spatial flatness is one of the predictions that make inflation so appealing. While it is technically possible to construct models that result in a non-flat Universe today [96, 97], the existing models are somewhat contrived and fine-tuned [98]. This subset will discard these models.

C. A 9-parameter set with $\Omega_\mathcal{K} = \alpha_\mathrm{S} = 0$. This reduced set corresponds to the large subset of the zoo of inflationary models that predict negligible running, including large field chaotic inflation models [99].

1.2. Energy Content

We take the matter and energy content of the Universe to be specified by the following parameters: the curvature $\Omega_\mathcal{K} = 1 - \Omega_\mathrm{m} - \Omega_\Lambda$, the physical dark matter density $\omega_\mathrm{c} = \Omega_\mathrm{c} h^2$, the baryon density $\omega_\mathrm{b} = \Omega_\mathrm{b} h^2$, the neutrino fraction $f_\nu = \Omega_\nu / \Omega_\mathrm{c}$, and the dark energy equation of state parameter w.

Other parameters not included here, but which could have an observable effect, include a time-dependent dark energy equation of state, nonstandard interactions in any of the dark sectors (cold dark matter, neutrinos, or dark energy), etc.

2. Data Analysis

We consider the following data sets in our analysis:

- Cosmic microwave background:
  We use CMB data from the WMAP experiment after three years of observation [23–26]. The data analysis is performed using the likelihood calculation package provided by the WMAP team on NASA's `LAMBDA` web page [100].

- Large scale structure:
  The large scale structure power spectrum of luminous red galaxies has been measured by the Sloan Digital Sky Survey. We use the same analysis technique on this data set as advocated by the SDSS team [19,82], with analytic marginalisation over the bias b' and the nonlinear correction parameter $Q_{\rm nl}$.

- Baryon acoustic oscillations (BAO):
  In addition to the power spectrum data we use the measurement of baryon acoustic oscillations in the two-point correlation function [20]. The analysis is performed following the procedure described in Ref. [20] (see also Ref. [101]), including analytic marginalisation over the bias b', and nonlinear corrections with the `halofit` [102] package.

- Type Ia supernovæ:
  We use the luminosity distance measurements of distant type Ia supernovæ provided by the Supernova Legacy Survey (SNLS) [83].

We do not include data from the Lyman-α forest in our analysis. These data were used in some previous studies that found very strong bounds on various cosmological parameters [85]. However, the strength of these bounds is mainly due to the fact that the Lyman-α analysis used in [85] leads to a much higher normalisation of the small-scale power spectrum than that obtained from the WMAP data. Other analyses of the same SDSS Lyman-α data find a lower normalisation, in better agreement with the WMAP result [103–105]. This kind of discrepancy between different analyses of the same data hints at unresolved systematic issues, and for this reason we prefer not to use the Lyman-α data at all.

For a large part of the analysis we use two different combinations of data sets, one consisting of WMAP and SDSS data only, and one which uses in addition data from SNIa (SNLS) and BAO. We will refer to the latter as "the full data set".

We perform the data analysis using the publicly available `cosmomc` package [13], modified to include the BAO likelihood calculations.

3. Results

In this section we will present the results of our analysis, starting with the most general model and constraints on the parameters describing the energy content, before moving on to a discussion of the inflationary sector.

3.1. Curvature, Dark Matter and Dark Energy

Since spatial flatness is one of the key predictions of inflationary cosmology, it is important to check whether also in the context of our more general model A, $\Omega_{\mathcal{K}} = 0$ is still compatible with the data. A glance at Figure IV.1 will, rather reassuringly, tell us that the answer is yes: we find $-0.022 < \Omega_{\mathcal{K}} < 0.026$ at 95% confidence level, this interval is only slightly larger than the one inferred by the SDSS collaboration from a vanilla+$\Omega_{\mathcal{K}}$ fit ($-0.015 < \Omega_{\mathcal{K}} < 0.023$, 95% C.L.) [19]; flatness is not compromised. In the discussion of the inflationary parameter space we will therefore assume $\Omega_{\mathcal{K}} = 0$, i.e., limit ourselves to parameter sets B and C.

What about the dark side of the Universe? The dark matter density $\Omega_c h^2$ is a crucial input in dark matter model building. A prime example of this is given by models with low energy SUSY where the dark matter particle is usually either the neutralino or the

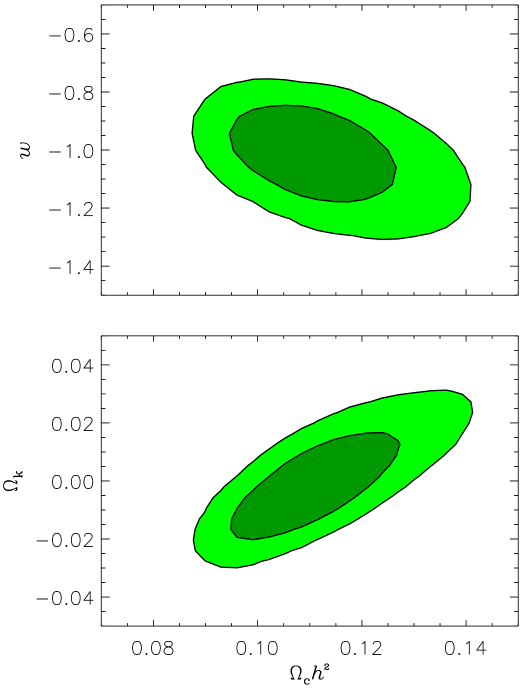

Figure IV.1.: Two-dimensional marginalised 68% and 95% C.L. contours for $\Omega_c h^2$, $\Omega_{\mathcal{K}}$, and w, using parameter set A and the full data set.

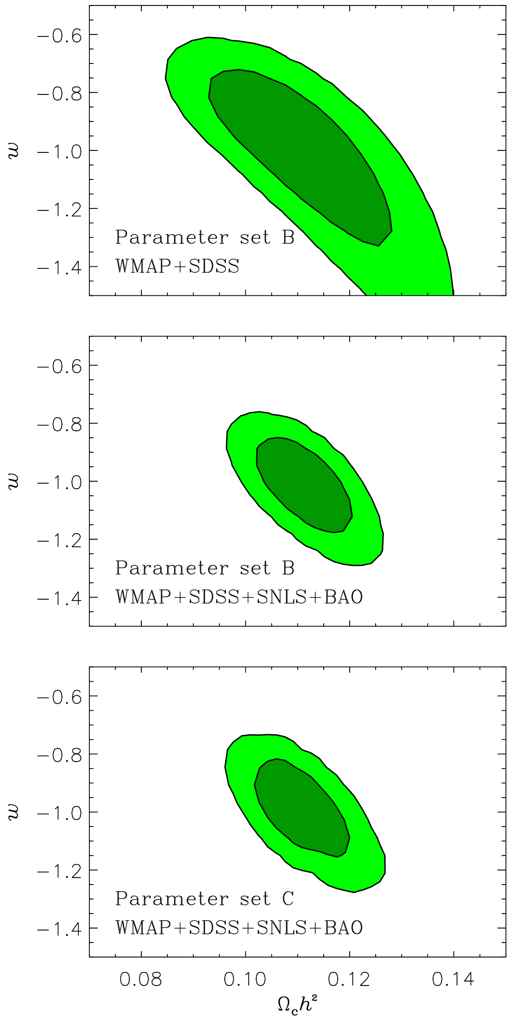

Figure IV.2.: Two-dimensional marginalised 68% and 95% C.L. contours for $\Omega_c h^2$ and w, using various parameter and data sets. *Top*: Parameter set B, WMAP+SDSS. *Centre*: Parameter set B, the full data set. *Bottom*: Parameter set C, the full data set.

gravitino. Large regions in parameter space in these models have been excluded by the fact that the predicted dark matter density is too high or too low [106–108].

In the vanilla model $\Omega_c h^2$ is a very well constrained quantity, with WMAP+SDSS giving a 68% C.L. limit of $\Omega_c h^2 = 0.1050^{+0.0041}_{-0.0040}$ [19]. This corresponds to a relative uncertainty of $\sigma(\Omega_c h^2)/\Omega_c h^2 \simeq 0.04$. The SDSS Collaboration also provide bounds on $\Omega_c h^2$ in extended models in which one additional parameter is added to the vanilla parameter set [19]. In most cases the bound on $\Omega_c h^2$ does not change significantly. However, when either $\Omega_\mathcal{K}$ or w is allowed to vary, $\sigma(\Omega_c h^2)/\Omega_c h^2$ increases to about 0.06 [19].

We have taken this investigation further by calculating the bound on $\Omega_c h^2$ for our

Parameter set	Data set	$\Omega_c h^2$	w
A	WMAP+SDSS+SNLS+BAO	$0.094 \rightarrow 0.136$	$-1.19 \rightarrow -0.88$
B	WMAP+SDSS	$0.092 \rightarrow 0.136$	$-1.44 \rightarrow -0.76$
B	WMAP+SDSS+SNLS+BAO	$0.100 \rightarrow 0.123$	$-1.12 \rightarrow -0.87$
C	WMAP+SDSS+SNLS+BAO	$0.100 \rightarrow 0.123$	$-1.11 \rightarrow -0.86$
vanilla	WMAP+SDSS	$0.097 \rightarrow 0.113$	-1

Table IV.1.: The 1D marginalised 95% C.L. allowed ranges for $\Omega_c h^2$ and w for various parameter and data sets.

various parameter and data sets. In Figure IV.2 we show the joint 2D marginalised constraints on $\Omega_c h^2$ and w for three different cases using parameter sets B and C. If only WMAP and SDSS data are used, a very strong degeneracy between $\Omega_c h^2$ and w weakens the bounds on both parameters. This degeneracy is broken when SNIa and BAO are included (as is also the case with the degeneracy between f_ν and w), yielding strong constraints on both parameters.

When spatial curvature is also allowed to vary, the bound on $\Omega_c h^2$ does change considerably. Figure IV.1 shows the 2D marginalised contours for $\Omega_c h^2, \Omega_\mathcal{K}$ and w, using parameter set A and the full data set. Here we find $\sigma(\Omega_c h^2)/\Omega_c h^2 \simeq 0.1$, so that $0.094 < \Omega_c h^2 < 0.136$, and $-1.19 < w < -0.88$ (1D at 95% C.L.). We note also that even though the allowed range for $\Omega_c h^2$ increases sizably with the inclusion of $\Omega_\mathcal{K}$, the same is not true for the dark energy equation of state parameter w; we do not find any evidence for deviations from a cosmological constant. In Table IV.1, we summarise the 1D 95% constraints on $\Omega_c h^2$ and w from Figs. IV.1 and IV.2. Let us stress that caution should be applied whenever the dark matter density is used as an input to constrain models such as the MSSM. Parameter regions that are excluded in the simplest vanilla model can easily be allowed in more general models, even without the introduction of more exotic features such as, e.g., isocurvature modes. If one is to take one single number inferred from cosmological observations as an input to constrain particle physics models, then the safest approach is to allow for the possibility that cosmology is not described by the vanilla model, but by something more general. From our calculations, we recommend taking the value derived in our model A: $0.094 < \Omega_c h^2 < 0.136$ (95% C.L.), but we warn that even this may not be the most conservative estimate.

3.2. Inflationary Parameters

Almost all inflationary models predict $\Omega_\mathcal{K}$ to be zero. As we have seen in the previous section, this prediction is also supported by our analysis of parameter set A. Therefore, in this section, we will work with the reduced 10-parameter set B, in which $\Omega_\mathcal{K}$ is already fixed at zero. Figure IV.3 shows the 2D likelihood contours for the parameters n_S, r and α_S using the full data set and parameter set B. These contours are obtained by marginalising over the other $(10-2)$ parameters not shown in the plot.

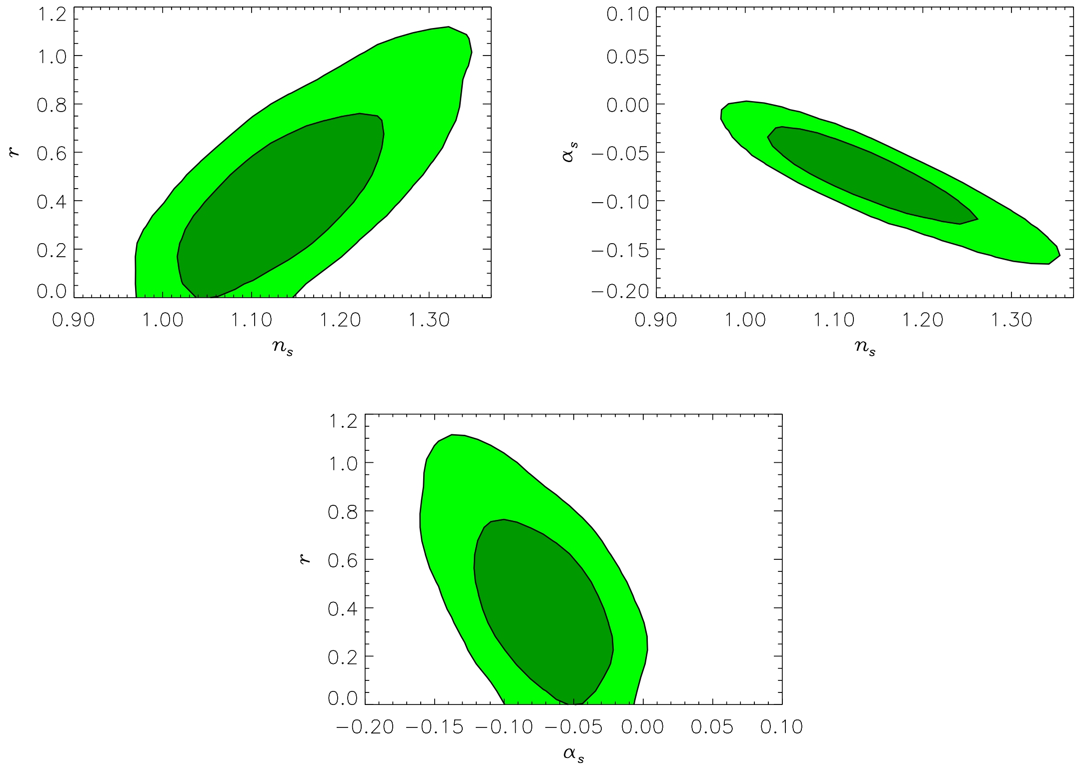

Figure IV.3.: Two-dimensional marginalised 68% and 95% C.L. contours for the inflationary parameters n_S, r, and α_S, using the full data set and parameter set B.

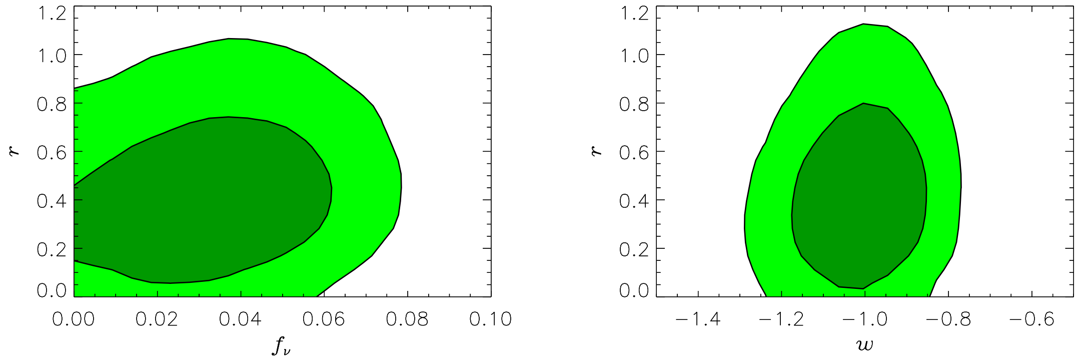

Figure IV.4.: Degeneracies between r and f_ν, w for the full data set and parameter set B, marginalised over $(10-2)$ parameters.

Parameter	WMAP+SDSS	Full data set
n_S	$0.97 \rightarrow 1.35$	$0.98 \rightarrow 1.28$
r	$0 \rightarrow 1.05$	$0 \rightarrow 0.81$
α_S	$-0.140 \rightarrow -0.005$	$-0.135 \rightarrow -0.004$

Table IV.2.: The 95% C.L. allowed ranges for n_S, r and α_S for parameter set B, marginalised over the other $(10-1)$ parameters.

Figure IV.3 should be compared with, e.g., Figs. 2 and 3 of Kinney *et al.* [109], which use data from WMAP and SDSS, and a parameter set similar to our set B but with f_ν and w fixed at 0 and -1, respectively. The comparison reveals that the two sets of likelihood contours are roughly similar, but with one important exception: the allowed range for the tensor-to-scalar ratio r in our case is much larger despite our use of additional data!

In order to understand this effect we plot the 2D likelihood contours for r and our additional parameters f_ν and w in Figure IV.4. Interestingly, a substantial degeneracy exists between r and the neutrino fraction f_ν, which in turn allows r to extend to much higher values. Table IV.2 displays the 1D 95% C.L. allowed ranges for n_S, r, and α_S, assuming parameter set B and using both WMAP+SDSS only and the full data set. It appears that the data slightly prefer a strong negative running of the spectral index, combined with a large tensor amplitude. This result, however, is strongly prior dependent and one could argue that a flat prior on α_S is not realistic, since generic slow roll inflation models predict a running no larger than $\mathcal{O}(10^{-3})$. Apart from that, the actual improvement in χ^2_{eff} is only about 3, so $\alpha_S = 0$ lying outside the 95% C.L. interval is likely a marginalisation effect (see also the top right plot of Figure IV.3) and should not be considered conclusive evidence against power-law models.

3.3. Chaotic Inflation

Single field inflation models with polynomial potentials generally predict negligible running. These models are thus represented by our 9-parameter set C in which $\alpha_S = 0$.

For the purpose of illustration we will indicate the predictions of the two simplest chaotic inflation models, $V = \frac{1}{2}m^2\phi^2$ and $V = \lambda\phi^4$, in our figures. As we saw in Chapter II, the calculation of the initial perturbation spectra requires one to evaluate the relevant functions of the potential at the time a given scale (e.g., today's horizon scale) leaves the horizon during inflation. The identification of scales today with scales during inflation, however, is not trivial, since we do not know the complete expansion history of the Universe. In particular, the duration of the radiation dominated phase, which depends on the inflationary model and on details of the reheating mechanism, is unknown. This ignorance can be expressed in terms of an additional parameter: the number of e-foldings of inflation N_H that took place after the scale we identify with

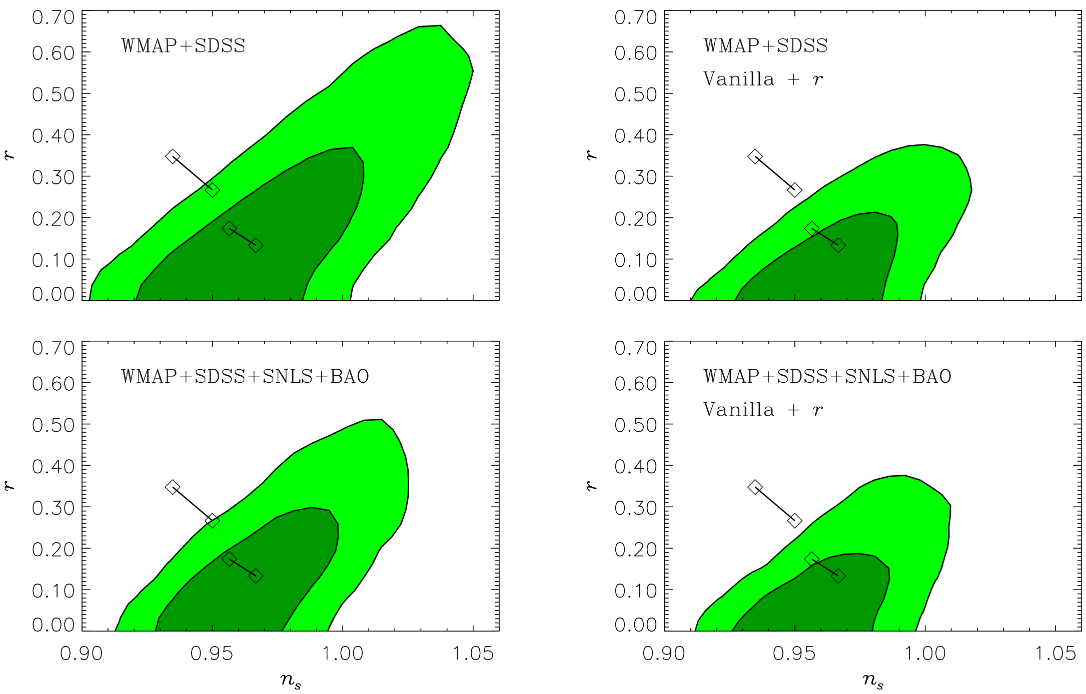

Figure IV.5.: Two-dimensional marginalised 68% and 95% C.L. contours for n_S and r. *Left:* parameter set C (consistent with predictions of chaotic inflation). *Right:* vanilla+r. The upper diagrams use WMAP+SDSS data and the lower ones the full data set. The two short black lines with boxes at the ends correspond to predictions of the $\lambda\phi^4$ (top left) and $m^2\phi^2$ models of inflation, with 46 to 60 e-foldings (left to right).[3]

today's Hubble scale left the horizon during inflation:

$$N_H \simeq 62 - \ln \frac{10^{16}\ \text{GeV}}{V^{1/4}} + \ln \frac{V^{1/4}}{V_{\text{end}}^{1/4}} - \tfrac{1}{3} \ln \frac{V_{\text{end}}^{1/4}}{T_{\text{rh}}}\,, \tag{IV.2}$$

where T_{rh} is the reheating temperature, V_{end} is the potential energy at the end of inflation and V is evaluated at horizon exit of today's Hubble scale [12]. Plugging in typical values, we obtain N_H in the range of 40 → 60.

Since the observable inflationary parameters are generally not independent of N_H, the prediction of a concrete inflationary model does not correspond to a point in parameter space, but rather a line. These are depicted for n_S and r, along with the corresponding 2D likelihood contours, in Figure IV.5. For comparison, we also plot the constraints obtained from a fit to the vanilla+r model, this nicely illustrates how going to a more general model can relax bounds on parameters.

---

[3] When taking into account one-loop effects in the chaotic inflationary scenario, the model lines in these plot may actually be smeared at the percent level [110].

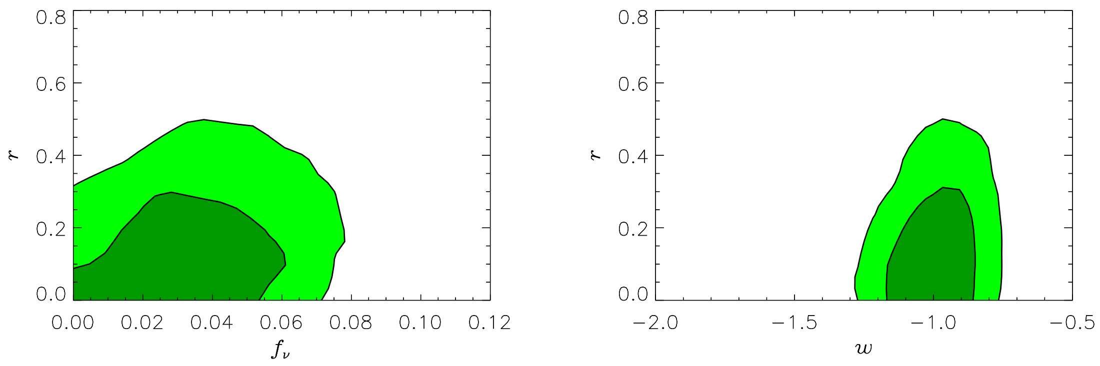

Figure IV.6.: Degeneracies between between r and f_ν, w for the full data set and parameter set C, marginalised over $(9-2)$ parameters.

Figure IV.5 should also be compared with Figure 4 of Kinney *et al.* [109], with Figure 14 of Spergel *et al.* [23], and with Figure 19 of Tegmark *et al.* [19] (cf. also Ref. [111]). In all cases our WMAP+SDSS contours encompass a markedly larger region. In particular, even with the inclusion of SNIa and BAO data, we find that the simplest $\lambda\phi^4$ model is still allowed by the data, contrary to the conclusions of Refs. [19, 23, 109, 111].

Interestingly, $\lambda\phi^4$ is compatible with data only if the number of e-foldings is relatively large, or equivalently, if the reheating temperature is high [112, 113].

We find a 1D 95% C.L. upper bound of $r < 0.31$, while Tegmark *et al.* report an almost identical $r < 0.33$. Kinney *et al.* also found $r < 0.31$ for the same vanilla+r model [109], but from a combination of WMAP and the SDSS main galaxy samples (as opposed to SDSS LRG used in this work and in Ref. [19]).

Using additional data from the Lyman-α forest, Seljak *et al.* [85] derived an even stronger upper bound, $r < 0.22$, for the same model space. The reason for the improvement is a degeneracy between r and σ_8, such that a higher value of r leads to a smaller preferred value of σ_8. Since the Lyman-α data used in [85] prefer a high value of σ_8, a small r value is correspondingly favoured. In fact, from a parameter fitting point of view, a negative r would be even better. All these points conspire to give a much stronger upper bound on r. However, as noted in Section IV.2, this phenomenon likely points to a systematic uncertainty in the Lyman-α normalisation, rather than a genuinely strong constraint on r.

Again, the explanation for our enlarged (n_S, r) allowed region lies in our expanded model parameter space. In Figure IV.6, we see that the degeneracy between r and f_ν encountered earlier in parameter set B is present also in parameter set C, albeit to a somewhat smaller extent. If a neutrino fraction of $0.03 \to 0.05$ is allowed (corresponding roughly to $\sum m_\nu \sim 0.3 \to 0.5$ eV), new parameter space opens up for n_S and r. We note in passing that the converse is not true. Taking r as a free parameter does not change the upper bound on the neutrino mass significantly compared to r fixed at zero.

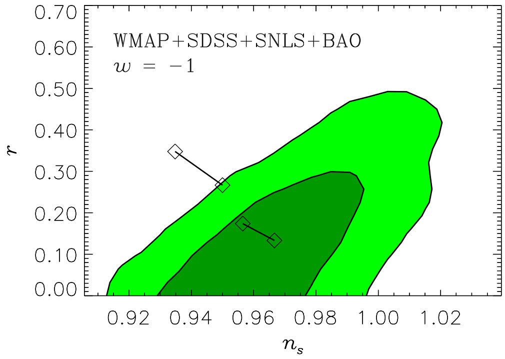

Figure IV.7.: Two-dimensional marginalised constraints on n_{S} and r for the parameter set C, but with the restriction $w = -1$.

Interestingly, this (f_ν, r) degeneracy also means that $\lambda\phi^4$ in its simplest form appears to prefer quasi-degenerate neutrino masses with a sum in the $0.3 \to 0.5$ eV range. This range is compatible with present laboratory limits from tritium beta decay experiments, $m_\nu < 2.2$ eV [27, 114, 115], as well as the claimed detection of neutrinoless double beta decay, and hence detection of the effective electron neutrino mass at $0.1 \to 0.9$ eV, by the Heidelberg–Moscow experiment [92–94]. The upcoming tritium beta decay experiment KATRIN will also probe neutrino masses to a comparable level of precision [116].

As an extra consistency check we present in Figure IV.7 the equivalent of the bottom left panel of Figure IV.5 (full data set and parameter set C), but with the additional restriction $w = -1$. Clearly, there is very little difference between the bottom left diagram of Figure IV.5 and Figure IV.7, since the combination of SNIa and BAO data effectively fixes w to -1 in the former case, as shown in Figure IV.6.

This confirms our conclusion that the difference between the allowed (n_{S}, r) regions in the left and right panels of Figure IV.5 is due to a degeneracy between r and the neutrino fraction f_ν. It should also be noted that the addition of SNIa and BAO data has very little impact on the vanilla+r model, because no strong parameter degeneracies are present in the WMAP+SDSS data. With SNIa and BAO included we find a 1D 95% C.L. bound of $r < 0.30$, instead of 0.31 for WMAP+SDSS alone.

3.4. The Effect of Nonlinearity

Up to this point we have used exactly the same analysis technique as the SDSS team when treating the LRG data. However, due to the large bias of the luminous red galaxies, the effects of nonlinear structure growth on the power spectrum will already set in at wavenumbers of approximately $k \sim 0.06 \to 0.07\ h\mathrm{Mpc}^{-1}$, and start dominating it for $k \geq 0.09\ h\mathrm{Mpc}^{-1}$ (see, for instance, Figure 9 of [19]). To test whether or not our results are subject to these effects, we perform the same analysis as in Figure IV.5, but retain

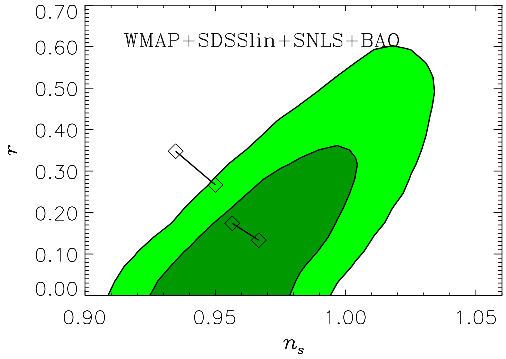

Figure IV.8.: Two-dimensional marginalised constraints on n_S and r for parameter set C. Only the linear part of the SDSS power spectrum has been used.

data only up to $k \sim 0.06\ h\mathrm{Mpc}^{-1}$ (band 11). We call this reduced data set SDSS$_\mathrm{lin}$, and the result is shown in Figure IV.8. Using only the linear part of the power spectrum data has no bearing on our conclusions. In fact, the 2D allowed region in (n_S, r) for parameter set C is only affected in the region where $n_\mathrm{S} > 1$. The SDSS data probes n_S more precisely when all data points are included, and this in turn leads to a truncation of the allowed region at high n_S.

In Table IV.3, we summarise the 1D marginalised constraints on n_S and r for parameter set C and its subsets.

Parameter set	Data set	n_S	r
C	WMAP+SDSS	0.927 → 1.038	0 → 0.51
C	WMAP+SDSS+SNLS+BAO	0.932 → 1.018	0 → 0.41
C	WMAP+SDSS$_\mathrm{lin}$+SNLS+BAO	0.931 → 1.025	0 → 0.47
C, w fixed	WMAP+SDSS+SNLS+BAO	0.933 → 1.019	0 → 0.40
C, w, f_ν fixed	WMAP+SDSS	0.931 → 1.011	0 → 0.31
C, w, f_ν fixed	WMAP+SDSS+SNLS+BAO	0.931 → 1.010	0 → 0.30

Table IV.3.: The 1D marginalised 95% C.L. allowed ranges for n_S and r for parameter set C and its subsets.

4. Discussion

We have performed a detailed study of cosmological parameter estimation in the context of extended models that encompass a larger model parameter space than the standard, flat ΛCDM cosmology. Using the vanilla model as a basis, we have included in addition a number of physically well-motivated parameters, such as a nonzero neutrino mass. We have considered an eleven-parameter model and subsets thereof to derive 1D constraints on the parameters, in contrast with the vanilla+1 approach adopted in most previous analyses which treats one extra parameter at a time.

In this more general framework, we find that in the context of standard slow roll inflation, constraints on dark matter and dark energy can be substantially altered. If only CMB and LSS data are used, the larger parameter space introduces new, strong degeneracies, e.g., between the physical dark matter density $\Omega_c h^2$ and the dark energy equation of state w. These degeneracies can be broken to a large extent by adding type Ia supernova and baryon acoustic oscillation data to the analysis. However, even with this expanded data set, we find that the bound on the physical dark matter density $\Omega_c h^2$ is relaxed by more than a factor of two compared to the vanilla model constraint. On the other hand, the more general model does not seem to lead to a significant bias on the spatial curvature $\Omega_{\mathcal{K}}$ or w, and our results are consistent with spatial flatness and a cosmological constant.

In the same spirit, we have studied how bounds on the inflationary parameters n_S, r, and α_S are affected by the introduction of extra parameters in the analysis. Unlike other recent analyses [19, 23, 109], we find that the simplest $\lambda\phi^4$ model of inflation is still marginally compatible with all present data at the 95% level. The source of this apparent discrepancy is a degeneracy between the tensor to scalar ratio r and the neutrino fraction f_ν, the latter of which was fixed at zero in the analyses of [19, 23, 109]. Reversing the argument, if $\lambda\phi^4$ is the true model of inflation, then it strongly favours a sum of quasi-degenerate neutrino masses between 0.3 and 0.5 eV, a range compatible with present data from laboratory experiments. This represents a clear example of how neutrino masses well within laboratory limits can bias conclusions about other, seemingly unrelated cosmological parameters.

V. Inflation with a Step

As we saw in Chapter II, we had to make a number of simplifying assumptions in the derivation of the power-law form of the initial power spectra, namely

- **Single field inflation**:
  Only one scalar field is of dynamical importance during inflation and is responsible for the exponential expansion as well as generation of the initial perturbations.
- **Initial conditions of the mode functions**:
  The initial conditions for the u_k are given by Equation (II.24), the Bunch-Davies vacuum state.
- **Slow roll**:
  The potential is flat and smooth, the slow roll parameters $\epsilon_{(H)}$, $\eta_{(H)}$ and their derivatives are much smaller than unity when the relevant scales leave the horizon.

At this point one might want to ask about the consequences of relaxing any one of these requirements and their theoretical motivations. Let us address this question point by point:

- **Many fields?**
  Considering that the most popular candidates for extensions of the Standard Model of particle physics predict a plethora of scalar fields, it may seem presumptuous to consider only one field. A careful analysis of perturbation theory in multi-field settings, however, shows that for straight trajectories in field space, there is no difference to the single field case. Only if the trajectory is curved do differences arise in the form of isocurvature modes [117]. If spontaneous symmetry breaking is involved, one may also face topological defects [118], which defy perturbative treatment.
- **No Bunch-Davies vacuum?**
  Unless inflation started just when today's observable scales left the horizon (in which case the assumption of the Universe being the vacuum state would be somewhat arbitrary), the observable scales will have had physical wavelengths shorter than the Planck scale at some point during inflation. Lacking a thorough understanding of quantum gravity, the state in which the mode functions are prepared when they enter the regime where the contributions of new physics become negligible, is somewhat speculative. This issue is known as the trans-Planckian problem of inflation [119] (see also Ref. [120] and references therein). Depending on the energy scale of inflation, such effects might for instance manifest themselves in oscillations superimposed on the usual power-law spectra [121].

- **Fast roll?**
  Recall that the slow roll conditions $\epsilon \ll 1$ and $|\eta| \ll 1$ were introduced to ensure that inflation lasts the required ~ 50 e-foldings. Demanding that all higher slow roll parameters are negligible is certainly a sufficient condition to achieve this goal. It is, however, not a necessary condition; instead of one, uninterrupted, phase of slow roll inflation, one might as well have several, shorter, periods of inflation, interrupted by non-inflationary or non-slow-roll-inflationary phases. A prime example is given by thermal inflation [122], where a second inflationary phase of ~ 10 e-foldings follows after an initial stage of ordinary inflation and reheating.

  Of course, an interruption of slow roll inflation is particularly interesting if it happens roughly 50 e-foldings before the end of inflation. In that event, there may be detectable traces in the observed anisotropy spectra. There are various theoretically well-motivated scenarios in which this happens to be the case, ranging from simple renormalisable polynomial potentials [123] over resonant particle production during inflation [124, 125], phase transitions in a supergravity model [126] or a second order phase transition in double D-term inflation [127] to the dynamics of M5-branes in heterotic M-theory [128, 129]. Depending on the model, the power spectrum can be adorned with various unusual features, e.g., cutoffs, steps or oscillations.

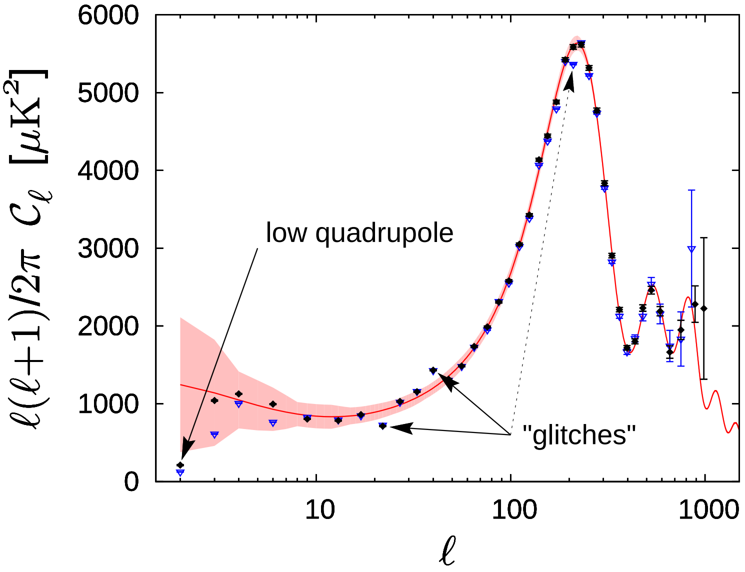

Figure V.1.: This figure shows the angular power spectrum of temperature anisotropies measured by the WMAP satellite after three years of observation, plotted as black diamonds, and the vanilla best fit to this data set. The arrows indicate the curiously low quadrupole and the two glitches at $\ell = 22$ and 40. The first year data (blue triangles) had another outlying point near $\ell = 210$, which has stepped back in line in the three year data.

Another motivation to look at deviations from power-law spectra lies in the data itself. A comparison of the 3 year WMAP temperature anisotropy data with the theoretical predictions of the best fit vanilla model reveals that a few data points are slightly at odds with the expectations of the vanilla model, see Figure V.1. With the second data release of the WMAP team after three years of observation, a few peculiarities have vanished, but two prominent features remain: two "glitches" at $\ell = 22$ and 40 and the low-lying quadrupole $\ell = 2$. The latter in particular has spawned a lot of speculation whether it might indicate new physical effects (see, e.g., Refs. [130–133]). Furthermore, the results of a reconstruction of the primordial power spectrum via a deconvolution of the observed angular power spectrum [134, 135] seem to imply that a spectrum with features may provide a significantly better fit to the data when compared with a power-law spectrum.

Whatever the reason, any evidence for departures from a power-law behaviour of the initial perturbations would be a very exciting discovery, for it would bear information about new physics at energy scales that are far beyond the reach of accelerator experiments. The remainder of this chapter will be devoted to the discussion and analysis of models in which the slow roll phase is briefly interrupted, leading to non-standard features in the spectrum, based on our Refs. [4, 6].

1. Chaotic Inflation Step Model

Let us examine in more detail the effect a violation of the slow roll conditions will have on the initial power spectrum. We assume that initially, the inflaton field will be in the slow roll regime. This is a helpful assumption because, as we have seen in Section 1.2, it fixes the initial conditions, removing the arbitrariness of having to choose them by hand. We shall also stipulate that the system returns to a state of slow roll inflation eventually. Otherwise, it would not be likely that inflation lasts long enough.

A particular model which fulfils these requirements, but can violate the slow roll conditions briefly, was considered in Ref. [136], (see also Ref. [137]). The inflaton potential is given by

$$V(\phi) = \tfrac{1}{2} m^2 \phi^2 \left(1 + c \tanh\left(\frac{\phi - b}{d}\right)\right). \tag{V.1}$$

This potential describes standard $m^2\phi^2$ chaotic inflation [99] with a step centred around $\phi = b$. The height of the step is determined by the parameter c, its gradient by d. For positive c, the potential will steepen at $\phi = b$, i.e., we have a “downward” step (see Figure V.2). While negative values of c are in principle possible, one might end up with a local minimum of the potential where the inflaton field can get stuck. We will see below that the phenomenological effects on the spectrum are similar, independent of the sign of c, so we limit our analysis to positive c. Furthermore, we will take the step to be small, $c \lesssim 0.1$ to avoid a dominance of the inflaton field's kinetic energy over its potential energy.

We should remark that this is only a toy model; it is to be regarded as an effective field theory description of a more fundamental theory. One could, for instance, imagine a (multi-field) scenario with a Lagrangian that contains terms like $\lambda_i \chi_i^2 \phi^2$, where the χ_i are scalar fields and the λ_i dimensionless coupling constants. The effective mass squared of ϕ is then determined by the bare mass $m_{\phi,0}$ of ϕ and the vacuum expectation values

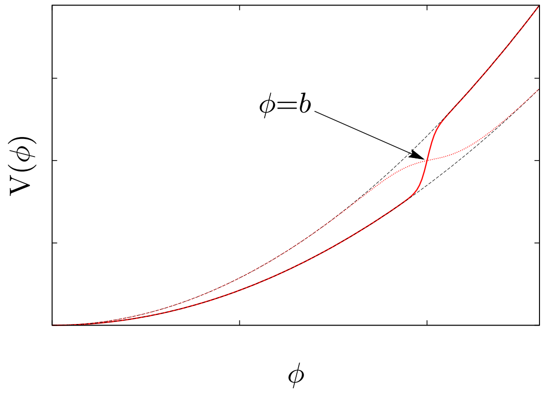

Figure V.2.: Schematic representation of the step potential. The solid red line has $c > 0$, the dotted red line $c < 0$, and the dashed black lines are the asymptotes of the potential.

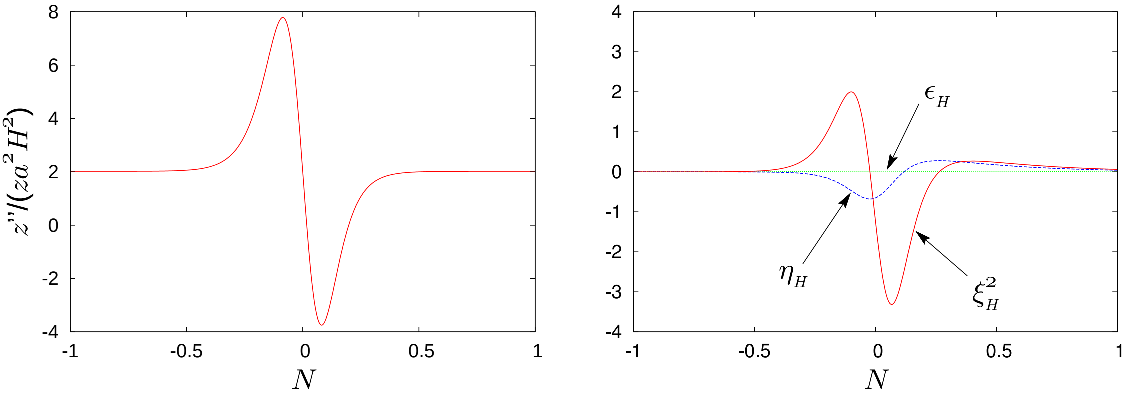

Figure V.3.: *Left*: z''/z divided by a^2H^2 for $b = 14$, $c = 10^{-3}$ and $d = 2 \times 10^{-2}$ versus the number of e-foldings. N is set to zero for $\phi = b$. It takes the inflaton field roughly half an e-folding to roll over the step. *Right*: Hubble slow roll parameters at the step, $\epsilon_{\rm H}$ (dotted green line) remains negligible throughout, while $\eta_{\rm H}$ (dashed blue line) and $\xi^2_{\rm H}$ (solid red line) violate the slow roll conditions.

(VEVs) of the χ_i:

$$m^2_{\phi,\rm eff} = m^2_{\phi,0} + \sum_i \lambda_i \langle \chi_i^2 \rangle . \tag{V.2}$$

If one of the χ_i undergoes a phase transition, its VEV will change and the inflaton field will feel it as a sharp drop (or increase, depending on the sign of λ_i) in its effective mass.

Let us now examine the consequences of adding a step to the potential: As pointed out in Section 1.3, the eventual spectrum crucially depends on the dynamics of z''/z, which can easily be deduced from the solution of Equations (II.17) and (II.18).

For a typical choice of parameters, we plot the numerical solution in the left panel of Figure V.3. Generically, we find that the step causes $z''/(za^2H^2)$ to peak before the inflaton field reaches b, followed by a dip shortly afterwards. It will return to the asymptotic slow roll value of ~ 2 after about one e-folding. Comparison with the Hubble slow roll parameters (right panel of Figure V.3, cf. Equation (II.32)) reveals that this behaviour is mainly caused by $\eta_{\rm H}$ and $\xi^2_{\rm H}$, while $\epsilon_{\rm H}$ remains small (this is a consequence of the condition $c \ll 1$). Beware that the potential slow roll approximation, Equation (II.39), will in general not work for this potential since the contribution of higher derivative terms can be large [138]. The smallness of ϵ_H (and hence ϵ) also implies that there will not be any sizable deviations from a power-law for the spectrum of tensor perturbations.

So, how will this particular behaviour of z''/z influence the solution for u_k and eventually the spectrum compared to a model with no step? From Equation (II.11) it is evident that modes with $k^2 \gg \text{Max}|z''/z|$, i.e., modes that are well within the horizon at the time of the step, will not be affected at all and u_k will remain in the oscillatory regime. For $k^2 \lesssim \text{Max}|z''/z|$, on the other hand, the peak in z''/z will result in a boost of exponential growth for u_k, reverting to oscillations when z''/z goes negative and eventually return to the growing solution. We portray the motion of u_k in the complex plane in Figure V.4.

When an oscillatory phase is preceded by a growing phase, the initial circle will be distorted to an ellipse. As the growth sets in again, the mode will be suppressed or enhanced, depending on the phase of the oscillation, which itself is k-dependent. In the spectrum, this can be observed as oscillations. This mechanism will be most effective for modes that are just leaving the horizon. If $k^2 \ll \mathrm{Max}|z''/z|$, the phase difference will

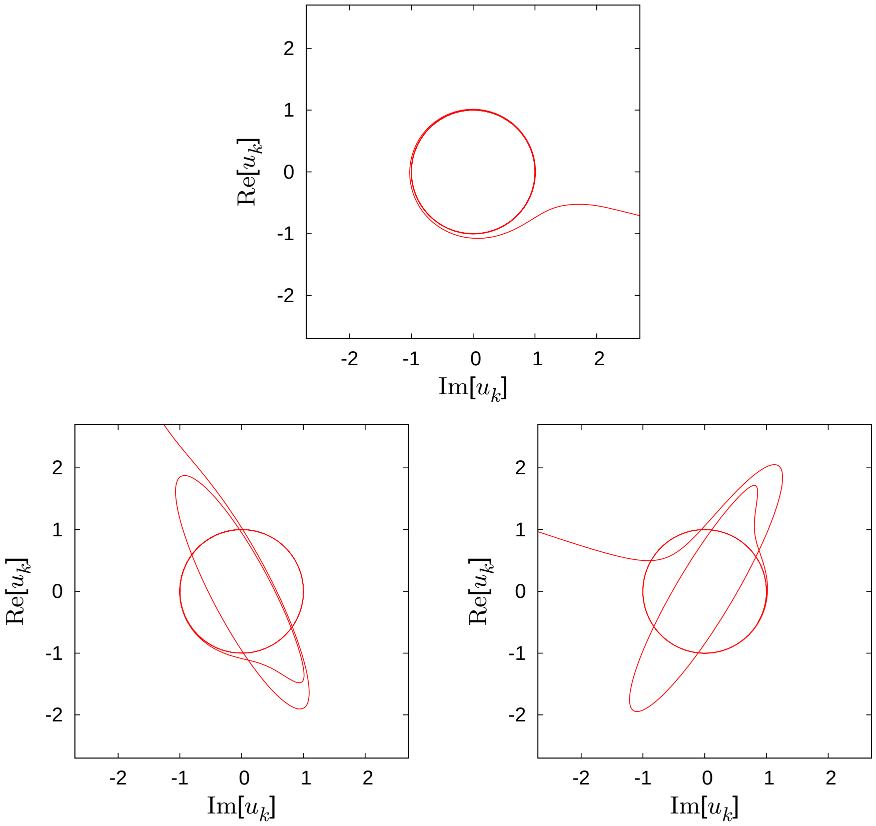

Figure V.4.: These plots show an illustration of the evolution of u_k in the complex plane. We have normalised u_k to one in the oscillating limit. The choice of initial conditions (II.26,II.27) ensures that the motion will initially be circular. The top plot shows a mode that is not affected by the feature; the circular oscillation goes straight into a growing motion. In the other two plots the circle gets deformed by an intermittent phase of growth triggered by the peak of z''/z, to be followed by another phase of elliptic oscillations (caused by the dip of z''/z) until finally the modes leave the horizon and start growing. Whether a mode is suppressed or enhanced by this mechanism depends on the phase of the oscillation when the growth sets in. Growth along the semi-major axis will lead to an enhancement (bottom left), whereas growth along the semi-minor axis entails a suppression (bottom right) with respect to the modes of the corresponding featureless model.

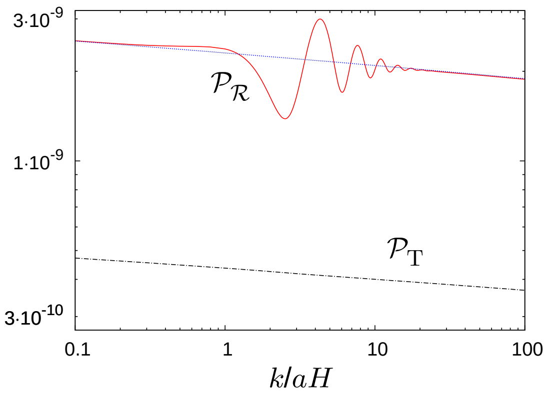

Figure V.5.: Primordial power spectrum for a model with $m = 7.5 \times 10^{-6}$, $b = 14$, $c = 10^{-3}$ and $d = 2 \times 10^{-2}$ (solid red line) with wavenumber k given in units of $aH|_{\phi=b}$. The dot-dashed black line depicts the spectrum of tensor perturbations. For comparison we also plot the scalar spectrum of the same model with c set to zero (dotted blue line).

be negligible.

The key to obtaining an oscillating spectrum is that $z''/(za^2H^2)$ has a local minimum; the deeper, the more pronounced the effect. This forces growing modes back into the oscillating regime, thus inducing a phase dependence when growth sets in again at a later time. In that sense potentials with an "upward" step (i.e., $c < 0$), where $z''/(za^2H^2)$ also exhibits a minimum, will lead to qualitatively similar spectra, provided the field does not get stuck.

Hence, a localised feature in the potential will lead to a localised "burst" of oscillations in the scalar spectrum (see also Ref. [139]), while large and small scales will remain unchanged with respect to the spectra of the asymptotic background models. This is shown in Figure V.5 along with the spectrum of tensor perturbations, which does not appreciably differ from a power-law, due to the smallness of ϵ. Note that the wavelengths affected by the feature are those that are about to leave the horizon as the inflaton field reaches the centre of the step. In particular, the frequency of the oscillations of the spectrum is proportional to this scale.

What remains is to identify the horizon size at the step with a physical scale today. This connection can be made if one knows the total number of e-foldings N_* of inflation that took place after a known physical scale k_* left the horizon. Technically, we evolve the background equations (II.17,II.18) until the end of inflation N_end, (defined by $\ddot{a}(N_\mathrm{end}) = 0$). The scale k_* can then be determined in units of $aH|_{\phi=b}$ via

$$k_* \leftrightarrow \frac{a(N_\mathrm{end} - N_*)\, H(N_\mathrm{end} - N_*)}{aH|_{\phi=b}}. \tag{V.3}$$

As long as the spectrum of the $c = 0$ background model is only mildly scale dependent, there will be a strong degeneracy between N_* and b: shifting the feature in the potential will have the same effect as shifting the scale of k. In the following we will therefore not treat N_* as a free parameter, but set $N_* = 50$ for $k_* = 0.05\ \mathrm{Mpc}^{-1}$. If we want the feature to affect scales that are within reach of current observations, this will require b to lie in the interval $14 \lesssim b \lesssim 15$.

2. Model Dependence

Having analysed a specific step potential in the previous section, let us now address the question of model dependence: Will we arrive at different conclusions if we modify the background inflationary model (e.g., $\lambda\phi^4$ instead of $m^2\phi^2$) or the parameterisation of the step?

We argue that a more general potential

$$V(\phi) = V_0 + f(\phi)\ S(\phi - b) \tag{V.4}$$

leads to a qualitatively similar spectrum as the potential (V.1). Here, $V_{\mathrm{bg}}(\phi) \equiv V_0 + f(\phi)$ is the background potential, which fulfils the slow roll conditions with f and V_0 positive definite. The function $S(\phi)$ parameterises the step, and should monotonically asymptote to $1 \pm c$ ($c \ll 1$) for $\phi \gg b$ and $\phi \ll b$, respectively, with $S(0) = 1$.

As we have seen above, the derivatives of the potential are crucial to determining the spectrum. In general, the derivatives of V are given by a simple application of the Leibniz rule

$$V^{(n)}(\phi) = \sum_{i=0}^{n} \binom{n}{i} f^{(i)}(\phi)\ S^{(n-i)}(\phi). \tag{V.5}$$

Far away from the step, the derivatives of S will be negligible and the potential and its derivatives are approximately

$$V(\phi) \simeq V_0 + f(\phi)(1 \pm c) \simeq V_{\mathrm{bg}}(\phi), \tag{V.6}$$

$$V^{(n)}(\phi) \simeq f^{(n)}(\phi)(1 \pm c) \simeq V_{\mathrm{bg}}^{(n)}(\phi). \tag{V.7}$$

Since the slow roll conditions hold here, the spectrum will be given by Equation (II.41) with

$$A_{\mathrm{S}} \simeq A_{\mathrm{S}}^{\mathrm{bg}} \left(1 \pm c\left(\frac{3f}{V_0 + f} - 2\right) + \mathcal{O}(c^2)\right), \tag{V.8}$$

$$n_{\mathrm{S}} \simeq n_{\mathrm{S}}^{\mathrm{bg}} \pm c\left(\frac{2V_0}{V_0 + f}\left(\eta^{\mathrm{bg}} - 6\epsilon^{\mathrm{bg}}\right)\right) + \mathcal{O}(c^2). \tag{V.9}$$

In the special case $V_0 = 0$, we have exactly $A_{\mathrm{S}} = A_{\mathrm{S}}^{\mathrm{bg}}(1 \pm c)$ and $n_{\mathrm{S}} = n_{\mathrm{S}}^{\mathrm{bg}}$. If $V_0 \neq 0$, there are additional corrections of order c to the normalisation and also corrections to the tilt, which are suppressed by c and the slow roll parameters of the background model. In both cases, one asymptotically recovers the spectrum of the background model in the limit $c \ll 1$.

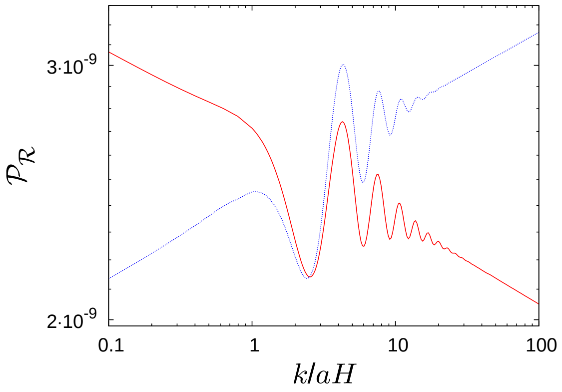

Figure V.6.: Primordial power spectra of a hybrid inflation type step model (V.10) with $V_0 = 3.7 \times 10^{-14}$, $m = 3.2 \times 10^{-8}$, $b = 0.0125$, $c = 10^{-3}$ and $d = 5 \times 10^{-5}$ (dotted blue line), and of potential (V.11) with parameters $\lambda = 6 \times 10^{-14}$, $b = 21$, $c = 5 \times 10^{-4}$ and $d = 0.02$ (solid red line). The hybrid inflation background model has $n_\mathrm{S} > 1$, i.e., an excess of power on small scales (in analogy to visible light spectra, this is also known as a "blue" spectrum), while the $\lambda\phi^4$ model has $n_\mathrm{S} < 1$ with more power on large scales ("red" spectrum).

Near the step, however, the derivatives of V will have a contribution from the derivatives of S. If the step is sharp enough, the nth derivative of V will be dominated by the nth derivative of S, since the other terms are suppressed with factors of the order of the slow roll parameters of the background model. Hence, the dynamics of z''/z near the step is practically independent of the background, but is determined by the form of S. On the other hand, any S that gives a z''/z which roughly shows a behaviour like the one depicted in Figure V.3, will lead to a burst of oscillations in the power spectrum. The similarities between spectra of different background models are illustrated in Figure V.6, where we plot the spectra of a hybrid inflation type potential

$$V(\phi) = V_0 + \tfrac{1}{2}\,m^2\phi^2 \left(1 + c\tanh\left(\frac{\phi - b}{d}\right)\right), \tag{V.10}$$

and another monomial potential with a different form of the step function

$$V(\phi) = \lambda\,\phi^4 \left(1 + c\arctan\left(\frac{\phi - b}{d}\right)\right). \tag{V.11}$$

Note that despite the difference in background models and step functions, the maxima and minima of the oscillations occur at the same wavelengths.

To alleviate the model dependence of the analysis when confronting theory with experiment, we choose a phenomenological approach and define the spectrum of a generalised step model as

$$\mathcal{P}_\mathcal{R}{}^{\mathrm{gsm}} = \mathcal{P}_\mathcal{R}{}^{\mathrm{step}} \left(\frac{k}{k_0}\right)^{n_\mathrm{S} - n_\mathrm{S}^{\mathrm{step}}}. \tag{V.12}$$

Here, ${\mathcal{P}_{\mathcal{R}}}^{\text{step}}$ is the spectrum obtained from the potential (V.1) and $n_{\text{S}}^{\text{step}} = 0.96$ is the spectral index of the $\frac{1}{2}m^2\phi^2$ model. The quantity n_{S} then describes the overall effective tilt of the spectrum. Spectra of this type will arise from potentials of the form of Equation (V.4). While the fine details of particular models may differ slightly from this approximation, Equation (V.12) will nevertheless capture the broad features of a large class of background models, since, as argued above, the shape of the burst of oscillations is largely independent of the background model. Minor differences would likely be washed out in the angular power spectrum of the CMB anyway [140]. The asymptotic behaviour of models with $V_0 = 0$ will be reproduced exactly; for $V_0 > 0$ it will be approximate, with errors of order c.

There is a catch however: in this analysis the parameters b, c and d will be bereaved of their meaning as parameters of the potential. Instead, they should be interpreted as phenomenological parameters which describe the spectrum. This does not preclude us from deriving meaningful constraints, though. We argue that the shape of the modulation of the spectrum is largely independent of the background, so similar modulations should be the consequence of similar step dynamics. A useful quantity in this context is the maximum value the slow roll parameters ϵ, η and ξ^2 can reach at the step. For the potential (V.1), we can estimate $\epsilon_{\max}$, $\eta_{\max}$ and $\xi^2_{\max}$ in terms of b, c and d:

$$\epsilon_{\max} \simeq \epsilon^{\text{bg}} + \frac{c^2}{2d^2} + \frac{2c}{bd}, \tag{V.13}$$

$$\eta_{\max} \simeq \eta^{\text{bg}} + 0.77\,\frac{c}{d^2}, \tag{V.14}$$

$$|\xi^2_{\max}| \simeq 2\,\frac{c^2}{d^4} + 4\,\frac{c}{bd^3}, \tag{V.15}$$

assuming $c < 1$, $d < 1$ and $b > 1$. Note that $\xi^2 = 0$ for the background model.

Along the same lines, one can replace b with k_{s}, corresponding to today's wavenumber of the perturbations that left the horizon during inflation when $\phi = b$. Another interesting quantity one may want to constrain is the number of e-foldings ΔN_{s} it takes the field to roll across the step. To get a rough order of magnitude estimate we can approximate the potential to be linear around b and define ΔN_{s} as the time it takes the field to roll from $\phi_0 \equiv b + \chi d$ to $\phi_0 \equiv b - \chi d$, i.e., the time it takes to traverse a fraction $\tanh\chi$ of the step's height. Assuming that the attractor condition is fulfilled, we obtain

$$\Delta N_{\text{s}} \approx \frac{2\chi^2 b^2 d^2}{2\chi bd + c\,(b^2 + \chi^2 d^2)\tanh\chi}\,. \tag{V.16}$$

We shall set $\chi = 3/2$, corresponding to $\tanh\chi \simeq 0.9$.

3. Data Analysis

Having explored the phenomenology of step-type inflaton potentials, we now turn our attention to the observational side. We would like to derive limits on the step parameters and, ultimately, find out if current cosmological data hold any evidence for features in the spectrum. The analysis of the previous chapter has shown that adding extra free parameters to a model can significantly affect the errors of the standard parameters. We will therefore also pay attention to possible biasing of cosmological parameter estimates in step models.

3.1. Models

We consider three different models for our analysis which differ only in the shape of the initial power perturbations. The energy content is taken to be parameterised by the minimal two free parameters $\Omega_\mathrm{b}h^2$ and $\Omega_\mathrm{c}h^2$. In addition to that, we allow the ratio of sound horizon and angular diameter distance at decoupling θ_s,[1] and the optical depth to reionisation, τ, to vary. The initial spectra of the models are determined by:

- **Step model**:
  The spectrum is calculated from the inflaton potential (V.1), the free parameters are b, c, d, and the overall normalisation A_S (which is directly proportional to m^2).
- **Generalised step model**:
  The spectrum is given by Equation (V.12). This model uses the same parameters as above, with the effective tilt n_S added as additional parameter.
- **Vanilla model**:
  For comparison, we consider the vanilla model with free parameters A_S and n_S.

We restrict our analysis to scalar perturbations. While tensor perturbations may, in principle, give a subdominant contribution, their spectrum will be smooth in the class of models studied here, so we do not expect any major degeneracies with the step parameters. The same holds true for other possible extensions of the vanilla model, such as massive neutrinos.

3.2. Data Sets

If there are oscillations present in the spectrum of primordial curvature perturbations, we can expect them to leave traces in the temperature and E-polarisation anisotropies of the CMB (but not in the B-polarisation!) as well as in the matter distribution. To

---

[1] It is advantageous to use this parameter instead of H_0 in MCMC analyses, since it has less degeneracies with other cosmological parameters when considering CMB data. Once the posterior in θ_s is known, it is straightforward to reconstruct the limits on H_0 [61].

assess the influence of different data on the constraints, we perform the analysis for each of the models using three basic combinations of data sets:

1. **WMAP only**:
   We use the WMAP three year temperature and polarisation anisotropy data [23–26]. The likelihood is determined using the WMAP likelihood code available at the LAMBDA website [100].

2. **CMB+LSS**:
   In addition to WMAP3 we use small scale CMB temperature anisotropy data from the ACBAR [141], BOOMERANG [142] and CBI [143] experiments, the SNLS SNIa data [83], plus the LSS power spectrum of galaxies from:

   a) the 2003 data release of the Sloan Digital Sky Survey [144].
   b) the 2dF Galaxy Redshift Survey [145].
   c) the LRG sample of the SDSS [19].

   In all cases we treat the luminous to dark matter bias parameter b' as a nuisance parameter that will be analytically marginalised over. All three data sets include points at scales which can no longer be treated in linear perturbation theory. Since the typically applied nonlinear corrections are calibrated under the assumption of a smooth spectrum, we prefer to limit the analysis to scales where the nonlinear corrections do not play an important rôle, i.e., $k \lesssim 0.2\ h\mathrm{Mpc}^{-1}$ for a) and b), and $k \lesssim 0.09\ h\mathrm{Mpc}^{-1}$ for c).

3. **CMB+BAO**:
   Same as data set 2c, plus two-point correlation function data from the SDSS LRG [20]. The correlation function is, in principle, particularly well suited to detecting even small amplitude oscillations of the spectrum: since the correlation function is essentially the Fourier transform of the power spectrum, oscillations of the spectrum will correspond to a peak in the correlation function. However, due to biasing and weakly nonlinear structure formation, this data is difficult to interpret and its application to our problem is not immediately straightforward. In particular, the effect of mode coupling which leads to a smoothing of features in the matter power spectrum needs to be considered. We treat this issue by marginalising over a smoothing parameter. For the technical details we refer the reader to Appendix A.3.

3.3. Priors

In the next section we will see that the likelihood function in this case has a rather unfortunate property: In certain parameter directions it does not go to zero, but rather converges to a constant value. Take for instance the parameter d: in the limit $d \to \infty$, the step will become arbitrarily smooth and the resulting spectrum will not be distinguishable from a power-law spectrum, which, as we know, gives a reasonable fit to the data and hence has a non-negligible likelihood. A similar reasoning can be applied to the parameters b and c. So, no matter how we choose the limits of our priors, we will always chop off a part of parameter space which contains a non-negligible fraction of the total volume of the likelihood function. Therefore, all confidence limits we derive from the resulting posteriors, and particularly those in the b-, c-, and d-subspaces of parameter space will be subject to how we set the limits of our priors and should only be taken as rough indicators. Having said that, let us now attempt to tackle the delicate task of assigning the priors:

b: The parameter b determines the scale at which the oscillations appear. To make sure this scale is in the range of sensitivity of our data, we choose a flat prior: $b \in (14, 15)$.

c: A flat prior on c itself would give a bias to large values, so we take a flat prior on the logarithm of c with $\log c \in (-6, -1)$. For $c \lesssim 10^{-5}$, the spectrum is virtually featureless.

d: It is advisable to choose a logarithmic prior for this parameter as well. Instead of $\log d$, we will impose a flat prior on $\log c/d^2 \in (-5, 3)$ (cf. Equation (V.14)). We also exclude steps that are too sharp or too shallow: $-2.5 \leq \log d \leq -0.5$.

The priors on the remaining parameters are of less impact, and we will take the default selection: apart from the hard-coded priors of `cosmomc` on the Hubble parameter ($40\ \mathrm{km\ s^{-1}Mpc^{-1}} < H_0 < 100\ \mathrm{km\ s^{-1}Mpc^{-1}}$) and the age of the Universe, A_{U}, ($10\ \mathrm{Gyr} < A_{\mathrm{U}} < 20\ \mathrm{Gyr}$), we impose flat priors on $\ln A_{\mathrm{S}}$, τ, θ_s, $\Omega_{\mathrm{b}}h^2$ and $\Omega_{\mathrm{c}}h^2$.

4. Results

4.1. Chaotic Inflation Step Model

Let us start with the simplest combination of model and data: the step model and WMAP data only. For now, we will set our sights to the step parameters, b, c and d. In the bottom panel of Figure V.7 we show the mean likelihood in the (b, c)-plane.

It is evident that the likelihood is far from Gaussian in these directions of parameter space. Its shape can best be described in terms of geographical features: we find two local peaks with $b \simeq 14.3$ and $b \simeq 14.8$, respectively. Towards the sides and the bottom edges, the likelihood has a plateau: if we make the step height too small ($c \to 0$) or push the feature out of the observable range of wavelengths (b near the edge), the spectrum will reduce to that of a $m^2\phi^2$ inflation model, so the likelihood will tend towards a constant value. Finally, we find a large and deep valley extending from the centre to

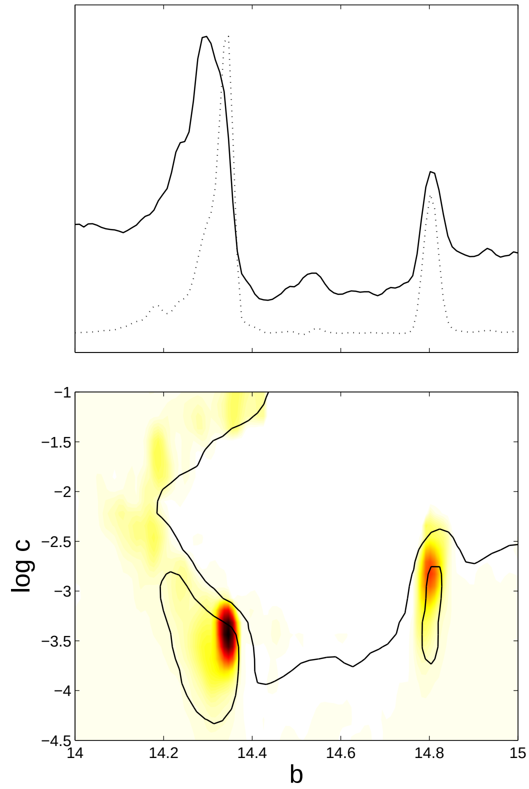

Figure V.7.: This diagram shows the results of a fit of the step model to the WMAP data. *Top*: posterior distribution (solid line) and mean likelihood (dotted line) for parameter b. *Bottom*: The closed lines represent the 15% C.L. contour, the open line is the 99% C.L. confidence contour derived from the posterior distribution in the (b, c)-directions of parameter space. We also show the mean likelihood (colour coded, darker means more likely).

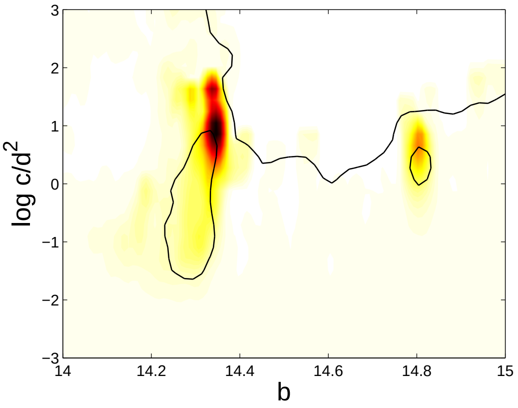

Figure V.8.: Mean likelihood, 15% and 99% C.L. contours in the $(b, c/d^2)$-directions of parameter space for the step model and WMAP data only.

the top right corner, where the fit to the data is extremely bad. The black lines are the 15% and 99% C.L. contours of the marginalised posterior distribution. With our choice of priors, the peaks comprise only a small fraction of the volume of the posterior.

This can also be seen in the top panel of Figure V.7, where we show the marginalised posterior distribution and mean likelihood for b. The peaks are much more prominent in the mean likelihood; this is due to a volume effect from integrating over the plateau, e.g., in the c direction. The parameter d is not very well constrained by the data, due to a degeneracy with c. Instead, we consider the constraints in the $(b, c/d^2)$ parameter space (Figure V.8). This parameter is as well constrained as c and, again, the presence of two maxima for b is evident. Also, the maxima are at values of c/d^2 of order one, where the slow roll conditions are strongly violated; values of $c/d^2 \ll 0.1$ correspond, on the other hand, to the usual slow roll $m^2\phi^2$ inflation and cannot be excluded by the data.

Without a doubt, the two peaks are the most interesting features of the likelihood. Since they occur at different values of b, they will correspond to oscillations at different scales. The larger b, the earlier the affected scales will have left the horizon during inflation and the later they reenter the horizon during radiation or matter domination. Hence, models with large b have features at large scales (low multipoles ℓ), while for models with smaller b the feature will be at smaller scales (higher ℓ). In Figure V.9 we show the primordial power spectra corresponding to the two local maxima. The best fit of the step model improves the fit of the vanilla model by $\Delta\chi^2_{\text{eff}} \sim -7$ and has oscillations at small scales. For the lower peak of the likelihood we find $\Delta\chi^2_{\text{eff}} \sim -5$ with oscillations at larger scales. This improvement comes exclusively from the temperature anisotropies. Given current experimental errors, neither the E-polarisation anisotropies nor the cross-correlation have the power to constrain steps.

In Figure V.10 we compare the angular power spectra of the two peaks of the step

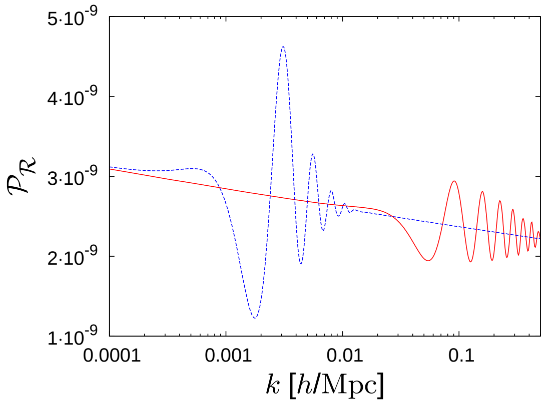

Figure V.9.: Primordial power spectra of the best fit step model (solid red line, $b = 14.34$, $\log c = -3.4$, $\log d = -2.21$) and of the best fit model in the smaller peak (dashed blue line, $b = 14.8$, $\log c = -3.05$, $\log d = -1.76$), using data set 1.

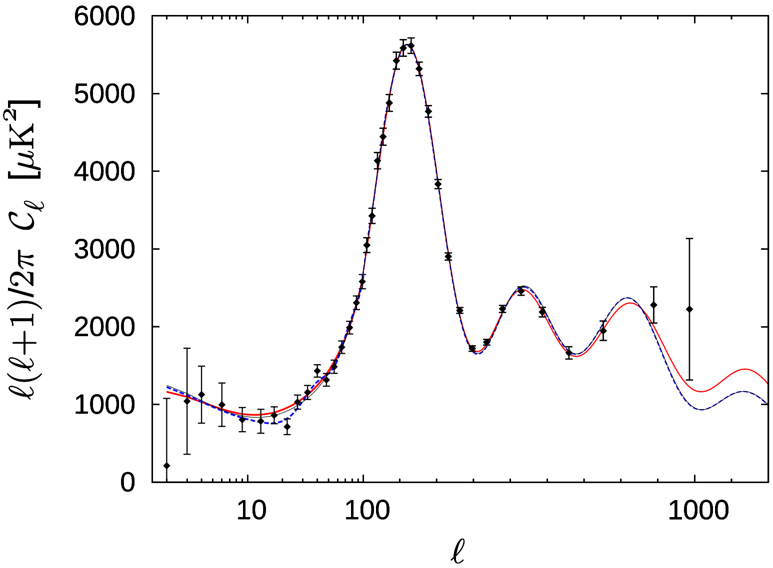

Figure V.10.: Angular power spectrum of CMB temperature anisotropies. Here we combine experimental error (i.e., noise) and cosmic variance in the error bars of the data points. The thin black line is the best fit of the vanilla reference model, the thick red line corresponds to the best fit of the step model (maximum near $b \simeq 14.3$) and the thick dashed blue line represents the $b \simeq 14.8$ maximum of the step model.

model with the best fit of the vanilla model. Over wide ranges of multipoles the curves are virtually identical. The $b \simeq 14.3$ spectrum shows deviations from vanilla near the third peak, the $b \simeq 14.8$ one differs between $\ell \sim 10 \rightarrow 50$. Incidentally, this is precisely where the binned data has two glitches, so we can attribute the improvement in χ^2_{eff} to the ability of step models with $b \simeq 14.8$ to explain the presence of these outlying data points. It also coincides with the maximum found in Ref. [84], where the same model was analysed with fixed cosmological parameters. As it happens, the small scale maximum corresponds to oscillations of the primordial power spectrum that have roughly the same wavelength as the anisotropies induced from acoustic oscillations of the plasma. This explains why there is no apparent superimposed oscillation in the angular power spectrum, as is the case with the large scale maximum.

Moreover, the $b \simeq 14.3$ maximum occurs at the extreme small scale end of the sensitivity of the WMAP data. This is not totally unexpected; at the outer ends of the data one has more freedom to modify the spectra, since a large part of the modification can be hidden at unobservable wavenumbers. So it is not unlikely to find a spectrum that gains its advantage from fitting just the last few data points better than the vanilla prediction. Hence, it is only natural to be sceptical about the significance of this result and ask for further evidence, considering also that the WMAP signal is rather noisy at scales near the third peak.

To either confirm or refute the small scale feature we repeat the analysis with addi-

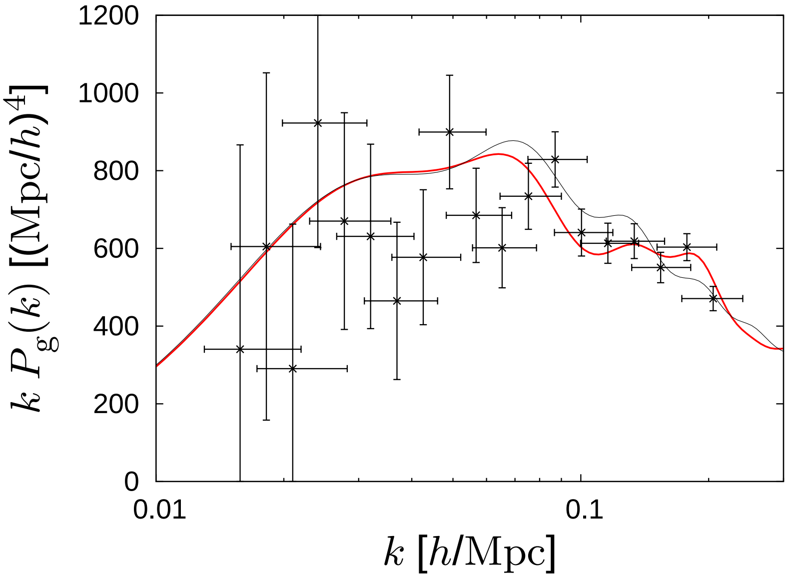

Figure V.11.: Galaxy power spectrum of the best fit step model (thick red line, using data set 2a) and the best fit vanilla model (thin black line) plotted against the 2003 SDSS main sample data. Note that a "chi by eye" estimate can be misleading here; this representation does not take into account that the spectra have to be convoluted with the window functions and that the normalisation will be marginalised over.

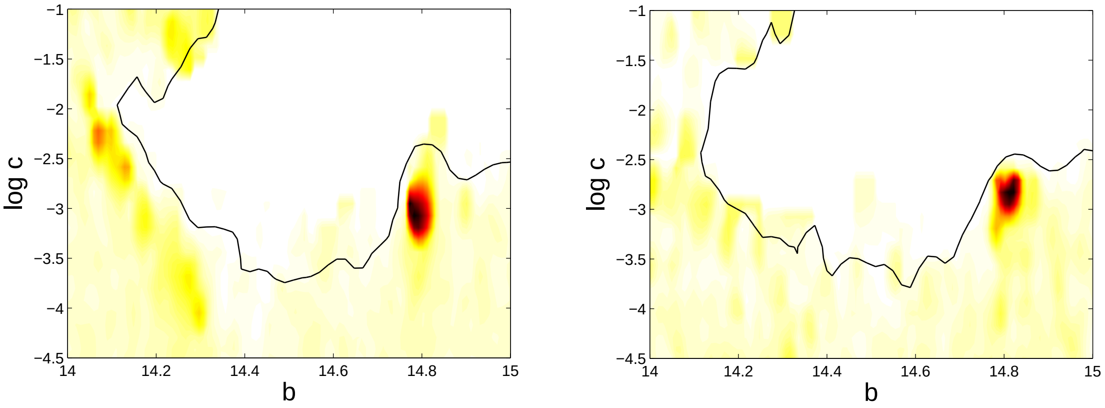

Figure V.12.: Mean likelihood and 99% confidence level contours for the step model in the (b, c)-plane of parameter space with data sets 2c (left) and 3 (right).

tional data sets that are sensitive to smaller scales, i.e., data sets 2 and 3. Using data set 2a, we could indeed confirm a feature with $b \simeq 14.3$, the fit even improved to a $\Delta\chi^2_{\text{eff}}$ of ~ -15 with respect to the vanilla model. Most of this improvement can be attributed to the SDSS data; we show the best fit matter power spectrum in Figure V.11.

However, using data sets 2b, 2c, or 3, we do not find such an enhancement of the $\Delta\chi^2_{\text{eff}}$ anymore. In fact, the SDSS LRG data even appear to disfavour a large feature near the third peak. Given the better quality of the LRG power spectrum data and the fact that the BAO data also does not seem to support this effect, it is likely that the improvement in the fit was just a fluke. The disappearance of this maximum of the likelihood function is illustrated in Figure V.12 (compare with the bottom panel of Figure V.7), where we show the mean likelihood and the 99% confidence level of the marginalised posterior in the (b, c) plane of parameter space. The inclusion of LRG and BAO data considerably tightens the constraints on features at small scales corresponding to values of b between $\sim$14.1 and $\sim$14.4, while for larger values of b, i.e., features at larger scales, the contours remain roughly the same. Figure V.13 shows the consequences of having an oscillatory feature in the power spectrum: the correlation function will have an additional peak at scales corresponding to the wavelength of the oscillations. With the current quality of LSS data, detecting a peak in the correlation function is comparatively easier than finding small amplitude oscillations in the power spectrum.

The feature at large scales ($b \simeq 14.8$), on the other hand, remains untouched by the addition of these data sets, its $\Delta\chi^2_{\text{eff}}$ neither improving nor worsening. This is plausible since the added data is not sensitive to features appearing at $\ell \simeq \mathcal{O}(10)$.

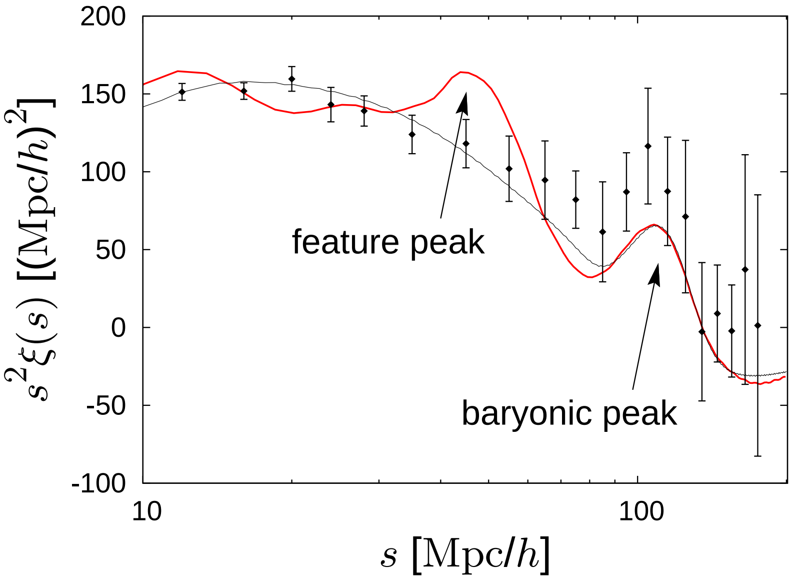

Figure V.13.: Galaxy correlation function for the vanilla model (thin black line) and a step model with $b = 14.3$, $\log c = -3.5$ and $\log d = -2$. The peak at a comoving separation of $s \sim 100$ Mpc/h is due to baryon oscillations, the one at $s \sim 45$ Mpc/h is caused by the oscillations in the primordial spectrum. The scale of this second peak is mainly determined by b. Note that the data points are correlated and the normalisation will be marginalised over. For the step correlation function, we additionally marginalise over the smoothing effect due to mode coupling. The stronger the smoothing, the lower the amplitude of the feature peak.

4.2. Generalised Step Model

So far we have only discussed the simple chaotic inflation step model. Let us now turn our attention to the generalised step model. This model is a direct subset of the vanilla model, e.g., for $c = 0$ their predictions coincide. As we have seen in Chapter IV, in the context of a model-dependent analysis it is an important question how the choice of model will affect the estimates of the parameters, particularly if the models are nested, as is the case here. Possible degeneracies between "standard" and newly introduced parameters can bias means as well as errors. In Figure V.14, we plot the marginalised likelihood distributions for the vanilla parameters for all three models with data set 1. We notice small differences between the step model and its generalised sibling for $\Omega_b h^2$, τ and the normalisation. These arise due to the fact that the tilt of the spectrum is fixed in the chaotic inflation step model. There is a well-known degeneracy between these parameters and the spectral index; fixing the tilt near the best fit value will reduce the errors on the parameters it is degenerate with.

The distributions for the generalised step and vanilla model exhibit a remarkable similarity which leads us to conclude that the presence of a feature will not have any statistically significant influence on the results for the parameters of the vanilla model. This conclusion remains unchanged if we consider the other data sets.

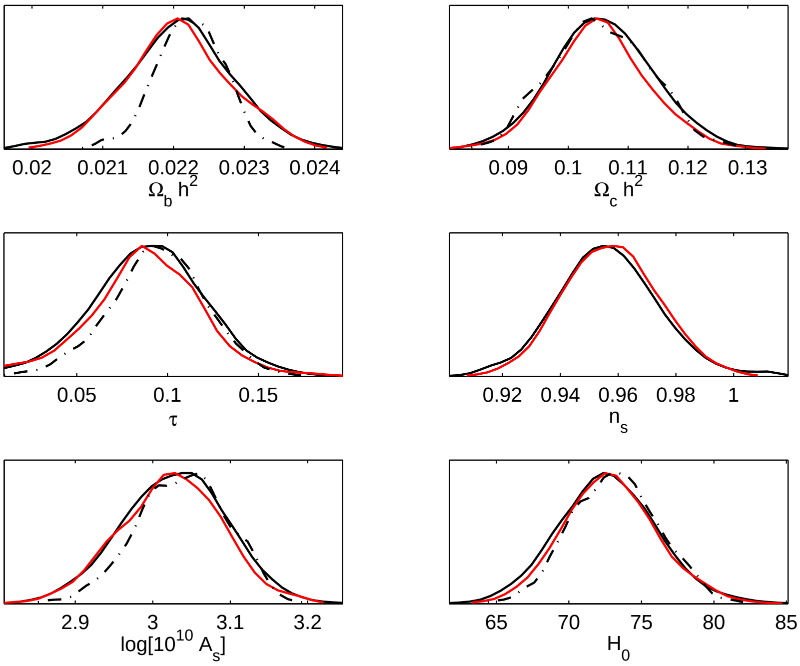

Figure V.14.: Marginalised posterior distributions for the vanilla model (solid red line), the step model (dot-dashed line) and the generalised step model (solid black line).

For the generalised step model and data set 1, the large scale feature maximum likelihood point is at ($b = 14.8$, $c = 0.001$, $d = 0.02$, $\Omega_\mathrm{b}h^2 = 0.0216$, $\Omega_\mathrm{c}h^2 = 0.102$, $\tau = 0.11$, $n_\mathrm{S} = 0.952$, $\ln[10^{10}A_\mathrm{S}] = 3.05$ and $H_0 = 72.7$), which lies near the maximum of the marginalised 1D posteriors of the vanilla model in Figure V.14. This is a further indication that the presence of a feature at large scales will not affect the estimates of the other parameters.

As we noted earlier, the parameters b, c and d lose their meaning in this phenomenological model, but we can still find constraints on derived quantities which may even be more intuitive. An example is the wavenumber k_s of the perturbations that left the horizon when the inflaton field passed the step (i.e., at the moment when $\phi = b$), a measure for the scales affected by the oscillations. We show the marginalised posterior and mean likelihood for k_s in Figure V.15. Qualitatively, the results are very similar to what we found for the simpler step model (cf. Figure V.7). Again, we can see how the inclusion of data sets sensitive to smaller scales reduces the evidence for a feature at scales $\gtrsim \mathcal{O}(10^{-2})$ Mpc^{-1}, with only a feature at large scales remaining under the most restrictive data set.

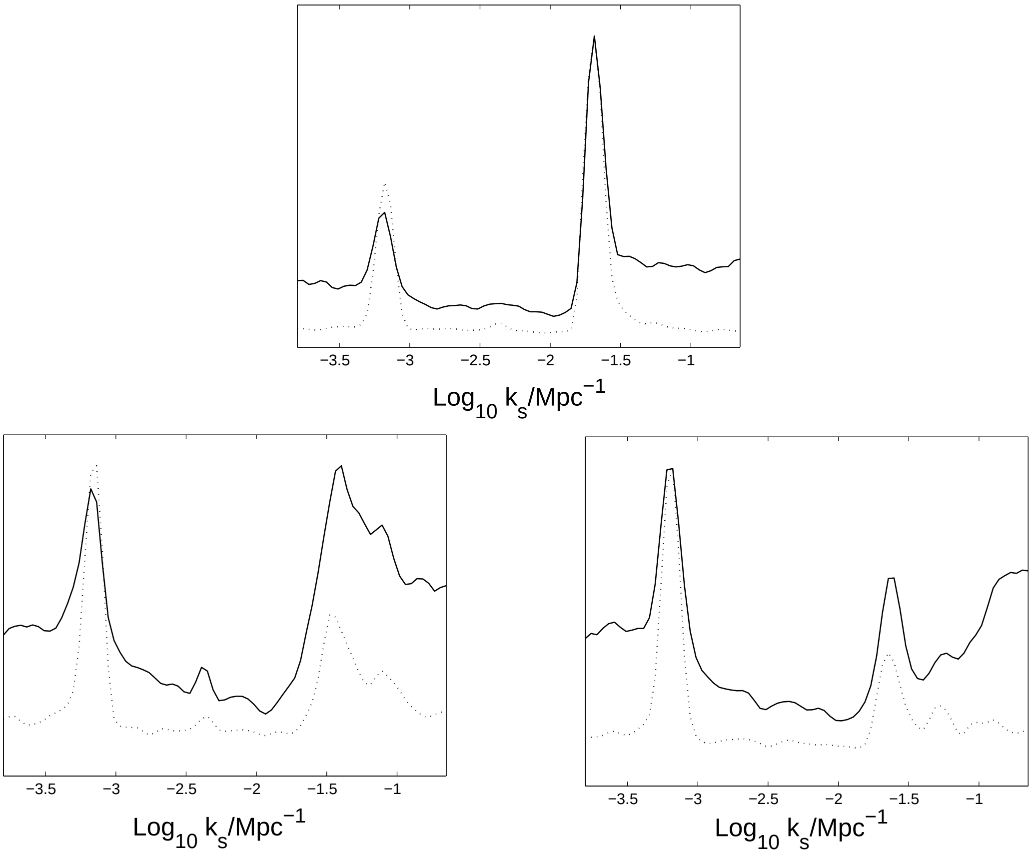

Figure V.15.: Marginalised posterior (solid line) and mean likelihood (dotted line) for parameter k_s in the generalised step model, indicating at which wavelengths a feature is likely to happen. *Top*: data set 1. *Bottom left*: data set 2c. *Bottom right*: data set 3.

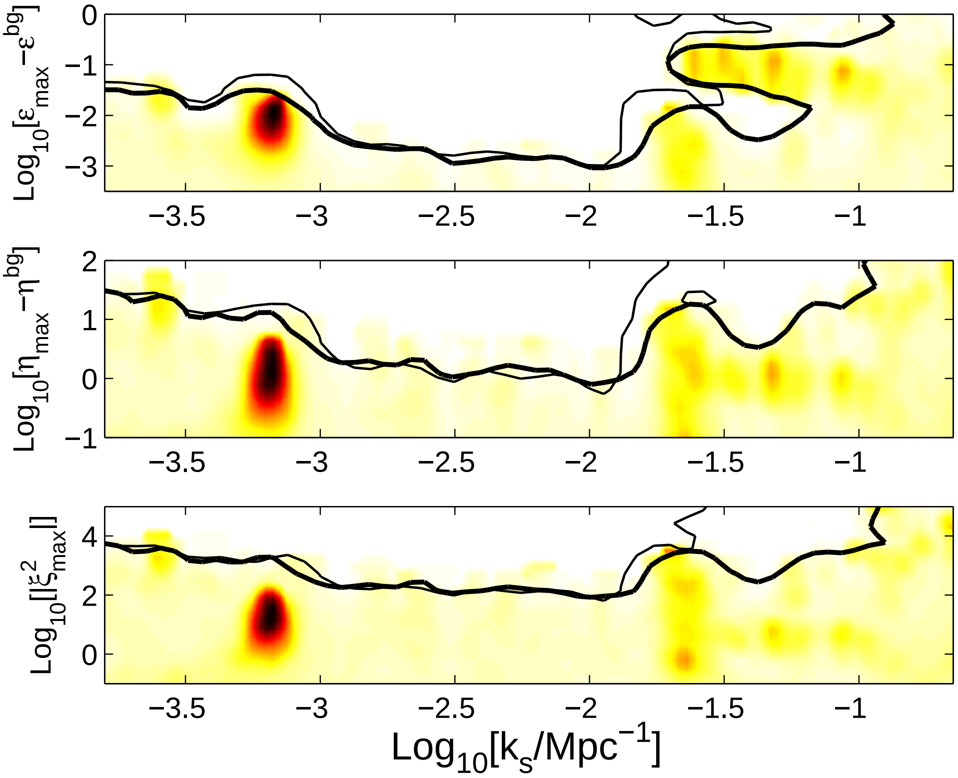

Figure V.16.: This plot shows the constraints on the peak values of the slow roll parameters during the step for the generalised step model. The thick lines denote the 99% confidence level for data set 3, the thin lines correspond to data set 1.

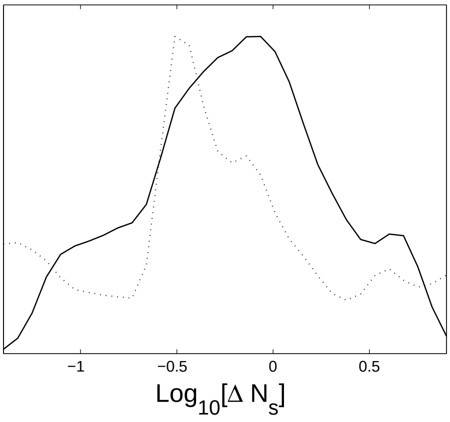

Figure V.17.: Marginalised posterior (solid line) and mean likelihood of the number of e-foldings it takes the inflaton field to traverse 90% of the step, using data set 3. Note that the prior is not flat in ΔN_{s}.

We show the constraints on the maximum values of the slow roll parameters of the step function in Figure V.16. While the WMAP3 data alone is only sensitive to features up to a wavelength of $\sim 10^{-2}$ Mpc^{-1}, the large scale structure data extends the sensitivity by almost a factor ten in k. We find fairly strong bounds on the maximum value of ϵ for the step function. In conjunction with Equation (II.55), this implies that the spectrum of tensor perturbations is unlikely to experience an oscillatory modulation like the scalar spectrum, since that would require ϵ to be of order one.

For the higher order slow roll parameters, values up to a few (for η) and up to a few hundred (for ξ^2) are still allowed by the data. Note, however, that these bounds depend on the parameterisation (they assume a tanh-form of the step), and, for $\eta \gtrsim 1$, not only ξ^2, but also higher order potential slow roll parameters will be non-negligible.

It is also interesting to look at how long it takes the inflaton field to roll over the step. We defined a rough estimate for the number of e-foldings in Equation (V.16), its marginalised posterior and mean likelihood are displayed in Figure V.17 for data set 3. There seems to be a slight preference for ΔN_{s} of $\mathcal{O}(10^{-1} \rightarrow 1)$, which is a realistic order of magnitude for the duration of phase transitions, but the difference is not big enough to be conclusive. The posterior in particular should be taken with care, since our choice of priors does not correspond to a flat prior on ΔN_{s}, thereby possibly introducing a systematic bias.

4.3. Evidence for Features?

Finally, let us address the question whether there is any evidence for the data to prefer a spectrum with features over a smooth one. The step model introduces three new parameters and fixes one (n_{S}) with respect to vanilla. In the generalised step model

	Step Model		Generalised Step Model		
Data Set	1	2a	1	2c	3
$\Delta\chi^2_{\text{eff}}$	-7	-15	-8	-7	-7
ΔAIC	-3	-11	-2	-1	-1
ΔBIC	10	2	18	19	19

Table V.1.: Akaike and Bayesian information criteria for a comparison of the two step models with the vanilla model.

we reinstate n_{S}, which yields, depending on the data set used, an additional $\Delta\chi^2_{\text{eff}}$ improvement of $1-2$. We did not expect a major improvement here, since the spectral tilt of the $m^2\phi^2$ model lies fairly close to the best fit value of the vanilla model with a freely varying n_{S}.

We summarise the results and the approximate ΔAIC and ΔBIC in Table V.1. Negative values indicate a preference for the step models. We find that for almost all data sets, except the peculiar set 2a, the AIC shows only a tiny preference for the step models, whereas the BIC appears to rule out decisively the necessity of including a step. This can be explained with the large number of data points; the WMAP data alone consists of 5249 data points (pixel data and pseudo-$\mathcal{C}_\ell$), giving a large penalty to extra parameters. Doubtlessly, a more stringent model comparison as outlined in Section III.4 would also overwhelmingly favour the simplicity of the vanilla model, due to the large amount of fine-tuning involved in the step models.

5. Discussion

We have performed a detailed analysis of the phenomenology of inflationary models with a small step-like feature in the inflaton potential.

Generically, the resulting spectrum of scalar perturbations will resemble that of the stepless background model with a superimposed burst of oscillations whose shape is determined by the form of the step only. We have confronted the theoretical predictions for the spectrum of a specific chaotic inflation model with a step with recent cosmological data to find out whether the data require the presence of such a feature and whether it may actually bias the estimates of other cosmological parameters such as, e.g., the baryon density. We have also repeated the same analysis for a more empirical but less model dependent spectrum, which might be expected from a step in an arbitrary inflationary background model.

With a combination of different data sets, a large chunk of the step model parameter space can be ruled out; only spectra with a very modest oscillation amplitude are still consistent with observations. The BAO data, in particular, prove to be a very sensitive probe for oscillating spectra.

Compared to the six-parameter vanilla model, using the most constraining data set, we find an improvement of the best fit χ^2_{eff} of about -5 for the chaotic inflation step model which comprises two extra parameters, and $\Delta\chi^2_{\text{eff}} \sim -7$ for the generalised step model, which has three extra parameters.

The best fit region of parameter space consists of models which show oscillations at wavelengths corresponding to multipoles $\ell \simeq \mathcal{O}(10)$, where the temperature-temperature correlation data of the CMB shows some glitches. Interestingly, the time it would take the inflaton field to traverse the step in these models is of the order of a few tenths of an e-folding, which is what one would expect for the time of a phase transition in more realistic multi-field models.

Whether the glitches are just statistical flukes and can be attributed to cosmic variance or stem from a physical effect, such as a feature in the inflaton potential, cannot be conclusively decided until we have better measurements of the E-polarisation spectra from experiments like PLANCK [48] or, in the more distant future, projects like the Inflation Probe [146]. An additional consistency check can be provided by an analysis of the bispectrum of CMB fluctuations, since the interruption of slow roll may also induce sizable non-Gaussianities [147, 148].

Conclusions and Outlook

In this work, we have scrutinised two important aspects of extracting information about the physics of the early Universe from data garnered from cosmological observations.

On the one hand, there is the question about the values of cosmological parameters and their uncertainties. We have shown that the conventional "vanilla(+1)" approach can lead to a systematic underestimation of errors and a possible bias due to the very restrictive assumptions inherent to the vanilla model. Our analysis of the data in a more general framework results in a relaxation of existing bounds by up to a factor of two, due to newly introduced degeneracies between standard and additional parameters. In particular, we find a degeneracy between the neutrino mass fraction f_ν and the tensor to scalar ratio r, opening up regions in parameter space that had previously been considered ruled out by other authors. For instance, we find that the $\lambda\phi^4$ chaotic inflation model is still marginally allowed by the data and that it prefers neutrinos with degenerate masses. While our analysis leads to more conservative estimates, it is certainly possible to generalise it even further and include other well-motivated effects, such as isocurvature modes or a time-dependent dark energy equation of state. This may lead to the identification of additional, hitherto unidentified parameter degeneracies, resulting in even weaker bounds.

On the other hand, improvements in the quality of data present an invitation to search for the subtle traces of more unusual physical effects not covered by the concordance model. We have investigated a class of models in which the usual slow roll inflationary regime is briefly interrupted due to a step-like feature in the inflaton potential, possibly triggered by a phase transition during inflation. The resulting brief interlude of fast roll manifests itself by a burst of oscillations in the primordial power spectrum of curvature perturbations. We find very restrictive bounds on the parameters describing these features; in the step model, only very modest deviations from the standard power-law spectrum are allowed by the data. If such features were indeed present in the initial spectrum, they would most likely appear at scales corresponding to multipoles $\ell \sim \mathcal{O}(10)$, offering an explanation for the "glitches" seen in the data. Our results illustrate the amazing robustness of the slow roll inflationary paradigm and how the data noose is tightening around the necks of more exotic models.

In the coming years we will continue to increase our knowledge about our Universe, as new satellite and ground based experiments will start to gather data. We will see a refinement of existing observations as well as data from completely new sources, such as weak gravitational lensing or 21 cm hydrogen spin flip surveys. It will be interesting to see whether new findings will confirm or challenge our current understanding of the Universe. Who knows, maybe Nature still holds a few surprises for us in store!

A. Appendix

1. Perturbations on a Sphere

In analogy to a Fourier expansion in Euclidean space, any square-integrable function f defined on the unit sphere can be expanded in terms of spherical harmonics $Y_{\ell m}$,

$$f(\theta,\varphi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}\, Y_{\ell m}(\theta,\varphi). \tag{A.1}$$

The coefficients $a_{\ell m}$ can be expressed in terms of f via

$$a_{\ell m} \equiv \int \mathrm{d}\Omega\; Y^{*}_{\ell m}(\theta,\varphi)\, f(\theta,\varphi), \tag{A.2}$$

and the $Y_{\ell m}$ are defined by

$$Y_{\ell m}(\theta,\varphi) \equiv \sqrt{\frac{2\ell+1}{4\pi}\,\frac{(\ell-m)!}{(\ell+m)!}}\; P^{\ell}_{m}(\cos\theta)\, e^{\mathrm{i}m\varphi}, \tag{A.3}$$

where P^{ℓ}_{m} is an associated Legendre polynomial [149].

It is often convenient to work with the dimensionless version of f,

$$F(\theta,\varphi) = \frac{f(\theta,\varphi)}{\langle f(\theta,\varphi)\rangle} - 1, \tag{A.4}$$

where $\langle\cdot\rangle$ denotes the ensemble average, and use this to define the *angular correlation function*

$$c(\theta) = \langle F(\theta_1,\varphi_1)\, F(\theta_2,\varphi_2)\rangle. \tag{A.5}$$

If F is an isotropic random field, the right hand side will only depend on $\theta = |\theta_1 - \theta_2|$ and the $a_{\ell m}$ of F satisfy

$$\langle a_{\ell m}\, a_{\ell' m'}\rangle = \delta_{\ell\ell'}\, \delta_{mm'}\, \mathcal{C}_{\ell}. \tag{A.6}$$

The $\mathcal{C}_{\ell}$ are known as the *angular power spectrum* of F. Note that $\mathcal{C}_{\ell}$ is defined as an ensemble average. An observation of the CMB temperature fluctuations, for instance, only gives us one realisation of the random field though. So if one wants to estimate the $\mathcal{C}_{\ell}$ from this observation, for each multipole ℓ one will only have $2\ell+1$ values of m to average over. This induces an additional uncertainty (apart from errors due to experimental noise), which is known as *cosmic variance*

$$\frac{\Delta\mathcal{C}_{\ell}}{\mathcal{C}_{\ell}} = \sqrt{\frac{2}{2\ell+1}}\,. \tag{A.7}$$

This assumes full sky coverage of the data. If only a fraction f_{sky} of the sphere is covered, the cosmic variance increases to

$$\frac{\Delta \mathcal{C}_\ell}{\mathcal{C}_\ell} = \sqrt{\frac{2}{(2\ell+1)\, f_{\mathrm{sky}}}}\,. \tag{A.8}$$

2. Evaluation of the Likelihood

Consider an experiment measuring the values of n observables $\boldsymbol{d}^{\mathrm{obs}} = \left(d_1^{\mathrm{obs}}, ..., d_n^{\mathrm{obs}}\right)$. This is to be compared with the predictions of a model $\mathcal{M}$ for these quantities, given by $\boldsymbol{d}^{\mathrm{th}} = \left(d_1^{\mathrm{th}}, ..., d_n^{\mathrm{th}}\right)$.

We will assume for now that the d_i^{th} are free variables and define a vector

$$\boldsymbol{x} \equiv \boldsymbol{d}^{\mathrm{obs}} - \boldsymbol{d}^{\mathrm{th}}. \tag{A.9}$$

We can now write the likelihood in terms of $\boldsymbol{x}$:

$$\mathcal{L}(\boldsymbol{d}^{\mathrm{obs}}|\boldsymbol{d}^{\mathrm{th}}) = \mathcal{L}(\boldsymbol{x} + \boldsymbol{d}^{\mathrm{th}}|\boldsymbol{d}^{\mathrm{th}}). \tag{A.10}$$

Clearly, at $\boldsymbol{x} = 0$, we have

$$\mathcal{L}(\boldsymbol{d}^{\mathrm{th}}|\boldsymbol{d}^{\mathrm{th}}) = 1, \tag{A.11}$$

hence we must be at a local maximum $\mathcal{L}_{\mathrm{max}}$. Let us now define

$$\chi^2_{\mathsf{eff}} = -2\ln\left(\mathcal{L}/\mathcal{L}_{\mathrm{max}}\right) = -2\ln\mathcal{L} \tag{A.12}$$

and expand this quantity in the x_i:

$$\chi^2_{\mathsf{eff}}(\boldsymbol{x}) = \chi^2_{\mathsf{eff}}(\boldsymbol{x} = 0) + \sum_i^n \left.\frac{\partial\chi^2_{\mathsf{eff}}}{\partial x_i}\right|_{\boldsymbol{x}=0} x_i + \tfrac{1}{2}\sum_i^n\sum_j^n \left.\frac{\partial^2\chi^2_{\mathsf{eff}}}{\partial x_i \partial x_j}\right|_{\boldsymbol{x}=0} x_i\, x_j + \mathcal{O}(3). \tag{A.13}$$

The first term is zero because of Equation (A.11), the second vanishes because $\boldsymbol{x} = 0$ is a local minimum of χ^2_{eff}, so we have

$$\chi^2_{\mathsf{eff}}(\boldsymbol{x}) = \tfrac{1}{2}\sum_i^n\sum_j^n \left.\frac{\partial^2\chi^2_{\mathsf{eff}}}{\partial x_i \partial x_j}\right|_{\boldsymbol{x}=0} x_i\, x_j + \mathcal{O}(3), \tag{A.14}$$

or, in matrix notation

$$\chi^2_{\mathsf{eff}}(\boldsymbol{x}) = x_i\, C^{-1}_{ij} x_j + \mathcal{O}(3) = \boldsymbol{x}^{\mathrm{T}}\, \boldsymbol{C}^{-1}\, \boldsymbol{x} + \mathcal{O}(3), \tag{A.15}$$

where $\boldsymbol{C}$ is the covariance matrix which basically describes the curvature of χ^2_{eff} around its minimum and its off-diagonal entries describe the correlation between variables. If we neglect the higher order corrections and assume the data to be uncorrelated, the covariance matrix will be diagonal and defining $C_{ii} = \sigma_i^2$ we recover the classical textbook definition of χ^2 [55]:

$$\chi^2_{\mathsf{eff}}(\boldsymbol{d}_i^{\mathrm{th}}) \simeq \sum_i \left(\frac{d_i^{\mathrm{obs}} - d_i^{\mathrm{th}}}{\sigma_i}\right)^2 = \chi^2(\boldsymbol{d}^{\mathrm{th}}), \tag{A.16}$$

corresponding to a multivariate Gaussian likelihood pdf

$$\mathcal{L}(\boldsymbol{d}^{\mathrm{obs}}|\boldsymbol{d}^{\mathrm{th}}) \simeq \exp\left[-\tfrac{1}{2}\, \boldsymbol{x}^{\mathrm{T}}\, \boldsymbol{C}^{-1}\, \boldsymbol{x}\right] = \prod_i^n \frac{1}{\sqrt{2\pi\sigma_i^2}}\, e^{-x_i^2/2\sigma_i^2}. \tag{A.17}$$

For correlated observables that follow individual Gaussian likelihood distributions, we would need a data vector $\boldsymbol{d}^{\mathrm{obs}}$ and a covariance matrix $\boldsymbol{C}$, supplied by the experiment, to evaluate the overall likelihood for a given prediction of our model.

If the errors are non-Gaussian, as is for instance the case with CMB data, a Taylor expansion of χ^2_{eff} will not be sufficient and a more careful treatment is needed, see for instance Refs. [150–152] for a discussion.

3. Galaxy Correlation Function with Oscillating Spectra

Oscillations in the dark matter power spectrum due to acoustic oscillations in the plasma prior to decoupling result in an single peak in the two-point correlation function of the distribution of galaxies $\xi(r)$. In Ref. [20], the authors report the detection of such a peak and identify it as corresponding to the baryonic oscillations of the matter power spectrum.
Since any oscillation of the spectrum, regardless of its origin, will lead to a feature in the correlation function, this data set is particularly well suited to constraining oscillations in the initial power spectrum as well, provided that the features are not completely washed out through subsequent evolution.

The correlation function is related to the matter power spectrum $P(k)$ via a Fourier transform:

$$\xi(r) \propto \int_0^\infty \mathrm{d}k\, k^2 P(k) \frac{\sin kr}{kr}. \tag{A.18}$$

Technically, the upper limit of the integral would be some ultraviolet cutoff k_{UV}, chosen such that the error in ξ is small ($\ll 1\%$). For the scales covered by the SDSS data, i.e., comoving separations between 12 and 175 h^{-1}Mpc, this requires a momentum cutoff $k_{UV} > 1\, h/\mathrm{Mpc}$. At these wavenumbers, however, nonlinear effects cannot be neglected anymore, which makes the theoretical prediction of ξ somewhat tricky.

The standard procedure is outlined in section 4.2 of Ref. [20] and involves corrections for redshift space distortion, nonlinear clustering, scale dependent bias, and a smoothing of features on small scales due to mode coupling. All of these methods were calibrated with nonlinear simulations in a vanilla cosmology setting and it is not obvious that they should be applicable to our case. With the exception of the smoothing, however, the effect of these corrections on the correlation function is smaller than 10% and will only be noticeable at scales $< 40h^{-1}$ Mpc (see Figure 5 of Ref. [20]). So even if we assume a large uncertainty in the nonlinear corrections, the accuracy of the theoretical correlation function will still be of order a few percent, that is smaller than the error bars of the data.

Let us look at the smoothing procedure in a bit more detail. In the usual case, the dewiggled transfer function T_{dw} is a weighted interpolation between the linear transfer function T_{lin} and the Eisenstein-Hu [153] no-wiggle transfer function T_{nw}

$$T_{\mathrm{dw}}(k) = w(k)\, T_{\mathrm{lin}}(k) + (1 - w(k))\; T_{\mathrm{nw}}, \tag{A.19}$$

with a weight function $w(k) = \exp\left[-(ak)^2\right]$ and $a = 7h^{-1}$ Mpc. This is related to the dewiggled spectrum by

$$P_{\mathrm{dw}}(k) = k\, T^2_{\mathrm{dw}}(k)\, \mathcal{P}_{\mathcal{R}}(k). \tag{A.20}$$

In the case of a non-smooth primordial power spectrum $\mathcal{P}_{\mathcal{R}}(k)$, one should of course also dewiggle the initial features. In order to recover the standard procedure for power-law spectra, we will instead smooth the quantity

$$\hat{T}(k) \equiv (P(k)/k)^{1/2} = T(k)\sqrt{\mathcal{P}_{\mathcal{R}}(k)}. \tag{A.21}$$

The use of the no-wiggle transfer function rests on the assumption that at small scales, mode coupling will totally erase all structure, which is reasonable as long as the amplitude of features is of the same order as that of the baryon oscillations. For much larger oscillations, mode coupling might not be efficient enough to erase all structure; it is likely that some residual oscillations will remain. So instead of a no-wiggle $\hat{T}_{\rm nw}$, we will use a smoothed $\hat{T}_{\rm s}$ defined by

$$\hat{T}_{\rm s}(k,q) = \exp\left[\frac{1}{q}\int\limits_{\ln k-q/2}^{\ln k+q/2} {\rm d}\ln k' \,\ln\left[\hat{T}_{\rm lin}(k')\right]\right], \tag{A.22}$$

i.e., a convolution of $\hat{T}_{\rm lin}$ with a top hat function of width q in log-log space. The dewiggled power spectrum is then given by

$$P_{\rm dw}(k,q) = k\left(w(k)\,\hat{T}_{\rm lin}(k) + (1-w(k))\,\hat{T}_{\rm s}(k,q)\right)^2. \tag{A.23}$$

Without turning to N-body simulations it would be hard to estimate how much the spectrum will have to be smoothed, though. Therefore, we will determine the BAO likelihood $\mathcal{L}_{\rm BAO}$ by marginalising over q:

$$\mathcal{L}_{\rm BAO} = \int {\rm d}q\,\mathcal{L}(q)\pi(q). \tag{A.24}$$

We take the prior $\pi(q)$ to be a top hat function between $q = 0$ (i.e., no smoothing at all) and an upper value $q_{\rm max}$, chosen such that it lies in a region where $\mathcal{L}(q)$ is flat in q, corresponding to a complete smoothing. In our numerical code, we approximate the integral (A.24) by averaging the likelihood over N values of q:

$$-\ln\mathcal{L}_{\rm BAO} \simeq {\rm Min}(-\ln\mathcal{L}(q_i)) - \ln\left[\frac{1}{N}\sum_{i=1}^{N}\frac{\mathcal{L}(q_i)}{{\rm Max}(\mathcal{L}(q_i))}\right], \tag{A.25}$$

with $q_i = (i-1)\,q_{\rm max}/(N-1)$.

Bibliography

[1] A. H. Guth, Phys. Rev. **D23**, 347 (1981).

[2] A. H. Guth and S. Y. Pi, Phys. Rev. Lett. **49**, 1110 (1982).

[3] A. A. Starobinsky, Phys. Lett. **B117**, 175 (1982).

[4] L. Covi, J. Hamann, A. Melchiorri, A. Slosar, and I. Sorbera, Phys. Rev. **D74**, 083509 (2006), astro-ph/0606452.

[5] J. Hamann, S. Hannestad, M. S. Sloth, and Y. Y. Y. Wong, Phys. Rev. **D75**, 023522 (2007), astro-ph/0611582.

[6] J. Hamann, L. Covi, A. Melchiorri, and A. Slosar, (2007), astro-ph/0701380.

[7] V. Mukhanov, *Physical Foundations of Cosmology* (Cambridge University Press, 2005).

[8] M. Tegmark, (1995), astro-ph/9511148.

[9] R. Trotta, *Cosmic Microwave Background Anisotropies: Beyond Standard Parameters* (PhD Thesis, 2004), astro-ph/0410115.

[10] C.-P. Ma and E. Bertschinger, Astrophys. J. **455**, 7 (1995), astro-ph/9506072.

[11] W. T. Hu, *Wandering in the background: A Cosmic microwave background explorer* (PhD Thesis, 1995), astro-ph/9508126.

[12] A. R. Liddle and D. H. Lyth, *Cosmological Inflation and Large-Scale Structure* (Cambridge University Press, 2000).

[13] A. Lewis and S. Bridle, Phys. Rev. **D66**, 103511 (2002), astro-ph/0205436.

[14] M. Doran, JCAP **0510**, 011 (2005), astro-ph/0302138.

[15] U. Seljak and M. Zaldarriaga, Astrophys. J. **469**, 437 (1996), astro-ph/9603033.

[16] R. A. C. Croft, D. H. Weinberg, N. Katz, and L. Hernquist, Astrophys. J. **495**, 44 (1998), astro-ph/9708018.

[17] M. McQuinn, O. Zahn, M. Zaldarriaga, L. Hernquist, and S. R. Furlanetto, Astrophys. J. **653**, 815 (2006), astro-ph/0512263.

[18] A. Lewis and A. Challinor, (2007), astro-ph/0702600.

[19] M. Tegmark *et al.*, Phys. Rev. **D74**, 123507 (2006), astro-ph/0608632.

[20] SDSS, D. J. Eisenstein *et al.*, Astrophys. J. **633**, 560 (2005), astro-ph/0501171.

[21] M. Bartelmann and P. Schneider, Phys. Rept. **340**, 291 (2001), astro-ph/9912508.

[22] H. Hoekstra *et al.*, Astrophys. J. **647**, 116 (2006), astro-ph/0511089.

[23] D. N. Spergel *et al.*, (2006), astro-ph/0603449.

[24] G. Hinshaw *et al.*, (2006), astro-ph/0603451.

[25] L. Page *et al.*, (2006), astro-ph/0603450.

[26] N. Jarosik *et al.*, (2006), astro-ph/0603452.

[27] W.-M. Yao *et al.*, Journal of Physics G **33**, 1+ (2006).

[28] M. M. Phillips, Astrophys. J. **413**, L105 (1993).

[29] Supernova Search Team, J. L. Tonry *et al.*, Astrophys. J. **594**, 1 (2003), astro-ph/0305008.

[30] Supernova Cosmology Project, R. A. Knop *et al.*, Astrophys. J. **598**, 102 (2003), astro-ph/0309368.

[31] G. A. Tammann, (2005), astro-ph/0512584.

[32] W. L. Freedman *et al.*, Astrophys. J. **553**, 47 (2001), astro-ph/0012376.

[33] E. W. Kolb and M. S. Turner, *The Early Universe* (Addison Wesley, 1989).

[34] C. Wetterich, Astron. Astrophys. **301**, 321 (1995), hep-th/9408025.

[35] N. J. Cornish, D. N. Spergel, and G. D. Starkman, Class. Quant. Grav. **15**, 2657 (1998), astro-ph/9801212.

[36] J. D. Bekenstein, Phys. Rev. **D70**, 083509 (2004), astro-ph/0403694.

[37] S. M. Carroll *et al.*, Phys. Rev. **D71**, 063513 (2005), astro-ph/0410031.

[38] J. M. Bardeen, Phys. Rev. **D22**, 1882 (1980).

[39] V. F. Mukhanov, H. A. Feldman, and R. H. Brandenberger, Phys. Rept. **215**, 203 (1992).

[40] E. D. Stewart and D. H. Lyth, Phys. Lett. **B302**, 171 (1993), gr-qc/9302019.

[41] J. M. Bardeen, J. R. Bond, N. Kaiser, and A. S. Szalay, Astrophys. J. **304**, 15 (1986).

[42] V. F. Mukhanov, Sov. Phys. JETP **67**, 1297 (1988).

[43] M. Sasaki, Prog. Theor. Phys. **76**, 1036 (1986).

[44] T. S. Bunch and P. C. W. Davies, Proc. Roy. Soc. Lond. **A360**, 117 (1978).

[45] A. R. Liddle, P. Parsons, and J. D. Barrow, Phys. Rev. **D50**, 7222 (1994), astro-ph/9408015.

[46] D. H. Lyth and A. Riotto, Phys. Rept. **314**, 1 (1999), hep-ph/9807278.

[47] A. R. Liddle and D. H. Lyth, Phys. Lett. **B291**, 391 (1992), astro-ph/9208007.

[48] `http://www.rssd.esa.int/index.php?project=Planck`.

[49] M. Zaldarriaga, D. N. Spergel, and U. Seljak, Astrophys. J. **488**, 1 (1997), astro-ph/9702157.

[50] M. Amarie, C. Hirata, and U. Seljak, Phys. Rev. **D72**, 123006 (2005), astro-ph/0508293.

[51] Y.-S. Song and L. Knox, Phys. Rev. **D68**, 043518 (2003), astro-ph/0305411.

[52] L. Hui and W. H. Kinney, Phys. Rev. **D65**, 103507 (2002), astro-ph/0109107.

[53] R. D. Cousins, Am. J. Phys. **63**, 398 (1995).

[54] J. Neyman, Phil. Trans. Roy. Soc. London, Series A **236**, 333 (1937).

[55] M. G. Kendall and A. Stuart, *The Advanced Theory of Statistics* (Charles Griffin & Company Limited, 1961).

[56] R. T. Cox, Am. J. Physics **14**, 1 (1946).

[57] M. Bucher, J. Dunkley, P. G. Ferreira, K. Moodley, and C. Skordis, Phys. Rev. Lett. **93**, 081301 (2004), astro-ph/0401417.

[58] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, J. Chem. Phys. **21**, 1087 (1953).

[59] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, *Markov Chain Monte Carlo in Practice* (Chapman & Hall/CRC, 1996).

[60] N. Christensen, R. Meyer, L. Knox, and B. Luey, Class. Quant. Grav. **18**, 2677 (2001), astro-ph/0103134.

[61] A. Kosowsky, M. Milosavljevic, and R. Jimenez, Phys. Rev. **D66**, 063007 (2002), astro-ph/0206014.

[62] M. Doran and C. M. Müller, JCAP **0409**, 003 (2004), astro-ph/0311311.

[63] A. Lewis, A. Challinor, and A. Lasenby, Astrophys. J. **538**, 473 (2000), astro-ph/9911177.

[64] J. Dunkley, M. Bucher, P. G. Ferreira, K. Moodley, and C. Skordis, Mon. Not. Roy. Astron. Soc. **356**, 925 (2005), astro-ph/0405462.

[65] A. Gelman and D. B. Rubin, Statist. Sci. **7**, 457 (1992).

[66] A. E. Raftery and S. M. Lewis, How many iterations for the Gibbs sampler?, in *Bayesian Statistics*, edited by J. M. Bernado, p. 765, OUP, 1992.

[67] H. Jeffreys, *Theory of Probability* (Oxford University Press, 1961).

[68] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press, 2003).

[69] R. Trotta, (2005), astro-ph/0504022.

[70] A. R. Liddle, P. Mukherjee, and D. Parkinson, Astron. Geophys. **47**, 4.30 (2006), astro-ph/0608184.

[71] D. Parkinson, P. Mukherjee, and A. R. Liddle, Phys. Rev. **D73**, 123523 (2006), astro-ph/0605003.

[72] M. Beltran, J. Garcia-Bellido, J. Lesgourgues, A. R. Liddle, and A. Slosar, Phys. Rev. **D71**, 063532 (2005), astro-ph/0501477.

[73] A. Niarchou and A. Jaffe, (2007), astro-ph/0702436.

[74] P. Mukherjee, D. Parkinson, P. S. Corasaniti, A. R. Liddle, and M. Kunz, Mon. Not. Roy. Astron. Soc. **369**, 1725 (2006), astro-ph/0512484.

[75] C. Pahud, A. R. Liddle, P. Mukherjee, and D. Parkinson, Phys. Rev. **D73**, 123524 (2006), astro-ph/0605004.

[76] C. Pahud, A. R. Liddle, P. Mukherjee, and D. Parkinson, (2007), astro-ph/0701481.

[77] J. Skilling, Nested sampling, in *Bayesian inference and maximum entropy methods in science and engineering, AIP Conference Proceeedings 735*, edited by R. Fischer, R. Preuss, and U. von Toussaint, pp. 395–405, American Institute of Physics, 2004.

[78] P. Mukherjee, D. Parkinson, and A. R. Liddle, Astrophys. J. **638**, L51 (2006), astro-ph/0508461.

[79] A. R. Liddle, (2007), astro-ph/0701113.

[80] H. Akaike, IEEE Transactions on automatic control **19**, 716 (1974).

[81] G. Schwarz, Annals of Statistics **6**, 461 (1978).

[82] W. J. Percival *et al.*, (2006), astro-ph/0608636.

[83] The SNLS, P. Astier *et al.*, Astron. Astrophys. **447**, 31 (2006), astro-ph/0510447.

[84] H. V. Peiris *et al.*, Astrophys. J. Suppl. **148**, 213 (2003), astro-ph/0302225.

[85] U. Seljak, A. Slosar, and P. McDonald, JCAP **0610**, 014 (2006), astro-ph/0604335.

[86] Super-Kamiokande, Y. Fukuda *et al.*, Phys. Rev. Lett. **82**, 2644 (1999), hep-ex/9812014.

[87] Super-Kamiokande, Y. Fukuda *et al.*, Phys. Rev. Lett. **82**, 1810 (1999), hep-ex/9812009.

[88] J. Lesgourgues and S. Pastor, Phys. Rept. **429**, 307 (2006), astro-ph/0603494.

[89] S. Hannestad, Phys. Rev. Lett. **95**, 221301 (2005), astro-ph/0505551.

[90] A. De La Macorra, A. Melchiorri, P. Serra, and R. Bean, (2006), astro-ph/0608351.

[91] C. Zunckel and P. G. Ferreira, (2006), astro-ph/0610597.

[92] H. V. Klapdor-Kleingrothaus, A. Dietz, H. L. Harney, and I. V. Krivosheina, Mod. Phys. Lett. **A16**, 2409 (2001), hep-ph/0201231.

[93] H. V. Klapdor-Kleingrothaus, I. V. Krivosheina, A. Dietz, and O. Chkvorets, Phys. Lett. **B586**, 198 (2004), hep-ph/0404088.

[94] H. V. Klapdor-Kleingrothaus, (2005), hep-ph/0512263.

[95] M. Cortês, A. R. Liddle, and P. Mukherjee, (2007), astro-ph/0702170.

[96] A. D. Linde, Phys. Rev. **D59**, 023503 (1999), hep-ph/9807493.

[97] A. D. Linde, M. Sasaki, and T. Tanaka, Phys. Rev. **D59**, 123522 (1999), astro-ph/9901135.

[98] A. Linde, JCAP **0305**, 002 (2003), astro-ph/0303245.

[99] A. D. Linde, Phys. Lett. **B129**, 177 (1983).

[100] `http://lambda.gsfc.nasa.gov/`.

[101] A. Goobar, S. Hannestad, E. Mörtsell, and H. Tu, JCAP **0606**, 019 (2006), astro-ph/0602155.

[102] The Virgo Consortium, R. E. Smith *et al.*, Mon. Not. Roy. Astron. Soc. **341**, 1311 (2003), astro-ph/0207664.

[103] M. Viel, M. G. Haehnelt, and V. Springel, Mon. Not. Roy. Astron. Soc. **367**, 1655 (2006), astro-ph/0504641.

[104] M. Viel and M. G. Haehnelt, Mon. Not. Roy. Astron. Soc. **365**, 231 (2006), astro-ph/0508177.

[105] M. Viel, M. G. Haehnelt, and A. Lewis, Mon. Not. Roy. Astron. Soc. Lett. **370**, L51 (2006), astro-ph/0604310.

[106] J. R. Ellis, K. A. Olive, Y. Santoso, and V. C. Spanos, Phys. Lett. **B565**, 176 (2003), hep-ph/0303043.

[107] J. R. Ellis, K. A. Olive, Y. Santoso, and V. C. Spanos, Phys. Lett. **B588**, 7 (2004), hep-ph/0312262.

[108] F. D. Steffen, JCAP **0609**, 001 (2006), hep-ph/0605306.

[109] W. H. Kinney, E. W. Kolb, A. Melchiorri, and A. Riotto, Phys. Rev. **D74**, 023502 (2006), astro-ph/0605338.

[110] M. S. Sloth, Nucl. Phys. **B748**, 149 (2006), astro-ph/0604488.

[111] J. Martin and C. Ringeval, JCAP **0608**, 009 (2006), astro-ph/0605367.

[112] S. Dodelson and L. Hui, Phys. Rev. Lett. **91**, 131301 (2003), astro-ph/0305113.

[113] A. R. Liddle and S. M. Leach, Phys. Rev. **D68**, 103503 (2003), astro-ph/0305263.

[114] V. M. Lobashev, Nucl. Phys. **A719**, 153 (2003).

[115] C. Kraus *et al.*, Eur. Phys. J. **C40**, 447 (2005), hep-ex/0412056.

[116] KATRIN, G. Drexlin, Nucl. Phys. Proc. Suppl. **145**, 263 (2005).

[117] S. Groot Nibbelink and B. J. W. van Tent, Class. Quant. Grav. **19**, 613 (2002), hep-ph/0107272.

[118] A. Vilenkin and E. P. S. Shellard, *Cosmic Strings and Other Topological Defects* (Cambridge University Press, 1994).

[119] J. Martin and R. H. Brandenberger, Phys. Rev. **D63**, 123501 (2001), hep-th/0005209.

[120] K. Schalm, G. Shiu, and J. P. van der Schaar, AIP Conf. Proc. **743**, 362 (2005), hep-th/0412288.

[121] J. Martin and R. Brandenberger, Phys. Rev. **D68**, 063513 (2003), hep-th/0305161.

[122] D. H. Lyth and E. D. Stewart, Phys. Rev. **D53**, 1784 (1996), hep-ph/9510204.

[123] H. M. Hodges, G. R. Blumenthal, L. A. Kofman, and J. R. Primack, Nucl. Phys. **B335**, 197 (1990).

[124] D. J. H. Chung, E. W. Kolb, A. Riotto, and I. I. Tkachev, Phys. Rev. **D62**, 043508 (2000), hep-ph/9910437.

[125] O. Elgaroy, S. Hannestad, and T. Haugboelle, JCAP **0309**, 008 (2003), astro-ph/0306229.

[126] J. A. Adams, G. G. Ross, and S. Sarkar, Nucl. Phys. **B503**, 405 (1997), hep-ph/9704286.

[127] J. Lesgourgues, Nucl. Phys. **B582**, 593 (2000), hep-ph/9911447.

[128] K. Becker, M. Becker, and A. Krause, Nucl. Phys. **B715**, 349 (2005), hep-th/0501130.

[129] A. Ashoorioon and A. Krause, (2006), hep-th/0607001.

[130] J. M. Cline, P. Crotty, and J. Lesgourgues, JCAP **0309**, 010 (2003), astro-ph/0304558.

[131] G. Efstathiou, Mon. Not. Roy. Astron. Soc. **343**, L95 (2003), astro-ph/0303127.

[132] J. P. Luminet, J. Weeks, A. Riazuelo, R. Lehoucq, and J. P. Uzan, Nature. **425**, 593 (2003), astro-ph/0310253.

[133] C. Gordon and W. Hu, Phys. Rev. **D70**, 083003 (2004), astro-ph/0406496.

[134] A. Shafieloo and T. Souradeep, Phys. Rev. **D70**, 043523 (2004), astro-ph/0312174.

[135] A. Shafieloo, T. Souradeep, P. Manimaran, P. K. Panigrahi, and R. Rangarajan, (2006), astro-ph/0611352.

[136] J. A. Adams, B. Cresswell, and R. Easther, Phys. Rev. **D64**, 123514 (2001), astro-ph/0102236.

[137] P. Hunt and S. Sarkar, Phys. Rev. **D70**, 103518 (2004), astro-ph/0408138.

[138] S. Chongchitnan and G. Efstathiou, JCAP **0701**, 011 (2007), astro-ph/0611818.

[139] C. P. Burgess, J. M. Cline, F. Lemieux, and R. Holman, JHEP **02**, 048 (2003), hep-th/0210233.

[140] M. Kawasaki, F. Takahashi, and T. Takahashi, Phys. Lett. **B605**, 223 (2005), astro-ph/0407631.

[141] ACBAR, C.-L. Kuo *et al.*, Astrophys. J. **600**, 32 (2004), astro-ph/0212289.

[142] C. J. MacTavish *et al.*, Astrophys. J. **647**, 799 (2006), astro-ph/0507503.

[143] A. C. S. Readhead *et al.*, Astrophys. J. **609**, 498 (2004), astro-ph/0402359.

[144] SDSS, M. Tegmark *et al.*, Astrophys. J. **606**, 702 (2004), astro-ph/0310725.

[145] The 2dFGRS, S. Cole *et al.*, Mon. Not. Roy. Astron. Soc. **362**, 505 (2005), astro-ph/0501174.

[146] `http://universe.nasa.gov/program/probes/inflation.html`.

[147] X. Chen, R. Easther, and E. A. Lim, (2006), astro-ph/0611645.

[148] D. Battefeld and T. Battefeld, (2007), hep-th/0703012.

[149] K. Semendyayev and I. Bronshtein, *Handbook of Mathematics* (Springer, 2004).

[150] J. R. Bond, A. H. Jaffe, and L. Knox, Phys. Rev. **D57**, 2117 (1998), astro-ph/9708203.

[151] J. R. Bond, A. H. Jaffe, and L. E. Knox, Astrophys. J. **533**, 19 (2000), astro-ph/9808264.

[152] A. H. Jaffe, J. R. Bond, P. G. Ferreira, and L. E. Knox, (2003), astro-ph/0306506.

[153] D. J. Eisenstein and W. Hu, Astrophys. J. **496**, 605 (1998), astro-ph/9709112.

Acknowledgments

Time to say thank you!
I should like to extend my gratitude to

- Laura Covi, for her advice, and getting me interested in this exciting topic,
- Prof. Wilfried Buchmüller, for his co-supervision,
- Toffi Lüdeling and Jochen Zahn, for sharing their wisdom with me,
- Markus Ahlers and Kai Schmidt-Hoberg, for careful proof-reading of the manuscript and helpful suggestions,
- Dirk Brömmel and Jörgen Samson, for their substantial contributions towards curing my computer illiteracy,
- all the folks of the DESY theory group, it's been a fun time,
- my collaborators, Steen Hannestad, Alessandro Melchiorri, Anže Slosar, Martin Sloth and Irene Sorbera for fruitful discussions and, well, collaborating with me,
- my parents, for always making sure I don't go hungry,
- and Yvonne ♡

Cheers!