DISSERTATION
submitted to the
Combined Faculties of the Natural Sciences and Mathematics
of the Ruperto-Carola-University of Heidelberg, Germany
for the degree of
DOCTOR OF NATURAL SCIENCES

Put forward by
Elena Sellentin
born in: Karlsruhe, Germany
Oral examination: 17th May 2016

Statistical Inference in Cosmology

Referees:
Prof. Dr. Luca Amendola
Prof. Dr. Matthias Bartelmann

STATISTICAL INFERENCE IN COSMOLOGY

Analysis of cosmic data is the only way to determine whether General Relativity is the law of gravity also on the largest scales in our Universe. The current standard model of cosmology, ΛCDM, is based on General Relativity, and fits all currently available data flawlessly. However, theoretical dissatisfaction with ΛCDM exists: cosmological data probe gravitational interactions, and ΛCDM fits the data only because it introduces two components of startling gravitional behaviour, the cosmological constant, Λ, and cold dark matter (CDM). The cosmological constant has a suspiciously small value when regarded from the perspective of quantum field theories, and cold dark matter has so far not been detected in any experiment of particle physics.

This thesis examines the cosmological standard model from the vantage point of statistics. A non-Gaussian likelihood approximation is presented and the need of an unbiased mechanism for dealing with estimated covariance matrices is addressed. Concerning neutrinos, a previously existent parameterization bias in the analysis of the cosmic microwave background is resolved. Using weak lensing and type Ia supernova data of the next generation, it is estimated how much can be learned about dark energy from these future data sets.

INFERENZSTATISTIK IN DER KOSMOLOGIE

Kosmologische Daten sind die einzigen aus welchen abgeleitet werden kann, ob die Allgemeine Relativitätstheorie auch auf den größten Skalen unseres Universums die Gravitation richtig beschreibt. Das momentane Standardmodell der Kosmologie, ΛCDM, baut auf der Allgemeinen Relativitätstheorie auf und fittet bisherige kosmische Daten problemlos. Es besteht jedoch eine theoretisch motivierte Skepsis bezüglich ΛCDM: Kosmische Daten untersuchen Auswirkungen der Gravitation, und ΛCDM fittet die Daten nur, weil es mit der kosmologischen Konstanten (Λ) und der kalten Dunklen Materie (CDM) zwei neue Substanzen mit erstaunlichen gravitativen Eigenschaften einführt. Aus quantenfeldtheoretischer Sicht hat die kosmologische Konstante jedoch einen verdächtig kleinen Wert, und Dunkle Materie ließ sich bisher in Teilchenexperimenten nicht nachweisen.

Diese Dissertation untersucht das kosmologische Standardmodell gemäß statistischer Gesichtspunkte. Eine nichtgaussche Näherung für Wahrscheinlichkeitsverteilungen wird vorgestellt, sowie eine erwartungstreue Inferenzmethode für den Fall geschätzter Kovarianzmatrizen. Das Vorkommen gewöhnlicher Neutrinos wird durch eine vorurteilsfreie Analyse bestätigt. Anhand künstlicher Weak-Lensing- und Supernova-Ia-Datensätze künftiger Beobachtungskampagnen wird ermittelt, wie präzise diese Datensätze Eigenschaften der Dunklen Energie einschränken können.

PREFACE

Why is the topic of statistical inference in cosmology so pressing an issue, that I dedicated my entire thesis to it? To understand this, let us begin by taking stock. During this thesis, Albert Einstein's description of gravity, General Relativity, turned 100 years old. Einstein published his theory on the 25th of November 1915, and apart from being able to explain the precession of Mercury's perihelion, the first experimental support for General Relativity came only four years later in 1919, when an expedition led by Sir Arthur Eddington measured the deflection of light by the Sun during a solar eclipse. Since then, further experiments have shown Einstein's theory of gravity to improve upon its Newtonian predecessor. The redshift of light was measured in the Earth's gravitational field, the Shapiro time-delay was measured in the Solar System with the Cassini spacecraft, and the orbit of the binary pulsar PSR 1913+16 was found to decay at the rate predicted by General Relativity. Further support for General Relativity has just recently been published on the 11th February 2016, when the LIGO collaboration announced the detection of the gravitational waves pattern of two merging black holes with a significance greater than $5.1\,\sigma$.

All of these experiments are perfectly in agreement with General Relativity. None requires the introduction of new, strange parameters. However, all of these experiments assess also rather small scales – from the size of the Earth to approximately the extent of the Solar System. Likewise, these experiments also investigate comparatively high-density environments because pulsars and black holes are certainly amongst the most dense objects of the Universe, but also the Earth with its density of $5.5\ \mathrm{g\,cm^{-3}}$ is many orders of magnitude more dense than the cosmic mean of a mere $10^{-29}\ \mathrm{g\,cm^{-3}}$.

Correspondingly, none of the aforementioned environments in which General Relativity was found to be valid is representative of our Universe as a whole. Nonetheless, if we combine General Relativity with the observationally well-established assumption of isotropy, and the assumption of homogeneity, which is advisable for reasons of modesty, then a rather successful description of the Universe can be derived. This has led to the current standard model of cosmology, ΛCDM, whose name is composed of its two main ingredients, the cosmological constant Λ, and cold dark matter (CDM). Baryons seem to contribute only a minor percentage to the Universe's density and are not even mentioned in the cosmic standard model's name.

So what is ΛCDM? It is the attempt to reproduce the success of General Relativity also on cosmic scales, and it includes the sobering insight that this is only possible if we introduce two new quantities: dark matter and dark energy. If we now interpret cosmology as a test of General Relativity and rank it amongst the other tests already mentioned, why is it then, that none but cosmology demands new parameters to fit Einstein's law of gravity?

This problem is not solved yet, and while theoretical cosmologists currently investigate a multitude of competitors for ΛCDM, it is clear that the final answer has to be filtered out from accurate data, when subjecting them to a highly reliable and sensitive statistical inference. This thesis is therefore closest in spirit to several other works in the literature which dedicate themselves to a statistical analysis of cosmic data.

The reader may be wondering who hides behind the curious ‘we’ in this thesis. In many grammatical constructs, ‘we’ refers to ‘the reader and me, the author’. In the vast majority of cases it refers however to me and my colleagues or coauthors who have all made important intellectual contributions to the work presented here. Let me therefore present the individual persons behind this ‘we’ in more detail and express my gratitude to them.

I take pride in having completed my doctorate under a scientifically and historically literate supervisor, Luca Amendola. I am glad he has cast his knowledge into black and white by publishing his book on dark energy, which runs in our cosmology group under the nickname ‘The Bible’. His book can however not replace Luca when it comes to generous support and spurring ambitions. Luca’s scientific vision has provided the fruitful silver thread for this thesis and for all this I am very grateful.

Furthermore, numerous coauthors have accompanied me during the last three years, broadening my scientific and cultural horizon alike. Ruth Durrer has added multiple razor-sharp arguments to my education. Alan Heavens has shared his vast experience on sophisticated statistical methods. The unfamiliar creativity of the untamable Brazilian chaos was impressively demonstrated by Miguel Quartin. Björn Malte Schäfer’s natural curiosity has triggered many deeper thoughts.

Four unidentified referees have provided welcome opinions on our papers. Three identified referees are welcome to further bestow on me their humorous remarks without the comforting shelter of anonymity. Further stimulating discussions with Martin Kilbinger, Tom Kitching, Martin Kunz, Ignacy Sawicki, Dominik Schwarz, Valeria Pettorino and Christof Wetterich are appreciated.

What must have been a never-ceasing and tiring avalanche of bureaucracy was kept hidden from me due to the assiduous work of Eduard Thommes and our secretaries. Including also Elmar Bittner and his reliable management of the ITP computer network, I thank for all this administrative support.

The Heidelberg Graduate School of Fundamental Physics (HGSFP) has given me the opportunity to attend specialized lectures, and provided valuable funds. Further, this thesis has received financial support through the RTG ‘Particle Physics beyond the Standard Model’ (DFG fund 1904), and the transregional collaborative research centre TR 33 ‘The Dark Universe’ of the German Science Foundation.

Finally, I am very pleased that my defence will be amended by the second referee, Matthias Bartelmann, and the third and fourth examiner, Volker Springel and Hans-Christian Schultz-Coulon. The defence stands in the shadow of a laborious man, whose recent cancer diagnosis is an unfair return for his long-term commitment to the department. I wish him all my best, hoping that modern medicine will grant him a full recovery and many more years of untroubled curiosity.

Contents

Table 1: Conventions used in this thesis

t	cosmic time
η	conformal time
$\dot{a} = \frac{\mathrm{d}a}{\mathrm{d}t}$	dot: derivative with respect to cosmic time, if not explicitely stated otherwise, e.g. a coordinate time or proper time.
$a' = \frac{\mathrm{d}a}{\mathrm{d}\eta}$	prime: derivative with respect to conformal time
Φ, Ψ	sign convention of the Bardeen potentials: following Amendola & Tsujikawa (2010a). (Note that Durrer (2008) uses $\Phi \to -\Phi$.)
Δ	density contrast
∇^2	Laplacian
δ	a perturbation or Dirac's delta function

Furthermore,

- we assume the Universe to be spatially flat throughout the entire thesis.
- we use the sign convention $(-,+++)$ for the metric.
- The terms 'accuracy' and 'precision', as well as their respective adjectives 'accurate' and 'precise', are often used synonymically within this thesis – but care has been taken that this never leads to problematic statements in the statistical sense. The reason is that accuracy and precision are colloquial synonyms already within a single language. With physics being an international environment, the additional complication of non-unique translations arises. Dictionaries of German, French, Italian and English do not uniquely distinguish between precision and accuracy and even add to the list of possible translations further closely related nouns, such as the French 'exactitude' and 'justesse' or the Italian 'esattezza'.

  Within statistical literature, precision and accuracy are however often contrasted, by defining precision as the remaining statistical scatter, and accuracy as the lack of biases or systematic errors. Obviously, these definitions imply that all statements about precision and accuracy, in the statistical sense, can always be rephrased in terms of statistical scatter and biases. There is no such ambiguity about scatter and biases, as there is about the potentially colloquially meant precision or accuracy. Within this thesis, we will therefore rather speak of biases and statistical scatter whenever the distinction between precision and accuracy, in the statistical sense, is of importance.

Chapter 1

The standard model of cosmology

The current standard model of cosmology, ΛCDM, is derived from its underlying theory, General Relativity. As this dissertation was written in a dark energy group, I will indicate with footnotes where in the derivation of ΛCDM the standard assumptions are exchanged against alternative assumptions such that well-known explanations for dark energy arise.

1.1 General Relativity

General Relativity gives up the Newtonian concept of a vectorial gravitational force, and instead describes the effects of gravity by a dynamical spacetime through the metric tensor $g^{\mu\nu}$. It follows the construction principle that gravity can locally by transformed away by transiting into a freely falling restframe (the strong equivalence principle[1]).

Such freely falling observers follow geodesics γ through spacetime, defined by (Straumann, 2013)

$$\nabla_{\dot{\gamma}}\dot{\gamma} = 0, \tag{1.1}$$

meaning a geodesic is locally autoparallel to itself, and where ∇ is the Riemannian connection that is sensitive to spacetime curvature due to

$$\nabla_{\partial_i}(\partial_j) \equiv \Gamma^k_{ij}\partial_k, \tag{1.2}$$

i.e. it describes the change of the tangent space's basis vectors ∂_i for infinitesimal displacements. The functions Γ^k_{ij} are known as Christoffel symbols.

A freely falling observer will find that their metric is locally Minkowskian, $g^{\mu\nu} = \eta^{\mu\nu} = \mathrm{diag}(-1, 1, 1, 1)$, and that derivatives of the metric are locally also zero $g^{\mu\nu}{}_{,\lambda} = 0$ due to the free fall.

Globally, however, gravity cannot be transformed away, leading to relative accelerations of different freely falling observers. For example, two bodies that fall freely next to each other above the Earth, will also approach each other and hence experience a relative acceleration. These relative accelerations are described by the *curvature* of spacetime, which depends on

---

[1] This strong equivalence principle is given up in models of coupled dark energy - dark matter.

the second derivatives of the metric. In local coordinates and for a Riemannian connection ∇, the Riemannian curvature tensor is

$$R^i_{jkl} = \partial_k \Gamma^i_{lj} - \partial_l \Gamma^i_{kj} + \Gamma^m_{lj}\Gamma^i_{km} - \Gamma^m_{kj}\Gamma^i_{lm}, \tag{1.3}$$

where the Christoffel symbols are

$$\Gamma^\mu_{\nu\lambda} = \frac{1}{2} g^{\mu\alpha}(g_{\alpha\nu,\lambda} + g_{\alpha\lambda,\nu} - g_{\nu\lambda,\alpha}). \tag{1.4}$$

In the free fall, the first order derivatives of the metric will be zero, such that the Christoffel symbols then vanish. Often, this is rephrased as 'the possibility to transform gravity locally away'.

The contraction of the Riemann tensor $R_{jl} = R^i_{jil}$ with the metric is the Ricci tensor, which can further be contracted down to the scalar curvature $R = R^i_i$. Being a scalar, R is an invariant of the space time curvature under transformations, and therefore is the natural invariant from which an action principle for GR can be constructed. Indeed, standard GR in the vacuum follows from the Einstein-Hilbert action[2]

$$S_{EH} = \frac{1}{16\pi G} \int \sqrt{-g}\,(R - 2\Lambda)\,\mathrm{d}^4x, \tag{1.5}$$

where $\sqrt{-g}$ is the square root of the metric determinant, which is needed to have an invariant measure. The constant Λ is known as the cosmological constant, and is allowed due to the Lovelock-Theorem[3]. The variation δS_{EH} with respect to the metric tensor leads to the field equations of gravity in the vacuum, $G_{\mu\nu} + \Lambda g_{\mu\nu} = 0$, meaning it describes the spacetime part of gravity. The Einstein tensor is given by

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R, \tag{1.6}$$

and satisfies $\nabla_\nu G^{\mu\nu} = 0$ due to the second contracted Bianchi identity (Straumann, 2013). In order to include present forms of matter and energy, their Lagrangian $\mathcal{L}_M$ needs to be added[4] to the action, leading to (Amendola & Tsujikawa, 2010a; Straumann, 2013)

$$S_{GR} = \int \sqrt{-g}\left(\frac{1}{16\pi G}(R - 2\Lambda) + \mathcal{L}_M\right)\mathrm{d}^4x \tag{1.7}$$

which is the action of standard General Relativity. The variation δS_{GR} now leads to Einstein's field equations with a source term

$$G_{\mu\nu} = 8\pi G T_{\mu\nu} - \Lambda g_{\mu\nu}, \tag{1.8}$$

---

[2] More complicated functions of R could also be used, leading to $f(R)$-gravity.

[3] We will revisit this statement in Sect. 2.4 when disentangling a classical and a quantum contribution to the cosmological constant.

[4] So called *non-minimally* coupled explanations of dark energy go beyond a simple addition.

where the energy-momentum tensor[5][6] $T_{\mu\nu}$ stems from the variation of the matter Lagrangian $\mathcal{L}_M$ with respect to the metric. In the reference frame of an observer moving with four-velocity u_μ the energy-momentum tensor reads (Amendola & Tsujikawa, 2010a)

$$T_{\mu\nu} = (\rho + p)u_\mu u_\nu + p g_{\mu\nu} + \Sigma_{\mu\nu}, \tag{1.9}$$

where ρ and p are the energy density and the pressure of the cosmic matter fields and $\Sigma_{\mu\nu}$ are viscous stresses. Densities, pressures and stresses are not Lorentz-invariant hence the observed energy-momentum tensor has to be projected into the reference frame of the observer as specified by his eigenvelocity u^μ.

The appearance of the pressure in Einstein's field equations is of special interest: while we are used from Newtonian gravity to the fact that pressure can counteract gravity, we now see that in a fully relativistic theory of gravity, the pressure itself is a source of gravity. This is because pressure contributes to the energy density and consequently gravitates.

Taking care of symmetries in the Einstein and energy-momentum tensors, Einstein's field equations are ten non-linear and coupled partial differential equations of second order in the metric. This means the esteemed superposition principle from Newtonian gravity does not carry over to GR[7], and the field equations cannot be solved by recursion to a Green's function. Since they are of second order in the metric, their solution furthermore needs the specification of two initial or boundary conditions.

Analytical solutions to the full field equations are however possible in case of highly symmetric spacetimes, and for simple forms of the energy-momentum tensor. This leads to analytical solutions such as the Schwarzschild- or Kerr metric for black holes, and according to observational evidence, also the spacetime of our Universe seems to possess enough symmetries, such that an analytical solution of cosmological relevance can be derived from GR. This will be the subject of the next section.

1.2 FLRW cosmology

On the largest scales, our Universe provides such symmetries: observations of the cosmic microwave background, galaxy catalogues or other sky maps show that the Universe is isotropic around our position. If we further believe that the Earth's position in the Universe is in no way special, then this implies that the Universe is also isotropic around any other point in it. This in turn implies spatial homogeneity.

---

[5] Explanations for dark energy that add hypothetical constituents to $T^{\mu\nu}$ are e.g. quintessense and k-essence.

[6] Gravity only probes the *total* energy-momentum tensor of all matter components in the Universe. This makes a distinction between dark energy and dark matter – via their gravitational effects only – impossible. See Sect. 2.3 for this dark degeneracy.

[7] It has been mused that the usually adopted split in a background metric and perturbations around it and solving these two contributions independently of each other, may give rise to late-time cosmic acceleration. This hypothesis is called backreaction of cosmic structures on the expansion, but it seems that the effect is too small.

Time and space are not invariantly separable in General Relativity. If an observer sees spatial homogeneity, this means his state of motion and his time-coordinate are fixed. This is the origin of the cosmic time: the cosmic time is that time, that an observer on spatially homogenous hypersurfaces sees. Under these observationally motivated symmetry assumptions[8], the curvature of the spatial hypersurfaces must be constant. Within this thesis, it is furthermore assumed that the curvature is constant because of observationally supported spatial flatness. The metric then takes the shape

$$g_{\mu\nu} = -dt^2 + a^2(t)[\mathrm{d}r^2 + r^2(\mathrm{d}\theta^2 + \sin^2\theta \mathrm{d}\phi^2)] \tag{1.10}$$

which is called the spatially flat Friedmann-Lemaitre-Robertson-Walker metric (FLRW-metric). Here a is the scale factor that only depends on time and we normalize $a = 1$ at the present epoch. The angles θ and ϕ are the two position angles on the celestial sphere. The cosmic time is t and refers to the proper time of an observer at rest with respect to the coordinates (r, θ, ϕ). Often it is convenient to switch to the so-called *conformal* time, given by

$$\eta = \int \frac{1}{a}\mathrm{d}t. \tag{1.11}$$

In concordance with isotropy and homogeneity, the energy momentum tensor as seen by an observer comoving with the cosmic flow cannot include any large-scale currents, T^{0i} and stresses $T^{ij}\ \forall\ i \neq j$. For such a comoving observer, the energy-momentum tensor is then

$$T_\mu^\nu = \mathrm{diag}(-\rho(t), p(t), p(t), p(t)). \tag{1.12}$$

Here, it is important to note that ρ and p refer to the total energy density and pressure of all matter constituents because all of them couple to gravity.
Einsteins field equations for the FLRW-metric then reduce to their $(0,0)$- and (i,i)-component, which were renamed the Friedmann-equations (Straumann, 2013)

$$\begin{aligned} (0,0) &\to H^2 = \frac{8\pi G}{3}\rho + \frac{\Lambda}{3} \\ (i,i) &\to \frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3p) + \frac{\Lambda}{3}, \end{aligned} \tag{1.13}$$

where $H = \dot{a}/a$ is called the Hubble function. The conformal Hubble function is

$$\mathcal{H} = \frac{\mathrm{d}a}{\mathrm{d}\eta}\frac{1}{a}. \tag{1.14}$$

It is convenient to split the total density ρ into different subspecies,

$$\rho(t) = \sum_i \rho_i(t), \tag{1.15}$$

---

[8] Looking out into the cosmos, we automatically look back in time. Consequently, we actually cannot distinguish between a timelike evolution of the cosmos and a spatially isotropic variation. The Lemaitre-Tolman-Bondi (LTB) metric restricts itself to isotropy but allows for a radial inhomogeneity. By allowing for an underdense region around the Milky Way, it can in principle provide an explanation for dark energy. However, LTB is in tension with the latest data (Amendola & Tsujikawa, 2010a).

where for each subspecies, we define an equation of state parameter w_i by the ratio of its pressure to the density

$$w_i = \frac{p_i}{\rho_i}. \tag{1.16}$$

This is convenient because although each particle species will contribute an energy density ρ_i, not all species will have a pressure. For non-relativistic particles, we have $p \ll \rho$, for relativistic particles like photons we have $w = 1/3$. Particles of small mass may be relativistic in the early Universe, and then become non-relativistic as they cool with the expansion. Their equation of state parameter then changes with time. The cosmological constant Λ has $w = -1$ but also other candidates for dark energy can achieve $w = -1$.
The energy densities ρ_i will decrease with expansion in a way that depends on their equation of state parameter. The Friedmann equations contain the automatically satisfied conservation equation $\nabla_\nu T^{0\nu} = 0$, which produces (Amendola & Tsujikawa, 2010a)

$$\dot{\rho} + 3H(\rho + p) = 0. \tag{1.17}$$

Using the equation of state parameter, the conservation equation (1.17) can be recast into the form

$$\frac{\mathrm{d}}{\mathrm{d}t}(a^3\rho) + w\rho\frac{\mathrm{d}}{\mathrm{d}t}(a^3) = 0, \tag{1.18}$$

which describes adiabatic cooling as the volume a^3 increases with expansion. If integrated, this leads to

$$\rho_i = \rho_0^i a^{-3(1+w_i)}, \tag{1.19}$$

where ρ_0^i is conventially used to be the density of a certain species today, at $a = 1$.
From (1.19), we infer that non-relativistic matter with $w = 0$ dilutes as $\rho(a) = \rho_0 a^{-3}$ and relativistic matter dilutes faster as not only its density drops, but also its pressure decreases, leading to $\rho(t) = \rho_0 a^{-4}$. The energy density contributed by a matter species with $w = -1$ such as the cosmological constant stays constant.
If we insert (1.19) into the Friedmann equations, and demand that the expansion of the Universe shall be dominated by a single species with w_i only, then the scale factor evolves as

$$a \propto t^{2/3(1+w_i)}, \tag{1.20}$$

if w_i is constant and $w_i \neq -1$. For $w_i = -1$, the solution of the Friedmann equations is

$$a \propto \exp(Ht). \tag{1.21}$$

Returning with (1.19) for multiple species to the Friedmann equations (1.13), the Hubble function can be written as

$$H^2(a) = H_0^2 E^2(a), \quad E^2(a) = \sum_i \Omega_{0;i} a^{-3(1+w_i)} \tag{1.22}$$

where the subscript zero denotes again quantities measured at $a = 1$. Consequently, H_0 is the Hubble constant. The density parameters $\Omega_{0;i}$ are defined as

$$\Omega_{0,\Lambda} = \frac{\Lambda}{3H_0^2} \tag{1.23}$$

for the cosmological constant, and for usual matter constituents like neutrinos, baryons or photons,

$$\Omega_{0,i} = \frac{\rho_{0,i}}{\rho_{0,crit}} \tag{1.24}$$

where the critical density is

$$\rho_{0,crit} = \frac{3H_0^2}{8\pi G} \tag{1.25}$$

which is the density needed today for the expansion of the Universe to come to a halt in the future.

Dark energy research started with the observation that the Universe's expansion has been accelerating in the approximately past 5 billion years. Such an acceleration means $\ddot{a} > 0$. From the Friedmann equations we obtain (Amendola & Tsujikawa, 2010a)

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}\rho(1+3w), \tag{1.26}$$

and consequently an accelerated expansion in the late Universe is only possible if there exists a dominant energy form with equation of state parameter $w < -1/3$.

1.3 Redshift and distance measures

The basis of all distance measures in cosmology is the redshift z. It appears because as a photon travels through the expanding spacetime, its wavelength is stretched by the scale factor. So if a photon is emitted at time t_e, to which a scale factor $a(t_e)$ corresponds, then

$$\frac{\lambda_o}{\lambda_e} = \frac{a(t_o)}{a(t_e)} \tag{1.27}$$

holds, where λ_o is the wavelength that the photon has at the epoch where the scale factor is $a(t_o)$. As the expansion decreases the energy of the photon, the wavelengths will increase, which – when applied to visible light – corresponds to a reddening. One therefore calls the wavelength shift the redshift z which is defined by

$$z = \frac{\lambda_o - \lambda_e}{\lambda_e}, \tag{1.28}$$

where λ_o is the wavelength that a photon has when it arrives at the observer. The redshift is therefore a measurable quantity if a spectrum of a cosmological source is taken and if the emission wavelength is known.

If one uses a cosmological model and specifies all parameters which the Hubble function depends upon, then one can convert such measured redshifts into distances. Such derived distances are therefore model-dependent and parameter-dependent quantities and one sometimes finds in the literature the formulation that a certain observation 'probes the distance-redshift relation'.

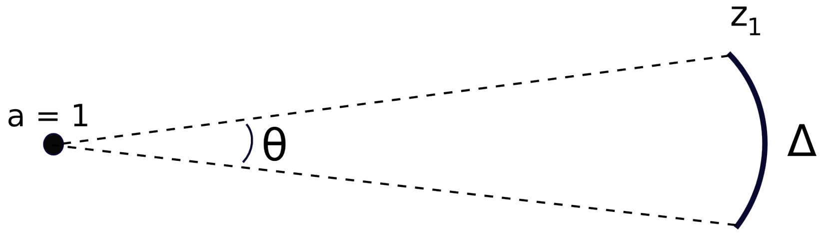

Figure 1.1: Derivation of the angular diameter distance, following Durrer (2008). The observer is at $a = 1$, where $z = 0$. The object lies in a hypersurface of constant time, as seen on the sky at a certain redshift z_1. Photons leaving the ends of the object simultaneously, travel along the dashed lines to us and the angular diameter distance is then $D_{\mathrm{ang}} = \Delta/\theta$.

Additionally, the concept of *a* unique distance does not exist in a curved spacetime. In a Euclidean, non-relativistic space, the unique distance measures the length of the shortest connection between two points at equal time. If a distance shall be measured in a spacetime with noticable curvature, this simultaneity gets lost and it is necessary to describe in more detail how a certain distance measure was derived. Of relevance for this thesis are the following distance measures.

The comoving distance between two redshifts integrates up the distance travelled by a photon. If we write the FLRW metric as $\mathrm{d}s^2 = -c^2\mathrm{d}t^2 + a^2(t)\mathrm{d}r^2$, then light travels on the lightcone $\mathrm{d}s^2 = 0$ of which the relevant root for distance measures is $\mathrm{d}r = -c\mathrm{d}t/a(t)$. The comoving distance is then the integral of $\mathrm{d}\chi$

$$\begin{aligned} D_{\mathrm{com}}(z_1, z_2) = \int_{\chi_1}^{\chi_2} \mathrm{d}r &= -\int_{t_0}^{t_1} \frac{c}{a(t)}\mathrm{d}t \\ &= \frac{c}{H_0}\int_{a(z_1)}^{a(z_2)} \frac{\mathrm{d}a}{a^2 E(a)} \\ &= c\int_{z_1}^{z_2} \frac{\mathrm{d}z'}{H(z')}. \end{aligned} \tag{1.29}$$

The angular diameter distance is defined such that if an object of size Δ subtends an angle θ on the sky, see Fig. 1.1, then we can define a distance D_{ang} such that in the Euclidean relation

$$D_{\mathrm{ang}} = \frac{\Delta}{\theta}, \tag{1.30}$$

holds. To derive it, we imagine the object has the physical arc length Δ and lies in a hypersurface of constant time as seen in Figure 1.1. Then we can measure the angle it subtends from two light signals that arrive at us today, and that left the both ends of the object at the same time. These light signals will have travelled along the comoving distance D_{com}. The physical arc length is then $\Delta = a(z)D_{\mathrm{com}}\theta$. Consequently, (1.30) shows that $D_{\mathrm{ang}} = a(z)D_{\mathrm{com}}$ is the angular diameter distance to us (Durrer, 2008). If it is measured between two hypersurfaces of constant z_1 and z_2 where neither has $a(z) = 1$, then the angular diameter distance between two redshifts is given by

$$D_{\mathrm{ang}}(z_1, z_2) = \frac{a(z_2)}{a(z_1)} D_{\mathrm{com}}(z_1, z_2). \tag{1.31}$$

This angular diameter distance is e.g. the distance of importance in gravitational lens theory, and for measuring the size of Baryonic Acoustic Oscillations (BAO) or the sound horizon on the CMB surface.

The luminosity distance is defined such that it reproduces the decrease of the flux of a source as in a Euclidean geometry. If a source emits per second the energy E, then it has the absolute luminosity $L_s = E/\mathrm{s}$. The flux of such a source is then the energy received by a detector per second and surface area. If the source radiates its photons isotropically, then the flux will be constant on spherical shells centered on the source. The luminosity distance D_{lum} is then defined via the flux F as measured on these shells

$$F = \frac{L_s}{4\pi D_{\mathrm{lum}}^2}, \tag{1.32}$$

where $4\pi D_{\mathrm{lum}}^2$ is the surface of the sphere over which the source's photons are distributed. The luminosity distance is in all spacetimes related to the angular diameter distance via the Etherington relation (Bartelmann & Schneider, 2001)

$$D_{\mathrm{lum}}(z_1, z_2) = \left(\frac{a(z_1)}{a(z_2)}\right)^2 D_{\mathrm{ang}}(z_1, z_2). \tag{1.33}$$

This comes about as follows: the source at redshift z_1 has a proper time interval $\mathrm{d}\tau_1 = a(z_1)\mathrm{d}t$. An observer at redshift z_0 has however a proper time interval $\mathrm{d}\tau_0 = a(z_0)\mathrm{d}t$. The observer will then measure a flux with respect to another proper time interval than that of the source. The total energy emitted by the source per its proper time interval is $L_s\mathrm{d}\tau$. As the photons travel along the comoving distance towards us, they are additionally redshifted by $a(z_1)/a(z_0)$. The flux is then (Durrer, 2008)

$$F = \frac{La^2(z_1)}{4\pi a^4(z_0) D_{\mathrm{com}}^2}, \tag{1.34}$$

from which we identify $D_{\mathrm{lum}} = a^2(z_0)/a^2(z_1) D_{\mathrm{com}}$ and therefore (1.33) holds.

Over small redshift differences, $\Delta z \ll 1$, the curvature of spacetime is not important, and in this limit all distance measures reduce in first order to a unique distance

$$d = \frac{cz}{H_0} + \mathcal{O}(z^2), \tag{1.35}$$

where $cz = v$ is the recession velocity of nearby galaxies whose detection demonstrated that the Universe is not static.

As all of these distances depend on H_0^{-1}, and as the Hubble constant is a measured quantity with a non-zero error, it became conventional to write $H_0 = 100h$ km/s/Mpc and to keep the dependence on h explicit when speaking of length scales. One therefore often finds expressions like $8h^{-1}$Mpc.

An important concept of relativistic gravity is that in contrast to Newtonian gravity, there exists a horizon beyond which a causal connection is impossible. As light travels along the comoving distance, this horizon r_H is (Longair, 2008)

$$r_H = \int_{t_i}^{t_0} = \frac{c\mathrm{d}t}{a(t)}, \tag{1.36}$$

where the lower integration bound t_i is the cosmic time when a photon was emitted, and the upper integration bound t_0 is the current epoch. For $t_i = 0$, r_H is the horizon as calculated from the Big Bang onwards.

The formation of cosmic structures is not yet contained in the FLRW-metric, but can only be described by a pertubative approach around it. This will follow in Sect. 4 but beforehand, we need some statistical definitions for the emerging structures.

1.4 The Poisson equation in an expanding universe

A useful relation for the weak lensing presented later is the Poisson equation of a density perturbation on an expanding background. Let us write the total gravitational potential as $\Phi' = \bar{\Phi} + \Phi$, where $\bar{\Phi}$ is the background value and Φ is a local perturbation. We now wish to relate the gravitational potential to the local density $\rho = (1 + \Delta(\boldsymbol{x}))\bar{\rho}(z)$, where the background density $\bar{\rho}(z)$ will decrease with cosmic expansion although the density contrast $\Delta(\boldsymbol{x})$ will typically grow.

The Poisson equation then describes how the local matter sources the gravitational potential (Bartelmann & Schneider, 2001)

$$\nabla_r^2 \Phi' = 4\pi G\rho, \tag{1.37}$$

where ∇_r^2 is the gradient with respect to physical coordinates. To single out the expansion of the Universe, we change to comoving coordiantes x, such that $\nabla_x = a\nabla_r$. We then have

$$\nabla_x^2(\bar{\Phi} + \Phi) = 4\pi G a^2 (1 + \Delta)\bar{\rho}. \tag{1.38}$$

If we assume that the superposition principle holds, then the background equation $\nabla_x^2 \bar{\Phi} = 4\pi G a^2 \bar{\rho}$ has to be seperately fulfilled, and we can single out the perturbative part

$$\nabla_x^2 \Phi = 4\pi G a^2 \Delta\bar{\rho}. \tag{1.39}$$

This statement is far from trivial, and became famous as the *Jeans swindle*. Demanding that the background solution obeys homogeneity, implies a constant gravitational potential. This in turn implies a vanishing density due to the Poisson equation. However, the mean density of the Universe is finite, so there exists a contradiction. This contradiction is luckily resolved in an expanding universe because the effect of the finite background density is cancelled by the Hubble expansion such that the perturbation (1.39) can indeed be analysed separately from the background (Falco et al., 2013).

Working our way forward from (1.39), the cosmic mean density drops as $\bar{\rho} = \rho_0 a^{-3}$ and we can relate ρ_0 to the usual density parameters $\Omega_0 = \rho_0/\rho_{\rm crit}$, where $\rho_{\rm crit} = 3H_0^2/8\pi G$. Therefore (Bartelmann & Schneider, 2001)

$$\nabla_x^2 \Phi = \frac{3H_0^2}{2a}\Omega_0 \Delta, \tag{1.40}$$

which is the Poisson equation rephrased in terms of the usual cosmological parameters H_0 and Ω_0 and where ∇_x is the gradient with respect to comoving coordinates.

Chapter 2

Brinks of the cosmic standard model

Studying a model also implies understanding its limitations. Consequently, this chapter will outline a few theoretical and observational brinks of ΛCDM.

The current state of cosmology may be summarized by stating that the abundance of the different cosmic matter constitutents is *quantitatively* quite well under control – the error bars on $\Omega_b, \Omega_r, \Omega_m, \Omega_\Lambda$ are down to a few percent. However, a *qualitative* explanation for dark matter and dark energy is still missing. So what does it mean from an observer's point of view when we speak about dark matter and dark energy?

On the lowest level, it means that both are free parameters, needed to fit the background observables like type Ia supernovae and BAO, and perturbation observables like the CMB and the cosmic large-scale structure. All of these observables are in turn convolved expressions of the Hubble function $H(z)$, and the growth function. Future measurements will not seriously alter our inference about the Hubble function or the growth function – what they may alter, however, is our qualitative understanding of dark matter and dark energy, assigning them the status of physically well understood phenomena, instead of mere fitting parameters.

This summary of the cosmic standard model may appear mildly heretic in comparison to the general tenor in cosmology, so let us back it up with further arguments. For dark energy, multiple *candidates* are in the game, but they cannot yet be discriminated with the current data. For example, the $f(R)$-model by Hu & Sawicki (2007) does not employ a cosmological constant but has a limit in which it reduces to the phenomenology of ΛCDM. Likewise, for dark matter, experiments like DAMA/LIBRA (Bernabei et al., 2014), XENON (Brown & Xenon Collaboration, 2014) or LUX (Faham & for the LUX Collaboration, 2014) etc. hunt for dark matter candidates, but yet have neither detected any WIMPs nor axions. Further, no constraints from particle physics about the *abundance* of a dark component exist – the standard model of particle physics does not include a dark sector.

Yet, it seems currently impossible that dark matter could be composed of standard model particles. Neither the six leptons, nor the four gauge bosons have the properties required of cold dark matter. A baryonic origin of dark matter is excluded as well: the abundance of baryons in the Universe can be considered fixed already a few minutes after the Big Bang. At this time the primordial nucleosynthesis (BBN) is completed and while subsequent astrophysical processes change the chemical composition of the baryons, they do not significantly alter their total abundance. BBN constraints therefore provide a solid upper limit on the total baryon density in our Universe, which turns out to be $\Omega_b h^2 = 0.0223 \pm 0.002$ (Steigman, 2006). Using $h = 0.68$, we then arrive at $\Omega_b = 0.048 \pm 0.0043$, which is about a tenth of the dark matter abundance required in ΛCDM. It is therefore impossible that dark matter consits of baryons which got locked-up after their creation into a electromagnetically hardly interacting state, such as Massive Compact Halo Objects (MACHOs), neutron stars, black holes, or other compact objects.

Prior to nucleosynthesis, a lock-up of standard model particles into primordial black holes could however produce a dark matter candidate which does not require the postulation of a dark particle. As this production mechanism precedes primordial nucleosynthesis, it also evades the BBN constraints on the baryon density. Although not yet definitely ruled out, the generation of primordial black holes requires initial conditions that are not favoured by the cosmic microwave background (Young & Byrnes, 2015). In total, we can conclude that dark matter is not understood yet.

For baryons themselves, the situation is somewhat better: the standard model of particle physics seems to explain this sector sufficiently well. However, we then need to admit that we do not know where the baryons come from cosmologically. The observed baryon-antibaryon asymmetry in our Universe is still a subject of intense investigation (Steigman, 2006).

The remainder of this chapter studies a few brinks of ΛCDM in more detail.

2.1 The origin from inflation

There is strong need for the Universe to have undergone a very early phase of accelerated expansion, called inflation. The reason is that the Friedmann equations are differential equations and therefore require initial conditions. Without inflation, these initial conditions appear highly fine-tuned and unnatural.

For example, in the remainder of this thesis, the Universe is always considered to be spatially flat – however, spatial flatness is an unstable solution of the Friedmann equations. In order to illustrate this, we allow for a spatial curvature K in this section only. The Friedmann equations in conformal time are then (Durrer, 2008)

$$\begin{aligned} \mathcal{H}^2 &= \frac{8\pi G}{3} a^2 \rho - K, \\ \mathcal{H}' &= -\frac{4\pi G}{3} a^2 (\rho + 3p), \end{aligned} \tag{2.1}$$

where we neglected the cosmological constant, or likewise absorbed it into ρ and where

$$\mathcal{H} = Ha. \tag{2.2}$$

Hence, the critical density is $\rho_{\text{crit}} = 3\mathcal{H}^2/(8\pi G a^2)$. Dividing the first Friedmann equation by $\mathcal{H}^2$ then yields

$$\Omega(t) - 1 = \frac{K}{\mathcal{H}^2}. \tag{2.3}$$

The time derivative of the deviation from spatial flatness $\Omega(t) - 1$ is then (Durrer, 2008)

$$(\Omega(t) - 1)' = (\Omega(t) - 1)\frac{8\pi G a^2}{3}\left(\rho\frac{1+3w}{\mathcal{H}}\right), \tag{2.4}$$

which shows that if there is no accelerated expansion with $w < -1/3$, then spatial flatness is an unstable solution. Furthermore, solving (2.4) numerically, it turns out that in order to have a spatial curvature of $\Omega \approx 1$, the spatial flatness about 5 minutes after the Big Bang (when nucleosynthesis took place), has to be on the order $|\Omega(t_{\text{nucleo}}) - 1| < 10^{-15}$, i.e. either the spatial flatness is extremely fine-tuned, or there should exist a convincing mechanism that produces such a negligible spatial curvature in the early Universe. An early inflationary era can provide this.

The second problem that an inflationary era can resolve, is the horizon problem. If we assume the Universe before CMB decoupling was dominated by radiation and matter only, and calculate from this the horizon (1.36) with (1.20) then this results in a horizon scale of ≈ 1° in the CMB. The CMB is however uniform over much larger scales. Again, an early period of accelerated expansion can explain this uniformity, as it increases the horizon scale. Furthermore, the inflationary framework provides a convenient mechanism for the seeds of cosmic structures. An expanding Universe can convert quantum fluctuations into classical particles, if it drives spontaneously created particles out of causal contact before they can annihilate. If the expansion is approximately exponential, a nearly scale invariant power-spectrum of the initial density perturbations is the consequence. The detailed mechanism is e.g. described in Mukhanov et al. (1992), and we shall only describe observations of such an inflationary structure initialization in Sect. 2.2.

2.2 Accelerated expansion from a scalar field

The Universe currently undergoes a phase of accelerated expansion, and in Sect. 2.1 we have seen that there is good reason to believe that there was previous phase of accelerated expansion in the early Universe. A cosmological constant can provide an exponential expansion; but other mechanisms can as well.

For example, a scalar field ϕ can produce an equation of state $w = -1$ and therefore mimic a cosmological constant. As a canonical scalar field has no internal degrees of freedom, it can be regarded as the 'next to minimal' model beyond a cosmological constant. In the framework of inflation, such a hypothetical scalar field is usually called the 'Inflaton', whereas a hypothetical scalar field responsible for the late-time cosmic acceleration is referred to as 'Quintessence'. The mechanism is however exactly the same.

Starting from the Lagrangian

$$\mathcal{L}_\phi = -\frac{1}{2}g^{\mu\nu}\partial_\mu\phi\partial_\nu\phi - V(\phi), \tag{2.5}$$

we see that the behaviour of the field is fixed once the potential $V(\phi)$ is specified. The energy density $\rho_\phi = -T^0_0$ and the pressure $p_\phi = T^i_i/3$ then follow from the energy momentum tensor

$$T_{\mu\nu} = -\frac{2}{\sqrt{-g}}\frac{\delta(\sqrt{-g}\mathcal{L})}{\delta g^{\mu\nu}}, \tag{2.6}$$

such that

$$\begin{aligned}\rho_\phi &= \frac{1}{2}\dot{\phi}^2 + V(\phi),\\ p_\phi &= \frac{1}{2}\dot{\phi}^2 - V(\phi).\end{aligned} \tag{2.7}$$

The equation of state is then

$$w_\phi = \frac{p_\phi}{\rho_\phi} = \frac{\frac{1}{2}\dot{\phi}^2 - V(\phi)}{\frac{1}{2}\dot{\phi}^2 + V(\phi)}, \tag{2.8}$$

such that $w = -1$ is reached if $\dot{\phi} = 0$, and for slowly varying ϕ small deviations from $w = -1$ exist. During the cosmic evolution, the field ϕ will also evolve, such that the equation of state will additionally vary with redshift. The field ϕ adapts to expansion by following the conservation equation $\nabla_\nu T^{0\nu} = 0$ from (1.17), which then reads (Amendola & Tsujikawa, 2010a)

$$\ddot{\phi}(t) + 3H(t)\dot{\phi}(t) + V_{,\phi} = 0. \tag{2.9}$$

If we further demand that this scalar field dominates the Hubble expansion, and omit a cosmological constant, then the first Friedmann equation (1.13) is

$$H^2 = \frac{1}{3M_{pl}^2}\left(\frac{1}{2}\dot{\phi}^2 + V(\phi)\right), \tag{2.10}$$

where the energy density ρ_ϕ from (2.7) was substituted and $1/M_{pl}^2 = 8\pi G$, where M_{pl} is the Planck mass to match the usual terminology of inflationary physics. We therefore see, that if a scalar field was responsible for inflation, then the Hubble expansion $H(t)$ is fixed by the choice of the potential $V(\phi)$, and vice versa: if the Hubble expansion during inflation could be reconstructed from measurements, then the potential $V(\phi)$ were known.

In fact, this reconstruction is being attempted. In the framework of inflation, when we restrict ourselves to a slow-roll evolution of the field and neglect the kinetic term, then (2.9) and (2.10) reduce to (Planck Collaboration et al., 2014)

$$\begin{aligned}3H(t)\dot{\phi} &\approx -V_{,\phi}\,,\\ H^2 &\approx \frac{V(\phi)}{3M_{pl}^2}.\end{aligned} \tag{2.11}$$

Then one can attempt to measure the conventionally defined slow-roll parameters ϵ_V and η_V (Planck Collaboration et al., 2014)

$$\epsilon_V = \frac{M_{pl}^2 V_{,\phi}^2}{2V^2}, \qquad \eta_V = \frac{M_{pl}^2 V_{,\phi\phi}}{V}, \tag{2.12}$$

which describe on one hand deviations of the Hubble expansion from a purely exponential expansion for which $w_\phi = -1$. From (2.7) we had seen that $w = -1$ is reached if $\dot{\phi} = 0$. Due to the conservation equation (2.9), this can however not be maintained, and the first order deviation from $w_\phi = -1$ in the slow-roll is then (Amendola & Tsujikawa, 2010a)

$$w_\phi \approx \frac{2}{3}\epsilon_V - 1. \tag{2.13}$$

On the other hand, the slow-roll parameters ϵ_V and η_V allow to reconstruct the first terms of the potential's Taylor series via its derivatives $V_{,\phi}$ and $V_{,\phi\phi}$. For dark energy research, the possible evolution of w with redshift is hoped to be measured soon, which would then allow to reconstruct the potential $V(\phi)$. For inflation, no time-resolved measurement is possible, and instead the potential needs to be reconstructed from the Hubble expansion's indirect effect on the structure growth, which then imprints itself onto the power spectrum in the CMB. From the CMB, we can infer the primordial power spectrum of scalar perturbations

$$P(k) = A_s \left(\frac{k}{k_0}\right)^{n_s - 1 + \frac{1}{2}\frac{\mathrm{d}n_s}{\mathrm{d}\log k}\log(k/k_0) + \frac{1}{6}\frac{\mathrm{d}^2 n_s}{\mathrm{d}\log k^2}(\log(k/k_0))^2 + \ldots}, \tag{2.14}$$

where $k_0 = 0.05\mathrm{Mpc}^{-1}$ is the pivot scale that can be well resolved by Planck and A_s is the amplitude. The scalar spectral index is n_s and a possible running of it is $\mathrm{d}n_s/\mathrm{d}\log k$. In the slow-roll approximation one has (Planck Collaboration et al., 2014)

$$\begin{aligned} n_s - 1 &\approx 2\eta_V - 6\epsilon_V, \\ \frac{\mathrm{d}n_s}{\mathrm{d}\log k} &\approx 16\epsilon_V\eta_V - 24\epsilon_V^2 - 2\xi_V^2, \end{aligned} \tag{2.15}$$

where ξ_V^2 depends on the third derivative of the potential as

$$\xi_V^2 = \frac{M_{pl}^4 V_{,\phi} V_{,\phi\phi\phi}}{V^2}. \tag{2.16}$$

If the data were good enough, a running of the running $\mathrm{d}^2 n_s/\mathrm{d}\log k^2$ and higher equivalents could also be measured, which would give access to ever higher derivatives of the inflaton potential. The current constraints by Planck are $n_s = 0.968 \pm 0.006$ and the running is well compatible with zero, $\mathrm{d}n_s/\mathrm{d}\log k = -0.003 \pm 0.007$ (Planck Collaboration et al., 2015). For the future it is being hoped that measurements of the tensor spectral index n_t will put additional competitive constraints on the slow-roll parameters for inflation. Similarly, upcoming surveys which target late-time observables like weak lensing or type Ia supernovae, hunt for a potential redshift evolution or a deviation from $w = -1$ of the equation of state parameter of dark energy.

2.3 Dark degeneracy

It is true that ΛCDM fits the cosmological data so well, that until now, no deviation from it has been detected. Yet, if one fits a model to data, one will always get constraints on the model parameters. This does not yet mean that the model describes the underlying physics correctly. The model ΛCDM then appears in a very interesting light if we consider the work by Kunz (2009). The results of Kunz (2009) shall be presented here, since they are an important wake-up call to what happens if a parameterized model is fitted to the data.

ΛCDM is a successful fit to the cosmological data because it introduces two separate dark components, dark matter and dark energy in the form of a cosmological constant Λ. Dark matter has not yet been detected in any particle physics experiment, and the cosmological constant is – although a natural degree of freedom in GR – still subject of strong debates due to the problems described in Sect. 2.4. One may then ask, whether this split into a dark component of $w = 0$ and one dark component $w = -1$, may be arbitrary. As Kunz (2009) shows, this seems to be the case indeed.

Starting from the argument that the Einstein equations only probe the *total* energy-momentum tensor of all cosmic matter constituents, and that the split

$$T^{\mu\nu}_{\text{dark}} = T^{\mu\nu}_{w_1} + T^{\mu\nu}_{w_2}, \tag{2.17}$$

of the dark energy-momentum tensor $T^{\mu\nu}_{\text{dark}}$ into two subcomponents with equations of state w_1 and w_2 is arbitrary from the point of view of the Einstein equations, Kunz (2009) derives the relation

$$w_{DE}(z) = \frac{H^2(z) - \frac{2}{3}(1+z)H(z)\frac{\mathrm{d}}{\mathrm{d}z}H(z)}{H_0^2\Omega_m(1+z)^3 - H^2(z)} \tag{2.18}$$

for an FLRW metric. Here Ω_m refers to the abundance of a dark matter component with $w = 0$, and w_{DE} is the equation of state parameter for dark energy. The expansion rate $H(z)$ is what background observations probe. Increasing the measurement sensitivities means that the errors on $H(z)$ will go down, making the relation (2.18) less noisy. However, even in the absence of any measurement uncertainties on $H(z)$, we see from (2.18) that the dark matter abundance Ω_m and the dark energy equation of state w_{DE} are degenerate. Fixing one, and inserting the measured constraints on $H(z)$, the other is given by (2.18). In this sense, dark energy and dark matter are not directly observable via cosmological constraints only, which Kunz (2009) demonstrates for a combination of supernova- and BAO data. If particle physicists could find dark matter by their experiments, and also put constraints on its abundance, the degeneracy (2.18) could be broken.

In Kunz (2009), the author then proceeds to show that also the (in 2009 available) CMB-data of WMAP cannot break the degeneracy (2.18), even though these are perturbations, instead of background observations.

In summary, the work by Kunz (2009) demonstrates that ΛCDM suffers from what is called a parameterization bias. To be fair, it should be pointed out that models which introduce multiple separate fluids for dark energy on top of a dark matter, will also inherit the problem of splitting (2.17), only that then, more energy-momentum tensors appear. Consequently, Kunz (2009) argues that also such alternatives to ΛCDM have the problem of parameterization biases. For the future of cosmology, it is therefore mandatory to keep a close eye on on direct dark matter detections by particle physicists.

2.4 The cosmological constant problem

The cosmological constant is a problem – independent of its value. Therefore, let us here not orbit around the usual statement that the observed value of Λ is 122 orders of magnitude too low when compared to expectations, especially as we have seen in Sect. 2.3 that cosmology has a problem when measuring Λ. Let us rather comment *why* Λ is an interesting problem on its own.

From the point of view of General Relativity, Λ is a classical degree of freedom which is allowed basically due to a mathematical indeterminacy. This can be e.g. stated in the shape of the following theorem.

> The Lovelock theorem: If the Einstein tensor in four dimensions shall satisfy $\nabla_\nu G^{\mu\nu} = 0$ and only depend on the metric $g_{\mu\nu}$ and its first and second derivatives (as is the case in GR), then $G^{\mu\nu}$ is unique, apart from a possible addition of a constant Λ as in
>
> $$G_{\mu\nu} \rightarrow G_{\mu\nu} + \Lambda g_{\mu\nu}, \tag{2.19}$$
>
> where the value of Λ is not specified by GR any further (Straumann, 2013).

Of course, one could relax the assumptions of the Lovelock theorem, and also allow for more than 4 dimensions or allow $G^{\mu\nu}$ to also depend on higher order derivatives of the metric or other elements of a gravitational theory. All this would go beyond standard General Relativity but at least within the framework of standard GR, Lovelock's theorem shows that Λ makes a natural appearance. It is true, Λ has the same dimension as the Ricci scalar R, and therefore sets a second scale to the curvature of spacetime, and it is not clear where this scale should come from, but it is generally admissible in GR.

Let us here however point out a problem. It is established beyond any doubt that matter fields exist in our Universe, and these fields have interaction potentials. For the sake of simplicity, let us assume a scalar field as in Sect. 2.2 and add a constant V_0 to its potential

$$V(\phi) \rightarrow V(\phi) + V_0. \tag{2.20}$$

This change of zero-point energy does not influence any other interactions, but it affects gravity. The equation of state (2.8) shows that the addition of V_0 is equivalent to introducing a cosmological constant. The puzzle of what determines the level of a potential then translates directly into the cosmological constant problem (Durrer, 2008).

Further, being a classical field theory from 1915, GR precedes quantum field theory. Quantum field theory, in turn, treats a free field as a collection of infinitely many quantum mechanical oscillators. As a quantum mechanical oscillator with kinetic term $\dot{\phi}$ and potential $V(\phi)$ cannot be at rest due to the uncertainty principle, this implies an infinite zero-point energy for the free field (Martin, 2012)

$$\langle 0|T_{\mu\nu}|0\rangle = -\rho_{\text{vac}} g_{\mu\nu}, \tag{2.21}$$

which is in quantum field theory addressed by adding or subtracting a constant to the potential $V(\phi)$ as in (2.20). The problem of the zero-point energy therefore cannot be ignored when quantum field theory shall include gravity. Additionally, quantum field theory predicts spontaneous creation of particle-antiparticle pairs out of the vacuum. Experiments on the Casimir effect demonstrate that these vacuum fluctuations exist, and it can be expected that they gravitate normally, and therefore should also influence the value of an observed cosmological constant (Martin, 2012).

Of course, a consistent theory of quantum gravity is still missing, but we can maintain that however unproblematic Λ may be within the classical theory of GR alone, it does become problematic when going beyond gravity. A very throrough review about the current state of research on topic is given by Martin (2012). We therefore conclude that even if cosmology stops to wonder about the cosmological constant at some point, other fields of research, especially quantum field theory, will continue working on the cosmological constant problem.

Chapter 3

Supernovae of type Ia and clustering statistics

In this chapter, the cosmological relevance of type Ia supernovae is described, and the statistical basis for the description of data sets from cosmological structures are given. The physics of cosmological structure formation itself is then described in chapter 4.

3.1 Supernovae of type Ia

Supernovae of type Ia are extremely luminous events that can outshine their host galaxies. Historically, it was the analysis of supernova data under the assumption of a flat Universe, that lead in 1998/1999 to the discovery of the Universe's accelerated late-time expansion (Riess et al., 1998; Perlmutter et al., 1999). The confirmation of the spatial flatness was then brought in 2000, when the CMB experiments BOOMERANG and MAXIMA were able to resolve the first acoustic peak in the temperature-fluctuation power spectrum (de Bernardis et al., 2000; Balbi et al., 2000).

In order to understand supernovae, a little detour in stellar evolution is needed. A stable star maintains the hydrostatic equilibrium (Maoz, 2007)

$$\frac{\mathrm{d}P(r)}{\mathrm{d}r} = -\frac{GM(r)\rho(r)}{r^2}, \tag{3.1}$$

where P, M and ρ are the pressure, the mass and the density respectively, and r is the radius as measured from the center of the star. So a star is in equilibrium when the pressure of each shell equilibrates the gravity of the outer stellar layers that press down on this shell. This pressure can be provided by thermal pressure, or – in the case of stellar remnants – the quantum mechanical Fermi pressure of either electrons or neutrons.

The main two factors that determine a star's evolution are its mass and its metallicity. All stars have in common that they initially burn hydrogen in their cores, synthesizing helium. If hydrogen is exhausted, the core of the star contracts under its self-gravity, until the fusion of successively heavier elements begins, which again provide enough thermal pressure to counterbalance the self-gravity.

If the star is massive enough, about $M > 8M_\odot$ (Maoz, 2007; Voigt, 2012), its core will subsequently undergo H-, He-, C-, Ne-, and Si- burning. Then, nuclear fusion stalls because heavier elements than iron rather release when decaying, instead of when being synthesized. After having completed the Si-burning, the stellar core can then no longer produce enough thermal pressure to counteract the self-gravity. It then collapses in free fall, until it has reached a density where the quantum mechanical Fermi pressure of its neutrons stabilizes it. This end result is called a neutron star.

While successivley heavier elements are burned in the core of such a massive star, the stellar shells still contain lighter elements. When the core faces the iron catastrophe and collapses, the outer shells of the star are not supported from below any more, and also collapse. When they hit the solid surface of the newly formed neutron star, they are reflected and forced again through the layers that fell after them. As these consist of combustible light elements, a very luminous explosion results which can then be observed as a supernova. The remaining hydrogen of the outer shells gives its presence away by absorption lines in the supernova spectrum. The end of such a massive star therefore produces a type II supernova, where the type II is defined by the presence of hydrogen lines in the spectrum (Amendola & Tsujikawa, 2010a).

In contrast, supernova of type Ia are thought to be thermonuclear disruption of a white dwarf. A white dwarf is the remnant of a star with less than about eight solar masses. Depending on the exact mass and metallicity of the star, its outer shells can become convective and start to pulsate during its evolution. It is then possible, that the star blows away its own shells before the core has reached the carbon-burning stage. The escaping shells can be seen for a few 10.000 years as planetary nebulae before they disperse completely, leaving behind the old stellar core (Maoz, 2007). This exposed core consists primarily of carbon and oxygen, with some helium – in principle a highly combustible mixture, if put under sufficient pressure. However, usually, the old core is supported by too low a pressure for combustion: it does not yet burn carbon because it is stabilized against gravitational collapse by the Fermi pressure of its degenerate electron gas. Such a stellar remnant is called a white dwarf.

This balance between Fermi pressure and self-gravity inside a white dwarf can be tipped, if the white dwarf accretes further mass. From hydrostatic equilibrium, it follows that the pressure scales as

$$P \propto \frac{M\rho}{r} \tag{3.2}$$

and since for a spherical dwarf $\rho \propto M/r^3$ it is

$$P \propto \frac{M^2}{r^4}. \tag{3.3}$$

In a non-relativistic degenerate electron gas, pressure and density scale as (Maoz, 2007)

$$P \propto \rho^{5/3}. \tag{3.4}$$

Additionally, in a degenerate Fermi gas, the pressure is *independent* of the temperature – it is solely provided by the quantum-mechanical Pauli principle for fermions.

Eliminating the pressure in (3.3), we have

$$r \propto M^{-1/3}, \tag{3.5}$$

meaning that when a white dwarf accretes further mass, it needs to *shrink* in order to provide enough pressure to counteract the additional self-gravity.

Pressure, however, is proportional to the energy density ε_{e^-} of the electrons. In the non-relativistic regime it therefore scales as

$$P \propto \varepsilon_{e^-} \propto v_{e^-}^2, \tag{3.6}$$

i.e. the pressure scales with the squared electron velocity v_{e^-}. In the relativistic regime, we have a linear relation instead

$$\varepsilon_{e^-} \propto v_{e^-}. \tag{3.7}$$

Increasing the pressure therefore means increasing the electron velocities. However, these are limited by the speed of light $v_{e^-} < c$, so obviously, there will come a moment when the Fermi pressure of the electrons cannot support the white dwarf against its own self-gravity anymore. This moment is reached when the white dwarf exceeds the Chandrasekhar mass of about $1.4 M_\odot$ (Maoz, 2007).

A white dwarf more massive than $1.4 M_\odot$ will collapse and thereby exceed the critical pressure needed for the fusion of its carbon and oxygen. As the pressure in the degenerate white dwarf is independent of the temperature, burning carbon and oxygen cannot increase the Fermi pressure. The thermal pressure of the heavy nuclei themselves is additionally subdominant to the Fermi pressure. Therefore, a sufficient pressure to stop the gravitational collapse cannot be built up and instead the nuclear fusion undergoes a runaway process which then disrupts the white dwarf.

Forcing a white dwarf beyond its Chandrasekhar mass is currently believed to occur either by accreting mass from a companion star in a double system (the single-degenerate scenario), or by two white dwarves merging (the double-degenerate scenario). As it is always about one Chandrasekhar mass that ignites, the luminosity of such a type Ia supernova scatters only little. Spectroscopically, such a supernova is classified as Ia if the spectrum contains the line of singly ionized silicon and no hydrogen lines (Amendola & Tsujikawa, 2010a).

Cosmologically, supernovae of type Ia are of special importance because their peak luminosity scatters little anyway, and allows to be standardized further. Phillips (1993) and Hamuy et al. (1996) have shown that less luminous type Ia supernovae fade faster. Consequently, observations of type Ia supernovae typically sample the supernova's lightcurve for about 40 days in order to correct for this peak luminosity - lightcurve width effect (Riess et al., 1998). Additionally, extinction is being corrected for, as well as the fact that depending on the redshift of the supernova, its brightness in different astronomical filters such as the V and B band will vary; this correction became known as K-correction (Perlmutter et al., 1999). After these corrections, it can then be assumed that the observed supernovae of type Ia have an intrinsic brightness that scatters only very little.

The only directly measurable quantity that is related to the intrinsic brightness of an object, is the flux F which we receive from it. For reasons of historical compatibility with the ancient astronomers Hipparchus and Ptolemy, an astronomical convention is to then convert these measured fluxes into *apparent magnitudes* defined by (Voigt, 2012)

$$m_1 - m_2 = -\frac{5}{2}\log_{10}\left(\frac{F_1}{F_2}\right) \tag{3.8}$$

if F_1 and F_2 are fluxes of two different objects. The zero-point of this scale has traditionally been the brightness of the star Vega (αLyrae) but nowadays there exist hints that Vega belongs to the class of stars that exhibit slight periodic variations of their brightness. Consequently, the zero-point of the magnitude system is now calibrated on the 96 stars of the *North Polar Sequence* and Vega has now a magnitude of $m = 0.03$ mag (Voigt, 2012). As the perceived brightness of an object depends on its distance to us, an *absolute* magnitude is defined as the apparent magnitude that the object would have if it lay in a luminosity distance of 10 parsec:

$$m - M = 5\log_{10}\left(\frac{D_{\mathrm{lum}}}{10\mathrm{pc}}\right). \tag{3.9}$$

Cosmologically more relevant are distances in megaparsec, so from (3.9) follows

$$m - M = 5\log_{10}\left(\frac{D_{lum}}{10^6 pc}\right) + 25. \tag{3.10}$$

The quantity $m - M$ is also known as distance modulus μ. From calibrations on Cepheid stars, it is known that the absolute magnitude of nearbye type Ia supernovae is approximately $M = -19.5$ in the astronomical B-band, which is as bright as an entire galaxy (Voigt, 2012). However, because the distance to nearby Cepheid stars scales with the Hubble constant, the absolute magnitude of the supernovae is degenerate with the determination of the Hubble constant (Riess et al., 1998). The remaining scatter around this absolute magnitude after all corrections is about 0.21mag (Phillips, 1993; Hamuy et al., 1996).

Via the luminosity distance the distance modulus depends on the background expansion and therefore on the cosmological parameters. The need of Perlmutter et al. (1999); Riess et al. (1998) to constrain themselves to a spatially flat Universe arises because supernova data alone cannot distinguish well between a cosmological constant and curvature. The CMB is needed to break this degeneracy. Riess et al. (1998) point out that CMB and supernova data yield nearly orthogonal constraints in the $(\Omega_m, \Omega_\Lambda)$-plane: While the SNe Ia are approximately sensitive to the combination $\Omega_m - \Omega_\Lambda$, the angular scale of the first acoustic peak is sensitive to $\Omega_m + \Omega_\Lambda$ (Riess et al., 1998).

The fact that the absolute brightness of type Ia supernovae needs to be determined from Cepheid observations under the assumption of a Hubble constant (Riess et al., 1998), does however not challenge the discovery of the acceleration: The Hubble constant H_0 is proportional to the derivative $\dot{a}$ of the scale factor today; the acceleration is the second derivative $\ddot{a}$ of the scale factor with respect to cosmic time, such that it measures *changes* in the Hubble parameter throughout cosmic history. These changes can be measured independently of H_0.

3.2 Random fields and power spectra

The density fields of our Universe are a single realization of a random process. Therefore, only the statistics of given sky maps are studied. If $f(\boldsymbol{x})$ is such a sky map, then its 2-point correlation function is given by

$$\xi(y) = \langle f(\boldsymbol{x}) f(\boldsymbol{x}+\boldsymbol{y}) \rangle, \tag{3.11}$$

where ξ can only depend on the modulus y in an isotropic and homogenous Universe as long as $\boldsymbol{x}$ and $\boldsymbol{y}$ lie in a spatial hyperplane of constant time. Due to the attractive nature of gravity, the correlation function of galaxy positions is typically positive on small scales. In order to integrate to zero, this implies a negative correlation on very large scales, however this effect is typically not seen in cosmic data due to the presence of noise. Note, that the angular brackets actually denote an ensemble average over different realizations of the random process. In cosmology, this ensemble average is necessarily replaced by a spatial averaging, as we have only one sky to observe. If the coherence length L_{coh} is much smaller than the horizon, this replacement is well motivated. However, in the presence of gravitational lensing, minute differences between directional averages and ensemble averages exist (Bonvin et al., 2015).
Further, if the Fourier transform is given by

$$\begin{aligned} \hat{f}(\boldsymbol{k}) &= \int \mathrm{d}^3 x f(\boldsymbol{x}) e^{i\boldsymbol{k}\boldsymbol{x}}, \\ f(\boldsymbol{x}) &= \frac{1}{(2\pi)^3} \int \mathrm{d}^3 k \hat{f}(\boldsymbol{k}) e^{-i\boldsymbol{k}\boldsymbol{x}}, \end{aligned} \tag{3.12}$$

then one defines a power spectrum $P(k)$ by

$$\langle \hat{f}(\boldsymbol{k}) \hat{f}^*(\boldsymbol{k'}) \rangle = (2\pi)^3 \delta_D(\boldsymbol{k}-\boldsymbol{k'}) P(k). \tag{3.13}$$

The Wiener-Khinchin theorem then states that

$$\xi(r) = \frac{1}{(2\pi)^3} \int P(k) \exp(-i\boldsymbol{k}\cdot\boldsymbol{r}) \mathrm{d}^3 k, \tag{3.14}$$

such that the 2-point correlation function and the power spectrum are Fourier conjugates. As $\xi(r)$ is a real valued function, and as the cosmological power spectrum only depends on the modulus k, equation (3.14) can be further simplified to

$$P(k) = 4\pi \int \xi(r) \frac{\sin kr}{kr} r^2 \mathrm{d}r. \tag{3.15}$$

Conventionally, the power spectrum of the late-time large-scale structure is normalized by (Amendola & Tsujikawa, 2010a; Bartelmann & Schneider, 2001)

$$\sigma_R^2 = \frac{1}{2\pi^2} \int P(k) W_R^2(k) k^2 \mathrm{d}k, \tag{3.16}$$

evaluated at $R = 8h^{-1}$Mpc, and where $W_R(k)$ is a window function that smoothes out structures on scales smaller than R.
If the statistics of a field is non-Gaussian, then higher-order correlation functions exist, e.g. the three-point function

$$\langle f(\boldsymbol{x}) f(\boldsymbol{x}+\boldsymbol{y}) f(\boldsymbol{x}+\boldsymbol{z}) \rangle. \tag{3.17}$$

The Fourier transforms of these quantities are called polyspectra and are of cosmological importance in basically two cases: for studying primordial non-Gaussianities (Amendola, 2002; Fergusson & Shellard, 2007) and then for studying the non-linear structure growth in the late Universe.

Chapter 4

Cosmological structure formation

Slight over- and underdensities Δ of the cosmic matter distribution

$$\Delta(\boldsymbol{x}) = \frac{\rho(\boldsymbol{x}) - \bar{\rho}}{\bar{\rho}}, \tag{4.1}$$

where $\bar{\rho}$ is the background density, can become gravitationally unstable, and thereby enhance with cosmic time, leading to the growth of cosmic structures. If the relative over- or underdensity with respect to the cosmic mean density is small enough, then the growth of these structures can be described with linear perturbation theory around the cosmic background solution. About 380.000 years after the Big Bang, these growing structures of the early Universe have imprinted themselves as temperature fluctuations on the nowadays observable cosmic microwave background (CMB). These fluctuations in the temperature T have a relative amplitude of $\Delta T/T \approx 10^{-5}$, demonstrating that the assumption of linearity is well justified for CMB physics (Durrer, 2008).

In the subsequent evolution of the Universe, a web of dark matter has formed from these early perturbations, and galaxies have formed along this cosmic web, allowing its indirect observation. It is assumed that the initial conditions that lead to the primary perturbations $\Delta^{\mathrm{ini}}(k)$ of the cosmic matter field, were set by inflation as described in Sect. 2.1.

If the statistics of these initial perturbations is Gaussian, $\Delta^{\mathrm{ini}}(k) \sim \mathcal{G}(0, P(k))$, then they can be completely described by a power spectrum $P(k)$, that quantifies the variance of the cosmic over- and underdensities as a function of their scale. We shall assume that all matter species in our Universe followed the same initial power spectrum which is the so-called 'adiabatic mode' of inflation.

However, since different matter species possess different interactions, they will develop differing power spectra during the cosmic evolution. The nowadays observed power spectrum of photons is highly oscillatory due to the pressure in the coupled photon-baryon plasma in the early Universe. In contrast, the pressureless cold dark matter simply collapses under gravity, leading to a relatively featureless power spectrum, only adorned by very mild wiggles, that are an imprint of the baryonic acoustic oscillations onto the dark matter power spectrum, because both couple to gravity. The following calculations will quickly become technically demanding, so a short summary of the main physical effects seems due.

The linear growth of these structures can be described by perturbing Einstein's field equations, and splitting them into a background solution and perturbations around it. Due to diffeomorphism invariance, the definition of the background is however arbitrary, leading to questions of gauge dependence (Mukhanov et al., 1992; Amendola & Tsujikawa, 2010a; Durrer, 2008) which need to be adressed.

The perturbed Einstein equations will then depend on a perturbed energy momentum tensor. On the microscopic level, $T^{\mu\nu}$ arises as second moment of a phase space distribution function f and the perturbations $\delta T^{\mu\nu}$ therefore originate from perturbations of this distribution function. The evolution of the phase space distribution function is governed by the Liouville equation or the Boltzmann equation, where the latter also needs the definition of a collision term. We therefore see, how quickly the complexity of fully relativistic structure formation grows.

Fortunately, under certain circumstances, the relativistic and microscopic approach can be significantly simplified: the equations of ideal hydrodynamics follow from the Boltzmann equation by taking moments. Also, on subhorizon scales, gravity can be approximately described as Newtonian. Consequently, there exists a limit for cosmic structure formation, in which it is sufficient to describe the gravitational instability of a perturbed ideal fluid under Newtonian gravity. This limit is well satisfied for the late-time growth of dark matter perturbations.

For describing the structure growth in the early Universe, this limit is however not sufficient: The horizon scale at decoupling corresponds to about 1 degree on the observed CMB map, which corresponds to the size of a thumbnail when held up against the sky. We can therefore see structures on much larger scales than their cosmic horizon, and therefore have to take into account horizon effects, to which Newtonian gravity is insensitive. Also, the ideal fluid approximation is not always well fulfilled in the pre-CMB area: obviously, during recombination, the mean free path of photons grew from initially zero (which corresponds to an ideal fluid), to infinity (which corresponds to free streaming). In order to account for this growing mean free path and the decreasing collision rate of photons during recombination, the photons must be described with the Boltzmann equation, including a collision term that describes the Thomson scattering of photons on electrons. Neutrinos, if approximated as having a zero cross section with matter, also free stream and consequently cannot be an ideal fluid.

Since all matter components couple to the metric, they will influence the evolution of each other's perturbations by gravitational interaction. The complexity of this early-universe physics is therefore usually handled with Boltzmann codes like CLASS (Audren et al., 2013a; Blas et al., 2011) or CAMB (Lewis & Challinor, 2011). If one is only interested in late-time cosmology, such codes can compress all the complex effects from the early Universe into a *transfer function*, which – if multiplied with an initial inflationary power spectrum – translates this into the post-CMB power spectrum. The further evolution of this post-CMB power spectrum can then be approximately described in the framework of Newtonian gravity.

The following sections will cast these considerations into mathematical form.

4.1 Relativistic structure growth: Perturbing Einstein's field equations

In chapter 1, we have solved the Einstein equations $G^{\mu\nu} = 8\pi G T^{\mu\nu}$ under the symmetry assumptions of homogeneity and isotropy, leading to the FLRW metric. The matter in our Universe is however not perfectly homogeneously distributed, as pronounced structures like voids and filamentary overdensities exist. We can study the growth of these structures by perturbing the FLRW metric.

In general, a metric can exhibit scalar, vector and tensorial perturbations. In an expanding universe, vector perturbations decay because they are redshifted. Tensorial perturbations give rise to gravitational waves, and collapsing structures originate from scalar perturbations of the metric. At linear order, scalar, vector and tensor perturbations decouple and the following analysis will therefore only focus on scalar perturbations.

We use the ansatz

$$G^\mu_\nu = \bar{G}^\mu_\nu + \delta G^\mu_\nu \quad \text{and} \quad T^\mu_\nu = \bar{T}^\mu_\nu + \delta T^\mu_\nu, \tag{4.2}$$

where quantities marked by an overbar are the background quantities that lead to the FLRW solution and δG^μ_ν and δT^μ_ν are the perturbations. The background solution leads to the FLRW cosmology and what remains to be solved are the perturbed Einstein equations,

$$\delta G^\mu_\nu = 8\pi G \delta T^\mu_\nu, \tag{4.3}$$

meaning the evolution of the background is not affected by the evolution of the perturbations. As Einstein's field equations are however non-linear coupled differential equations, the superposition principle that is so often used for linear differential equations, does not hold and perturbations may affect the evolution of the background. This means that first averaging over the large-scale inhomogeneities of our Universe and then solving the Einstein equations (which leads to the FLRW solution) and then working our its perturbations, does not need to produce the same result as first working out the Einstein equations of the inhomogeneous Universe, and then average, and identify perturbations on top of the average. This effect is known as *back reaction*. The final proof of whether back reaction is important in ΛCDM or not, is yet to be worked out, but the current consensus in cosmology is that back reaction seems to be negligible, and the ansatz (4.2) is then well justified in cosmology.

Accepting (4.2), the perturbations $\delta G^{\mu\nu}$ stem from perturbations of the metric, for which we set

$$g_{\mu\nu} = \bar{g}_{\mu\nu} + \delta g_{\mu\nu}, \tag{4.4}$$

where $\bar{g}_{\mu\nu}$ is the FLRW metric.

The inversion of a 4 × 4-metric is a highly non-linear operation but we want to restrict ourselves to a linear perturbation analysis. Therefore, $\delta g^{\mu\nu}$ is not simply the matrix inverse of $\delta g_{\mu\nu}$. Instead, demanding that the splitted total metric 4.4 satisfies $g_{\alpha\gamma} g^{\gamma\beta} = \delta^\beta_\alpha$ leads to (Amendola & Tsujikawa, 2010a)

$$\delta g^{\mu\nu} = -\delta g_{\alpha\beta} \bar{g}^{\alpha\mu} \bar{g}^{\beta\nu}. \tag{4.5}$$

The split into a background and its perturbations is gauge-dependent. Before we choose a gauge in the following calculations, let us first point out what the issue with gauges is: General Relativity describes gravity by a curvature of spacetime, yet, we cannot associate a meaning to spacetime curvature in *particular* coordinates, because the reference frame is subject to choice. Mathematically, this means that General Relativity is invariant under diffeomorphisms φ. A Diffeomorphism is a bijective and smoothly differentiable map, and the inverse map is so as well. Consequently, a diffeomorphism allows to map smoothly and uniquely back and forth between two manifolds. Diffeomorphism invariance of General Relativity then means, that if $\{g, \psi_m\}$ are the metric and the matter fields that solve the Einstein equations, and φ^* is the pull-back of a diffeomorphism φ, then $\{\varphi^*(g), \varphi^*(\psi_m)\}$ solve the Einstein equations as well. In other words: solutions that can be mapped onto each other by a diffeomorphism, are physically equivalent (Straumann, 2013).

We can then use diffeomorphisms to gauge: this will not change the physical setup, but the freedom to choose the background coordinates may simplify calculations (Mukhanov et al., 1992). For example, let us transform old coordinates x^μ onto new coordinates $\tilde{y}^\mu = \varphi(x^\mu)$, then the metric changes as (Amendola & Tsujikawa, 2010a)

$$\tilde{g}_{\mu\nu} = \frac{\partial x^\alpha}{\partial \tilde{x}^\mu}\frac{\partial x^\beta}{\partial \tilde{x}^\nu} g_{\alpha\beta}. \tag{4.6}$$

In order to work out the theory of a physical effect, e.g. structure growth, it is sufficient to work consistently within one gauge, because physically, this gauge is as good as any other. Only when observables shall be compared to actually measurements, then gauge-invariant observables are mandatory such that the physical interpretation is unambiguous.

In this introduction to relativistic structure growth, we will work in the longitudinal gauge in which the Bardeen potential Φ is equal to the Newtonian gravitational potential (Mukhanov et al., 1992). For a very thorough comparison and derivation of the equations for structure formation in both gauges, see Ma & Bertschinger (1995). In the longitudinal gauge, the line element of the spatially flat perturbed FLRW metric, with scalar perturbations only, is

$$\mathrm{d}s^2 = a^2(\eta)[-(1+2\Psi)\mathrm{d}\eta^2 + (1+2\Phi)\delta_{ij}\mathrm{d}x^i\mathrm{d}x^j] \tag{4.7}$$

where we use the sign convention for the Bardeen potentials Ψ and Φ from (Amendola & Tsujikawa, 2010a), and η is the conformal time. Derivatives with respect to η will be denoted by a prime. We then have

$$\delta g_{\mu\nu} = \begin{pmatrix} -2a^2\Psi & 0 \\ 0 & 2a^2\Phi\delta_{ij} \end{pmatrix} \qquad \rightarrow \qquad \delta g^{\mu\nu} = \begin{pmatrix} 2\Psi/a^2 & 0 \\ 0 & -2\Phi/a^2\delta_{ij} \end{pmatrix}, \tag{4.8}$$

where (4.5) was used for the inversion. To first order, the perturbed Christoffel symbols are then

$$\delta\Gamma^\mu_{\nu\lambda} = \frac{1}{2}\delta g^{\mu\alpha}(\bar{g}_{\alpha\nu,\lambda} + \bar{g}_{\alpha\lambda,\nu} - \bar{g}_{\nu\lambda,\alpha}) + \frac{1}{2}\bar{g}^{\mu\alpha}(\delta g_{\alpha\nu,\lambda} + \delta g_{\alpha\lambda,\nu} - \delta g_{\nu\lambda,\alpha}) \tag{4.9}$$

which follow to be (Amendola & Tsujikawa, 2010a)

$$\begin{aligned}\delta\Gamma^0_{00} &= \Psi' \\ \delta\Gamma^0_{ij} &= \delta_{ij}[\Phi' + 2\mathcal{H}(\Phi - \Psi)] \\ \delta\Gamma^0_{0i} &= \delta\Gamma^i_{00} = \partial_i\Psi \\ \delta\Gamma^i_{j0} &= \Phi'\delta^i_j.\end{aligned} \tag{4.10}$$

The meaning of the perturbed Christoffel symbols $\delta\Gamma^i_{jk}$ is that an observer, comoving with the background expansion, will see local gravitational fields due to the small-scale inhomogeneities.

The linear perturbation of the metric is then propagated through to the Einstein tensor for which we need the perturbed Ricci tensor

$$\delta R_{\mu\nu} = \delta\Gamma^\alpha_{\mu\nu,\alpha} - \delta\Gamma^\alpha_{\mu\alpha,\nu} + \delta\Gamma^\alpha_{\mu\nu}\Gamma^\beta_{\alpha\beta} + \Gamma^\alpha_{\mu\nu}\delta\Gamma^\beta_{\alpha\beta} - \delta\Gamma^\alpha_{\mu\nu}\Gamma^\beta_{\alpha\nu} - \Gamma^\alpha_{\mu\beta}\delta\Gamma^\beta_{\alpha\nu}. \tag{4.11}$$

The Ricci scalar $R = R^\mu_\mu = g^{\mu\beta}R_{\beta\mu}$ has then the perturbation

$$\delta R = \delta g^{\mu\alpha}R_{\alpha\mu} + \bar{g}^{\mu\alpha}\delta R_{\alpha\mu}, \tag{4.12}$$

where the background FLRW metric raises the indices of the perturbed Ricci tensor such that no terms quadratic in the perturbation appear. Following the same logic also for the perturbed Einstein tensor one has (Amendola & Tsujikawa, 2010a)

$$\delta G_{\mu\nu} = \delta R_{\mu\nu} - \frac{1}{2}\delta g_{\mu\nu}R - \frac{1}{2}\bar{g}_{\mu\nu}\delta R, \tag{4.13}$$

and the perturbed Einstein tensor with one index up is

$$\delta G^\mu_\nu = \delta g^{\mu\alpha}G_{\alpha\nu} + \bar{g}^{\mu\alpha}\delta G_{\alpha\nu}, \tag{4.14}$$

which shows that the perturbations pick up information about the background evolution as encoded in $G_{\alpha\nu}$. This is the reason why, for example in the cosmic microwave background, different structures emerge depending on the previous Hubble expansion of the Universe, which in turn depends on the present matter constituents via (1.22).

For the scalar perturbations of the FLRW metric we now obtain the non-diagonal perturbed Einstein tensor (Amendola & Tsujikawa, 2010a)

$$\begin{aligned}\delta G^0_0 &= \frac{2}{a^2}[3\mathcal{H}(\mathcal{H}\Psi - \Phi') + \nabla^2\Phi] \\ \delta G^0_i &= \frac{2}{a^2}(\Phi' - \mathcal{H}\Psi)_{|i} \\ \delta G^i_j &= \frac{2}{a^2}[(\mathcal{H}^2 + 2\mathcal{H}')\Psi + \mathcal{H}\Psi' - \Phi'' - 2\mathcal{H}\Phi']\delta^i_j \\ &\quad + \frac{1}{a^2}[\nabla^2(\Psi + \Phi)\delta^i_j - (\Psi + \Phi)^i_{|j}]\end{aligned} \tag{4.15}$$

where the Laplacian is meant to be covariant $\nabla^2 f = f_{;\mu}^{;\mu}$ and covariant derivatives with respect to some spatial coordinate i of the background metric are denoted by $f_{|i}$. Equation (4.15) describes the perturbations of the geometrical sector in the Einstein equations. In order to derive the full perturbed Einstein equations, this has to be combined with perturbations of the energy-momentum tensor.

For an observer with four-velocity u_μ, the energy-momentum tensor takes the form (1.9), where it is the question of how p, ρ and $\Sigma_{\mu\nu}$ and their perturbations arise. Later, in Sect. 4.2, we will describe a microscopic derivation of the perturbed energy-momentum tensor from a phase space distribution. Here, we only want to describe how the Einstein equations react to perturbations of the energy-momentum tensor once $T^{\mu\nu}$ is given. To have a clearer physical picture also here, we could e.g. imagine that the energy-momentum tensor is derived from a Lagrangian $\mathcal{L}$ by

$$T_{\mu\nu} = \frac{-2}{\sqrt{-g}} \frac{\delta(\sqrt{-g}\mathcal{L})}{\delta g^{\mu\nu}}, \tag{4.16}$$

which is for example of interest when studying dark energy candidates for which the Lagrangian is given, e.g. for quintessence field φ, where,

$$\mathcal{L}_\varphi = -\frac{1}{2} g^{\mu\nu} \partial_\mu \varphi \partial_\nu \varphi - V(\varphi). \tag{4.17}$$

The symmetry assumptions of the FLRW background then demand the background energy-momentum tensor to be diagonal

$$\bar{T}^{\mu\nu} = (\bar{\rho} + \bar{p}) \bar{u}^\mu \bar{u}^\nu + \bar{p} \bar{g}^{\mu\nu}, \tag{4.18}$$

where $\bar{u}^\mu$ is the velocity of an observer comoving with the cosmic expansion, meaning $\bar{u}^\mu \propto \delta_0^\mu$ such that in this reference frame, only the dilution with the passing of cosmic time appears.

In the perturbed FLRW metric (4.7), the velocity of an observer is of course perturbed by local gravitational potentials. These can introduce motions of matter with respect to the comoving coordinates x^i; these motions are usually called peculiar velocities and are given by

$$v^i = a \frac{\mathrm{d}x^i}{\mathrm{d}\eta} \tag{4.19}$$

where η is the conformal time and the $\mathrm{d}x^i$ are comoving as given in the metric (4.7). The velocity of an observer in this perturbed metric is then $u^\mu = \mathrm{d}x^\mu / \mathrm{d}s$ such that at first order (Amendola & Tsujikawa, 2010a)

$$u^0 = \frac{1 - \Psi}{a}, \qquad u^i = \frac{v^i}{a}. \tag{4.20}$$

Lowering the indices, we have $u_\mu = \bar{g}_{\mu\nu} u^\nu$ reducing to (Amendola & Tsujikawa, 2010a)

$$u_0 = -a(1 + \Psi), \qquad u_i = a v_i. \tag{4.21}$$

These perturbed velocities lead to a non-diagonal energy-momentum tensor in a universe that deviates from FLRW. Similary, stresses can appear, such that a stress-energy tensor Σ^i_j can be added to the energy-momentum tensor.
Projecting onto the velocity (4.20), we find the perturbations of the energy-momentum tensor to be (Amendola & Tsujikawa, 2010a; Mukhanov et al., 1992)

$$\begin{aligned} \delta T^0_0 &= -\delta\rho \\ \delta T^0_i &= -\delta T^i_0 = (\rho + p)v^i = (1+w)\rho v^i \\ \delta T^i_j &= \delta p \delta^i_j + \Sigma^i_j . \end{aligned} \tag{4.22}$$

The perturbations δp of the pressure can be related to a perturbation of the density for a barotropic fluid, where the equation of state fixes the pressure if the density is given and vice versa. Usually, this is parameterized by a sound speed $\delta p = c_s^2 \delta\rho$.
We now have all ingredients to set up the perturbed Einstein equations, but as usual, solving differential equations is easier in Fourier space where the derivatives are replaced by multiplications with the wave vectors; in our context, *comoving spatial* wave vectors are meant since the differential operators work on the expanding background and time derivatives are kept explicitely to give access to the dynamics. Understanding the following Φ, Ψ, ρ, θ and Σ^i_j as Fourier amplitudes, the perturbed Einstein equations read (Amendola & Tsujikawa, 2010a; Mukhanov et al., 1992)

$$k^2\Phi + 3\mathcal{H}(\Phi' - \mathcal{H}\Psi) = 4\pi G a^2 \delta\rho \tag{4.23}$$

$$k^2(\Phi' - \mathcal{H}\Psi) = -4\pi G a^2 (1+w)\rho\theta \tag{4.24}$$

$$\Phi'' + 2\mathcal{H}\Phi' - \mathcal{H}\Psi' - (\mathcal{H}^2 + 2\mathcal{H}')\Psi = -4\pi G a^2 c_s^2 \delta\rho \tag{4.25}$$

$$-k^2(\Phi + \Psi) = \frac{3a^2}{2}\left(\frac{k_i k_j}{|k|^2} - \frac{1}{3}\delta_{ij}\right)\Sigma^i_j . \tag{4.26}$$

This is the system that in principle needs to be solved for all matter components in the Universe. If $\Sigma^i_j = 0$, then (4.26) reduces to

$$\Phi = -\Psi . \tag{4.27}$$

Numerical codes such as CLASS (Blas et al., 2011; Audren et al., 2013a) however use extensively the fact that the structure formation must additionally satisfy the conservation equation $\nabla_\nu T^{\mu\nu} = 0$. Just like in hydrodynamics, the $\mu = 0$ component of this conservation equation will be a continuity equation and $\mu = i$ gives rise to the Euler equation. However, since $\nabla_\nu T^{\mu\nu} = 0$ is a covariant equation, the continuity equation and the Euler equation will also be covariant and we have (Amendola & Tsujikawa, 2010a)

$$\mu = 0 \implies \quad \Delta' + 3\mathcal{H}(c_s^2 - w)\Delta = -(1+w)(\theta + 3\Phi') \tag{4.28}$$

$$\mu = i \implies \quad \theta' + \left[\mathcal{H}(1-3w) + \tfrac{w'}{1+w}\right]\theta = -\nabla^2\left(\tfrac{c_s^2}{1+w}\Delta + \Psi\right), \tag{4.29}$$

where Δ is the density constrast (4.1) and $\theta = \boldsymbol{i k} \cdot \boldsymbol{v}$ is the velocity divergence. These equations describe the evolution of relativistic fluids on an expanding background. If a non-relativistic fluid is being studied, the equations simplify due to $w = 0$ and on a static background $\mathcal{H} = 0$. In CLASS, these equations are used as 'short-cuts' when solving the structure formation equations (4.25) and (4.26) which also allows a cross-check of whether the finite differencing schemes produce energy or matter due to numerical inaccuracies.

In this section, we have taken special care to describe the gravitational perturbations of the spacetime through $\delta G^{\mu\nu}$. The perturbed Einstein equations then also need a perturbed energy-momentum tensor. Here, we have used a macroscopic definition of the energy-momentum tensor (4.16). This may however be an insufficient description if particle interactions have to be modelled more carefully. In the following, we describe how perturbations $\delta T^{\mu\nu}$ can be calculated from kinetic theory.

4.2 Perturbations of the energy-momentum tensor from microscopic physics

Hydrodynamics of fluids is an effective theory, gained from kinetic theory by averaging over phase space cells. In an ideal fluid, the mean free path between collisions is $\lambda_{\text{free}} = 0$. However, we will later study the cosmic microwave background, which was released exactly because the mean free path of photons suddenly became $\lambda_{\text{free}} \gg 0$. Therefore, prior to the CMB release, there must have been a transition phase where the mean free path began to increase, and in this phase, the photons can no longer be described as a fluid. Let us therefore go back to kinetic theory at this point.

We imagine an ensemble of particles, whose trajectories through phase space are described by some setup-specific Hamiltonian equations of motion. We introduce coordinates x^μ and denote in this section by an overdot $\dot{x}^\mu$, their derivative with respect to proper time s. The canonically conjugate momenta p^μ are then

$$p_\mu = \frac{\partial \mathcal{L}(x^\mu, \dot{x}^\mu)}{\partial \dot{x}^\mu}, \tag{4.30}$$

where $\mathcal{L}$ is the Lagrange function of the ensemble.

The relativistic phase space is spanned by the seven (x^μ, p^i). The zero-component p^0 of the momentum is not an independent phase space coordinate, since the particles must obey

$$g_{\mu\nu}(x)p^\mu p^\nu = -m^2. \tag{4.31}$$

So if the particles live on a spacetime $\mathcal{M}$, and we introduce coordinates x and the tangent space $T\mathcal{M}$ of the spacetime, then the constraint (4.31) restricts the particles to the seven-dimensional subspace (Durrer, 2008)

$$P_m = \{(x, p) \in T\mathcal{M} | g_{\mu\nu}(x)p^\mu p^\nu = -m^2\}, \tag{4.32}$$

which is called the mass shell. Particles of different masses m are confined to different mass shells P_m. If $T_x\mathcal{M}$ is the tangent space at a fixed x, the fibre of the mass-bundle is then (Durrer, 2008)

$$P_m(x) = \{p \in T_x\mathcal{M} | g_{\mu\nu}(x)p^\mu p^\nu = -m^2\}. \tag{4.33}$$

The evolution of the ensemble can be described by the evolution of a phase space distribution function $f(x^\mu, p^i)$, which shall describe the probability with which certain position- and momentum states in phase space are occupied at a given time. The invariant measure $\mathrm{d}\mu$ of probability in phase space can be constructed from the invariant measure $|g|\mathrm{d}^4x\mathrm{d}^4p$ by multiplying it with the delta-function $2\delta_D(g_{\mu\nu}p^\mu p^\nu + m^2)$ in order to restrict it to the mass shell (Durrer, 2008). The factor 2 will disappear by simplifying the delta function via

$$\delta_D[h(x)] = \sum_i \frac{\delta_D(x_i)}{|h'(x_i)|} \quad \text{with roots} \quad h(x_i) = 0. \tag{4.34}$$

Expanding the mass shell constraint as

$$\delta_D[g_{\mu\nu}p^\mu p^\nu + m^2] = \delta_D[p_0^2 + 2p_0(g_{0i}p^i) + g_{ij}p^ip^j + m^2], \tag{4.35}$$

and defining that $P_0(x, \boldsymbol{p})$ shall be the zero component of the four vector p for which the argument of the delta function vanishes, the delta function reduces to

$$\delta_D[g_{\mu\nu}p^\mu p^\nu + m^2] = \frac{\delta_D[P_0(x,\boldsymbol{p})]}{2|P_0(x,\boldsymbol{p})|}. \tag{4.36}$$

The invariant measure on the seven-dimensional mass shell is then

$$\begin{aligned} d\mu &= |g|2\delta_D(g_{\mu\nu}p^\mu p^\nu + m^2)\mathrm{d}^4x\mathrm{d}^4p \\ &= |g|\frac{\delta_D[P_0(x,\boldsymbol{p})]}{|P_0(x,\boldsymbol{p})|}\mathrm{d}^4x\mathrm{d}^4p \\ &= \frac{\sqrt{|g|}\mathrm{d}^4x\ \sqrt{|g|}\mathrm{d}^3p}{|P_0(x,\boldsymbol{p})|}. \end{aligned} \tag{4.37}$$

In the last line, the transition to the seven-dimensional mass shell by integrating over p^0 has been made.

For Einstein's field equations, we need an energy-momentum tensor. On the microscopic level, this is the second moment of the phase space distribution function (Durrer, 2008)

$$T^{\mu\nu}(x) = \int_{P_m(x)} \frac{\sqrt{|g(x)|}}{|P_0(x,\boldsymbol{p})|} p^\mu p^\nu f(x,\boldsymbol{p})\mathrm{d}^3p. \tag{4.38}$$

Consequently, perturbations of the energy-momentum tensor can be calculated from a perturbed phase space distribution function and then be plugged into the perturbed Einstein equations (4.25). We therefore have to study the evolution of a perturbed phase space distribution function.

4.3 Evolution of the phase space distribution function

The Liouville equation states that the total time derivative of the phase space distribution function has to vanish; in a relativistic setting, the proper time s is meant and we have

$$\frac{\mathrm{d}}{\mathrm{d}s}f(x,p) = 0. \tag{4.39}$$

As the system evolves in time, the probability with which different phase space cells $\mathrm{d}x^\mu \mathrm{d}p^i$ are occupied, changes. Including also the implicit time dependences, this expands to (Durrer, 2008)

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}s} f &= \frac{\partial x^\mu}{\partial s} \frac{\partial f}{\partial x^\mu} + \frac{\partial p^i}{\partial s} \frac{\partial f}{\partial p^i} \\ &= \left[\dot{x}^\mu \partial_\mu + \dot{p}^i \frac{\partial}{\partial p^i} \right] f = 0 \end{aligned} \tag{4.40}$$

where the derivative with respect to p^0 is left out, since we operate on the mass shell. The external forces, under which the phase space distribution function develops, enter the Liouville equation through the derivatives of the conjugate momenta $\dot{\boldsymbol{p}}$. However, General Relativity describes gravity not by a force, but by the curvature of spacetime. In the generally relativistic setting, we therefore have to relate the conjugate momenta to the spacetime curvature, which is achieved as follows: given some metric $g_{\mu\nu}$ that solves the Einstein equations $G^{\mu\nu} = 8\pi G T^{\mu\nu}$, the Lagrange function of a freely falling particle in this spacetime is

$$\mathcal{L}(x, \dot{x}) = \frac{m}{2} g_{\mu\nu}(x) \dot{x}^\mu \dot{x}^\nu, \tag{4.41}$$

where m is the mass of a particle, if it has one, or in case of a massless particle, some arbitrary constant that does not affect the equations of motion.
The Euler-Lagrange equations are

$$\frac{\mathrm{d}}{\mathrm{d}s} \frac{\partial \mathcal{L}}{\partial \dot{x}^\lambda} - \frac{\partial \mathcal{L}}{\partial x^\lambda} = 0. \tag{4.42}$$

For the Lagrange function (4.41), the Euler-Lagrange equations yield (Straumann, 2013)

$$\ddot{x}^\alpha + \frac{1}{2} g^{\alpha\lambda} (2 g_{\lambda\nu,\mu} - g_{\mu\nu,\lambda}) \dot{x}^\mu \dot{x}^\nu = 0 \tag{4.43}$$

The prefactor of $\dot{x}^\mu \dot{x}^\nu$ is an oddly written Christoffel symbol that can be identified if the summation over all μ and ν is carried out explicitly. Consequently, the Euler-Lagrange equation of motion for the freely falling particle is (Straumann, 2013)

$$\ddot{x}^\alpha + \Gamma^\alpha_{\mu\nu} \dot{x}^\mu \dot{x}^\nu = 0 \tag{4.44}$$

which is the geodesic equation. From the definition of the conjugate momenta (4.30), we find for the freely falling particle

$$p^\mu = m \dot{x}^\mu \tag{4.45}$$

and if we finally insert this into the geodesic equation, we find the derivative of the conjugate momentum to be (Durrer, 2008)

$$m \dot{p}^\mu = -\Gamma^\mu_{\nu\alpha} p^\alpha p^\nu. \tag{4.46}$$

We can now insert this $\dot{p}^\mu$ into the Liouville equation in order to include gravity (Durrer, 2008)

$$\left[p^\mu \partial_\mu - \Gamma^i_{\mu\nu} p^\mu p^\nu \frac{\partial}{\partial p^i} \right] f = 0 \tag{4.47}$$

where we have also replaced $m\dot{x}^\mu$ by p^μ in the first term.

The Liouville equation describes particles for which collisions can be neglected. It can be specialized to an FLRW-universe by plugging in the Christoffel symbols of the FLRW-metric and by demanding that the phase space distribution function must respect the symmetry assumptions of FLRW, i.e. f shall only depend on the absolute momentum, but not its direction. The Liouville equation for the FLRW-metric then shows that the phase space distribution function of photons adapts to expansion only by redshifting the momenta of the particles under consideration. The reason is, that if photons were in thermal equilibrium in the early Universe, then their phase space distribution function was the black-body spectrum

$$f = \frac{1}{\exp(ap/T) - 1}, \tag{4.48}$$

where p is the absolute value of the comoving momentum, and hence ap is the physical photon energy. A chemical potential has been omitted as supported by observations. For neutrinos, the distribution function is a Fermi-Dirac distribution,

$$f = \frac{1}{\exp(ap/T) + 1}. \tag{4.49}$$

For both distributions, the expansion redshifts the momentum, which can be absorbed into a redefined temperature

$$T(a) = \frac{a_{\mathrm{ini}}}{a} T_{\mathrm{ini}}, \tag{4.50}$$

such that the black-body spectrum stays a black-body spectrum and the Fermi-Dirac distribution stays a Fermi-Dirac distribution, only the temperature changes. After decoupling, when the particles under consideration cease to be tightly interacting, this ‘temperature’ gets detached from its thermodynamical definition in the previous equilibrium state, but stays the parameter that uniquely defines the distributions (Durrer, 2008). If collisions which might bring a distribution out of thermal equilibrium during a decoupling process are neglected, and only the evolution in a homogeneously expanding universe shall be studied, then it is sufficient to parameterize the evolution of the phase space density as $f = f(v)$, where $v = ap$ is the redshift corrected momentum.

If particles have a non-vanishing cross section and collide, the exchange of momentum during a collision, as well as the divertion of the particle propagation directions, will additionally affect the evolution of the distribution function. This can be taken into account by adding a collision term to the Liouville equation. Under the assumption that only 2-body collisions are relevant, and that there is no correlation in the position of particles, the Liouville equation can be replaced by the Boltzmann equation (Durrer, 2008; Bartelmann, 2013)

$$\frac{\mathrm{d}}{\mathrm{d}s} f = C[f] = \frac{\mathrm{d}f_+(x, \boldsymbol{p})}{\mathrm{d}t} - \frac{\mathrm{d}f_-(x, \boldsymbol{p})}{\mathrm{d}t} \tag{4.51}$$

where $\mathrm{d}f_+/\mathrm{d}t$ denotes the particles being scattered into the phase space $f_+(x, \boldsymbol{p})$ per unit time, and $f_-(x, \boldsymbol{p})$ denotes the depopulation of phase space cells due to collisions. The exact details of the collision term $C[f]$ depend on the nature of the particles and anisotropic differential collisional cross sections are common. For example for photons, the main scattering process is Thomson scattering off electrons, which is highly anisotropic and also polarization dependent.

In the context of structure formation, we are interested in how the phase space distribution function develops small perturbations, being exposed to a slightly perturbed background metric. If we denote background quantities by an overbar and make the ansatz $f = \bar{f}(v) + F(x^\mu, v, \theta, \phi)$ and $\Gamma^\alpha_{\mu\nu} = \bar{\Gamma}^\alpha_{\mu\nu} + \delta\Gamma^\alpha_{\mu\nu}$, where $\delta\Gamma^\alpha_{\mu\nu}$ results from perturbations of the metric, then to first order the perturbed Liouville equation is (Durrer, 2008)

$$p^\mu \partial_\mu f - \bar{\Gamma}^i_{\mu\nu} p^\mu p^\nu \frac{\partial f}{\partial p^i} - \delta\Gamma^i_{\mu\nu} p^\mu p^\nu \frac{\partial \bar{f}}{\partial p^i} = 0. \tag{4.52}$$

The solution of the perturbed Liouville and Boltzmann equations quickly becomes technical, especially because the background-perturbation split is gauge dependent, and because the details of the solution differ between the individual particle species, depending on whether they are massive or massless, and depending on their collision term. The usual CMB reviews therefore omit these technicalities completely. However, the upcoming chapter 7 depends on one of these details, so let us describe the general procedure of how the perturbed Liouville equation (4.52) or the Boltzmann equation can be solved.

The motivation to solve (4.52) is two-fold: On the one hand, the CMB is composed of photons, and the perturbations in the phase space density of the photons are responsible for the observed CMB temperature fluctuations. We therefore need to derive a theoretical prediction of the CMB power spectrum from (4.52) together with a collision term for Thomson scattering.

On the other hand, all particle species couple to the metric, and therefore, the perturbations of all other particle species will influence the perturbations of the photons via the Einstein equations. These depend on the perturbed energy-momentum tensor and we know that for each particle species, the energy-momentum tensor is on the microscopic level related to the particles phase space density by the moment (4.38). Therefore, (4.52) has to be solved for all other particle species as well, in order to get the time evolution of their energy-momentum tensor and couple them gravitationally to the photons via Einstein's equations.

Let us therefore describe the solution of the perturbed Liouville or Boltzmann equation.

4.4 Solving the perturbed Liouville or Boltzmann equation

Following Durrer (2008), the following equations are derived in the longitudinal gauge, however, the sign convention of the Bardeen potentials Φ and Ψ from Amendola & Tsujikawa (2010a) will be used. As in the upcoming chapter 7, we will restrict ourselves to scalar perturbations only.

The FLRW-metric with scalar perturbations in this convention is (Amendola & Tsujikawa, 2010a; Durrer, 2008)

$$\mathrm{d}s^2 = -a^2(\eta)(1+2\Psi)\mathrm{d}\eta^2 + a^2(1+2\Phi)\gamma_{ij}\mathrm{d}x^i\mathrm{d}x^j, \tag{4.53}$$

where γ_{ij} is the metric of the spatial hypersurfaces. From this metric the perturbed Christoffel symbols can be derived by (4.9).

For the metric (4.53), the perturbed Christoffel symbols are functions of the gravitational potentials Ψ and Φ and therefore depend on the perturbations of all particle species such that they are gravitationally coupled. Additionally, they also depend on the Hubble expansion of the background solution.

We had dissected the phase space density into $f = \bar{f}(v) + F(x^\mu, v, \theta, \phi)$, where $v = ap$, is the redshift-corrected momentum. If the phase space density is uniquely defined through a temperature, as is the case for the Bose-Einstein distribution (4.48) and the Fermi-Dirac distribution (4.49), then the fluctuations $F(x^\mu, v, \theta, \phi)$ can be translated into fluctuations of this temperature.

To stick to the notation of Durrer (2008), we transform the fluctuations onto

$$\mathcal{F} = F + \Phi v(\mathrm{d}\bar{f}/\mathrm{d}v), \tag{4.54}$$

where the correction $\Phi v(\mathrm{d}\bar{f}/\mathrm{d}v)$ originates from the gauge-invariant formulation of perturbation variables in Durrer (2008), but else does not affect the physical interpretation of the CMB power spectrum since it only contributes a non-measureable monopole.

The perturbations $\mathcal{F}(x^\mu, v, \theta, \phi)$ still carry the full information about the momenta. For the CMB, we are however only interested in how spatial variations of $\mathcal{F}$ evolve with time η, and how they map onto the sky seen in direction $\boldsymbol{n}$. We therefore write x^μ as $(\eta, \boldsymbol{x})$ and drop the information on the momenta by integrating them out (Durrer, 2008)

$$\mathcal{M}(\eta, \boldsymbol{x}, \boldsymbol{n}) = \frac{\pi}{a^4\bar{\rho}} \int v^3 \mathcal{F} \mathrm{d}v, \tag{4.55}$$

where $\bar{\rho}$ is the mean density. For photons, where the temperature is the only variable in the distribution function, these perturbations are the temperature fluctuations

$$\mathcal{M}(\eta, \boldsymbol{x}, \boldsymbol{n}) = \frac{\delta T}{T}(\eta, \boldsymbol{x}, \boldsymbol{n}). \tag{4.56}$$

If there exists a collisional integral, it is also integrated over momentum

$$C[\mathcal{M}] = \frac{\pi}{a^4\bar{\rho}} \int v^3 \left(\frac{\mathrm{d}f_+}{\mathrm{d}\eta} - \frac{\mathrm{d}f_-}{\mathrm{d}\eta}\right) \mathrm{d}^3 v, \tag{4.57}$$

where we see from the prefactor $1/(a^4\bar{\rho})$ that the collisions become less frequent with progressing expansion.

The Liouville or Boltzmann equation of the full phase space distribution function f can now be translated into a differential equation for the reduced phase space function $\mathcal{M}$, which we do not need to cite here since it is still the Liouville or Boltzmann equation. Differential equations are easier to solve in Fourier space, since a Fourier transformation turns differential operators into algebraic ones. Therefore, one conventionally dissects $\mathcal{M}(\eta, \boldsymbol{x}, \boldsymbol{n})$ into spatial Fourier modes, keeping the time dependence explicitely

$$\mathcal{M}(\eta, \boldsymbol{k}, \boldsymbol{n}) = \int \mathrm{d}^3 x e^{-i\boldsymbol{k}\boldsymbol{x}} \mathcal{M}(\eta, \boldsymbol{x}, \boldsymbol{n}). \tag{4.58}$$

Likewise $C[\mathcal{M}]$ is also dissected into Fourier modes. The transformation to Fourier space has the further advantage that in the linear growth studied here, different Fourier modes evolve independently of each other. This gives the possibility for approximations which is heavily used in the Boltzmann solver CLASS (Blas et al., 2011) because sub- and superhorizon modes evolve differently.
The CMB is however observed on the sky, and it is therefore easiest to not dissect it into plane waves, but into spherical harmonics instead

$$\frac{\Delta T}{T}(\boldsymbol{x}_0, \boldsymbol{n}, t_0) = \sum_{l,m} a_{lm} Y_{lm}(\boldsymbol{n}), \tag{4.59}$$

and the usual coefficients C_ℓ are

$$\langle a_{lm}, a^*_{lm} \rangle = C_\ell \delta_{\ell\ell'} \delta_{mm'}. \tag{4.60}$$

Both, the Fourier modes $e^{-i\boldsymbol{k}\boldsymbol{x}}$ and the spherical harmonics Y_{lm} are a complete set of orthonormal basis functions. A convenient bridge between these two sets is provided by a third orthonormal set, the Legendre polynomials P_ℓ, which can be analytically expanded in spherical harmonics and Fourier modes alike

$$\begin{aligned} P_\ell(\boldsymbol{n} \cdot \boldsymbol{n}') &= \frac{4\pi}{2\ell+1} \sum_{m=-\ell}^{\ell} Y_{\ell m}(\boldsymbol{n}) Y_{\ell m}(\boldsymbol{n}'), \\ e^{i\boldsymbol{k}\cdot\boldsymbol{n}t} &= \sum_{l=0}^{\infty} (2\ell+1) i^\ell j_\ell(kt) P_\ell(\mu). \end{aligned} \tag{4.61}$$

In order to approach the CMB spectrum C_ℓ, we therefore expand the Fourier modes $\mathcal{M}(\eta, \boldsymbol{k}, \boldsymbol{n})$ in Legendre polynomials. We call the direction cosine to the line of sight $\mu = \boldsymbol{k} \cdot \boldsymbol{n}/|\boldsymbol{k}|$, such that $\mathcal{M}(\eta, \boldsymbol{k}, \boldsymbol{n}) = \mathcal{M}(\eta, \boldsymbol{k}, \mu)$ and since $\mu \in [-1, 1]$, the expansion in Legendre polynomials is now possible, leading to (Durrer, 2008)

$$\mathcal{M}(t, \boldsymbol{k}, \mu) = \sum (2\ell+1)(-i)^\ell \mathcal{M}_\ell(t, \boldsymbol{k}) P_\ell(\mu). \tag{4.62}$$

The expansion coefficients are then

$$\mathcal{M}_\ell(\eta, \boldsymbol{k}) = \frac{i}{2} \int_{-1}^{1} \mathrm{d}\mu \mathcal{M}(\eta, \boldsymbol{k}, \mu) P_\ell(\mu). \tag{4.63}$$

As the perturbations grow, the $\mathcal{M}_\ell$ will develop a power spectrum M_ℓ, defined by (Durrer, 2008)

$$\langle \mathcal{M}_\ell(\eta, \boldsymbol{k}), \mathcal{M}_{\ell'}(\eta, \boldsymbol{k}^*) \rangle = \mathsf{M}_\ell(\eta, k)(2\pi)^3 \delta^3(\boldsymbol{k} - \boldsymbol{k}')\delta_{\ell\ell'} \tag{4.64}$$

where $\delta^3(\boldsymbol{k} - \boldsymbol{k}')\delta_{\ell\ell'}$ reflects that the different Fourier modes evolve independently, and the different ℓ-modes are also orthonormal.

The power spectrum of the CMB of course mixes the different k-modes, so the k-dependent power spectrum M_ℓ relates to the CMB power spectrum on the sphere by the integration (Durrer, 2008)

$$C_\ell = \frac{2}{\pi} \int \mathrm{d}k k^2 \mathsf{M}_\ell(k). \tag{4.65}$$

So we see that in order to calculate the CMB power spectrum, we need to solve the evolution equation for the different $\mathcal{M}_\ell$. For photons, this is given by the hierarchy

$$\dot{\mathcal{M}}_\ell + k\frac{\ell+1}{2\ell+1}\mathcal{M}_{\ell+1} - k\frac{\ell}{2\ell+1}\mathcal{M}_{\ell-1} + \dot{\kappa}\mathcal{M}_\ell = \delta_{\ell 0}\dot{\kappa}\mathcal{M}_0 + \frac{1}{3}\delta_{\ell 1}[-k(\Phi+\Psi) + \dot{\kappa}V^{(b)}] + \dot{\kappa}\frac{1}{10}\delta_{\ell 2}\mathcal{M}_2, \tag{4.66}$$

where $V^{(b)}$ is the velocity of the baryons, and $\kappa(t_1, t_2) = \int_{t_1}^{t2} a\sigma_T n_e \mathrm{d}t$ is the optical depth that includes the Thomson cross section $\sigma_T = 6.65 \cdot 10^{-25}\mathrm{cm}^2$ for collisions between photons and electrons. The factor n_e is the number density of free electrons that decreases with expansion and then drops rapidly during recombination.

For neutrinos, there is no optical depth and if we treat them as massless relativistic particles, the hierarchy is (Durrer, 2008)

$$\dot{\mathcal{N}}_\ell + \frac{k}{2\ell+1}\left[(\ell+1)\mathcal{N}_{\ell+1} - \ell\mathcal{N}_{\ell-1}\right] = \frac{-k}{3}\delta_{\ell 1}(\Phi + \Psi), \tag{4.67}$$

where $\mathcal{N}_\ell$ is the neutrino equivalent of $\mathcal{M}_\ell$. Equations (4.66) and (4.67) are nothing less than the Boltzmann equation and the Liouville equation, only heavily disguised by the consecutive transformation into Fourier space and then into P_ℓ space. The physics is however the same as discussed in Sect. 4.3 and especially, since it is a differential equation, initial conditions have to be specified, for which we use an initial power spectrum as generated by inflation.

Notice that (4.67) is a hierarchy: different multipoles $\mathcal{N}_\ell$ couple to their neighbouring multipoles $\mathcal{N}_{\ell+1}$ and $\mathcal{N}_{\ell-1}$, if the latter exists. This coupling between the multipoles describes the free-streaming of the neutrinos, which are treated as collisionless in the standard approach. In chapter 7, we will modify this hierarchy in order to replace freestreaming neutrinos by a relativistic fluid.

This section has served a twofold purpose: first, we saw how perturbations of an energy-momentum tensor can be derived from kinetic theory if the hydrodynamical, or macroscopic description of a matter species via its Lagrangian is insufficient. We have also seen, why perturbations of the phase space distribution are observable as temperature fluctuations in the cosmic microwave background (4.60).

Next, we derive the Newtonian limit of structure growth, as this applies to dark matter on subhorizon scales in the late Universe, which we will need for studying weak lensing.

4.5 Newtonian limit and fluid approximation

Above, we have studied cosmological structure formation in a fully relativistic setting. Clearly, on scales where spacetime curvature can be neglected, this will reduce to a Newtonian limit. But why should it be possible to treat even structures of cosmological extent still within the Newtonian approximation, as done in simulations like EAGLE (Schaye et al., 2015) or the Millenium simulations (Boylan-Kolchin et al., 2009; Springel, 2000)? The reason is that although General Relativity provides a consistent treatment and a theoretically sound framework, the equations for the Universe's expansion and its forming structures can also partly be constructed from Newtonian gravity and hydrodynamics – if one is willing to add in by hand correction terms to account for relativistic effects and if one ignores some problematic theoretical interpretations of the Newtonian treatment as described below.

Let us review how the Friedmann equations (1.13) without a cosmological constant can be 'derived' from Newtonian gravity, with the following toy model.

The observed Hubble flow tells us that galaxies escape radially from us today. We can then assume spherical symmetry and at first order model the Universe as a homogeneously filled, expanding and self-gravitating sphere.

The equation of motion of a test mass m (e.g. a galaxy) on the surface of a sphere of radius r and mass M is then

$$m\ddot{r} = -\frac{GMm}{r^2} = -\frac{4\pi G}{3}\rho r \tag{4.68}$$

where ρ is the mass density. We now add in by hand that in a relativistic treatment, the pressure of a fluid contributes an energy density $\rho_p = 3p$, leading to

$$\ddot{r} = -\frac{GMm}{r^2} = -\frac{4\pi G}{3}(\rho + 3p)r. \tag{4.69}$$

Furthermore, let us assume that such an isotropically expanding sphere will be subject to adiabatic cooling. We had seen in Sect. 1.2 that the two Friedmann equations automatically contain adiabatic cooling, and that the second Friedmann equation can be replaced by the first law of thermodynamics. So please note that assuming adiabatic cooling here is equivalent to already taking one of the two Friedmann equations as given. Denoting the internal energy as $U = \frac{4}{3}\pi r^3 \rho$, where the pressure is left out, and the volume as V, we have

$$\mathrm{d}U + p\mathrm{d}V = 0 \Rightarrow 3r^2\dot{r}\rho + r^3\dot{\rho} = 3pr^2\dot{r}. \tag{4.70}$$

Combining (4.69) and (4.70) we find (Longair, 2008)

$$2\dot{r}\ddot{r} = \frac{8\pi G}{3}(2r\dot{r}\rho + r^2\dot{\rho}), \tag{4.71}$$

where $2\dot{r}\ddot{r}$ is the time derivative of $\dot{r}^2$ and the term in brackets on the right hand side is the time derivative of ρr^2. Hence, the time integral of (4.71) is

$$\left(\frac{\dot{r}}{r}\right)^2 = \frac{8\pi G}{3}\rho + \frac{\mathrm{const}}{r^2} \tag{4.72}$$

which is the first Friedmann equation with an integration constant which cannot be used as a cosmological constant, because it has the wrong scaling with r. Instead, this constant represents the spatial curvature.

We therefore see that a hand-waving construction of the Friedmann equations is possible. However, the result of this derivation should be interpreted exactly within the setup of its derivation: If we assume the initial conditions were set by some explosion-like mechanism, then it describes the peculiar motion of particles on a static background which evolve under Newtonian gravity. This is not a valid derivation of the *cosmological* Friedmann equations: since we measure escape velocities larger than the speed of light, we know that the escape results from the background expansion, instead of a superluminal galaxy motion with respect to a static background. Furthermore, the explosion-like mechanism that set the initial conditions, should not be confused with the Big Bang: the mentioned explosion would take place at a given point in time and at a given special spatial point on a static background. In contrast, the cosmological Big Bang corresponds to the point in *time* where all physical distances are compressed to a singularity because the scale factor goes as $a \to 0$. Consequently, the cosmological Big Bang is not a point *in* space, it is a contraction *of* space, and in the cosmological standard model, the Big Bang has happened today at *each* point about 14 billion years ago.

Notwithstanding these inconsistencies of interpretation, we have nonetheless an equation that describes the escape of galaxies. Similarly, a semi-Newtonian derivation of cosmological structure formation exists. The late-time growth of cosmic structures is often modelled with numerical simulations which treat the cosmic matter fields as perfect fluids that develop instabilities under Newtonian gravity. In chapters 5 and 9, we will use a growth function that results from such a treatment, so a presentation of it seems adequate.

We begin by taking the Newtonian limit of gravity and replace the tensorial Einstein equations by the Poisson equation

$$\nabla^2\phi = 4\pi G\rho. \tag{4.73}$$

Next, we derive the non-relativistic limit of hydrodynamics from the kinetic theory presented in Sect. 4.2. Working with a fixed Newtonian background, the determinant factor $\sqrt{-g}$ in the energy-momentum tensor (4.38) is identical to unity. The transition from microscopic theory to fluid dynamics then starts by taking moments of the phase space distribution function f with respect to p^μ in order to lose the microscopic information on the p^μ. The first moment is the current (Bartelmann, 2013; Durrer, 2008)

$$J^\mu(t,\boldsymbol{x}) = \int \frac{\mathrm{d}^3p}{P_0} f(t,\boldsymbol{x},\boldsymbol{p})p^\mu = (1,\boldsymbol{v}) \tag{4.74}$$

where $\boldsymbol{v}$ is the velocity and the energy is $E = P_0$. The second moment of f is the energy-momentum tensor on the static background

$$T^{\mu\nu} = \int \frac{\mathrm{d}^3p}{P_0} f(t,\boldsymbol{x},\boldsymbol{p})p^\mu p^\nu. \tag{4.75}$$

Before, we had seen the conservation law $\nabla_\nu T^{\mu\nu} = 0$ in a general relativistic setting. If we had not restricted ourselves to FLRW-Universes, we would also have encountered the conservation law for the current, $\nabla_\nu J^\nu = 0$. Depriving the covariant derivative ∇ of a flexible spacetime, it reduces to $\nabla_\nu \to \partial_\nu$ such that the (special)-relativistic fluid equations read

$$\partial_\nu J^\nu = 0, \quad \text{and} \quad \partial_\nu T^{\mu\nu} = 0. \tag{4.76}$$

Both of these conservation laws arise on the microscopic level as momentum integrals of the Boltzmann equation: upon introduction of a global coordinate system with $x = (ct, x^1, x^2, x^3)$, and the dot being the derivative with respect to the coordinate time t, the Boltzmann equation (4.51) simplifies to (Bartelmann, 2013)

$$\partial_t f + \dot{\boldsymbol{x}}\frac{\partial f}{\partial \boldsymbol{x}} + \dot{\boldsymbol{p}}\frac{\partial f}{\partial \boldsymbol{p}} = C[f]. \tag{4.77}$$

The phase space distribution function f in (4.77) still carries the entire information on the momenta – however, hydrodynamics is only interested in what a fluid does when observed on a macroscopic level, i.e. how its velocity field, its density and its pressure evolve, and consequently, the microscopic information on the individual momenta can be marginalized over.

Let us begin by taking momenta over the collision term

$$\int C[f]\mathrm{d}^3p, \quad \int C[f]p^\mu \mathrm{d}^3p, \tag{4.78}$$

where $C[f]$ describes the population and depopulation of phase space cells due to scattering. If the collisions are elastic, then they will conserve energy and momentum and particle number. Hence, if only enough particles are averaged over, then it is true that there is a *redistribution* of energy, momentum and particle number but the net change is zero due to the conservation. The integrals (4.78) are then zero. The defining property of a perfect fluid is now that there exist microscopic fluid elements, which compose the entire fluid, and which consist internally of enough particles, such that averaging over them the collision integrals (4.78) vanish. The collision term of collisionless dark matter vanishes by definition.

The equations of hydrodynamics now follow by taking moments of the Boltzmann equation. The mean is given by

$$\int \partial_t f \ \mathrm{d}^3p + \int \dot{\boldsymbol{x}}\frac{\partial f}{\partial \boldsymbol{x}} \ \mathrm{d}^3p + \int \dot{\boldsymbol{p}}\frac{\partial f}{\partial \boldsymbol{p}} \ \mathrm{d}^3p = \int C[f] \ \mathrm{d}^3p, \tag{4.79}$$

which is $\partial_\nu J^\nu$. In the non-relativistic limit this is the usual continuity equation (Bartelmann, 2013; Longair, 2008)

$$\partial_t \rho + \nabla \cdot (\rho \boldsymbol{v}) = 0. \tag{4.80}$$

From the first moment

$$\int \partial_t f \ p^\mu \mathrm{d}^3p + \int \dot{\boldsymbol{x}}\frac{\partial f}{\partial \boldsymbol{x}} \ p^\mu \mathrm{d}^3p + \int \dot{\boldsymbol{p}}\frac{\partial f}{\partial \boldsymbol{p}} \ p^\mu \mathrm{d}^3p = \int C[f] \ p^\mu \mathrm{d}^3p, \tag{4.81}$$

follows $\partial_\nu T^{\mu\nu} = 0$, and if we take the non-relativistic limit, this produces the Euler equation for momentum conservation (Bartelmann, 2013; Longair, 2008)

$$\partial_t \boldsymbol{v} + (\boldsymbol{v} \cdot \boldsymbol{\nabla})\boldsymbol{v} = -\frac{1}{\rho}\nabla p - \nabla\phi. \tag{4.82}$$

The above demonstrates how generally relativistic structure formation with a microscopic description of the energy-momentum tensor, can be reduced to a hydrodynamical description under Newtonian gravity. The cosmic structures will then be fluid instabilities that grow under gravity.

4.6 Newtonian structure formation

For cosmological structure formation equations(4.73,4.80,4.82) are applied to the cosmic matter densities. One then introduces the usual splitting

$$\begin{aligned} \rho &= \bar{\rho} + \delta\rho, \\ p &= \bar{p} + \delta p, \\ \phi &= \bar{\phi} + \delta\phi, \\ \boldsymbol{v} &= \bar{\boldsymbol{v}} + \delta\boldsymbol{v}, \end{aligned} \tag{4.83}$$

where background quantities are denoted by an overbar. These backgroud quantities evolve with the FLRW solution and due to homogeneity $\nabla\bar{\rho} = \nabla\bar{p} = 0$ on hypersurfaces of constant cosmic time. The first order perturbations of the continuity equation (4.80) is then (Longair, 2008)

$$\dot{\Delta} = -\nabla \cdot \delta\boldsymbol{v}, \tag{4.84}$$

where the dot denotes the derivative with respect to cosmic time. The perturbation of the Euler equation is (Longair, 2008)

$$\delta\dot{\boldsymbol{v}} + (\delta\boldsymbol{v} \cdot \nabla)\bar{\boldsymbol{v}} = -\frac{1}{\bar{\rho}}\nabla\delta p - \nabla\delta\phi. \tag{4.85}$$

Finally, the linearly perturbed Poisson equation is (Longair, 2008)

$$\nabla^2\delta\phi = 4\pi G\delta\rho, \tag{4.86}$$

which is a very nice result, since in Newtonian gravity, the superposition principle holds exactly.

Let us consider the meaning of these equations. In cosmology, we actually have to work on an expanding background, but the non-relativistic hydrodynamics employed here uses a static and absolute background. This means all gradients ∇ are taken with respect to what would in cosmology be called physical coordinates. We can then introduce the expanding Universe by hand, if we transform to comoving coordinates. We set

$$\boldsymbol{x} = a(t)\boldsymbol{r}, \tag{4.87}$$

where $\boldsymbol{r}$ is in comoving coordinates and $\boldsymbol{x}$ is in physical coordinates. The gradients ∇ in (4.84,4.85,4.86) are then $\nabla = \nabla_x = \frac{1}{a}\nabla_r$. Additionally, the velocity $\boldsymbol{v}$ splits into

$$\begin{aligned} \boldsymbol{v} = \dot{\boldsymbol{x}} &= \dot{a}\boldsymbol{r} + a\dot{\boldsymbol{r}} \\ &= \bar{\boldsymbol{v}} + \delta\boldsymbol{v}, \end{aligned} \tag{4.88}$$

meaning that the background flow $\bar{\boldsymbol{v}} = \dot{a}\boldsymbol{r}$ is given by the Hubble flow at a fixed comoving coordinate $\boldsymbol{r}$, and that there exists an additional migration along comoving coordinates $\delta\boldsymbol{v} = a\dot{r}$, which are the peculiar velocities, and we denote them as $\delta\boldsymbol{v} = a\boldsymbol{u}$.
This reparameterization allows us to include the expansion of the Universe by hand and the comoving Euler equation then follows to be

$$\dot{\boldsymbol{u}} + 2H\boldsymbol{u} = -\frac{1}{\bar{\rho}a^2}\nabla_r\delta p - \frac{1}{a^2}\nabla_r\delta\phi. \tag{4.89}$$

For adiabatic perturbations, we have $c_{\mathrm{s}}^2 = \partial p/\partial\rho$ at constant entropy. The continuity equation (4.84) and the comoving Euler equation (4.85) can then be combined to the usual wave equation by taking another time derivative of the continuity equation and the comoving divergence of the Euler equation, and then eliminating $\nabla_r \cdot \boldsymbol{u}'$ that appears in both equations. This elimination produces a differential equation for the evolution of the density contrast Δ (Longair, 2008)

$$\ddot{\Delta} + 2H\dot{\Delta} = \frac{c_s^2}{a^2}\nabla_r^2\Delta + 4\pi G\bar{\rho}\Delta. \tag{4.90}$$

This result may appear suprisingly simple in comparison to the generally relativistic description of structure formation. The simplicity results however mainly from the fact that using a scalar Newtonian potential is indeed a massive simplification in comparison to a fully tensorial description of gravity as in GR.
The wave equation (4.90) is solved by Fourier modes

$$\Delta \propto \exp(i[\boldsymbol{k}_c \cdot \boldsymbol{r} - \omega t]) \tag{4.91}$$

such that the comoving divergence ∇_r^2 is replaced by a multiplication with the comoving wave vector k_c, producing (Longair, 2008)

$$\ddot{\Delta} + 2H\dot{\Delta} = \Delta(4\pi G\bar{\rho} - k^2c_s^2), \tag{4.92}$$

where $\boldsymbol{k}_c = a\boldsymbol{k}$. We therefore see that the growth of the density contrast Δ is determined by the equation of an oscillator, damped by the cosmic expansion and driven by the gravitational attraction $4\pi G\bar{\rho}\Delta$ of already overdense regions. The term $-k^2c_s^2\Delta$ originates from the pressure; obviously it is able to counteract the gravitational attraction, meaning that fluids with a pressure do not enhance primordial density contrasts as easily as pressureless fluids do.

In fact, the growth of cosmic structures can be understood better, if we reduce (4.92) to an undamped harmonic oscillator by excluding the expansion term $2H\Delta'$, corresponding to a static Universe. The eigenfrequency of the oscillator that solves the ansatz (4.91) is then

$$\omega^2 = c_s^2 k^2 - 4\pi G\bar{\rho}. \tag{4.93}$$

Obviously, if ω is real, the ansatz (4.91) describes an oscillation in time, where the sign of ω determines whether the phase rotates clockwise or counterclockwise in the complex plane. This happens if $\omega^2 > 0$, i.e. $c_s^2 k^2 > 4\pi G\bar{\rho}$ meaning that the pressure of the fluid provides a sufficiently strong restoring force in order to support the fluid against gravitational collapse. This is e.g. true for the baryon-photon plasma in the pre-CMB Universe and these oscillations have imprinted themselves onto the CMB power spectrum.

On the other hand, if the pressure is smaller than the gravitational attraction, then ω is imaginary, which cancels the imaginary prefactor of the phase in (4.91). We can then have $\Delta \propto \exp \pm|A|t$, where $|A|$ is the growth rate. The positive solution then corresponds to an exponential growth of the structures.

Of course, such an exponential growth never occurs in the real Universe: first of all, as soon as the growth becomes non-linear, equation (4.92) is not an adequate description anymore. Secondly, our Universe expands, such that the eigenfequency (4.93) evolves with time as $\bar{\rho}$ reduces with expansion. Additionally, the Hubble-drag term $2H\Delta'$ damps the growth of structures, leading to an algebraic growth instead of exponential instability. For example, during matter domination, the growth of dark matter follows $D_+ \propto a$.

In the extreme case of a vanishing pressure, ω^2 will always be negative. This is true for pressureless cold dark matter, which is the reason for why it forms structures much more efficiently than baryons whose structure formation is hindered by a pressure.

Obviosuly, whether a fluid that possesses some pressure oscillates or not, also depends on the scale $\lambda \propto k^{-1}$, indicating that collapse can set in on the largest scales, even if the smallest scales are still supported by pressure[1].

The function $D_+(t)$ that is the growing-mode solution of the second-order differential equation (4.92) is usually called *growth function*. As the expansion of the Universe influences the structure growth by providing the Hubble drag, and as the expansion depends via Einstein's equation on the matter content of the Universe, we see directly that the shape of D_+ carries information about the cosmological parameters. Mapping the growth of structures therefore provides a second window onto the cosmological parameters, apart from geometrical probes such as Supernovae of Type Ia and BAOs.

Let us finish this section with a warning: The result of such a Newtonian calculation cannot be sensitive to horizon- and retardation effects. This is especially important for structure

---

[1] The wavelength that separates growth and oscillation is called the Jeans wavelength and is known e.g. from the context of star formation. There, a similar analysis is carried out that describes why a molecular cloud can collapse under its self-gravity if it loses enough thermal energy through radiation. This collapse is facilitated if the molecular gas has a high metallicity, which enable a strong radiative cooling and thereby reduce the pressure very efficiently until collapse sets in. In cosmology, such radiation processes are not the dominant trigger of collapse. Instead, the scaling with k, the Hubble drag and the evolution of ω with the expansion are the dominating effects.

growth in the early Universe, where the standard model predicts that with the progress of time, modes re-enter the horizon that had previously been driven out of the horizon by inflation. Such effects can be correctly accounted for by solving the full perturbed Einstein equations. A typical procedure in modern cosmology is to run a Boltzmann solver like CLASS or CAMB in order to model the early structure formation, and to let this code calculate a *transfer function* $\mathcal{T}(k, z)$ that translates an initial inflationary power spectrum into a power spectrum at lower redshift of approximately $z \approx 100$ or $z \approx 10$. Until $z \approx 100$, the horizon has increased to such an extent that modelling structure formation on the observable scales with the Newtonian approximation seems to be sufficiently precise.

Chapter 5

Cosmic shear

We have seen that the growth of cosmic structures depends on the cosmological parameters. In the late Universe, the dominant clustering matter component is dark matter, whose power spectrum can be calculated from the formalism presented in chapter 4. However, as dark matter does at best only emit a tiny amount of radiation when decaying into standard model particles, the dark matter fields of the Universe cannot be directly observed.

Due to their gravity, overdensities of dark matter however attract baryons. It can therefore be expected that the baryons in our Universe trace the dark matter web. If the baryonic gas cools enough when falling into the dark matter potential wells, galaxy formation can set in. Therefore, observing galaxies allows to infer the distribution of the underlying dark matter. This inference is however complicated by the issue of biases: what triggers galaxy formation is not yet understood in all detail, but it is clear that the galaxies form stochastically along the overdensities of the dark matter web, tracing it in a biased way. The hope is that this bias will be better understood in the future, such that it can be subtracted from data or be included into the total uncertainty on the data.

We can therefore assert that mapping galaxies will allow to infer the dark matter power spectrum. One way to achieve this, is via galaxy clustering analysis, e.g. (Gil-Marín et al., 2015) for which the issue of biases is quite important. Another mechanism that is less sensitive with respect to biasing is weak lensing on cosmic scales, also called cosmic shear. The following sections shall describe this mechanism in more detail.

5.1 Deflection in gravitational potentials

Light of distant galaxies needs to propagate through the inhomogeneous gravitational potentials between its origin and the Earth. As light propagates along geodesics and as geodesics adapt to gravitational fields, the inhomogeneity of the gravitational potentials deflects lightrays and imprints itself onto the images of the distant galaxies, leading to magnification or demagnification and distortion. This effect is known as gravitational lensing. The observed deflection will build up continuously during the photon propagation, such that weak lensing is an integrated effect. Typically, a magnification or distortion of about one percent builds up on cosmic scales (Kilbinger, 2015).

In this section, it will be handy to denote the speed of light in vacuum explicitely by c. We will use a coordinate time t, and a spatial vector $\boldsymbol{x}$. We will additionally assume there is no gravitational slip, such that the Bardeen potentials are related by $\Psi = \Phi$.[1] If a lens is confined to a spatial extent that is much smaller than the Universe's curvature radius, then its metric can be approximated by the post-Minkowskian line element

$$\mathrm{d}s^2 = -c^2\left(1 + \frac{2\Phi}{c^2}\right)\mathrm{d}t^2 + \left(1 - \frac{2\Phi}{c^2}\right)\mathrm{d}\boldsymbol{x}^2. \tag{5.1}$$

Light rays that pass this lens, run on the light cone of the metric and have $\mathrm{d}s = 0$. If we additionally demand that $\Phi/c^2 \ll 1$ and Taylor expand to first order, then we find from $\mathrm{d}s = 0$ that there exists an effective speed of light that is smaller than the vacuum value (Narayan & Bartelmann, 1996)

$$c' = \left|\frac{\mathrm{d}\boldsymbol{x}}{\mathrm{d}t}\right| = c\left(1 + \frac{2\Phi}{c^2}\right)\ . \tag{5.2}$$

Of course, the light does not actually slow down when passing the potential Φ. Instead, clocks go slower in the vicinity of masses – but since we used a *coordinate* time t, which is adequate for a distant observer, we are insensitive to this and instead see in our coordinates an effectively lower speed of light.

We can then associate an effective index of refraction n with the lens, defined by $c' = c/n$, leading to (Narayan & Bartelmann, 1996)

$$n = \frac{c}{c'} = 1 - \frac{2\Phi}{c^2}, \tag{5.3}$$

to lowest order in Φ/c^2.

As demonstrated by Perlick (2000), Fermat's principle carries over to the problem studied here, such that the light rays will follow paths that extremize the total propagation time $\tau = \int c'\mathrm{d}t$. The extremization $\delta\tau = 0$ then leads to the deflection angle

$$\tilde{\boldsymbol{\alpha}} = -\frac{2}{c^2}\int \nabla_\perp \Phi \mathrm{d}\lambda, \tag{5.4}$$

where $\mathrm{d}\lambda$ is the line-of-sight integration and $\nabla_\perp\Phi$ that component of the gravitational potential's gradient that is perpendicular to the light path, meaning $\nabla_\perp$ is a two-dimensional operator. Usually, it is sufficient to employe Born's approximation and integrate along a radial coordinate w_L instead of the perturbed light path λ.

For suitably aligned, and sufficiently compact and strong lenss such as galaxies or galaxy clusters, the extremization $\delta\tau$ can often lead to multiple solutions, depending on the

---

[1] If $\Phi \neq \Psi$ is explicitely kept in the lensing formalism, and if additionally to lensing the motion of solid particles is studied, then the presence of gravitational slip can be measured. This is for example possible on lensing galaxy clusters.

symmetries and the structure of the potential Φ and the direction of the initially incident light ray. Consequently, multiply lensd sources are known e.g. the famous Einstein cross (Dai et al., 2003). For lensing by the cosmic large-scale structure, such multiple images do not appear, and one speaks of weak lensing.

The deflection angle of a light ray can only be determined when the initial position of the source is known, which is typically not true. What is measurable however, is the *differential* deflection of neighboring light rays. We interpret the deflection angle as a map on the sphere, introducing spherical angles $\boldsymbol{\theta} = (\theta, \phi)$, and investigate variations of the deflection with $\boldsymbol{\theta}$. Note, what this implies: as $\tilde{\boldsymbol{\alpha}}$ is a gradient field by (5.4), it cannot have a curl, $\nabla \times \tilde{\boldsymbol{\alpha}} = 0$. The weak lensing observables derived in the next section can therefore only have E-modes, no B-modes.

5.2 Convergence maps

We have seen that the angular diameter distance D_{ang} is defined such that if an object of physical size d subtends the angle $\Delta\theta$, then $D_{\mathrm{ang}}\Delta\theta = d$. The weak lensing formalism is easier to derive in *comoving* angular diameter distances w, defined by

$$D_{\mathrm{ang}}(z) = a(z)w(z). \tag{5.5}$$

We can then transform the gradient $\nabla_\perp$ perpendicular to the line of sight to a gradient with respect to the spherical angles (θ, ϕ) by

$$\nabla_\perp = \frac{1}{w_L}\nabla_\theta, \tag{5.6}$$

where w_L is the comoving angular diameter distance to the lens L that was responsible for the deflection.

From the redshifts z, we can infer distances $w(z)$ once we specify a cosmological model. This distance information for source and lens is then incorporated into the so-called reduced deflection angle (Bartelmann & Schneider, 2001)

$$\boldsymbol{\alpha}(\boldsymbol{\theta}, w_s) = \frac{2}{c^2}\int_0^{w_s} \mathrm{d}w_L \frac{w_s - w_L}{w_s} \nabla_\perp \Phi[w_L\boldsymbol{\theta}, w_L], \tag{5.7}$$

where w_s is the distance to the source, w_L are the distances to the lenss that lie between observer and source and that are integrated over, and $\boldsymbol{\theta}$ is the position on the sky where the source is seen after it was deflected by $\boldsymbol{\alpha}(\boldsymbol{\theta})$. Below, in the *lens equation*, we will see that $\boldsymbol{\alpha}$ is the deflection angle on the observer's sky.

An isotropic focusing or defocusing of lightrays is described by the divergence of $\boldsymbol{\alpha}$ and we define

$$\nabla_{\boldsymbol{\theta}} \cdot \boldsymbol{\alpha}(\boldsymbol{\theta}, w) =: 2\kappa(\boldsymbol{\theta}, w), \tag{5.8}$$

where κ is the convergence. From (5.7) it follows that a source at distance w_s has (Bartelmann & Schneider, 2001)

$$\kappa(\boldsymbol{\theta}, w_s) = \frac{1}{c^2}\int_0^{w_s} \mathrm{d}w_L \frac{(w_s - w_L)w_L}{w_s}\left(\frac{\partial^2}{\partial x_1^2} + \frac{\partial^2}{\partial x_2^2}\right)\Phi[\boldsymbol{\theta}w_L, w_L], \tag{5.9}$$

where the derivative with respect to the angles on the sphere was converted into derivatives with respect to the anguar separation $\boldsymbol{x}$ via the small angle approximation $\boldsymbol{x}(\boldsymbol{\theta}, w') = w'\boldsymbol{\theta}$. The focusing or defocusing of lightrays leads to a magnification or demagnification of single sources and can already be of direct cosmological interest, if standardizable candles, such as type-Ia supernovae, are investigated. Also, very faint sources can make it above the limiting flux of a telescope if they are sufficiently magnified, which can allow to study extremely highly redshifted sources.

For cosmology, the interest in sources is however usually limited to their property of letting themselves be lensd – what is being studied intensively, are the statistical properties of the lensing matter fields in order to probe the growth of structure.

From (5.9), we see that the 2-dimensional Laplacian $(\partial_1^2 + \partial_2^2)\Phi$ of the gravitational potential appears. If this were the 3-dimensional Laplacian, the Newtonian Poisson equation would allow to substitute it by the density contrast, whose growth can be derived as in Sect. 4.

In order to achieve this, we use that a given lens is more or less effective, depending on its position between source and observer: from (5.7), we see that if the lens is positioned very close to the source, $w_L \approx w_s$, then it can barely deflect the source. Likewise, if the lens is positioned very close to us $w_L \approx 0$, then the integral will again be approximately zero and no deflection occurs. Consequently, only such lenss are effective that lie at a distance $0 < w_{\text{eff}} < w_s$, i.e. well in between us and the sources. More quantitatively, the efficiency of a lens is modulated by the distance prefactor $(w_s - w_L)w_L$ in (5.9). This is a parabola which peaks at $w_L/2$, such that lenss that lie half way between observer and source are geometrically most efficient (Kilbinger, 2015).

If we imagine such an effective lens provides the gravitational potential $\Phi^{w_{\text{eff}}}(w)$, then we would like to add its second derivative along the line of sight, $\partial_z^2\Phi^{w_{\text{eff}}}$, to (5.9). We can safely do so, because this second derivative will be

$$\int_0^{w_s} \frac{\partial^2\Phi^{w_{\text{eff}}}(w)}{\partial w^2}\mathrm{d}w = \left.\frac{\partial\Phi^{w_{\text{eff}}}(w)}{\partial w}\right|_{w=0}^{w_s} \approx 0, \tag{5.10}$$

meaning it averages out along the line of sight (Bartelmann & Schneider, 2001). The reason is that $\partial_w\Phi^{w_{\text{eff}}}$ is a gravitational force, and the most effective lenss will be so far away, that their gravitational attraction on source and observer can safely be neglected.

We can consequently add the term $\partial^2\Phi/\partial x_3^2$ to the two-dimensional Laplacian in (5.9), and then exploit that the 3-dimensional Laplacian $\nabla^2\Phi$ is related to the cosmic density contrast via (1.4)

$$\nabla^2\Phi = \frac{3H_0^2\Omega_0}{2a}\Delta. \tag{5.11}$$

The convergence field $\kappa(\boldsymbol{\theta}, w_s)$ is then given by (Bartelmann & Schneider, 2001)

$$\kappa(\boldsymbol{\theta}, w_s) = \frac{3H_0^2\Omega_0}{2c^2}\int_0^{w_s} \mathrm{d}w_L \frac{(w_s - w_L)w_L}{w_s}\frac{\Delta[w_L\boldsymbol{\theta}, w_L]}{a(w_L)}. \tag{5.12}$$

where the term $\Delta[w_L\boldsymbol{\theta}, w_L]$ describes the growing structures as a function of their distance w_L to the observer. The integral then includes the inhomogeneous matter fields in front of the source plane which perturb the photon propagation.

Up until here, we have worked with sources at a single, fixed distance w_s. In order to include all galaxies in the Universe that can act as sources, the single-redshift convergence $\kappa(\boldsymbol{\theta}, w_s)$ needs to be averaged over the source redshift distribution.

If $p(z)\mathrm{d}z$ is the number of galaxies in the redshift slice $\mathrm{d}z$, then we can convert this into a source-distance distribution $G(w)$ by

$$G(w_s)\mathrm{d}w_s = S(z)p(z)\mathrm{d}z. \tag{5.13}$$

Here, $S(z)$ is meant to be a selection function. It can either refer to all galaxies that made it into the catalogue at all, or it can be a more complicated function. For example, it could tag either blue or red galaxies. It could also fold in a redshift-dependent selection function in order to split the galaxies into distinct populations which then allows weak lensing tomography[2]. Widely spread selection functions for tomography are the Heaviside function, leading to rectangular redshift bins, or

$$S(z) = \frac{1}{\sqrt{2\pi}\Delta z}\exp\left(-\frac{1}{2}\frac{(z - z_c)^2}{\Delta z^2}\right), \tag{5.14}$$

leading to Gaussian redshift bins. The convergence weighted by the source distribution is then

$$\bar{\kappa}(\boldsymbol{\theta}) = \int_0^{w_{\mathrm{lim}}} \mathrm{d}w_s G(w_s)\kappa(\boldsymbol{\theta}, w_s), \tag{5.15}$$

where w_{lim} is the maximum distance, given by the limiting flux of the observing instrument. Conventionally, the above average is often also expressed in terms a weighting function $\bar{W}(w)$ which is defined by

$$\bar{W}(w) = \int_w^{w_{\mathrm{lim}}} \mathrm{d}w_s G(w_s)\frac{w_s - w}{w_s}, \tag{5.16}$$

which combines the geometrical efficiency with the number of lensd sources, and is therefore also known as lensing efficiency function (Kilbinger, 2015). In terms of (5.16), the convergence as weighted with the source distribution is then (Bartelmann & Schneider, 2001)

$$\bar{\kappa}_G(\boldsymbol{\theta}) = \frac{3H_0^2\Omega_0}{2c^2}\int_0^{w_{\mathrm{lim}}} \mathrm{d}w_L \bar{W}(w_L)\frac{w_L}{a(w_L)}\Delta[w_L\boldsymbol{\theta}, w_L], \tag{5.17}$$

where we kept explicit that the averaged convergence depends on the chosen function G.

---

[2] Tomography can only be achieved by partitioning the source distribution $p(z)\mathrm{d}z$. All matter in front of the sources will automatically act as lenss, such that they cannot be meaningfully binned.

5.3 Convergence, shear and the lensing potential

The reduced deflection angle $\boldsymbol{\alpha}$ from equation (5.7) is in general a complicated integral, as it depends on the mass distribution of the lens. Its interpretation is however straightforward: it is the amount by which the detected image of a source shifts due to intervening masses, although the source actually lies at position $\boldsymbol{\beta}$, i.e.

$$\boldsymbol{\beta} = \boldsymbol{\theta} - \boldsymbol{\alpha}(\boldsymbol{\theta}). \tag{5.18}$$

This equation is known as the *lens equation.* Mapping between images and sources can then be described locally by the Jacobian matrix

$$\mathcal{A}(\boldsymbol{\theta}) = \frac{\partial \boldsymbol{\beta}}{\partial \boldsymbol{\theta}}, \tag{5.19}$$

which includes a derivative of $\boldsymbol{\alpha}$ and therefore describes in linear order the distortions of a light bundle originating from $\boldsymbol{\beta}$ and being mapped to $\boldsymbol{\theta}$. In components, the Jacobian is

$$\mathcal{A}(\boldsymbol{\theta}) = \left(\delta_{ij} - \frac{\partial \alpha_i}{\partial \theta_j}\right). \tag{5.20}$$

The choice of the coordinates $\boldsymbol{\theta}$ is arbitrary. Conventionally, one first introduces a lensing potential ψ defined by

$$\boldsymbol{\alpha}(\boldsymbol{\theta}) = \nabla_\theta \psi(\boldsymbol{\theta}). \tag{5.21}$$

This can be achieved by employing $\nabla_\perp = w_L^{-1} \nabla_\theta$ in equation (5.7) and then using that since ∇_θ is perpendicular to the line of sight, it can be pulled out of the line of sight integral such that the lensing potential follows to be (Kilbinger, 2015)

$$\psi(\boldsymbol{\theta}, w_s) = \frac{2}{c^2} \int_0^{w_s} \mathrm{d}w_L \frac{w_s - w_L}{w_s w_L} \Phi[w_L \boldsymbol{\theta}, w_L]. \tag{5.22}$$

This lensing potential is the combination of all effects that determine the deflection of lightrays that start from the source distance w_s and are seen by the observer in direction $\boldsymbol{\theta}$. As it contains the integral over the gravitational potential Φ of the lensing mass distribution, we see already here that for lensing by the cosmic large-scale structure, ψ will be a stochastic quantity and it will inherit the power spectrum from the density fluctuations in the Universe. In terms of the scalar lensing potential, the local mapping (5.20) reads (Bartelmann & Schneider, 2001)

$$\mathcal{A}(\boldsymbol{\theta}) = \left(\delta_{ij} - \frac{\partial^2 \psi(\boldsymbol{\theta})}{\partial \theta_j \partial \theta_i}\right). \tag{5.23}$$

This mapping can be split up into a scalar contribution, by taking the trace, that then describes the isotropic focusing or defocusing of the mapping. The remaining anisotropic contribution after subtracting the trace from $\mathcal{A}$ is then a shear matrix.[3]

---

[3] This is completely analogous to calculating the pressure of an energy-momentum tensor from the trace, and defining that the trace-less part is then the anisotropic stress tensor which leads to shearing forces.

The trace of $\mathcal{A}$ is

$$\mathrm{Tr}\,(\mathcal{A}) = (2 - \nabla^2\psi) = 2(1-\kappa), \tag{5.24}$$

where

$$\kappa(\boldsymbol{\theta}) = \frac{1}{2}\nabla^2\psi(\boldsymbol{\theta}), \tag{5.25}$$

is the convergence due to (5.21). We decompose the Jacobi matrix as $\mathcal{A} = (1-\kappa)\mathcal{I} - \Gamma$, and consequently the shear matrix is

$$\Gamma = \begin{pmatrix} \gamma_1 & \gamma_2 \\ \gamma_2 & -\gamma_1 \end{pmatrix}, \tag{5.26}$$

where the shear *components* are

$$\begin{aligned} \gamma_1 &= \frac{1}{2}\left(\frac{\partial^2\psi}{\partial\theta_1^2} - \frac{\partial^2\psi}{\partial\theta_2^2}\right), \\ \gamma_2 &= \frac{\partial^2\psi}{\partial\theta_1\partial\theta_2}. \end{aligned} \tag{5.27}$$

The Jacobi matrix is then

$$\mathcal{A} = (1-\kappa)\mathcal{I} - \Gamma = \begin{pmatrix} 1-\kappa-\gamma_1 & -\gamma_2 \\ -\gamma_2 & 1-\kappa+\gamma_1 \end{pmatrix} \tag{5.28}$$

From the two tensorial shear components γ_1 and γ_2 *the* shear

$$\gamma = \sqrt{\gamma_1^2 + \gamma_2^2}, \tag{5.29}$$

can be constructed by taking the square root of the determinant of $-\Gamma$. While the two shear components γ_1 and γ_2 still depend on how the coordinates $\boldsymbol{\theta}$ were chosen, the scalar shear γ is independent of the coordinate choice and can therefore be regarded as a genuine attribute of the lensd image.

Let us understand what shear and convergence do. We imagine a circular source $\delta\boldsymbol{\beta}$ represented in polar coordinates

$$\delta\boldsymbol{\beta} = |\beta| \begin{pmatrix} \sin(a) \\ \cos(a) \end{pmatrix}, \tag{5.30}$$

such that a is the phase and $|\beta|$ is the original radius of the source. The mapping of sources $\delta\boldsymbol{\beta}$ onto images $\delta\boldsymbol{\theta}$ is given by the inverse Jacobian

$$\delta\boldsymbol{\theta} = \mathcal{A}^{-1}\delta\boldsymbol{\beta}. \tag{5.31}$$

For $\gamma = 0$, we have

$$\delta\boldsymbol{\theta} = \begin{pmatrix} \frac{1}{1-\kappa} & 0 \\ 0 & \frac{1}{1-\kappa} \end{pmatrix} \delta\boldsymbol{\beta}, \tag{5.32}$$

such that the mapping is a multiple of the identity matrix and the convergence only focuses or defocuses isotropically but it does not distort. If there exists a shear γ, then the mapping can best be worked out in the eigenbasis of $\mathcal{A}^{-1}$

$$\delta\boldsymbol{\theta} = \begin{pmatrix} \lambda_+ & 0 \\ 0 & \lambda_- \end{pmatrix} \delta\boldsymbol{\beta}, \tag{5.33}$$

where the eigenvalues are $\lambda_\pm = 1/(1-\kappa\mp\gamma)$. We then see, that if we choose two orthogonal axes in the intrinsically circular source, e.g.

$$\boldsymbol{\beta}_1 = \begin{pmatrix} |\beta| \\ 0 \end{pmatrix}, \quad \boldsymbol{\beta}_2 = \begin{pmatrix} 0 \\ |\beta| \end{pmatrix} \tag{5.34}$$

then these will be originally of equal length, but then be mapped onto

$$\boldsymbol{\theta}_1 = \lambda_+\boldsymbol{\beta}_1, \quad \boldsymbol{\theta}_2 = \lambda_-\boldsymbol{\beta}_2, \tag{5.35}$$

which means that the circular source gets distorted to an ellipse with semi-major axis $a = \lambda_+|\beta|$ and semi-minor axis $b = \lambda_-|\beta|$. We therefore conclude that the shear γ of weak lensing is responsible for distorting circular sources to ellipses, and the convergence increases or decreases images.

The shear is the usual quantity that is being assessed in weak lensing surveys. The ellipticity ϵ is defined in terms of the semi-major and semi-minor axes as

$$\epsilon = \frac{a-b}{a+b}, \tag{5.36}$$

and consequently, the ellipticity of an image gives away the shear at first order (Bartelmann & Schneider, 2001; Kilbinger, 2015)

$$\epsilon = \frac{\lambda_+ - \lambda_-}{\lambda_+ + \lambda_-} = \frac{\gamma}{1-\kappa} \approx \gamma \tag{5.37}$$

where the last approximation holds for $\kappa \ll 1$ which is typically true for weak lensing. We therefore see that if we have a resolved image of an intrinsically circular source, then we can measure the ellipticity that it acquired by lensing from (5.36) and compare this to the theoretical prediction (5.37). If the convergence κ is known, the last approximation does not need to be carried out; $\gamma/(1-\kappa)$ is in any case the directly observable quantity and also known as *reduced shear* (Kilbinger, 2015).

An observational complication arises because weak lensing catalogues typically use galaxies as lensd sources, and galaxies are not intrinsically circular. Instead, their shapes vary significantly, which leads to an intrinsic source ellipticity ϵ_s at lowest order. This has to be added to the ellipticity induced by lensing (Kilbinger, 2015)

$$\epsilon = \frac{\gamma}{1-\kappa} + \epsilon_s \approx \gamma + \epsilon_s. \tag{5.38}$$

Observations indicate that the intrinsic ellipcities of galaxies follow rather well a Gaussian distribution with standard deviation $\sigma_\epsilon \approx 0.2$ (Kilbinger, 2015). One can then suppress the intrinsic ellipticities by averaging over N galaxies, such that

$$\sigma_{\epsilon;N} = \frac{\sigma_\epsilon}{\sqrt{N}}. \tag{5.39}$$

So for $N \to \infty$, the shot noise $\sigma_{\epsilon;N}$ of the intrinsic ellipticities vanishes, as the differently oriented sources average out to a circular source. The signal-to-noise ratio is then

$$\frac{S}{N} = \frac{\gamma}{\sigma_\epsilon/\sqrt{N}}. \tag{5.40}$$

As typical values for the shear γ are in the regime of a few percent, one needs to average over a few hundred galaxies, in order to achieve a signal to noise above unity. As typically galaxies within a solid angle element are averaged over, this will limit the angular resolution of the catalogue way below the point spread function of the observing instrument.

The measurement of ellipticities is in reality complicated further due to pixelisation and the point spread function of the observing instrument, as well as potentially irregular shapes of the original galaxies. Often, spatial moments of the galaxy surface brightness yield a useful measure of ellipticities, see Bartelmann & Schneider (2001) for an overview.

Although weak lensing has so far been observed in infrared and visible light, for example with CFHTLenS (Heymans et al., 2013; Kilbinger et al., 2013) and DES (Chang et al., 2015), the future of cosmology will likely see a routinely executed weak lensing analysis in radio frequencies as well. The reason is the resolution of a galaxy, such that its ellipticity can be measured: infrared and visible light is emitted by the stellar population of a galaxy. This is only the innermost core of the galaxy, leading to a small picture. Galaxies are however surrounded by a vast halo of neutral hydrogen. The 21-cm line of this gas can then be detected with radio observatories like the upcoming SKA (Takahashi et al., 2015).

5.4 Weak lensing power spectra

As a final step, we now need to take into account that the cosmic density field is a random field, and cosmological models can predict its statistics, but not its single realization on the sky. Consequently, even though we measure just a shear or convergence map (5.17), which originates from a projected overdensity map $\Delta[w_L\boldsymbol{\theta}, w_L]$, only the statistics of this map can be predicted by theoretical models, and we hence have to derive the power spectrum of the shear or convergence map.

Let us imagine we were given a measured convergence map $\kappa(\boldsymbol{\theta})$. Then we could directly measure the correlation function

$$\mathcal{C}_\kappa(|\boldsymbol{\phi}|) = \langle \kappa(\boldsymbol{\theta})\kappa(\boldsymbol{\theta}+\boldsymbol{\phi})\rangle, \tag{5.41}$$

where it was used that the sky-projected correlation function in an FLRW-universe needs to be radially symmetric and where the angular brackets actually denote an ensemble average, which is in practice however replaced by a spatial average. The angular power spectrum of this measured convergence is then

$$
\begin{aligned}
\langle \hat{\kappa}(\boldsymbol{\ell})\hat{\kappa}^*(\boldsymbol{\ell}')\rangle &= \int \mathrm{d}^2\theta e^{i\boldsymbol{\theta}\boldsymbol{\ell}} \int \mathrm{d}^2\phi e^{-i(\boldsymbol{\theta}+\boldsymbol{\phi})\boldsymbol{\ell}'} \mathcal{C}_\kappa(|\boldsymbol{\phi}|) \\
&= (2\pi)^2 \delta_D(\boldsymbol{\ell}-\boldsymbol{\ell}') \int \mathrm{d}^2\phi e^{-i\boldsymbol{\phi}\boldsymbol{\ell}'} \mathcal{C}_\kappa(|\boldsymbol{\phi}|) \\
&= (2\pi)^2 \delta_D(\boldsymbol{\ell}-\boldsymbol{\ell}') C_\kappa(\ell),
\end{aligned}
\tag{5.42}
$$

where $C_\kappa(\ell)$ is the angular power spectrum of the convergence. Likewise an angular power spectrum $\langle \hat{\gamma}(\boldsymbol{\ell})\hat{\gamma}^*(\boldsymbol{\ell}')\rangle$ of the shear exists.

The measured power spectrum (5.41) then needs to be related to a theoretical prediction. Such a theoretical power spectrum can be calculated in the following way.

We need to include that lensing is an effect integrated along the line of sight, such that we see a 2-dimensional projection of an inherently 3-dimensional power spectrum. The integration over a redshift can be achieved with Limber's approximation that converts a 3-dimensional power spectrum into a 2-dimensional one. Limber's approximation states that if a 2-dimensional stochastic quantity $X(\boldsymbol{\theta})$ is the projection of a 3-dimensional stochastic quantity $Y(w\boldsymbol{\theta}, w)$ meaning

$$
X(\boldsymbol{\theta}) = \int_0^{w_{\mathrm{lim}}} Q(w)\, Y(w\boldsymbol{\theta}, w)\mathrm{d}w, \tag{5.43}
$$

where $Q(w)$ is a radial weighting function, then the 2d angular power spectrum $P_x(\ell)$ of X is related in the flat-sky approximation to the 3-dimensional Cartesian power spectrum $P_y(k)$ of Y as (Bartelmann & Schneider, 2001)

$$
P_x(\ell) = \int_0^{w_{\mathrm{lim}}} Q^2(w) P_y(k = \ell/w)\mathrm{d}w. \tag{5.44}
$$

For a constant weight function Q, this equation is rather intuitive, as the marginal of a Gaussian is still a Gaussian and hence the powerspectra of the two Gaussians are related. The Limber approximation (5.44) generalizes this result by demanding that correlations of Y exist only on scales that are smaller than typical scales on which the weight function Q varies.

As the convergence (5.17) is a twodimensional projection of the density contrast Δ, with the weight function

$$
Q(w) = \frac{3H_0^2\Omega_0}{2c^2}\bar{W}(w)\frac{w}{a}, \tag{5.45}
$$

it follows that the angular power spectrum $P_\kappa(\ell)$ of the convergence is related to $P_\Delta(k)$ via (Bartelmann & Schneider, 2001)

$$
C_\kappa(\ell, w_b) = \frac{9H_0^4\Omega_0^2}{4c^4} \int_0^{w_b} \mathrm{d}w \frac{\bar{W}(w)}{a^2(w)} P_\Delta\left(\ell/w, w\right). \tag{5.46}
$$

This integrates up the power spectrum $P_\Delta\left(\ell/w, w\right)$ which evolves with w, due to the finite lookback time and can then be related to measured power spectra.
As a last step, we need to clarify why convergence spectra are typically calculated, but shear spectra are measured. The reason is, that the ellipticity is easier to measure, whereas calculating with the scalar convergence instead of with the spin-2 shear is easier. Conveniently, the shear and the convergence have the same power spectrum. To show this, we Fourier transform the lensing potential

$$\hat{\psi}(\boldsymbol{\ell}) \propto \int \psi(\boldsymbol{\theta}) e^{i\boldsymbol{\theta}\cdot\boldsymbol{\ell}}. \tag{5.47}$$

From equations (5.25) and (5.27) we then find that the Fourier-transformed convergence and shear are (Bartelmann & Schneider, 2001)

$$\begin{aligned} \hat{\kappa} &= -\frac{1}{2}(l_1^2 + l_2^2)\hat{\psi}(\ell) \\ \hat{\gamma}_1 &= -\frac{1}{2}(\ell_1^2 - \ell_2^2)\hat{\psi}(\ell), \quad \hat{\gamma}_2 = -\ell_1\ell_2\hat{\psi}(\ell) \\ \hat{\gamma}^2 &= \hat{\gamma}_1^2 + \hat{\gamma}_2^2 = \left[\frac{1}{4}(\ell_1^2 - \ell_2^2)^2 + (\ell_1\ell_2)^2\right]\hat{\psi}^2(\ell). \end{aligned} \tag{5.48}$$

From the power spectrum 5.42, we know that we only need to care about $\boldsymbol{\ell} = \boldsymbol{\ell}'$. Hence proving that $\hat{\kappa}^2$ and $\hat{\gamma}^2$ are identical also proves that the power spectra are identical, and indeed

$$4|\hat{\gamma}|^2 = \left[(l_1^2 - l_2^2)^2 + 4l_1^2 l_2^2\right]|\hat{\psi}|^2 = (l_1^2 + l_2^2)^2|\hat{\psi}|^2 = 4|\hat{\kappa}|^2, \tag{5.49}$$

which completes the proof that the power spectra of shear and convergence are identical. This is not too surprising, as both, shear and convergence, inherit their statistics from the second-order derivatives statistics of the lensing potential.
Weak lensing is sensitive to cosmological parameters because it probes the conversion of redshifts into distances, see e.g. (5.16), and on the other hand the growth of cosmic structures by integrating over the evolving power spectrum of the gravitational potentials. In chapter 9, we will use this to forecast the constraints on cosmological parameters from a Euclid-like survey.

Chapter 6

Introduction to multivariate statistical methods

Being physicists, we assume that certain natural laws govern our Universe, and that these laws can be cast into a mathematical form. In other words, we construct a mathematical model M of the Universe. This model will depend on parameters which we want to measure, so $M = M(\boldsymbol{\theta})$, where $\boldsymbol{\theta} = (\theta_1, ..., \theta_n)$ are the n model parameters. The true values of these parameters are however inaccessible to us, due to measurement errors, degeneracies between the different θ_i, or limited data sets. We therefore have to estimate or infer such parameters. This is done within the framework of statistics. The better the data, and the better the statistical analysis, the more precisely can the parameters of a model and the model itself be inferred from cosmic data.

The first concession that statistics makes, is to admit that we have to deal with defined *estimators* instead of *true* quantities. An estimator is a random variable. A random variable x is drawn from a probability density function $f(x)$ and we write

$$x \sim f(x). \tag{6.1}$$

If $f(x)$ has a broad width, then x has a great uncertainty; if $f(x)$ is sharply peaked, then x does not scatter much.

A true quantity is not a random variable. If $g(x)$ is a function of a random variable x, then $g(x)$ is also a random variable. A random vector $\boldsymbol{x}$ is a multivariate random variable, and again drawn from a probability density function

$$\boldsymbol{x} \sim f(\boldsymbol{x}). \tag{6.2}$$

Measured data $\boldsymbol{X}$ are on the one hand produced by an underlying model that uses some true parameters $\underline{\boldsymbol{\theta}}$, but are on the other hand subject to random noise and therefore also random variables. Consequently, as functions of random variables return again random variables, we can only infer *estimators* $\boldsymbol{\theta}$ of the true parameters $\underline{\boldsymbol{\theta}}$, where the $\boldsymbol{\theta}$ are random.

A probability density function $f(x)$ is defined such that the probability P of finding x in the interval $[a, b]$ is $P(x) = \int_a^b f(x)\mathrm{d}x$. Probabilities P are defined to fullfil the Kolmogorov axioms

1. $P \geq 0$
2. $\sum_i P_i = 1$, especially $P(A) + P(\bar{A}) = 1$
3. $P(A, B|C) = P(A|C)P(B|A, C) = P(B|C)P(A|B, C)$,

where $P(A)$ is the probability that event A occurs, and $P(\bar{A})$ is the probability that it does not occur. $P(A|B)$ is the probability of A occuring, given that B has already occured. $P(A, B)$ is the probability that A and B occur simultaneously. For continous probabilities, the second rule implies $\int_{-\infty}^{\infty} f(x)\mathrm{d}x = 1$.
From the last rule follows directly Bayes' theorem

$$P(A|B, C) = \frac{P(A|C)P(B|A, C)}{P(B|C)}, \tag{6.3}$$

which shows that probabilities are not symmetric under exchange of the premises and the implications. This is of central importance if A is interpreted as a model $M(\boldsymbol{\theta})$ and B is a data set $\boldsymbol{X}$. We then have

$$P(M(\boldsymbol{\theta})|\boldsymbol{X}) = \frac{P(M(\boldsymbol{\theta}))P(\boldsymbol{X}|M(\boldsymbol{\theta}))}{P(\boldsymbol{X})}. \tag{6.4}$$

We will return to this equation below, after having clarified some further aspects of probability densities.
A probability density can be defined in terms of its moments, provided they exist. The n-th moment m_n of the multivariate random variable $\boldsymbol{x}$, that follows the distribution function $P(\boldsymbol{x})$ is (Mardia et al., 1979)

$$m_n = \int [\boldsymbol{x} \otimes \boldsymbol{x} \otimes \ldots \otimes \boldsymbol{x}] P(\boldsymbol{x}) \mathrm{d}^n x, \tag{6.5}$$

where $\otimes$ is the outer tensor product. Note, that we write m_n, although in case of a multivariate variable the zeroth-order moment is a scalar (unity), the first moment is a vector (the mean $\boldsymbol{\mu}$), the second moment a matrix, and the higher order moments tensors of respective rank.
Central higher order moments m_n^c follow after the mean has been subtracted from the data, leading to (Gregory, 2005)

$$m_n^c = \int [(\boldsymbol{x} - \boldsymbol{\mu}) \otimes (\boldsymbol{x} - \boldsymbol{\mu}) \otimes \ldots \otimes (\boldsymbol{x} - \boldsymbol{\mu})] P(\boldsymbol{x} - \boldsymbol{\mu}) \mathrm{d}^n x. \tag{6.6}$$

The second central moment is called the covariance matrix

$$\boldsymbol{C} = \int P(\boldsymbol{x}) \left[(\boldsymbol{x} - \boldsymbol{\mu}) (\boldsymbol{x} - \boldsymbol{\mu})^T \right] \mathrm{d}^n x, \tag{6.7}$$

and is a first estimate of the probability distribution's width. The higher order moments then correct this width. If – apart from the mean – only even moments exist, then the

probability distribution is symmetric about the mean. If also odd moments exist, the probability distribution is said to be skewed, typically possessing one long tail. Dealing with a skewed probability distribution can be quite counterintuitive at first, because the event that occurs on *average* is then not the event that occurs *most likely.*
Not all probability distributions have all moments. The only independent moments of a Gaussian distribution are its mean and its covariance. Once the mean and the covariance are given, all higher moments of a Gaussian are fixed, due to Wick's theorem. A Cauchy distribution does not have a mean or a covariance, or any other moments because the integrals (6.6) do not converge.

6.1 Moment-generating function

If the moments of a probability density function $P(\boldsymbol{x})$ exist, then the moment-generating function $m_x(t)$ is given by the Laplace transform of $P(\boldsymbol{x})$ (Gregory, 2005)

$$m_x(\boldsymbol{t}) = \langle e^{\boldsymbol{t}\cdot\boldsymbol{x}}\rangle = \int_{-\infty}^{+\infty} e^{\boldsymbol{t}\cdot\boldsymbol{x}} P(\boldsymbol{x}) \mathrm{d}^n x, \tag{6.8}$$

where $\boldsymbol{t}$ is the Laplace conjugate to $\boldsymbol{x}$. If the moment-generating functions of two probability density functions exist and are identical, then the probability density functions are also identical. The moments m_n then follow from the moment-generating function by differentiation

$$\begin{aligned} \frac{\mathrm{d}^r}{\mathrm{d}t^r} m_x(t)|_{t=0} =& \frac{\mathrm{d}^r}{\mathrm{d}t^r}\langle e^{\boldsymbol{t}\cdot\boldsymbol{x}}\rangle|_{t=0} \\ =& \left\langle \frac{\mathrm{d}^r}{\mathrm{d}t^r} e^{\boldsymbol{t}\cdot\boldsymbol{x}} \right\rangle |_{t=0} \\ =& \langle[\boldsymbol{x}\otimes ... \otimes \boldsymbol{x}] e^{\boldsymbol{t}\cdot\boldsymbol{x}}\rangle_{t=0} \\ =& m_n \end{aligned} \tag{6.9}$$

where in the second to third line, it was used that the differentiation with respect to t and the integration over x commute. Equivalently, there exists a central moment-generating function (Gregory, 2005)

$$m_{x-\mu}(t) = \langle e^{\boldsymbol{t}\cdot(\boldsymbol{x}-\boldsymbol{\mu})}\rangle. \tag{6.10}$$

The moments are then coefficients in a Taylor expansion of the moment-generating function because

$$\langle e^{\boldsymbol{t}(\boldsymbol{x}-\boldsymbol{\mu})}\rangle = \left\langle 1 + \boldsymbol{t}\cdot(\boldsymbol{x}-\boldsymbol{\mu}) + \frac{\boldsymbol{t}^2(\boldsymbol{x}-\boldsymbol{\mu})^2}{2!} + \frac{\boldsymbol{t}^3(\boldsymbol{x}-\boldsymbol{\mu})^3}{3!} ... \right\rangle. \tag{6.11}$$

where the tensor product is hidden in the powers for notational brevity[1]. We therefore see that an rth derivative with respect to t produces the rth moment.

---

[1] The meaning is unambiguous if one writes everything in components.

6.2 Reparameterization

Under a reparameterization from the variable $\boldsymbol{x} \sim f(\boldsymbol{x})$ to a variable $\boldsymbol{y}(\boldsymbol{x})$, the distribution $g(\boldsymbol{y})$ can be calculated from (Anderson, 2003)

$$f(\boldsymbol{x})\mathrm{d}^n x = g(\boldsymbol{y})\mathrm{d}^n y \tag{6.12}$$

meaning

$$g(\boldsymbol{y}) = f(\boldsymbol{x} = \tilde{\boldsymbol{y}}(\boldsymbol{x}))|J| \tag{6.13}$$

where $\tilde{\boldsymbol{y}}$ is the inverse function of $\boldsymbol{y}(\boldsymbol{x})$ and the Jacobi determinant $|J|$ of the transformation is

$$|J| = \det\left(\frac{\partial \boldsymbol{x}}{\partial \boldsymbol{y}}\right). \tag{6.14}$$

If N random samples $\boldsymbol{X}_i$ are drawn from their distribution $f(\boldsymbol{X}_i)$, then sample *estimators* of the distribution's mean and covariance can be calculated. The estimated mean is

$$\bar{\boldsymbol{X}} = \frac{1}{N}\sum_{i=1}^{N}\boldsymbol{X}_i, \tag{6.15}$$

which is the sample mean. The sample covariance matrix $\boldsymbol{S}$ is also an estimator of the true covariance matrix of f

$$\boldsymbol{S} = \frac{1}{N-1}\sum_{i=1}^{N}(\boldsymbol{X}_i - \bar{\boldsymbol{X}})(\boldsymbol{X}_i - \bar{\boldsymbol{X}})^T. \tag{6.16}$$

The factor $1/(N-1)$ in the definition of the sample covariance matrix accounts for the fact that due to estimating the mean $\bar{\boldsymbol{X}}$ from the data as well, one degree of freedom is lost. Both, $\bar{\boldsymbol{X}}$ and $\boldsymbol{S}$ are then unbiased estimators, where the definition of a bias $\mathcal{B}$ is (Anderson, 2003)

$$\mathcal{B}(\boldsymbol{\theta}) = \langle\boldsymbol{\theta}\rangle - \underline{\boldsymbol{\theta}}, \tag{6.17}$$

i.e. a bias is defined as the difference between the estimated mean[2] $\langle\boldsymbol{\theta}\rangle$ and the true quantity $\underline{\boldsymbol{\theta}}$. An unbiased estimator has $\mathcal{B}(\boldsymbol{\theta}) = 0$, and while a linear function of an unbiased estimator is again unbiased, a non-linear function in general is not.

6.3 The Gaussian distribution and the Central Limit Theorem (CLT)

The moment-generating function allows to prove the central limit theorem, which demonstrates why one can produce a Gaussianly distributed data set out of an intrinsically

---

[2] To be specific, this is the definition of a mean-bias. Definition of biases with respect to other statistical quantities, as e.g. the median, also exist. Within this thesis, a 'bias' will always refer to the mean-bias.

non-Gaussian data set, provided there are enough data points to average over. We prove the central limit theorem (CLT) for the univariate case following Gregory (2005).
Assume that there are n independently and identically distributed random variables x_i, meaning they are drawn from from the same distribution $p(x)$. We now average over them, producing the mean

$$\bar{x} = \frac{1}{n}\sum_{i}^{n} x_i. \tag{6.18}$$

Further, $p(x)$ must have a finite mean μ and variance σ^2 but is apart from these requirements arbitrary. Under these prerequisites, the sample average $\bar{x}$ for $n \to \infty$ will be

$$\bar{x} \sim \mathcal{G}(\mu, \sigma^2/n), \tag{6.19}$$

where $\mathcal{G}(\mu, \sigma^2/n)$ is the Gaussian of mean μ and variance σ^2/n. Equivalently, the rescaled variable Y follows

$$Y = \frac{\bar{x} - \mu}{\sigma/\sqrt{n}} \sim \mathcal{N}(0, 1) \tag{6.20}$$

where $\mathcal{N}(0, 1)$ is the normal distribution.
This can be proven as follows. Without loss of generality, apart from assuming a finite variance, we rescale the random variables by subtracting the mean and dividing out the variance

$$z_i = \frac{x_i - \mu}{\sigma}. \tag{6.21}$$

By the definition of σ and μ, we then have (Gregory, 2005)

$$\langle z_i \rangle = 0, \quad \langle z_i^2 \rangle = 1. \tag{6.22}$$

In terms of the z_i, the variable Y is then

$$Y = \frac{1}{\sqrt{n}}\sum_{i=1}^{n} z_i, \tag{6.23}$$

and its moment-generating function is then

$$m_Y(t) = \langle e^{tY} \rangle = \left\langle \exp\left(t\sum_{i=1}^{n}\frac{z_i}{\sqrt{n}}\right)\right\rangle = \left\langle \exp\left(\frac{tz_i}{\sqrt{n}}\right)\right\rangle^{n}. \tag{6.24}$$

Note the power of n in the last equality, which arises because independent probabilities multiply, and the samples x_i are independent draws from $p(x)$. The average can then be expanded into a series (Gregory, 2005)

$$\left\langle \exp\left(\frac{tz_i}{\sqrt{n}}\right)\right\rangle = \left\langle 1 + \frac{tz_i}{\sqrt{n}} + \frac{t^2 z_i^2}{2n} + \frac{t^3 z_i^3}{3! n^{3/2}} + ... \right\rangle, \tag{6.25}$$

where the second term is zero due to $\langle z_i \rangle = 0$, and in the third term $\langle z_i^2 \rangle = 1$, because of the rescaling by σ. The power of n in the moment-generating function can then be expanded as

$$\begin{aligned} m_Y(t) &= \left[1 + \frac{t^2}{2n} + \frac{t^3 \langle z_i^3 \rangle}{3! n^{3/2}} + ...\right]^n \\ &= \left[1 + \frac{1}{n}\left(\frac{t^2}{2} + \frac{t^3 \langle z_i^3 \rangle}{3! n^{1/2}} + ...\right)\right]^n . \end{aligned} \tag{6.26}$$

If we now introduce a shorthand for the series expansion

$$u = \frac{t^2}{2} + \frac{t^3 \langle z_i^3 \rangle}{3! n^{1/2}} + ..., \tag{6.27}$$

then

$$m_Y(t) = \left[1 + \frac{u}{n}\right]^n \tag{6.28}$$

If we now draw ever more samples, such that $n \to \infty$, then

$$u \to \frac{t^2}{2}, \tag{6.29}$$

because the other fractions will be supressed by the powers of n in their denominator. As we saw previously in equation (6.11), that each summand gives rise to one moment, we therefore see that the higher moments die out, if we average over increasingly more samples n. At the same time, for $n \to \infty$, we have

$$\lim_{n\to\infty} \left[1 + \frac{u}{n}\right]^n = e^u . \tag{6.30}$$

Consequently, in the limit of $n \to \infty$, the moment-generating function of Y is (Gregory, 2005)

$$m_Y(t) = e^{\frac{t}{2}}, \tag{6.31}$$

which is the Laplace transform of the standard normal distribution $\mathcal{N}$. Hence, the central limit theorem is proven. Most importantly, we also see the limit in which it arises: if n is large but finite, higher moments of the initial distribution $p(x)$ can survive – if they exist.

6.4 Mahalanobis distances and p-values

From General Relativity, we are familiar with a non-trivial metric in order to assign distances. In multivariate statistics, a Mahalanobis distance is the equivalent concept and assesses the statistical distance between random variables (Anderson, 2003).

If $\boldsymbol{x} \sim P(\boldsymbol{\mu}, \boldsymbol{C})$ and $\boldsymbol{y} \sim P(\boldsymbol{\mu}, \boldsymbol{C})$ are two multivariate random samples, drawn from the distribution $P(\boldsymbol{\mu}, \boldsymbol{C})$, and if $\boldsymbol{\mu}$ and $\boldsymbol{C}$ are the mean and the covariance matrix of P, then a Mahalanobis distance between $\boldsymbol{x}$ and $\boldsymbol{y}$ is (Mardia et al., 1979)

$$d_C = \sqrt{(\boldsymbol{x} - \boldsymbol{y})^T \boldsymbol{C}^{-1} (\boldsymbol{x} - \boldsymbol{y})}. \tag{6.32}$$

where the subscript C indicates that the distance is measured in units of the covariance $\boldsymbol{C}$, but other choices could have been made. The meaning of the Mahalanobis distance becomes clear when rescaling

$$x_i \to \frac{x_i - \mu_i}{\sigma_i}, \qquad y_i \to \frac{y_i - \mu_i}{\sigma_i} \tag{6.33}$$

such that the rescaled covariance matrix is then the identity matrix. Then, the Mahalanobis distance would simply be the usual Euclidean distance. Therefore, the Mahalanobis distance measures how many standard deviations $\boldsymbol{x}$ and $\boldsymbol{y}$ lie away from each other, in other words whether they are compatible within the expected noise. Consequently, if this distance exceeds a certain critical value, one would typically begin to wonder (a) whether $\boldsymbol{x}$ and $\boldsymbol{y}$ are really drawn from the same distribution, (b) whether the covariance matrix has been correctly assessed or (c) whether the distribution P may have surprisingly large higher moments, such that the covariance matrix alone does not well capture the statistical scatter. However, having statistical noise means that sometimes, a data set must be a very unlikely random sample from the tails of its probability distribution. The Mahalanobis distance will then be surprisingly large, even though the two samples do indeed stem from the same distribution. In order to quantify how often this will be the case, the statistical distribution of Mahalanobis distances themselves needs to be worked out. Due to mathematical simplifications, usually the distribution of the *squared* Mahalanobis distance is worked out. A famous example of a squared Mahalanobis distance is (Anderson, 2003)

$$d^2_{C;\mu} = \chi^2 = (\boldsymbol{x} - \boldsymbol{\mu})^T \boldsymbol{C}^{-1} (\boldsymbol{x} - \boldsymbol{\mu}), \tag{6.34}$$

where $\boldsymbol{x}$ follows the multivariate Gaussian $\mathcal{G}(\boldsymbol{\mu}, \boldsymbol{C})$, and (Mardia et al., 1979)

$$d^2_{S;\boldsymbol{\mu}} = T^2 = (\boldsymbol{x} - \boldsymbol{\mu})^T \boldsymbol{S}^{-1} (\boldsymbol{x} - \boldsymbol{\mu}), \tag{6.35}$$

where a sample estimated covariance matrix has been used instead of the true covariance matrix. The cosmologically relevant version of the T^2 distribution has been worked out during this thesis and can be found in chapter 10.

The distribution of χ^2 is indeed the χ^2-distribution, which arises as follows. If $x_1, ..., x_\nu$ are ν independent standard normally distributed samples, $x_i \sim \mathcal{N}(0, 1)$, then the sum

$$y = \sum_{i=0}^{\nu} x_i^2, \tag{6.36}$$

follows the χ^2 distribution, given by

$$\chi^2(y, \nu) = \frac{1}{2^{\nu/2}\Gamma(\nu/2)} y^{\frac{\nu-2}{2}} \exp\left(-\frac{y}{2}\right), \tag{6.37}$$

where ν are the degrees of freedom.

In case of the Mahalanobis distance (6.34), one has the sum

$$d^2_{C;\mu} = \chi^2 = (x_i - \mu_i)(C^{-1}_{ij})(x_j - \mu_j). \tag{6.38}$$

So if $\boldsymbol{C}$ is indeed the correct covariance matrix, and $\boldsymbol{\mu}$ is the correct mean of the data, such that

$$\frac{x_i - \mu_i}{\sigma_i} \sim \mathcal{N}(0, 1), \tag{6.39}$$

then (6.38) follows indeed the χ^2-distribution. If however, the mean $\boldsymbol{\mu}$ is a parametric model that introduces some systematic error, then the Mahalanobis distance (6.38) does not follow the χ^2-distribution. Likewise, if $\boldsymbol{C}$ is not the correct covariance matrix, then (6.38) also does not follow the χ^2-distribution.

If $\boldsymbol{C}$ and $\boldsymbol{\mu}$ are indeed the correct covariance and mean, and we average the distance (6.38) over the data, we find

$$\langle (x_i - \mu_i)(C_{ij}^{-1})(x_j - \mu_j) \rangle = \mathrm{Tr}\left(\boldsymbol{C}\boldsymbol{C}^{-1}\right) = \dim(\boldsymbol{x}), \tag{6.40}$$

i.e. the result is the dimension of the data set, which is then identical to the degrees of freedom in the χ^2 distribution. This is the origin of the usual rule-of-thumb that a $\chi^2/\mathrm{degF} \approx 1$ is a good fit. However, as soon as one needs to estimate $\boldsymbol{\mu}$ or $\boldsymbol{C}$, the degrees of freedom will reduce due to the additionally estimated quantities. For a discussion of the degrees of freedom, see e.g. (Andrae et al., 2010).

Further, the distribution of χ^2 has a certain width – consequently, one will not always find a $\chi^2/\mathrm{degF} \approx 1$ even if the fit is correct. As mentioned before, sometimes, the realization of the statistical noise will necessarily be a very unlikely one, and the found value of χ^2 is then surprisingly large. How often this happens, is quantified by p-values. The p-value for any statistic S is the probability P, that under the realization of the noise, this statistic is equal to, or exceeds a certain critical threshold S_c (Sellke et al., 2001)

$$p = P(S \geq S_c). \tag{6.41}$$

However, the realization of the noise will depend on the underlying *true* model M that generated this noise. Writing this explicitely, the p-value is (Lyons, 2014)

$$p = P(S \geq S_c | M). \tag{6.42}$$

The p-value really needs to be taken literally: It is the fraction of times, that the realization of noise around the true model M produces a measured statistic S that is as extreme or more extreme than a chosen threshold S_c.

In case of the test statistic S being the measured χ^2, the p-value is then the upper-tail integral of the χ^2-distribution $P(\chi^2)$,

$$p = \int_{\chi_c^2}^{\infty} P(\chi^2) \mathrm{d}\chi^2. \tag{6.43}$$

These values are usually tabulated at the end of statistics books, e.g. (Mardia et al., 1979). From (6.43) together with (6.39), we see directly which error we make when we use a wrong model to explain the statistical noise: if mean and covariance are not the correct ones, then the Mahalanobis distance does not follow the χ^2-distribution. If the errors are e.g.

overerstimated, then one will find too often that $\boldsymbol{x}$ and the mean $\boldsymbol{\mu}$ are compatible within their errors, meaning that one will find too often a low Mahalanobis distance. This will lead to distortions of the χ^2 distribution. If we nontheless calculate the p-value (6.43), then the resulting p-value will be based on wrong assumptions.

In general, the p-value is *not* the probability of the model being wrong. Converting a p-value into such a statement would need a prior on M, see Sellke et al. (2001) for a thorough discussion of how often a correct hypothesis is rejected and how often a wrong hypothesis is accepted when using p-values.

This also explains why a p-value is usually only adopted to *reject* a null-hypothesis H_0. Say we have reason to believe that H_0 is true, and then a measurements produces a statistic S whose p-value is $p = P(S > S_c|H_0) < 10^{-5}$, then the interpretation of this p-value were: if we had done everything correctly, and H_0 were correct, then we would measure such small a p-value only $1 : 10^{-5}$ of the times – it seems so unlikely, that the current measurement is this very rare case, so we think instead that there must be an error somewhere, and possibly H_0 is wrong.

6.5 Parameter estimation

If a data vector $\boldsymbol{x}$ has been observed, and there exists a theoretical model that can make a parameterized prediction for the mean $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{\theta})$, then the parameters $\boldsymbol{\theta}$ can be estimated by minimizing a Mahalanobis distance. For a Gaussian distribution of $\boldsymbol{x}$, minimizing χ^2 is equivalent to maximizing

$$L(\boldsymbol{x}|\boldsymbol{\theta}) = \frac{1}{\sqrt{|2\pi \boldsymbol{C}|}} \exp\left(-\frac{1}{2}\chi^2\right), \tag{6.44}$$

which is called the *likelihood*. Accordingly, as a minimal Mahalanobis distance means that the observed data vector is most compatible with the mean $\boldsymbol{\mu}(\boldsymbol{\theta})$, those parameters $\boldsymbol{\theta}$ that maximize the likelihood are a good estimate of the true parameters. Consequently, this procedure is known as maximum likelihood estimation.

Obviously, if we include only enough parameters into $\boldsymbol{\mu}(\boldsymbol{\theta})$, then we can decrease the Mahalanobis distance to arbitrarily small values. This is called overfitting, and can be detected by an *evidence* calculation, see Sect. 6.8, or a *cross-validation* technique. There are many forms of cross-validation, but all have in common that they hide some data points from the analysis, and then fit to the remaining points only. This will detect a join-the-dots like overfitting, because as soon as the hidden points are reincluded in the analysis, their Mahalanobis distance to an overfitted model will be much larger than the other points' distances. In the extreme case, it will appear as if the hidden points stem from another distribution. In contrast, if the fit lies smoothly within the statistical scatter, then the hidden points will be statistically compatible with the used points and the fit.

6.6 The scaled-χ^2 distribution and the Wishart distribution

Paving the way for the Wishart distribution, we also mention the scaled χ^2-distribution. Let $x_1, ..., x_n$ be n samples drawn from a Gaussian distribution, $x_i \sim \mathcal{G}(0, \sigma^2)$, with the true covariance $\sigma^2 \neq 1$. Then, we can study the distribution of their averaged square

$$s^2 = \frac{1}{n} \sum_i x_i^2. \tag{6.45}$$

If we divide out the variance, we have

$$\frac{s^2}{\sigma^2} = \frac{1}{n} \sum_i \frac{x_i^2}{\sigma^2}. \tag{6.46}$$

We can then introduce the auxiliary variables $m_i = x_i/\sigma$, such that $m_i \sim \mathcal{N}(0, 1)$ again. But then, we know that

$$y = \sum_{i=1}^{n} m_i^2 \sim \chi^2(n), \tag{6.47}$$

by the definition of the χ^2 distribution.
Consequently, we can derive the distribution of s^2 from

$$\frac{s^2}{\sigma^2} = \frac{1}{n} y \tag{6.48}$$

by transforming variables,

$$\frac{\mathrm{d}y}{\mathrm{d}s^2} = \frac{n}{\sigma^2}. \tag{6.49}$$

The distribution of s^2 is then

$$\begin{aligned} s^2 \sim\ & \chi^2(y, n) \frac{\mathrm{d}y}{\mathrm{d}s^2} \\ \propto & \left(\frac{ns^2}{\sigma^2}\right)^{\frac{n}{2}-1} \exp\left(-\frac{1}{2}\frac{ns^2}{\sigma^2}\right) n\sigma^{-2} \\ \propto & \sigma^{-n} \exp\left(-\frac{1}{2}\frac{ns^2}{\sigma^2}\right). \end{aligned} \tag{6.50}$$

The last line is the unnormalized scaled-χ^2 distribution, where the scaling with σ accounts for the non-unit variance.
The Wishart distribution is the multivariate generalization of the scaled-χ^2-distribution. It replaces the scalar samples $x_i \sim \mathcal{G}(0, \sigma^2)$ by multivariate samples $\boldsymbol{X}_i \sim \mathcal{G}(0, \boldsymbol{\Sigma})$, and the ordinary square by the tensor product. The density of

$$\boldsymbol{A} = \sum_{i=1}^{n} \boldsymbol{X}_i \boldsymbol{X}_i^T, \tag{6.51}$$

is then the Wishart distribution (Anderson, 2003)

$$\mathcal{W}(\boldsymbol{A}|\boldsymbol{\Sigma}, n) = \frac{|\boldsymbol{A}|^{\frac{n-p-1}{2}} \exp\left(-\frac{1}{2}\mathrm{Tr}\left(\boldsymbol{A}\boldsymbol{\Sigma}^{-1}\right)\right)}{2^{pn/2}|\boldsymbol{\Sigma}|^{n/2}\Gamma_p(n/2)}, \tag{6.52}$$

where p is the dimension of the $\boldsymbol{X}_i$ and Γ_p the multivariate Gamma function

$$\Gamma_p(t) = \pi^{p(p-1)/4}\prod_{i=1}^{p}\Gamma[t-(i-1)/2]. \tag{6.53}$$

6.7 Bayesian and frequentist statistics

Statistics can be broadly grouped into the classes of Bayesian statistics, and frequentist statistics. The prerequisites described in the previous Sect. 6 are common to both. The two classes differ however in whether they assume data $\boldsymbol{X}$ to be given, wanting to infer the likelihood that a certain model $\mathcal{M}$ is true (the Bayesian perspective), or whether they assume a model $\mathcal{M}$ to be true a priori, and then calculating the likelihood of measuring the data $\boldsymbol{X}$, given this model (the frequentist perspective.)

Particle physicists are traditionally frequentists as this branch of statistics relies heavily on repetitive executions of an experiment. Particle physicists can achieve this by letting e.g. their accelerators run for longer, producing more collisions between particles. Current cosmological research tends to use Bayesian statistics instead, as we have only one Universe to observe, which fundamentally limits the amount of data that we can ultimately get. Also, Bayesian statistics allows to address the question of model selection better than frequentist statistics does.

The point where frequentist statistics and Bayesian statistics begin to differ, is in how they interpret the term ‘probability’. To a frequentist, a probability is a frequency: given that the outcome of an experiment can be $(e_1, ..., e_n)$ different events, the probability of event e_i is the frequency with which e_i occurs if the experiment is repeated an infinite number of times.

This is certainly one good definition of probability, but another good definition is to interpret a probability as a *credibility*. For example, one could want to investigate nuclear reactor accidents. The data will then be actually occured accidents – and one certainly does not want to repeat such accidents for a large number of times in order to be able to speak of a frequency. Yet, from the general understanding of reactors, and the general understanding of a cause-and-reaction chain, one will *believe* that certain models explain reactor accidents quite well, while other models seem contrived and unlikely. Bayesian statistics provides a mathematical framework to quantify this assessment of probability as a degree of *belief*. In total, Bayesian statistics accounts more flexibly for the fact that the interpretation of a data set might be ambiguous, as e.g. multiple models could explain the taken data.

The relation between frequentist and Bayesian statistics is encoded in Bayes theorem (6.4)

$$P(\boldsymbol{\theta}_M|\boldsymbol{X}) = \frac{P(\boldsymbol{\theta}_M)P(\boldsymbol{X}|\boldsymbol{\theta}_M)}{P(\boldsymbol{X})}, \tag{6.54}$$

where $P(\boldsymbol{\theta}_M|\boldsymbol{X})$ is the *posterior* likelihood that quantifies the degree of belief that we have in M being true and having the parameters $\boldsymbol{\theta}$, given that the data $\boldsymbol{X}$ were observed. We now indicate by the subscript M that the parameters are to be evaluated within the framework of the model M. This is done to avoid the confusion of whether $M(\boldsymbol{\theta}_1)$ and $M(\boldsymbol{\theta}_2)$ are to be called different models, because the parameters take different values – or whether different models are really different from first principles, such that model $M_1(\boldsymbol{\theta})$ makes different predictions than $M_2(\boldsymbol{\theta})$, even though the parameters take the same values. Within this thesis, the latter point of view was adopted.

Further, $P(\boldsymbol{X}|\boldsymbol{\theta}_M)$ is the likelihood of getting a data set $\boldsymbol{X}$, assuming the model M is true and its parameters take the value $\boldsymbol{\theta}$. This is the quantity that frequentists evaluate, and will in the following be rewritten as $L(\boldsymbol{X}|\boldsymbol{\theta}_M)$. Usually, the model is then called a hypothesis. The term $P(\boldsymbol{\theta}_M)$ is called a prior for the model parameters. We will rewrite priors as $\mathcal{P}(\boldsymbol{\theta}_M)$. Tracing the origin of the priors back to Bayes theorem, we see that they are mandatory in order to interchange the order of premises and consequences. They therefore cannot be omitted. A prior quantifies the credibility distribution that we assign to a parameter, before data were taken. A good prior is a prior that can be overruled by the data. It is however non-trivial to come up with such a prior. Various concepts are currently in use:

- **Theory-motivated priors.** Often, a theory makes precise predictions for the upper and lower boundary of a parameter. Typical choices for the prior are then to use a uniform distribution between these boundaries, or a powerlaw distribution.

- **Conjugated priors.** Conjugate priors stem from the same family of distributions as the posteriors. For example, convolving a Gaussian with a Gaussian produces again a Gaussian, so a Gaussian prior is conjugate to a Gaussian posterior. In general, conjugate priors usually allow for an analytical expression for the posterior. They are therefore quite popular, if the prior is only needed to parameterize some uncertainty, and where the exact shape of the prior is not very important (Anderson, 2003).

- **Objective priors.** Objective priors try to quantify the prior belief that without having taken data, we have no means from which to draw information. Consequently, objective priors try to be uninformative. In low dimensions, this can be achieved quite well, but in higher dimensions, priors that were uninformative for low dimensions, can become informative. Consider for example the flat (uniform) prior for a Euclidean distance r to the origin. In one dimension, the uniform prior is completely uninformative. In high dimensions, r is the radius of a hypersphere. As the volume of an $n-$sphere is $V_n(r) = \pi^{n/2}r^n/\Gamma(n/2+1)$, the flat prior on r suddenly becomes very informative. A famous prior, that is invariant under reparameterizations, and therefore regarded as uninformative, is the Jeffreys prior $\sqrt{|\mathcal{I}_F(\boldsymbol{\theta})|}$, which is the square root of the Fisher information's determinant (Jeffreys, 1961b). Another method of constructing an objective prior are to maximize the entropy between likelihood and posterior, for example through an evaluation of the Kullback-Leibler divergence (Sun & Berger, 2006). The idea behind this is, that if the entropy is maximized, the prior has the least possible effect on the posterior.

- **Exact frequentist matching priors.** Using such priors, the credibility contours of the Bayesian posterior map exactly to the likelihood contours of a frequentist coverage. I.e. if the experiment, that was analyzed in a Bayesian way, could be reproduced arbitrarily many times, then the frequentist analysis would find the same likelihood contours as the Bayesian credibility levels. Usually, such priors need to be explicitely constructed (Sun & Berger, 2006).

- **Priors from other experiments.** Often, if two experiments are combined, one finds the formulation that 'a prior of the other experiment is added'. What is meant, is that the two likelihoods of the experiments are combined, and one is interpreted as a prior that updates the constraining power of the old data set to a new, more constraining posterior likelihood. Such priors are not to be confused with the purely constructed Bayesian priors described above. Instead, such priors mimic the procedure of learning by taking new data. Care needs to be taken, that these priors are not used twice or hidden. For example, baryonic feedback models could include baryonic parameters that are to be determined, and they are determined such that the model reproduces the expectations from ΛCDM. If this model, with fixed baryonic parameters, is then used to analyze real data, the remaining free cosmological parameters will be biased towards ΛCDM.

Note also, that a prior parameterizes whether one has measured parameters or not. If the posterior peaks sharply within the prior range, then the data are informative and provide a good measurement of the parameters. If the posterior covers the entire prior space, then we have not learned anything from the data – we are just stuck on our prior beliefs.
The term $P(\boldsymbol{X})$ is the likelihood that the parametric model $M(\boldsymbol{\theta})$ can produce the data at all, independent of its parameter values. $P(\boldsymbol{X})$ is consequently given by (Trotta, 2008)

$$P(\boldsymbol{X}) = \int L(\boldsymbol{X}|\boldsymbol{\theta}_M)\mathcal{P}(\boldsymbol{\theta}_M)\mathrm{d}^n\theta, \tag{6.55}$$

Equation (6.55) therefore completely marginalizes over all parameters, and $P(\boldsymbol{X})$ is therefore also known as marginal likelihood or *evidence*. In the following, we will write the evidence as ε. When inferring parameters of a model, ε is an unimportant normalization constant. As it integrates out the parameters, it obviously cannot affect the parameter inference any more. However, when comparing different models $M_1, ..., M_n$, then ε is of central importance.
Having clarified these concepts, Bayes' theorem as used within Bayesian statistics is

$$P(\boldsymbol{\theta}_M|\boldsymbol{X}) = \frac{\mathcal{P}(\boldsymbol{\theta}_M)L(\boldsymbol{X}|\boldsymbol{\theta}_M)}{\varepsilon}. \tag{6.56}$$

6.8 Occam's razor and model comparisons

The evidence has so far been presented as a mere normalization constant for the posterior likelihood, but it has more interesting qualities, and these shall be mentioned.

Evaluating the evidence is numerically cumbersome and special codes like Nested Sampling were developed for this purpose (Skilling, 2004). Evaluating ε is however often well worth the pains, as the evidence penalizes models that are too complicated in the light of current data, and it further allows to compare different models.

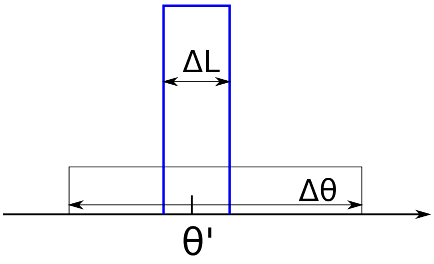

Figure 6.1: Toymodel for the evaluation of the evidence.

The penalization of too complicated models is also called *Occam's razor*. The principle of Occam's razor exploits that the evidence will sink with wide prior ranges which can be understood from the toy model in Fig. 6.1: shown in black is a wide prior on a parameter θ. As the prior is normalized, we have $\int \mathcal{P}(\theta)\mathrm{d}\theta = 1 = \mathcal{P}(\theta)\Delta\theta$. The likelihood is depicted in blue and has the peak height $L(\theta')$, and width ΔL. The evidence is then the overlap of prior and likelihood

$$\begin{aligned}\varepsilon = \int L(\boldsymbol{X}|\boldsymbol{\theta}_M)\mathcal{P}(\boldsymbol{\theta}_M)\mathrm{d}^n\theta &= L(\theta')\mathcal{P}(\theta')\Delta L \\ &= L(\theta')\frac{\Delta L}{\Delta\mathcal{P}}.\end{aligned} \tag{6.57}$$

In the last line, the normalization of the prior was used. We therefore see, that the evidence includes a factor $L(\theta')$, which is the goodness of the best fit. Typically, one has $L \propto \exp(-\frac{1}{2}\chi^2)$ and if the inclusion of a new parameter leads to a significant improvement in the goodness of the fit, then $L(\theta')$ will strongly increase the evidence. If, however, including the new parameter does not significantly improve the fit to the data, then the term $L(\theta')$ will stay approximately constant. However, including a new parameter enforces that a new prior interval $\Delta\mathcal{P}$ needs to be defined. If this interval is wide, the term $1/\Delta\mathcal{P}$ will bring the evidence down, thereby penalizing the inclusion of the new parameter.

Frequentist statistics does not have a similary good mechanism to assess how many parameters the data warrant. Increasing the number of parameters will lead to a slow and monotonic decrease of χ^2/degF and it cannot well be seen when fitting to the noise begins.

The evidence can then be used to compare different models to each other, where e.g. model M_1 may be the standard model of cosmology, and model M_2 adds one more parameter to it. If for one special value of the new parameter, the old model M_1 is re-established, then one speaks of a *nested* model. Model comparisons can however also be conducted on models that differ entirely.

The question of 'How likely is my model M_1 in light of the data?' cannot be answered, as likelihoods are normalized, and in order to assess the likelihood of one model, all other ever possible models would need to be known. However, one can still gain relative statements, such as M_1 is with the current data more likely than M_2 with the following mechanism:

The evidence $\varepsilon = L(\boldsymbol{X}|M_1)$ was the likelihood of the model M_1 being able to produce the data $\boldsymbol{X}$ at all. If we invert this, $L(M_1|\boldsymbol{X})$ is the likelihood of the model, given the observed

Table 6.1: Jeffreys scale

$\lvert \log\left(\frac{\varepsilon(M_1)}{\varepsilon(M_2)}\right) \rvert$	odds	interpretation	prob. of favoured model
≤ 1.0	3:1	better data is needed	≤ 0.75
≤ 2.5	12:1	weak evidence	0.923
≤ 5.0	$\leq$ 150:1	moderate evidence	0.993
≥ 5.0	$>$ 150:1	strong evidence	> 0.993

data. The inversion is done with Bayes theorem

$$L(M_1|\boldsymbol{X}) = L(\boldsymbol{X}|M_1)\frac{\mathcal{P}(M_1)}{\mathcal{P}(\boldsymbol{X})}, \tag{6.58}$$

where $\mathcal{P}(M_1)$ is a prior on the model, and $\mathcal{P}(\boldsymbol{X})$ is a prior on getting the data, whatever model may be true. The term $\mathcal{P}(\boldsymbol{X})$ will again appear when we invert the evidence of another model M_2. It will then drop out if we take the ratio

$$\begin{aligned} \frac{L(M_1|\boldsymbol{X})}{L(M_2|\boldsymbol{X})} &= \frac{\mathcal{P}(M_1)L(\boldsymbol{X}|M_1)}{\mathcal{P}(M_2)L(\boldsymbol{X}|M_2)} \\ &= \frac{\mathcal{P}(M_1)}{\mathcal{P}(M_2)}\frac{\varepsilon_1}{\varepsilon_2}. \end{aligned} \tag{6.59}$$

The last line contains two ratios: first, the ratio of our prior beliefs about the models. This quantifies which model we prefer for various reasons. The second ratio quantifies which model the *data* prefer. Being natural scientists, it is the second ratio that counts more. So on this basis, we see that the evidence allows to assess in a quantitative manner which model is more likely the one that produced the data. The interpretation of the evidence ratios is usually done with the Jeffreys scale, see Tab. 6.1 (Trotta, 2008).

6.9 The Fisher matrix

The Fisher matrix F is a quick way to map out the posterior likelihood. It is defined as (Tegmark et al., 1997; Amendola & Tsujikawa, 2010a)

$$F_{\alpha\beta} = \langle \partial_\alpha \partial_\beta \mathcal{L} \rangle \tag{6.60}$$

such that it is the averaged Hessian matrix in parameter space. The approximated posterior is then the Gaussian

$$P = N \exp(-\frac{1}{2}F_{\alpha\beta}\Delta p_\alpha \Delta p_\beta). \tag{6.61}$$

Consequently, if the diagonal elements of the Fisher matrix are much larger than the off-diagonal elements, then the covariances between parameters will be small, and the measurement will be precise.

As Gaussians allow for many analytical results, the Fisher matrix allows for a recipe-like manipulation of posteriors which shall be described here.

If two experiments constrain certain parameters independently, then the posterior likelihoods can be multiplied. This corresponds to an addition of the Fisher matrices. If, however, there exists a cross-correlation between these experiments, then these must be taken into account too.

If one changes the parameterization of a model, this corresponds to a basis change in the posterior space. If one changes from the parameter set $\boldsymbol{p}$ to the parameter set $\boldsymbol{q}$, then the Jacobian matrix for this basis exchange is

$$J = \frac{\partial \boldsymbol{p}}{\partial \boldsymbol{q}}. \tag{6.62}$$

The Fisher matrix then transforms as any other linear map under a basis change as

$$\boldsymbol{F}_q = \boldsymbol{J}^{-1} \boldsymbol{F}_p \boldsymbol{J} \tag{6.63}$$

A degenerate Fisher matrix means that one parameter cannot be determined.

Fixing a parameter p_i at its best fit value, also called *maximizing* that parameter, means that Δp_{p_i} is identically to zero. Equivalently, the row and column from the Fisher matrix that represent this parameter can be discarded.

Marginalizing a parameter, i.e. integrating it out, can be achieved by dropping a row and column from the inverse Fisher matrix. This result is a consequence of the following analytical integral for Gaussians:

$$\int_{-\infty}^{\infty} e^{-ax^2} \mathrm{d}x = \sqrt{\pi/a} \quad (\text{for } a > 0). \tag{6.64}$$

So if a multidimensional Gaussian is marginalized, the result will be again a Gaussian, and the integral (6.64) just contributes a constant factor.

For example, if we marginalize the 2-dimensional Gaussian posterior

$$P_{2D} \propto \exp(-\frac{1}{2} p_i F_{ij} p_j) \tag{6.65}$$

where we assume without loss of generality that the best fit values lie at zero, and where the Fisher matrix is

$$F = \begin{pmatrix} F_{11} & F_{12} \\ F_{12} & F_{22} \end{pmatrix} \tag{6.66}$$

then the recipe of marginalizing the parameter p_2 is to first invert, leading to

$$F^{-1} = \frac{1}{F_{11}F_{22} - F_{12}^2} \begin{pmatrix} F_{22} & -F_{12} \\ -F_{12} & F_{11} \end{pmatrix}. \tag{6.67}$$

After striking out the second row and colum, the remaining entry of the Fisher matrix is then $F_{1D} = F_{11} - F_{12}^2/F_{22}$, leading to the one-dimensional posterior

$$P_{1D} \propto \exp\left(-\frac{1}{2}(F_{11} - F_{12}^2/F_{22}) p_1^2\right). \tag{6.68}$$

On the other hand, carrying out the marginalization analytically, we have

$$\begin{aligned}\int P_{2D}\mathrm{d}p_2 &= \int \exp(-\frac{1}{2}p_iF_{ij}p_j)\mathrm{d}p_2\\ &= \int \exp\left(-\frac{1}{2}(F_{11}-F_{12}^2/F_{22})p_1^2\right)\exp\left(-\frac{1}{2}\left[\frac{F_{12}}{F_{22}}p_1-p_2\right]^2\right)\mathrm{d}p_2.\end{aligned} \tag{6.69}$$

Here, only the second term depends on p_2 and its integral is given by (6.64) which returns a constant. We therefore see that only the first term remains, and is identical to the result (6.68) that we obtained by dropping a row and a column. As the marginalization of one parameter returns a lower-dimensional Gaussian, this calculation can be reapeated for an arbitrary number of parameters.

If the Fisher matrix is seen as a linear map from the data space to the parameter space, then it describes by how much the data set $\boldsymbol{X}$ is compressed into the parameter set $\boldsymbol{\theta}$. The determinant of a linear map describes how the hypervolume changes during the mapping. A useful figure of merit (FoM) is therefore given by the determinant of the Fisher matrix, and if only two parameters are targeted, by the determinant of the 2 × 2 submatrix

$$FoM = \sqrt{|\boldsymbol{F}_{2\times 2}|}. \tag{6.70}$$

For example, dark energy research could be interested in the (w_a, w_0)-submatrix of the Fisher matrix, if the CPL-parameterization is used (Chevallier & Polarski, 2001).

Often, it is claimed that the Fisher matrix satisfies the Cramer-Rao bound (Cramer, 1946). However, in Sect. 9.4, we will show why the Cramer-Rao bound is in general not a quantity of interest in cosmology, and why the Fisher matrix can even violate it.

6.10 Metropolis-Hastings and Hamilton Monte Carlo Sampling

If a likelihood is non-Gaussian, analytical results like marginals are typically not possible. If one then wants to map out the posterior, it is necessary to take refuge to numerical methods. If the number of parameters to be determined is n_p, and if n_p is small, then one can simply evaluate the posterior on a grid. Typically, the posterior will however be of a roundish shape and for increasing n_p, the volume of the empty edges in the box will increase relative to the volume of appreciable posterior likelihood. Grid methods are therefore quickly limited by numerical speed.

The solution is then to switch to a sampling method of which the Metropolis-Hastings algorithm and the Hamilton Monte Carlo shall be described here.

A Markov process borrows the idea of transition probabilities between different states as known from thermodynamics or quantum mechanics. If there exist multiple discrete probability levels $\mathcal{P}_i$, and the possibility of a transition from state $\mathcal{P}_i$ to state $\mathcal{P}_j$ has the rate probability $r_{i\to j}$, then an equilibrium between the occupation of the different states $\mathcal{P}_i$ is reached if

$$r_{i\to j}\mathcal{P}_i = r_{j\to i}\mathcal{P}_j, \tag{6.71}$$

meaning that the population and depopulation occurs at transition rates whose ratio is identical to the probability of the states. This principle is sometimes called *detailed balance*. If applied to atoms, it regulates the absorption and emission of photons between the different electron shells. For example, the phenomenon of phosphoresence occurs for substances that absorb photons in order to reach equilibrium with an external light source – once put in the dark, they are brought out of equilibrium and begin to re-emitt photons in order to now equilibrate with the darkness.

A Monte Carlo Markov Chain now mimicks the occupation of different probability states by sampling: discrete evaluations $P(\boldsymbol{\theta}_i)$ of the posterior at i discrete points in parameter space serve as the $\mathcal{P}_i$.

The Metropolis-Hasting algorithm for sampling a posterior $P(\boldsymbol{\theta})$ is given by the following procedure (Hajian, 2007a):

Metropolis-Hastings algorithm

1. Come up with a good/better guess for a Gaussian approximation $\mathcal{G}_P(\boldsymbol{\theta})$ to the posterior. Obviously, the Fisher matrix can provide this, but a crudely estimated parameter covariance matrix from potential previous analyses also fulfills the purpose.

2. For $i = 0$ to N_{MCMC}
   if $i = 0$ evaluate the posterior P at some point $\boldsymbol{\theta}_0$ in parameter space, preferentially one with a presumed high likelihood, else use the current $\boldsymbol{\theta}_i$ of the chain.

3. Draw a random step in parameter space $\Delta\boldsymbol{\theta}_i \sim \mathcal{G}_P(\boldsymbol{\theta})$.

4. Calculate $P(\boldsymbol{\theta}_i + \Delta\boldsymbol{\theta}_i)$ and $R = \frac{P(\boldsymbol{\theta}_i+\Delta\boldsymbol{\theta}_i)}{P(\boldsymbol{\theta}_i)}$.

5. If $R > 1$, then the posterior probability at the new point $\boldsymbol{\theta}_i + \Delta\boldsymbol{\theta}_i$ is larger than the old probability; the new point is then accepted as $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i + \Delta\boldsymbol{\theta}_i$.

6. If $R < 1$, then draw $\alpha \sim \mathrm{Uniform}[0, 1]$.
   If $\alpha > R$, then $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i$, i.e. the point $\boldsymbol{\theta}_i + \Delta\boldsymbol{\theta}_i$ is rejected because it has too low a probability.
   If, however, $\alpha < R$, then $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i + \Delta\boldsymbol{\theta}_i$, i.e. the trial point is accepted because it has still a fairly high probability.

7. Store all points $\boldsymbol{\theta}_i$. These then build up the Monte Carlo Markov Chain.

If the chain samples the posterior badly, its parameter covariance should be determined, and a new run beginning at point (1) should then be started. Here, points (4-6) mimic the detailed balance. Rejecting a sample should therefore rather be regarded as the compound process of first accepting it testwise on the state of lower posterior likelihood, and then re-emitting it in order to reach the balance condition (6.71).

Point (1) of this algorithm is a means of optimization: if the Gaussian approximates the posterior well, then the random steps in point (3) will approximately follow the orientation of the posterior. If the Gaussian approximation is bad, the random steps will guide the sampler towards regions of very low posterior likelihood. This will then result in a high rejectance rate in step (6).

If the starting point $\boldsymbol{\theta}_0$ of a chain is a very unlikely combination of parameters, then the chain will at the beginning first run up the posterior gradient due to point (5). This phase is called burn-in, and corresponds to a beginning population of so far completely unpopulated states $\mathcal{P}_i$[3]. Consequently, overproportionally many points will be accepted. Only when the sampler has reached the peak region it will begin to equilibrate with the underlying probability distribution. In equilibrium, the density n of accepted points is then proportional to the posterior likelihood

$$n(\boldsymbol{\theta}) \propto P(\boldsymbol{\theta}), \tag{6.72}$$

which is why the sampler can be used to assess the posterior likelihood. For parameter inference, the proportionality is sufficient, as all probability distributions can be normalized retrospectively, if needed.

Marginalizing parameters of an MCMC run is achieved by histograming the chain as a function of all parameters that shall remain. This is structurally equivalent to what usual Monte Carlo integration routines like VEGAS (Lepage, 1978, 1980) or MISER do (Press & Farrar, 1990).

As known from thermodynamics, it will typically take time before the emission rates between the different states have reached the equilibrium condition (6.71). For the Monte Carlo Markov chain this means that is should not be aborted too early; typical length for chains used in this thesis were 10^5 points, in case of DALI often 10^7 points, simply because the sampler was so fast and allowed for it.

In the terminology of MCMC-runs, reaching equilibrium is usually referred to a chain having converged. Whether a chain has converged can best be decided by comparing it to other chains. If those chains find the same best fit and parameter covariance (within the sampling errors) and do not show any offsets when plotted, and if all chains sample the entire parameter space sufficiently, then usually equilibrium has been reached (Audren et al., 2013a; Allison & Dunkley, 2014). In equilibrium, an efficient sampling is usually reached if the ratio between accepted and rejected points is about 0.3. If the acceptance ratio is much closer to one, it means that the sampler only makes very small steps $\Delta\boldsymbol{\theta}_i$ and consequently $R \approx 1$. Such small steps inhibit the sampler from propagating efficiently through the entire parameter space. In the opposite limit of a very small acceptance ratio, the sampler uses too large steps and gets rejected because it leaves the region of appreciable posterior likelihood.

---

[3] In the analogy to the phosphoresent substance above, the burn-in corresponds to taking the substance out of the dark and putting it into the light to ‘charge up’.

These considerations imply that the Metropolis-Hastings algorithm presented above works best for approximately Gaussian likelihoods. In case of strong non-Gaussianities, large steps $\Delta\boldsymbol{\theta}_i \sim \mathcal{G}_P$ will often guide the sampler into the wrong direction, thereby bringing down the acceptance ratio, and very small steps will sample the non-Gaussian posterior inefficiently. These problems can be solved by updating to a Hamilton Monte Carlo (HMC) Sampler.
An HMC sampler is actually nothing else than a Metropolis-Hasting sampler – only that it increases the distance between subsequent samples in a very intelligent way. It begins by introducing a potential energy $U(\boldsymbol{\theta})$ as (Hajian, 2007a)

$$U(\boldsymbol{\theta}) = -\log P(\boldsymbol{\theta}). \tag{6.73}$$

It then introduces further a kinetic energy

$$K(\boldsymbol{u}) = \boldsymbol{u}^T\boldsymbol{u}/2, \quad \boldsymbol{u} \sim \mathcal{N}(0, \mathcal{I}), \tag{6.74}$$

where the velocity $\boldsymbol{u}$ is a random variable drawn from the multivariate normal distribution. Kinetic and potential energy are then combined into a Hamiltonian

$$H(\boldsymbol{\theta}, \boldsymbol{u}) = U(\boldsymbol{\theta}) + K(\boldsymbol{u}). \tag{6.75}$$

The exponential of the Hamiltonian, which corresponds to the Boltzmann factor in thermodynamics, is then related to the posterior likelihood $P(\boldsymbol{\theta})$ as

$$\exp\left(-H(\boldsymbol{\theta}, \boldsymbol{u})\right) = P(\boldsymbol{\theta})\mathcal{N}(0, \mathcal{I}) \tag{6.76}$$

So if the auxiliary velocities $\boldsymbol{u}$ are marginalized over, which is equivalent to not protocolling them in the chain, then sampling $\exp(-H)$ allows to sample $P(\boldsymbol{\theta})$.
The HMC algorithm now uses the Hamiltonian (6.75) in order to increase the distance between two Metropolis-Hastings steps. The Hamiltonian equations of motion are

$$\dot{\boldsymbol{\theta}} = \boldsymbol{u}, \quad u_i = -\frac{\partial H}{\partial \theta_i}. \tag{6.77}$$

These can be solved numerically and as they will later only provide auxiliary steps, it is sufficient to apply the leapfrog algorithm (Hajian, 2007a)

$$\begin{aligned} u_i(t + \frac{\epsilon}{2}) &= u_i(t) - \frac{\epsilon}{2}\left(\frac{\partial U}{\partial \theta_i}\right)_{\boldsymbol{\theta}(t)} \\ \theta_i(t + \epsilon) &= \theta_i(t) + \epsilon u_i(t + \epsilon/2) \\ u_i(t + \frac{\epsilon}{2}) &= u_i(t) - \frac{\epsilon}{2}\left(\frac{\partial U}{\partial \theta_i}\right)_{\boldsymbol{\theta}(t+\epsilon)}. \end{aligned} \tag{6.78}$$

The HMC sampler now alternates between leapfrog steps and Metropolis-Hastings steps. Its algorithm can be represented as (Hajian, 2007a)

Hamilton Monte Carlo algorithm

1. For $i = 0$ to N_{MCMC}
   if $i = 0$, choose a starting point $\boldsymbol{\theta}_0$,
   else use the current $\boldsymbol{\theta}_i$ of the chain.
2. Draw a random velocity $\boldsymbol{u}_i \sim \mathcal{N}(0, \mathcal{I})$.
   Leapfrog loop
   (a) Use $\boldsymbol{\theta}_i$ and $\boldsymbol{u}_i$ as initial conditions for the Hamiltonian equations of motions.
   (b) For $j = 0$ to N_L
   make leapfrog steps that update $(\boldsymbol{\theta}_j, \boldsymbol{u}_j) \rightarrow (\boldsymbol{\theta}_{j+1}, \boldsymbol{u}_{j+1})$
3. Having arrived at $(\boldsymbol{\theta}_{N_L}, \boldsymbol{u}_{N_L})$, calculate $R = \exp[-H(\boldsymbol{\theta}_i, \boldsymbol{u}_i) + H(\boldsymbol{\theta}_{N_L}, \boldsymbol{u}_{N_L})]$.
4. If $R > 1$, the new point is accepted, $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i$.
5. If $R < 1$, draw $\alpha \sim \mathrm{Uniform}[0, 1]$.
   If $\alpha > R$, then $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i$, i.e. the trial point $\boldsymbol{\theta}_{N_L}$ is rejected.
   If $\alpha < R$, then $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_{N_L}$, i.e. the trial point is accepted.

We therefore see that the HMC algorithm includes the Metropolis-Hastings algorithm and its rule of detailed balance, in order to decide over acceptance and rejectance. However, it replaces the rule by which $\Delta\boldsymbol{\theta}_i = \boldsymbol{\theta}_i - \boldsymbol{\theta}_{i-1}$ is calculated. While Metropolis-Hastings draws $\Delta\boldsymbol{\theta}_i \sim \mathcal{G}_P$, HMC calculates it via the leapfrog loop. By solving the Hamiltonian equations of motion, the sampler is guided along isocontours of H as Hamiltonian dynamics conserves energy. Once the sampler has found regions of high likelihood, it will then be automatically guided towards regions of comparably high likelihood. Drawing a random kinetic energy $\boldsymbol{u}_i$ before each leapfrog loop ensures that different levels of likelihood are covered.

While calculating the leapfrog moves, the gradient of the potential energy must be evaluated. As the potential energy is $U(\boldsymbol{\theta}) = -\log P(\boldsymbol{\theta})$, this is typically the most expensive step in the entire algorithm. The leapfrog loop is however a purely auxilliary feature in the algorithm, used to guide the sampler through parameter space. A less accurate evaluation of this loop does not harm the entire algorithm. Therefore, a speed-up can be achieved by supplying additionally an approximate posterior that is much faster to evaluate than the real posterior.

During this thesis, the DALI approximation to non-Gaussian posteriors was used as approximate potential for the HMC sampler. If $P(\boldsymbol{\theta})$ is the real posterior, and $D_P(\boldsymbol{\theta})$ is its DALI approximation, then we can define two Hamiltonians

$$H_P(\boldsymbol{\theta}, \boldsymbol{u}) = U(\boldsymbol{\theta}) + K(\boldsymbol{u}), \qquad H_D(\boldsymbol{\theta}, \boldsymbol{u}) = -\log D_P(\boldsymbol{\theta}) + K(\boldsymbol{u}) \tag{6.79}$$

H_P behing the Hamiltonian of the real posterior, and H_D being the Hamiltonian of the DALI approximation. While the Metropolis-Hastings decision rule over acceptance and rejectance in the HMC algorithm must use the full Hamiltonian H_P, the leapfrog moves can be calulated along the DALI contours

$$
\begin{aligned}
u_i(t+\frac{\epsilon}{2}) &= u_i(t) - \frac{\epsilon}{2}\left(\frac{\partial(-\log D_P)}{\partial\theta_i}\right)_{\boldsymbol{\theta}(t)} \\
\theta_i(t+\epsilon) &= \theta_i(t) + \epsilon u_i(t+\epsilon/2) \qquad (6.80) \\
u_i(t+\frac{\epsilon}{2}) &= u_i(t) - \frac{\epsilon}{2}\left(\frac{\partial(-\log D_P)}{\partial\theta_i}\right)_{\boldsymbol{\theta}(t+\epsilon)} .
\end{aligned}
$$

The algorithm for a DALI-aided HMC sampler as developed during this thesis is then:

DALI-aided HMC algorithm

1. FOR $i = 0$ TO N_{MCMC}
   if $i = 0$, choose a starting point $\boldsymbol{\theta}_0$,
   else use the current $\boldsymbol{\theta}_i$ of the chain.
2. Draw a random velocity $\boldsymbol{u}_i \sim \mathcal{N}(0, \mathcal{I})$.
   DALI Leapfrog loop
   (a) Use $\boldsymbol{\theta}_i$ and $\boldsymbol{u}_i$ as initial conditions for the Hamiltonian equations of motions.
   (b) For $j = 0$ to N_L
   make leapfrog steps along DALI, following (6.80) in order to update $(\boldsymbol{\theta}_j, \boldsymbol{u}_j) \rightarrow (\boldsymbol{\theta}_{j+1}, \boldsymbol{u}_{j+1})$
3. Having arrived at $(\boldsymbol{\theta}_{N_L}, \boldsymbol{u}_{N_L})$, calculate $R = \exp[-H_P(\boldsymbol{\theta}_i, \boldsymbol{u}_i) + H_P(\boldsymbol{\theta}_{N_L}, \boldsymbol{u}_{N_L})]$ from the real posterior $P(\boldsymbol{\theta})$ – not (!) from DALI's $H_D(\boldsymbol{\theta})$.
4. IF $R > 1$, the new point is accepted, $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i$.
5. IF $R < 1$, draw $\alpha \sim \text{Uniform}[0, 1]$.
   IF $\alpha > R$, then $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i$, i.e. the trial point $\boldsymbol{\theta}_{N_L}$ is rejected.
   IF $\alpha < R$, then $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_{N_L}$, i.e. the trial point is accepted.

In summary, the DALI-HMC algorithm *runs* along DALI, but it *equilibrates* with the real posterior $P(\boldsymbol{\theta})$.
The chain that this algorithm builds up will peak where the actual posterior likelihood peaks, not where the fiducial for DALI lay. Accordingly, DALI can not only be used to forecast parameter constraints but also to analyze real data.

Chapter 7

Testing for neutrinos in the cosmic microwave background

This chapter presents a work that has been executed with Ruth Durrer during a two-month stay at the cosmology group in Geneva in 2014. The presentation follows closely our publication (Sellentin & Durrer, 2015). This work contains rather standard statistical methods but deals with multiple intricacies of how the statistical results shall be interpreted. Most importantly, a model comparison is carried out and it is shown how conclusions drawn from a data set can be adulterated if an unphysical model is fitted to the data; this is sometimes referred to as 'parameterization bias'.

We addressed the question of how certain we can be that the massive relativistic particles detected by the CMB are neutrinos. The problem is somewhat subtle: from the point of view of a particle physicist, a certain particle type is defined by its properties as its mass, and its interactions with other particles. Standard neutrinos, as known from Earth-bound experiments, have a tiny mass, are highly relativistic, and can be very well approximated as free-streaming. The claim of having detected neutrinos in the CMB is traditionally based on the detection of a particle species which is relativistic and which has a nearly zero mass. These could be neutrinos, if these particles were also free-streaming. But it could be other light and relativistic particles, if they were not free-streaming. We therefore tested explicitely for free-streaming in order to test for another defining property of neutrinos. We have shown that free-streaming particles are preferred over a relativistic perfect fluid with $\Delta\chi^2 \simeq 21$. We have also studied the possibility of replacing the neutrinos by a viscous fluid and found that also a relativistic viscous fluid with either the standard values $c^2_{\mathrm{eff}} = c^2_{\mathrm{vis}} = 1/3$ or best-fit values for c^2_{eff} and c^2_{vis} has $\Delta\chi^2 \simeq 20$ and thus cannot provide a good fit to present CMB data either.

7.1 Neutrinos in the early Universe

The cosmic microwave background (CMB) is the most precious cosmological dataset which we analyse to determine the content of our Universe because its physics is linear and hence

can be very precisely described. Alone and in combination with other data like Type Ia supernovae, the CMB has been used to infer that our Universe is presently dominated by dark energy which may be in the form of a cosmological constant Λ contributing a density parameter of $\Omega_\Lambda \simeq 0.7$, and pressure less matter which is dominated by cold dark matter with $\omega_m = \Omega_m h^2 = \Omega_{\rm cdm} h^2 + \Omega_b h^2 = \omega_{\rm cdm} + \omega_b \simeq 0.14$ where the contribution from baryons is $\omega_b = \Omega_b h^2 \simeq 0.022$, see Planck Collaboration et al. (2015) for the latest values.
Furthermore, there are the photons which make up the CMB and which contribute $\Omega_\gamma h^2 = 2.48 \times 10^{-5}$ and there are cosmic neutrinos. In the standard model of 3 massless neutrino species, they contribute a density parameter of $\Omega_\nu h^2 = 1.69 \times 10^{-5}$. Taking into account neutrino masses, in the minimal model with normal hierarchy and with a maximal neutrino mass of 0.056eV (Forero et al., 2014), one obtains $\Omega_\nu h^2 \simeq 0.5 \times 10^{-3}$.
These are very small numbers. Nevertheless, during the radiation dominated epoch at temperatures above about 1eV, neutrinos and photons are the dominant constituents of the Universe, and the neutrinos contribute a fraction $f_{\rm rad} = \Omega_\nu/(\Omega_\gamma + \Omega_\nu) \simeq 0.4$ to the total energy density of the Universe. At recombination, $z_{\rm dec} \simeq 1100$ they still contribute

$$f_{\rm dec} \simeq \Omega_\nu / \left(\Omega_\gamma + \Omega_\nu + \Omega_m/(1+z_{\rm dec})\right) \simeq 0.1\,,$$

i.e., 10% to the total energy density of the Universe.
The first indication that cosmic neutrinos are really present in the Universe in thermal abundance came from nucleosynthesis calculations. The abundance of primordial helium-4 is very sensitive to the expansion rate at temperature $T_{\rm nuc} \simeq 0.08$MeV, which is determined via the Friedman equation by the energy density of the Universe. At this temperature the energy density is dominated by photons and neutrinos such that these species determine the background expansion rate. The observed helium-4 abundance requires $N_{\rm eff} \simeq 3 \pm 1$ species of neutrinos (Steigman, 2012) in order to provide the correct expansion rate during nucleosynthesis. Somewhat more stringent results have been obtained from the recent Planck data, $N_{\rm eff} = 2.99 \pm 0.4$, see Planck Collaboration et al. (2015).
However, the nucleosynthesis results really only require a relativistic component with the given energy density in order to provide the correct background expansion. However, neutrinos are not only relativistic in the early universe, but they possess additional particle properties to which the background expansion alone is insensitive. Especially, neutrinos are collisionless below $T_\nu \simeq 1.4$MeV corresponding to the redshift $z_\nu \simeq 10^{10}$ where they decouple from the cosmic fluid. The CMB data is not only sensitive to the presence of relativistic components in addition to photons via their contribution to the background, but it also allows us to study their perturbations which are sensitive to additional particle properties of this relativistic species which we can then compare with those expected from true neutrinos.
In the CMB, neutrinos are usually modeled as collisionless particles in order to mimic the neutrino free-streaming, i.e. the standard CMB analyses assume by default that if there exist light relativistic particles, then they will be neutrinos, and their abundance is then constrained by the CMB data. We will lift this question to a higher level and ask *whether* the detected relativistic particles are indeed neutrinos or could we also fit the CMB with a relativistic fluid instead?

There is no doubt that the CMB detects three relativistic species apart from the photons. We will refer to these as 'neutrinos' for brevity, although at first, it is not clear that they are the standard model neutrinos. We shall assume an agnostic point of view and not choose any model of neutrino interaction. We just study whether a perfect fluid or a viscous fluid of relativistic particles can fit the present CMB data. More detailed studies of constraints of neutrino properties with cosmological data, where the latter are given by concrete non-standard particle physics models and where not only CMB but also large scale structure data is considered can be found in Basboll et al. (2009); Archidiacono & Hannestad (2014). Another model for cosmic neutrinos which we shall call the 'viscous free-streaming model' has been studied before (Hu, 1998; Hu et al., 1999; Trotta & Melchiorri, 2005; Archidiacono et al., 2011; Gerbino et al., 2013; Audren et al., 2014), however, we shall argue that this model uses an imperfect parameterization and is neither a perfect nor a viscous fluid. The model adopted in these works therefore seems unnatural to us.

We first compare the standard CMB-anisotropy calculation with a computation where the relativistic particles are modeled as a perfect fluid. We use the package Monte Python (Audren et al., 2013a), using standard cosmological parameters $\omega_b, \omega_{\rm cdm}, h, n_s, A_s, z_{\rm reio}$ along with the Planck nuisance parameters, see Ade et al. (2014).

We fix the background density of the neutrinos to the Planck-2013 best fit since the presence of relativistic particles has been accurately measured and we just want to investigate how precisely the particle properties of the detected relativistic particles can be measured. This means we keep $N_{\rm eff}$ fixed. The primordial helium fraction is then determined by the value of ω_b. We have checked that allowing the mass of the neutrinos to vary, does not make any difference to our results. Therefore, for the sake of simplicity, we show the results for one massive eigenstate with the close to minimal mass of 0.06 eV, the other two eigenstates are treated as massless. We do not adopt any additional priors. Allowing also $N_{\rm eff}$ and with it the primordial helium fraction to vary as a function of $N_{\rm eff}$, does not alter any of the following results. It leads of course to larger error bars on the other parameters.

We find that treating neutrinos as collisionless particles fits the data significantly better than a simple relativistic perfect fluid.

Next we show that neutrinos can neither be modeled as a viscous fluid. We also compare our results with the approach which is found in previous literature (Trotta & Melchiorri, 2005; Archidiacono et al., 2011; Gerbino et al., 2013).

7.2 Neutrinos in the CMB

In standard CMB computations one assumes that neutrinos are massless, freely streaming particles and one solves the Liouville equation for them, see, e.g., Durrer (2008).

$$\begin{aligned}
\dot{\mathcal{N}}_0 + k\mathcal{N}_1 &= 0\,, && (7.1)\\
\dot{\mathcal{N}}_1 + \frac{k}{3}\left[2\mathcal{N}_2 - \mathcal{N}_0\right] &= \frac{k}{3}(\Phi+\Psi)\,, && (7.2)\\
\dot{\mathcal{N}}_\ell + \frac{k}{2\ell+1}\left[(\ell+1)\mathcal{N}_{\ell+1} - \ell\mathcal{N}_{\ell-1}\right] &= 0\,, \quad \ell > 1\,. && (7.3)
\end{aligned}$$

Here $\mathcal{N}_\ell$ is the ℓth multipole moment of the energy integrated neutrino distribution function in Fourier space and k is the wave number, while Φ and Ψ are the Bardeen potentials. The moments 0 to 2 are related to the neutrino density perturbation, δ_ν, the potential of the velocity perturbation, V_ν,and the anisotropic stress, Π_ν, in longitudinal gauge by

$$\begin{aligned} \delta_\nu &= 4(\mathcal{N}_0 + \Phi)\,, && (7.4)\\ V_\nu &= 3\mathcal{N}_1\,, && (7.5)\\ \Pi_\nu &= 12\mathcal{N}_2\,. && (7.6) \end{aligned}$$

One truly only needs these first three moments of the distribution function since only they enter the energy momentum tensor which couples to the gravitational field and affects the evolution of the CMB photons. Nevertheless, in the Liouville equation each mode $\mathcal{N}_\ell$ is coupled by free-streaming to $\mathcal{N}_{\ell+1}$ and $\mathcal{N}_{\ell-1}$ such that power between the different multipoles is exchanged. Especially, power that would enter δ_ν, V_ν and Π_ν via $\mathcal{N}_0, \mathcal{N}_1$ and $\mathcal{N}_2$ without free-streaming, can be propagated into higher order multipoles by free-streaming, and thereby vanish from the energy-momentum tensor and the Einstein equations. Therefore, in order to obtain $\mathcal{N}_0$, $\mathcal{N}_1$ and $\mathcal{N}_2$ with sufficient precision in case of free-streaming, one usually solves the neutrino hierarchy up to $\ell^\nu_{\rm max} \sim 10 - 20$. The minimum of 10 results from the need to avoid problems from so called numerical 'reflections' when solving the differential equations (7.3) with a finite differences method.

Below we shall consider treating neutrinos as a relativistic perfect fluid, which corresponds to cutting the hierarchy at $\ell^\nu_{\rm max} = 1$ or as a relativistic viscous fluid, which corresponds to cutting the hierarchy at $\ell^\nu_{\rm max} = 2$. We shall also consider a viscous fluid with arbitrary sound speed $c^2_{\rm eff}$ and viscosity $c^2_{\rm vis}$. This corresponds to cutting the hierarchy at $\ell^\nu_{\rm max} = 2$ and replacing (7.1) to (7.3) by the following system of equations, see Hu (1998); Hu et al. (1999); Trotta & Melchiorri (2005):

$$\begin{aligned} \dot{\mathcal{N}}_0 + k\mathcal{N}_1 &= \mathcal{H}(1 - 3c^2_{\rm eff})\mathcal{N}_0 && (7.7)\\ \dot{\mathcal{N}}_1 + \frac{k}{3}\left[2\mathcal{N}_2 - 3c^2_{\rm eff}\mathcal{N}_0\right] &= -\mathcal{H}(1 - 3c^2_{\rm eff})\mathcal{N}_1 + \frac{k}{3}(3c^2_{\rm eff}\Phi + \Psi) && (7.8)\\ \dot{\mathcal{N}}_2 + k\left[\frac{3}{5}\mathcal{N}_3 - 3c^2_{\rm vis}\frac{2}{5}\mathcal{N}_1\right] &= 0 && (7.9) \end{aligned}$$

and the higher multipoles stream freely with

$$\dot{\mathcal{N}}_\ell + \frac{k}{2\ell + 1}\left[(\ell + 1)\mathcal{N}_{\ell+1} - \ell\mathcal{N}_{\ell-1}\right] = 0\,, \quad \ell > 2\,. \tag{7.10}$$

In addition, the fact that the perturbations have to be evaluated in the rest frame of the fluid, leads to subtle changes of Eqs. (7.7-7.9) as described in (Audren et al., 2014). We include these for our modeling of the massive and massless neutrinos because otherwise the results would be gauge dependent. In our viscous fluid model we differ from the treatment in the above mentioned papers by cutting the hierarchy at $\ell^\nu_{\rm max} = 2$ in the and we consistently also set $\mathcal{N}_3 \equiv 0$.

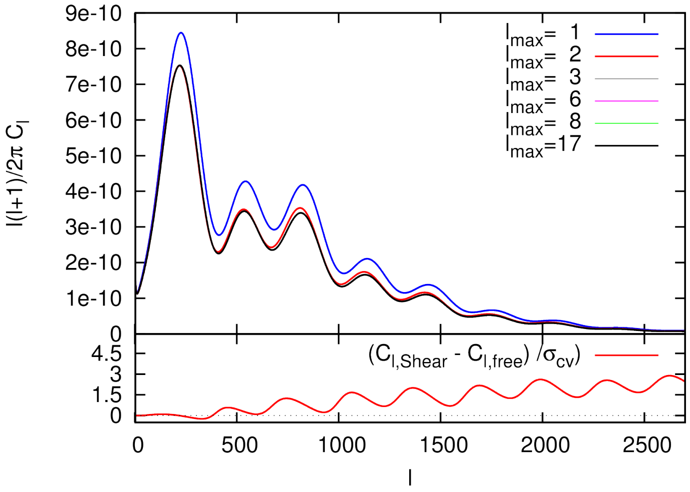

Figure 7.1: The temperature anisotropy spectra at fixed cosmological parameters for different values of the cutoff in the neutrino hierarchy, $\ell^{\nu}_{\rm max} = 1, 2, 3, 6, 8, 17$. For $\ell^{\nu}_{\rm max} > 2$ the result changes very little. In the bottom panel, the difference between C_ℓ for the viscous neutrino fluid ($\ell^{\nu}_{\rm max} = 2$) and free-streaming neutrinos, is compared to cosmic variance, which roughly corresponds to the Planck error out to $\ell \simeq 2000$. For low ℓ cosmic variance does not allow to discriminate between the viscous fluid model and free-streaming neutrinos, but for high ℓ the difference between these two models is up to three times larger than cosmic variance $\sigma_{cv} = C^{\rm free}_\ell \cdot \sqrt{2/(2\ell+1)f_{sky}}$. (Note that here $\ell^{\nu}_{\rm max}$ denotes the maximal neutrino multipole while ℓ refers to the CMB multipoles.)

In Fig. 7.1, we show the CMB anisotropy power spectrum for fixed cosmological parameters by modeling the neutrino hierarchy up to $\ell^{\nu}_{\rm max}$. Already for $\ell^{\nu}_{\rm max} = 2$ the difference between the standard calculation setting $\ell^{\nu}_{\rm max} = 17$ becomes very small. Nevertheless, as is visible from the lower panel, for a cosmic variance limited experiment, like Planck for $\ell \lesssim 2000$, this difference for $\ell^{\nu}_{\rm max} = 2$ is highly significant.

In Fig. 7.2, we show how the fit to the data improves as a function of the maximally allowed neutrino multipole: truncating at $\ell^{\nu}_{\rm max} = 1, 2$ leads to the deteriorated fits of the ideal and relativistic viscous fluid. Truncating at $\ell^{\nu}_{\rm max} = 5$ leads actually to a slightly better fit than solving the Boltzmann hierarchy up to $\ell^{\nu}_{\rm max} = 17$. At the same time, it is evident from Fig. 7.4 that the best-fit values and the 1σ contours of the cosmic parameters do not change when cutting the hierarchy anywhere between $\ell^{\nu}_{\rm max} = 4$ and $\ell^{\nu}_{\rm max} = 17$. The negative $\Delta\chi^2 = -1.29$ for $\ell^{\nu}_{\rm max} = 5$ therefore does not seem to stem from physics but might be due to numerical inaccuracies and to the modeling of the experimental uncertainties (see also the discussion about model 'evidence'). But even if truncating at different $\ell^{\nu}_{\rm max}$ may lead to typical changes in $\Delta\chi^2$ on the order of unity, the conclusion that the ideal and the standard viscous fluid are worse fits than free-streaming neutrinos which we shall draw, remains valid since their $\Delta\chi^2$ is much higher.

We have investigated whether neutrinos can be modelled by a relativistic perfect or viscous fluid. For this, we have replaced the massless and massive neutrinos by a (relativistic)

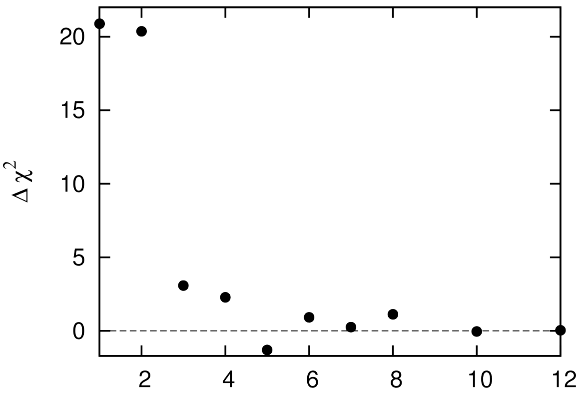

Figure 7.2: The improvement of the fit, $\Delta\chi^2 = \chi^2(\ell^\nu_{\text{max}}) - \chi^2(\ell^\nu_{\text{max}} = 17)$, as a function of the maximal considered neutrino multipole ℓ^ν_{max} in the Liouville equations for the neutrinos. Free-streaming neutrinos correspond to $\Delta\chi^2 = 0$.

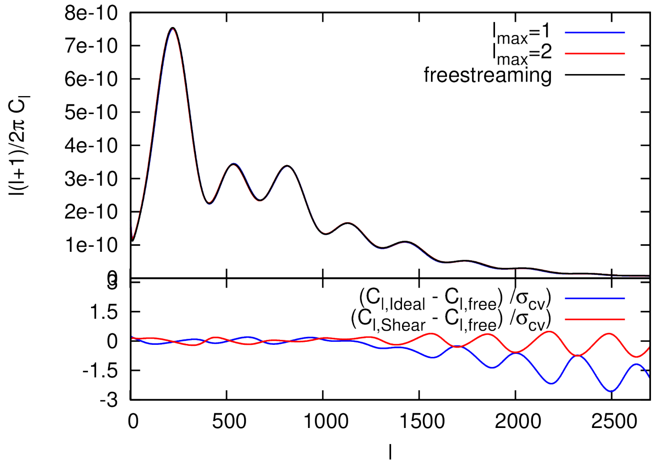

Figure 7.3: The temperature anisotropy spectra for best-fit parameters modeling neutrinos as a perfect fluid (blue), a relativistic viscous fluid (red) and standard free-streaming neutrinos (black) are shown. The difference of the best-fit spectra is not visible by eye. However, the bottom plot shows the difference in units of the cosmic variance error and here it becomes clear, that the Planck experiment can distinguish the spectra.

perfect fluid or a (relativistic) viscous fluid and run the modified CMB code CLASS (Blas et al., 2011; Audren et al., 2013b) in combination with Monte Python (Audren et al., 2013a) to find best-fit values of the standard cosmological parameters from the Planck data. In Fig 7.3 we compare the spectra obtained in this way with the spectrum from free-streaming neutrinos and in Fig. 7.4 we show the best fit parameters. By eye, the curves look identical. But when considering the difference in units of the cosmic variance, it becomes clear that a cosmic variance limited experiment like Planck can measure the small difference.
Not only are several of the cosmological parameters significantly different, see Fig. 7.4, but

the fit is also much worse. The $\Delta\chi^2$ for both fluid approximations increases by:

$$\Delta\chi^2_{\text{ideal}} \approx 21\,, \qquad \Delta\chi^2_{\text{visc}} \approx 20\,. \tag{7.11}$$

Let us formulate this in terms of the Bayes factor, K (see Jeffreys (1961a); Kass & Raftery (1995)) which indicates whether model M_1 (for us free-streaming neutrinos) is favoured over model M_2 (for us either ideal fluid or viscous fluid neutrinos). The Bayes factor is defined as (Jeffreys, 1961a)

$$K \equiv \frac{P(D|M_1)}{P(D|M_2)} = \frac{P(M_1|D)P(M_2)}{P(M_2|D)P(M_1)} = \frac{P(M_1|D)}{P(M_2|D)}\,. \tag{7.12}$$

The last equality is due to the fact that in our case, both models have the same parameters so the model spaces are identical. The difference is only that in model 1 the neutrinos are the standard free-streaming neutrinos while in model 2 they are an ideal or a viscous fluid. Since we want to test exactly this hypothesis (or rather we want to see whether the CMB data is sensitive to this hypothesis) we cannot give model 2 a smaller model probability. We therefore set $P(M_1) = P(M_2)$. Here $P(M|D)$ is the probability of a model given the data while $P(D|M)$ is the probability of the data given the model. The prior, $P(D)$ drops out in the ratio K. But the $P(M|D)$ are exactly the likelihoods which we determine in our MCMC code, so that

$$2\log(K) = \Delta\chi^2\,. \tag{7.13}$$

According to Ref. Kass & Raftery (1995), while $-2 < 2\log(K) < 2$ is ‘not worth mentioning’ a value $2 < 2\log(K) < 6$ can be interpreted as ‘positive’ but not strong evidence while $10 < 2\log(K)$ is ‘decisive’ evidence (see also Appendix B in Jeffreys (1961a)). First we conclude that all the fluctuations in $\Delta\chi^2$ for $\ell^\nu_{\text{max}} \geq 4$ are ‘not worth mentioning’. The contrary holds for $\ell^\nu_{\text{max}} \leq 2$, in this case the evidence in favour of the free-streaming model is truly ‘decisive’.

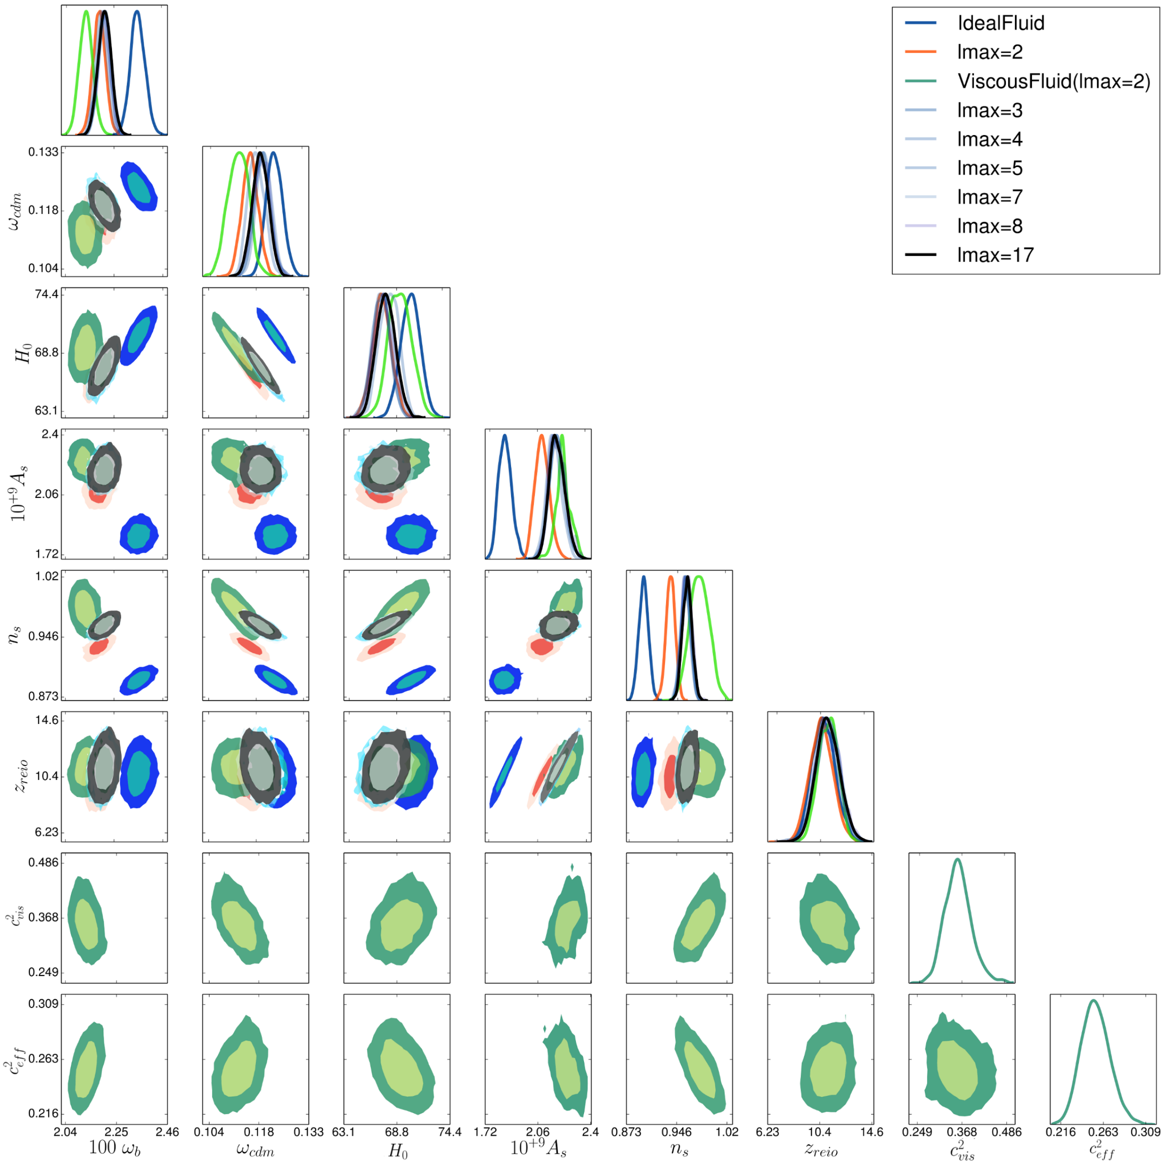

Figure 7.4: The best fit parameters for neutrinos modeled as a perfect fluid (blue), as a relativistic viscous fluid (orange), as a viscous fluid with arbitrary sound speed c^2_{eff} and viscosity c^2_{vis} (green), and as standard free-streaming neutrinos (black) are shown. We also show free-streaming neutrinos with different $\ell^{\nu}_{\mathrm{max}}$ as indicated in the legend in light blue to grey shades. The best-fit values of several parameters for the perfect fluid and the free-streaming model differ significantly. The best-fit values of most parameters for the viscous fluid and the free-streaming model are similar, they all agree within 1.5σ apart from n_s which for the relativistic viscous fluid model differs by more than 2σ. Truncating the Boltzmann hierarchy for the neutrinos at other maximally allowed $\ell \geq 3$ (indicated in different shades of blue) leads to parameter constraints that are indistinguishable from the standard Planck fit, consistent with Fig. 7.1. Allowing also N_{eff} to vary does not change these results, as can be seen in Fig. 7.7.

This shows that cosmic neutrinos cannot be modelled either by a relativistic perfect fluid or by a viscous fluid. Nevertheless, one may be surprised that the ideal fluid model is not much more strongly excluded than the viscous fluid. Also, our value of $\Delta\chi^2_{\rm ideal}$ is significantly smaller than other values published in the literature, see e.g. Archidiacono & Hannestad (2014). One reason for this at first surprising finding is that we include all Planck nuisance parameters in our MCMC analysis and allow for the priors suggested by the Planck collaboration. As an illustration of how this affects the results we show the 1d likelihoods of the nuisance parameters in Fig. 7.6. As one can see there, the fluid models prefer different values for some of the nuisance parameters, especially the best fit amplitudes of the inferred kinetic SZ-effect, A_{ksz} and of the cosmic infrared background at 146GHz, A_{cbi143} are very different for the ideal fluid model than for the viscous fluid or free-streaming model, but as is evident from Fig. 7.6, these parameters are very badly constrained by the data. The increase in $\Delta\chi^2_{\rm ideal}$ in Ref. Archidiacono & Hannestad (2014), however is not only due to a less conservative CMB analysis but also to the inclusion of large-scale structure data. We avoid the inclusion of large-scale structure data here, in order to keep the properties of neutrinos at high energies (i.e. in the early universe) separate from late-time cosmology effects, where properties of neutrinos can also be used to explain dark energy phenomenology (Ayaita et al., 2013). Furthermore, the topic of this work is whether the free-streaming of neutrinos can be detected in the CMB.

To test the importance of the nuisance parameters we have also run a chain where we fixed them to their best-fit values in the free-streaming model. This, of course, reduces the size of the model space significantly. In this case the increase in $\Delta\chi^2$ for both, the ideal and the viscous fluid model in comparison to standard LCDM is larger namely

$$\Delta\chi^2_{\rm ideal} \approx 32\,, \qquad \Delta\chi^2_{\rm visc} \approx 30\,, \tag{7.14}$$

with fixed nuisance parameters.

However, these increases in $\Delta\chi^2$ are not trivial to interpret: fixing the nuisance parameters to the best-fit values of the Planck collaboration is a form of including knowledge about which parameter values the current CMB data prefer, when fitted with free-streaming neutrinos. Therefore, the self-consistent values are those given in Eq. (7.11).

Let us also compare this analysis with previous work (Trotta & Melchiorri, 2005; Archidiacono et al., 2011; Gerbino et al., 2013) on neutrino clustering properties, where a somewhat different standpoint has been taken. There, eqs. (7.1) to (7.3) are replaced by eqs. (7.7) to (7.10). A similar, non-perfect-fluid treatment has already been suggested in Hu (1998); Hu et al. (1999). However, eqs. (7.7) to (7.10) describe neither a perfect nor an imperfect fluid since the higher moments, $\ell \geq 3$, are not damped by collisions but evolve like those of free-streaming particles. We dub this mixture model 'viscous free-streaming model'.

The advantage of the viscous free-streaming model is that it is 'nested' inside the standard model of free-streaming neutrinos with two additional parameters which take the values $c^2_{\rm eff} = c^2_{\rm vis} = 1/3$ in the standard model. The viscous free-streaming can then be regarded as a phenomenological model that allows for a transition between free-streaming and viscosity. Previous works, especially Gerbino et al. (2013) have found that the preferred values of these

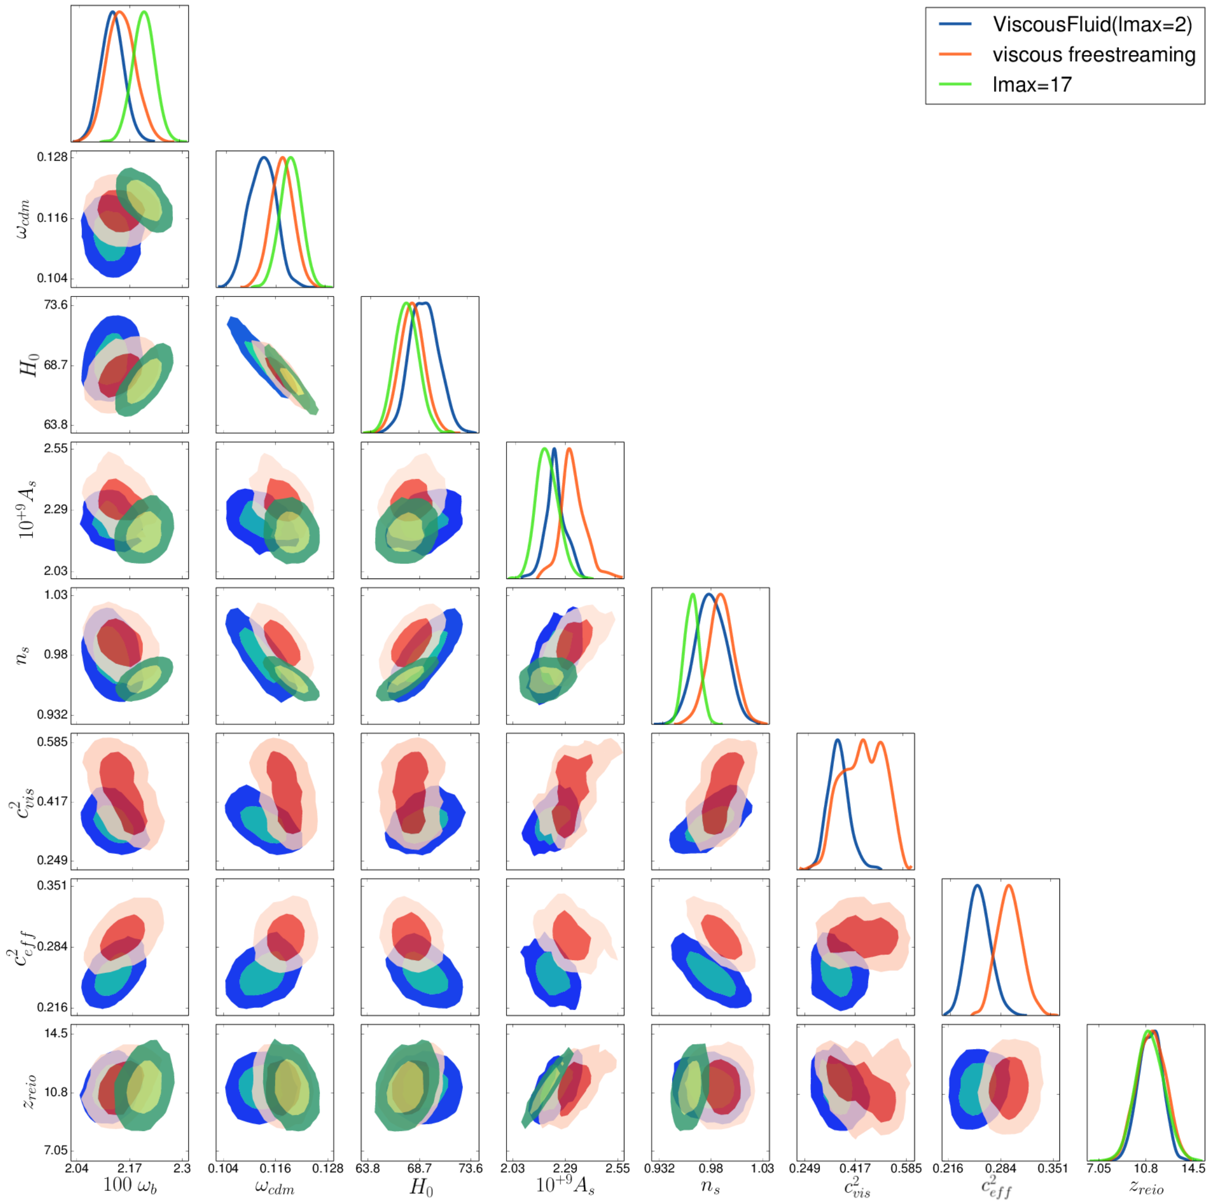

Figure 7.5: The best-fit parameters for neutrinos modeled as a viscous fluid with variable $c^2_{\rm eff}$ and $c^2_{\rm vis}$ and $\ell^{\nu}_{max=2}$ (blue), for neutrinos modeled as a viscous free-streaming fluid with $\ell^{\nu}_{max} = 17$(orange), and for standard free-streaming neutrinos (green) are shown. The best-fit values of most parameters for the two different viscosity models differ by about one standard deviation.

parameters are indeed close to the standard relativistic ones. Nevertheless, the physical meaning of $c_{\rm eff}$ and $c_{\rm vis}$ remains unclear in the viscous free-streaming since only the evolution of the first and second moment but not higher moments are affected by collisions in this model. This seems unphysical to us and we are not aware of a physical example which leads to such a behaviour. Viscosity damps out all higher moments and thereby inhibits free-streaming of all higher moments. Nonetheless, it has been found that the viscous free-streaming model succeeds in fitting also the latest Planck Data (Planck Collaboration et al., 2015), preferring again the standard values for $c_{\rm eff}^2 = c_{\rm vis}^2 = 1/3$ which represent in fact the only case in which the inconsistency between viscosity and free-streaming vanishes. It is not too surprising that the data prefer the only parameter values for which the unphysical model turns into a physically sound one. For a more detailed discussion of how the cutting of ℓ-modes and the effective fluid parameters $c_{\rm eff}^2$ and $c_{\rm vis}^2$ map to particle properties, see Oldengott et al. (2015).

However, modelling a true viscous fluid, not only requires the introduction of the new parameters $c_{\rm eff}^2$ and $c_{\rm vis}^2$ but also either cutting the neutrino hierarchy at $\ell = 2$ or describing the evolution of the higher moments with a collision term, as e.g. in Basboll et al. (2009). Accordingly, in our model of a viscous fluid we set $\mathcal{N}_\ell \equiv 0$ for all $\ell \geq 3$ and fit for $c_{\rm eff}^2$ and $c_{\rm vis}^2$. As we have discussed above, this model with $c_{\rm eff}^2 = c_{\rm vis}^2 = 1/3$, i.e., the relativistic viscous fluid, provides a bad fit to the observed CMB anisotropies. Before concluding that the three relativistic particles in the CMB are indeed free-streaming neutrinos, we need, however, to check whether another value of $c_{\rm eff}^2$ and $c_{\rm vis}^2$ might provide a better fit.

Introducing two new free parameters, will of course improve the fit, and we find that the difference in χ^2 for the best fit with respect to the relativistic viscous fluid is: $\Delta\chi^2 \approx -0.58$, with the negative sign indicating an improvement of the fit. The standard value of $c_{\rm eff}^2 = 1/3$ is excluded at 3σ whereas $c_{\rm vis}^2 = 1/3$ is compatible within one standard deviation. Nevertheless, in this case, the model space of the new model 2 is increased which enhances the Bayes factor in favour of model 1. A rule of thumb is that each new parameter has to improve $\Delta\chi^2$ by at least 1 in order to compensate the 'Occam's razor factor' $P(M_2)/P(M_1)$ in Eq. (7.12). Hence the modest improvement of the best fit by 0.58 after the introduction of two new parameters, leads to the conclusion that these two new parameters are not justified.

For the viscous free-streaming model, we find that the fit improves by $\Delta\chi^2 \approx -3.72$ with respect to free-streaming neutrinos. This improvement of the fit is somewhat stronger than what is usually expected when adding two additional parameters, $\Delta\chi^2 \approx -2$, but this can be a usual statistical fluctuation in $\Delta\chi^2$ as it leads to a $2\log(K) \simeq \Delta\chi^2 + 2 = -0.72$, hence 'not worth mentioning'.

In Fig. 7.4 we have compared the parameter values obtained by replacing neutrinos by a perfect fluid, a relativistic viscous fluid or by a viscous fluid with arbitrary effective sound speed $c_{\rm eff}^2$ and viscosity, $c_{\rm vis}^2$ with the results for standard neutrinos. For completeness, we compare in Fig. 7.5 also the parameter constraints for variable viscosity parameters $c_{\rm eff}^2 = c_{\rm vis}^2$ for a viscous free-streaming fluid and for a true viscous fluid that cannot build up moments with order higher than $\ell = 2$.

7.3 Likelihoods of the nuisance parameters and $N_{\rm eff}$

In Fig 7.6 we show the marginalised 1-parameter distribution of the nuisance parameters used in our analysis of the Planck data. We use the priors as suggested by the Planck Collaboration (Ade et al., 2014). As one sees in the figure, several of these parameters are not well constrained by the data. Also, in the perfect fluid model several nuisance parameters take quite different values than for the viscous fluid or the free-streaming model. This allows the perfect fluid model to fit the data not significantly worse than the viscous fluid model. Fixing these parameters leads to a somewhat larger value of $\Delta\chi^2_{\rm ideal} - \Delta\chi^2_{\rm visc} \simeq 2$, see Eq. (7.14).

In Fig. 7.7 we compare the two-dimensional likelihoods of both, the ideal fluid model and the viscous fluid model with and without varying $N_{\rm eff}$. Even though the error bars of course increase when including this additional parameter, the main results are unchanged. While the peak value of $N_{\rm eff}$ for the viscous model changes by less than one standard deviation, the ideal fluid would actually prefer a lower $N_{\rm eff}$. Note also that the width of the distribution of $N_{\rm eff}$ in the viscous fluid model is very similar to the standard one, (see Ade et al. (2014), Fig. 21) while for the ideal fluid model $N_{\rm eff}$ is somewhat more constrained. Nevertheless, the increase in $\Delta\chi^2$, which is the main point of this study, remains stable.

7.4 Conclusions

We have studied how neutrinos are detected in the CMB and have shown that they are not only relevant as additional relativistic degrees of freedom, but CMB anisotropies and polarisation are also very sensitive to their clustering properties. While the Planck-2013 data is in good agreement with free-streaming neutrinos, it cannot be fitted by neutrinos modeled as a relativistic perfect fluid. The best-fit model with perfect fluid neutrinos leads to a $\Delta\chi^2 = 21$ with respect to the best fit free-streaming neutrinos. Even including anisotropic stress, i.e. allowing for a relativistic viscous fluid cannot fit the data. The increase in χ^2 with respect to the best-fit models with free-streaming neutrinos is $\Delta\chi^2 \approx 20$. Using the evidence scale introduced by Jeffrey (Jeffreys, 1961a), this result can be considered as 'decisive evidence' for neutrino free-streaming in the CMB.

The fit can be improved only slightly when allowing arbitrary values for the effective sound speed and the viscosity, $c^2_{\rm eff}$ and $c^2_{\rm vis}$. But including these two additional parameters and truncating the neutrino Boltzmann hierarchy at $\ell^{\nu}_{\rm max} = 2$ in order to consistently model a viscous fluid, the fit improves by $\Delta\chi^2 = -0.58$ w.r.t. the best relativistic viscous fluid model with fixed values $c^2_{\rm eff} = c^2_{\rm vis} = 1/3$. Therefore, the introduction of these additional parameters is not favoured.

Using observations of the cosmic microwave background anisotropies and polarization we have not only found that there are 3 species of light particles, but we can also infer that these relativistic particles are freely streaming, making it plausible that these particles are indeed neutrinos. These results are robust under the variation of $N_{\rm eff}$ as an additional parameter and $N_{\rm eff}$ is found to peak at the standard value of three neutrino species.

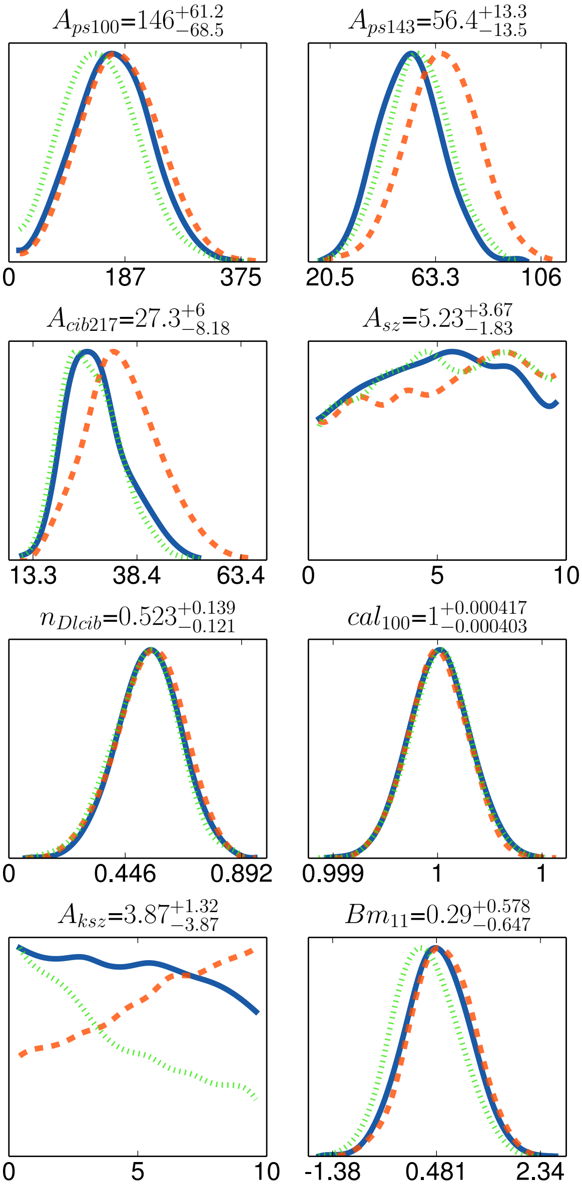

Figure 7.6: The 1d likelihoods of the Planck nuisance parameters for free-streaming neutrinos (solid, blue), the viscous fluid model (dotted, green) and the ideal fluid model (dashed, orange). The best-fit values indicated on top of each panel are those of the viscous model. Especially the best-fit amplitudes of the kinetic SZ-effect, A_{ksz} and of the cosmic infrared background at 146GHz, A_{cbi146} are not well constrained and are very different for the ideal fluid model and for the viscous fluid or free-streaming model.

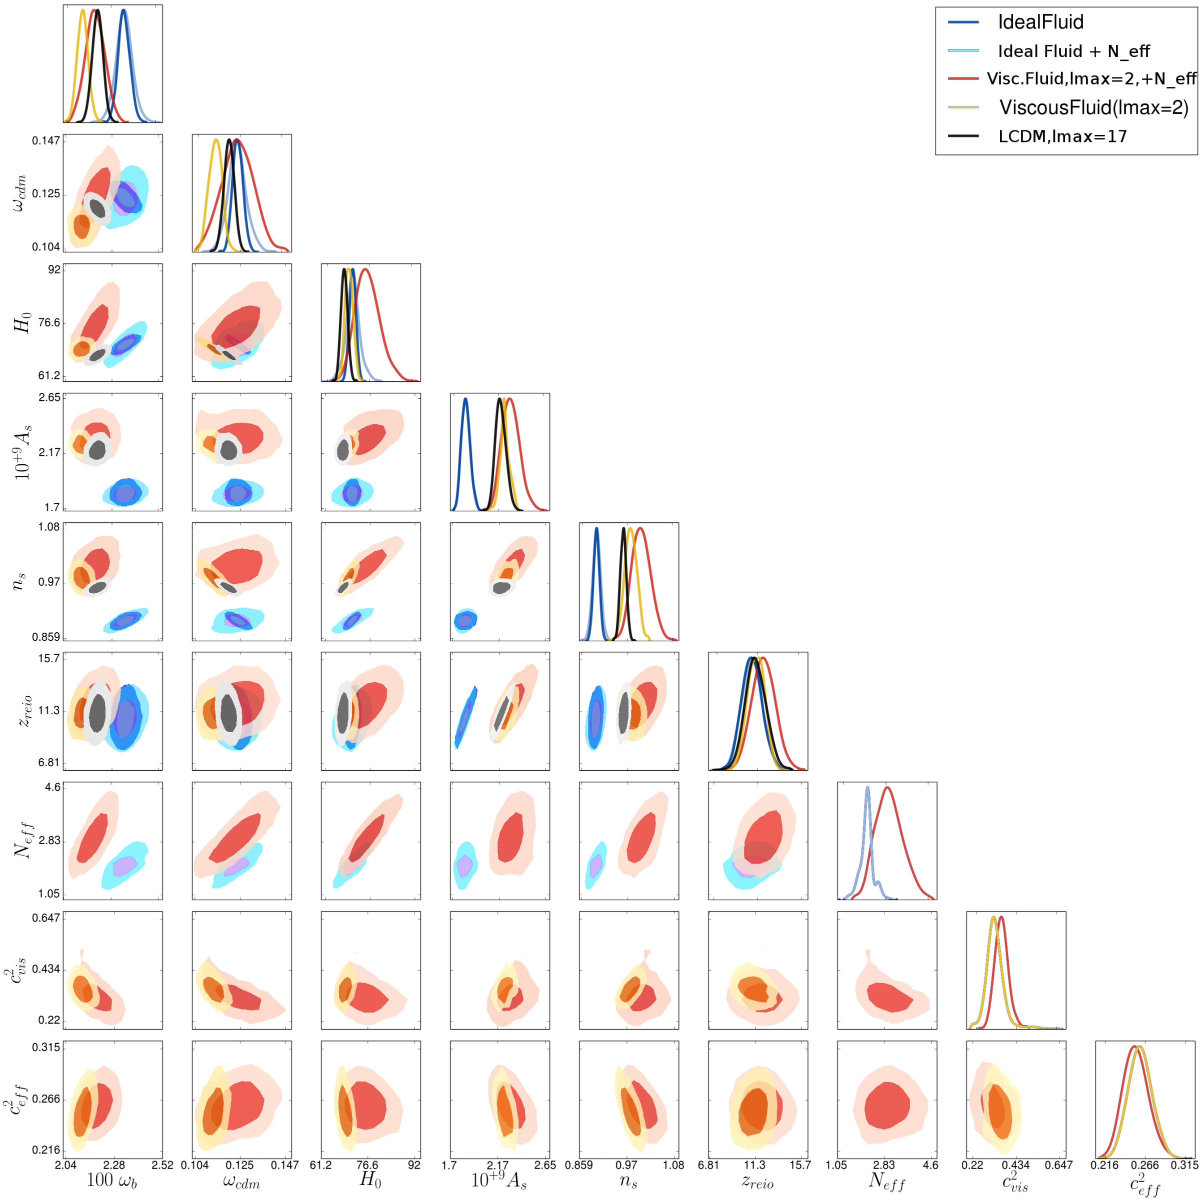

Figure 7.7: Comparing the results of Fig. 7.4, with the results for a varying N_{eff}. The inclusion of the parameter N_{eff} does not significantly alter the best-fit values of the other cosmological parameters, and N_{eff} itself peaks at around the standard value of three neutrino species. The 1-parameter likelihoods of the ideal fluid model remain virtually unchanged. Those of the viscous fluid model change mainly by widening. The best-fit parameters move by less than one standard deviation.

Chapter 8

DALI: Derivative Approximation for LIkelihoods

In this chapter, our new method DALI (Derivative Approximation for LIkelihoods) is first derived, and then applications of it are presented. In section 8.1, the general framework of DALI is described. In section 9.3, it is described how DALI can not only be used for forecasting, but also for the analysis of real data: the DALI series allows to drive a Hamilton Monte Carlo sampler, which can read in real data (or mock data), such that the unapproximated likelihood of real (or mock) data can be speedily sampled. Section 9.4 contains considerations about the applicability and validity of the Cramer-Rao inequality. This chapter follows rather loosely the publications Sellentin et al. (2014); Sellentin (2015); Sellentin & Schäfer (2016) as further insights were gained between publication of these papers and the submission of this thesis.

8.1 The non-Gaussian extension of the Fisher matrix

Evaluating a multidimensional likelihood can be a computationally costly procedure. If speed matters, often a good approximation of the likelihood is required. A widely used approximation of likelihoods is the Fisher matrix approximation, which singles out the Gaussian part of a likelihood (Tegmark et al., 1997). Because many analytical results for Gaussians are available, such as the position of the 1-σ confidence contours and higher-order equivalents, the Fisher matrix approximation is fast to evaluate. It has also become widely used as it allows for the easy computation of figures of merit, simple determinants of the matrix elements and manipulations thereof, that can be used to evaluate the expected performance of an experiment, for example as introduduced to dark energy research by Albrecht et al. (2006).

The alternatives to the Gaussian approximation are grid evaluations of the likelihood, or sampling techniques such as Monte Carlo Markov Chains (MCMC), Nested Sampling (Audren et al., 2013a; Allison & Dunkley, 2014; Skilling, 2004), and Population Monte Carlo that uses iterative updates of a mixture model to capture non-Gaussianities (Kilbinger et al., 2010; Wraith et al., 2009). These methods tackle the challenge of characterising non-Gaussian likelihoods by using sophisticated numerical algorithms. Gram-Charlier and Edgeworth-type expansions can also be used to capture non-Gaussianities, but suffer from regions in the parameter space where the approximated likelihood turns negative, thereby violating the Kolmogorov axioms for a probability (Cramer, 1946).

Nonetheless, likelihood approximations are urgently needed throughout the physical sciences, whenever evaluating a full likelihood is numerically too costly, e.g. when forecasting parameter constraints of a future experiment, where many different configurations need to be simulated, see e.g. Pillepich et al. (2012); Laureijs et al. (2011a). A quick check of the resulting likelihood is also desirable when optimizing a data analysis pipeline, or when establishing novel observables and testing how precisely they can constrain model parameters, see e.g. Chantavat et al. (2014).

Non-Gaussian likelihood approximations, that maintain positive definiteness and normalizability, whilst rivaling the Fisher matrix in terms of speed, have recently become a focus of research. Transformations of the likelihood to Gaussianity are one way of tackling this problem (Joachimi & Taylor, 2011). During this doctorate, we have developed another approximation, named DALI: Derivative Approximation for LIkelihoods. It was published in the papers Sellentin et al. (2014) and Sellentin (2015) and the application to weak lensing has been published in Sellentin & Schäfer (2016). DALI assumes that the data are Gaussianly distributed but is else independent of the application, i.e. it would work for all observables to which it would be specified.

This is achieved by building on the Fisher matrix, but expanding the posterior to higher orders. If the posterior P really is Gaussian in the parameters, the higher order derivatives of $\log P$ will be zero, such that the extended method falls back onto the Fisher matrix and nothing is lost. If the higher order derivatives are non-zero, a gain in shape fidelity is to be expected. As many posteriors have a smooth shape and resemble often a “surrealistic” version of an ellipse, i.e. the ellipses are slightly curved, flexed or otherwise distorted, already the inclusion of just a few higher-order derivatives promises good improvements. We begin the presentation of DALI by discussing why Gaussian data can produce non-Gaussian parameter likelihoods.

8.2 Gaussianity

DALI assumes a data set $\boldsymbol{X}$ with Gaussian errors, leading to the unapproximated likelihood

$$L(\boldsymbol{X}|\boldsymbol{p}) = \frac{1}{\sqrt{(2\pi)^d|\boldsymbol{C}|}} \exp\left(-\frac{1}{2}(\boldsymbol{X}-\boldsymbol{\mu})^T\boldsymbol{C}^{-1}(\boldsymbol{X}-\boldsymbol{\mu})\right), \tag{8.1}$$

where $\boldsymbol{p}$ is a vector of p parameters. The mean of the data $\boldsymbol{\mu}$ and the covariance matrix $\boldsymbol{C}$ are predicted by a parameterized physical model and can in general both depend on the p parameters. These parameters shall be constrained by maximizing the likelihood using the data which is collected in the data vector $\boldsymbol{X}$. The number of data points is d and $|\boldsymbol{C}|$ is the determinant of the covariance matrix. The covariance matrix is given by

$$\boldsymbol{C}(\boldsymbol{p}) = \langle (\boldsymbol{X} - \boldsymbol{\mu})\,(\boldsymbol{X} - \boldsymbol{\mu})^T \rangle, \tag{8.2}$$

such that for a linear model, the parameters enter already quadratically in the covariance matrix.

The corresponding log-likelihood $\mathcal{L} = -\ln(L)$ of the Gaussian Eq. (8.1) is

$$\mathcal{L} = \frac{1}{2}\mathrm{Tr}\left(\ln(\boldsymbol{C}) + \boldsymbol{C}^{-1}\langle \boldsymbol{D} \rangle\right), \tag{8.3}$$

where we neglected the 2π factors of the normalization, and where $\boldsymbol{D} = (\boldsymbol{X} - \boldsymbol{\mu})(\boldsymbol{X} - \boldsymbol{\mu})^T$ is the data matrix. Angular brackets denote averaging over realizations of the data.

The numerical cost of evaluating this likelihood will increase with the number of data points, the complexity of calculating the model predictions $\boldsymbol{\mu}$ and the calculation of the covariance matrix under variation of the parameters. In case of Bayesian inference, the likelihood could be updated to a posterior by multiplying with priors and normalizing by the corresponding evidence. Consequently, we will in the following use the terms likelihood and posterior rather interchangeably.

The assumption of Gaussian errors is not a severe constraint, since due to the central limit theorem, all data that stem from a distribution of finite variance, can be rebinned into a data set with Gaussian errors – if enough data points are available. However, having Gaussian errors in the *data* space does not mean that the resulting likelihood will be Gaussian in the *parameter* space. This is only exactly fulfilled if all parameters are linear, and if there are no degeneracies between the parameters. Therefore, the mathematical tools available to exploit Gaussian likelihoods, such as their analytical marginalization over nuisance parameters, cannot automatically be used in the parameter space. However, a Gaussian likelihood can also be expected if the data set is constraining enough, such that essentially a linear Taylor approximation of the model and the covariance matrix around the best-fit point is sufficient. This explains why the Fisher matrix has become so popular in forecasting the performance of precision experiments, which were designed to tightly constrain targeted parameters.

In contrast, achieving extremely constraining data with a new experiment cannot be expected by default if for example extensions to a standard model are to be investigated and new parameters measured for the very first time. If the forecasted data is not expected to be extremely constraining, the likelihood will not be peaked so sharply around the best-fit that a linear Taylor approximation of the model, and the covariance matrix alone may not be good enough. This already hints at why the following non-Gaussian likelihood approximation needs to build on higher order derivatives.

We specialize to two opposite cases: First, if the model predicts the mean of the data as a function of some parameters, $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{p})$, and the covariance matrix is parameter independent. Then, we specialize to the opposite case of the model predicting the covariance as a function of the parameters, $\boldsymbol{C} = \boldsymbol{C}(\boldsymbol{p})$, and the mean of the data being constant.

The latter occurs for example, when the mean is zero but fluctuations around that mean can be of different amplitudes, and this is encoded in the covariance. Examples are a measurement of pure noise, which clearly has mean zero, but where the covariance of the noise depends on parameters. Another example is any kind of mode decomposition, where again it is clear that a mode has mean zero. A cosmological example is the galaxy power spectrum, which arises from density fluctuations around the cosmic mean value, and where the mean overdensity must be zero, due to mass conservation. The power spectrum can then be used as the covariance in the following framework, where it is the covariance of the Fourier amplitudes of the overdensity field.[1]

8.3 Problems when approximating likelihoods

Approximating a likelihood is more complicated than approximating a more general function because one typically wishes the likelihood to be positive semi-definite at all orders; otherwise negative probabilities occur, which are non-sensical. Positive semi-definiteness is a strong constraint and not automatically fulfilled by a usual Taylor series approximation of the likelihood. For example, Taylor approximating a standard normal distribution of some variable x yields,

$$\exp(-x^2) = 1 - x^2 + \frac{1}{2}x^4 + \mathcal{O}(x^5). \tag{8.4}$$

If truncated at second order, this approximation becomes negative at 2-σ from the best-fit, or begins rising to infinity at about 2-σ when truncated at fourth-order. This divergence makes the likelihood approximation not normalizable, such that no measure for relative likelihoods can be defined. Both, second- and fourth-order approximation of the standard normal distribution therefore violate defining properties of a likelihood. Obviously, a continuation of the Taylor approximation Eq. (8.4) to very high orders would remedy both of these issues but this would be a cumbersome approach. It is well known that Taylor approximating the log-likelihood instead, reconstructs the Gaussian likelihood much more quickly

$$\exp(-x^2) = \exp(-\mathcal{L}) = \exp(-\mathcal{T}(\mathcal{L})), \tag{8.5}$$

where $\mathcal{T}(\mathcal{L})$ denotes the Taylor series of the log-likelihood. If this Taylor series is evaluated at the maximum of the standard normal distribution then already the first- and second-order terms of this series recovers the Gaussian likelihood completely, and all higher orders of the series are identically zero. The approximation schemes Eq. (8.4) and Eq. (8.5) are both mathematically valid ways of approximating the standard normal distribution, even though they lead to entirely different Taylor series. The scheme outlined in Eq. (8.5) is however much more advantageous because it leads already at second order to the desired approximation, and negative likelihoods then do not appear at all, since the exponential function is always positive. Therefore, we see that the choice of which quantity shall be approximated influences decisively how quickly the approximation recovers the shape

---

[1] Often, however, such analyses are carried out by comparing a measured power spectrum to a parameterized power spectrum, which is then treated as the mean. In these cases the covariance matrix would then be the covariance *of* the power spectrum (a four-point function) instead of *the* powerspectrum.

of the original function, and whether unwanted artifacts appear when truncating the approximation at low orders.

The choice of Taylor approximating the log-likelihood, instead of the likelihood, to second order in multiple dimensions yields a Hessian matrix whose expectation value is the Fisher (or information) matrix. Denoting partial derivatives by $\partial_\alpha f = f_{,\alpha}$, the Fisher matrix of Eq. (8.1) can be written as

$$\begin{aligned} F_{\alpha\beta} =& \langle \mathcal{L}_{,\alpha\beta} \rangle |_{\hat{\boldsymbol{p}}} \\ =& \frac{1}{2} \mathrm{Tr} \left(\boldsymbol{C_0}^{-1} C_{,\alpha} \boldsymbol{C_0}^{-1} C_{,\beta} \right) + \boldsymbol{\mu}_{,\alpha} \boldsymbol{C_0}^{-1} \boldsymbol{\mu}_{,\beta} \end{aligned} \tag{8.6}$$

where the derivatives are evaluated at the maximum likelihood point $\hat{\boldsymbol{p}}$ and summation over repeated indices is implied and $\boldsymbol{C_0}$ is evaluated at the maximum likelihood point and therefore constant and cannot be derived with respect to parameters.

The corresponding likelihood approximation is then given by

$$L(\boldsymbol{X}|\boldsymbol{p}) \approx N \cdot \exp(-\frac{1}{2} F_{\alpha\beta} \Delta p_\alpha \Delta p_\beta) \tag{8.7}$$

where the $\Delta p_\alpha = p_\alpha - \hat{p}_\alpha$ are the offsets from the best-fit point $\hat{p}_\alpha$ and N is a normalization constant.

The Fisher approximation results in the usual ellipsoidal, multi-variate correlated Gaussian confidence contours, which often do not recover the shape of a non-Gaussian likelihood distribution well. A continuation of the Taylor series is then desirable in order to capture these non-Gaussianites. This wish for a continuation of the series is predicated on the requirement to solve the issue of normalizability and positive-definiteness at all orders. Also, it is preferrable to recover the essential shape of the likelihood with as little additional terms as possible for computational efficiency. Clearly, just as there exist multiple ways in approximating the likelihood Eq. (8.4), there will exist multiple ways for continuing the approximation from that given by the Fisher matrix. These extended approximations will pick up the desired information about the likelihood's shape with different efficiencies. Let us first demonstrate why a usual continuation of the log-likelihood Taylor approximation does not produce the desired result.

8.3.1 Parameter-dependent mean

We consider a posterior that depends on n parameters p_α, where α takes the values $1 \dots n$. Denoting with P the posterior distribution, we expand the log-likelihood $\mathcal{L} = -\log(P)$ as a function of the parameters p_α in Taylor series around the maximum of the likelihood, indicated by the subscript 0:

$$\begin{aligned} -\mathcal{L} \equiv \log P \approx \log P_0 + \frac{1}{2} \left(\log P \right)_{,\alpha\beta} \Big|_0 \Delta p_\alpha \Delta p_\beta \\ + \frac{1}{3!} \left(\log P \right)_{,\alpha\beta\gamma} \Big|_0 \Delta p_\alpha \Delta p_\beta \Delta p_\gamma \\ + \frac{1}{4!} \left(\log P \right)_{,\alpha\beta\gamma\delta} \Big|_0 \Delta p_\alpha \Delta p_\beta \Delta p_\gamma \Delta p_\delta, \end{aligned} \tag{8.8}$$

where summation over repeated indices is implied. The first order derivatives vanish because we are at the maximum of the posterior. Expanding to the second order yields the Fisher approximation. From the third order onwards, non-Gaussianities are taken into account, that correct for misestimates of the posterior by the Fisher matrix, and thereby lead to a deformation of its shape. We can write the approximation

$$\begin{aligned} P = N \exp \Bigg[& -\frac{1}{2} F_{\alpha\beta} \Delta p_\alpha \Delta p_\beta - \frac{1}{3!} S_{\alpha\beta\gamma} \Delta p_\alpha \Delta p_\beta \Delta p_\gamma \\ & - \frac{1}{4!} Q_{\alpha\beta\gamma\delta} \Delta p_\alpha \Delta p_\beta \Delta p_\gamma \Delta p_\delta - \mathcal{O}(5) \Bigg] , \end{aligned} \tag{8.9}$$

where N is a normalization constant,

$$\begin{aligned} F_{\alpha\beta} &= \mathcal{L}_{,\alpha\beta} \, , \\ S_{\alpha\beta\gamma} &= \mathcal{L}_{,\alpha\beta\gamma} \, , \\ Q_{\alpha\beta\gamma\delta} &= \mathcal{L}_{,\alpha\beta\gamma\delta} . \end{aligned} \tag{8.10}$$

Working out the derivatives, we find

$$\begin{aligned} S_{\alpha\beta\gamma} &= \langle \mathcal{L}_{,\alpha\beta} \mathcal{L}_{,\gamma} \rangle + \langle \mathcal{L}_{,\gamma\alpha} \mathcal{L}_{,\beta} \rangle + \langle \mathcal{L}_{,\beta\gamma} \mathcal{L}_{,\alpha} \rangle \\ &= \boldsymbol{\mu}_{,\alpha\beta} M \boldsymbol{\mu}_{,\gamma} + \text{cycl}., \end{aligned} \tag{8.11}$$

The fourth-order tensor is

$$\begin{aligned} Q_{\alpha\beta\gamma\delta} = {} & \boldsymbol{\mu}_{,\alpha\gamma\delta} M \boldsymbol{\mu}_{,\beta} + \boldsymbol{\mu}_{,\delta\gamma} M \boldsymbol{\mu}_{,\beta\alpha} \\ & + \boldsymbol{\mu}_{,\alpha\beta\delta} M \boldsymbol{\mu}_{,\gamma} + \boldsymbol{\mu}_{,\delta\beta} M \boldsymbol{\mu}_{,\gamma\alpha} \\ & + \boldsymbol{\mu}_{,\alpha\gamma\beta} M \boldsymbol{\mu}_{,\delta} + \boldsymbol{\mu}_{,\beta\gamma} M \boldsymbol{\mu}_{,\delta\alpha} \\ & + \boldsymbol{\mu}_{,\delta\gamma\beta} M \boldsymbol{\mu}_{,\alpha} \, . \end{aligned} \tag{8.12}$$

Here, the $n \times n$ matrix $F_{\alpha\beta}$ is the usual Fisher matrix. We dub the $n \times n \times n$ tensor $S_{\alpha\beta\gamma}$ the flexion tensor and define the scalar

$$S \equiv S_{\alpha\beta\gamma} \Delta p_\alpha \Delta p_\beta \Delta p_\gamma . \tag{8.13}$$

Likewise we call $Q_{\alpha\beta\gamma\delta}$ the quarxion tensor and define its scalar

$$Q \equiv Q_{\alpha\beta\gamma\delta} \Delta p_\alpha \Delta p_\beta \Delta p_\gamma \Delta . p_\delta \tag{8.14}$$

Now we see, that if we truncate after the second order, i.e. at the Fisher matrix, then the exponential Eq (8.9) contains only a quadratic form with a negative sign. The argument of the exponential function is consequently always negative, which ensures that the probability stays finite. This handy feature is not necessarily true for the quarxions and never true for the flexions: The flexion is cubic in the Δp and will therefore always become negative at large enough Δp. Whenever negative flexion and quarxions terms become larger than the term from the Fisher matrix, the argument of the exponential becomes positive and the higher order Taylor approximation of the log-likelihood leads to a likelihood which diverges at large Δp.

8.3.2 Parameter-dependent covariance

For a constant mean, and a parameter-dependent covariance matrix, the log-likelihood's Taylor-approximation has

$$\begin{aligned} S_{\alpha\beta\gamma} =& \mathcal{L}_{,\alpha\beta\gamma}|_{\hat{\boldsymbol{p}}} \\ &= -2\mathrm{Tr}\left(\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\gamma}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\beta}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\alpha}\right) \\ &+ \frac{3}{2}\mathrm{Tr}\left(\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\gamma}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\alpha\beta}\right), \end{aligned} \tag{8.15}$$

and

$$\begin{aligned} Q_{\alpha\beta\gamma\delta} =& \mathcal{L}_{,\alpha\beta\gamma\delta}|_{\hat{\boldsymbol{p}}} \\ &= 9\ \mathrm{Tr}\left(\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\delta}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\gamma}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\beta}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\alpha}\right) \\ &+ \frac{3}{2}\mathrm{Tr}\left(\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\gamma\delta}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\alpha\beta}\right) \\ &- 12\ \mathrm{Tr}\left(\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\gamma\delta}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\beta}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\alpha}\right) \\ &+ 2\ \mathrm{Tr}\left(\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\gamma}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\alpha\beta\delta}\right) \end{aligned} \tag{8.16}$$

which gives the Taylor series of the log-likelihood up to fourth order, after being averaged over realizations of the data.

Just like in case of a parameter-dependent mean, this approximation will in general be unnormalizable since it will diverge somewhere in parameter space. This is partly due to the odd powers of Δp, which will clearly become negative on one side of the fiducial point (about which the expansion is made) if they are positive on the other side of the fiducial point; the argument of the exponential function will then become positive even for small displacements from the best-fit point, and the approximation will begin to diverge. Also the summation over even powers of Δp can lead to divergences, as terms of the structure $\Delta p_1 \Delta p_1 \Delta p_1 \Delta p_2$ will appear that are linear and cubic in Δp.

Secondly, we observe that even if only first-order derivatives of the covariance matrix were non-vanishing, the above series would still not terminate after the Fisher approximation. The first lines of Eq. (8.15) and Eq. (8.16) contain only first order derivatives of the covariance matrix and make it clear that at the n-th Taylor order a term of the shape

$$\mathrm{Tr}\left(\left(\boldsymbol{C_0}^{-1}\boldsymbol{C}_{,\alpha}\,\Delta p_\alpha\right)^n\right) \tag{8.17}$$

appears, where we have expressed the repeated multiplication of the same matrices as a power. As new information on the parameter dependence of the covariance matrix is encoded in its higher order derivatives, the terms Eq. (8.17) do not add any of the new information which we target; they simply stem from the slowly convergent Taylor series of the logarithm.

Therefore we see that a Taylor approximation of the log-likelihood beyond second order is a valid but laborious way to include non-Gaussian behaviour: the log-likelihood would need to be approximated to much higher than the 4th order, before it can be expected to be normalizable for a physical application.

8.4 The DALI principle of construction

We have learned in the previous sections that a Taylor expansion of the log-likelihood to higher orders will in general lead to divergences. These divergences can only be avoided in an application-independent way if the argument of the exponential is negative everywhere in parameter space. One way of achieving this is to demand of the approximation to have the shape

$$L \approx N \exp\left(-\mathcal{Q}\right), \tag{8.18}$$

where $\mathcal{Q}$ is a quadratic function of the parameters and therefore always positive definite. The expansion Eq. (8.9) of the log-likelihood does not have this shape.

Demanding the approximation to take the form (8.18), obviously raises the question how to choose $\mathcal{Q}$. If any arbitrary positive definite $\mathcal{Q}$ is constructed from the components of $S_{\alpha\beta\gamma}$ and $Q_{\alpha\beta\gamma\delta}$, then the approximation will in general not possess a good shape fidelity, meaning that it will not converge towards the true underlying distribution in a transparent way.

In order to remedy this problem, and to ensure a controlled convergence of the approximation to the true likelihood, we therefore demand that $\mathcal{Q}$ shall be a function of a Taylor approximation of some quantity in the log-likelihood. Then we can achieve normalizability and shape fidelity at the same time: The Taylor approximation will guarantee the convergence of the approximation towards the underlying likelihood, and the quadratic form will guarantee the normalizability.

In the following, we will indeed construct the DALI series from a Taylor approximation of the mean $\boldsymbol{\mu}(\boldsymbol{p})$ and the covariance $\boldsymbol{C}(\boldsymbol{p})$. The convergence criterion for this approximation is given in Sec. (8.6).

8.5 Beyond Gaussianity with DALI

8.5.1 Parameter-dependent mean

We consider now cases in which the parameters appear only in a theoretical model $\boldsymbol{\mu}$ that is compared to a data set, and not in the covariance matrix. With $M = C^{-1}$ being the inverse of the parameter-independent and positive-definite covariance matrix in the data space, the log-likelihood is then

$$\mathcal{L} = -\frac{1}{2}\left(\boldsymbol{X} - \boldsymbol{\mu}(\boldsymbol{p})\right) M \left(\boldsymbol{X} - \boldsymbol{\mu}(\boldsymbol{p})\right). \tag{8.19}$$

The log-likelihood therefore is already a quadratic form in $\boldsymbol{\mu}$ and we can satisfy the DALI principle from Sec. (8.4) if we Taylor expand the mean around the best-fit $\hat{p}_\alpha$

$$\boldsymbol{\mu}(\boldsymbol{p}) = \hat{\boldsymbol{\mu}} + \boldsymbol{\mu}_{,\alpha}\Delta p_\alpha + \frac{1}{2}\boldsymbol{\mu}_{,\alpha\beta}\Delta p_\alpha \Delta p_\beta + ..., \tag{8.20}$$

where $\hat{\boldsymbol{\mu}} \equiv \boldsymbol{\mu}(\hat{\boldsymbol{p}})$. The log-likelihood then expands to

$$\begin{aligned} -\mathcal{L} \approx \quad & \frac{1}{2}[\boldsymbol{m}-\hat{\boldsymbol{\mu}}]M[\boldsymbol{m}-\hat{\boldsymbol{\mu}}] - \Big(\boldsymbol{m}-\hat{\boldsymbol{\mu}}\Big)M\Big(\boldsymbol{\mu}_{,\alpha}\Delta p_\alpha + \frac{1}{2}\boldsymbol{\mu}_{,\alpha\beta}\Delta p_\alpha \Delta p_\beta\Big) \\ & + \frac{1}{2}\boldsymbol{\mu}_{,\alpha}M\boldsymbol{\mu}_{,\beta}\Delta p_\alpha \Delta p_\beta + \frac{1}{2}\boldsymbol{\mu}_{,\alpha}M\boldsymbol{\mu}_{,\beta\gamma}\Delta p_\alpha \Delta p_\beta \Delta p_\gamma \\ & + \frac{1}{8}\boldsymbol{\mu}_{,\alpha\beta}M\boldsymbol{\mu}_{,\gamma\delta}\Delta p_\alpha \Delta p_\beta \Delta p_\gamma \Delta p_\delta + \ldots . \end{aligned} \tag{8.21}$$

The first term on the right hand side is an irrelevant constant that can be absorbed in the normalization; the second term averages out to zero, while the remaining terms include non-Gaussianities into the approximation and we have arrived at our aim. We can now show that this DALI series is just a resummation of the Taylor expansion of the log-likelihood. It is obvious from their definition Eq. (8.10) that flexions and quarxions posess some symmetries under index permutation. When taking the full flexion or quarxion term, all the distinct terms of the same type in Eqs. (8.11) and (8.12) become indistinguishable. For instance, by renaming the indexes,

$$\boldsymbol{\mu}_{,\alpha\beta}M\boldsymbol{\mu}_{,\gamma}\Delta p_\alpha \Delta p_\beta \Delta p_\gamma = \boldsymbol{\mu}_{,\gamma\alpha}M\boldsymbol{\mu}_{,\beta}\Delta p_\gamma \Delta p_\alpha \Delta p_\beta \tag{8.22}$$

$$= \boldsymbol{\mu}_{,\beta\gamma}M\boldsymbol{\mu}_{,\alpha}\Delta p_\beta \Delta p_\gamma \Delta p_\alpha \, . \tag{8.23}$$

Therefore we can simplify

$$\begin{aligned} S &= 3\boldsymbol{\mu}_{,\alpha\beta}M\boldsymbol{\mu}_{,\gamma}\Delta p_\alpha \Delta p_\beta \Delta p_\gamma \, , \\ Q &= (4\boldsymbol{\mu}_{,\alpha\gamma\delta}M\boldsymbol{\mu}_{,\beta} + 3\boldsymbol{\mu}_{,\delta\gamma}M\boldsymbol{\mu}_{,\beta\alpha})\Delta p_\alpha \Delta p_\beta \Delta p_\gamma \Delta p_\delta \, . \end{aligned} \tag{8.24}$$

Now, as anticipated, the expansion can be arranged in order of derivatives. That is, to second order in the $\boldsymbol{\mu}$ derivatives we have

$$\begin{aligned} P = N \exp \Bigg[& -\frac{1}{2}\boldsymbol{\mu}_{,\alpha}M\boldsymbol{\mu}_{,\beta}\Delta p_\alpha \Delta p_\beta \\ & - \Bigg(\frac{1}{2}\boldsymbol{\mu}_{,\alpha\beta}M\boldsymbol{\mu}_{,\gamma}\Delta p_\alpha \Delta p_\beta \Delta p_\gamma \\ & + \frac{1}{8}\boldsymbol{\mu}_{,\delta\gamma}M\boldsymbol{\mu}_{,\beta\alpha}\Delta p_\alpha \Delta p_\beta \Delta p_\gamma \Delta p_\delta\Bigg) + \mathcal{O}(3)\Bigg] , \end{aligned} \tag{8.25}$$

which is exactly the normalizable and positive definite DALI series (8.21). Likewise, at third order we have

$$
\begin{aligned}
P = N & \\
& \exp\Bigg[-\frac{1}{2}\boldsymbol{\mu}_{,\alpha}M\boldsymbol{\mu}_{,\beta}\Delta p_\alpha\Delta p_\beta - \Bigg(\frac{1}{2}\boldsymbol{\mu}_{,\alpha\beta}M\boldsymbol{\mu}_{,\gamma}\Delta p_\alpha\Delta p_\beta\Delta p_\gamma \\
& + \frac{1}{8}\boldsymbol{\mu}_{,\delta\gamma}M\boldsymbol{\mu}_{,\beta\alpha}\Delta p_\alpha\Delta p_\beta\Delta p_\gamma\Delta p_\delta\Bigg) \\
& - \Bigg(\frac{1}{6}\boldsymbol{\mu}_{,\delta}M\boldsymbol{\mu}_{,\beta\alpha\gamma}\Delta p_\alpha\Delta p_\beta\Delta p_\gamma\Delta p_\delta \\
& + \frac{1}{3!2!}\boldsymbol{\mu}_{,\alpha\beta\delta}M\boldsymbol{\mu}_{,\gamma\tau}\Delta p_\alpha\Delta p_\beta\Delta p_\gamma\Delta p_\delta\Delta p_\tau \\
& + \frac{1}{3!3!2!}\boldsymbol{\mu}_{,\alpha\beta\gamma}M\boldsymbol{\mu}_{,\delta\tau\sigma}\Delta p_\alpha\Delta p_\beta\Delta p_\gamma\Delta p_\delta\Delta p_\tau\Delta p_\sigma\Bigg) + \mathcal{O}(4)\Bigg].
\end{aligned}
\tag{8.26}
$$

The approximated posteriors (8.25)–(8.26) represent true distributions and are the second- and third-derivative correction, respectively, over the Fisher approximation.

8.5.2 Parameter-dependent covariance

In order to derive the DALI series for a parameter-dependent covariance matrix and a constant mean, we express the variation of the covariance matrix over the parameter space by its Taylor series and single out the constant zeroth-order term

$$
\boldsymbol{C}(\boldsymbol{p}) = \boldsymbol{C_0} + \mathcal{T}^1_{(C)}, \tag{8.27}
$$

where $\boldsymbol{C_0}$ is the constant covariance matrix evaluated at the likelihood maximum, and

$$
\mathcal{T}^1_{(C)} = \sum_{n=1}^{\infty}\frac{\boldsymbol{C}^{(n)}|_{\hat{\boldsymbol{p}}}}{n!}(p_\alpha - \hat{p}_\alpha)...(p_n - \hat{p}_n) \tag{8.28}
$$

is the p-dimensional Taylor series of the covariance matrix, beginning at the first derivative $\boldsymbol{C}^{(1)}$. The derivatives are chosen to be evaluated at the maximum likelihood point, denoted by $\hat{\boldsymbol{p}}$. For $\boldsymbol{\mu} \equiv 0$ the data matrix is $\boldsymbol{D} = \boldsymbol{X}\boldsymbol{X}^T$ which is parameter independent. The log-likelihood is then given by

$$
\begin{aligned}
\mathcal{L} &= \frac{1}{2}\mathrm{Tr}\left(\ln(\boldsymbol{C})\right) + \frac{1}{2}\mathrm{Tr}\left(\langle\boldsymbol{X}\boldsymbol{X}^T\rangle\boldsymbol{C}^{-1}\right) \\
&= \frac{1}{2}\mathrm{Tr}\left(\ln\left(\boldsymbol{C_0}[1 + \boldsymbol{C_0}^{-1}\mathcal{T}^1_{(C)}]\right) + \langle\boldsymbol{X}\boldsymbol{X}^T\rangle\left(\boldsymbol{C_0} + \mathcal{T}^1_{(C)}\right)^{-1}\right),
\end{aligned}
\tag{8.29}
$$

where angular brackets denote averaging over the data and $\langle\boldsymbol{X}\boldsymbol{X}^T\rangle$ is kept explicitely, in order to emphasize that it does not depend on parameters, although it will later average out to be the measured covariance matrix. So far, the covariance matrix has only been

rewritten, but no approximation has been made. We will now approximate, in order to achieve a quadratic form.
If the Taylor series $\mathcal{T}^1_{(C)}$ is evaluated only sufficiently close to the maximum likelihood point, then we will have $\mathcal{T}^1_{(C)} \ll \boldsymbol{C}_{\mathbf{0}}$ and we can consistently approximate Eq. (8.29) up to second order in $\mathcal{T}^1_{(C)}$. This leads to the targeted shape Eq. (8.18). We therefore approximate by applying the matrix inversion identity (also known as Woodbury identity)

$$(\boldsymbol{A}+\boldsymbol{B})^{-1} = \boldsymbol{A}^{-1} - \boldsymbol{A}^{-1}\left(1+\boldsymbol{B}\boldsymbol{A}^{-1}\right)^{-1}\boldsymbol{B}\boldsymbol{A}^{-1} \tag{8.30}$$

to find an approximation for the inverted covariance matrix

$$\begin{aligned}\left(\boldsymbol{C}_{\mathbf{0}} + \mathcal{T}^1_{(C)}\right)^{-1} &= \boldsymbol{C}_0^{-1} + \boldsymbol{C}_0^{-1}\mathcal{T}^1_{(C)}\,\boldsymbol{C}_0^{-1} \\ &\quad - \boldsymbol{C}_0^{-1}\mathcal{T}^1_{(C)}\,\boldsymbol{C}_0^{-1}\mathcal{T}^1_{(C)}\,\boldsymbol{C}_0^{-1} + \mathcal{O}(3),\end{aligned} \tag{8.31}$$

where the approximation was truncated at second order since we target the shape Eq. (8.18). The quadratic term of the logarithm's Taylor expansion is,

$$\ln(1+x) = x - \frac{x^2}{2} + \mathcal{O}(x^3). \tag{8.32}$$

The quadratic approximation of the log-likelihood then becomes

$$\begin{aligned}\mathcal{L} \approx{}& \frac{1}{2}\mathrm{Tr}\left(\ln(\boldsymbol{C}_{\mathbf{0}}) + \boldsymbol{C}_0^{-1}\mathcal{T}^1_{(C)} - \frac{1}{2}\boldsymbol{C}_0^{-1}\mathcal{T}^1_{(C)}\,\boldsymbol{C}_0^{-1}\mathcal{T}^1_{(C)}\right) \\ &+ \frac{1}{2}\mathrm{Tr}\left[\langle \boldsymbol{X}\boldsymbol{X}^T\rangle\left(\boldsymbol{C}_0^{-1} - \boldsymbol{C}_0^{-1}\mathcal{T}^1_{(C)}\boldsymbol{C}_0^{-1}\right.\right. \\ &\qquad\qquad\left.\left. + \boldsymbol{C}_0^{-1}\mathcal{T}^1_{(C)}\boldsymbol{C}_0^{-1}\mathcal{T}^1_{(C)}\boldsymbol{C}_0^{-1} + \mathcal{O}(3)\right)\right].\end{aligned} \tag{8.33}$$

Applying $\langle \boldsymbol{X}\boldsymbol{X}^T\rangle = \boldsymbol{C}_{\mathbf{0}}$ the likelihood approximation simplifies to

$$\begin{aligned}L &\approx N\exp(-\mathcal{L}) \\ &= N\exp\left(-\frac{1}{4}\mathrm{Tr}\left(\boldsymbol{C}_0^{-1}\mathcal{T}^1_{(C)}\boldsymbol{C}_0^{-1}\mathcal{T}^1_{(C)}\right) + \mathcal{O}(3)\right) \\ &= N\exp\left(-\frac{1}{4}\mathrm{Tr}\left[\boldsymbol{C}_0^{-1}(\boldsymbol{C}_{,\alpha}\,\Delta p_\alpha + \frac{1}{2}\boldsymbol{C}_{,\alpha\beta}\,\Delta p_\alpha \Delta p_\beta + \ldots)\right.\right. \\ &\qquad\qquad\left.\left.\boldsymbol{C}_0^{-1}(\boldsymbol{C}_{,\alpha}\,\Delta p_\alpha + \frac{1}{2}\boldsymbol{C}_{,\alpha\beta}\,\Delta p_\alpha \Delta p_\beta + \ldots)\right] + \mathcal{O}(3)\right),\end{aligned} \tag{8.34}$$

where $\ln(\boldsymbol{C}_{\mathbf{0}})$ and $\boldsymbol{C}_{\mathbf{0}}\boldsymbol{C}_0^{-1} = 1$ are constants and were absorbed into the normalization constant N. In the last step, a repeated multiplication of the same terms appears. This can be rewritten as

$$\begin{aligned}L &\approx \\ &= N\exp\left(-\frac{1}{4}\mathrm{Tr}\left(\left(\boldsymbol{C}_0^{-1}(\boldsymbol{C}_{,\alpha}\,\Delta p_\alpha + \frac{1}{2}\boldsymbol{C}_{,\alpha\beta}\,\Delta p_\alpha \Delta p_\beta + \ldots)\right)^2\right)\right),\end{aligned} \tag{8.35}$$

where the repeated multiplication of the same matrices in the trace was made more explicit by denoting it as a square.

We therefore have arrived at an approximation of the shape Eq. (8.18) that includes higher order derivatives of the covariance matrix. This approximation will consequently remain normalizable everywhere in parameter space. This result generalizes the usual Fisher matrix in a straightforward way: if the Taylor-approximation $\mathcal{T}^1_{(C)}$ is truncated at first order, the usual Fisher matrix approximation Eq. (8.6) of the likelihood is obtained and the higher order corrections are then

$$\begin{aligned} L \approx N \exp\Bigg(&-\frac{1}{4}\mathrm{Tr}\left(\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\alpha}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\beta}\right)\Delta p_\alpha \Delta p_\beta \\ &-\frac{1}{4}\mathrm{Tr}\left(\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\alpha}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\beta\gamma}\right)\Delta p_\alpha \Delta p_\beta \Delta p_\gamma \\ &-\frac{1}{16}\mathrm{Tr}\left(\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\alpha\beta}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\gamma\delta}\right)\Delta p_\alpha \Delta p_\beta \Delta p_\gamma \Delta p_\delta \\ &-\frac{1}{24}\mathrm{Tr}\left(\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\alpha\beta}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\gamma\delta\epsilon}\right)\Delta p_\alpha \Delta p_\beta \Delta p_\gamma \Delta p_\delta \Delta p_\epsilon \\ &-\frac{1}{144}\mathrm{Tr}\left(\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\alpha\beta\gamma}\,\boldsymbol{C}_0^{-1}\boldsymbol{C}_{,\delta\epsilon\phi}\right)\Delta p_\alpha \Delta p_\beta \Delta p_\gamma \Delta p_\delta \Delta p_\epsilon \Delta p_\phi \Bigg), \end{aligned} \tag{8.36}$$

where we have chosen to truncate the Taylor expansion of the covariance matrix at third order for brevity; the continuation to fourth and higher orders of the covariance matrix is however obvious from Eq. (8.35). The terms that are cubic and quintic in the Δp can become negative and thereby decrease the likelihood estimate in regions, where it had been overerstimated by the even-order terms. In total however, the terms combine to a quadratic form, and thereby the approximation is known to not diverge anywhere in parameter space.

8.6 Criteria of applicability

Non-Gaussianity can arise from at least two sources. For example if the data has only little constraining power then even the likelihood for a model with only mildly non-linear parameters will develop non-Gaussianities. In contrast, if the data are very constraining non-Gaussianites will still occur if parameters are degenerate with each other over a finite range. In this case the non-Gaussianities can be recovered by DALI.

The approximation of Eq. (8.35) is strictly valid if the following criteria are fulfilled:

- The data set $\boldsymbol{X}$ must be so constraining that the likelihood is confined to a region Δp where the second-order Taylor approximations Eq. (8.32) and Eq. (8.31) dominate over their higher orders.
- Approximating the log in Eq. (8.33) requires

$$\mathrm{Tr}\left(\boldsymbol{C}_0^{-1}\mathcal{T}^1_{(C)}\right) \ll 1 \tag{8.37}$$

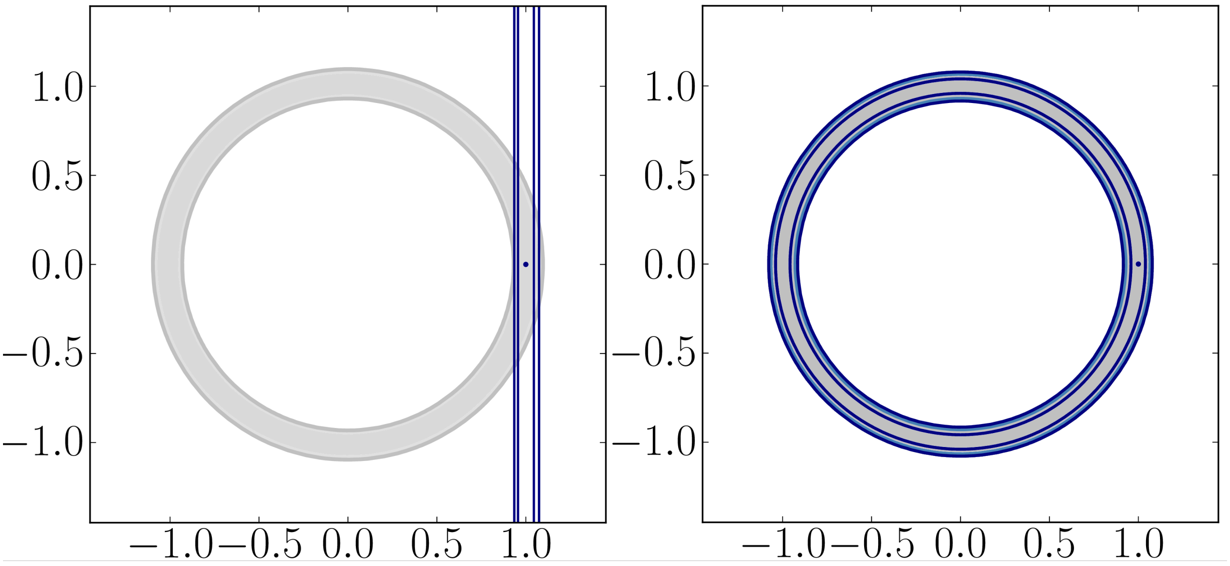

Figure 8.1: The unapproximated likelihood of Eq. (8.39) is depicted in grey. Since Eq. (8.39) is the equation of a circle, the likelihood has a ring shape. Left: The Fisher approximation in blue, with fiducial point indicated by the small blue dot. The Fisher matrix is singular and therefore appears as a set of parallel lines. Right: in blue the DALI approximation using second-order derivatives of the covariance matrix Eq. (8.39).

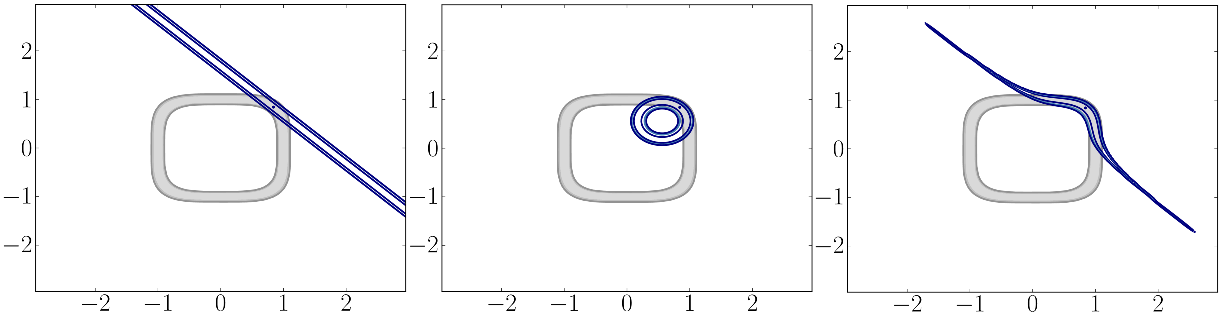

Figure 8.2: Like Fig.(8.1) only for a likelihood using the covariance matrix Eq. (8.40). The likelihood is again depicted by the empty grey ring, and the different approximations are depicted in blue: Fisher matrix (left), DALI with second order derivatives of the covariance matrix (middle), DALI with third order derivatives (right).

which can be solved for parameter offsets Δp

$$\Delta p_\alpha \ll \frac{1}{\mathrm{Tr}\left(\boldsymbol{C}_0^{-1}(\boldsymbol{C}_{,\alpha} + \frac{1}{2}\boldsymbol{C}_{,\alpha\beta}\,\Delta p_\beta + ...)\right)}. \tag{8.38}$$

The last requirement will be fulfilled if the data set confines the preferred parameter space to an area within which the Taylor approximation captures well the variation of the covariance matrix throughout the parameter space. DALI is therefore expected to work well in case of rather constraining data and degenerate parameters, while a good recovery of non-Gaussianities for weakly constraining data and mild non-linear dependences on the parameters would require Taylor-approximating the log-likelihood to much higher orders with the corresponding difficulties detailed in Sect. (8.3). If the condition Eq. (8.38) is only marginally fulfilled, the DALI approximations will still converge although they will not pick up all the shape information of the likelihood. Mismatches between the shape of the approximation and the real likelihood shape will then be observed. This is already known from the Fisher matrix, and expected to be more mild in DALI since the higher order derivatives will correct upon the Fisher matrix misestimates.

8.7 Numerical complexity and marginalization

The one incontrovertible advantage of the Fisher matrix is its speed. A quick estimate of the complexity of the DALI approximation can be obtained by observing that the expensive computations needed for the matrices are the evaluations of the vectors of the derivatives. For n parameters, there are n possible first derivatives, so the complexity rises linearly with n. It rises with $n(n+1)/2$ when second-order derivatives are included, because there are $n(n+1)/2$ distinct derivatives then. Likewise, it rises as $(n^3+3n^2+2n)/6$ when including third-order derivatives. Since every numerical derivative of order p requires (at lowest accuracy) $p+1$ evaluations of the posterior, the complexity for large n goes like n^2 and $(2/3)n^3$ for second and first order, respectively. In comparison, grids or MCMC-routines evaluate the full likelihood (which implies generating theoretical predictions of the data at every point in parameter space) typically thousands of times already for e.g. four parameters. Therefore only for $\mathcal{O}(1000)$ [$\mathcal{O}(100)$] parameters does DALI with second [third] order derivatives require roughly the same $\mathcal{O}(10^6)$ evaluations of a typical Monte Carlo run in large parameter spaces. In practice the evaluation of the posterior is thus significantly faster with DALI, as most forecasts in cosmology rely on less than a dozen of free parameters, and the posterior can be numerically costly to compute. Notice however, that only Gaussian posteriors are again Gaussians with fewer dimensions if they are marginalized. This analytical result makes marginalizations with Fisher matrices extremely fast. For non-Gaussian posteriors, for which DALI is interesting, there exists no general analytical marginalization. Therefore DALI will be slower in this respect than Fisher matrix approximations. Let us point out that this is not a flaw of DALI – is necessarily the price that one has to pay, if the non-Gaussianity of a posterior shall be captured because with the loss of Gaussianity, also its analytical results are lost. The public DALI code takes care of the needed numerical replacements like the numerical marginalization.

8.8 Parameter-dependent covariance matrix and mean

We now have two DALI approximations, one dealing with constant covariance matrices and a parameter-dependent mean, and one dealing with a constant mean and a parameter-dependent covariance. An interesting question is whether the two approximations can be combined to approximate a likelihood where *both* mean and covariance matrix depend on data. So far, if both, $\boldsymbol{\mu}$ and $\boldsymbol{C}$ depend on the same parameters, no quadratic form that is at the same time a Taylor series has been found due to the appearance of cross-terms[2] between derivatives of $\boldsymbol{\mu}$ and $\boldsymbol{C}$, e.g. $\boldsymbol{\mu}_{,\alpha}\,\boldsymbol{C}_{,\beta}\,\boldsymbol{\mu}_{,\gamma}$. Neglecting these cross-terms will produce a DALI expansion that is a simple multiplication of Eq. (8.35) and Eq. (8.26). This may be a good approximation in many cases, e.g. when the covariance matrix depends strongly on some parameters but not on those on which the mean depends. However, due to omitting the cross-terms, in general this expansion will not be able to recover all information and therefore it may not yield a good approximation.

8.9 Illustrative test cases

The strength of this new approximation scheme was tested on two toy models of particularly severe non-Gaussianities which arise from degeneracies. Both toy models are two-dimensional and have $\boldsymbol{\mu} \equiv 0$. The data set consists of 50 data points. The covariance matrix of the first is diagonal and given by

$$C_{ij}(\boldsymbol{p}) = (p_1^2 + p_2^2)\delta_{ij}, \tag{8.39}$$

with the Kronecker delta δ_{ij}. Since $p_1^2 + p_2^2 = \text{const}$ is the equation of a circle, this model produces a ring-shaped unapproximated likelihood, with the interior of the ring being a region of zero likelihood. All points which lie exactly on the circle will maximize the likelihood and any of them could be chosen as fiducial point for evaluating the DALI approximation. Taking more than 50 data points would decrease the thickness of the ring but would never be able to lift the degeneracy, even for an infinite number of measurements. Such likelihoods appear for example in particle physics for measurements of the Cabibbo-Kobayashi-Maskawa matrix (Charles et al., 2015).
The covariance matrix of the second toy model is,

$$C_{ij}(\boldsymbol{p}) = (p_1^4 + p_2^4)\delta_{ij}, \tag{8.40}$$

which again possesses a closed degeneracy line of a somewhat boxy ring shape. Again, each point along the line $p_1^4 + p_2^4 = \text{const}$ can serve as a fiducial point for the DALI approximation.

---

[2] In Sellentin et al. (2014), the appendix contained an approximation for parameter-dependent mean and covariance, but this approximation is not a strict Taylor expansion and not always positive definite.

The unapproximated likelihoods of these two models are depicted as grey shades in Fig. (8.1-8.2), where the two shades indicate the 68% and 95% confidence contours. Both toy models were then approximated by Eq. (8.35), truncated at different orders. The Fisher matrix of both cases is degenerate and appears as parallel non-closing lines. Changing the evaluation point of the derivatives cannot break this degeneracy. The second-order DALI approximation already finds the full circle, since no higher than second-order derivatives exist in this case. For the second toy-model, a complete recovery of the likelihood would require the calculation of fourth-order derivatives. Although this could be done analytically in the case at hand, in general such a calculation would need a numerical solution. We therefore maintain the truncation of the expansion Eq. (8.35) at third-order, as implemented in the public code DALI. The resulting approximation can be seen in Fig. (8.2, right) for third-order derivatives, or second-order derivatives in Fig. (8.2, middle). The degeneracy of the Fisher matrix is lifted in both cases, and the improvement in shape fidelity can easily be seen. As typical applications of this method would not posess such strong parameter degeneracies, it can be expected that the DALI method will reconstruct the likelihood contours with great accuracy.
A C++ code for DALI is public at DALI. However, due to the structural similarity with the Fisher matrix, any already existing Fisher code can easily be upgraded to a DALI code by adding the higher order derivatives.

Chapter 9

Scientific results from DALI

Having established the framework of DALI in the previous chapter, we proceed here to present various scientific uses of DALI. This chapter follows partly our publications Sellentin et al. (2014) and Sellentin & Schäfer (2016).

9.1 Forecasts for type Ia supernovae

We consider now an application of our method to supernova data. The measurable quantity is the distance modulus, which is related to the dimensionless luminosity distance by,

$$\mu_i = 5 \log \hat{d}(z_i), \tag{9.1}$$

where the index i denotes the dependence on a given redshift and $\hat{d}$ is the dimensionless luminosity distance. The log-likelihood for the supernovae after marginalization of the Hubble constant and the absolute luminosity is (Amendola & Tsujikawa, 2010b)

$$\mathcal{L} = -\log L = \frac{1}{2}\left(S_2 - \frac{S_1^2}{S_0}\right), \tag{9.2}$$

where the sums are

$$S_n = \sum_i \frac{(m_i - \mu_i)^n}{\sigma_i^2}, \tag{9.3}$$

where m_i is a measurement at redshift z_i and the corresponding theoretical mean μ_i. The log-likelihood can be written as

$$\mathcal{L} = \frac{1}{2} X_i M_{ij} X_j, \tag{9.4}$$

where $X_i = m_i - \mu_i$ and the inverse covariance matrix is,

$$M_{ij} = s_i s_j \delta_{ij} - \frac{s_i^2 s_j^2}{S_0}, \tag{9.5}$$

where $s_i = 1/\sigma_i$ and only here we de not imply a summation over the repeated indices. If one assumes $s_i = 1/\sigma$ (constant) then the covariance matrix is

$$M_{ij} = \sigma^{-2}\left(\delta_{ij} - \frac{1}{N}\right). \tag{9.6}$$

So finally we have

$$\begin{aligned} F^{\mathrm{SN}}_{\alpha\beta} &= \left\langle \left(\frac{\partial \mu_i}{\partial p_\alpha} M_{ij} X_j \right)^2 \right\rangle, \\ &= 25 \, \frac{\partial \log \hat{d}_i}{\partial p_\alpha} M_{ij} \frac{\partial \log \hat{d}_j}{\partial p_\beta} \,. \end{aligned} \tag{9.7}$$

Similarly, the flexion and quarxion tensors and the DALI expansion are then obtained by replacing μ_i with $5 \log \hat{d}_i$.

Notice that a parameter that appears additively in μ_i, like the offset, will not enter the DALI terms; therefore, the analytic marginalization of the posterior affects only the Fisher term and remains analytic also in DALI.

In order to demonstrate the potential of DALI, we show how accurately it can recover the “banana-shaped” posterior of the Supernova Union2.1 catalogue (Amanullah et al., 2010). This catalogue comprises the distance moduli of 580 SNeIa, which we use for the data points $m_i(z_i)$ of Eq. (9.3), together with their respective errors σ_i. We compare this data set with the distance moduli obtained from a flat wCDM cosmology with the Chevallier-Polarski-Linder parameterization for the dark energy equation of state (Chevallier & Polarski, 2001; Linder, 2003),

$$w(a) = w_0 + w_a(1 - a). \tag{9.8}$$

We choose the fiducial parameters to be the best fit parameters of the supernova posterior found in (Amanullah et al., 2010) for the wCDM model and evaluate the distance moduli at the redshifts of the Union2.1 catalogue.

In Figure 9.1 we depict in grey solid contours the non-approximated posterior (obtained with a grid method), which we will frequently refer to as the “full” posterior.

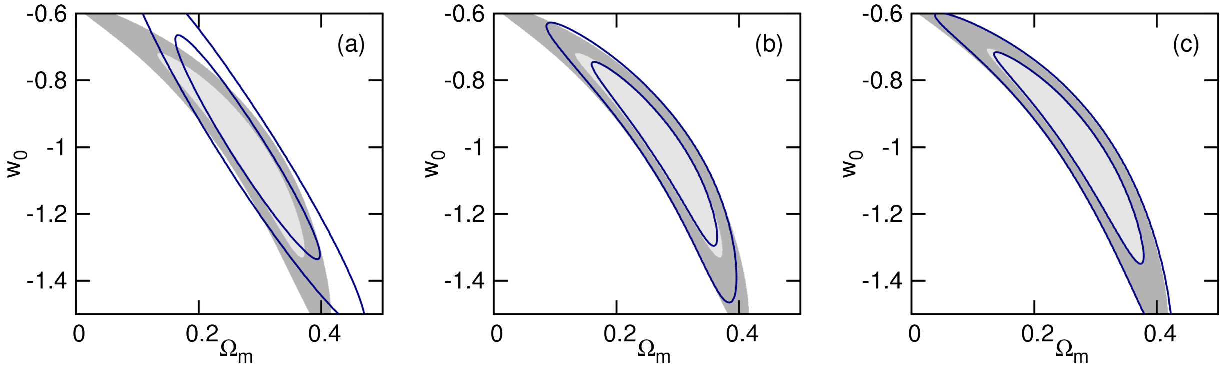

Figure 9.1: Un-approximated posterior of the Union2.1 catalogue (grey) and approximations (dark-blue) for $w_a = 0$. The confidence contours are drawn at the 1 and 2σ confidence levels. Panel (a): The Fisher Matrix; panel (b): Eq. (8.25), DALI to second order; panel (c): Eq. (8.26), DALI to third order.

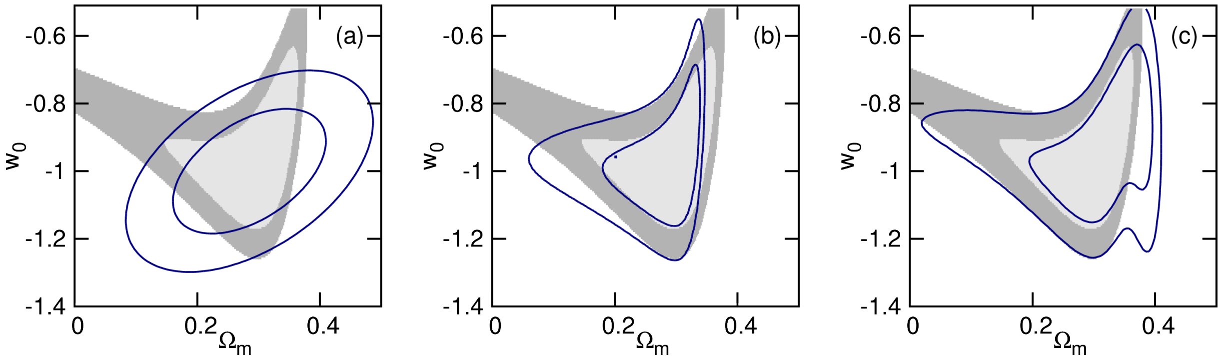

Figure 9.2: As Fig. 9.1 but for a mock catalogue of 1000 SNe and marginalizing over w_a.

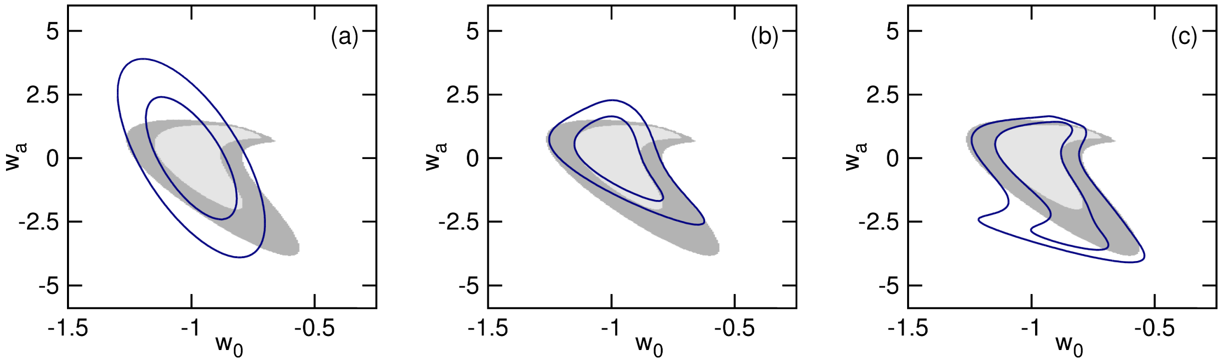

Figure 9.3: As Fig. 9.2 but this time marginalized over Ω_m in the interval $[0, 1]$. Note that the upper half of the Fisher-ellipse covers parameter ranges with high w_a, indicateing that the Fisher matrix does not capture the underlying physics well, whereas DALI does.

Here and in all other figures the contours are drawn at 1 and 2σ (we follow standard procedures and use 1σ and 2σ as shorthand notation for 68.3 and 95.4% confidence levels). The improvement of the shape fidelity by successively adding higher order derivatives to the posterior can be seen as one inspects panel (a) which is the Fisher matrix, (b) DALI with second-order derivatives and (c) DALI with third-order derivatives in Figure 9.1.

As the observational campaigns for dark energy proceed and more data are collected the posteriors are expected to become more and more Gaussian. To investigate the use of DALI in this respect, we mock a future supernova catalogue with 1000 SNeIa, uniformly distributed in the redshift range $0 < z < 2$. We use a flat CPL-cosmology with $\Omega_m^{\mathrm{fid}} = 0.285$, $w_0^{\mathrm{fid}} = -1$ and $w_a^{\mathrm{fid}} = 0$ as fiducial. In the (Ω_m, w_0)–plane, such a catalogue yields a posterior of similar shape to Figure 9.1, which DALI can recover nicely.

We further demonstrate the potential of our method on a posterior with higher non-Gaussianity: in Figure 9.2, we marginalized the posterior of Figure 9.1 over w_a in the range $(-\infty, +\infty)$, and our method can recover the shape of this heavily non-Gaussian posterior quite accurately. In Figure 9.3, we marginalized instead over Ω_m. Notice that the upper half of the Fisher ellipse extends far into the the parameter space of positive w_a, which corresponds to a completely different expansion history of the universe, one that is ruled out at many σ by supernova data.

The leakage of the DALI contours out of the true likelihood when including third-order derivatives in Figs. 9.2 and 9.3 is a numerical artefact. Already from the Fisher matrix approach it is known that numerical derivatives have to be estimated accurately since they determine size and orientation of the error contours. Figs. 9.2 and 9.3 had been generated with an old version of the DALI code which used a numerically fast but rough algorithm for calculating third derivatives: it took another derivative of precomputed and splined second derivatives. The public version of DALI uses a more accurate routine: it calculates third derivatives by using finite differences on the original function, not any already derived quantities. Redoing the analysis with the improved code results in Fig. 9.4, demonstrating that with the more carefully conducted estimate of third derivatives, the erroneous leakage of the DALI likelihoods disappears.

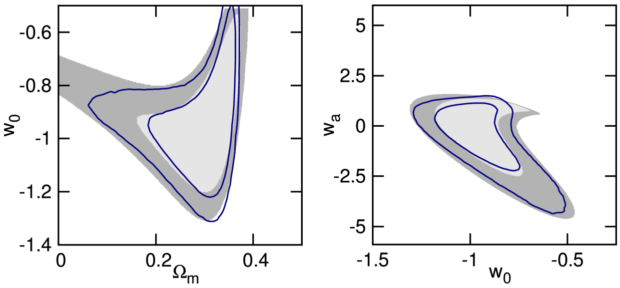

Figure 9.4: Replacement of the third-order derivatives of Fig. 9.2 and Fig. 9.3, using the public version of the DALI code where third-order derivatives are calculated more accurately. This demonstrates that the previously observed mismatch between the true likelihood and the DALI contours was only due to numerical inaccuracies.

9.2 Forecasts for Euclid weak lensing

In this section, we use DALI on a weak lensing data set in order to test how constraining it is, if no external priors are added. Including priors into a data analysis can mask the information content of a given data set alone. However, since the information content of a data set is usually estimated with the Fisher matrix, priors are added to enforce an approximately Gaussian likelihood. Here, we estimate the information content of a Euclid-like weak lensing data set with and without priors. Additionally, we demonstrate how DALI allows switching to a speedy Hamiltonian Monte Carlo sampling of a highly curved likelihood with acceptance rates of ≈ 0.5. This shows how quick forecasts can be upgraded to accurate forecasts whenever needed.

Weak cosmic lensing is currently a field of intense focus: It allows the measurement of the cosmological parameters especially in the late Universe and is therefore an ideal probe for dark energy physics or models of modified gravity. Weak gravitational lensing has been observed with increasing singificance by e.g. CFHTLenS (Kilbinger et al., 2013; Heymans et al., 2013), allowing the determination of cosmological parameters.

In the future, weak lensing will be investigated on about a third of the sky with the upcoming Euclid satellite (Laureijs et al., 2011b). While the Euclid data set is not yet available, its constraining power on different extensions of the current cosmological standard model is being forecasted, see e.g. Amendola & Tsujikawa (2010c). Also, statistical techniques are being improved, or the data analysis is being refined, for example by switching from a two-dimensional weak lensing analysis to weak lensing tomography (Hu, 1999, 2002) and 3d cosmic shear (Heavens, 2003; Castro et al., 2005; Heavens et al., 2006), or by including higher-order polyspectra of the weak lensing shear (Munshi et al., 2010), or by combining lensing with other tracers of cosmological structure growth. There will be large-scale lensing surveys on the way to Euclid with an emphasis on dark energy, for instance the Kilo-degree Survey (KidS) (de Jong et al., 2013) and the Dark-Energy survey (DES) (Melchior et al., 2015).

All these different surveys need a tool in order to assess the information content of the data set under a specific analysis, i.e. a way to forecast the expected constraining power of the analysis. In principle, Monte Carlo Markov Chains (MCMC), Nested Sampling or grid-based likelihood evaluations are a well suited tool for these aim, but they are very time consuming, such that the Fisher matrix is used instead.

However, using the Fisher matrix, demands that the posterior can be well approximated by a multivariate Gaussian. As weak lensing possesses strong parameter degeneracies, this is not automatically guaranteed and usually, the weak lensing data set is combined with external probes or priors such that parameter degeneracies are broken. This removes non-Gaussianities at the cost of mixing weak lensing with other data sets.

Another question is how sensitive weak lensing is *alone*. In order to answer this question, existing non-Gaussianities have to be captured accurately. Obviously, DALI is a well-suited tool for this and we shall in the following compare it to the Fisher matrix approximation and to MCMC-evaluations.

Our cosmological parameter set consists of $\boldsymbol{\theta} = (\Omega_m, \sigma_8, n_s, h, w)$ which are the density of cold dark matter today, the normalization of the power spectrum, the primordial spectral index, the Hubble constant and a redshift-independent dark energy equation of state parameter. Our fiducial cosmology is $\Omega_m = 0.25, \sigma_8 = 0.8, n_s = 0.96, h = 0.7, w = -0.98$. We keep the density of baryons fixed to $\Omega_b = 0.04$. We model Euclid's weak lensing survey (Laureijs et al., 2011b) to reach out to a median redshift of 0.9 and to yield $\bar{n} = 4.8 \times 10^8$ galaxies per steradian. We assume that the shape measurement produces a Gaussian noise with standard deviation $\sigma_\epsilon = 0.3$. We use a sky fraction of $f_{\mathrm{sky}} = 0.35$ and a multipole range of 30 to 3000. For the redshift distribution of the galaxies, we use the common parameterization (Refregier & the DUNE collaboration, 2008)

$$p(z)\mathrm{d}z \propto \left(\frac{z}{z_0}\right)^2 \exp\left(-\left(\frac{z}{z_0}\right)^\beta\right) \mathrm{d}z. \tag{9.9}$$

We model the processing of an initial density fluctuation power spectrum during the early phases of the Universe with the ansatz

$$P_\delta(k, z) \propto k^{n_s} T^2(k, z), \tag{9.10}$$

where $T(k, z)$ is the transfer function for which we use the fit by Bardeen et al. (1986). We also include the additional variance of the cosmic density field on nonlinear scales as described by Smith et al. (2003). The growth of the power spectrum is then modelled by $D_+(a) = \delta(a)/\delta(1)$. The spectrum is normalised to the value σ_8,

$$\sigma_8^2 = \int_0^\infty \frac{k^2 \mathrm{d}k}{2\pi^2} W^2(8\ \mathrm{Mpc}/h \times k)\, P_\delta(k), \tag{9.11}$$

with a Fourier-transformed spherical top-hat $W(x) = 3j_1(x)/x$ as the filter function, where $j_1(x)$ is the spherical Bessel function of the first kind.
With the comoving Poisson equation

$$\Delta\Phi = 3\Omega_m/(2\chi_H^2)\delta, \tag{9.12}$$

we convert between the density contrast δ and its gravitational potential Φ such that from the spectrum of the CDM density perturbations the spectrum of the gravitational potential Φ follows to be

$$P_\Phi(k) \propto \left(\frac{3\Omega_m}{2\chi_H^2}\right)^2 k^{n_s-4}\, T(k)^2. \tag{9.13}$$

To include redshift information, one divides the galaxy sample into n_{bin} redshift intervals and computes the lensing potential ψ at the position θ for each redshift bin i separately,

$$\psi_i(\theta) = \int_0^{\chi_H} \mathrm{d}\chi\, W_i(\chi)\Phi, \tag{9.14}$$

hence $\psi_i(\theta)$ is related to the gravitational potential Φ by projection with the weight function $W_i(\chi)$

$$W_i(\chi) = 2\frac{D_+(a)}{a}\frac{G_i(\chi)}{\chi}. \tag{9.15}$$

Modes $\psi_{\ell m,i}$ of the lensing potential $\psi_i(\theta)$ are obtained by the decomposition $\psi_{\ell m,i} = \int \mathrm{d}\Omega\ \psi_i(\theta) Y^*_{\ell m}(\theta)$ into spherical harmonics $Y_{\ell m}(\theta)$. The distribution $p(z)\mathrm{d}z$ of the lensed galaxies in redshift is incorporated in the lensing efficiency function $G_i(\chi)$,

$$G_i(\chi) = \int_{\min(\chi,\chi_i)}^{\chi_{i+1}} \mathrm{d}\chi'\, p(\chi') \frac{\mathrm{d}z}{\mathrm{d}\chi'} \left(1 - \frac{\chi}{\chi'}\right) \tag{9.16}$$

with $\mathrm{d}z/\mathrm{d}\chi' = H(\chi')/c$ and the bin edges χ_i and χ_{i+1}, respectively. Angular spectra $C_{\psi,ij}(\ell)$ of the tomographic weak lensing potential can be written as the variance $\langle \psi_{\ell m,i} \psi^*_{\ell' m',j} \rangle = \delta_{\ell\ell'} \delta_{mm'} C_{ij}(\ell)$, which we approximate by the corresponding flat-sky expression,

$$C_{\psi,ij}(\ell) = \int_0^{\chi_H} \frac{\mathrm{d}\chi}{\chi^2}\, W_i(\chi) W_j(\chi)\, P_\Phi(k = \ell/\chi). \tag{9.17}$$

The convergence κ and the shear γ follow by double differentiation of the lensing potential with respect to angles, $\kappa = \ell^2\psi/2$, therefore their spectra are equal to $\ell^4 C_{\psi,ij}(\ell)/4$. Observed spectra of the weak lensing shear will contain a constant contribution $\sigma_\epsilon^2 n_{\mathrm{bin}}/\bar{n}$ known as shape noise, which translates into

$$\hat{C}_{\psi,ij}(\ell) = C_{\psi,ij}(\ell) + \sigma_\epsilon^2 \frac{n_{\mathrm{bin}}}{\bar{n}} \times \ell^4\, \delta_{ij}, \tag{9.18}$$

which will be at the same time the covariance matrix for measurements of modes $\psi_{\ell m,i}$. The covariance between different redshift bins, $C_{\psi,ij}(\ell)$, is non-zero for $i \neq j$ because light rays share the section between the observer and the closer tomography bin, such that they contain partially the same statistical information. We will mainly work with 2d-weak lensing, but for an additional comparison with 2-bin tomography, we choose bins such that they contain the same number $\bar{n}$ of galaxies.

Weak gravitational lensing derives its sensitivity on cosmological parameters from a combination of the amplitudes of gravitational potentials and geometry: The amplitude of the spectrum is set by $(\Omega_m \sigma_8)^2$. Additionally, the growth function $D_+(a)$, the peak of the power spectrum and the lensing efficiency function $G_i(\chi)$ carry a weaker dependence on Ω_m.

Therefore, the hyperbolic degeneracy line between these two parameters will approximately follow $\Omega_m \propto \sigma_8^{-1}$. The growth of dark matter fluctuations depends also on the background expansion and therefore also on the equation of state w for dark energy. The dark energy equation of state also influences the conversion of redshifts into distance measures. The shape of the spectrum $P(k)$ is mainly determined by n_s, h and Ω_m, which is reflected in the spectra $C_{\psi,ij}(\ell)$ more weakly due to the weighting functions $W_i(\chi)$. The shape parameter of the power spectrum is $\Gamma = \Omega_m h$, which introduces a degeneracy between Ω_m and h which will again be a hyperbolic line.

On the full sky, $2\ell + 1$ independent m-modes contribute to each ℓ-mode. The likelihood for the ℓ-modes when averaged over the m-modes is then

$$\mathcal{L}\left(\{\psi_{\ell m,i}\}\right) = \prod_\ell \mathcal{L}\left(\psi_{\ell,i} | \hat{C}_{\psi,ij}(\ell)\right)^{2\ell+1}. \tag{9.19}$$

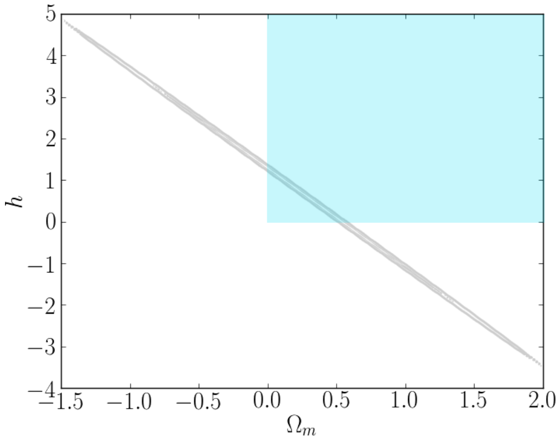

Figure 9.5: The marginalized Fisher matrix (grey) in the Ω_m, h-plane. The blue rectangle indicates the area bounded by the constraints $h > 0$, $\Omega_m > 0$ which might be interpreted as minimal priors that could be applied to foster the constraining power of the weak lensing data set.

Observational issues like an incomplete sky coverage and a point spread function lead to a coupling of different ℓ. Although of course relevant for analyses of real data sets, this is often omitted for forecasting (Hannestad et al., 2006) because non-diagonal or non-block diagonal covariance matrices are much harder to invert. For a Monte Carlo sampler, this inversion would be necessary for each sample, and we therefore assume that the different ℓ-modes decouple for the sake of speed.
The likelihood for each observed mode $\psi_{\ell m,i}$ if the theory predicts a covariance $\hat{C}_{\psi,ij}(\ell)$, is Gaussian in the data,

$$\mathcal{L}\left(\psi_{\ell m,i}\right) = \frac{1}{\sqrt{(2\pi)^{n_{\mathrm{bin}}}\mathrm{det}\hat{C}_\psi(\ell)}} \exp\left(-\frac{1}{2}\psi_{\ell m,i}(\hat{C}_\psi(\ell)^{-1})_{ij}\psi_{\ell m,j}\right), \tag{9.20}$$

due the fact that both the cosmic structures as well as the noise are approximately Gaussian random fields. Consequently, the logarithmic likelihood $L = -\ln\mathcal{L}$ is up to an additive constant equal to

$$L = \sum_\ell \frac{2\ell+1}{2} \left(\mathrm{Tr}\,(\ln)\,\hat{C}_\psi + (\hat{C}_\psi^{-1})_{ij}\;\psi_{\ell m,i}\psi_{\ell m,j}\right) \tag{9.21}$$

We will often refer to Eq. (9.21) as the true likelihood since no approximations apart from physical approximations such as the flat-sky approximation and integrating the lensing signal along a straight line were used so far, and the assumption that the projected lensing potential has nearly Gaussian fluctuation statistics.
In the following, we compare which parameter space is preferred by the Fisher matrix, DALI and MCMC and their Figure of Merits (FoM): The Fisher matrix allows a convenient definition of a FoM via the determinant of 2×2 submatrices

$$\mathrm{FoM} = \sqrt{\mathrm{det}\boldsymbol{F}_{2\times2}}. \tag{9.22}$$

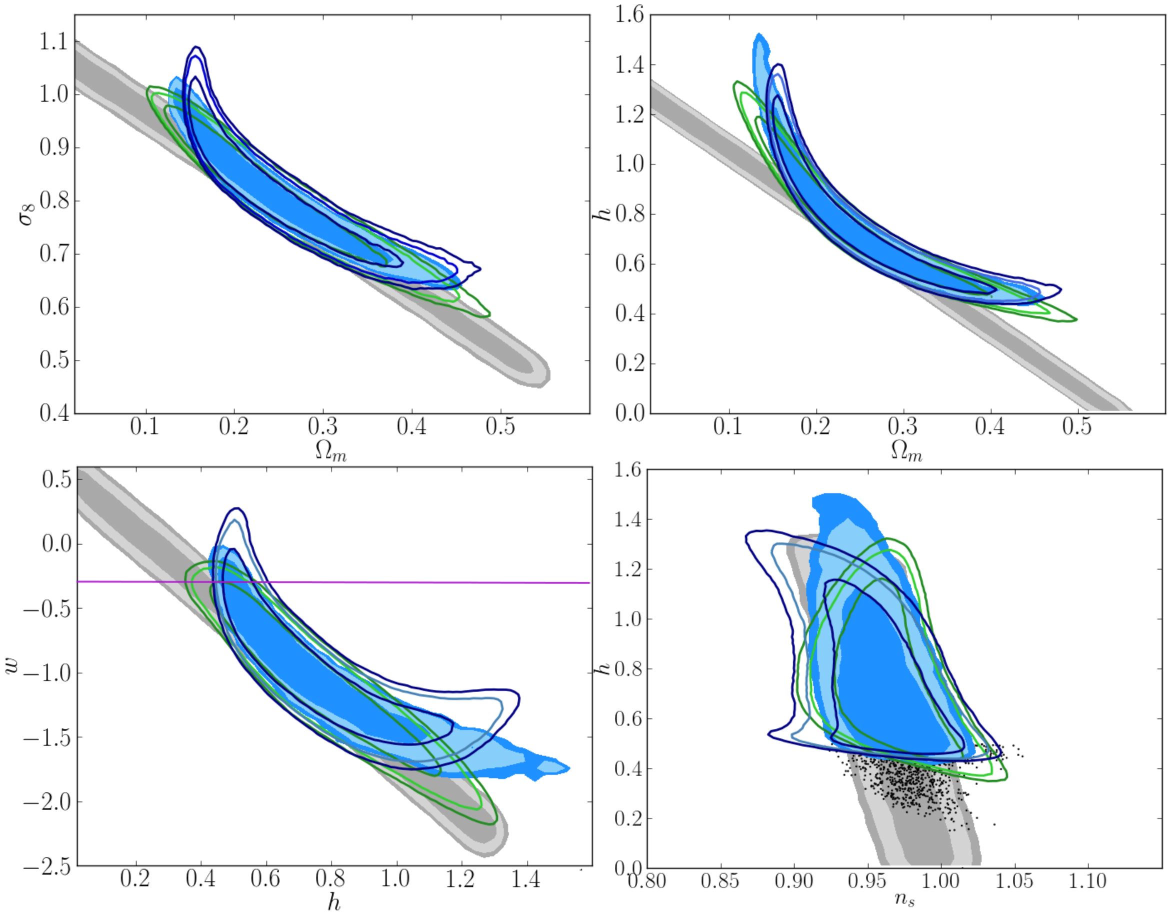

Figure 9.6: Comparison of the different likelihood approximations with the MCMC-sampled likelihood (depicted in blue) for 2d-weak lensing. The contours enclose 68%, 95% and 99% of the likelihood. Solid grey: Fisher approximation combined with the additional constraints $\Omega_m > 0$ and $h > 0$ (implemented by sampling from the Fisher matrix and discarding all unphysical samples before marginalizing). DALI with second-order derivatives of the covariance matrix is shown in open green contours, DALI with second and third derivatives is shown in open blue contours. The dots in the bottom right panel are samples drawn from the Fisher matrix approximated likelihood and are predicted to be points of high likelihood by the Fisher matrix. However, when calculating the unapproximated likelihood of these samples, they turn out to be extremely unlikely parameter combinations. This demonstrates that the sharp cutoff towards lower h in the bottom right panel is correct. In the top left and bottom right panel, the likelihood asymptotes roughly towards $h \approx 0.4$, and the cutoff in the n_s, h-plane is just a different projection of this behaviour. The purple line in the bottom left panel indicates the constraint $w < -1/3$ for accelerated expansion, and it can be seen that over $\sim 90\%$ of the MCMC and DALI contours fall within the parameter space of accelerated expansion, thereby indicating strongly the presence of a dark energy, whereas $\sim 30\%$ of the Fisher matrix cover parameter regions that would not lead to accelerated expansion.

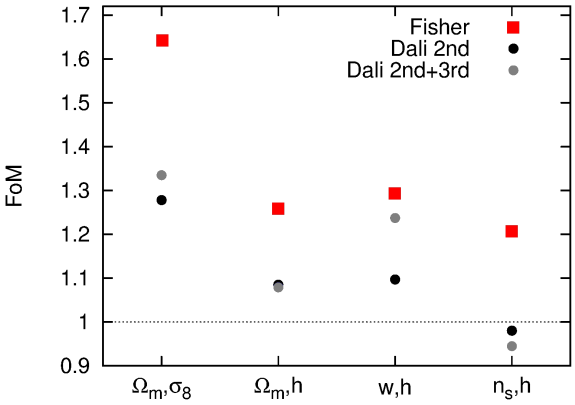

Figure 9.7: Figure of merit from the different approximations, relative to the MCMC figure of merit. The non-Gaussian DALI approximations always perform better than the Fisher matrix, although no clear trend can be made out. However, the DALI FoMs differ by maximally $\sim 30\%$ from the MCMC FoM, whereas the Fisher FoM differs about two times more, namely by up to $\sim 65\%$.

This corresponds to using the area enclosed by a chosen confidence contour in a given parameter plane as a FoM. We generalize this concept to our non-Gaussian forecasts by defining that the FoM shall be the area enclosed by the 95%-confidence contour.
We begin by evaluating the Fisher matrix for this setup. Fig. 9.5 shows the marginalized Fisher matrix approximated likelihood in the Ω_m, h-plane. Clearly, the Fisher matrix reaches far into unphysical regions of negative Ω_m and h. Sensitivity with respect to the Hubble constant enters weak lensing through the shape parameter $\Gamma = \Omega_m h$ which is a length scale and must therefore be positive definite. Negative h and Ω_m are therefore non-sensical and these negative values must be excluded. This shows, that the Fisher matrix cannot be used for a 2d weak lensing analysis for the Euclid satellite without enforcing by priors that the shape parameter has to be positive definite. In Sect. 9.4, we discuss why the Cramer-Rao inequality does not need to hold if unphysical parameter ranges are covered by the Fisher matrix.
For the comparison of the Fisher matrix with DALI and MCMC, we therefore augment the Fisher matrix with the prior knowledge $\Omega_m > 0, h > 0$. In practice, we implement this by drawing samples from the Fisher matrix approximated likelihood, and discarding all samples that fall into the unphysical regions. The introduction of these sharp cutoffs in Ω_m and h leads to a non-Gaussian likelihood approximation. This approximation is depicted in grey in Fig. 9.6 and was also used for comparing FoMs in Fig. 9.7.
A comparison with MCMC-sampled likelihoods shows that the data are actually more constraining than predicted by the Fisher matrix, and a DALI evaluation of the likelihood contours reveals that the problem is entirely due to non-Gaussianities and degeneracies between non-linear parameters.
In Fig. 9.6, a comparison between the Fisher matrix, DALI and MCMC samples of the

likelihood is shown. For MCMC and DALI, no prior constraints like $\Omega_m > 0$ were used. Highly curved degeneracy lines and asymmetric likelihood shapes are evident. These curved degeneracy lines are well approximated by DALI, although not perfectly. As the likelihood asymptotes to $h \approx 0.4$ in the Ω_m, h-plane and in the h, w-plane, negative h are excluded without the use of any priors. This shows that the 2d weak lensing analysis is able to predict a lower bound of $h \geq 0.5$ on its own. Also due to the highly curved likelihood shapes, Ω_m does not become negative but stays confined to the physical region. These strong changes in the allowed range of Ω_m and h in comparison to the Fisher matrix, propagate into the constraints of the remaining parameters σ_8, n_s, w. For dark energy, the curved DALI approximation predicts $0.3 > w > -2.0$. In contrast, the Fisher matrix allows much smaller and even positive w. This is interesting for the forecasting of dark energy constraints: An accelerated expansion of the universe requires $w < -1/3$. About one third of the Fisher matrix covers however the parameter space $w > -1/3$, and only two thirds fall into the parameter range of accelerated expansion. In contrast, DALI and MCMC both favour the accelerated expansion by a much larger degree: about 90% of their preferred parameter range corresponds to an accelerating universe. Note, that the fact that the Fisher matrix also covers parameter regions of decelerated expansion stems from it being by construction symmetric around the best fit point.

Already in case of the type Ia supernovae in Sect.9.1, we observed that this high symmetry leads to the Fisher matrix covering parameter ranges of decelerated expansion, whereas the real likelihood did not. The symmetry problem is especially evident in the bottom right panel of Fig. (9.6). There, we drew samples from the Fisher matrix (depicted as black dots) that are predicted by the Fisher matrix approximation to have a high likelihood, but that turn out to be extremely unlikely parameter combinations when inserted into the unapproximated likelihood.

Fig. 9.7, demonstrates that the DALI FoMs are closer to the MCMC FoM than the Fisher FoM. This illustrates that DALI is the more trustworthy tool for forecasting in the presence of non-Gaussianities and that modelling parameter degeneracies correctly, can remove the necessity to impose priors in order to establish a Gaussian likelihood. This allows the scientific return of a single data set to be optimized independently of external priors.

We will now test which non-Gaussianities remain if a prior on h is introduced to a 2d-weak lensing analysis or if the analysis is improved by updating to a 2-bin tomography. This testing for non-Gaussianities is not possible with the Fisher matrix, but of high interest, as the usual forecasts use the Fisher matrix nonetheless. We use a Gaussian prior on h with standard deviation $\sigma_h = 0.03$, roughly corresponding to the precision of current local constraints on the Hubble constant (Riess et al., 2011). Fig. 9.8 shows that even when using this prior, the posterior likelihood is not peaked sharply enough but non-Gaussianities remain. It also shows that switching to 2-bin tomography outperforms the inclusion of this prior into a 2d weak lensing analysis. Figure 9.9 shows a comparison of the Fisher and DALI forecasts and MCMC for 2-bin tomography. The likelihood is then well approximated by a multivariate Gaussian.

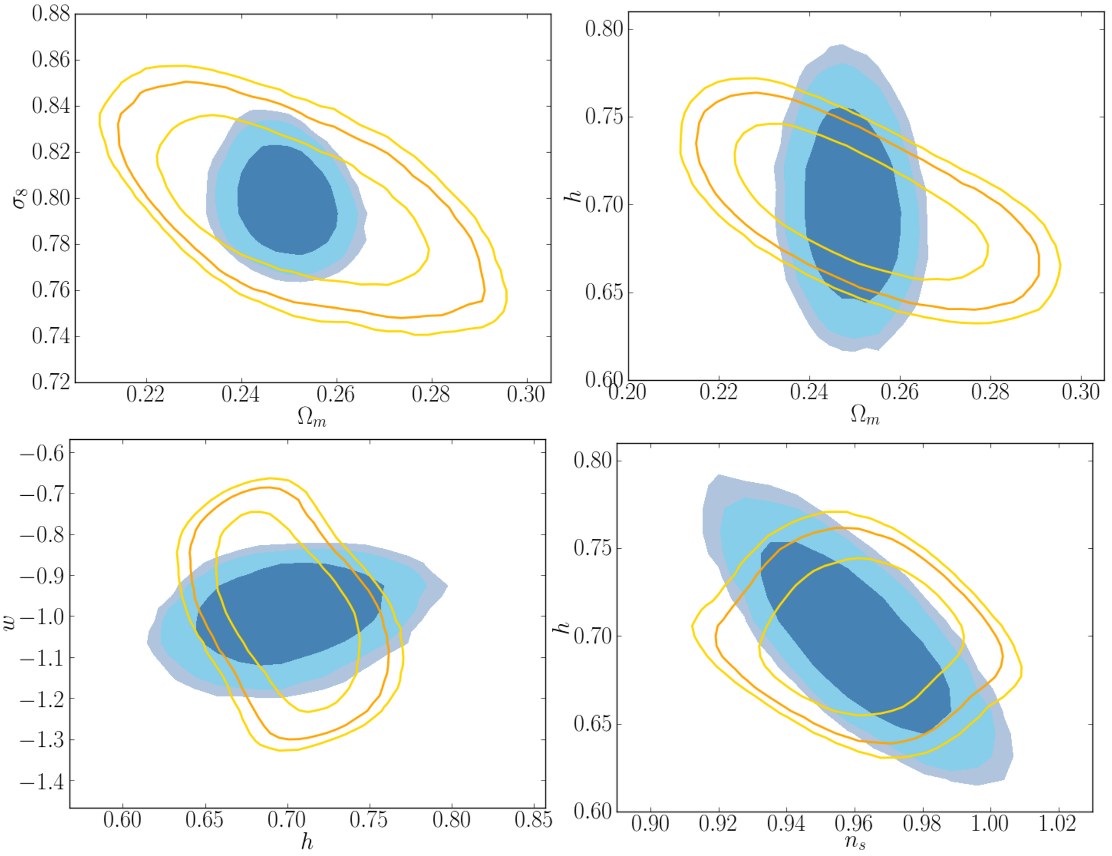

Figure 9.8: Yellow: DALI constraints for 2d weak lensing with a Gaussian prior on h from local measurements. Since DALI with second-order derivatives and DALI with second and third-order derivatives agree very well, only the latter are shown. Blue: MCMC-sampled likelihood for a 2-bin tomography analysis of the same data set without a prior on h.

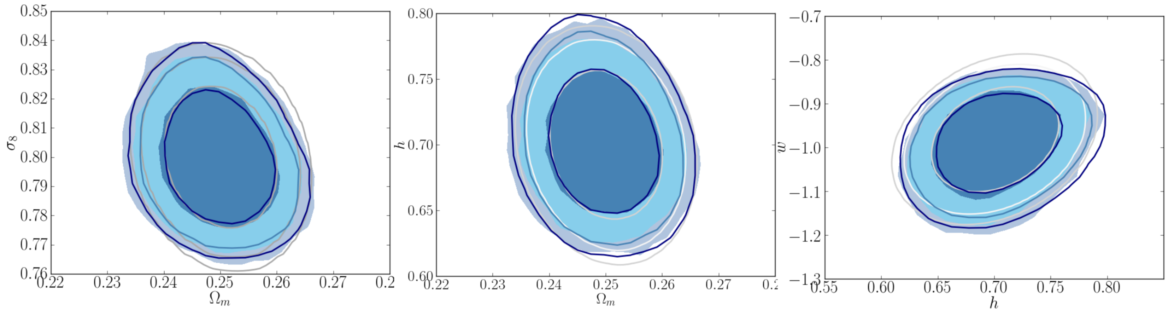

Figure 9.9: 2-bin tomography for the same weak lensing survey increases the constraining power enough such that the Fisher matrix (grey contours) agree well with the MCMC contours (solid blue). DALI finds nearly the same confidence contours as the Fisher matrix since only minor non-Gaussianity is present.

9.3 Accelerating Hamilton Monte Carlo with DALI

In this section, we use DALI to aid MCMC-sampling the unapproximated weak lensing likelihood. Since MCMC chains need a data set for calculating the likelihood, we generated a ΛCDM mock data set which our sampler reads in.

The Metropolis-Hastings algorithm (Metropolis, 1985) works well for approximately multivariate Gaussian likelihoods but has problems with following highly curved likelihoods. We therefore employed a Hamilton-Monte-Carlo (HMC) sampler, which uses Hamiltonian dynamics for describing a random walk on a potential P corresponding to the logarithmic likelihood,

$$P(x_\mu) = -\ln(\mathcal{L}), \tag{9.23}$$

and a kinetic energy to introduce the randomness needed for sampling.

The algorithm takes multiple leapfrog steps along contours of approximately constant likelihood before performing a Metropolis-Hastings step by which it decides whether the new point is accepted or rejected (Hajian, 2007b). For each leapfrog step, the HMC sampler takes derivatives of the logarithmic likelihood and follows these, thereby adjusting well to curved likelihoods, see Sect. 6.10.

Calculating derivatives of the true log-likelihood can be numerically costly. A gain in performance can then be achieved if the log-likelihood is replaced by an approximation which is fast to evaluate, such as DALI. Consequently, we do not use the log-likelihood of Eq. (9.21) for the sampler, but its DALI approximation for the leapfrog steps along the potential. Calculating the true weak-lensing likelihood is then only needed in the Metropolis-Hastings steps. This procedure speeds up the performance of our sampler by a factor ranging between 30 and 100, depending on how many leapfrog steps were done in each iteration of the HMC algorithm.

A potential issue with using DALI contours to guide an HMC sampler is that DALI might exclude regions of the parameter space that are actually preferred by the true likelihood. In order to avoid this problem, we introduce a temperature to widen the potential Eq. (9.23),

$$\ln P(x_\mu) \to \ln P(x_\mu)/T. \tag{9.24}$$

If the temperature is set too high, the contours of the potential Eq. (9.24) will not generate samples that follow the true likelihood well. This leads to a reduction of the acceptance rate. We find that $T = 3$ leads to an acceptance rate between 0.3 and 0.5 while still giving the sampler the possibility to reach all regions in parameter space that are erroneously not covered by the DALI approximation. In Fig. 9.10 we plot samples of such an MCMC-chain, demonstrating that the sampler has indeed been able to cover the true likelihood fully: The accepted samples are surrounded by a rim of rejected samples. This rim shows that the sampler had the chance to explore regions of parameter space with low likelihood.

Note, that this procedure is not limited to forecasting: the DALI contours can also be used to accelerate an HMC sampler that calculates the likelihood conditioned on real data, as long as the peak of the real likelihood lies in the vicinity of the fiducial for which DALI was calculated. All that would need to be changed is that the HMC sampler reads in the actual

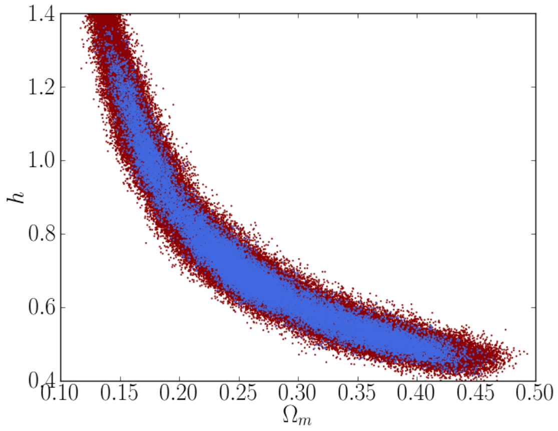

Figure 9.10: Samples from a Hamilton Monte Carlo Chain: For the approximate potential, the tempered DALI likelihood Eq. (9.23) was used. Rejected samples are depicted in red, accepted samples in blue. Clearly visible is a red rim where accepted and rejected samples do not mix, demonstrating that the sampler was able to reach all points in parameter space that are preferred by the data. In the other two-dimensional planes, there also exists a ring of rejected samples.

data, instead of the mock data. We therefore advise to store DALI tensors or even complete DALI-driven HMC codes from forecasts, as they can be recycled in the data analysis once the real data are available.
In contrast, using the Fisher matrix as an approximate potential for the HMC sampler has proven ineffective: since it does not capture the curvature of the likelihood, the sampler is frequently guided towards regions of extremely low likelihood if it follows the isocontours of the Fisher approximation. Consequently, even after adjusting the number of leapfrog steps, no higher acceptance rate than 0.02 in our application could be gained, while many regions of the preferred parameter space were not sampled (in an acceptable time) at all.

9.4 Revisiting the Cramer-Rao inequality

The Fisher matrix has also become widely used in cosmology, as it represents the lower bound on the variance of parameter estimators. This statement is known as Cramer-Rao inequality: if $\sigma^2(\theta_i)$ is the variance of an unbiased parameter estimator θ_i, then the Cramer-Rao bound states (Tegmark et al., 1997)

$$\sigma^2(\theta_i) \geq 1/\sqrt{(F_{ii})}, \tag{9.25}$$

if all other parameters θ_j are fixed. If other parameters θ_j are to be inferred jointly with θ_i, then the minimum achievable variance increases to

$$\sigma^2(\theta_i) \geq (\boldsymbol{F}^{-1})_{ii}^{1/2}, \tag{9.26}$$

where it was used that the inversion marginalizes out the other parameters.
The tenor in cosmology is, that the Cramer-Rao inequality guarantees that the Fisher matrix leads to the tightest achievable parameter constraints. However, let us have a closer look at the assumptions: first, the Cramer-Rao inequality makes a statement about the *variance* of a parameter. This will only be the identical to the 68%-confidence contour, if the posterior is Gaussian. In general, cosmologists are however interested in the confidence contours, not in the variance, so care has to be taken. More elaborate considerations about the Cramer-Rao inequality can be found below.

9.4.1 Frequentist and Bayesian Fisher matrix

Suppose there exists an observable $\mathbf{m} = [m_1, ..., m_n]$ to which a theoretical prediction by a model μ corresponds that is a function of a parameter set: $\boldsymbol{\mu} = \boldsymbol{\mu}(p_1, ..., p_m)$. In the Fisher matrix formalism the observed outcome is the mean values of the observables assumed as the null hypothesis. This method allows a quick way to estimate errors on cosmological parameters, given errors in observable quantities. The frequentist Fisher matrix is defined as the Hessian of the log-likelihood function $\mathcal{L} = -\log(P)$,

$$F_{\alpha\beta} = \langle -(\log P(\mathbf{m}, \boldsymbol{\mu}))_{,\alpha\beta} \rangle \,. \tag{9.27}$$

This can be simplified as follows

$$\begin{aligned} F_{\alpha\beta} &= \langle -(\log P(\mathbf{m}, \boldsymbol{\mu}))_{,\alpha\beta} \rangle \\ &= \left\langle -\frac{P_{,\alpha\beta}}{P} + (\log P)_{,\alpha} (\log P)_{,\beta} \right\rangle \\ &= \langle (\log P)_{,\alpha} (\log P)_{,\beta} \rangle \end{aligned} \tag{9.28}$$

since

$$\left\langle \frac{P_{,\alpha\beta}}{P} \right\rangle = \int \frac{P_{,\alpha\beta}}{P} P d^n x = \partial_{\alpha\beta} \int P d^n x = 0 \,. \tag{9.29}$$

In the case of Gaussian data, the likelihood for n data points is

$$P = \frac{1}{(2\pi)^{n/2}\sqrt{|C|}} e^{-\frac{1}{2}(m_i - \mu_i) C^{-1}_{ij} (m_j - \mu_j)} . \tag{9.30}$$

The Fisher matrix is then (suppressing the data indices i, j and implicitly summing over them unless otherwise specified) (Tegmark et al., 1997)

$$F_{\alpha\beta} = \frac{1}{2} \mathrm{Tr}\Big[C_{,\alpha} C^{-1} C_{,\beta} C^{-1}\Big] + \boldsymbol{\mu}_{,\alpha} C^{-1} \boldsymbol{\mu}_{,\beta} \,. \tag{9.31}$$

By the Cramer-Rao inequality, a model parameter p_α cannot have a variance smaller than $1/(F_{\alpha\alpha})^{1/2}$ (evaluated for unbiased estimators) when all other parameters are fixed, or a precision $(F^{-1})^{1/2}_{\alpha\alpha}$ when all other parameters are marginalized over. Notice however that

the Cramer-Rao inequality concerns variances and does not say anything about the relative size of the confidence regions.
The purely Bayesian definition of the Fisher matrix is instead:

$$F^B_{\alpha\beta} = -\log(P(\mathbf{m},\boldsymbol{\mu}))_{,\alpha\beta}\Big|_{BF}, \tag{9.32}$$

where the derivatives have to be evaluated at the best fit values of the parameters, i.e. for parameters such that

$$P_{,\alpha} = 0\,. \tag{9.33}$$

This definition makes no reference to the average over the data, which in the Bayesian context are fixed once and for all by the current experiment. Expressions (9.27) and (9.32) are however in general different and the Cramer-Rao inequality does not hold in general for F^B. We can also write

$$\begin{aligned} F^B_{\alpha\beta} &= -\log(P(\mathbf{m},\boldsymbol{\mu}))_{,\alpha\beta}|_{BF} \\ &= -\frac{P_{,\alpha\beta}}{P}\Big|_{BF} + (\log P)_{,\alpha}(\log P)_{,\beta}|_{BF} \\ &= -\frac{P_{,\alpha\beta}}{P}\Big|_{BF} \end{aligned} \tag{9.34}$$

due to Eq. (9.33) .
We show now that the only cases in which (9.27) (evaluated at the best fit) and (9.32) coincide are *a*) when the data are Gaussian and the parameters enter in a linear way in the mean and in the variance and *b*) in the case of forecasting.
In fact we have

$$(\log P)_{,\alpha} = -\frac{1}{2}\mathrm{Tr}\Big[C_{,\alpha}C^{-1} + CD_{,\alpha} - C^{-1}C_{,\alpha}C^{-1}D\Big]\,, \tag{9.35}$$

where we defined the data matrix with its components

$$D_{ij} = X_i X_j \tag{9.36}$$

and the data vector

$$\mathbf{X} \equiv \mathbf{m} - \boldsymbol{\mu}\,. \tag{9.37}$$

Notice that

$$D_{ij,\alpha} = -2\mu_{i,\alpha}X_j\,. \tag{9.38}$$

The best fit condition $(\log L)_{,\alpha} = 0$ gives

$$\mathrm{Tr}[C^{-1}C_{,\alpha}C^{-1}D] = \mathrm{Tr}[C_{,\alpha}C^{-1} + C^{-1}D_{,\alpha}]\,. \tag{9.39}$$

If C does not depend on the parameters, the best fit equation becomes

$$D_{,\alpha} = 0\,. \tag{9.40}$$

Finally we have

$$\begin{aligned} F^B_{\alpha\beta} = -(\log P)_{,\alpha\beta}|_{BF} = \\ \frac{1}{2}\mathrm{Tr}\Big[C^{-1}C_{,\alpha\beta}(I - C^{-1}D) - C^{-1}C_{,\beta}(C^{-1}C_{,\alpha} \\ + C^{-1}D_{,\alpha}) + C^{-1}C_{,\alpha}C^{-1}C_{,\beta}C^{-1}D \\ + C^{-1}C_{,\beta}C^{-1}C_{,\alpha}C^{-1}D - 2\boldsymbol{\mu}_{,\alpha\beta}C^{-1}\mathbf{X} \\ + 2\boldsymbol{\mu}_{,\alpha}C^{-1}\boldsymbol{\mu}_{,\beta} - C^{-1}C_{,\alpha}C^{-1}D_{,\beta}\Big]_{BF}. \end{aligned} \tag{9.41}$$

Inserting the best fit condition (9.39) we obtain

$$\begin{aligned} F^B_{\alpha\beta} &= -(\log P)_{,\alpha\beta}|_{BF} \\ &= \frac{1}{2}\mathrm{Tr}\Big[C^{-1}C_{,\alpha\beta}(I - C^{-1}D) + C^{-1}C_{,\beta}C^{-1}C_{,\alpha} \\ &- 2\boldsymbol{\mu}_{,\alpha\beta}C^{-1}\mathbf{X} + 2\boldsymbol{\mu}_{,\alpha}C^{-1}\boldsymbol{\mu}_{,\beta}\Big]_{BF} \\ &= F_{\alpha\beta(BF)} + \Sigma_{\alpha\beta(BF)}\,, \end{aligned} \tag{9.42}$$

where

$$\Sigma_{\alpha\beta} = \frac{1}{2}\mathrm{Tr}\Big[C^{-1}C_{,\alpha\beta}(I - C^{-1}D)\Big] - \boldsymbol{\mu}_{,\alpha\beta}C^{-1}\mathbf{X}\,. \tag{9.43}$$

This shows that if the parameters enter linearly in C and in μ, the two Fisher matrices coincide (always assuming Gaussian data). The matrix Σ expresses the difference between frequentist and Bayesian Fisher matrices. The first one is the one that ensures the Cramer-Rao inequality. The second one is the matrix that approximates the posterior.
Now, when we do forecasts, we generate mock data with variance given by C and mean given by μ. If we evaluate the average Fisher matrix for many mock data, then we obtain

$$\langle D \rangle = C \tag{9.44}$$

$$\langle X \rangle = 0 \tag{9.45}$$

so that

$$\langle \Sigma_{\alpha\beta} \rangle = 0\,. \tag{9.46}$$

When forecasting, we in general identify the two Fisher matrices, or rather we can say that the generation of mock data implements the frequentist approach. Analysing real data, however, one should use the Bayesian Fisher matrix, because this is the approximation to the posterior.

9.4.2 Fisher matrix versus Fisher information

We demonstrate that the Cramer-Rao inequality does not need to apply if the Fisher matrix covers unphysical parameter ranges. To illustrate this, we imagine a distribution function $f(\boldsymbol{X}, \theta)$, where $\boldsymbol{X}$ is the data set, and for simplicity only one parameter θ shall be estimated (else, one would simply need to marginalize over the other parameters).

The Cramer-Rao inequality actually holds for the Fisher information $\mathcal{I}$ (Cramer, 1946), which is the averaged squared gradient of the log-likelihood,

$$\mathcal{I} = \int f(\boldsymbol{X}, \theta)[\partial_\theta \log(f(\boldsymbol{X}, \theta))]^2 \mathrm{d}^n x. \tag{9.47}$$

In contrast, the diagonal Fisher matrix elements are the averaged curvature of the negative log-likelihood

$$F_{\theta\theta} = -\int f(\boldsymbol{X}, \theta)\partial_\theta\partial_\theta \log(f(\boldsymbol{X}, \theta))\mathrm{d}^n x. \tag{9.48}$$

Explicitely calculating the second derivatives in Eq. 9.48 shows that the Fisher matrix and the Fisher information are related by

$$F_{\theta\theta} = \mathcal{I} - \int \partial_\theta\partial_\theta f(\boldsymbol{X}, \theta)\mathrm{d}^n x. \tag{9.49}$$

In order for the Fisher matrix to be identical to the Fisher information, the second term needs to vanish, which will be the case if the differentiation with respect to the parameters and the averaging over the data interchange. This however requires that the distribution $f(\boldsymbol{X}, \theta)$ and its derivatives exist for all combinations of the data and the parameters. This will usually be the case, but not always, because physical parameters often have bounds beyond which they become non-sensical (e.g. negative masses). So if the Fisher matrix reaches into unphysical domains of the parameter space, the Cramer-Rao inequality can break down. We found this to be true in the case of the Fisher matrix from weak lensing in Fig. 9.5.

9.5 Analytical inversion with the Sherman-Morrison-Woodbury formula

Having taken a deeper look at the Fisher matrix in (9.4.1) and (9.4.2), let us here also include a new result derived in Amendola & Sellentin (2016) about how to invert the Fisher matrix analytically under the addition of priors. This allows to optimize the constraints on interdependent parameters much more effectively than by simple trial-and-error inversion under ad hoc assumed priors. As the outcome of inverting a big matrix cannot be foreseen intuitively, the following formula is a great help in understanding better which parameters affect the uncertainties of other parameters the most.

The idea is entirely based on the Sherman–Morrison–Woodbury formula (Sherman & Morrison, 1950; Woodbury, 1950). This formula states that if M is a square matrix and $\boldsymbol{u}, \boldsymbol{v}$ are vectors then

$$(\mathsf{M} + \boldsymbol{u}\boldsymbol{v}^\mathsf{T})^{-1} = \mathsf{M}^{-1} - \frac{\mathsf{M}^{-1}\boldsymbol{u}\boldsymbol{v}^\mathsf{T}\mathsf{M}^{-1}}{1 + \boldsymbol{v}^\mathsf{T}\mathsf{M}^{-1}\boldsymbol{u}}, \tag{9.50}$$

where T denotes transposition. The formula allows to quickly find the inverse of M when a matrix $\boldsymbol{u}\boldsymbol{v}^T$ is added to M. If we demand that M is a non-degenerate Fisher matrix, i.e. the inverse of the covariance matrix of parameters, then M is symmetric and positive definite. If we further set $\boldsymbol{u} = \boldsymbol{v}$, we can construct a prior matrix $\mathsf{P} = \boldsymbol{u}\boldsymbol{u}^\mathsf{T}$ where the vector $\boldsymbol{u}$ is

$$\boldsymbol{u} = \{0, 0, ...p_k^{-1}, ...\} = p_k^{-1}\hat{\boldsymbol{u}}, \tag{9.51}$$

where p_k is the prior standard deviation on the control parameter θ_k, that shall quantify the expected improvement on θ_k in a feasible future experiment. The vector $\hat{\boldsymbol{u}}$ is the k-th basis vector. Introducing a matrix $\hat{\mathsf{P}} = \hat{\boldsymbol{u}}\hat{\boldsymbol{u}}^T$ i.e. a matrix whose elements are all zero except for $\hat{\mathsf{P}}_{kk} = 1$, we can write $\mathsf{P} = p_k^{-2}\hat{\mathsf{P}}$, and Eq. (9.50) then specializes to

$$(\mathsf{M} + \mathsf{P})^{-1} = \mathsf{M}^{-1} - \frac{\mathsf{M}^{-1}\mathsf{P}\mathsf{M}^{-1}}{1 + \mathrm{Tr}(\mathsf{P}\mathsf{M}^{-1})}. \tag{9.52}$$

The variance σ_i^2 of our target parameter θ_i, after having marginalized over all other parameters, is given by $\sigma_i^2 = \mathsf{M}_{\mathrm{ii}}^{-1}$. Its improved variance $\sigma_{i,\mathrm{new}}^2$ after having added a prior to the k-th control parameter is then

$$\begin{aligned}
\sigma_{i,\mathrm{new}}^2 &= (\mathsf{M} + \mathsf{P})_{ii}^{-1} && (9.53)\\
&= \mathsf{M}_{ii}^{-1} - \frac{(\mathsf{M}^{-1}\mathsf{P}\mathsf{M}^{-1})_{ii}}{1 + \mathrm{Tr}(\mathsf{P}\mathsf{M}^{-1})} && (9.54)\\
&= \mathsf{M}_{ii}^{-1} - \frac{p_k^{-2}(\mathsf{M}^{-1}\hat{\mathsf{P}}\mathsf{M}^{-1})_{ii}}{1 + p_k^{-2}\mathrm{Tr}(\hat{\mathsf{P}}\mathsf{M}^{-1})} && (9.55)\\
&= \sigma_i^2 - \frac{p_k^{-2}(\mathsf{M}^{-1}\hat{\mathsf{P}}\mathsf{M}^{-1})_{ii}}{1 + p_k^{-2}\sigma_k^2} && (9.56)\\
&= \sigma_i^2 - \frac{(\mathsf{M}^{-1}\hat{\mathsf{P}}\mathsf{M}^{-1})_{ii}}{p_k^2 + \sigma_k^2}. && (9.57)
\end{aligned}$$

Now we use the fact that the fully marginalized variance of the control parameter is $\mathrm{Tr}(\hat{\mathsf{P}}\mathsf{M}^{-1}) = \sigma_k^2$ and

$$(\mathsf{M}^{-1}\hat{\mathsf{P}}\mathsf{M}^{-1})_{ii} = \rho_{ik}^2\sigma_i^2\sigma_k^2, \tag{9.58}$$

where $\rho_{ik} = \mathsf{M}_{ik}^{-1}/\sqrt{\mathsf{M}_{ii}^{-1}\mathsf{M}_{kk}^{-1}}$ is the correlation coefficient and one has $|\rho_{ik}| \leq 1$ because of positive-definiteness of M. So we derive our main result

$$\sigma_{i,\mathrm{new}}^2 = \sigma_i^2 - \frac{\rho_{ik}^2\sigma_i^2\sigma_k^2}{p_k^2 + \sigma_k^2}, \tag{9.59}$$

which describes directly and transparently how the variance of the target parameter θ_i decreases if we measure better the parameter θ_k. This formula therefore is a highly convenient replacement for the typically needed numerical inversion where no intuition can be built.

Also, if the prior on the control parameter k is very weak, i.e. $p_k \to \infty$, the error σ_i of the target parameter, does not change. This equation can be trivially applied even when the control parameter coincides with the target parameter, by putting $i = k$ and the self correlation $\rho_{ik} = 1$.

From the previous equation, the decrease $\Delta\sigma_i^2 = \sigma_{i,\mathrm{new}}^2 - \sigma_i^2$ follows to be

$$\frac{\Delta\sigma_i^2}{\sigma_i^2} = -\frac{\rho_{ik}^2}{1+\varepsilon}, \tag{9.60}$$

where $\varepsilon = p_k^2/\sigma_k^2$. This tells us that if we add a prior to the error on the control parameter which is ε times the current error, then the target parameter constraint decreases by a fraction $\rho_{ik}^2/(1+\varepsilon)$. At most, the fractional decrease is then $\rho_{ik}^2 \leq 1$. So the very simple recipe for choosing the most convenient control parameter to improve the estimation of the target parameter, is to select the most correlated one. This of course was to be entirely expected; our formula (9.60) quantifies the effect in a very simple way as a function of the correlation coefficient and of the ratio ε.

A generalization to several control or target parameters is described in Amendola & Sellentin (2016).

Chapter 10

Parameter inference in case of an uncertain covariance matrix

All of the previous works have assumed that the statistical fluctuations of the data are Gaussian, and that this Gaussian has a precisely known covariance matrix. In the following, we will relax this assumption and demand that the covariance matrix can only be estimated, such that it will have a non-zero uncertainty itself. This chapter follows closely our publication Sellentin & Heavens (2016), but amends it by more detailed calculations.

10.1 Estimated covariance matrices

A very common problem in statistical inference concerns data that are Gaussian-distributed. The likelihood of the observed data $\boldsymbol{X}_\mathrm{o}$ is a multivariate Gaussian, characterised only by a mean data vector $\boldsymbol{\mu}$ and a covariance matrix $\boldsymbol{\Sigma}$:

$$G(\boldsymbol{X}_\mathrm{o}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{\sqrt{|2\pi\boldsymbol{\Sigma}|}} \exp\left[-\frac{1}{2}(\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})\right]. \tag{10.1}$$

The posterior probability of the parameters is proportional to the likelihood, now treated as a function of the parameters (through the dependence of the mean and the covariance matrix), multiplied by a suitable prior. Ideally one has analytic expressions for the mean and covariance in terms of the model parameters, but in many cases these are not available, and one or both may need to be estimated from simulated data which mimic the experiment that is to be analysed (e.g., Semboloni et al. (2006); Heymans et al. (2013)), or from the data themselves (e.g., Budavári et al. (2003)). However, although an unbiased simulated covariance matrix $\boldsymbol{S}$ can be constructed, its inverse is not an unbiased estimator of the inverse (or precision) matrix $\boldsymbol{\Sigma}^{-1}$, which is what is needed in the likelihood Eq. (10.1). One can construct an unbiased estimator of $\boldsymbol{\Sigma}^{-1}$ by a rescaling of $\boldsymbol{S}$ (Anderson, 2003), as advocated by Hartlap et al. (2007). This widens up the credible intervals. If simulations are computationally cheap, then one can generate a large number N of simulated datasets and obtain an accurate estimate of the covariance matrix. This asymptotic regime occurs only when N far exceeds the size of the data vector, p. In many practical cases this is not possible, and the number of simulated datasets is small, with the consequence that statistical noise in the precision matrix propagates into errors in the parameters (Taylor et al., 2013; Dodelson & Schneider, 2013; Hamimeche & Lewis, 2009). However, there is a more fundamental difficulty with the approach adopted, as it assumes that the likelihood is still Gaussian, albeit with a different precision matrix, whereas in fact it is not.

A principled way to tackle the problem is to recognise that the simulated data provide *samples* of the covariance matrix, so $\boldsymbol{S}$ is itself a random object, based on a number of simulations. For Gaussian data, we have the advantage that the sample distribution of $\boldsymbol{S}$ is known, for a given true covariance matrix $\boldsymbol{\Sigma}$, and we can exploit this, with a suitable prior, by constructing the probability of $\boldsymbol{\Sigma}$ conditional on the sample $\boldsymbol{S}$, and then marginalising over the unknown covariance matrix $\boldsymbol{\Sigma}$. This can be done analytically for our preferred choice of Jeffreys prior for $\boldsymbol{\Sigma}$. As a consequence, we properly propagate the uncertainty in the covariance matrix into the final inference, computing the quantity we want, i.e., the likelihood given the *simulated* covariance matrix $\boldsymbol{S}$ and the number of samples N on which is it based: $P(\boldsymbol{X}_\mathrm{o}|\boldsymbol{\mu}, \boldsymbol{S}, N)$. This object, where we keep the dependence on the number of simulated datasets N explicit to emphasize its importance, is the main result of this work. It is not Gaussian, but rather follows a modified version of the multivariate t-distribution. In practical terms, it is no more expensive to compute than the Hartlap-scaled Gaussian likelihood, but statistically sound, and can be retrospectively applied to many analyses that have used a different likelihood function by appropriate re-weighting of points, provided that the chains adequately sample the parameter space that the t-distribution favours.

10.2 Replacing a true covariance matrix by an estimator

We adopt the usual position in which cosmologists find themselves when inferring model parameters $\boldsymbol{\theta}$ from a data set: observing epochs of the Universe, we usually have just one observation $\boldsymbol{X}_\mathrm{o}$, which is a single realization of some statistical process. The errors on the data vector $\boldsymbol{X}_\mathrm{o}$ are often not known analytically, but there may be reason to believe that $\boldsymbol{X}_\mathrm{o}$ stems from a multivariate Gaussian process

$$\boldsymbol{X}_\mathrm{o} \sim \mathcal{N}_p\left(\boldsymbol{\mu}(\boldsymbol{\theta}), \boldsymbol{\Sigma}(\boldsymbol{\theta})\right), \tag{10.2}$$

and the specification of a covariance matrix $\boldsymbol{\Sigma}(\boldsymbol{\theta})$ is therefore sufficient to describe the statistical scatter of $\boldsymbol{X}_\mathrm{o}$. Here, p is the dimension of the data vector and its covariance matrix. Both, the mean of the data, and the covariance matrix may depend on the parameters $\boldsymbol{\theta}$ that are to be inferred. In the following, we do not denote this dependence on $\boldsymbol{\theta}$ any more, but it is still implied.

Under these conditions, and if $\boldsymbol{\Sigma}$ were known precisely, the likelihood would be a multivariate Gaussian Eq. (10.1). In the following we show that this Gaussian likelihood should be replaced, if the true $\boldsymbol{\Sigma}$ must be replaced by an estimator of it. One method – viable for

Frequentists – of estimating the covariance matrix, is to draw further independent data vectors from the distribution of $\boldsymbol{X}_\mathrm{o}$ and to calculate their sample covariance. Sadly, these repeated independent measurements are usually impossible in cosmology.

However, if we believe that we can simulate the observation then we are able to generate further samples $\boldsymbol{X}_i$ that are statistically equivalent to $\boldsymbol{X}_\mathrm{o}$. The covariance matrix can then be estimated from these simulations.

If we run N independent simulations, then

$$\bar{\boldsymbol{X}} = \frac{1}{N}\sum_{i=1}^{N}\boldsymbol{X}_i \tag{10.3}$$

is the average, and an unbiased estimator of $\boldsymbol{\Sigma}$ is

$$\boldsymbol{S} = \frac{1}{N-1}\sum_{i=1}^{N}(\boldsymbol{X}_i - \bar{\boldsymbol{X}})(\boldsymbol{X}_i - \bar{\boldsymbol{X}})^T. \tag{10.4}$$

We would now like to replace the inverse covariance matrix $\boldsymbol{\Sigma}$ by its estimator $\boldsymbol{S}^{-1}$. In the following, we derive an analytical replacement for Eq. (10.1) after this exchange. We then study the effects that an estimated covariance matrix introduces to parameter estimation.

10.3 Derivation of the multivariate t-distribution

If $\boldsymbol{\Sigma}$ is not exactly known, it must be marginalized over, adopting a fiducial distribution for it that captures as much information about $\boldsymbol{\Sigma}$ as we have access to. Given an estimate $\boldsymbol{S}$ from simulations, this fiducial distribution can be derived as follows.

Any matrix of the type $\boldsymbol{M} = \sum_{i=1}^{m}\boldsymbol{Y}_i\boldsymbol{Y}_i^T$ is by construction a Wishart matrix (Anderson, 2003; Mardia et al., 1979), if $\boldsymbol{Y}$ is drawn from a multivariate Gaussian. We therefore know that $\boldsymbol{S}$ follows the Wishart distribution

$$\mathcal{W}(\boldsymbol{S}|\boldsymbol{\Sigma}/n, n) = \frac{|\boldsymbol{S}|^{\frac{n-p-1}{2}}\exp\left(-\frac{1}{2}n\mathrm{Tr}\left(\boldsymbol{\Sigma}^{-1}\boldsymbol{S}\right)\right)}{2^{\frac{pn}{2}}|\boldsymbol{\Sigma}/n|^{\frac{n}{2}}\Gamma_p\left(\frac{n}{2}\right)}, \tag{10.5}$$

where we call $n = N - 1$ the degrees of freedom and Γ_p is the p-dimensional Gamma function (Anderson, 2003).

We can invert this distribution to yield a fiducial distribution $P(\boldsymbol{\Sigma}|\boldsymbol{S})$ of $\boldsymbol{\Sigma}$ conditioned on the estimator $\boldsymbol{S}$ by using Bayes' Theorem

$$\mathcal{W}(\boldsymbol{S}|\boldsymbol{\Sigma})\pi(\boldsymbol{\Sigma}) = P(\boldsymbol{\Sigma}|\boldsymbol{S})\pi(\boldsymbol{S}), \tag{10.6}$$

and adopting priors π. Since the determinant of the positive-definite covariance matrix is strictly positive, it is a scaling parameter, and we therefore assume the independence-Jeffreys prior (Jeffreys, 1961b; Sun & Berger, 2006)

$$\pi(\boldsymbol{\Sigma}) \propto |\boldsymbol{\Sigma}|^{-\frac{p+1}{2}}. \tag{10.7}$$

This is by construction invariant under reparametrizations, and can therefore be regarded as uninformative, independent of the choice of parameters. The power $(p+1)/2$ also leads to $N-1$ degrees of freedom in the inverse Wishart distribution, which is an intuitive result. Another power would only change the degrees of freedom, showing that the influence of the prior can be lessened by increasing the number of simulations N. We therefore have

$$\begin{aligned} P(\boldsymbol{\Sigma}|\boldsymbol{S}) &\propto \mathcal{W}(\boldsymbol{S}|\boldsymbol{\Sigma}/n, n)\pi(\boldsymbol{\Sigma}) \\ &\propto |\boldsymbol{\Sigma}|^{-\frac{n+p+1}{2}} \exp\left(-\frac{1}{2}n\mathrm{Tr}\left(\boldsymbol{\Sigma}^{-1}\boldsymbol{S}\right)\right) \\ &\propto \mathcal{W}^{-1}(\boldsymbol{\Sigma}|n\boldsymbol{S}, n), \end{aligned} \tag{10.8}$$

so $\boldsymbol{\Sigma}$ follows an inverse Wishart distribution

$$\mathcal{W}^{-1}(\boldsymbol{\Sigma}|\boldsymbol{C}, n) = \frac{|\boldsymbol{C}|^{\frac{n}{2}}|\boldsymbol{\Sigma}|^{-\frac{n+p+1}{2}} \exp\left(-\frac{1}{2}\mathrm{Tr}\left(\boldsymbol{\Sigma}^{-1}\boldsymbol{C}\right)\right)}{2^{\frac{np}{2}}\Gamma_p\left(\frac{n}{2}\right)}, \tag{10.9}$$

where we have abbreviated $\boldsymbol{C} = n\boldsymbol{S}$. Increasing the number of samples, n, will make this distribution more sharply peaked, meaning the true covariance $\boldsymbol{\Sigma}$ is better estimated. In general however, the inverse Wishart distribution is highly skewed for low n, and it has a biased expectation value.

Given the distribution Eq. (10.9), we can now marginalize the multivariate Gaussian likelihood over the unknown covariance

$$\begin{aligned} &\int_0^\infty G(\boldsymbol{X}|\boldsymbol{\mu}, \boldsymbol{\Sigma})\mathcal{W}^{-1}(\boldsymbol{\Sigma}|n\boldsymbol{S}, n)\mathrm{d}\boldsymbol{\Sigma} \\ &\propto |\boldsymbol{\Sigma}|^{-1/2}|\boldsymbol{\Sigma}|^{-\frac{n+p+1}{2}} \times \\ &\exp\left(-\frac{1}{2}n\mathrm{Tr}\left(\boldsymbol{\Sigma}^{-1}\boldsymbol{S}\right) - \frac{1}{2}(\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})\right)\mathrm{d}\boldsymbol{\Sigma} \\ = &\int |\boldsymbol{\Sigma}|^{-\frac{N+p+1}{2}} \exp\left(-\frac{1}{2}\mathrm{Tr}\left(\boldsymbol{\Sigma}^{-1}\boldsymbol{Q}\right)\right)\mathrm{d}\boldsymbol{\Sigma}, \end{aligned} \tag{10.10}$$

if we define $\boldsymbol{Q} = n\boldsymbol{S} + (\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})(\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})^T$. The last line is structurally the integration over an unnormalized inverted Wishart distribution $\mathcal{W}^{-1}(\boldsymbol{\Sigma}|\boldsymbol{Q}, N)$. So we know the result is the normalization constant as in Eq. (10.9). Therefore, the marginalization results in

$$\begin{aligned} &\int_0^\infty G(\boldsymbol{X}|\boldsymbol{\mu}, \boldsymbol{\Sigma})\mathcal{W}^{-1}(\boldsymbol{\Sigma}|n\boldsymbol{S}, n)\mathrm{d}\boldsymbol{\Sigma} \\ &\propto |\boldsymbol{Q}|^{-\frac{N}{2}}. \end{aligned} \tag{10.11}$$

Resubstituting $\boldsymbol{Q}$ we have

$$\begin{aligned} |\boldsymbol{Q}|^{-\frac{N}{2}} &= |n\boldsymbol{S} + (\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})(\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})^T|^{-\frac{N}{2}} \\ &= |n\boldsymbol{S}\left(1 + \frac{(\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})^T\boldsymbol{S}^{-1}(\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})}{n}\right)|^{-\frac{N}{2}} \\ &= |n\boldsymbol{S}|^{-\frac{N}{2}}\left(1 + \frac{(\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})^T\boldsymbol{S}^{-1}(\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})}{n}\right)^{-\frac{N}{2}}. \end{aligned} \tag{10.12}$$

We therefore arrive at an unnormalized multivariate t-distribution

$$\begin{aligned}
&\int_0^\infty G(\boldsymbol{X}|\boldsymbol{\mu},\boldsymbol{\Sigma})\mathcal{W}^{-1}(\boldsymbol{\Sigma}|n\boldsymbol{S},n)\mathrm{d}\boldsymbol{\Sigma} \\
&\propto \left(1+\frac{(\boldsymbol{X}_\mathrm{o}-\boldsymbol{\mu})^T\boldsymbol{S}^{-1}(\boldsymbol{X}_\mathrm{o}-\boldsymbol{\mu})}{N-1}\right)^{-\frac{N}{2}}.
\end{aligned} \tag{10.13}$$

It is however not the standard expression of a multivariate t-distribution. The standard expression assumes that there are N data vectors that determine the data mean (i.e. where the likelihood peaks) and the data covariance (i.e. the width of the likelihood). In contrast, we have assumed that there is exactly one data vector that determines where the likelihood will peak – and N simulated vectors from which we estimate the covariance.
The standard expression for the t-distribution is

$$\mathcal{T}(\boldsymbol{X},\mu,\boldsymbol{\Sigma}) = \frac{c_p|\boldsymbol{S}|^{-1/2}}{\left(1+\frac{(\boldsymbol{X}-\boldsymbol{\mu})^T\boldsymbol{S}^{-1}(\boldsymbol{X}-\boldsymbol{\mu})}{\nu}\right)^{\frac{\nu+p}{2}}}, \tag{10.14}$$

where ν are the degrees of freedom and p is the dimensionality of the data set, and the normalization constant is,

$$c_p = \frac{\Gamma\left(\frac{\nu+p}{2}\right)}{(\pi\nu)^{p/2}\Gamma\left(\frac{\nu}{2}\right)}. \tag{10.15}$$

Here, the Γ are now the usual one-dimensional Gamma functions.
We can work out the normalization of our t-distribution by transforming Eq. (10.14) onto our cosmological setup of just one observed data vector. We demand $\nu = N - p$ in order to get the same powers in the denominator. For $\nu = N - p$ Eq. (10.14) reads

$$\mathcal{T}(\boldsymbol{X},\mu,\boldsymbol{\Sigma}) = \frac{c_p'|\boldsymbol{S}|^{-1/2}}{\left(1+\frac{(\boldsymbol{X}-\boldsymbol{\mu})^T\boldsymbol{S}^{-1}(\boldsymbol{X}-\boldsymbol{\mu})}{N-p}\right)^{\frac{N}{2}}}, \tag{10.16}$$

with the normalization

$$c_p' = \frac{\Gamma\left(\frac{N}{2}\right)}{(\pi(N-p))^{p/2}\,\Gamma\left(\frac{N-p}{2}\right)}. \tag{10.17}$$

The normalization means that

$$\int \mathcal{T}(\boldsymbol{X},\mu,\boldsymbol{\Sigma})\mathrm{d}^p x = 1, \tag{10.18}$$

i.e. it is a normalization with respect to the integration over the p elements of $\boldsymbol{X}$. We can get to our parameterization Eq. (10.13) by demanding,

$$\boldsymbol{X} = \boldsymbol{X}_\mathbf{o}\sqrt{\frac{N-p}{N-1}}, \tag{10.19}$$

which implies for consistency that

$$\boldsymbol{\mu} = \boldsymbol{\mu_o}\sqrt{\frac{N-p}{N-1}}. \tag{10.20}$$

which is just a redefinition which we shall not explicitely distinguish in notation. The multivariate transformation reads,

$$\begin{aligned} &\mathcal{T}(\boldsymbol{X}, \mu, \boldsymbol{\Sigma})\mathrm{d}^p x = \mathcal{T}(\boldsymbol{X}_\mathrm{o}, \mu, \boldsymbol{\Sigma})\mathrm{d}^p x_0 \\ &\to \mathcal{T}(\boldsymbol{X}_\mathrm{o}, \mu, \boldsymbol{\Sigma}) = \mathcal{T}(\boldsymbol{X}, \mu, \boldsymbol{\Sigma}) \ \det\left(\frac{d\boldsymbol{X}}{d\boldsymbol{X}_\mathrm{o}}\right). \end{aligned} \tag{10.21}$$

The variable transformation of $\boldsymbol{X}$ to $\boldsymbol{X}_\mathrm{o}$ is a $p \times p$ matrix because all p elements of $\boldsymbol{X}$ must be derived with respect to all p elements of $\boldsymbol{X}_\mathrm{o}$. Fortunately, this is a diagonal matrix in our case

$$\frac{\mathrm{d}\boldsymbol{X}}{\mathrm{d}\boldsymbol{X}_\mathrm{o}} = \mathrm{diag}_{p\times p}\left(\sqrt{\frac{N-p}{N-1}}\right) \tag{10.22}$$

and its determinant is

$$\det\left(\frac{d\boldsymbol{X}}{d\boldsymbol{X}_\mathrm{o}}\right) = \left(\frac{N-p}{N-1}\right)^{p/2}. \tag{10.23}$$

Combining this determinant with the normalization Eq. (10.17), we therefore arrive at a normalized cosmologist's version of a t-distribution

$$\mathcal{T}(\boldsymbol{X}_\mathrm{o}, \mu, \boldsymbol{\Sigma}) = \frac{\bar{c}_p|\boldsymbol{S}|^{-1/2}}{\left(1 + \frac{(\boldsymbol{X}_\mathrm{o}-\boldsymbol{\mu})^T\boldsymbol{S}^{-1}(\boldsymbol{X}_\mathrm{o}-\boldsymbol{\mu})}{N-1}\right)^{\frac{N}{2}}}, \tag{10.24}$$

with the normalization

$$\bar{c}_p = \frac{\Gamma\left(\frac{N}{2}\right)}{(\pi(N-1))^{p/2}\,\Gamma\left(\frac{N-p}{2}\right)}. \tag{10.25}$$

The normalization depends on the dimension p and the number of simulations N such that we must demand $N > p$. If a covariance matrix must be replaced by an estimator from simulations, the likelihood Eq. (10.24) replaces the multivariate Gaussian Eq. (10.1).

10.4 Attempting to debias a Gaussian likelihood

Instead of using the t-distribution Eq. (10.24) it has become standard in cosmology to follow a procedure outlined by Hartlap et al. (2007), where the authors propose to stick with a Gaussian likelihood, and only to replace the true inverse covariance matrix by a scaled inverse sample covariance matrix

$$\boldsymbol{\Sigma}^{-1} \to \alpha\boldsymbol{S}^{-1}, \tag{10.26}$$

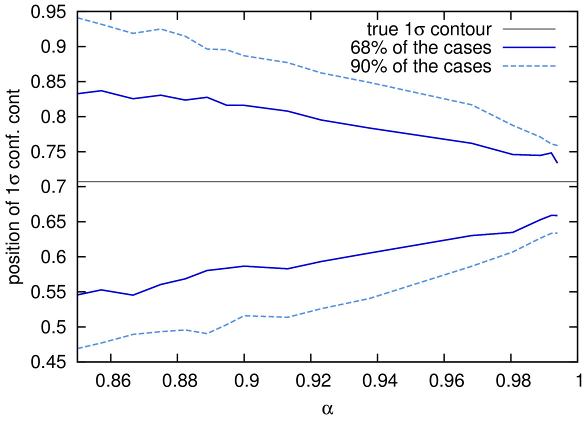

Figure 10.1: The 1σ-confidence contour of a one-dimensional normal distribution lies at $1/\sqrt{2} \approx 0.707$. However, if the covariance is estimated from simulations, its random scatter will make the estimated likelihood randomly too narrow or too broad. In 68% (90%) of the estimated covariances, the then deduced 1σ-contour falls into the area bordered by the dark blue (dashed blue) lines. The number of simulations increases with α from Eq. (10.27).

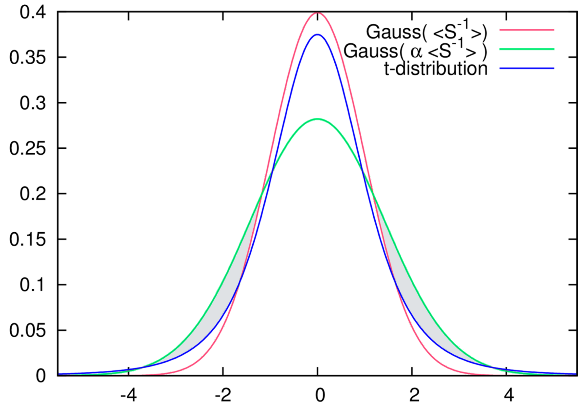

Figure 10.2: Comparison of the two Gaussian likelihoods and the t-distribution for a *particular* estimated $\boldsymbol{S}$, using $N = 5, p = 1, \alpha = 0.5$ which are examples. The grey shaded areas indicate the heavy and short wings of the Hartlap-scaled likelihood.

with

$$\alpha = \frac{N-p-2}{N-1}. \tag{10.27}$$

This is motivated by the fact that $\boldsymbol{S}^{-1}$ follows an inverse Wishart distribution, which has a biased expectation value $\langle \boldsymbol{S}^{-1} \rangle = \alpha^{-1}\boldsymbol{\Sigma}^{-1}$ as shown in Anderson (2003). Here, the angular brackets denote averaging over the inverse Wishart distribution.

Hartlap et al. (2007) argue that this debiased inverse covariance matrix will remove all biases from parameter inference. However, the situation is more complex. In a Bayesian analysis one would not necessarily define an estimator $\hat{\boldsymbol{\theta}}$, but if one does, the bias is $b_\theta = \langle \hat{\boldsymbol{\theta}} \rangle - \boldsymbol{\theta}$ where the angular brackets now denote the average over the likelihood of the parameters. Adopting the wrong sampling distribution will yield incorrect posterior distributions, with biased parameter estimates (should they be made) and incorrect errors, even if the inverse covariance matrix itself has been debiased.

We compare univariate examples of the likelihoods and the modified t-distribution Eq. (10.24) in Fig. 10.2: the Hartlap-scaled and the unscaled Gaussian only differ in width, whereas the t-distribution has a more sharply peaked central region but broader extreme wings than a Gaussian, allowing for more scatter away from the peak.

Additionally, the scaling in Eq. (10.26) implies a sharp mapping between the estimator $\boldsymbol{S}^{-1}$ and $\boldsymbol{\Sigma}^{-1}$, which does not account for the randomness of $\boldsymbol{S}^{-1}$, due to the finite width of the inverse Wishart distribution. Therefore, $\alpha\boldsymbol{S}^{-1}$ applied to a *single* given $\boldsymbol{S}^{-1}$ should not be interpreted as a reliable ‘debiasing’ but rather a scaling that widens up the Gaussian likelihood Eq. (10.1) in an essentially random way. This randomness will propagate through the parameter inference and introduce a scatter of the likelihood contours of which we show a simple example in Fig. 10.1. This scatter can only be reduced by estimating the inverse covariance matrix more precisely.

10.5 Comparison of the distributions

We assess whether the procedures Eq. (10.24) and Eq. (10.26) introduce errors into the parameter inference by a Frequentist experiment based on the cumulative probability $P(< x)$ of a random variable x. By simulating the estimation of a data mean, given many different estimated $\boldsymbol{S}^{-1}$ and realizations of the data, we can generate the true cumulative distribution. We can then compare with the cumulative distributions that assume either the Hartlap-scaled Gaussian or the modified t-distribution to this truth. In Fig. 10.3, we find that the t-distribution correctly reproduces the actual cumulative distribution – the line is straight with a slope of unity. The Hartlap-scaled likelihood does not capture the scatter around the peak correctly, which will lead to a misestimate of the parameter errors, even on average. As expected, the discrepancy decreases as more simulations are included in the estimation of $\boldsymbol{S}$ (i.e., as $\alpha \to 1$).

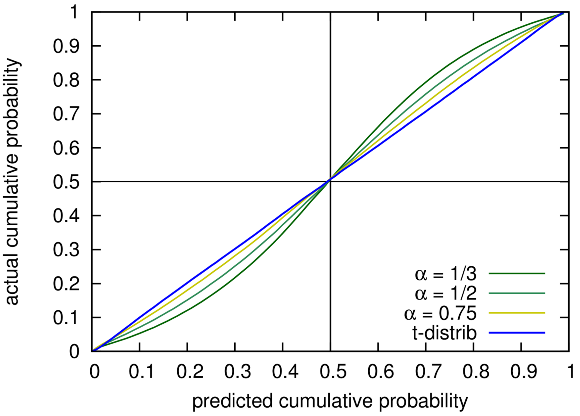

Figure 10.3: Predicted cumulative probability versus the actual cumulative probability. The latter was calculated by evaluating the Hartlap-scaled Gaussian likelihood and the t-distribution for 150 estimated covariance matrices and $10,000$ data sets. The t-distribution follows the diagonal line of unit slope, meaning it predicts correctly the shape of the likelihood, whereas the Hartlap-scaled likelihood is typically too broad.

10.6 Assessment of confidence in higher dimensions

The issue at hand can be studied in higher dimensions by investigating the distribution of the following quantities:

$$\chi^2 = (\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})^T\, \boldsymbol{\Sigma}^{-1}\, (\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})\,, \tag{10.28}$$

which is the true χ^2; the same quantity but with the estimated $\boldsymbol{S}$ replacing $\boldsymbol{\Sigma}$,

$$T^2 = (\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})^T\, \boldsymbol{S}^{-1}\, (\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})\,; \tag{10.29}$$

and the Hartlap-scaled version

$$H^2 = (\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})^T\, \alpha\boldsymbol{S}^{-1}\, (\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})\,. \tag{10.30}$$

By construction, we have $\langle H^2\rangle = \langle\chi^2\rangle$, meaning the Hartlap-scaling does indeed debias the expectation value. It does however underestimate statistical scatter, as we shall show in the following.

χ^2 follows the χ^2_p-distribution, which only arises if the covariance is precisely known and indeed the correct covariance of $\boldsymbol{X}_\mathrm{o}$. The quantity T^2 will not follow the χ^2_p-distribution, because it contains not only a random vector $\boldsymbol{X}_\mathrm{o} \sim \mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, but additionally the random estimate of the covariance matrix that follows the Wishart distribution $\mathcal{W}(\boldsymbol{\Sigma}/n, n)$. T^2 therefore follows

$$\frac{T^2(n-p+1)}{pn} \sim F_{p,n-p+1}, \tag{10.31}$$

where $n = N - 1$, and the $F_{p,n-p+1}$ is the F-distribution of p and $n-p+1$ degrees of freedom (Anderson, 2003). Consequently, a change of variables shows that,

$$T^2 \sim \frac{\Gamma\left(\frac{n+1}{2}\right)}{\Gamma(p/2)\Gamma(r/2)}\, \frac{n^{-p/2}(T^2)^{p/2-1}}{(T^2/n+1)^{\frac{n+1}{2}}}. \tag{10.32}$$

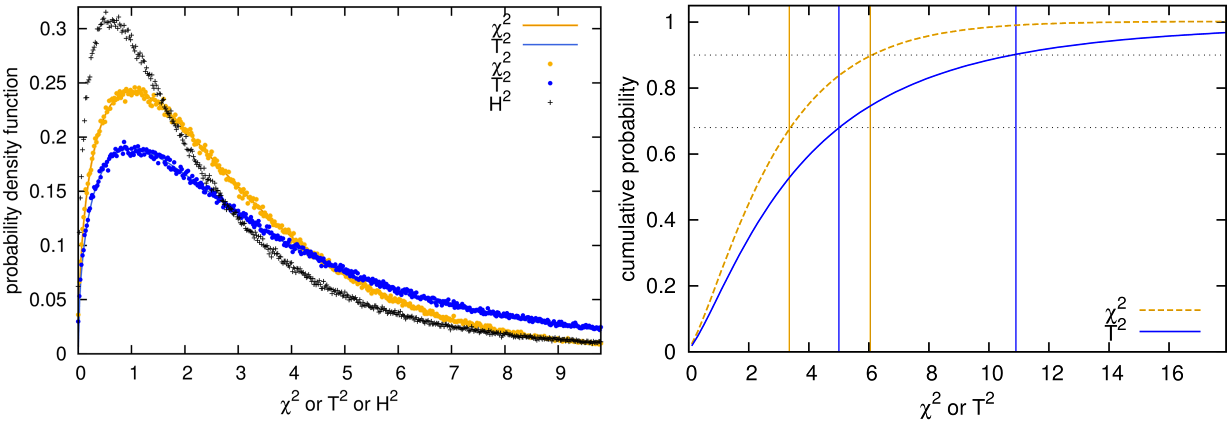

Figure 10.4: **Left:** The distribution of different interpretations for $(\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})^T \boldsymbol{S}^{-1} (\boldsymbol{X}_\mathrm{o} - \boldsymbol{\mu})$, using $p = 3$, $N = 10$. Dots represent simulations, solid lines are the analytical χ_p^2- and T^2-distribution. For $N \gg p$, the T^2-distribution approximates the χ_p^2-distribution. The closer N is to p, the more differs the T^2-distribution from the χ_p^2-distribution, being broader than χ_p^2, leading to a cumulative distribution that rises more slowly. The Hartlap-scaled H^2 follows the black distribution which is more sharply peaked than the χ_p^2, although the χ_p^2-distribution is the minimal scatter that one can achieve; this means the Hartlap-scaled H^2 underestimates the joint scatter of $\boldsymbol{X}_\mathrm{o}$ and $\boldsymbol{S}^{-1}$. **Right:** The cumulative distributions of χ^2 and T^2 from the left. The vertical lines mark the 68% and 90% confidence limits.

instead of $T^2 \sim \chi_p^2$. Only for $N \to \infty$ will the Wishart distribution tend towards a delta function, and the distribution of T^2 will then tend towards a χ_p^2-distribution.

The distribution of the Hartlap-scaled H^2 is more sharply peaked than that of χ^2, thereby suggesting that the experiment has less statistical scatter than the χ_p^2 distribution on average. This is impossible since the χ_p^2 distribution is subject to scatter of the random vector $\boldsymbol{X}_\mathrm{o}$ only.

The cumulative probabilities $P_\mathrm{c}(\chi_\mathrm{c}^2)$ or $P_\mathrm{c}(T_\mathrm{c}^2)$ give our confidence that the mean $\boldsymbol{\mu}$ of the multivariate vector $\boldsymbol{X}_\mathrm{o}$ is enclosed within an ellipsoid bounded by χ_c^2 or T_c^2. The more slowly rising cumulative distribution function of T^2 therefore shows that we need $T^2 > \chi^2$ in order to achieve the same confidence that the mean is captured within the confidence contours. In parameter space, this will lead to an increase of the Bayesian confidence intervals.

10.7 Reweighting an MCMC chain that sampled from a Gaussian likelihood.

We have shown above that T^2, χ^2 and H^2 follow different distributions, which will affect parameter inference. Often, the error of confusing a T^2 with a χ^2 or H^2 can retrospectively be undone with very little numerical effort by reweighting an existing MCMC chain.

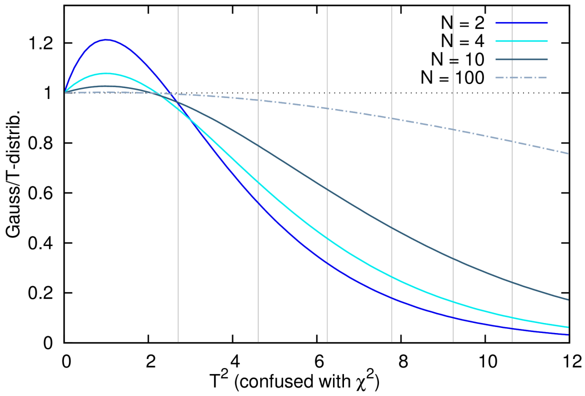

Figure 10.5: Unnormalized weights $G(\boldsymbol{X}_{\rm o}, \boldsymbol{\mu}, \boldsymbol{S}^{-1})/P(\boldsymbol{X}_{\rm o}, \boldsymbol{\mu}, \boldsymbol{S}^{-1}, n)$ for mapping between a Gaussian likelihood and a t-distribution. The normalization depends on the dimensionality of the data set, and leads to an offset along the y-axis, that is however independent of theoretical parameters. The number of simulations in the covariance matrix is N. The vertical lines depict the χ^2 values (2.71,4.61,6.25,7.78,9.24,10.64) that enclose 90% confidence for a multivariate Gaussian.

In Fig 10.5 we plot weights for reweighting a chain that sampled from $\exp(-\chi^2/2)$. If a Hartlap-scaling has been applied, it would additionally need to be removed.
We note that the maximum of the t-distribution in the full parameter space coincides with the maximum of χ^2 (and also of H^2), but once any parameters are marginalised over, the resulting parameter posteriors will not in general peak in the same place.

Chapter 11

Conclusions

As cosmologists, our principal goal is to determine the physical model that describes our Universe, and the values that the parameters of this model take. Implications of this goal are a theoretical satisfaction with the found solution, and a strong support by cosmological data. Ultimately, one could hope to establish a cosmological standard model, that is theoretically sound and firmly rooted in fundamental physics, and which is additionally indisputably preferred over all contender models by the cosmic data. These goals and hopes shall provide the criteria by which we draw the conclusions from this thesis.

11.1 On ΛCDM

Let us first answer whether a satisfactory cosmological model, as described above, has already been established. After all, the current standard model, ΛCDM, fits all currently available data on cosmic scales, and its parameters are well constrained. Further, ΛCDM is derived from the consistent framework of General Relativity, which is experimentally well supported on sub-galactic scales. Therefore, if ΛCDM did not require dark matter and the cosmological constant, it would be a theoretically satisfactory model. However, with dark matter escaping direct detections so far, and with the cosmological constant being observationally indistinguishable from a vacuum energy that is expected to be order of magnitudes larger, doubt-raising questions of fine-tuning and unnaturalness arise as has been argued in Sect. 2.4.

These doubts about ΛCDM could be defused (although not resolved), by relegating the problematic aspects of dark matter and dark energy to other areas of physics. The problem of vacuum energy could be passed to the domain of quantum field theories, stating that although it *manifests* itself in cosmological observations, it actually *originates* from another branch of physics and is therefore not a shortcoming of ΛCDM itself.

Concerning dark matter, primordial nucleosynthesis rules out that dark matter consists of standard model particles, as argued in chapter 2. One could therefore promote the dark matter problem to a postulation of a dark particle and thereby re-attribute it to particle physics. Of course, the inverse is also true: the null-results of currently executed direct dark matter searches could be explained by claiming there exists a fundamental problem with gravity on galactic and super-galactic scales instead.

Hence, as long as the dark components of ΛCDM remain unexplained, ΛCDM does not yet satisfy all criteria demanded of a convincing cosmological model, and potential alternatives have to be discussed.

11.2 On deviations from ΛCDM

Having argued why ΛCDM does not yet fulfil all requirements demanded of a satisfactory cosmologcial model, let us next answer whether there is any serious reason to doubt ΛCDM. Disproving ΛCDM could be achieved by proving one of its assumptions wrong. Large-scale isotropy is however observationally so well established, that it cannot be seriously contested. Homogeneity is not observable, as our observations are restricted to the backwards lightcone, but is demanded by the paradigm that our Earth's position in the Universe is not preferred. We have therefore good reason to believe that the background Universe can be well described with a FLRW metric, and structure growth with perturbations of FLRW.

Remaining reasons for doubt are then the question of whether General Relativity describes gravity on cosmic scales correctly and whether the ingredients to the Universe's energy-momentum tensor are accurately known. At the time of submitting this thesis, a recent recapitulatory study of these doubts is given in Planck Collaboration et al. (2015), where models of modified gravity and dark energy are confronted with the latest data. Planck Collaboration et al. (2015) studies constraints on a variable equation of state w for dark energy. For scalar field types of dark energy, Planck Collaboration et al. (2015) puts constraints on the field's potential via a parameterization similar to the Taylor expansion from Sect. 2.2. Furthermore, Planck Collaboration et al. (2015) studies observational consequences of modified gravity via the gravitational slip

$$\eta = -\frac{\Phi}{\Psi}, \tag{11.1}$$

and potential modifications $\mu(a, \boldsymbol{k})$ of the Poisson equation

$$-k^2\Psi = 4\pi G a^2 \mu(a, \boldsymbol{k}) \rho \Delta \tag{11.2}$$

where deviations of η and μ from unity indicate deviations from General Relativity.

In Planck Collaboration et al. (2015), no significant deviations of w, η or μ from their ΛCDM values are found. However, in chapter 6, we have argued that a parameter can be considered measured, if its standard deviation is much smaller than the parameter's allowed prior range. As long as a cosmological model does not predict the prior ranges of w, η and μ, it is therefore unclear whether the data used in Planck Collaboration et al. (2015) have already acquired the needed precision to measure w, η and μ.

This problem of lacking theoretically motivated prior ranges cannot be resolved by demanding ever more accurate cosmological data. However, the inverse is true: the better cosmic data and their analyses are, the more likely is a serendipitous detection of potential deviations from ΛCDM.

We therefore conclude that given current data, the Universe is observationally compatible with the phenomenology of ΛCDM, although this does not necessarily mean that ΛCDM describes our Universe qualitatively correctly. We will next outline which steps need to be undertaken to arrive at a robust model of the Universe.

11.3 Towards a robust cosmological model

If ΛCDM succeeds to also fit the future cosmic data, and its inherent problems could be eliminated by a direct detection of dark matter and a resolution of the cosmological constant problem, then ΛCDM would almost certainly advance to the incontestable model of cosmology.

While it is still thinkable that ΛCDM may not be the correct explanation of the Universe, further theoretical progress on alternative cosmological models is needed, as well as an improvement of cosmological data and their analysis.

This thesis has contributed to the latter, and apart from having detected and resolved a parameterization bias concerning neutrinos in the cosmic microwave background, two omnipresent statistical tools of cosmology were examined critically: the Fisher matrix and the scaling of estimated covariance matrices as proposed by Hartlap et al. (2007).

Concerning the Fisher matrix, we found that it is often used too optimistically, trusting in the Cramer-Rao inequality, and trusting in the usually untested assumption of the underlying posterior being well approximated by a multivariate Gaussian distribution. This thesis reassesses the significance of the Cramer-Rao inequality for cosmology and furthermore contains a Derivative Approximation of LIkelihoods ('DALI') which extends the Fisher matrix approximation by including non-Gaussianities. DALI is not restricted to cosmology, but can also be applied in any other scientific context where a parameterized model is compared to data, as e.g. in particle physics. As DALI approximates likelihoods, it can be used for speedy forecasts (see Sect. 9), but it can also guide and speed up a Hamiltonian Monte Carlo sampler. DALI can thereby also be used for the analysis of real data (see Sect. 6.10 and Sect. 9.3).

Concerning the parameter inference with estimated covariance matrices, we disproved the claim of Hartlap et al. (2007) that their method is unbiased, and replaced it by a truly unbiased method. A wide variety of current and future surveys is directly affected by this: weak lensing surveys like CFHTLenS (Heymans et al., 2013; Kilbinger et al., 2013), KiDS (de Jong et al., 2013) and DES (Chang et al., 2015) and galaxy clustering surveys like BOSS (Gil-Marín et al., 2015) have so far all used the method of Hartlap et al. (2007) and therefore all suffer from the mentioned bias. Future analyzes should employ the unbiased inference from chapter 10. As the differences between the Hartlap-scaled Gaussian likelihood and the t-distribution here presented are especially important in the wings of the likelihoods, a strong impact on evidence calculations and model selection is expected, which shall be investigated in future works as well.

Bibliography

Ade P., et al., 2014, A&A, 571, A16, 1303.5076

Albrecht A., Bernstein G., Cahn R., Freedman W. L., Hewitt J., Hu W., Huth J., Kamionkowski M., Kolb E. W., Knox L., Mather J. C., Staggs S., Suntzeff N. B., 2006, ArXiv Astrophysics e-prints, astro-ph/0609591 , ADS

Allison R., Dunkley J., 2014, MNRAS, 437, 3918, 1308.2675 , ADS

Amanullah et al., 2010, Ap.J., 716, 712, 1004.1711 , ADS

Amendola L., 2002, ApJ, 569, 595, astro-ph/0107527 , ADS

Amendola L., Sellentin E., 2016, ArXiv e-prints, 1602.01746 , ADS

Amendola L., Tsujikawa S., 2010a, Dark Energy - Theory and Observations. Cambridge University Press.

Amendola L., Tsujikawa S., 2010b, Dark Energy: Theory and Observations, ADS

Amendola L., Tsujikawa S., 2010c, Dark Energy: Theory and Observations. Cambridge University Press, ADS

Anderson T. W., 2003, An Introduction to Multivariate Statistical Analysis. Wiley Series in Probability and Statistics, 3rd ed.

Andrae R., Schulze-Hartung T., Melchior P., 2010, ArXiv e-prints, 1012.3754 , ADS

Archidiacono M., Calabrese E., Melchiorri A., 2011, Phys. Rev., D84, 123008, 1109.2767

Archidiacono M., Hannestad S., 2014, JCAP, 1407, 046, 1311.3873

Audren B., Bellini E., Cuesta A. J., Gontcho S. G. A., Lesgourgues J., Niro V., Pellejero-Ibanez M., Pérez-Ràfols I., Poulin V., Tram T., Tramonte D., Verde L., 2014, ArXiv e-prints, 1412.5948 , ADS

Audren B., Lesgourgues J., Benabed K., Prunet S., 2013a, JCAP, 2, 1, 1210.7183 , ADS

Audren B., Lesgourgues J., Benabed K., Prunet S., 2013b, JCAP, 2, 1, 1210.7183 , ADS

Ayaita Y., Weber M., Wetterich C., 2013, Phys. Rev. D, 87, 043519, 1211.6589 , ADS

Balbi A., Ade P., Bock J., Borrill J., Boscaleri A., De Bernardis P., Ferreira P. G., Hanany S., Hristov V., Jaffe A. H., Lee A. T., Oh S., Pascale E., Rabii B., Richards P. L., Smoot G. F., Stompor R., Winant C. D., Wu J. H. P., 2000, Astrophys. J. L., 545, L1

Bardeen J. M., Bond J. R., Kaiser N., Szalay A. S., 1986, Astrophysical Journal, 304, 15, ADS

Bartelmann M., 2013, Theoretical Astrophysics. WILEY-VCH.

Bartelmann M., Schneider P., 2001, Physics Reports, 340, 291, ADS

Basboll A., Bjaelde O. E., Hannestad S., Raffelt G. G., 2009, Phys.Rev., D79, 043512, 0806.1735

Bernabei R., Belli P., Cappella F., Caracciolo V., Cerulli R., Dai C. J., d'Angelo A., d'Angelo S., Di Marco A., He H. L., Incicchitti A., Ma X. H., Montecchia F., Prosperi D., Sheng X. D., Wang R. G., Ye Z. P., 2014, Nuclear Instruments and Methods in Physics Research A, 742, 177, 1403.1404 , ADS

Blas D., Lesgourgues J., Tram T., 2011, JCAP, 1107, 034, 1104.2933

Bonvin C., Clarkson C., Durrer R., Maartens R., Umeh O., 2015, JCAP, 7, 040, ADS

Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., Lemson G., 2009, MNRAS, 398, 1150, 0903.3041 , ADS

Brown A., Xenon Collaboration 2014, in American Institute of Physics Conference Series Vol. 1604 of American Institute of Physics Conference Series, XENON dark matter searches: Results and the future. pp 313–318, ADS

Budavári T., Connolly A. J., Szalay A. S., Szapudi I., Csabai I., Scranton R., Bahcall N. A., Brinkmann J., Eisenstein D. J., Frieman J. A., Fukugita M., Gunn J. E., Johnston D., Kent S., Loveday J. N., 2003, ApJ, 595, 59, astro-ph/0305603

Castro P. G., Heavens A. F., Kitching T. D., 2005, Phys. Rev. D, 72, ADS

Chang C., Vikram V., Jain B., Bacon D., Amara A., Becker M. R., Bernstein G., Bonnett C., Bridle S., Brout D., Busha M., Frieman J., Gaztanaga E., Hartley W., Jarvis M., Kacprzak T., et al. 2015, Physical Review Letters, 115, 051301, 1505.01871 , ADS

Chantavat T., Sawangwit U., Sutter P. M., Wandelt B. D., 2014, 1409.3364 , ADS

Charles J., Deschamps O., Descotes-Genon S., Lacker H., Menzel A., Monteil S., Niess V., Ocariz J., Orloff J., Perez A., Qian W., Tisserand V., Trabelsi K., Urquijo P., Vale Silva L., CKMfitter Group 2015, PRD, 91, 073007, 1501.05013 , ADS

Chevallier M., Polarski D., 2001, International Journal of Modern Physics D, 10, 213, ADS

Chevallier M., Polarski D., 2001, Int.J.Mod.Phys., D10, 213, gr-qc/0009008

Cramer H., 1946, Mathematical Methods of Statistics. Princeton University Press.

Dai X., Chartas G., Agol E., Bautz M. W., Garmire G. P., 2003, ApJ, 589, 100, astro-ph/0301592 , ADS

de Bernardis P., Ade P. A. R., Bock J. J., Bond J. R., Borrill J., Boscaleri A., Coble K., Crill B. P., De Gasperis G., Farese P. C., Ferreira P. G., Ganga K., Giacometti M., Hivon E., Hristov V. V., Iacoangeli A., Jaffe A. H., 2000, NATURE, 404, 955, astro-ph/0004404

de Jong J. T. A., Kuijken K., Applegate D., Begeman K., Belikov A., Blake C., Bout J., Boxhoorn D., et al. 2013, The Messenger, 154, 44

Dodelson S., Schneider M. D., 2013, Phys. Rev. D, 88, 063537, 1304.2593

Durrer R., 2008, The Cosmic Microwave Background. Cambridge University Press, ADS

Faham C., for the LUX Collaboration 2014, ArXiv e-prints, 1405.5906 , ADS

Falco M., Hansen S. H., Wojtak R., Mamon G. A., 2013, MNRAS, 431, L6, 1210.3363 , ADS

Fergusson J. R., Shellard E. P. S., 2007, Phys. Rev. D, 76, ADS

Forero D., Tortola M., Valle J., 2014, Phys.Rev., D90, 093006, 1405.7540

Gerbino M., Di Valentino E., Said N., 2013, Phys. Rev., D88, 063538, 1304.7400

Gil-Marín H., Percival W. J., Brownstein J. R., Chuang C.-H., Grieb J. N., Ho S., Kitaura F.-S., Maraston C., Prada F., Rodríguez-Torres S., Ross A. J., Samushia L., Schlegel D. J., Thomas D., Tinker J. L., Zhao G.-B., 2015, ArXiv e-prints, 1509.06386 , ADS

Gregory P., 2005, Bayesian Logical Data Analysis for the Physical Sciences. Cambridge University Press

Hajian A., 2007a, PRD, 75, 083525, astro-ph/0608679 , ADS

Hajian A., 2007b, Phys. Rev. D, 75, 083525, astro-ph/0608679 , ADS

Hamimeche S., Lewis A., 2009, Phys. Rev. D, 79, 083012, 0902.0674

Hamuy M., Phillips M. M., Suntzeff N. B., Schommer R. A., Maza J., Aviles R., 1996, AJ, 112, 2391, astro-ph/9609059 , ADS

Hannestad S., Tu H., Wong Y. Y., 2006, JCAP, 6, 25, astro-ph/0603019 , ADS

Hartlap J., Simon P., Schneider P., 2007, A,A, 464, 399, astro-ph/0608064

Heavens A., 2003, MNRAS, 343, 1327, astro-ph/0304151 , ADS

Heavens A. F., Kitching T. D., Taylor A. N., 2006, Phys. Rev. D, 373, 105, ADS

Heymans C., Grocutt E., Heavens A., Kilbinger M., Kitching T. D., Simpson F., Benjamin J., Erben T., Hildebrandt 2013, MNRAS, 432, 2433, 1303.1808

Heymans C., Grocutt E., Heavens A., Kilbinger M., Kitching T. D., Simpson F., Benjamin J., Erben T., Hildebrandt H., Hoekstra H., Mellier Y., Miller L., Van Waerbeke L., Brown M. L., 2013, MNRAS, 432, 2433, 1303.1808

Hu W., 1998, ApJ, 506, 485, astro-ph/9801234

Hu W., 1999, ApJL, 522, L21, ADS

Hu W., 2002, Phys. Rev. D, 66, 083515, ADS

Hu W., Eisenstein D. J., Tegmark M., White M. J., 1999, Phys. Rev., D59, 023512, astro-ph/9806362

Hu W., Sawicki I., 2007, Phys. Rev. D, 76, 064004, 0705.1158 , ADS

Jeffreys H., 1961a, The Theory of Probability. Oxford University Press

Jeffreys H., 1961b, Theory of Probability (OUP)

Joachimi B., Taylor A., 2011, Mon.Not.Roy.Astron.Soc., 416, 1010, 1103.3370

Kass R., Raftery A., 1995, Journal of the Americal Statistical Association, 90, 773

Kilbinger M., 2015, Reports on Progress in Physics, 78, 086901, 1411.0115 , ADS

Kilbinger M., Fu L., Heymans C., Simpson F., Benjamin J., Erben 2013, MNRAS, 430, 2200, 1212.3338

Kilbinger M., Wraith D., Robert C. P., Benabed K., Cappé O., Cardoso J.-F., Fort G., Prunet S., Bouchet F. R., 2010, MNRAS, 405, 2381, 0912.1614 , ADS

Kunz M., 2009, Phys. Rev. D, 80, 123001, astro-ph/0702615 , ADS

Laureijs R., Amiaux J., Arduini S., Augueres J. ., Brinchmann J., Cole R., Cropper M., Dabin C., Duvet L., Ealet A., et al. 2011a, ArXiv e-prints, 1110.3193 , ADS

Laureijs R., Amiaux J., Arduini S., Auguères J. ., Brinchmann J., Cole R., Cropper M., Dabin C., Duvet L., Ealet A., et al. 2011b, ArXiv e-prints, 1110.3193 , ADS

Lepage G. P., 1978, Journal of Computational Physics, 27, 192

Lepage G. P., 1980, Cornell preprint CLNS

Lewis A., Challinor A., , 2011, CAMB: Code for Anisotropies in the Microwave Background, Astrophysics Source Code Library, 1102.026 , ADS

Linder E. V., 2003, Phys. Rev. Lett., 90, 091301, astro-ph/0208512

Longair M., 2008, Galaxy Formation. Springer, secod edition

Lyons L., 2014, ArXiv e-prints, 1409.1903 , ADS

Ma C.-P., Bertschinger E., 1995, ApJ, 455, 7, astro-ph/9506072 , ADS

Maoz D., 2007, Astrophysics in a Nutshell. Princeton University Press.

Mardia K. V., Kent J. T., Bibby J. M., 1979, Multivariate analysis. Probability and Mathematical Statistics, London: Academic Press

Martin J., 2012, Comptes Rendus Physique, 13, 566, 1205.3365 , ADS

Melchior P., Suchyta E., Huff E., Hirsch M., Kacprzak T., Rykoff E., Gruen D., Armstrong R., et al. 2015, MNRAS, 449, 2219, 1405.4285 , ADS

Metropolis N., 1985, in Alcouffe R., Dautray R., Forster A., Ledonois G., Mercier B., eds, Lecture Notes in Physics, Berlin Springer Verlag Vol. 240 of Lecture Notes in Physics, Berlin Springer Verlag, Monte-Carlo: In the Beginning and Some Great Expectations. p. 62, ADS

Mukhanov V. F., Feldman H. A., Brandenberger R. H., 1992, Pysics reports., 215, 203, ADS

Munshi D., Smidt J., Heavens A., Coles P., Cooray A., 2010, ArXiv e-prints, 1003.5003 , ADS

Narayan R., Bartelmann M., 1996, ArXiv Astrophysics e-prints, astro-ph/9606001 , ADS

Oldengott I. M., Rampf C., Wong Y. Y. Y., 2015, JCAP, 4, 16, 1409.1577 , ADS

Perlick V., 2000, Ray Optics, Fermat's Principle, and Applications to General Relativity. Springer-Verlag, ADS

Perlmutter S., et al., 1999, ApJ, 517, 565, astro-ph/9812133

Phillips M. M., 1993, APJL, 413, L105, ADS

Pillepich A., Porciani C., Reiprich T. H., 2012, MNRAS, 422, 44, 1111.6587 , ADS

Planck Collaboration Ade P. A. R., Aghanim N., Armitage-Caplan C., Arnaud M., Ashdown M., Atrio-Barandela F., Aumont J., Baccigalupi C., Banday A. J., et al. 2014, A.&A., 571, A22, 1303.5082

Planck Collaboration Ade P. A. R., Aghanim N., Arnaud M., Arroja F., Ashdown M., Aumont J., Baccigalupi C., Ballardini M., Banday A. J., et al. 2015, ArXiv e-prints, 1502.02114 , ADS

Planck Collaboration Ade P. A. R., Aghanim N., Arnaud M., Ashdown M., Aumont J., Baccigalupi C., Banday A. J., Barreiro R. B., Bartlett J. G., et al. 2015, ArXiv e-prints, 1502.01589 , ADS

Planck Collaboration Ade P. A. R., Aghanim N., Arnaud M., Ashdown M., Aumont J., Baccigalupi C., Banday A. J., Barreiro R. B., Bartolo N., et al. 2015, ArXiv e-prints, 1502.01590 , ADS

Press W. H., Farrar G. R., 1990, Computers in Physics, pp 190–195

Refregier A., the DUNE collaboration 2008, ArXiv 0802.2522, 802, 0802.2522 , ADS

Riess A. G., et al., 1998, AJ, 116, 1009, astro-ph/9805201

Riess A. G., Macri L., Casertano S., Lampeitl H., Ferguson H. C., Filippenko A. V., Jha S. W., Li W., Chornock R., 2011, ApJ, 730, 119, 1103.2976 , ADS

Schaye J., Crain R. A., Bower R. G., Furlong M., Schaller M., Theuns T., Dalla Vecchia C., Frenk C. S., McCarthy I. G., Helly J. C., Jenkins A. e. a., 2015, MNRAS, 446, 521, 1407.7040 , ADS

Sellentin E., 2015, MNRAS, 453, 893, 1506.04866 , ADS

Sellentin E., Durrer R., 2015, Phys. Rev. D, 92, 063012, 1412.6427 , ADS

Sellentin E., Heavens A. F., 2016, MNRAS, 456, L132, 1511.05969 , ADS

Sellentin E., Quartin M., Amendola L., 2014, MNRAS, 441, 1831, 1401.6892 , ADS

Sellentin E., Schäfer B. M., 2016, MNRAS, 456, 1645, 1506.05356 , ADS

Sellke T., Bayarri M., Berger J. O., 2001, American Statistician, 55, 62

Semboloni E., Mellier Y., van Waerbeke L., Hoekstra H., Tereno I., Benabed K., Gwyn S. D. J., Fu L., Hudson M. J., Maoli R., Parker L. C., 2006, A&A, 452, 51, astro-ph/0511090

Sherman J., Morrison W. J., 1950, Ann. Math. Statist., 21, 124

Skilling J., 2004, in Fischer R., Preuss R., Toussaint U. V., eds, American Institute of Physics Conference Series Vol. 735 of American Institute of Physics Conference Series, Nested Sampling. pp 395–405, ADS

Smith R. E., Peacock J. A., Jenkins A., White S. D. M., Frenk C. S., Pearce F. R., Thomas P. A., Efstathiou G., Couchman H. M. P., 2003, MNRAS, 341, 1311, ADS

Springel V., , 2000, GADGET-2: A Code for Cosmological Simulations of Structure Formation, Astrophysics Source Code Library, 0003.001 , ADS

Steigman G., 2006, International Journal of Modern Physics E, 15, 1, astro-ph/0511534 , ADS

Steigman G., 2012, Adv.High Energy Phys., 2012, 268321, 1208.0032

Straumann N., 2013, General Relativity, 2nd Ed., Springer (2013)

Sun D., Berger J., 2006, Proc. Valencia/ISBA 8th World Meeting on Bayesian Statistics

Takahashi K., Brown M., Burigana C., Jackson C., Jarvis M., Kitching K. D. T. D., Kneib J. P., Masamune Oguri M., Prunet S., Shan H., Starck J. L., Yamauchi D., 2015, Advancing Astrophysics with the Square Kilometre Array (AASKA14), p. 159, 1501.03859 , ADS

Taylor A., Joachimi B., Kitching T., 2013, MNRAS, 432, 1928, 1212.4359

Tegmark M., Taylor A., Heavens A., 1997, Astrophys.J., 480, 22, astro-ph/9603021

Trotta R., 2008, Contemporary Physics, 49, 71, 0803.4089 , ADS

Trotta R., Melchiorri A., 2005, Phys.Rev.Lett., 95, 011305, astro-ph/0412066

Voigt H.-H., 2012, Abriss der Astronomie. WILEY-VCH. Edited by H. J. Röser and W. Tscharnuter.

Woodbury M., 1950, Statistical Research Group, Memo Rep. No. 42

Wraith D., Kilbinger M., Benabed K., Cappé O., Cardoso J.-F., Fort G., Prunet S., Robert C. P., 2009, Phys. Rev. D, 80, 023507, 0903.0837 , ADS

Young S., Byrnes C. T., 2015, JCAP, 4, 034, 1503.01505 , ADS