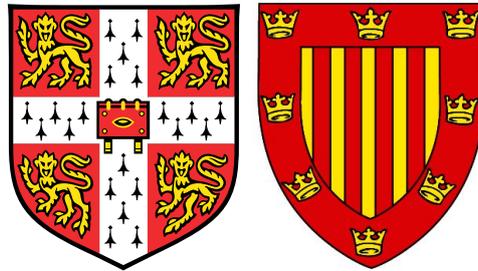


Constraining the kinetically dominated Universe

Bayesian methods and primordial cosmology



Lukas Tobias Hergt

Cavendish Astrophysics Group
University of Cambridge
Peterhouse

This dissertation is submitted for the degree of
Doctor of Philosophy

December 2020

*dedicated to the loving memory of my father,
Simon Erdmann Hergt*

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface and specified in the text.

It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the preface and specified in the text.

It does not exceed the prescribed word limit for the Physics Degree Committee.

Lukas Tobias Hergt
December 2020

Summary

Constraining the kinetically dominated Universe

Bayesian methods and primordial cosmology

Lukas Tobias Hergt

In this thesis a systematic approach is taken to the study of the generation of primordial perturbations from a finite amount of cosmic inflation. It is demonstrated how a pre-inflationary stage that is kinetically dominated causes features in the primordial power spectrum. These features can be associated with cosmological parameters, which are then constrained using Bayesian methods and CMB data from the Planck satellite.

Chapter 1 places this thesis into the context of the current state of cosmology and introduces the needed theoretical background and notation. The remainder of the thesis presents both my published and ongoing work.

Chapter 2 starts with an overview of Bayesian methods which are then applied to the question of how a uniform versus logarithmic prior distribution affect parameter estimations and model comparisons. Two examples of extensions to the current standard model of cosmology are investigated in detail, the tensor-to-scalar ratio of primordial perturbations and three massive neutrinos.

Chapter 3 was published in Hergt et al. (2019) [[I & chapter 3](#)] and makes a case for a pre-inflationary phase of kinetic dominance. To that end the phase-space trajectories for three representative inflationary potentials are inspected. Comparing different priors on the initial conditions of the trajectories, a significant fraction of trajectories are found to start out in a kinetically dominated stage.

Chapter 4 was published in Hergt et al. (2019) [[II & chapter 4](#)]. Assuming a spatially flat cosmology, it contrasts the effects from slow-roll inflation to those from a finite amount of inflation with an earlier kinetically dominated stage. CMB data from the Planck satellite are used to compare these two models and constrain the parameters pertaining to the primordial Universe. There is no clear preference for either model. Any Occam penalty from the higher complexity of a model with a kinetically dominated stage is balanced out by a better fit to the large-scale lack of power in the CMB.

Chapter 5 extends the previous analysis, allowing for a variable amount of spatial curvature. Even a small amount of present-day curvature significantly limits the total amount of inflation, rendering the effects of kinetic dominance more prominent in the observable window of the CMB. A model comparison of various single-field inflation models leads to similar results to those of a flat cosmology, driven mostly by constraints from the tensor-to-scalar ratio and from reheating.

I conclude with potential directions of future research.

Table of contents

Declaration	v
Summary	vii
Acknowledgements	xiii
Conventions	xv
1 Introduction	1
1.1 The evolution of our Universe	1
1.2 The Cosmic Microwave Background (CMB)	3
1.3 Other Observables	4
1.4 The (current) standard model of cosmology	5
1.5 The homogeneous and isotropic, expanding Universe	9
1.5.1 Einstein’s equations of general relativity	9
1.5.2 FLRW metric of the homogeneous and isotropic Universe	10
1.5.3 Conformal time	11
1.5.4 Particle horizon	11
1.5.5 Friedmann equations	11
1.6 Cosmic inflation	13
1.6.1 Definitions of inflation	14
1.6.2 The Horizon Problem	15
1.6.3 The Flatness Problem	15
1.6.4 Seeds for anisotropies	16
1.7 Correlation functions and power spectra	17
1.7.1 Primordial power spectra	18
1.7.2 CMB angular power spectra	18
1.8 Conclusion	21
2 Uniform vs logarithmic priors	23
2.1 Introduction	23
2.2 Methods	25
2.2.1 Bayesian inference	25
2.2.2 Kullback–Leibler divergence and Occam’s razor	26
2.2.3 Cosmological models	27

2.2.4	Data	27
2.2.5	Statistical and cosmological software	28
2.3	Tensor-to-scalar ratio	29
2.3.1	Tensor-to-scalar ratio: Posteriors	29
2.3.2	Tensor-to-scalar ratio: Evidence and Kullback–Leibler divergence	32
2.3.3	Mock example	35
2.4	Neutrino masses	38
2.4.1	Neutrino masses: Posteriors	39
2.4.2	Neutrino masses: Evidence and Kullback–Leibler divergence	42
2.4.3	Neutrino hierarchy	44
2.5	Discussion	45
3	A case for kinetic dominance	47
3.1	Introduction	47
3.2	The Planck surface in phase space	48
3.3	Inflaton potentials	49
3.3.1	Quadratic potential	49
3.3.2	Plateau potentials	50
3.3.3	Hilltop potentials	50
3.4	Inflaton Phase-Space	50
3.5	Priors on initial conditions	52
3.6	Conclusion	54
4	Constraints on kinetic dominance	55
4.1	Introduction	55
4.2	Background Evolution during Kinetic Dominance	56
4.2.1	Potentials	57
4.2.2	Kinetic Dominance initial conditions	59
4.3	Primordial Power Spectrum	60
4.3.1	Number of e-folds	65
4.3.2	Equation of state parameter	67
4.4	CMB power spectrum	68
4.5	MCMC analysis	68
4.5.1	Posteriors and priors on model parameters	70
4.5.2	Evidences	74
4.5.3	Power spectrum predictive posteriors and Kullback–Leibler divergences	75
4.6	Conclusions	77
5	Finite inflation in curved space	79
5.1	Introduction	79
5.2	Methods	82
5.2.1	Bayesian inference and nested sampling	82

5.2.2	ODE solvers	83
5.2.3	Statistical and cosmological software	84
5.2.4	Data	84
5.3	Inflationary background equations	84
5.4	Finite inflation	87
5.4.1	Kinetic dominance (KD)	87
5.4.2	Curvature domination	88
5.4.3	Inhomogeneities	90
5.5	Initial Conditions for inflation	90
5.5.1	Setting initial conditions at the start of inflation	90
5.6	Linking primordial to present-day scales	92
5.6.1	Calibration of the present-day scale factor	94
5.6.2	Calibration of the wavenumber of primordial perturbations	94
5.7	The comoving Hubble horizon	94
5.8	Conformal time	96
5.9	Reheating	101
5.10	The primordial power spectrum (PPS)	103
5.10.1	Power-law PPS	103
5.10.2	The slow-roll approximate PPS	104
5.10.3	The full numerical PPS	105
5.11	Inflationary potentials and slow-roll predictions	107
5.11.1	Monomial potential	108
5.11.2	Natural potential	110
5.11.3	Double-Well potential	111
5.11.4	Hilltop potential	111
5.11.5	Starobinsky potential	112
5.12	Choice of parametrisation	112
5.13	Nested sampling results	119
5.13.1	Λ CDM extensions	119
5.13.2	Inflation models	125
5.14	Conclusion	140
	Conclusion	143
	Future work	144
	Final remarks	145
	References	147

Acknowledgements

First and foremost, I would like to express how grateful I am to my supervisors, Anthony Lasenby, Mike Hobson and Will Handley. I feel incredibly lucky to have had the three of you. You have been a remarkable source of all things Bayesian, cosmology and computing and you have stayed patient with me when I lost patience with myself. Thank you for sharing your invaluable experience with me and thank you for all the guidance and support over these past four years. It has truly been a privilege and pleasure working with you.

Thank you also to my Master supervisors Adam Amara, Robert Brandenberger and Alexandre Réfrégier, who have paved the way towards this PhD.

I would also like to thank the Isaac Newton Trust, the STFC and the Cavendish Laboratory for their support, and all the people working behind the scenes in the Cavendish Astrophysics Group, the Kavli Institute for Cosmology and the Institute of Astronomy, who have helped me throughout these years.

A huge thanks goes to my fellow astro PhD students for sharing this journey with me. Thanks go to my office mates, Björn, Pablo, Ed, Fruzi and Roger for helpful discussions and welcome distractions. A special thanks to Björn for a good start into the working day with green tea and to Fruzi for our fruitful collaborations and for helping me to stay fit by teaming up for runs and work outs. Thanks go to my cosmology group mates from office K34, Carina, Do Young, Yun Cherng and Will, for our mutual exchange of worries, advice and words of encouragement. Do Young, I doubt I shall ever again have such a talented accompanist on the piano. It was a joy practising with you, I hope we can team up again some day. Further thanks go to my Churchill lunch group, Andrin, Íñigo, James, Lewis and Roger. Thanks to you, lunch was always a welcome break with lively discussions of physics, politics and beyond. A highlight each week was meeting up with my fellow astro boardgame geeks, Nick, Jasleen, Matteo, Sophie and Rob. This was especially true during my final months of writing up while being trapped inside in times of lockdown.

I would also like to express my gratitude to all my friends from outside of work. Thank you to my house mates, James, Tom and Stefan, for taking me in during a global pandemic. I am very grateful to everyone from Peterhouse and especially my Peterhouse pals, Carolina, Erika, Íñigo, Monique, Simone and Theresa, for the welcoming and caring atmosphere and the stimulating cross-disciplinary discussions. An unforgettable part of my life in Cambridge was the Peterhouse chapel choir. Thank you, Adam, Ben, Jenni, Joanne, Rachel, Seth and Simon, for the rehearsals, tours and concerts that have helped me balance out my working life. A huge thank you also to Alfonso, JJ, Marios and the whole Queerleader gang. Having some like-minded friends has played an important part for my mental well-being during my time here in Cambridge.

I would never have got to this point if it were not for the unconditional love and support from my closest friends and family at home, for which I am deeply grateful. Suse and Rainer, you have been there through all the difficult times we have faced as a family. Mum, Dad, Jojo, Jaja and Papus, you have always had my back and encouraged me, particularly in times of self-doubt, and I cannot thank you enough for this.

Conventions

Constants

The constants used in the numerical calculations of the primordial universe are based on the `scipy.constants` [1] package which in turn is based on the 2018 CODATA recommended values of the Fundamental Physical Constants as published by the National Institute of Standards and Technology (NIST) [2]:

$$\text{speed of light in vacuum} \quad c \equiv 299\,792\,458 \text{ m s}^{-1}, \quad (1)$$

$$\text{reduced Planck constant} \quad \hbar = 1.054\,571\,817 \dots \times 10^{-34} \text{ J s}, \quad (2)$$

$$\text{gravitational constant} \quad G = (6.674\,30 \pm 0.000\,15) \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}, \quad (3)$$

$$\text{Boltzmann constant} \quad k_B = 1.380\,649 \times 10^{-23} \text{ J K}^{-1}. \quad (4)$$

However, if not specified otherwise, I will be working in reduced Planck units with $c = \hbar = 1$.

Units

Based on the IAU 2012 Resolution B2 on the re-definition of the astronomical unit of length [3] and on the IAU 2015 Resolution B2 [4] implicitly defining the parsec we have:

$$1 \text{ au} \equiv 149\,597\,870\,700 \text{ m}, \quad (5)$$

$$1 \text{ pc} \equiv 648\,000 \pi^{-1} \text{ au} \approx 3.085\,677\,581\,491\,367 \times 10^{16} \text{ m} \quad (6)$$

For numerical calculations of the primordial Universe, we primarily work in *reduced* Planck units with the reduced Planck mass, time and length given by:

$$m_p \equiv \sqrt{\frac{\hbar c}{8\pi G}} \quad \approx 2.435 \times 10^{18} \text{ GeV } c^{-2} \quad \approx 4.341 \times 10^{-9} \text{ kg}, \quad (7)$$

$$t_p \equiv \sqrt{\frac{8\pi G \hbar}{c^5}} = \frac{\ell_p}{c} = \frac{\hbar}{m_p c^2} \quad \approx 4.106 \times 10^{-19} \hbar/\text{GeV} \quad \approx 2.703 \times 10^{-43} \text{ s}, \quad (8)$$

$$\ell_p \equiv \sqrt{\frac{8\pi G \hbar}{c^3}} = c t_p = \frac{\hbar}{m_p c} \quad \approx 4.106 \times 10^{-19} \hbar c/\text{GeV} \quad \approx 8.103 \times 10^{-35} \text{ m}. \quad (9)$$

Thus the units for some frequently used parameters with $c = \hbar = 1$ are:

$$\text{scale factor: } [a] = \ell_{\text{p}} = \frac{1}{m_{\text{p}}}$$

$$\text{Hubble parameter: } [H] = \frac{1}{t_{\text{p}}} = m_{\text{p}}$$

$$\text{energy density: } [\rho] = \frac{m_{\text{p}}}{\ell_{\text{p}}^3} = m_{\text{p}}^4$$

$$\text{inflaton field: } [\phi] = \frac{1}{\ell_{\text{p}}} = m_{\text{p}}$$

$$\text{inflaton potential: } [V(\phi)] = \frac{m_{\text{p}}}{\ell_{\text{p}}^3} = m_{\text{p}}^4$$

Initialisms

Λ CDM	cold dark matter with a cosmological constant Λ
CMB	cosmic microwave background
PPS	primordial power spectrum
SR	slow-roll
KD	kinetic dominance
NH	normal neutrino hierarchy
IH	inverted neutrino hierarchy

Other conventions

- The same metric signature $(-, +, +, +)$ is used as in [5].
- Repeated indices of tensorial equations are to be summed over. Greek indices run over all four components of spacetime tensors with 0 denoting the time component, whereas Roman indices run only over the three spatial components.
- The following Fourier convention is used:

$$\tilde{f}(\mathbf{k}) = \int_{-\infty}^{\infty} d^3\mathbf{x} f(\mathbf{x}) e^{-i\mathbf{k}\cdot\mathbf{x}} \quad f(\mathbf{x}) = \frac{1}{(2\pi)^3} \int_{-\infty}^{\infty} d^3\mathbf{k} \tilde{f}(\mathbf{k}) e^{i\mathbf{k}\cdot\mathbf{x}} \quad (10)$$

Chapter 1

Introduction

This chapter aims to provide the reader with a brief overview of key concepts of modern cosmology. Whilst it cannot compete with more extensive introductions, it serves as a review of the necessary theoretical background and introduces notation used later in the thesis. The books [6–10] and reviews [11, 12] were found very useful in its preparation.

1.1 The evolution of our Universe

At the heart of modern cosmology lies the assumption that our place in the Universe is in no way special and that at the largest scales any observer would see all directions as equivalent no matter where they are. This simple, yet powerful idea is called the *cosmological principle* and it is the continuous progression that we have taken in our historical views of the cosmos, from the geocentric to the heliocentric principle and beyond to the cosmological principle today. Throughout the early 20th century it indeed only was a principle, but particularly over the course of the last decades, observations of the sky on the largest scales have managed to back up the cosmological principle with empirical data, confirming the homogeneity and isotropy on the largest scales.

Another piece of the puzzle that is our Universe was the discovery that our Universe must be expanding. This led to what is now referred to as the hot Big Bang model, which proposes that our Universe originated about 14 billion years ago in a very hot and dense state and has been expanding and thus cooling ever since. This picture has been successful in explaining observations such as the relation of the receding velocity versus distance of far away galaxies called the Hubble diagram, the abundance of chemical elements in the Universe, and the black-body spectrum of the cosmic microwave background (CMB) radiation.

“A picture is worth a thousand words” the saying goes, and [fig. 1.1](#) conveys in a useful visual way:

- the evolution of our Universe with its key stages,
- how we know about it, i.e. what observational probes we have,
- the associated time and energy scales.

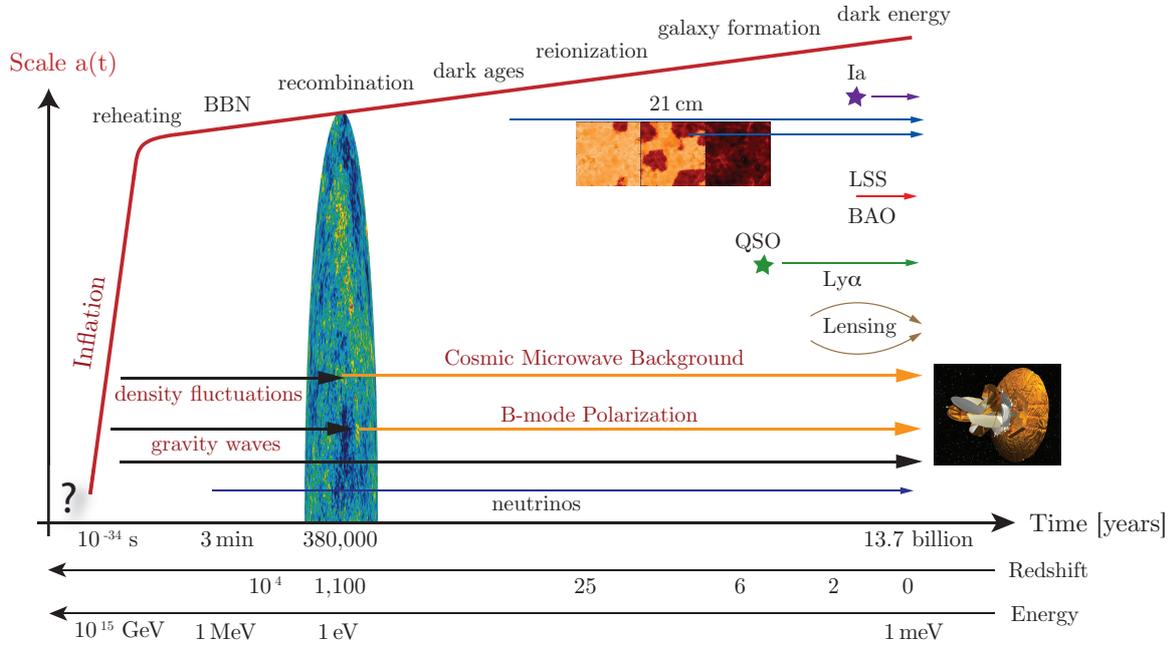


Figure 1.1 Schematic evolution of the Universe and how we know about it, with arrows indicating astrophysical observables (figure taken from [11]): Our current standard model of cosmology starts at the Big Bang marked by the question mark, highlighting our lack of a unified theory of quantum gravity to explain physics beyond the Planck scale. Since then our initially hot and dense Universe has expanded, illustrated by the red line roughly tracking the scale, and cooled, indicated by the x -axis for the Energy. The gradual cooling of the Universe causes interaction rates between the particles of the primordial plasma to drop until some interactions cease entirely, once their rate becomes smaller than the expansion rate of the Universe. In a tiny fraction of a second after the Big Bang, we assume that cosmic inflation caused an exponential expansion, ensuring that our observed Universe is isotropic to one part in 10^5 and close to flat. This was followed by neutrino decoupling and Big Bang Nucleosynthesis (BBN) where protons and neutrons started to combine to form light elements (mostly Hydrogen, Helium, and some Lithium). At recombination, protons and electrons start to combine and form atoms, such that the Universe becomes neutral. The leftover radiation is no longer scattered, but can now freely stream. It is called the cosmic microwave background (CMB) and is mostly isotropic, but has small observable anisotropies, the result of initial primordial density fluctuations. Similarly, primordial gravitational waves imprint an observable polarisation pattern onto the CMB. During the dark ages the Universe is essentially neutral such that no new light is being emitted, giving the epoch its name. Eventually the small initial density perturbations grow via gravitational instability to form the large scale structure (LSS). This gravitational collapse then leads to the formation of stars, galaxies and quasi stellar objects (QSO) that reionise the Universe. When the first stars run out of nuclear fuel, they explode as supernovae (e.g. type Ia supernovae), providing a mechanism for the formation of heavy particles.

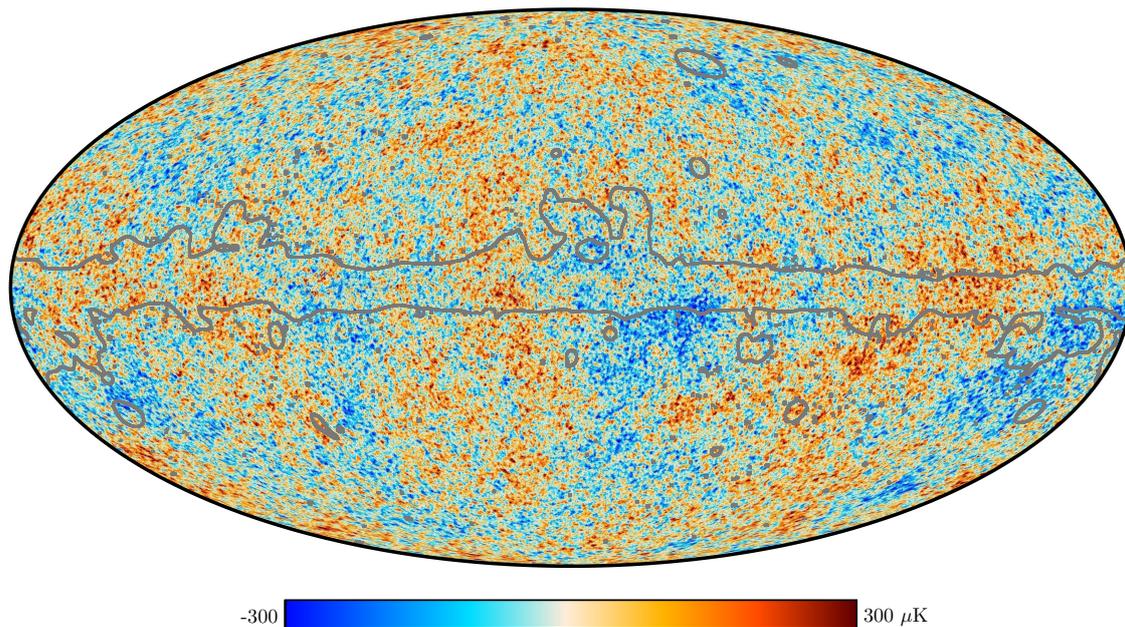


Figure 1.2 Map of the cosmic microwave background (CMB) anisotropies as observed by the Planck satellite [16] and produced with the SMICA component separation algorithm. The grey outline shows the extent of the confidence mask used to exclude the Milky Way and other foreground contamination from the computation of the CMB power spectrum.

1.2 The Cosmic Microwave Background (CMB)

When Wilson and Penzias first detected the cosmic microwave background (CMB) in 1965 [13, 14], they actually stumbled upon it more by accident. They planned to use their radiometer for radio astronomy, but after attempting to eliminate or account for every contribution to the noise temperature they still ended up with an excess noise of (3.5 ± 1.0) K. In his Nobel lecture, Wilson tells the anecdote of how they even suspected a pair of pigeons to be the potential culprits, but a thorough clean-up of their droppings did not rid the antenna of its residual noise [15]. Only after communication with Dicke did they become aware of the possibility that this residual noise might actually be the black-body background radiation left over from a hot and dense early universe.

Later, in 1990, the Cosmic Background Explorer (COBE) confirmed that the CMB indeed has a near-perfect black-body spectrum and revealed that there are very faint anisotropies across the sky. This last finding was further improved by the Wilkinson Microwave Anisotropy Probe (WMAP) [17–20] and by the Planck satellite [16, 21, 22].

These CMB observations are key pieces of evidence for the hot Big Bang model and have been crucial in transforming cosmology into a precision science. The CMB is the left-over radiation from the Big Bang that decoupled from ordinary matter at the epoch of recombination when the Universe became neutral roughly 380 000 years after the Big Bang. Since then it travelled freely through space until reaching us today. With a near-perfect black-body spectrum, the CMB has an almost uniform temperature of $T_{\text{CMB}} = (2.725\,48 \pm 0.000\,57)$ K [23] across the whole sky. Only tiny fluctuations can be observed, ranging from -300 μK to 300 μK , as shown on the full sky map in fig. 1.2 that was observed by the Planck satellite [16].

1.3 Other Observables

The Planck satellite has provided us with exquisite data of the CMB, allowing us to make inferences to unprecedented accuracy. Although we focus on data from the Planck satellite in this thesis, we should acknowledge that there are other probes that can be complementary to the CMB. This is particularly useful for the exploration of potential systematic errors or parameter degeneracies. Combining different data sets requires care, though. When two datasets disagree significantly on one or more parameters, then this is an indication of unaccounted systematic errors or of something missing in the underlying model, such that the combination of both datasets will not lead to a meaningful result. The currently most prominent example for this is the so-called “Hubble tension” referring to a discrepancy in the Hubble parameter between measurements of the early (CMB) and late (local) Universe [24].

CMB lensing

The CMB light that originated at the epoch of recombination and reaches us today had to pass through the large scale structure which deflected (or lensed) it weakly from its straight path via gravitational lensing. These small deflections lead to tiny distortions and blurring of the CMB pattern in temperature and polarisation maps of the sky that are related to the gravitational potential integrated along the line of sight. The effect of this lensing was measured to very high precision in both temperature and polarisation maps by the Planck satellite [25].

Further, this can be used to reconstruct the so-called lensing potential, i.e. the gravitational potential integrated along the line of sight, essentially a map of all the matter in the Universe in projection. This map can be used as data additional to the temperature and polarisation data.

Baryon Acoustic Oscillations (BAO)

Initial density fluctuations in the primordial plasma lead to counteracting forces of gravity and radiation-induced pressure. This causes oscillations similar to sound waves in air. While the (non-interactive) dark matter perturbations grow in place at the centre of these sound waves, the baryonic* perturbations are carried out in an expanding spherical wave. By the epoch of recombination when photons and baryons decouple, the radius of this wave has grown to a radius of about 150 Mpc and is commonly referred to as the sound horizon. The resulting dominant overdensities at the centre and at the sound horizon of that initial perturbation lead to a resonant peak in the power spectrum at the corresponding wavelength. This characteristic acoustic peak can be used as a so-called “standard ruler” measurable at a number of redshifts and thus allowing to determine the rate of expansion of the Universe (as parametrised by the Hubble parameter H_0) on geometric grounds [26].

*Note that in Astrophysics the term “baryons” is applied to all ordinary (i.e. non dark) matter. More details in [section 1.4](#).

Thus, BAOs provide a complementary approach to determining the Hubble parameter and curvature of the Universe that could help break the degeneracy between these parameters present in CMB data. However, there have been concerns about combining CMB with BAO data due to a potential tension between the datasets that might hint towards systematic errors or new physics [27–29].

Supernovae

Type Ia supernovae have a tight relation between peak absolute magnitude and the width of their light curve. This makes it possible to use them as so-called “standard candles”, meaning that the discrepancy of the observed brightness and the surface brightness inferred from the width of the light curve can be used to constrain the luminosity distance. This led to the 2011 Nobel prize “for the discovery of the accelerating expansion of the Universe through observations of distant supernovae” [30].

Distance measurements

A model-independent way of measuring the expansion rate of the Universe (as opposed to with the CMB where some realisation of a cosmological model needs to be assumed) is by calibrating the distances of supernovae with e.g. Cepheid variable stars [31] or stars from the tip of the red giant branch (TRGB) [32]. The Cepheid variables or TRGBs themselves are calibrated using parallax measurements in the Milky Way or through detached eclipsing binaries in the Large Magellanic Cloud or through masers in nearby galaxies [31]. This successive calibration of distance scales is also referred to as the cosmological distance ladder. The Hubble constant inferred from these measurements tends to be higher than from CMB or BAO measurements, and in case of the calibration with Cepheid variable stars the resulting Hubble constant of $H_0 = (74.03 \pm 1.42) \text{ km s}^{-1} \text{ Mpc}^{-1}$ is over 4σ in tension with CMB results. This so-called “Hubble tension” is currently the source of much discussion (e.g. at the 2019 conferences “Consistency of cosmological datasets: evidence for new physics?” in Cambridge, “Tensions between the Early and the Late Universe” in Santa Barbara, or “Cosmic Controversies” in Chicago).

1.4 The (current) standard model of cosmology

Thanks to its success in explaining the Hubble diagram (i.e. the farther away a galaxy is from us the faster it recedes from us), the abundance of light elements in the Universe and the uniform black-body radiation of the cosmic microwave background (CMB), the “Hot Big Bang” model has long been accepted as the standard model of cosmology. Over the course of the last couple decades, data from supernovae and particularly from the CMB have furthered our picture and the currently best-fitting model to describe our Universe is often known as Λ CDM: A universe that is dominated today by some form of dark energy behaving like a vacuum energy (i.e. like the cosmological constant Λ from general relativity) with density parameter $\Omega_\Lambda \approx 0.7$ indicating roughly 70% of the overall energy content, and where the

Table 1.1 Parameter constraints on the six Λ CDM sampling parameters and some derived parameters from Planck 2018 temperature and polarisation data [33]. The third column additionally uses data from Planck lensing and from baryon acoustic oscillations (BAOs).

Parameter	TT,TE,EE+lowE	TT,TE,EE+lowE+lensing+BAO
	68 % limits	68 % limits
$\Omega_b h^2$	0.02236 ± 0.00015	0.02242 ± 0.00014
$\Omega_c h^2$	0.1202 ± 0.0014	0.11933 ± 0.00091
$100 \theta_s$	1.04109 ± 0.00030	1.04101 ± 0.00029
τ_{reio}	$0.0544^{+0.0070}_{-0.0081}$	0.0561 ± 0.0071
$\ln(10^{10} A_s)$..	3.045 ± 0.016	3.047 ± 0.014
n_s	0.9649 ± 0.0044	0.9665 ± 0.0038
H_0	67.27 ± 0.60	67.66 ± 0.42
Ω_Λ	0.6834 ± 0.0084	0.6889 ± 0.0056
Ω_m	0.3166 ± 0.0084	0.3111 ± 0.0056
$10^9 A_s$	$2.101^{+0.031}_{-0.034}$	2.105 ± 0.030

matter $\Omega_m \approx 0.3$ behaves like a cold and dark fluid, i.e. it is essentially collisionless and does not interact electromagnetically. For a more in-depth review I recommend the (skeptical) guide to the Λ CDM model by Scott (2018) [12] as an excellent reference.

The Λ CDM model is typically described through 6 (arguably 7) free parameters. The Planck satellite has measured these to high precision [33–35]. Table 1.1 lists some parameter constraints from Planck 2018 temperature and polarisation data [33]. The free parameters are:

Baryon density $\omega_b = h^2 \Omega_b$: The baryonic matter density Ω_b is made up of the ordinary (non-relativistic) matter in the Universe, i.e. atoms, ions, protons, electrons etc. Note that this is the astronomical definition of baryons and as such differs from the concept of baryons in particle physics. Often a rescaling with the dimensionless Hubble parameter h is included making it proportional to the physical energy density: $\omega_b = h^2 \Omega_b$.

Cold dark matter density $\omega_c = h^2 \Omega_c$: The cold dark matter density Ω_c makes up the remaining part (after subtracting the baryon density) of the overall matter density Ω_m . Here, too, a rescaling with the dimensionless Hubble parameter h is commonly included: $\omega_c = h^2 \Omega_c$.

Angular sound horizon $100 \theta_s$: The distance sound waves could have travelled in the time before recombination is known as the sound horizon r_s . Together with the distance d to the surface of last scattering of photons (at recombination where the Universe became neutral) it determines the angular size of the sound horizon $\theta_s \approx \frac{r_s}{d}$. It is commonly referred to multiplied by a factor of 100.

Reionisation optical depth τ_{reio} : The epoch when matter sufficiently collapsed under gravity such that the first stars formed which (re-)ionised the Universe is referred to as the epoch of reionisation. The optical depth to reionisation τ_{reio} quantifies how much the CMB photons are scattered by that reionised medium at late times.

Scalar power spectrum amplitude $\ln(10^{10} A_s)$: The amplitude A_s of the initial scalar perturbations (also referred to as primordial density perturbations) is typically measured at the pivot scale $k_* = 0.05 \text{ Mpc}^{-1}$ and given in its logarithmic form. These initial perturbations serve as seeds for the CMB anisotropies which in turn are the seeds for the eventual formation of the large scale structure.

Scalar spectral index n_s : The spectral index n_s of primordial density perturbations governs the slope of the initial power spectrum and, like the amplitude A_s , is measured at the pivot scale $k_* = 0.05 \text{ Mpc}^{-1}$. In an early attempt to explain the observed structures in the Universe, Harrison (1970) [36], and Peebles and Yu (1970) [37] independently suggested a scale-invariant initial power spectrum with $n_s = 1$ from adiabatic initial conditions, which was further picked up by Zeldovich (1972) [38]. Cosmic inflation (more on this in [section 1.6](#)) predicts the scalar spectral index to be slightly below (albeit very close to) scale invariance ($n_s \lesssim 1$). One of the major successes of inflationary cosmology was the confirmation of this departure from scale invariance at the 8σ level (see [table 1.1](#)) with CMB data.

CMB temperature T_{CMB} : Today’s temperature of the CMB has been measured to extremely high precision to be $T_{\text{CMB}} = (2.725\,48 \pm 0.000\,57) \text{ K}$ [23]. As such, it is often considered to be fixed and not treated as a free parameter.

Beyond these 6 (or 7) free parameters, there are various assumptions that go into the Λ CDM model. Many of them can be tested for their relevance through parameter extensions of the Λ CDM model, e.g. extensions with the curvature density parameter $\Omega_{K,0}$, the effective number of relativistic species N_{eff} , the sum of neutrino masses $\sum m_\nu$ of the three neutrino mass eigenstates, the running n_{run} of the scalar spectral index, or the tensor-to-scalar ratio r of primordial perturbations. [Table 1.2](#) shows the constraints on these parameters [33]. In the Λ CDM model these parameters are accounted for as follows:

Spatial curvature $\Omega_{K,0}$: The global spatial geometry of the Universe is assumed to be flat $\Omega_{K,0} = 0$ in the base Λ CDM cosmology. Since the curvature constraints are degenerate with some other parameters (e.g. the present-day Hubble H_0 , matter density Ω_m and Λ density parameter Ω_Λ), tighter constraints can be achieved through combination of different types of data. A joint analysis of data from CMB, CMB lensing and baryon acoustic oscillations (BAO) leads to $\Omega_{K,0} = 0.0007 \pm 0.0037$ [33], perfectly consistent with spatial flatness. However, there have been concerns about possible tension between CMB and lensing or BAO datasets when including non-zero curvature [27–29]. CMB data on its own actually prefers a closed universe (see also [table 1.2](#)). The subscript “0” refers to the present-day value of the parameter and, different to the other density parameters

(e.g. Ω_b and Ω_c), it is retained for the curvature density parameter $\Omega_{K,0}$ in this thesis to clearly set it apart from the also frequently used time-variable parameter Ω_K .

Effective number of relativistic species N_{eff} : The standard Λ CDM model assumes an effective number of relativistic species of $N_{\text{eff}} \approx 3.046$ [39, 40]. The 3 is for the three standard model neutrinos. The effective number is expected to be slightly larger than three, because the neutrinos were not completely decoupled when electron-positron annihilation began. Hence, the neutrinos were heated up by decaying electron-positron pairs [33, 41]. Results from an extension of the Λ CDM model with N_{eff} are perfectly consistent with the base Λ CDM assumption.

Neutrino masses m_ν : Since neutrino oscillation experiments measure two non-zero mass-squared differences [42–44], we know that there must be three discrete mass eigenstates. But for the Λ CDM cosmology this is typically simplified with the assumption of two massless and one massive neutrino with $\sum m_\nu = 0.06$ eV. Currently CMB and BAO data only provide a rather loose upper bound. Future CMB experiments aim at pushing this bound further down, such that the ordering of the mass squared splittings from oscillation experiments can be resolved [45–47].

Running of the scalar spectral index n_{run} : The running $n_{\text{run}} \equiv dn_s / d \ln k$ of the scalar spectral index n_s is assumed to be negligible. Most inflation models from a single scalar field predict a slightly negative value for n_{run} . The current best constraints from CMB data are consistent with no running and with all major inflation models with a single scalar field.

Tensor-to-scalar ratio r : Primordial gravitational waves are assumed to be negligible. Most inflation models from a single scalar field predict a non-zero value. A possible detection of such tensor modes in the primordial power spectrum is often referred to as a “smoking gun” of inflation, providing the best indirect evidence together with the scalar spectral index n_s . Lowering the constraints on this parameter is a prime target in most upcoming CMB experiments such as the Simons Observatory [45], the LiteBIRD satellite [46] or CMB-S4 [47].

Table 1.2 Parameter constraints on single-parameter extensions of the Λ CDM model. The second column shows constraints from Planck 2018 temperature and polarisation data [33]. The third column additionally uses data from Planck lensing and from baryon acoustic oscillations (BAOs). Note that the tensor-to-scalar ratio $r_{0.002}$ is reported for the pivot scale $k_* = 0.002 \text{ Mpc}^{-1}$ in the Planck paper, differing from the pivot scale otherwise used in this thesis of $k_* = 0.05 \text{ Mpc}^{-1}$.

Parameter	TT,TE,EE+lowE	TT,TE,EE+lowE+lensing+BAO
	95 % limits	95 % limits
$\Omega_{K,0}$	$-0.044^{+0.033}_{-0.034}$	0.0007 ± 0.37
N_{eff}	$2.92^{+0.36}_{-0.37}$	$2.99^{+0.34}_{-0.33}$
$\sum m_\nu$	$< 0.257 \text{ eV}$	$< 0.120 \text{ eV}$
$dn_s / d \ln k$..	-0.006 ± 0.013	-0.004 ± 0.013
$r_{0.002}$	< 0.107	< 0.106

1.5 The homogeneous and isotropic, expanding Universe

This section reviews the mathematical background on which cosmic inflation is built. A combination of Einstein’s theory of general relativity (section 1.5.1) with the Friedmann–Lemaître–Robertson–Walker (FLRW) metric (section 1.5.2) leads to the Friedmann equations of cosmology (section 1.5.5), which govern the dynamics of the expansion of the Universe.

1.5.1 Einstein’s equations of general relativity

Einstein’s theory of general relativity [48] is rooted in the equivalence principle, which states that the *inertial* mass of a particle is equal to the *gravitational* mass of that particle in a gravitational field. Equivalently we can say that the acceleration of a particle in a gravitational field is independent of the particle’s nature. This removes gravity from being one of the fundamental forces and turns it into a property of spacetime itself. The other forces (the electromagnetic, strong or weak force) still act on objects on that spacetime background. However, in the absence of other forces, objects move along their geodesics, the generalisation of the notion of a “straight line” in a curved spacetime. The spacetime itself in turn is distorted by the presence of massive objects.

The formalism of general relativity can be effectively summarised with the variational action principle:

$$S = \int d^4x \sqrt{-g} (\mathcal{L}_{\text{EH}} + \mathcal{L}_{\text{matter}}), \quad (1.1)$$

where g corresponds to the determinant of the metric tensor $g_{\mu\nu}$ and where the Lagrangian density \mathcal{L} is split into a matter part $\mathcal{L}_{\text{matter}}$ and into a gravitational part $\mathcal{L}_{\text{EH}} = \frac{1}{2}m_{\text{p}}^2 R$, governed by the Ricci scalar R . The integral of \mathcal{L}_{EH} on its own is known as the Einstein–

Hilbert action. The matter Lagrangian $\mathcal{L}_{\text{matter}}$ represents the Universe's matter (or energy) content and is connected to the energy-momentum tensor

$$T_{\mu\nu} \equiv \frac{-2}{\sqrt{-g}} \frac{\delta}{\delta g^{\mu\nu}} \left(\sqrt{-g} \mathcal{L}_{\text{matter}} \right). \quad (1.2)$$

The Euler–Lagrange equations will then lead to Einstein's field equations of general relativity, given by:

$$R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R = \frac{1}{m_{\text{p}}^2} T_{\mu\nu}. \quad (1.3)$$

The Ricci scalar R and tensor $R_{\mu\nu}$ encapsulate the curvature of the spacetime of the Universe. The Ricci scalar is the contraction of the Ricci tensor with the metric, while the Ricci tensor itself is given by:

$$R_{\mu\nu} = \Gamma_{\mu\nu,\alpha}^{\alpha} - \Gamma_{\mu\alpha,\nu}^{\alpha} + \Gamma_{\beta\alpha}^{\alpha} \Gamma_{\mu\nu}^{\beta} - \Gamma_{\beta\nu}^{\alpha} \Gamma_{\mu\alpha}^{\beta}, \quad (1.4)$$

with the Christoffel symbols

$$\Gamma_{\alpha\beta}^{\mu} = \frac{1}{2} g^{\mu\nu} \left(g_{\alpha\nu,\beta} + g_{\beta\nu,\alpha} - g_{\alpha\beta,\nu} \right), \quad (1.5)$$

where a subscript comma denotes a partial derivative, i.e. $A_{,\alpha} = \frac{\partial}{\partial x^{\alpha}} A$.

In principle one can add a cosmological constant term $\Lambda g_{\mu\nu}$ to the left hand side of eqs. 1.3. However, this term can also be interpreted as the vacuum energy density ρ_{vac} and thus included in the energy-momentum tensor $T_{\mu\nu}$ on the right hand side, which is the approach adopted in this thesis.

1.5.2 FLRW metric of the homogeneous and isotropic Universe

Assuming *homogeneity* and *isotropy* on the largest scales leads to the Friedmann–Lemaître–Robertson–Walker (FLRW) metric of spacetime [49–54]. Using comoving, spherical coordinates r , θ , and ϕ , it is defined through the line element:

$$ds^2 = g_{\mu\nu} dx^{\mu} dx^{\nu} = -dt^2 + a^2(t) \left[\frac{dr^2}{1 - Kr^2} + r^2 \left(d\theta^2 + \sin^2\theta d\phi^2 \right) \right] \quad (1.6)$$

where the scale factor $a(t)$ characterises the relative size and K the global curvature of the spatial hypersurfaces of constant time*. The spatial curvature parameter K can take the following values:

$$K = \begin{cases} +1, & \text{spherical or closed,} \\ 0, & \text{Euclidean or flat,} \\ -1, & \text{hyperbolic or open.} \end{cases} \quad (1.7)$$

*Note, that there are different conventions in the treatment of the spatial curvature parameter K . Here, we absorb any arbitrariness in the magnitude of K into the radial coordinate and the scale factor, such that K only takes one of $\{-1, 0, +1\}$ (see e.g. [8] for more details). Because of the sign of the spatial curvature parameter K , another common way of referring to the geometry of the Universe is as positive and negative curvature. Note, however, that this can be easily confused with the sign of the curvature density parameter Ω_K (see definition 1.19), which has the opposite sign of K .

Note that for the dynamics in flat ($K = 0$) universes the overall normalisation of the scale factor $a(t)$ is irrelevant. All that matters is the ratio of $a(t)$ at different times. In [chapter 5](#) we will see, though, how in curved ($K \neq 0$) universes the overall normalisation of $a(t)$ plays an important role in linking the scale of the Universe across different epochs in the evolution of the Universe.

1.5.3 Conformal time

It is often convenient to define conformal time η as an alternative time coordinate to physical time t :

$$d\eta \equiv \frac{dt}{a(t)} \quad (1.8)$$

Writing the FRW metric with respect to conformal time, it becomes conformally equivalent to Minkowski space (hence the name):

$$ds^2 = a^2(\eta) \left(-d\eta^2 + dX^2 \right), \quad (1.9)$$

where we have abbreviated the comoving spatial part of the FRW metric in [eq. \(1.6\)](#) as dX^2 . Thus, and since $ds^2 = 0$ for photons, the conformal time η is also equal to the comoving distance light could have traveled in a time interval t . In other words, the trajectories for photons are straight lines at 45° angles in spacetime diagrams with η and X as axes (where we assume that our coordinate system points in the radial direction of the photon's motion).

1.5.4 Particle horizon

Having linked conformal time to the comoving distance of radially propagating photons in the preceding section, it is now straightforward to define the maximum comoving distance χ that light can propagate between an initial time t_i and some later time t , which we call the particle horizon:

$$\chi_p(\eta) = \eta - \eta_i = \int_{t_i}^t \frac{d\tilde{t}}{a(\tilde{t})}. \quad (1.10)$$

1.5.5 Friedmann equations

Applying the cosmological principle also to the contents of the universe, we assume the energy-momentum tensor to be that of a perfect isotropic fluid as seen by a comoving observer: $T^\mu{}_\nu = \text{diag}(-\rho, p, p, p)$ with an energy density ρ and a pressure p .

Together with the FRW metric from [eq. \(1.6\)](#) we can then condense the dynamics from the Einstein [eqs. 1.3](#) to two equations, the Friedmann equations, governing the dynamics of

Table 1.3 Dynamics of the scale factor a , energy density ρ and comoving Hubble horizon $(aH)^{-1}$ for different dominating components to the matter-energy content of the Universe expressed through their equation-of-state parameter w . Note that curvature domination will lead to gravitational collapse for closed universes. Hence, taken on its own, the proportionalities for curvature in the last row are only valid solutions for open universes.

domination	equation-of-state	scale factor	energy density	comoving Hubble horizon
general	$w \neq -1$	$a = t^{\frac{2}{3(1+w)}}$	$\rho = a^{-3(1+w)}$	$(aH)^{-1} \propto a^{\frac{1+3w}{2}}$
radiation	$w = +\frac{1}{3}$	$a = t^{1/2}$	$\rho = a^{-4}$	$(aH)^{-1} \propto a^1$
matter	$w = 0$	$a = t^{2/3}$	$\rho = a^{-3}$	$(aH)^{-1} \propto a^{1/2}$
cosmological constant Λ dark energy inflation	$w \approx -1$	$a = e^{Ht}$	$\rho = \text{const.}$	$(aH)^{-1} \propto a^{-1}$
curvature	$w = -\frac{1}{3}$	$a = t^1$	$\rho = a^{-2}$	$(aH)^{-1} \propto \text{const.}$

the expansion of the smooth universe in form of the scale factor $a(t)$:

$$H^2 \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{1}{3m_{\text{p}}^2}\rho - \frac{K}{a^2} \quad (1.11)$$

$$\dot{H} + H^2 = -\frac{1}{6m_{\text{p}}^2}(\rho + 3p), \quad (1.12)$$

where a dot denotes differentiation with respect to cosmic time ($\dot{f} \equiv \frac{df}{dt}$) and where $H = \frac{\dot{a}}{a}$ is the Hubble parameter that quantifies the expansion rate. Combining eqs. (1.11) and (1.12) or equivalently using the conservation of the energy-momentum tensor $\nabla_{\mu}T^{\mu}_{\nu} = 0$ we can derive the continuity equation:

$$\dot{\rho} = -3H(\rho + p). \quad (1.13)$$

For a general fluid with an equation of state parameter $w_i = p_i/\rho_i$, we can then derive the dynamics of the energy density as:

$$\rho_i \propto a^{-3(1+w_i)}, \quad (1.14)$$

$$a \propto \begin{cases} t^{2/3(1+w_i)} & w \neq -1 \\ e^{H_0 t} & w = -1 \end{cases} \quad (1.15)$$

Table 1.3 lists the solutions for the dynamics of the scale factor and energy density for some of the commonly assumed fluids.

In a somewhat more realistic scenario, one assumes a mixture of multiple fluids that do not interact except through their mutual gravitation, but that otherwise each behave like a

perfect fluid contributing to the energy density and pressure from eqs. (1.11) to (1.13). In that case we need to sum over all components:

$$\rho = \sum_i \rho_i \quad \text{and} \quad p = \sum_i p_i, \quad (1.16)$$

where the index runs over all relevant contributions. In our Universe, we observe contributions from radiation ($w_r = \frac{1}{3}$), matter ($w_m = 0$), and a cosmological constant ($w_\Lambda = -1$). We can also incorporate the contribution from curvature to the dynamics of the scale factor in form of a cosmological fluid with equation of state parameter $w_K = -\frac{1}{3}$.

It is common to define the dimensionless energy density parameter in form of the ratio of the energy density relative to the critical energy density $\rho_{\text{crit}} = 3 m_{\text{p}}^2 H_0^2$:

$$\Omega_i \equiv \frac{\rho_i}{\rho_{\text{crit}}}. \quad (1.17)$$

This allows us to rewrite eq. (1.11) as

$$\left(\frac{H}{H_0}\right)^2 = \sum_i \Omega_{i,0} \left(\frac{a}{a_0}\right)^{-3(w_i+1)}, \quad (1.18)$$

with the curvature density parameter defined as

$$\Omega_K \equiv -\frac{K}{(aH)^2} \quad (1.19)$$

and with the consistency relation

$$\sum_i \Omega_i = 1. \quad (1.20)$$

1.6 Cosmic inflation

Despite its great success, there are some shortcomings to the hot Big Bang model in explaining the initial state of the Universe. These shortcomings are the horizon problem, the flatness problem and the question of the seed of anisotropies that lead to our observed large scale structure of the Universe. We outline these problems in sections 1.6.2 to 1.6.4.

The leading theory to address these questions is through the phenomenon of *cosmic inflation* that was first developed in the early 1980s by Starobinskiĭ (1979) [55], Guth (1981) [56], Linde (1982) [57] and Albrecht and Steinhardt (1982) [58]. Cosmic inflation assumes that the Universe underwent a period of exponential expansion very early in the Universe's history (somewhere between 10^{-36} and 10^{-32} seconds after the Big Bang), before the start of the standard hot Big Bang evolution. Intriguingly, inflation offers an explanation to both the large-scale homogeneity and flatness, as well as the small scale fluctuations that grew to become structures like our own Milky Way.

The simplest way of implementing such an early accelerated expansion is via a single scalar field ϕ , which we call the inflaton, and its potential $V(\phi)$. Assuming that the inflaton field is the only contribution to the matter-energy content of the hot and dense early Universe, the matter Lagrangian from eq. (1.1) describing the dynamics of the inflaton field minimally coupled to gravity is given by:

$$\mathcal{L}_{\text{matter}} = \mathcal{L}_\phi = \frac{1}{2} g^{\mu\nu} \partial_\mu \phi \partial_\nu \phi - V(\phi). \quad (1.21)$$

With the FLRW metric from eq. (1.6) and restricting to the case of a homogeneous field $\phi(t, \mathbf{x}) = \phi(t)$ this leads to an energy-momentum tensor of the form of a perfect fluid with energy density ρ_ϕ and pressure p_ϕ given by:

$$\rho_\phi = -\frac{1}{2} \dot{\phi}^2 + V(\phi), \quad p_\phi = -\frac{1}{2} \dot{\phi}^2 - V(\phi). \quad (1.22)$$

Inserting these into the Friedmann eqs. (1.11) and (1.12) and the continuity eq. (1.13) we obtain the background equations of inflation governing the expansion of the Universe and the equation of motion for the inflaton field:

$$H^2 = \frac{1}{3 m_{\text{p}}^2} \left(\frac{1}{2} \dot{\phi}^2 + V(\phi) \right) - \frac{K}{a^2}, \quad \dot{H} = -\frac{1}{2 m_{\text{p}}^2} \dot{\phi}^2 + \frac{K}{a^2}, \quad \ddot{\phi} + 3H\dot{\phi} + V'(\phi) = 0. \quad (1.23)$$

1.6.1 Definitions of inflation

Accelerated expansion: Cosmic inflation is most commonly described as a period of exponential or accelerated expansion, motivating the definition through a positive second derivative of the scale factor:

$$\ddot{a} > 0. \quad (1.24)$$

Shrinking comoving Hubble horizon: To address various shortcomings of the hot Big Bang model (see e.g. section 1.6.2), it is often more useful to work with the equivalent definition of a decreasing comoving Hubble horizon $(aH)^{-1}$:

$$\frac{d}{dt} \left(\frac{1}{aH} \right) < 0. \quad (1.25)$$

This can be easily shown to lead to an accelerated expansion as in definition 1.24 by explicitly carrying out the time derivative.

Negative pressure: Definition 1.25 in turn can be related to the second Friedmann eq. (1.12) resulting in a third option of defining inflation

$$p < -\frac{1}{3}\rho \quad \Rightarrow \quad w \equiv \frac{p}{\rho} < -\frac{1}{3}. \quad (1.26)$$

Thus, inflation implies negative pressure.

Potential dominated scalar field: With the expressions for energy density ρ_ϕ and pressure p_ϕ from eq. (1.22) we can turn the definition 1.26 into one directly relating to the inflaton field:

$$V(\phi) > \dot{\phi}^2 \quad (1.27)$$

Thus, when the inflaton's potential dominates over the kinetic term ($V(\phi) \gg \dot{\phi}^2$), then the Universe undergoes an exponential expansion.

1.6.2 The Horizon Problem

The comoving *Hubble* horizon $(aH)^{-1}$ is usually described as the length scale within which physical processes operate coherently. It is also the length scale beyond which general relativistic effects become important. It should not be confused with the comoving *particle* horizon

$$\eta \equiv \int_0^t \frac{dt}{a(t)} = \int_0^a \frac{da}{a^2 H} = \int_0^a \frac{d \ln a}{aH}, \quad (1.28)$$

which is the maximum comoving distance travelled by light since the beginning of the Universe. This tells us that objects separated by comoving distances greater than η today could never have been in causal contact.

This poses a problem for the conventional hot Big Bang model. With the comoving Hubble horizon growing monotonically during radiation and matter domination (cf. table 1.3), there was not enough time before the epoch of recombination to allow for a dynamical process to explain the near-homogeneity of the CMB. Indeed we expect that a causally connected patch on the CMB subtends an angle of only about 1.6° [59], far from the uniformity that we observe across the whole sky. This problem is traditionally referred to as the *horizon problem*.

However, this problem can be solved by cosmic inflation. Since the comoving Hubble horizon $(aH)^{-1}$ is shrinking during inflation (cf. definition 1.25), it must have been much greater in the past. Since the comoving Hubble horizon is in the integrand of eq. (1.28), it is possible that η was much greater than $(aH)^{-1}$ at the epoch of recombination and that scales that were outside the comoving Hubble horizon then, might have been in causal contact well before.

1.6.3 The Flatness Problem

Although current data still leave room for a non-zero present-day curvature density parameter $\Omega_{K,0}$, it seems safe to say that this must be small (see also table 1.2). Since the curvature density parameter Ω_K from definition 1.19 is proportional to the square of the comoving Hubble horizon, and since the comoving Hubble horizon is growing monotonically during radiation and matter domination of the hot Big Bang evolution (cf. table 1.3), this means that the curvature density, already small today, must have been very finely tuned to near zero in the past. This tuning of the curvature density is referred to as the *flatness problem*. Although

there is no reason why the curvature density parameter should not have been very small in the past, a theory that generically predicts this would be more desirable from a fine-tuning perspective.

If the radiation-dominated epoch of the hot Big Bang was preceded by a sufficiently long period of cosmic inflation this would solve the flatness problem. Since by [definition 1.25](#) the comoving Hubble horizon decreases during inflation, this naturally drives the universe toward flatness (keep in mind [definition 1.19](#)).

1.6.4 Seeds for anisotropies

The most remarkable feature of cosmic inflation is that it explains simultaneously the large-scale smoothness and the small-scale deviation from that said smoothness. It provides a mechanism for the origin of the tiny anisotropies that can be observed in the CMB and that through gravitational collapse eventually grow and form the large scale structure of galaxies, stars, and planets.

For a detailed review of the calculation of the primordial power spectra (PPS) generated by quantum fluctuations, I refer the interested reader to the lecture notes by Baumann (2009) [\[11\]](#) or to their favourite book (e.g. one of [\[7–10\]](#)). The general idea is to use the Einstein [eqs. 1.3](#) and relate perturbations of the inflaton field (entering via the energy-momentum tensor)

$$\phi \rightarrow \phi + \delta\phi, \quad (1.29)$$

to perturbations of the metric $g_{\mu\nu}$ from [eq. \(1.6\)](#):

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu = -(1 + 2\Phi) dt^2 + 2aB_i dx^i dt + a^2 [(1 - 2\Psi)\delta_{ij} + E_{ij}] dx^i dx^j. \quad (1.30)$$

$$B_i \equiv \partial_i B - S_i, \quad E_{ij} \equiv 2\partial_{ij} E + \partial_i F_j + \partial_j F_i + h_{ij} \quad (1.31)$$

The term in square brackets consists of the *spatial curvature perturbation* Ψ and the *spatial shear* E_{ij} , which is symmetric ($E_{ij} = E_{ji}$) and traceless ($E_i^i = 0$). The term B_i is called the *shift function* and Φ the *lapse function*. The vector perturbations S_i and F_i decay in an expanding universe and are irrelevant for simple inflation models. The tensor perturbations h_{ij} are traceless ($h_i^i = 0$) and divergenceless ($\partial^i h_{ij} = 0$) (or transverse).

In order not to confuse physical perturbations with coordinate perturbations, it is more convenient to work with gauge-invariant quantities, i.e. variables that are insensitive to perturbative coordinate changes ($x^\mu \rightarrow x^\mu + \delta x^\mu$). One such gauge-invariant scalar is the *comoving curvature perturbation*, which for the inflaton field ϕ is given by

$$\mathcal{R} = \Psi + \frac{H}{\dot{\phi}} \delta\phi. \quad (1.32)$$

Applying the Einstein [eqs. 1.3](#) to the metric and inflaton perturbations leads to the second-order differential equation of motion for the Fourier transform \mathcal{R}_k of the comoving curvature

perturbation* [60]:

$$\ddot{\mathcal{R}}_k + \left(\frac{\dot{\phi}^2}{m_{\text{p}}^2 H} + \frac{2\ddot{\phi}}{\dot{\phi}} + 3H \right) \dot{\mathcal{R}}_k + \frac{k^2}{a^2} \mathcal{R}_k = 0, \quad (1.33)$$

with the wavenumber k . Once a scale comes to exceed the shrinking comoving Hubble horizon, i.e. once $k \ll aH$, the comoving curvature perturbation freezes and is conserved until it re-enters in the more recent past during radiation or matter domination.

Besides the scalar perturbations, which are associated with density fluctuations, there are also tensor perturbations in eq. (1.30). The gauge-invariant tensor perturbations h_{ij} are associated with primordial gravitational waves and can be specified by two scalars, h^+ and h^\times , corresponding to the two possible polarisation states. The equivalent expression to eq. (1.33) for both tensor modes h_k^+ and h_k^\times is given by

$$\ddot{h}_k + 3H\dot{h}_k + \frac{k^2}{a^2} h_k = 0. \quad (1.34)$$

1.7 Correlation functions and power spectra

The two-point (auto-)correlation function of a field $\varphi(\mathbf{x})$ is a measure for the similarity (or difference) between two points \mathbf{x} separated by a distance \mathbf{r} . It is defined as

$$\xi_\varphi(\mathbf{r}) = \langle \varphi(\mathbf{x})\varphi(\mathbf{x} + \mathbf{r}) \rangle. \quad (1.35)$$

Often it is more convenient to work in Fourier space. The Fourier transform of the correlation function $\xi(\mathbf{r})$ is called the power spectrum $P(\mathbf{k})$ and equivalently defined as

$$(2\pi)^3 \delta^{(3)}(\mathbf{k} - \mathbf{k}') P(\mathbf{k}) = \langle \varphi(\mathbf{k})\varphi(\mathbf{k}') \rangle \quad (1.36)$$

with the Fourier convention

$$P_\varphi(\mathbf{k}) = \int d^3r \xi_\varphi(\mathbf{r}) e^{-i\mathbf{k}\cdot\mathbf{r}}. \quad (1.37)$$

In case of an isotropic and homogeneous Universe both correlation function and power spectrum simplify and depend only on radial distance $\xi_\varphi(\mathbf{r}) = \xi_\varphi(r)$ or wavenumber magnitude $P_\varphi(\mathbf{k}) = P_\varphi(k)$. This allows us to rewrite eqs. (1.35) to (1.37) as

$$\sigma_\varphi \equiv \langle \varphi(x) \rangle = \xi_\varphi(0) = \int \frac{d^3k}{(2\pi)^3} P_\varphi(k) = \int d(\log k) \mathcal{P}_\varphi(k) \quad (1.38)$$

*For simplicity we only show the equation of motion for a flat universe here. The corresponding one including non-zero curvature is discussed in chapter 5.

with the dimensionless power spectrum defined as

$$\mathcal{P}_\varphi(k) \equiv \frac{k^3}{2\pi^2} P_\varphi(k). \quad (1.39)$$

1.7.1 Primordial power spectra

Since both the comoving curvature perturbations \mathcal{R}_k as well as the tensor perturbations h_k are conserved on super-horizon scales (see also [section 1.6.4](#)), their primordial power spectra can be read off after horizon crossing:

$$\mathcal{P}_{\mathcal{R}}(k) = \frac{k^3}{2\pi^2} |\mathcal{R}_k|^2, \quad (1.40)$$

$$\mathcal{P}_t(k) = 2 \cdot \frac{k^3}{2\pi^2} |h_k|^2, \quad (1.41)$$

where the factor of 2 accounts for the two possible polarisation states of the tensor modes.

Slow-roll single-field inflation leads to an almost scale invariant scalar power spectrum. The standard Λ CDM model accounts for this with the phenomenological power-law power spectrum

$$\mathcal{P}_{\mathcal{R}}(k) = A_s \left(\frac{k}{k_*} \right)^{n_s - 1}, \quad (1.42)$$

where the power amplitude A_s and the spectral index n_s (previously mentioned in [section 1.4](#) and [table 1.1](#)) are defined at the pivot scale k_* . If not stated differently the pivot scale $k_* = 0.05 \text{ Mpc}^{-1}$ is used in this thesis. The deviation of the spectral index n_s from scale invariance has been confirmed with CMB data to almost 8σ precision.

1.7.2 CMB angular power spectra

In [fig. 1.2](#) we have shown the full-sky map of the CMB temperature anisotropies. Since these fluctuations are measured across the celestial sphere, a two-dimensional surface, it is more convenient to expand the temperature deviations $\Delta T(\hat{\mathbf{n}})$ in a direction given by the unit vector $\hat{\mathbf{n}}$ in spherical harmonics $Y_{\ell,m}(\hat{\mathbf{n}})$ rather than to work with the three-dimensional Fourier spectrum (see [\[61\]](#) for more details):

$$\Delta T(\hat{\mathbf{n}}) \equiv T(\hat{\mathbf{n}}) - \langle T \rangle = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}). \quad (1.43)$$

In general, no cosmological theory can tell us the particular values of the coefficients $a_{\ell m}$, which would imply knowledge about the exact position of fluctuations on the sky. Of cosmological interest is instead the statistical distribution of the coefficients $a_{\ell m}$. The mean value is zero, but the variance will have some non-zero value:

$$\langle a_{\ell m} \rangle = 0, \quad \langle a_{\ell m} a_{\ell' m'}^* \rangle = \delta_{\ell\ell'} \delta_{mm'} C_\ell. \quad (1.44)$$

If the perturbations ΔT are Gaussian, then all their statistical information is encapsulated in the multipole coefficients C_ℓ . The uncertainty on the C_ℓ is given by:

$$\frac{\Delta C_\ell}{C_\ell} = \sqrt{\frac{2}{2\ell + 1}}. \quad (1.45)$$

This fractional difference is known as *cosmic variance*. Fortunately it decreases with the multipole moment ℓ^* . However for small ℓ (i.e. on large scales) cosmic variance limits the accuracy with which the C_ℓ can be measured.

Figures 1.3 to 1.5 show the power spectra of the CMB anisotropies as measured by the Planck satellite [62] together with the corresponding best-fit Λ CDM model using the parameter values from table 1.1. All CMB spectra are shown with the commonly used rescaling

$$D_\ell^{XX} \equiv \frac{1}{2\pi} \ell(\ell + 1) C_\ell^{XX}, \quad (1.46)$$

where $X \in \{T, E, B\}$ is a placeholder for temperature, or E - or B -mode polarisation respectively. E -mode polarisation arises from Thomson scattering in the (slightly) anisotropic primordial plasma. B -mode polarisation can result from primordial gravitational waves or from lensed E -modes leaking into B -modes.

*At first glance it might be surprising that the error bars for multipoles $\ell = 3$ and $\ell = 5$ are greater than the ones at $\ell = 2$ and $\ell = 4$ in fig. 1.3 even though according to eq. (1.45) the error bars should decrease for larger multipoles. However, note that the ΔC_ℓ are also proportional to the C_ℓ , which are relatively high for the multipoles 3 and 5 and relatively low for 2 and 4.

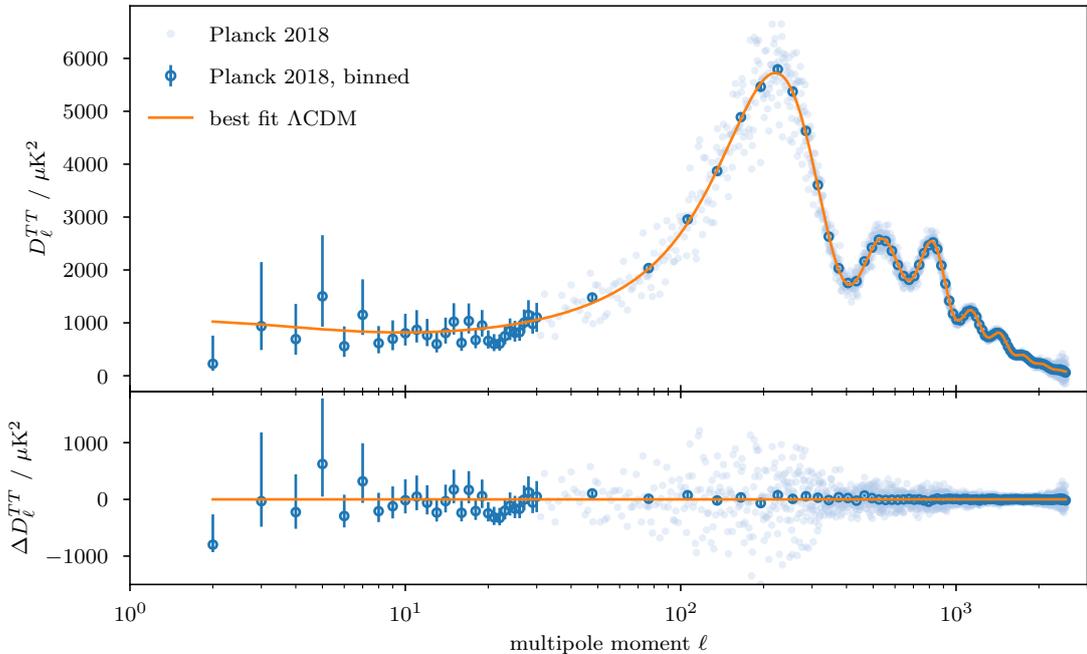


Figure 1.3 Temperature (TT) power spectrum of the CMB anisotropies. We show the data from the Planck 2018 legacy and the corresponding curve for the best-fit Λ CDM model. The lower panel shows the residuals with respect to the best-fit model.

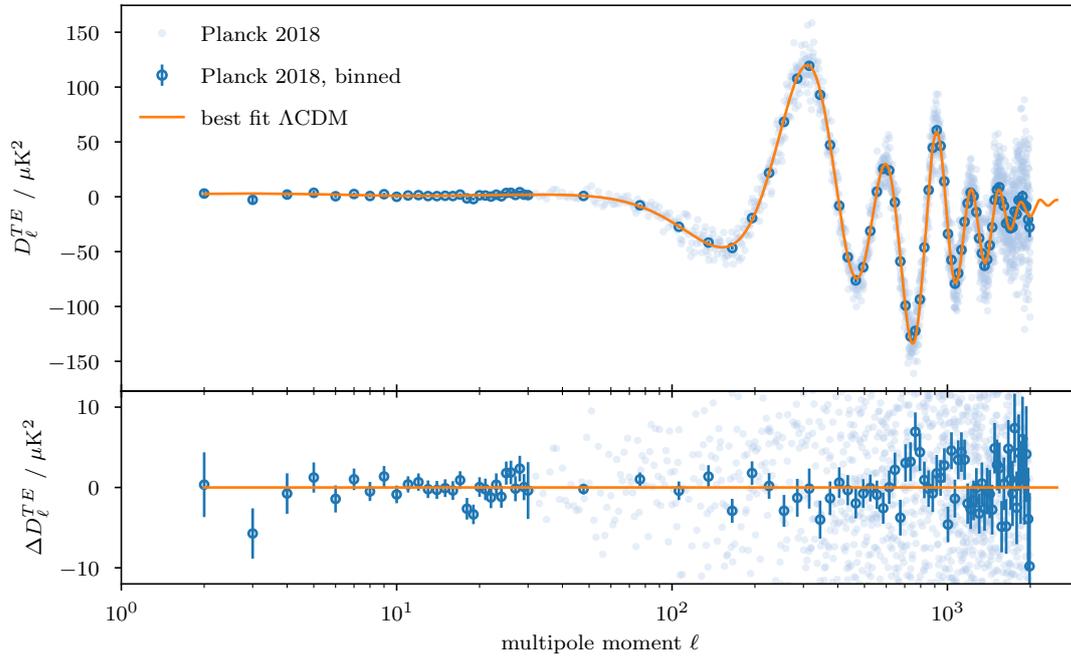


Figure 1.4 Temperature and polarisation (TE) cross-spectrum of the CMB anisotropies. We show the data from the Planck 2018 legacy and the corresponding curve for the best-fit Λ CDM model. The lower panel shows the residuals with respect to the best-fit model.

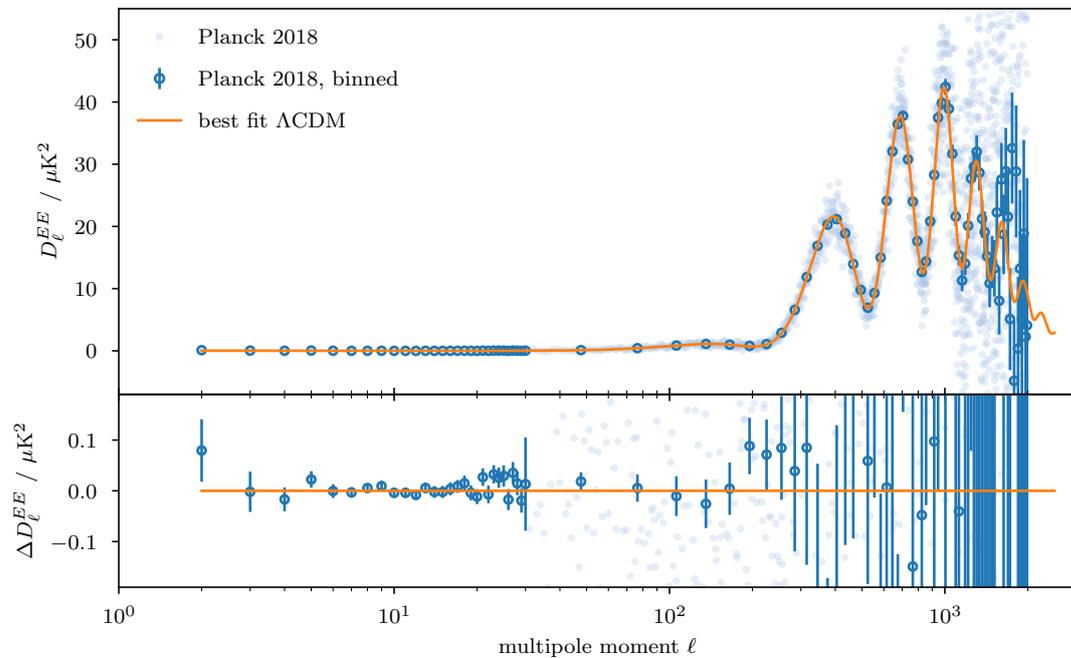


Figure 1.5 E-mode polarisation (EE) power spectrum of the CMB anisotropies. We show the data from the Planck 2018 legacy and the corresponding curve for the best-fit Λ CDM model. The lower panel shows the residuals with respect to the best-fit model.

The Planck likelihood function is derived using a hybrid approach, splitting the data into a low and a high multipole regime. The details differed for the three data releases and can be found in the corresponding likelihood papers [62–64]. Differences include e.g. the choice of low-to-high multipole threshold, of frequency channels, and of foreground masks. At low multipoles the distribution of the C_ℓ cannot be well approximated by a Gaussian due to the limited degrees of freedom per multipole ℓ . At the map level the CMB signal \mathbf{s} and instrumental noise \mathbf{n} are individually nearly Gaussian distributed and thus also the actually observed map $\mathbf{m} = \mathbf{s} + \mathbf{n}$ [63]. Under that assumption the low- ℓ likelihood function is given by:

$$\mathcal{L}(C_\ell) = \Pr(\mathbf{m}|C_\ell) = \frac{1}{2^{n/2}\pi^{n/2}|\mathbf{M}|^{1/2}} \exp\left(-\frac{1}{2}\mathbf{m}^\top \mathbf{M}^{-1}\mathbf{m}\right), \quad (1.47)$$

where n is the number of observed pixels and $\mathbf{M}(C_\ell) = \mathbf{S}(C_\ell) + \mathbf{N}$ is the data covariance matrix made up of the CMB and noise covariance matrices respectively (see eq. (21) in [63] or eq. (3) in [64]). For high multipoles there are sufficient degrees of freedom per multipole ℓ such that a Gaussian approximation for the distribution of the C_ℓ is possible:

$$-\ln \mathcal{L}(\hat{\mathbf{C}}|\mathbf{C}(\boldsymbol{\theta})) = \frac{1}{2} \left[\hat{\mathbf{C}} - \mathbf{C}(\boldsymbol{\theta}) \right]^\top \Sigma^{-1} \left[\hat{\mathbf{C}} - \mathbf{C}(\boldsymbol{\theta}) \right] + \text{const.}, \quad (1.48)$$

where $\hat{\mathbf{C}}$ is the data vector, $\mathbf{C}(\boldsymbol{\theta})$ is the model prediction for parameter values $\boldsymbol{\theta}$, and Σ is the covariance matrix (see eq. (13) in [64] or eq. (20) in [62]).

1.8 Conclusion

We have now presented the bigger picture and covered the necessary theoretical background on inflationary cosmology. The next chapter is concerned with prior considerations for half-constrained parameter extensions of Λ CDM. The following chapters discuss the effects of a finite amount of inflation (just enough to solve the horizon and flatness problem) on the primordial and CMB power spectrum. This is first motivated in [chapter 3](#) using phase-space diagrams for the inflaton field. In [chapters 4](#) and [5](#) a Bayesian analysis of such inflationary universes is performed for flat and curved cosmologies respectively.

Although all the chapters are self-contained and can in principle be read in any order, it is instructive to continue with [chapter 2](#), which starts out with a brief review of the Bayesian methods that we apply to CMB data throughout this thesis.

Chapter 2

Bayesian evidence for the tensor-to-scalar ratio r and neutrino masses m_ν : Effects of uniform vs logarithmic priors

In this chapter we review the effect that the choice of a uniform or logarithmic prior has on the Bayesian evidence and hence on Bayesian model comparisons when data provide only a one-sided bound on a parameter. We investigate two particular examples: the tensor-to-scalar ratio r of primordial perturbations and the mass of individual neutrinos m_ν , using the cosmic microwave background temperature and polarisation data from Planck 2018 and the NuFIT 5.0 data from neutrino oscillation experiments. We argue that the Kullback–Leibler divergence, also called the relative entropy, mathematically quantifies the Occam penalty. We further show how the Bayesian evidence stays invariant upon changing the lower prior bound of an upper constrained parameter. Using a logarithmic prior instead of a uniform prior frees the Λ CDM extensions from their additional Occam penalty bringing the Bayesian evidence up to par with the Λ CDM model without extension. For all prior options we get only a very slight preference for the normal over the inverted neutrino hierarchy with Bayesian odds of about 3 : 2.

2.1 Introduction

The “principle of insufficient reason” (Bernoulli [65]) or “principal of indifference” (renamed by Keynes [66]) states that in the event of multiple, mutually exclusive, possible outcomes and in the absence of any relevant evidence, we should assign the same probability to all outcomes [67]. In a Bayesian analysis, this is generalised to continuous parameters in the form of uninformative priors. Complete prior ignorance about a location parameter is represented by assigning a uniform distribution to the prior. Ignorance about a scale parameter on the other hand is represented by assigning a logarithmic prior, i.e. a uniform distribution on the logarithm of the parameter [67]. However, it is not always clear whether a parameter should be treated as a location or scale parameter. This is quite commonly discussed when faced with a strictly positive parameter such as a mass or an amplitude that is very small, yet still unconstrained. In general, the decision whether to use a uniform or logarithmic prior has

effects on credibility bounds and on the Bayesian evidence, i.e. on both levels of Bayesian inference: parameter estimation and model comparison. Under the reasoning that you can set the lower bound to zero and thus incorporate all possible small values, the uniform prior is often preferred, whereas the logarithmic prior is criticised for a lack of an unambiguous lower bound, and because the ultimate choice of the lower bound might affect a 95 % credibility bound and the Bayesian evidence.

In this paper we show that the very last statement is typically not true and that the choice of a lower bound for such a logarithmic prior is less problematic than commonly assumed. To that end we will look at two cosmological examples in particular: the tensor-to-scalar ratio r of primordial perturbations as well as the neutrino masses m_ν , where both uniform and logarithmic priors have been applied historically (for the tensor-to-scalar ratio, see e.g. [33, 68–72]; and for the neutrino masses see e.g. [33, 73–81]). Since there is no theoretical preference for one scale over another in either case, we believe that a logarithmic prior would indeed be better motivated from a Bayesian perspective. However, we concede that from a sampling perspective and regarding the description of posterior credibility bounds it can be simpler to work with uniform priors.

The best constraints on the tensor-to-scalar ratio $r_{0.05} \lesssim 0.06$ come from joint analyses of cosmic microwave background (CMB) data, CMB lensing, and baryon acoustic oscillations (BAO) [33, 71], where a uniform prior on r was adopted. A common goal of upcoming CMB experiments such as the Simons Observatory [45], the LiteBIRD satellite [46] and the next-generation “Stage-4” ground-based CMB experiment (CMB-S4) [47] is to push to a tensor-to-scalar ratio of $r \sim 10^{-3}$. In pushing to such small values of r , the question of whether to adopt a uniform or logarithmic prior in one’s analysis becomes more pertinent.

Since neutrino oscillation experiments measure non-zero mass differences, we can conclude that two or more neutrinos must have mass. However, the absolute scale of the individual neutrino masses m_i cannot be measured by the oscillation experiments, but only the mass-squared splittings $\Delta m_{ij}^2 = m_i^2 - m_j^2$. The strongest bound on the absolute neutrino mass scales is currently provided again by combined CMB and BAO data, limiting the sum of the neutrino masses to $\sum m_\nu \lesssim 0.12 \text{ eV}$ at 95 % confidence [80] (see also [78, 81] for other recent analyses).

When investigating the three discrete neutrino mass eigenstates, the question of uniform vs logarithmic priors arises again. Note, however, that given the known mass splittings from oscillation experiments, the three neutrino mass scales are linked. If one mass scale is known, then the others can be inferred from the mass squared splittings. Hence, only one mass scale is truly unknown and assuming scale invariant (i.e. logarithmic) priors on all three neutrino masses simultaneously would unduly favour smaller neutrino masses and thus a normal neutrino hierarchy (NH) with $m_1 < m_2 \ll m_3$ compared to an inverted neutrino hierarchy (IH) with $m_3 \ll m_1 < m_2$ (for more on this see also discussions in [73, 74]).

This paper is structured as follows: In [section 2.2](#) we will start by giving a brief description of our Bayesian analysis framework, including the data and base cosmological model used, as well as the means of computing the Bayesian evidence. In [section 2.3](#) we apply this to the

tensor-to-scalar ratio r and compare to a theoretical mock example. In [section 2.4](#) we perform the equivalent analysis for the neutrino masses and contrast the results for the two neutrino hierarchies. We conclude in [section 2.5](#).

2.2 Methods

2.2.1 Bayesian inference

There are two levels to Bayesian inference: parameter estimation and model comparison (see e.g. [\[67, 82\]](#)). Both these levels are based on Bayes' theorem which relates inference inputs (likelihood and prior) to yielded outputs (posterior and evidence):

$$\begin{aligned} \Pr(\theta|D, M) \times \Pr(D|M) &= \Pr(D|\theta, M) \times \Pr(\theta|M), \\ \text{Posterior} \times \text{Evidence} &= \text{Likelihood} \times \text{Prior}, \\ \mathcal{P}_M(\theta) \times \mathcal{Z}_M &= \mathcal{L}_M(\theta) \times \pi_M(\theta). \end{aligned} \quad (2.1)$$

The posterior \mathcal{P} is the main quantity of interest in a parameter estimation, representing our state of knowledge about the parameters θ in a given model M , inferred from our prior information π and the likelihood \mathcal{L} of the parameters under the data D . The evidence \mathcal{Z} is pivotal for model comparisons.

Were we interested only in parameter estimation, then it would be sufficient to care only about the proportionality of the posterior to the product of likelihood and prior and the Bayesian evidence could be neglected as a mere normalisation factor. However, for the comparison of say two models A and B the evidence becomes important with the posterior odds ratio of the two models given by:

$$\frac{\Pr(B|D)}{\Pr(A|D)} = \frac{\Pr(B)}{\Pr(A)} \times \frac{\mathcal{Z}_B}{\mathcal{Z}_A}. \quad (2.2)$$

Typically models are assigned the same prior preference such that the first term on the right-hand side becomes unity, leaving simply the evidence ratio $\mathcal{Z}_B/\mathcal{Z}_A$, which can be interpreted as betting odds for the two models. We typically quote this in terms of the log-difference of evidences between two models $\Delta \ln \mathcal{Z} = \ln(\mathcal{Z}_B/\mathcal{Z}_A)$.

The evidence is the marginal likelihood

$$\mathcal{Z}_M = \int \mathcal{L}_M(\theta) \pi_M(\theta) d\theta = \langle \mathcal{L}_M \rangle_\pi, \quad (2.3)$$

and can be numerically approximated with Laplace's method [\[83\]](#), estimated from a posterior distribution attained e.g. from a Monte Carlo Markov Chain (MCMC) via the Savage–Dickey density ratio (SDDR) [\[84–87\]](#) or via a nearest-neighbour approach [\[88, 89\]](#) or computed more directly with nested sampling [\[90–96\]](#).

If the posterior distribution and the evidence have both been determined, then as a byproduct one can also compute the Kullback–Leibler (KL) divergence, also called the relative

entropy:

$$\mathcal{D}_{\text{KL},M} = \int \mathcal{P}_M(\theta) \ln \left(\frac{\mathcal{P}_M(\theta)}{\pi_M(\theta)} \right) d\theta = \left\langle \ln \frac{\mathcal{P}_M}{\pi_M} \right\rangle_{\mathcal{P}}, \quad (2.4)$$

which quantifies the overall compression from prior to posterior distribution.

2.2.2 Kullback–Leibler divergence and Occam’s razor

It should be noted that the Bayesian evidence naturally incorporates the so-called Occam’s razor that penalises models for unnecessary complexity. It can be formulated as the principle to “Accept the simplest explanation that fits the data” [82]. This can be neatly demonstrated with the example of a model A with a uniform prior and Gaussian likelihood (see e.g. [82, 97]). Assuming a uniform prior in the range $a \in [a_{\min}, a_{\max}]$ and a Gaussian likelihood on a with mean a_0 and standard deviation σ_a , the Bayesian evidence decomposes into two terms:

$$\mathcal{Z}_A = \mathcal{L}_A(a_0) \times \frac{\sigma_a \sqrt{2\pi}}{a_{\max} - a_{\min}} \quad (2.5)$$

The first term on the right-hand side is the maximum likelihood point of model A . With additional parameters, this term would only increase and therefore can only favour model A . The second term incorporates the ratio of posterior to prior uncertainty of model A . Since the posterior uncertainty σ_a is generally smaller than the prior uncertainty $(a_{\max} - a_{\min})$, this term penalises model A for each of its parameters and thus embodies its Occam penalty. Note that the posterior and prior uncertainties appear *inversely* in the normalisation factor of the actual distributions.

More generally, the KL-divergence can actually be used as an estimator of the Occam penalty, which becomes clearer when rewriting the log-evidence according to*:

$$\begin{aligned} \ln \left(\int \mathcal{L}_M \pi_M d\theta \right) &= \int \mathcal{P}_M \ln \mathcal{L}_M d\theta - \int \mathcal{P}_M \ln \left(\frac{\mathcal{P}_M}{\pi_M} \right) d\theta, \\ (\text{log-})\text{evidence} &= \text{parameter fit} - \text{Occam penalty}, \\ \ln \mathcal{Z}_M &= \langle \ln \mathcal{L}_M \rangle_{\mathcal{P}} - \mathcal{D}_{\text{KL},M}, \end{aligned} \quad (2.6)$$

where we have dropped the dependence on θ to save space. Analogous to the example from eq. (2.5), the first term on the right-hand side encapsulates the fit of the model, while the KL-divergence is the average log-ratio of posterior to prior distribution (see also the last equality in eq. (2.4)), thus identifying it as the Occam penalty. This has been intuitively applied e.g. in the third figure in [27], but as far as the authors are aware eq. (2.6) is the first time this has been stated analytically in this form.

While known to experts, a widely unappreciated fact is that the evidence stays unaffected by an *unconstrained* parameter, i.e. when the data provide no information for that parameter [98]. In terms of eq. (2.6) this is reflected in an invariant likelihood $\langle \ln \mathcal{L} \rangle_{\mathcal{P}}$ and a zero KL-divergence

*Note that proving eq. (2.6) becomes surprisingly straight-forward when going from right to left and making use of Bayes’ theorem (2.1).

$\mathcal{D}_{\text{KL}}(\pi = \mathcal{P}) = 0$. Using the alternative labelling we can rephrase this: Adding an unconstrained parameter does not affect the fit, but also does not incur an additional Occam penalty and hence also leaves the evidence unaffected.

A popular measure for an effective number of constrained parameters is the Bayesian model complexity [99]. However, this quantity relies on the use of a point estimator such as the posterior mean or mode, which is why we prefer using the Bayesian model dimensionality d in the following sections (see [100] for a more detailed discussion on Bayesian complexities/dimensionalities). The Bayesian model dimensionality can be computed straightforwardly from the posterior distribution as the posterior *variance* of the log-likelihood:

$$\frac{d_M}{2} = \int \mathcal{P}_M(\theta) \left(\ln \frac{\mathcal{P}_M(\theta)}{\pi_M(\theta)} - \mathcal{D}_{\text{KL},M} \right)^2 d\theta \quad (2.7)$$

$$= \left\langle (\ln \mathcal{L}_M)^2 \right\rangle_{\mathcal{P}} - \left\langle \ln \mathcal{L}_M \right\rangle_{\mathcal{P}}^2. \quad (2.8)$$

Note the connection to eq. (2.6), where we used the posterior *average* of the log-likelihood. As such, these two quantities provide an interesting additional perspective to that of the $(\ln \mathcal{Z}, \mathcal{D}_{\text{KL}})$ pair. The posterior average of the log-likelihood informs us about the parameter fit and the posterior variance of the log-likelihood measures the models' complexity in the form of the number of constrained parameters.

2.2.3 Cosmological models

In the following sections we perform Bayesian model comparisons on one-parameter extensions to the Λ CDM model (universe dominated today by a cosmological constant Λ and by cold dark matter), which we parametrise with the standard 6 cosmological parameters listed in table 2.1 with their corresponding prior ranges.

In section 2.3 we extend the Λ CDM model by the tensor-to-scalar ratio r of primordial perturbations, which is set to $r = 0$ in Λ CDM. In section 2.4 we extend the base model by allowing for three distinct neutrino masses. In the Λ CDM model these are typically fixed to two massless neutrinos and a single massive neutrino with $m_\nu = 0.06$ eV.

2.2.4 Data

We use the 2018 temperature and polarisation data from the Planck satellite [62], which we abbreviate as “TT,TE,EE+lowE”. Note that this is the same abbreviation as in the corresponding Planck publication itself. The specific use of “lowE” but lack of “lowT” might lead to the conclusion that only E-mode and no temperature data were used at low multipoles. However, this is *not* the case. Both high- ℓ and low- ℓ temperature auto-correlation data are implied in that abbreviation.

In section 2.4 we additionally use the NuFIT 5.0 (2020) data from neutrino oscillation experiments [42–44] to set Gaussian priors on the mass squared splittings δm^2 and Δ_m^2 .

Table 2.1 Cosmological parameters of the base Λ CDM cosmology the way they are sampled in our Bayesian analysis. The second column shows their corresponding prior ranges. The third column lists their mean and 68 % limits from our base Λ CDM nested sampling run with TT,TE,EE+lowE data from Planck 2018 and is in almost perfect agreement with table 2 in [33].

Parameter	Prior range	68 % limits
$\omega_b = h^2\Omega_b$	[0.019, 0.025]	0.02236 ± 0.00015
$\omega_c = h^2\Omega_c$	[0.025, 0.471]	0.1199 ± 0.0014
$100\theta_s$	[1.03, 1.05]	1.04191 ± 0.00029
τ_{reio}	[0.01, 0.40]	$0.0540^{+0.0073}_{-0.0084}$
$\ln(10^{10}A_s)$	[2.5, 3.7]	$3.043^{+0.015}_{-0.016}$
n_s	[0.885, 1.040]	0.9641 ± 0.0042

2.2.5 Statistical and cosmological software

We explore the posterior distributions of cosmological and nuisance parameters using `Cobaya` [101], which provides both the MCMC sampler developed for `CosmoMC` [102, 103] with a “fast dragging” procedure described in [104] and also the nested sampling code `PolyChord` [95, 96], tailored for high-dimensional parameter spaces, which can simultaneously determine the Bayesian evidence. Both samplers are interfaced with the cosmological Boltzmann code `CLASS` [105–107], which computes the theoretical CMB power spectra for temperature and polarisation modes.

`PolyChord` samples the entire prior volume and thereby allows us to infer both the posterior distribution and the Bayesian evidence. This is done with multiple (nested) iterations i of uniformly sampling a volume X_i with a sample number n_{live} of so-called “live points”. This gives the evidence according to:

$$\mathcal{Z} = \sum_i (X_{i-1} - X_i) \mathcal{L}_i. \quad (2.9)$$

Since the volumes X_i are not known exactly, we can only make inferences about the evidence in terms of a probability distribution $\text{Pr}(\mathcal{Z})$. Different to most other ways of determining the Bayesian evidence, eq. (2.9) also allows us to estimate the sampling error on the evidence. The evidence is approximately log-normally distributed, so in practice it is best to report the mean and variance of $\log \mathcal{Z}$. The size of the sampling error scales with the number of live points according to:

$$\Delta \log \mathcal{Z} \sim \mathcal{O}\left(n_{\text{live}}^{-1/2}\right). \quad (2.10)$$

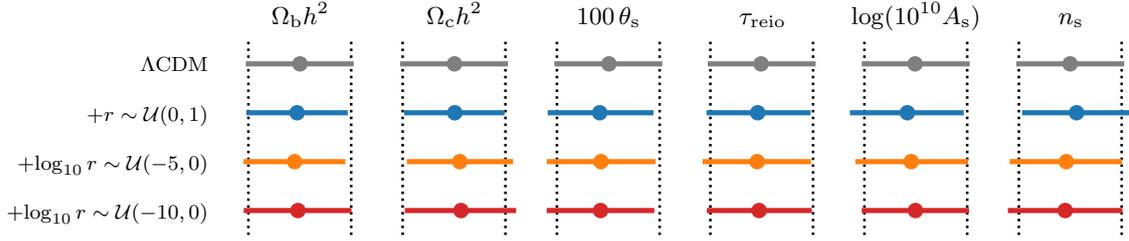


Figure 2.1 Stability of the cosmological parameters for the tensor-to-scalar ratio extension of the base Λ CDM cosmology with different priors on r : uniform in blue, logarithmic with lower bound -5 in orange and logarithmic with lower bound -10 in red. For each parameter we show the mean and the extent from quantile 0.16 to 0.84, i.e. the inner 68% limits.

The run-time depends linearly on the number of live points, but parallelization with n_{procs} processes gives a speed-up of:

$$\text{Speedup}(n_{\text{procs}}) \sim n_{\text{live}} \log\left[1 + \frac{n_{\text{procs}}}{n_{\text{live}}}\right]. \quad (2.11)$$

A rule of thumb for the number of live points to use in order to get meaningful results on evidence estimates is roughly $n_{\text{live}} \sim 25d$, where d is the dimension of the sampling space, i.e. the number of sampling parameters. Thus for a cosmology run we will typically require about 500 to 1000 live points. For more details on the nested sampling calculation of the Bayesian evidence and its errors we refer the reader to [95].

We use `GetDist` [108] to generate the data tables of marginalised parameter values. The post-processing of the nested sampling output for the computation of Bayesian evidence, KL-divergence and Bayesian model dimensionality, as well as the plotting functionality for posterior contours is performed using the python module `anesthetic` [109].

2.3 Tensor-to-scalar ratio

The tensor-to-scalar ratio r quantifies what fraction of primordial perturbations is in the form of gravitational waves, produced e.g. during cosmic inflation and potentially detectable in their contribution to CMB B-modes.

So far, the major experiments probing the contribution of tensor modes to the CMB power spectrum have adopted a uniform prior on r [33, 71]. However, the common target of $r \sim 10^{-3}$ for many upcoming CMB experiments such as the Simons Observatory, the LiteBIRD satellite or CMB-S4, warrants the question as to whether a scale invariant prior might be better to handle such low values. This question frequently brings up arguments of the ambiguity of the lower bound to a logarithmic prior and its potential effect on the Bayesian evidence.

2.3.1 Tensor-to-scalar ratio: Posteriors

Figure 2.1 gives an overview of the stability of the cosmological base parameters across different priors for r and compares them to the Λ CDM base model by showing their mean and 68%

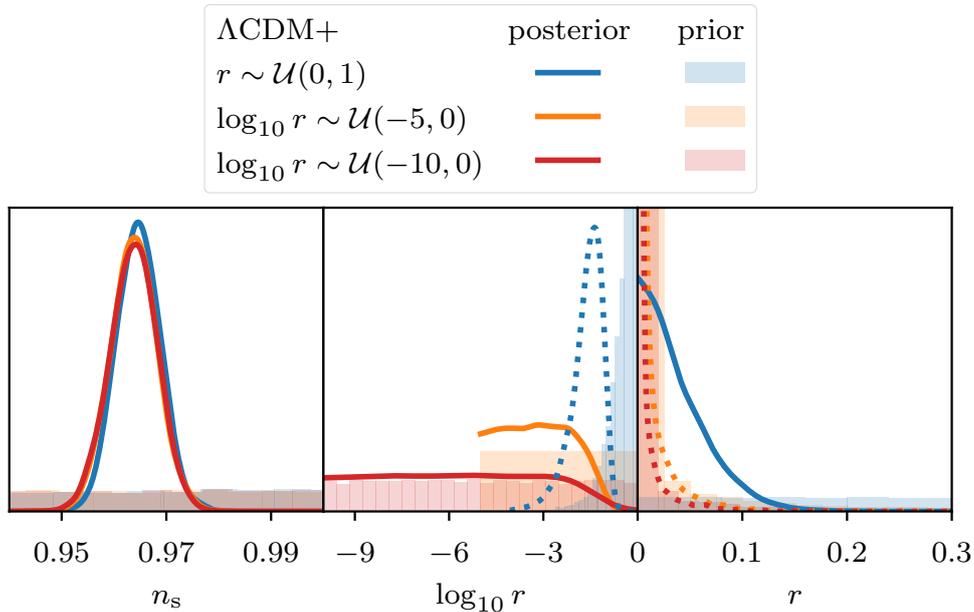


Figure 2.2 Normalised one-dimensional posterior distributions for Planck 2018 TT,TE,EE+lowE data for the spectral index n_s and the tensor-to-scalar ratio r of primordial perturbations, contrasting the difference between using a uniform (blue) or logarithmic (orange and red) prior on r . The shaded histograms illustrate the prior distributions. Note that the dotted lines show the inferred parameters r and $\log_{10} r$ in the respective opposite domain. This is done only to provide a more direct visual comparison. However, these dotted contours are not data-driven parameter constraints. In particular the blue dotted line results purely from a lack of small prior samples when sampling uniformly over r , and does not in fact constitute a lower bound on the tensor-to-scalar ratio.

ranges. In addition to the Λ CDM base run, we have taken nested sampling runs with both a uniform prior on the tensor-to-scalar ratio $r \sim \mathcal{U}(0, 1)$ and with two logarithmic priors with different lower bounds, $\log_{10} r \sim \mathcal{U}(-5, 0)$ and $\log_{10} r \sim \mathcal{U}(-10, 0)$. The near perfect alignment across different setups reflects how little the tensor-to-scalar ratio correlates with the other parameters.

In [fig. 2.2](#) we focus on the spectral index n_s and the tensor-to-scalar ratio r (or $\log_{10} r$) in particular by showing their one-dimensional marginalised posterior distributions. [Figure 2.3](#) shows the corresponding two-dimensional joint probability contours of the 68% and 95% levels for n_s and r (or $\log_{10} r$). We have included shaded histograms in the 1d plots and scatter points in the 2d plots to give a notion of the prior distributions.

As already expected from [fig. 2.1](#), the marginalised posterior for the spectral index is near identical, irrespective of the prior on r . The tensor-to-scalar ratio in the right panel of [fig. 2.2](#) drops off exponentially from $r = 0$ to larger values, thereby significantly compressing the prior, which spans up to unity. When sampling logarithmically the posterior levels off towards small scales and shows a step-like behaviour at the upper bound.

We have included the kernel density estimate from the uniform r -samples in the $\log_{10} r$ plot and vice versa (dotted lines). This allows us to compare more directly what sort of numerical values were actually used in those two cases. At a first naive glance one might be

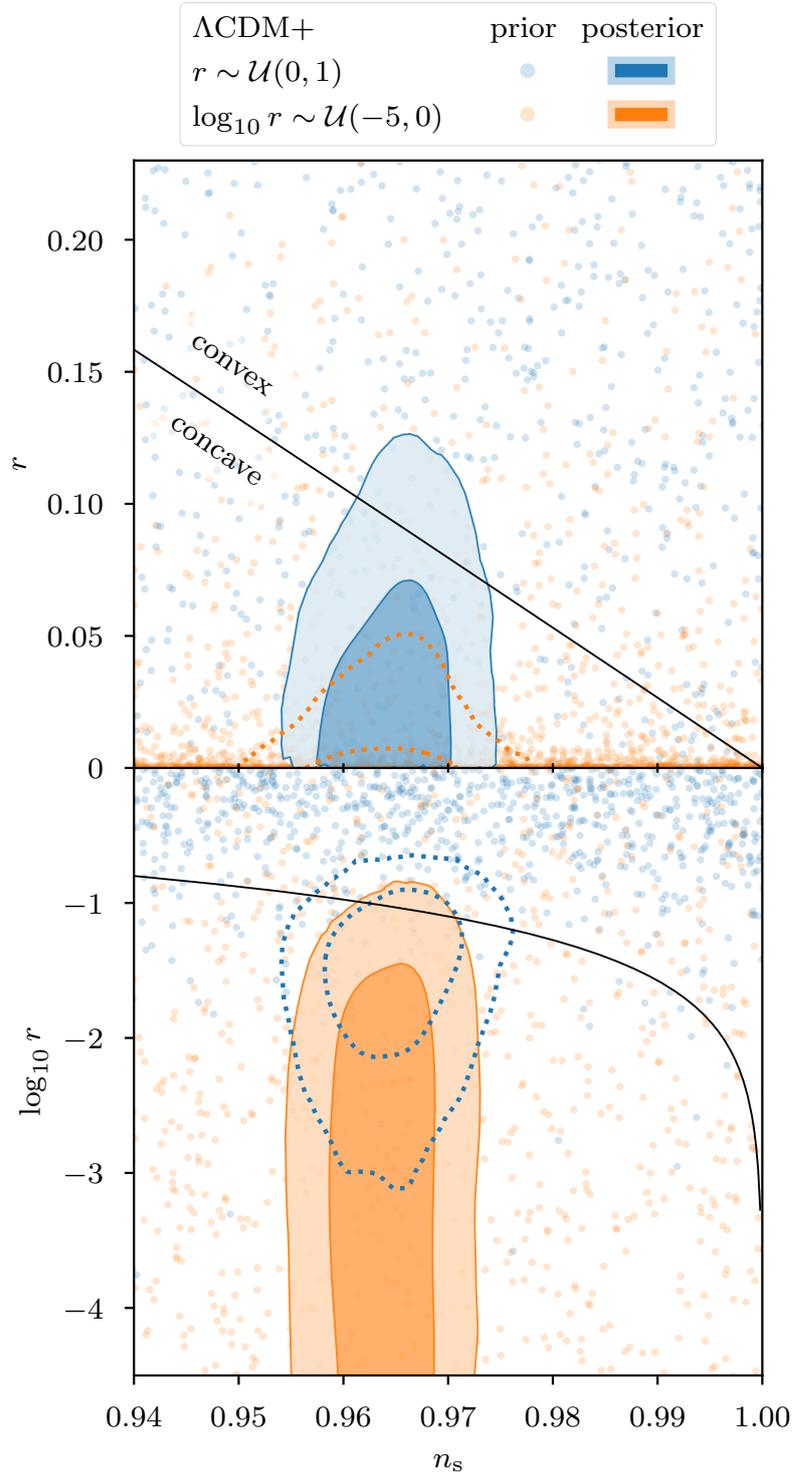


Figure 2.3 Two-dimensional version of [fig. 2.2](#) showing the 68% and 95% levels of the posterior contours for Planck 2018 TT,TE,EE+lowE data for spectral index n_s and tensor-to-scalar ratio r , where again a uniform (blue) or logarithmic (orange) prior on r was used. The scattered dots give a notion of that prior distribution. Note that the dotted lines are not true constraints as explained in [fig. 2.2](#). The thin black line divides the n_s - r parameter space into regions of convex and concave inflationary potentials.

concerned that the dotted blue line actually indicates a lower bound, however, looking at the blue shaded histogram in the 1d plot or the blue scatter points in the 2d plot it becomes clear that this is entirely prior driven and reflects that uniform sampling of r does not reach such low values (see also [110] on a related discussion about the importance of adjusting the density when setting the x -scale to ‘log’). With a target of $r \sim 10^{-3}$ this highlights how the parameter space is sampled rather inefficiently at those low values of interest when applying a uniform prior, which would be an argument for adopting a logarithmic prior in the future.

One problem to be aware of with the unconstrained posteriors from a logarithmic prior is that upper bounds in form of e.g. 95% limits will change with the lower bound on the logarithmic parameter: the smaller the lower prior bound, the smaller also the upper posterior bound. This lack of a stable posterior bound is a result of the definition via percentiles, a notion inspired by a normal distribution. For other types of distributions, such as the step-like posteriors seen in the middle panel of fig. 2.2, percentiles of that sort are not the ideal measure for an upper bound. For such a step-like posterior a better alternative would be to quantify the position of the step directly, e.g. where the posterior drops to some fraction of its plateau value. In the case that an exponential distribution provides a good fit to the non-logarithmic parameter (see the mock example in the following section 2.3.3), the parameter value where the posterior is $1/e$ times its maximum turns out to be a stable choice, which corresponds to the mean of the exponential distribution. Indeed, using that $1/e$ measure for the step position we get roughly the same upper bound on the tensor-to-scalar ratio for all prior options:

$$\begin{aligned} r &< 0.06, \\ \log_{10} r &< -1.22. \end{aligned} \tag{2.12}$$

Note, that these are *not* the habitually quoted 95% upper bounds on the tensor-to-scalar ratio. For the uniform sampling run of r , this limit in this case is closer to roughly an 80% upper bound. Note further that the choice of the $1/e$ fraction provides a particularly stable bound, because of the connection to the mean of the exponential distribution.

2.3.2 Tensor-to-scalar ratio: Evidence and Kullback–Leibler divergence

Nested sampling provides us with distributions for log-evidence $\ln \mathcal{Z}$, KL-divergence \mathcal{D}_{KL} and Bayesian model dimensionality d in the same way as for the posterior of free model parameters, which can be calculated straightforwardly using `anesthetic`’s analysis tools for nested sampling output [109]. Figure 2.4 shows the contours for those quantities in a triangle plot. We have normalised all quantities with respect to the base Λ CDM model, such that e.g. for the log-evidence we have:

$$\Delta \ln \mathcal{Z} = \ln \mathcal{Z} - \ln \mathcal{Z}_{\Lambda\text{CDM}}. \tag{2.13}$$

Table 2.2 lists the summary statistics for the quantities from fig. 2.4.

The marginalised plot for the difference in log-evidence (topmost panel) with $\Delta \ln \mathcal{Z} = -2.95 \pm 0.25$ for the r -extension of Λ CDM shows that it is considerably disfavoured when

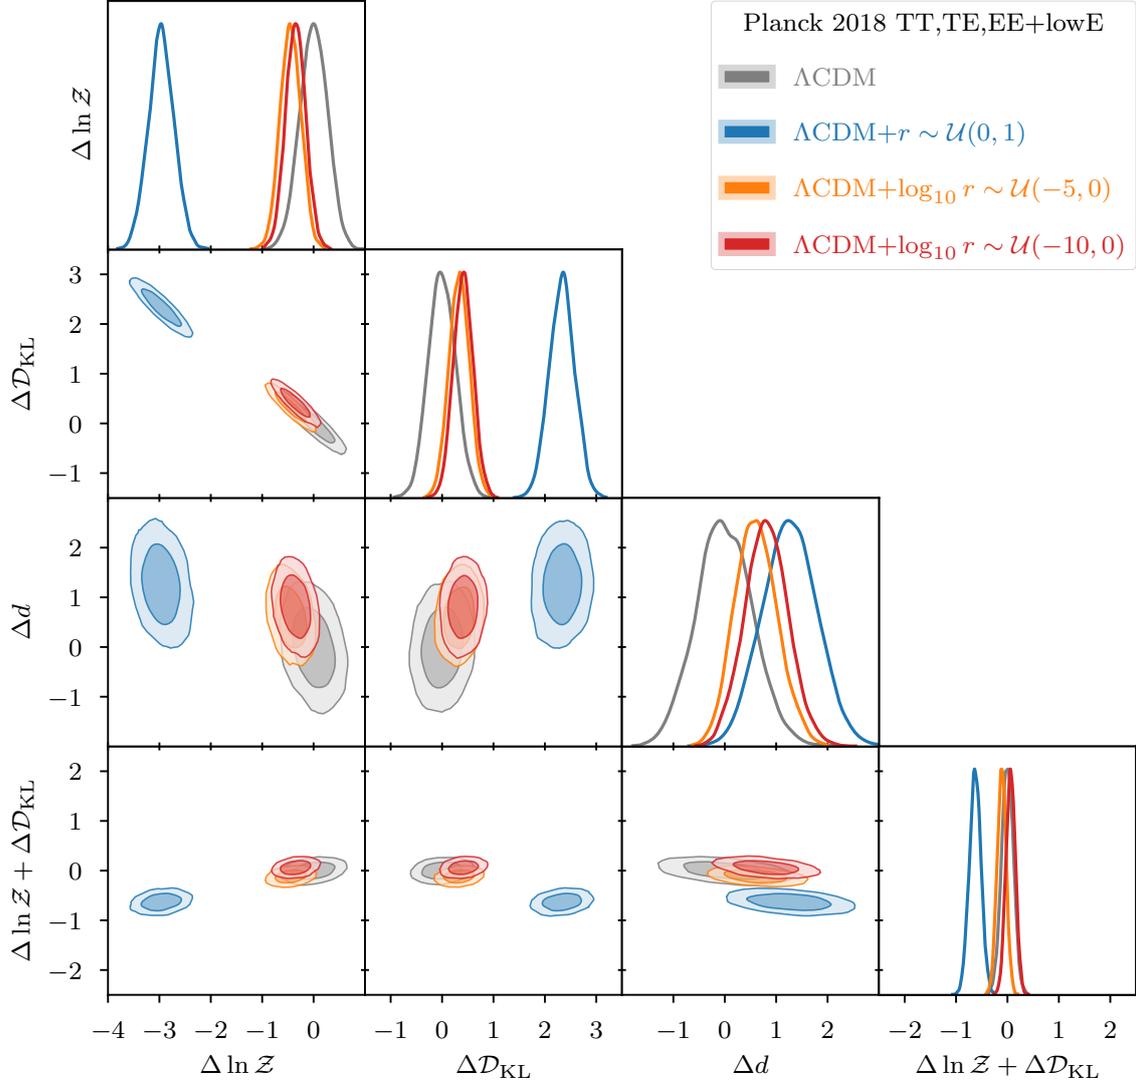


Figure 2.4 Effect of uniform vs logarithmic priors on Bayesian model comparison for the tensor-to-scalar ratio r : log-evidence $\Delta \ln \mathcal{Z}$, Kullback–Leibler divergence \mathcal{D}_{KL} (in nats), Bayesian model dimensionality d , and posterior average of the log-likelihood $\langle \ln \mathcal{L} \rangle_{\mathcal{P}} = \ln \mathcal{Z} + \mathcal{D}_{\text{KL}}$. The probability distributions represent errors arising from the nested sampling process. In the limit of infinite life points these distributions would become point statistics, in contrast to posterior distributions. We normalise with respect to the Λ CDM model without r (i.e. with $r = 0$). Note, how switching from uniform to logarithmic sampling of r (i.e. from blue to orange/red) moves the contours along their $\ln \mathcal{Z}$, \mathcal{D}_{KL} degeneracy line, i.e. relative entropy is traded in for evidence. Note further by comparison of the orange and red lines, how changing the lower bound of the logarithmic sampling interval (by 5 log-units) barely affects the contours.

Table 2.2 Mean and standard deviation of the log-evidence $\ln \mathcal{Z}$, Kullback–Leibler divergence \mathcal{D}_{KL} and Bayesian model dimensionality d of the base Λ CDM cosmology and its r extension from Planck 2018 TT,TE,EE+lowE data [62]. The Δ indicates normalisation with respect to the base Λ CDM model.

Model	betting odds	$\Delta \ln \mathcal{Z}$	$\Delta \mathcal{D}_{\text{KL}}$	Δd
Λ CDM	1	-0.00 ± 0.27	0.00 ± 0.26	0.00 ± 0.56
Λ CDM+ $r \sim \mathcal{U}(0, 1)$	1 : 20	-2.95 ± 0.25	2.33 ± 0.25	1.25 ± 0.53
Λ CDM+ $\log_{10} r \sim \mathcal{U}(-5, 0)$	5 : 8	-0.45 ± 0.20	0.33 ± 0.20	0.61 ± 0.42
Λ CDM+ $\log_{10} r \sim \mathcal{U}(-10, 0)$	5 : 7	-0.34 ± 0.20	0.41 ± 0.20	0.81 ± 0.43

applying a uniform prior. However, switching from a uniform to a logarithmic prior negates the difference in log-evidence completely, such that the $\log r$ extension ends up almost on par with the base Λ CDM model.

Changing the lower bound for the logarithmic prior, on the other hand, barely affects the evidence value at all. We have performed a run with a lower bound of $\log_{10} r = -5$ and another with $\log_{10} r = -10$, i.e. five orders of magnitude difference in r , and the $\ln \mathcal{Z}$ distributions are almost identical and, hence, their means well within one standard deviation of one another. As explained in [section 2.2.1](#), this is due to $\log r$ being unconstrained below a certain threshold and the Bayesian evidence picking up only on *constrained* parameters. This can seem counter-intuitive, since the Bayesian evidence is generally understood to automatically penalise additional parameters. The key point is that the Occam penalty essentially enters into the Bayesian evidence in the form of the ratio of posterior to prior volume. If both volumes are the same, then they divide out and do not contribute to the Occam penalty.

The last point becomes clearer by also taking into account the KL-divergence and recalling [eq. \(2.6\)](#), where we identified \mathcal{D}_{KL} as a measure for the Occam penalty. Looking at the correlation plot between log-evidence and KL-divergence makes it clear that there is a trade-off happening between those two quantities when switching between uniform and logarithmic priors. While the evidence increases for the logarithmic prior, the KL-divergence decreases, as expected from the posterior plots in [fig. 2.2](#), which shows how the change from prior to posterior happens only at about $\log_{10} r \gtrsim -2$. This is further reflected in the Bayesian model dimensionality d , which shows a clear growing trend from about $d = 16$ for the base Λ CDM model via a $\log r$ extension to about $d = 17$ for the r extension reflecting the one additional sampling parameter. Note that the total number of sampled parameters consists of 6 base cosmological parameters (+1 for the r extension) and 21 nuisance parameters from the Planck likelihood.

Because of the trade-off between log-evidence and KL-divergence it is interesting also to look at their sum, which from [eq. \(2.6\)](#) we know turns out to be the posterior *average* of the log-likelihood:

$$\ln \mathcal{Z} + \mathcal{D}_{\text{KL}} = \langle \ln \mathcal{L} \rangle_{\mathcal{P}}. \quad (2.14)$$

This makes for an interesting pairing with the Bayesian model dimensionality, since $d/2$ is the posterior *variance* of the log-likelihood. As such, these two quantities provide an alternative perspective to that of the evidence and KL-divergence. The posterior average and variance of the log-likelihood are a measure of the fit and complexity respectively. $\langle \ln \mathcal{L} \rangle_{\mathcal{P}}$ is shown in the last panel in [fig. 2.4](#), where we indeed see that the line for uniform sampling of r has moved much closer to the other lines, which is to be expected, since r and $\log r$ are fundamentally the same parameter and therefore lead to a similar goodness of fit.

This behaviour can also be understood analytically, which we explore in the following section in a one-dimensional mock example, simulating the r vs $\log r$ result.

2.3.3 Mock example

To illustrate further the role of a uniform vs a logarithmic prior on a Bayesian model comparison, we propose the following mock example, which is loosely based on the pedagogical example by Sivia and Skilling [97] explaining the effect of an additional (although in that case *constrained*) parameter, which we already outlined in [section 2.2.1](#).

Here, we will not assume a Gaussian likelihood that ultimately fully constrains a parameter, but rather we will assume an exponential distribution as our likelihood on a strictly positive parameter (which is the maximum-entropy distribution when only a mean is known):

$$\mathcal{L}(a) = P_0 e^{-a/\mu}, \quad (2.15)$$

where $P_0 = \Pr(D | a = 0)$ is the maximum likelihood value for the data D at $a = 0$ and where μ is the mean of the likelihood distribution describing the data. Thus, the likelihood is constrained only on one side, providing an upper bound, as shown in the left panel of [fig. 2.5](#).

We will assume a model A , where we sample the parameter a uniformly in the interval $[a_1, a_2]$. Furthermore, we will assume a model B , where we uniformly sample the parameter $b = \log_{10} a$ in the interval $[b_1, b_2]$, corresponding to logarithmically sampling the parameter a . Since both models are fundamentally governed by the same quantity and will use the same likelihood, any difference in Bayesian inference quantities will be purely prior driven.

We will make the following assumptions on the ordering of the prior limits:

$$0 = a_1 < 10^{b_1} \ll \mu \ll 10^{b_2} = a_2 = 1. \quad (2.16)$$

This ordering is motivated as follows: For the upper limit we require that the likelihood has essentially dropped to zero. Hence, without loss of generality, we can set the upper limit to one and require $\mu \ll 1$. The lower limit for the positive parameter a can be explicitly set to zero when sampling uniformly. However, when sampling logarithmically we need to pick some finite lower limit, which we require to be in the region $10^{b_1} \ll \mu$, where the likelihood has essentially saturated with respect to b (see right panel in [fig. 2.5](#)). The dependence of Bayesian quantities such as the evidence \mathcal{Z} or the Kullback–Leibler divergence \mathcal{D}_{KL} on the prior choice on the one hand and on this lower limit b_1 on the other is the goal of this mock example.

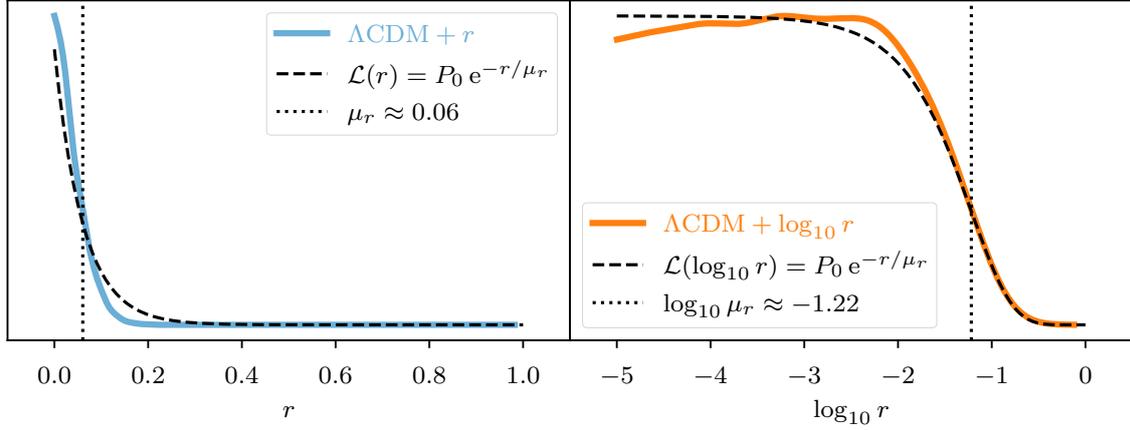


Figure 2.5 Exponential likelihood distribution from our mock example in eq. (2.15) compared to Planck 2018 temperature and polarisation data (TT,TE,EE+lowE) on the tensor-to-scalar ratio r with uniform sampling of r on the left and logarithmic sampling of r on the right. Note how the mean μ_r fulfills the ordering required by eq. (2.16) and how the lower limit on $\log_{10} r$ is well into the saturation plateau of posterior/likelihood.

The corresponding priors for models A and B can thus be written as:

$$\pi_A(a) = \frac{1}{a_2 - a_1} \Theta(a - a_1) \Theta(a_2 - a), \quad (2.17)$$

$$\pi_B(b) = \frac{1}{b_2 - b_1} \Theta(b - b_1) \Theta(b_2 - b), \quad (2.18)$$

where $\Theta(x)$ is the Heaviside step function.

We can compute the evidence and Kullback–Leibler divergence for models A and B as:

$$\mathcal{Z}_A = \int \mathcal{L}(a) \pi_A(a) da = \frac{P_0 \mu}{a_2 - a_1} \left(e^{-a_1/\mu} - e^{-a_2/\mu} \right), \quad (2.19)$$

$$\mathcal{Z}_B = \int \mathcal{L}(10^b) \pi_B(b) db = \frac{P_0}{b_2 - b_1} \frac{1}{\ln(10)} \left[\text{Ei} \left(-\frac{10^{b_2}}{\mu} \right) - \text{Ei} \left(-\frac{10^{b_1}}{\mu} \right) \right], \quad (2.20)$$

$$\begin{aligned} \mathcal{D}_{\text{KL},A} &= \int \frac{\mathcal{L}(a) \pi_A(a)}{\mathcal{Z}_A} \ln \left(\frac{\mathcal{L}(a)}{\mathcal{Z}_A} \right) da \\ &= \ln \frac{P_0}{\mathcal{Z}_A} - 1 - \frac{P_0}{\mathcal{Z}_A} \frac{1}{a_2 - a_1} \left[a_1 \exp \left(-\frac{a_1}{\mu} \right) - a_2 \exp \left(-\frac{a_2}{\mu} \right) \right], \end{aligned} \quad (2.21)$$

$$\begin{aligned} \mathcal{D}_{\text{KL},B} &= \int \frac{\mathcal{L}(10^b) \pi_B(b)}{\mathcal{Z}_B} \ln \left(\frac{\mathcal{L}(10^b)}{\mathcal{Z}_B} \right) db \\ &= \ln \frac{P_0}{\mathcal{Z}_B} - \frac{P_0}{\mathcal{Z}_B} \frac{1}{b_2 - b_1} \frac{1}{\ln(10)} \left[\exp \left(-\frac{10^{b_1}}{\mu} \right) - \exp \left(-\frac{10^{b_2}}{\mu} \right) \right], \end{aligned} \quad (2.22)$$

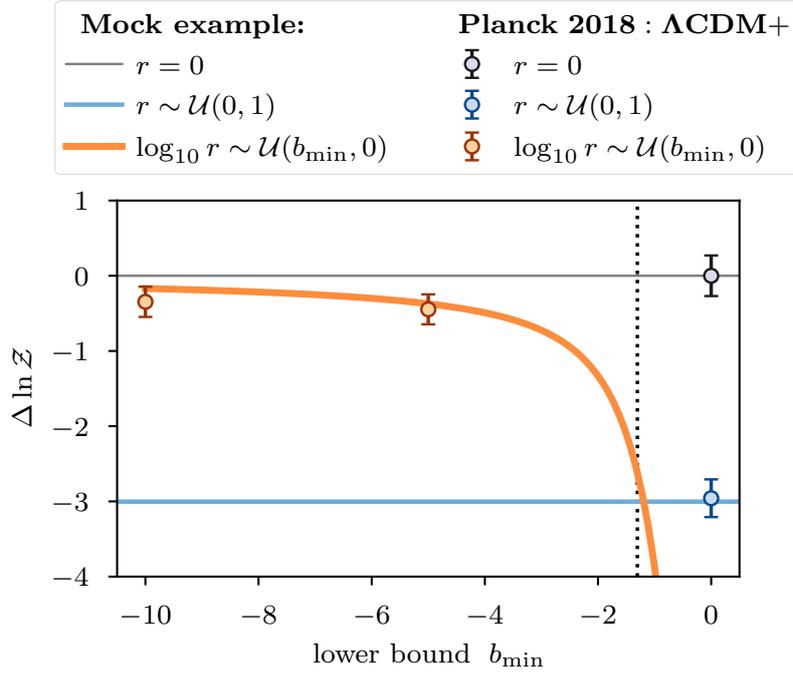


Figure 2.6 Dependence of the log-evidence on the lower prior bound b_{\min} : Comparison of the results in eqs. (2.19) and (2.20) for the one-dimensional mock example to the nested sampling results from table 2.2. The vertical dotted line corresponds to the mean used in the likelihood distribution (cf. fig. 2.5).

where Ei refers to the exponential integral. With the ordering from eq. (2.16) we can then approximate these to give:

$$\Delta \ln \mathcal{Z}_A \approx \ln \mu \quad \sim -3, \quad \Delta \mathcal{D}_{\text{KL},A} \approx -\Delta \ln \mathcal{Z}_A - 1 \quad \sim 2, \quad (2.23)$$

$$\Delta \ln \mathcal{Z}_B \approx \ln \left(1 - \frac{\log_{10} \mu}{b_1} \right) \quad \sim 0, \quad \Delta \mathcal{D}_{\text{KL},B} \approx -\Delta \ln \mathcal{Z}_B \quad \sim 0, \quad (2.24)$$

where we normalise with respect to a base model O with $a = 10^b = 0$ fixed, such that $\mathcal{Z}_O = P_0$ and $\mathcal{D}_{\text{KL},O} = 0$. The numerical values assume $\mu \sim 0.06$, which is roughly the posterior mean of the tensor-to-scalar ratio under uniform sampling in the preceding section. Hence, we can compare these zeroth-order numerical approximations to the results in fig. 2.4 and table 2.2, which indeed match.

Figure 2.6 makes this comparison more thoroughly, comparing the results from our one-dimensional mock example in eqs. (2.19) and (2.20) with the nested sampling results from table 2.2 for a variable lower bound b_{\min} of the logarithmic prior. The mean $\ln \mathcal{Z}$ of the base model with $r = 0$ for both the mock example and for the base Λ CDM nested sampling run are zero by definition of our normalisation. They serve only as calibration for the models with uniform (blue) and logarithmic (orange) priors. All three nested sampling runs agree well with the prediction from the mock example within their margins of errors.

Figure 2.6 illustrates how the evidence levels off with regards to the choice of the lower bound of the logarithmic prior (orange line) also reflected in the near equal evidences of the

nested sampling runs with lower prior bounds of -5 and -10 respectively. Note that the good agreement between mock example and data in [fig. 2.6](#) is due to the fact that the tensor-to-scalar ratio is almost completely uncorrelated with the other cosmological parameters, with the biggest (yet still small) correlation coming from the spectral index n_s (cf. [fig. 2.3](#)).

So far we have only discussed the effect of changing the *lower* prior bound, which is the main point of this paper. For completeness we should also address the effects of changing the upper prior bound. The results from using a logarithmic prior are not very sensitive to the upper prior bound as long as it is not pushed into the posterior bulk, i.e. not below roughly -1 in our mock example. For a uniform prior decreasing the upper prior bound will increase the Bayesian evidence (and vice versa). However, as long as the upper prior bound stays above the posterior bulk, the evidence would always remain below that belonging to the logarithmic prior.

2.4 Neutrino masses

In Planck’s baseline cosmology, the neutrinos are assumed to be comprised of two massless neutrinos and one massive neutrino with mass $m_\nu = 0.06$ eV with the effective number of neutrino species set slightly larger than 3 to $N_{\text{eff}} = 3.046$ [[33](#), [39](#), [40](#)]. [Figure 2.7](#) shows the dependence of the CMB power spectrum on the neutrino mass m_ν while adjusting the dark energy density Ω_Λ to keep to a flat geometry. The most visible effect is a shift of the acoustic peaks, due to the change in the epoch of matter-radiation equality and thus the distance to the last scattering surface. For heavier neutrinos the peaks are shifted to larger scales. In addition to the shift in peaks, heavier neutrinos slightly suppress power both on the largest scales and towards small scales. This becomes even more evident in the large scale structure of the Universe, which makes probes such as BAOs or CMB lensing complementary for the determination of neutrino masses (see e.g. [[111](#)]). Accordingly the best constraints on the sum of neutrino masses of about $\sum m_\nu < 0.12$ eV as mentioned earlier are provided by analyses combining CMB with CMB lensing and BAOs [[33](#)].

Upcoming CMB experiments such as the Simons Observatory, LiteBIRD or CMB-S4 and large scale structure (LSS) experiments such as Euclid will allow us to fully constrain the sum of neutrino masses $\sum m_\nu$. However, even under the most optimistic assumptions, it will not be possible to disentangle the individual contributions of the three neutrino flavours with cosmological data alone [[79](#)]. To achieve that, we need additional data from solar, atmospheric, reactor and accelerator experiments as summarised in NuFIT 5.0 (2020) [[43](#), [44](#)] that provide us with the mass square splittings:

$$\delta m^2 = 7.42^{+0.21}_{-0.20} \times 10^{-5} \text{ eV}^2 \quad (\text{NH \& IH}), \quad (2.25)$$

$$\Delta m^2 = \begin{cases} 2.517^{+0.026}_{-0.028} \times 10^{-3} \text{ eV}^2 & (\text{NH}), \\ 2.498^{+0.028}_{-0.028} \times 10^{-3} \text{ eV}^2 & (\text{IH}), \end{cases} \quad (2.26)$$

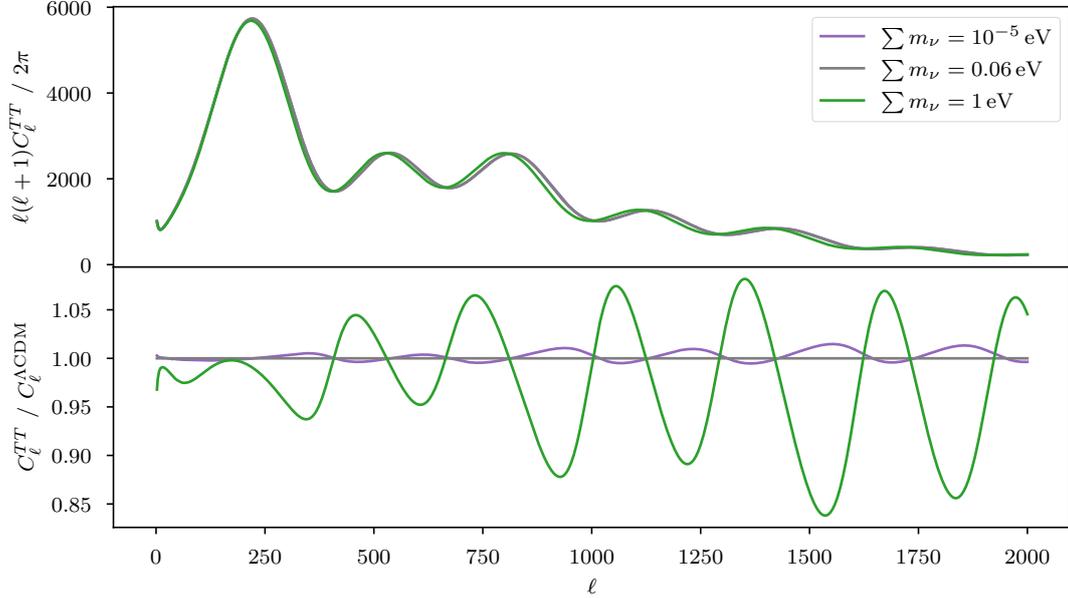


Figure 2.7 Dependence of the unlensed CMB temperature power spectrum on neutrino mass, where we assumed a single massive neutrino. A mass of 0.06 eV corresponds to the baseline used by the Planck collaboration [33].

where δm^2 is the smaller squared mass splitting between the light and the medium neutrino mass for the normal neutrino hierarchy (NH) and between the medium and the heavy neutrino mass for the inverted neutrino hierarchy (IH), and Δm^2 is the larger squared mass splitting between the light and the heavy neutrino mass in both cases.

With the knowledge of the two squared mass splittings, the remaining uncertainty lies mostly with the scale of the lightest neutrino. In the following Bayesian analysis we therefore apply Gaussian priors according to eqs. (2.25) and (2.26) and vary over the lightest neutrino mass.

2.4.1 Neutrino masses: Posteriors

We have taken nested sampling runs for an extension of the base Λ CDM cosmology with three individual neutrino masses, where we have used both a uniform prior $m_{\text{light}} \sim \mathcal{U}(0, 1)$ and logarithmic priors with different lower bounds, $\log_{10} m_{\text{light}} \sim \mathcal{U}(-5, 0)$ and $\log_{10} m_{\text{light}} \sim \mathcal{U}(-10, 0)$, on the lightest neutrino mass. The other two neutrino masses are then derived from m_{light} together with δm^2 and Δm^2 from eqs. (2.25) and (2.26):

$$m_{\text{medium}}^2 = \begin{cases} m_{\text{light}}^2 + \delta m^2 & \text{(NH),} \\ m_{\text{light}}^2 + \Delta m^2 - \delta m^2 & \text{(IH),} \end{cases} \quad (2.27)$$

$$m_{\text{heavy}}^2 = m_{\text{light}}^2 + \Delta m^2. \quad (2.28)$$

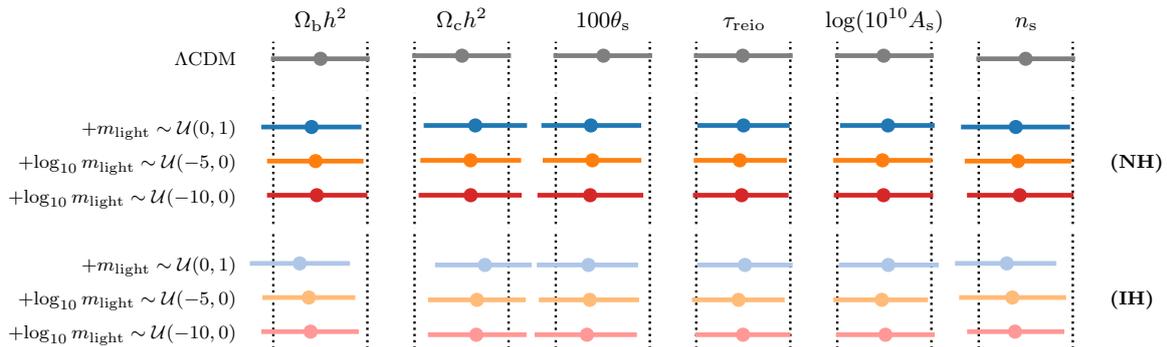


Figure 2.8 Stability of the cosmological parameters for the 3-neutrino extension of the base Λ CDM cosmology for different priors on m_{light} : uniform in blues, logarithmic with lower bound of -5 in oranges and logarithmic with lower bound of -10 in reds. The darker set of colors corresponds to the normal neutrino hierarchy (NH) and the lighter set to the inverted hierarchy (IH). For each parameter we show the mean and the extent from quantile 0.16 to 0.84, i.e. the inner 68% limits.

Figure 2.8 gives an overview of the stability of the cosmological base parameters across the different priors for m_{light} and compares them to the Λ CDM base model by showing their mean and 68% ranges. Compared to fig. 2.1 for the tensor-to-scalar ratio there are some clear parameter shifts visible in relation to the base Λ CDM model, but all shifts stay well within the 68% bounds. Comparing across prior models, the parameters are perfectly stable.

Figure 2.9 shows the one-dimensional marginalised posterior distributions for the three individual neutrino masses m_{light} , m_{medium} , and m_{heavy} , as well as the sum of all three $\sum m_{\nu}$ for both the normal and the inverted hierarchy. We have included shaded histograms to give a notion of the prior distributions. The vertical black dotted lines indicate roughly the lower bound for the medium and heavy neutrino mass as determined from the mass squared splittings under the assumption where the light neutrino mass is zero.

When looking at the lightest neutrino mass in the first row, the picture is very similar to that for the tensor-to-scalar ratio before, and most of what we have said in section 2.3.1 applies here, too. One has an almost exponential drop-off from zero when sampling uniformly over the mass (left column), significantly compressing the prior, which turns into a more step-like behaviour with respect to the logarithm of the mass when sampling the mass logarithmically (right column).

Note that the medium and heavy mass from rows 2 and 3 as well as the sum of all masses in the bottom row are *derived* quantities and therefore do not show the same prior behaviour visible for the light neutrino mass. This is not so apparent for the derived masses, when sampling uniformly over the light neutrino mass, although one can see a slight step for the heavy mass in the NH case and for both medium and heavy mass in the IH case. However, when sampling logarithmically over the light neutrino mass, then the picture is much clearer. The probability density for medium and heavy neutrino mass bulks up around their rough lower minimum set by the smaller and larger mass square splitting respectively.

There are two perspectives that one can adopt here. On one hand, one could criticise the choice of a logarithmic prior for being ultimately too prior (or theory) driven and not

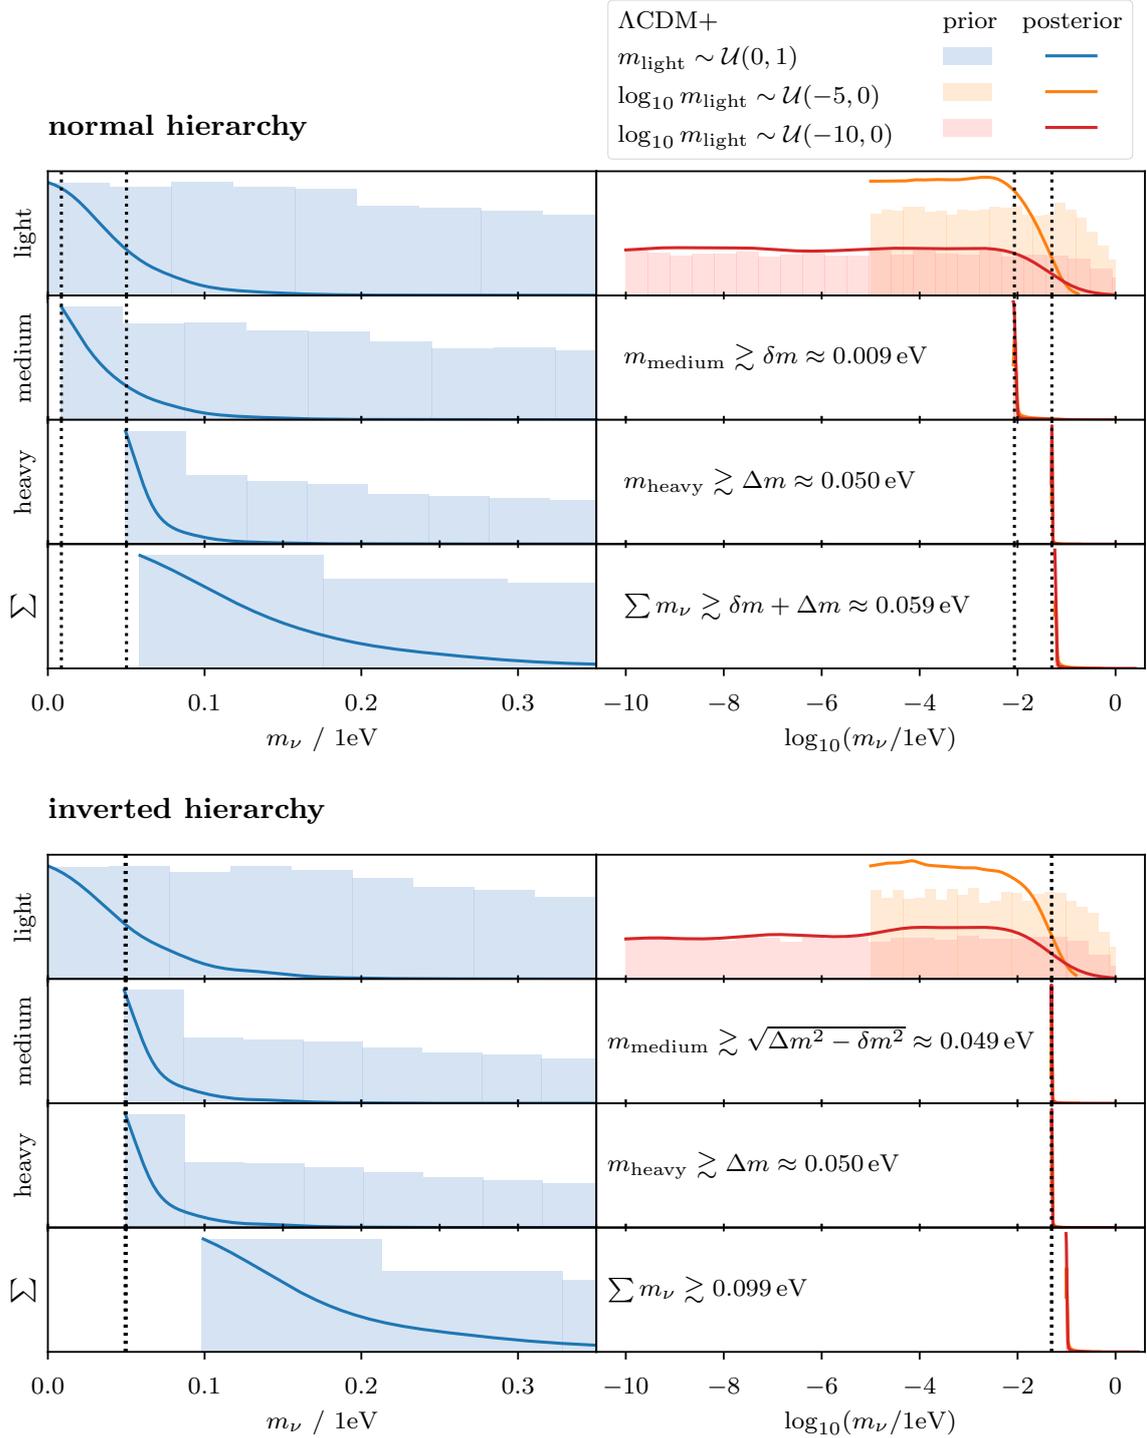


Figure 2.9 One-dimensional posterior distributions of neutrino masses with normal hierarchy (NH) in the top panel and with inverted hierarchy (IH) in the bottom panel for TT,TE,EE+lowE data from Planck 2018 and neutrino oscillation data on the mass squared splittings from NuFIT 5.0 (2020). The vertical black dotted lines give the rough lower limit on medium and heavy mass that is set by the mass squared splittings δm^2 and Δm^2 . For the inverted hierarchy these dotted lines appear almost on top of each other. The rows show the posteriors for the light, medium, and heavy neutrino mass and sum of all neutrino masses, respectively. The columns contrast the difference between using a uniform (blue, left) or logarithmic (orange and red, right) prior on the light neutrino mass m_{light} . The shaded histograms give a notion of that prior distribution.

Table 2.3 Mean and standard deviation of the log-evidence $\ln \mathcal{Z}$, Kullback–Leibler divergence \mathcal{D}_{KL} and Bayesian model dimensionality d of the base ΛCDM cosmology and its 3-neutrino extension from Planck 2018 TT,TE,EE+lowE data [62]. The second block of rows shows the results from the normal neutrino hierarchy and the third block for the inverted hierarchy. The Δ indicate normalisation with respect to the base ΛCDM model.

Model	$\Delta \ln \mathcal{Z}$	$\Delta \mathcal{D}_{\text{KL}}$	Δd	
ΛCDM	-0.00 ± 0.27	0.00 ± 0.26	0.00 ± 0.55	
normal	$\Lambda\text{CDM}+ m_{\text{light}} \sim \mathcal{U}(0, 1)$	-2.40 ± 0.25	1.47 ± 0.25	3.31 ± 0.60
	$\Lambda\text{CDM}+ \log_{10} m_{\text{light}} \sim \mathcal{U}(-5, 0)$	-0.13 ± 0.24	0.27 ± 0.24	0.12 ± 0.50
	$\Lambda\text{CDM}+ \log_{10} m_{\text{light}} \sim \mathcal{U}(-10, 0)$	-0.93 ± 0.25	0.73 ± 0.24	1.27 ± 0.53
inverted	$\Lambda\text{CDM}+ m_{\text{light}} \sim \mathcal{U}(0, 1)$	-2.90 ± 0.25	1.88 ± 0.24	2.29 ± 0.57
	$\Lambda\text{CDM}+ \log_{10} m_{\text{light}} \sim \mathcal{U}(-5, 0)$	-0.39 ± 0.24	0.06 ± 0.24	0.22 ± 0.49
	$\Lambda\text{CDM}+ \log_{10} m_{\text{light}} \sim \mathcal{U}(-10, 0)$	-0.44 ± 0.24	0.05 ± 0.24	1.01 ± 0.51

reflective of the data. On the other hand, one could say that this is the natural result of our state of knowledge of the mass square splittings and our true ignorance about the scale of the lightest neutrino mass.

We wonder whether this very last statement could be contested, e.g. could we say that we would expect the lightest neutrino mass to be of a magnitude similar to that of the medium neutrino mass in the NH? However, this is not the case, when checking for precedence by looking at the other set of leptons, the electron, muon and tauon, where we have roughly around 2 orders of magnitude between their masses [112].

Comparing the two hierarchies with one another, we can see that the major difference lies in the medium neutrino mass (and therefore also the sum of all neutrino masses), which is restricted to larger masses in the inverted hierarchy compared to the normal hierarchy, as expected from the mass square splitting (black dotted lines).

2.4.2 Neutrino masses: Evidence and Kullback–Leibler divergence

In [fig. 2.10](#) we show the results from our nested sampling runs for the log-evidence $\ln \mathcal{Z}$, KL-divergence \mathcal{D}_{KL} , Bayesian model dimensionality d and posterior average of the log-likelihood $\langle \ln \mathcal{L} \rangle$. We again normalise with respect to the base ΛCDM model. [Table 2.3](#) lists the summary statistics for these quantities. As already the case for the posterior, the picture here is again similar to the one for the tensor-to-scalar ratio in [section 2.3.2](#).

Looking at the distributions for the log-evidence (topmost diagonal panel) shows that the addition of the neutrino parameters with uniform sampling over the light neutrino mass (either hierarchy) is disfavoured with over 2 log-units compared to the base ΛCDM model with a single massive neutrino of fixed mass (and 2 massless). Since the mass square splittings enter on the prior level in our analysis and remain essentially unconstrained by the cosmological data, any change to the evidence is almost entirely driven by the light neutrino mass parameter.

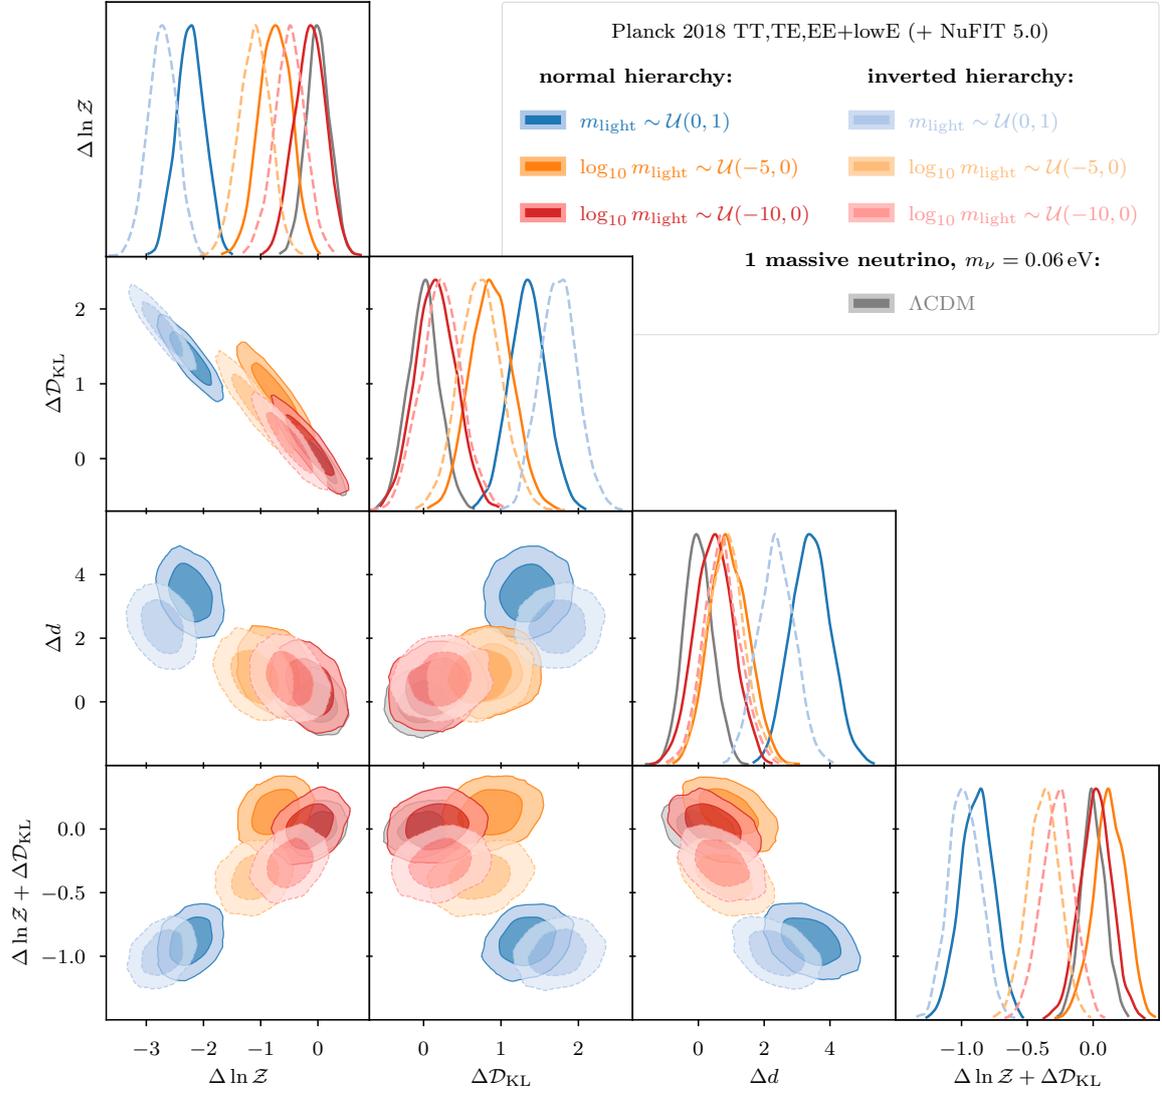


Figure 2.10 Effect of uniform vs logarithmic priors on the light neutrino mass m_{light} for Bayesian model comparison: log-evidence $\Delta \ln \mathcal{Z}$, Kullback–Leibler divergence \mathcal{D}_{KL} , Bayesian model dimensionality d , and posterior average of the log-likelihood $\langle \ln \mathcal{L} \rangle_{\mathcal{P}} = \ln \mathcal{Z} + \mathcal{D}_{\text{KL}}$. The probability distributions represent errors arising from the nested sampling process. In the limit of infinite life points these distributions would become point statistics, in contrast to posterior distributions. We normalise with respect to the ΛCDM model with two massless and only one massive neutrino with $m_\nu = 0.06$ eV. Note, how switching from uniform to logarithmic sampling of m_{light} moves the contours along their $\ln \mathcal{Z}$, \mathcal{D}_{KL} degeneracy line, i.e. relative entropy is traded in for evidence. Note further by comparison of the orange and red lines, how changing the lower bound of the logarithmic sampling interval (by 5 log-units!) barely affects the contours (bar some expected statistical fluctuation).

Hence, it is not surprising that upon switching to a logarithmic prior on m_{light} the log-evidence increases again while the KL-divergence drops close to the level of the Λ CDM model. We need to keep in mind that since this is an extension to the Λ CDM model, it has in principle a better chance of fitting the data, such that any difference in the Bayesian evidence can be attributed to an Occam penalty, which the shift between uniform and logarithmic sampling confirms.

As expected from our investigations for the tensor-to-scalar ratio and especially with regards to our mock example from [section 2.3.3](#), changing the lower bound for the logarithmic prior does not affect the Bayesian evidence. We have again performed runs with two different lower bounds of $\log m_{\text{light}} = -5$ and $\log m_{\text{light}} = -10$, i.e. five orders of magnitude apart. With both of these bounds well into the area of the posterior (see top right panel of [fig. 2.9](#)) where it has levelled off, we do not expect much change to the evidence value, which is confirmed in [table 2.3](#) and [fig. 2.10](#) for both hierarchies.

Looking at the normalised posterior average of the log-likelihood $\langle \ln \mathcal{L} \rangle_{\mathcal{P}} = \ln \mathcal{Z} + \mathcal{D}_{\text{KL}}$ we again confirm

$$\Delta \ln \mathcal{Z}_{\text{uni}} + \Delta \mathcal{D}_{\text{KL,uni}} \approx -1, \quad (2.29)$$

$$\Delta \ln \mathcal{Z}_{\text{log}} + \Delta \mathcal{D}_{\text{KL,log}} \approx 0, \quad (2.30)$$

matching our mock results from [eqs. \(2.23\)](#) and [\(2.24\)](#), independent from the mock parameter μ .

2.4.3 Neutrino hierarchy

A Bayesian model comparison of the normal vs the inverted neutrino hierarchy is beyond the scope of this paper and has been done before with more stringent data [[80](#), [113](#), [114](#)]. However, with posteriors and evidences at hand, we shall briefly discuss the situation here.

There have been claims to a strong preference of the normal over the inverted neutrino hierarchy [[73](#)], however, such strong evidence can typically be traced back to prior volume effects [[74](#)], i.e. the effect of a reduced sampling space for the inverted hierarchy. In other words, we need to watch out and properly distinguish to what extent any Bayesian preference is assigned already on the prior level and to what extent is that preference indeed driven by the data.

In our analysis both hierarchies start out on an equal footing. With the same prior on the light neutrino mass and equivalent Gaussian priors on the mass squared splittings from neutrino oscillation experiments, the prior volume for both hierarchies is essentially the same. Note that although the means for the larger mass squared splitting Δm^2 are slightly different in the two hierarchies, its standard deviations are essentially the same.

With an evidence difference of only about half a log-unit (and thereby not much bigger than the error on the evidence estimates) any preference of the normal hierarchy is meagre at best in case of uniform sampling of m_{light} , and even less when switching to logarithmic sampling, as visible in [fig. 2.10](#). It should be noted, though, that we have used only CMB temperature and polarisation data here. Adding data from CMB lensing or baryon acoustic

oscillations would have further shrunk the constraints on the sum of neutrino masses and thereby possibly strengthened the case for the normal hierarchy.

2.5 Discussion

We demonstrate how switching between a uniform and a logarithmic prior on some single-bounded model parameter results in a trade-off between Bayesian evidence and Kullback–Leibler divergence (or relative entropy). The common scenario is that of insufficient data sensitivity, leading to a one-sided bound on a parameter. For a location parameter this typically causes an exponential drop-off, which translates to a step-like behaviour when turned into the corresponding scale parameter. We show that the ambiguity of the lower bound of the scale parameter does not affect a Bayesian model comparison, provided the lower bound is chosen sufficiently far into the likelihood plateau.

We demonstrate this behaviour for two cases of parameter extensions to the Λ CDM model of cosmology, namely for the tensor-to-scalar ratio of primordial perturbations and for the case of three non-degenerate neutrino masses. In both cases we confirm that switching from a uniform prior to a logarithmic prior will get rid of (most of) the Occam penalty associated with that parameter, since unconstrained parameters do not affect the Bayesian evidence. Thus the Bayesian evidence is roughly on par with the un-extended (base) model, with the only difference in the form of an uninformative parameter. Furthermore and for the same reason, the exact choice of the lower bound for the logarithmic prior does not change the Bayesian evidence. When the likelihood levels off, e.g. due to insufficient sensitivity in the data, then so does the Bayesian evidence.

We should note at this point that we discovered that, different from checks with a Markov Chain Monte Carlo (MCMC) sampler, the nested sampler (needed for determining the Bayesian evidence) did not apply a multivariate Gaussian prior on the two nuisance parameters associated with the Sunyaev–Zeldovich effect. This seems to slightly affect the posteriors on the lightest neutrino mass, but tests already indicate that the qualitative results of the model comparison stay the same. Following correction of the nested sampler, new runs will be carried out and results from these will go into a paper which is being prepared on the basis of this chapter. We expect the effects on the results to be negligible compared to the statistical uncertainty, however.

Chapter 3

A case for kinetically dominated initial condition for inflation

This chapter was published in Hergt et al. (2019) [I & chapter 3] and makes a case for setting initial conditions for inflation at the Planck epoch in the kinetically dominated regime. For inflationary potentials with a plateau or a hill, i.e. potentials that are bounded from above within a certain region of interest, we cannot claim complete ignorance of the energy distribution between kinetic and potential energy, and equipartition of energy at the Planck epoch becomes questionable. We analyse different classes of potentials in phase-space and quantify the fraction of the Planck surface that is kinetically dominated. Considering bounded potentials with very small amplitudes as favoured by current data and restricting ourselves to the domains of phase-space that are of interest to cosmic inflation, we find that initial conditions of the inflaton field should be set in the kinetically dominated regime regardless of any choice of prior.

3.1 Introduction

Handley et al. (2014) [115] show under broad assumptions that classical inflationary universes generically emerge in a regime where the kinetic energy of the inflaton dominates over the potential. In contrast, the traditional procedure for setting initial conditions for inflation defines them at the Planck epoch with a total energy density of the order of $\rho \sim m_{\text{p}}^4$ [116–119] and, lacking any further prior constraints, partitions inflaton energy equally between kinetic and potential energy [117–121].

In this article we show how the choice of certain potentials gives additional prior constraints. We make a case for why initial conditions at the Planck time should be set using kinetic dominance as opposed to assuming that the potential holds half the energy.

In another work we explore the effects of kinetic dominance initial conditions on cosmological parameters [II & chapter 4]. When exploring observational consequences, the traditional approach is to assume the inflaton to be in the slow-roll regime, where the potential dominates the inflaton energy, as this is an attractor solution to the inflaton's equation of motion. However, this assumption is only reasonable for large amounts of inflation, where all

observable modes have already been well within the comoving Hubble horizon at the start of inflation. Indeed there are several cases that predict fewer e-folds of inflation, e.g. for certain inflation models [122], for closed universes [123], or through scenarios such as “finite inflation” [124, 125] or “just enough inflation” [126, 127]. In all these cases, kinetic dominance initial conditions will cause oscillations and a large scale cutoff in the primordial power spectra of scalar and tensor perturbations that translate through to the level of the angular power spectrum of the cosmic microwave background (CMB) [II & chapter 4, 121, 126–132], which could explain the observed lack of power on large scales in the case of the standard Λ cold dark matter (Λ CDM) model of cosmology [35, 133, 134].

3.2 The Planck surface in phase space

Energies beyond the Planck scale require a quantum theory of gravity. Only when $\rho \lesssim m_{\text{p}}^4$ may we set initial conditions for any *classical* evolution [116–119]. For the inflaton field ϕ in a spatially-flat universe this means starting from the Planck circle where, neglecting spatial inhomogeneities,

$$\frac{1}{2}\dot{\phi}^2 + V(\phi) = m_{\text{p}}^4. \quad (3.1)$$

In the parametrisation $x \equiv \text{sgn}(\phi)\sqrt{V(\phi)}$ and $y \equiv \dot{\phi}/\sqrt{2}$ this indeed is a circle with radius $r = m_{\text{p}}^2$. However, in the $(\phi, \dot{\phi})$ phase-space the Planck surface is not, in general, circular. Depending on the potential $V(\phi)$, the shape of the surface in the $(\phi, \dot{\phi})$ phase-space changes, leading to features such as local maxima or asymptotes in $\dot{\phi}$. In the (x, y) parametrisation this manifests as part of the circle becoming effectively excluded.

We define the inflating region via $\ddot{a} > 0$, or equivalently,

$$\dot{\phi}^2 < V(\phi). \quad (3.2)$$

Similarly, we conservatively define the kinetically dominated (KD) region by the condition that kinetic energy dominates the potential by 2 orders of magnitude:

$$\dot{\phi}^2 > 100 V(\phi). \quad (3.3)$$

Figure 3.1 shows the Planck circle in the (x, y) parametrisation as a black dotted line. The inflating and KD regions are sectors centred along the x - and y -axis respectively. This picture holds irrespective of the choice of potential. However, the choice of potential can effectively exclude parts of this circle: e.g. for the plateau and hilltop potentials defined in eqs. (3.5) and (3.6) below, this is illustrated by the grey shaded regions covering smaller values of x as the amplitude parameter Λ decreases. Note how for very small Λ^2 the relevant part of the circle is reduced to within the KD region [135]. In fig. 3.2, discussed more fully below, we see that the picture changes significantly upon changing to $(\phi, \dot{\phi})$ phase-space and for

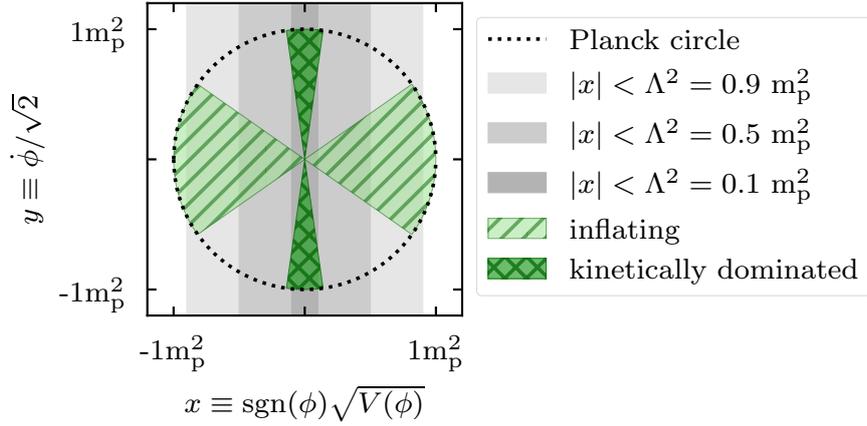


Figure 3.1 The Planck circle in the (x, y) parametrisation is depicted by the black dotted line. The stripy and checked hatched regions correspond to the inflating region and the kinetically dominated region, respectively, [see eqs. (3.2) and (3.3)]. The grey shading highlights the effective x -range relevant for three given values of Λ^2 for potentials such as plateau or hilltop potentials [e.g. the Starobinsky and Landau-Ginzburg potentials defined in eqs. (3.5) and (3.6), respectively].

alternative potentials. In the following section we review our three representative potential classes: quadratic, plateau, and hilltop.

3.3 Inflaton potentials

3.3.1 Quadratic potential

The simplest potential typically used for single field inflation models is given by

$$V(\phi) = m^2 \phi^2, \quad (3.4)$$

where m is the mass of the inflaton field. This quadratic potential is often defined with a multiplicative factor $\frac{1}{2}$, which is irrelevant for this analysis and omitted here for reasons of compatibility with other power law potentials. In general, $m \sim 10^{-6} m_p$ in order to produce an appropriate primordial power spectrum amplitude [35]. Though quadratic potentials are disfavoured by current data because of their excess production of gravitational waves [131], we include them for comparison with other models and traditional theoretical discussion.

The quadratic potential plays a special role in the phase-space representation because (x, y) is equivalent to $(\phi, \dot{\phi})$ phase-space up to a factor of m . Because of this close relationship for quadratic potentials, the distinction between the two spaces is often overlooked. In the following sections we discuss the effects of considering alternative potentials.

3.3.2 Plateau potentials

Some high-energy models give rise to plateau-like potentials [136–138]. A popular example is the Starobinsky potential [136]:

$$V(\phi) = \Lambda^4 \left[1 - e^{-\sqrt{2/3}\phi/m_p} \right]^2, \quad (3.5)$$

where the scale $\Lambda^2 \sim 10^{-6} m_p^2$ is comparable to the inflaton mass m of the quadratic potential. Unlike the quadratic potential, due to concavity the plateau potential produces very few gravitational waves making it preferred by Planck data [131].

For small amplitudes, $\Lambda < m_p$, the asymptote in the potential for $\phi \rightarrow \infty$ translates through to $(\phi, \dot{\phi})$ phase-space, where $\dot{\phi}$ converges to $\sqrt{2(m_p^4 - \Lambda^4)}$ from its global maximum at $\sqrt{2}m_p^2$. In (x, y) space this asymptote manifests as a cut from the Planck circle as an unattainable region. No value of $\phi > 0$ can reach the region $x = \sqrt{V} \geq \Lambda^2$. Note that on the exponentially growing side $\phi < 0$, the potential and thus x increase beyond the Planck scale irrespective of the amplitude Λ . However, for cosmic inflation only the plateau side of the potential is of interest to us.

3.3.3 Hilltop potentials

The double-well (Landau-Ginzburg) potential is given by

$$V(\phi) = \Lambda^4 \left[1 - \left(\frac{\phi}{\mu} \right)^2 \right]^2. \quad (3.6)$$

This potential has been frequently studied in the context of cosmic inflation and spontaneous symmetry breaking [139–146]. It is a particular realisation of the family of quadratic hilltop potentials [147] that ensures positivity (required here for the calculation of the Planck surface). As for plateau potentials, the concavity in the hilltop region ensures low production of primordial gravitational waves [131].

The inflaton rolls away from a local maximum at $\phi = 0$ to a minimum at $\phi = \pm\mu$. We can identify the region of interest $-\mu < \phi < \mu$ where the potential is bounded from above ($|x| = \sqrt{V} \leq \Lambda^2$), which effectively cuts the Planck circle as shown in fig. 3.1.

3.4 Inflaton Phase-Space

In fig. 3.2 we plot the quadratic, plateau, and hilltop models side by side. The top row shows the potentials from eqs. (3.4) to (3.6). The other panels are $(\phi, \dot{\phi})$ phase-space plots for decreasing values of the inflaton mass m , or amplitude Λ from top to bottom.

For the quadratic potential, the shape of the Planck surface is ellipsoidal and thus closed. For the plateau potential this ellipse is drawn out to $\phi \rightarrow \infty$ because of the asymptote in the potential $V(\phi \rightarrow \infty) = \Lambda^4$. Only for $\Lambda > m_p$ does the Planck surface become closed. The double-well potential consists of two such deformed ellipsoids which are connected for $\Lambda \leq m_p$.

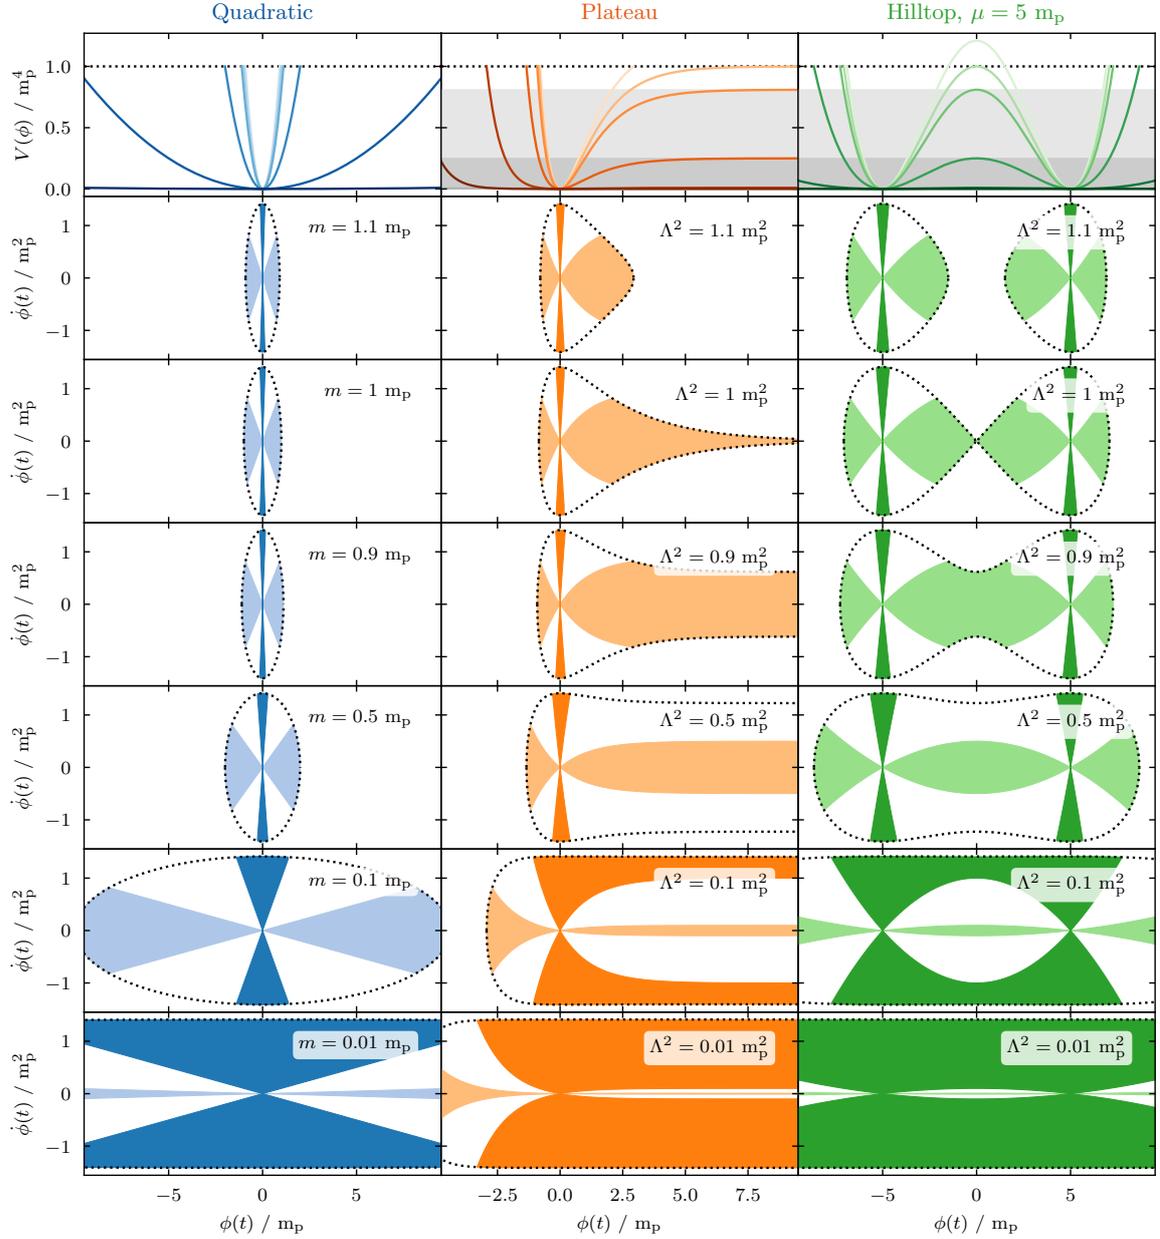


Figure 3.2 The top row shows $V(\phi)$ for the quadratic potential from eq. (3.4) in blue (left), for the plateau potential from eq. (3.5) in orange (middle), and for the hilltop potential from eq. (3.6) in green (right). The lines in each potential plot grow darker with decreasing m or Λ^2 . The grey shaded regions correspond to those in fig. 3.1 delimiting the maximum value of the plateau or hill region for selected values of Λ^2 . The other panels depict $(\phi, \dot{\phi})$ phase-space diagrams for the corresponding potentials with values for m and Λ^2 decreasing from the second row downwards. The enveloping dotted line denotes the Planck surface. Light shading corresponds to the inflating regions where $\dot{\phi}^2 < V(\phi)$. Dark shading corresponds to the kinetically dominated regions where $\dot{\phi}^2 > 100 V(\phi)$. Note that realistic values for m and Λ^2 , i.e. conforming with data for the primordial amplitude A_s from e.g. the Planck satellite [35], are orders of magnitude smaller (i.e. continuing the rows further down) than the values picked here for reasons of visualisation.

Of particular interest is the kinetically dominated (KD) regime that only covers a tiny part at the poles of the Planck surface when m or Λ are large. When decreasing m in the quadratic model, the KD regime expands in proportion to the Planck surface. For the plateau model, the KD region can continue to stretch while the Planck surface already extends to infinity. Analogously, in the hilltop case the Planck surface already spans from $-\mu$ to μ and KD eventually spreads to cover the whole range for very small Λ .

3.5 Priors on initial conditions

When setting initial conditions at the Planck scale, the traditional approach is to assume equipartition between kinetic and potential energy [117–121]. It then is concluded that an initial value for the inflaton ϕ_0 may take any value respecting the condition that $V(\phi_0) \lesssim m_{\text{p}}^4$. With only this constraint, for very small inflaton masses $m \sim 10^{-5} m_{\text{p}}$ this means that ϕ_0 is allowed to take a huge range of values spanning many orders of magnitude. Faced with complete ignorance regarding the scale of ϕ_0 it is arguably more reasonable to choose a logarithmic prior for the initial value of the inflaton field.

In [fig. 3.3](#) the trajectories of the inflaton field in $(\phi, \dot{\phi})$ phase-space are plotted using different priors for the generation of initial conditions on the Planck surface. [Figure 3.4](#) shows the fraction of trajectories starting in the kinetically dominated (KD) region for each prior respectively. We compare uniform and logarithmic priors in ϕ , a uniform prior on the energy distribution between potential and kinetic energy, and a prior uniform on arc-length on the Planck surface. In general, we can see that equipartition of energy pushes the trajectories outside the KD window whereas a scale-invariant logarithmic prior has the opposite effect.

For the quadratic potential, prior choice plays a significant role. The KD fractions for priors in ϕ or energy are independent of the inflaton mass m . Only for the uniform prior on arc-length do we find a decrease of the KD fraction at large mass, when the Planck surface spans a very small ϕ -range. For energy equipartition we get a KD fraction of only 2% and for a uniform prior on ϕ about 14%. A logarithmic prior on ϕ raises the fraction to about 79%, where we start at $\phi_{\text{start}} = 10^{-4} m_{\text{p}}^2/m$, which is roughly the scale from which a sufficient amount of e-folds of inflation is produced for realistic values of m .

Unlike the quadratic potential, the KD fraction for the plateau and hilltop potentials does depend on Λ^2 . This is related to the changing shape of the Planck surface illustrated in [fig. 3.2](#) above. For $\Lambda > m_{\text{p}}$ the Planck surface is bounded and the KD fractions display a similar behaviour to the quadratic model. The fraction drops towards smaller values of Λ , as the Planck surface stretches across a greater ϕ -range, whereas the KD regime does not. Crossing $\Lambda = m_{\text{p}}$ reverses this behaviour in the hilltop case. As the Planck surface spans across $\phi = 0$, while the KD regime continues to spread out, the KD fraction starts to rise. Finally for $\Lambda^2 < m_{\text{p}}^2/\sqrt{51}$, the KD fraction levels off. For the Starobinsky model, where the Planck surface stretches to infinity, the fraction stays at zero at intermediate amplitudes, and jumps to unity once the KD regime also spreads to infinity.

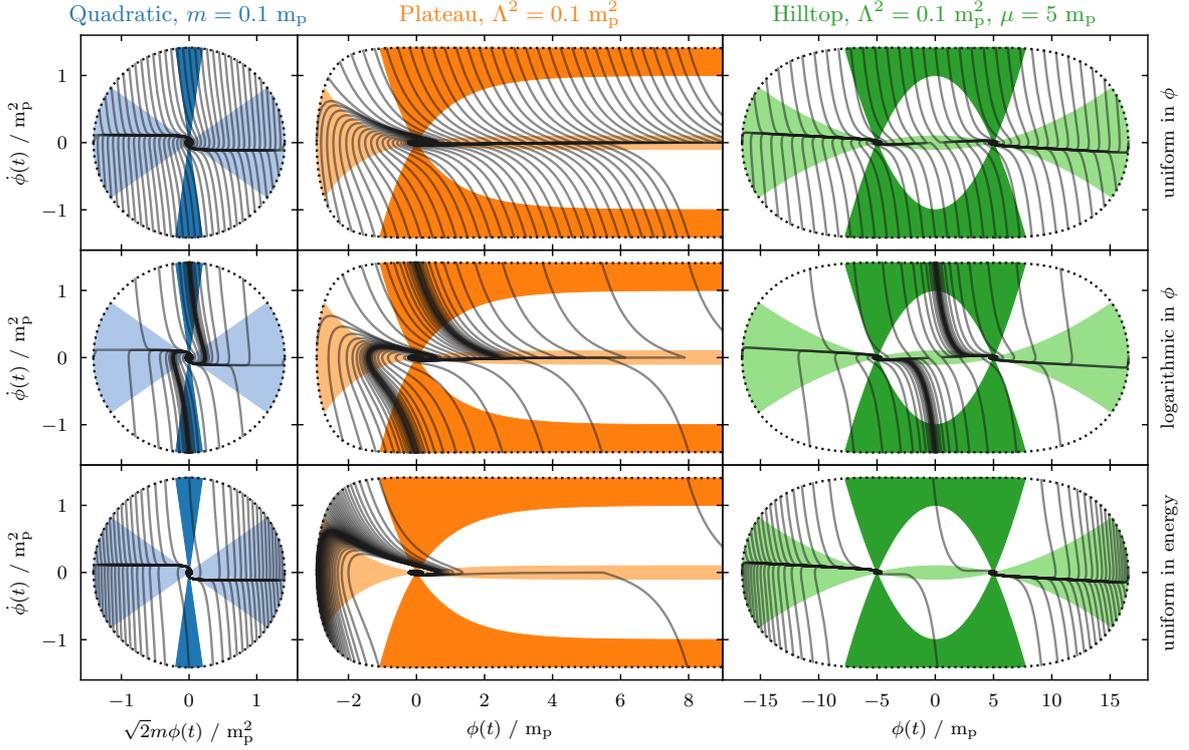


Figure 3.3 $(\phi, \dot{\phi})$ phase-space diagrams with inflaton trajectories for the quadratic potential in blue (left), the plateau potential in orange (middle), and the hilltop potential in green (right). Light shading correspond to the inflating regions where $\dot{\phi}^2 < V(\phi)$. Dark shading corresponds to the kinetically dominated regions where $\dot{\phi}^2 > 100 V(\phi)$. The enveloping dotted line denotes the Planck surface. In the top row initial conditions were set with a uniform prior on ϕ , with a logarithmic prior on ϕ in the middle row, and with a uniform prior on the potential energy in the range $[0, m_{\text{pl}}^4]$ in the bottom row.

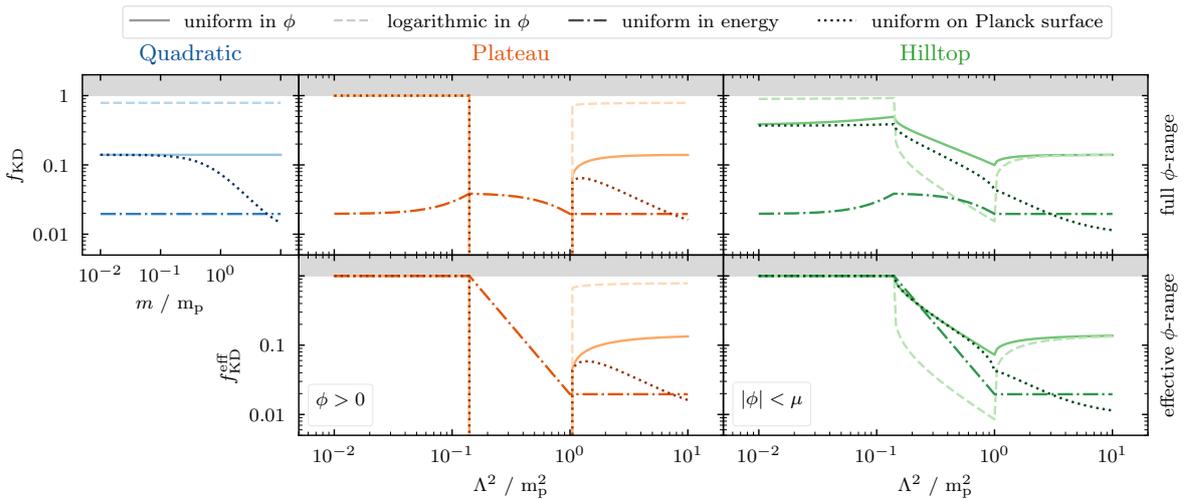


Figure 3.4 Fraction f_{KD} of the kinetically dominated region to the full Planck region in $(\phi, \dot{\phi})$ phase-space corresponding to the prior choices from fig. 3.3. The second row shows how these fractions change when we only consider the ϕ -range effectively necessary for the type of inflation we are interested in, i.e. $\phi > 0$ for the plateau potential and $|\phi| < \mu$ in the hilltop case.

For both the plateau and hilltop potentials we note how a uniform prior on the energy distribution causes only a small number of trajectories to start in the regions of interest: $\phi > 0$ and $-\mu < \phi < \mu$, respectively. Once we restrict ourselves to the ϕ -ranges of importance, all trajectories start in the KD region irrespective of the prior choice. This change from the full ϕ -range to a restricted ϕ -range is illustrated by going from the first to the second row of panels in [fig. 3.4](#). Thus, for potentials with an upper bound in the region of interest (i.e. most concave models) initial conditions for inflation should naturally be set in the KD regime.

3.6 Conclusion

We have shown how phase-space considerations can be significantly more complicated in comparison with the simplistic Planck circle picture when considering models other than quadratic inflation. This is particularly relevant for potentials with features such as plateaus or hills, i.e. concave and bounded from above in the region of interest for slow-roll inflation.

For such models we find that the regime of kinetic dominance covers the entire Planck surface in the regions of interest and for realistically small values of the potential amplitude. Thus, all trajectories in the phase-space of the inflaton field start out kinetically dominated, independent of the prior choice for the initial conditions. Therefore, for these cases initial conditions for inflation at the Planck epoch should be set in the kinetically dominated regime and not assuming energy equipartition.

Chapter 4

Constraining the kinetically dominated Universe

In this chapter, which was published in Hergt et al. (2019) [II & chapter 4], we present cosmological constraints from Planck 2015 data for a universe that is kinetically dominated at very early times. We perform a Markov chain Monte Carlo analysis to estimate parameters and use nested sampling to determine the evidence for a model comparison of the single-field quadratic and Starobinsky inflationary models with the standard Λ CDM cosmology. In particular we investigate how different amounts of inflation before and after horizon exit affect the primordial power spectrum and subsequently the power spectrum of the cosmic microwave background. We find that the model using kinetically dominated initial conditions for inflation performs similarly well in terms of Bayesian evidence as a model directly starting out in the slow-roll phase, despite having an additional parameter. The data show a slight preference for a cutoff at large scales in the primordial and temperature power spectra.

4.1 Introduction

Inflation was first introduced in the 70s and 80s (see [55–57] for some of the original papers and section 2 of [148] for a more extensive introduction) and plays an important role in today’s standard model of cosmology (Λ CDM). Besides solving issues such as the horizon and flatness problems, it provides a mechanism for generating primordial perturbations that can serve as seeds for the formation of cosmic structure, which in turn generate the observed temperature anisotropies in the cosmic microwave background (CMB) [149].

Typically, a slow-roll (SR) inflation model is assumed, whereby the kinetic energy of a single scalar field ϕ is dominated by its potential $V(\phi)$ and hence the inflaton “slowly rolls down” the potential. Generically, the slow-roll condition is an attractor solution so even from a position in phase space where slow-roll is not satisfied, the inflaton will rapidly lose speed and approach a slow-roll regime [I & chapter 3, 115–117, 120, 121, 129, 150].

High-precision measurements of the CMB, first through WMAP [151] then through Planck [21, 22], have significantly contributed to the success of the standard Λ CDM model of cosmology. Nonetheless, the data also revealed features in the CMB angular power spectrum

hinting at potential additional physics [133, 134, 152, 153]. These features include the low-multipole lack of power and a small dip at multipoles of approximately 20–25. These features may be caused by corresponding features in the primordial power spectrum (PPS), which recently has led to many investigations of PPS with a cutoff [121, 126–132]. One might be inclined to brush off any significance of the low-multipole lack of power as an effect to be attributed to cosmic variance. However, it has been shown to be driving parameter extensions (e.g. spatial curvature, see [33]) despite the large error bars owing to cosmic variance at these small multipoles.

Previous findings [123] have suggested that the Universe should have emerged from an initial singularity with the energy density dominated by the kinetic term $\dot{\phi}$ of the inflaton field. It can indeed be shown that such a kinetically dominated (KD) early universe emerges generically under rather broad assumptions [115]. This is particularly relevant for inflationary potentials that have an upper limit in the inflaton range of interest, such as plateau or hilltop potentials [I & chapter 3].

If the total amount of inflation is limited, the effects of KD become observable. Indeed there are several cases that predict fewer e-folds of inflation, e.g. for certain inflation models [122], for closed universes [123], or through scenarios such as “finite inflation” [124, 125] or “just enough inflation” [126, 127]. Another way of motivating KD is through the “just enough inflation” scenario [126, 127].

We review how KD initial conditions result in oscillations and a cutoff towards large scales in the PPS and consequently also in the CMB angular power spectrum. We show how these features depend mainly on the amount of inflation happening before or after horizon exit of a given mode k . We confirm previous constraints from [130] for cosmological parameters given KD initial conditions using a different Markov chain Monte Carlo (MCMC) sampler and expand on the analysis by comparing the Bayesian evidences for the different models.

We start out by summarising the inflationary background evolution in section 4.2, and by introducing two inflationary potentials, the quadratic and the Starobinsky potential, which we will use throughout this paper. In section 4.2.2 we review the kinetic dominance regime that provides us with the initial conditions for the numerical integration of the inflaton equations of motion and the mode equations for the primordial perturbations, which lead us to the analyses of the PPS in section 4.3 and the CMB angular power spectrum in section 4.4. Finally, in section 4.5 we present the results from our MCMC analysis and conclude in section 4.6.

4.2 Background Evolution during Kinetic Dominance

We focus on single-field inflationary models as determined by an inflaton field $\phi(t)$ in a spatially flat universe. Assuming the inflaton dominates all other species early in the history of the Universe, the background dynamics are governed by the Friedmann and continuity equations

for the inflaton

$$H^2 = \frac{1}{3 m_{\text{p}}^2} \left(\frac{1}{2} \dot{\phi}^2 + V(\phi) \right), \quad (4.1a)$$

$$\dot{H} = -\frac{1}{2 m_{\text{p}}^2} \dot{\phi}^2, \quad (4.1b)$$

$$\ddot{\phi} + 3H\dot{\phi} + V'(\phi) = 0, \quad (4.1c)$$

where a dot denotes differentiation with respect to cosmic time, $\dot{f} \equiv \frac{df}{dt}$. For convenience we set $c = \hbar = 1$ and use the reduced Planck mass $m_{\text{p}} = \sqrt{\frac{\hbar c}{8\pi G}}$.

Inflation is defined as a positive acceleration of the scale factor $\ddot{a} > 0$, or equivalently as a shrinking comoving Hubble horizon $\frac{d}{dt} \left(\frac{1}{aH} \right) < 0$. Using eqs. (4.1a) to (4.1c) we can recast this condition for inflation in terms of the inflaton field ϕ

$$\dot{\phi}^2 < V(\phi), \quad (4.2)$$

or in terms of the equation-of-state parameter $w \equiv \frac{p}{\rho}$ relating pressure p and energy density ρ of the inflaton

$$w_{\phi} = \frac{p_{\phi}}{\rho_{\phi}} = \frac{\frac{1}{2} \dot{\phi}^2 - V(\phi)}{\frac{1}{2} \dot{\phi}^2 + V(\phi)} < -\frac{1}{3}. \quad (4.3)$$

The amount of inflation from some time t to the end of inflation t_{end} can be measured in terms of the number of e-folds of the scale factor $a(t)$

$$N(a) \equiv \ln \left(\frac{a_{\text{end}}}{a} \right), \quad (4.4)$$

where $a_{\text{end}} = a(t_{\text{end}})$.

4.2.1 Potentials

To perform numerical integrations of the background dynamics in eqs. (4.1a) to (4.1c) we have focused on two specific potentials in particular: the quadratic potential and the Starobinsky potential shown in fig. 4.1.

Quadratic potential

The quadratic potential is defined by

$$V(\phi) = m^2 \phi^2, \quad (4.5)$$

where m is the mass of the inflaton field. This quadratic potential is often defined with a multiplicative factor $\frac{1}{2}$, omitted here for reasons of compatibility with other power law potentials. Though disfavoured by the Planck data, we are considering the quadratic potential here as the conceptually simplest implementation of a single scalar inflaton field. Using the slow-roll (SR) approximation $\dot{\phi}^2 \ll V(\phi)$ we can predict the spectral index and the tensor to

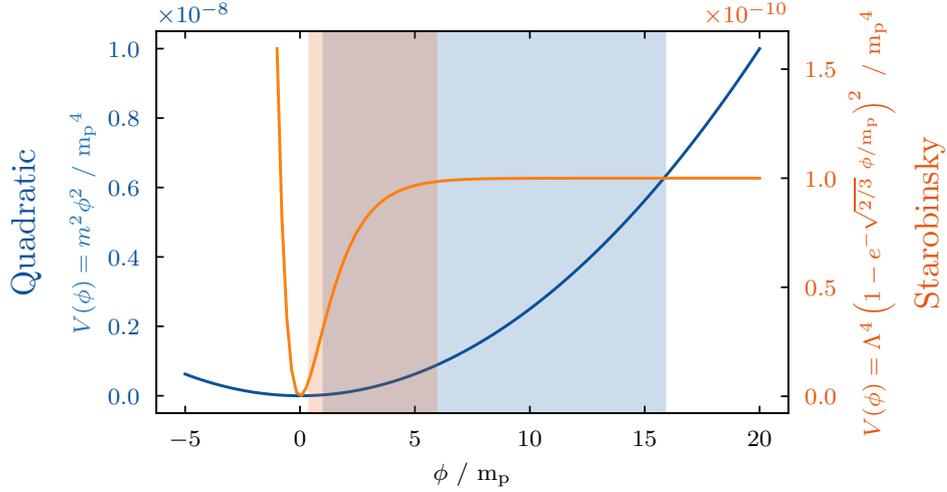


Figure 4.1 Quadratic (blue) and Starobinsky (orange) potentials as functions of the inflaton field ϕ , where $m = 5 \times 10^{-6} m_p$ for the quadratic potential and $\Lambda^2 = 10^{-5} m_p^2$ for the Starobinsky potential. The shaded regions mark the start and end of inflation in the case of kinetic dominance initial conditions.

scalar ratio to be

$$n_s \approx 1 - \frac{2}{N_*}, \quad r \approx \frac{8}{N_*}, \quad (4.6)$$

where N_* is the observable amount of inflation from horizon exit of a given pivot scale k_* to the end of inflation. Thus for $N_* = 55$ e-folds we expect $n_s \approx 0.964$ and $r \approx 0.145$.

Starobinsky potential

The Starobinsky potential is the potential representation in the Einstein frame of an $(R + R^2)$ modified theory of gravity first proposed by [136] and is given by

$$V(\phi) = \Lambda^4 \left[1 - \exp\left(-\sqrt{\frac{2}{3}} \frac{\phi}{m_p}\right) \right]^2. \quad (4.7)$$

Unlike quadratic inflation, the Starobinsky model gives rise to a low tensor-to-scalar ratio r , as is preferred by current data [131]. In the same manner as quadratic inflation, we can determine the spectral index and the tensor to scalar ratio using the slow-roll approximation

$$n_s \approx 1 - \frac{2}{N_*}, \quad r \approx \frac{12}{N_*^2}. \quad (4.8)$$

Thus for $N_* = 55$ e-folds we expect $n_s \approx 0.964$ and $r \approx 0.004$.

4.2.2 Kinetic Dominance initial conditions

The initial conditions for the integration of the background eqs. (4.1a) to (4.1c) are usually chosen according to the slow-roll (SR) regime, satisfying

$$\dot{\phi}^2 \ll V(\phi). \quad (4.9)$$

However, we do not need to place ourselves (somewhat artificially) directly into the period of SR inflation. As observed previously [116, 117], the expansion of the Universe acts as a damping term in the equation of motion (4.1c). This means the SR condition is an attractor solution, such that no matter where we start out in the $(\phi, \dot{\phi})$ phase-space we will end up on the SR attractor (provided one assumes an appropriate inflationary potential). Indeed, Handley et al. [115] show under broad assumptions that classical inflationary universes generically emerge from an initial singularity ($a \rightarrow 0$) with the kinetic energy of the inflaton dominating its potential energy [115], which we will refer to as kinetic dominance (KD)

$$\dot{\phi}^2 \gg V(\phi). \quad (4.10)$$

In a recently submitted paper [I & chapter 3], we make a case for kinetically dominated initial conditions for inflation through a $(\phi, \dot{\phi})$ phase-space exploration. This is particularly relevant in cases where the potential is bounded from above, e.g. hilltop or plateau potentials.

In case of limited amounts of inflation, such as suggested for certain inflation models [122], for closed universes [123], or through scenarios such as “finite inflation” [124, 125] or “just enough inflation” [126, 127, 154], effects of KD would be expected to be observable.

In the KD limit we can use the first terms of a series expansion of the background variables to generate a set of initial conditions for a sufficiently early starting time t_0 of the numerical integration

$$\phi_0 \equiv \phi(t_0) = \phi_p - \sqrt{\frac{2}{3}} m_p \ln \frac{t_0}{t_p}, \quad (4.11a)$$

$$\dot{\phi}_0 \equiv \dot{\phi}(t_0) = -\sqrt{\frac{2}{3}} \frac{m_p}{t_0}, \quad (4.11b)$$

$$H_0 \equiv H(t_0) = \frac{1}{3t_0}, \quad (4.11c)$$

$$a_0 \equiv a(t_0) = a_p \left(\frac{t_0}{t_p} \right)^{1/3}, \quad (4.11d)$$

where a_p and t_p can be set to unity as the exact value does not matter here due to rescaling symmetries [115]. ϕ_p controls the total number of e-folds of inflation $N_{\text{tot}} = N(a_{\text{start}})$, i.e. from the start of inflation $a_{\text{start}} = a(t_{\text{start}})$ to its end.

Figure 4.2 shows the evolution of the background variables ϕ , $\ln a$, w_ϕ , and H respectively, integrated using both SR and KD initial conditions and using the chaotic potential from eq. (4.5). For this figure, the initial conditions were set at the cosmic time $t_0 = 1$ and chosen such that $N_{\text{tot}} = 60$ e-folds are produced during inflation. For comparison, the end

of inflation in the SR case was shifted such that it matches the KD case. The inflaton mass of $m = 5 \times 10^{-6} m_{\text{p}}$ was chosen to produce an amplitude A_s of the primordial power spectrum close to the observational value. In all cases we see how the evolution begins differently depending on whether SR or KD initial conditions were chosen, but eventually the KD solution converges towards the SR solution. We can estimate the time of transition by equating the SR approximation of the Friedmann eq. (4.1a) with the approximation of the Hubble parameter during KD from eq. (4.11c), which ends up with an expression dependent on the inflaton mass m_ϕ (or equivalently the potential amplitude) and the initial field value: $t_i \approx 2 m_{\text{p}} / (m_\phi \phi_i) \sim 10^{-4} t_{\text{p}}$, where the subscript “i” indicates the moment of inflation start.

To distinguish between the different regimes it is useful to look at the equation-of-state parameter w_ϕ for the inflaton field and comparing with eq. (4.3)

$$w_\phi \begin{cases} \approx 1 & \text{kinetic dominance, } \dot{\phi}^2 \gg V(\phi), \\ > -\frac{1}{3} & \text{no inflation,} \\ < -\frac{1}{3} & \text{(fast-roll) inflation,} \\ \approx -1 & \text{slow-roll inflation, } \dot{\phi}^2 \ll V(\phi). \end{cases} \quad (4.12)$$

The equation-of-state parameter w_ϕ illustrates how in the SR case we directly start out in the inflationary epoch, whereas for the KD case we can specify a start and end point of inflation where w_ϕ crosses the $-1/3$ mark. For reasons of clarity, the evolution of w_ϕ was cut off at the end of inflation, after which it starts oscillating rapidly.

It is worth mentioning at this point, how the evolution of a universe with an equation of state parameter of *close to* -1 is very similar to the evolution in a universe dominated by a cosmological constant Λ with an equation of state parameter of *exactly* -1 (also referred to as a de Sitter universe and probably our future fate). Physically these two states are very different, though, the first being governed by a scalar field ϕ and the latter by what can be interpreted as a vacuum energy density ρ_Λ .

Figure 4.3 shows the evolution of the comoving Hubble horizon as a function of the logarithm of the scale factor. As expected it shrinks during inflation. However, during KD the comoving Hubble horizon initially grows until the onset of inflation where it meets the SR solution and starts decreasing. Thus, in a universe initially going through a phase of KD there exists a maximum to the comoving Hubble horizon and consequently there are very large scales that have never been within the horizon before the start of inflation.

4.3 Primordial Power Spectrum

For the evolution of the primordial perturbations we work directly with the primordial curvature perturbations \mathcal{R} and the tensor perturbations h as functions of cosmic time and for

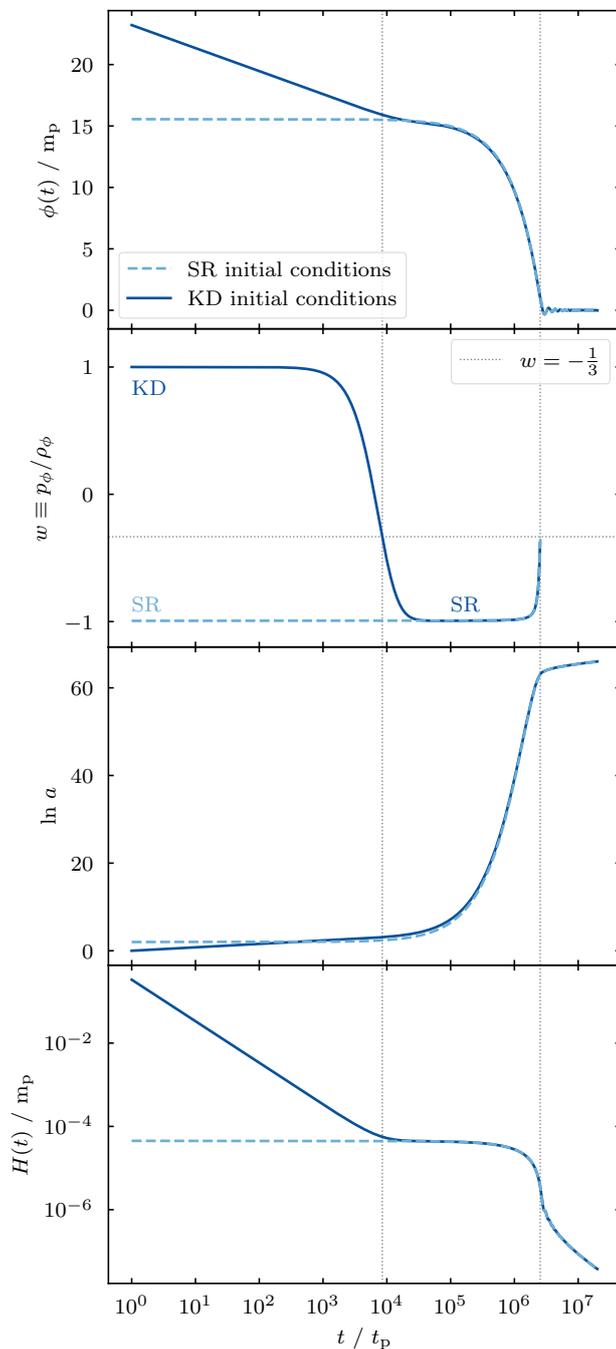


Figure 4.2 Evolution of the inflaton field $\phi(t)$, equation-of-state parameter $w_\phi(t)$, scale factor $a(t)$, and Hubble parameter $H(t)$ respectively for the quadratic potential from eq. (4.5). The inflaton mass was taken to be $m = 5 \times 10^{-6} m_p$. The light dashed line starts out directly in the slow-roll (SR) regime, whereas the dark solid line starts out during kinetic dominance (KD) and then later joins the SR attractor. The initial conditions were set such that $N_{\text{tot}} = 60$ e-folds of inflation are produced. The equation-of-state parameter w_ϕ is useful in determining the start and end of inflation in the KD case (dotted lines).

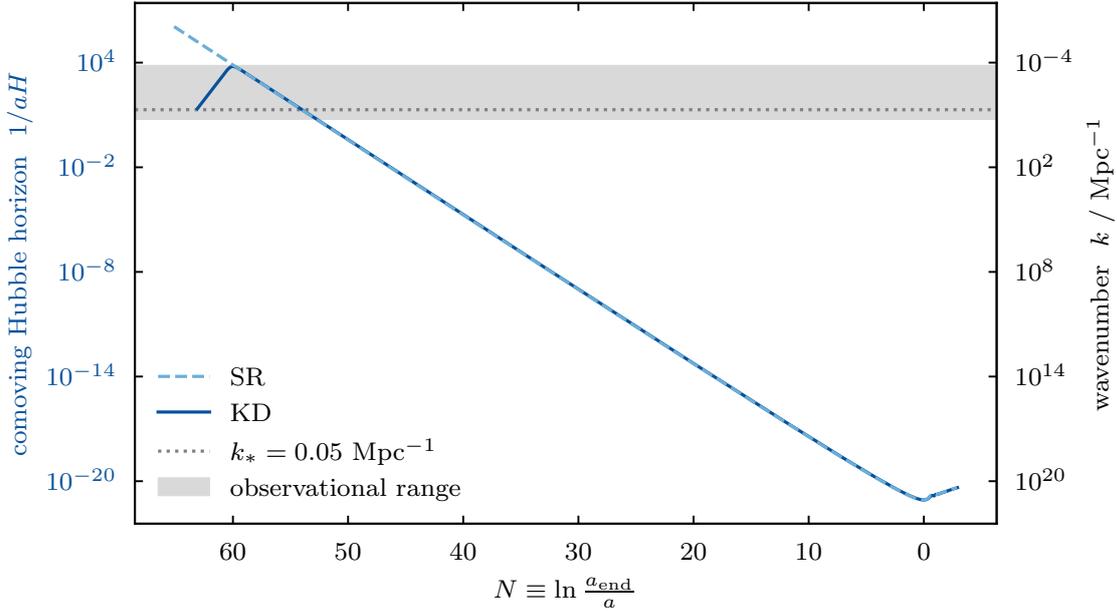


Figure 4.3 The comoving Hubble horizon is plotted here as a function of N . For the dashed line slow-roll (SR) initial conditions were used and for the solid line kinetic dominance (KD) initial conditions. Note how the comoving Hubble horizon falls monotonically for the SR solution, whereas it has a local maximum at the start of inflation for the KD solution, which is not evident in [fig. 4.2](#). Thus, in the KD case, there are scales $k^{-1} > (aH)^{-1}$ that were never within the comoving Hubble horizon.

a given mode k [[8](#), [155](#)]:

$$\ddot{\mathcal{R}}_k + \left(\frac{\dot{\phi}^2}{m_{\text{p}}^2 H} + \frac{2\ddot{\phi}}{\dot{\phi}} + 3H \right) \dot{\mathcal{R}}_k + \frac{k^2}{a^2} \mathcal{R}_k = 0, \quad (4.13)$$

$$\ddot{h}_k + 3H\dot{h}_k + \frac{k^2}{a^2} h_k = 0, \quad (4.14)$$

where the dot again refers to the derivative with respect to cosmic time.

For the numerical integration of the differential equations we loosely follow the scheme outlined in [[153](#), [155](#), [156](#)]. We reduce the differential equations into a first-order system and superimpose two orthogonal solutions. We start out by only evolving the background [eqs. \(4.1a\) to \(4.1c\)](#). To reduce computation time it is useful to delay the integration of the mode [eqs. \(4.13\) and \(4.14\)](#) until the comoving Hubble horizon has shrunk to a scale of about 100 times the size of a given mode. This therefore requires any given mode to be sufficiently well within the comoving horizon before horizon crossing which is the case for slow-roll (SR) initial conditions.

This approach also works for very *small* modes (large k) for kinetic dominance (KD) initial conditions. However, during KD the comoving Hubble horizon is still growing and reaches its maximum at the start of inflation (see [fig. 4.3](#)). Thus, for *large* modes (small k) it is impossible to start “well within” the comoving Hubble horizon $k^{-1} \ll (aH)^{-1}$. To maximise

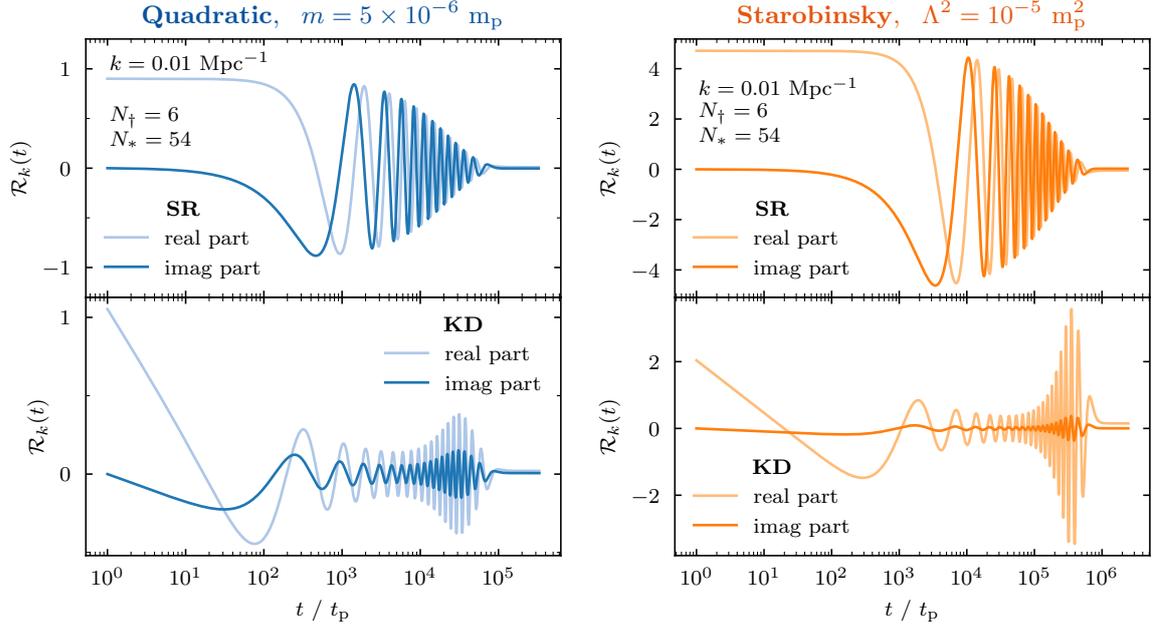


Figure 4.4 Evolution of the primordial curvature perturbations $\mathcal{R}_k(t)$ from eq. (4.13) for the given mode $k = 0.01 \text{ Mpc}^{-1}$ for quadratic inflation on the left with an inflaton mass of $m = 5 \times 10^{-6} m_p$ and for Starobinsky inflation with an amplitude of $\Lambda^2 = 10^{-5} m_p^2$. The background variables were set up using slow-roll (SR) initial conditions in the top plots and using kinetic dominance (KD) initial conditions in the bottom plots such that a total number of $N_{\text{tot}} = 60$ e-folds were produced. In terms of the quantities defined in section 4.3.1, they are split into $N_{\dagger} = 6$ e-folds *before* and $N_* = 54$ e-folds *after* horizon exit of the pivot scale $k_* = 0.05 \text{ Mpc}^{-1}$.

the number of modes starting within the comoving Hubble horizon we choose to start the evolution for those large modes synced up at the onset of inflation.

The initial conditions for the mode equations (note, these are not the same as the initial conditions for the inflaton, i.e. not SR or KD initial conditions) are set through the definition of the quantum vacuum. For SR initial conditions for the inflaton field, typically, the Bunch-Davies vacuum is chosen, which defines the quantum vacuum via Hamiltonian diagonalisation. For KD initial conditions, on the other hand, the vacuum choice becomes relevant, see e.g. [128, 157, 158]. In this paper we limit ourselves to the Bunch-Davies vacuum, leaving the exploration of alternative vacua to a later work.

We apply the Bunch-Davies vacuum on a linear combination of two orthogonal solutions. The real and imaginary parts of the curvature perturbation $\mathcal{R}_k(t)$ are plotted in fig. 4.4, using SR and KD initial conditions for the inflaton. For a good visualisation we use an inflaton mass of $m = 5 \times 10^{-6} m_p$ for the quadratic potential and an amplitude of $\Lambda^2 = 10^{-5} m_p^2$ for the Starobinsky potential, and the mode $k = 0.01 \text{ Mpc}^{-1}$. Higher k -values would result in increasingly more oscillations.

We read off the frozen values of the primordial perturbations after horizon exit and obtain the scalar and tensor power spectra

$$\mathcal{P}_{\mathcal{R}}(k) = \frac{k^3}{2\pi^2} |\mathcal{R}_k|^2, \quad (4.15)$$

$$\mathcal{P}_t(k) = 2 \cdot \mathcal{P}_h(k) = 2 \cdot \frac{k^3}{2\pi^2} |h_k|^2, \quad (4.16)$$

where the factor 2 in the tensor spectrum comes from the two possible polarization states of gravitational waves.

In order to compare our results to CMB data, we need to calibrate the perturbation scales. Calculations of the evolution of the universe from the end of inflation until today constrain the (observable) number of e-folds remaining during inflation after a given pivot scale k_* exited the Hubble horizon, to roughly within $50 \lesssim N_* \lesssim 60$ [5, 159, 160]. In accordance with Planck [131] we choose $k_* = 0.05 \text{ Mpc}^{-1}$ for our pivot scale. We then calibrate our k -axis by determining the value $a_* H_*$ (cf. fig. 4.3) for which N_* e-folds of inflation remain after horizon exit

$$k \mapsto \frac{aH}{a_* H_*} k_*. \quad (4.17)$$

In fig. 4.5 we have plotted the numerical solutions of the PPS for quadratic and Starobinsky inflation, each with SR and KD initial conditions for the inflaton. In agreement with eqs. (4.6) and (4.8) quadratic and Starobinsky inflation show a very similar spectral index n_s and a tensor-to-scalar ratio r differing by about two orders of magnitude. As expected, the choice of SR or KD initial conditions does not affect small scales, since smaller scales freeze out later in the inflationary history when the slow-roll approximation is fully applicable for both cases. For larger scales we see oscillations and a cutoff towards small k .

The existence of the cutoff can be attributed to the preceding kinetically dominated phase and the brief period of fast-roll inflation [120, 150]. The larger modes spent less time within the horizon and the largest modes have actually *never* been inside the horizon (scales greater than the maximum of the Hubble horizon in fig. 4.3).

With ongoing work we see that the amplitude and frequency of the oscillations as well as the slope of the cutoff depend on the choice of the quantum vacuum, and consequently on the initial conditions for the curvature perturbations. The cutoff position, however, is caused by the breakdown of the SR approximation and, thus, is independent of the vacuum choice. Since the cutoff position appears to be the driving force for model comparison between SR and KD, we expect the vacuum choice to only have negligible effects. However, the vacuum choice could become more relevant should the oscillations manage to sink into the dip in power at multipoles ℓ of approximately 20–25, but as mentioned earlier we leave the detailed exploration of alternative vacua such as those proposed in [157, 158] to a later work.

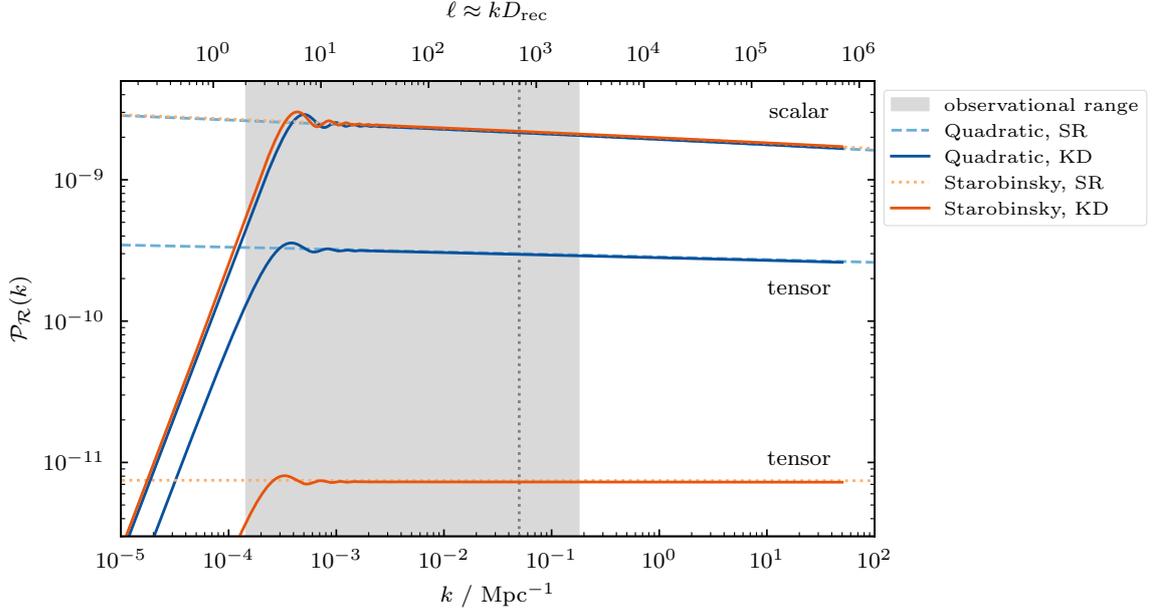


Figure 4.5 Primordial Power Spectra: The upper four lines correspond to the scalar power spectra, i.e. from primordial curvature perturbations. The lower four lines correspond to the tensor power spectra, i.e. from gravitational waves. On small scales the power spectra from slow-roll (SR) and from kinetic dominance (KD) initial conditions agree well with one another and show the characteristic power-law behaviour. Towards larger scales the SR power spectra continue along the power-law slope whereas the KD power spectra start oscillating and eventually show a cutoff. These very large scales are the ones that were never within the comoving Hubble horizon in an initially kinetically dominated universe (cf. fig. 4.3). The dotted vertical line marks the pivot scale $k_* = 0.05 \text{ Mpc}^{-1}$, used for the calculation of spectral index n_s and tensor-to-scalar ratio r as well as for the calibration of the k -axis. The parameters governing n_s , r and the cutoff position were set to the best-fit values from the MCMC analysis in section 4.4. The shaded region gives a rough estimate of the observational window in the CMB power spectrum.

4.3.1 Number of e-folds

The exact position of the cutoff in the PPS for KD initial conditions depends on the initial value ϕ_p in eq. (4.11a). This is also related to the number of e-folds *before* horizon crossing which we denote by N_{\dagger} as opposed to the e-folds N_* *after* horizon crossing. Together they make up the total number of inflationary e-folds

$$N_{\text{tot}} \equiv \ln \left(\frac{a_{\text{end}}}{a_{\text{start}}} \right) = N_{\dagger} + N_*. \quad (4.18)$$

It is very hard to *a-priori* constrain the total number of e-folds N_{tot} . Assuming inflation started after the Planck epoch, an upper bound on N_{tot} can be set. For a quadratic potential with a roughly realistic inflaton mass of $m = 5 \times 10^{-6} m_{\text{p}}$ such a bound is of an order of about $\max(N_{\text{tot}}) \sim 10^{10}$ e-folds [122].

Assuming the inflaton underwent a kinetically dominated phase before inflation, i.e. where $\dot{\phi} \gg V(\phi)$, we expect a significantly smaller number of e-folds, $N_{\text{tot}} \ll 10^{10}$ e-folds. Stronger claims on the total amount of inflation have been made in the context of “finite inflation” [124,

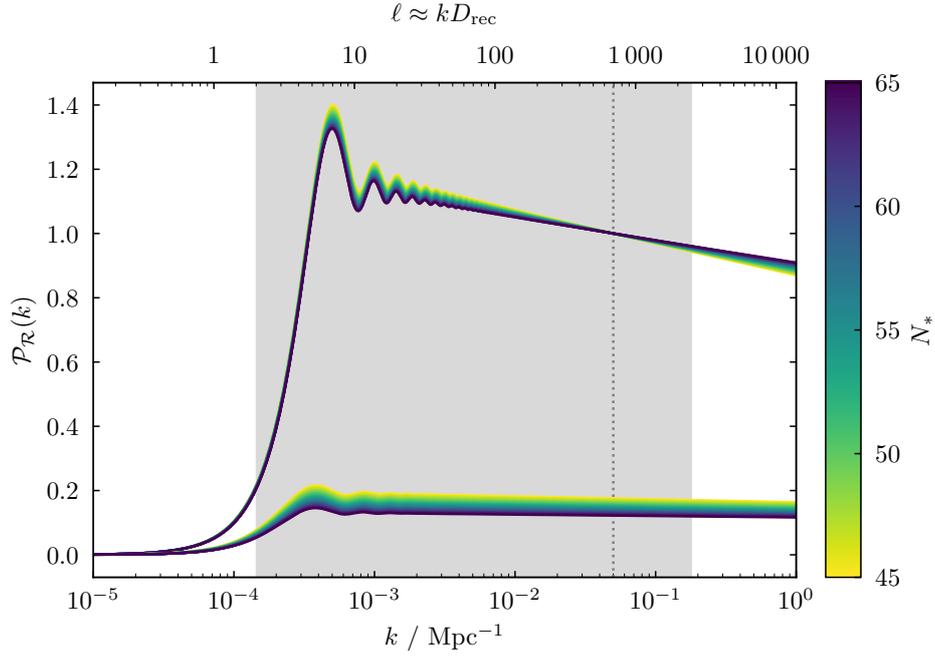
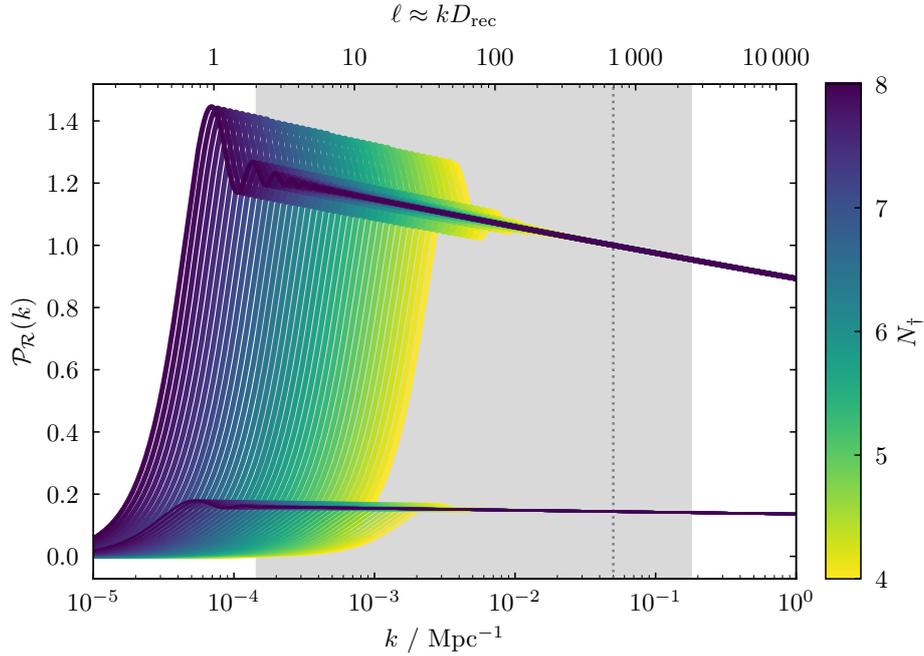
(a) $N_* \equiv$ e-folds of inflation *after* horizon crossing.(b) $N_{\dagger} \equiv$ e-folds of inflation *before* horizon crossing.

Figure 4.6 PPS for scalar (upper) and tensor (lower) perturbations. Plot (a) varies the number of observable e-folds N_* while keeping a fixed value of $N_{\dagger} = 6$ e-folds. As denoted in eqs. (4.6) and (4.8), N_* governs the spectral index n_s , i.e. the slope of the power spectrum, and the tensor-to-scalar ratio r . The larger N_* the larger also n_s and thus a smaller slope (closer to scale invariance). The tensor-to-scalar ratio on the other hand decreases with increasing N_* . Plot (b) varies the number of e-folds N_{\dagger} before horizon exit for a fixed value of $N_* = 55$ e-folds. Both n_s and r stay unaffected in this case. N_{\dagger} instead governs the low- k cutoff position of the PPS pushing it to ever smaller k -values as N_{\dagger} grows.

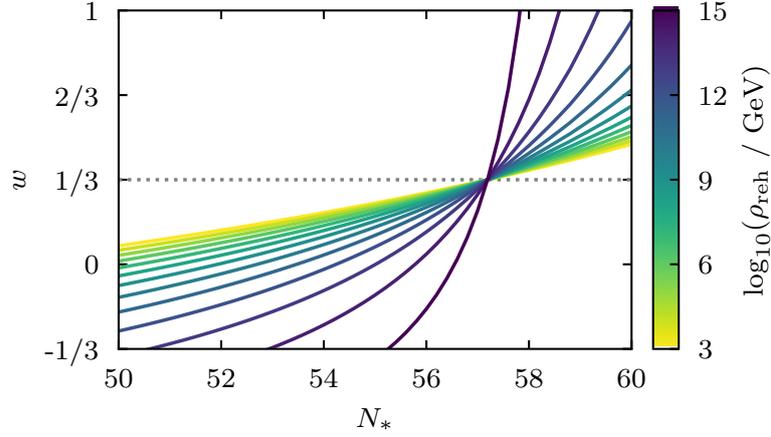


Figure 4.7 Equation-of-state parameter w after inflation as a function of N_* for a range of given values of ρ_{reh} . We get w here by inverting eq. (4.19) and using the quadratic potential and the best fit values specified below in table 4.1, and by setting $g_{\text{reh}} = 10^3$.

125] or “just enough inflation” [126, 127, 154], where $N_{\text{tot}} \gtrsim N_*$. Also, the expected amount of inflation can drop significantly depending on the choice of potential. While $\langle N_{\text{tot}} \rangle \sim 10^{10}$ for the quadratic potential, it can turn out to be as low as $\langle N_{\text{tot}} \rangle \sim 10^1$ or 10^2 for natural inflation depending on the symmetry breaking parameter f as shown in [122].

Figure 4.6 shows the effect of N_{\dagger} and N_* on the primordial power spectrum $\mathcal{P}_{\mathcal{R}}(k)$ for the quadratic potential. The behaviour is very similar for the Starobinsky potential with the major difference being a significantly smaller tensor-to-scalar ratio for the Starobinsky model as can already be inferred from eqs. (4.6) and (4.8). As those equations suggest, we find that N_* governs both the spectral index n_s and the tensor-to-scalar ratio r . On the other hand, N_{\dagger} leaves both these parameters invariant. Instead it shifts the cutoff position along the k -axis. More total e-folds N_{tot} , i.e. a longer period of inflation, and thus a larger N_{\dagger} pushes the cutoff to ever smaller k -values (larger scales). Thus, large scale CMB data will help us to constrain N_{\dagger} and N_{tot} .

Note that for SR initial conditions there is no clear start to inflation. One may therefore consider SR to correspond to the $N_{\dagger}, N_{\text{tot}} \rightarrow \infty$ limit of KD.

4.3.2 Equation of state parameter

The observable number of e-folds N_* can be related to the energy density of reheating ρ_{reh} and the effective equation of state parameter w from the end of inflation to reheating via the matching equation [159–162]:

$$\begin{aligned}
 N_* = & 55.99 - \ln \frac{10^{16} \text{ GeV}}{V_*^{1/4}} + \ln \frac{V_*^{1/4}}{V_{\text{end}}^{1/4}} - \frac{1}{12} \ln g_{\text{reh}} \\
 & - \frac{1-3w}{3(1+w)} \ln \frac{V_{\text{end}}^{1/4}}{\rho_{\text{reh}}^{1/4}} + \frac{1}{3(1+w)} \ln \frac{2}{3},
 \end{aligned} \tag{4.19}$$

where V_* and V_{end} correspond to the inflaton potential at horizon crossing of the pivot scale and at the end of inflation respectively, and g_{reh} is the effective number of degrees of freedom at reheating. To determine the numerical value in the first term we here have used the CMB temperature $T_0 = 2.725 \text{ K}$, the pivot scale $k_* = 0.05 \text{ Mpc}^{-1}$, and today's effective number of degrees of freedom $g_0 = 43/11$. [Figure 4.7](#) shows the relationship between w and N_* for given ρ_{reh} . Post-inflationary oscillations of the inflaton in a quadratic potential can be associated with an equation of state parameter of $w = 0$ [[163](#)]. For the Starobinsky potential we similarly expect $0 \lesssim w < 1/3$.

4.4 CMB power spectrum

To translate the primordial power spectra (PPS) from [eqs. \(4.15\) and \(4.16\)](#) through to the angular power spectrum of the cosmic microwave background (CMB) we make use of the Boltzmann solver `CAMB` [[164–166](#)], which we modify such that it takes our PPS. To that end we first modify our input PPS such that they are normalised at the pivot scale k_* and the desired amplitude is then given by the `CAMB` parameter A_s

$$\mathcal{P}_{\mathcal{R}}(k) \mapsto A_s \cdot \frac{\mathcal{P}_{\mathcal{R}}(k)}{\mathcal{P}_{\mathcal{R}}(k_*)}. \quad (4.20)$$

We can do this, because the background [eqs. \(4.1a\) to \(4.1c\)](#) are invariant under a simultaneous rescaling of the time coordinate and the inflaton potential

$$t \mapsto \sigma^{-1}t, V(\phi) \mapsto \sigma^2 V(\phi) \Rightarrow \mathcal{P}_{\mathcal{R}}(k) \mapsto \sigma^2 \mathcal{P}_{\mathcal{R}}(k), \quad (4.21)$$

effectively making a substitution of $\sigma^2 = 1/m^2$ or $\sigma^2 = 1/\Lambda^4$ to get a PPS $\mathcal{P}_{\mathcal{R}}(k, \sigma)$ independent of the potential amplitude. The PPS amplitude can then be linked to any desired mass m_0 or amplitude Λ_0 through $A_s = m_0^2 \mathcal{P}_{\mathcal{R}}(k, \sigma)$. The same results are obtained using the alternative Boltzmann solver `Class` [[105–107, 167–169](#)].

[Figure 4.8](#) shows the CMB angular temperature power spectrum for the Planck data [[64](#)], for its Λ CDM best-fit model [[35](#)], and for the quadratic inflation model with kinetic dominance (KD) initial conditions. The characteristic features of the KD initial conditions: low- k cutoff and oscillation, are still apparent although diluted from convolution with the transfer functions. As for the PPS, the cutoff position depends on the number of e-folds before horizon exit N_{\dagger} . For a sufficiently small value, the cutoff sinks into the low- ℓ lack of power found in the Planck data. The oscillations, however, are too heavily smoothed to follow the dip at multipoles ℓ at approximately 20–25. This is in line with the findings in [[130](#)].

4.5 MCMC analysis

We performed a Markov chain Monte Carlo (MCMC) analysis to extract the cosmological parameters of extended Λ CDM models alongside the kinetic dominance (KD) initial conditions. To that end we used `CAMB`'s MCMC extension `CosmoMC` [[102, 103](#)] in conjunction with Planck's

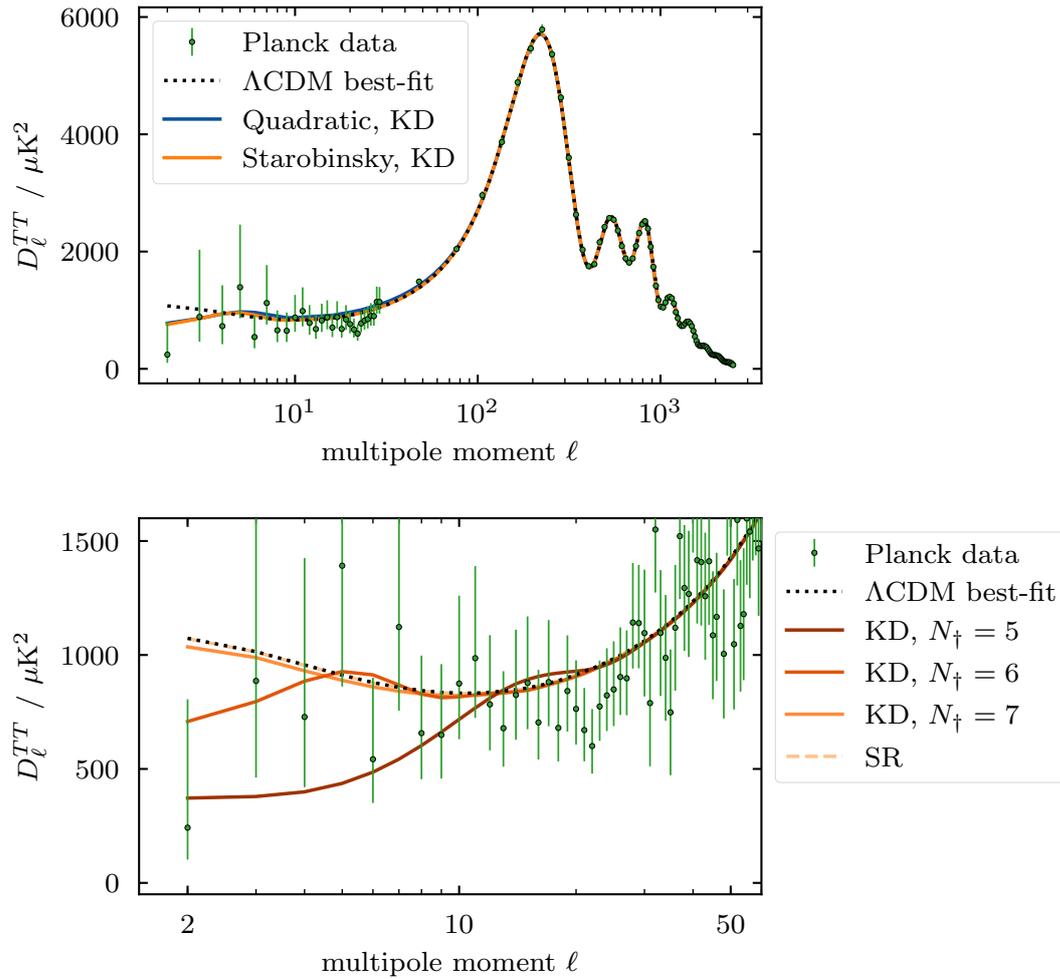


Figure 4.8 CMB angular power spectrum $D_{\ell}^{TT} \equiv \ell(\ell+1)\mathcal{C}_{\ell}^{TT}/(2\pi)$: The upper panel shows the best-fit lines (obtained from the MCMC analysis in [section 4.4](#)) for the Λ CDM model, and the quadratic and Starobinsky model with kinetic dominance (KD) initial conditions respectively. The lower panel zooms in on the low- ℓ region and shows how the cutoff and oscillations from the PPS have translated through to the CMB power spectrum, where the cutoff position still depends on the value of N_{\dagger} . Large values become indistinguishable to the slow-roll (SR) case as the cutoff moves out of the observable region.

temperature and low- ℓ polarisation data (TT+lowP) and corresponding likelihood code [64]. Additionally we perform a model comparison using `CosmoChord` which is a `PolyChord` [95, 96] plug-in for `CosmoMC`. `PolyChord` is a Bayesian inference tool for the simultaneous calculation of evidences and sampling of posterior distributions, and allows us to calculate the Bayes' factor of models. It performs well even on moderately high-dimensional posterior distributions, and can cope with arbitrary degeneracies and multi-modality. As such it is the successor to `MultiNest` [92–94], a variation of Nested Sampling [170].

For our parameter estimation we added N_{\dagger} and N_* as new parameters in place of n_s . We put a flat prior within the range of $50 < N_* < 60$ in accordance with the expected number of observable e-folds [5, 131, 160, 171], which roughly covers the range from $w = 0$ to $w = 1/3$ even for low values of ρ_{reh} . For N_{\dagger} we chose a range from 4 to 15. We choose to cut values greater than $N_{\dagger} = 15$ e-folds as the PPS becomes observationally indistinguishable from the slow-roll (SR) case. We retained the amplitude parameter A_s to multiply our normalised PPS by, as already detailed in eq. (4.20). With these three parameters in place, the PPS is fully parametrised. Both the spectral index n_s and the tensor-to-scalar ratio r turn into derived parameters inferred from the input PPS (cf. fig. 4.6). The remaining standard cosmological parameters were varied as for the Λ CDM case, namely the baryon density parameter $\Omega_b h^2$, the mass density parameter $\Omega_c h^2$, the optical depth τ , and the ratio of the sound horizon to the angular diameter distance θ_{MC} . Figure 4.9 shows a triangle plot (created using `GetDist` [108]) of all these parameters and tables 4.1 and 4.2 list the means of the marginalised parameters and their uncertainties.

The models considered are the standard Λ CDM model, $r\Lambda$ CDM which is a one-parameter extension by the tensor-to-scalar ratio r , and the quadratic and Starobinsky inflation models each with SR and KD initial conditions.

4.5.1 Posteriors and priors on model parameters

We begin by considering the constraints on the spectral index n_s and tensor-to-scalar ratio r , detailed in the third and fourth rows and columns of fig. 4.9, and highlighted in fig. 4.10. We plot only the $r\Lambda$ CDM model and the quadratic and Starobinsky model with KD initial conditions in Figures 4.9 and 4.10 as the Λ CDM model and the SR inflation models are visually very similar to their counterparts for the shared parameters. The major difference lies in the additional parameters: the tensor-to-scalar ratio r for $r\Lambda$ CDM and N_{\dagger} for the KD inflation models.

As expected, we find significant differences for the amount of tensor modes, as r is significantly larger for the quadratic inflation model than for the Starobinsky model, and larger even than the 68% upper bound of the $r\Lambda$ CDM model.

Both inflationary models exhibit cut-off effects in their posterior contours. This is due to the relationship between N_* , n_s and r from eqs. (4.6) and (4.8). The flat prior on N_* leads to an induced prior on n_s and r that is much narrower than the traditional Λ CDM or $r\Lambda$ CDM priors. This constraint is then projected onto the other parameters.

Table 4.1 Marginalized parameter values at 68 % limits for different models and using CMB temperature data with low- ℓ polarization (TT+lowP). For the number of e-folds N_{\dagger} before and N_* after horizon exit we additionally provide the best-fit values as the data does not clearly delimit these parameters and they differ considerably from the 68 % limits. The posterior distribution for N_{\dagger} is essentially flat causing the mean to fall roughly in the middle of the defined prior range. However, there is a peak at small values causing the different best-fit value (cf. fig. 4.9). N_* is driven to high values for the quadratic model due to its correlation with the tensor-to-scalar ratio r , and rather unconstrained for the Starobinsky model.

Parameters	N_{\dagger}		N_*		n_s	r
	[4, 15]		[50, 60]		[0.885, 1.040]	[0, 1]
TT+lowP	68 % limits	best-fit	68 % limits	best-fit	68 % limits	68 % limits
Λ CDM					0.9655 ± 0.0063	
$r\Lambda$ CDM					0.9665 ± 0.0062	< 0.0504
Quadratic, SR			> 55.4	60.00	$0.9641^{+0.0022}_{-0.00069}$	$0.1425^{+0.0027}_{-0.0086}$
Quadratic, KD	$9.8^{+3.4}_{-4.4}$	6.02	> 55.4	57.48	$0.9642^{+0.0022}_{-0.00067}$	$0.1407^{+0.0026}_{-0.0085}$
Starobinsky, SR			—	57.98	$0.9644^{+0.0028}_{-0.0013}$	$0.00365^{+0.00026}_{-0.00055}$
Starobinsky, KD	> 7.92	6.09	—	57.07	$0.9649^{+0.0027}_{-0.0011}$	$0.00356^{+0.00021}_{-0.00053}$

Table 4.2 Marginalized parameter values at 68 % limits for different models and using temperature data with low- ℓ polarization (TT+lowP). Note, how the parameter values stay relatively similar across different models while the errors go down for models with explicit inflationary models (quadratic and Starobinsky) which can also be seen in the narrower contours in fig. 4.9. This is attributed to the prior on N_* setting an effective, very narrow prior on the spectral index n_s and the tensor-to-scalar ratio r . Comparing slow-roll (SR) and kinetic dominance (KD) models, we find that these parameters do not distinguish between them at all.

Parameters	$\Omega_b h^2$	$\Omega_c h^2$	τ	H_0	$\ln(10^{10} A_s)$
Prior ranges	[0.019, 0.025]	[0.095, 0.145]	[0.01, 0.40]	[1.03, 1.05] on $100\theta_{MC}$	[2.5, 3.7]
TT+lowP	68 % limits	68 % limits	68 % limits	68 % limits	68 % limits
Λ CDM	0.02223 ± 0.00023	0.1197 ± 0.0022	0.078 ± 0.019	67.3 ± 1.0	3.089 ± 0.037
$r\Lambda$ CDM	0.02224 ± 0.00023	0.1195 ± 0.0022	0.076 ± 0.019	67.42 ± 0.98	3.086 ± 0.036
Quadratic, SR	0.02215 ± 0.00019	0.1203 ± 0.0013	0.067 ± 0.016	67.02 ± 0.57	3.069 ± 0.032
Quadratic, KD	0.02216 ± 0.00019	0.1202 ± 0.0013	0.068 ± 0.016	67.04 ± 0.57	3.071 ± 0.032
Starobinsky, SR	0.02222 ± 0.00019	0.1201 ± 0.0013	0.076 ± 0.016	67.18 ± 0.56	3.088 ± 0.033
Starobinsky, KD	0.02222 ± 0.00019	0.1199 ± 0.0013	0.077 ± 0.017	67.24 ± 0.57	3.089 ± 0.034

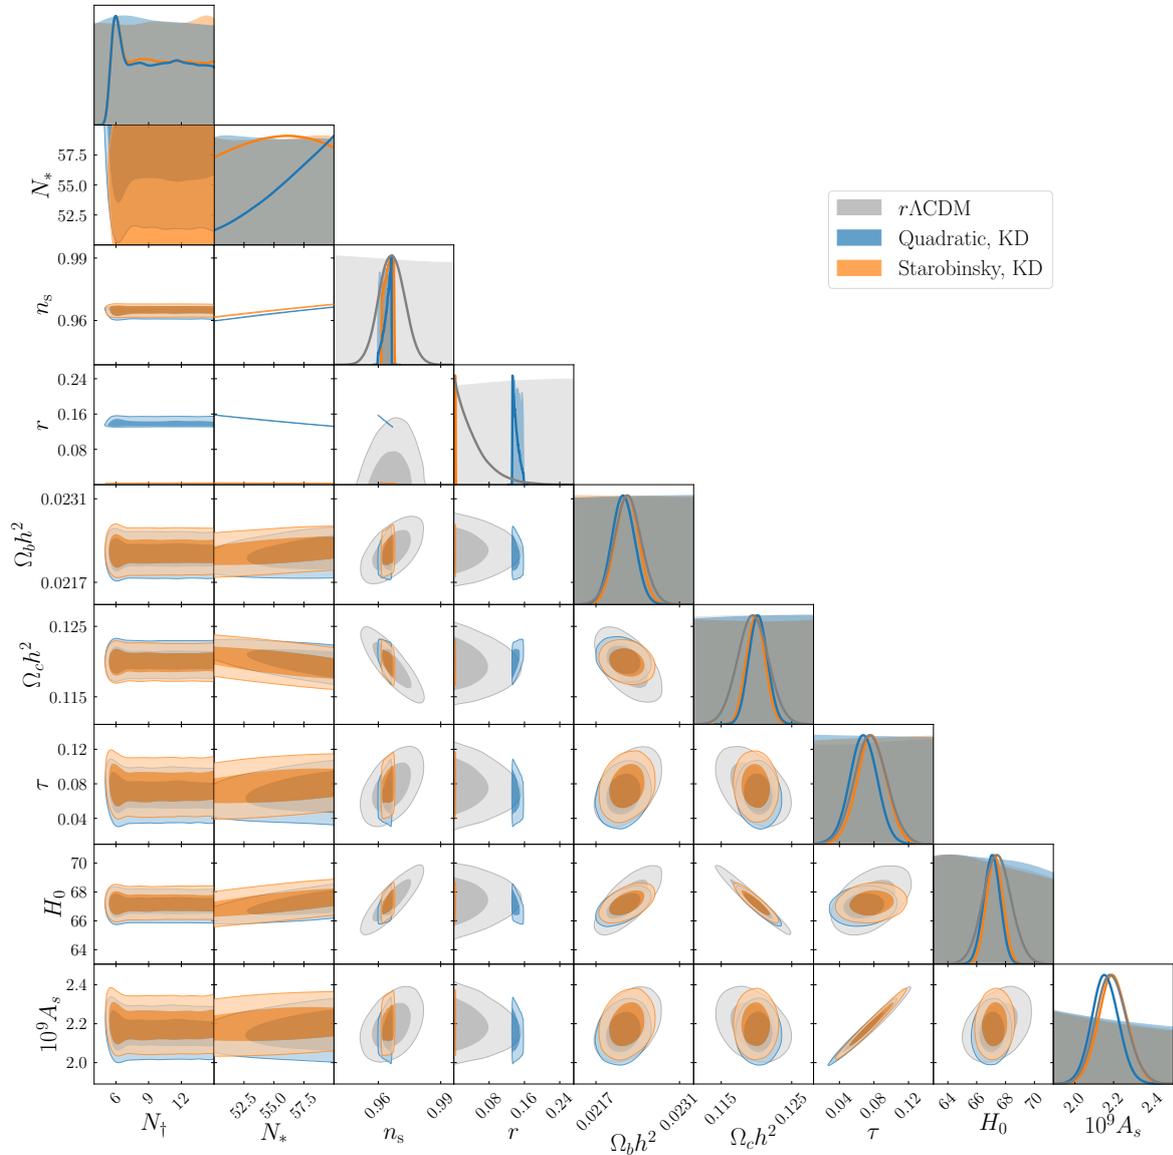


Figure 4.9 Triangle plot of the parameters: The number of e-folds N_{\dagger} before and N_* after horizon exit, the spectral index n_s , the tensor-to-scalar ratio r , the baryon density parameter Ω_b , the mass density parameter Ω_c , the optical depth τ , the Hubble parameter H_0 , and the primordial amplitude of scalar perturbations A_s . The contours delimit the 0.68, and 0.95 levels. The shaded regions in the 1D plots on the diagonal correspond to the flat priors on the input parameters or the derived parameters in the case of n_s , r , H_0 and A_s . Note that the apparent darker shade of grey is really just the overlay of the priors of all models. The thicker solid lines in the 1D plots are the posterior distributions. [Figure 4.10](#) shows the n_s - r -plane separately and enlarged. For the mean and standard deviation of marginalised parameters see [tables 4.1](#) and [4.2](#).

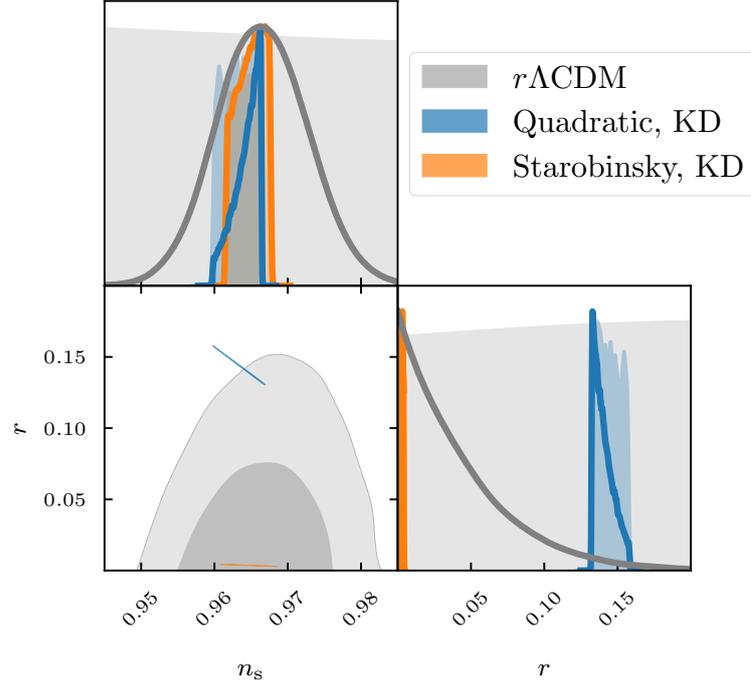


Figure 4.10 Spectral index n_s vs tensor-to-scalar ratio r triangle plot (zoomed in from [fig. 4.9](#)). The shaded regions in the 1D plots denote the flat input prior for $r\Lambda\text{CDM}$ and the derived priors for the inflationary models. The thick solid lines are the posteriors.

Given this *a-priori* predictivity in r and n_s , one might object at this point that the prior range chosen for N_* is too narrow. However, the broad prior ranges for n_s and r in the $r\Lambda\text{CDM}$ model may be viewed as a phenomenological model-averaging over a wide class of inflationary models. It allows $r\Lambda\text{CDM}$ to represent and compare many inflation models in an n_s - r -plot ([fig. 4.10](#)). Thus, it is only natural that specific models give narrower priors on parameters such as the spectral index or the tensor-to-scalar ratio, and it is this which eventually allows the falsification of different inflationary models.

Consider now the marginalised posteriors involving the number of e-folds before and after horizon exit (N_{\dagger} and N_*), detailed in the first and second rows and columns of [fig. 4.9](#), and best-fit values in [tables 4.1](#) and [4.2](#).

Neither N_* nor N_{\dagger} are clearly constrained for either model. For quadratic inflation, N_* is driven to high values in order to decrease the tensor-to-scalar ratio r and thus we get a lower bound for the 68% limits. For the Starobinsky model, N_* is essentially only constrained through the prior choice which was here taken to be $50 < N_* < 60$. A small amount of constraining power comes from the correlation with n_s . N_{\dagger} on the other hand behaves very similarly for both inflation models. While very low values are clearly ruled out by the data, the posterior plateaus for larger values, the exception being a single peak at about $N_{\dagger} = 6$ roughly a factor 2 above the plateau. Low values will push the power spectrum cutoff unfavourably far into the data. The best-fit value manages to position the cutoff such that it aligns with the low- ℓ lack of power. Once the cutoff is pushed out of the observable region, KD is equivalent

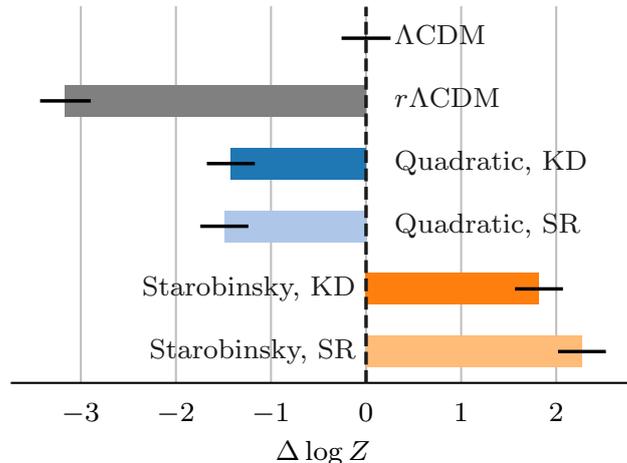


Figure 4.11 Difference of log evidences $\Delta \log Z$ with respect to the Λ CDM model as reference. The evidence Z follows a log-normal distribution with location parameter given by the plotted value and scale parameter given by the short black line. The errors are roughly a quarter log unit throughout.

to SR, there is no change to the CMB power spectrum, and all large values of N_{\dagger} become equally likely.

Finally, from the remaining rows and columns of [fig. 4.9](#), and [tables 4.1](#) and [4.2](#) one can see that all of the standard cosmological parameters ($\Omega_b h^2$, $\Omega_c h^2$, τ , H_0 , A_s , n_s) for all additional models are consistent with the values obtained for the standard Λ CDM model.

4.5.2 Evidences

From [fig. 4.10](#) we already judged the Starobinsky model to perform better than the quadratic model, since the line for the Starobinsky model sits in the middle of the $r\Lambda$ CDM contour, whereas the line for the quadratic model lies on the outer edge of the 95% contour. For a proper model comparison we calculate and compare their respective Bayesian evidences. Using `CosmoChord` we calculated the evidences $Z \equiv P(\mathcal{D}|\mathcal{M})$ for a given model \mathcal{M} using the Planck data \mathcal{D} . [Figure 4.11](#) visualizes the Bayes' factors, i.e. the difference of log evidences $\Delta \log Z$ where we use the Λ CDM model as a reference model. The prior ranges used in the model comparison are listed in [tables 4.1](#) and [4.2](#).

Comparing Λ CDM and $r\Lambda$ CDM shows the effect a single additional parameter can have. Though $r\Lambda$ CDM has an additional parameter and thus can make a greater variety of predictions, it also has to spread its predictive probability over a greater volume of parameter space and thus more thinly. This penalizes $r\Lambda$ CDM considerably. For the comparison here we have chosen a prior range of $r \in [0, 1]$, which reads as an assumption that the tensor modes are smaller than the scalar modes.

As in the standard Λ CDM model, the SR models vary a total of six parameters. One of those parameters N_* replaces the spectral index n_s , which becomes a derived parameter (as does the tensor-to-scalar ratio r). The inflation models with KD initial conditions introduce one additional input parameter N_{\dagger} , resulting in a total of seven parameters.

As expected the quadratic model is disfavoured compared to the Starobinsky model with a difference of about 3 log units, mainly driven by the high tensor-to-scalar ratio r in the quadratic model. Due to their reduced parameter space, or equivalently their increased predictivity, they both outperform the very general $r\Lambda$ CDM model. Only the Starobinsky model with its very low tensor modes manages to do better than the standard Λ CDM model, which effectively conditions r to be zero. When comparing SR initial conditions to KD initial conditions the data do not show a clear preference towards one model or the other. Note again that we are investigating the Bayesian evidence here, i.e. a measure that takes into account both fit and Occam penalty of a model. Considering that an additional parameter is used for the KD case, the model manages to make up for the associated Occam penalty factor with a slightly better fit to the data. Phrased the other way around, the better fit that easily might be expected when applying a “chi-by-eye” to [fig. 4.8](#) is being penalised for the needed additional parameter, such that ultimately there is no clear preference for SR or KD.

One could question whether increasing the upper limit for the prior on N_{\dagger} would dilute the Bayesian evidence for the KD cases. However, one can show that if the posterior density levels off to a constant value at some point as for N_{\dagger} here, then the evidence will also be constant when varying the upper limit of a flat prior from that point onward, which we also confirmed numerically. In other words the Bayesian evidence ratio only penalises models for *constrained* parameters (see also [chapter 2](#) and [\[98\]](#)). From [figs. 4.6b](#) and [4.8](#) we can indeed expect the plateau in the posterior of N_{\dagger} to stretch to infinity in this case. Thus, the Bayesian evidence will not change for a higher upper limit on N_{\dagger} .

4.5.3 Power spectrum predictive posteriors and Kullback–Leibler divergences

The major observable differences between the SR and the KD cases are the low- ℓ cutoff and oscillations in the power spectrum. In the upper half of [fig. 4.12](#) we show the prior and posterior densities of MCMC samples for both PPS and CMB power spectra for the Starobinsky model with KD initial conditions. The low- k and low- ℓ cutoff from KD is not pushed out by the data but stays at the lower end of the observable region. We calculated the relative entropy or Kullback-Leibler divergence D_{KL} going from the prior distribution to the posterior distribution (bottom plots in [fig. 4.12](#)). While the information gain throughout most of the spectrum is rather high and roughly constant, it drops off to roughly a fourth of its value towards the largest observable scales due to cosmic variance.

[Figure 4.13](#) additionally includes the divergence for the quadratic model and for SR initial conditions. The quadratic model shows a higher information gain than the Starobinsky model, which is most prominent for the tensor modes of the PPS. This is related to the tensor-to-scalar ratio being driven to small values. Assuming the quadratic model was the correct model, one knows that N_* would need to be high in order to get a sufficiently low tensor-to-scalar ratio. The higher information gain at large scales in case of SR initial conditions is attributed to the rigidity of the model. Assuming this time that SR initial conditions are correct, the data

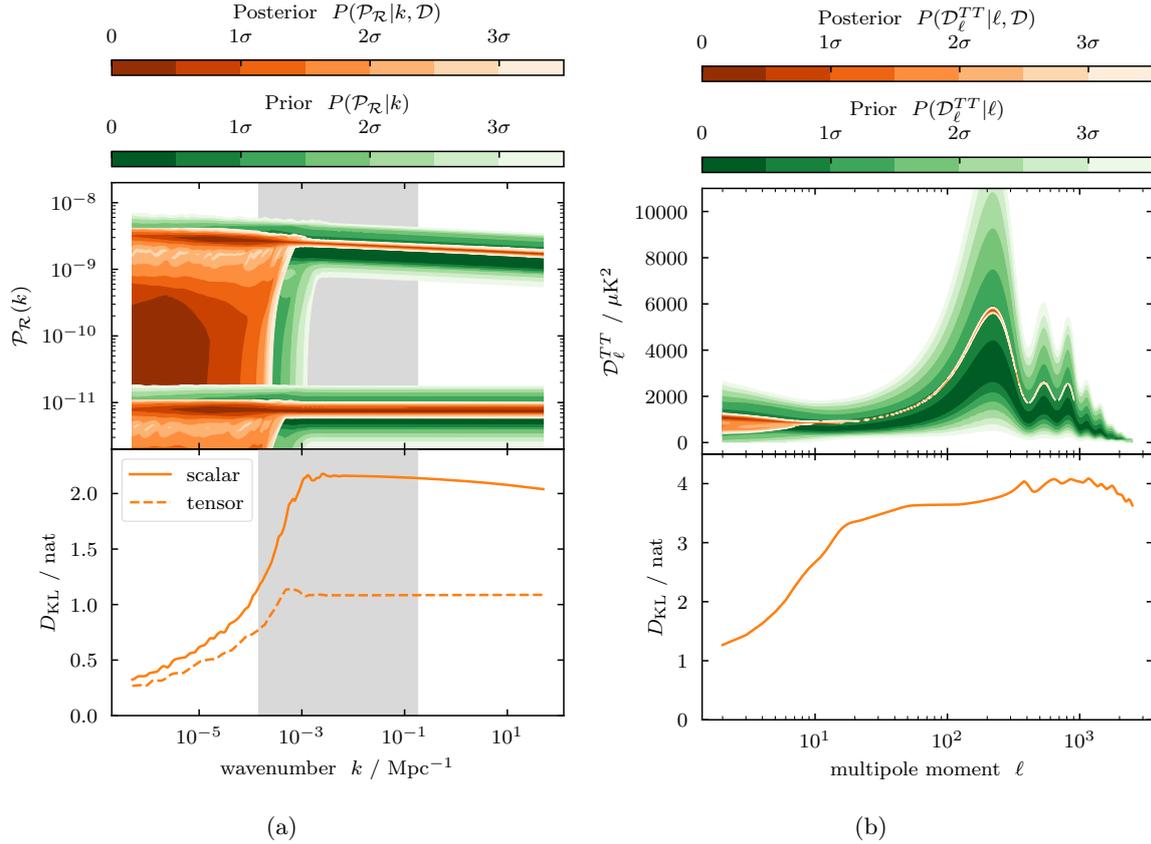


Figure 4.12 The upper panels show density plots of the primordial power spectra for scalar $\mathcal{P}_{\mathcal{R}}(k)$ and tensor $\mathcal{P}_t(k)$ perturbations in (a) and the CMB angular temperature power spectrum $\mathcal{D}_{\ell}^{TT} \equiv \ell(\ell+1)\mathcal{C}_{\ell}^{TT}/(2\pi)$ in (b) for parameter samples from the prior (green) and posterior (orange) distributions of the MCMC runs for the Starobinsky model. The lower panels show the corresponding plots for the relative entropy or Kullback-Leibler divergence D_{KL} when going from the prior to the posterior distribution. The data \mathcal{D} are very constraining for large wavenumbers $k \gtrsim 10^{-3} \text{ Mpc}^{-1}$ and multipoles $\ell \gtrsim 10$, and drive up the information gain accordingly in those domains. From there the relative entropy plummets to roughly a fourth its previous value towards larger scales reflecting the lack of constraining power of the data. This is where the power spectrum cutoff can sink in. (Figure created using `fgivenx` [172].)

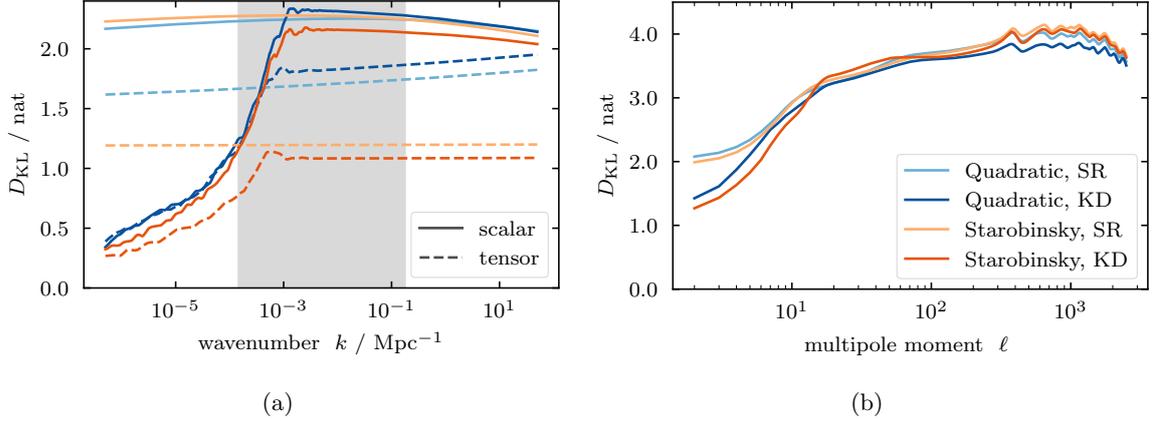


Figure 4.13 Kullback-Leibler divergence comparing quadratic vs Starobinsky inflation, and KD vs SR initial conditions. Plot (a) shows the Kullback-Leibler divergence for the PPS of scalar $\mathcal{P}_{\mathcal{R}}(k)$ and tensor $\mathcal{P}_t(k)$ perturbations. Plot (b) shows the Kullback-Leibler divergence for the CMB angular temperature power spectrum $\mathcal{D}_\ell^{TT} \equiv \ell(\ell + 1)\mathcal{C}_\ell^{TT}/(2\pi)$.

constrain the amplitude at small scales and the SR model then tells us that there must be a similar amplitude at large scales.

4.6 Conclusions

We have shown that using kinetically dominated (KD) initial conditions instead of slow-roll (SR) initial conditions for homogeneous and isotropic single-field inflation causes oscillations and a cutoff towards large scales in the primordial power spectrum (PPS). The position of oscillations and cutoff is governed by the amount of inflation N_{\dagger} preceding horizon exit for any given pivot mode. The amount of inflation N_* after horizon exit determines the scalar spectral index n_s and the tensor-to-scalar ratio r .

We illustrate how these features carry through to the CMB power spectrum, where the cutoff in the PPS can sink into the low- ℓ lack of power in the CMB. The oscillations get washed out going from the PPS to the CMB such that they are not strong enough to model the dip in CMB power at multipoles ℓ of approximately 20–25.

We perform an MCMC analysis and find that all standard cosmological parameters (Ω_b , Ω_c , τ , H_0 , A_s , n_s) for all the models taken into consideration ($r\Lambda$ CDM, quadratic inflation with SR and KD initial conditions, Starobinsky inflation with SR and KD initial conditions) are consistent with the standard Λ CDM model. As expected, we find significant differences for the amount of tensor modes, favouring Starobinsky over quadratic inflation. Both the e-folds N_{\dagger} and N_* cannot be clearly estimated. The amount of inflation before horizon exit can be constrained from below and shows a peak at about $N_{\dagger} = 6$ e-folds. From there it rapidly drops off to about half the peak amplitude and plateaus. This reflects that KD initial conditions are indistinguishable from SR initial conditions for large values of N_{\dagger} . The amount of observable inflation N_* is essentially unconstrained, hence any constraints are mostly driven by the choice of prior.

In a model comparison the Starobinsky model performs better and the quadratic model worse than the standard Λ CDM model. They both perform significantly better than the $r\Lambda$ CDM model. Although we do not find a significant difference between the use of SR or KD initial conditions in terms of evidence, it is intriguing that the KD model manages to balance the penalty for an additional parameter with a slightly improved fit for small N_{\dagger} at low multipoles, due to its effect on the overall power level in this region.

Finally, in an analysis of the posterior density and the Kullback-Leibler divergence, we confirm that most of the information gain from the data happens on small scales, i.e. for large multipoles. It will be interesting to consider in future work whether including large scale polarisation data from a future cosmic variance limited CMB experiment can help to discriminate more definitively between SR and KR conditions in terms of their effects on low- ℓ CMB power.

Chapter 5

Finite inflation in curved space

In this chapter, we consider the effects of non-zero spatial curvature on cosmic inflation in the light of cosmic microwave background (CMB) anisotropy measurements from the Planck 2018 legacy release (P18) and from BICEP2 and the Keck array (BK15). Even a small percentage of non-zero curvature today would significantly limit the total number of e-folds of the scale factor during inflation, rendering just-enough inflation scenarios with a kinetically dominated stage preceding slow-roll inflation more likely.

5.1 Introduction

Measurements of the cosmic microwave background (CMB) radiation anisotropies with the Planck satellite [16, 21, 22] have allowed us to refine our view of the Universe to unprecedented precision. This has led to what can be viewed as the current standard model of cosmology, called Λ CDM for the two dominant contributions to the overall energy density today: a cosmological constant Λ and cold dark matter (effectively collisionless with no electromagnetic interactions). The Λ CDM model is characterised through six free parameters, most of which can by now be given to percent-level precision [33–35]. It assumes the Friedmann–Lemaître–Robertson–Walker (FLRW) metric [49–54] with flat spatial geometry.

Cosmic inflation was originally [55–58] developed as a mechanism to explain the observed homogeneity and flatness of our Universe and is the currently preferred description of the primordial universe. In the Λ CDM model it is summarised in the form of two parameters, the scalar power amplitude A_s and spectral index n_s , characterising a power-law primordial power spectrum (PPS) of density anisotropies. The small deviation from unity of the scalar spectral index, and thus from scale invariance of the PPS, has been confirmed to 8σ precision and is one of the prime successes of cosmic inflation [33]. In its simplest form, the accelerated expansion of the Universe during cosmic inflation is driven by a single scalar field ϕ with a standard kinetic term that slowly rolls down a potential $V(\phi)$. This slow-roll scenario in a flat Universe has been investigated extensively for various inflationary potentials [131, 148, 173–178].

Despite the success of flat Λ CDM there has been a persistent preference for positive curvature (closed universes) in CMB temperature and polarisation data [33–35, 134, 151,

179, 180]. The latest data release from the Planck satellite [33] in particular has sparked discussion around evidence for spatial curvature in the CMB and about a possible discordance or tension with measurements from other sources such as baryon acoustic oscillations (BAO) or luminosity distance data, e.g. type Ia supernovae or Cepheid variables [27–29, 180, 181].

Non-zero spatial curvature affects the CMB anisotropy spectrum on two levels. First, curvature terms in the Boltzmann equations will modify the transfer functions, which encode the evolution of linear perturbations through the standard Big Bang epochs of radiation, matter and Λ domination until today. This makes up the curvature effect, which is commonly studied through an extension of the base Λ CDM model with a variable curvature density parameter $\Omega_{K,0}$ as mentioned in the abstract. Note that we use the subscript 0 to refer to present-day quantities. Second, a detection of present-day non-zero curvature would strongly limit the total amount of slow-roll inflation (measured in form of e-foldings of the scale factor a) and thereby would affect the PPS, particularly on large scales [123]. This could explain some of the unexplained features in the CMB angular power spectra, such as the lack of power on the largest scales [182]. Additionally, non-zero curvature renders scenarios of finite inflation more likely, including those with a phase of fast-roll inflation or kinetic dominance (KD) (where the inflaton’s kinetic energy dominates over its potential energy) preceding slow-roll. See e.g. [116, 118] for early accounts on the generality of slow-roll inflation as an attractor solution and possible preceding stages of KD and fast-roll with and without curvature. Other motivations for KD or fast-roll include holographic bounds [124, 125, 183, 184] or certain potentials that render a preceding phase of KD or fast-roll more likely [I & chapter 3] or that predict fewer e-folds of inflation [122]. This form of a finite amount of inflation is often also referred to as just-enough inflation [185]. The consequences of a preceding KD or fast-roll stage have been mostly studied assuming a flat cosmology, see e.g. [II & chapter 4, 115, 120, 121, 126–130, 150, 154, 186–189]. There have been a few previous studies of the large scale curvature effects on the PPS and how they translate to the CMB anisotropy spectra [123, 190–193], but these did not go beyond a phenomenological study of some specific parameter combinations.

In this paper we build on the Λ CDM extension with a non-zero curvature density parameter $\Omega_{K,0}$, which already accounts for curvature effects on the transfer functions of the Boltzmann equations of cosmology. We further investigate the other implications of non-zero curvature on cosmic inflation, the PPS of scalar and tensor perturbations, and thus the CMB temperature and polarisation angular power spectra. To that end we compute the PPS numerically in order to include large scale curvature effects where the slow-roll approximation of inflation no longer holds. We interface our numerical PPS with the Boltzmann code CLASS [105] to compute the CMB anisotropy spectra, which we then use for parameter estimation and model comparison in a fully Bayesian analysis, making use of Cobaya’s [101] interface with the nested sampling code PolyChord [95, 96]. Using CLASS with its fully quantised treatment for closed universes [169] addresses the concerns raised in [181] about the proper treatment of the power spectra on large scales with non-zero spatial curvature

in other cosmological codes. For the post-processing of the nested sampling chains we use `Anesthetic` [109].

Note that there are two different perspectives one can adopt when applying a prior to the curvature density parameter $\Omega_{K,0}$. On the one hand, one can claim ignorance about the spatial curvature of the universe and apply a uniform prior over some range, typically symmetric about $\Omega_{K,0} = 0$. This is the approach taken e.g. in [33–35] and also in this paper. A preference for non-zero curvature by the data will then limit the total amount of inflation as described in the preceding paragraph. On the other hand, one can take the view that *a priori* one expects inflation to produce a large number of e-folds of the scale factor, in which case it has been suggested that one should instead adopt a prior that is strongly peaked at $\Omega_{K,0} = 0$, with tails extending to non-zero curvature values [181]. We leave the exploration of such a prior preference for spatial flatness to future work.

The structure of this paper is as follows. We first review our statistical and computational methods in [section 5.2](#). In the following [sections 5.3 to 5.5](#) we review the necessary theoretical background of inflation, including reasons why it might be finite, and the initial conditions we use in such cases of finite inflation. These sections also serve as an introduction of our notation. In [section 5.7](#) we present our numerical results for the evolution of the comoving Hubble horizon prior to and during inflation, with a specific focus on the effects of curvature at the very start of inflation. [Section 5.8](#) focuses on the conformal time passing before versus after the end of inflation. This places an important constraint on primordial parameters, especially the primordial curvature, in order to solve the horizon problem. Next, in [section 5.9](#), we investigate another constraint originating from the reheating epoch following inflation, which is particularly relevant for the total duration of inflation. In [section 5.10](#) we review the computations going into the generation of the PPS. In [section 5.11](#) we present some popular choices of inflationary potentials and their predictions for slow-roll parameters such as the scalar spectral index n_s , its running n_{run} , and the tensor-to-scalar ratio r of primordial perturbations. We compare the slow-roll predictions for these parameters to their corresponding one-parameter extensions of Λ CDM, while also allowing for non-zero spatial curvature. Much of the theory presented up to this point is well-covered in the literature for the flat case. However, since the curved case exhibits sufficient complexity when all these pieces are put together, it warrants this lengthier exposition in this paper. [Section 5.12](#) gives an overview of our choice of parametrisation used for our nested sampling runs, the results of which we present in [section 5.13](#). We start out by presenting results only from parameter extensions of Λ CDM with the curvature density parameter $\Omega_{K,0}$ and the tensor-to-scalar ratio r . We then present results for various single scalar field inflation models using the fully numeric computation of the PPS. We conclude in [section 5.14](#).

5.2 Methods

We make use of the same tools as in our preceding analysis in [chapter 2](#) on the effects of half-constrained parameters on the Bayesian evidence. Hence, this section follows closely the corresponding “methods” section therein.

5.2.1 Bayesian inference and nested sampling

For the estimation of the probability distribution for a set of model parameters θ and for the comparison of various models M with one another, we make use of Bayesian methods rooted in Bayes’ theorem. The theorem describes how to update a prior belief $\pi_M(\theta)$ with the likelihood $\mathcal{L}_M(\theta)$ of the parameters θ in light of new data D :

$$\begin{aligned} \Pr(\theta | D, M) \times \Pr(D | M) &= \Pr(D | \theta, M) \times \Pr(\theta | M), \\ \text{Posterior} \times \text{Evidence} &= \text{Likelihood} \times \text{Prior}, \\ \mathcal{P}_M(\theta) \times \mathcal{Z}_M &= \mathcal{L}_M(\theta) \times \pi_M(\theta). \end{aligned} \tag{5.1}$$

The posterior $\mathcal{P}_M(\theta)$ is the main quantity of interest in a parameter estimation, representing our state of knowledge about the parameters θ given a model M and the data D . The evidence \mathcal{Z}_M is pivotal for model comparison.

Were we interested only in parameter estimation, then it would be sufficient to consider the unnormalised posterior, which is proportional to the product of likelihood and prior, and the Bayesian evidence could be neglected as a mere normalisation factor. However, for the comparison of two models A and B the evidence becomes important. Putting the two models on the same footing a priori, i.e. $\Pr(A) = \Pr(B)$, the posterior probability ratio is equal to the evidence ratio of the two models:

$$\frac{\Pr(A|D)}{\Pr(B|D)} = \frac{\Pr(D|A)}{\Pr(D|B)} = \frac{\mathcal{Z}_A}{\mathcal{Z}_B} \tag{5.2}$$

This ratio can be interpreted as betting odds for the two models. We typically quote this in terms of the log-difference of evidences $\Delta \ln \mathcal{Z} = \ln(\mathcal{Z}_A/\mathcal{Z}_B)$.

The Bayesian evidence is sometimes also referred to as the marginal likelihood and thereby can be interpreted as the prior average of the likelihood:

$$\mathcal{Z}_M = \int \mathcal{L}_M(\theta) \pi_M(\theta) d\theta = \langle \mathcal{L}_M \rangle_{\pi}. \tag{5.3}$$

It is notoriously difficult to calculate because it requires one to take the whole parameter space into account, unlike the posterior distribution, which typically spans only a small fraction of the sampled parameter space. On the other hand, it takes the complexity of a model into account by automatically applying an Occam penalty to over-parametrised models. We use `PolyChord` [95, 96], a nested sampler that is designed to efficiently sample high-dimensional

parameter spaces with a speed hierarchy allowing to oversample the many nuisance parameters that come with experiments such as Planck [62] or the BICEP and Keck Array [71].

Another useful quantity to investigate is the Kullback–Leibler (KL) divergence defined as

$$\mathcal{D}_{\text{KL},M} = \int \mathcal{P}_M(\theta) \ln \left(\frac{\mathcal{P}_M(\theta)}{\pi_M(\theta)} \right) d\theta. \quad (5.4)$$

It is also referred to as the relative entropy, describing its role in quantifying the information gain when going from the prior to the posterior distribution. Splitting up the log-evidence into KL-divergence and posterior average of the log-likelihood highlights how the KL-divergence can also be viewed as a measure for the Occam penalty that goes into the Bayesian evidence, with the posterior average of the log-likelihood being a measure for the fit of the model (see also [chapter 2](#)):

$$\begin{aligned} \ln \left(\int \mathcal{L}_M \pi_M d\theta \right) &= \int \mathcal{P}_M \ln \mathcal{L}_M d\theta - \int \mathcal{P}_M \ln \left(\frac{\mathcal{P}_M}{\pi_M} \right) d\theta, \\ (\text{log-})\text{evidence} &= \text{parameter fit} - \text{Occam penalty}, \\ \ln \mathcal{Z}_M &= \langle \ln \mathcal{L}_M \rangle_{\mathcal{P}} - \mathcal{D}_{\text{KL},M}. \end{aligned} \quad (5.5)$$

Besides the posterior *average* of the log-likelihood, we note that the posterior *variance* of the log-likelihood gives us the Bayesian model dimensionality d_M , a measure of the number of parameters *constrained* by the data, which typically differs from the total number of free sampling parameters [100]:

$$\frac{d_M}{2} = \int \mathcal{P}_M(\theta) \left(\ln \frac{\mathcal{P}_M(\theta)}{\pi_M(\theta)} - \mathcal{D}_{\text{KL},M} \right)^2 d\theta \quad (5.6)$$

$$= \langle (\ln \mathcal{L}_M)^2 \rangle_{\mathcal{P}} - \langle \ln \mathcal{L}_M \rangle_{\mathcal{P}}^2. \quad (5.7)$$

5.2.2 ODE solvers

For the integration of the ordinary differential equations (ODEs) of the primordial inflationary background, to be introduced in [section 5.3](#) in [eqs. \(5.14\) to \(5.17\)](#), we use `scipy`'s Runge–Kutta (RK) integrator [1, 194]. The mode [eqs. \(5.47\) and \(5.54\)](#) of primordial perturbations are highly oscillatory and therefore pose a challenge for standard RK integrators, as they need to trace every oscillation in the solution. Instead, we use `oscode` [195, 196] for these equations, whose adaptive algorithm [197] switches automatically between conventional RK steps based on a fourth and fifth order Taylor approximation and WKB steps which make use of the Wentzel–Kramers–Brillouin approximation and are better suited for tracking oscillatory solutions. This allows the integrator to skip over multiple wavelengths in a single step, decreasing the computation time over conventional methods by several orders of magnitude and improving the accuracy.

5.2.3 Statistical and cosmological software

We explore posterior distributions of cosmological and nuisance parameters using `Cobaya` [101], which interfaces the sampling with the theory codes and provides both the MCMC sampler developed for `CosmoMC` [102, 103] with a “fast dragging” procedure described in [104] and the nested sampling code `PolyChord` [95, 96], tailored for high-dimensional parameter spaces, which can determine the Bayesian evidence simultaneously. For the cosmological theory code we use the Boltzmann solver `CLASS` [105, 106, 168, 169], which computes the theoretical CMB power spectra for temperature and polarisation modes. We extend `CLASS` with our own code computing the primordial power spectrum for various inflationary potentials, making use of the aforementioned ODE solver `oscode` [195].

We use `GetDist` [108] to generate the data tables of marginalised parameter values. The post-processing of the nested sampling output for the computation of Bayesian evidence, KL-divergence and Bayesian model dimensionality, as well as the plotting functionality for posterior contours is performed using the python module `anesthetic` [109].

5.2.4 Data

We use the 2018 temperature and polarisation data from the Planck satellite [62], abbreviated as “ $TT, TE, EE + \text{low}E$ ” in the Planck publication. Note that the specific use of “ $\text{low}E$ ” but lack of “ $\text{low}T$ ” might mislead one to the conclusion that only E -mode and no temperature data were used at low multipoles. However, this is *not* the case. The abbreviation implies the inclusion of both high- ℓ and low- ℓ temperature auto-correlation data. To save space we will frequently refer to this Planck data release as “P18”.

Additionally, we use B -mode data from the 2015 observing season of BICEP2 and the Keck Array CMB experiments [71] (2018 data release), which we abbreviate as “BK15”.

Note that we have chosen not to include data from CMB lensing and from Baryon Acoustic Oscillations (BAOs), which have been shown to be in some tension for closed universes [27–29]. The reasons for this are unclear, but may be related to the same issue as with the lensing parameter A_{lens} or the fact that the corresponding likelihoods assume a fiducially flat cosmology.

5.3 Inflationary background equations

The Friedmann equations [49] and the related continuity equation are derived from the Einstein equations of general relativity [48] assuming the homogeneous and isotropic FLRW metric [49–54]. They govern the dynamics of the Universe in form of the scale factor a and

its energy content given by the energy density ρ and pressure p :

$$H^2 + \frac{Kc^2}{a^2} = \frac{8\pi G}{3}\rho + \frac{\Lambda c^2}{3}, \quad (5.8)$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}\left(\rho + \frac{3p}{c^2}\right) + \frac{\Lambda c^2}{3}, \quad (5.9)$$

$$0 = \dot{\rho} + 3H\left(\rho + \frac{p}{c^2}\right), \quad (5.10)$$

where G is Newton's gravitational constant, $H = \frac{\dot{a}}{a}$ is the Hubble parameter, $K \in \{-1, 0, +1\}$ is the sign of the spatial curvature* corresponding to open, flat and closed universes respectively, and Λ is the cosmological constant. Dots denote derivatives with respect to cosmic time.

Equation (5.8) can be re-expressed in terms of density parameters Ω by introducing the critical density

$$\rho_{\text{crit}} \equiv \frac{3H^2}{8\pi G}, \quad \Omega_i \equiv \frac{\rho_i}{\rho_{\text{crit}}}, \quad \Omega_K \equiv -\frac{Kc^2}{(aH)^2}, \quad (5.11)$$

$$1 = \sum_i \Omega_i, \quad i \in \{r, m, \Lambda, K, \phi\}, \quad (5.12)$$

where the index i runs over all relevant types of fluids such as radiation r , matter m , dark energy (cosmological constant) Λ , a scalar inflaton field ϕ or curvature K . Depending on the mix of such fluids various realisations of universes with different geometries, different matter content, and different evolutions are possible. Given a vanishing cosmological constant the evolution of a universe is tightly linked to the spatial curvature. A closed universe will eventually recollapse, whereas an open universe would expand forever. For a non-zero cosmological constant this is not generally true. For a negative cosmological constant all universes will eventually recollapse, while for a positive cosmological constant most universes will in fact expand forever irrespective of the spatial curvature.

We will now switch to *reduced* Planck units with the Planck mass $m_{\text{p}} \equiv \sqrt{\frac{\hbar c}{8\pi G}}$ and with $c = \hbar = 1$. We assume that a single, scalar field ϕ , which we call the inflaton, dominates over all other species early in the history of the Universe, with the possible exception of curvature. The energy density ρ and pressure p for the inflaton field ϕ are

$$\rho_\phi = \frac{1}{2}\dot{\phi}^2 + V(\phi), \quad p_\phi = \frac{1}{2}\dot{\phi}^2 - V(\phi), \quad (5.13)$$

where $V(\phi)$ is the potential of the inflaton field. Inserting energy density and pressure into eqs. (5.8) to (5.10) and switching to reduced Planck units we get the background equations for the evolution of our inflaton field early on, before contributions from radiation and matter

*Note, that there are different conventions in the treatment of the spatial curvature parameter K . Here, we absorb any arbitrariness in the magnitude of K into the radial coordinate and the scale factor a , such that K only takes one of $\{-1, 0, +1\}$ (see e.g. [198] for more details). Consequently we generally have $a_0 \neq 1$, in contrast to the flat case.

become relevant:

$$H^2 + \frac{K}{a^2} = \frac{1}{3m_{\text{p}}^2} \left(\frac{1}{2} \dot{\phi}^2 + V(\phi) \right), \quad (5.14)$$

$$\frac{\ddot{a}}{a} = \dot{H} + H^2 = -\frac{1}{3m_{\text{p}}^2} \left(\dot{\phi}^2 - V(\phi) \right), \quad (5.15)$$

$$0 = \ddot{\phi} + 3H\dot{\phi} + V'(\phi). \quad (5.16)$$

Equations (5.14) and (5.15) can be combined to give a slightly simpler potential-independent expression:

$$\dot{H} = -\frac{1}{2m_{\text{p}}^2} \dot{\phi}^2 + \frac{K}{a^2}. \quad (5.17)$$

Inflation is commonly defined as a period of accelerated expansion $\ddot{a} > 0$ from which we can derive the following equivalent definitions of inflation:

$$\ddot{a} > 0, \quad (5.18)$$

$$\frac{d}{dt}(aH)^{-1} < 0, \quad (5.19)$$

$$V(\phi) > \dot{\phi}^2. \quad (5.20)$$

Equation (5.19) is a direct consequence of eq. (5.18) and provides a more practical way of defining inflation, since it directly relates to the horizon and flatness problem which motivated inflation in the first place. Equation (5.20) can be derived from the previous definitions together with the second Friedmann eq. (5.15). It links the definition of inflation to the inflaton field ϕ and its time derivative $\dot{\phi}$, which we will use in the following sections for setting initial conditions at the start of inflation.

The equation of state for the inflaton field is

$$w_{\phi} \equiv \frac{p_{\phi}}{\rho_{\phi}} = \frac{\frac{1}{2}\dot{\phi}^2 - V(\phi)}{\frac{1}{2}\dot{\phi}^2 + V(\phi)} \quad (5.21)$$

and can be used to differentiate between different regimes such as kinetic dominance (KD) or slow-roll (SR):

$$w_{\phi} \begin{cases} \approx 1 & \text{kinetic dominance, } \dot{\phi}^2 \gg V(\phi), \\ > -\frac{1}{3} & \text{no inflation,} \\ < -\frac{1}{3} & \text{(fast-roll) inflation,} \\ \approx -1 & \text{slow-roll inflation, } \dot{\phi}^2 \ll V(\phi). \end{cases} \quad (5.22)$$

5.4 Finite inflation

In this paper we seek to further investigate the effect of a finite amount of inflation on various quantities such as the comoving Hubble horizon, conformal time, or the primordial power spectrum of curvature perturbations and gravitational waves. Where there is only a finite amount of inflation, there is also a start to inflation, which will influence our choice of initial conditions, as we will discuss further in [section 5.5](#).

There are various mechanisms that can prevent inflation at early times. In this paper we focus on two components: kinetic dominance (KD) and spatial curvature. Their effects on the energy density ρ and the comoving Hubble horizon $a_0/(aH)$ are shown in [figs. 5.1](#) and [5.2](#) and detailed in the following [sections 5.4.1](#) and [5.4.2](#). Similarly, inhomogeneities in the inflaton field might prevent inflation at early times, as sketched in [fig. 5.3](#). We briefly comment on inhomogeneities in [section 5.4.3](#), but for simplicity we neglect them for the remainder of this paper and focus on KD and curvature as the regimes preceding inflation.

5.4.1 Kinetic dominance (KD)

In [Hergt et al. \(2019\) \[II & chapter 4\]](#) we worked in a spatially flat universe and compared the effect of setting the initial conditions for inflation during SR with $\dot{\phi}^2 \ll V(\phi)$, or KD with $\dot{\phi}^2 \gg V(\phi)$ (cf. [eq. \(5.22\)](#)). In both regimes the initial conditions for the background [eqs. \(5.14\) to \(5.17\)](#) can be conveniently expressed in analytic form [[II & chapter 4, 115](#)].

In [fig. 5.1](#) (orange line) we schematically illustrate the role of both the KD and SR regime in the overall evolution of the Universe. Note in particular that the energy density scales as

$$\rho_{\text{KD}} \propto a^{-6} \quad \text{and} \quad \rho_{\text{SR}} \propto a^0,$$

and the comoving Hubble horizon as

$$(aH)_{\text{KD}}^{-1} \propto a^2 \quad \text{and} \quad (aH)_{\text{SR}}^{-1} \propto a^{-1}.$$

With the inclusion of spatial curvature, we prefer setting the initial conditions at the start of inflation, i.e. at the turnover point where the comoving Hubble horizon becomes maximal and changes from growing during KD to shrinking during SR. This prevents (at the prior level) running into spatially closed universes that collapse (the comoving Hubble horizon diverges) even before inflation has actually started. Nevertheless, we can integrate the inflationary background [eqs. \(5.14\) to \(5.17\)](#) both forwards and backwards in time to recover the SR and the KD regime respectively. Forwards in time, SR inflation is an attractor solution, meaning that regardless of where in $(\phi, \dot{\phi})$ phase-space the inflaton starts from, it will end up in the SR regime, where the inflaton “slowly rolls down the potential”, motivating its name (see also [[I & chapter 3, 116, 118, 135, 178, 199](#)]). Integrating backwards in time, the attractor solution is kinetic dominance [[I & chapter 3, 115](#)].

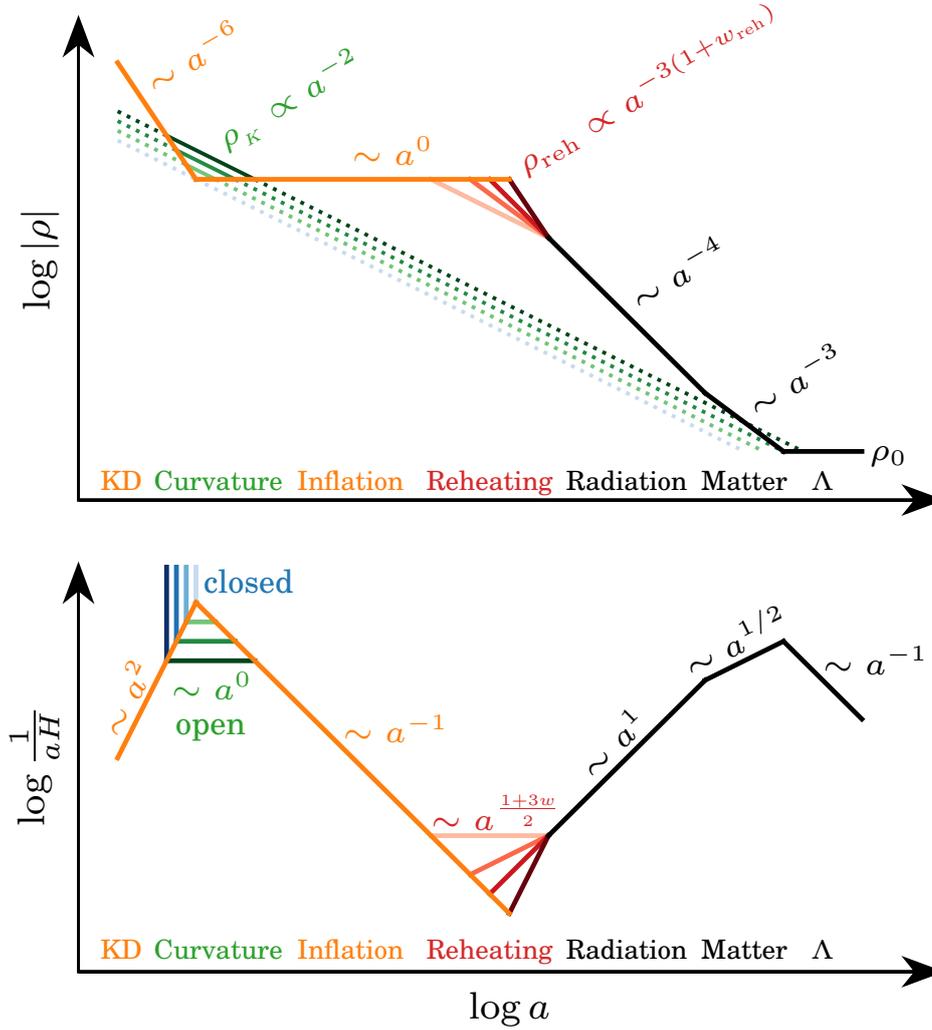


Figure 5.1 Sketch of the evolution of the energy density ρ and the comoving Hubble horizon $(aH)^{-1}$. We highlight the possible effect of different levels of curvature in green. Darker shades of green correspond to a higher curvature density. The dotted lines indicate how the curvature becomes increasingly less relevant during inflation. If curvature were to dominate over the sum of all other components (inflaton, radiation, matter, dark energy) in a closed universe, this will lead to collapsing universes (diverging comoving Hubble horizon), as sketched in blue. Orange corresponds to the evolution in the absence of curvature or where curvature is negligible, starting out in kinetic dominance (KD) and then transitioning into inflation. Black corresponds to the standard Big Bang evolution from radiation, to matter and roughly today to Λ (or dark energy) domination. The red lines correspond to different reheating scenarios parametrised by the equation of state parameter of reheating w_{reh} .

5.4.2 Curvature domination

If the energy density of spatial curvature ρ_K is of the same order as that of the inflaton field, then this leads to visible effects in the comoving Hubble horizon in the transition region between KD and SR. This is shown schematically in [figs. 5.1](#) and [5.2](#).

[Figure 5.2](#) highlights the symmetry between open (green) and closed (blue) universes around the flat case (orange). During inflation and going backwards in time, the comoving Hubble horizon grows in proportionality to the scale factor a . Due to the direct connection

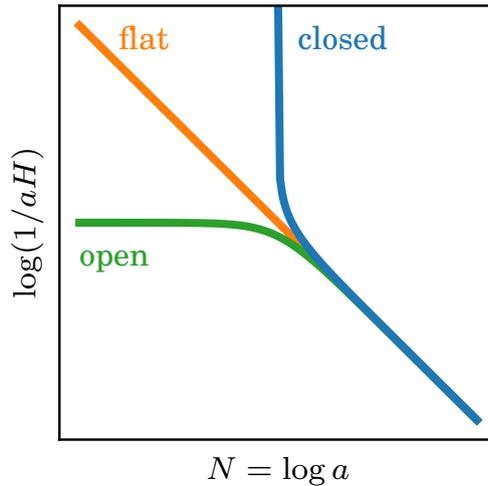


Figure 5.2 Schematic view of effects from significant amounts of spatial curvature on the evolution of the comoving Hubble horizon preceding and during inflation, where we assume slow-roll conditions, such that any term involving $\dot{\phi}$ is neglected and the potential V is constant. The comoving Hubble horizon then becomes $(aH)^{-1} = (V e^{2N} - 3 m_{\text{p}}^2 K)^{-1/2}$. Open universes are constrained by eq. (5.31). For significantly closed universes the comoving Hubble horizon diverges.

between comoving Hubble horizon $(aH)^{-1}$ and curvature density parameter Ω_K (by definition of Ω_K in eq. (5.11)), this means that the curvature density grows as well. Once the curvature density comes to dominate, the comoving Hubble horizon diverges for a closed universe and levels off for an open universe. Thus, we can only really refer to a phase of curvature *domination* in the open case. Nonetheless, effects of curvature on the comoving Hubble horizon can already be seen for lower (i.e. large but not dominating) levels of curvature, as we show in more detail in fig. 5.6. An actual divergence in the closed case could have its origin in a coasting or bouncing universe. Note that bouncing universe models typically involve some modification of general relativity, a formulation of quantum gravity or additional assumptions about the inflaton field (see [200] for a review). Viewed from the opposite end by starting in kinetic dominance, a diverging comoving Hubble horizon would in almost all cases correspond to a collapsing universe. The plateau in the open case could in principle reach back indefinitely, if the universe started exactly with $\Omega_K = 1$. Curvature below that value would mean a preceding phase of KD.

Figure 5.1 illustrates schematically the role of curvature in the early evolution of the Universe. With $\rho_K \propto a^{-2}$ the curvature density drops slower than the kinetic energy density from the inflaton field and therefore becomes more relevant in the vicinity of the local maximum of the comoving Hubble horizon, causing the horizon to flatten or sharpen in the open or closed case, respectively. Once the inflaton potential comes to dominate over the kinetic term, its energy density, which scaled as a^{-6} during KD, becomes constant during inflation, thereby quickly exceeding the curvature density.

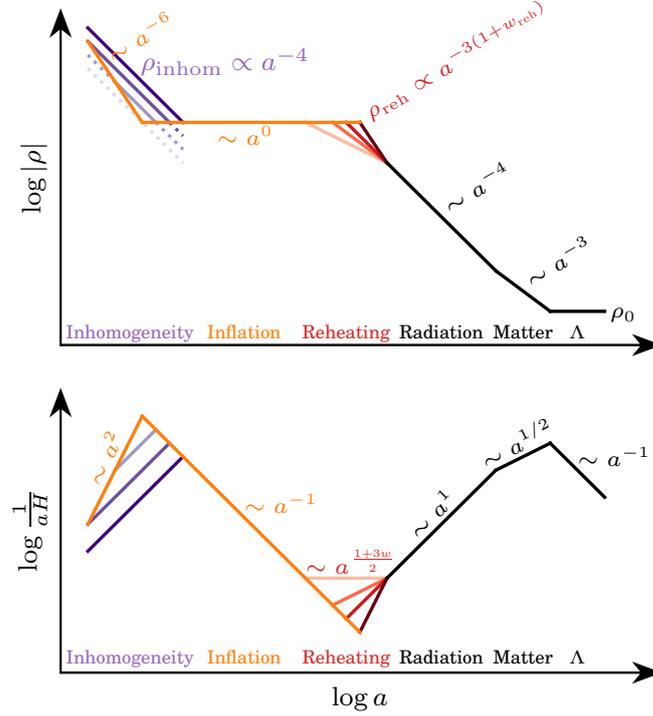


Figure 5.3 Sketch of the evolution of the energy density ρ and the comoving Hubble horizon $(aH)^{-1}$ as in fig. 5.1. This time we highlight in purple the possible effect of different levels of inhomogeneities prior to the start of inflation.

5.4.3 Inhomogeneities

Various analyses [178, 199, 201–204] have found that inhomogeneities cause a pre-inflationary phase where the energy density scales as

$$\rho_{\text{inhom}} \propto a^{-4}. \quad (5.23)$$

This is sketched in fig. 5.3, which also demonstrates how high levels of primordial inhomogeneities might hide a period of kinetic dominance. However, like kinetic dominance, inhomogeneities will lead to a finite start to inflation.

5.5 Initial Conditions for inflation

5.5.1 Setting initial conditions at the start of inflation

In section 5.4.1 we have already mentioned that KD leads to analytic expressions for the evolution of background variables, which may serve as initial conditions for the inflaton field ϕ . However, it is equally possible to set the initial conditions at the start of inflation and then numerically integrate both backwards in time towards kinetic dominance and forwards in time into inflation. This gives us a better handle on new parameters that are introduced through the addition of curvature such as the primordial curvature density parameter at the start of

inflation $\Omega_{k,i}$. We will label parameters referring to the start of inflation with a subscript “i”, mnemonic for “initial” or “inflation”.

At the start of inflation we can set the inequalities in the definitions for inflation from eqs. (5.18) to (5.20) and (5.22) to equality, leading to the following expressions:

$$\ddot{a}_i = 0, \quad (5.24)$$

$$\frac{d}{dt}(a_i H_i)^{-1} = 0, \quad (5.25)$$

$$V(\phi_i) = \dot{\phi}_i^2, \quad (5.26)$$

$$w_{\phi,i} = -\frac{1}{3}, \quad (5.27)$$

This simplifies the background eq. (5.14), and together with eq. (5.11) for the curvature density parameter allows us to relate the scale factor a_i , the inflaton field ϕ_i and the curvature density parameter $\Omega_{K,i}$ at the start of inflation:

$$\begin{aligned} H_i^2 &\stackrel{(5.14)}{=} \frac{V(\phi_i)}{2 m_p^2} - \frac{K}{a_i^2}, \\ &\stackrel{(5.11)}{=} \frac{-K}{a_i^2 \Omega_{K,i}}. \end{aligned} \quad (5.28)$$

Numerically integrating the background eqs. (5.14) and (5.15) and the equation of motion for the inflaton field eq. (5.16) requires initial values for the variables $\{N, \phi, \dot{\phi}\}$, where $N = \ln a$ is the number of e-folds of the scale factor and a is measured in reduced Planck units. Equation (5.26) links the initial value for the time derivative of the inflaton field $\dot{\phi}_i$ to the potential of the inflaton field $V(\phi_i)$ at the start of inflation. This leaves the e-folds N_i and the inflaton field ϕ_i as free parameters. The initial value of the curvature density parameter $\Omega_{K,i}$ can be derived using eq. (5.28) and hence could be varied in place of either N_i or ϕ_i :

$$\Omega_{K,i} = \left[1 - \frac{V(\phi_i)}{2 m_p^2 K} e^{2N_i} \right]^{-1}, \quad (5.29)$$

$$N_i = \frac{1}{2} \ln \left[\frac{2 m_p^2 K}{V(\phi_i)} \left(1 - \Omega_{K,i}^{-1} \right) \right]. \quad (5.30)$$

From this we can derive the condition that the primordial curvature density needs to be

$$\Omega_{K,i} < 1 \quad (5.31)$$

in order for inflation to start after the Big Bang. Equality would correspond to $N_i \rightarrow -\infty$ or $a_i = 0$, i.e. inflation starting at the Big Bang.

The initial value for the inflaton field ϕ_i determines the amount of e-folds of inflation. Hence, it can be useful to infer ϕ_i from a desired number of e-folds of inflation. In the following

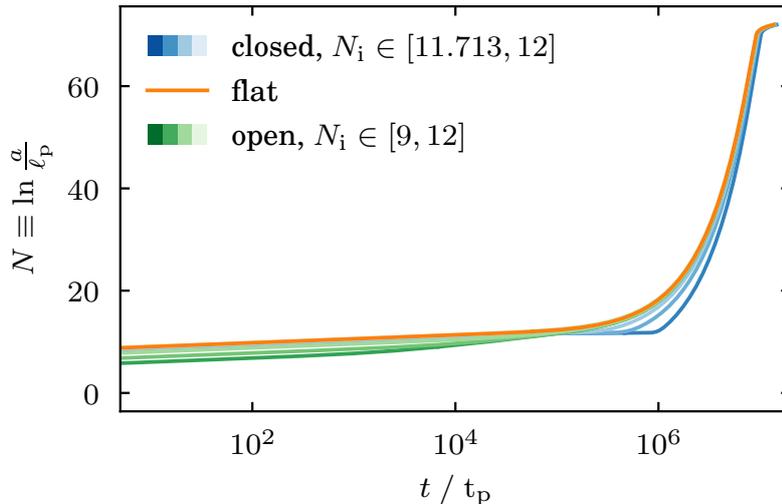


Figure 5.4 Evolution of the logarithm of the scale factor $N \equiv \ln(a/\ell_p)$ with respect to physical time t for closed ($\Omega_K < 0$), flat ($\Omega_K = 0$) and open ($\Omega_K > 0$) universes. Both of these variables are used as independent variables in [fig. 5.5](#) where we also detail the generation of these curves. The scale factor and physical time are given in the reduced Planck units for length ℓ_p and time t_p respectively.

we will consider the total e-folds of inflation N_{tot} and the e-folds of inflation N_{\dagger} *before* and N_* *after* the pivot scale $k_* = 0.05 \text{ Mpc}^{-1}$ crosses the comoving Hubble horizon.

Using these initial conditions we can integrate [eqs. \(5.14\) to \(5.17\)](#) forwards and backwards with respect to cosmic time t or with respect to the number of e-folds N of the scale factor. The connection between the independent variables t and N is shown in [fig. 5.4](#), illustrating the exponential growth of the universe during inflation. We can also compute various other quantities such as the inflaton field ϕ , its time derivative $\dot{\phi}$, the equation of state parameter w_ϕ or the Hubble parameter H , all of which are shown in [fig. 5.5](#) with respect to cosmic time in the left column and with respect to e-folds in the right column. We show the solutions for a flat universe in orange. In green and blue we show the slightly different evolution of open and closed universes respectively. We show these curved cases for different amounts of primordial curvature, which we achieve by varying the starting point N_i . For a clean visualisation we chose the initial conditions such that inflation ends at $N_{\text{end}} = 70$.

5.6 Linking primordial to present-day scales

In order to link primordial to present-day scales we need to first calibrate the scale factor a and the wavenumber k associated with curvature perturbations, which we briefly review in this section.

Curved universes have an advantage over flat universes when discussing scales in that [eq. \(5.11\)](#) provides a direct link between the curvature density parameter and the scale factor. Given today's curvature density parameter $\Omega_{K,0}$, this allows a calibration of the scale factor without any knowledge of the evolution of the universe. Otherwise, as is the case for flat universes, we would have to make additional assumptions, e.g. by introducing a free parameter

Background variables for Starobinsky potential

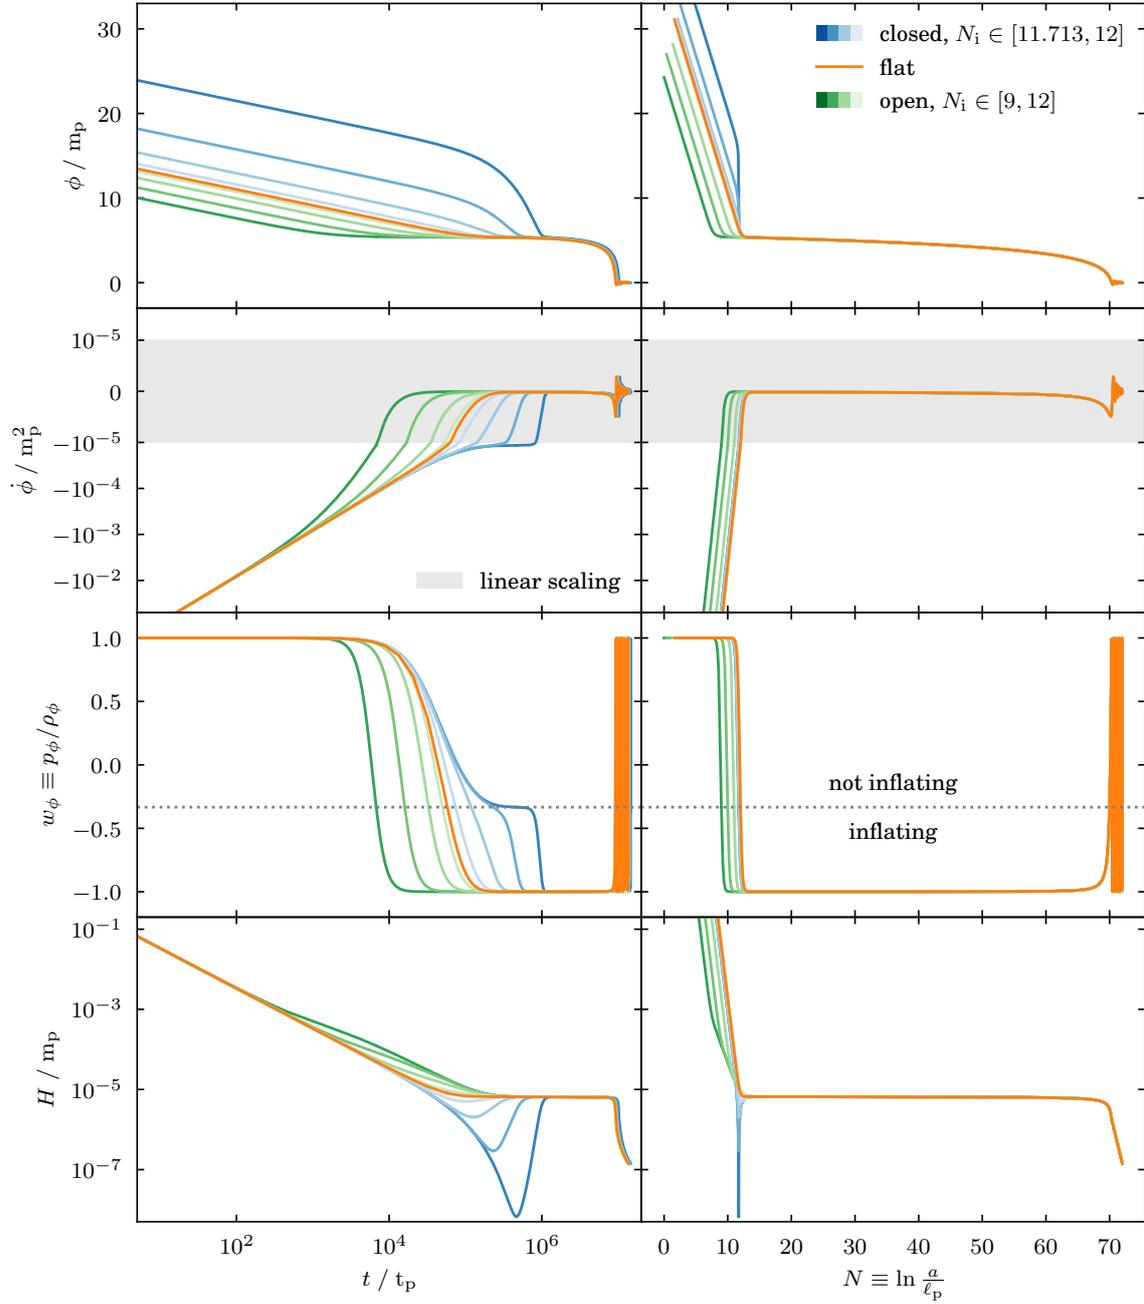


Figure 5.5 Evolution of background variables during inflation for closed ($\Omega_K < 0$), flat ($\Omega_K = 0$) and open ($\Omega_K > 0$) universes. The rows show the inflaton field ϕ , its first time derivative $\dot{\phi}$ with a semi-logarithmic scaling with threshold at $|\dot{\phi}| = 10^{-5}$, the equation-of-state parameter w_ϕ with the horizontal dotted line indicating the threshold $w_\phi = -\frac{1}{3}$ between inflating and non-inflating, and the Hubble parameter H , respectively. The left column is with respect to physical time t and the right column with respect to the natural logarithm of the scale factor $N \equiv \ln a$ (see also fig. 5.4 for their interdependence). The curves were initialised at the start of inflation where $V(\phi_i) = \dot{\phi}_i^2$ and from there integrated backwards and forwards in time. For visualisation purposes inflation was specified to end at $N_{\text{end}} = 70$ and the start of inflation was varied uniformly for the closed case within $N_i \in [9.4, 10]$ and for the open case within $N_i \in [7, 10]$. These ranges can be converted to the primordial curvature density parameter corresponding roughly to $\Omega_{K,i} \in [-300, -3]$ and $\Omega_{K,i} \in [0.997, 0.451]$, respectively. Note that these plots were generated using the Starobinsky potential, but the general picture remains qualitatively mostly the same independent of the choice of potential.

on the observable e-folds N_* from horizon crossing to the end of inflation or by making specific assumptions about the evolution of the Universe between the end of inflation and before the standard Big Bang evolution, i.e. about the epoch of reheating.

5.6.1 Calibration of the present-day scale factor

In order to calculate the comoving Hubble horizon or the primordial power spectrum, we need to first calibrate the scale factor a , which in this paper we do by deriving the present-day scale factor a_0 from the present-day curvature density parameter $\Omega_{K,0}$ and Hubble parameter H_0 . This follows directly from eq. (5.11):

$$a_0 = \frac{c}{H_0} \sqrt{\frac{-K}{\Omega_{K,0}}}. \quad (5.32)$$

5.6.2 Calibration of the wavenumber of primordial perturbations

As is standard practice, we formulate the condition for horizon crossing in terms of the comoving Hubble horizon (as opposed to the particle horizon). We will evolve the gauge-invariant curvature perturbations \mathcal{R}_k for a given wavenumber k . Its reciprocal $1/k$ (ignoring possible factors of 2π that, one could argue, should be introduced) can be thought of as the comoving wavelength scale of the perturbation itself [60]. Whilst the length-scale $1/k$ of perturbations is smaller than the comoving Hubble horizon, the curvature perturbations oscillate. From the definition for inflation in eq. (5.19) we know that the comoving Hubble horizon shrinks during inflation. Once it drops below $1/k$, the oscillations stop and the curvature perturbations “freeze”, as the corresponding modes have become larger than the characteristic length-scale over which physical processes operate coherently. We use the transition point, which we refer to as horizon crossing, to link any given curvature perturbation observable today to the comoving Hubble horizon:

$$k = \frac{aH}{a_0}. \quad (5.33)$$

This allows us to draw the dotted line in fig. 5.6 representing the pivot scale $k_* = 0.05 \text{ Mpc}^{-1}$.

5.7 The comoving Hubble horizon

Figure 5.6 contrasts the evolution of the comoving Hubble horizon in closed and open universes for varying amounts of primordial curvature. For visualisation purposes, we calibrate today’s scale factor a_0 as described in section 5.6.1 by fixing today’s curvature density parameter $|\Omega_{K,0}| = 0.01$. Fixing $\Omega_{K,0}$ yields a linear relation between the evolution of the comoving Hubble horizon and the evolution of the curvature density parameter, and we therefore plot both, on opposite y -axes. This makes apparent how inflation solves the flatness problem, as the shrinking comoving Hubble horizon during inflation (by definition in eq. (5.19)) corresponds to the shrinking of the curvature density parameter, such that the standard Big

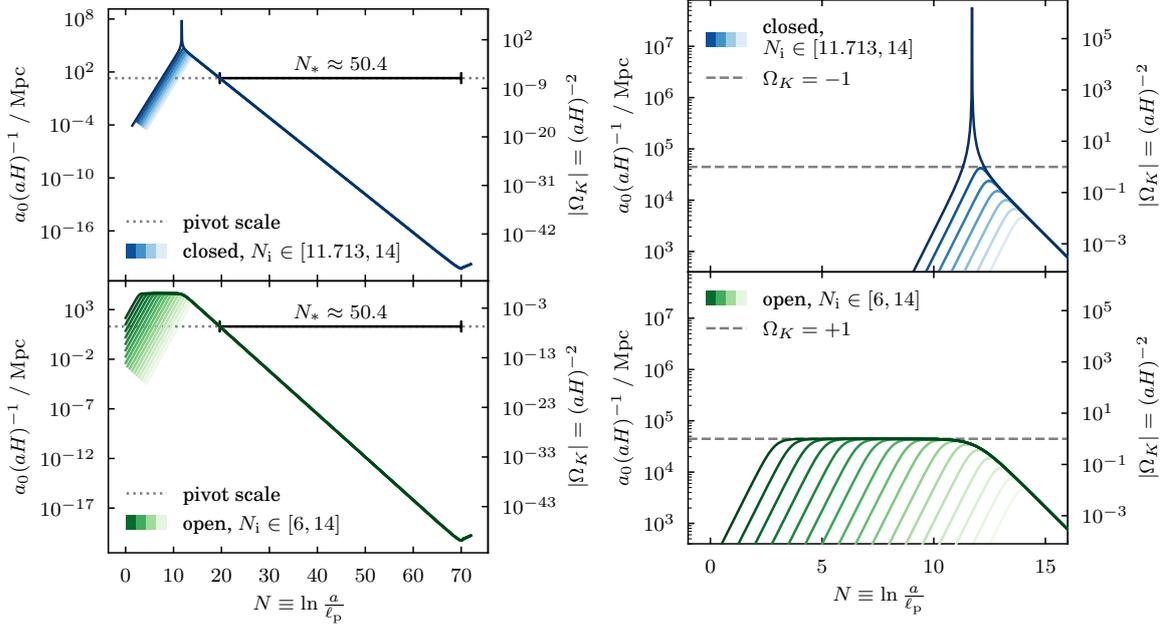


Figure 5.6 Evolution of the comoving Hubble horizon $ca_0(aH)^{-1}$ with respect to number of e-folds of the scale factor $N \equiv \ln a$. The secondary y-axis relates the comoving Hubble horizon to the absolute value of the curvature density parameter $|\Omega_K| = (aH)^{-2}$ for curved universes. As in fig. 5.5, the initial conditions were set such that inflation ends at $N_{\text{end}} = 70$ and with a varying start of inflation N_i . For the comoving Hubble parameter we also need to specify today's scale factor a_0 , which can be derived from today's curvature density parameter $\Omega_{K,0}$ and eq. (5.32). Here, we have set $|\Omega_{K,0}| = 0.01$. This effectively fixes the number of e-folds N_* from horizon crossing of the pivot scale $k_* = 0.05$ Mpc to the end of inflation. The right hand plot is a zoom-in into the region of the start of inflation where the shape for closed universes (upper panels) differs to the one of open universes (lower panels). The dashed line marks the limit of the constraint for open universes $\Omega_K < 1$ from eq. (5.31). Note that these plots were generated using the Starobinsky potential, but the general picture remains qualitatively the same independent of the choice of potential.

Bang evolution thereafter starts out with a sufficiently small curvature density parameter. For any given total number of e-folds of inflation N_{tot} the choice of $|\Omega_{K,0}|$ decides how many e-folds pass *before* versus *after* horizon crossing of the pivot scale k_* . We refer to these numbers of e-folds with N_{\dagger} and N_* respectively. The e-folds N_* (i.e. after horizon crossing of the pivot scale) we also call the *observable* number of e-folds of inflation because of their direct connection to primordial cosmological parameters (e.g. the scalar spectral index n_s) in flat slow-roll inflation models, where the total number of e-folds is typically assumed to be much larger but ultimately unknown.

As in figs. 5.4 and 5.5 we vary the start of inflation N_i while keeping the end of inflation fixed to $N_{\text{end}} = 70$ in fig. 5.6. This effectively also fixes the number of e-folds after horizon crossing of the pivot scale to $N_* \approx 50.4$. The total number of e-folds on the other hand shrinks with larger N_i as $N_{\text{tot}} = N_{\text{end}} - N_i$ and thus the initial value for the inflaton field ϕ_i decreases as well.

Equation (5.29) links the primordial curvature density parameter $\Omega_{K,i}$ to the e-folds N_i at inflation start. Thus, a smaller N_i means a larger $|\Omega_{K,i}|$ and in turn a larger comoving

Hubble horizon at inflation start. For open universes this gets capped by the constraint from eq. (5.31), meaning for very early starts of inflation the primordial curvature density parameter tends to unity:

$$\Omega_{K,i} \rightarrow +1 \quad \iff \quad N_i \rightarrow -\infty. \quad (5.34)$$

For closed universes on the other hand, the primordial curvature density parameter diverges as N_i is pushed to earlier times:

$$\Omega_{K,i} \rightarrow -\infty \quad \iff \quad N_i \rightarrow \frac{1}{2} \ln \left(\frac{2 m_p^2 K}{V(\phi)} \right). \quad (5.35)$$

Note how a small amplitude of the inflationary potential (as expected from data) in eq. (5.35) pushes inflation start until *after* the Planck epoch for closed universes: $N_i > N_p = 0$. For small levels of primordial curvature, the shape of the curve in fig. 5.6 is the same for open and closed universes and matches that of a flat universe. With increasing primordial curvature the curve becomes flatter for open universes and pointier for closed universes, moulding to the limits expressed in eqs. (5.34) and (5.35) and sketched out in fig. 5.2.

It has frequently been proposed that it would be more natural to count the number of e-folds during inflation in terms of the comoving Hubble horizon $(aH)^{-1}$ instead of the scale factor a , because of its direct relation with the flatness and the horizon problem [160, 205, 206]. Where curvature effects are negligible, i.e. where $\Omega_{K,i} < 1$, these measures are actually closely related due to the comoving Hubble horizon scaling as a^{-1} during slow-roll inflation, independently of the geometry of the universe. With the primordial curvature $\Omega_{K,i}$ approaching unity at the start of inflation, this common scaling breaks down. The behaviour of the comoving Hubble horizon in a closed universe is inverse to that in an open universe, as shown in fig. 5.2.

For closed universes, the number of e-folds of the scale factor a are in fact more informative when it comes to the effect of KD features on the primordial power spectrum. Specifically, the position of the cutoff in the primordial power spectrum caused by finite inflation (more on this later in section 5.10) is governed by the number of e-folds N_{\dagger} of the scale factor before horizon crossing of the pivot scale. So, although for a late inflation start this number is closely related to the maximum of the comoving Hubble horizon at inflation start or the ratio $f_i = \Omega_{K,i}/\Omega_{K,0}$ of primordial to present-day curvature, in the limit from eq. (5.35) even a very large change in $(a_i H_i)^{-1}$ and f_i will hardly affect N_{\dagger} and the cutoff position will cease shifting.

5.8 Conformal time

In order to solve the horizon problem, the conformal time during inflation has to match or exceed the conformal time passing thereafter until today (see also [207]). Conformal time can be expressed in terms of the comoving Hubble horizon $(aH)^{-1}$ and the e-folds of the scale

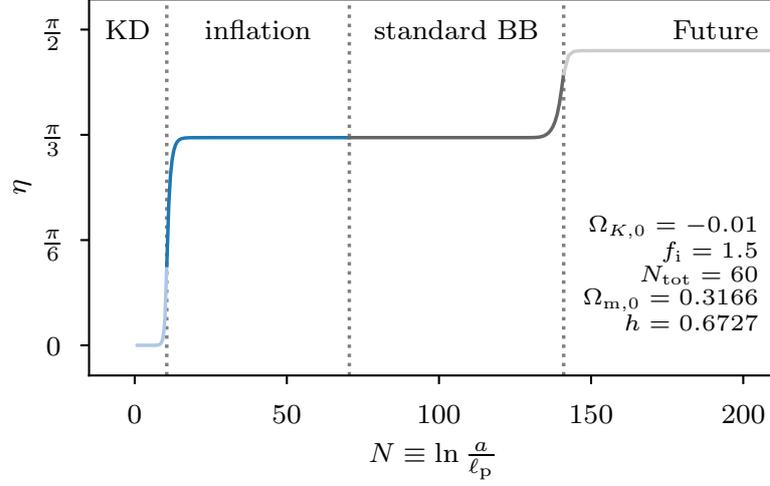


Figure 5.7 Evolution of conformal time η through different stages of the universe: kinetic dominance (KD) in light blue, inflation in blue, standard Big Bang evolution (radiation, matter and Λ domination) in grey, and from today onwards in light grey.

factor $N \equiv \ln a$:

$$\eta = \int \frac{dt}{a(t)} = \int \frac{d \ln a}{aH}. \quad (5.36)$$

Comparing this expression to [fig. 5.6](#) and bearing in mind that the comoving Hubble horizon in the figure is shown on a logarithmic scale, it is clear that the largest contribution to the conformal time passing prior to the end of inflation comes from the peak around the start of inflation. Analogously, the majority of conformal time passing after the end of inflation comes from the peak around the present-day comoving Hubble horizon. This is clear from the jumps in [fig. 5.7](#), which shows the accumulation of conformal time from before inflation start until some future time. The regions where the conformal time plateaus correspond to the regions where either the inflaton ϕ or the cosmological constant Λ have made the comoving Hubble horizon shrink so much that there is almost no contribution to the integral in [eq. \(5.36\)](#). This also holds for the post-inflationary epoch of reheating, which consequently can be neglected with regards to conformal time, which we will do throughout this section.

In this section we will be focusing on two quantities in particular: the total amount of conformal time η_{total} passing from the Big Bang prior to inflation up to the future conformal boundary, and the ratio $\eta_{\text{before}}/\eta_{\text{after}}$ of conformal time passing before to after the end of inflation.

The total amount of conformal time η_{total} is important for the closed universe theory described by Lasenby and Doran in [\[123\]](#). Here a natural boundary condition on the transition to the final asymptotic de Sitter state is found that requires the total conformal time available to the universe to equal $\pi/2$.

Also, as stated earlier, the ratio $\eta_{\text{before}}/\eta_{\text{after}}$ is important for addressing the horizon problem. In order to solve the horizon problem we require that more conformal time has

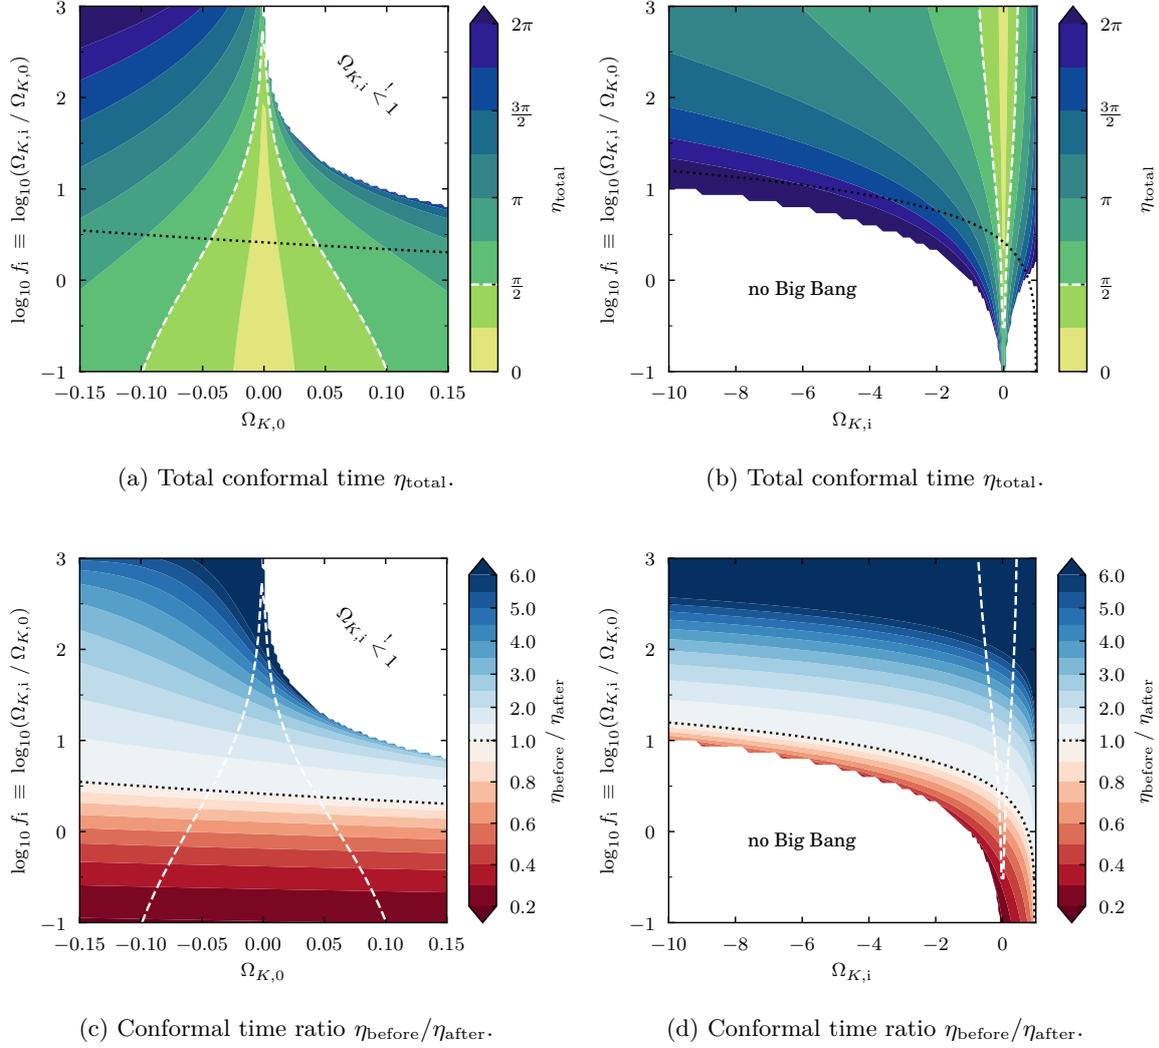


Figure 5.8 The total conformal time η_{total} from before inflation until the future conformal boundary and the ratio of conformal time *before* to *after* the end of inflation are shown dependent on the primordial and present-day density parameters, $\Omega_{K,i}$ and $\Omega_{K,0}$ respectively. The **white dashed line** highlights the value of $\eta_{\text{total}} = \frac{\pi}{2}$, which in some closed universe theories [123] is predicted to be a constraint. The **black dotted line** indicates where $\eta_{\text{after}} = \eta_{\text{before}}$. Thus, the **blue area** highlights where the horizon problem is solved and the **red area** where inflation was insufficient in order to solve the horizon problem. In all these cases the following parameters were fixed: the total number of e-folds of inflation $N_{\text{tot}} = 60$, today's matter density parameter $\Omega_{\text{m},0} = 0.3166$ and the dimensionless Hubble parameter $h = 0.6727$.

passed before than after the end of inflation, which we will refer to as the horizon constraint:

$$\frac{\eta_{\text{before}}}{\eta_{\text{after}}} > 1. \quad (5.37)$$

Figure 5.8 illustrates how both the total amount of conformal time η_{total} and the ratio $\eta_{\text{before}}/\eta_{\text{after}}$ depend on the primordial and present-day curvature density parameters, $\Omega_{K,i}$ and $\Omega_{K,0}$ respectively. We plot both these parameters against the ratio of primordial to present-day curvature

$$f_i \equiv \frac{\Omega_{K,i}}{\Omega_{K,0}}. \quad (5.38)$$

The parameter f_i will prove useful also later on for decoupling the effects of primordial and present-day curvature on the primordial power spectrum. Here, it is useful when looking at the black dotted line, which separates the plots into red regions where inflation was insufficient to solve the horizon problem, and blue regions where it was, i.e. the dotted line corresponds to $\eta_{\text{before}}/\eta_{\text{after}} = 1$. As is particularly clear in **fig. 5.8c**, this separation depends primarily on f_i when considering a prior range of $\Omega_{K,0} \in [-0.15, 0.15]$. From this we can infer that in order to solve the horizon problem we require

$$\log_{10} f_i \gtrsim 0.5. \quad (5.39)$$

The white regions in the left panels with $\log_{10} f_i$ versus $\Omega_{K,0}$ correspond to the constraint for open universes from **eq. (5.31)**, also seen in **fig. 5.6**. The white regions in the right panels correspond to universes that would have collapsed in the past (labelled “no Big Bang”) or that would collapse in the future before reaching the future conformal boundary.

The possible constraint of a total conformal time of $\eta_{\text{total}} = \pi/2$ (white dashed lines) can be satisfied while also resolving the horizon problem, as part of the white dashed line lies in the blue region. This would push the present day universe close to flat (see **Figure 5.8a**) and the primordial curvature density parameter close to unity (see **fig. 5.8b**).

Besides the primordial and present-day curvature density parameters, there are some other parameters (the number of e-folds N_{tot} , matter density $\Omega_{m,0}$ and Hubble parameter H_0) that enter into the calculation of both the total conformal time as well as the conformal time ratio. However, their contribution to conformal time is negligible compared to that of the curvature parameters as seen in **fig. 5.9**.

For the inflationary part of the calculation we additionally need to consider the mass of the inflaton and the duration of inflation. The mass of the inflaton (or the amplitude of the inflationary potential) can be mapped to the amplitude of the primordial power spectrum, which, as we will see in **fig. 5.13**, has no effect on the comoving Hubble horizon and is thus irrelevant for the calculations of conformal time. The total inflationary e-folds N_{tot} only influence the comoving Hubble horizon towards the end of inflation where N_{tot} becomes minimal (see also **fig. 5.13**). At that point the comoving Hubble horizon is many orders of magnitude smaller than at its start and consequently this contribution to the integral for

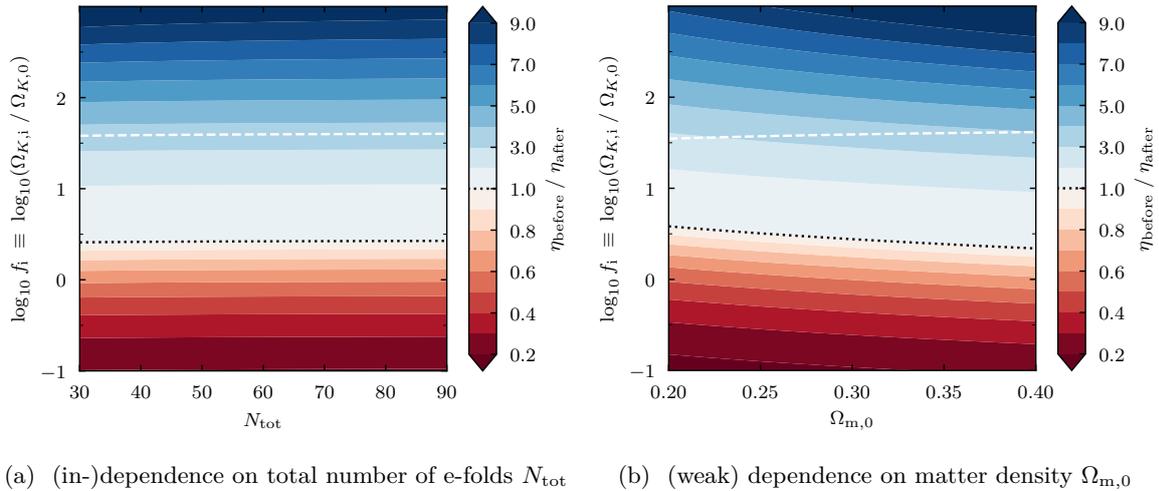


Figure 5.9 Conformal time ratio $\eta_{\text{before}}/\eta_{\text{after}}$ as in fig. 5.8, but here exploring the (in-)dependence on the total number of e-folds of inflation N_{tot} (left) and the present-day matter density parameter $\Omega_{\text{m},0}$ (right). The following parameters were fixed: today’s curvature density parameter $\Omega_{K,0} = -0.01$, the dimensionless Hubble parameter $h = 0.6727$, today’s matter density parameter $\Omega_{\text{m},0} = 0.3166$ (only in left panel), and the total number of e-folds of inflation $N_{\text{tot}} = 60$ (only in right panel).

conformal time in eq. (5.36) is negligible. Figure 5.9a illustrates how the conformal time ratio is almost independent of the total amount of inflation N_{tot} . Note that this goes against the rule of thumb of requiring order 60 e-folds of inflation to solve the horizon problem, which is valid when the start of inflation is fixed. Figure 5.9a tells us that we can solve the horizon problem equally well for only 30 e-folds. The essential thing is that the comoving Hubble horizon (or the curvature density parameter) needs to have been sufficiently large at the start of inflation compared to today, reinforcing the requirement from eq. (5.39) that $\log_{10} f_i \gtrsim 0.5$. Note, however, that we are investigating conformal time completely isolated from other possible constraints from reheating, here. We will investigate constraints from reheating in the following section 5.9 and later in section 5.13.2.

For the calculation of conformal time after the end of inflation and throughout radiation, matter and Λ domination we need to further consider today’s matter density parameter $\Omega_{\text{m},0}$ and Hubble parameter H_0 . Through eq. (5.32) the Hubble parameter mostly serves as a normalisation factor to the scale factor a_0 and therefore primarily only shifts the comoving Hubble horizon along $\ln a$, which does not affect the integral for conformal time in eq. (5.36). Increasing the present-day matter density parameter $\Omega_{\text{m},0}$ increases the matter contribution to the comoving Hubble horizon, which therefore becomes larger during matter domination in general and at the end of matter domination in the late-time Universe in particular, when the comoving Hubble horizon peaks (see also fig. 5.1). Thus, there is a dependence of the conformal time ratio $\eta_{\text{before}}/\eta_{\text{after}}$ on the matter density parameter. However, for the range of $\Omega_{\text{m},0} \in [0.2, 0.4]$ this dependence is weak compared to the dependence on f_i , as seen in fig. 5.9b.

5.9 Reheating

While the start of inflation plays a crucial role for considerations of the conformal time and thereby also the horizon problem, the end of inflation is important for the period of reheating. Reheating links the primordial evolution of the Universe to the standard Big Bang evolution, comprised of radiation, matter, and Λ domination. Going back to [fig. 5.1](#), we can see this schematically for the energy density ρ and the comoving Hubble horizon $(aH)^{-1}$. For the reheating period we plot four characteristic equation-of-state parameters w : the lower limit $w_{\text{reh}} = -1/3$, matter domination $w_{\text{reh}} = 0$, radiation domination $w_{\text{reh}} = 1/3$ and the upper limit $w_{\text{reh}} = 1$.

The link between primordial and standard Big Bang evolution becomes particularly important in the case of the Universe having non-zero curvature, since the latter informs us about the overall scale of the Universe, as established in [sections 5.6.1](#) and [5.7](#). This dependence on the curvature density parameter $\Omega_{K,0}$ is illustrated in [fig. 5.10](#) showing the comoving Hubble horizon (upper panel). Different from [fig. 5.2](#), in [fig. 5.10](#) we fix both start N_i and end N_{end} of inflation while allowing today's curvature density parameter $\Omega_{K,0}$ to vary. Consequently we can no longer show the evolution of the comoving Hubble horizon $a_0/(aH)$ and the curvature density parameter Ω_K in one plot, since today's curvature density parameter $\Omega_{K,0}$ serves as calibrator for today's scale factor a_0 which in turn calibrates the comoving Hubble horizon. We therefore also plot an uncalibrated version $(aH)^{-1}$, where the primordial evolution collapses onto a single line, whereas the standard Big Bang evolution shifts vertically with $\Omega_{K,0}$.

We parametrise the epoch of reheating through an effective equation-of-state parameter w_{reh} , a duration ΔN_{reh} and an energy scale ρ_{reh} where thermalisation is guaranteed to have occurred (see also [\[162\]](#) for more details on this reheating parametrisation). Note first that w_{reh} is an effective parameter. During the inflaton's oscillations around a potential minimum at the end of inflation, the equation-of-state parameter also oscillates rapidly between ± 1 . For the *effective* equation-of-state parameter we consider the time-averaged value. For a monomial potential with exponent p this gives $w_{\text{reh}} = (p - 2)/(p + 2)$, e.g. we have $w_{\text{reh}} = 0$ for a quadratic and $w_{\text{reh}} = 1/3$ for a quartic potential. Since reheating is by definition a post-inflationary epoch, we at the very least expect that on average $w_{\text{reh}} > -1/3$ (otherwise we would have more inflation, cf. [eq. \(5.22\)](#)). Additionally, the equation of state is typically capped at $w_{\text{reh}} < 1$ to avoid a super-luminal sound speed [\[208\]](#). Together this leads to our first reheating constraint:

$$-\frac{1}{3} < w_{\text{reh}} < 1. \quad (5.40)$$

Second, we note that there is little information on the energy scale of thermalisation ρ_{th} . Hence, we also view the energy scale ρ_{reh} as an effective parameter by which thermalisation must have happened, but not necessarily equal to ρ_{th} . Thus, the case $\rho_{\text{th}} > \rho_{\text{reh}}$ will effectively be reflected in the equation-of-state parameter w_{reh} incorporating part of the radiation dominated epoch and thereby tending towards $w = 1/3$. In order for reheating not to affect any confirmed observations of the standard Big Bang cosmology, we require at the very least that the epoch

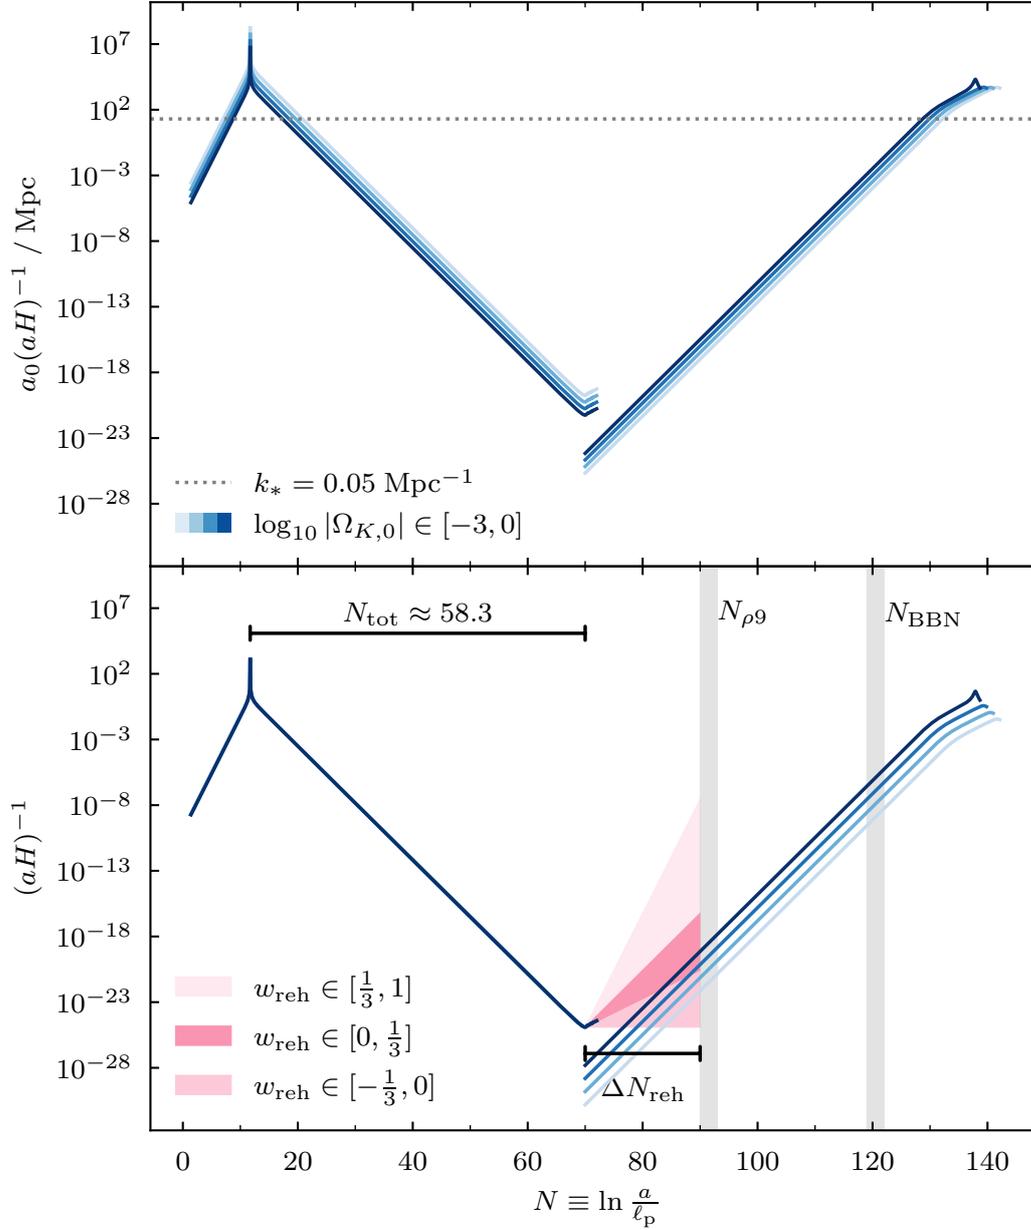


Figure 5.10 The top panel shows the comoving Hubble horizon before, during and after inflation. Different from fig. 5.6, here, we vary today’s curvature density parameter $\Omega_{K,0}$ which governs the calibration between primordial and today’s scales. Therefore we cannot simultaneously draw a y -axis for the evolution of the curvature density parameter anymore. The axis would be different for every $\Omega_{K,0}$. If we forego the linking to today’s scales, we can actually map the curves during inflation onto one another as done in the lower panel, which allows for a better visualisation of the reheating period after inflation and before the standard Big Bang evolution (radiation, matter and Λ domination). The shaded regions span the range of allowed values for the effective equation-of-state parameter w_{reh} during reheating. The duration of reheating ΔN_{reh} is bounded at least by the requirement that reheating should have finished before Big Bang nucleosynthesis (BBN) happens. Taking a stricter view, one might require that reheating has ended by the time $N_{\rho 9}$, when the energy density has dropped to $\rho_{\text{reh}}^{1/4} = 10^9 \text{ GeV}$ (see also fig. 5.1 for a sketch on the evolution of the energy density in parallel to the comoving Hubble horizon).

of reheating must have happened before Big Bang nucleosynthesis (BBN):

$$N_{\text{reh}} < N_{\text{BBN}} = \ln\left(\frac{a_0}{1 + z_{\text{BBN}}}\right), \quad (5.41)$$

where a_0 is inferred from the present-day curvature density parameter (see [section 5.6.1](#)) and where we use $z_{\text{BBN}} = 10^9$ as a rough estimate of the epoch of BBN.

The pink shaded regions in [fig. 5.10](#) subdivide the range for w_{reh} from [eq. \(5.40\)](#), with the dividing lines given by $w_{\text{reh}} = 0$ corresponding to a matter dominated epoch of reheating and $w_{\text{reh}} = 1/3$ corresponding to a radiation dominated epoch of reheating. Requiring matter domination exactly, i.e. fixing $w_{\text{reh}} = 0$, is an often used model for reheating, because most single-field inflationary potentials can be approximated by the quadratic potential close to their minimum, and thus predict $w_{\text{reh}} = 0$. The duration of reheating ΔN_{reh} (or equivalently the energy scale of reheating ρ_{reh}) is still a free parameter in this case. Radiation domination, on the other hand, would seamlessly continue into the standard Big Bang evolution and is therefore also referred to as instant reheating and often used as the most restrictive but simplest case of reheating, since it leads to $\Delta N_{\text{reh}} = 0$ and $\rho_{\text{reh}} = V_{\text{end}} = V(\phi_{\text{end}})$.

From the marginal variation of the standard Big Bang evolution owing to curvature, we can already deduce that the latter will barely affect the equation of state of reheating. Much more important is the role of the total amount of inflation N_{tot} , which determines whether inflation ends before or after the primordial curve crosses the radiation domination line in [fig. 5.10](#). The crucial role curvature plays in this scenario is through the linking of scales between primordial and standard Big Bang evolution.

In the very permissive scenario outlined by [eqs. \(5.40\)](#) and [\(5.41\)](#) linking primordial and late-time evolution will practically always be possible if inflation ends early, before crossing the radiation domination line. Otherwise an equation-of-state parameter $w_{\text{reh}} > 1/3$ will be required to catch up in time with the standard Big Bang evolution. In more restrictive settings such an equation-of-state parameter is typically excluded at the prior level [[131](#), [148](#), [173](#)], which we will explore further in [section 5.13.2](#).

5.10 The primordial power spectrum (PPS)

5.10.1 Power-law PPS

In the base Λ CDM cosmological model the primordial power spectrum of scalar curvature perturbations \mathcal{R} is phenomenologically described via two of its six free parameters in form of a simple power law:

$$\mathcal{P}_{\mathcal{R}}(k) = A_s \left(\frac{k}{k_*}\right)^{n_s - 1}, \quad (5.42)$$

where the power amplitude A_s and the spectral index $n_s \equiv 1 + d \ln \mathcal{P}_{\mathcal{R}}(k_*) / d \ln k$ are the free parameters with the subscript “s” referring to scalar perturbations and where k_* is a pivot

scale in the window of observable wavenumbers k . We choose to work with the commonly used pivot scale of $k_* = 0.05 \text{ Mpc}^{-1}$. Note that we are using the primordial power spectrum in its dimensionless form.

One of the prime successes of large-field inflation so far is the prediction of a spectral index n_s slightly smaller than unity, where unity would correspond to a scale-invariant power spectrum. This deviation from scale-invariance has been confirmed by the measurements of the Planck satellite to high precision [33–35].

Another major prediction of inflation is the presence of primordial gravitational waves, typically parametrised by a (non-zero) tensor-to-scalar ratio $r = A_t/A_s$. The PPS for gravitational waves is defined analogously to eq. (5.42) but the tensor spectral index n_t is typically given without the ‘ -1 ’ in the exponent:

$$\mathcal{P}_t(k) = r A_s \left(\frac{k}{k_*} \right)^{n_t} \quad (5.43)$$

For the tensor-to-scalar ratio r and the tensor spectral index n_t we assume the inflation consistency relation for a single scalar field with a standard kinetic term [131, 209]:

$$n_t = -\frac{r}{8} \left(2 - \frac{r}{8} - n_s \right). \quad (5.44)$$

Besides the tensor-to-scalar ratio, another common extension to eq. (5.42) is an expansion to higher orders in $d \ln k$, introducing the running of the spectral index

$$n_{\text{run}} = \frac{dn_s}{d \ln k} = \frac{d^2 \ln \mathcal{P}_{\mathcal{R}}(k_*)}{d \ln k^2}. \quad (5.45)$$

In section 5.11 we introduce various large-field inflation models and their predictions for the spectral index n_s , its running n_{run} and the tensor-to-scalar ratio r .

5.10.2 The slow-roll approximate PPS

Using just the background quantities a , H , and $\dot{\phi}$ from the solution to the background eqs. (5.14) to (5.16) we can compute a slow-roll (SR) approximation to the primordial power spectrum (see e.g. [210] for a derivation):

$$\mathcal{P}_{\mathcal{R}}(k) \approx \left(\frac{H^2}{2\pi\dot{\phi}} \right)_{k=aH}^2, \quad (5.46)$$

where the subscript expresses that the quantities need to be evaluated where each mode crosses the comoving Hubble horizon, i.e. where $k = aH$. This approximation is accurate on sufficiently small scales (large k), where the PPS takes the form of an almost scale-invariant power-law, which motivates the phenomenological power-law spectrum from eq. (5.42).

In models of finite inflation this approximation breaks down on large scales (small k), where the modes have not started out from sufficiently well within the comoving Hubble horizon, and where the primordial power spectrum then exhibits a cutoff towards large scales. This

cutoff behaviour can already be observed qualitatively in the approximate PPS, however to properly quantify this cutoff, we need to perform a full numerical integration of the primordial perturbations.

Figure 5.13 shows the approximate PPS and its dependence on various input parameters, which will be discussed in **section 5.12** in more detail.

5.10.3 The full numerical PPS

In order to solve the PPS numerically, we need to integrate the Mukhanov–Sasaki equation for the curvature perturbation \mathcal{R}_k , which can be written as a damped harmonic oscillator with respect to cosmic time (adapted from [193]):

$$0 = \ddot{\mathcal{R}}_k + 2\gamma\dot{\mathcal{R}}_k + \omega^2\mathcal{R}_k, \quad (5.47)$$

$$\text{with damping } 2\gamma = (3 + \xi)H, \quad (5.48)$$

$$\text{and frequency } \omega^2 = \frac{\kappa^2}{a^2} - \frac{K}{a^2}(1 + \xi), \quad (5.49)$$

where damping and frequency share the term

$$\xi = \frac{2\kappa^2}{\kappa^2 + K\mathcal{E}} \left(\mathcal{E} + \frac{\ddot{\phi}}{H\dot{\phi}} + \Omega_K \right),$$

with

$$\mathcal{E} = \frac{\dot{\phi}^2}{2H^2}, \quad \frac{\ddot{\phi}}{H\dot{\phi}} = -3 - \frac{V'(\phi)}{H\dot{\phi}}, \quad \Omega_K = -\frac{K}{(aH)^2}.$$

For the wavenumber k we use the following effective expression in curved spaces:

$$\kappa^2 = k^2 + kK(K + 1) - 3K, \quad (5.50)$$

$$\text{with } \begin{cases} k \in \mathbb{R}, k > 0 & \text{if } K = 0, -1, \\ k \in \mathbb{Z}, k > 2 & \text{if } K = +1. \end{cases}$$

We get the expression for the effective wavenumber κ from Fourier transforming the $\nabla_i\nabla^i$ operator in curved space [193]. Note how the wavenumber becomes discrete with $k > 2$ for positively curved (closed) universes.

In the small-scale limit $k \rightarrow \infty$ or in the absence of curvature $K = 0$ we recover the better known terms from the flat universe case (compare e.g. with equation 16.45 in [60]):

$$\kappa^2 \longrightarrow k^2, \quad (5.51)$$

$$\omega^2 \longrightarrow \frac{k^2}{a^2}, \quad (5.52)$$

$$2\gamma \longrightarrow \frac{\dot{\phi}^2}{H} + 3H + \frac{2\ddot{\phi}}{\dot{\phi}}. \quad (5.53)$$

For tensor modes, the modification of the Mukhanov–Sasaki equation from the flat case to curvature is much simpler. The equivalent form of eq. (5.47) for tensor modes is

$$\ddot{h} + 3H\dot{h} + \left(\frac{\kappa^2}{a^2} + \frac{5K}{a^2} \right) h = 0, \quad (5.54)$$

which again reduces as expected to eqs. (5.51) and (5.52) in the limit of small scales.

Using `oscode`'s [195] efficient algorithm for oscillatory ordinary differential equations, we can integrate eqs. (5.47) and (5.54) for each mode k from the start of inflation until well past horizon crossing for that given mode, where the frozen values of the primordial perturbations can be read off. We can then compute the primordial power spectra for scalar and tensor perturbations according to:

$$\mathcal{P}_{\mathcal{R}}(k) = \frac{k^3}{2\pi^2} |\mathcal{R}_k|^2, \quad (5.55)$$

$$\mathcal{P}_t(k) = 2 \cdot \mathcal{P}_h(k) = 2 \cdot \frac{k^3}{2\pi^2} |h_k|^2, \quad (5.56)$$

where the factor 2 in the tensor spectrum comes from the two possible polarisation states of gravitational waves.

We initialise \mathcal{R}_k and h_k in their vacuum state defined as the state which minimises energy density via the renormalised stress-energy tensor [158, 193]:

$$\mathcal{R}_{k,i} = \frac{1}{z\sqrt{2k}}, \quad \dot{\mathcal{R}}_{k,i} = -\frac{ik}{a}\mathcal{R}_{k,i}, \quad (5.57)$$

$$h_{k,i} = \frac{1}{a}\sqrt{\frac{2}{k}}, \quad \dot{h}_{k,i} = -\frac{ik}{a}h_{k,i}. \quad (5.58)$$

We prefer these initial conditions over similar formulations such as the commonly used Bunch Davies vacuum [211], because their predictions are stable across different choices of dependent or independent variables, i.e. they are invariant under canonical transformations [212].

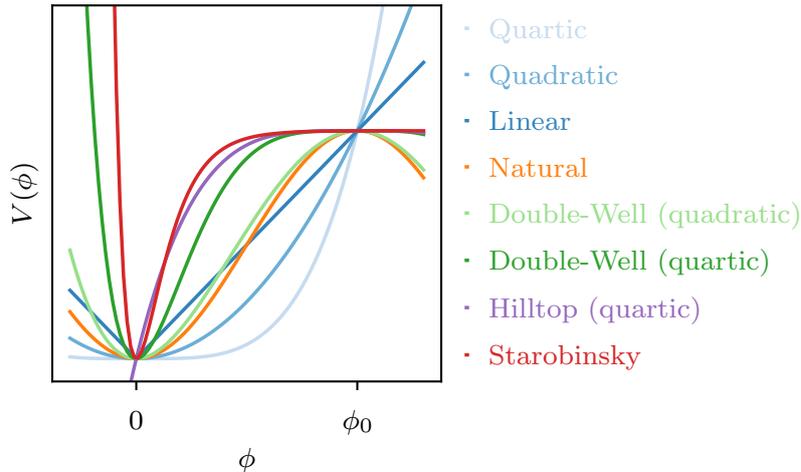


Figure 5.11 Schematic view of various inflationary potentials that we investigate. If necessary, potentials are shifted such that the potential minimum and hence the end stage after inflation lies at the origin at $\phi = 0$. Potentials with a local maximum are defined such that that maximum lies at the potential parameter ϕ_0 to the right of the origin. The potentials were plotted with the same potential amplitude Λ , except for the quartic, quadratic and linear potentials which were rescaled for visualisation purposes such that they meet the other potentials in the point $V(\phi_0)$. In this form, the linear potential illustrates the categorisation into convex and concave potentials, which is frequently used in n_s - r plots, e.g. in [fig. 5.12](#).

5.11 Inflationary potentials and slow-roll predictions

In this section we briefly review a few scalar single-field inflation models. [Figure 5.11](#) shows a schematic view of the various inflationary potentials used in this paper. To ease the computation of the inflation models and their comparison with one another, we try to unify the notation by rewriting traditional formulations as follows. They will share a potential amplitude parameter Λ (not to be confused with the cosmological constant Λ) in units of the reduced Planck mass, $[\Lambda] = m_p$. The potential minimum $V = 0$ will be located at the origin $\phi = 0$ and any potential local maximum (for the natural, hilltop and double-well potentials) will be located at a parameter ϕ_0 . The amplitude parameter, common to all inflationary potential, is linked directly to the power amplitude A_s of scalar primordial perturbations.

Note that we only consider large-field inflation in this paper, i.e. models where the field excursion of the inflaton takes values greater than the Planck scale. Small-field inflation predicts a tensor-to-scalar ratio so small that it will remain unobservable for the near future [\[213\]](#).

[Figure 5.12](#) illustrates the slow-roll (SR) predictions for the tensor-to-scalar ratio r , the spectral index n_s and its running $n_{\text{run}} = dn_s/d\log k$. Due to the wide dynamic range predicted for the tensor-to-scalar ratio, we show r scaled both linearly in the upper left and logarithmically in the lower left plot. In the Planck inflation papers [\[131, 148, 173\]](#) a linear scaling in r was preferred, however with upcoming CMB experiments such as the Simons Observatory [\[45\]](#), the LiteBIRD satellite [\[46\]](#) or CMB-S4 [\[47\]](#) pushing to a tensor-to-scalar ratio of about $r \sim 10^{-3}$, a logarithmic scaling of r allows for better visualisation and sampling of the smaller scales. For a recent discussion on uniform versus logarithmic priors on r and

their effects on Bayesian model comparison see [chapter 2](#). The slow-roll predictions for the running of the spectral index from all inflation models considered here, on the other hand, only span a small fraction of the posterior distribution, which can be seen in the upper right plot of [fig. 5.12](#). For a better comparison of the predictions of the individual models we zoom in on the n_{run} -range in the lower right plot. This highlights how the uncertainty of the running of the spectral index is far too large for the purpose of distinguishing between these inflation models.

Note that the Planck and Planck+BK15 contours in [fig. 5.12](#) come from an extension of the base Λ CDM cosmology not only with the tensor-to-scalar ratio and the running of the spectral index, but also with the spatial curvature parameter $\Omega_{K,0}$. As such, the contours differ from those in the Planck inflation papers [[131](#), [148](#), [173](#)] and in the BK15 paper [[71](#)], which we elaborate on in our results in [section 5.13.1](#). Note further that we only look at the SR predictions in this section, comparing them to the n_s - r and n_s - n_{run} contours. The results from our nested sampling analysis follow later in [section 5.13.2](#).

5.11.1 Monomial potential

The monomial potentials* are one of the simplest classes of inflationary potentials, given by:

$$V(\phi) = \Lambda^4 \left(\frac{\phi}{m_{\text{p}}} \right)^p \quad (5.59)$$

Inflation ends when the value of the inflaton field drops to

$$\phi_{\text{end}} = \frac{m_{\text{p}} p}{\sqrt{2}}. \quad (5.60)$$

With that we can approximate the number of e-folds of inflation from some ϕ until the end of inflation to be

$$N(\phi) \simeq \frac{\phi^2}{2 m_{\text{p}}^2 p} - \frac{p}{4}, \quad (5.61)$$

showing that the number of e-folds of inflation grow quadratically with the inflaton field value for monomial potentials. This scaling carries through to the spectral index, to the tensor-to-scalar ratio and to the running of the spectral index, where we get to leading order in $1/N_*$:

$$n_s \simeq 1 - \frac{2+p}{2N_*}, \quad r \simeq \frac{4p}{N_*}, \quad n_{\text{run}} \simeq -\frac{2+p}{2N_*^2}. \quad (5.62)$$

*Inflation models with a monomial potential are also referred to as large-field inflation (e.g. in the encyclopædia inflationaris [[175](#)]), however, since large-field displacements are not unique to the monomial potential, we prefer naming the inflation models after their potential shapes, here. Yet another name often associated with the monomial potential is chaotic inflation [[214](#)], but similarly chaotic inflation at its core actually pertains to the idea that the inflaton started from a chaotic initial state varying wildly from one place to another, rather than describing a specific potential model. See [[215](#)] for a helpful discussion of terminology.

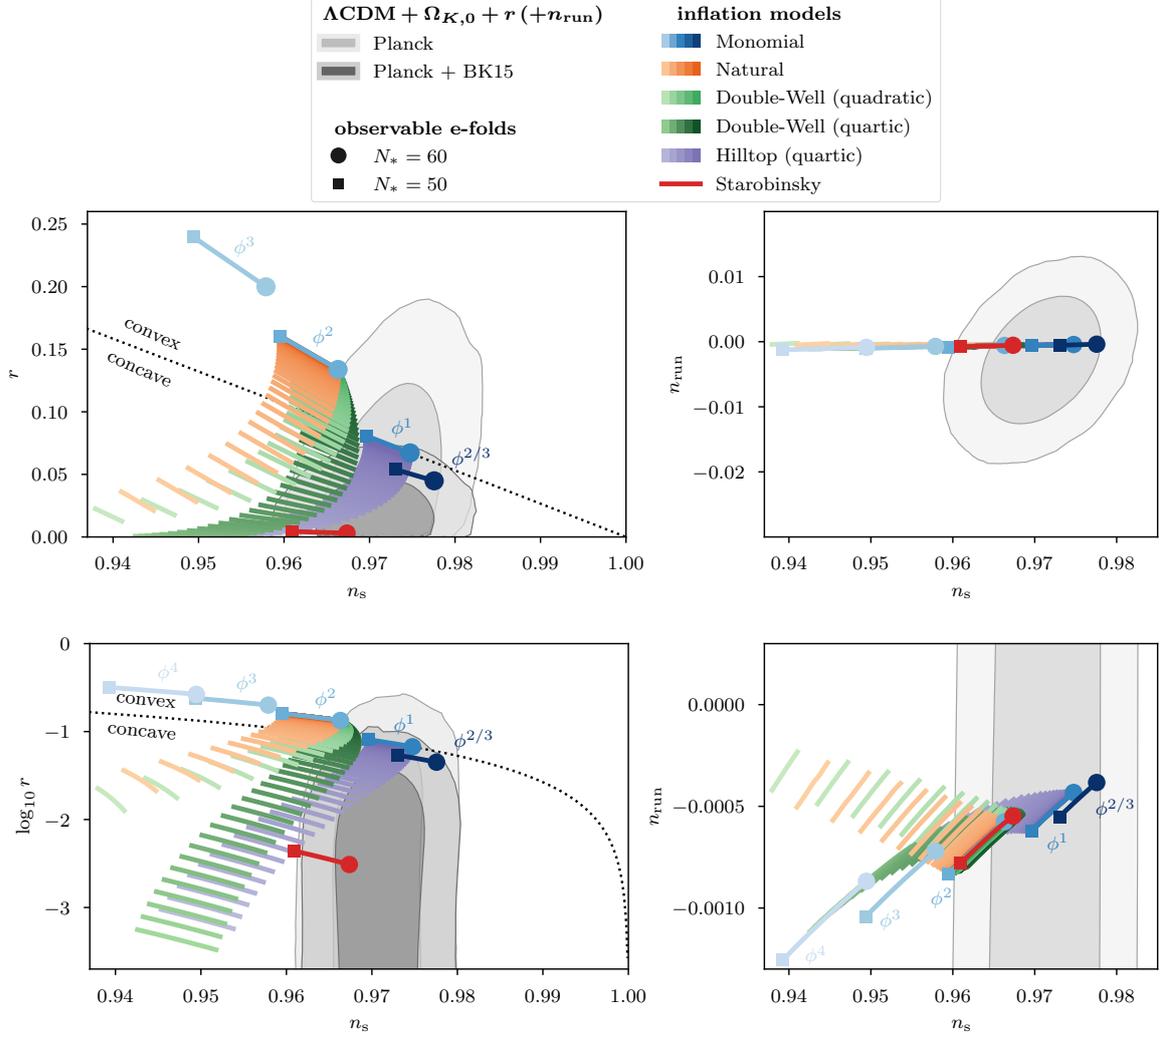


Figure 5.12 Comparison of slow-roll inflation model predictions for the spectral index n_s against the tensor-to-scalar ratio r on the left and the running of the spectral index n_{run} on the right. The top left plot is linear and the bottom left plot logarithmic in r . The top right plot shows the full joint contour for n_s and n_{run} and the bottom right plot a zoom-in into the region relevant for the inflationary potentials. Note that for all these plots the curvature density parameter $\Omega_{K,0}$ is one of the sampling parameters. For each inflation model we show the line(s) delimited by the requirement of producing 50 to 60 e-folds of observable inflation N_* , which is the rough range needed for viable reheating scenarios. In blue we show the slow-roll predictions for the monomial potentials from eq. (5.59) and in red that of Starobinsky (or R^2) inflation from eq. (5.68). For the other inflation models we show a range of predictions for a range of values of the potential hill parameter ϕ_0 , spaced logarithmically. We show natural inflation from eq. (5.65) in orange, quartic hilltop inflation from eq. (5.67) in purple and both quadratic and quartic double-well potentials from eq. (5.66) in light and dark green respectively.

In [fig. 5.12](#) we show in blue the slow-roll predictions for the quartic, cubic, quadratic and linear monomial as well as for the monomial with $p = 2/3$.

Quartic and cubic potentials: Similarly to the flat case in the Planck inflation papers, the predictions for quartic and cubic inflation in light blue lie far outside the 95% contours from the Λ CDM extensions in grey.

Quadratic potential: The quadratic potential with $p = 2$ has long been used as the simplest realisation of single-field inflation. It is often given in the following (slightly different) form with an additional pre-factor of one half:

$$V(\phi) = \frac{1}{2}m^2\phi^2, \quad (5.63)$$

where m is referred to as the inflaton mass, which can be related to the potential amplitude Λ in [eq. \(5.59\)](#) directly. Although allowing for spatial curvature to vary in [fig. 5.12](#), which significantly stretches the Planck contours to larger r , this stretching coincides with a shift to larger n_s such that the prediction for quadratic inflation ends up just outside the 95% Planck contours, just like in the flat case [\[173\]](#). With the addition of BK15 data, the SR prediction lies far outside the contour irrespective of any curvature effects.

Linear and $p = 2/3$ potential: These two potentials are motivated by axion monodromy [\[216, 217\]](#) and agree better with the Planck and Planck+BK15 contours. Both in fact profit from the shift to a larger spectral index that comes with varying curvature.

5.11.2 Natural potential

Natural inflation is motivated by particle physics considerations [\[218\]](#) to naturally accommodate the very flat potentials required for inflation. It is given by the periodic potential

$$V(\phi) = \Lambda^4 \left[1 + \cos\left(\frac{\phi}{f}\right) \right], \quad (5.64)$$

where f corresponds to the global symmetry-breaking scale and governs the slope of the potential.

We can rewrite the potential such that the local maximum lies at $\phi_0 = \pi f$ and is given by the potential amplitude $V(\phi_0) = \Lambda^4$. From this unstable maximum the inflaton rolls down to the minimum at the origin $\phi = 0$ (see also [fig. 5.11](#)):

$$V(\phi) = \frac{\Lambda^4}{2} \left[1 - \cos\left(\pi \frac{\phi}{\phi_0}\right) \right]. \quad (5.65)$$

In order to produce sufficient e-folds of (large-field) inflation, we require a potential hill parameter $\phi_0 \gtrsim 10 m_p$ (or correspondingly for f). In the limit of very large $\phi_0 \gtrsim 100 m_p$, the

spectral index and tensor-to-scalar ratio of the natural potential tend to those of the quadratic potential.

While natural inflation still overlaps with the Planck contours in a flat universe [173], because of its ability to accommodate a smaller tensor-to-scalar ratio, it only touches the 95 % contours in the curved case due to the shift in the spectral index.

5.11.3 Double-Well potential

Similarly to the natural potential, we define the double-well potential such that the local maximum lies at ϕ_0 with the maximum potential value given by the potential amplitude $V(\phi_0) = \Lambda^4$ (see also fig. 5.11):

$$V(\phi) = \Lambda^4 \left[1 - \left(\frac{\phi - \phi_0}{\phi_0} \right)^p \right]^2. \quad (5.66)$$

where p can in principle take any positive value. We will consider the quadratic ($p = 2$) and quartic ($p = 4$) double-well in particular.

Double-well potentials are typically associated with small-field inflation. However, inflation with large field displacements $\phi_0 > m_p$ is also possible. In that case the spectral index and tensor-to-scalar ratio tend to that of the quadratic potential for very large $\phi_0 \gtrsim 100 m_p$, similarly to natural inflation and irrespective of the parameter p . For smaller values of the potential hill parameter $\phi_0 \gtrsim 10 m_p$, both the spectral index and tensor-to-scalar ratio decrease. The SR predictions for the quadratic double-well are very close to those of the natural potential, which is to be expected considering their similar shapes (cf. fig. 5.11). The quartic double-well with its flatter hill leads to a faster drop in r , and therefore a greater overlap with the Planck and Planck+BK15 contours.

5.11.4 Hilltop potential

Closely related to double-well potentials, hilltop potentials are given by:

$$V(\phi) = \Lambda^4 \left[1 - \left(\frac{\phi - \phi_0}{\phi_0} \right)^p + \dots \right], \quad (5.67)$$

in which only the first order in ϕ^p is retained and higher order terms (indicated by the ellipsis) are neglected, since the latter only become relevant towards the end of inflation. For small values of the potential hill parameter ϕ_0 the SR predictions are close to those of the double-well potential, but for larger values the spectral index, running, and tensor-to-scalar ratio will tend towards those of the linear potential (monomial with $p = 1$) instead of the quadratic potential, since that is what the eq. (5.67) approximates to close to $V(\phi) = 0$. However, this asymptotic behaviour would have to be different if the higher order terms were present, which are required to ensure the positiveness of the potential. We therefore prefer to work with the double-well potential for the scope of this paper. The asymptotic behaviour does, however,

mean that the SR predictions agree better with the Planck and Planck+BK15 contours than those of the double-well potentials.

5.11.5 Starobinsky potential

The Starobinsky potential, given in the Einstein frame by

$$V(\phi) = \Lambda^4 \left[1 - e^{-\sqrt{\frac{2}{3}} \frac{\phi}{m_{\text{p}}}} \right]^2, \quad (5.68)$$

was the first proposed inflationary potential and motivated by an extension of the Einstein–Hilbert action with a term quadratic in the Ricci tensor [136, 148]. Therefore this model of inflation is frequently also referred to as R^2 inflation.

It can be shown that inflation generated by the Higgs field of the (particle physics) standard model can be reduced to the potential from eq. (5.68) in the Einstein frame, where all parameters connected to the Higgs boson are included in the amplitude parameter Λ [175]. This motivates yet another name for this type of potential: Higgs inflation.

Due to the shift in spectral index from varying curvature, the SR predictions for the Starobinsky potential no longer lie as spot-on in the centre of the 68% contour lines as in the flat case. Nevertheless, they remain in excellent agreement with the Planck and Planck+BK15 contours.

5.12 Choice of parametrisation

Table 5.1 lists the sampling parameters used in our Bayesian analysis together with their prior range and fiducial values which they are fixed to for visualisation purposes in some figures, such as the figs. 5.13 and 5.14.

For the base Λ CDM model, we use the usual six sampling parameters also used in [33]:

- $\omega_{\text{b}} = h^2 \Omega_{\text{b}}$: Baryon density today
- $\omega_{\text{c}} = h^2 \Omega_{\text{c}}$: Cold dark matter density today
- $100 \theta_{\text{s}}$: Angular size of sound horizon at last scattering
- τ_{reio} : Optical depth to reionization
- $\ln(10^{10} A_{\text{s}})$: Scalar power spectrum amplitude
- n_{s} : Scalar spectral index

Additionally we consider the following parameter extensions to the base Λ CDM model:

- r : Tensor-to-scalar power ratio
- $\Omega_{K,0}$: Spatial curvature parameter today

Table 5.1 Overview of the cosmological parameters with their fiducial values used for visualisation purposes (e.g. in [figs. 5.13](#) and [5.14](#)) and their prior ranges used in our Bayesian analysis. We list the six base parameters of the Λ CDM cosmology in the first block, the Λ CDM extension parameters in the second block, and primordial parameters pertaining to a full inflationary analysis in the third block. Note that the prior range and the fiducial values for the inflaton field ϕ_i is dependent on the choice of the inflationary potential. We explore the result of different constraints on this parameter in [section 5.13.2](#).

Parameter	fiducial value	Prior range	Definition
$\omega_b = h^2\Omega_b$	0.022632	$0.019 < \omega_b < 0.025$	Baryon density today
$\omega_c = h^2\Omega_c$	0.11792	$0.025 < \omega_c < 0.471$	Cold dark matter density today
θ_s	0.01041338	$1.03 < 100\theta_s < 1.05$	Angular size of sound horizon at last scattering
τ_{reio}	0.0495	$0.01 < \tau_{\text{reio}} < 0.40$	Optical depth to reionization
A_s	2×10^{-9}	$2.5 < \ln(10^{10}A_s) < 3.7$	Amplitude of the scalar power spectrum
n_s	0.97235	$0.885 < n_s < 1.040$	Scalar spectral index
r		$0 < r < 1$ $-5 < \log_{10} r < 0$	Tensor-to-scalar power ratio
$\Omega_{K,0}$	-0.01	$-0.15 < \Omega_{K,0} < 0.15$	Curvature density today
ϕ_i potential dependent		Inflaton field at start of inflation
$f_i \equiv \frac{\Omega_{K,i}}{\Omega_{K,0}}$	5	$-1 < \log_{10} f_i < 5$	Fraction of primordial to present-day curvature

All primordial parameters $\{A_s, n_s, r\}$ are taken at the pivot scale of $k_* = 0.05 \text{ Mpc}^{-1}$.

The primordial parameters A_s and n_s refer to the simplified power-law spectrum from [eq. \(5.42\)](#). When we derive the primordial power spectrum from an inflationary potential as outlined in the previous sections [sections 5.10.1](#) and [5.10.2](#), these parameters and also r turn into derived parameters. In addition to the base parameters $\omega_b, \omega_c, \theta_s, \tau_{\text{reio}}$ and A_s , as well as the spatial curvature parameter $\Omega_{K,0}$, we sample over the following parameters in our analysis of individual inflationary potentials:

- $\log_{10} \phi_i$: Inflaton field value at the start of inflation.
- $\log_{10} f_i \equiv \log_{10} \frac{\Omega_{K,i}}{\Omega_{K,0}}$:
Fraction of primordial to present-day curvature.

Instead of the power amplitude A_s , we could have used the amplitude parameter for the inflationary potential Λ . The two are directly related to one another, since the background [eqs. \(5.14\)](#) to [\(5.16\)](#) are invariant under a simultaneous rescaling of the time coordinate and the inflaton potential:

$$\begin{aligned}
 t &\mapsto \sigma^{-1}t \\
 V(\phi) &\mapsto \sigma^2 V(\phi) \\
 \Rightarrow \mathcal{P}_{\mathcal{R}}(k) &\mapsto \sigma^2 \mathcal{P}_{\mathcal{R}}(k).
 \end{aligned}
 \tag{5.69}$$

However, we prefer sampling over A_s , as it allows for a more direct comparison with the base Λ CDM model and its extensions. Also, as opposed to Λ , the power amplitude A_s is not correlated with the other primordial parameters that affect the comoving Hubble horizon or the e-folds of inflation. This becomes very clear in the top row of [fig. 5.13](#) showing the variation of the comoving Hubble horizon with respect to the logarithm of A_s on the left, and that of the primordial power spectrum (PPS) on the right. While A_s governs the amplitude of the PPS by definition, it leaves the comoving Hubble horizon invariant. In inflation models such as natural or double-well inflation, with a local maximum separated from the global minimum by ϕ_0 , the potential amplitude Λ is also strongly correlated with ϕ_0 , and sampling A_s instead of Λ avoids having to navigate that degeneracy.

Similarly, we choose to use the present-day spatial curvature parameter $\Omega_{K,0}$ for a better comparison with the $\Omega_{K,0}$ -extension of Λ CDM. Alternatively, one could use reheating parameters to track the evolution of energy densities in the universe and infer the present-day scale factor a_0 and curvature density. We defer exploring this option to future work.

There is considerable freedom in the choice between the primordial parameters at the start of inflation, i.e. parameters related to the inflaton field ϕ_i , the e-folds N_i , or the primordial curvature $\Omega_{K,i}$, which are all linked via [eqs. \(5.29\)](#) and [\(5.30\)](#). Other parameters that are connected to these are e.g. the total number of e-folds of inflation N_{tot} or the e-folds of inflation before (N_{\dagger}) and after (N_*) horizon crossing of the pivot scale.

The fraction f_i of primordial curvature is a useful sampling parameter for two reasons: firstly, due to it governing the ratio of conformal time passing before and after the end of inflation which we explored in the previous [section 5.8](#) (see especially [fig. 5.8c](#)), and secondly because of its major role in governing the cutoff position in the primordial power spectrum, which we explore in more detail in the following. [Figures 5.13](#) and [5.14](#) show the effect of our sampling parameters on the comoving Hubble horizon (left column of [fig. 5.13](#)), on the slow-roll approximation of the primordial power spectrum (right column of [fig. 5.13](#)), on the fully numerically integrated primordial power spectrum (left column in [fig. 5.14](#)) and on the CMB power spectrum (right column in [fig. 5.14](#)).

We already mentioned the straightforward effect of varying the power amplitude (top row) on the PPS and the CMB power spectrum, and that the comoving Hubble horizon and the e-folds of inflation are unaffected by the power amplitude.

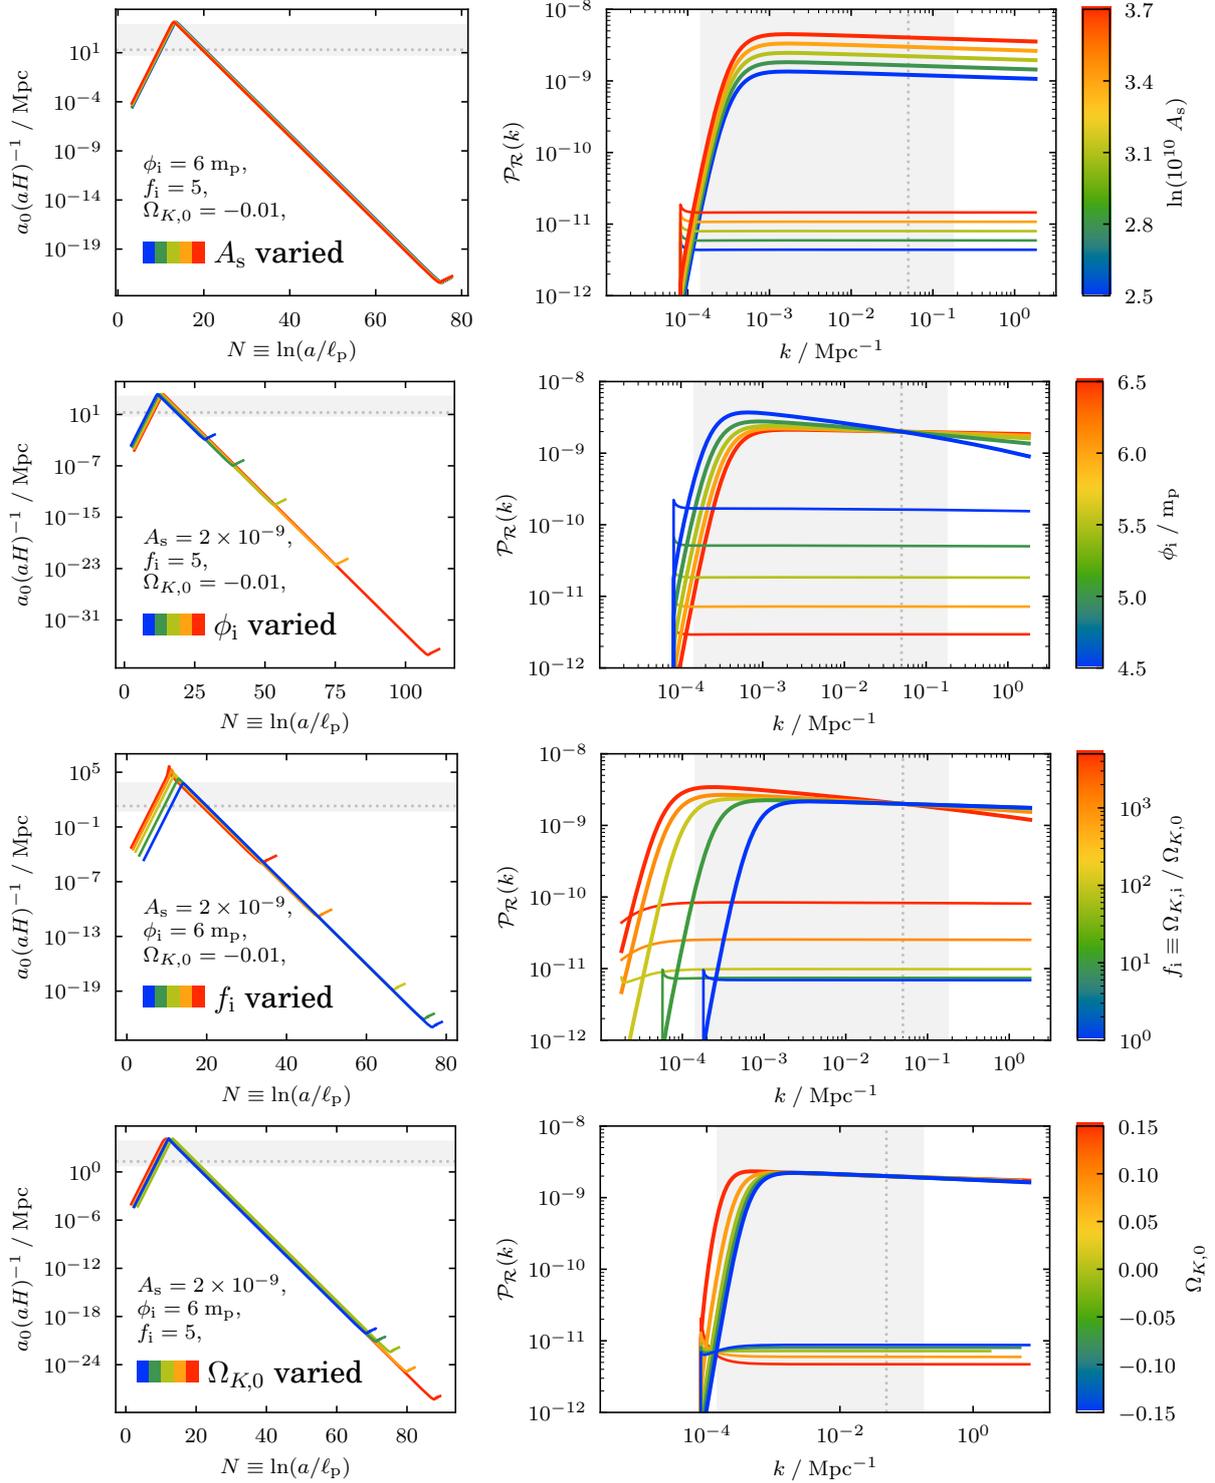


Figure 5.13 Parameter dependence of the comoving Hubble horizon in the left column and of the slow-roll approximation of the primordial power spectrum (PPS) in the right column on the sampling parameters for our Bayesian analysis: power amplitude A_s , inflaton field value at the start of inflation ϕ_i , fraction of primordial to present-day curvature f_i and present-day curvature density $\Omega_{K,0}$, where one parameter is varied in each row, while the others stay fixed. The upper and heavier lines in the PPS plots correspond to scalar, the lower and thinner lines to tensor perturbations. We used the Starobinsky potential to generate these plots, explaining the fairly big gap between scalar and tensor modes. The corresponding plots for the fully numerically integrated PPS and for the CMB power spectrum are shown in [fig. 5.14](#).

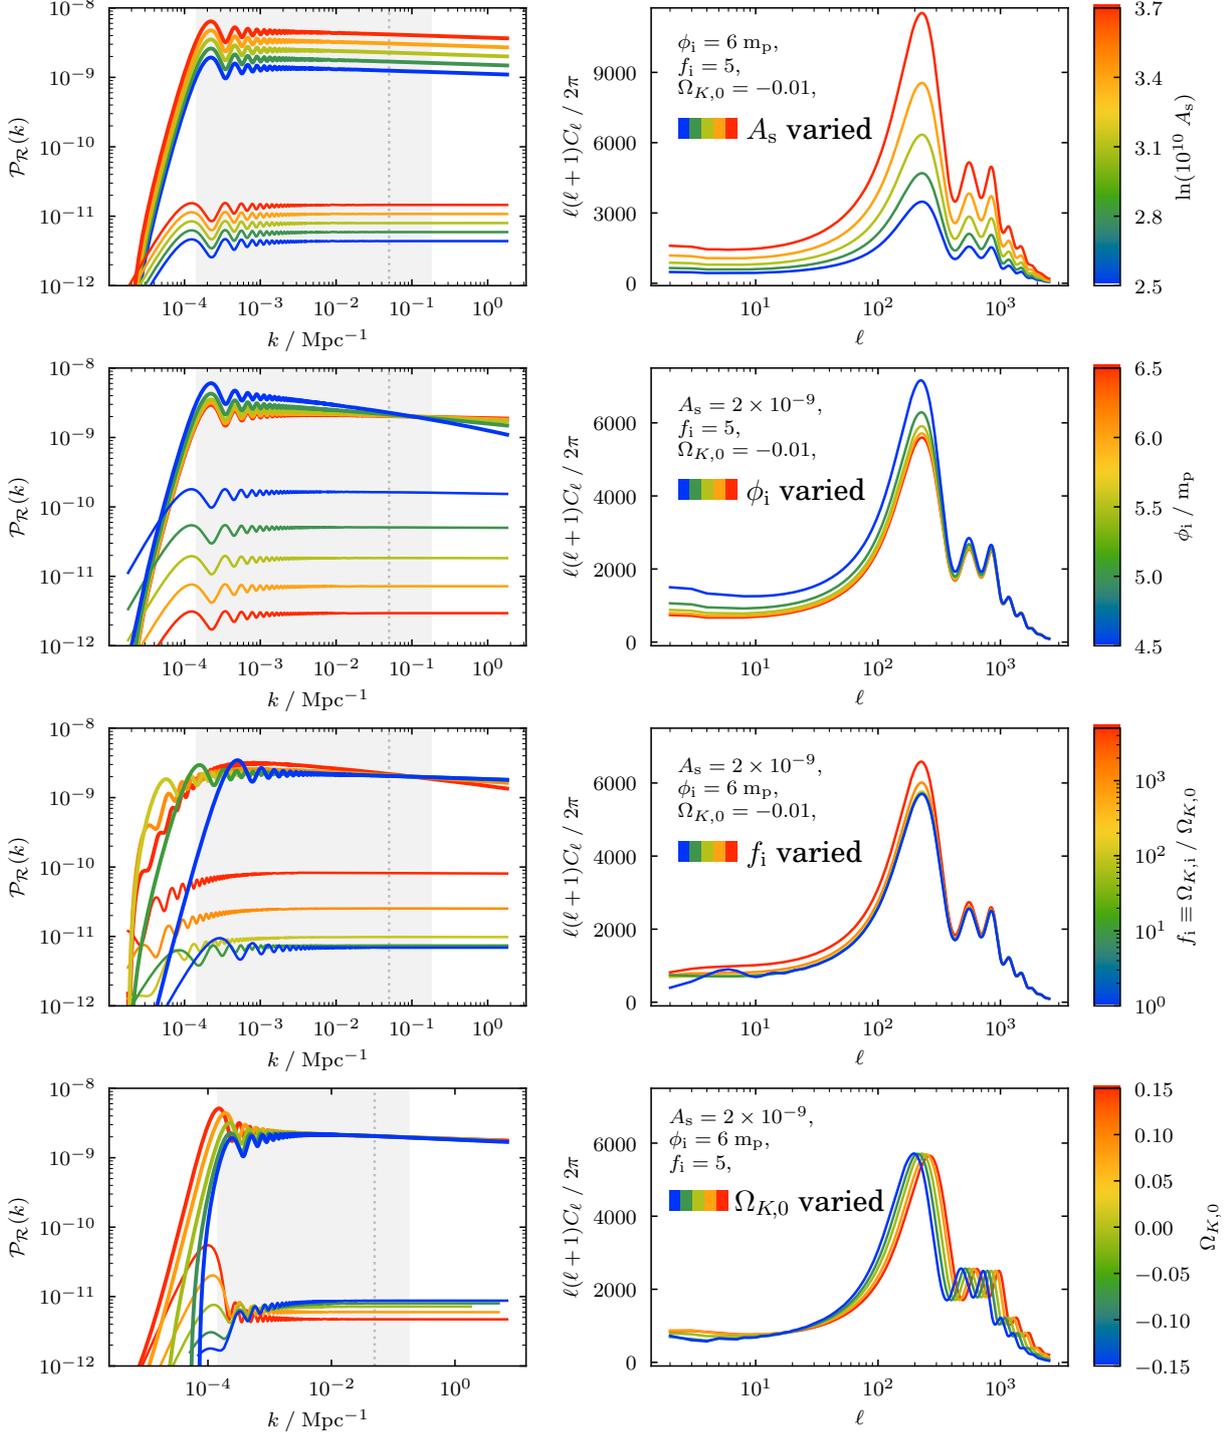


Figure 5.14 Parameter dependence of the numerically integrated primordial power spectrum (PPS) in the left column and of the CMB power spectrum in the right column on the sampling parameters for our Bayesian analysis: power amplitude A_s , inflaton field value at the start of inflation ϕ_i , fraction of primordial to present-day curvature f_i and present-day curvature density $\Omega_{K,0}$, where one parameter is varied in each row, while the others stay fixed. The upper and heavier lines in the PPS plots correspond to scalar, the lower and thinner lines to tensor perturbations. We used the Starobinsky potential to generate these plots, explaining the fairly big gap between scalar and tensor modes. The corresponding plots for the comoving Hubble horizon and for the slow-roll approximation of the PPS are shown in [fig. 5.13](#). For computational simplicity of the CMB spectra, we fix (alongside ω_b , ω_{cdm} , τ_{reio}) the Hubble parameter $h = 0.5409$ instead of the angular size θ_s of the sound horizon at last scattering.

The second row of [figs. 5.13](#) and [5.14](#) shows the variation of the comoving Hubble horizon, PPS, and CMB power spectrum with respect to the initial inflaton field value ϕ_i at the start of inflation. Looking at the variation of the comoving Hubble horizon, one can observe how ϕ_i is closely linked to the number of e-folds during inflation, e.g. the total number of e-folds of inflation N_{tot} or the observable number of e-folds N_* from horizon crossing of the pivot scale $k_* = 0.05 \text{ Mpc}^{-1}$ until the end of inflation. Hence, we could have used one of those as a sampling parameter instead of ϕ_i , which would have had the advantage that the number of e-folds are more independent of the choice of inflationary model, and that they are directly related to spectral index n_s and tensor-to-scalar ratio r . The initial inflaton field value is correlated with n_s and r as well, but [fig. 5.15](#) shows how fixing the total number of e-folds actually also fixes spectral index and tensor-to-scalar ratio, whereas both still vary slightly with f_i and $\Omega_{K,0}$ when only fixing ϕ_i . In other words, fixing the number of e-folds decorrelates the spectral index and the tensor-to-scalar ratio from the other sampling parameters. However, from a computational perspective it is much more stable and faster to work with ϕ_i , since both N_{tot} and N_* would require an additional optimisation loop, which may not be guaranteed to be solvable for non-trivial parameter combinations. Besides, using ϕ_i allows us to compare the various inflation models on grounds of how much inflation they predict in the first place, and how this holds up when we apply constraints from reheating, which we will explore further in [section 5.13.2](#).

In the third row of [figs. 5.13](#) and [5.14](#), the comoving Hubble horizon, PPS, and CMB power spectrum are varied with respect to the fraction f_i . For a fixed *present-day* curvature density this is equivalent to a variation of the *primordial* curvature density $\Omega_{K,i}$ and thus also to a variation of the initial size of the comoving Hubble horizon at the start of inflation. This variation with respect to f_i is also shown in [fig. 5.15](#) with a fixed amount of inflation. Since fixing N_{tot} decorrelates f_i from the spectral index and the tensor to scalar ratio, this isolates the effect of f_i on the large-scale (small k) cutoff position and the shape of the PPS, independent of slopes or amplitudes. [Figure 5.15](#) contrasts this behaviour for closed (in blue) and open (in green) universes, in which f_i affects the shape differently. In both cases increasing f_i initially (i.e. for small f_i) shifts the PPS cutoff to larger scales, out of the CMB observable window. However, once $\Omega_{K,i}$ gets close to or exceeds unity, this shift is replaced by a suppression of perturbation modes in the closed case and an amplification in the open case, for large scales just about smaller than the PPS cutoff. Note how looking at the slow-roll approximation of the PPS only may be misleading when trying to gauge the effect of f_i on the PPS. The shift of the cutoff to larger scales is similar in both the approximate and full numerical PPS as long as the curvature density is comparably small. However, once primordial curvature plays a significant role, the approximate PPS stops changing. The secondary, geometry-dependent effects on large scales are only visible in the fully numerically integrated PPS. This is not surprising considering that the slow-roll approximation is only valid for modes that were well within the comoving Hubble horizon at the start of inflation.

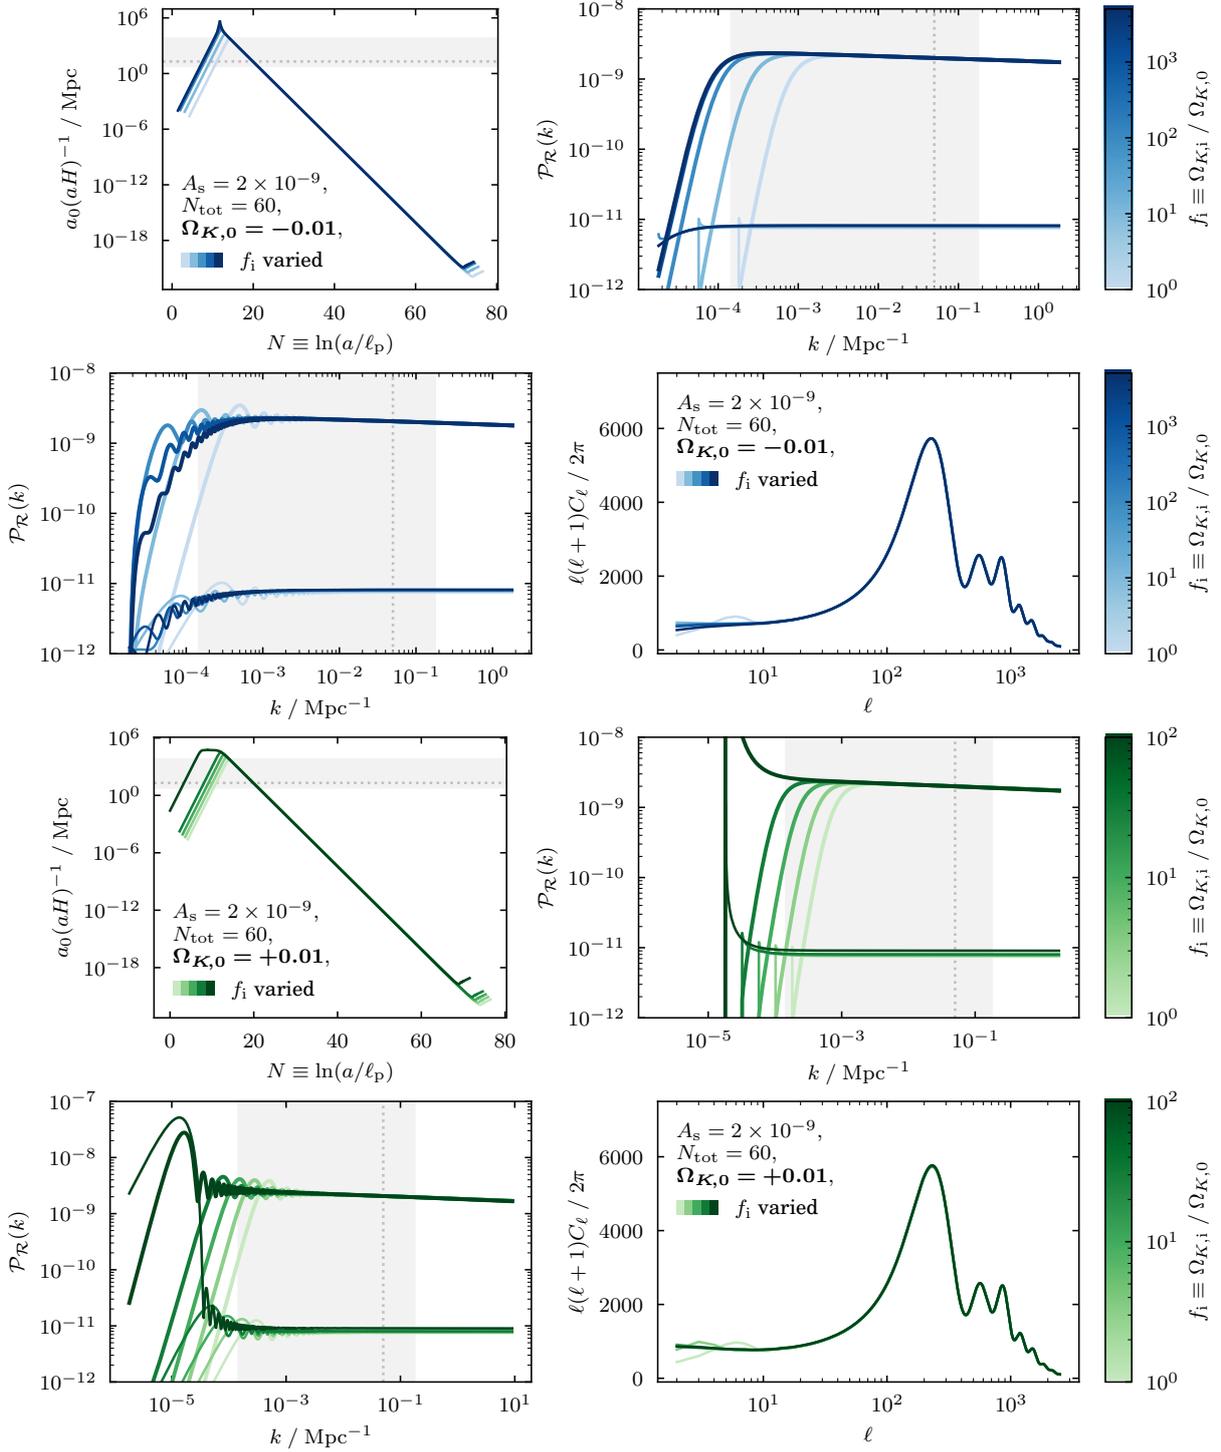


Figure 5.15 Similar to [figs. 5.13](#) and [5.14](#) we show the parameter dependence of the comoving Hubble horizon, the slow-roll approximation of the primordial power spectrum (PPS), the fully numerically integrated PPS and the CMB power spectrum on the fraction of primordial to present-day curvature $f_i \equiv \Omega_{K,i}/\Omega_{K,0}$. Instead of fixing the initial inflaton field value ϕ_i , in this case we fix the total number of e-folds of inflation N_{tot} . These are directly linked to the spectral index n_s and the tensor-to-scalar ratio r , which therefore are fixed here as well. In the upper two rows in blue the present-day curvature was fixed to $\Omega_{K,0} = -0.01$, in the lower two rows in green to $\Omega_{K,0} = +0.01$, thus showing the effects of a closed and open universe respectively. We used the Starobinsky potential to generate these plots, explaining the fairly big gap between scalar and tensor modes. The upper and heavier lines in the PPS plots correspond to scalar, the lower and thinner lines to tensor perturbations. For computational simplicity of the CMB spectra, we fix (alongside ω_b , ω_{cdm} , τ_{reio}) the Hubble parameter $h = 0.5409$ instead of the angular size θ_s of the sound horizon at last scattering.

The different shapes of the PPS for closed and open universes can be seen more directly in the last row of [figs. 5.13](#) and [5.14](#), where the present-day curvature density is varied, showing the transition from large-scale power *suppression* for closed universes to *amplification* for open universes. The ability of positive curvature to suppress large-scale power is particularly interesting in light of the lack of power on large scales found in full-sky CMB data (see [\[182\]](#) for an early discussion on this). The horizontal shift of the CMB power spectrum with $\Omega_{K,0}$ is not to be attributed to the PPS but to a curvature effect already present in the Λ CDM model (an effect on the transfer function, not the PPS).

The cutoff and oscillations in the PPS towards large scales (small k) are features of a kinetically dominated or fast-roll stage prior to inflation, already known and studied for flat universes [[II & chapter 4, 130](#)]. However, in flat universes these features can easily be pushed outside the observable window by large amounts of inflation, which is no longer the case with non-zero spatial curvature.

5.13 Nested sampling results

In this section we present the results from our Bayesian analysis using nested sampling. We start by investigating one- and two-parameter extensions to the base Λ CDM model in [section 5.13.1](#). In [section 5.13.2](#) we then change the phenomenological description of the primordial power spectrum from the Λ CDM model to that of specific inflationary models with a full numerical integration of the mode [eqs. \(5.47\)](#) and [\(5.54\)](#).

5.13.1 Nested sampling results: Λ CDM extensions

Since this paper focuses on cosmic inflation in curved universes, one obvious extension of the base Λ CDM model to investigate is an extension with the present-day curvature density parameter $\Omega_{K,0}$. Cosmic inflation governs the primordial Universe. Therefore we additionally look at extensions with primordial parameters. Possible parameter extensions to the base Λ CDM model include the running of the spectral index n_{run} and the tensor-to-scalar ratio r . In the following sections we focus on the tensor-to-scalar ratio r , which is more strongly constrained by current datasets than the running n_{run} . Hence, in what follows we present the results of a Bayesian analysis of the Λ CDM model extended by r and $\Omega_{K,0}$, both independently and jointly.

Posteriors of Λ CDM extensions

In [fig. 5.16](#) we show the one-dimensional and the pairwise joint two-dimensional posterior distributions for the present-day curvature density parameter $\Omega_{K,0}$, the spectral index n_s and the tensor-to-scalar ratio r . We present the results using both Planck 2018 $TT, TE, EE + \text{low}E$ data only, and using the data from the BICEP and Keck Array in addition.

Because of the importance of the degree of compression from prior to posterior distribution for model comparison, we also include the prior distributions in grey in [fig. 5.16](#), which are the same for all models. For visualisation purposes we illustrate the two-dimensional

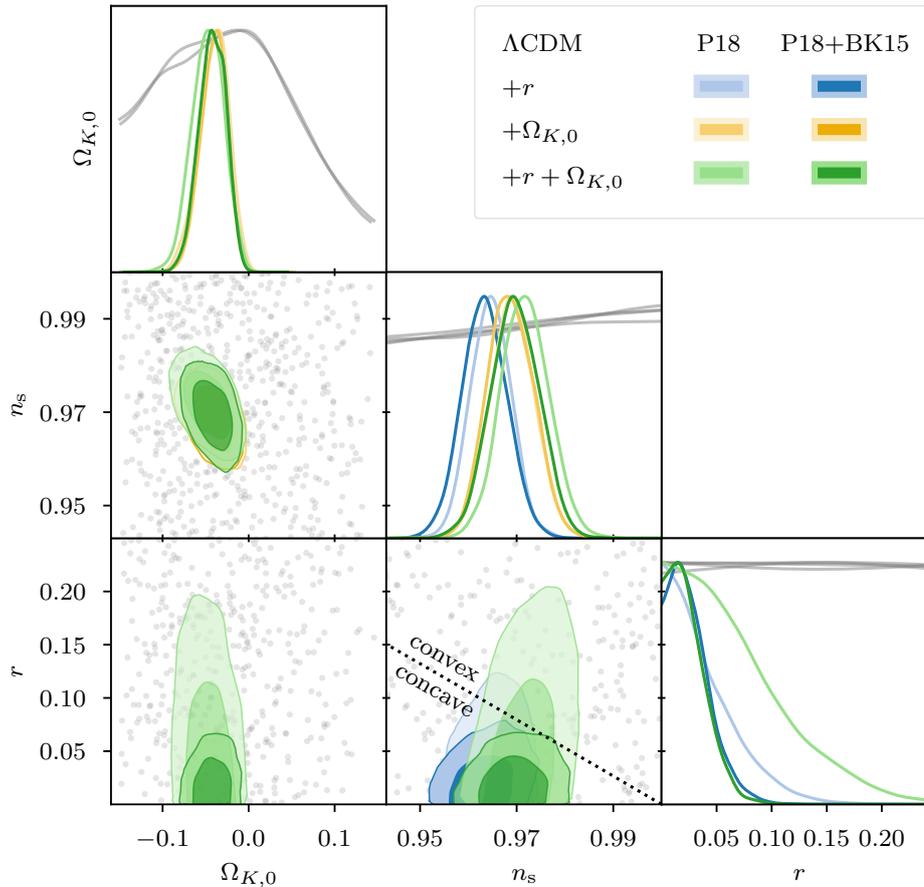


Figure 5.16 Posterior constraints on the Λ CDM extension parameters, the tensor-to-scalar ratio r and the curvature density parameter $\Omega_{K,0}$. We also include the spectral index n_s , since the n_s - r plot is typically the main plot of interest when investigating inflation models. The dotted line splits the n_s - r plot into the regions of convex or concave inflationary potentials. The lighter hue corresponds to using Planck 2018 $TT, TE, EE + \text{low}E$ data only [62], whereas the darker hue corresponds to additionally using data from the BICEP and Keck Array [71]. The grey lines and dots illustrate the prior distributions.

prior distribution in form of scatter points, as contours are better suited for constrained distributions. Note that while in principle all three parameters are sampled uniformly across their prior range, some parameter combinations need to be excluded at the prior level in order to compute a viable cosmological model, e.g. parameter combinations with large dark energy density Ω_Λ and small matter density Ω_m leading to universes that had no Big Bang in the first place. This leads to effectively non-uniform priors, the non-uniformity being somewhat visible for the spectral index n_s and very clear for the curvature density parameter $\Omega_{K,0}$ with a clear prior preference of flat universes.

CMB results for the one-parameter extensions have been investigated thoroughly in previous analyses [33, 71], giving a mostly closed universe for the $\Omega_{K,0}$ extension and an upper bound on the tensor-to-scalar ratio for the r -extension. Joint analyses of r and $\Omega_{K,0}$ have been briefly discussed in [131, 148, 173]. While the curvature parameter is little affected by the tensor-to-scalar ratio, the inverse is not true. The uncertainty on the tensor-to-scalar ratio

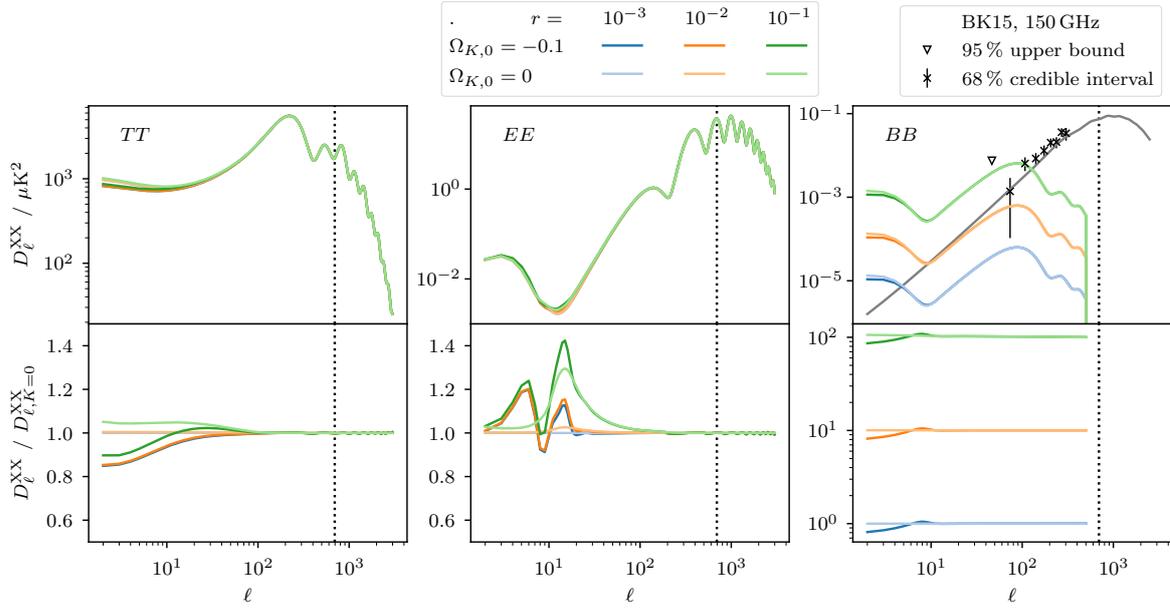


Figure 5.17 CMB power spectra of the temperature (TT) and polarisation (EE in the middle and BB on the right) anisotropies for different parameter values of the curvature density parameter $\Omega_{K,0}$ and the tensor-to-scalar ratio r . The lighter hue corresponds to a flat universe, whereas the darker hue assumes a closed universe with $\Omega_{K,0} = -0.1$. The bottom plots show the relative difference, where we use the spectrum with $\Omega_{K,0} = 0$ and $r = 10^{-3}$ as reference in the denominator. Note that we are using the *unlensed* spectra for visualisation of the effects of the tensor-to-scalar ratio, here. The black line in the B -mode power spectrum on the right is the contribution of *lensed* E -modes to the B -mode spectrum for $\Omega_{K,0} = 0$, $r = 0$.

increases considerably when allowing non-zero curvature with Planck data. This difference vanishes when BK15 data is taken into consideration, though, giving essentially the same distribution on r as without curvature. However, the shift in the spectral index n_s from curvature is retained when allowing for a non-zero tensor-to-scalar ratio. This is important for inflation models and reheating bounds, as will be explored in the later [section 5.13.2](#).

Note that according to current data from Baryon Acoustic Oscillations (BAOs) these constraints would be pulled towards a flat universe, e.g. combining CMB, CMB lensing and BAO measurements leads to $\Omega_{K,0} = 0.0007 \pm 0.0037$, which would also undo the shift in the spectral index. However, there has been concern over the combination of BAO with CMB data for curved universes [27–29], which is why in the present study we restrict ourselves to CMB data only, and leave a more involved analysis including BAOs for future work. The same applies (albeit to a lesser extent) to CMB lensing.

At first glance it might be surprising that the posterior distribution of the tensor-to-scalar ratio changes so significantly upon including non-zero curvature when computed from Planck data only, but remains essentially unchanged when including BK15 data. This phenomenon may be explained by the BK15 data offering an additional observable, the B -mode polarisation, which is much more sensitive to changes in the tensor-to-scalar ratio than the temperature data or the E -modes. This can be seen in [fig. 5.17](#) which shows the CMB temperature (TT) and polarisation (EE and BB) power spectra respectively for combinations

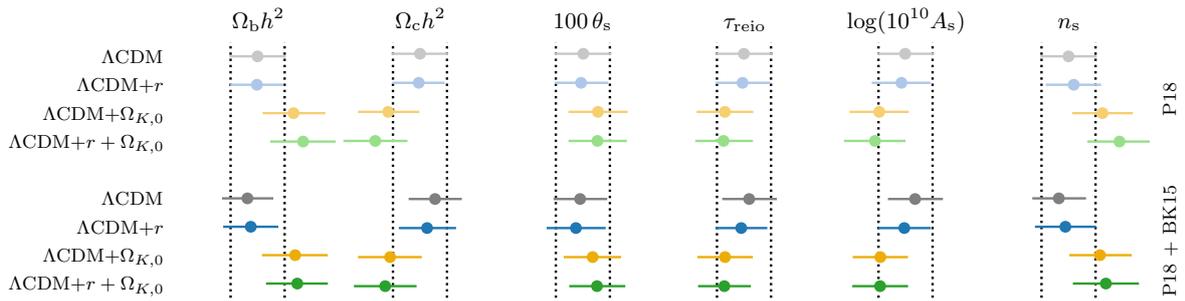


Figure 5.18 Parameter (in-)stability for extensions to the Λ CDM base model using Planck 2018 $TT, TE, EE + lowE$ [62] data only (top four, light hue) and data from both Planck and from the BICEP and Keck Array [71] (bottom four, dark hue). The dot represents the parameter mean and the error bars correspond to one standard deviation (1σ). We show the results for extensions with the tensor-to-scalar ratio r and/or the curvature density parameter $\Omega_{K,0}$. The vertical dotted lines serve as visual references and are the 1σ -boundaries for the Λ CDM model from Planck data only (top line).

of $r = \{10^{-3}, 10^{-2}, 10^{-1}\}$ with and without curvature. While in the case of TT and EE spectra different values of the tensor-to-scalar ratio are negligible in comparison to the effects of curvature, this behaviour is reversed in case of the BB spectra, where r shows a significantly stronger influence. Hence, including the B -mode data from BK15 results in essentially the same upper bound on the tensor-to-scalar ratio regardless of whether the universe is assumed to be curved or flat.

Nevertheless, it should be noted that the addition of BK15 data pushes the posterior constraints slightly further away from flatness.

Figure 5.18 summarises how the six cosmological base parameters change across the various extensions as regards their mean and standard deviation. Differences owing to the addition of BK15 data are negligible, with the parameter constraints all lying well within one standard deviation of one another across all models. Hence, both P18 and combined P18+BK15 exhibit the same trends when comparing different models with one another. Extension with the tensor-to-scalar ratio leaves the cosmological base parameters essentially invariant. Adding curvature on the other hand shifts all these parameters by roughly one standard deviation. The biggest shift is in the spectral index n_s .

Model comparison of Λ CDM extensions

For the comparison of the Λ CDM extensions we investigate the log-evidence $\ln \mathcal{Z}$, Kullback-Leibler divergence \mathcal{D}_{KL} , Bayesian model dimensionality d , and the posterior average of the log-likelihood $\langle \ln \mathcal{L} \rangle_{\mathcal{P}} = \ln \mathcal{Z} + \mathcal{D}_{KL}$. Figure 5.19 shows these quantities for Planck 2018 $TT, TE, EE + lowE$ data in the upper triangle plot and additionally using data from the BICEP and Keck Array in the lower triangle plot. Note that since these different datasets result in fundamentally different likelihood values, their absolute evidence values are not directly comparable and are therefore plotted separately. The relative differences of the various models, on the other hand, are comparable.

Including the tensor-to-scalar ratio (blue) is disfavoured with a log-evidence of $\ln \mathcal{Z} = -3.0 \pm 0.3$ which translates to betting odds of about 1 : 20 against the r -extension. This is

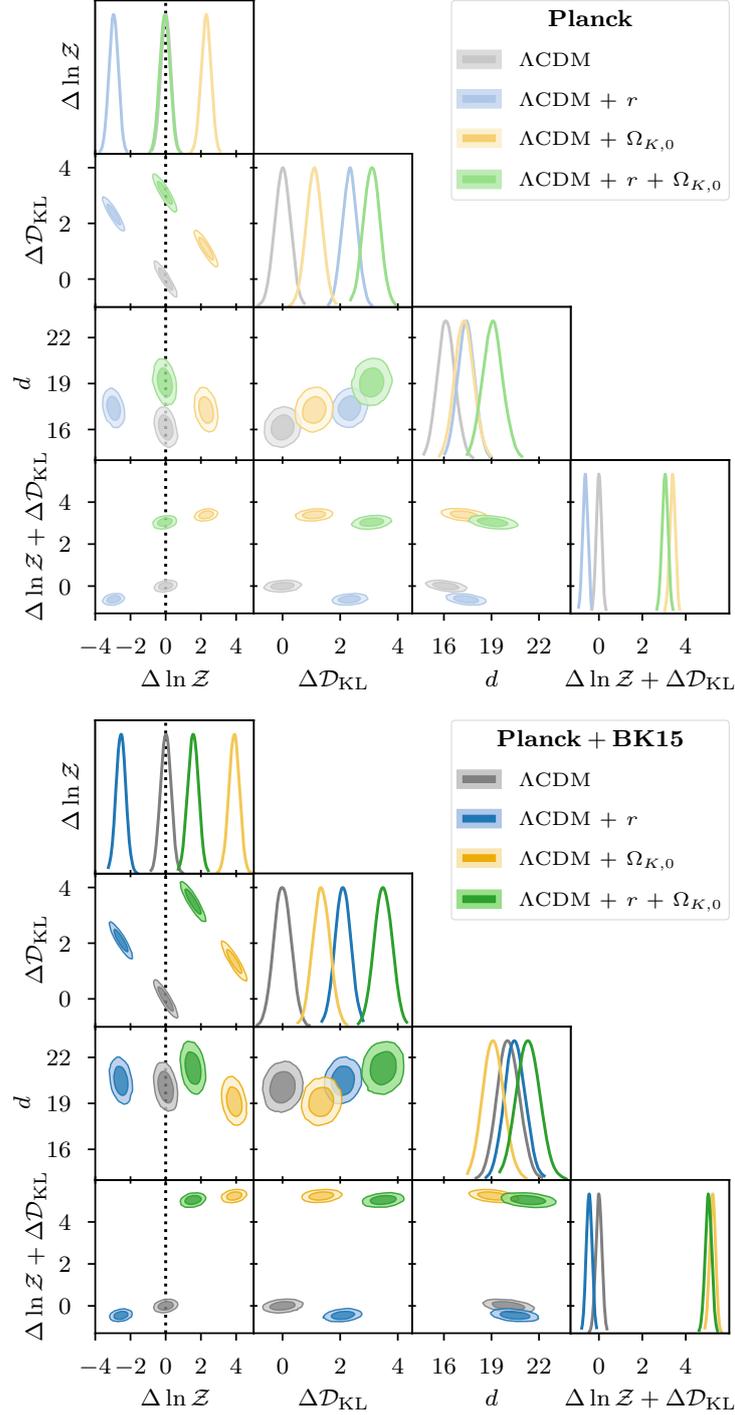


Figure 5.19 Bayesian model comparison for extensions to the base Λ CDM model with the tensor-to-scalar ratio r and/or the curvature density parameter $\Omega_{K,0}$. In the upper plot in a light hue we show the results using Planck 2018 $TT, TE, EE + \text{low}E$ data only, whereas in the lower plot in a darker hue we show the results when additionally including data from the BICEP and Keck Array. We show the log-evidence $\ln \mathcal{Z}$, Kullback-Leibler divergence \mathcal{D}_{KL} (in nats), Bayesian model dimensionality d , and the posterior average of the log-likelihood $\langle \ln \mathcal{L} \rangle_{\mathcal{P}} = \ln \mathcal{Z} + \mathcal{D}_{\text{KL}}$. The Δ denotes normalisation with respect to the base Λ CDM model without extensions (i.e. with $r = 0$ and $\Omega_{K,0} = 0$) for easier comparison. The probability distributions represent errors arising from the nested sampling process. In the limit of infinite life points these distributions would become point statistics, in contrast to posterior distributions.

entirely driven by the Occam penalty for the additional parameter and because of the lack of any clear B -mode signal. Sampling the tensor-to-scalar ratio logarithmically would leave $\log r$ mostly unconstrained and therefore the Bayesian evidence essentially invariant compared to the base Λ CDM model. Consequently the KL-divergence, which is effectively a measure of the Occam penalty, would be much smaller, too, such that switching between uniform and logarithmic prior corresponds roughly to moving contours along the $\ln \mathcal{Z} - \mathcal{D}_{\text{KL}}$ degeneracy line (see [chapter 2](#)).

Including the present-day curvature density parameter $\Omega_{K,0}$ (yellow), on the other hand, is favoured with a log-evidence of $\ln \mathcal{Z} = 2.3 \pm 0.3$ compared to the base Λ CDM model, which translates to betting odds of about 10 : 1 in favour of the curvature extension. Note that this is smaller compared to findings in [\[27, 28\]](#) which give odds of 50 : 1 and a log-evidence of about 3.3 respectively. However, in those cases the upper bound on the flat prior on $\Omega_{K,0}$ was chosen to be 0.05 or 0, whereas we have chosen our prior symmetrically around zero as $[-0.15, 0.15]$. We have tested that increasing the prior range beyond this barely affects the evidence. This is because of theoretical constraints provided by the Boltzmann solver CLASS that are put on the nominally uniform prior distribution. The resulting effective prior on $\Omega_{K,0}$ is shown in the top panel of [fig. 5.16](#). The prior distribution quickly drops to zero beyond these bounds, such that those regions do not contribute to the Bayesian evidence.

In a two-parameter extension with tensor modes *and* spatial curvature, their individual one-parameter effects cancel and the log-evidence is essentially equal to that of the base Λ CDM model. This only holds for the evidence, however, with the KL-divergence effectively adding up to represent the large Occam penalty from two additional parameters.

It is worth noting how the tensor-to-scalar ratio and the curvature density have an almost orthogonal effect on the Λ CDM model in the $\ln \mathcal{Z}, \mathcal{D}_{\text{KL}}$ plane. While the tensor-to-scalar ratio shifts the contour along the $\ln \mathcal{Z} + \mathcal{D}_{\text{KL}}$ line, the curvature density shifts along the $\ln \mathcal{Z} - \mathcal{D}_{\text{KL}}$ line. The former corresponds to a shift caused mostly by an Occam factor (quantified by $\Delta \mathcal{D}_{\text{KL}}$). The latter corresponds to a shift mostly driven by a better fit, which can be quantified by the posterior average of the log-likelihood, related to Bayesian evidence and KL-divergence as: $\ln \mathcal{Z} + \mathcal{D}_{\text{KL}} = \langle \ln \mathcal{L} \rangle_{\mathcal{P}}$ (cf. [eq. \(5.5\)](#)). Hence, this can be seen directly in [fig. 5.19](#) by looking at the joint contours of $\Delta \mathcal{D}_{\text{KL}}$ and $\langle \ln \mathcal{L} \rangle_{\mathcal{P}}$.

The third parameter in [fig. 5.19](#), the Bayesian model dimensionality, typically comes with a large error and is therefore not very specific in its count of the number of constrained parameters. For Planck data only, it does show the expected ordering, with the dimensionality increasing with additional parameters. For combined Planck and BK15 data, the dimensionality of the Λ CDM base model is not the smallest, however, with a standard deviation of about 0.7 on the dimensionality, it is not too concerning that a parameter difference of 1 is not properly resolved. Especially, when one takes into account that the full number of *sampling* parameters consists of 6 cosmology parameters (plus one or two for extensions), 21 Planck nuisance parameters and another 7 nuisance parameters for BK15.

Comparing upper (only Planck) and lower (Planck+BK15) plot in [fig. 5.19](#), we note firstly that the addition of BK15 data only increases the log-evidence of the r -extension (blue) slightly

to $\ln \mathcal{Z} = -2.5 \pm 0.3$, but is still disfavoured with odds of about 1 : 12 compared to the Λ CDM base model (grey), telling us that the BK15 data does not require a non-zero tensor-to-scalar ratio for a sufficiently good fit. Secondly, the model comparison of Λ CDM + $\Omega_{K,0}$ (yellow) with Λ CDM + r + $\Omega_{K,0}$ (green) (i.e. again adding tensor modes) also stays mostly unaffected by the addition of BK15 data. Thirdly, comparisons between models with and without curvature do change. Regardless of the tensor-to-scalar ratio, adding BK15 data increases the evidence for models including curvature. Looking at the posterior average of the log-likelihood shows that this is linked to a better fit of curved models. This means that the BK15 data has a greater effect on model comparisons involving the curvature parameter $\Omega_{K,0}$ than involving the tensor-to-scalar ratio r , which is curious in light of our previous discussion in [section 5.13.1](#) of the changes to the posterior. There the role was inverted, with the BK15 data having a greater effect on the posterior of the tensor-to-scalar ratio r than on the curvature density parameter $\Omega_{K,0}$. On the other hand, the difference in posteriors is better reflected in the KL-divergence, where the difference between models with (green) and without (blue) curvature increases upon addition of BK15 data, reflecting the greater compression in [fig. 5.16](#) taking place for the curvature case when going from Planck only to Planck+BK15. In fact, the KL-divergence for the flat r -extension stays the same between Planck and Planck+BK15.

5.13.2 Nested sampling results: Inflation models

In the following we present the results from our nested sampling runs with fully numerically integrated primordial power spectra for the inflation models considered in [section 5.11](#). For results that are mostly independent of the choice of potential we only show plots from the Starobinsky potential, representative of all potentials. Similarly, we only show results from the combined data of Planck and BK15 for some parameters, when there are no clearly visible differences with or without the BK15 data.

In addition to the prior bounds specified in [table 5.1](#) and to standard constraints from Λ CDM cosmology, we also enforce the curvature constraint for open universes from [eq. \(5.31\)](#), the horizon constraint from [eq. \(5.37\)](#) and the reheating constraints from [eqs. \(5.40\)](#) and [\(5.41\)](#) at the prior level.

Posteriors of primordial sampling parameters

In [section 5.12](#) we introduced the amplitude of scalar density perturbations A_s , the inflaton field value at the start of inflation ϕ_i , the fraction of primordial curvature $f_i \equiv \Omega_{K,i}/\Omega_{K,0}$ and the present day curvature density parameter $\Omega_{K,0}$ as our primordial sampling parameters. In [fig. 5.20](#) we show their prior (in grey) and posterior (in red) constraints for the Starobinsky model from Planck and BK15 data.

The picture is very similar across all potentials considered here. The only notable difference between inflation models is in the inflaton field parameter owing to the different scaling of e-folds of the scale factor during inflation. The Starobinsky model with its exponential scaling is restricted to much lower values of ϕ_i than most other models, whose scaling is closer to

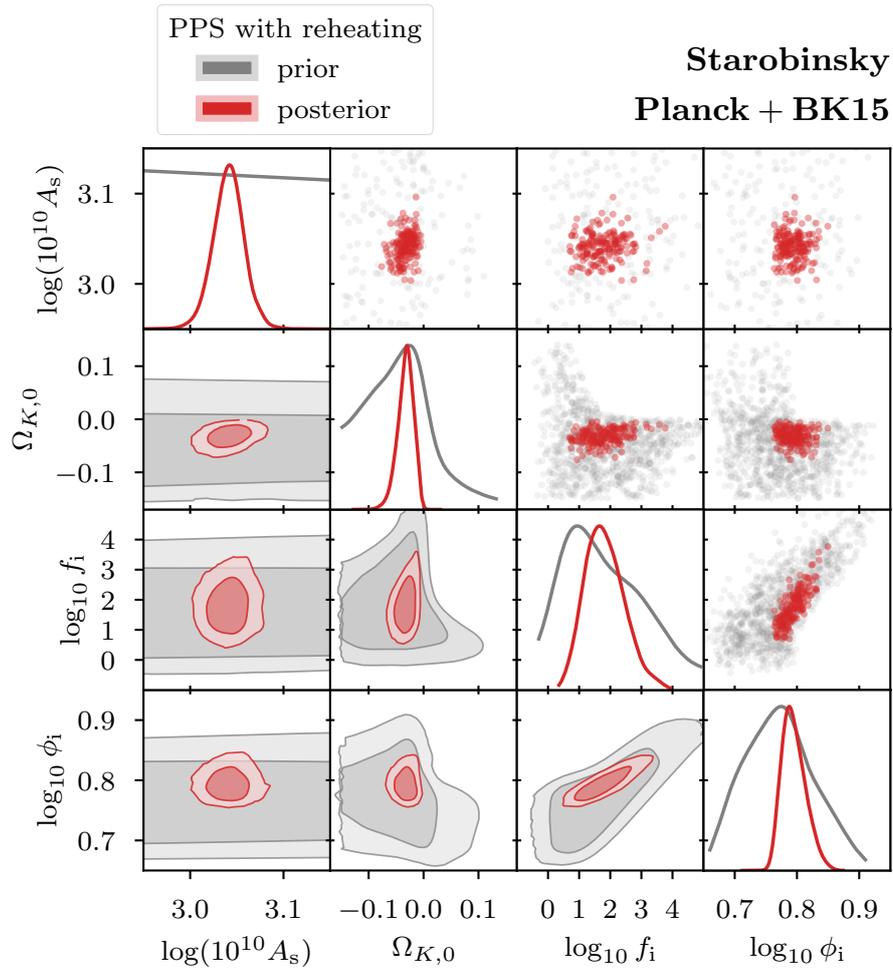


Figure 5.20 Prior (in grey) and posterior (in red) distributions of the parameters used to compute the primordial power spectrum with the Starobinsky potential. Note that all the priors shown here are initially set as uniform priors, but deviate from uniformity owing to additional constraints from curvature, reheating and horizon considerations.

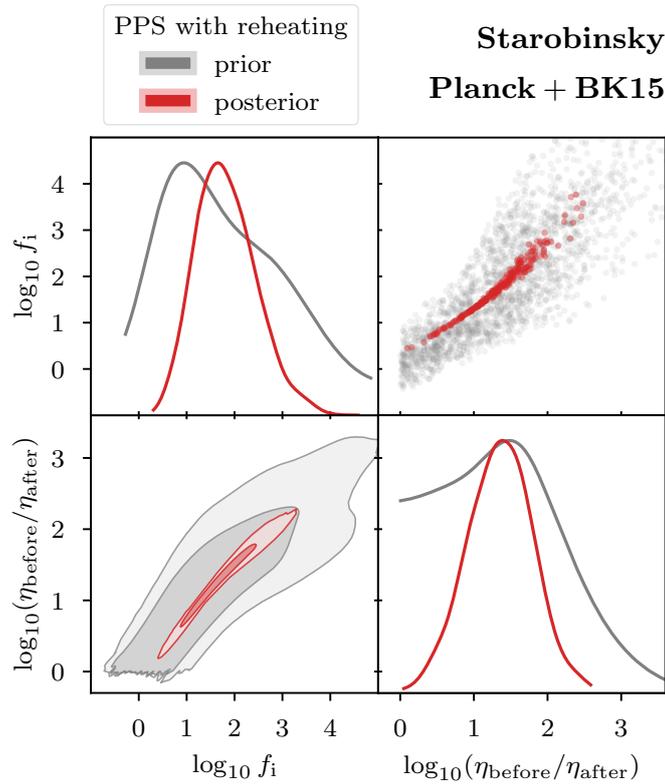


Figure 5.21 Prior (in grey) and posterior (in red) distributions of the primordial curvature fraction f_i and the conformal time ratio $\eta_{\text{before}}/\eta_{\text{after}}$.

quadratic. This is primarily driven by the reheating constraints limiting the total number of e-folds of inflation. We investigate this in more detail in [section 5.13.2](#).

The amplitude of the scalar primordial power spectrum A_s is by far the best constrained of these parameters. This comes as no surprise, considering that it is also one of the six parameters in the base Λ CDM cosmology.

The prior distribution of the primordial curvature fraction f_i is the joint result of the curvature constraints and the horizon constraint from [eq. \(5.37\)](#) specifying that the conformal time that passed before the end of inflation needs to be greater than thereafter in order to solve the horizon problem.

The sharp drop in the prior towards low values of the primordial curvature fraction f_i is driven by the horizon constraint as expected from our analysis in [section 5.8](#). This is confirmed in [fig. 5.21](#), where we show the prior (in grey) and posterior (in red) distributions of the primordial curvature fraction f_i and conformal time ratio $\eta_{\text{before}}/\eta_{\text{after}}$. Indeed the correlation between f_i and the ratio of $\eta_{\text{before}}/\eta_{\text{after}}$ together with the cut of $\eta_{\text{before}} > \eta_{\text{after}}$ excludes low values of f_i . On the posterior level, the correlation reduces almost to a one-to-one correspondence. The data actually prefer a conformal time ratio of about 10 to 100 and reject smaller values such that the horizon problem is indeed solved.

Towards larger values of f_i the prior is reduced owing to curvature constraints. First, the curvature constraint for open universes from [eq. \(5.31\)](#) only allows comparatively small

values of f_i . The prior is reduced further for too large curvature in closed universes, as these universes lack a Big Bang in the first place. It is useful in this context to look at the joint distribution with the present-day curvature and also at the corresponding plots from the conformal time analysis in [fig. 5.8](#). This also indicates that a total elapse of conformal time of $\eta_{\text{total}} = \pi/2$ from pre-inflationary Big Bang to the future conformal boundary, as proposed in the closed universe theory discussed in [\[123\]](#), is consistent with the data.

The correlation between f_i and the initial value for the inflaton field ϕ_i is needed to prevent closed universes from collapsing, and is a direct consequence of the first Friedmann [eq. \(5.14\)](#).

Posteriors of derived parameters

Computing the primordial power spectrum from the inflationary background [eqs. \(5.14\)](#) to [\(5.17\)](#) and the mode [eqs. \(5.47\)](#) and [\(5.54\)](#) of scalar and tensor perturbations turns the phenomenological spectral index n_s and tensor-to-scalar ratio r into derived parameters. These parameters mostly depend on the observable number of e-folds of inflation N_* from horizon crossing of the pivot scale k_* until the end of inflation. In a curved universe this dependence is equivalent to a dependence on the total number of e-folds of inflation N_{tot} . As discussed in [section 5.9](#), the equation-of-state parameter of reheating w_{reh} is also mostly driven by the amount of inflation.

[Figure 5.22](#) shows all these parameters for various potential models in two triangle plots, the first using Planck data only, and the other combining Planck with BK15 data. The prior distributions are shown in a lighter hue as histograms for the one-dimensional distributions and as scatter plots for the pairwise joint distributions. Posterior distributions are plotted with a darker hue. For n_s and r we also show the posterior for the Λ CDM extension with curvature $\Omega_{K,0}$ and with r (sampled logarithmically). We show the tensor-to-scalar ratio on a logarithmic scale to better visualise the large difference in the inflation models' predictions. We show results from the quadratic potential in blue, natural potential in orange, quartic double-well potential in green, and Starobinsky potential in red.

The two-dimensional distributions (both prior and posterior) show the degeneracy lines between all these parameters for the various inflation models clearly. Only the models with an extra inflationary parameter, such as natural and double-well inflation, show a slightly greater spread, which would be more apparent on a linear scale in r .

The reheating parameter $w_{\text{reh,BBN}}$ was allowed to vary from $-1/3$ to 1 , here, thereby placing a theoretical upper limit on the spectral index and lower limit on the tensor-to-scalar ratio. This is particularly apparent for the quadratic and natural potential and when including BK15 data. The data prefer a lower tensor-to-scalar ratio, but the reheating prior limits how far down the posterior contours can be pushed. Note that this is a very permissive reheating prior. In [section 5.13.2](#) we show the effects of different reheating priors in more detail.

The results for the quadratic and natural potential are very similar. The preferred number of observable e-folds is roughly $N_* \approx 60$, even larger with BK15 data included. This is somewhat larger than the more commonly quoted 50 to 60 e-folds owing to the larger spectral index when including curvature. The effective reheating parameter $w_{\text{reh,BBN}}$ is centred

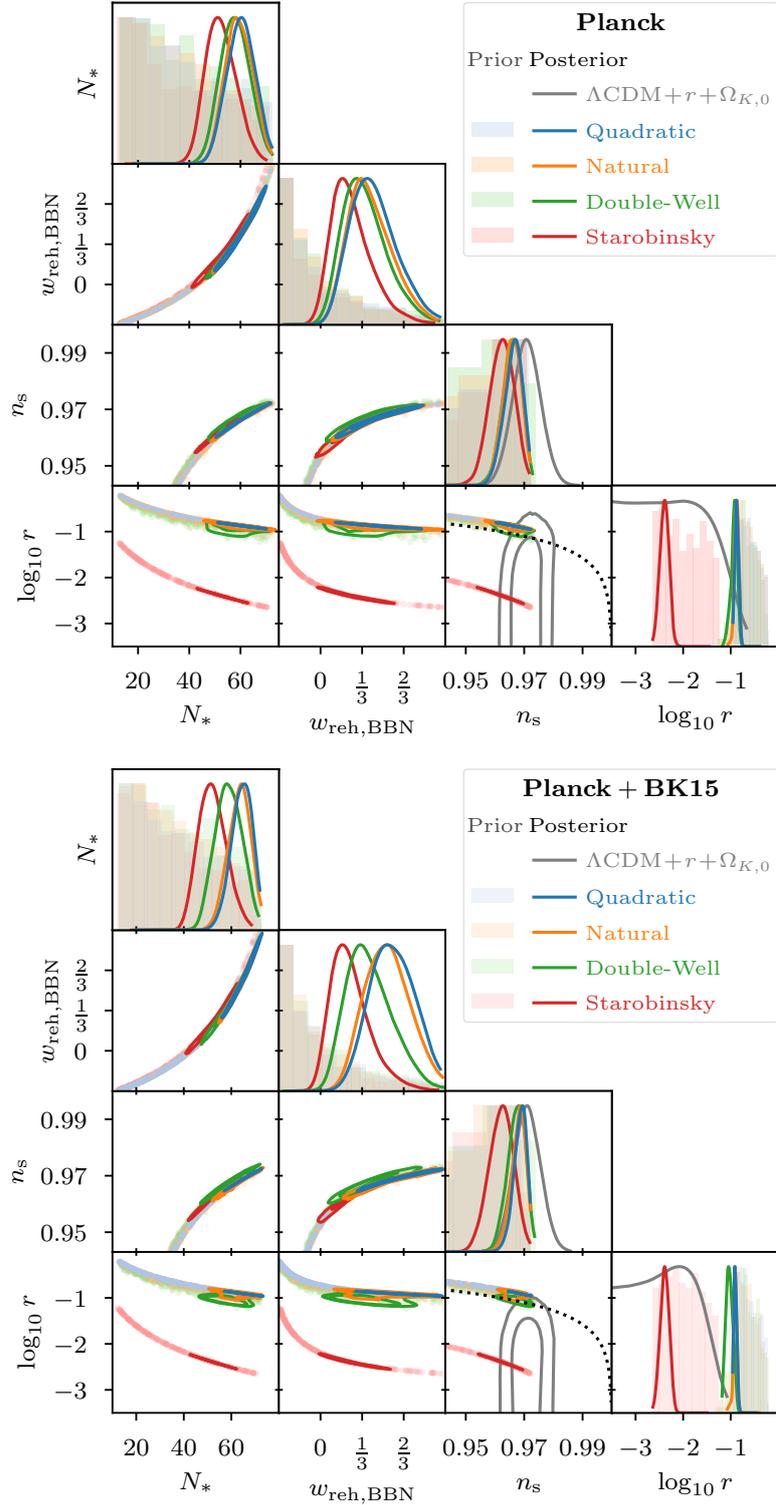


Figure 5.22 Posterior distributions of the observable number of e-folds of inflation N_* , the equation-of-state parameter of reheating until BBN $w_{\text{reh, BBN}}$, the spectral index n_s and the tensor-to-scalar ratio r for various inflation models. We show quadratic inflation in blue, natural inflation in orange and Starobinsky inflation in red. In green we show the contours for a power-law primordial power spectrum. The one-dimensional histogram plots and the two-dimensional scatter plots in a lighter hue illustrate the prior distributions of the corresponding parameters, which are non-uniform here since all parameters are derived parameters. The visible cutoffs of both prior and posterior distribution are driven by the (permissive) reheating constraint $-1/3 < w_{\text{reh, BBN}} < 1$.

around $1/3$ for Planck data only. This agrees well with the effective nature of the parameter tending towards $1/3$ in cases where thermalisation would have happened much earlier than Big Bang Nucleosynthesis as previously discussed in [section 5.9](#). Both the spectral index and tensor-to-scalar ratio are prior constrained. The models maximise their likelihood by pushing n_s to its upper and r to its lower prior bound.

It might seem surprising that the results for the natural potential are so similar to those of the quadratic potential. One could have expected natural inflation's ability to accommodate for a smaller tensor-to-scalar ratio via a smaller potential hill parameter ϕ_0 to pull the posterior away from quadratic inflation, which indeed is slightly visible when including BK15 data. However, there is a trade-off between the spectral index and the tensor-to-scalar ratio in natural inflation. The data push simultaneously to smaller r and larger n_s , whereas natural inflation gives a smaller r only for a smaller n_s .

The posterior of the Starobinsky model shows a preference for fewer observable e-folds N_* than for the other models, despite the exponential dependence on the initial inflaton field value ϕ_i . This also yields a lower effective reheating parameter $w_{\text{reh,BBN}}$ and spectral index n_s and can be attributed to the Starobinsky model's generally lower tensor-to-scalar ratio r . Where the other models push to the limit set by the reheating prior ($w_{\text{reh,BBN}} < 1$) to try and accommodate as small a tensor-to-scalar ratio as possible, the posterior for the Starobinsky model is well within the unconstrained plateau region of the likelihood on $\log r$, which remains the case when including BK15 data. Hence, there is no pressure towards smaller r for the Starobinsky model and since the BK15 data essentially only improve the sensitivity on r , the posterior of the Starobinsky model remains unaffected.

Best-fit power spectra

In [fig. 5.23](#) we show the best-fit primordial power spectra (PPS) that enter the computation of the angular TT , TE , and EE auto- and cross-spectra, plotted on top of the corresponding Planck 2018 data in [figs. 5.24](#) to [5.26](#) respectively. We use the usual normalisation of the angular CMB power spectra according to:

$$D_\ell^{XX} \equiv \frac{\ell(\ell+1)}{2\pi} C_\ell^{XX}. \quad (5.70)$$

We compare three representative best-fit PPS. Black and yellow correspond to power-law spectra with the best-fit parameters from the base Λ CDM model and from its extension with the curvature density parameter Λ CDM+ $\Omega_{K,0}$ respectively. Red corresponds to the scalar and tensor spectra that were numerically integrated using the Starobinsky potential. Except for the amplitude of tensor modes, the power spectra from other inflationary potentials are very similar. Hence, we only show the best-fit result for the Starobinsky model.

Comparing the Λ CDM model and its curvature extension shows that the major difference arising from the addition of $\Omega_{K,0}$ is a slightly larger spectral index n_s as previously shown in [figs. 5.16](#) and [5.18](#). This results in a little less power on large scales, i.e. for small wavenumber k and multipole ℓ .

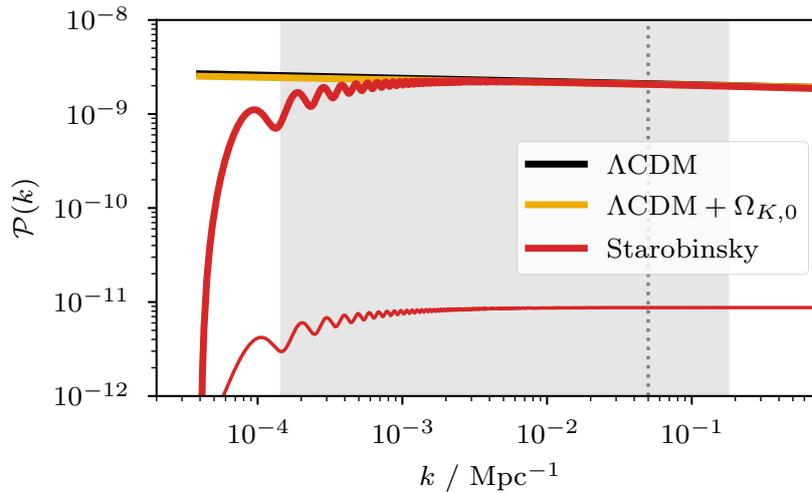


Figure 5.23 Best-fit primordial power spectra for the power-law spectra from Λ CDM in black and its extension with curvature $\Omega_{K,0}$ in yellow, and for the fully numerically integrated scalar (heavy upper line) and tensor (thin lower line) spectra from Starobinsky inflation in red. The vertical dotted line corresponds to the pivot scale $k_* = 0.05 \text{ Mpc}^{-1}$, where power amplitude A_s , spectral index n_s and tensor-to-scalar ratio r are measured. The grey shaded region illustrates roughly the CMB observable window. The corresponding CMB spectra are shown in [figs. 5.24 to 5.26](#).

The PPS for the Starobinsky model shows the typical cutoff and oscillations towards large scales (small k) that are associated with kinetic dominance initial conditions. With a best-fit curvature parameter of roughly $\Omega_{K,0} \approx -0.05$, there is a further overall suppression of power towards large scales from positive curvature (negative $\Omega_{K,0}$, cf. [fig. 5.15](#)).

While the oscillations visible in the PPS in [fig. 5.23](#) do not propagate through to the CMB spectra in [figs. 5.24 to 5.26](#), the additional suppression of power is visible in the TT , TE and EE spectra. This is added on top of the effect of a larger spectral index when including positive curvature, which also effectively decreases large-scale power, although to a much smaller degree than the curvature effect on the PPS.

Apart from these differences on large scales, all models agree on small scales, driven mainly by the high precision on the power amplitude A_s and by the good agreement between a power-law spectrum and the slow-roll predictions from inflation on small scales.

Model comparison of inflation models

As in [section 5.13.1](#), we investigate the log-evidence $\ln \mathcal{Z}$, Kullback-Leibler divergence \mathcal{D}_{KL} , Bayesian model dimensionality d , and the posterior average of the log-likelihood $\langle \ln \mathcal{L} \rangle_{\mathcal{P}} = \ln \mathcal{Z} + \mathcal{D}_{\text{KL}}$ for the various inflation models. We again show these quantities in [fig. 5.27](#) in two triangle plots, for Planck 2018 $TT, TE, EE + \text{low}E$ data in the upper plot and additionally using data from the BICEP and Keck Array in the lower plot.

We normalise with respect to the base Λ CDM model (vertical and horizontal grey lines marking zero). We also show the results for the Λ CDM+ r + $\Omega_{K,0}$ model with both a uniform and logarithmic prior on the tensor-to-scalar ratio r . Switching between these two priors roughly indicates the extent of the Occam penalty incurred by including r (see also [chapter 2](#)

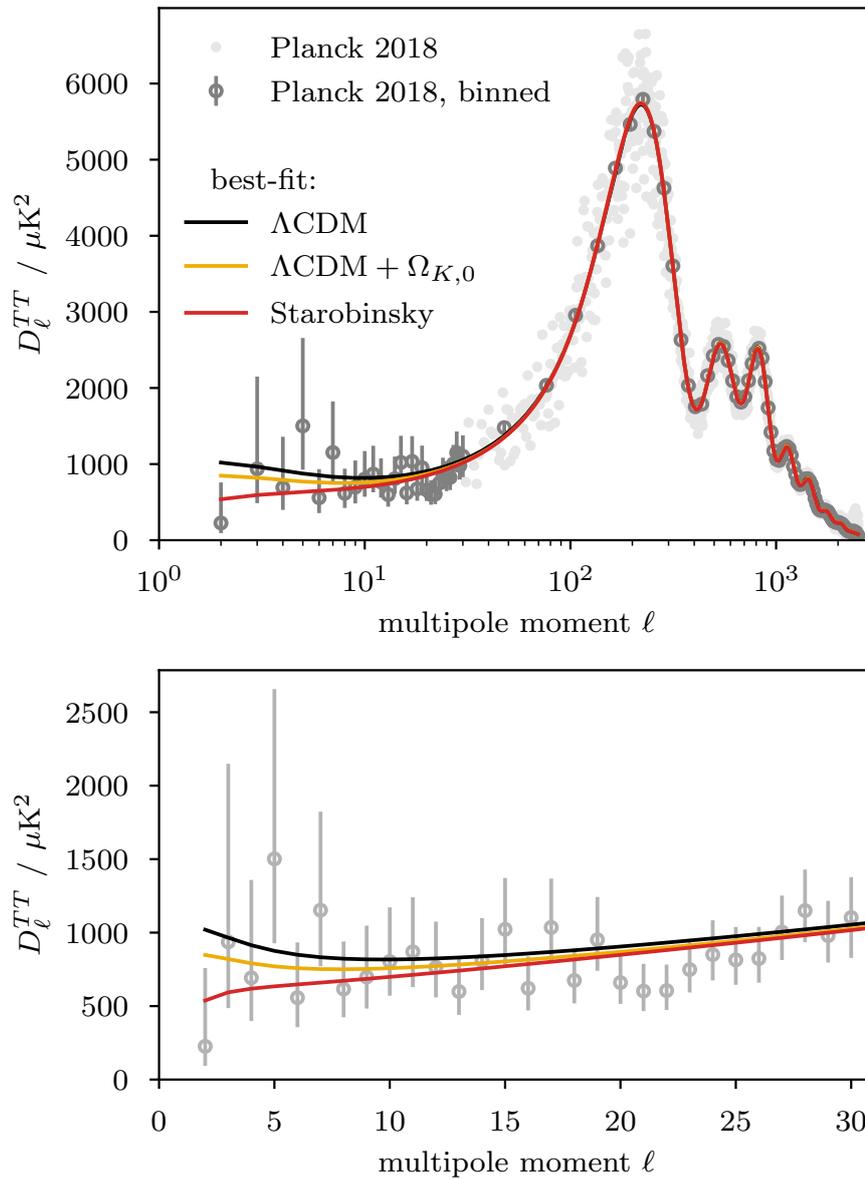


Figure 5.24 Comparison of the best-fit angular TT power spectra to the Planck 2018 TT data. The lower plot zooms in on the low- ℓ multipole range from the upper plot. The best-fit spectra were computed using the corresponding primordial power spectra from [fig. 5.23](#).

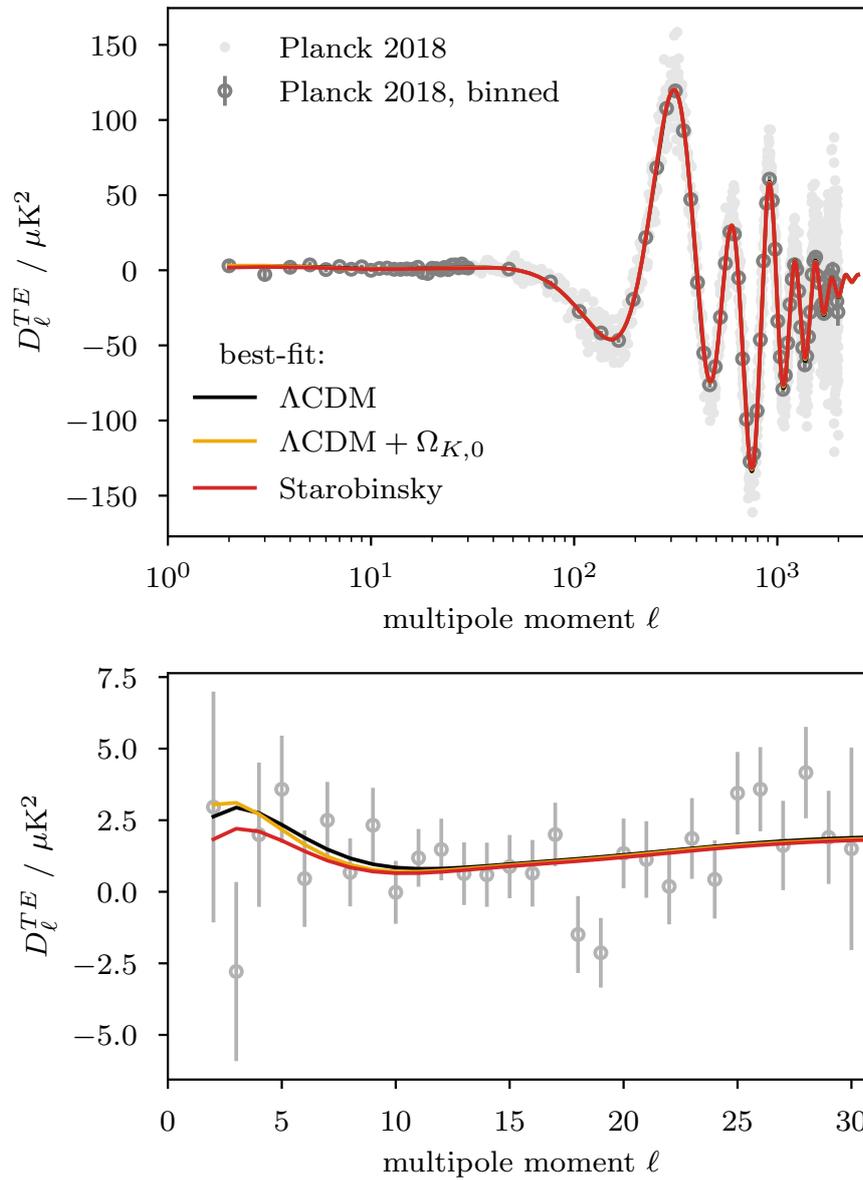


Figure 5.25 Comparison of the best-fit angular TE power spectra to the Planck 2018 TE data. The lower plot zooms in on the low- ℓ multipole range from the upper plot. The best-fit spectra were computed using the corresponding primordial power spectra from [fig. 5.23](#).

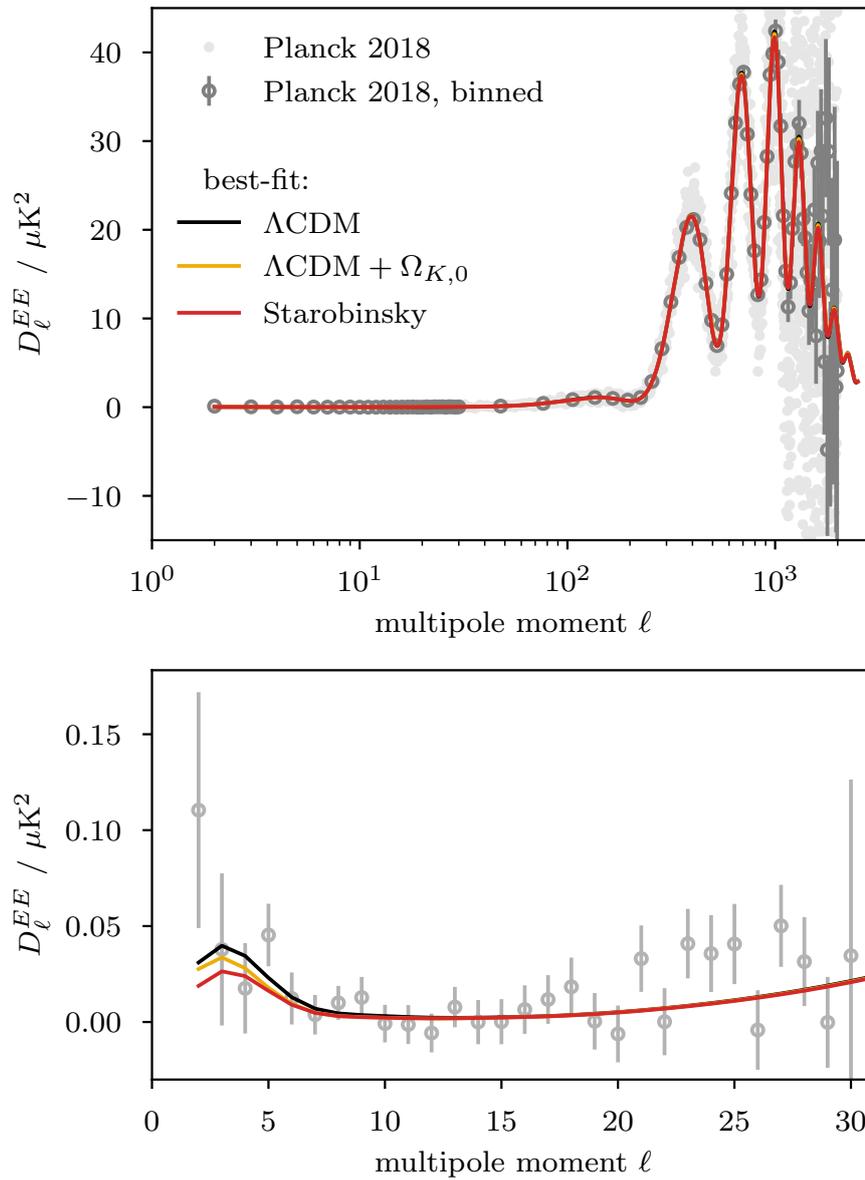


Figure 5.26 Comparison of the best-fit angular EE power spectra to the Planck 2018 EE data. The lower plot zooms in on the low- ℓ multipole range from the upper plot. The best-fit spectra were computed using the corresponding primordial power spectra from [fig. 5.23](#).

for a more complete discussion). It should be noted, though, that because of the different sampling parameters and their priors it is difficult to compare the very phenomenological description of the primordial Universe as expressed by power-law parameters A_s , n_s and r , to the much more specific generation of the PPS from inflation models. We have tried to mitigate this problem by using the same prior on the power amplitude A_s for the inflation models. However, this cannot be done for the spectral index and the tensor-to-scalar ratio. Instead, these become derived parameters, mostly dependent on the initial inflaton field value ϕ_i . The prior bounds on $\log_{10} \phi_i$ are naturally motivated by physical considerations. As already seen in [fig. 5.20](#), the priors on $\log_{10} \phi_i$ are limited by external constraints from reheating and from needing to solve the horizon problem. This is a feature of curved universes providing an absolute scale for the Universe and thereby a limit on the total amount of inflation N_{tot} (which is linked to ϕ_i), not just the observable amount of inflation N_* .

From the upper triangle plot in [fig. 5.27](#) we see that when only taking into account the data from Planck 2018, all inflation models considered perform similarly. This is in contrast to known results for flat universes [[II & chapter 4, 131, 148, 173, 208, 219](#)] and is due to the looser constraints on the tensor-to-scalar ratio when varying the present-day curvature density $\Omega_{K,0}$ simultaneously. Looking at the relative evidence $\Delta \ln \mathcal{Z}$, there is only a weak preference for inflation models over Λ CDM, and a comparison with Λ CDM+ $\log r + \Omega_{K,0}$ indicates that this might be only due to the CMB's preference for including non-zero positive curvature. Indeed, including the KL-divergence in our analysis shows that the inflation models provide a similar fit, $\langle \ln \mathcal{L} \rangle_{\mathcal{P}}$, and Occam penalty, $\Delta \mathcal{D}_{\text{KL}}$, to the tensor and curvature extension of Λ CDM. Taking into account its large uncertainty, the Bayesian model dimensionality is very similar for all inflation models and roughly $\Delta d \approx 1$ to 2 greater than for Λ CDM, reflecting the addition of two sampling parameters.

Including data from BICEP and the Keck array changes the results significantly, as seen in the lower triangle plot of [fig. 5.27](#). Where the contours of the various inflation models were bunching up at very similar evidence and KL-divergence values in the P18 only case, they spread out in the P18+BK15 case. Quadratic and natural inflation do not manage to provide a sufficiently small tensor amplitude (while preserving a sufficiently high scalar spectral index) under reheating constraints, and therefore are disfavoured with Bayesian odds of about 1 : 20 compared to Λ CDM, and roughly 1 : 2000 compared to the Starobinsky model. The quartic double-well potential fares better and is roughly on par with the Λ CDM model, which is to be attributed to the lower tensor-to-scalar ratio. The Starobinsky model on the other hand remains a strong competitor. The compression from prior to posterior measured by the KL-divergence has actually decreased, while the fit has increased, resulting in a higher log-evidence overall than even the Λ CDM+ $\log r + \Omega_{K,0}$ model. Although this is not a clear preference when taking the uncertainty on $\ln \mathcal{Z}$ into account, this result goes to show that the Starobinsky model manages to accommodate all the phenomenological requirements of the data for the PPS.

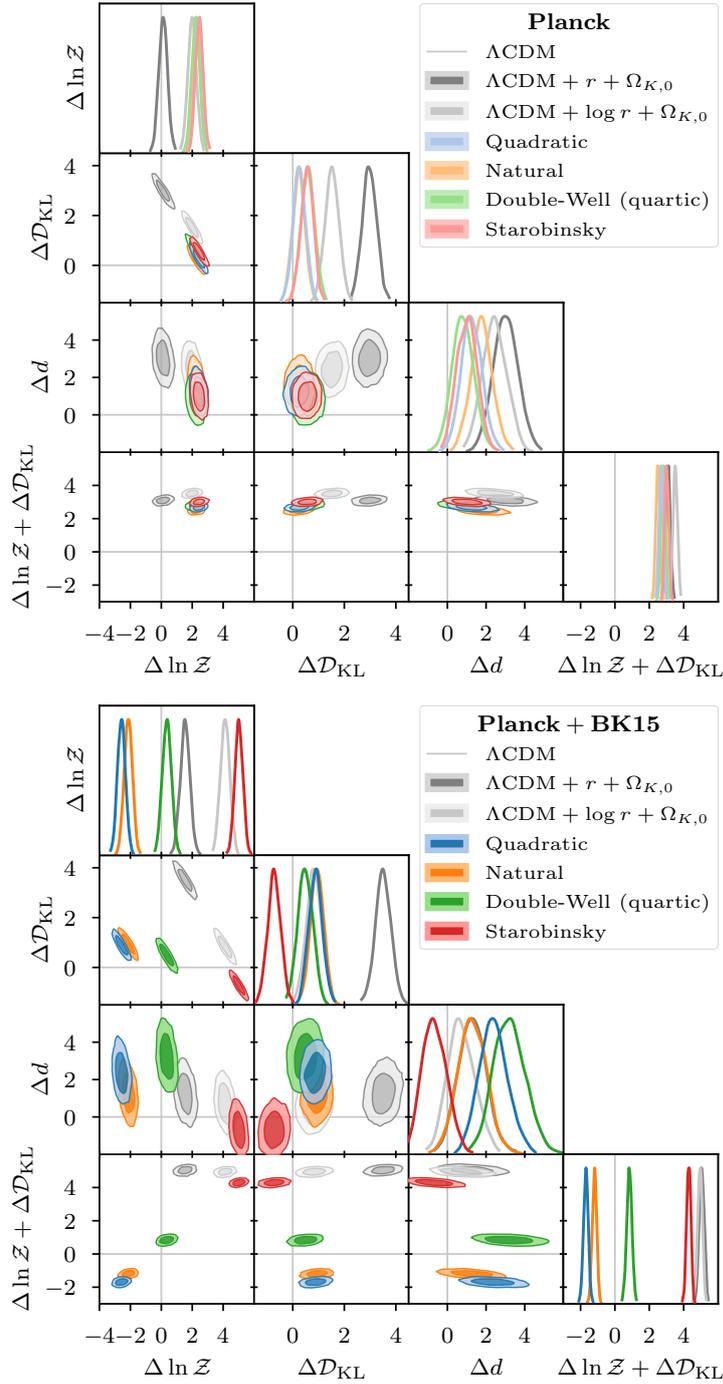


Figure 5.27 Bayesian model comparison for various inflation models. We show quadratic inflation in blue, natural inflation in orange, results from the quartic double-well potential in green and the Starobinsky model in red. In the upper plot in a light hue we show the results using Planck 2018 $TT, TE, EE + lowE$ data only, whereas in the lower plot in a darker hue we show the results when additionally including data from the BICEP and Keck Array. We show the log-evidence $\ln \mathcal{Z}$, Kullback-Leibler divergence \mathcal{D}_{KL} (in nats), Bayesian model dimensionality d , and the posterior average of the log-likelihood $\langle \ln \mathcal{L} \rangle_{\mathcal{P}} = \ln \mathcal{Z} + \mathcal{D}_{KL}$. The Δ denotes normalisation with respect to the base Λ CDM model without extensions (i.e. with $r = 0$ and $\Omega_{K,0} = 0$) indicated by the vertical and horizontal grey lines. As an additional reference we also include the results for an extension of Λ CDM with tensor modes and curvature in grey. The probability distributions represent errors arising from the nested sampling process. In the limit of infinite life points these distributions would become point statistics, in contrast to posterior distributions.

Effect of reheating constraints on evidences

In [section 5.9](#) we introduced the constraints from reheating on the end of inflation. In this section we will contrast the following two reheating scenarios:

$$\text{(permissive)} \quad N_{\text{reh}} = N_{\text{BBN}}, \quad -\frac{1}{3} < w_{\text{reh}} < 1, \quad (5.71)$$

$$\text{(restrictive)} \quad \rho_{\text{reh}}^{1/4} = 10^9 \text{ GeV} \quad -\frac{1}{3} < w_{\text{reh}} < \frac{1}{3}. \quad (5.72)$$

Similar categories can be found in [\[148\]](#).

[Figure 5.28](#) shows the posterior distributions of the derived equation-of-state parameter w_{reh} for the permissive scenario in the upper panel and the restrictive one in the lower panel. We again present the results in blue, orange, green and red for quadratic, natural, quartic double-well and Starobinsky inflation respectively. In grey we illustrate the underlying prior distribution, which is derived from the prior distributions listed in [table 5.1](#). Note how this favours small w_{reh} values a priori (driven by a degeneracy with ϕ_i and related to those seen in [fig. 5.16](#)), but is clearly overcome by the data. As already remarked in [section 5.13.2](#), the addition of BK15 data (from dashed to solid lines) has little effect on the Starobinsky model, but significantly increases the reheating parameter for quadratic and natural inflation. This is driven by the smaller tensor-to-scalar ratio required by the BK15 data, which is mostly independent of the reheating scenario used.

Comparing both reheating scenarios overall shows how the posterior for w_{reh} is diluted away from instant reheating at $w_{\text{reh}} = 1/3$ the shorter the duration of reheating, i.e. for an earlier (stricter) end to reheating at a higher energy density. Phrased the other way round, the longer reheating is allowed to last, the more the posterior on the *effective* equation-of-state parameter gets concentrated around $w_{\text{reh}} = 1/3$, which is equal to the equation-of-state parameter during the subsequent epoch of radiation domination (see [figs. 5.1](#) and [5.10](#) for a visual aid). At a first glance it might appear counter-intuitive that a permissive reheating scenario should result in tighter constraints on w_{reh} . However, the way to read this is that for the permissive reheating scenario essentially all posterior samples fall into the acceptable range of w_{reh} , which is not the case for the restrictive reheating scenario.

Working with P18 data only, quadratic and natural inflation are peaked at $w_{\text{reh}} = 1/3$, which corresponds to instant reheating. With BK15 data, the bulk of the posterior mass is at values greater than $1/3$, meaning that the comoving Hubble horizon needs to grow faster during reheating than during radiation domination to catch up with the standard Big Bang evolution. Note that this result is in stark contrast to the analytic prediction of matter dominated reheating, i.e. $w_{\text{reh}} \approx 0$, from the time averaged oscillations of the inflaton field around its potential minimum (see also [section 5.9](#)). The dilution of the posterior with strict reheating somewhat reconciles these models with matter dominated reheating, but this shows that any such oscillations can only last for a short time in case of quadratic or natural inflation.

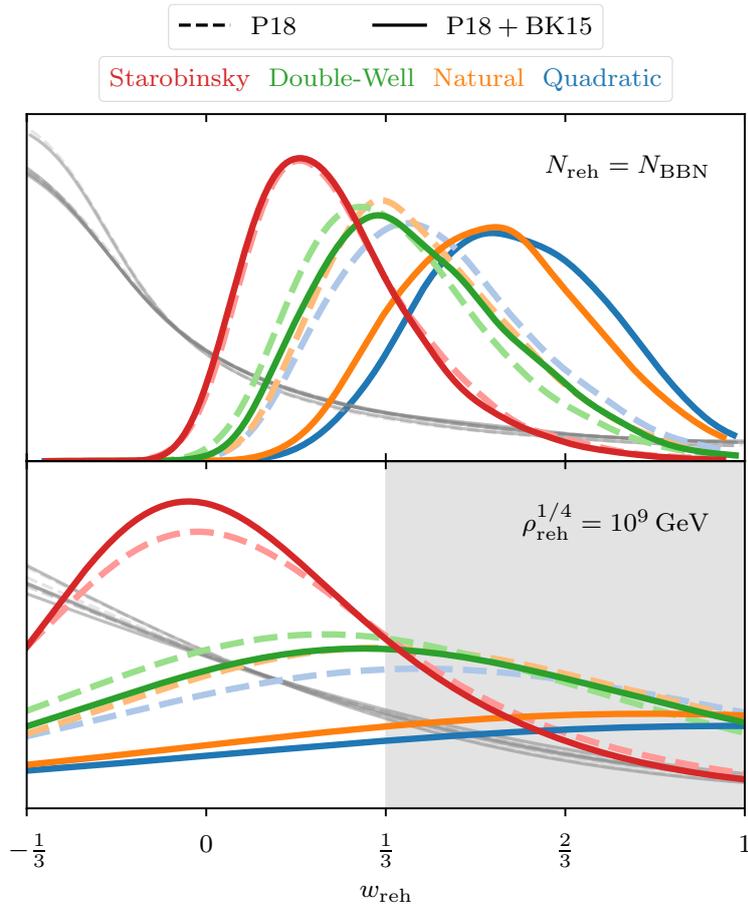


Figure 5.28 Posterior distributions for the derived equation-of-state parameter of reheating w_{reh} for the permissive (upper panel) and restrictive (lower panel) reheating scenario from eqs. (5.71) and (5.72) respectively. The grey shading in the lower panel highlights that under the restrictive reheating scenario only $-1/3 < w_{\text{reh}} < 1/3$ is allowed. We show quadratic inflation in blue, natural inflation in orange, results from the quartic double-well potential in green and the Starobinsky model in red. The grey lines illustrate the sampled prior distributions. Dashed lines correspond to P18 data only, and solid lines to P18 and BK15 data combined.

The Starobinsky model peaks right in-between 0 and $1/3$ in case of permissive reheating and roughly at 0 for restrictive reheating. Thus, for Starobinsky inflation, matter dominated oscillations around the potential minimum agree very well with the data, further adding to the success of the model, which it already accumulated on the level of the spectral index and the tensor-to-scalar ratio (although we recognise that these are all connected).

For the nested sampling runs presented in the previous sections and in figs. 5.22 and 5.27 in particular, we only used the permissive reheating scenario from eq. (5.71) as a prior constraint. In order to infer the evidence \mathcal{Z} and Kullback–Leibler divergence \mathcal{D}_{KL} with the restrictive scenario as prior constraint, we use `anesthetic`'s [109] importance sampling feature for nested samples. This frees us from the need to recompute entire nested sampling runs. However, as with importance sampling of MCMC chains, it only works well provided sufficient coverage of the importance sampled subspace of the original parameter space. Hence, the uncertainties tend to increase, which is especially the case for quadratic and natural inflation, for which

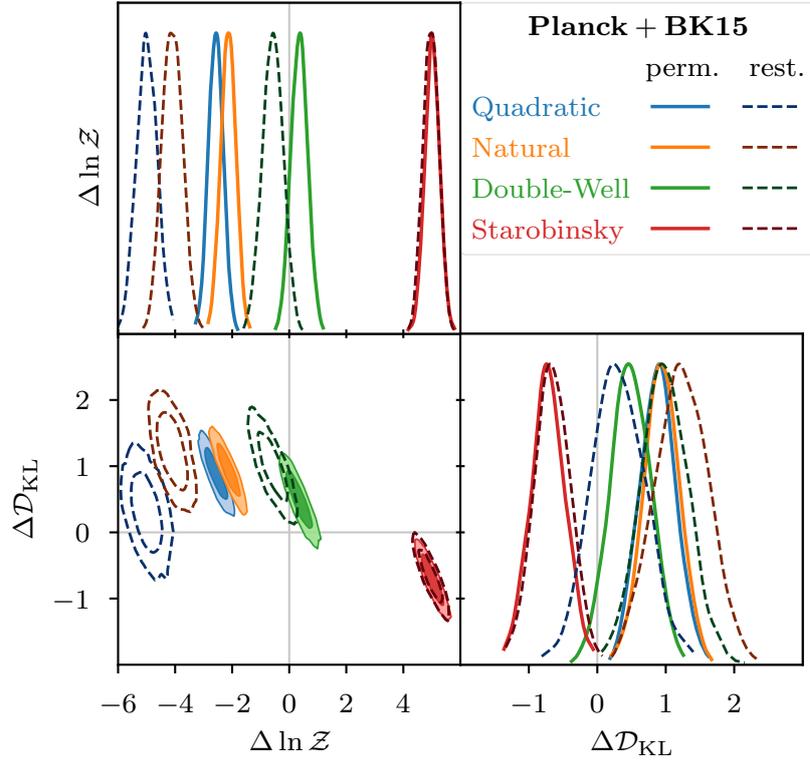


Figure 5.29 Bayesian model comparison of permissive (solid lines, cf. eq. (5.71)) and restrictive (dashed lines, cf. eq. (5.72)) reheating scenarios for various inflation models using P18 and BK15 data combined. We show quadratic inflation in blue, natural inflation in orange, results from the quartic double-well potential in green and the Starobinsky model in red. We show the log-evidence $\ln \mathcal{Z}$ and Kullback-Leibler divergence \mathcal{D}_{KL} (in nats). The Δ denote normalisation with respect to the base Λ CDM model indicated by the vertical and horizontal grey lines.

most sample points belong to the excluded region of parameter space with $w_{\text{reh}} > 1/3$, as is clearly visible in fig. 5.28.

In fig. 5.29 we compare the log-evidence $\ln \mathcal{Z}$ and KL-divergence \mathcal{D}_{KL} of the permissive with those of the restrictive reheating scenario for combined P18 and BK15 data. We omit the corresponding plot for P18 data on its own, because of the agglomeration of the contours on top of one another similar to the upper triangle plot in fig. 5.27. The filled contours in fig. 5.29 for the permissive reheating case from eq. (5.71) are the same as in the lower triangle plot in fig. 5.27. The unfilled contours with dashed lines come from the importance sampling with the restrictive reheating case from eq. (5.72).

For quadratic and natural inflation, the evidence shrinks further by two to three log-units. The contours for the quartic double-well model shift by only about one log-unit and the evidence for the Starobinsky model remains unchanged. This increases the gap between quadratic and Starobinsky inflation to Bayesian odds of overwhelming 1 : 10000.

5.14 Conclusion

Despite the success of flat Λ CDM there has been a persistent tendency towards positive curvature (closed universes) in CMB temperature and polarisation data. The latest data release from Planck in particular has sparked some discussion of possible evidence for spatial curvature in the CMB. In this paper we have investigated what the presence of such non-zero spatial curvature means for inflation.

We have reviewed how curvature links primordial and late-time scales, and how the detection of non-zero late-time curvature limits the total amount of inflation, thereby placing a bound on the comoving Hubble horizon, which becomes maximal at the onset of inflation. This sets tight constraints on initial conditions for inflation in order to solve the horizon and flatness problems, which we have folded into a Bayesian comparison of various inflationary potentials.

Note that, like in [chapter 2](#), the nested sampler (needed for determining the Bayesian evidence) did not apply a multivariate Gaussian prior on the two nuisance parameters associated with the Sunyaev–Zeldovich effect. Since this affects all models equally, we expect this to have only a negligible effect on the *comparison* of models. We will publish any quantitative changes in a paper based on this chapter.

We have computed the primordial power spectra from these inflation models numerically, revealing oscillations and a cutoff towards large scales, which are common features of finite inflation. Additionally we have shown how curvature leads to an additional suppression or to an amplification of power on large scales for closed and open universes respectively, which holds for both scalar and tensor perturbations.

In our Bayesian analysis we have used CMB data from the Planck 2018 legacy archive and from the 2015 observing season of BICEP2 and the Keck array. Nested sampling runs of the base Λ CDM cosmology and its extensions with the present-day curvature density parameter $\Omega_{K,0}$ (applying a uniform prior) and/or the tensor-to-scalar ratio r , presented in [fig. 5.16](#), have confirmed that the inclusion of curvature significantly weakens the bounds on the tensor-to-scalar ratio when only taking temperature and E -mode polarisation into account. This fails to hold, however, when including B -mode data, in which case the bounds match those of a flat cosmology. Note that we have adopted a uniform prior on the curvature density parameter $\Omega_{K,0}$. The CMB constraints on the spectral index n_s , on the other hand, point to a persistently larger value, roughly one standard deviation greater with than without curvature, albeit with roughly 6σ still clearly below scale invariance. Nevertheless, this changes the picture of slow-roll predictions from various inflation models, as shown in [fig. 5.12](#). We have computed the Bayesian evidence and Kullback–Leibler divergence for the various extensions and confirmed previous findings of the CMB having a preference for closed cosmologies. This preference is reduced when tensor modes are included, which come with a significant Occam penalty. Interestingly, the addition of B -mode data further strengthens the preference for closed universes. The details of this model comparison are presented in [fig. 5.19](#).

Using the aforementioned numerically integrated primordial power spectra, we have also computed the Bayesian posteriors and evidence from four single-field inflationary potentials: the quadratic, natural, quartic double-well and Starobinsky potential. We have found prior constraints on the primordial curvature, giving a lower bound from horizon considerations and upper bounds from considerations of an open or closed global geometry. Similarly, there are prior constraints on an initial inflaton field value. These are the combined effect of curvature linking the primordial to the late-time universe and of possible reheating scenarios.

As in previous findings considering curvature or finite inflation, an improved fit to CMB data is achieved via a suppression of power and smoothing of peaks on the largest scales*. In the absence of B -mode data all inflation models considered perform similarly well, a result of the weaker bound on the tensor-to-scalar ratio. However, with B -mode data taken into account, we obtain similar results to the flat case, with the Starobinsky model significantly outperforming the other inflation models, as seen in [fig. 5.27](#).

All four inflation models, with the exception of the Starobinsky model, are reheating constrained, which becomes very clear when looking at the spectral index and the tensor-to-scalar ratio as shown in [fig. 5.22](#). Both the high spectral index from the inclusion of non-zero curvature and the low tensor-to-scalar ratio from B -mode data push those inflation models to the edges of their prior constraints, with the limits set by the equation-of-state parameter of reheating. In the first instance we have only used very permissive reheating constraints. We have then used importance sampling to explore stricter reheating constraints, presented in [figs. 5.28](#) and [5.29](#), which has significantly penalised quadratic and natural inflation while keeping any inference on the Starobinsky model essentially invariant.

*This is analogous to effects of the artificial lensing parameter A_{lens} .

Conclusion

This thesis was concerned with quantifying previous findings that there might be observable imprints in the primordial power spectrum from a generically expected phase of kinetic dominance preceding inflation. The motivation for this was illustrated in [chapter 3](#) in the form of phase-space plots of the inflaton field. It was shown that the inflaton trajectories for concave and bounded potentials in particular start out kinetically dominated, independent of the prior choice for initial conditions.

In [chapter 4](#) a flat spatial geometry of the universe was assumed and a Bayesian analysis of the inflationary Universe was performed, comparing the standard slow-roll initial conditions with those from kinetic dominance. The cutoff in power on large scales from kinetic dominance was shown to lead to a better fit to the Planck data. However, in a flat Universe there is no limit to the total number of e-folds leading to unconstrained posterior distributions with a plateau reaching up to the prior boundaries. [Chapter 2](#) addressed such issues in light of their effects on the Bayesian evidence. It demonstrated that the specific choice of a prior bound is in fact not problematic for an unconstrained parameter, but leaves the Bayesian evidence invariant. With that in mind a comparison of slow-roll vs kinetic dominance initial conditions showed no decisive preference for either model. Any Occam penalty from the higher complexity of a model with a kinetically dominated stage was balanced out by a better fit to the large-scale lack of power in the CMB.

[Chapter 5](#) was an extension of the analysis from [chapter 4](#) with the added complexity of a non-zero spatial curvature. In many ways curved universes provide useful features. For one, they provide a scale that can be used to link the primordial to the standard Big Bang evolution without the need for further assumptions on an intermediate period of reheating. In other words the addition of the curvature density as a sampling parameter relieves us of the need to also sample over additional reheating parameters. Instead these reheating parameters become derived parameters. Broadly speaking, fewer sampling parameters means a smaller model complexity and therefore also a smaller Occam penalty on the model. More accurately, however, [chapter 2](#) clarifies that only *constrained* sampling parameters add to a model's complexity and Occam penalty. Any additional sampling parameter will add to the computational complexity and run-time, though, which is also referred to as the “curse of dimensionality”.

Future work

Computational predictions for the total amount of inflation

Little effort has been invested so far into an expectation value of the total amount of inflation. The best resource on that question is the paper by Remmen and Carroll (2014) [122]. However, they only consider the quadratic and natural potential in a flat cosmology, both of which are essentially ruled out by Planck data. With the tensor-to-scalar ratio of the Starobinsky potential being one of the prime targets of future CMB missions, it would be interesting to extend the analysis therewith. This will require some thought, though, since the Planck surface or any other choice of initial surface is no longer closed in phase-space for the Starobinsky model, potentially complicating a viable choice of phase-space measure. In light of the limiting effect of spatial curvature on inflation, including this in an analysis would be interesting as well. However, primordial curvature effectively adds another dimension to the phase-space trajectories, which again breaks the closure of phase-space.

Forecasts and updates with upcoming CMB experiments

The increased sensitivity to B -modes of upcoming CMB experiments such as the Simons Observatory [45], CMB Stage-4 [47] or the LiteBIRD satellite [46] will allow pushing to a tensor-to-scalar ratio as low as $r \sim 10^{-3}$ (making [chapter 2](#) even more relevant). This will significantly improve constraints on inflation models. In the absence of any detection, it will allow us to rule out the entire large-field inflation branch.

Like Planck, the LiteBIRD satellite will observe the entire sky, but with a significantly improved sensitivity on CMB polarisation, i.e. both E - and B -modes. This means that in contrast to ground based experiments, LiteBIRD provides data also on the largest scales, down to multipoles of $\ell = 2$. Thus it will be able to access both the reionization bump at $\ell \simeq 4$ and the recombination bump at $\ell \simeq 80$, rendering a detection of a B -mode signal substantially more likely.

Since a finite amount of inflation only affects the primordial power spectrum on very large scales, it will be particularly interesting to see whether the LiteBIRD satellite will further substantiate the large-scale lack of power found so far in CMB data. In the meantime it would be interesting to explore this question with forecasts of these experiments, assuming fiducial models generated both with and without a limited amount of inflation.

Inhomogeneities

As already remarked in [chapter 5](#), primordial inhomogeneities are another possible cause of finite inflation. Their energy density scales as $\rho_{\text{inh}} \sim a^{-4}$, while that of the homogeneous component is approximately constant. Thus the universe grows more inhomogeneous backwards in time, but forwards in time it becomes homogeneous, the energy density and the Hubble parameter become constant and inflation starts. It would be interesting to investigate whether this leads to similar effects in the primordial power spectrum as from a kinetically

dominated phase. How this might tie in with a kinetically dominated phase still requires further investigation.

Reheating priors

In [chapter 5](#) the curvature density parameter was assigned a uniform prior (although effects from other constraints skew this prior toward flatness), which set implicit priors on the derived reheating parameters. It would be interesting to turn this dependence around and assume uniform priors on the equation of state parameter and energy density of reheating instead, such that the curvature density becomes a derived parameter. This would allow longer periods of inflation a priori and thus impose a stronger prior preference towards spatial flatness. It would then be interesting to see, whether this prior preference will still be overpowered by the preference of the CMB data for a closed Universe. Similarly, one could also envision using a prior peaked at zero as suggested e.g. in [\[182\]](#).

Final remarks

The Planck satellite has brought an unprecedented precision to cosmology and solidified the standing of the current “standard” Λ CDM model. However, there are many questions that still remain open. Excitingly, some of these questions should be answered in the near future, with improved data. On the one hand, as already mentioned, we soon should be able to either detect CMB B -modes or rule out the entire large-field inflation branch. On the other hand, we soon might have clear cosmological constraints on neutrino masses, confirming the true neutrino hierarchy. Both of these results would go beyond a mere improvement of constraints, but actually answer concrete questions. And with additional data on the largest accessible scales, we will be better equipped in answering the question whether there is a large scale lack of power associated with an early kinetically dominated universe.

References

- [I & chapter 3] L. T. Hergt, Will Handley, M. P. Hobson, and A. N. Lasenby. “[Case for kinetically dominated initial conditions for inflation](#)”. *Physical Review D* 100.2 (2019), p. 023502. arXiv: [1809.07185](#). [This paper corresponds to [chapter 3](#) of this thesis.] (Cit. on pp. [vii](#), [47](#), [55](#), [56](#), [59](#), [80](#), [87](#)).
- [II & chapter 4] L. T. Hergt, Will Handley, M. P. Hobson, and A. N. Lasenby. “[Constraining the kinetically dominated universe](#)”. *Physical Review D* 100.2 (2019), p. 023501. arXiv: [1809.07737](#). [This paper corresponds to [chapter 4](#) of this thesis.] (Cit. on pp. [vii](#), [47](#), [48](#), [55](#), [80](#), [87](#), [119](#), [135](#)).
- [1] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, and Paul van Mulbregt. “[SciPy 1.0: fundamental algorithms for scientific computing in Python](#)”. *Nature Methods* 17.3 (2020), pp. 261–272. arXiv: [1907.10121](#) (cit. on pp. [xv](#), [83](#)).
- [2] *CODATA Recommended Values of the Fundamental Physical Constants: 2018*. URL: <https://physics.nist.gov/cuu/Constants/> (visited on October 21, 2020) (cit. on p. [xv](#)).
- [3] *IAU 2012 Resolution B2 on the re-definition of the astronomical unit of length*. International Astronomical Union, 2012 (cit. on p. [xv](#)).
- [4] *IAU 2015 Resolution B2 on recommended zero points for the absolute and apparent bolometric magnitude scales*. International Astronomical Union, 2015 (cit. on p. [xv](#)).
- [5] Scott Dodelson and Lam Hui. “[Horizon Ratio Bound for Inflationary Fluctuations](#)”. *Physical Review Letters* 91.13 (2003), p. 131301. arXiv: [0305113](#) [[astro-ph](#)] (cit. on pp. [xvi](#), [64](#), [70](#)).
- [6] Viatcheslav Mukhanov. “[Physical Foundations of Cosmology](#)”. Cambridge University Press, 2005, pp. 1–421 (cit. on p. [1](#)).
- [7] Scott Dodelson. “[Modern Cosmology](#)”. Elsevier, 2003, pp. 1–440 (cit. on pp. [1](#), [16](#)).
- [8] M. P. Hobson, G. Efstathiou, and A. N. Lasenby. “[General Relativity: An Introduction for Physicists](#)”. Cambridge University Press, 2006, pp. 1–572 (cit. on pp. [1](#), [10](#), [16](#), [62](#)).
- [9] Steven Weinberg. “[Cosmology](#)”. Oxford University Press, 2008, pp. 1–593 (cit. on pp. [1](#), [16](#)).
- [10] David Hilary Lyth and Andrew R. Liddle. “[The Primordial Density Perturbation: Cosmology, Inflation and the Origin of Structure](#)”. Cambridge University Press, 2009, pp. 1–497 (cit. on pp. [1](#), [16](#)).
- [11] Daniel Baumann. “[TASI Lectures on Inflation](#)” (2009). arXiv: [0907.5424](#) (cit. on pp. [1](#), [2](#), [16](#)).
- [12] Douglas Scott. “[The Standard Model of Cosmology: A Skeptic’s Guide](#)” (2018). arXiv: [1804.01318](#) (cit. on pp. [1](#), [6](#)).
- [13] A. A. Penzias and R. W. Wilson. “[A Measurement of Excess Antenna Temperature at 4080 Mc/s.](#)” *The Astrophysical Journal* 142 (1965), p. 419 (cit. on p. [3](#)).

- [14] R. H. Dicke, P. J. E. Peebles, P. G. Roll, and D. T. Wilkinson. “Cosmic Black-Body Radiation.” *The Astrophysical Journal* 142 (1965), p. 414 (cit. on p. 3).
- [15] Robert Woodrow Wilson. *The Cosmic Microwave Background Radiation*. Nobel Lecture. 1978 (cit. on p. 3).
- [16] Planck Collaboration. “Planck 2018 results: I. Overview and the cosmological legacy of Planck”. *Astronomy & Astrophysics* 641 (2020), A1. arXiv: [1807.06205](https://arxiv.org/abs/1807.06205) (cit. on pp. 3, 79).
- [17] C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, D. N. Spergel, G. S. Tucker, E. Wollack, E. L. Wright, C. Barnes, M. R. Greason, R. S. Hill, E. Komatsu, M. R. Nolta, N. Odegard, H. V. Peiris, L. Verde, and J. L. Weiland. “First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Preliminary Maps and Basic Results”. *The Astrophysical Journal Supplement Series* 148.1 (2003), pp. 1–27. arXiv: [0302207](https://arxiv.org/abs/0302207) [[astro-ph](https://arxiv.org/archive/astro)] (cit. on p. 3).
- [18] G. Hinshaw, J. L. Weiland, R. S. Hill, N. Odegard, D. Larson, C. L. Bennett, J. Dunkley, B. Gold, M. R. Greason, N. Jarosik, E. Komatsu, M. R. Nolta, L. Page, D. N. Spergel, E. Wollack, M. Halpern, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. L. Wright. “Five-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Data Processing, Sky Maps, and Basic Results”. *The Astrophysical Journal Supplement Series* 180.2 (2009), pp. 225–245. arXiv: [0803.0732](https://arxiv.org/abs/0803.0732) (cit. on p. 3).
- [19] N. Jarosik, C. L. Bennett, J. Dunkley, B. Gold, M. R. Greason, M. Halpern, R. S. Hill, G. Hinshaw, A. Kogut, E. Komatsu, D. Larson, M. Limon, S. S. Meyer, M. R. Nolta, N. Odegard, L. Page, K. M. Smith, D. N. Spergel, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. “Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Sky Maps, Systematic Errors, and Basic Results”. *The Astrophysical Journal Supplement Series* 192.2 (2011), p. 14. arXiv: [1001.4744](https://arxiv.org/abs/1001.4744) (cit. on p. 3).
- [20] C. L. Bennett, D. Larson, J. L. Weiland, N. Jarosik, G. Hinshaw, N. Odegard, K. M. Smith, R. S. Hill, B. Gold, M. Halpern, E. Komatsu, M. R. Nolta, L. Page, D. N. Spergel, E. Wollack, J. Dunkley, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. L. Wright. “Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results”. *The Astrophysical Journal Supplement Series* 208.2 (2013), p. 20. arXiv: [1212.5225](https://arxiv.org/abs/1212.5225) (cit. on p. 3).
- [21] Planck Collaboration. “Planck 2013 results: I. Overview of products and scientific results”. *Astronomy & Astrophysics* 571 (2014), A1. arXiv: [1303.5062](https://arxiv.org/abs/1303.5062) (cit. on pp. 3, 55, 79).
- [22] Planck Collaboration. “Planck 2015 results: I. Overview of products and scientific results”. *Astronomy & Astrophysics* 594 (2016), A1. arXiv: [1502.01582](https://arxiv.org/abs/1502.01582) (cit. on pp. 3, 55, 79).
- [23] D. J. Fixsen. “The Temperature of the Cosmic Microwave Background”. *The Astrophysical Journal* 707.2 (2009), pp. 916–920. arXiv: [0911.1955](https://arxiv.org/abs/0911.1955) (cit. on pp. 3, 7).
- [24] Adam G. Riess. “The expansion of the Universe is faster than expected”. *Nature Reviews Physics* 2.1 (2020), pp. 10–12. arXiv: [2001.03624](https://arxiv.org/abs/2001.03624) (cit. on p. 4).
- [25] Planck Collaboration. “Planck 2018 results: VIII. Gravitational lensing”. *Astronomy & Astrophysics* 641 (2020), A8. arXiv: [1807.06210](https://arxiv.org/abs/1807.06210) (cit. on p. 4).
- [26] Daniel J. Eisenstein, Idit Zehavi, David W. Hogg, Roman Scoccimarro, Michael R. Blanton, Robert C. Nichol, Ryan Scranton, Hee-Jong Seo, Max Tegmark, Zheng Zheng, Scott F. Anderson, Jim Annis, Neta Bahcall, Jon Brinkmann, Scott Burles, Francisco J. Castander, Andrew Connolly, Istvan Csabai, Mamoru Doi, Masataka Fukugita, Joshua A. Frieman, Karl Glazebrook, James E. Gunn, John S. Hendry, Gregory Hennessy, Zeljko Ivezic, Stephen Kent, Gillian R. Knapp, Huan Lin, Yeong-Shang Loh, Robert H. Lupton, Bruce Margon, Timothy A. McKay, Avery Meiksin, Jeffery A. Munn, Adrian Pope, Michael W. Richmond, David Schlegel, Donald P. Schneider, Kazuhiro Shimasaku, Christopher Stoughton, Michael A. Strauss, Mark SubbaRao, Alexander S. Szalay, Istvan Szapudi, Douglas L. Tucker, Brian Yanny,

- and Donald G. York. “Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies”. *The Astrophysical Journal* 633.2 (2005), pp. 560–574. arXiv: 0501171 [astro-ph] (cit. on p. 4).
- [27] Will Handley. *Curvature tension: evidence for a closed universe*. 2019. arXiv: 1908.09139 (cit. on pp. 5, 7, 26, 80, 84, 121, 124).
- [28] Eleonora Di Valentino, Alessandro Melchiorri, and Joseph Silk. “Planck evidence for a closed Universe and a possible crisis for cosmology”. *Nature Astronomy* (2019), pp. 1–8. arXiv: 1911.02087 (cit. on pp. 5, 7, 80, 84, 121, 124).
- [29] Eleonora Di Valentino, Alessandro Melchiorri, and Joseph Silk. *Cosmic Discordance: Planck and luminosity distance data exclude Λ CDM*. 2020. arXiv: 2003.04935 (cit. on pp. 5, 7, 80, 84, 121).
- [30] Saul Perlmutter. *Measuring the Acceleration of the Cosmic Expansion Using Supernovae*. Nobel Lecture. 2011 (cit. on p. 5).
- [31] Adam G. Riess, Stefano Casertano, Wenlong Yuan, Lucas M. Macri, and Dan Scolnic. “Large Magellanic Cloud Cepheid Standards Provide a 1% Foundation for the Determination of the Hubble Constant and Stronger Evidence for Physics beyond Λ CDM”. *The Astrophysical Journal* 876.1 (2019), p. 85. arXiv: 1903.07603 (cit. on p. 5).
- [32] Wendy L. Freedman, Barry F. Madore, Dylan Hatt, Taylor J. Hoyt, In Sung Jang, Rachael L. Beaton, Christopher R. Burns, Myung Gyoon Lee, Andrew J. Monson, Jillian R. Neeley, M. M. Phillips, Jeffrey A. Rich, and Mark Seibert. “The Carnegie-Chicago Hubble Program. VIII. An Independent Determination of the Hubble Constant Based on the Tip of the Red Giant Branch”. *The Astrophysical Journal* 882.1 (2019), p. 34. arXiv: 1907.05922 (cit. on p. 5).
- [33] Planck Collaboration. “Planck 2018 results: VI. Cosmological parameters”. *Astronomy & Astrophysics* 641 (2020), A6. arXiv: 1807.06209 (cit. on pp. 6–9, 24, 28, 29, 38, 39, 56, 79–81, 104, 112, 120).
- [34] Planck Collaboration. “Planck 2013 results: XVI. Cosmological parameters”. *Astronomy & Astrophysics* 571 (2014), A16. arXiv: 1303.5076 [astro-ph.CO] (cit. on pp. 6, 79, 81, 104).
- [35] Planck Collaboration. “Planck 2015 results: XIII. Cosmological parameters”. *Astronomy & Astrophysics* 594 (2016), A13. arXiv: 1502.01589 (cit. on pp. 6, 48, 49, 51, 68, 79, 81, 104).
- [36] E. R. Harrison. “Fluctuations at the Threshold of Classical Cosmology”. *Physical Review D* 1.10 (1970), pp. 2726–2730 (cit. on p. 7).
- [37] P. J. E. Peebles and J. T. Yu. “Primeval Adiabatic Perturbation in an Expanding Universe”. *The Astrophysical Journal* 162 (1970), p. 815 (cit. on p. 7).
- [38] Y. B. Zeldovich. “A Hypothesis, Unifying the Structure and the Entropy of the Universe”. *Monthly Notices of the Royal Astronomical Society* 160.1 (1972), 1P–3P (cit. on p. 7).
- [39] Gianpiero Mangano, Gennaro Miele, Sergio Pastor, Teguyco Pinto, Ofelia Pisanti, and Pasquale D. Serpico. “Relic neutrino decoupling including flavour oscillations”. *Nuclear Physics B* 729.1-2 (2005), pp. 221–234. arXiv: 0506164 [hep-ph] (cit. on pp. 8, 38).
- [40] Pablo F. de Salas and Sergio Pastor. “Relic neutrino decoupling with flavour oscillations revisited”. *Journal of Cosmology and Astroparticle Physics* 2016.07 (2016), pp. 051–051. arXiv: 1606.06986 (cit. on pp. 8, 38).
- [41] Lars Husdal. “On Effective Degrees of Freedom in the Early Universe”. *Galaxies* 4.4 (2016), p. 78 (cit. on p. 8).
- [42] Ivan Esteban, M. C. Gonzalez-Garcia, Alvaro Hernandez-Cabezudo, Michele Maltoni, and Thomas Schwetz. “Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ_{23} , δ CP, and the mass ordering”. *Journal of High Energy Physics* 2019.1 (2019), p. 106. arXiv: 1811.05487 (cit. on pp. 8, 27).

- [43] Ivan Esteban, M. C. Gonzalez-Garcia, Michele Maltoni, Thomas Schwetz, and Albert Zhou. *The fate of hints: updated global analysis of three-flavor neutrino oscillations*. 2020. arXiv: 2007.14792 (cit. on pp. 8, 27, 38).
- [44] *NuFIT / NuFIT 5.0*. <http://www.nu-fit.org/>. Accessed: 2020-08-13. 2020 (cit. on pp. 8, 27, 38).
- [45] The Simons Observatory Collaboration. “The Simons Observatory: science goals and forecasts”. *Journal of Cosmology and Astroparticle Physics* 2019.02 (2019), pp. 056–056. arXiv: 1808.07445 (cit. on pp. 8, 24, 107, 144).
- [46] Adrian Lee et al. “Astro2020: APC White Papers. LiteBIRD: an all-sky cosmic microwave background probe of inflation”. *Bulletin of the AAS* 51.7 (2019) (cit. on pp. 8, 24, 107, 144).
- [47] CMB-S4 Collaboration. *CMB-S4 Science Case, Reference Design, and Project Plan*. 2019. arXiv: 1907.04473 (cit. on pp. 8, 24, 107, 144).
- [48] A. Einstein. “Die Grundlage der allgemeinen Relativitätstheorie”. *Annalen der Physik* 354.7 (1916), pp. 769–822 (cit. on pp. 9, 84).
- [49] A. Friedman. “Über die Krümmung des Raumes”. *Zeitschrift für Physik* 10.1 (1922), pp. 377–386 (cit. on pp. 10, 79, 84).
- [50] Georges Lemaître. “L’Univers en expansion”. *Annales de la Société Scientifique de Bruxelles* A53 (1933), p. 51 (cit. on pp. 10, 79, 84).
- [51] H. P. Robertson. “Kinematics and World-Structure”. *The Astrophysical Journal* 82 (1935), p. 284 (cit. on pp. 10, 79, 84).
- [52] H. P. Robertson. “Kinematics and World-Structure II.” *The Astrophysical Journal* 83 (1936), p. 187 (cit. on pp. 10, 79, 84).
- [53] H. P. Robertson. “Kinematics and World-Structure III.” *The Astrophysical Journal* 83 (1936), p. 257 (cit. on pp. 10, 79, 84).
- [54] A. G. Walker. “On Milne’s Theory of World-Structure *”. *Proceedings of the London Mathematical Society* s2-42.1 (1937), pp. 90–127 (cit. on pp. 10, 79, 84).
- [55] A. A. Starobinskiĭ. “Spectrum of relict gravitational radiation and the early state of the universe”. *Journal of Experimental and Theoretical Physics Letters* 30 (1979), p. 682 (cit. on pp. 13, 55, 79).
- [56] Alan H. Guth. “Inflationary universe: A possible solution to the horizon and flatness problems”. *Physical Review D* 23.2 (1981), pp. 347–356 (cit. on pp. 13, 55, 79).
- [57] A.D. Linde. “A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems”. *Physics Letters B* 108.6 (1982), pp. 389–393 (cit. on pp. 13, 55, 79).
- [58] Andreas Albrecht and Paul J. Steinhardt. “Cosmology for Grand Unified Theories with Radiatively Induced Symmetry Breaking”. *Physical Review Letters* 48.17 (1982), pp. 1220–1223 (cit. on pp. 13, 79).
- [59] Steven Weinberg. “Inflation: Three puzzles”. *Cosmology*. Oxford University Press, 2008. Chap. 4.1, pp. 202–208 (cit. on p. 15).
- [60] M. P. Hobson, G. Efstathiou, and A. N. Lasenby. “Inflationary cosmology: Classical evolution of curvature perturbations”. *General Relativity: An Introduction for Physicists*. Cambridge University Press, 2006. Chap. 16.15, pp. 449–452 (cit. on pp. 17, 94, 106).
- [61] Steven Weinberg. “Primary fluctuations in the microwave background: A first look”. *Cosmology*. Oxford University Press, 2008. Chap. 2.6, pp. 135–148 (cit. on p. 18).
- [62] Planck Collaboration. “Planck 2018 results: V. CMB power spectra and likelihoods”. *Astronomy & Astrophysics* 641 (2020), A5. arXiv: 1907.12875 (cit. on pp. 19, 21, 27, 34, 42, 83, 84, 120, 122).
- [63] Planck Collaboration. “Planck 2013 results: XV. CMB power spectra and likelihood”. *Astronomy & Astrophysics* 571 (2014), A15. arXiv: 1303.5075 (cit. on p. 21).

- [64] Planck Collaboration. “Planck 2015 results: XI. CMB power spectra, likelihoods, and robustness of parameters”. *Astronomy & Astrophysics* 594 (2016), A11. arXiv: [1507.02704](#) (cit. on pp. 21, 68, 70).
- [65] J. Bernoulli. “Ars conjectandi”. Thurnisiorum, Basilae, 1713 (cit. on p. 23).
- [66] John Maynard Keynes. “A Treatise On Probability”. Macmillan and Co., 1921 (cit. on p. 23).
- [67] D. S. Sivia and J. Skilling. “Data Analysis: A Bayesian Tutorial”. Oxford University Press, 2006, pp. 1–246 (cit. on pp. 23, 25).
- [68] King Lau, Jia-Yu Tang, and Ming-Chung Chu. “Cosmic microwave background constraints on the tensor-to-scalar ratio”. *Research in Astronomy and Astrophysics* 14.6 (2014), pp. 635–647. arXiv: [1305.3921](#) (cit. on p. 24).
- [69] Gabriela Barenboim and Wan-Il Park. *On the tensor-to-scalar ratio in large single-field inflation models*. 2015. arXiv: [1509.07132](#) (cit. on p. 24).
- [70] Paolo Creminelli, Sergei Dubovsky, Diana López Nacir, Marko Simonović, Gabriele Trevisan, Giovanni Villadoro, and Matias Zaldarriaga. “Implications of the scalar tilt for the tensor-to-scalar ratio”. *Physical Review D* 92.12 (2015), p. 123528. arXiv: [1412.0678](#) (cit. on p. 24).
- [71] The Keck Array and BICEP2 Collaborations. “Constraints on Primordial Gravitational Waves Using Planck, WMAP, and New BICEP2/Keck Observations through the 2015 Season”. *Physical Review Letters* 121.22 (2018), p. 221301. arXiv: [1810.05216](#) (cit. on pp. 24, 29, 83, 84, 108, 120, 122).
- [72] Koichi Hirano. *Inflation with very small tensor-to-scalar ratio*. 2019. arXiv: [1912.12515](#) (cit. on p. 24).
- [73] Fergus Simpson, Raul Jimenez, Carlos Pena-Garay, and Licia Verde. “Strong Bayesian evidence for the normal neutrino hierarchy”. *Journal of Cosmology and Astroparticle Physics* 2017.06 (2017), pp. 029–029. arXiv: [1703.03425](#) (cit. on pp. 24, 44).
- [74] T. Schwetz, K. Freese, M. Gerbino, E. Giusarma, S. Hannestad, M. Lattanzi, O. Mena, and S. Vagnozzi. *Comment on “Strong Evidence for the Normal Neutrino Hierarchy”*. 2017. arXiv: [1703.04585](#) (cit. on pp. 24, 44).
- [75] Allen Caldwell, Alexander Merle, Oliver Schulz, and Maximilian Totzauer. “Global Bayesian analysis of neutrino mass data”. *Physical Review D* 96.7 (2017), p. 073001. arXiv: [1705.01945](#) (cit. on p. 24).
- [76] Francesco Capozzi, Eleonora Di Valentino, Eligio Lisi, Antonio Marrone, Alessandro Melchiorri, and Antonio Palazzo. “Global constraints on absolute neutrino masses and their ordering”. *Physical Review D* 95.9 (2017), p. 096014. arXiv: [1703.04471](#) (cit. on p. 24).
- [77] Arthur Loureiro, Andrei Cuceu, Filipe B. Abdalla, Bruno Moraes, Lorne Whiteway, Michael McLeod, Sreekumar T. Balan, Ofer Lahav, Aurélien Benoit-Lévy, Marc Manera, Richard P. Rollins, and Henrique S. Xavier. “Upper Bound of Neutrino Masses from Combined Cosmological Observations and Particle Physics Experiments”. *Physical Review Letters* 123.8 (2019), p. 081301. arXiv: [1811.02578](#) (cit. on p. 24).
- [78] Francesco Capozzi, Eleonora Di Valentino, Eligio Lisi, Antonio Marrone, Alessandro Melchiorri, and Antonio Palazzo. “Addendum to “Global constraints on absolute neutrino masses and their ordering””. *Physical Review D* 101.11 (2020), p. 116013. arXiv: [2003.08511](#) (cit. on p. 24).
- [79] Maria Archidiacono, Steen Hannestad, and Julien Lesgourgues. *What will it take to measure individual neutrino mass states using cosmology?* 2020. arXiv: [2003.03354](#) (cit. on pp. 24, 38).
- [80] Shouvik Roy Choudhury and Steen Hannestad. “Updated results on neutrino mass and mass hierarchy from cosmology with Planck 2018 likelihoods”. *Journal of Cosmology and Astroparticle Physics* 2020.07 (2020), pp. 037–037. arXiv: [1907.12598](#) (cit. on pp. 24, 44).

- [81] Patrick Stöcker, Csaba Balázs, Sanjay Bloor, Torsten Bringmann, Tomás E. Gonzalo, Will Handley, Selim Hotinli, Cullan Howlett, Felix Kahlhoefer, Janina J. Renk, Pat Scott, Aaron C. Vincent, and Martin White. *Strengthening the bound on the mass of the lightest neutrino with terrestrial and cosmological experiments*. 2020. arXiv: [2009.03287](#) (cit. on p. 24).
- [82] David J. C. MacKay. “Model Comparison and Occam’s Razor”. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003. Chap. 28, pp. 343–355 (cit. on pp. 25, 26).
- [83] David J. C. MacKay. “Laplace’s Method”. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003. Chap. 28, pp. 341–342 (cit. on p. 25).
- [84] James M. Dickey. “The Weighted Likelihood Ratio, Linear Hypotheses on Normal Location Parameters”. *The Annals of Mathematical Statistics* 42.1 (1971), pp. 204–223 (cit. on p. 25).
- [85] Roberto Trotta. “Applications of Bayesian model selection to cosmological parameters”. *Monthly Notices of the Royal Astronomical Society* 378.1 (2007), pp. 72–82. arXiv: [0504022 \[astro-ph\]](#) (cit. on p. 25).
- [86] Roberto Trotta. “Forecasting the Bayes factor of a future observation”. *Monthly Notices of the Royal Astronomical Society* 378.3 (2007), pp. 819–824. arXiv: [0703063 \[astro-ph\]](#) (cit. on p. 25).
- [87] Licia Verde, Stephen M. Feeney, Daniel J. Mortlock, and Hiranya V. Peiris. “(Lack of) Cosmological evidence for dark radiation after Planck”. *Journal of Cosmology and Astroparticle Physics* 2013.09 (2013), pp. 013–013. arXiv: [1307.2904](#) (cit. on p. 25).
- [88] Alan Heavens, Yabebal Fantaye, Arrykrishna Mootoovaloo, Hans Eggers, Zafirah Hosenie, Steve Kroon, and Elena Sellentin. *Marginal Likelihoods from Monte Carlo Markov Chains*. 2017. arXiv: [1704.03472](#) (cit. on p. 25).
- [89] Alan Heavens, Yabebal Fantaye, Elena Sellentin, Hans Eggers, Zafirah Hosenie, Steve Kroon, and Arrykrishna Mootoovaloo. “No Evidence for Extensions to the Standard Cosmological Model”. *Physical Review Letters* 119.10 (2017), p. 101301. arXiv: [1704.03467](#) (cit. on p. 25).
- [90] John Skilling. “Nested sampling for general Bayesian computation”. *Bayesian Analysis* 1.4 (2006), pp. 833–859 (cit. on p. 25).
- [91] D. S. Sivia and J. Skilling. “Nested sampling”. *Data Analysis: A Bayesian Tutorial*. Oxford University Press, 2006. Chap. 9, pp. 181–208 (cit. on p. 25).
- [92] Farhan Feroz and M. P. Hobson. “Multimodal nested sampling: an efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analyses”. *Monthly Notices of the Royal Astronomical Society* 384.2 (2008), pp. 449–463. arXiv: [0704.3704](#) (cit. on pp. 25, 70).
- [93] F. Feroz, M. P. Hobson, and M. Bridges. “MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics”. *Monthly Notices of the Royal Astronomical Society* 398.4 (2009), pp. 1601–1614. arXiv: [0809.3437](#) (cit. on pp. 25, 70).
- [94] Farhan Feroz, Michael P. Hobson, Ewan Cameron, and Anthony N. Pettitt. “Importance Nested Sampling and the MultiNest Algorithm”. *The Open Journal of Astrophysics* 2.1 (2019). arXiv: [1306.2144](#) (cit. on pp. 25, 70).
- [95] W. J. Handley, M. P. Hobson, and A. N. Lasenby. “PolyChord: nested sampling for cosmology”. *Monthly Notices of the Royal Astronomical Society: Letters* 450.1 (2015), pp. L61–L65. arXiv: [1502.01856](#) (cit. on pp. 25, 28, 29, 70, 80, 82, 84).
- [96] W. J. Handley, M. P. Hobson, and A. N. Lasenby. “PolyChord: next-generation nested sampling”. *Monthly Notices of the Royal Astronomical Society* 453.4 (2015), pp. 4385–4399. arXiv: [1506.00171](#) (cit. on pp. 25, 28, 70, 80, 82, 84).
- [97] D. S. Sivia and J. Skilling. “Model selection”. *Data Analysis: A Bayesian Tutorial*. Oxford University Press, 2006. Chap. 4, pp. 78–102 (cit. on pp. 26, 35).

- [98] Will Handley and Pablo Lemos. “Quantifying tensions in cosmological parameters: Interpreting the DES evidence ratio”. *Physical Review D* 100.4 (2019), p. 043504. arXiv: [1902.04029](#) (cit. on pp. [26](#), [75](#)).
- [99] David J. Spiegelhalter, Nicola G. Best, Bradley P. Carlin, and Angelika van der Linde. “Bayesian measures of model complexity and fit”. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64.4 (2002), pp. 583–639 (cit. on p. [27](#)).
- [100] Will Handley and Pablo Lemos. “Quantifying dimensionality: Bayesian cosmological model complexities”. *Physical Review D* 100.2 (2019), p. 023512. arXiv: [1903.06682](#) (cit. on pp. [27](#), [83](#)).
- [101] Jesus Torrado and Antony Lewis. *Cobaya: Code for Bayesian Analysis of hierarchical physical models*. 2020. arXiv: [2005.05290](#) (cit. on pp. [28](#), [80](#), [84](#)).
- [102] Antony Lewis and Sarah Bridle. “Cosmological parameters from CMB and other data: A Monte Carlo approach”. *Physical Review D* 66.10 (2002), p. 103511. arXiv: [0205436](#) [[astro-ph.CO](#)] (cit. on pp. [28](#), [68](#), [84](#)).
- [103] Antony Lewis. “Efficient sampling of fast and slow cosmological parameters”. *Physical Review D* 87.10 (2013), p. 103529. arXiv: [1304.4473](#) [[astro-ph.CO](#)] (cit. on pp. [28](#), [68](#), [84](#)).
- [104] Radford M. Neal. *Taking Bigger Metropolis Steps by Dragging Fast Variables*. 2005. arXiv: [0502099](#) [[math](#)] (cit. on pp. [28](#), [84](#)).
- [105] Julien Lesgourgues. *The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview*. 2011. arXiv: [1104.2932](#) (cit. on pp. [28](#), [68](#), [80](#), [84](#)).
- [106] Diego Blas, Julien Lesgourgues, and Thomas Tram. “The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes”. *Journal of Cosmology and Astroparticle Physics* 2011.07 (2011), pp. 034–034. arXiv: [1104.2933](#) (cit. on pp. [28](#), [68](#), [84](#)).
- [107] Julien Lesgourgues and Thomas Tram. “The Cosmic Linear Anisotropy Solving System (CLASS) IV: efficient implementation of non-cold relics”. *Journal of Cosmology and Astroparticle Physics* 2011.09 (2011), pp. 032–032. arXiv: [1104.2935](#) (cit. on pp. [28](#), [68](#)).
- [108] Antony Lewis. *GetDist: a Python package for analysing Monte Carlo samples*. 2019. arXiv: [1910.13970](#) (cit. on pp. [29](#), [70](#), [84](#)).
- [109] Will Handley. “anesthetic: nested sampling visualisation”. *Journal of Open Source Software* 4.37 (2019), p. 1414. arXiv: [1905.04768](#) (cit. on pp. [29](#), [32](#), [81](#), [84](#), [138](#)).
- [110] John C. Forbes. *A PDF PSA, or Never gonna set_xscale again – guilty feats with logarithms*. 2020. arXiv: [2003.14327](#) (cit. on p. [32](#)).
- [111] Martina Gerbino and Massimiliano Lattanzi. “Status of Neutrino Properties and Future Prospects—Cosmological and Astrophysical Constraints”. *Frontiers in Physics* 5.FEB (2018), p. 70 (cit. on p. [38](#)).
- [112] Particle Data Group. “Review of Particle Physics”. *Progress of Theoretical and Experimental Physics* 2020.8 (2020) (cit. on p. [42](#)).
- [113] Steen Hannestad and Thomas Schwetz. “Cosmology and the neutrino mass ordering”. *Journal of Cosmology and Astroparticle Physics* 2016.11 (2016), pp. 035–035. arXiv: [1606.04691](#) (cit. on p. [44](#)).
- [114] Alan F. Heavens and Elena Sellentin. “Objective Bayesian analysis of neutrino masses and hierarchy”. *Journal of Cosmology and Astroparticle Physics* 2018.04 (2018), pp. 047–047. arXiv: [1802.09450](#) (cit. on p. [44](#)).
- [115] Will Handley, S. D. Brechet, A. N. Lasenby, and M. P. Hobson. “Kinetic initial conditions for inflation”. *Physical Review D* 89.6 (2014), p. 063505. arXiv: [1401.2253](#) (cit. on pp. [47](#), [55](#), [56](#), [59](#), [80](#), [87](#)).
- [116] V.A. Belinsky, L.P. Grishchuk, I.M. Khalatnikov, and Ya.B. Zeldovich. “Inflationary stages in cosmological models with a scalar field”. *Physics Letters B* 155.4 (1985), pp. 232–236 (cit. on pp. [47](#), [48](#), [55](#), [59](#), [80](#), [87](#)).
- [117] A.D. Linde. “Initial conditions for inflation”. *Physics Letters B* 162.4–6 (1985), pp. 281–286 (cit. on pp. [47](#), [48](#), [52](#), [55](#), [59](#)).

- [118] V. A. Belinsky, H. Ishihara, I. M. Khalatnikov, and H. Sato. “On the Degree of Generality of Inflation in Friedmann Cosmological Models with a Massive Scalar Field”. *Progress of Theoretical Physics* 79.3 (1988), pp. 676–684 (cit. on pp. 47, 48, 52, 80, 87).
- [119] Andrei Linde. “Inflationary Cosmology”. *Inflationary Cosmology*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 1–54. arXiv: 0705.0164 (cit. on pp. 47, 48, 52).
- [120] D. Boyanovsky, H. J. de Vega, and N. G. Sanchez. “CMB quadrupole suppression. II. The early fast roll stage”. *Physical Review D* 74.12 (2006), p. 123007 (cit. on pp. 47, 52, 55, 64, 80).
- [121] C. Destri, H. J. de Vega, and N. G. Sanchez. “CMB quadrupole depression produced by early fast-roll inflation: Monte Carlo Markov chains analysis of WMAP and SDSS data”. *Physical Review D* 78.2 (2008), p. 023013. arXiv: 0804.2387 (cit. on pp. 47, 48, 52, 55, 56, 80).
- [122] Grant N. Remmen and Sean M. Carroll. “How Many e -Folds Should We Expect from High-Scale Inflation?”. *Physical Review D* 90.6 (2014), p. 063517. arXiv: 1405.5538 (cit. on pp. 48, 56, 59, 65, 67, 80, 144).
- [123] Anthony Lasenby and Chris Doran. “Closed universes, de Sitter space, and inflation”. *Physical Review D* 71.6 (2005), p. 063502. arXiv: 0307311 [astro-ph] (cit. on pp. 48, 56, 59, 80, 97, 98, 128).
- [124] T. Banks and W. Fischler. *An upper bound on the number of e -foldings*. 2003. arXiv: 0307459 [astro-ph] (cit. on pp. 48, 56, 59, 65, 80).
- [125] Daniel Phillips, Andrew Scacco, and Andreas Albrecht. “Holographic bounds and finite inflation”. *Physical Review D* 91.4 (2015), p. 043513. arXiv: 1410.6065 (cit. on pp. 48, 56, 59, 65, 80).
- [126] Erandy Ramirez and Dominik J. Schwarz. “Predictions of just-enough inflation”. *Physical Review D* 85.10 (2012), p. 103516. arXiv: 1111.7131 (cit. on pp. 48, 56, 59, 67, 80).
- [127] Erandy Ramirez. “Low power on large scales in just-enough inflation models”. *Physical Review D* 85.10 (2012), p. 103517. arXiv: 1202.0698 (cit. on pp. 48, 56, 59, 67, 80).
- [128] Carlo R. Contaldi, Marco Peloso, Lev Kofman, and Andrei Linde. “Suppressing the lower multipoles in the CMB anisotropies”. *Journal of Cosmology and Astroparticle Physics* 2003.07 (2003), pp. 002–002. arXiv: 0303636 [astro-ph] (cit. on pp. 48, 56, 63, 80).
- [129] C. Destri, H. J. de Vega, and N. G. Sanchez. “Preinflationary and inflationary fast-roll eras and their signatures in the low CMB multipoles”. *Physical Review D* 81.6 (2010), p. 063520. arXiv: 0912.2994 (cit. on pp. 48, 55, 56, 80).
- [130] Andrew Scacco and Andreas Albrecht. “Transients in finite inflation”. *Physical Review D* 92.8 (2015), p. 083506. arXiv: 1503.04872 (cit. on pp. 48, 56, 68, 80, 119).
- [131] Planck Collaboration. “Planck 2015 results: XX. Constraints on inflation”. *Astronomy & Astrophysics* 594 (2016), A20. arXiv: 1502.02114 (cit. on pp. 48–50, 56, 58, 64, 70, 79, 103, 104, 107, 108, 120, 135).
- [132] Simony Santos da Costa, Micol Benetti, and Jailson Alcaniz. “A Bayesian analysis of inflationary primordial spectrum models using Planck data”. *Journal of Cosmology and Astroparticle Physics* 2018.03 (2018), pp. 004–004. arXiv: 1710.01613 (cit. on pp. 48, 56).
- [133] G. Hinshaw, D. N. Spergel, L. Verde, R. S. Hill, S. S. Meyer, C. Barnes, C. L. Bennett, M. Halpern, N. Jarosik, A. Kogut, E. Komatsu, M. Limon, L. Page, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. “First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: The Angular Power Spectrum”. *The Astrophysical Journal Supplement Series* 148.1 (2003), pp. 135–159. arXiv: 0302217 [astro-ph] (cit. on pp. 48, 56).

- [134] D. N. Spergel, L. Verde, H. V. Peiris, E. Komatsu, M. R. Nolta, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. “First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters”. *The Astrophysical Journal Supplement Series* 148.1 (2003), pp. 175–194. arXiv: [0302209 \[astro-ph\]](#) (cit. on pp. [48](#), [56](#), [79](#)).
- [135] Swagat S. Mishra, Varun Sahni, and Alexey V. Toporensky. “Initial conditions for inflation in an FRW universe”. *Physical Review D* 98.8 (2018), p. 083538. arXiv: [1801.04948](#) (cit. on pp. [48](#), [87](#)).
- [136] A.A. Starobinsky. “A new type of isotropic cosmological models without singularity”. *Physics Letters B* 91.1 (1980), pp. 99–102 (cit. on pp. [50](#), [58](#), [112](#)).
- [137] Gia Dvali and S.-H. Henry Tye. “Brane inflation”. *Physics Letters B* 450.1-3 (1999), pp. 72–82. arXiv: [9812483 \[hep-ph\]](#) (cit. on p. [50](#)).
- [138] M. Cicoli, C.P. Burgess, and F. Quevedo. “Fibre inflation: observable gravity waves from IIB string compactifications”. *Journal of Cosmology and Astroparticle Physics* 2009.03 (2009), pp. 013–013. arXiv: [0808.0691](#) (cit. on p. [50](#)).
- [139] A. Zee. “Broken-Symmetric Theory of Gravity”. *Physical Review Letters* 42.7 (1979), pp. 417–421 (cit. on p. [50](#)).
- [140] Frank S. Accetta, David J. Zoller, and Michael S. Turner. “Induced-gravity inflation”. *Physical Review D* 31.12 (1985), pp. 3046–3051 (cit. on p. [50](#)).
- [141] F. Lucchin, S. Matarrese, and M.D. Pollock. “Inflation with a non-minimally coupled scalar field”. *Physics Letters B* 167.2 (1986), pp. 163–168 (cit. on p. [50](#)).
- [142] David I. Kaiser. “Primordial spectral indices from generalized Einstein theories”. *Physical Review D* 52.8 (1995), pp. 4295–4306 (cit. on p. [50](#)).
- [143] Fedor Bezrukov and Mikhail Shaposhnikov. “The Standard Model Higgs boson as the inflaton”. *Physics Letters B* 659.3 (2008), pp. 703–706. arXiv: [0710.3755](#) (cit. on p. [50](#)).
- [144] A. Cerioni, F. Finelli, A. Tronconi, and G. Venturi. “Inflation and reheating in induced gravity”. *Physics Letters B* 681.5 (2009), pp. 383–386 (cit. on p. [50](#)).
- [145] Daniel Burns, Sotirios Karamitsos, and Apostolos Pilaftsis. “Frame-covariant formulation of inflation in scalar-curvature theories”. *Nuclear Physics B* 907 (2016), pp. 785–819. arXiv: [1603.03730](#) (cit. on p. [50](#)).
- [146] Nilay Bostan, Ömer Güleriyüz, and Vedat Nefer Şenoğuz. “Inflationary predictions of double-well, Coleman-Weinberg, and hilltop potentials with non-minimal coupling”. *Journal of Cosmology and Astroparticle Physics* 2018.05 (2018), pp. 046–046. arXiv: [1802.04160](#) (cit. on p. [50](#)).
- [147] Lotfi Boubekur and David. H. Lyth. “Hilltop inflation”. *Journal of Cosmology and Astroparticle Physics* 2005.07 (2005), pp. 010–010. arXiv: [0502047 \[hep-ph\]](#) (cit. on p. [50](#)).
- [148] Planck Collaboration. “Planck 2013 results: XXII. Constraints on inflation”. *Astronomy & Astrophysics* 571 (2014), A22. arXiv: [1303.5082](#) (cit. on pp. [55](#), [79](#), [103](#), [107](#), [108](#), [112](#), [120](#), [135](#), [137](#)).
- [149] V. F. Mukhanov, H. A. Feldman, and R. H. Brandenberger. “Theory of cosmological perturbations”. *Physics Reports* 215.5-6 (1992), pp. 203–333 (cit. on p. [55](#)).
- [150] D. Boyanovsky, H. J. de Vega, and N. G. Sanchez. “CMB quadrupole suppression. I. Initial conditions of inflationary perturbations”. *Physical Review D* 74.12 (2006), p. 123006 (cit. on pp. [55](#), [64](#), [80](#)).
- [151] G. Hinshaw, D. Larson, E. Komatsu, D. N. Spergel, C. L. Bennett, J. Dunkley, M. R. Nolta, M. Halpern, R. S. Hill, N. Odegard, L. Page, K. M. Smith, J. L. Weiland, B. Gold, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, E. Wollack, and E. L. Wright. “Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results”. *The Astrophysical Journal Supplement Series* 208.2 (2013), p. 19 (cit. on pp. [55](#), [79](#)).

- [152] H. V. Peiris, E. Komatsu, L. Verde, D. N. Spergel, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, E. Wollack, and E. L. Wright. “First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Implications For Inflation”. *The Astrophysical Journal Supplement Series* 148.1 (2003), pp. 213–231. arXiv: [0302225 \[astro-ph\]](#) (cit. on p. 56).
- [153] Michael J. Mortonson, Cora Dvorkin, Hiranya V. Peiris, and Wayne Hu. “CMB polarization features from inflation versus reionization”. *Physical Review D* 79.10 (2009), p. 103519. arXiv: [0903.4920](#) (cit. on pp. 56, 62).
- [154] Michele Cicoli, Sean Downes, Bhaskar Dutta, Francisco G. Pedro, and Alexander Westphal. “Just enough inflation: power spectrum modifications at large scales”. *Journal of Cosmology and Astroparticle Physics* 2014.12 (2014), pp. 030–030. arXiv: [1407.1048](#) (cit. on pp. 59, 67, 80).
- [155] Jennifer Adams, Bevan Cresswell, and Richard Easther. “Inflationary perturbations from a potential with a step”. *Physical Review D* 64.12 (2001), p. 123514. arXiv: [0102236 \[astro-ph\]](#) (cit. on p. 62).
- [156] Michael J Mortonson, Hiranya V Peiris, and Richard Easther. “Bayesian analysis of inflation: Parameter estimation for single field models”. *Physical Review D* 83.4 (2011), p. 043505. arXiv: [1007.4205](#) (cit. on p. 62).
- [157] C Armend riz-Pic n and Eugene A. Lim. “Vacuum choices and the predictions of inflation”. *Journal of Cosmology and Astroparticle Physics* 2003.12 (2003), pp. 006–006. arXiv: [0303103 \[hep-th\]](#) (cit. on pp. 63, 64).
- [158] W. J. Handley, A. N. Lasenby, and M. P. Hobson. “Novel quantum initial conditions for inflation”. *Physical Review D* 94.2 (2016), p. 024041. arXiv: [1607.04148](#) (cit. on pp. 63, 64, 106).
- [159] Andrew R Liddle and David H Lyth. “The cold dark matter density perturbation”. *Physics Reports* 231.1-2 (1993), pp. 1–105. arXiv: [9303019 \[astro-ph\]](#) (cit. on pp. 64, 67).
- [160] Andrew R Liddle and Samuel M Leach. “How long before the end of inflation were observable perturbations produced?” *Physical Review D* 68.10 (2003), p. 103503. arXiv: [0305263 \[astro-ph\]](#) (cit. on pp. 64, 67, 70, 96).
- [161] Peter Adshead and Richard Easther. “Constraining inflation”. *Journal of Cosmology and Astroparticle Physics* 2008.10 (2008), p. 047. arXiv: [0802.3898](#) (cit. on p. 67).
- [162] Peter Adshead, Richard Easther, Jonathan Pritchard, and Abraham Loeb. “Inflation and the scale dependent spectral index: prospects and strategies”. *Journal of Cosmology and Astroparticle Physics* 2011.02 (2011), pp. 021–021. arXiv: [1007.3748](#) (cit. on pp. 67, 101).
- [163] Michael S. Turner. “Coherent scalar-field oscillations in an expanding universe”. *Physical Review D* 28.6 (1983), pp. 1243–1247 (cit. on p. 68).
- [164] Antony Lewis, Anthony Challinor, and Anthony Lasenby. “Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models”. *The Astrophysical Journal* 538.2 (2000), pp. 473–476. arXiv: [9911177 \[astro-ph\]](#) (cit. on p. 68).
- [165] Antony Lewis. “CMB anisotropies from primordial inhomogeneous magnetic fields”. *Physical Review D* 70.4 (2004), p. 043011. arXiv: [0406096 \[astro-ph\]](#) (cit. on p. 68).
- [166] Anthony Challinor and Antony Lewis. “Lensed CMB power spectra from all-sky correlation functions”. *Physical Review D* 71.10 (2005), p. 103010. arXiv: [0502425 \[astro-ph\]](#) (cit. on p. 68).
- [167] Julien Lesgourgues. *The Cosmic Linear Anisotropy Solving System (CLASS) III: Comparison with CAMB for LambdaCDM*. 2011. arXiv: [1104.2934](#) (cit. on p. 68).
- [168] Thomas Tram and Julien Lesgourgues. “Optimal polarisation equations in FLRW universes”. *Journal of Cosmology and Astroparticle Physics* 2013.10 (2013), pp. 002–002. arXiv: [1305.3261](#) (cit. on pp. 68, 84).

- [169] Julien Lesgourgues and Thomas Tram. “Fast and accurate CMB computations in non-flat FLRW universes”. *Journal of Cosmology and Astroparticle Physics* 2014.09 (2014), pp. 032–032. arXiv: [1312.2697](#) (cit. on pp. [68](#), [80](#), [84](#)).
- [170] John Skilling. “Nested Sampling”. *AIP Conference Proceedings*. Ed. by R Fischer, R Preuss, and U.~V. Toussaint. Vol. 735. AIP, 2004, pp. 395–405 (cit. on p. [70](#)).
- [171] Laila Alabidi and David H Lyth. “Inflation models and observation”. *Journal of Cosmology and Astroparticle Physics* 2006.05 (2006), pp. 016–016. arXiv: [0510441 \[astro-ph\]](#) (cit. on p. [70](#)).
- [172] Will Handley. “fgivenx: A Python package for functional posterior plotting”. *Journal of Open Source Software* 3.28 (2018), p. 849. arXiv: [1908.01711](#) (cit. on p. [76](#)).
- [173] Planck Collaboration. “Planck 2018 results: X. Constraints on inflation”. *Astronomy & Astrophysics* 641 (2020), A10. arXiv: [1807.06211](#) (cit. on pp. [79](#), [103](#), [107](#), [108](#), [110](#), [111](#), [120](#), [135](#)).
- [174] Anna Ijjas, Paul J. Steinhardt, and Abraham Loeb. “Inflationary paradigm in trouble after Planck2013”. *Physics Letters B* 723.4-5 (2013), pp. 261–266. arXiv: [1304.2785](#) (cit. on p. [79](#)).
- [175] Jérôme Martin, Christophe Ringeval, and Vincent Vennin. “Encyclopædia Inflationaris”. *Physics of the Dark Universe* 5-6 (2014), pp. 75–235. arXiv: [1303.3787](#) (cit. on pp. [79](#), [108](#), [112](#)).
- [176] Andrei Linde. *Inflationary Cosmology after Planck 2013*. 2014. arXiv: [1402.0526](#) (cit. on p. [79](#)).
- [177] Anna Ijjas, Paul J. Steinhardt, and Abraham Loeb. “Inflationary schism”. *Physics Letters B* 736 (2014), pp. 142–146. arXiv: [1402.6980](#) (cit. on p. [79](#)).
- [178] Debika Chowdhury, Jérôme Martin, Christophe Ringeval, and Vincent Vennin. “Assessing the scientific status of inflation after Planck”. *Physical Review D* 100.8 (2019), p. 083537. arXiv: [1902.03951](#) (cit. on pp. [79](#), [87](#), [90](#)).
- [179] Jean-Philippe Uzan, Ulrich Kirchner, and George F. R. Ellis. “Wilkinson Microwave Anisotropy Probe data and the curvature of space”. *Monthly Notices of the Royal Astronomical Society* 344.4 (2003), pp. L65–L68 (cit. on p. [79](#)).
- [180] George Efstathiou and Steven Gratton. *A Detailed Description of the CamSpec Likelihood Pipeline and a Reanalysis of the Planck High Frequency Maps*. 2019. arXiv: [1910.00483](#) (cit. on p. [80](#)).
- [181] George Efstathiou and Steven Gratton. “The evidence for a spatially flat Universe”. *Monthly Notices of the Royal Astronomical Society: Letters* 496.1 (2020), pp. L91–L95. arXiv: [2002.06892](#) (cit. on pp. [80](#), [81](#)).
- [182] G. Efstathiou. “Is the low cosmic microwave background quadrupole a signature of spatial curvature?” *Monthly Notices of the Royal Astronomical Society* 343.4 (2003), pp. L95–L98 (cit. on pp. [80](#), [119](#), [145](#)).
- [183] Andreas Albrecht. “de Sitter equilibrium as a fundamental framework for cosmology”. *Journal of Physics: Conference Series* 174 (2009), p. 012006. arXiv: [0906.1047](#) (cit. on p. [80](#)).
- [184] Andreas Albrecht. “Cosmic Curvature from de Sitter Equilibrium Cosmology”. *Physical Review Letters* 107.15 (2011), p. 151102 (cit. on p. [80](#)).
- [185] Dominik J. Schwarz and Erandy Ramirez. *Just enough inflation*. 2009. arXiv: [0912.4348](#) (cit. on p. [80](#)).
- [186] Andrei Linde. “Fast-Roll Inflation”. *Journal of High Energy Physics* 2001.11 (2001), pp. 052–052. arXiv: [0110195 \[hep-th\]](#) (cit. on p. [80](#)).
- [187] Erandy Ramirez and Dominik J. Schwarz. “ ϕ^4 inflation is not excluded”. *Physical Review D* 80.2 (2009), p. 023525 (cit. on p. [80](#)).
- [188] Louis Lello and Daniel Boyanovsky. “Tensor to scalar ratio and large scale power suppression from pre-slow roll initial conditions”. *Journal of Cosmology and Astroparticle Physics* 2014.05 (2014), pp. 029–029. arXiv: [1312.4251](#) (cit. on p. [80](#)).

- [189] H. V. Ragavendra, Debika Chowdhury, and L. Sriramkumar. “Unique Contributions to the Scalar Bispectrum in ‘Just Enough Inflation’”. *Workshop on Frontiers in High Energy Physics 2019*. Ed. by Anjan Giri and Rukmani Mohanta. Springer, Singapore, 2020, pp. 39–47. arXiv: [1906.03942](#) (cit. on p. 80).
- [190] Martin White and Douglas Scott. “Why Not Consider Closed Universes?” *The Astrophysical Journal* 459 (1996), p. 415. arXiv: [9508157 \[astro-ph\]](#) (cit. on p. 80).
- [191] Béatrice Bonga, Brajesh Gupta, and Nelson Yokomizo. “Inflation in the closed FLRW model and the CMB”. *Journal of Cosmology and Astroparticle Physics* 2016.10 (2016), pp. 031–031. arXiv: [1605.07556](#) (cit. on p. 80).
- [192] Béatrice Bonga, Brajesh Gupta, and Nelson Yokomizo. “Tensor perturbations during inflation in a spatially closed Universe”. *Journal of Cosmology and Astroparticle Physics* 2017.05 (2017), pp. 021–021. arXiv: [1612.07281](#) (cit. on p. 80).
- [193] Will Handley. “Primordial power spectra for curved inflating universes”. *Physical Review D* 100.12 (2019), p. 123517. arXiv: [1907.08524](#) (cit. on pp. 80, 105, 106).
- [194] J.R. Dormand and P.J. Prince. “A family of embedded Runge-Kutta formulae”. *Journal of Computational and Applied Mathematics* 6.1 (1980), pp. 19–26 (cit. on p. 83).
- [195] F. J. Agocs, W. J. Handley, A. N. Lasenby, and M. P. Hobson. “Efficient method for solving highly oscillatory ordinary differential equations with applications to physical systems”. *Physical Review Research* 2.1 (2020), p. 013030. arXiv: [1906.01421](#) (cit. on pp. 83, 84, 106).
- [196] F. J. Agocs, M. P. Hobson, W. J. Handley, and A. N. Lasenby. *Dense output for highly oscillatory numerical solutions*. 2020. arXiv: [2007.05013](#) (cit. on p. 83).
- [197] W. J. Handley, A. N. Lasenby, and M. P. Hobson. *The Runge-Kutta-Wentzel-Kramers-Brillouin Method*. 2016. arXiv: [1612.02288](#) (cit. on p. 83).
- [198] M. P. Hobson, G. Efstathiou, and A. N. Lasenby. “The Friedmann–Robertson–Walker metric”. *General Relativity: An Introduction for Physicists*. Cambridge University Press, 2006. Chap. 14.6, p. 572 (cit. on p. 85).
- [199] Dalia S. Goldwirth and Tsvi Piran. “Initial conditions for inflation”. *Physics Reports* 214.4 (1992), pp. 223–292 (cit. on pp. 87, 90).
- [200] Diana Battefeld and Patrick Peter. “A critical review of classical bouncing cosmologies”. *Physics Reports* 571 (2015), pp. 1–66. arXiv: [1406.2790](#) (cit. on p. 89).
- [201] Dalia S. Goldwirth and Tsvi Piran. “Inhomogeneity and the onset of inflation”. *Physical Review Letters* 64.24 (1990), pp. 2852–2855 (cit. on p. 90).
- [202] Dalia S. Goldwirth. “Inhomogeneous initial conditions for inflation”. *Physical Review D* 43.10 (1991), pp. 3204–3213 (cit. on p. 90).
- [203] William E. East, Matthew Kleban, Andrei Linde, and Leonardo Senatore. “Beginning inflation in an inhomogeneous universe”. *Journal of Cosmology and Astroparticle Physics* 2016.09 (2016), pp. 010–010. arXiv: [1511.05143](#) (cit. on p. 90).
- [204] Katy Clough, Eugene A. Lim, Brandon S. DiNunno, Willy Fischler, Raphael Flauger, and Sonia Paban. “Robustness of inflation to inhomogeneous initial conditions”. *Journal of Cosmology and Astroparticle Physics* 2017.09 (2017), pp. 025–025. arXiv: [1608.04408](#) (cit. on p. 90).
- [205] Andrew R. Liddle, Paul Parsons, and John D. Barrow. “Formalizing the slow-roll approximation in inflation”. *Physical Review D* 50.12 (1994), pp. 7222–7232. arXiv: [9408015 \[astro-ph\]](#) (cit. on p. 96).
- [206] Matthew Civiiletti and Brandon Delacruz. “Natural inflation with natural number of e-foldings”. *Physical Review D* 101.4 (2020), p. 043534. arXiv: [2004.05238](#) (cit. on p. 96).
- [207] Jingwei Liu and Fulvio Melia. “Viability of slow-roll inflation in light of the non-zero k_{\min} measured in the cosmic microwave background power spectrum”. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 476.2239 (2020), p. 20200364. arXiv: [2006.02510](#) (cit. on p. 96).

- [208] Richard Easther and Hiranya V. Peiris. “Bayesian analysis of inflation. II. Model selection and constraints on reheating”. *Physical Review D* 85.10 (2012), p. 103533. arXiv: [1112.0326](#) (cit. on pp. [101](#), [135](#)).
- [209] James E. Lidsey, Andrew R. Liddle, Edward W. Kolb, Edmund J. Copeland, Tiago Barreiro, and Mark Abney. “Reconstructing the inflaton potential—an overview”. *Reviews of Modern Physics* 69.2 (1997), pp. 373–410. arXiv: [9508078](#) [[astro-ph](#)] (cit. on p. [104](#)).
- [210] David Hilary Lyth and Andrew R. Liddle. “Generating ζ at horizon exit”. *The Primordial Density Perturbation: Cosmology, Inflation and the Origin of Structure*. Cambridge University Press, 2009. Chap. 25, pp. 404–423 (cit. on p. [104](#)).
- [211] T. S. Bunch and P. C. W. Davies. “Quantum field theory in de Sitter space: renormalization by point-splitting”. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 360.1700 (1978), pp. 117–134 (cit. on p. [106](#)).
- [212] F. J. Agocs, L. T. Hergt, W. J. Handley, A. N. Lasenby, and M. P. Hobson. “Quantum initial conditions for inflation and canonical invariance”. *Physical Review D* 102.2 (2020), p. 023507. arXiv: [2002.07042](#) (cit. on p. [106](#)).
- [213] David H. Lyth. “What Would We Learn by Detecting a Gravitational Wave Signal in the Cosmic Microwave Background Anisotropy?” *Physical Review Letters* 78.10 (1997), pp. 1861–1863. arXiv: [9606387](#) [[hep-ph](#)] (cit. on p. [107](#)).
- [214] A. D. Linde. “Chaotic inflation”. *Physics Letters B* 129.3-4 (1983), pp. 177–181 (cit. on p. [108](#)).
- [215] Alexander Vilenkin. *Eternal inflation and chaotic terminology*. 2004. arXiv: [0409055](#) [[gr-qc](#)] (cit. on p. [108](#)).
- [216] Eva Silverstein and Alexander Westphal. “Monodromy in the CMB: Gravity waves and string inflation”. *Physical Review D* 78.10 (2008), p. 106003. arXiv: [0803.3085](#) (cit. on p. [110](#)).
- [217] Liam McAllister, Eva Silverstein, and Alexander Westphal. “Gravity waves and linear inflation from axion monodromy”. *Physical Review D* 82.4 (2010), p. 046003. arXiv: [0808.0706](#) (cit. on p. [110](#)).
- [218] Katherine Freese, Joshua A. Frieman, and Angela V. Olinto. “Natural inflation with pseudo Nambu-Goldstone bosons”. *Physical Review Letters* 65.26 (1990), pp. 3233–3236 (cit. on p. [110](#)).
- [219] Jérôme Martin, Christophe Ringeval, and Roberto Trotta. “Hunting down the best model of inflation with Bayesian evidence”. *Physical Review D* 83.6 (2011), p. 063524. arXiv: [1009.4157](#) (cit. on p. [135](#)).

