


# Primordial evolution of cosmological perturbations:

Theory and computation



**Fruzsina Julia Agocs**

Cavendish Astrophysics Group  
University of Cambridge  
Gonville & Caius College

This thesis is submitted for the degree of  
*Doctor of Philosophy*

September 2021



# 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.

I further state that no substantial part of my thesis 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.

Fruzsina Julia Agocs,  
September 2021



# Abstract

## Primordial evolution of cosmological perturbations: Theory and computation

Fruzsina Julia Agocs

This thesis discusses results in theoretical cosmology related to the primordial evolution of cosmological perturbations with the help of numerical analysis. Primordial fluctuations are thought to be the seeds of large-scale structure, therefore modelling their initial conditions and evolution is key to understanding how structure forms. Recent tensions in the inferred values of cosmological parameters have raised interest in models that deviate from the currently accepted standard model of cosmology, the  $\Lambda$  cold dark matter model, but these alternative models can form the computational bottleneck of cosmological inference. I first present a numerical method (and associated open-source software) for solving a class of highly oscillatory ordinary differential equations efficiently, which speeds up the forward-modelling step of cosmological inference significantly by enabling fast numerical evolution of primordial fluctuations. I discuss other uses of the numerical routine in the physical sciences and report on its latest application to more accurately constrain the universe's spatial curvature.

The evolution of primordial perturbations cannot be fully determined without initial conditions. I therefore inspect popular methods for setting initial conditions from the perspective of their behaviour under canonical transformations, and find only one set of initial conditions invariant under such transformations. I demonstrate the possible observational consequences of canonical non-invariance of the initial conditions and argue that an invariant set should be used in models that retain memory of the initial conditions. I discuss preliminary investigations into generalising the canonically invariant initial conditions to universes with non-zero spatial curvature before concluding with a summary of future research avenues I view as worthy of exploration.

**Chapters 3** and **6** of this thesis are based on the publications titled *Efficient method for solving highly oscillatory ordinary differential equations with applications to physical systems*, in Physical Review Research [1], and *(py)oscode: fast solutions of oscillatory ODEs*, published in The Journal of Open Source Software [2]. **Chapter 4** is based on work published as a pre-print under the title *Dense output for highly oscillatory numerical solutions* on the arXiv [3]. Finally, **Chapter 7** is based on the manuscript *Quantum initial conditions for inflation and canonical invariance* published in Physical Review D [4]. I am the sole or leading author and contributor to all of the above publications.



*Azoknak, akik engem tanítottak*  
*To all those who taught me*



# Acknowledgements

First of all, I would like to sincerely thank my supervisors and advisers, Anthony Lasenby, Mike Hobson, and Will Handley. I am grateful for the invaluable insights, mentorship and patience you provided throughout these four years and I feel honoured to have had the opportunity to learn from you.

I would like to extend my gratitude to my collaborators: Lukas Hergt, Julien Lesgourgues, Nils Schöneberg, Imre Fekete, and András Molnár, for many fruitful discussions and for broadening my knowledge in statistics, cosmology, and numerical analysis.

I had the exceptional opportunity to spend 6 months at the British Antarctic Survey to learn about the Earth's climate, and to further my computing (and in particular, machine learning) skills. I thank my BAS supervisor, Scott Hosking, and members of the BAS AI lab for making my placement a productive and pleasant one. I would like to further thank the STFC and CDT for making the placement (and the rest of my PhD) possible through their sponsorship and organising.

Studying at Cambridge has had a profound impact on me. I am indebted to my undergraduate supervisors and tutors who guided me through my studies: John Ellis, Laura McMahon, Will Handley, Derek Barnes, Rafi Blumenfeld, and Ryan MacDonald. My thanks goes to Ranjan Vasudevan, who first introduced me to research in astronomy. I am grateful to my excellent teachers from Áldás, Eötvös, Fazekas, and AISB: Bognár Anikó, Elekes Sándor, Dancsó Éva, Rubóczky György, Albert Attila, Chris Mortenson and Patty Devlin.

I spent my entire adult life (so far) in Cambridge, and during that time I could not have asked for a better home than Gonville & Caius college. I thank them for their support and for letting me supervise their students.

I am lucky to have had the support of an amazing group of friends during my PhD. Emma, Adam, and Josh: thank you for the many boardgames, sports, and for practically adopting me. My IoA and Cavendish friends, Lukas, Sid, Sophie, Amy, Anjali, Jasleen, Nick, Dan, Peter, and Roger, for your company at lunches and in the office, and the unforgettable Lake District trip. Special thanks go to Lukas for keeping me company during the pandemic and teaming up with me to cook Christmas dinner; to Sophie for her help with the citizenship and postdoc applications; and to Sid for the many burritos and inviting me to your Christmas party (including all the discussions that followed). The original IoA PhD team: Douglas, Sophie R,

and Michael, thank you for letting me see what it's like to be a PhD student, and for introducing me to climbing. Austen, Dom, Augustin, and Tim: thank you for being great college (and house-) mates and D&D party members.

Completing my undergraduate degree and PhD would not have been possible without the unwavering support of my parents. Your sacrifices have allowed me to do what I enjoy doing. I cannot thank you enough for your unconditional love. To my sister, Nina: thank you for always being there for me, and always making sure I'm not alone (it is no overstatement that I developed social skills thanks to your efforts in kindergarten).

Finally, I would like to say thank you to my other half, Ant. You kept me sane during the application and thesis writing process with your love, support, and incredible patience. I would not have dared to dream about having a partner like you.

# Table of contents

Notation & conventions	xv
List of figures	xvii
List of tables	xix
<b>1 Outline</b>	<b>1</b>
1.1 Motivation and aims . . . . .	1
1.2 How this thesis is structured . . . . .	3
<b>2 Integration of Ordinary Differential Equations</b>	<b>5</b>
2.1 Ordinary Differential Equations . . . . .	5
2.1.1 Definitions . . . . .	5
2.1.2 Oscillatory ODEs . . . . .	6
2.2 Numerical methods for the solution of ODEs . . . . .	8
2.2.1 Euler’s method and its error properties . . . . .	8
2.2.2 Generalisations of Euler’s method . . . . .	10
2.2.3 Runge–Kutta methods . . . . .	12
2.2.4 Specialised methods for oscillatory ODEs . . . . .	16
2.3 Requirements of a practical numerical routine . . . . .	22
2.3.1 Stability . . . . .	22
2.3.2 Adaptive stepsize control . . . . .	24
2.3.3 Dense output . . . . .	30
<b>3 An efficient method for solving highly oscillatory ODEs</b>	<b>33</b>
3.1 Introduction . . . . .	34
3.2 Computational strategy . . . . .	35
3.2.1 The Wentzel–Kramers–Brillouin approximation . . . . .	36
3.3 Implementation . . . . .	37
3.3.1 Numerical integration and differentiation . . . . .	37
3.3.2 Explicit Runge–Kutta formulae based on Gauss–Lobatto stencil points . . . . .	39

3.3.3	Defining $\omega(t)$ and $\gamma(t)$ . . . . .	40
3.3.4	Estimating the error in RK and WKB steps . . . . .	41
3.3.5	Stepping procedure . . . . .	41
3.3.6	Algorithm summary . . . . .	44
3.4	Example applications . . . . .	45
3.4.1	Airy equation . . . . .	45
3.4.2	Burst equation . . . . .	45
3.4.3	Schrödinger equation . . . . .	49
3.5	Limitations . . . . .	50
3.6	Conclusions . . . . .	52
<b>4</b>	<b>Dense output</b> . . . . .	<b>55</b>
4.1	Introduction . . . . .	56
4.2	Dense output from the WKB expansion . . . . .	57
4.2.1	Gaussian quadrature . . . . .	57
4.2.2	Dense output from Gauss–Lobatto integration . . . . .	58
4.2.3	Error bound on dense output from Gauss–Lobatto integration . . . . .	59
4.3	Dense output from a Runge-Kutta method based on Gauss–Lobatto nodes . . . . .	60
4.4	Examples . . . . .	65
4.5	Conclusions . . . . .	67
<b>5</b>	<b>Cosmology of primordial perturbations</b> . . . . .	<b>69</b>
5.1	The universe at zeroth order . . . . .	69
5.1.1	Gravity . . . . .	70
5.1.2	Dynamics of a homogeneous, isotropic, expanding universe . . . . .	71
5.2	Observations . . . . .	75
5.2.1	The Cosmic Microwave Background . . . . .	75
5.2.2	Other observational probes . . . . .	78
5.2.3	Happy accidents . . . . .	79
5.3	Cosmic inflation . . . . .	80
5.3.1	Single-field inflation . . . . .	80
5.4	Perturbations . . . . .	84
5.4.1	Gauge choice . . . . .	84
5.4.2	Perturbed matter and metric . . . . .	85
5.4.3	Subhorizon evolution: the Mukhanov–Sasaki equation . . . . .	87
5.4.4	Quantisation and initial conditions . . . . .	89
5.4.5	Connection to CMB anisotropies . . . . .	89
5.5	Inference in cosmology . . . . .	91
5.5.1	$\Lambda$ CDM . . . . .	94

5.5.2	Tensions . . . . .	97
<b>6</b>	<b>Oscode in cosmology</b>	<b>99</b>
6.1	The Mukhanov–Sasaki equation . . . . .	100
6.2	Comparison with BINGO . . . . .	101
6.3	A model using kinetic dominance . . . . .	104
6.4	A closed universe model . . . . .	106
6.5	Conclusions . . . . .	109
<b>7</b>	<b>Quantum initial conditions for primordial perturbations</b>	<b>111</b>
7.1	Introduction . . . . .	112
7.2	Background . . . . .	114
7.2.1	Canonical transformations . . . . .	114
7.2.2	Dynamics of primordial perturbations . . . . .	116
7.2.3	Vacuum choices . . . . .	118
7.3	Methods . . . . .	122
7.3.1	Field redefinition . . . . .	122
7.3.2	Surface terms . . . . .	123
7.4	Results . . . . .	123
7.4.1	Initial conditions . . . . .	123
7.4.2	Primordial power spectra . . . . .	129
7.4.3	Cosmic Microwave Background . . . . .	136
7.5	Conclusions . . . . .	137
	Appendix 7.A Hamiltonian diagonalisation under field redefinitions and addition of surface terms . . . . .	139
	7.A.1 Field redefinition . . . . .	140
	7.A.2 Surface terms . . . . .	141
	Appendix 7.B Minimising the renormalised stress–energy tensor under field redefinitions . . . . .	142
<b>8</b>	<b>Conclusions and future outlook</b>	<b>145</b>
8.1	Computational tools . . . . .	145
8.1.1	Future work . . . . .	146
8.2	Theoretical results . . . . .	151
8.2.1	Future work . . . . .	152
	<b>References</b>	<b>155</b>



# Notation & conventions

## Notation

$t$  (cosmic) time

$\eta$  conformal time

$\dot{f}$   $\frac{df}{dt}$ , differentiation with respect to time

$f'$  differentiation with respect to a function's argument, but in a cosmological context it may mean  $\frac{df}{d\eta}$ , differentiation with respect to conformal time

- Differential operators  $dx, dy, \dots$ , Euler's constant  $e$ , and the imaginary unit  $i$  all appear italicised.
- A variable printed in boldface, e.g.  $\mathbf{v}$ , denotes a vector or a matrix, as specified in the text. Vectors and tensors may, however, not always appear in bold.

## Abbreviations

$\Lambda$ CDM cold dark matter model with a cosmological constant

CMB cosmic microwave background

FLRW Friedmann–Lemaître–Robertson–Walker

KD kinetic dominance

HD Hamiltonian diagonalisation

ODE ordinary differential equation

PPS primordial power spectrum

SR slow-roll

RK Runge–Kutta

RST renormalised stress–energy tensor

WKB Wentzel–Kramers–Brillouin

## Mathematical and physical conventions

- Fourier transforms and synthesis are defined as

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

- Spacetime indices are denoted with Greek letters  $\mu, \nu, \sigma, \dots$ . The index 0 is reserved for the temporal component, and spatial components are denoted with Latin indices  $i, j, k, \dots$ .
- In a cosmological context, an upper index denotes components of a contravariant vector (a vector), and a lower index denotes those of a covariant vector (a covector). Under a change of basis, they transform accordingly.
- This thesis makes use of the Einstein summation convention: if an index is repeated and appears once in the upper and once in the lower position, it is summed over.
- Throughout this thesis, the metric has a positive signature  $(+, -, -, -)$ .
- For simplicity, formulae are written in natural units,

$$c = \hbar = G = k_B = 1,$$

and unless stated otherwise, the Planck mass is retained,

$$m_P = \frac{1}{\sqrt{8\pi G}} = \frac{1}{\sqrt{8\pi}}.$$

# List of figures

2.1	Classification of numerical ODE solvers . . . . .	11
2.2	Instability in the numerical solution of a stiff ODE. . . . .	23
2.3	Adaptive stepsize control visualised as a feedback loop. . . . .	27
2.4	Numerical Van der Pol equation with classical and PI3333 stepsize control . . . .	30
3.1	Error progression in RK and WKB steps and <code>oscode</code> 's stepsize-update algorithm	43
3.2	<code>oscode</code> 's numerical solution and error when solving the Airy equation . . . . .	46
3.3	<code>oscode</code> 's numerical solution and error when solving the burst equation . . . . .	47
3.4	<code>oscode</code> 's number of oscillations traversed per step when solving the burst equation	47
3.5	<code>oscode</code> 's relative error at varying tolerance thresholds when solving the burst equation . . . . .	48
3.6	<code>oscode</code> 's runtime as a function of total oscillations in the burst equation . . . .	48
3.7	<code>oscode</code> 's runtime as a function of tolerance at varying number of oscillations in the burst equation . . . . .	49
3.8	Breakdown of <code>oscode</code> 's step-types when solving the burst equation against the total number of oscillations . . . . .	49
3.9	<code>oscode</code> 's numerical solution of the quantum harmonic oscillator . . . . .	51
4.1	Log-residuals of partial Gauss–Lobatto quadrature . . . . .	61
4.2	Residuals of partial Gauss–Lobatto quadrature . . . . .	61
4.3	Dense output from <code>oscode</code> solving the Airy equation . . . . .	66
4.4	Residuals of dense output from <code>oscode</code> solving the Airy equation . . . . .	66
4.5	Dense output from <code>oscode</code> solving the burst equation . . . . .	67
5.1	Evolution of the FLRW background . . . . .	83
5.2	Seeding of fluctuations in inflation . . . . .	85
5.3	Geometric meaning of a gauge choice . . . . .	86
5.4	Schematic evolution of a single perturbation mode relative to the comoving Hubble sphere . . . . .	91

6.1	Reference primordial power spectra computed with <code>oscode</code> and BINGO for comparing runtimes . . . . .	102
6.2	BINGO v. <code>oscode</code> runtime ratio as a function of the comoving wavevector . . . . .	103
6.3	<code>oscode</code> 's runtime variation with comoving wavenumber . . . . .	103
6.4	Example numerical evolution of a single primordial perturbation mode computed with BINGO and <code>oscode</code> . . . . .	104
6.5	An example primordial power spectrum of a kinetically dominated universe, as computed by <code>oscode</code> . . . . .	105
6.6	Numerically computed primordial power spectra of closed universes with varying degrees of initial curvature . . . . .	108
6.7	Numerical and analytic PPS comparison for closed universes . . . . .	109
6.8	Statistical results concerning curved universes and finite inflation from [159, 160] . . . . .	110
7.1	Schematic evolution of a single perturbation mode relative to the comoving Hubble horizon . . . . .	130
7.2	Comparison of PPS and CMB power spectra resulting from HD and RST initial conditions . . . . .	132
7.3	Variation of features in the PPS with changing parametrisation of HD initial conditions . . . . .	133
7.4	Magnified central region of Figs. 7.3b and 7.3d . . . . .	135
7.5	PPS and CMB TT power spectra arising from selected parametrisations of HD initial conditions . . . . .	136
7.6	CMB TE and EE spectra arising from selected parametrisations of HD initial conditions . . . . .	137
7.7	Low- $\ell$ region of the CMB angular power spectra from Figs. 7.5 and 7.6 . . . . .	138
7.8	Prior and posterior of the parameters appearing in HD initial conditions from Planck 2018 lowT data . . . . .	139

# List of tables

2.1	Butcher tableau for a generic explicit Runge–Kutta method . . . . .	13
2.2	Attainable order of explicit Runge–Kutta methods . . . . .	14
3.1	4-stage, 4th order Runge–Kutta formula used in <code>oscode</code> . . . . .	39
3.2	6-stage, 5th order Runge–Kutta formula used in <code>oscode</code> . . . . .	40
3.3	Reference energy eigenvalues of the quantum harmonic oscillator with quartic anharmonicity . . . . .	52
4.1	6-stage, 5th order Runge–Kutta formula used in <code>oscode</code> . . . . .	63
4.2	Modified $b$ -coefficients in <code>oscode</code> 's Runge–Kutta formula for obtaining a 4th order estimate mid-step . . . . .	63
4.3	$P$ -matrix for generating dense output from <code>oscode</code> 's Runge–Kutta method . . . . .	65
5.1	Quantities appearing in the FLRW metric . . . . .	72
5.2	Dynamics of a FLRW universe dominated by a given fluid . . . . .	74
7.1	Notation used in <a href="#">Chapter 7</a> . . . . .	113
7.2	Fiducial values of Planck 2018 cosmological parameters . . . . .	134



# Chapter 1

## Outline

### 1.1 Motivation and aims

At the time of writing this thesis, particle accelerators are able to reach collision energies of 14 TeV, which match the characteristic energy scale of the universe at  $10^{-10}$  seconds after the Big Bang. Therefore from this point onwards, the conditions in the Universe can be reproduced in a laboratory setting and the fundamental laws on high-energy physics can be experimentally tested. While  $10^{-10}$  seconds is impressively close to the moment of the Big Bang, in order to understand the origins of cosmic structure we must “go back” even further in time – to  $10^{-34}$  seconds, when quantum-scale fluctuations were first thought to be generated. In lieu of experiments and direct observation, physicists must rely on numerical simulations to probe the early Universe indirectly via inference. A numerical simulation will typically start from some carefully chosen initial conditions set at early times and run until the output of the simulation can be compared with observations. Since our position as observers is fixed in time (at least relative to cosmic timescales) to now,  $\approx 13.7$  billion years after the Big Bang, the numerical simulations have to follow the evolution of the observables for an extremely long time. While for the majority of this time the observables may barely change (due to the growing distances between astrophysical objects brought on by the expansion of the Universe), the simulation must still account for a wide range of physical processes over its runtime.

It is immediately obvious that in order to infer the state of the universe at early times correctly, the models encoded in the numerical simulation have to be chosen carefully: the right set of physical interactions have to be selected which dominate over all others. It is less apparent – and perhaps a little overlooked – that choosing the appropriate numerical methods to evaluate models is equally important. In inference, the same numerical simulation often has to be run millions of times with a slightly different set of model parameters to paint a detailed enough picture of their likelihood. Even for the simplest models, the breadth of cosmological simulations forces this task to be performed on high-performance computing clusters, where it

may take a few days<sup>1</sup> to complete. At this level of computational complexity, using efficient numerical methods clearly matters, since not only will we see the results sooner, the carbon footprint of our research is reduced due to the smaller electricity consumption of the cluster. Importantly, more efficient numerical methods also allow for a wider variety of models to be tested, through extending the limit on (computational) complexity and allowing for some simplifying assumptions to be relaxed.

The current standard model of cosmology,  $\Lambda$ CDM (where  $\Lambda$  denotes a cosmological constant and the rest of the letters stand for cold dark matter) predicts the large-scale statistics of the Universe with spectacular accuracy and with only a few free parameters. Recent tensions in the inferred values of these cosmological parameters have, however, raised interest in models that deviate from this standard model [5–8], which has been largely unchanged for the last two decades. These alternative models are often more computationally demanding to evaluate, and unlike  $\Lambda$ CDM, may necessitate the numerical evolution of primordial perturbations. One example is models involving non-zero spatial curvature, which may be favoured by data from the Planck satellite [9–11]. The primordial evolution of cosmological perturbations is described by an ordinary differential equation whose solution is highly oscillatory for the majority of the interval of interest, but changes slowly in some regions. There are currently no numerical routines available from mainstream, open-source numerical libraries that are equipped to solve this type of equation efficiently. While there exist some specialised solvers for special cases of highly oscillatory differential equations, they tend not to be general enough to tolerate the frequency of oscillations varying over time, are only valid in the highly oscillatory regime, or (if they can do both) have not been implemented for public use. These factors motivated the development of `oscocode` [1–3], an open-source numerical routine for solving a class of highly oscillatory ordinary differential equations (including that which describes the primordial evolution of cosmological perturbations), which I present in this thesis, alongside examples of its successful and potential applications in physics.

The predictions of cosmological models are only as robust as the assumptions underlying them. An example is the set of quantum initial conditions set for the perturbations at primordial times, which influence the predictions of models with reduced amounts of inflation such as the closed universe models from the paragraph above. These initial conditions are set using principles from quantum field theory on a static spacetime, but these do not translate well to the inflating early universe. I discuss a key weakness in the traditional methods for setting initial conditions which arises from their behaviour under *canonical transformations*, a set of transformations used in classical mechanics to find a better parametrisation of a system without changing its physics. I found [4] that while most standard methods produce initial conditions which change under canonical transformations, one method [12] produces invariant

---

<sup>1</sup>or up to a couple of months for more intricate models

initial conditions and therefore unambiguous large-scale statistics predictions. I argue for using this set of initial conditions above others.

In summary, this thesis discusses the theoretical framework and computational tools I developed for more robust and efficient evaluation of cosmological models, and details examples of their applications in research.

## 1.2 How this thesis is structured

Logic dictates that a description of the (numerical) methodology precedes its applications. I therefore choose to start with [Chapter 2](#) introducing concepts in numerical analysis before a discussion of `oscode`'s algorithm in [Chapter 3](#). Some of `oscode`'s applications can be understood without familiarity with the physics of primordial perturbations; these are also detailed in [Chapter 3](#). Given that one of the main uses of `oscode` in cosmology is closely tied to the primordial evolution of perturbations, I opted to delay outlining `oscode`'s applications in cosmology until [Chapter 6](#), after an introduction of the relevant physics in [Chapter 5](#). [Chapter 7](#) then reports my findings in connection with the quantum initial conditions for perturbations for spatially flat universes. I draw conclusions and discuss future research ideas in [Chapter 8](#).



## Chapter 2

# Integration of Ordinary Differential Equations

A matematikában az ember nem megérti a dolgokat, hanem megszokja.

In mathematics, one does not come to understand things, but gets used to them.

---

*Neumann János*

In this chapter, I introduce key definitions and concepts from numerical analysis. To put the numerical work in [Chapters 3](#) and [4](#) into context, I also review past and ongoing efforts to construct numerical solvers for oscillatory ordinary differential equations and identify some criteria a practically useful numerical method should fulfil.

## 2.1 Ordinary Differential Equations

### 2.1.1 Definitions

We first review some definitions and theorems that will be relevant to our discussion of ordinary differential equations.

Let  $X$  be an  $N$ -dimensional (real) vector space, and  $I$  an interval of  $\mathbb{R}$ . Given a function  $\mathbf{f} : I \times X \mapsto X$ , a first-order ordinary differential equation (ODE) system is defined as

$$\mathbf{y}'(x) = \mathbf{f}(x, \mathbf{y}(x)), \tag{2.1}$$

where prime refers to differentiation with respect to  $x$ . The solution of [Eq. \(2.1\)](#) is a function  $\mathbf{y} \in C^1(I, X)$ , the set of continuously differentiable  $X$ -valued functions defined on  $I$ , such that [Eq. \(2.1\)](#) is satisfied for  $\forall x \in I$  and any subsidiary conditions are satisfied. A possible

generalisation of Eq. (2.1) is an  $n$ th-order ODE, defined as

$$\mathbf{y}^{(n)}(x) = \mathbf{f}(x, \mathbf{y}(x), \mathbf{y}'(x), \dots, \mathbf{y}^{(n-1)}(x)), \quad (2.2)$$

where  $\mathbf{f} : I \times X^n \mapsto X$  is given and the solution is  $\mathbf{y} \in C^n(I, X)$ .

Subsidiary conditions are specified because the solution to Eq. (2.1) is, in general, not unique. We distinguish between *initial value problems* and *boundary value problems* based on whether the subsidiary condition is the value of  $\mathbf{y}$  and all its derivatives given at the same point (at the start of the interval), or at the integration boundaries. In general, the subsidiary conditions can be nonlinear algebraic constraints on the values of  $\mathbf{y}$  and its derivatives.

If  $\mathbf{f}$  is of the form

$$\mathbf{f}(x, \mathbf{y}) = \mathbf{A}(x)\mathbf{y} + \mathbf{b}(x), \quad (2.3)$$

with  $\mathbf{A}(x)$  being a matrix-valued function and  $\mathbf{b}$  a vector-valued function, the differential equation is *linear*. If  $\mathbf{b}(x)$  is the zero function, the ODE is also *homogeneous*. Finally, a system for which  $\mathbf{f}$  does not depend on the independent variable  $x$  explicitly, is *autonomous*. In the following chapters we will restrict our attention to linear equations, since the equations of interest for this thesis arise from cosmological perturbation theory and are thus linearised. When discussing Runge–Kutta methods, we will often only consider first-order autonomous systems. This comes without a loss of generality, because the general  $n$ th order system Eq. (2.2) can always be converted to a first-order, autonomous one:

$$\mathbf{u}' = \mathbf{g}(\mathbf{u}), \quad (2.4)$$

$$\begin{bmatrix} x \\ \mathbf{y} \\ \mathbf{y}' \\ \dots \\ \mathbf{y}^{(n-1)} \end{bmatrix}' = \begin{bmatrix} 1 \\ \mathbf{y}' \\ \mathbf{y}'' \\ \dots \\ \mathbf{f}(x, \mathbf{y}, \mathbf{y}', \dots, \mathbf{y}^{(n-1)}) \end{bmatrix}. \quad (2.5)$$

### 2.1.2 Oscillatory ODEs

Of particular interest to physics is the numerical solution of oscillatory ODEs, but neither the general form of an oscillatory ODE nor what constitutes an adequate solution can be defined precisely, as both depend on the field of application. In general, the solution of an oscillatory system is characterised by (at least one) fast component of the solution varying periodically about a slow component. This behaviour can arise in a range of differently structured equation systems: they may be first- or second-order ODEs, differential algebraic systems (DAEs) that include constraints, linear or nonlinear, homogeneous or inhomogeneous, or partial differential

equations (PDEs). For instance, the system

$$y'' + \sin(y) = \cos(\omega x), \quad x \geq 0, \quad y(0) = y_0, \quad y'(0) = y'_0 \quad (2.6)$$

has a solution comprising of a slow overall oscillation (the solution to the homogeneous part of the problem,  $y'' + \sin(y) = 0$ ), and forcing oscillations of frequency  $\omega$  winding around it. The amplitude of the secondary oscillations decreases with  $\omega$ . The time-independent, one-dimensional Schrödinger equation is of the form

$$y'' + g(x)y = 0, \quad (2.7)$$

where  $g(x)$  is large and usually slowly varying. The solution oscillates with a large, slowly modulated amplitude. Finally, the system

$$\begin{bmatrix} x \\ y \\ u \\ v \end{bmatrix}' = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -\lambda & 0 & 0 & 0 \\ 0 & -\lambda + 1/y & 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ u \\ v \end{bmatrix}, \quad \varepsilon^2 \lambda = \frac{\sqrt{x^2 + y^2} - 1}{\sqrt{x^2 + y^2}} \quad (2.8)$$

describes a stiff pendulum with unit mass, length, and gravity, a (large) spring constant  $1/\varepsilon^2$ , coordinates  $(x, y)$  and corresponding velocity components  $(u, v)$ . The solution is superficially similar to that of [Eq. \(2.6\)](#) (low-amplitude oscillations superimposed on a slow solution), but while the forced pendulum is described by an inhomogeneous, second-order ODE, the stiff pendulum's equation of motion is a DAE system with four first-order, linear differential equations and an algebraic constraint.

What constitutes an adequate numerical solution may also be situation-dependent. In some cases, it is sufficient to find the slowly-varying solution, or the envelope of high-frequency oscillations. Other times, the phase of oscillations may also carry important information and has to be reconstructable from the numerical solution. In molecular or orbital dynamics, the emphasis is often on conserving a quantity (such as energy or angular momentum).

In the following chapters, we will focus on a particular type of oscillatory ODE of the form

$$y'' + 2\gamma(x)y' + \omega^2(x)y = 0, \quad (2.9)$$

which may be recast as a system of first-order, autonomous differential equations. The numerical methods introduced will be able to recover the full solution (not just the amplitude). They will prioritise efficiency by using stepsizes which are large compared to the timescales of the fast solution, and do this by exploiting the global properties of the problem.

## 2.2 Numerical methods for the solution of ODEs

### 2.2.1 Euler's method and its error properties

Let us approximate the solution of an initial value problem of the form<sup>1</sup>

$$y' = f(y), \quad y(x_0) = y_0 \in X \quad (2.10)$$

(which includes non-autonomous, higher-order ODE systems) by making a sequence of values of the independent variable starting from  $x_0$ ,  $(x_0, x_1, x_2, \dots)$ , which we will call steps. The associated stepsizes,  $x_1 - x_0, x_2 - x_1, \dots$  are denoted  $h_1, h_2, \dots$ . The approximate numerical solution takes the values  $y_0, y_1, y_2, \dots$  at these points (with  $y_0$  known exactly). Euler's method [13] is summarised as

$$y_n = y_{n-1} + h_n f(y_{n-1}), \quad n = 1, 2, \dots, \quad (2.11)$$

meaning that the method approximates the solution at the end of a timestep as if  $y_{n-1}$  were the exact solution at  $x_{n-1}$ , and  $y'$  were constant on  $[x_{n-1}, x_n]$ . *One-step* methods such as Euler's possess the general form

$$y_n = \Phi_h(x_{n-1}, y_{n-1}, y_n), \quad (2.12)$$

where the name comes from  $\Phi_h$  only being a function of the numerical result at the previous step,  $y_{n-1}$ . If  $\Phi_h$  is further independent of  $y_n$ , then the method is *explicit*, since the solution at each step of the method can be computed in an iterative manner, as opposed to by solving a system of algebraic equations (in *implicit* methods). A more general formulation of the explicit Euler's method gives a continuous approximation for the solution on  $[x_{n-1}, x_n]$  as

$$Y(x) = y_{n-1} + (x - x_{n-1})f(y_{n-1}). \quad (2.13)$$

Given  $f$  is well-behaved, we can expect the approximation in Eq. (2.11) to get better with smaller  $h$ . How quickly we gain accuracy by reducing the stepsize, and how quickly error accumulates over several steps can be quantified via the local and global error properties of the method. We define these key quantities below.

**Definition 2.2.1** (Global truncation error). *The global error of a method is defined by*

$$e_n \equiv y_n - y(x_n), \quad (2.14)$$

or, if it is defined to be continuous, as

$$e(x) \equiv Y(x) - y(x). \quad (2.15)$$

---

<sup>1</sup>from here onwards, the vector notation will be dropped from  $y$  and  $f$  for simplicity, but it is to be understood that the formulae generalise to systems of ODEs where  $y$  and  $f$  are vector-valued.

**Definition 2.2.2** (Local truncation error). *The local error of a numerical method is*

$$l_n \equiv \Phi_h(x_{n-1}, y(x_{n-1}), y(x_n)) - y(x_n). \quad (2.16)$$

Eq. (2.16) shows how the approximate solution  $Y$  diverges from the actual solution  $y$  within a single step, and is reset at every step by virtue of evaluating  $\Phi_h$  at  $y(x_{n-1}), y(x_n)$  rather than  $y_{n-1}, y_n$ . The continuous generalisation of the local truncation error for Euler's method is

$$l(x) = y(x) - y(x_{n-1}) - (x - x_{n-1})y'(x_{n-1}), \quad (2.17)$$

where  $f(x_{n-1})$  has been replaced by  $y'(x_{n-1})$ . If  $f$  is continuously differentiable, we can place a bound on the local error by considering Taylor's theorem with its remainder written in Lagrangian form,

$$y(b) - y(a) - (b - a)y'(a) = \frac{1}{2}(b - a)^2 y''(\xi), \quad \xi \in (a, b), \quad (2.18)$$

and then choosing  $a = x_{n-1}$ ,  $b = x$  such that

$$\|l(x)\| \leq \frac{1}{2}(x - x_{n-1})^2 \max_{\xi \in [x_{n-1}, x]} \|y''(\xi)\|, \quad (2.19)$$

where  $\|\cdot\|$  refers to the Euclidean norm. In simpler terms, the local error is asymptotically  $\mathcal{O}(h^2)$ . A related definition is that of the consistency order.

**Definition 2.2.3** (Consistency order). *If the local truncation error of a numerical method takes the form*

$$\|l_n\| = \varphi_n h^{p+1} + \mathcal{O}(h^{p+2}), \quad (2.20)$$

*then the method is said to be of consistency order  $p$ .*

When a numerical method is referred to as having order  $p$ , this is understood to be referring to their consistency order. Constructing a bound on the global error is more difficult in general, and at first glance, the two errors defined above seem to have little in common. They may, however, be connected via [Lemma 2.2.1](#), for which the following definition is necessary.

**Definition 2.2.4** (Upper logarithmic norm). *The upper logarithmic norm of a function  $f$  of a single variable is defined as*

$$M[f] = \sup_{u \neq v} \frac{\langle u - v, f(u) - f(v) \rangle}{\langle u - v, u - v \rangle}, \quad (2.21)$$

*where the angled brackets refer to the inner product.*

The following lemma can then be shown [14].

**Lemma 2.2.1** (Upper bound on the global error).

$$\|e_n\| \lesssim \max_{m \leq n} \frac{\|l_m\|}{h} \frac{e^{M[f]x_n} - 1}{M[f]}. \quad (2.22)$$

The lemma proves an important point in numerical analysis: the global error can be kept below a given tolerance by controlling  $l(x)/h$ , the local error per unit stepsize. If the latter is kept near some tolerance limit, the global error will be proportional to that limit with a constant of proportionality only dependent on the right-hand-side of the ODE and the choice of steps  $x_n$ . Applying Eq. (2.22) to Euler's method yields

$$\|e_n\| \leq C(x_n)h, \quad (2.23)$$

i.e. the global error is asymptotically  $\mathcal{O}(h)$ . Since  $\lim_{h \rightarrow 0} e_n = 0$ , Euler's method is thus shown to be *convergent*. Its rate of convergence is summarised by its *convergence order*.

**Definition 2.2.5** (Convergence order). *If the function  $f$ , the right-hand-side of an ODE, is sufficiently smooth, and that the global error of a numerical method solving the ODE behaves as*

$$\lim_{h \rightarrow 0} \|e_n\| = \mathcal{O}(h^p), \quad (2.24)$$

*the numerical method is said to be of convergence order  $p$ .*

## 2.2.2 Generalisations of Euler's method

The explicit Euler's method introduced in the previous section is the simplest example in a large family of numerical methods built on the same principles. Its steps are quick and computationally cheap to evaluate, and it has been shown to be convergent, but due to its low rate of convergence ( $p = 1$ ), a significant reduction in stepsize needs to be made in order to gain accuracy. A common motivation behind developing novel numerical methods has been to achieve asymptotic errors that behave like some high power of  $h$  (*high-order* methods), in order to gain accuracy faster. One might also wish to improve the numerical method's stability, which is discussed further in Section 2.3.1. The accuracy gain and better stability properties have to be balanced against a potential increase in computational cost.

There are three mainstream independent avenues through which the order or stability of a method can be increased. These are described below and their relationships are plotted in Fig. 2.1.

**More evaluations per step:** Referring to further evaluations of  $f$  (with different arguments) in a single step, it is possible to increase the consistency order of a method, for example

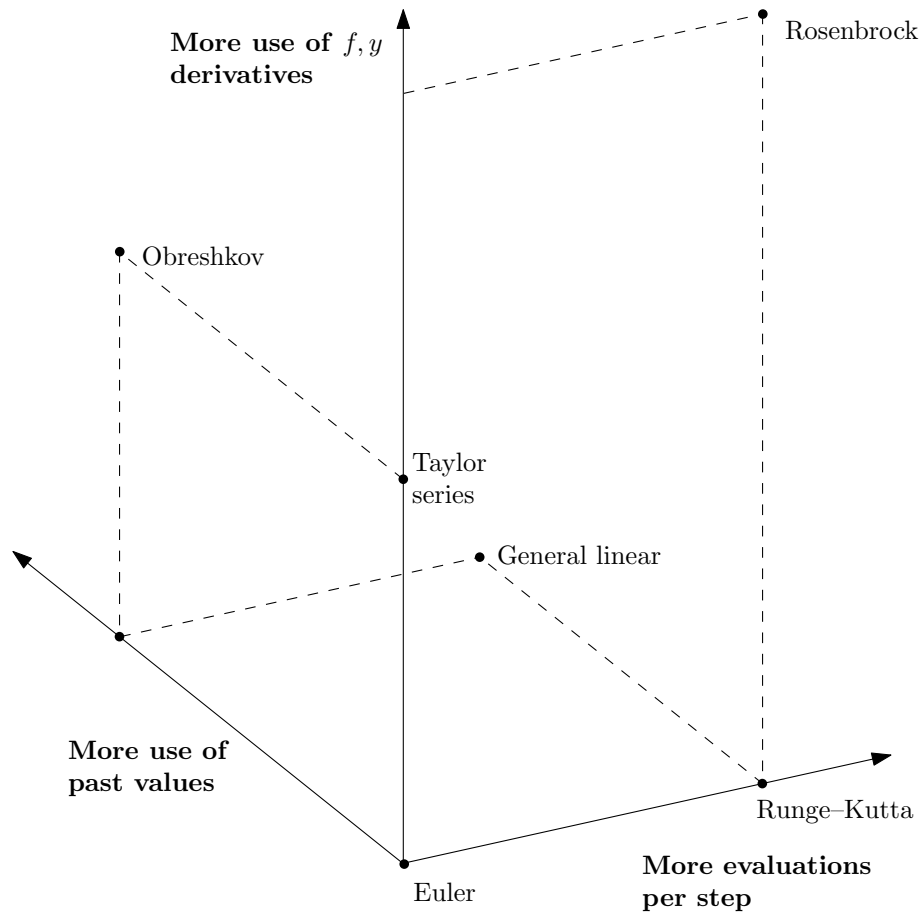


Fig. 2.1 Generalisations of Euler's method. The order of Euler's method can be improved in three independent ways: by using more evaluations of the right-hand-side,  $f$ , to approximate the solution at the end of the timestep; by making use of the values of the approximate solution  $y$  at previous timesteps; and by using higher derivatives such as  $f'$ . This figure is based on Fig. 224(i) of [15].

the method described by the formula

$$y_n = y_{n-1} + \frac{h}{2} [f(y_{n-1}) + f(y_{n-1} + hf(y_{n-1}))] \quad (2.25)$$

has an  $\mathcal{O}(h^3)$  local error. The above is an example of an explicit *Runge-Kutta* method, which will be further discussed in [Section 2.2.3](#).

**Dependence on additional previous values:** Higher-order errors may also be achieved by including the values  $y_{n-i}$ ,  $f(y_{n-i})$  with  $i = 2, \dots, k$ , giving these method the name *k-step* or *multistep*. A simple example is the explicit method (with uniform stepsize  $h$ )

$$y_n = y_{n-1} + \frac{h}{2} [3f(y_{n-1}) - f(y_{n-2})], \quad (2.26)$$

whose local error is  $\mathcal{O}(h^3)$ . If the dependence on past values is linear, the method in question is a *linear multistep* method. The first (explicit) methods of this type were proposed by Adams and Bashforth [16], while their implicit counterparts were refined by Moulton [17]. Implicit linear multistep methods are often evaluated iteratively by taking a sequence of predictor and corrector steps, the development of which is associated with Milne [18–20]. Backwards differentiation formulae (BDF) are a type of implicit linear multistep methods especially designed for the solution of stiff sets of differential equations. Adams and BDF methods are commonly found in modern numerical libraries [21, 22].

**Usage of higher derivatives:** A natural improvement on the local error of Euler’s method, based on Eq. (2.17), could come from including more terms in the Taylor expansion it uses<sup>2</sup>, e.g.

$$y_n = y_{n-1} + hf(y_{n-1}) + \frac{h^2}{2} f'(y_{n-1})f(y_{n-1}), \quad (2.27)$$

where  $f'$  is the Jacobian. Methods of this kind are named after Rosenbrock [23], who proposed  $f'$  to be used directly in the design of implicit Taylor-series-based methods to cut down computational costs.

### 2.2.3 Runge–Kutta methods

The idea to allow multiple function evaluations within a step as a generalisation of Euler’s method was first proposed by Runge in 1895 [24], and further developed by Heun [25], Kutta [26], and others. Until the 1950s [27, 28], only methods of order 5 and lower were considered, and efforts were initially focused on explicit methods, whose implicit counterparts are now deemed to be a good option for solving stiff sets of equations. A systematic approach for the analysis of Runge–Kutta methods was developed by Butcher [29] and is now ubiquitous in the field.

Explicit Runge–Kutta methods for the non-autonomous system Eq. (2.1) follow the algorithm below when evaluating the solution. First, approximations of the dependent variable are evaluated at  $s$  intermediate points between  $x_{n-1}$  and  $x_n$  according to

$$Y_i = y_{n-1} + h \sum_{j=1}^{i-1} a_{ij} F_j, \quad i = 1, 2, \dots, s, \quad (2.28)$$

where

$$F_i = f(x_{n-1} + c_i h, Y_i), \quad (2.29)$$

---

<sup>2</sup>Runge–Kutta methods make use of the same idea, but without using the Jacobian.

0					
$c_2$	$a_{21}$				
$c_3$	$a_{31}$	$a_{32}$			
$\vdots$	$\vdots$	$\vdots$	$\ddots$		
$c_s$	$a_{s1}$	$a_{s2}$	$\cdots$	$a_{s,s-1}$	
	$b_1$	$b_2$	$\cdots$	$b_{s-1}$	$b_s$

Table 2.1 Butcher tableau for an explicit Runge–Kutta method.

are approximations of  $y'$  at the same intermediate points. Then the  $F_i$  are linearly combined to form a higher-order approximate solution:

$$y_n = y_{n-1} + h \sum_{i=1}^s b_i F_i. \quad (2.30)$$

The methods is said to have  $s$  stages, since  $f$  is evaluated  $s$  times during a step. The constants  $a_{ij}$ ,  $c_i$ , and  $b_i$  characterise the method fully and are often summarised in a Butcher tableau, shown in [Table 2.1](#) Implicit Runge–Kutta methods would have non-zero  $a_{ij}$  values for  $j \geq i$ , and for autonomous systems, the  $c_i$  can be disregarded. By considering an autonomous system in which  $x$  is a component of  $y$  (such as [Eqs. \(2.4\)](#) and [\(2.5\)](#)), one can derive the *consistency condition*,

$$\sum_{j=1}^s b_j = 1, \quad (2.31)$$

which has to hold true in order for the  $x_n$  as forecast by the method and the actual  $x_n = x_{n-1} + h$  to be equal. By the same reasoning, for the intermediate points at which the  $Y_i$  are evaluated in [Eq. \(2.28\)](#) to match those at which the  $F_i$  are evaluated in [Eq. \(2.29\)](#) we need

$$c_i = \sum_{j=1}^{i-1} a_{ij} \quad (2.32)$$

to hold for each  $i = 1, 2, \dots, s$ . In this section, we will focus on explicit Runge–Kutta methods, and refer the reader interested in implicit methods to [\[30\]](#).

It is clear that the coefficients  $a_{ij}$  and  $b_i$  will determine the local error of a given Runge–Kutta method. In order for the numerical solution [Eq. \(2.30\)](#) to match the Taylor expansion of  $y(x_{n-1} + h)$  up to a given power of  $h$ , the Butcher coefficients need to satisfy a number of necessary and sufficient constraints called *order conditions*. The order conditions can be derived algebraically by Taylor expanding each  $Y_i$  and  $F_i$  in [Eq. \(2.30\)](#) around  $x_{n-1}$  and  $y_{n-1}$ , and matching the coefficients of powers of  $h$  up to a given order. Keeping track of terms (and all the contributions to the coefficient of a given power of  $h$ ) quickly becomes difficult and error-prone as  $s$  is increased. Butcher offers a more elegant and compact derivation with the

$p$	# order conditions	$s_{\min}$	$s_{\min}(s_{\min} + 1)/2$
1	1	1	1
2	2	2	3
3	4	3	6
4	8	4	10
5	17	6	21
6	37	7	28
7	85	9	45
8	200	11	66

Table 2.2 The table shows the minimum number of stages required ( $s_{\min}$ ) for an explicit Runge–Kutta method to achieve order  $p$ , the number of order conditions needed to be satisfied at that order, and the number of free parameters the method has when operating at the minimum number of stages (equal to  $s_{\min}(s_{\min} + 1)/2$ ).

help of graph theory, in which *rooted trees*<sup>3</sup> are associated with polynomials in the  $b_i$  and  $a_{ij}$ , with each tree resulting in a condition. All conditions associated with trees with no more than  $p$  vertices have to be satisfied for the method to achieve order  $p$ . A complete description of the approach can be found in [31].

An  $s$ -stage explicit Runge–Kutta method has  $s(s+1)/2$  free parameters, whereas the number of order conditions it has to satisfy equals the number of rooted trees with no more than  $p$  vertices, which scales as shown in Table 2.2. The table also shows the minimum number of stages required to achieve a given order, and the number of degrees of freedom the method has at the minimum number of stages. The first few rows might lead one to conjecture that the Runge–Kutta method has to have at least as many degrees of freedom as there are order conditions, but this is proven false at higher orders, where it is sufficient to have fewer free parameters than the number of order conditions would imply (if they were independent).

In practice, Runge–Kutta methods of order 4 or 5 are the default setting in numerical initial value problem solver routines [21, 22], but often formulae up to 8th order are available. Butcher tableaux such as Table 2.1, with only one row of  $b_i$  coefficients are rare however, instead they frequently appear as *embedded pairs*, with at least one additional row of coefficients to compute a lower (or higher) order solution with. This is due to the need for estimating the dominant term in the local truncation error of the method used to propagate the solution, i.e.  $\varphi_n h^{p+1}$  in Eq. (2.20), which is most commonly performed one of the following ways.

**Richardson extrapolation:** A method devised by Richardson [32] is based on repeating the numerical step from  $x_{n-1}$  to  $x_n$  with two steps of size  $h/2$ , and using the difference between the one full-step and the two half-step estimates ( $\tilde{y}_n$ ) to deduce the error. Taking

<sup>3</sup>a non-rooted tree is a connected graph with no circuits. A rooted tree is a tree in which a particular vertex is distinguished as the root, the purpose of which is to reduce the number of isomorphic rooted trees to those that map the root of one to the root of the other.

two half-steps is equivalent to carrying out a Runge–Kutta step of order  $p$  in  $2s$  stages whose local errors sum to

$$2^{-p}\varphi(x_n)h^{p+1} + \mathcal{O}(h^{p+2}), \quad (2.33)$$

(as opposed to  $\varphi(x_n)h^{p+1} + \mathcal{O}(h^{p+2})$ ), therefore

$$(1 - 2^{-p})(y_n - \tilde{y}_n) \quad (2.34)$$

is an asymptotically correct estimate of the local error. One might criticise this approach for obtaining the error estimate at three times the computational effort, but one might argue that the excess work is only  $(3s - 1)/2s$ -fold, due to evaluations common to both steps and in comparison to propagating the solution with the  $2s$ -stage method (i.e. replacing  $y_n$  with  $\tilde{y}_n$ ).

**Embedded pairs:** A more popular choice is to devise methods with two sets of  $b_i$  coefficients, of which the set  $b_i$  gives an order  $p$ , and  $b_i^*$  yields an order  $\tilde{p}$  solution. With the coefficients chosen appropriately, the difference between the two is then used as an estimate of the local error, and one of the two methods propagates the solution. Embedded methods have fewer free parameters to satisfy the order conditions, therefore they often use additional stages, e.g. it is not possible to find a set of coefficients that yields a 4th order method and an error estimator in 5 stages. The extra stages don't always come at increased computational cost, because the additional stage is often chosen to be an evaluation at  $c_s = 1$ , i.e. at the end of a step. This evaluation would be necessary anyway for the next step, provided the current step is accepted, and is called the first same as last (FSAL) approach. The earliest embedded pairs were proposed by Merson [33], and improved upon by Fehlberg [34–37]. These were both of the type  $p(\tilde{p})$  with  $p < \tilde{p}$ , meaning that the lower order method is used to propagate the solution. Fehlberg's popular 4(5) pair in particular has been optimised such that the error coefficients of the lower order result are minimal, which means that the error estimate sometimes underestimates the error of the higher order mode, which therefore *should not* be used as the solution. Modern codes almost exclusively use *local extrapolation*, i.e. they have been designed to use the higher order mode for the solution. The most ubiquitous Runge–Kutta pairs are those proposed by Dormand and Prince [38], which make heavy use of the FSAL approach. Some Dormand–Prince pairs have been compared with those of Fehlberg by Shampine [39], among others.

**Reusing past solution values:** It would be ideal if computing an error estimate didn't come at any additional  $f$ -evaluations. Stoller and Morrison [40] proposed the idea that if successive steps  $y_{n-3}, y_{n-2}, y_{n-1}, y_n$  can be forced to have a given stepsize-ratio which would place them at the nodes of a quadrature method, e.g. to be at  $x_{n-1}, x_{n-1} + (6 - \sqrt{6})h/10, x_{n-1} + (6 + \sqrt{6})h/10, x_{n-1} + h$ , then, if the Runge–Kutta method

gives fourth order estimates of  $y$  and fifth order estimates of  $y'$  at these points, they may be linearly combined using Radau quadrature with 3 nodes [41] to give an estimate of  $y(x_n + h) - y(x_n)$ , and in turn the local error on  $y_n$ ,

$$\frac{h}{36} \left[ 4f(y_{n-3}) + (16 + \sqrt{6})f(y_{n-2}) + (16 - \sqrt{6})f(y_{n-1}) \right] + y_{n-3} - y_n. \quad (2.35)$$

Ceschino and Kuntzmann [42] later extended the method not to require uneven spacing between steps, and to handle methods of order up to five.

### 2.2.4 Specialised methods for oscillatory ODEs

We have seen in [Section 2.1.2](#) that oscillatory solutions in ODEs can arise in many ways and different applications prioritise different properties of a numerical solver. As a result, methods for solving oscillatory ODEs are highly specific to the application. Due to the size of the literature, we focus our review of past efforts to solve oscillatory ODEs that are capable of solving equations of the form [Eq. \(2.9\)](#) *efficiently*, without dampening or averaging over the oscillations, and which can resolve the oscillations. Often when oscillatory behaviour arises in a system, there are signals or data varying at highly different scales (“slow” and “fast” solutions), with the slow signal usually carrying important information and the fast, oscillatory “noise” being superimposed onto it. Another large class of problems is characterised by a set of coupled ODEs whose solution has highly oscillatory and slow components. In the examples of oscillatory problems in cosmology we pursue here the separation of scales is more subtle than in the *multiscale* problems described above: it is the slowly changing amplitude of the fast oscillations that carries useful information. The extensively researched field of multiscale methods provides numerical solvers appropriate for multiscale problems, but these solvers are not suitable for the single linear oscillator with high, time-varying frequency [Eq. \(2.9\)](#). We recommend [43–45] (and references therein) for reviews of a wide range of oscillatory (including multiscale) problems not discussed here and their associated numerical methods.

As a remark on the efficiency of numerical methods for oscillatory problems, Linda Petzold notes [44] that “...in general, one should not expect to be able to numerically solve nonlinear highly oscillatory problems using stepsizes which are large relative to the timescale of the fast solution” and notes that a major difference between standard and oscillatory numerical methods is that while the former make use of local information about the problem only, the latter must use implicit or explicit assumptions about the global behaviour or mathematical structure of the problem to achieve large stepsizes. This opinion is echoed by Iserles [46], who points out that standard numerical methods rely on a polynomial approximation of the solution, most often via Taylor’s theorem, (e.g. in Runge–Kutta and many linear multistep methods), and while this approximation holds up in sufficiently small neighbourhoods around  $x_n$ , it is inadequate for distances beyond  $\sim 1$  wavelength for an oscillatory solution. This forces the stepsizes of Taylor’s theorem-based methods to be small relative to the timescales of the problem, rendering them

inefficient by construction. So far, no approximation as versatile and widely applicable as the Taylor expansion has been found for highly oscillatory problems, and the field is characterised by many distinct approaches being trialled in parallel. We provide details of these research avenues below.

### Asymptotic methods

Iserles [45, 46] advocates for numerical methods that exploit *asymptotic* expansions. An asymptotic relation between two functions in some limit, denoted

$$f(x) \sim g(x), \quad x \rightarrow x_0, \quad (2.36)$$

means that the relative error between them goes to 0 in the appropriate limit, or, equivalently, that their ratio tends to unity:

$$f(x) - g(x) \ll g(x), \quad x \rightarrow x_0 \quad \text{or} \quad \lim_{x \rightarrow x_0} f(x)/g(x) = 1. \quad (2.37)$$

Asymptotic methods seek to construct the successive terms of a formal power series expansion to approximate the behaviour of the solution in a given limit, often with the substitution

$$y(x) = e^{S(x)}. \quad (2.38)$$

Using the method of *dominant balance*, terms in the differential equation for  $S(x)$  can be identified as negligible in the asymptotic limit, and dropped, turning the equation into an asymptotic relation. The asymptotic relation is then replaced by an equality sign, and the resulting approximate equation is solved. The solution will be asymptotically correct (although not unique), and is then checked against the assumptions that allowed terms in the original first-order ODE for  $S(x)$  arising from Eq. (2.38) to be neglected. The power of the substitution in Eq. (2.38) cannot be overestimated: it reduces an  $n$ th order linear differential equation for  $y(x)$  to an approximate first-order one for  $S(x)$  valid in the asymptotic limit (e.g. around an irregular singular point). For example, to derive an approximate solution for the equation

$$y'' + p(x)y' + q(x)y = 0, \quad (2.39)$$

around an irregular singular point  $x = x_0$ , using the substitution in Eq. (2.38), we get

$$S'' + (S')^2 + p(x)S' + q(x) = 0. \quad (2.40)$$

Around an irregular singular point, the  $S''$  term can be dropped since  $S'' \ll (S')^2$  usually holds true. Thus Eq. (2.40) reduces to

$$(S')^2 \sim -p(x)S' - q(x), \quad x \rightarrow x_0, \quad (2.41)$$

which, once the asymptotic relation is turned into an equality sign, is much easier to solve.

Asymptotic relations can be derived for oscillatory functions, a well-known example in quantum mechanics is the Wentzel–Kramers–Brillouin (WKB) approximation which is of the form

$$y(x) \sim \exp\left(\frac{1}{\varepsilon} \sum_{i=0}^N \varepsilon^i S_i(x)\right), \quad (2.42)$$

where  $\varepsilon$  is a small parameter the above power series is expanded in terms of, and the  $S_i$ , which depend on  $\omega(x)$  and  $\gamma(x)$ , can be found recursively. The WKB approximation is introduced in more detail in Chapter 3 and used heavily in the numerical methods described in this thesis.

Arnold et al. [47] devised an efficient numerical method (named the *WKB-marching method*) for the time-independent, one-dimensional Schrödinger equation,

$$\varepsilon^2 \varphi''(x) + a(x)\varphi(x) = 0, \quad (2.43)$$

where  $0 < \varepsilon \ll 1$  is very small, and  $a(x) \geq a_0 > 0$  is a smooth function, therefore  $\varphi(x)$  is highly oscillatory. They propose a method that functions on a coarse spatial grid with stepsizes larger than the associated wavelength,  $h > \lambda = (2\pi\varepsilon)/\sqrt{a(x)}$ . The method starts with a preprocessing step, the transformation of Eq. (2.43) into a smoother problem via a series of substitutions. The resulting ODE system is

$$\frac{dZ}{dx} = \varepsilon N_\varepsilon Z, \quad (2.44)$$

where  $N_\varepsilon$  is a  $2 \times 2$  matrix that depends on  $x$  (its exact form is unimportant for the present discussion), and  $Z = (z_1, z_2)^T$  is a 2-component vector. The transformation to take one back to  $\varphi(x)$  from  $Z$ , in the asymptotic limit  $\varepsilon \rightarrow 0$ , reduces to

$$\varphi(x) = \frac{1}{\sqrt{2}(a(x))^{\frac{1}{4}}} \left( -ie^{\frac{i}{\varepsilon}\phi^\varepsilon(x)} z_1 + ie^{-\frac{i}{\varepsilon}\phi^\varepsilon(x)} z_2 \right), \quad (2.45)$$

where

$$\phi^\varepsilon(x) = \int_0^x d\tau \left( \sqrt{a(\tau)} - \varepsilon^2 \beta(\tau) \right). \quad (2.46)$$

Eq. (2.45) is a linear combination of two Wentzel–Kramers–Brillouin solutions truncated after the second term (i.e.  $N = 2$  in the sum of Eq. (2.42)). The purpose of the chain of transformations is to arrive at Eq. (2.44), whose solution still oscillates, but now with a small,  $\mathcal{O}(\varepsilon^2)$  amplitude around a smoothly varying function. Solving Eq. (2.44) on a coarse grid ( $h \gg \varepsilon$ ) therefore incurs at most  $\mathcal{O}(\varepsilon^2)$  error. The numerical solution of Eq. (2.44) involves the

numerical calculation of oscillatory integrals of the form

$$\int_x^{x+h} \beta(y) \exp\left(\frac{2i}{\varepsilon}\phi(y)\right) dy, \quad (2.47)$$

where  $\phi$  and  $\beta$  are slowly-varying functions of  $x$ . The literature (e.g. [48]) offers several asymptotic methods to carry out the integrals, but these often focus on achieving high  $\varepsilon$ -order (rather than  $h$ -order) in their error estimates. The authors therefore derive their own asymptotic method inspired by those described in [48, 49]. Their novel quadrature method has  $\mathcal{O}(h^2)$ , which they argue is better for the solution of an ODE than  $\mathcal{O}(h)$ , as one gains more accuracy by reducing the stepsize. In a later work by Körner et al. [50], the WKB marching method is complemented with a Runge–Kutta–Fehlberg routine to switch to when the solution approaches a *turning point*, where it changes between oscillatory and evanescent. It also gains an *adaptive stepsize algorithm* (to be discussed in Section 2.3.2) in the oscillatory regime, which is based on comparing the output of the WKB marching method when using an  $\mathcal{O}(h^2)$  and  $\mathcal{O}(h)$  quadrature formula for the required integrals. The resulting numerical algorithm is, however, limited to the  $\gamma(x)$ -free form of Eq. (2.9).

### Envelope-following methods

In the field of circuit simulation, problems typically have a highly oscillatory solution with a single, high frequency, and information about the system is carried in the variation of the amplitude of oscillations. Methods that exploit this property are built to reconstruct the *quasi-envelope* of the solution, defined to be a continuous function connecting an initial value on the solution to values of the solution at a sequence of points separated by the period  $T$ . In between the samples, the quasi-envelope is interpolated. It is different from the conventional definition of an envelope in that the quasi-envelope depends on what  $x_0$  the initial value was sampled at, and is therefore not unique. These types of methods aim to take stepsizes that are larger than the wavelength of oscillations by using the fact that the quasi-envelope is slowly-varying, and are called *envelope-following* methods. The earliest examples include *multirevolution* methods used in astronomy to calculate the orbits of artificial satellites [51–53], often requiring a physical reference point (e.g. an apogee or perigee). Since the quasi-envelope coincides with the solution of the ODE at multiples of the period  $T$ , the idea is that the solution can be recovered by integrating the ODE from a point on the quasi-envelope for at most a cycle. The method assumes that the period  $T$  is known a priori, but it need not be constant (in which case, it assumes knowledge of the initial period). Petzold [54] proposed a method which, given an initial estimate, will compute  $T$ , and re-compute it periodically if necessary. They pose the problem as a system of first-order ODEs,

$$y'(x) = f(y, x), \quad y(0) = y_0, \quad 0 \leq x \leq L, \quad (2.48)$$

where  $y(x)$  may be a vector whose components all oscillate with period  $T$ , or vary slowly (but at least one component should be oscillatory). Eq. (2.9) may be rewritten in this form. They then define the quasi-envelope  $z(x)$  as

$$z(x + T) = z(x) + Tg(z(x), x), \quad 0 \leq t \leq L - T, \quad (2.49)$$

with

$$g(z, x) = \frac{1}{T} [\tilde{y}(x + T, x) - \tilde{y}(x, x)] \quad (2.50)$$

and

$$\frac{d}{ds} \tilde{y}(x + s, x) = f(\tilde{y}(x + s, x), x + s), \quad \tilde{y}(x, x) = z. \quad (2.51)$$

This way, if  $z(0) = y(0)$ , then  $z(nT) = y(nT)$ , i.e.  $z$  coincides with  $y$  at multiples of the period. If  $y$  is close to being periodic, then  $z$  will vary slowly. The aim is then to solve the difference equation Eq. (2.49) approximately, as this will be equivalent to solving Eq. (2.48) (since  $g$  is obtained by integrating Eq. (2.48) over one period). For solving Eq. (2.49), several methods have been derived: with stepsizes  $H \gg T$ , Petzold [54] proposed a generalised Adams-Moulton method [55], but other standard explicit or implicit ODE methods can be used as well. The case when  $T$  is not constant is dealt with by transforming the independent variable such that it becomes constant. Re-computing a slowly-changing period, however, involves solving the original problem for a cycle, which incurs computational cost.

### The phase function method

Recent work by Bremer, Heitman, and others [56–58] offers a numerical method to solve

$$y''(x) + \lambda^2 q(x)y(x) = 0, \quad (2.52)$$

where  $q$  is positive and  $\lambda$  is real and large, therefore  $y(x)$  is highly oscillatory. Although this form is less general than Eq. (2.9), with a transformation of the independent or dependent variable, any equation of the form Eq. (2.9) can be simplified to Eq. (2.52) (although the new independent-dependent variable pair may not be computationally advantageous). They propose an algorithm whose running time is independent of the parameter  $\lambda$ . They do so by finding the non-oscillatory phase function  $\alpha(x)$  such that

$$u(x) = \frac{\cos(\alpha x)}{\sqrt{|\alpha'|}}, \quad v(x) = \frac{\sin(\alpha x)}{\sqrt{|\alpha'|}} \quad (2.53)$$

form a basis in the space of solutions of Eq. (2.52).  $\alpha$  is a phase function if and only if it satisfies *Kummer's equation* (in its logarithm form),

$$r'' - \frac{1}{4}(r')^2 + 4\lambda^2(e^r - q) = 0, \quad (2.54)$$

with

$$\alpha = \lambda \int_a^x e^{\frac{r(u)}{2}} du. \quad (2.55)$$

Kummer’s equation Eq. (2.54) may not possess a non-oscillatory solution  $\alpha(x)$ , but if  $q(x)$  is non-oscillatory, there exists a non-oscillatory function  $\alpha(x)$  such that the basis functions  $u(x)$ ,  $v(x)$  from Eq. (2.53) approximate the exact solutions of Eq. (2.52) with  $\mathcal{O}((\mu\lambda)^{-1} \exp(-\mu\lambda))$ , where the constant  $\mu$  depends on  $q$ , but is independent of  $\lambda$ . The resulting functions  $r$  and  $\alpha$  are non-oscillatory in that they can be represented accurately via series expansions (e.g. the authors’ open-source code employs Chebyshev polynomials for this) with  $\mathcal{O}(1)$  terms. This should be compared to the super- and hyperasymptotic expansions in [59, 60] which have  $\mathcal{O}(\exp(-\rho\lambda))$  accuracy and require  $\mathcal{O}(\lambda)$  terms. The authors propose a method for constructing a solution for Eq. (2.54), which differs from the non-oscillatory approximate solution described above by  $\mathcal{O}(\exp(-\frac{1}{3}\mu\lambda))$ , which yields an error in the solution of Eq. (2.52) comparable to machine precision in the high-frequency regime. This process includes solving stiff ODEs, for which they use a specialised, highly stable (spectral deferred correction) method from [61]. In the extremely high frequency limit, computing  $u$  and  $v$  involves the evaluation of trigonometric functions with very large arguments, which limits the accuracy (when using finite precision arithmetic). Another serious limitation of the method is that it is only applicable in the large- $\lambda$  limit, it therefore needs to be paired with an appropriate alternative in regions where Eq. (2.52) becomes non-oscillatory.

### oscode in the oscillatory solver landscape

oscode was born out of the need for a fast and efficient solver for oscillatory ODEs of the form Eq. (2.9), more specifically those relevant to primordial cosmology, in order to speed up the forward modelling stage of e.g. a Monte–Carlo exploration of a parameter space. This area of application requires the solver to be able to deal with a general, second-order, linear ODE, i.e. Eq. (2.9) with a non-zero  $\gamma$  term. In principle, any ODE of the form Eq. (2.9) can be transformed to one with the  $\gamma$  term absent by a rescaling of  $x$  or  $y$ , but in practice, this may introduce numerical instabilities and there may be a preferred choice for the independent-dependent variable pair. An important example comes from primordial cosmology, where the well-known friction-free form of the oscillatory Mukhanov–Sasaki equation, introduced in Eq. (5.60), is numerically unstable: the dependent variable decreases exponentially for some of the solution range, and the independent variable “saturates” in the same region (during inflation). The former is problematic because a numerical method will pick up the exponentially increasing solution in the exponential decay phase (although this could potentially be remedied by running the solver backwards in time), and in the latter case time-steps of the solver become infinitesimally small. There exists a numerically stable form of the Mukhanov–Sasaki equation, Eq. (5.61), but it contains a first-derivative term.

It was therefore imperative that `oscode` be able to accept a non-zero  $\gamma$  term, and out of the methods reviewed above, Petzold’s envelope-following method [54] is the only other one to do so. Unfortunately, there is no open-source code available that implements this method. Bremer’s phase function approach [58] is the only other method with an open-source implementation available, with `oscode` being unique as a peer-reviewed software package for oscillatory ODEs.

An important feature `oscode` possesses is the ability to switch to a different, more suitable solver for non-oscillatory regions of the ODE. The only other method capable of doing this is [50], which for the relevant section cites the work presented here [1] as its inspiration.

The lack of open-source implementations of oscillatory ODE solvers makes benchmarking `oscode` difficult. The only comparison we could draw would be with the phase function method, but such a comparison would not be fair as the two methods operate within significantly different tolerance regimes. The phase function method – by design – operates at machine precision<sup>4</sup>, while `oscode` was designed with a relative tolerance range of  $10^{-4} - 10^{-6}$  in mind. This is because the measurement of physical quantities relevant to primordial cosmology would likely not warrant the numerical ODE solver to be more accurate than about 1 in  $10^6$ , and unnecessary computation time is best avoided in a parameter search setting.

## 2.3 Requirements of a practical numerical routine

### 2.3.1 Stability

Stability generally means *robustness* against small disturbances. In the context of numerical methods, it more specifically refers to a numerical method tracing the correct solution within the tolerance requirements when integrating over a large interval. An example for instability is shown in Fig. 2.2, where the numerical solver has been applied to a *stiff* system of ODEs [63]: one which has multiple, very different scales on which the solution is changing. Explicit methods typically have to reduce their stepsizes when solving such equations to below what would be reasonable at the given tolerance level, in order to avoid picking up any (transient) solutions not being sought.

By convention, the stability properties of a numerical method are determined by applying it to one or more test equations. The simplest of them is the *Dahlquist test equation*,

$$y'(x) = \lambda y(x), \tag{2.56}$$

where  $\lambda \in \mathbb{C}$  is a constant, and  $y(x)$  is a scalar-valued function. The analytic solution is of course  $Ae^{\lambda x}$ . Defining  $z = h\lambda$ , the analytic solution is bounded for  $\{z \in \mathbb{C} : \Re(z) \leq 0\}$ , i.e. on the non-positive half-plane. The stability of a numerical method is defined via the region of the complex plane within which its solution estimate for the Dahlquist equation stays bounded.

<sup>4</sup>Or the highest possible precision allowed by the conditioning of the problem at hand.




Fig. 2.2 Visualising instability in the numerical solution of a stiff ODE. The ODE has two solutions varying on different timescales: a transient (dashed), and an equilibrium (solid line) solution. Even though the transient solution decays away quickly, and the initial conditions would give the equilibrium solution, care needs to be taken for the numerical method not to pick up any of the other solution, and produce the behaviour pictured. To avoid this phenomenon, explicit solvers are forced to keep their stepsizes much shorter than the accuracy requirements would allow. Figure reproduced from [62].

Even though this does not determine the precise quantitative behaviour of the numerical method, it does qualitatively describe how well the method is able to trace the closed-form solution.

**Definition 2.3.1** (Stability function). *Consider the numerical solution of the Dahlquist equation with parameter  $\lambda \in \mathbb{C}$ , and suppose the numerical method took a step of size  $h$  from  $x_{n-1}$  to  $x_n$ . If the numerical solution is of the form*

$$y_n = R(\lambda h)y_{n-1}, \quad (2.57)$$

*then the function  $R$  is the method's stability function.*

**Definition 2.3.2** (Stability region). *A Runge–Kutta method with stability function  $R$  is stable in the region*

$$S = \{z \in \mathbb{C} : |R(z)| \leq 1\}. \quad (2.58)$$

The stability function of a Runge–Kutta method is clearly a polynomial in  $z$ , which leads to its stability region being bounded. For example, Euler’s method has

$$y_n = (1 + \lambda h)y_{n-1}, \quad (2.59)$$

and therefore its stability function is  $R(z) = 1 + z$ . The region for which  $|1 + z| \leq 1$  on the complex plane is a disc of unit radius centred on  $z = -1$ . The stability region of a more complex method, e.g. a higher-order Runge–Kutta, consists of lobes which may stretch farther into the negative real axis, making the method less prone to blowing up when encountering a quickly decaying solution.

While explicit methods are not to be used in practice on problems known to be stiff, stability may be important in cases as simple as  $y \sim e^{-\lambda x}$ . It is therefore worth investigating the stability properties of any novel numerical method.

### 2.3.2 Adaptive stepsize control

An essential feature of any good numerical method is having control over its estimated error by adapting its stepsize, i.e. keeping said error below a user-specified absolute or relative tolerance. While this may add a little computational overhead (e.g. from rejecting and repeating a step because the estimated error was too large), it is well worth the effort, as the algorithm will be able to take larger steps in ‘uneventful’ regions of integration. Adaptive stepsize algorithms rely on having access to past stepsizes, past local truncation errors, and tolerance settings. The novel numerical method to be introduced in [Chapter 3](#) has the ability to trace the solution of an ODE in both its oscillatory and slowly-varying phases. It is efficient thanks to its ability to take large steps covering many periods in the oscillatory region, but it also has to be able to reduce the stepsize quickly upon exiting an oscillatory regime. This means that choosing its stepsize-update algorithm appropriately will be essential.

#### Classical adaptive stepsize

Textbooks [55, 62] often start the discussion of adaptive stepsize by introducing the following, classical approach. Consider a method of order  $p$ , whose local error  $l$  we wish to keep below a tolerance limit  $\text{tol}$ . To account for ODE systems where  $l$  is a vector quantity, we can take its (Euclidean) norm or select the largest component. Proceeding with the former approach, we write

$$\|l\| = Ch^{p+1} + \mathcal{O}(h^{p+2}) \leq \text{tol}. \quad (2.60)$$

In practice, we may wish to allow the user to control both the relative ( $\text{rtol}$ ) and absolute ( $\text{atol}$ ) tolerance<sup>5</sup> and define  $\text{tol}$  as

$$\text{tol} = \text{atol} + \|y\| \cdot \text{rtol}. \quad (2.61)$$

<sup>5</sup>It is also possible to set  $\text{atol}$  and  $\text{rtol}$  separately for each component of the solution.

If we assume that  $h$  stays small relative to the limit set by the problem and method being used, then the last term in Eq. (2.60) can safely be neglected. Setting sufficiently low values for the tolerances we can always ensure that this approximation stays valid, which forms the basis of the so-called *asymptotic error model*. Among the stepsizes that would keep the local error within tolerance, we choose the largest for the sake of efficiency. Therefore we wish to set  $h_n$  such that

$$\text{tol} \approx Ch_n^{p+1}. \quad (2.62)$$

Using our asymptotic error model,  $\|l_n\| \approx Ch_{n-1}^{p+1}$ , we then obtain

$$\rho_{n-1} = \frac{h_n}{h_{n-1}} = \left( \frac{\text{tol}}{\|l_n\|} \right)^{\frac{1}{p+1}}, \quad h_n = \rho_{n-1} h_{n-1}. \quad (2.63)$$

Eq. (2.63) defines how much the current stepsize can safely be increased for the next step, should the current one be accepted. The current step should only be accepted if the anticipated error is below tolerance, i.e. if  $\rho_{n-1} > 1$ . Otherwise, the current step has to be re-attempted with a smaller stepsize as suggested by Eq. (2.63).

A weakness of the above algorithm is Eq. (2.60) being an approximation only, since it only takes into account the leading order term in the error. When used in practice, therefore, a number of additional precautions need to be taken:

- A safety factor  $0.8 \leq S \leq 0.9$  is included to slightly overestimate the local error:  $h_n = S\rho_{n-1}h_{n-1}$ .
- A minimum and maximum factor by which the stepsize can decrease and increase is added (or a hard minimum and maximum value for  $h$  is implemented).
- When decreasing the stepsize after a failed step, one may wish to replace the exponent  $1/(p+1)$  with  $1/p$  in Eq. (2.63) in order to ensure that the local error of the new attempted step is within tolerance.

Another, significant weakness of the classical approach is the behaviour of successive stepsizes when the stepsize is being limited by the numerical method's stability. The stepsizes are gradually increased until the edge of the stability region is reached, at which point the current step is rejected and the stepsize is decreased drastically (due to the sudden, large increase in local error). Since the stepsize was decreased beyond what would be reasonable, it is then increased rapidly again, taking the method out of the stability region, at which point the cycle repeats, resulting in oscillations in  $h$  and many rejected steps. This behaviour can be mitigated against using methods from *control theory*. We summarise basic concepts in control theory below, and refer to [64–66] for further details.

### Control theory for asymptotic error models

The basic idea underlying control theory is interpreting the interaction of the numerical method, the local error, and the sequence of stepsizes as a signal feedback loop, and constructing a *controller* to filter out certain components of the signal to optimise the behaviour of the system. The numerical method is a *process*, which takes a trial stepsize  $h$  as the *input signal* and generates the local error estimate as the *output signal*. The stepsize update algorithm is the *controller*. Much like in digital signal processing, the controller aims to determine the frequency response of the system, which we investigate with the  $z$ -transform (also known as the discrete Laplace transform), a generalisation of the Fourier transform for the complex plane:

**Definition 2.3.3** (z-transform). *Let  $x \in \mathbb{C}^N$  be a sequence. The  $z$ -transform of  $x$ , denoted  $\hat{x}$  is then*

$$\hat{x} = \hat{x}(z) = \sum_{k=0}^{\infty} x_k z^{-k}. \quad (2.64)$$

Since  $z$  can be written as  $Re^{i\phi}$ , Eq. (2.64) can be thought of as a generalised version of a discrete Fourier transform which is not restricted to the complex unit circle ( $R = 1$ ), and the  $x_k$  can be thought of as frequency components of the signal  $x$ . It is then not unsurprising that acting on a sequence  $x$  with the forward shift operator

$$T : \mathbb{C}^N \mapsto \mathbb{C}^N; \{x_n\}_{n=0}^{\infty} \mapsto \{x_{n+1}\}_{n=0}^{\infty} \quad (2.65)$$

amounts to a multiplication in frequency space,

$$\widehat{Tx}(z) = \sum_{k=0}^{\infty} x_{k+1} z^{-k} = z \sum_{k=0}^{\infty} x_{k+1} z^{-(k+1)} + \text{const.} = z\hat{x}(z) + \text{const.} \quad (2.66)$$

With the help of the  $z$ -transform, we can analyse the classical adaptive stepsize algorithm and improve upon it. Our *process* can be written

$$\begin{aligned} l_n &= \varphi_n h_n^p, \\ \ln l_n &= p \ln h_n + \ln \varphi_n, \end{aligned} \quad (2.67)$$

and the controller is expressed as

$$\begin{aligned} h_{n+1} &= \left(\frac{\text{tol}}{l_n}\right)^{\frac{1}{p}} h_n, \\ \ln h_{n+1} &= \ln h_n + \frac{1}{p}(\ln \text{tol} - \ln l_n). \end{aligned} \quad (2.68)$$

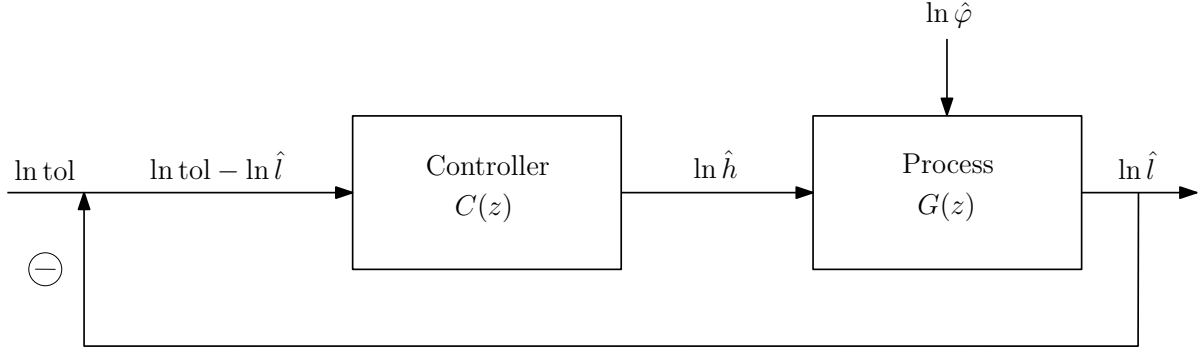


Fig. 2.3 Adaptive stepsize control interpreted as a feedback system. The *process* is the numerical method solving an ODE iteratively, taking as input a stepsize  $\ln \hat{h}$  and producing a local error estimate  $\ln \hat{l}$ . The ODE determines the error coefficient  $\ln \hat{\varphi}$ , which is present as an external disturbance. The system subtracts the error from the desired tolerance and attempts to minimise the result, the *control error signal*. The controller adapts the stepsize for the process based on the control error. The process and controller are represented by transfer functions and together form the closed loop transfer function, which must be stable, i.e. it must not let fluctuations in the external disturbance  $\ln \hat{\varphi}$  increase the control error  $\ln \text{tol} - \ln \hat{l}$ . Figure reproduced from [64].

We then apply the  $z$ -transform to Eqs. (2.67) and (2.68), denoting the  $z$ -transform of a series  $\{\ln x_n\}_{n=0}^{\infty}$  as  $\ln \hat{x}$ , and get

$$\ln \hat{l} = p \ln \hat{h} + \ln \hat{\varphi} = G(z) \ln \hat{h} + \ln \hat{\varphi}, \quad (2.69)$$

$$\ln \hat{h} = \frac{1}{p} \frac{1}{z-1} (\ln \text{tol} - \ln \hat{l}) = C(z) (\ln \text{tol} - \ln \hat{l}), \quad (2.70)$$

where  $G(z)$  and  $C(z)$  are called the *process* and *control transfer functions*, respectively. Eqs. (2.69) and (2.70) can be further simplified by noting that setting  $\text{tol} = 1$  does not reduce generality, since it only causes the rescaling of stepsizes. We can then write

$$\ln \hat{h} = -\frac{C(z)}{1 + C(z)G(z)} \ln \hat{\varphi} = H_{\varphi}(z) \ln \hat{\varphi}, \quad (2.71)$$

$$\ln \hat{l} = \frac{1}{1 + C(z)G(z)} \ln \hat{\varphi} = L_{\varphi}(z) \ln \hat{\varphi}. \quad (2.72)$$

This pair of equations is referred to as *closed loop dynamics*, since the system they describe can be expressed in a circuit-like fashion as illustrated in Fig. 2.3. The quantities  $H_{\varphi}(z)$ <sup>6</sup> and  $L_{\varphi}(z)$  are called the *stepsize transfer map* and *error transfer map*, respectively. The former is determined by the adaptive stepsize algorithm we choose, and the latter by the error model (which we view as fixed, for now). In this model  $\ln \varphi_n$  plays the role of an external signal coming from the numerical solution of the ODE.  $H_{\varphi}$  and  $L_{\varphi}$  can then be thought of as digital

<sup>6</sup>It is customary to work with the scaled stepsize transfer map,  $pH_{\varphi}(z)$  instead, since this makes  $pH_{\varphi}(1) = 1$ , regardless of the method's order.

filters that determine how the system (the stepsizes and the local error estimates) responds to the numerical error arising during the solution of the ODE (i.e. the external disturbance  $\ln \hat{\varphi}$ ). The transfer maps are rational functions (of  $z$ ) whose poles determine the stability<sup>7</sup> of the system. More specifically, the system is stable if all poles of  $H_\varphi$  are located strictly inside the (complex) unit circle. Assuming a signal of a single frequency,  $\ln \varphi_n = e^{i\omega n}$ , the response of the stepsizes is

$$\ln h_n = H_\varphi(e^{i\omega})e^{i\omega n}. \quad (2.73)$$

Since the mapping from  $\ln \hat{\varphi}$  to  $\ln \hat{h}$  is linear, the filter  $H_\varphi$  cannot change the frequency (or spectral content) of the input signal, only its amplitude and phase. The frequency response of the system is then the mapping  $\omega \mapsto A(\omega)$ , where the *attenuation*  $A(\omega)$  is

$$A(\omega) = |H_\varphi(e^{i\omega})|. \quad (2.74)$$

Substituting in the  $C(z)$  appropriate for the classical adaptive stepsize algorithm, we get  $\omega \mapsto 1$  for the frequency response of  $pH_\varphi$ , i.e. this controller does not attenuate the amplitude of any input signal, including that of high-frequency noise. This can definitely be improved. Let us reconsider the adaptive stepsize algorithm and write it in the more general multiplicative form

$$h_{n+1} = \rho_n h_n, \quad (2.75)$$

which encodes the importance of the previous step. Taking logarithms and applying the  $z$ -transform yields

$$\ln \hat{h} = \frac{1}{z-1} \ln \hat{\rho} = -C(z) \ln \hat{l} = -\frac{1}{z-1} F(z) \ln \hat{l}, \quad (2.76)$$

where  $F(z) = P(z)/Q(z)$  is a digital filter represented by a rational function. Writing the polynomials as

$$P(z) = \sum_{j=0}^M \beta_{M-j} z^j, \quad \text{and} \quad Q(z) = \sum_{j=0}^M \alpha_{M-j} z^j \quad (2.77)$$

and applying the inverse  $z$ -transform, we get a formula for the most general adaptive stepsize algorithm:

$$h_{n+1} \prod_{j=1}^M \left( \frac{h_{n-j+1}}{h_{n-j}} \right)^{\alpha_j} = h_n \prod_{j=0}^M \left( \frac{\text{tol}}{l_{n-j}} \right)^{\frac{\beta_j}{p}}. \quad (2.78)$$

By choosing  $P(z)$  and  $Q(z)$  carefully, better stability properties can be achieved (by virtue of filtering out high-frequency noise, for example). We list a couple of special cases of [Eq. \(2.78\)](#) below.

<sup>7</sup>In this context, the system is stable if it produces a finite output signal from a finite input.

**Integrating (I) controller:** this class of controllers have  $\beta_0 \neq 0$ , with all other parameters zero. This results in a slight modification of the classical adaptive stepsize algorithm:

$$h_{n+1} = \left(\frac{\text{tol}}{l_n}\right)^{\frac{\beta_0}{p}} h_n. \quad (2.79)$$

The word “integrating” comes from taking logarithms and writing

$$\ln h_{n+1} = \ln h_0 + \sum_{j=0}^n n \frac{\beta_0}{p} (\ln \text{tol} - \ln l_j), \quad (2.80)$$

and recognising that the last term is a discrete integral (of the control error signal).

**Proportional integrating (PI) controller:** A more stable controller is achieved by taking into account the error of the previous step by adding a term proportional to it (giving its name),

$$h_{n+1} = \left(\frac{\text{tol}}{l_n}\right)^{\frac{\beta_0}{p}} \left(\frac{\text{tol}}{l_{n-1}}\right)^{\frac{\beta_1}{p}} h_n. \quad (2.81)$$

A particularly effective PI controller for embedded Runge–Kutta methods is the PI3333, which has  $\beta_0 = 2/3$  and  $\beta_1 = -1/3$ . The different signs effectively filter out high-frequency oscillations in the stepsizes by not letting the stepsize react too drastically to an increase in the local error.

An appropriately chosen PI controller can drastically improve the efficiency of explicit Runge–Kutta methods on stiff sets of equations, and is therefore well worth considering as a “less bad” alternative to implicit methods. To demonstrate this, take the Van der Pol problem,

$$\begin{aligned} y_1' &= y_2, \\ y_2' &= \mu(1 - y_1^2)y_2 - y_1, \\ y_1(x=0) &= 1.5, \quad y_2(x=0) = 3.0, \end{aligned} \quad (2.82)$$

with  $\mu = 3000$ , on the interval  $[0, 0.5]$ . We solve this with a (4, 5)th order “custom” Runge–Kutta pair embedded in `oscode` and to be introduced in [Chapter 3](#), and set a relative tolerance of  $10^{-6}$ . We compare the stepsize-progression of the solver while using the classical stepsize-update algorithm and a PI3333 controller in [Fig. 2.4](#). The PI3333 controller can be seen to filter out high-frequency oscillations in the stepsize-progression, which otherwise result in a large number of rejected steps in the classical algorithm. This also ensures better control of the global error in problems where the scale of the stepsize needs to vary over the course of the solution, e.g. problems with regions of stiffness or oscillations. [Fig. 2.4](#) depicts preliminary results: while the current implementation of `oscode` uses the classical stepsize-update formula, in the future its




Fig. 2.4 Step size progression during the numerical solution of the Van der Pol problem. In the upper panel, the classical stepsize-update algorithm was used, and in the lower panel, a PI3333 controller. The underlying numerical method in both cases was a pair of (4, 5)th order Runge–Kutta formulae used in `oscodel` and introduced in [Chapter 3](#). By virtue of filtering out high-frequency oscillations in the signal (the stepsize-progression), the PI3333 controller achieves better a better acceptance ratio and stability.

Runge–Kutta method is planned to be updated with a set of stepsize controllers the user can choose from.

### 2.3.3 Dense output

The feature of producing dense output (also referred to as the *continuous extension* of a numerical method) is the ability of the numerical method to output solution values at user-specified values of the independent variable. Due to using adaptive stepsize control, a numerical method would evaluate the solution at the end of the integration range, and at intermediate points it determines based on its local error. The user may wish to enquire about the solution at different, and perhaps a large number of points, e.g. to get a better idea of its behaviour, for plotting purposes, to look for specific events, or to treat discontinuities. An option is to force stepsizes to be of certain lengths in order to naturally involve these points in the integration process, but this would increase the computational overhead significantly if the number of desired outputs is large. This problem is even more apparent if the user-specified tolerance is

large, the numerical method is high-order or particularly efficient, if the solution is smooth, all resulting in naturally large stepsizes. One therefore relies on *interpolation* methods to generate an approximate solution mid-step. This ideally should not involve more than a couple (or any!) extra evaluations of the right-hand-side of the ODE, to minimise additional computational effort, but should be of similar order of accuracy as the results obtained at the ends of steps. For linear multistep methods, *piecewise polynomials* are most commonly used, an example for which we review in detail in [Chapter 4](#). Methods that specialise in the efficient solution of highly oscillatory ODEs do not have a well-established set of tools to perform interpolation, and polynomials are unsuitable for this task if the method achieves stepsizes much larger than the wavelength of an oscillation (for the same reason Taylor-series based methods cannot increase their stepsizes beyond approximately one wavelength when solving the same ODE). Identifying the slowly-varying parts of the highly oscillatory solution as computed by a WKB-expansion based numerical method (the  $S_i(x)$  terms in [Eq. \(2.42\)](#)) and performing interpolation on these, we introduce an efficient continuous extension for a class of asymptotic solvers in [Chapter 4](#).



## Chapter 3

# An efficient method for solving highly oscillatory ODEs

Floyd could imagine a dozen things that could go wrong; it was little consolation that it was always the thirteenth that actually happened.

---

*Arthur C. Clarke, 2010: Odyssey Two*

This chapter discusses a novel numerical routine and open-source software (`oscodel`) for the efficient solution of one-dimensional, second-order, ordinary differential equations with rapidly oscillating solutions. The method is based on a Runge–Kutta-like stepping procedure that makes use of the Wentzel–Kramers–Brillouin (WKB) approximation to skip regions of integration where the characteristic frequency varies slowly. In regions where this is not the case, the method is able to switch to a made-to-measure Runge–Kutta integrator that minimises the total number of function evaluations. Three applications of the method are discussed in detail: the Airy equation and an equation exhibiting a burst of oscillations illustrate the error properties of the method, and the Schrödinger equation provides a physical example where the eigenvalues of a non-analytic quantum system are found via a shooting method.

The development of this method was inspired by the need to solve the Mukhanov–Sasaki equation in cosmology, which describes the evolution of cosmological perturbations, and has to be solved in order to compute the primordial power spectrum of the perturbations in some cosmological models of interest. This application will be discussed in [Chapter 6](#), after an introduction of the relevant background material in [Chapter 5](#).

### 3.1 Introduction

Runge–Kutta (RK) methods are powerful tools for numerically solving systems of first-order ordinary differential equations, and as such are often the default option in numerical routines for this task. There are cases however when more efficient methods are needed than Runge–Kutta, such as where the solution exhibits rapid oscillations. Problems classified as oscillatory are common in physics, yet the set of tools available to solve oscillatory systems efficiently is small, and problems are often treated on a case-by-case basis, using analytic approximations such as the Wentzel–Kramers–Brillouin (WKB) method [67].

There has been a proposal for an algorithm in pre-print [68] that generalises the Runge–Kutta stepping procedure, but uses the WKB approximation to forecast the solution instead of a Taylor expansion when the solution is highly oscillatory. The proposed algorithm was named RKWKB, and while it served as the theoretical foundation of the numerical method presented here, there are a number of key differences. Most importantly, I extended the algorithm so that it can be applied to damped oscillators, and to equations without closed-form frequency and first-derivative terms. I also made significant adjustments to the adaptive stepsize algorithm and the method to evaluate whether the WKB approximation is applicable at the current timestep. These modifications and other details of implementation are outlined in [Section 3.3](#).

The method was developed to deal with differential equations of the form

$$\ddot{x}(t) + 2\gamma(t)\dot{x}(t) + \omega^2(t)x(t) = 0, \quad (3.1)$$

where  $\gamma$  and  $\omega$  may or may not be expressed as a closed-form function of time. If they cannot be, but depend on time through a set of ‘background’ variables that can be obtained numerically, they may be supplied to the solver as array-like data structures sampled over time, as detailed in [Section 3.3.3](#). Since the efficiency of the solver relies on the WKB approximation being valid for a portion of the integration range, the solver is intended for problems where the frequency is slowly varying (relative to the timescales of the problem) for a part of the integration range. The numerical solver, `oscode`, is available on [GitHub](#)<sup>1</sup> and [PyPI](#)<sup>2</sup>, and can be accessed through a `C++` or `Python` interface.

This chapter is structured as follows. [Section 3.2](#) is a high-level overview of the algorithm, the implementation details of which are discussed in [Section 3.3](#). This is followed by example applications of the method in [Section 3.4](#). [Section 3.5](#) discusses factors the user needs to be aware of that might limit the performance of the solver, and [Section 3.6](#) concludes the chapter with a short summary.

---

<sup>1</sup><https://github.com/fruzsinaagocs/oscode>

<sup>2</sup><https://pypi.org/project/pyoscode/>

## 3.2 Computational strategy

The basis for this solver is the generalised stepping approach detailed in [68], which is summarised below. Having a numerical estimate for the solution  $x$  and its derivative  $\dot{x}$  at time  $t$ , the solution at a later time  $t + h$  is obtained. Then, using an error estimate on the proposed step, the stepsize  $h$  is updated such that the error estimate stays within a local tolerance limit. Such adaptive control of the stepsize is a requirement for robust numerical solvers. Starting from two functions  $f_{\pm}(t)$  that form an appropriate basis set for the true solution of the second-order differential equation, and are linearly independent at all  $t$ , the solver matches the correct solution and its derivative by linearly combining  $f_{\pm}$  and their derivatives:

$$x(t + h) = A_+ f_+(t + h) + A_- f_-(t + h), \quad (3.2)$$

and

$$\dot{x}(t + h) = B_+ \dot{f}_+(t + h) + B_- \dot{f}_-(t + h), \quad (3.3)$$

where

$$A_{\pm} = \frac{\dot{x}(t) f_{\mp}(t) - x(t) \dot{f}_{\mp}(t)}{\dot{f}_{\pm}(t) f_{\mp}(t) - \dot{f}_{\mp}(t) f_{\pm}(t)}, \quad (3.4)$$

and

$$B_{\pm} = \frac{\ddot{x}(t) \dot{f}_{\mp}(t) - \dot{x}(t) \ddot{f}_{\mp}(t)}{\ddot{f}_{\pm}(t) \dot{f}_{\mp}(t) - \ddot{f}_{\mp}(t) \dot{f}_{\pm}(t)}. \quad (3.5)$$

In the above,  $\ddot{x}(t)$  may be obtained from the differential equation itself, using  $x(t)$  and  $\dot{x}(t)$ . It is shown in [68] that the above procedure reduces to Euler's method in the limit of vanishing stepsize  $h$  and with the appropriate choice of  $f_{\pm}$ . The above approach therefore allows one to pick trial solutions  $f_{\pm}$  that approximate the true solution well over a larger range than an  $n^{\text{th}}$  order polynomial, which would be the choice for  $f_{\pm}$  in the case of an  $n^{\text{th}}$  order Runge–Kutta method.

As [68] suggests, the WKB method can be used to derive an analytic approximation to the true solution of a single oscillator on timescales much shorter than  $\frac{\omega}{\dot{\omega}}$ , the timescale on which the frequency changes. The WKB solutions, detailed in the following subsection, are ideal candidates for  $f_{\pm}$  over such timescales.

In general however, the frequency cannot be expected to vary slowly over the entire range of integration, and the WKB solutions might not always be a good choice for  $f_{\pm}$ . To counter this, a dynamic switching mechanism is included in the solver, which consists of attempting two steps of size  $h$  simultaneously. First, a Runge–Kutta step of order 5 is calculated (a ‘RK step’ hereafter), then a step using Eqs. (3.2) to (3.5), with  $f_{\pm}$  set to the WKB solutions (a ‘WKB step’). Based on the error estimates on each of these, the next stepsize,  $h^*$ , is computed. The step with the larger next predicted stepsize is chosen. This is to minimise the number of steps the solver needs to take to achieve a given local accuracy, and hence minimise runtime.

The step with the chosen method may be accepted or rejected, and the stepsize  $h$  increased or decreased.

The two methods are described in greater detail in the subsections that follow, with their error estimates discussed in [Section 3.3.4](#).

### 3.2.1 The Wentzel–Kramers–Brillouin approximation

Starting from the equation

$$\ddot{x}(t) + 2\gamma(t)\dot{x}(t) + \omega^2(t)x(t) = 0, \quad (3.6)$$

we wish to derive asymptotic expansions of the two independent solutions, in the limit that  $\omega$  is slowly varying relative to  $x$ , but  $\gamma$  need not be so. In the absence of a first-derivative term the derivation starts by introducing a power-counting parameter  $T$ , with  $T \gg 1$ :

$$\ddot{x} + T^2\omega^2x = 0. \quad (3.7)$$

If we now insert  $\gamma$  and allow it to vary on shorter timescales than  $\omega$ , the equivalent equation to consider is

$$\ddot{x} + 2\gamma\dot{x} + T^2\omega^2x = 0. \quad (3.8)$$

Following [\[67\]](#), one can then seek asymptotic approximations in the form of an exponential power series<sup>3</sup>

$$x(t) \sim \exp\left(T \sum_{n=0}^{\infty} S_n(t)T^{-n}\right). \quad (3.9)$$

Substituting [Eq. \(3.9\)](#) into [Eq. \(3.8\)](#), setting coefficients of powers of  $T$  to zero, one arrives at the recursion

$$\begin{aligned} \dot{S}_0(t) &= \pm i\omega, \\ \dot{S}_i(t) &= -\frac{1}{2\dot{S}_0} \left( \ddot{S}_{i-1} + 2\gamma\dot{S}_{i-1} + \sum_{j=1}^{i-1} \dot{S}_j\dot{S}_{i-j} \right). \end{aligned} \quad (3.10)$$

---

<sup>3</sup>Note that the following expression appears erroneously in [\[68\]](#), in that  $T$  should be replaced with  $T^{-1}$ .

The first four terms in the asymptotic series in the presence of a first-derivative term are

$$\begin{aligned}
S_0 &= \pm i \int \omega dt, \\
S_1 &= -\frac{1}{2} \ln \omega - \int \gamma dt, \\
S_2 &= \pm i \int -\frac{1}{2} \frac{\gamma^2}{\omega} - \frac{1}{2} \frac{\dot{\gamma}}{\omega} + \frac{3}{8} \frac{\dot{\omega}^2}{\omega^3} - \frac{1}{4} \frac{\ddot{\omega}}{\omega^2} dt, \\
S_3 &= \frac{1}{4} \frac{\gamma^2}{\omega^2} + \frac{1}{4} \frac{\dot{\gamma}}{\omega^2} - \frac{3}{16} \frac{\dot{\omega}^2}{\omega^4} + \frac{1}{8} \frac{\ddot{\omega}}{\omega^3}.
\end{aligned} \tag{3.11}$$

As [67] states, the WKB series is a singular perturbative expansion. The sum in Eq. (3.9) is usually divergent (unless it truncates) and needs to be truncated at some term in order to be a good approximation of  $x(t)$ . To use Eq. (3.9) as an approximate solution to Eq. (3.6), one needs to set  $T = 1$ . This is allowed despite having assumed  $T \gg 1$  (see, e.g. [67]) as long as the asymptotic inequalities

$$TS_0(t) \gg S_1(t) \gg \dots \gg T^{1-n} S_n(t) \tag{3.12}$$

hold uniformly within the interval  $[t, t+h]$ . If the asymptotic inequalities are satisfied and the first term not included in the asymptotic WKB series is small,

$$T^{-n} S_{n+1}(t) \ll 1, \tag{3.13}$$

$x(t) \sim \exp\left(T \sum_{i=0}^n T^i S_i(t)\right)$  is a good approximation.

To utilise the WKB solutions, we set  $f_{\pm}(t)$  to  $x(t)$  according to Eq. (3.9) with  $T = 1$ . Computing a WKB step from  $t$  to  $t+h$  thus involves

$$S_i(t+h) - S_i(t) = \int_t^{t+h} \dot{S}_i(t') dt'. \tag{3.14}$$

If the solver enters an integration region suitable for being approximated by WKB solutions, the stepsize  $h$  is expected to increase, and the error on the integrals Eq. (3.14) is expected to dominate the error on  $x$  and  $\dot{x}$  in WKB steps. Although in these regions  $\omega$  changes slowly, care needs to be taken to evaluate the integrals accurately. As  $\omega(t)$  and  $\gamma(t)$  may not be available in closed form, the integrals are computed numerically.

## 3.3 Implementation

### 3.3.1 Numerical integration and differentiation

I chose to calculate the integrals Eq. (3.14) using a method from the Gaussian quadrature family [69], Gauss–Lobatto integration [70]. Gaussian quadrature formulae work by modelling the integrand as a linear combination of appropriately chosen mutually orthogonal polynomials.

As a side effect, a certain class of integrands make the integral exact (in the Gauss–Lobatto case, polynomials of degree  $2n - 3$ , where  $n$  is the number of abscissas). Typically the remainder in such methods is proportional to a higher order ( $2n - 2$  for Gauss–Lobatto) derivative of the integrand, which one can expect to be small in integration regions of interest, where WKB is a good approximation and several oscillations can be stepped over, such that  $h \gg \frac{2\pi}{\omega}$ . This property makes Gaussian quadrature superior to integrating the  $\dot{S}_i$  with a Runge–Kutta step, as the latter would approximate the integral from  $t$  to  $t + h$  with a Taylor expansion around  $t$ , with an error as some power of  $h$ . Gaussian quadrature methods are also desirable because they converge exponentially fast with  $n$ , due to the order of the method increasing with  $n$  as well as the density of points of evaluation [62]. This makes them a better choice than Newton–Cotes methods with equally spaced abscissas, such as the trapezoidal rule or Simpson’s method.

Gauss–Lobatto integration with  $n = 6$  was chosen in particular because the abscissas it uses include the beginning and endpoints of integration,  $t$  and  $t + h$ . This makes it a FSAL (first same as last) method, and one could design a 5<sup>th</sup> order, 6-stage Runge–Kutta formula based on the same abscissas, minimising the number of evaluations of  $\omega(t)$  and  $\gamma(t)$  during a single step of the algorithm (WKB and RK). The remainder on a Gauss–Lobatto integral is given analytically, but since it involves the  $(2n - 2)$ <sup>th</sup> derivative of the integrand, it is more common to be estimated as the difference between the results with  $n$  and  $n - 1$  abscissas.

The integrands in Eq. (3.11) contain derivatives of  $\omega$  and  $\gamma$ , which may not be available in closed form, and hence will also be calculated numerically. Since  $\omega$  and  $\gamma$  are already to be evaluated at 9 distinct points (the Gauss–Lobatto abscissas) per step for the integrals in  $S_0$  and  $S_1$ , it is worth re-using these values and derive finite difference formulae using them as stencil points [71], by solving

$$\begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix} = \frac{1}{h^D} \begin{bmatrix} s_1^0 & \dots & s_n^0 \\ \vdots & \ddots & \vdots \\ s_1^{n-1} & \dots & s_n^{n-1} \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ \vdots \\ D! \\ \vdots \\ 0 \end{bmatrix}, \quad (3.15)$$

where the  $s_i$  define the stencil such that the points of evaluation are  $t_i = t + s_i h$ ,  $D$  is the order of derivative desired, and the  $D!$  is the  $(D + 1)$ <sup>th</sup> entry in the vector on the right-hand-side. The  $w_i$  are the resulting weights of the function evaluations:

$$\left. \frac{d^D f}{dt^D} \right|_t = \sum_{i=1}^n w_i f(t_i). \quad (3.16)$$

The finite difference formulae above work by cancelling the first  $D$  terms in the Taylor expansion of  $f$  around  $t$  and setting the coefficient of the  $(D + 1)$ <sup>th</sup> term to 1. These give  $D + 1$  constraints, but since we are free to choose the weights of  $n_s$  evaluations of  $f$  (where  $n_s$  is the number of

stencil points), we can cancel a further  $n_s - D - 1$  terms in the Taylor series and thus get a result that is accurate to  $\mathcal{O}(h^{n_s - D})$ . While  $h$  is expected to be large in a region of integration where WKB is a good approximation, the coefficient multiplying  $h^{n_s - D}$  is expected to be small (since the derivatives of  $\omega$  are expected to be small), and therefore we argue that finite differences is an acceptable method for estimating the derivatives of  $\omega$  and  $\gamma$ .

### 3.3.2 Explicit Runge–Kutta formulae based on Gauss–Lobatto stencil points

For the present problem, the system to be solved can be rewritten as

$$\dot{\mathbf{y}} = \mathbf{f}(\mathbf{y}, t), \quad (3.17)$$

with

$$\mathbf{y} = (x, \dot{x}) = (y_1, y_2), \quad (3.18)$$

$$\mathbf{f} = (y_2, -\omega^2(t)y_1 - 2\gamma(t)y_2), \quad (3.19)$$

which is suitable for solution via a Runge–Kutta method. We aim to construct a pair of *explicit* Runge–Kutta formulae for advancing the solution of Eq. (3.18) and estimating its error, since explicit methods can be evaluated iteratively.

For most combinations of the number of stages  $s$  of a Runge–Kutta method and its desired order of accuracy  $N$ , the order constraints do not pin down all entries in the Butcher tableau, and the leftover degrees of freedom are often fixed by minimising the coefficient of the leading-order term in the local error. An efficient embedded (4,5) pair developed by [72] demonstrates this, and is used in `rksuite` [73] (used by the `NAG Library` [21]) as one of the possible Runge–Kutta formulations. For the present problem, we are interested in solving the order constraints of a 6-stage, 5<sup>th</sup> order method<sup>4</sup>, with the  $c_i$  set to the Gauss–Lobatto abscissas for  $n = 6$ , and a 4-stage, 4<sup>th</sup> order method with its  $c_i$  equal to the Gauss–Lobatto abscissas for  $n = 5$  with the exception of the midpoint. This way we can recycle the evaluations of  $\omega$  and  $\gamma$  at the abscissas to calculate the integrals in Eq. (3.14), estimate their errors, take a Runge–Kutta step in  $x, \dot{x}$

<sup>4</sup>The highest order a 6-stage method can achieve is 5, as proven in [31].

0				
$\frac{1}{2} \left(1 - \sqrt{\frac{3}{7}}\right)$	$\frac{1}{2} \left(1 - \sqrt{\frac{3}{7}}\right)$			
$\frac{1}{2} \left(1 + \sqrt{\frac{3}{7}}\right)$	$-\frac{1}{4} \left(3 + 5\sqrt{\frac{3}{7}}\right)$	$\frac{1}{4} (5 + \sqrt{21})$		
1	$-\frac{1}{4} (3 + 7\sqrt{21})$	$-\frac{1}{4} (21 + 5\sqrt{21})$	$\frac{1}{14} \left(-1 + \sqrt{\frac{3}{7}}\right)$	
	$-\frac{1}{12}$	$\frac{7}{12}$	$-\frac{1}{12}$	$-\frac{1}{12}$

Table 3.1 Butcher tableau for the 4-stage, 4<sup>th</sup> order Runge–Kutta method used in the solver, based on 4 out of 5 stencil points of a Gauss–Lobatto quadrature with  $n = 5$  stencil points.

$0$						
$\frac{1}{2} \left(1 - \sqrt{\frac{1}{3} + \frac{2\sqrt{7}}{21}}\right)$	0.117472					
$\frac{1}{2} \left(1 - \sqrt{\frac{1}{3} - \frac{2\sqrt{7}}{21}}\right)$	-0.186247	0.543632				
$\frac{1}{2} \left(1 + \sqrt{\frac{1}{3} - \frac{2\sqrt{7}}{21}}\right)$	-0.606430	1	0.249046			
$\frac{1}{2} \left(1 + \sqrt{\frac{1}{3} + \frac{2\sqrt{7}}{21}}\right)$	2.899356	-4.368525	2.133806	0.217890		
$1$	18.679963	-28.850577	10.720534	1.414741	-0.964661	
	0.112755	0	0.506557	0.048300	0.378474	-0.046089

Table 3.2 Butcher tableau for the 6-stage, 5<sup>th</sup> order Runge–Kutta method used by the solver. The timepoints of evaluation are the 6 stencil points used for Gauss–Lobatto quadrature with  $n = 6$ .

and get their error estimates all at the same time. The order constraints for this system can be solved symbolically with no leftover degrees of freedom, demonstrated in [74]. The resulting coefficients are summarised in the form of Butcher tableaux: Table 3.2 contains the coefficients of the 4-stage, 4<sup>th</sup> order method, and Table 3.2 contains those of the 6-stage, 5<sup>th</sup> method.

### 3.3.3 Defining $\omega(t)$ and $\gamma(t)$

In many problems of interest, the frequency and the friction term will not be explicit functions of time, but functions of variables that depend on time through a set of differential equations that may only be solved numerically. The algorithm requires the values of  $\omega(t)$  and  $\gamma(t)$  to be known at 9 distinct points in each step along the solution, but is otherwise blind to how the functions are defined. In order for the solver (and in particular Gauss–Lobatto integration) to work reliably, the frequency and friction terms need to be known at any timepoint within the integration range to high (at least 1 in  $10^9$ ) accuracy. High accuracy is required here as during a WKB step (where  $h, \omega \gg 1$ ), any interpolation error in  $\omega(t), \gamma(t)$  is amplified by roughly a factor of the stepsize:

$$x(t) \sim e^{i \int_t^{t+h} \omega(t) dt} \approx e^{ih\omega}, \quad (3.20)$$

$$\left| \frac{\Delta x}{x} \right| \approx \left| \frac{ih\Delta\omega e^{ih\omega}}{e^{ih\omega}} \right| = h\Delta\omega. \quad (3.21)$$

For convenience the solver has been set up such that the user can provide values of the functions (or their natural logarithms) as vectors evaluated on an evenly spaced, monotonically increasing grid over time. It will then carry out linear interpolation whenever a function evaluation is required. The even spacing in the independent variable is a requirement for the sake of speed, as it simplifies the search for the nearest gridpoints ahead of the interpolation.

If evaluation on an evenly spaced grid is not possible or the grid cannot be made fine enough for linear interpolation to be sufficiently accurate, the user may define  $\omega(t)$  and  $\gamma(t)$  as interpolated functions using a suitable interpolation method.

### 3.3.4 Estimating the error in RK and WKB steps

The RK and WKB steps each give  $x$  and  $\dot{x}$  (referred to by their subscripts), and the difference between the 4<sup>th</sup> and 5<sup>th</sup> order RK steps gives an error on them. Estimating the error on a WKB step is less straightforward, and the next paragraph explains how two distinct error estimates can arise which dominate in different regions of the solution. The present method calculates these both and uses whichever value is larger.

The obvious equivalent error estimate on WKB steps,  $\Delta x_{\text{WKB}}$  and  $\Delta \dot{x}_{\text{WKB}}$ , is the difference between an  $N^{\text{th}}$  and  $(N - 1)^{\text{th}}$  order estimate, where  $N$  refers to the highest-order  $S$ -term in Eq. (3.11) included in the WKB expansion. This estimate is a good proxy for the validity of the WKB approximation because it can signal the breakdown of the relations Eq. (3.12), but in a region where they hold, it is expected that the numerical error in the  $S_i(t)$  will dominate  $\Delta x$  and  $\Delta \dot{x}$ . We can therefore estimate the error on the WKB step arising from the imperfect numerical integration of  $\dot{S}_i(t)$  as

$$\Delta x_{\text{WKB}} = A_+ \Delta f_+ + A_- \Delta f_-, \quad (3.22)$$

$$\Delta f_{\pm} = f_{\pm} \sum_{i=0}^n \Delta [S_i]_t^{t+h}, \quad (3.23)$$

and

$$\Delta \dot{x}_{\text{WKB}} = B_+ \Delta \dot{f}_+ + B_- \Delta \dot{f}_-, \quad (3.24)$$

$$\Delta \dot{f}_{\pm} = \Delta f_{\pm} \frac{\dot{f}_{\pm}}{f_{\pm}}. \quad (3.25)$$

Note that in the above,  $f$  and its derivatives are evaluated at  $t + h$  according to Eqs. (3.4) and (3.5), and that it is assumed that the numerical integration of  $\dot{S}_i$  are the *only* sources of error, i.e. the  $\dot{S}_i$  can be acquired perfectly.

### 3.3.5 Stepping procedure

Let us summarise the different error estimates:

- $\Delta x_{\text{RK}}, \Delta \dot{x}_{\text{RK}}$ : error on RK step,
- $\Delta x_{\text{WKB}}, \Delta \dot{x}_{\text{WKB}}$ : error on WKB step from computing  $[S_i]_t^{t+h}$  numerically,
- $\Delta x_{\text{WKB}}^t, \Delta \dot{x}_{\text{WKB}}^t$ : error on WKB step from truncation of WKB series.

In order for the solver to switch successfully to the most suitable method dynamically, and adapt the stepsize to stay within the error bound required, it has to determine two things:

1. Which step (RK or WKB) to choose that yields the largest possible next stepsize within acceptable tolerance?
2. What should the size of the next step be?

The answer to 1. requires forecasting the error progression of both methods with the stepsize, i.e. requires knowledge of  $\Delta x(h)$  and  $\Delta \dot{x}(h)$ . For the RK step this behaviour is known to be a power-law, and for WKB steps we shall assume two separate power-laws with different exponents  $n_{\text{WKB}}$  and  $n_{\text{WKB}}^t$ , for when the dominant error on WKB steps arises from the numerical integrals and the truncation of the asymptotic series, respectively. First, the dominant error on each type of step is determined,

$$\Delta_{\text{RK}} = \max(\epsilon, \Delta x_{\text{RK}}, \Delta \dot{x}_{\text{RK}}), \quad (3.26)$$

$$\Delta_{\text{WKB}} = \max(\epsilon, \Delta x_{\text{WKB}}, \Delta \dot{x}_{\text{WKB}}, \Delta x_{\text{WKB}}^t, \Delta \dot{x}_{\text{WKB}}^t), \quad (3.27)$$

where  $\epsilon$  is a small number close to machine precision, for safety. The type of dominant error on the WKB step, ‘*truncation*’, or ‘*integral*’ is recorded. Starting from a current stepsize  $h$ , the largest possible steps within the error bound  $tol$  are then

$$h_{\text{RK}} = h \times \left( \frac{tol}{\Delta_{\text{RK}}} \right)^{1/n_{\text{RK}}},$$

$$h_{\text{WKB}} = h \begin{cases} \left( \frac{tol}{\Delta_{\text{WKB}}} \right)^{1/n_{\text{WKB}}^t}, & \text{if ‘truncation’,} \\ \left( \frac{tol}{\Delta_{\text{WKB}}} \right)^{1/n_{\text{WKB}}}, & \text{otherwise.} \end{cases} \quad (3.28)$$

The step with the larger stepsize will then be chosen as a trial step, but is not yet accepted. The next stepsize is then predicted. If the chosen method is RK, this next stepsize is simply

$$h_{\text{next}} = h_{\text{RK}}. \quad (3.29)$$

If the chosen method is WKB however (i.e. the truncated WKB series was deemed sufficient to approximate the solution), the error arising from truncation of the WKB series will be ignored:

redefine  $\Delta_{\text{WKB}}$  as

$$\Delta_{\text{WKB}} = \max(\epsilon, \Delta x_{\text{WKB}}, \Delta \dot{x}_{\text{WKB}}), \quad (3.30)$$

$$\text{then } h_{\text{next}} = h \left( \frac{tol}{\Delta_{\text{WKB}}} \right)^{1/n_{\text{WKB}}}. \quad (3.31)$$




Fig. 3.1 Schematic plot of the assumed error progression in RK and WKB steps with increasing stepsize  $h$ . After the steps have been calculated from  $t$  to  $t + h_{\text{current}}$ , the errors of each method are shown by points  $A$ ,  $B$  and  $C$ , the latter two arising from the truncation of the WKB asymptotic series and the numerical integrals present in the series, respectively. The dominant type of error on the WKB step in this case is the ‘truncation’. Assuming power-law behaviour in the errors for both steps with different exponents for each type of error, the next largest stepsize within the required tolerance ‘tol’ would be  $h_{\text{RK}}$  and  $h_{\text{WKB}}$ , marked by points  $A'$  and  $B'$ . Since  $h_{\text{WKB}} > h_{\text{RK}}$ , the algorithm in this case chooses the WKB step. The size of the next step is therefore determined solely on the basis of the ‘integral’ error, marked by  $C$ , and is going to be the projection of  $C'$ . Since this next stepsize is larger than the previous, the step is accepted.

This process is illustrated in Fig. 3.1. Finally, if  $h_{\text{next}} > h$ , the current error did not exceed the tolerance limit and the step is accepted. Otherwise the step is rejected and one must ensure that the step is re-attempted with sufficiently small  $h$ . The new stepsize in both cases is calculated via

$$h_{\text{next}} = h \times \begin{cases} \left(\frac{\text{tol}}{\Delta_{\text{RK}}}\right)^{1/(n_{\text{RK}}-1)}, \\ \left(\frac{\text{tol}}{\Delta_{\text{WKB}}}\right)^{1/(n_{\text{WKB}}-1)} & \text{if ‘integral’,} \\ \left(\frac{\text{tol}}{\Delta_{\text{WKB}}^t}\right)^{1/(n_{\text{WKB}}^t-1)} & \text{if ‘truncation’}. \end{cases} \quad (3.32)$$

This ensures  $h$  is decreased after rejected steps and increased following accepted ones.

The slopes of the errors as functions of  $h$  in Fig. 3.1 were not chosen at random. The error in a 6-stage, 5<sup>th</sup>-order RK method goes as  $h^5$  for  $h < 1$  and  $h^6$  for  $h > 1$ . The error arising

from truncation of the WKB series, upon entering a region well-approximated by the WKB expansion, is expected to be proportional to  $h$ . This can be understood starting from the relative error on  $x$  based on [67],

$$\frac{\Delta x}{x} \sim S_{N+1}, \quad (3.33)$$

for an  $N^{\text{th}}$  order WKB estimate. For  $N \geq 1$ ,  $S_{N+1}$  is a numerical integral of a small and nearly constant quantity, and is therefore  $\propto \dot{S}_{N+1}h$ . The error on WKB steps arising from the integrals  $[S_i]_t^{t+h}$  are on the other hand expected to go roughly as the errors on the integrals themselves (see Eq. (3.22)). Although more difficult to predict, this is expected to be dominated by the imperfect evaluations of the integrands, which contain numerical derivatives. The largest of these are the first derivatives, which will have an error  $\propto h^{n_s-1}$ . Since the algorithm uses the  $n = 6$  Gauss–Lobatto evaluations to calculate all derivatives, we set  $n_s = 6$ .

By the above reasoning, the algorithm by default has

$$n_{\text{RK}} = 5, \quad n_{\text{WKB}} = 5, \quad n_{\text{WKB}}^t = 2, \quad (3.34)$$

but the user can set these parameters to better fit the problem in question. For example, for optimal step acceptance/rejection ratio, for all burst examples in Section 3.4.2  $n_{\text{WKB}} = 8$ ,  $n_{\text{WKB}}^t = 1$  was used.

### 3.3.6 Algorithm summary

In summary, the algorithm goes through the following steps:

1. Stepping from  $t$  to  $t + h$ , evaluate  $\omega$  and  $\gamma$  at the Gauss–Lobatto stencil points for  $n = 6$  and  $n = 5$ , a total of 9 different points.
2. Use the Butcher tableaux Table 3.1 to construct a RK step in  $x$  and  $\dot{x}$ , and use the difference as the error  $\Delta x_{\text{RK}}$ .
3. Use finite difference methods to evaluate all necessary derivatives of  $\omega$  and  $\gamma$  needed for the derivatives of terms  $(\dot{S}_i)$  in the WKB series Eq. (3.11).
4. Use Gauss–Lobatto quadrature with  $n = 6$  and  $n = 5$  to evaluate the terms in the WKB series and their errors, taken as the difference.
5. Construct an  $N^{\text{th}}$  and  $(N - 1)^{\text{th}}$  order WKB step in  $x$  and  $\dot{x}$ .
6. Compute the ‘truncation’ error on WKB steps as the difference between the  $N^{\text{th}}$  and  $(N - 1)^{\text{th}}$  order estimates, and the ‘integral’ from Eq. (3.24).
7. Find the dominant error and its type based on Eq. (3.28).
8. Predict the next stepsize,  $h_{\text{next}}$ , based on Eq. (3.31).

9. **If**  $h_{\text{next}} > h$ , accept the step, and update  $x$ ,  $\dot{x}$ , and  $t$ .
10. **Otherwise**, reject the step and calculate  $h_{\text{next}}$  according to [Eq. \(3.32\)](#).
11. Update  $h$  to  $h_{\text{next}}$ .
12. Repeat steps 1–11.

## 3.4 Example applications

### 3.4.1 Airy equation

One of the simplest non-trivial equations the method's efficiency can be tested on is the Airy equation,

$$\ddot{x} + tx = 0, \quad (3.35)$$

which has the solutions  $\text{Ai}(-t)$  and  $\text{Bi}(-t)$ . We set initial conditions such that the solution becomes  $\text{Ai}(-t) + i\text{Bi}(-t)$ . All derivatives of  $\omega$  decrease with time, hence the algorithm is expected initially to perform RK steps and at a later time switch to WKB, making the Airy equation an ideal example to test both the accuracy of the WKB steps and the stepsize-update procedure. This behaviour is illustrated qualitatively in [Fig. 3.2](#). The second panel in [Fig. 3.2](#) then details the error properties of the RK and WKB phases, and shows that whilst the global relative error grows in RK steps, it levels off once the WKB phase is entered. In contrast, for a pure RK method, the stepsize decreases whilst the relative error continues growing. The RKWKB-based solver (`oscode`) has no difficulty stepping through the Airy solution until times as late as  $10^8$ , at which point the stepsize becomes too large to store  $[S_i]_t^{t+h}$  with the required precision. This limitation is discussed in [Section 3.5](#).

### 3.4.2 Burst equation

The numerical solution of the following equation illustrates the switching mechanism between RK/WKB steps especially well. Consider

$$\ddot{x} + \frac{n^2 - 1}{(1 + t^2)^2} x = 0, \quad (3.36)$$

which satisfied by

$$x(t) = \frac{\sqrt{1 + t^2}}{n} (\cos(n \arctan t) + i \sin(n \arctan t)), \quad (3.37)$$

a solution characterised by a burst of approximately  $n/2$  oscillations in the region  $|t| < n$ . The exact solution and numerical estimates of the solution at the steps taken by `oscode` are shown in [Fig. 3.3](#).




Fig. 3.2 Numerical solution of the Airy equation obtained with the solver (dots and triangles), overlaid on the true solution as computed by the `boost` maths library. The algorithm exhibits a clear switch from taking RK steps to WKB steps at around  $t \approx 4$ , as expected. Despite the  $t$ -axis being logarithmic, the stepsize-increase is clearly visible as time increases, and the rate of change of  $\omega$  decreases. Also shown is the accumulation of relative error during the numerical solution of the Airy equation until late times, showing the difference between a purely RK-based approach and RKWKB (`oscode`). The relative tolerance was set to be  $10^{-4}$ , which the RKWKB solution does not exceed, but navigates such that the largest possible steps are taken whilst staying within this limit. In contrast, a solver taking only RK steps quickly decreases its steps whilst accumulating error.

[Fig. 3.3](#) also shows the error accumulated in the numerical solution of this example. This clearly shows that once the burst of oscillations is encountered, taking WKB steps becomes more efficient, and the solver allows the stepsize to grow until the local error reaches its tolerance limit. It then keeps the local error at this limit whilst traversing as many oscillations as possible. The global error is also seen to level off, at a slightly higher value than the local tolerance. To demonstrate that as many oscillations are stepped over as possible, [Fig. 3.4](#) shows the number of oscillations traversed during a single step of the solver as a function of time having a sharp peak near  $t = 0$ , where it is able to leap through  $10^4$  oscillations.

The robustness of the algorithm was tested by monitoring the numerical error as a function of time for relative tolerances ranging from  $10^{-6}$  to  $10^{-4}$ , shown in [Fig. 3.5](#). The global error in




Fig. 3.3 Numerical solution of the burst equation with  $n = 40$ . A relative tolerance of  $10^{-3}$  was set, and the equation was solved from  $t = -2n$  to  $t = 2n$ . For the sake of resolving the oscillations around  $t = 0$  on the top panel, a zoomed-in view of the solution is shown.




Fig. 3.4 Number of oscillations stepped over in a single step, while solving the burst equation with  $n = 10^5$ , and a relative tolerance of  $10^{-4}$ .




Fig. 3.5 Progression of the relative error in the burst equation with  $n = 10^5$ , with different settings of the local relative tolerance ‘rtol’.




Fig. 3.6 Relative runtime of solving the burst equation from  $t = -2n$  to  $t = 2n$ , with  $n$  varying from  $10^1$  to  $10^{10}$ , and the relative tolerance, ‘rtol’ from  $10^{-4}$  to  $10^{-6}$ . The runtimes are referenced to the median of the  $n$ -range and relative tolerance of  $10^{-4}$ , as indicated by the dotted lines.

all of the above examples reaches a constant value of  $\sim 10 \times \text{rtol}$  by the end of the oscillatory phase.

Finally, **Figs. 3.6** and **3.7** show that the algorithm is efficient over a range of values of  $n$  (which determine the total number of oscillations) and tolerances. The algorithm shows a slow, 4-fold runtime increase over 9 orders of magnitude change in the number of oscillations, which is due to the increase in WKB steps needed to traverse the oscillatory region, shown in **Fig. 3.8**. **Fig. 3.7** also reveals that the algorithm is most efficient in the relative tolerance range of  $10^{-6} - 10^{-4}$ . For tolerances lower than this, a 4-5<sup>th</sup> order RK pair is not generally recommended.

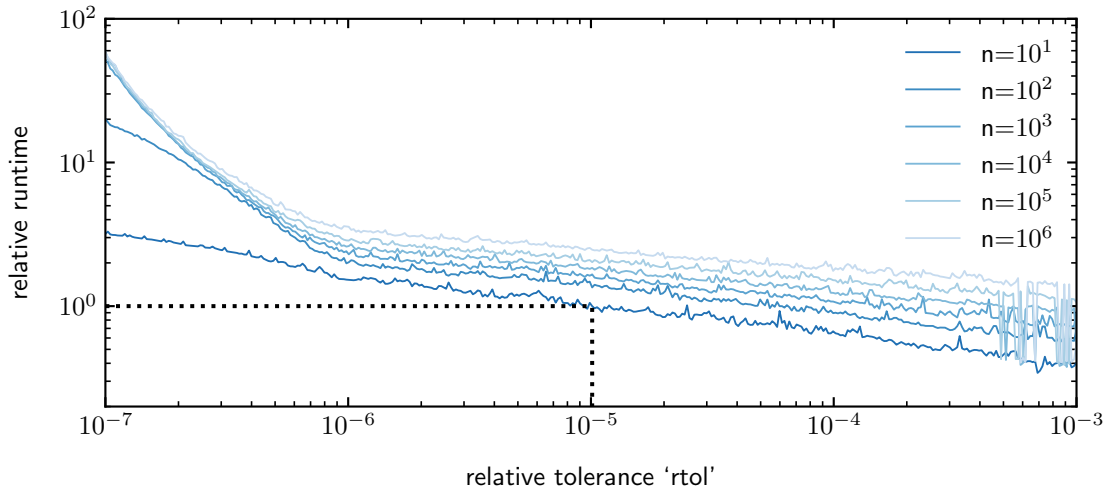


Fig. 3.7 Relative runtime of solving the burst equation, as a function of the relative tolerance set (with the absolute tolerance set to 0). The different curves show different values of  $n$ , proportional to the total number of oscillations traversed.




Fig. 3.8 Step breakdown in solving the burst equation from  $t = -2n$  to  $t = 2n$ , with  $n$  varying from  $10^1$  to  $10^{10}$ , and the relative tolerance, 'rtol' set to  $10^{-4}$ .

### 3.4.3 Schrödinger equation

The one-dimensional time-independent Schrödinger equation for a potential  $V(x)$  takes the form

$$\Psi''(x) + 2m(E - V(x))\Psi(x) = 0, \quad (3.38)$$

where  $\hbar$  has been set to 1. The WKB method's original use was to compute approximate solutions of Eq. (3.38), which suggests that `oscode` can be used as an alternative to traditional methods (such as the Numerov method [75]) to calculate fast numerical solutions. Starting with

an analytic example, [Fig. 3.9](#) shows the numerical evaluation of the energy eigenfunction  $\Psi_n$  for the  $n^{\text{th}}$  energy level in a harmonic potential well  $V(x) = x^2$ , for a range of  $n$ -s including high-energy excited states. [Fig. 3.9](#) clearly shows that `oscde` only needs to take a few steps once inside the potential well, suggesting that computation time is greatly reduced relative to purely Runge–Kutta based approaches. In this example the analytic solution for the eigenfunctions were available and were used to set the values of  $\Psi$  and  $\Psi'$  at the integration boundaries.

In a general potential well, analytic solutions are not accessible and the energy eigenvalues are unknowns to be computed. Shooting methods [\[76\]](#) are frequently used to estimate the eigenvalues in such cases. We shall use one such method to find the energies of the quantum harmonic oscillator with quartic anharmonicity, which has the potential

$$V(x) = x^2 + \lambda x^4. \quad (3.39)$$

An initial guess for the eigenvalue,  $E$ , is made. We start integration from points  $\pm x_0$  outside the potential on either side of  $x = 0$ , where  $E \ll V(x)$ , using the initial conditions  $\Psi(\pm x_0) = 0$  and  $\Psi'(\pm x_0) = 1$ . We integrate towards the inside of the potential well in order to avoid contamination of the exponentially decaying solution by the growing mode when one integrates away from the well. The first initial condition is a good approximation far outside the potential well, and  $\Psi'$  can be chosen arbitrarily as it accounts to a choice of normalisation. The two numerical solutions,  $\Psi_L$  and  $\Psi_R$  meet at an intermediate point  $x_1$ . At  $x_1$ , both  $\Psi$  and  $\Psi'$  must be continuous if  $E$  is an eigenvalue. Therefore the normalisation-independent quantity

$$\left| \frac{\Psi'_L}{\Psi_L} - \frac{\Psi'_R}{\Psi_R} \right| \quad (3.40)$$

is minimised as a function of  $E$ . A few examples of the eigenvalues thus computed are presented in [Table 3.3](#), alongside their matching values from [\[77\]](#). Note that in order to get equivalent eigenvalues, we need to set  $m = 0.5$ . The eigenvalues are in good agreement up to highly excited states.

### 3.5 Limitations

When applying `oscde` to a problem, it is worth considering whether the solver's performance would be limited by strict accuracy requirements or by a non-ideal choice of independent or dependent variable.

As shown in [Fig. 3.7](#), the algorithm's runtime scales up gently in the relative tolerance range  $[10^{-4}, 10^{-6}]$ . The user is therefore advised to use the solver if such accuracies are acceptable for the problem in question. If the problem requires  $\text{rtol} < 10^{-6}$ , one would need a higher-order Runge–Kutta pair as an alternative solver to WKB, such as a (7,8) pair used by the NAG Library.




Fig. 3.9 Energy eigenfunctions in a harmonic potential well. In units of  $m = \hbar = 1$  and with a potential  $V(x) = x^2$ , the  $n^{\text{th}}$  level has energy  $\sqrt{2}(n - 1/2)$ . The wavefunctions in this potential are given analytically in terms of the Hermite polynomials, and are plotted in black. Numerical integration was started from both sides of  $x = 0$ , from well outside the potential (where  $E \ll V(x)$ ), until  $x = 0.5$ . The initial conditions were set using the analytic solution for  $\Psi$  and  $\Psi'$ . The relative tolerance was set to be  $10^{-3}$ .

$n$	$E_n$	$E_n^*$
0	1.392353	1.392352
1	4.648815	4.648813
2	8.6550501	8.6550500
3	13.156806	13.156804
4	18.0577	18.0576
15	88.6104	88.6103
16	96.1291	96.1296
17	103.793	103.795
18	111.6025	111.6020
19	119.5440	119.5442
50	417.05620	417.05626
100	1035.5440	1035.5442
1000	21932.7848	21932.7840
10000	471103.81	471103.80

Table 3.3 Energy eigenvalues of the quantum harmonic oscillator with quartic anharmonicity. The left-hand column  $E_n$  shows the eigenvalues found by the shooting method based on `oscode`, to be compared with the right-hand column  $E_n^*$ , which lists the (rounded) results of [77].

As mentioned in Section 3.4.1, the solver will not be able to fulfil the accuracy requirements if at any time the integral(s)  $[S_i]_t^{t+h}$  exceed  $\sim 10^{12}$ . Care needs to be taken especially with the first term in the WKB series,  $\pm i \int_t^{t+h} \omega dt$ , as this is expected to be the largest in a region where the WKB approximation is appropriate. The reason underlying this limit is that the solver needs to compute the exponential of this large imaginary term, which requires large accuracy modulo  $2\pi$ . Storing such large numbers accurately is limited by machine (double) precision, and the solver might start accumulating error. In Section 3.4.1 the stepsize and frequency become so large at  $t > 10^8$  that this limit is reached.

Finally, the solver is only efficient if in some region the frequency is slowly varying, therefore care needs to be taken to choose an appropriate independent-dependent variable pair if the problem so allows.

## 3.6 Conclusions

This chapter introduced a novel numerical solver for second-order, ordinary differential equations that can be written in the form of a one-dimensional oscillator, with a time-varying frequency and friction term that do not necessarily have a closed form. It was illustrated qualitatively through several examples that the solver is extremely efficient at traversing oscillatory regions of the solution if the frequency of oscillations varies slowly, by exploiting the WKB approximation in these regions. Through Fig. 3.8 we explored this feature quantitatively: we varied the characteristic frequency of oscillations in the burst equation Eq. (3.36), and measured the number of steps the solver took, broken down into RK/WKB steps. One can read off that the number of steps, proportional to the solver's runtime, scales with  $n$ , the approximate frequency of oscillations, roughly as  $\mathcal{O}(n^{1/15})$ . In theory the solver could achieve  $\mathcal{O}(1)$  runtime (a runtime

independent of frequency of oscillations) since the accuracy of the WKB approximation is independent of the magnitude of the frequency, but in practice this is not achieved due to finite precision representation, which e.g. limits the maximum length of a step, the accuracy of numerical derivatives and integrals, etc. For comparison, the scaling of *any* linear multistep method is  $\mathcal{O}(n)$ , since they all rely on a polynomial representation of the oscillatory solution. Bremer’s phase function method from [58] has a theoretical  $\mathcal{O}(1)$  runtime as well, but is only usable in the highly oscillatory regime of the ODE. In contrast, the solver presented here can also detect regions where the WKB approximation is not valid, and can dynamically switch to a Runge–Kutta integrator.

Due to the order of the RK method used and the number of terms retained in the WKB expansion, the method proved to be most efficient in the relative tolerance range  $10^{-6} \lesssim \text{rtol} \lesssim 10^{-4}$ . Within this range, for reasons listed above, `oscode` outperforms conventional (linear multistep) methods in highly oscillatory problems roughly by a factor of the frequency of oscillations.



# Chapter 4

## Dense output

This chapter presents a method for constructing a continuous extension (otherwise known as dense output) of a numerical routine in the special case of the numerical solution being a scalar-valued function exhibiting rapid oscillations. Such cases call for numerical routines that make use of the known global behaviour of the solution to forecast the solution at each step of the independent variable. An example is `oscde`, a numerical routine introduced in [Chapter 3](#) which uses the WKB approximation when the solution oscillates rapidly and otherwise behaves as a RK solver. Polynomial interpolation is not suitable for producing the solution at an arbitrary point mid-step, since efficient numerical methods based on the WKB approximation will step through multiple oscillations at once. The method proposed here is based on identifying the slowly-varying aspect of the highly oscillatory solution, and interpolating that instead. For a method based on the WKB approximation, the slowly-varying functions from which the full solution can be reconstructed are the amplitude and phase, i.e. the real and imaginary terms in the WKB expansion. These terms contain integrals that need to be evaluated numerically in each step, and are typically computed via Gaussian quadrature. This chapter takes Gauss–Lobatto quadrature as an example to show how Gaussian quadrature formulae may be used for interpolation, and estimates the error on this interpolation method. Finally, it draws attention to previous work on the continuous extension of Runge–Kutta formulae, and formulates the continuous extension of a RK method based on Gauss–Lobatto quadrature nodes, thus describing how to generate dense output from each of the methods underlying `oscde`.

While the examples presented in this chapter are specific to `oscde`, its underlying principles can be used to generate dense output from a more general class of methods that rely on asymptotic expansions.

## 4.1 Introduction

During the numerical solution of an ODE for an initial value problem, algorithms attempt to control the global error by adapting their stepsize - the spacing between values of the independent variable  $x_i$  at which they naturally choose to evaluate the solution  $y_i$  [62]. The user might, however, wish to *specify* the points at which the output is evaluated, for e.g. plotting the solution, event location, or treating discontinuities. The natural steps the numerical algorithm takes may be too large for such purposes. Artificially decreasing the steps (e.g. by integrating from one output point to the next) would be inefficient and increase computation time for a large number of outputs. One therefore relies on interpolation methods to generate an approximate solution mid-step. Dense output should be produced with minimal computational overhead, i.e. with as few additional evaluations of the ODE as possible, and at a similar level of accuracy as that achieved at the natural steps.

Interpolation of slowly changing functions in the context of dense output is well established [62, 69, 70], and often uses (piecewise) polynomials. There are instances, however, when polynomials are not applicable, one example being when the function to be interpolated undergoes several oscillations between two points of evaluation. There are several methods available to solve ODEs with highly oscillatory solutions efficiently (see, e.g. [54, 58, 78, 79]), which have in common that the global behaviour of the solution informs computation: they all exploit the prior knowledge that the solution is oscillatory. As a result, these algorithms may only ‘naturally’ evaluate the solution every couple of oscillations, greatly reducing the number of steps taken, and requiring a different approach to computing dense output.

The present approach is based on identifying the slowly varying terms in the ODE from which the oscillatory solution can be constructed, and performing interpolation (based on polynomials) on these terms. More specifically, this chapter discusses the generation of dense output for `oscode` [1], a numerical method based on the WKB approximation [67], but the methodology shown is applicable to any solver using asymptotic expansions with non-oscillatory terms. Since `oscode` uses the WKB expansion to trace the solution in its oscillatory regimes but relies on a RK method otherwise, results from the continuous extension of RK methods are also reviewed and an existing method due to Shampine [39] is adapted to compute dense output in the non-oscillatory regimes.

[Section 4.2](#) discusses dense output from steps when `oscode` uses the WKB expansion to forecast the solution, based on the continuous extension of Gaussian quadrature. [Section 4.3](#) contains a derivation of a continuous extension of `oscode`’s custom RK method. Examples of dense output from both the WKB and RK regimes are shown in [Section 4.4](#), followed by a brief summary in [Section 4.5](#).

Throughout this chapter the word integration is used to describe the process of solving an ODE, whereas the numerical evaluation of integrals will be referred to as quadrature.

## 4.2 Dense output from the WKB expansion

To construct the solution at an arbitrary point along the length of a step at the cost of just arithmetic operations (and no extra evaluations of terms in the ODE), I first review which numerical methods are used to construct the solution at the very end of the step, and which evaluations are available upon the completion of successful steps as a result.

`oscode` uses a WKB approximation expanded up to and including the  $S_3$  term. If a WKB step runs from  $x$  to  $x + h$ , the integrals in Eq. (3.11) will have those limits. For reasons stated in Section 3.3.1, the integrals are computed using a form of Gaussian quadrature called Gauss–Lobatto rules. The quadrature method operates with  $n$  evaluations of the integrand, two of which are always at the start and end of the step. To obtain an error estimate on the integrals, for each Gauss–Lobatto quadrature carried out with  $n$  nodes there is one computed with  $n - 1$  nodes, the difference between the two giving the error estimate. As will be discussed in Section 4.2.1, Gaussian quadrature fits an interpolant to the integrand based on evaluations of the integrand at the  $n$  nodes, hence it is possible to (1) evaluate this approximate integrand at any point, and (2) evaluate the integral itself at any point mid-step. This gives dense output on the  $S_0$ ,  $S_2$  terms, their derivatives, and the second term in  $S_1$  together with its derivative. Dense output from the derivatives of the  $S_i$  are required for constructing  $y'$  at any point.

With the strategy for obtaining dense output from the numerical integrals discussed, the next elements to consider are the various derivatives of  $\omega$  and  $\gamma$  appearing in the WKB expansion. Since these are not available through the ODE directly, `oscode` uses the same trick RK methods are based upon: it combines evaluations of terms in the ODE ( $\omega$ ,  $\gamma$ ) at various values of  $x < x_i < x + h$  such that when Taylor expanded, all terms lower than, and a maximal number of terms higher than the required derivative order vanish [1, 71]. This amounts to finding the coefficients for a finite difference equation which has fixed stencil points. As discussed, computation of the Gauss–Lobatto integrals  $\int \omega dx$  and  $\int \gamma dx$  requires  $2n - 3$  intermediate evaluations of  $\omega$  and  $\gamma$  per step, with `oscode` using  $n = 6$ . These 9 evaluations in total can be used to obtain numerical derivatives sufficiently accurately at the nodes of Gauss–Lobatto integration, which can then be fit with the Gauss–Lobatto interpolant, and evaluated at arbitrary points.

### 4.2.1 Gaussian quadrature

The strategy behind Gaussian quadrature is to mimic the integrand  $f(x)$  with a linear combination of orthogonal polynomials which have known integrals. The polynomials are chosen to best represent the integrand, different choices defining different quadrature rules. The linear combination  $F(x)$  is fit to the integrand using a number of evaluations of the latter, such that they match at the abscissas  $x = x_i$ :

$$\lim_{x \rightarrow x_i} F(x) = f(x_i). \quad (4.1)$$

The integral then takes the form

$$\begin{aligned}\tilde{I}(f) &= \int_a^b F(s)ds = \frac{b-a}{2} \int_{-1}^1 F(s(x))dx \\ &= \frac{b-a}{2} \sum_{i=1}^n w_i f(x_i)\end{aligned}\tag{4.2}$$

for a Gaussian method of  $n$  nodes, with  $w_i$  being the weights of the method. Note that the integration limits have been shrunk down to  $(-1, 1)$  by a linear transformation of the independent variable,

$$s = \frac{b-a}{2}x + \frac{b+a}{2}.\tag{4.3}$$

The abscissas and weights can be chosen such that the order of the method, which in the context of quadrature means the degree for which all polynomials are integrated exactly by Eq. (4.2), is much larger than  $n$ , e.g.  $2n - 1$  for Gauss–Legendre rules, and  $2n - 3$  for Gauss–Lobatto. Consequently, the remainder or error on the integral goes as a higher-order derivative of the integrand,

$$R(f) \propto f^{(m)}(\xi), \quad -1 < \xi < 1,\tag{4.4}$$

with  $m = 2n$  and  $m = 2n - 2$  for Gauss–Legendre and Gauss–Lobatto rules, respectively. This makes Gaussian quadrature an especially attractive choice for performing the various numerical integrals appearing in the WKB approximation: if the WKB approximation is valid, the derivatives of the integrands involved will generally be small, therefore the applicability of the WKB approximation and Gaussian quadrature align well.

#### 4.2.2 Dense output from Gauss–Lobatto integration

If the form of the interpolating polynomial  $F(x)$  used by a quadrature rule is known, the numerical integral can be carried out until an arbitrary point within the integration limits,  $a < c < b$ , straightforwardly. Abscissas for Gauss–Lobatto integration are chosen to be the integration limits themselves ( $x = \pm 1$ ), and the roots of the polynomials

$$P'_{n-1}(x),\tag{4.5}$$

where  $P_n(x)$  is the  $n$ th Legendre polynomial. The orthogonal polynomials associated with this quadrature rule then have to be

$$(1-x^2)P'_{n-1}(x),\tag{4.6}$$

From this, the interpolation polynomial can be uniquely constructed as

$$F(x) = \frac{(1-x^2)P'_{n-1}(x)}{2P'_{n-1}(-1)(1+x)}f(-1) + \frac{(1-x^2)P'_{n-1}(x)}{2P'_{n-1}(1)(1-x)}f(1) \quad (4.7)$$

$$+ \sum_{i=2}^{n-2} \frac{(1-x^2)P'_{n-1}(x)}{P''_{n-1}(x_i)(1-x_i^2)(x-x_i)}f(x_i) \quad (4.8)$$

Using l'Hôpital's rule one can verify that this indeed reduces to  $f(x_i)$  in the limit  $x \rightarrow x_i$ .

In the continuous extension of this method, the weights  $w_i$  become position-dependent. If one were to evaluate the solution at  $x = c$ , with  $a < c < b$ , Eq. (4.7) has to be integrated up until that point,

$$\begin{aligned} \tilde{I}(f, c) &= \int_a^c F(s)ds = \frac{b-a}{2} \int_{-1}^{\tilde{c}} F(s(x))dx \\ &= \frac{b-a}{2} \sum_{i=1}^n w_i(\tilde{c})f(x_i), \end{aligned} \quad (4.9)$$

where the weights are given by

$$w_1(\tilde{c}) = \frac{1}{2P'_{n-1}(-1)} \int_{-1}^{\tilde{c}} (1-x)P'_{n-1}(x)dx, \quad (4.10)$$

$$w_i(\tilde{c}) = \frac{1}{2P''_{n-1}(x_i)(1-x_i^2)} \int_{-1}^{\tilde{c}} \frac{(1-x^2)P'_{n-1}(x)}{x-x_i}dx, \quad (4.11)$$

$$w_n(\tilde{c}) = \frac{1}{2P'_{n-1}(1)} \int_{-1}^{\tilde{c}} (1+x)P'_{n-1}(x), dx \quad (4.12)$$

with  $2 \leq i \leq n-1$ .

### 4.2.3 Error bound on dense output from Gauss–Lobatto integration

Section 4.2.1 claimed that the remainder of Gauss–Lobatto quadrature is proportional to  $f^{(2n-2)}$ , where  $f(x)$  is the integrand and  $n$  nodes are used to fit the interpolant to  $f$ . This is based on Peano's error representation [80]:

**Theorem 4.2.1** (Peano's error bound). *Suppose  $R(P) = 0$  holds for all polynomials  $P \in \Pi_n$ , i.e. every polynomial of degree less than or equal to  $n$  is integrated exactly by the quadrature rules. Then for all functions  $f \in C^{n+1}[a, b]$ ,*

$$R(f) = \int_a^b f^{(n+1)}(t)K(t)dt, \quad (4.13)$$

where

$$K(t) = \frac{1}{n!} R_x[(x-t)_+^n], \quad (x-t)_+^n = \begin{cases} (x-t)^n & \text{if } x \geq t, \\ 0 & \text{if } x < t, \end{cases} \quad (4.14)$$

and

$$R_x[(x-t)_+^n] \quad (4.15)$$

is the remainder from the quadrature on  $(x-t)_+^n$  when the latter is considered as a function of  $x$ .

$K(t)$  is termed the Peano kernel of the remainder operator  $R$ . If the Peano kernel has constant sign on  $[a, b]$  (which is the case for Gauss–Lobatto rules), it follows from the mean-value theorem of integral calculus that

$$R(f) = f^{(n+1)}(\xi) \int_a^b K(t) dt \quad \text{for some } \xi \in (a, b). \quad (4.16)$$

The kernel  $K(t)$  is the same for all integrands  $f$  as long as the same quadrature rules are used. Therefore the integral  $\int_a^b K(t) dt$  can be evaluated for any  $f$ , the simplest choice being the polynomial  $f(x) = x^{n+1}$ . Substituting this into [Eq. \(4.16\)](#) and eliminating the kernel integral, one obtains

$$R(f) = \frac{R(x^{n+1})}{(n+1)!} f^{(n+1)}(\xi) \quad \text{for some } \xi \in (a, b). \quad (4.17)$$

The elements of the above derivation do not depend on the integration limits  $a, b$ , and so the expression [Eq. \(4.17\)](#) for the remainder holds when applied to dense output, with one notable difference. Gauss–Lobatto quadrature, when performed ‘in full’ on the interval  $[a, b]$ , integrates polynomials  $P \in \Pi_{2n-3}$  exactly (the proof of which can also be found in [\[80\]](#)), for which it is sometimes called a  $2n-3$  order method. However, the same is not true if the upper integral limit is changed to  $c < b$ , while the nodes stay the same, as done in our computation of dense output. The proof of Gauss–Lobatto rules being of order  $2n-3$  hinges on its basis polynomials being orthogonal on the interval  $[a, b]$ , which does not generally hold true on  $[a, c]$ . All one can say about the order of Gauss–Lobatto integration on the interval  $[a, c]$  is that it is uniformly  $n-1$ , since the interpolant [Eq. \(4.7\)](#) is a polynomial of degree  $n-1$ . Using [Eq. \(4.17\)](#), one can numerically (or otherwise) compute the Peano bound for the dense output from Gauss–Lobatto integration, which is seen in [Figs. 4.1](#) and [4.2](#).

### 4.3 Dense output from a Runge–Kutta method based on Gauss–Lobatto nodes

If it is suspected that the WKB approximation will not apply well in some regions of the integration range of an ODE, an alternative method is needed to efficiently traverse those regions. Runge–Kutta formulae are versatile and perform well in regions in which the solution

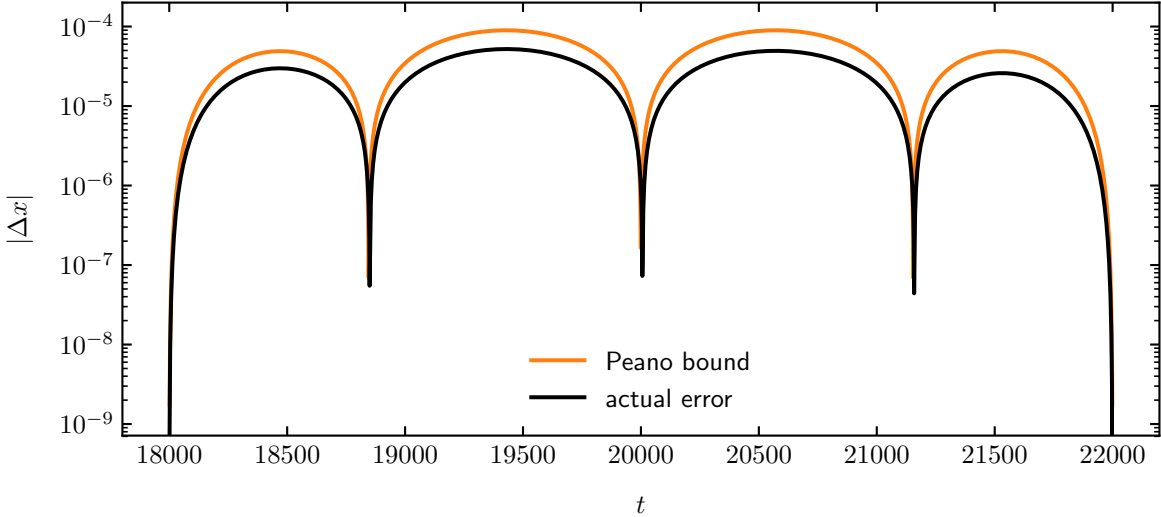


Fig. 4.1 Base-10 logarithm of the residuals from the partial Gauss-Lobatto integration of  $\omega = \sqrt{t}$  against the expected Peano error bound.




Fig. 4.2 Actual residuals from the partial Gauss-Lobatto integration of  $\omega = \sqrt{t}$  against the expected Peano error bound, shown with a linear  $y$ -axis..

is not oscillatory. For the sake of solving the ODE at hand with a Runge–Kutta method, it is rewritten as

$$y' = G(x, y). \quad (4.18)$$

In this work, a 6-stage, 5th order explicit method is considered which uses the nodes of the 6th order Gauss–Lobatto quadrature rules as its stages, i.e. the  $c_i$  shall coincide with the  $x_i$ . This is so that when one computes a WKB approximation of the solution from  $x = x_n$  to  $x = x_n + h$ , a Runge–Kutta step of the same size can be computed at the cost of a few arithmetic operations. The method however holds for any Runge–Kutta formula of the same  $s$  and  $n$ , and can be generalised to others.

An approach to extend a Runge–Kutta formula developed by Horn [81] is to take a hypothetical step from  $x_n$  to  $x_n + \sigma h$ :

$$\begin{aligned} k_0 &= G(x_n, y_n), \\ k_i^* &= G\left(x_n + c_i^* \sigma h, y_n + \sigma h \sum_{j=1}^{i-1} a_{ij}^* k_j^*\right), \quad i = 1..s^*, \\ y_{n+1}^* &= y_n + \sigma h \sum_{i=1}^s b_i^* k_i^*. \end{aligned} \quad (4.19)$$

The positions of  $G$ -evaluations can be made identical to those in the original Runge–Kutta formula if

$$\begin{aligned} c_i^* &= \frac{1}{\sigma} c_i, \\ a_{ij}^* &= \frac{1}{\sigma} a_{ij}. \end{aligned} \quad (4.20)$$

Depending on the required order  $n^*$  of the new solution  $y_{n+1}^*$ , it may be necessary to extend the Butcher tableau and add more stages to the formula,  $s^* > s$ . One can then derive a set of coefficients  $b_i^*(\sigma)$  for each  $\sigma$  that give the required order, and thus have a continuous extension. This approach, however, yields an answer such that

$$\lim_{\sigma \rightarrow 1} y_{n+1}^*(\sigma) \neq y_{n+1}, \quad (4.21)$$

that is, the dense output would be discontinuous across steps. This is an undesirable property as ideally the dense output should be  $C^1$  (continuously differentiable).

An alternative was proposed by Shampine in [39] which uses Horn’s interpolant to obtain a solution at one  $\sigma$  at a similar order as that at achieved at the end of the step, then performs local polynomial interpolation based on this intermediate solution  $y_{n+1}^*$ , and information available at both ends of a step after a successful step. Horn has shown that for the Runge–Kutta–Fehlberg 4(5) formula, an intermediate solution at order 4 is available at  $\sigma = 0.6$  with  $s^* = s$ , i.e. for free. She did so by deriving the necessary order constraints for  $b_i^*(\sigma)$  in a concise manner, which for

0						
$\frac{1}{2} \left(1 - \sqrt{\frac{1}{3} + \frac{2\sqrt{7}}{21}}\right)$	0.117472					
$\frac{1}{2} \left(1 - \sqrt{\frac{1}{3} - \frac{2\sqrt{7}}{21}}\right)$	-0.186247	0.543632				
$\frac{1}{2} \left(1 + \sqrt{\frac{1}{3} - \frac{2\sqrt{7}}{21}}\right)$	-0.606430	1	0.249046			
$\frac{1}{2} \left(1 + \sqrt{\frac{1}{3} + \frac{2\sqrt{7}}{21}}\right)$	2.899356	-4.368525	2.133806	0.217890		
1	18.679963	-28.850577	10.720534	1.414741	-0.964661	
	0.112755	0	0.506557	0.048300	0.378474	-0.046089

Table 4.1 Butcher tableau for the 6-stage, 5<sup>th</sup> order Runge-Kutta method used by `oscode` in [1]. The points of function evaluations coincide with the abscissas of Gauss-Lobatto quadrature with  $n = 6$ .

$b_1^*$	0.2089555395
$b_2^*$	0
$b_3^*$	0.7699501023
$b_4^*$	0.009438629906
$b_5^*$	-0.003746982422
$b_6^*$	0.01540271068

Table 4.2 Modified Butcher tableau entries for obtaining a 4th order estimate at  $x = x_n + \sigma h$ , with  $\sigma = 0.58665886817$ . The rest of the Butcher tableau entries match those in Table 4.1.

$n^* = 4$ ,  $s^* = s = 6$  are

$$\sum_{i=1}^6 b_i^* c_i^j = \frac{\sigma^j}{j+1} \quad \text{for } j = 0, 1, 2, 3, \quad (4.22)$$

$$\sum_{i=3}^6 b_i^* a_{i2} = 0.$$

Following Horn's procedure with the Butcher tableau entries derived in [1] and summarised in Table 4.1, one finds that for this custom Runge-Kutta formula there is a free 4th order solution available at

$$\sigma = 0.58665886817 \quad (4.23)$$

with the associated  $b_i^*$  coefficients summarised in Table 4.2. With this intermediate point, the following information is available:  $y_n$ ,  $y'_n$  through  $k_0$ ,  $y_{n+1}$  and  $y'_{n+1}$  through  $k_7$ , and  $y_{n+1}^*$ . With these five constraints it is possible to carry out local quartic interpolation, meaning interpolation would be carried out for each successful step separately. This procedure, by relying on the solution and its derivative at both ends of a step, is ensured to provide a piecewise polynomial interpolant that is globally  $C^1$ . This then ensures that the entire solution is  $C^1$  continuous, since the interpolant is clearly  $C^1$  in WKB steps, and stays true across step boundaries by virtue of the stepping scheme matching initial conditions (the value of the solution  $y(x)$  and its derivative  $y'(x)$ ) at step boundaries, as explained in Section 3.2.

In practice, local interpolation can be performed efficiently as follows. Let the interpolating polynomial take the form

$$y_{n+1} = y_n + a_1\sigma + a_2\sigma^2 + a_3\sigma^3 + a_4\sigma^4 \quad (4.24)$$

The constraints at the ends of a step and at the intermediate point  $x_n + \sigma h$  can be written purely in terms of the  $k_i$ , the Butcher tableau coefficients, and the  $a_1$ – $a_4$ :

$$\begin{aligned} k_0 &= a_1 \\ \sum_{i=1}^7 b_i k_i &= \sum_{i=1}^4 a_i, \\ k_7 &= \sum_{i=1}^4 i a_i, \\ \sigma \sum_{i=1}^7 b_i^* k_i &= \sum_{i=1}^4 \sigma^i a_i. \end{aligned} \quad (4.25)$$

From this one can extract the matrices  $M$  and  $S$  such that

$$\begin{aligned} MQ^T &= SK, \\ Q &= [a_1 \ a_2 \ a_3 \ a_4], \\ K^T &= [k_1 \ k_2 \ k_3 \ k_4 \ k_5 \ k_6 \ k_7]. \end{aligned} \quad (4.26)$$

Note that if [Eq. \(4.18\)](#) is a vector equation, the  $a_i$  and  $k_i$  are promoted to column vectors and  $Q$  and  $K$  become matrices, but the notation used here still holds. We can then extract a constant matrix  $P$  from  $Q$  via

$$Q = K^T P, \quad (4.27)$$

which lets us compute the dense output at the cost of the arithmetic operations underlying

$$y = hQZ, \quad (4.28)$$

with

$$Z = \begin{bmatrix} \sigma_1, \sigma_2, \dots, \sigma_N \\ \sigma_1^2, \sigma_2^2, \dots, \sigma_N^2 \\ \sigma_1^3, \sigma_2^3, \dots, \sigma_N^3 \\ \sigma_1^4, \sigma_2^4, \dots, \sigma_N^4 \end{bmatrix}, \quad (4.29)$$

$N$  being the number of points we require output at.  $P$  can be pre-computed for each Runge–Kutta formula, and in our case is given in [Table 4.3](#). This procedure is used in open-source scientific computing libraries such as `scipy` [\[82\]](#).

1	-2.48711376	2.42525041	-0.82538093
0	0	0	0
0	3.78546138	-5.54469086	2.26578746
0	-0.27734213	0.74788587	-0.42224334
0	-2.94848704	7.41087391	-4.08391191
0	0.50817346	-1.20070313	0.64644062
0	1.4193081	-3.8386162	2.4193081

Table 4.3 Pre-computed  $P$ -matrix for a 5th order, 6-stage Runge–Kutta formula based on 6 Gauss–Lobatto nodes, as defined in Eq. (4.27).

## 4.4 Examples

A few examples of dense output from the RK and WKB methods used in `oscode` are shown below.

The Airy equation,

$$y'' + xy = 0, \quad (4.30)$$

describes an oscillator with a frequency  $\omega(x) = \sqrt{x}$  that rises with increasing  $x$ , but at a decreasing rate. As  $x$  increases, therefore, it becomes more favourable for `oscode` to use the WKB approximation to forecast the solution, and it switches over to do so from using a RK method at around  $x \sim 5$ , as shown in Fig. 4.3. The figure shows dense output from both the initial RK and the late WKB phase on top of the analytic solution. The error on the dense output from the numerical solution of the Airy equation (for a longer integration range) is shown in Fig. 4.4, exhibiting a pattern similar to that seen in Fig. 4.1. The similarity is due to the leading term in the numerical error coming from the  $S_0 = \int \omega(x)dx$  term of the WKB expansion. At large values of  $x$  the WKB approximation is valid, and so successive terms decrease rapidly in the expansion.

Fig. 4.5 shows dense output from the numerical solution of the equation

$$y'' + \frac{n^2 - 1}{(1 + x^2)^2} y = 0 \quad (4.31)$$

with  $n = 40$ . The parameter  $n$  governs the number of oscillations the solution exhibits around  $x = 0$ . When `oscode` solves equation, it uses the WKB approximation around  $x = 0$  and uses RK otherwise, switching in a symmetrical manner. The error properties and performance of `oscode` are explored in [1] using the example of Eq. (4.31), but note that due to no extra evaluations of the terms in this ODE being made during the computation of dense output, the latter does not increase the overall computing time significantly.




Fig. 4.3 Dense output from `oscode` solving the Airy equation. As the frequency becomes larger but more slowly-changing, the method switches from using the RK method to the WKB approximation and the distance between natural steps (orange dots) of the algorithm increases. The first segment of solid, orange line on top of the analytic solution (black line) denotes dense output from the RK method used by `oscode`, with the second orange segment showing dense output throughout a WKB step.




Fig. 4.4 Absolute (top) and relative (bottom) residuals from dense output of the Airy equation, relative to the analytic solution. Note that the error at the ends of natural steps (orange dots) are much smaller than throughout the steps, due to Gauss–Lobatto integration being much higher ( $2n - 3$ ) order at the end of a step than mid-way ( $n - 1$ ).




Fig. 4.5 Dense output from `oscode` solving Eq. (4.31), an equation exhibiting a burst of oscillations.

## 4.5 Conclusions

Dense output (evaluation of the numerical solution of an ODE at user-specified points) cannot always be constructed by polynomial interpolation between the natural steps of a numerical algorithm. One example is the efficient solution of ODEs with highly oscillatory solutions, for in this case the algorithm may traverse many oscillations in a single step. Such equations are extremely common in physics, and often form computational bottlenecks in e.g. the forward-modelling phase of Bayesian inference, when tackled with conventional (Runge–Kutta-like) methods.

Out of the methods available to efficiently solve highly oscillatory ODEs, `oscode` uses the WKB approximation to forecast the solution many wavelengths ahead if the characteristic frequency of oscillations changes on a much longer timescale than the solution itself, and otherwise behaves as a RK solver. This chapter introduced procedures to generate dense output from each of the methods underlying `oscode`.


In a region where the one-dimensional solution of an ODE oscillates and the WKB approximation is valid, interpolation is performed with known methods on the slowly-changing frequency ( $\omega$ ) and damping ( $\gamma$ ) terms in the ODE, and construct the solution using the WKB approximation. The numerical integrals of  $\omega$  and  $\gamma$  appearing in the WKB expansion can be computed efficiently and to high accuracy with Gaussian quadrature methods. As an example, dense output from one method of the Gaussian family (Gauss–Lobatto quadrature) was summarised, and an error bound for the output was derived.

In regions where the RK method is more appropriate, I reviewed existing techniques to obtain a continuous extension of RK methods, and demonstrated them on the example of a RK method based on nodes of 6-point Gauss–Lobatto quadrature.

## Chapter 5

# Cosmology of primordial perturbations

This chapter serves as a brief introduction to topics in primordial cosmology that appear throughout this thesis, providing theoretical background and establishing notation. For reference, it relies heavily on the excellent books [83–88] and reviews [5, 89–91].



Randall Munroe, [xkcd](#)

### 5.1 The universe at zeroth order

On the largest lengthscales, the simplest assumption one could make about the distribution of structure in the universe is that it is uniform at all points in space and in all directions, i.e. it is *homogeneous* and *isotropic*. This is equivalent to assuming that no points in the universe are special, a philosophy known as the *cosmological principle*. Observations confirm that this assumption is remarkably good<sup>1</sup>, to the extent that the universe at large redshift – while perturbations are small – can be modelled exceptionally well with perturbation theory. The perturbative treatment first requires identifying a dominant, exactly solvable problem (the

<sup>1</sup>Although clearly not perfect, otherwise we would not exist.

isotropic, homogeneous, expanding universe), and then developing expressions for the perturbed system in terms of power series in some small parameter(s). This chapter follows the same procedure, and discusses the dynamics of a smooth universe before considering perturbations.

### 5.1.1 Gravity

The dominant interaction that dictates how structure in the universe forms and how spacetime itself evolves is gravity, described exceptionally well by the theory of general relativity<sup>2</sup>. It is fundamentally distinct from the other forces (electromagnetic, strong, and weak) in that one can select a frame of reference, a freely falling frame (which does not rotate and occupies a small region of spacetime), in which it ‘disappears’. In this frame, an object under the influence of gravity moves in a straight line with a constant velocity, and obeys the laws of special relativity. This thought experiment considered by Albert Einstein led him to the proposal of the *equivalence principle* and general relativity. In a nutshell, matter induces spacetime curvature, and objects follow extremal paths, *geodesics*, on this curved spacetime. This deviation from the straight path between two points in vacuum (Minkowski spacetime) is what we perceive as gravity.

The joint time-evolution of matter and spacetime geometry is found by extremising the action

$$S = \int d^4x \sqrt{g} \mathcal{L}, \quad (5.1)$$

where  $g = |\det g_{\mu\nu}|$ ,  $d^4x \sqrt{g}$  is the relativistic (Lorentz-invariant) volume element, and  $\mathcal{L}$  the relevant Lagrangian density. The Lagrangian density is usually split into a gravitational and a matter part,

$$\mathcal{L} = \mathcal{L}_G + \mathcal{L}_M, \quad (5.2)$$

where  $\mathcal{L}_G = \frac{1}{2} m_{\text{P}}^2 R$ , and  $R$  is the Ricci scalar, a function of the metric tensor and its derivatives. Without matter, the action  $S_{\text{EH}} = \int d^4x \sqrt{g} \mathcal{L}_G$  would be called the Einstein–Hilbert action, which is the simplest<sup>3</sup> valid action one could write down with the right mathematical properties. This splitting implicitly assumes *minimal coupling*: no direct coupling between fields representing the matter content of the universe and the metric. One could introduce coupling with terms such as  $\xi \phi_{\text{M}}^2 R$  with  $\xi$  a constant, but many such non-minimally coupled theories can be brought to the minimally coupled form by redefining the matter field(s)  $\phi_{\text{M}}$ . Extremising the action via the Euler–Lagrange equations yields Einstein’s field equations:

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R = \frac{1}{m_{\text{P}}^2} T_{\mu\nu}, \quad (5.3)$$

<sup>2</sup>Although modifications of the original theory which cannot be ruled out experimentally do exist.

<sup>3</sup>There is actually a simpler term one could add to the action, which is just the Lorentz-invariant volume element multiplied by a constant,  $S = \int d^4x \sqrt{g} (R - 2\Lambda)$ . The  $\Lambda$  is known as the *cosmological constant*, and it can be absorbed into a cosmological fluid with an effective energy density (or stress–energy tensor).

where  $G_{\mu\nu}$  is the Einstein tensor,  $R_{\mu\nu}$  the Ricci tensor, and  $T_{\mu\nu}$  the stress–energy (also referred to as energy–momentum) tensor. The Einstein equations capture the effect the presence of matter has on spacetime curvature: the left-hand-side (LHS) of [Eq. \(5.3\)](#) contains quantities related to spacetime curvature, and the stress–energy tensor is determined by the matter content via

$$T_{\mu\nu} = -\frac{2}{\sqrt{g}} \frac{\delta}{\delta g^{\mu\nu}} (\sqrt{g} \mathcal{L}_M). \quad (5.4)$$

In 4 dimensions and with  $g_{\mu\nu}$  being symmetric, the Einstein equations appear to give 10 distinct equations. These are not independent however, as symmetries known as the Bianchi constraints only leave 6 degrees of freedom. Moreover, the Bianchi identity leads to the conservation of energy–momentum:

$$\nabla_\mu G^{\mu\nu} = 0 \implies \nabla_\mu T^\mu_\nu = 0. \quad (5.5)$$

For any given matter configuration, the Einstein equations give us the metric, which in turn tells us about particle trajectories. The Einstein equations, however, can only be solved exactly in a select few cases. Fortunately, a homogeneous, isotropic universe – in line with the cosmological principle – is among them.

### 5.1.2 Dynamics of a homogeneous, isotropic, expanding universe

#### The metric tensor

The most general metric allowed by the symmetries of a homogeneous, isotropic, time-dependent universe is the Friedmann–Lemaître–Robertson–Walker metric, which has the form [\[92–96\]](#)

$$ds^2 = dt^2 - a(t)^2 \left[ d\chi^2 + S(\chi)^2 (d\theta^2 + \sin^2(\theta) d\phi^2) \right], \quad (5.6)$$

where  $S(\chi)$  is determined by the spatial curvature parameter  $K \in \{0, \pm 1\}$ ,

$$S(\chi) = \begin{cases} \sin(\chi) & \text{closed, } K = 1, \\ \chi & \text{flat, } K = 0, \\ \sinh(\chi) & \text{open, } K = -1. \end{cases} \quad (5.7)$$

The quantities appearing in this metric are defined in [Table 5.1](#). The scale factor  $a(t)$  multiplying the spatial part of the metric encodes expansion: any two objects that are at rest (have constant  $\chi$ ,  $\theta$ , and  $\phi$  coordinates) have a physical distance, as opposed to coordinate distance, that grows in proportion to  $a(t)$ . The rate of expansion is expressed via the Hubble function<sup>4</sup>, defined as

$$H \equiv \frac{\dot{a}}{a}, \quad (5.8)$$

<sup>4</sup>We choose to call  $H(t)$  the Hubble function rather than Hubble parameter to emphasise its time-dependence. Parameters are normally thought of as being constant, which could lead to confusion.  $H_0$ , the Hubble function's value today, may be referred to as a parameter.

symbol	meaning
$t$	cosmic time
$a(t)$	scale factor
$\chi$	comoving radial coordinate
$\theta$	polar angle
$\phi$	azimuthal angle
$K$	spatial curvature parameter

Table 5.1 Quantities appearing in the FLRW metric.

where the overdot denotes differentiation with respect to cosmic time  $t$ .

### Conformal time and horizons

A further useful quantity to introduce is conformal time  $\eta$ :

$$\eta \equiv \int^t \frac{d\sigma}{a(\sigma)}, \quad (5.9)$$

with which the metric becomes

$$ds^2 = a(\eta)^2 [d\eta^2 - dX^2], \quad (5.10)$$

where the comoving spatial part of the metric has been absorbed into  $dX^2$ . Note that a conformal transformation<sup>5</sup> can bring us back to the Minkowski metric, the two are therefore conformally equivalent. In conformal time, null geodesics (including the worldlines of photons, or lightcones) become straight lines, which makes conformal time an ideal coordinate for plotting spacetime diagrams in terms of.

In an FRLW universe that starts from a singularity, there is a maximum comoving radial distance particles could have travelled at any given time, given by the *particle horizon*:

$$\chi_p = \int_{t_i}^{t_0} \frac{d\sigma}{a(\sigma)}, \quad (5.11)$$

It is useful to relate this quantity to the *comoving Hubble radius*  $\chi_H$ :

$$\chi_p = \int_{\ln a_i}^{\ln a_0} (aH)^{-1} d \ln a = \int_{\ln a_i}^{\ln a_0} \chi_H d \ln a. \quad (5.12)$$

<sup>5</sup>A shrinking or stretching of the manifold via a transformation of the form  $g_{\mu\nu}(x) \rightarrow \bar{g}_{\mu\nu}(x) = \Omega^2(x)g_{\mu\nu}(x)$ , not to be confused with a coordinate transformation  $x^\mu \rightarrow x'^\mu$ .

A more intuitive explanation of the two horizons is that if two events are separated by a comoving distance greater than  $\chi_H$ , they are not causally connected at that time, whereas if the distance is greater than  $\chi_p$ , they could have never been causally connected. The comoving Hubble horizon (or radius) becomes a quantity of great importance when discussing cosmic inflation, because not only does it signal the start/end of inflation,

$$\ddot{a} > 0 \iff \dot{\chi}_H < 0 \iff \text{inflating}, \quad (5.13)$$

it also determines the behaviour of comoving curvature perturbations.

### Dynamics: Friedmann equations

To solve for the overall evolution of the smooth universe, one needs to solve the Einstein equations from Eq. (5.3) with the relevant metric being the FLRW, Eq. (5.6). One further needs to model the matter content of the universe and find  $T_{\mu\nu}$  from the point of view of some observer.

Let the observer have a worldline tangential to the timelike 4-velocity

$$u^\mu = \frac{dx^\mu}{d\tau}, \quad (5.14)$$

with  $\tau$  being the proper time of the observers, such that  $u_\mu u^\mu = g_{\mu\nu} u^\mu u^\nu = 1$ . For the sake of simplicity and in accordance with the cosmological principle, we can then assume that the universe is filled with a single perfect fluid that only interacts gravitationally<sup>6</sup> and thus has the stress–energy tensor

$$T_\nu^\mu = (p + \rho)u^\mu u_\nu - p\delta_\nu^\mu, \quad (5.15)$$

which in a frame comoving with the fluid (i.e. one with  $u^\mu = (1, 0, 0, 0)$ ) becomes

$$T_\nu^\mu = \begin{pmatrix} \rho & 0 & 0 & 0 \\ 0 & -p & 0 & 0 \\ 0 & 0 & -p & 0 \\ 0 & 0 & 0 & -p \end{pmatrix}. \quad (5.16)$$

When substituted into the Einstein equations Eq. (5.3) this then yields the Friedmann equations,

$$\frac{\ddot{a}}{a} = \dot{H} + H^2 = -\frac{1}{6m_{\text{p}}^2}(\rho + 3p), \quad (5.17)$$

$$\left(\frac{\dot{a}}{a}\right)^2 = H^2 = \frac{1}{3m_{\text{p}}^2}\rho - \frac{K}{a^2}. \quad (5.18)$$

---

<sup>6</sup>A generalisation would be to assume a mixture of multiple fluids.

fluid name	$w$	$a(t)$	$\rho(a)$	$(aH)^{-1}$
matter	0	$t^{2/3}$	$a^{-3}$	$t^{1/3}$
radiation	$\frac{1}{3}$	$t^{1/2}$	$a^{-4}$	$t^{1/2}$
dark energy	-1	$e^{H_0 t}$	const.	$e^{-H_0 t}$
curvature	$-\frac{1}{3}$	$t$	$a^{-2}$	const.

Table 5.2 Dynamics of the FLRW universe dominated by the energy density of a given fluid with equation of state parameter  $w$ . We express the overall dynamics via the scale factor as a function of cosmic time,  $a(t)$ , the energy density of the fluid as a function of the scale factor,  $\rho(a)$ , and the comoving Hubble radius as a function of cosmic time,  $(aH)^{-1}$ .

From these an additional equation can be derived,

$$\dot{\rho} = -3(\rho + p)\frac{\dot{a}}{a}, \quad (5.19)$$

called the continuity equation. One last piece of information we need is a link between the pressure  $p$  of the fluid and its energy density  $\rho$ . This is given by the fluid's equation of state,

$$p = w\rho, \quad (5.20)$$

with a constant equation of state parameter  $w$ . With this, Eqs. (5.18) and (5.19) can be solved to give

$$\rho \propto a^{-3(1+w)}, \quad (5.21)$$

$$a(t) \propto \begin{cases} t^{\frac{2}{3(1+w)}} & w \neq -1, \\ e^{Ht} & w = -1. \end{cases} \quad (5.22)$$

Table 5.2 contains the equation of state parameter of the most commonly assumed types of fluids (matter, radiation, curvature and dark energy) alongside the evolution of the scale factor  $a(t)$  and energy density  $\rho(a)$  in a universe dominated by that kind of fluid. Curvature and dark energy can be regarded as cosmological fluids in this context: the effect of the curvature term ( $-\frac{K}{a^2}$  in Eq. (5.18)) and a cosmological constant on the Friedmann equations Eqs. (5.17) and (5.18) can be treated as if there were a fluid present with an effective energy density, pressure, and equation of state parameter. Note that according to Eq. (5.21), the energy density of pressureless matter decays as  $a^{-3}$ , consistent with the conservation of the number of particles, but that of radiation decays more quickly, as  $a^{-4}$ . One can derive that the wavelength of a photon gets stretched  $\propto a(t)$  with the expansion, therefore it loses energy as  $a^{-1}$ . Comparing the rate of decay of energy densities from Table 5.2, it is clear that over time, the energy density

and hence dynamics of the universe was dominated by radiation-matter-(curvature)-dark energy, resulting in a variation of the rate of expansion. To treat the presence of more than one fluid species more rigorously, we just need to write the total energy density and pressure as the sum of specific energy densities and pressures of each (non-interacting) fluid component,

$$\rho = \sum_i \rho_i, \quad p = \sum_i p_i = \sum_i w_i \rho_i. \quad (5.23)$$

Upon substituting this into Eq. (5.17), dividing through by  $H_0^2$  (where the 0 subscript refers to taking the present-day value of the relevant quantity), and defining the dimensionless energy densities as

$$\Omega_i \equiv \frac{\rho_i}{3m_{\text{p}}^2 H_0^2} = \frac{\rho_i}{\rho_{\text{crit}}}, \quad \Omega_K = -\frac{K}{(aH)^2}, \quad (5.24)$$

we get

$$\left(\frac{H}{H_0}\right)^2 = \sum_i \Omega_{i,0} \left(\frac{a}{a_0}\right)^{-3(w_i+1)}, \quad (5.25)$$

which can be further simplified by taking the present-day scale factor to be  $a_0 = a(t_0) = 1$ . Evaluating Eq. (5.25) at  $t = t_0$ , we find that the dimensionless energy densities are related by

$$1 = \sum_i \Omega_{i,0}. \quad (5.26)$$

## 5.2 Observations

### 5.2.1 The Cosmic Microwave Background

One of the most powerful probes of the early universe is the Cosmic Microwave Background (CMB) radiation. This radiation is made of photons free-streaming through space that is now nearly transparent to them, coming from all directions. These photons are relics of an early, hot, and dense universe, from a time when the characteristic energy scale was high enough to keep electrons and protons from forming neutral hydrogen, thus existing as primordial plasma. Photons were constantly scattering off of electrically charged particles, making their mean free path short and the primordial universe opaque to radiation. The constant interaction ensured that thermal equilibrium was maintained between matter and radiation. The stretching of spacetime, as we have seen, also ‘stretches’ the wavelength of photons, they therefore lose energy and the universe cools over time. After about 380,000 years, the universe has reached sufficiently low energy scales to allow for protons and electrons to form neutral hydrogen. In this era of *recombination*<sup>7</sup>, photons were finally able to free-stream and from this point on interacted with matter only via gravity. We see this radiation today as the CMB, observable only in the microwave regime of the electromagnetic spectrum thanks to the continued stretching of the

<sup>7</sup>which should really be called combination, since it is the first time neutral hydrogen was formed.

photons' wavelength with the expansion. The CMB thus offers us a view of the universe far earlier than the light from the first stars.

The two most important observations one can immediately<sup>8</sup> make of the CMB are that

1. It has a nearly perfect blackbody spectrum associated with the temperature  $T_{\text{CMB}} = 2.72548 \pm 0.00057$  K,
2. It is very close to being isotropic; however fluctuations in its temperature of the magnitude 1 in  $10^5$  exist.

While these observations about the CMB 'as a whole' are already enough to challenge the standard Big Bang model (as will be discussed in [Section 5.2.3](#)), to compare competing models and models with observations, information at all scales in the CMB needs to be exploited.

### Measuring CMB anisotropies: the $C_\ell$ 's

While no theory will predict a concrete pattern of anisotropies in the CMB, they can predict their summary statistics. Out of the many possible statistics, most commonly used are two-point functions, or power spectra in Fourier space. Since we observe the CMB over the surface of a sphere (the sky), angular spectra are used. Suppose that the temperature of the universe is described by the field

$$T(\mathbf{x}, \hat{\mathbf{n}}, \eta) = T(\eta) [1 + \Theta(\mathbf{x}, \hat{\mathbf{n}}, \eta)], \quad (5.27)$$

where  $\mathbf{x}$  denotes position (of the observer),  $\hat{\mathbf{n}}$  is the unit direction of the incoming photon, and  $\eta$  is conformal time. Since we can only observe here (on Earth) and now<sup>9</sup>, the direction of incoming CMB photons completely characterises their position in spacetime. On the surface of a sphere, the spherical harmonic functions,  $Y_{\ell m}(\hat{\mathbf{n}})$ , form a complete basis set, therefore we decompose the anisotropies in terms of the  $Y_{\ell m}(\hat{\mathbf{n}})$ :

$$\Theta(\mathbf{x}, \hat{\mathbf{n}}, \eta) = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}(\mathbf{x}, \eta) Y_{\ell m}(\hat{\mathbf{n}}). \quad (5.28)$$

The  $\ell, m$  are integers ( $\ell$  often being called the multipole) and are conjugates to  $\hat{\mathbf{n}}$  the same way the wavevector  $\mathbf{k}$  is conjugate to  $\mathbf{x}$  in the Fourier transform. The  $a_{\ell m}$  completely specify the CMB, therefore no theory can predict them exactly. They can, however, estimate the underlying statistical distribution from which the  $a_{\ell m}$  are drawn. They will have a vanishing

<sup>8</sup>At the first detection of the CMB by Penzias and Wilson in 1965 [97, 98], the intensity of the radiation was only measured at a single wavelength, the blackbody spectrum was only confirmed later. Likewise, the anisotropies were not found for another 25 years after first detection of the CMB [99].

<sup>9</sup>today, or in the last couple of decades – this variation in time is small in comparison to the time over which the CMB varies, although with sufficient sensitivity, the real-time cooling of the CMB could be tracked and used to compute the local expansion rate [100].

mean but a non-zero variance:

$$\langle a_{\ell m} \rangle = 0, \quad \langle a_{\ell m} a_{\ell' m'}^* \rangle = \delta_{\ell\ell'} \delta_{mm'} C_\ell, \quad (5.29)$$

where the angled brackets denote an average over the sky<sup>10</sup>. When mapping the CMB, all  $2\ell + 1$   $a_{\ell m}$ 's are measured for any given  $\ell$ . They are samples drawn from the same underlying distribution for a given  $\ell$ , which has zero mean and a spread of  $\sqrt{C_\ell}$ , as expressed by Eq. (5.29). The larger the number of samples,  $2\ell + 1$ , the better an idea the  $a_{\ell m}$ 's give about the spread of the distribution. There is therefore an inherent uncertainty associated with the  $C_\ell$ 's, termed *cosmic variance*, stemming from the limited number of samples we have access to to estimate the spread. The uncertainty scales as the inverse squareroot of the number of samples<sup>11</sup>,

$$\frac{\Delta C_\ell}{C_\ell} = \sqrt{\frac{2}{2\ell + 1}}. \quad (5.30)$$

Cosmic variance naturally limits the accuracy with which the  $C_\ell$ 's at low multipoles can be determined.

Photons have two measurable properties: energy and polarisation. The polarisation is decomposed into  $E$ - and  $B$ -modes, named after their resemblance to electric and magnetic fields.  $E$ -modes have  $(-1)^\ell$  parity on the sphere, while  $B$ -modes have  $(-1)^{\ell+1}$ , meaning under  $\hat{\mathbf{n}} \rightarrow -\hat{\mathbf{n}}$ ,  $E$ -modes change sign for odd  $\ell$ , and  $B$ -modes for even  $\ell$ . In plots of the CMB angular spectra, the  $y$ -axis is typically labelled

$$D_\ell^{XX} = \frac{1}{2\pi} \ell(\ell + 1) C_\ell^{XX}, \quad (5.31)$$

where  $X \in \{T, E, B\}$ , standing for temperature and  $E$ - or  $B$ -mode polarisation. Out of all possible CMB anisotropy power spectra, most of the information comes from the  $TT$  spectrum, followed by  $TE$ ,  $EE$ , and finally  $BB$ , with the ordering being determined by the signal-to-noise (S/N) ratio of the corresponding measurements. Adding the polarisation data, however, can help break parameter degeneracies, and even provide new information<sup>12</sup>. While cosmological information can be gathered from many other sources (see Section 5.2.2), at the present day, the CMB represents the source with the largest statistical power. Since the anisotropies are assumed to be Gaussian, the total constraining power in the CMB is decided by how many  $\ell$ -modes one can measure. The latest space-based instrument, the *Planck* satellite, can measure anisotropies up to  $\ell \sim 1500$ , which equates to  $\sim 900\sigma$  of detection on its own. Adding the

<sup>10</sup>Technically, angled brackets denote an ensemble average, i.e. an average over the different realisations of the same random process. Since we only have one Universe, the ensemble average is replaced by averaging over widely separated parts of the CMB that are causally disconnected, arguing that the two are equivalent. This equivalence is a result of the ergodicity of the relevant fields which can be proven formally in some cases [101], but is understood as an axiom in cosmology.

<sup>11</sup>assuming the anisotropies are Gaussian.

<sup>12</sup>E.g. large-angle  $BB$  measurements can help compute the tensor-to-scalar ratio parameter,  $r$ , which is otherwise hard to constrain.

power of other experiments that measure polarisation (e.g. ACTPol [102], SPTpol [103]), the CMB has thus far provided enough S/N worth  $> 1000\sigma$  of detection. Future experiments may be able to measure up to  $\ell \sim 3000$  (and not much beyond due to damping and foregrounds), which is cause for optimism, but the information remaining in the CMB is finite.

### 5.2.2 Other observational probes

Needless to say, the CMB is not the only observable one can use to constrain cosmological models. Some alternatives are described below.

**Large-scale structure (LSS).** Structure on the largest scales can be probed with galaxy surveys like the Two Degree Field Galaxy Redshift Survey (2dF, [104]), or the Sloan Digital Sky Survey (SDSS, [105]). Matter perturbations on such large scales have grown linearly, unlike perturbations on small scales (e.g. the Earth), the evolution of which was not driven by cosmology.

**Gravitational lensing.** Although CMB photons have largely been free-streaming to reach us today since recombination, their trajectories have been geodesics in a spacetime curved by matter in and near the line-of-sight. Weak gravitational lensing thus shifts and distorts the CMB, and can change the polarisation pattern by converting  $E$ -modes to  $B$ -modes. The distortions can however be mapped and the lensing potential reconstructed, which again is telling of the formation of large-scale structure in the universe.

**Baryon Acoustic Oscillations (BAO).** Before recombination, baryonic matter, radiation, and dark matter were in thermal equilibrium in a primordial plasma. Overdensities in this plasma would try to collapse under gravity, countered by radiation pressure. This would cause spherical acoustic waves to propagate outwards from each overdensity. Since dark matter only interacts gravitationally, it stays in the centre while shells of overdense regions of baryons and radiation expand around them. The maximum radius to which the shell can grow before recombination is determined by the sound speed in the medium and is called the *sound horizon*. This configuration of overdensities is observable in the CMB power spectrum as excess power (an acoustic peak) at the wavelength corresponding to the sound horizon. Later, it causes more galaxies to form at separations equal to the sound horizon, which can be detected in galaxy clustering surveys. BAO can thus be used as a *standard ruler* to determine the expansion rate of the universe, independent of supernova measurements.

**Supernovae.** Type Ia supernovae occur when a white dwarf accretes enough of its binary companion's material to reach a critical mass, the Chandrasekhar mass limit. Since type Ia supernova explosions happen at approximately the same mass, the resulting supernovae's maximum absolute luminosities will be roughly equal and deducible from the

shape of their lightcurve, making the objects good *standard candles*. Since their intrinsic luminosity is well constrained, their measured flux can be used to infer their distance, once the latter is calibrated using distances to local objects, e.g. Cepheid variable stars. This way, a cosmic (distance) ladder can be built and the Hubble parameter can be inferred from Hubble’s law. [8, 106–108].

### 5.2.3 Happy accidents

The CMB tells us about the state of the universe at the time of last scattering. The picture it paints is, however, highly improbable within the framework of the conventional Big Bang theory presented in Section 5.1.2. An additional component, *inflation*, is necessary for the CMB we observe not to be the result of extremely fine-tuned initial conditions. Inflation explains the lack of curvature observed (the *flatness problem*), and the smoothness of the CMB (the *horizon problem*).

#### The flatness problem

Measurements show that the present-day curvature density parameter,  $\Omega_{K,0} = -\frac{K}{(a_0 H_0)^2}$ , is very close to zero. Considering the curvature density as a function of time, it is proportional to the square of the Hubble radius,  $\Omega_K \propto (aH)^{-2}$ , which, in the standard Big Bang theory, grows with time (see the last column of Table 5.2). In order to get as small a present-day value as we observe ( $\lesssim \mathcal{O}(1)$ ), the curvature density at would have therefore needed to be  $\mathcal{O}(10^{-16})$  at Big Bang Nucleosynthesis, and  $\mathcal{O}(10^{-61})$  when the scale factor was comparable to the Planck scale. In other words,  $\Omega = \sum_{i \neq K} \Omega_i = 1$  is an unstable fixed point, and the standard Big Bang universe is naturally driven away from flatness. While an extremely small initial  $\Omega_K$  is not impossible, the level of fine-tuning it requires makes our present universe look like a happy accident.

#### The horizon problem

The degree of smoothness in the CMB turns out to be highly unusual if we consider the size of causally connected patches at the time of its emission (at last scattering) as seen today. The maximum distance between two points that could have been causally connected in the past is given by the particle horizon  $\tau$ , which in the conventional Big Bang expansion varies as

$$\tau = \int^t \frac{d\sigma}{a(\sigma)} = \int_0^t \sigma^{-\frac{2}{3(1+w)}} d\sigma \propto a^{\frac{1+3w}{2}}. \quad (5.32)$$

In radiation- or matter-dominated eras (i.e. in the earliest times of the standard Big Bang evolution), the particle horizon therefore grew with time. It can further be shown that in the standard picture the particle horizon is proportional to the Hubble horizon,  $(aH)^{-1}$ . This means that CMB fluctuations of a given (constant) lengthscale entering the horizon today were

far outside the horizon at recombination (and therefore causally independent)<sup>13</sup>. To get the homogeneity we observe in the CMB on scales spanning many lengthscales is therefore only possible by chance in the standard Big Bang model.

### 5.3 Cosmic inflation

Introducing an early phase of accelerated expansion of the universe can resolve both the flatness and horizon problem [109–112]. *Inflation* is defined as accelerated growth, or equivalently, a shrinking Hubble horizon,

$$\ddot{a} > 0 \iff (aH)^{-1} \text{ decreases.} \quad (5.33)$$

This process is thought to have taken place at around  $10^{-34}$ s after the Big Bang, when the Universe was too hot for ordinary matter to have formed, and when the scale factor itself was quantum-scale. The most suitable phenomenology to use is therefore quantum field theory<sup>14</sup>. Below we introduce some key concepts from inflationary cosmology using the simplest family of inflationary models involving a single scalar field.

#### 5.3.1 Single-field inflation

The simplest inflationary models involve a single, time-dependent scalar field  $\phi \equiv \phi(t)$  and associated particle, the *inflaton*. This field is minimally ( $\xi = 0$ ) coupled to gravity and has some self-interaction described by the inflationary potential  $V(\phi)$ . The field is then associated with the Lagrangian density

$$\mathcal{L}_\phi = \frac{1}{2} \nabla^\mu \phi \nabla_\mu \phi - V(\phi). \quad (5.34)$$

The Lagrangian density can be substituted into Eq. (5.4) to obtain the fluid’s stress–energy tensor. The energy density  $\rho$  and pressure  $p$  can then be found by inspection of the 00 and  $ij$  component, respectively, giving

$$\rho_\phi = \frac{1}{2} \dot{\phi}^2 + V(\phi), \quad p_\phi = \frac{1}{2} \dot{\phi}^2 - V(\phi). \quad (5.35)$$

<sup>13</sup>It can be estimated that if the universe was filled with radiation prior to recombination, the size of causally connected patches in the CMB today would be  $\sim 1^\circ$ , roughly twice the size of the full moon. The CMB, however, appears homogeneous on much larger scales than that.

<sup>14</sup>Although models with sufficiently long periods of inflation run into the problem of crossing over to trans-planckian scales, where quantum field theoretical principles are probably not applicable [113, 114].

Substituting Eq. (5.35) into Eqs. (5.17) to (5.19) yields the background dynamics of a universe with the inflaton as the sole source of energy density:

$$\frac{\ddot{a}}{a} = \dot{H} + H^2 = -\frac{1}{3m_p^2}(\dot{\phi}^2 - V(\phi)), \quad (5.36)$$

$$H^2 = \frac{1}{3m_p^2} \left( \frac{1}{2}\dot{\phi}^2 + V(\phi) \right) - \frac{K}{a^2}, \quad (5.37)$$

$$0 = \ddot{\phi} + 3H\dot{\phi} + \frac{dV(\phi)}{d\phi}. \quad (5.38)$$

The equations above possess two rescaling symmetries. They are invariant under the rescaling of the scale factor,

$$a \rightarrow \alpha a, \quad (5.39)$$

and also a simultaneous rescaling of time, the Hubble function, and the potential:

$$\begin{aligned} t &\rightarrow \sigma^{-1}t, \\ H &\rightarrow \sigma H, \\ V(\phi) &\rightarrow \sigma^2 V(\phi). \end{aligned} \quad (5.40)$$

These symmetries reduce the degrees of freedom in the general solution and initial conditions of the Friedmann equations.

From Eq. (5.36) it is clear that the condition for inflation to occur is<sup>15</sup>

$$V(\phi) > \dot{\phi}^2. \quad (5.41)$$

At one extreme, if the potential of the inflaton dominates ( $V(\phi) \gg \dot{\phi}^2$ ), the field equations can be solved exactly and the Universe undergoes *slow-roll inflation*. The other extreme,  $\dot{\phi}^2 \gg V(\phi)$ , is also solvable analytically, and is called *kinetic dominance*.

### Slow-roll and kinetic dominance

Eq. (5.38) can be understood intuitively as the equation of motion of a block with position  $\phi(t)$  sliding down a potential well  $V(\phi)$ , and being slowed by a friction term  $\propto H$ . From Eq. (5.36) it is clear that for large values of  $V(\phi)$ , this friction term will be large. Once the inflaton has slowed down enough,

$$\dot{\phi}^2 \ll V(\phi) \quad (5.42)$$

becomes a good approximation. Consequently,

$$H(t) \sim \text{constant} \equiv H_0 \implies a(t) \sim e^{H_0 t}, \quad (5.43)$$

<sup>15</sup>More generally, by inspection of Eq. (5.17), the condition for inflation is  $\rho + 3p < 0$ .

the scale factor grows exponentially and the Universe inflates. One can express how much inflation took place via the *total number of e-folds*,

$$N = \ln a. \quad (5.44)$$

Due to the presence of the friction term  $3H\dot{\phi}$ , the inflaton in most cases will eventually lose enough energy to enter a slowly rolling phase: slow-roll inflation is an attractor solution for most inflationary potentials.

The slow-roll approximation Eq. (5.42) can be used to set initial conditions to Eqs. (5.36) to (5.38). However, this would imply that the universe had started from an (eternally) inflating phase, rather than a singularity. An alternative is offered by [115], which proves on theoretical grounds that in most single-field inflation models and under minimal assumptions, an attractor solution called *kinetic dominance*, in which the kinetic energy of the inflaton dominates, will be reached backwards in time.

Kinetic dominance is defined as  $\dot{\phi}^2 \gg V$ . In this limit, the field equations admit closed-form solutions with only one free parameter  $\phi_P$  (one can set  $a_P = 1$ , but it will be retained here),

$$\begin{aligned} a &= a_P t^{\frac{1}{3}}, \\ \phi &= \phi_P - \sqrt{\frac{2}{3}} m_{\text{P}} \ln t, \\ \dot{\phi} &= -\sqrt{\frac{2}{3}} m_{\text{P}} t^{-1}, \\ H &= \frac{1}{3t}, \end{aligned} \quad (5.45)$$

The results of [115] show that the above can be used to set initial conditions for the Friedmann equations. Fig. 5.1 plots their numerical solutions, clearly showing that slow-roll inflation begins after the initial kinetically dominated phase.

Models starting from kinetic dominance followed by a brief period of inflation require their primordial power spectra to be evaluated numerically which can get computationally expensive. These models can produce a suppression of power at low angular scales in CMB spectra [116], and may thus produce better fits than the  $\Lambda$ CDM model. Kinetic dominance has also been shown on statistical [117, 118] grounds to be a likely initial phase the Universe undergoes before the onset of inflation. As will be shown in Section 5.4.3, the numerical calculations required are ideal for `oscode`. Due to these factors, models involving kinetic dominance and ‘just enough inflation’ are used frequently in the cosmological examples presented in this thesis.

### The flatness and horizon problems revisited

Inflation offers a resolution to both problems discussed in Section 5.2.3. Since the comoving Hubble horizon shrinks during inflation and  $\Omega_K \propto (aH)^{-2}$ , the Universe is naturally driven

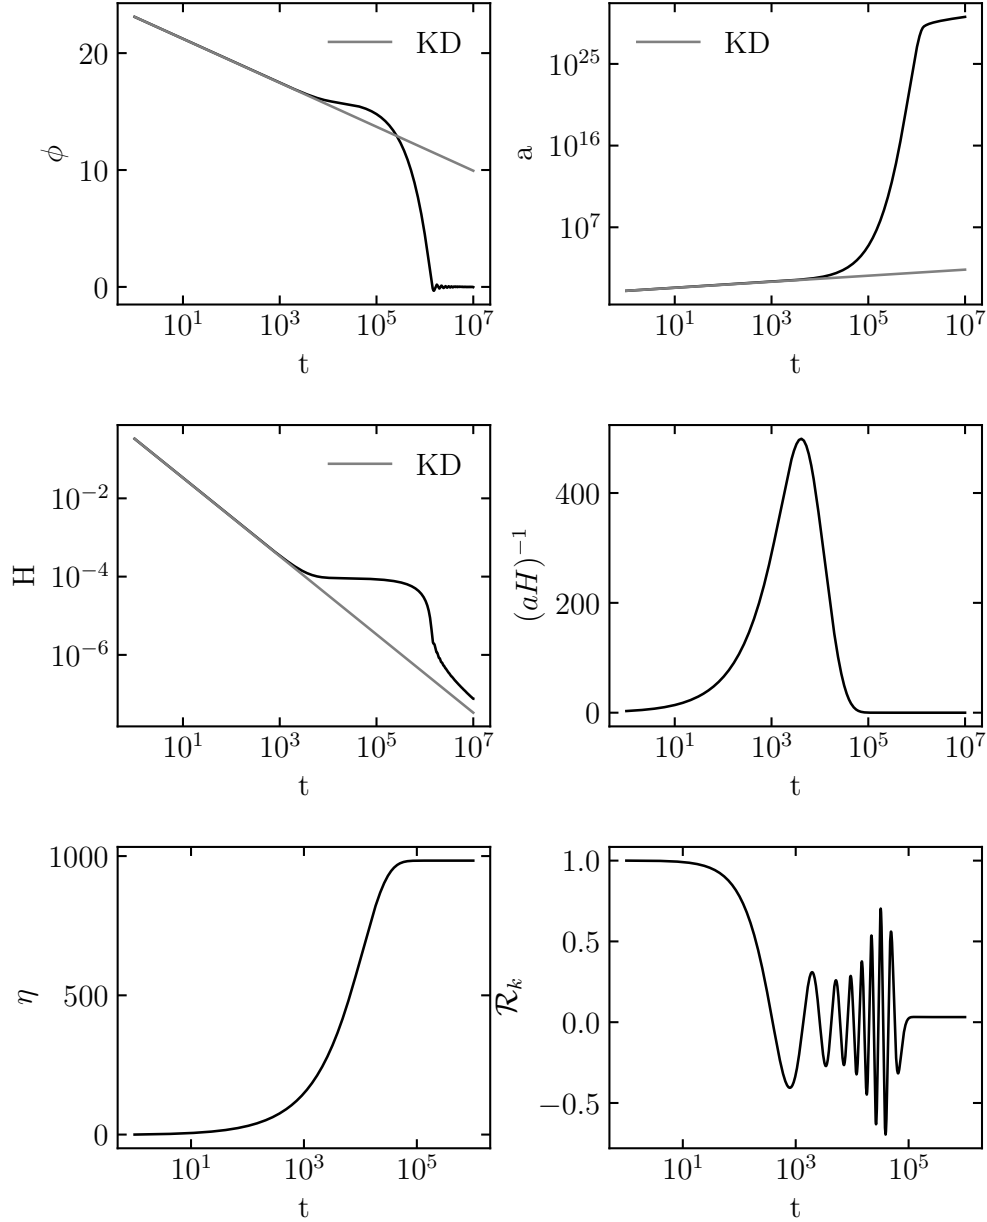


Fig. 5.1 Evolution of the ‘background’ FLRW Universe (inflaton field on top left, scale factor on top right, Hubble parameter on middle left), the comoving Hubble horizon (middle right), conformal time (bottom left), and a typical gauge-invariant curvature perturbation (bottom right, introduced in [Section 5.4](#)). Plotted in grey are the kinetically dominated limits of the first three quantities, which they emerge from, before entering inflation. The start of inflation can be seen by the decrease in the comoving Hubble radius. Well into inflation, when the example perturbation’s comoving lengthscale exits the Hubble horizon, its amplitude freezes.

towards flatness. It also follows that comoving lengthscales entering our horizon today will have had the chance to be inside the Hubble horizon during inflation, allowing them to come to thermal equilibrium and explaining the smoothness of the CMB.

As well as explaining the smoothness of the CMB, inflation also suggests a mechanism for quantum fluctuations to grow into the anisotropies we see in the CMB.

## 5.4 Perturbations

Consider that due to quantum fluctuations, patches develop in the field  $\phi$  that differ in value slightly from the homogeneous average. Since  $\phi_P$  directly determines the amount of expansion the universe undergoes whilst inflating, these patches will undergo slightly different amounts of inflation and will finish inflating at slightly different times, as illustrated in Fig. 5.2. This results in differences in matter density across the patches, and is how inflation can help amplify fluctuations and is thus a predictive theory. This section summarises the mathematics of getting CMB anisotropies from quantum fluctuations at primordial times.

Symmetries of the FLRW spacetime simplify the mathematics of perturbations considerably. Due to the (spatial) translation invariance of the linear equations of motion of the perturbations, we can treat their Fourier modes independently. We will further decompose the perturbations into scalar, vector, and tensor parts via a procedure called the SVT decomposition. Rotational invariance of the FLRW background ensures that these components also evolve independently.

### 5.4.1 Gauge choice

The large degree of isotropy observed in the CMB means a perturbative analysis of the quantities of interest is sufficiently accurate, and we can neglect any term higher than linear order in the perturbations. We split both the metric and the matter fields into a homogeneous background and a spatially dependent perturbation part:

$$X(t, \mathbf{x}) = \bar{X}(t) + \delta X(t, \mathbf{x}), \quad (5.46)$$

where  $\delta X$  is small. Perturbing the spacetime metric (or gravitational field potential) brings with it a fundamental difficulty. General relativity is a covariant theory: it involves manifolds with no preferred coordinate systems. If we were to perturb a velocity field,  $\delta v^\mu = v'^\mu - v^\mu$ , the two fields on the right-hand-side would be defined on different manifolds  $M'$  and  $M$ , but subtraction of vectors on different manifolds cannot be covariantly defined. Embedding  $M$  and  $M'$  in a higher-dimensional manifold  $N$  doesn't help: the two vector fields would still be defined at different points. We therefore require an identification map to decide which points on  $M$  and  $M'$  are equivalent, a *gauge choice*. If part of this mapping is left arbitrary, we are left with a remaining *gauge freedom*, which results in fictitious *gauge modes* among the perturbations. A gauge choice can be understood in geometrical terms [91], and is illustrated




Fig. 5.2 Patches of the Universe will start from different initial field values ( $\phi_P$ ), and hence will undergo different amounts (e-folds) of inflation.

in Fig. 5.3. Consider a one-parameter family of 4-manifolds  $M_\varepsilon$  embedded in a 5-manifold  $N$ . Out of these,  $M_0$  represents the unperturbed, background spacetime, and each of the  $M_\varepsilon$  is a perturbed spacetime. The correspondence between points on the  $M_\varepsilon$  is specified via a vector field  $X$  on  $N$ , which is everywhere transverse to the  $M_\varepsilon$ . The points which lie on the same integral curve  $\gamma$  are taken to be equivalent. A choice of a particular  $X$  is a gauge choice.

#### 5.4.2 Perturbed matter and metric

The most general form a perturbed FRLW metric can take is [90, 91]

$$ds^2 = -(1 + 2\Phi)dt^2 + 2aB_i dx^i dt + a^2 [(1 - 2\Psi)c_{ij} + E_{ij}] dx^i dx^j, \quad (5.47)$$

where  $c_{ij}$  denotes the spatial part of the unperturbed FLRW metric. The perturbations  $B_i$  and  $E_{ij}$  can be further decomposed into scalar, vector, and tensor components [119] via the SVT decomposition as

$$B_i = \partial_i B - S_i, \quad \text{with} \quad \partial^i S_i = 0, \quad (5.48)$$

and

$$E_{ij} = 2\partial_{ij} E + 2\partial_{(i} F_{j)} + h_{ij}, \quad \text{with} \quad \partial^i F_i = 0, \quad h_i^i = \partial^i h_{ij} = 0. \quad (5.49)$$




Fig. 5.3 The choice of gauge can be interpreted geometrically. The figure shows a family of manifolds  $M_\epsilon$  representing spacetimes, parametrised by a single parameter  $\epsilon$ . Out of these,  $M_0$  is the unperturbed FLRW spacetime, and each subsequent  $M$  has been perturbed. To identify equivalent points across the perturbed manifolds, a vector field  $X$  is defined which is transverse to the manifolds at those points. Equivalent points lie on the same integral curves  $\gamma_i(\epsilon)$  of  $X$ , which are the solutions of  $d\gamma_i(\epsilon)/d\epsilon = X$ . The choice of  $X$  is therefore a choice of gauge; a different vector field  $Y$  would select different points to be equivalent across the perturbed manifolds.

The parameters leave a total of 10 degrees of freedom, enough to construct any (symmetric) metric. The matter field (in a single-field model) is perturbed as

$$\phi(t, \mathbf{x}) = \bar{\phi}(t) + \delta\phi(t, \mathbf{x}). \quad (5.50)$$

Out of the perturbations introduced above, vectors decay with the expansion of the universe, and are not relevant to inflationary models. The tensor perturbations are already gauge-invariant, and the scalars can be combined to form a set of gauge-invariant variables, found by Bardeen [120]. Of particular interest is the *comoving curvature perturbation*  $\mathcal{R}$ , which during inflation takes the form

$$\mathcal{R} = \Psi + \frac{H}{\dot{\phi}} \delta\phi. \quad (5.51)$$

$\mathcal{R}$  gets its name from measuring the spatial curvature of constant- $\phi$  hypersurfaces. The linearised, perturbed Einstein equations can then be used to show that  $\mathcal{R}$  is constant on superhorizon scales for adiabatic perturbations,

$$\dot{\mathcal{R}} = -\frac{H}{\bar{\rho} + \bar{p}} \delta p_{\text{en}} + \frac{k^2}{(aH)^2} (\dots) \implies \dot{\mathcal{R}} \approx 0 \text{ if } \delta p_{\text{en}} = 0 \text{ and } k < aH. \quad (5.52)$$

where  $\delta p_{\text{en}} = \delta p - \frac{\bar{p}}{\bar{\rho}} \delta \rho$  is related to the entropy (also called isocurvature) perturbation and expresses perturbations in the relative number densities of the different particle types present. Perturbations of a system can either be adiabatic, for which  $\delta \rho_{\text{en}} = 0$ , or entropic. In [Eq. \(5.52\)](#),  $k$  denotes the Fourier wavenumber of the perturbation mode, and we say the mode is superhorizon when its lengthscale,  $1/k$ , is much larger than the comoving Hubble horizon,  $k < aH$ .

The naturally gauge-invariant tensor perturbations  $h_{ij}$  are often decomposed into eigenmodes of the spatial Laplacian,

$$\nabla^2 e_{ij} = -k^2 e_{ij}^{(+, \times)}, \quad (5.53)$$

with (comoving) wavevector  $k$ , amplitude  $h(t)$ , and the two possible polarisation states denoted  $(+, \times)$ .

The subhorizon evolution of the scalar and tensor perturbations is given by a wave equation. In the scalar case it is called the *Mukhanov–Sasaki* equation, and it will be revisited frequently in this thesis when discussing `oscode`'s applications in cosmology.

### 5.4.3 Subhorizon evolution: the Mukhanov–Sasaki equation

To derive the subhorizon equation of motion for the scalar curvature perturbations (originally performed by [\[121–125\]](#), reviewed in [\[90\]](#)), we start from the action describing the single-field, slow-roll inflation model from [Section 5.3.1](#):

$$S = \frac{1}{2} \int d^4x \sqrt{|g|} [R - \nabla_\mu \phi \nabla^\mu \phi - 2V(\phi)], \quad (5.54)$$

where  $R$  is the Ricci scalar, and we set the Planck mass to unity,  $m_p = (8\pi G)^{-1/2} = 1$ . We choose to work in the *comoving gauge*, defined by

$$\delta\phi = 0, \quad g_{ij} = a^2 [(1 - 2\mathcal{R})\delta_{ij} + h_{ij}], \quad \partial_i h_{ij} = h_i^i = 0, \quad (5.55)$$

where we implicitly restricted ourselves to a spatially flat Universe by replacing  $c_{ij}$  from [Eq. \(5.47\)](#) with  $\delta_{ij}$ . This action is then expanded to second order in  $\mathcal{R}$  in a lengthy derivation. The first order action vanishes by design. We get

$$S_{(2)} = \frac{1}{2} \int d^4x a^3 \frac{\dot{\phi}^2}{H^2} [\dot{\mathcal{R}}^2 - a^{-2} (\partial_i \mathcal{R})^2], \quad (5.56)$$

which we now need to transform into canonical form with a ‘bare’ kinetic term. This is achieved by switching to conformal time and introducing the Mukhanov variable  $v$ ,

$$z = \frac{a\dot{\phi}}{H}, \quad v \equiv z\mathcal{R}, \quad (5.57)$$

followed by integration by parts to get rid of the term  $\propto v'v$ , where prime denotes differentiation with respect to conformal time, and discarding the boundary term. This results in the action describing a canonically normalised scalar:

$$S_{(2)} = \frac{1}{2} \int d\tau d^3x \left[ (v')^2 + \frac{z''}{z} v^2 - (\partial_i v)^2 \right] \quad (5.58)$$

Writing the perturbation in terms of its Fourier modes

$$v(\tau, \mathbf{x}) = \int \frac{d^3k}{(2\pi)^3} v_{\mathbf{k}}(\tau) e^{i\mathbf{k}\cdot\mathbf{x}}, \quad (5.59)$$

and extremising the action yields the equation of motion

$$v_k'' + \left( k^2 - \frac{z''}{z} \right) v_k = 0. \quad (5.60)$$

Note that the vector notation has been dropped from the wavevector  $k$  as the the resulting equation only depends on its magnitude. Computationally it is often more convenient to write the Mukhanov–Sasaki equation in terms of  $\mathcal{R}_k$ , since this has a constant magnitude after horizon crossing, and cosmic time,

$$\ddot{\mathcal{R}}_k + 2 \left( \frac{H}{2} + \frac{\dot{z}}{z} \right) \dot{\mathcal{R}}_k + \frac{k^2}{a^2} \mathcal{R}_k = 0, \quad (5.61)$$

or with the number of e-folds  $N$  as the independent variable and the overdot denoting differentiation with respect to  $N$ ,

$$\ddot{\mathcal{R}}_k + 2 \left( \frac{\ddot{\phi}}{\dot{\phi}} - \frac{1}{2} \dot{\phi}^2 + \frac{3}{2} \right) \dot{\mathcal{R}}_k + \left( \frac{k}{aH} \right)^2 \mathcal{R}_k = 0. \quad (5.62)$$

In either case, the resulting equation of motion, by design, is one describing a harmonic oscillator with a time-dependent frequency and a potential damping term.

Deriving the equivalent of the Mukhanov–Sasaki equation for tensor perturbations proceeds in a similar way, but is relatively simple in comparison. Expanding the Einstein–Hilbert action to second order in the tensor perturbation yields an action for a massless scalar field on a FLRW background:

$$S_{(2)} \propto \frac{1}{2} \int d\tau dx^3 a^2 \left[ (h'_{ij})^2 - (\partial_l h_{ij})^2 \right], \quad (5.63)$$

which after Fourier expansion and the substitution  $v_{\mathbf{k}}^{(+,\times)} = ah_{\mathbf{k}}^{(+,\times)}$  becomes

$$S_{(2)} = \frac{1}{2} \int d\tau d^3\mathbf{k} \left[ (v_{\mathbf{k}}^{(+,\times)})^2 - \left( k^2 - \frac{a''}{a} \right) (v_{\mathbf{k}}^{(+,\times)})^2 \right], \quad (5.64)$$

and leads to the equation of motion

$$v_{\mathbf{k}}^{''(+,\times)} + \left( k^2 - \frac{a''}{a} \right) v_{\mathbf{k}}^{(+,\times)} = 0. \quad (5.65)$$

#### 5.4.4 Quantisation and initial conditions

The perturbations represented by the field  $\mathcal{R}(t, \mathbf{x})$  are quantum-scale. In order to study them, we therefore need to perform (canonical) quantisation: promote the relevant fields to operators, impose canonical commutation relations, and define a vacuum state to fix the mode functions completely. Quantisation is done by analogy with the quantum harmonic oscillator, with a key difference being the time-dependence of the frequency in the curvature perturbations' case. This causes the energy eigenstates (and indeed the vacuum state) to change with time, but we can set the mode functions to be in a given state at one given (initial) time regardless, which can serve as initial conditions of the Mukhanov–Sasaki equation. The change of variables and integration by parts between Eqs. (5.56) and (5.58) was necessary to bring the action to canonically normalised form and draw an analogy with the quantum harmonic oscillator, although this can be achieved with different independent-dependent variable pairs. Chapter 7 discusses in detail the conventional way of quantising the Mukhanov action and choosing the vacuum state, the potential pitfalls and limits of the conventional approach, as well as alternative methods for setting initial conditions.

#### 5.4.5 Connection to CMB anisotropies

Once the Hubble horizon has shrunk enough during inflation below the lengthscales of all perturbation modes of interest, none of them evolve any further. The distribution of power in perturbations over different scales can then be summarised by the *primordial power spectrum* (PPS)<sup>16</sup>,

$$\langle \mathcal{R}_{\mathbf{k}} \mathcal{R}_{\mathbf{k}'} \rangle = (2\pi)^3 \delta(\mathbf{k} + \mathbf{k}') P_{\mathcal{R}}(k), \quad \mathcal{P}_{\mathcal{R}}(k) = \frac{k^3}{2\pi^2} P_{\mathcal{R}}(k), \quad (5.66)$$

where the angled brackets denote taking an ensemble average over the perturbations, and the  $\mathcal{R}_{\mathbf{k}}$  are read off after horizon exit. The perturbations, however, will later (after inflation and reheating) re-enter the horizon and evolve again, under the influence of well-understood and tested laws of physics. The changes to the superhorizon  $\mathcal{R}_{\mathbf{k}}$  between an initial time  $\tau_*$

<sup>16</sup>We give the expression for the PPS of scalar curvature perturbations here, but an equivalent expression exists for the tensor modes, summed over the two polarisations

and a later time  $\tau$ , and the connection between  $\mathcal{R}_{\mathbf{k}}$  and the quantity  $\mathcal{Q}$  being measured are encompassed in the *transfer functions*  $T_{\mathcal{Q}}$ :

$$\mathcal{Q}_{\mathbf{k}}(\tau) = T_{\mathcal{Q}}(k, \tau, \tau_*) \mathcal{R}_{\mathbf{k}}(\tau_*). \quad (5.67)$$

The derivation of the  $T_{\mathcal{Q}}(k, \tau, \tau_*)$  in different cosmological scenarios is beyond the scope of this thesis and we refer the interested reader to Dodelson's *Modern Cosmology* [84]. We quote the important result linking the CMB anisotropies to the PPS:

$$C_{\ell}^{XY} = \frac{2}{\pi} \int dk k^2 P(k) \Delta_{X\ell}(k) \Delta_{Y\ell}(k), \quad (5.68)$$

$$\Delta_{X\ell}(k) = \int_0^{\tau_0} d\tau S_X(k, \tau) P_{X,\ell}(k[\tau_0 - \tau]), \quad (5.69)$$

where the  $P(k)$  refer to the PPS (of scalar and tensor perturbations),  $X$  and  $Y$  are placeholders for the indices  $T, E, B$  for the temperature and polarisation spectra, and the  $\Delta_{X,Y\ell}$  are the transfer functions written as line-of-sight integrals factorised into (physical) source terms  $S_{X,Y}(k, \tau)$  and geometric projection factors  $P_{X,Y\ell}(k[\tau_0 - \tau])$ . In practice, the  $C_{\ell}$ 's are computed numerically for the given cosmological model by one of the several Boltzmann codes available, e.g. CLASS [126–129] or CAMB [130].

Fig. 5.4 summarises the evolution of primordial perturbations from the setting of their initial conditions to them being linked to CMB anisotropies. Our computational strategy for computing CMB power spectra will in general consist of

1. Setting initial conditions to the background cosmology at some initial time  $t_i$ ,
2. Solving the Friedmann equations (numerically) from  $t_i$  through inflation until all perturbations of interest have exited the comoving Hubble horizon,
3. Setting initial conditions for the perturbations  $\mathcal{R}_{\mathbf{k}}$  (or the tensor modes) at a time  $t_{\text{start}}$ , which may be later than  $t_i$ , and (numerically) solving the Mukhanov–Sasaki equation to evolve them until all modes have exited the comoving Hubble horizon,
4. Well after horizon exit, reading off the now constant perturbation amplitudes and computing the primordial power spectrum  $\mathcal{P}_{\mathcal{R}}(k)$ ,
5. Plugging the primordial power spectra into a Boltzmann solver to compute CMB observables, such as  $C_{\ell}^{TT}$ ,
6. Comparing the  $C_{\ell}^{TT}$  predicted by the model at hand with observations to infer the values of model parameters or compare competing models.

The details of the last step are discussed in the following section.

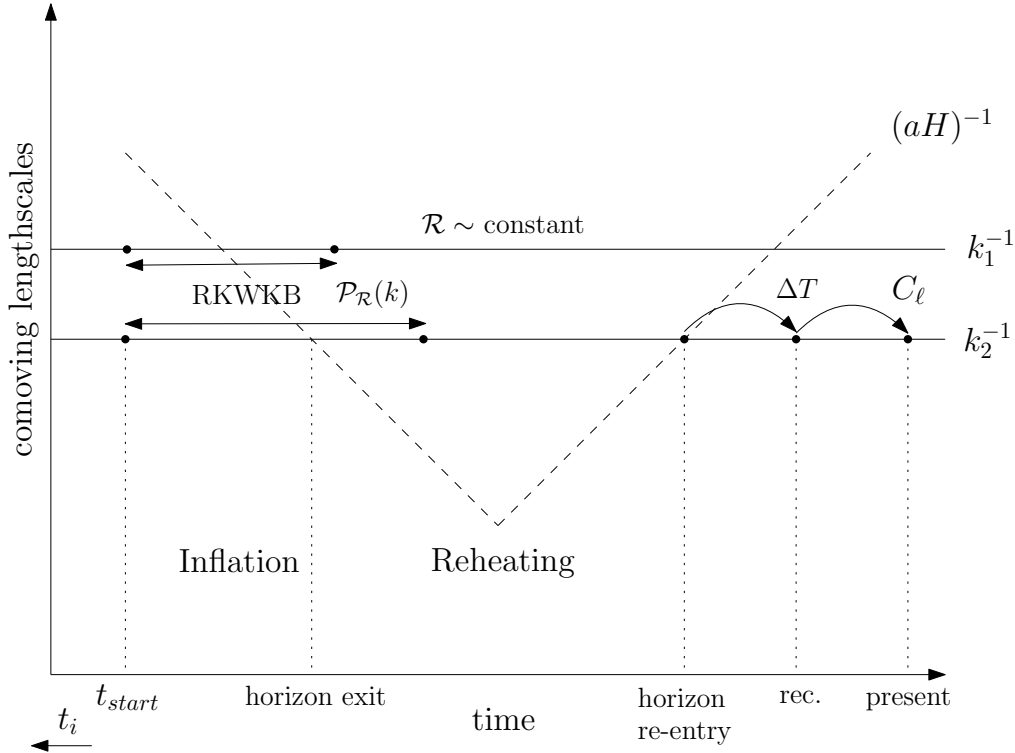


Fig. 5.4 Computational strategy for constructing the CMB temperature spectrum. We start by setting initial conditions to the background FLRW universe at  $t_i$ , then set separate conditions for the curvature perturbations  $\mathcal{R}_k$  at a later time,  $t_{\text{start}}$ . From there until each perturbation is well outside of the Hubble horizon, we solve the Mukhanov–Sasaki equation with the RKWKB method (`oscode`) from [Chapter 3](#). Then the primordial power spectrum  $\mathcal{P}_{\mathcal{R}}$  is constructed, which is unchanged until the modes re-enter. To get from the primordial spectrum to the  $C_\ell$ -s of the CMB, Boltzmann codes can be used.

## 5.5 Inference in cosmology

The  $\Lambda$ CDM (short for dark energy ( $\Lambda$ ) + cold dark matter) model is the current best fit model of the cosmos on large scales. The goodness of fit of this model can be compared to another’s via Bayesian *model comparison* and the values of the underlying cosmological parameters (with their uncertainties) can be quantified through *inference*. In the Bayesian context, the probability of an event refers to our degree of belief in its incidence, rather than the limiting relative frequency of it. Conditional probability, denoted  $P(A|B)$  is then the probability that event  $A$  happens given  $B$  occurred, defined as

$$P(A|B) = \frac{P(A \cap B)}{P(B)}. \quad (5.70)$$

Both inference and model comparison rely on Bayes' theorem linking a conditional probability to its reverse:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}. \quad (5.71)$$

In cosmological modelling, we typically construct models  $\mathcal{M}$  after observing some data  $\mathcal{D}$ . The model is characterised by a set of parameters  $\Theta_{\mathcal{M}}$ . When making inference based on some observations, knowledge of the experimental setup and errors associated with the measurement tells us the probability of obtaining the dataset we observed given the model and its parameters,

$$\mathcal{L} \equiv P(\mathcal{D}|\Theta_{\mathcal{M}}, \mathcal{M}), \quad (5.72)$$

which is termed the *likelihood*. However, we are interested in the reverse of this, the probability distribution of possible parameter values of a given model in light of the observations,

$$P(\Theta_{\mathcal{M}}|\mathcal{D}, \mathcal{M}), \quad (5.73)$$

the *posterior*. The two can be related via Bayes' theorem,

$$\mathcal{P} \equiv P(\Theta_{\mathcal{M}}|\mathcal{D}, \mathcal{M}) = \frac{P(\mathcal{D}|\Theta_{\mathcal{M}}, \mathcal{M})P(\Theta_{\mathcal{M}}|\mathcal{M})}{P(\mathcal{D}|\mathcal{M})} \equiv \frac{\mathcal{L}\pi}{\mathcal{Z}} \quad (5.74)$$

where  $\pi$  denotes the *prior*, our pre-experiment knowledge of the distribution of model parameters, and  $\mathcal{Z}$  is the *evidence* or *marginal likelihood*. With the help of the law of total probability, the evidence can be expressed as

$$\mathcal{Z} = P(\mathcal{D}|\mathcal{M}) = \int P(\mathcal{D}|\Theta_{\mathcal{M}}, \mathcal{M})P(\Theta_{\mathcal{M}}|\mathcal{M})d\Theta_{\mathcal{M}}. \quad (5.75)$$

$\mathcal{Z}$  is therefore the likelihood marginalised over the model parameters, and is independent of the parameters themselves.

If we wanted to enquire about the parameters of a given model in light of some observations, we need to calculate the posterior probability distribution  $P(\Theta_{\mathcal{M}}|\mathcal{D}, \mathcal{M})$ , which will be a function of all parameters of interest. It will have the same number of dimensions as the number of parameters, and is therefore often plotted using corner plots, or with all but one parameters marginalised out. For the purposes of parameter inference, i.e. comparing the relative probability of each  $\Theta_{\mathcal{M}}$ , the evidence can be regarded as a normalisation constant, and be neglected. We therefore only need  $\mathcal{L}\pi$  for each parameter, which reflects our prior knowledge  $\pi$  being updated with the likelihood of the observed data  $\mathcal{L}$ .

When there are several models  $\{\mathcal{M}_i\}$  to describe the data and the task is their comparison,  $\mathcal{Z}$  becomes the key quantity. The reverse of this expresses the probability of a model given the

data. Using Bayes' theorem, this posterior of each model can be broken down to

$$P(\mathcal{M}_i|\mathcal{D}) = \frac{P(\mathcal{D}|\mathcal{M}_i)P(\mathcal{M}_i)}{P(\mathcal{D})}, \quad (5.76)$$

with our prior belief in each model,  $P(\mathcal{M}_i)$ , and the evidence appearing on the right-hand-side. A good, non-informative choice is the uniform prior which weighs each model equally. This however means that the relative probability of models is the ratio of their evidences:

$$\frac{P(\mathcal{M}_i|\mathcal{D})}{P(\mathcal{M}_j|\mathcal{D})} = \frac{\mathcal{Z}_i}{\mathcal{Z}_j}. \quad (5.77)$$

At the heart of Bayesian model comparison is the Occam's razor mentality: the simplest model which describes the data wins. The evidence,  $\mathcal{Z}$ , encodes both the goodness of fit and a penalty for overfitting, i.e. using too many parameters. This can be seen from the following deconstruction of the log-evidence [131]:

$$\ln \left[ \int \mathcal{L} \pi d\theta \right] = \int \mathcal{P} \ln \mathcal{L} d\theta - \int \mathcal{P} \ln \frac{\mathcal{P}}{\pi} d\theta, \quad (5.78)$$

log-evidence = parameter fit – Occam penalty,

$$\ln \mathcal{Z} = \langle \ln \mathcal{L} \rangle_{\mathcal{P}} - \mathcal{D}_{\text{KL}}, \quad (5.79)$$

where we shortened  $\theta_{\mathcal{M}}$  to  $\theta$ , and defined the Kullback–Leibler divergence,  $\mathcal{D}_{\text{KL}}$ , which effectively estimates the Occam penalty.

The *Bayesian model dimensionality* [132],  $\mathfrak{d}$ , measures the effective number of constrained parameters (i.e. those the data provide information on), and is defined as the posterior variance of the log-likelihood,

$$\frac{\mathfrak{d}}{2} = \int \mathcal{P} \left( \ln \frac{\mathcal{P}}{\pi} - \mathcal{D}_{\text{KL}} \right)^2 d\theta = \langle (\ln \mathcal{L})^2 \rangle_{\mathcal{P}} - \langle \ln \mathcal{L} \rangle_{\mathcal{P}}^2. \quad (5.80)$$

The ensemble comprised of the (log-)evidence, the posterior average of the log-likelihood (the first term on the right-hand-side of Eq. (5.78)), the Kullback–Leibler divergence, and the Bayesian model dimensionality provides a detailed description of the performance of a given model.

In practice, the posterior will most likely not be one of the very few distributions computer software can sample in closed form, and the evidence will be difficult to evaluate numerically. Instead, the posterior is typically crawled by a Markov Chain Monte Carlo (MCMC) sampler which produces samples of the parameter vector  $\Theta_{\mathcal{M}}$  in proportion to their posterior probability. There are several software packages available for performing MCMC (e.g. `emcee` [133], `CosmoMC` [134]), and for the more direct nested sampling (e.g. `MultiNest` [135], `PolyChord` [136, 137]) which also provides evidences.

### 5.5.1 $\Lambda$ CDM

The current standard model of cosmology (SMC) earned its status by fitting data with high signal-to-noise ratio (and the CMB power spectra in particular) extremely well. This model has been around, and barely changed, since at least 1990 [138], but became particularly difficult to ignore in 1998, with the discovery of cosmic acceleration [139]. The model gets its name from suggesting that most of the matter content in the universe is *cold* and *dark* (CDM), meaning it is collisionless and does not interact electromagnetically, and that the dominant energy density in the universe is from a cosmological constant (also known as dark energy or vacuum energy),  $\Lambda$ .  $\Lambda$ CDM achieves its close fit to data with only 6 (or 7) free parameters. In addition to the 6, there are a number of assumptions underlying it, which are sometimes phrased as fixed parameters. Let us examine the free parameters and assumptions in turn.

#### Free parameters

**Physical baryon density,  $\Omega_b h^2$ .** Baryonic<sup>17</sup> matter density, scaled by the dimensionless Hubble parameter,  $h \equiv H_0/100 \text{ km s}^{-1} \text{ Mpc}^{-1}$ .

**Physical dark matter density,  $\Omega_c h^2$ .** The remaining cold dark matter density in the universe.

**Angular sound horizon,  $\theta_*$ .** The angular size of the sound horizon calculated as  $r_*/d$ , where  $r_*$  is the *sound horizon*, the maximum distance sound waves could have travelled until recombination, and  $d$  is the angular diameter distance to the surface of last scattering.

**Reionisation optical depth,  $\tau$ .** Following recombination, neutral matter collapsed into gravitational potential wells, and began forming increasingly complex structures. When the first stars were formed, their radiation partly re-ionised the neutral matter surrounding them, during the epoch of *cosmic dawn* or *reionisation*. As electrically charged matter interacts electromagnetically,  $\tau$  characterises how transparent the universe was to CMB photons at this time.

**Scalar primordial power spectrum amplitude,  $\ln 10^{10} A_s$ , and tilt,  $n_s$**  The primordial power spectrum of scalar curvature perturbations,  $\mathcal{P}_{\mathcal{R}}(k)$ , shows the amplitude squared of each Fourier mode of the gauge-invariant curvature perturbation,  $\mathcal{R}_k$ , after they exit the comoving Hubble horizon. This spectrum is assumed to have the form  $\mathcal{P}_{\mathcal{R}}(k) = A_s (k/k_*)^{(n_s-1)}$ , with two parameters describing it. The amplitude  $A_s$  is usually defined at the pivot scale  $k_* = 0.05 \text{ Mpc}^{-1}$ . The predecessor of  $\Lambda$ CDM,  $s$ CDM (for ‘standard’ CDM), involved no cosmological constant and assumed a scale-invariant or Harrison-Zeldovich-Peebles primordial power spectrum, with  $n_s = 1$ . Observations of galaxy

<sup>17</sup>Astronomers are known for their liberal definitions when it comes to matter, e.g. elements in the periodic table after hydrogen are routinely called metals, and in the context of cosmology, everything ordinary and non-relativistic is considered baryonic (including electrons).

clustering, however, demanded more power at large angular scales. Many models of inflation naturally predict  $n_s \lesssim 1$ , ensuring a better fit to data.

**Temperature of the CMB today,  $T_0$ .** This seventh ‘free’ parameter is often ignored and taken as fixed due to its value being determined to extremely high precision [5].

### Assumptions and fixed parameters

Below are some of the assumptions the SMC builds on. They have been tested through constraining the relevant parameters in extensions of the standard model, often using a combination of data, e.g. constraints from particle physics experiments, CMB lensing, and BAO.

**Spatial flatness on large scales:**  $\Lambda$ CDM assumes that the geometry of the universe on the largest scales is flat,  $\Omega_{K,0} = 0$ . An extension of the  $\Lambda$ CDM model involving a primordial ( $\Omega_{K,i}$ ) and a present-day curvature parameter ( $\Omega_{K,0}$ ) will be of particular interest to this thesis, as non-zero primordial spatial curvature ensures a ‘shorter’ period of inflation where the initial conditions of perturbations become important and their evolution needs to be computed numerically. CMB data on their own shows a slight preference for a closed universe, and combining different datasets (e.g. CMB + BAO) becomes highly non-trivial due to potential tensions [9, 10, 140].

**Dark energy equation of state parameter:** It is assumed that dark energy has  $w = -1$ , consistent with being vacuum energy or a cosmological constant. Extensions of  $\Lambda$ CDM would allow  $w$  to be different than  $-1$  or vary with time, e.g. quintessence models (see e.g. [141] and references therein).

**Number of relativistic species:** The *effective* number of relativistic species is assumed to be slightly larger than 3, the number of neutrinos in the Standard Model of particle physics. The effective number of species is defined via their effective number of degrees of freedom associated with number density,  $g_{*n}(T)$ , which is a temperature-dependent function. Due to the neutrinos being heated while decoupling from the primordial plasma – which happened much earlier than the decoupling of photons, about 1 s after the Big Bang, due to neutrinos only being weakly interacting –, the effective number of relativistic species is slightly larger than three, number of standard model neutrinos [142]. Since the individual masses and the mass hierarchy of neutrinos are not known,  $\Lambda$ CDM uses the simplifying assumption of one massive and two massless neutrinos, the massive one having the minimum total mass,  $\sum m_\nu = 0.06$  eV.

**Running of the scalar primordial power spectrum:** In base  $\Lambda$ CDM, the spectral index is assumed to be near-constant, and its so-called running,  $n_{\text{run}} = dn_s/d \ln k$  vanishingly

small. Inflationary models involving a single scalar field usually predict a small negative number for  $n_{\text{run}}$ .

**Density perturbations are adiabatic:** Perturbations in a system of multiple fluid components can be entropy/isocurvature or adiabatic in nature, as mentioned in [Section 5.4.2](#). These form a kind of ‘basis’, in that all perturbations can be decomposed into isocurvature and adiabatic parts. They represent distinct models of structure formation: adiabatic fluctuations begin with already established perturbations in matter density in the very early universe and are generated by inflation, while isocurvature modes may start from cosmological defects [\[143\]](#). The models lead to acoustic oscillations in the photon-baryon fluid that are  $\pi/2$  out of phase relative to each other, and thus predict the first (acoustic) peak in the CMB  $TT$  spectrum to be at a different  $\ell$  [\[144\]](#). Although there exist mechanisms which can wash out the acoustic signatures of both models, the observation of a clear acoustic peak at a location consistent with adiabatic perturbations and CMB polarisation data have ruled out a dominant isocurvature contribution, a subdominant one is still possible and predicted by e.g. multi-field inflation models [\[145\]](#).

**Initial conditions were Gaussian:** The field which sourced the primordial perturbations is assumed to have a Gaussian distribution initially. If this were the case, all information about the field (and the anisotropies) would be contained in the two-point correlation function or its Fourier-space equivalent, the power spectrum<sup>18</sup>. Most inflationary models, including the simplest single-field models do, however, predict some degree of non-Gaussianity which can thus be used to distinguish between them. Models that deviate from slow-roll inflation and predict primordial power spectra with features (i.e. not scale-invariant power spectra) generate large but not unreasonable non-Gaussianities. Since these models are also typically the ones which have to be evaluated numerically and may provide a better fit to data than  $\Lambda$ CDM’s scale-invariant spectrum, they are of significant interest of this thesis, and so is any deviation from the assumption of Gaussian initial conditions. The Planck and WMAP missions mapping the CMB provide upper limits on non-Gaussianities, but the sensitivity of current measurements is insufficient to determine which inflationary mechanism is preferred by data. Non-Gaussianities are quantified via the three-point correlation function of perturbations or its Fourier transform, the bispectrum. The latter depends on the length of three momenta,  $k_1, k_2, k_3$ , which form triangles. The overall extent of non-Gaussianities is described by the dimensionless amplitude of the bispectrum,  $f_{\text{NL}}$ . The value of  $f_{\text{NL}}$  for equilateral configurations of the three moments can be linked to an energy scale which determines the strength of interactions of the inflaton (and any additional fields), therefore measuring its magnitude would be decisive in ruling out large classes of inflation models.

<sup>18</sup>In a linear, Gaussian theory of primordial perturbations which start out from a vacuum state, the  $N$ -point correlation functions are identically zero for  $N > 2$ .

**Gravitational waves:** The contribution from gravitational waves is assumed to be negligible. Since primordial tensor modes manifest as gravitational waves in the late universe, the relevant parameter here is  $r$ , the tensor-to-scalar ratio, taken at some pivot scale, e.g.  $k_* = 0.002 \text{ Mpc}^{-1}$ . For each scalar field model of inflation one can compute the level of gravitational waves predicted, which range from large values already ruled out to small ones which are not yet detectable. Since gravitational waves can source  $B$ -mode polarisation in the CMB, the level of which may therefore provide new information on inflation.

### 5.5.2 Tensions

A tension is a small divergence of a model from observations not significant enough to be called a discrepancy. The tensions we mention here are at the  $\sim 3\sigma$  level, which should be compared to the  $\sim 1000\sigma$ 's worth of information in CMB data. There are a lot of potential tensions at the  $3\sigma$  level one could select from this amount of data, but the hope is to select those that have the potential to grow into significant disparities.

Currently, one of the most researched tensions is one concerning the present value of the Hubble parameter  $H_0$ : distance ladder measurements from type IA supernovae, calibrated by nearby Cepheids or tip of the red giant branch (TRGB) stars prefer a larger value than that obtained from the CMB. Another example is the tension in the matter power spectrum normalisation  $\sigma_8$ , which can be probed by galaxy cluster counts. CMB measurements on their own have shown a small, but persistent preference for closed universes, which is in disagreement with the flatness assumption of  $\Lambda$ CDM. Besides the tensions, one can identify anomalies (low-significance features) in CMB temperature data that seem to be in disagreement with  $\Lambda$ CDM, such as a lack of power at low multipoles, even-odd multipole asymmetry, the CMB Cold Spot, etc. While one should be sceptical of claims that argue strongly for beyond- $\Lambda$ CDM physics based on low-level anomalies or tensions, our theoretical and numerical toolset must be ready should the need to expand beyond the standard model arise [5, 8, 146–148]. More specifically, our numerical methods should be able to handle models which require analytic assumptions to be replaced with approximate numerical calculations efficiently, e.g. for the computation of the PPS of scalar and tensor perturbations, for computing their bispectrum, or the transfer functions. Numerical PPS are a necessity in models that not only wish to extend  $\Lambda$ CDM with a present-day curvature parameter but want to study the implications of primordial curvature; efficient bispectra would enable the wider exploration of primordial non-Gaussianities; and faster computation of the transfer functions would streamline the inference process overall (this being the current computational bottleneck). This need for fast numerical solvers motivates the work presented in the previous [Chapters 3 and 4](#). The numerical tools introduced therein are put to the test on computationally challenging, but promising alternative models in [Chapter 6](#). A large class of models which promise a better fit to CMB data (e.g. just enough inflation models)

retain memory of the initial conditions set to the primordial curvature perturbations. The method by which these conditions are set therefore gains importance, and must be theoretically robust. [Chapter 7](#) highlights this point and weighs sets of initial conditions against each other.

# Chapter 6

## Oscode in cosmology

This chapter explores the problem that motivated the development of `oscode`: solving the Mukhanov–Sasaki equation [90], the equation of motion of primordial curvature perturbations, within computationally challenging inflationary models. The Mukhanov–Sasaki equation has the form of a generalised oscillator with a time-dependent frequency and a first-order derivative term present, the frequency depending on the characteristic wavenumber of the perturbation. For inference in cosmology from the CMB, it is necessary either to assume an approximate form for the primordial power spectrum of curvature perturbations, or to solve the Mukhanov–Sasaki equation for a range of characteristic wavenumbers to compute a spectrum (see, e.g. [89] for a thorough review). In the case of single-field slow-roll inflation, most models lead to a scale-invariant primordial power spectrum [84] which can be obtained analytically [149], but models that introduce features in the primordial power spectrum can improve the fit to CMB observations [150]. In such cases when one relies on a numerical solution of the Mukhanov–Sasaki equation, the numerical method used often has to track the highly oscillatory solution over many wavelengths. Runge–Kutta-based methods are therefore generally impractical for this task. Some Runge–Kutta solvers such as BINGO [151] can, however, prove efficient for some single-field inflation models by taking a shortcut and not integrating the perturbation throughout its oscillatory phase [152]<sup>1</sup>.

**Section 6.1** briefly recaps the equations describing the evolution of primordial perturbations, and introduces existing tools for solving them. The focus will be on BINGO (BI-spectra and Non-Gaussianity Operator), a highly optimised code `oscode` will be compared to in **Section 6.2**. **Sections 6.3** and **6.4** showcase the power of `oscode` when it comes to computationally intensive and unexplored models: a universe starting from kinetic dominance, and a models with non-zero primordial spatial curvature.

---

<sup>1</sup>though not all models admit such shortcuts.

## 6.1 The Mukhanov–Sasaki equation

In the toy examples discussed in [Chapter 3](#), there was only a frequency,  $\omega$ -term present in the differential equation to be solved, and it was available to arbitrary precision. This may not always be the case, as (1) one may want to switch to a more physically meaningful independent-dependent variable pair, which can introduce a friction term  $\gamma$ , and (2) the frequency and friction terms might themselves be available only through numerically solving a set of differential equations. The Mukhanov–Sasaki equation illustrates both of these cases.

The Mukhanov–Sasaki equation describes the time-evolution of perturbations in a background universe with an FLRW metric. This background, in the simplest models, assumes the presence of a single time-dependent scalar field  $\phi(t)$ . The field has self-interactions described by the potential  $V(\phi)$ , and its dynamics are defined by the equations

$$H^2 + \frac{K}{a^2} = \frac{1}{3} \left( \frac{1}{2} \dot{\phi}^2 + V(\phi) \right), \quad (6.1)$$

$$\dot{H} + H^2 = -\frac{1}{3} \left( \dot{\phi}^2 - V(\phi) \right), \quad (6.2)$$

$$0 = \ddot{\phi} + 3H\dot{\phi} + V_{,\phi}, \quad (6.3)$$

out of which only two are independent. The quantities involved were defined in [Section 5.3.1](#). In this section, flat and closed universe models will be considered, starting with the flat case, therefore  $K = 0$  until stated otherwise. Perturbing the field and the metric, and introducing the gauge-invariant scalar  $\mathcal{R}$  (the comoving curvature perturbation) one can then arrive at the Mukhanov–Sasaki equation (as outlined in [Section 5.4.3](#)), which we write as

$$\ddot{\mathcal{R}}_k + 2 \left( \frac{\ddot{\phi}}{\dot{\phi}} - \frac{1}{2} \dot{\phi}^2 + \frac{3}{2} \right) \dot{\mathcal{R}}_k + \left( \frac{k}{aH} \right)^2 \mathcal{R}_k = 0. \quad (6.4)$$

In the above equation, the overdot denotes differentiation with respect to  $N = \ln a$ , and  $\mathcal{R}_k$  is the mode with wavenumber  $k$  in the Fourier decomposition of  $\mathcal{R}$ . During inflation,  $a(t) \sim e^{Ht}$  and  $H$  is approximately constant, hence  $N$  is a natural independent variable candidate.

In the limit of slow-roll inflation,  $\frac{1}{2} \dot{\phi}^2 \ll V(\phi)$ , the background [Eq. \(6.3\)](#) admit analytic solutions. They also do in kinetic dominance [[116](#)], the opposite limit  $\frac{1}{2} \dot{\phi}^2 \gg V(\phi)$ . Kinetic dominance has been shown to be a likely phase the Universe went through before inflation in most single-field models [[115](#), [117](#), [118](#)]. Both limits can thus be used to set initial conditions to [Eq. \(6.3\)](#), which can then be integrated numerically.

The Mukhanov–Sasaki equation in the flat case can also be solved analytically if for all  $k$ -modes of interest,  $k \gg aH$ . Since  $k^{-1}$  is the characteristic lengthscale of a perturbation mode, this means that all modes of interest are assumed to be well inside the Hubble horizon. Letting the Mukhanov–Sasaki equation emerge from this limit is equivalent to choosing a vacuum state (see [[86](#)]) which, together with a normalisation condition, are enough to provide initial

conditions for the mode functions  $\mathcal{R}_k$ . This choice of vacuum and the initial conditions are referred to as *Bunch–Davies*. With a different choice of vacuum, it is possible to set initial conditions on the  $\mathcal{R}_k$  in kinetic dominance, when modes are not necessarily inside the Hubble horizon. The form of  $\mathcal{R}_k$  in the kinetically dominated limit are derived in [12]. In the models to be investigated here, both slow-roll and kinetically dominated initial conditions (for the cosmological background variables and the perturbations alike) shall be considered.

The Mukhanov–Sasaki Eq. (6.4) is of the form of a generalised oscillator with a first-derivative  $\gamma$  term present, with both  $\gamma$  and the frequency  $\omega$  being (in general non-analytic) functions of time as they depend on the cosmological background. It follows that when a  $k$ -mode is inside the Hubble horizon,  $k > aH$ , it oscillates with some varying amplitude and frequency usually proportional to  $k$ , and one can show that the mode ‘freezes out’ once outside the Hubble horizon, meaning  $\mathcal{R}_k \sim \text{const}$ . The wavenumber-dependence of the frequency term makes this equation challenging to solve for large values of  $k$  without resorting to approximations.

The goal is to solve the Mukhanov–Sasaki for a range of  $k$ -modes (including large values of  $k$ ) until after horizon exit when each mode has a constant amplitude, in order to obtain the primordial power spectrum

$$\mathcal{P}_{\mathcal{R}}^2(k) = \frac{k^3}{2\pi^2} |\mathcal{R}_k|^2. \quad (6.5)$$

## 6.2 Comparison with BINGO

We shall first adopt the computational strategy employed by many solvers designed to compute primordial power spectra, for example BINGO [151], and ModeCode [153], `oscode`’s performance with the former. BINGO is a `Fortran`-based code for efficient evaluation of the scalar bispectrum which necessarily calculates the PPS of scalar perturbations on the way. In the comparison, BINGO’s input parameters are modified to only carry out the PPS calculation.

BINGO gets around the computational challenges posed by high-frequency oscillation by using a trick: it has been shown that in the case of a single-field inflationary model and assuming the universe emerges from slow-roll inflation, it is sufficient to evolve each curvature perturbation from a time they are well inside the Hubble horizon (from, say,  $k/aH = 100$ , see [152]), until the perturbation freezes out outside of the Hubble horizon ( $k/aH = 10^{-2}$ ). This avoids integrating the solution through the majority of its oscillatory phase. For the comparison to be fair, `oscode`’s calculations are set up to do the same. First, the cosmological background ( $\phi(N)$ ,  $\dot{\phi}(N)$ , ...) is integrated numerically, starting from slow-roll conditions, set such that the total number of e-folds of inflation,  $N_{\text{tot}} \sim 60$ . The inflationary model used in this example involves a quadratic potential,

$$V(\phi) = \frac{1}{2}m^2\phi^2, \quad (6.6)$$




Fig. 6.1 Primordial power spectra of the gauge-invariant scalar curvature perturbations, generated by BINGO and the `oscde`-based solver with identical parameters. The spectra have been computed up to extremely large values of  $k$  for the sake of comparing the runtimes of the two codes.

where the inflaton mass is set to one,  $m = 1$ . The initial scale factor is set such that a pivot mode, corresponding to  $k = 0.05 \text{ Mpc}^{-1}$  leaves the Hubble horizon when there are 50 e-folds of inflation left. For each mode, the  $N$  values corresponding to the start and end of integration are found, then `oscde` is given  $\omega(N)$  and  $\gamma(N)$  as one-dimensional arrays. The code performs linear interpolation on these arrays (for which the sampling of the arrays needs to be sufficiently fine - for the present example  $5 \times 10^5$  equally spaced points were used between  $N = 0$  and  $N = 75$ ). `Oscde` then proceeds to solve the mode evolution for each  $k$  starting from Bunch–Davies initial conditions. Identical parameters are given to BINGO and the `oscde`-based solver, including a relative tolerance of  $10^{-4}$  and an absolute tolerance of 0. The computed power spectra are identical, as shown in Fig. 6.1.

The computation time for the solvers to obtain  $\mathcal{R}_k$  is measured and plotted as a function of  $k$  in Figs. 6.2 and 6.3. The former shows the ratio of BINGO and `oscde`’s runtimes as a function of  $k$ , and the latter just that of `oscde`, relative to the median  $k$ . Together they show that BINGO’s runtime is logarithmic in  $k$ , while `oscde`’s is constant. They also show that `oscde` performs better than BINGO by at least a factor of two, and at most a factor of 4 in the  $k$ -range of interest.

This can be explained by looking at Eq. (6.4). The frequency term is the relative size of the perturbation with respect to the Hubble horizon, which is a fixed number at the start and end of integration of each  $k$ -mode, so it will no longer scale with  $k$ . The friction term is approximately constant during inflation. The integration range,  $\Delta N$ , is determined by the points where  $k/aH = c_0$  (a constant), which is also roughly constant during inflation. Therefore the number of oscillations the perturbations undergo over the range of wavenumbers in the PPS barely changes, so one can expect a WKB-based method to traverse the oscillations in constant




Fig. 6.2 Ratio of the runtime of BINGO and our solver during the evaluation of a scalar primordial power spectrum, as a function of wavevector.




Fig. 6.3 Progression of our solver's relative runtime with increasing wavenumber, whilst calculating a scalar primordial power spectrum. The times are referenced to the computation time corresponding to the median  $k$ -value, indicated by the dotted lines.

time. In reality, the integration range increases slowly with  $N$ , and small variations in the friction term cause the oscillations to change in shape, hence the slow increase in the runtime of BINGO. The two-fold runtime-difference present even at the smallest values of  $k$  can be explained by the difference in the number of steps taken. Fig. 6.4 shows the intermediate steps taken by RKSUITE, a numerical routine implementing efficient Runge–Kutta methods and used by BINGO, and the intermediate steps taken by `oscode`, whilst computing the time-evolution of a single  $k$ -mode. `oscode` is able to traverse the oscillatory region of the mode's evolution in significantly fewer steps than the Runge–Kutta method, giving a reduction in computing time.

In models where one has to start integrating the mode equation from deeper within the horizon, the starting frequency during the evolution of modes is larger, and the performance




Fig. 6.4 Comparison of BINGO and our solver in the evolution of a single perturbation with wavevector  $k = 10^{-5} \text{ Mpc}^{-1}$ . The black reference line is a dense solution generated with a Runge–Kutta (7,8)<sup>th</sup> order pair. On top of it the top panel shows the steps that RKSUITE’s (4,5)<sup>th</sup> order Runge–Kutta solver takes (a total of  $\sim 150$ ), the bottom panel the steps that our solver takes (a total of  $\sim 60$ ). The relative tolerance was set to  $10^{-4}$  for both methods.

difference between an RKWKB-based approach and a Runge–Kutta integrator is even more pronounced. Examples include universes emerging from kinetic dominance, axion monodromy models [154] or models with alpha vacua initial conditions [155].

### 6.3 A model using kinetic dominance

Inflationary models involving kinetic dominance have already been investigated [117, 118], but previous articles assumed the Bunch–Davies vacuum as the initial state of perturbations such that the perturbations start their evolution from a universe that just began inflating (i.e. the perturbations never evolve on a kinetically dominated background). In these works, the cosmological background in terms of  $N$  is integrated from the initial state

$$\phi = \phi_P - \sqrt{6} \ln N, \quad (6.7)$$

$$\dot{\phi} = -\sqrt{6 - \frac{2V}{H^2}}, \quad (6.8)$$

$$H = \frac{1}{3} e^{-3N}, \quad (6.9)$$

where  $\dot{\phi}$  contains a contribution from the potential in order to make the system numerically stable. In kinetic dominance, [12] obtains an analytic solution for the perturbation modes,




Fig. 6.5 Scalar primordial power spectrum of perturbations emerging from kinetic dominance. The mode equation was solved from a fixed, early time (well inside kinetic dominance) until long after horizon crossing, which is only feasible if the solver used is capable of traversing many oscillations at once. The relative tolerance was set to be  $10^{-4}$ .

which in terms of  $N$  takes the form

$$\mathcal{R}_k = \sqrt{\frac{3\pi}{8}} \frac{1}{z} e^N \left[ A_k H_0^{(1)} \left( \frac{3}{2} k e^{2N} \right) + B_k H_0^{(2)} \left( \frac{3}{2} k e^{2N} \right) \right], \quad (6.10)$$

$$\dot{\mathcal{R}}_k = \sqrt{\frac{27\pi}{8}} \frac{k}{z} e^{3N} \left[ A_k H_0^{(1)\prime} \left( \frac{3}{2} k e^{2N} \right) + B_k H_0^{(2)\prime} \left( \frac{3}{2} k e^{2N} \right) \right] + \left( 1 - \frac{\dot{z}}{z} \right) \mathcal{R}_k, \quad (6.11)$$

where  $H^{(1)}$  and  $H^{(2)}$  are Hankel functions of the first and second kind, and  $A_k, B_k$  are constants. We set  $A_k = 0$  and  $B_k = 1$ , and chose parameters such that the total number of e-folds during inflation,  $N_{\text{tot}} \approx 60$ , and the pivot scale corresponding to  $k = 0.05 \text{ Mpc}^{-1}$  today leaves the horizon when there are  $N_* \approx 54$  e-folds of inflation left. We set the initial conditions for the background at  $N = 0$ , and for the modes at a constant  $N = 1.1$ , and integrate until far after horizon exit, as in [Section 6.2](#).

The resulting primordial power spectrum is shown in [Fig. 6.5](#). Such computations are only possible if the solver used can trace oscillations in the solution extremely efficiently, and indeed we found that calculating a spectrum starting from kinetic dominance and from a fixed fraction of the horizon in slow-roll can be carried out on similar timescales using `oscode`. It is worth noting that a fast solver has been developed specifically for the Mukhanov–Sasaki equation [\[78\]](#) that works on the basis of using analytic approximations of the solution when the frequency is well approximated by an exponential or a linear polynomial. This gives a significant speed-up over Runge–Kutta methods, but relies on the Mukhanov–Sasaki equation to be transformable to

a form without a first-order derivative term. Closed universe models do not have this property, but can still be investigated with `oscode`.

## 6.4 A closed universe model

In this example we consider closed universe models ( $K = 1$ ) with primordial curvature. The cosmological background evolution Eqs. (6.1) to (6.3) can be cast into a system of linear ODEs in terms of a new dependent variable  $\Omega_k$ ,

$$\frac{d \ln |\Omega_k|}{dN} = 4 + |\Omega_k| (4K - 2a^2 V(\phi)), \quad (6.12)$$

$$\left( \frac{d\phi}{dN} \right)^2 = 6 + |\Omega_k| (6K - 2a^2 V(\phi)), \quad (6.13)$$

where  $\Omega_k = \frac{K}{(aH)^2}$ . We shall consider a cosmological background emerging from kinetic dominance, such that the Hubble horizon,  $(aH)^{-1} = \sqrt{\Omega_k}$ , grows until it reaches a maximum  $\sqrt{\Omega_k^i}$  at e-folds  $N_i$ . From this point the horizon shrinks, and inflation starts. The parameters  $(\Omega_k^i, N_i)$ , together with the requirement  $\dot{\Omega}_k(N_i) = 0$  fully fix the background evolution, and hence determine the amount of inflation,  $N_{\text{tot}}$ . Brent's method of root finding [156] was used to search for the  $N_i$  for a given  $\Omega_k^i$  that yields  $N_{\text{tot}} = 60$ . Hence the primordial power spectra have all other parameters fixed, with only  $\Omega_k^i$ , the initial curvature at the start of inflation, changing. Integration of the background is started from  $N_i$  and is performed forwards until the end of inflation (and backwards, if necessary) to cover the integration range of the perturbation modes.

The mode functions obey the generalised Mukhanov–Sasaki equation in the presence of non-zero spatial curvature,  $K \neq 0$ , with frequency and first-derivative terms given by [157]

$$\omega^2 = \Omega_k \left( (k_2 - K) - \frac{2Kk_2}{EK + k_2} \frac{\dot{E}}{E} \right), \quad (6.14)$$

$$2\gamma = K\Omega_k + 3 - E + \frac{k_2}{EK + k_2} \frac{\dot{E}}{E}, \quad (6.15)$$

where  $E = \frac{1}{2}\dot{\phi}^2$  and

$$k_2 = \begin{cases} k(k+2) - 3K, & \text{if } K > 0, \\ k^2 - 3K, & \text{otherwise.} \end{cases} \quad (6.16)$$

The modes are started from  $N = N_i$  using the Bunch–Davies conditions introduced in Section 6.1. Although the Bunch–Davies solution has been derived from the Mukhanov–Sasaki equation in a flat universe, its closed universe equivalent is not yet known. An important feature of closed universe primordial power spectra is that the values of the *comoving wavenumber*  $k$ , appearing in the above equations, are quantised to only take integer values, with the lowest possible value of  $k = 3$  [158]. The comoving wavenumber, measured in Planck units, is related to the physical

scale of the perturbation today via

$$k_{\text{today}} = \frac{k}{a_0}, \quad (6.17)$$

where  $a_0$  is the present day scale factor, given in terms of the present day reduced Hubble parameter,  $h \equiv H_0 / (100 \text{ km s}^{-1} \text{ Mpc}^{-1})$ , and the present day curvature density,  $\Omega_{k,0}$ , by

$$a_0 \approx 4.3 \times 10^4 \left( \frac{h}{0.7} \right)^{-1} \left| \frac{\Omega_{k,0}}{0.01} \right|^{-\frac{1}{2}} \text{ Mpc}. \quad (6.18)$$

**Fig. 6.6** shows the resulting (quantised) primordial power spectra for various values of initial curvature  $\Omega_k^i$ , each with an associated spectrum treating comoving  $k$  as a continuous variable plotted underneath. Calculating the spectra with `oscode` provided roughly three orders of magnitude reduction in computing time compared to Runge–Kutta-like methods.

A systematic analysis of universes with non-zero primordial and present day curvature has been carried out by [159, 160], which relies on `oscode` to compute primordial power spectra numerically. This work was motivated by the recurring preference for closed universes in CMB temperature and polarisation data [150, 161–166]. The possibility of non-zero present day curvature would significantly limit the number of e-folds of inflation, thus increasing the likelihood of just enough inflation models with a kinetically dominated phase preceding slow-roll inflation. Inference in this work is based on CMB data from the Planck 2018 legacy release [167] (temperature and  $E$ -mode polarisation) and from the 2018 data release<sup>2</sup> of BICEP2 and the Keck Array [168] ( $B$ -mode polarisation), but does not make use of CMB lensing data or BAOs due these being in tension for closed universes [9–11]. The analysis expands on previous studies by going beyond simply extending  $\Lambda$ CDM with a present-day curvature parameter  $\Omega_{K,0}$ . It also considers inflationary models with additional parameters describing the primordial evolution of perturbations, and uses fully numerically integrated primordial power spectra. The new parameters are  $\phi_i$  and  $f_i \equiv \Omega_{K,i} / \Omega_{K,0}$ , the value of the inflaton field at the start of inflation and the fraction of primordial to present day curvature density, respectively. The analysis considers inflaton potentials typically used in this type of comparisons the Planck inflation papers [169]: quadratic, double-well (quartic), Starobinsky, and natural [170]. The fully numerically computed primordial power spectra contain important details (oscillatory features at low- $k$ ) which are lost if they are to be approximated analytically, as demonstrated by **Fig. 6.7**. The study computed Bayesian posteriors and evidences for the four inflaton potentials and found an improved fit to CMB data due to a suppression of power and smoothing of peaks on large scales (low- $\ell$ ) in closed universes. The  $TT$ ,  $TE$ ,  $EE$  and low- $E$  data did not show particular preference for any of the four potentials, but with  $B$ -mode polarisation data included, the Starobinsky potential outperformed others significantly. These results are illustrated in **Fig. 6.8**. Overall, the paper confirmed that Planck 2018 data prefer a slightly closed universe,

<sup>2</sup>which corresponds to the 2015 observing season

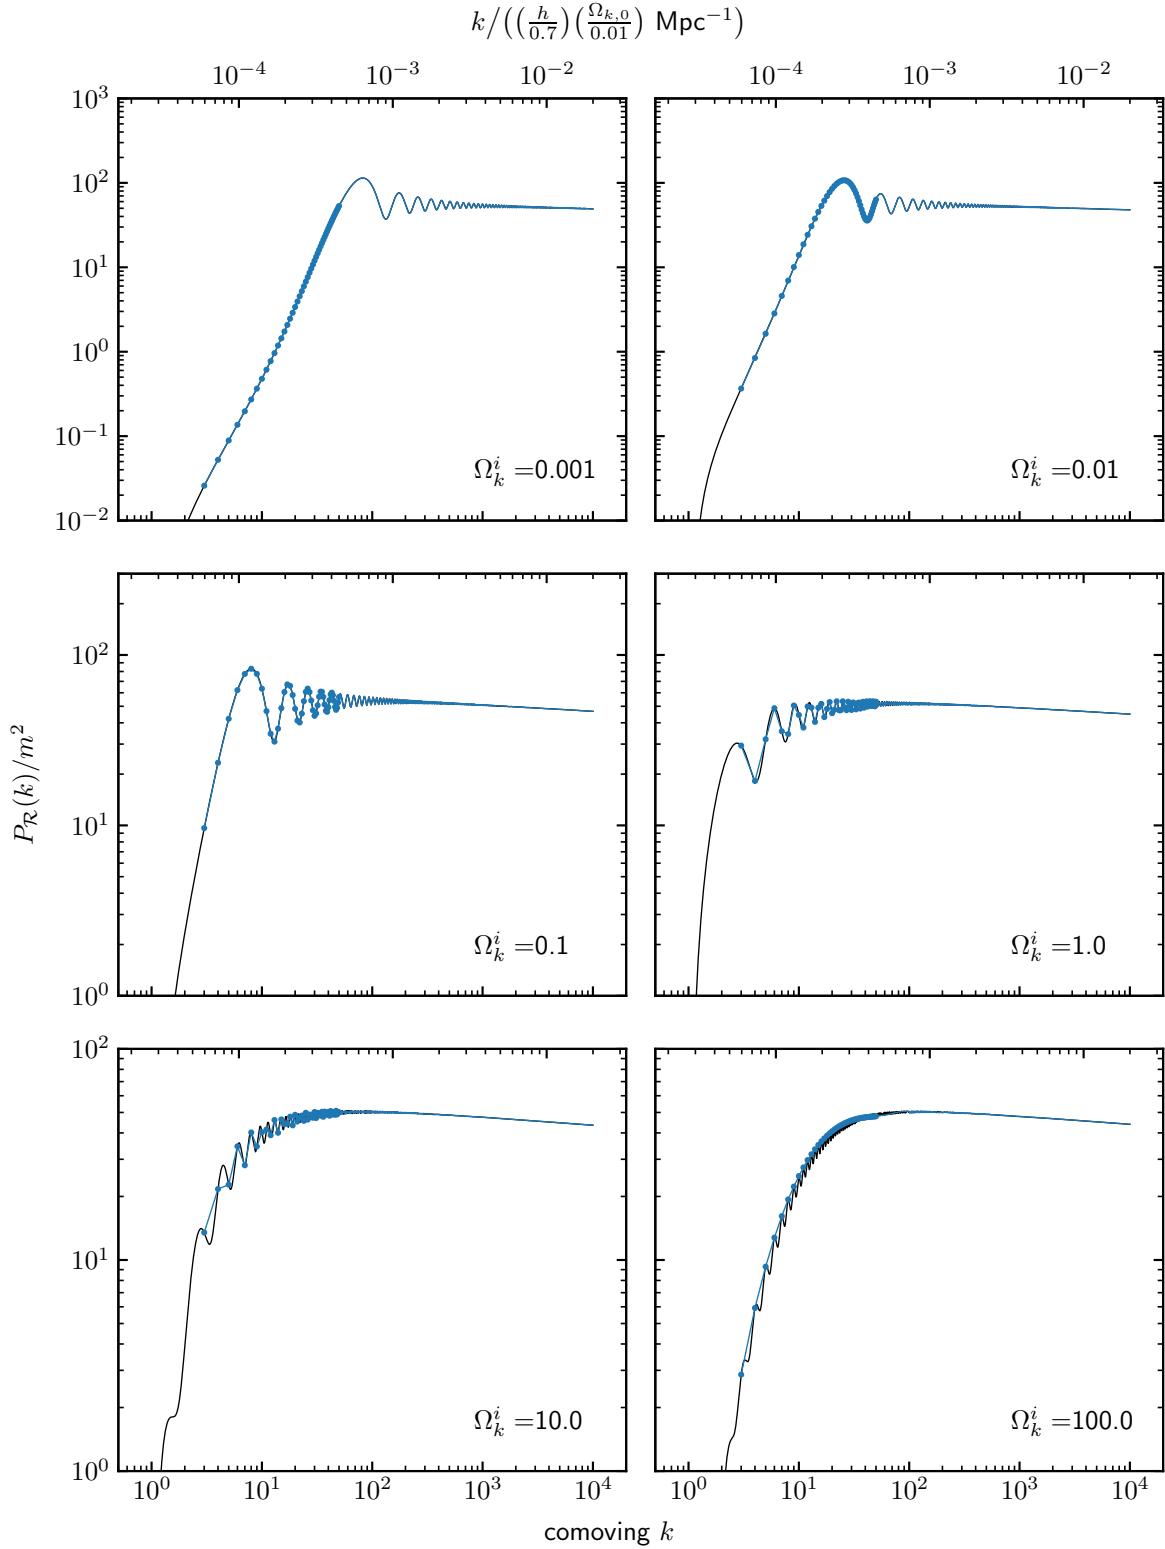


Fig. 6.6 Scalar primordial power spectra in universes with varying initial curvature. The start of inflation,  $N_i$  is adjusted to vary with the curvature at the start of inflation,  $\Omega_k^i$ , such that the total e-folds of inflation,  $N_{\text{tot}} = 60$  is constant. In curved universes, only integer values of comoving  $k$  are allowed, with  $k \geq 3$  (continuous line with  $k \leq 50$  highlighted), but for clarity we include the continuous spectrum (dashed line). The modes are started from the Bunch–Davies vacuum.




Fig. 6.7 The figure shows numerical (left) and approximate (right) primordial power spectra of closed universes with varying primordial-present day curvature ratio  $f_i \equiv \Omega_{K,i}/\Omega_{K,0}$ . The present day curvature density, is fixed to  $\Omega_{K,0} = -0.01$ , as are other parameters. The PPS of scalar perturbations are plotted with heavier lines and have a higher amplitude than those of the tensor perturbations, plotted with thin lines. The shaded area represents the range of observable perturbation wavenumbers, and the dotted, horizontal line signals the pivot scale  $k_* = 0.05 \text{ Mpc}^{-1}$ . To generate this figure, a background cosmology starting from kinetically dominated initial conditions was used, and the perturbations are initialised at the start of inflation from an initial state that minimises the 00-component of the renormalised stress–energy tensor (which will be discussed in [Section 7.4.1](#)). A Starobinsky inflationary potential was used. Figure adapted from [[159](#), [160](#)].

modelled either as an extension of  $\Lambda$ CDM with a present-day curvature parameter, or as a universe emerging from kinetic dominance with non-zero primordial curvature.

## 6.5 Conclusions

This chapter showed how `oscode` can be applied to solve Mukhanov–Sasaki equation, an equation of a generalised oscillator in which both the time-dependent frequency and friction term need to be computed numerically in advance. `oscode`’s performance was compared to that of `BINGO`, a highly efficient `Fortran` code that computes the scalar bispectrum by first computing a primordial power spectrum of scalar curvature perturbations using a fast Runge–Kutta solver available from `RKSUITE`. We measured for each wavenumber  $k$  how long each code takes to compute a solution to the Mukhanov–Sasaki equation from sub-horizon ( $k/aH = 100$ ) to super-horizon ( $k/aH = 0.01$ ) times with all parameters identical, and found that `oscode` takes constant time in  $k$ , being approximately twice as fast as `BINGO` in the observational range. Following from the theoretical considerations in [Section 3.6](#), since `BINGO` uses a Runge–Kutta method whose runtime scales with the frequency or number of oscillations linearly, the performance gap between `BINGO` and `oscode` would be increased further (roughly linearly with the frequency or number of oscillations) if oscillations have to be tracked for a

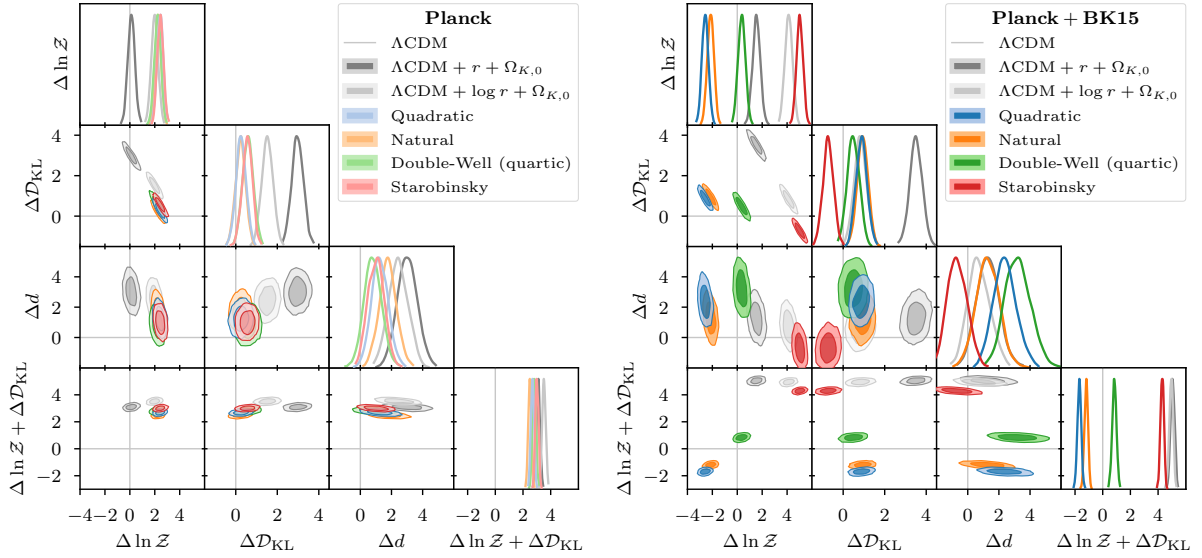


Fig. 6.8 Bayesian model comparison of inflation models within curved universes with finite inflation. The left side (in light hue) shows results using Planck 2018  $TT, TE, EE + lowE$  data only, while on the right side (in darker hue), data from the BICEP and Keck Array have also been included. The statistical quantities plotted are the log-evidence  $\ln \mathcal{Z}$ , the Kullback-Leibler divergence  $\mathcal{D}_{KL}$  (in nats), the 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  $\Delta$  signals that quantities have been normalised with respect to the base  $\Lambda$ CDM model without extensions (with  $r = 0$  and  $\Omega_{K,0} = 0$ ), marked by vertical and horizontal grey lines. As an additional reference, results for an extension of  $\Lambda$ CDM with tensor modes and curvature have also been included, in grey. The probability distributions represent uncertainty arising from the nested sampling process. Figure taken from [159, 160].

longer period, e.g. when the perturbation modes are initialised further inside the Hubble horizon. To show that `oscode` is agnostic to the frequency of oscillations, it was used to integrate each perturbation mode from a single fixed time through horizon entry and exit (rather than from a time each perturbation scale is a given fraction of the horizon), starting from kinetically dominated initial conditions for both the background variables and the perturbations. The last cosmological usecase of `oscode` presented here was provided by models with non-zero primordial and present-day spatial curvature. In this regime, the Mukhanov–Sasaki equation cannot be transformed into a first-derivative-free form (i.e. one with  $\gamma = 0$ ), which excludes other efficient numerical methods for oscillatory ODEs such as [78] and those introduced in Section 2.2.4. Results from [159, 160], a systematic study of curved universe models with finite inflation that uses `oscode` to generate numerically computed primordial power spectra, were summarised.

## Chapter 7

# Quantum initial conditions for primordial perturbations

Never trust an experimental result until it has been confirmed by theory.

---

*Sir Arthur Eddington*

We take a closer look at setting vacuum initial conditions for the perturbations appearing in the Mukhanov–Sasaki equation. The initial conditions will be investigated from the point of view of certain transformations of the associated action describing the perturbations. The transformations considered are classified as canonical transformations, which leave the equation of motion of the perturbations and the commutator structure intact. They have also been shown to preserve the time-evolution of expectation values of quantum operators, unless the vacuum state associated with the perturbations also changes under the transformation. These properties suggest that it would be of interest to find vacuum prescriptions (and therefore initial conditions) that also remain unchanged under canonical transformations. The chapter shows that initial conditions derived via minimising the vacuum expectation value of the Hamiltonian and those obtained using the Danielsson vacuum prescription are *not* invariant under these transformations, whereas those obtained by minimising the local energy density are invariant. The range of physically distinct initial conditions obtainable by Hamiltonian diagonalisation is derived, and the effect of choosing non-invariant initial conditions on the scalar primordial power spectrum and CMB observations is illustrated. I argue that the invariance of the vacuum prescription obtained by minimising the renormalised stress–energy tensor should make it the preferred procedure for setting initial conditions for primordial perturbations. All other procedures reviewed in this work yield ambiguous initial conditions, which is problematic both in theory and in practice.

## 7.1 Introduction

In [Chapter 6](#), the analysis of the inflationary models involving kinetic dominance and non-zero spatial curvature relied on the numerical solution of the Mukhanov–Sasaki equation of curvature perturbations. A particular solution to the Mukhanov–Sasaki requires two initial conditions for the Fourier modes of perturbations, which are usually motivated by quantum mechanical vacuum considerations. Often the initial quantum state for the primordial perturbations is chosen so as to minimise the Hamiltonian density. In an expanding spacetime, however, the Hamiltonian becomes time-dependent, leading to the ground state at a given time no longer being the ground state at a later time. The divergent Hamiltonian then yields infinite particle density at times other than the instant the initial state is set [[12](#), [171](#)]. Handley et al. [[12](#)] thus suggest minimisation of a quantity describing the local energy density instead, and derive a different set of vacuum conditions. The above both define the initial state as some minimal energy density state, but there are many others that instead define it as the state annihilated by annihilation operators, such as one proposed in [[172](#), [173](#)], the  $\alpha$ -vacua [[155](#), [174–177](#)], and others.

In this chapter, the aforementioned procedures for setting initial conditions are reviewed from the perspective of invariance under a set of transformations that are canonical. Canonical transformations leave the classical evolution and the commutator structure of the system invariant by construction. The behaviour of expectation values of quantum operators under canonical transformations is not as obvious: [[178](#), [179](#)] both argued that since the wavefunction only changes by a field-dependent phase, the expectation values must be invariant, but while [[178](#)] concludes that vacua arising from different choices of canonical variables are equivalent, [[179](#)] observes that the vacuum (the state that minimises the vacuum expectation value of the Hamiltonian) does in fact change. In [[180](#)], the authors then proved formally that if the two sets of canonical variables select out different vacuum states (as is the case with Hamiltonian diagonalisation), then there will be differences in expectation values and the observable differences depend only on the canonical transformation at the time when the initial state (e.g. the vacuum) is set. If canonical variables yield the same initial (vacuum) state, then the expectation values of operators will also match. This suggests that it is of interest to investigate vacuum prescriptions that also remain unchanged under canonical transformations. Such a prescription would unambiguously define initial conditions for the perturbation mode functions, and thus would be the preferred choice for setting the vacuum. This is because any ambiguity arising from other choices in the initial conditions, which do not share this invariance, would lead to uncertainty in the frozen-out amplitude of the mode functions, the primordial power spectrum, and ultimately in the power spectrum of CMB anisotropies.

As first observed by Fulling in 1979 [[171](#)], the initial conditions resulting from the popular Hamiltonian diagonalisation method (minimising the Hamiltonian density) do in fact change under canonical transformations, a fact that he used to argue against the Hamiltonian diagonal-

isation procedure. His observations will be used again here to suggest instead the minimisation of the local energy density via the renormalised stress–energy tensor for setting initial conditions. Weiss [181] then added field redefinitions to the group of canonical transformations considered by Fulling, but instead came to the conclusion that a preferred Hamiltonian can be picked out for its desirable mathematical properties. This preferred Hamiltonian is the one written in terms of conformal time  $\eta$  and the Mukhanov variable  $v$ . The preferred Hamiltonian coincides with the one considered conventionally, for its action describes a canonically normalised scalar field. A recent study [180] provides a thorough review of canonical transformations and their effect on scalar field fluctuations during inflation, observing that the transformations can be used to select out different vacuum states.

In addition to reinforcing Fulling’s findings that the Hamiltonian diagonalisation vacuum is a canonically non-invariant procedure and demonstrating how this property may affect observations, this chapter shows that the Danielsson vacuum suffers the same pathology, whereas minimising the local energy density through the renormalised stress–energy tensor does not. The latter two procedures were not available to Fulling at the time. Despite the potential theoretical shortcomings of some vacua, the choice of method for setting the vacuum state may not influence any observable quantities (such as CMB temperature and polarisation spectra) if the initial conditions are set when the perturbation modes of interest lie well within the Hubble horizon. Nonetheless, observable differences do occur in particular for ‘just enough inflation’ models. Such models may be particularly relevant for closed cosmologies [158, 162, 182–184], but the analysis of initial conditions in the presence of spatial curvature will be explored in a future study.

Section 7.2 of this chapter covers the relevant mathematical and physical background, with Section 7.2.2 summarising the classical theory of inflationary perturbations and Section 7.2.3 reviewing how the vacuum choices considered arise. Section 7.3 then explains what transformations of the vacuum-setting procedures are carried out in this work. Results are presented in Section 7.4, broken into subsections to show the effect of the canonical transformations considered, under each vacuum prescription, on the initial conditions, the primordial power spectrum of scalar curvature perturbations, and the temperature-temperature power spectrum

symbol	meaning
$t$	cosmic time
$\eta$	conformal time
$\tau$	an arbitrary timelike independent variable
$\dot{f}$	$\frac{\partial f}{\partial t}$
$f'$	$\frac{\partial f}{\partial \eta}$
$\partial_\tau f$	$\frac{\partial f}{\partial \tau}$

Table 7.1 Notation used in this chapter.

of fluctuations in the CMB, respectively, under the ‘just enough inflation’ model. [Section 7.5](#) summarises the findings of this chapter.

The notation used throughout this chapter is summarised in [Table 7.1](#). Note that besides using natural units as explained in [Notation & conventions](#), we set  $8\pi G = 1$ .

## 7.2 Background

### 7.2.1 Canonical transformations

To motivate the choice of transformations considered in this work, we briefly review ‘classical’ canonical transformations and a broader class of canonical transformations performed on an extended phase space. For a thorough review of canonical transformations in cosmology, see [\[180\]](#).

In the Lagrangian formulation of classical mechanics, all information about the dynamics of the system is carried by the Lagrangian  $L(q, \dot{q}, t)$ , a function of the  $n$  generalised coordinates  $q_i$  and their derivatives. The Euler–Lagrange equations yield the equations of motion even if one performs a reversible change of coordinates (or point transformation)

$$q_i \rightarrow Q_i(q, t). \quad (7.1)$$

The Hamiltonian formulation of mechanics puts the generalised coordinates  $q_i$  and their derivatives  $\dot{q}_i$  on an equal footing by considering generalised momenta

$$p_i = \frac{\partial L}{\partial \dot{q}_i} \quad (7.2)$$

and working in the  $2n$ -dimensional phase space  $(q_i, p_i)$  instead of configuration space  $(q_i, \dot{q}_i)$ . In the Hamiltonian formalism there exists a broader range of transformations than [Eq. \(7.1\)](#) that leave the Hamiltonian equations of motion invariant, called *canonical transformations*. Canonical transformations allow one to mix the generalised coordinates and momenta and to do so in a time-dependent manner

$$q_i \rightarrow Q_i(q_i, p_i, t), \quad p_i \rightarrow P_i(q_i, p_i, t), \quad (7.3)$$

whilst preserving the form of Hamilton’s equations of motion,

$$\begin{aligned} \dot{p}_i &= -\frac{\partial H}{\partial q_i}, & \dot{q}_i &= \frac{\partial H}{\partial p_i}, \\ \dot{P}_i &= -\frac{\partial K}{\partial Q_i}, & \dot{Q}_i &= \frac{\partial K}{\partial P_i}, \end{aligned} \quad (7.4)$$

where  $K$  is the new Hamiltonian. For the equations of motion to be conserved, the principle of extremal action has to be satisfied in both the initial and transformed system:

$$\delta \int_{t_1}^{t_2} (p_i \dot{q}_i - H(q, p, t)) dt = 0, \quad (7.5)$$

$$\delta \int_{t_1}^{t_2} (P_i \dot{Q}_i - K(Q, P, t)) dt = 0. \quad (7.6)$$

For Eq. (7.5) to imply Eq. (7.6), we require

$$\lambda(p_i \dot{q}_i - H) = P_i \dot{Q}_i - K + \dot{F}, \quad (7.7)$$

with  $\lambda = 1$ <sup>1</sup>. The Lagrangian of the system can at most change by a total derivative  $dF$ , with  $F$  being called the generating function of the transformation. All canonical transformations possess a generating function [185, 186]. The Hamiltonian then transforms as

$$K = H + \frac{\partial F}{\partial t}. \quad (7.8)$$

It can be shown that the Poisson bracket,

$$\{f, g\} = \sum_{i=1}^N \frac{\partial f}{\partial q_i} \frac{\partial g}{\partial p_i} - \frac{\partial f}{\partial p_i} \frac{\partial g}{\partial q_i}, \quad (7.9)$$

is invariant under canonical transformations, and that the converse is true: if the Poisson bracket structure is conserved such that

$$\{Q_i, Q_j\} = \{P_i, P_j\} = 0, \quad \{Q_i, P_j\} = \delta_{ij}, \quad (7.10)$$

then the transformation is canonical.

The (quantum) field theoretical equivalent of Eq. (7.7), adding a total differential to the Lagrangian, is the freedom to add a 4-divergence to the Lagrangian density,

$$\mathcal{L}' = \mathcal{L}(t, \phi, \partial_\mu \phi) + \nabla_\mu J^\mu(\phi, \partial_\mu \phi), \quad (7.11)$$

which, if the independent variable stays the same throughout, is equivalent to integrating the associated action by parts and discarding the surface term. One of the transformations considered in this work is the addition of such vanishing surface terms, which thus covers all canonical transformations. Commutators are the quantum counterparts of Poisson brackets, and are thus preserved under canonical transformations. In all transformations described in

<sup>1</sup>Transformations with  $\lambda \neq 1$  are called extended canonical transformations, and will not be considered in this work.

this work, it is ensured that the commutator structure of the destination system matches that of the original.

In the conventional formulation of the principle of least action,  $t$  is the Newtonian absolute time, and has a distinguished role. This is not always desirable and for relativistic considerations, one may wish to treat the generalised coordinates  $q_i$  and  $t$  on an equal footing. This can be achieved by extending the phase space considered with time  $t$  and its conjugate momentum, the negative Hamiltonian:  $(q_i, p_i, t, -\mathcal{H})$  [187–189]. In this extended space it is possible to define an extended set of Hamilton’s equations, and define canonical transformations that involve a redefinition of time,

$$(t, q_i, p_i, -\mathcal{H}) \rightarrow (T, Q_i, P_i, -\mathcal{K}). \quad (7.12)$$

The other type of transformations this work considers, simultaneous time and field redefinitions, are canonical transformations on the extended phase space. While in itself it does not include all canonical transformations on the extended space, it is motivated by practical considerations: one may wish to derive initial conditions ‘from first principles’ for a different field, related to the original one by a time-dependent rescaling, for its better stability properties in numerical simulations.

## 7.2.2 Dynamics of primordial perturbations

### The perturbed classical action

The dynamics of primordial curvature perturbations arise from perturbing the metric and matter fields around a homogeneous, isotropic, expanding background (for a thorough review, see e.g. [89, 90, 190]). In an inflationary model with a single scalar field  $\phi(t)$  and potential  $V(\phi)$  on a Friedmann–Robertson–Walker (FRW) spacetime, the system can be described via the action

$$S = \frac{1}{2} \int d^4x \sqrt{-g} [R - (\partial_\mu \phi)(\partial^\mu \phi) - 2V(\phi)], \quad (7.13)$$

where the metric  $g_{\mu\nu}$  is the spatially flat FRW metric. Perturbing this action in the comoving gauge

$$\delta\phi = 0, \quad g_{ij} = a^2[(1 - 2\mathcal{R})\delta_{ij} + h_{ij}], \quad \partial_i h_{ij} = h_i^i = 0, \quad (7.14)$$

yields, to second order in the gauge-invariant curvature perturbation  $\mathcal{R}$ ,

$$S_2 = \frac{1}{2} \int d^4x a z^2 [\dot{\mathcal{R}}^2 - a^{-2}(\partial_i \mathcal{R})^2]. \quad (7.15)$$

Traditionally this is then written in terms of the Mukhanov variable  $v = z\mathcal{R} = \frac{a\dot{\phi}}{H}\mathcal{R}$  and conformal time  $\eta$  and integrated by parts to give

$$S_2 = \frac{1}{2} \int d\eta d^3x [(v')^2 - (\partial_i v)^2 + \frac{z''}{z} v^2]. \quad (7.16)$$

The reason behind the variable choice and integration by parts is that in the resulting action  $v$  is canonically normalised, i.e. it has a kinetic term of the form  $\frac{1}{2}\partial_\mu v \partial^\mu v$ , and thus yields an equation of motion of a particular form. Writing  $v$  as a Fourier decomposition,

$$v(\eta, x) = \int \frac{d^3k}{(2\pi)^3} v_{\mathbf{k}}(\eta) e^{i\mathbf{k}\cdot\mathbf{x}}, \quad (7.17)$$

the action in Eq. (7.16) gives the equation of motion of an oscillator:

$$v_k'' + \left( k^2 - \frac{z''}{z} \right) v_k = 0, \quad (7.18)$$

where the vector notation on the wavevector  $k$  has been removed due to the isotropy of the field  $v$ . The variables  $(\eta, v)$  were thus chosen because they yield an oscillator's equation of motion (without a 'first-derivative term' proportional to  $v'$ ), allowing the classical field  $v$  to be quantised by analogy with a time-dependent quantum harmonic oscillator.

### Scalar fields in curved spacetime

In order to obtain the action in Eq. (7.15) describing the dynamics of the gauge-invariant curvature perturbation, action Eq. (7.13) was perturbed in both the metric and the inflaton field, but a gauge was then chosen to keep only the metric perturbations and leave the inflaton unperturbed. One can arrive at analogous dynamics (a similar equation of motion to Eq. (7.18)) by instead modelling the field perturbations as a new massless scalar field  $\varphi$  on a spacetime with an unperturbed metric. Starting from the action describing the dynamics of the field perturbations  $\varphi$ ,

$$S = \frac{1}{2} \int d^4x \sqrt{-g} (g^{\mu\nu} \partial_\mu \varphi \partial_\nu \varphi - m^2 \varphi^2), \quad (7.19)$$

which one traditionally considers in terms of the auxiliary field  $y = a\varphi$  and conformal time (for similar reasons as  $v$  was considered in the previous section), one derives the Fourier space equation of motion

$$y_k'' + \left( k^2 - \frac{a''}{a} + a^2 m^2 \right) y_k = 0. \quad (7.20)$$

The similarity between Eq. (7.18) and (Eq. (7.20)) means that one can *identically* map the modes  $y_k$  derived from a massless scalar field  $\varphi$  onto the perturbation modes  $v_k$ , provided one ensures that the two associated background spacetimes satisfy

$$\frac{a''}{a} = \frac{z''}{z}. \quad (7.21)$$

One particular solution of Eq. (7.21) is the case  $z \propto a$ , which holds true or becomes a good approximation in a number of cases. It is true exactly in power-law inflation models [191], and becomes a good approximation in slow-roll inflation [192], which is power-law inflation to first

order and is admitted by many models. The proportionality is also a good approximation in kinetic dominance [115], described in more detail in Section 7.4.2. Therefore, in these cases, one may treat the Mukhanov variable in an inflating universe as if it were a scalar field  $\varphi(\eta, x)/a(\eta)$  on the same background, obeying the action Eq. (7.19). It should be emphasised that one may always map the results for the test scalar field onto the primordial perturbations, but in general the mapping will not be the identity and so the equations of motion will not be form-identical. The advantage of working with action Eq. (7.19) over (Eq. (7.16)) is that the former is in covariant form (once each  $\partial_\mu$  has been replaced by  $\nabla_\mu$ ), and hence can be used to derive further covariant quantities such as the stress–energy tensor, the minimisation of which provides a definition of the ground state (see Section 7.2.3).

### 7.2.3 Vacuum choices

In the above models of primordial perturbations there are multiple ways to define a vacuum or ground state. These result in expressions for the Fourier modes of the perturbations considered ( $v_k$  or  $y_k$ ), which can then be used as initial conditions for the perturbation modes, and affect the form of the primordial power spectrum the modes admit. This section introduces three different definitions of the vacuum state: *Hamiltonian diagonalisation*, the *Danielsson vacuum*, and *minimising the renormalised stress–energy tensor*. We derive the initial conditions each vacuum definition gives for the mode functions if applied in the conventional way. We then examine how the initial conditions change under a field redefinition in the associated action and the addition of surface terms in Section 7.4.1.

Acknowledging that the literature on inflationary initial conditions is vast, Section 7.4.1 further considers the Bunch-Davies vacuum, the adiabatic vacuum and  $\alpha$ -vacua. The behaviour of some of these vacua under canonical transformations has been investigated in the past, but for others this is either impossible (due to them relying on a special choice of canonical variable pair) or remains to be carried out.

#### Hamiltonian diagonalisation

To obtain a quantum theory from the semiclassical action Eq. (7.16), the field  $v$  is promoted to an operator:

$$\hat{v}(\eta, x) = \int \frac{d^3k}{(2\pi)^3} \left[ \hat{a}_k v_k(\eta) e^{ik \cdot x} + \hat{a}_k^\dagger v_k^*(\eta) e^{-ik \cdot x} \right], \quad (7.22)$$

where  $\hat{a}_k^\dagger$  and  $\hat{a}_k$  are the creation and annihilation operators, respectively. The momentum conjugate to  $v$  is

$$\pi = \frac{\partial \mathcal{L}}{\partial v'} = v', \quad (7.23)$$

(with  $\mathcal{L}$  being the Lagrangian density associated with the action) and is quantised accordingly. To impose canonical commutation relations

$$[\hat{a}_k, \hat{a}_{k'}^\dagger] = (2\pi)^3 \delta(k - k'), \quad (7.24)$$

and to obey quantum dynamics, the time-dependent part of the mode functions in Eq. (7.22) must satisfy the equation of motion and a normalisation constraint:

$$v_k'' + \left(k^2 - \frac{z''}{z}\right) v_k = 0, \quad (7.25)$$

$$v_k' v_k^* - v_k^{*'} v_k = -i. \quad (7.26)$$

Eqs. (7.25) and (7.26) do not fully determine the mode functions  $v_k(\eta)$ . To fix the leftover degree of freedom, one needs to specify a vacuum, which obeys  $\hat{a}_k |0\rangle = 0$ . One popular definition of the ground state is that which minimises the vacuum expectation value of the Hamiltonian of the system. The choice of variables  $(\eta, v)$  results in the expression

$$\langle 0 | H | 0 \rangle \propto \int \frac{d^3k}{(2\pi)^3} \left[ |v_k'|^2 + \left(k^2 - \frac{z''}{z}\right)^2 |v_k|^2 \right]. \quad (7.27)$$

We must then minimise the contribution to the Hamiltonian separately for each  $k$ -mode, with respect to the mode functions belonging to that mode, subject to the constraint Eq. (7.26). This leads to the solutions

$$\begin{aligned} |v_k|^2 &= \frac{1}{2\sqrt{k^2 - \frac{z''}{z}}}, \\ v_k' &= -i\sqrt{k^2 - \frac{z''}{z}} v_k, \end{aligned} \quad (7.28)$$

which can be used to set initial conditions on the mode functions  $v_k$ .

### Danielsson vacuum

Danielsson [173] proposed that the vacuum should instead be chosen such that

$$\begin{aligned} |v_k|^2 &= (2k)^{-1}, \\ v_k' &= \left(-ik + \frac{a'}{a}\right) v_k. \end{aligned} \quad (7.29)$$

This result is derived from ‘first principles’, working in the Heisenberg picture. For detailed reviews of the following, see [193–195].

In the Heisenberg picture, the operators carry time-dependence. When quantising the massless field  $y$  appearing in the action [Eq. \(7.19\)](#) with  $m = 0$ , we may write in Fourier space

$$\hat{y}_k(\eta) = \frac{1}{\sqrt{2k}} \left[ \hat{a}_k(\eta) + \hat{a}_{-k}^\dagger(\eta) \right], \quad (7.30)$$

and for its conjugate momentum  $\pi_k = y'_k - \frac{a'}{a} y_k$ ,

$$\hat{\pi}_k(\eta) = -i\sqrt{\frac{k}{2}} \left[ \hat{a}_k(\eta) - \hat{a}_{-k}^\dagger(\eta) \right]. \quad (7.31)$$

The creation and annihilation operators mix over time via a Bogoliubov transformation

$$\begin{aligned} \hat{a}_k(\eta) &= \alpha_k(\eta)\hat{a}_k(\eta_0) + \beta_k(\eta)\hat{a}_{-k}^\dagger(\eta_0), \\ \hat{a}_{-k}^\dagger(\eta) &= \beta_k^*(\eta)\hat{a}_k(\eta_0) + \alpha_k^*(\eta)\hat{a}_{-k}^\dagger(\eta_0), \end{aligned} \quad (7.32)$$

where  $\alpha_k(\eta)$  and  $\beta_k(\eta)$  are time-dependent mixing coefficients. One can therefore isolate the time-dependent parts of the fields,

$$\begin{aligned} \hat{y}_k(\eta) &= f_k(\eta)\hat{a}_k(\eta_0) + f_k^*(\eta)\hat{a}_{-k}^\dagger(\eta_0), \\ i\hat{\pi}_k(\eta) &= g_k(\eta)\hat{a}_k(\eta_0) - g_k^*(\eta)\hat{a}_{-k}^\dagger(\eta_0), \end{aligned} \quad (7.33)$$

with

$$\begin{aligned} f_k(\eta) &= \frac{1}{\sqrt{2k}}(\alpha_k(\eta) + \beta_k^*(\eta)), \\ g_k(\eta) &= \sqrt{\frac{k}{2}}(\alpha_k(\eta) - \beta_k^*(\eta)), \end{aligned} \quad (7.34)$$

where the  $f_k$  now take the role of the mode functions  $y_k$  in the Schrödinger picture, because they carry all time dependence of the field operator. At  $\eta = \eta_0$ , the creation and annihilation operators are by definition unmixed, therefore

$$\beta_k(\eta_0) = 0 = \sqrt{\frac{k}{2}}f_k^*(\eta_0) - \frac{1}{\sqrt{2k}}g_k^*(\eta_0). \quad (7.35)$$

Identifying the mode function  $y_k$  with  $f_k$  and the conjugate momentum  $\pi_k$  with  $-ig_k$ , we obtain the Danielsson vacuum [Eq. \(7.29\)](#).

Immediately it is clear that since the Danielsson vacuum relates a field and its conjugate momentum, the initial conditions derived from it will generally change under transformations that change that relationship, e.g. the addition of a surface term to the action, described in [Section 7.3.2](#).

### Minimising the renormalised stress–energy tensor

[12] have proposed to determine the ground state by minimising the vacuum expectation value of the local energy density to avoid the excessive particle production of the Hamiltonian diagonalisation approach. The local energy density is computed as the 00 component of the stress–energy tensor of the system Eq. (7.19). General relativity requires a symmetric stress–energy tensor, as it appears on the right-hand side of the Einstein equations. It is defined as [86]

$$T_{\mu\nu} = -\frac{2}{\sqrt{-g}} \frac{\delta \mathcal{S}_m}{\delta g^{\mu\nu}}, \quad (7.36)$$

with the subscript m signalling the matter part of the action. Just like the expectation value of the Hamiltonian, the stress–energy tensor is divergent, and has to be renormalised to yield finite quantities. There exist several procedures of renormalisation of the stress–energy tensor, and in [12] the Hadamard point-splitting method is used (thoroughly described in [86]). In summary, this consists of first quantising the field  $y$  from Section 7.2.2, writing down an expression for the Hadamard Green function,

$$G^{(1)}(x, x') = \frac{1}{2} \langle 0 | \{ \varphi(x), \varphi(x') \} | 0 \rangle \quad (7.37)$$

then applying the bi-scalar derivative  $\mathcal{D}_{\mu\nu}$  to

$$G^{(1)}(x, x') - G_{\text{DS}}^{(1)}(x, x'), \quad (7.38)$$

where the second term denotes the de-Witt Schwinger geometrical terms. The geometrical terms do not depend on the variables with respect to which one minimises the stress–energy tensor, and are therefore summarised as  $\tilde{T}$  and ignored in the minimisation process. The coincidence limit  $x \rightarrow x'$  is then taken to yield  $\langle 0 | T_{\mu\nu} | 0 \rangle_{\text{ren}}$ . Altogether,

$$\langle 0 | T_{\mu\nu} | 0 \rangle_{\text{ren}} = \lim_{x \rightarrow x'} \frac{1}{2} \left[ (\nabla_\mu \nabla_{\nu'} + \nabla_{\mu'} \nabla_\nu) - g_{\mu\nu} \nabla_\alpha \nabla^{\alpha'} + g_{\mu\nu} m^2 \right] G^{(1)}(x, x') + \tilde{T}, \quad (7.39)$$

which is a functional to be minimised with respect to the mode functions  $\{y_k, y_k^*, y'_k, y_k^{*'}\}$  treated as independent variables, subject to the normalisation constraint on  $y_k$  arising from the canonical commutation relations. The minimisation process yields the solutions

$$\begin{aligned} |y_k|^2 &= \frac{1}{2\sqrt{(k^2 + m^2 a^2)}}, \\ y'_k &= \left( -i\sqrt{k^2 + m^2 a^2} + \frac{a'}{a} \right) y_k, \end{aligned} \quad (7.40)$$

which, under the circumstances described in [Section 7.2.2](#), can be used to set

$$\begin{aligned} |v_k|^2 &= (2k)^{-1}, \\ v'_k &= \left(-ik + \frac{z'}{z}\right) v_k \end{aligned} \tag{7.41}$$

for the Mukhanov variable.

### 7.3 Methods

In the theories of primordial perturbations reviewed in the previous section, the standard choices of variables was explained by the simple form of equation of motion they admitted, more specifically that the modes involved behaved like harmonic oscillators, so their quantisation was known.

However, as will be shown, one can always make a canonical transformation of the action of the system in the extended phase space by redefining the independent and dependent variables simultaneously such that the resulting equation of motion is first-derivative-free and appropriate commutation relations are satisfied. The new, redefined field will then still be quantisable by analogy with the harmonic oscillator. Apart from the conveniently ‘bare’  $k^2$  term in the associated equation of motion (originating from the canonically normalised scalar field in the action), there is nothing that makes the standard choice of variables  $(\eta, v)$  special. In fact there is no choice of variables for which the action is canonically normalised in the case of a non-flat universe [\[157\]](#). In addition to field redefinitions, one could always add a surface term to the Lagrangian, equivalent to performing an integration by parts, and obtain a dynamically equivalent system (so long as the boundary terms vanish). This section summarises the procedures described above by which the vacuum prescriptions will be transformed.

#### 7.3.1 Field redefinition

For all vacuum-setting methods involving the quantisation of a field, we consider quantising an alternative field related to the conventional choice by a time-dependent, scalar-valued, homogeneous function  $h$ , and the redefinition of time to a new independent variable  $\tau$ , whilst keeping in mind how the form of the metric changes. We then derive the vacuum conditions in analogy with the conventional procedures. For Hamiltonian diagonalisation this redefinition will thus be

$$t \rightarrow \tau(t), \quad \mathcal{R} \rightarrow \chi(x, \tau) \equiv \frac{\mathcal{R}}{h(\tau)}, \tag{7.42}$$

and for minimising the renormalised stress–energy tensor, we shall consider instead

$$t \rightarrow \tau(t), \quad \varphi \rightarrow \chi(x, \tau) \equiv \frac{\varphi}{h} = \frac{y}{ah}, \tag{7.43}$$

as the field being quantised is  $\varphi$ .

We shall derive a constraint linking  $\tau$  and  $h$  for each action considered to ensure it yields the equation of motion of an oscillator, as this is not guaranteed by the transformation Eq. (7.42). This generally leaves an unconstrained integration constant  $C_0$ . The field redefinition also does not necessarily conserve the commutator structure. For all transformations considered, this will be ensured via an additional constraint that needs to be satisfied during the minimisation of the vacuum expectation value of the Hamiltonian density or the 00-component of the renormalised stress–energy tensor.

### 7.3.2 Surface terms

In the conventional Hamiltonian diagonalisation approach, one performs an integration by parts to obtain the action Eq. (7.16). This is equivalent to adding the total derivative

$$\left(\frac{z'}{z}v^2\right)' \quad (7.44)$$

to the associated Lagrangian density, or a vanishing *surface or boundary term* to the action. Under the field redefinition of the previous section, the action changes such that the appropriate boundary term to add will be

$$-\partial_\tau \left(\chi^2 \frac{\partial_\tau h}{h}\right). \quad (7.45)$$

It is clear that (i) there are infinitely many choices of vanishing boundary terms one could add; and that (ii) the boundary terms modify the form of the action and in turn the field-conjugate momentum relationship. For Hamiltonian diagonalisation and the Danielsson vacuum we shall investigate how the addition of the boundary term Eq. (7.45) alters the initial conditions, and show that minimising the renormalised stress–energy tensor is invariant under the addition of boundary terms by construction.

Detailed calculations of initial conditions arising from systems subject to field redefinitions and addition of surface terms can be found in Sections 7.A and 7.B.

## 7.4 Results

### 7.4.1 Initial conditions

#### Hamiltonian diagonalisation

Under the field redefinition Eq. (7.42), minimising the vacuum expectation value of the Hamiltonian Eq. (7.27) gives the generalised initial conditions

$$\begin{aligned} |\chi_k|^2 &= (2C_0\omega_k)^{-1}, \\ \partial_\tau \chi_k &= -i\omega_k \chi_k, \end{aligned} \quad (7.46)$$

where  $\omega_k$  is the time-dependent frequency of the equation of motion in terms of  $(\tau, \chi_k)$ , and  $C_0$  is a constant.

As a sanity check, substituting  $h = z^{-1}$  and  $C_0 = 1$ , corresponding to quantising the Mukhanov variable in terms of conformal time, we recover the conventional initial conditions [Eq. \(7.28\)](#). However, changing variables to  $(\eta, v)$  in the generalised initial conditions [Eq. \(7.46\)](#) yields

$$\begin{aligned} |v_k|^2 &= \frac{1}{2\sqrt{k^2 + \frac{h''}{h} + 2\left(\frac{h'z'}{hz}\right)}}, \\ v'_k &= \left( -i\sqrt{k^2 + \frac{h''}{h} + 2\frac{h'z'}{hz}} + \frac{h'}{h} + \frac{z'}{z} \right) v_k, \end{aligned} \quad (7.47)$$

which carry the arbitrary function  $h$ . This means that one can derive a family of initial conditions for  $v_k(\eta)$  depending on the field in terms of which the action was written.

Moreover, under addition of the boundary term defined in [Section 7.3.2](#) and the field re-definition that led to [Eq. \(7.47\)](#), one can derive another set of conditions:

$$\begin{aligned} |v_k|^2 &= \frac{1}{2\sqrt{k^2 - \left(\frac{h'}{h}\right)^2}}, \\ v'_k &= \left( -i\sqrt{k^2 - \left(\frac{h'}{h}\right)^2} + \frac{h'}{h} + \frac{z'}{z} \right) v_k. \end{aligned} \quad (7.48)$$

Not only is this another family of solutions depending on the function  $h$ , it is a different set to [Eq. \(7.47\)](#)! This can be seen by noting that for all values of  $h'/h$  in [Eq. \(7.48\)](#) there will be a lower limit for  $k$  below which the expression under the square root becomes negative, which is not allowed (for the squared magnitude of the perturbation has to be real). This ‘forbidden region’ is different for [Eq. \(7.47\)](#). Since we could have chosen any action connected to [Eq. \(7.16\)](#) via a canonical transformation, we could have arrived at a range of different Hamiltonians yielding different quantum initial conditions for the perturbations. Hamiltonian diagonalisation thus gives ambiguous initial conditions depending on the choice of canonical variables.

The two example solution families obtained via Hamiltonian diagonalisation are parametrised by a time-dependent function  $h$ . If we choose to use the solutions as initial conditions, and the spacetime slice we set them on is chosen such that all modes (i.e. all  $k$ ) are set simultaneously at a conformal time  $\eta_0$ , the shape of the function  $h(\eta)$  only matters in the vicinity of  $\eta_0$ . We can determine how much freedom this gives by Taylor-expanding  $h$  near  $\eta_0$ ,

$$h(\eta) = \sum_{n=0}^{\infty} \frac{1}{n!} h_n (\eta - \eta_0)^n, \quad (7.49)$$

substituting into Eq. (7.47), then finally evaluating at  $\eta = \eta_0$ . This gives

$$|v_k|^2 = \frac{1}{2\sqrt{k^2 + \frac{h_2}{h_0} + 2\left(\frac{h_1 z'}{h_0 z}\right)}},$$

$$v'_k = \left( -i\sqrt{k^2 + \frac{h_2}{h_0} + 2\frac{h_1 z'}{h_0 z} + \frac{h_1}{h_0} + \frac{z'}{z}} \right) v_k, \quad (7.50)$$

showing that there are two real degrees of freedom in this set of initial conditions ( $h_2/h_0$  and  $h_1/h_0$ ), whereas Eq. (7.48) only depends on  $h_1/h_0$ :

$$|v_k|^2 = \frac{1}{2\sqrt{k^2 - 2\left(\frac{h_1}{h_0}\right)^2}},$$

$$v'_k = \left( -i\sqrt{k^2 - 2\left(\frac{h_1}{h_0}\right)^2 + \frac{h_1}{h_0} + \frac{z'}{z}} \right) v_k. \quad (7.51)$$

It is worth noting that there exist prescriptions that do not set initial conditions for all perturbations simultaneously: [173] for example chooses a finite  $\eta_0(k)$  for each  $k$  such that the physical momentum corresponding to the mode is given by some fixed scale, and it is common to initialise the perturbations when their lengthscales are a fixed fraction of the Hubble horizon. These prescriptions may lead to more degrees of freedom in the initial conditions, but such considerations are deferred to future work.

In [196], the authors do not attempt to find the preferred choice of vacuum (if that exists), rather they parametrise a general choice of initial conditions by two complex scalars,  $X$  and  $Y$  ( $Y$  being purely imaginary), that characterise the choice of vacuum:

$$v_k(\eta_0) = \frac{e^{i\phi_1}}{\sqrt{k}} \left[ 1 + \frac{X+Y}{2}\theta_0 + \mathcal{O}(\theta_0^2) \right], \quad (7.52)$$

$$v'_k(\eta_0) = -i\sqrt{k}e^{i\phi_2} \left[ 1 + \frac{Y-X}{2}\theta_0 + \mathcal{O}(\theta_0^2) \right]. \quad (7.53)$$

Their initial conditions are the solutions of the mode equation Eq. (7.18) in the limit where the metric is Minkowski, with added corrections due to expansion as a power series in the small, dimensionless parameter  $\theta_0$ :

$$\theta_0 = \left. \frac{aH}{k} \right|_{\eta=\eta_0}. \quad (7.54)$$

In terms of this parametrisation, the solutions Eq. (7.50) represent a subset with

$$X = -Y, \quad Y \text{ arbitrary}, \quad (7.55)$$

which can be seen by factoring out  $1/\sqrt{k}$  from the expression for  $v_k$ , expanding the rest in powers of  $1/k$ , and observing that there is no term linear in  $1/k$ .

### Danielsson vacuum

It is easily seen that Danielsson's result, for a different choice of field Eq. (7.43), generalises as

$$\pi_\chi(\tau) = -ik\chi(\tau), \quad (7.56)$$

where  $\pi_\chi$  is the momentum conjugate to  $\chi$ . Using the action Eq. (7.19) under the field redefinition (Eq. (7.43)) to compute the conjugate momentum, this gives for  $y_k$ :

$$\begin{aligned} |y_k|^2 &= \frac{(ah)^2}{2k}, \\ y'_k &= \left( -\frac{ik}{(ah)^2} + \frac{a'}{a} \right) y_k. \end{aligned} \quad (7.57)$$

Under the addition of the surface term Eq. (7.45) to the action (which has the effect of eliminating the  $\partial_\tau\chi$  term), the momentum conjugate to the field  $\chi$  changes, and the initial conditions become

$$\begin{aligned} |y_k|^2 &= \frac{(ah)^2}{2k}, \\ y'_k &= \left( -\frac{ik}{(ah)^2} + \frac{a'}{a} + \frac{h'}{h} \right) y_k. \end{aligned} \quad (7.58)$$

Since canonical transformations by addition of surface terms only change the relationship between the field and its conjugate momentum, infinitely many solution families like Eq. (7.57) and Eq. (7.58) exist that differ in  $y'_k(y_k)$ . There is no theoretical guidance as to which action one should consider, and hence which set of initial conditions is the correct one in the Danielsson prescription.

### Minimising the renormalised stress–energy tensor

Working from the action Eq. (7.19) under a canonical field redefinition, and minimising the expression Eq. (7.39) fixes the mode functions  $\chi_k$  fully as

$$\begin{aligned} |\chi_k|^2 &= \frac{1}{2\sqrt{h^4 a^4 (k^2 + m^2 a^2)}}, \\ \partial_\tau \chi_k &= \left( -\frac{i}{C_0} \sqrt{h^4 a^4 (k^2 + m^2 a^2)} - \frac{\partial_\tau h}{h} \right) \chi_k. \end{aligned} \quad (7.59)$$

Converting back to conformal time and  $y$  gives the solutions

$$\begin{aligned} |y_k|^2 &= \frac{1}{2\sqrt{(k^2 + m^2 a^2)}}, \\ y'_k &= \left( -i\sqrt{k^2 + m^2 a^2} + \frac{a'}{a} \right) y_k, \end{aligned} \quad (7.60)$$

independently of the choice of  $h(\tau)$  and  $\tau$ , and in agreement with solutions Eq. (7.40) derived using the conventional treatment.

From the definition Eq. (7.36), it can be shown that an added 4-divergence (involving the field and the metric) to the Lagrangian does not change the form of the stress–energy tensor as long as the added surface term does not break covariance of the existing Lagrangian. Not only has the stress–energy tensor been constructed in a covariant way, [197] have recently shown that one can define a ground state family, with respect to which its expectation value is covariantly conserved. This possibility was also pointed out in [171]. All that matters in the minimisation process is the form of the field-dependent part of  $T_{\mu\nu}$  (the geometrical terms do not come into play as they do not depend on the field), and therefore the initial conditions derived via minimisation of the renormalised stress–energy tensor are invariant under the addition of surface terms.

Minimising the 00-component of the renormalised stress–energy tensor is therefore invariant under canonical field redefinitions and addition of surface terms. The stress–energy tensor was defined to be covariant, therefore its invariance under the addition of surface terms comes as no surprise, but it has not been constructed to be invariant under field redefinitions. In taking the 00-component, however, covariance has been broken, and one must ensure that the same results hold when a coordinate-independent quantity is minimised.

To make sure this is the case, one should consider the eigenvalue belonging to a timelike eigenvector of the tensor. Let us therefore generate the off-diagonal components of  $\langle 0|T_{\mu\nu}|0\rangle$ , ignoring those arising from the geometrical terms:

$$\langle 0|T_{i0}|0\rangle_{\text{ren}} = \frac{1}{2} \int \frac{d^3k}{(2\pi)^3} \frac{k_i}{\dot{\tau} a z^2}, \quad (7.61)$$

$$\langle 0|T_{ij}|0\rangle_{\text{ren}} = \frac{1}{2} \int \frac{d^3k}{(2\pi)^3} 2h^2 \chi_k \chi_k^* k_i k_j, \quad (7.62)$$

with  $i \neq j$ . All the above components include integrals with integrands odd in  $k_i$ , and so they vanish. The de Witt–Schwinger terms in the off-diagonal elements can also be shown to vanish due to the metric  $g_{\mu\nu}$  being diagonal<sup>2</sup>. Consequently the expectation value of the renormalised stress–energy tensor is diagonal, and its eigenvalues are just the diagonal entries. The eigenvalue corresponding to the timelike eigenvector is the 00 element, confirming the results of this section.

### Generalised analogy between a scalar field on a curved spacetime and the curvature perturbation

In Section 7.2.3, initial conditions were derived for the Mukhanov variable in an inflating background from the solutions for  $y_k$  in Eq. (7.40) by drawing an analogy between their

<sup>2</sup>For an explicit expression for  $\langle 0|T_{\mu\nu}|0\rangle_{\text{ren}}$  that confirms this, see [86].

equations of motion, Eq. (7.18)) and (Eq. (7.20)). For the two equations to be the same,  $a''/a = z''/z$  was required. We now derive similar requirements for the two equations of motion to have the same form, after having performed the field redefinitions

$$\mathcal{R} \rightarrow \chi_1 = \frac{\mathcal{R}}{h_1} \quad (7.63)$$

and

$$\varphi \rightarrow \chi_2 = \frac{\varphi}{h_2} = \frac{y}{ah_2}, \quad (7.64)$$

respectively. The first, obvious requirement is that the redefined field in Eq. (7.20) has to be massless,  $m = 0$ . For the classical perturbed action we then have the equation of motion

$$0 = \partial_{\tau_1\tau_1}\chi_1 + \left[ \left( \frac{kh_1^2z^2}{C_1^2} \right)^2 + \frac{\partial_{\tau_1\tau_1}h_1}{h_1} - 2 \left( \frac{\partial_{\tau_1}h_1}{h_1} \right)^2 \right] \chi_1, \quad (7.65)$$

whereas from the massless scalar field we get

$$0 = \partial_{\tau_2\tau_2}\chi_2 + \left[ \left( \frac{kh_2^2a^2}{C_2^2} \right)^2 + \frac{\partial_{\tau_2\tau_2}h_2}{h_2} - 2 \left( \frac{\partial_{\tau_2}h_2}{h_2} \right)^2 \right] \chi_2. \quad (7.66)$$

We are free to choose  $h_1$ ,  $h_2$ , and the constants  $C_1$  and  $C_2$ . We first have to set

$$\frac{h_1z}{C_1} = \frac{h_2a}{C_2} \quad (7.67)$$

for the first term in the frequency to be the same, then we need to satisfy

$$\frac{h_1''}{h_1} - 2 \left( \frac{h_1'}{h_1} \right)^2 = \frac{h_2''}{h_2} - 2 \left( \frac{h_2'}{h_2} \right)^2, \quad (7.68)$$

where prime denotes differentiation with respect to the function's argument. This differential equation admits the solution

$$\frac{h_1}{h_2} = A \exp \left( \int dt \frac{h_2^2}{B - \int dt h_2^2} \right). \quad (7.69)$$

One can always choose  $h_1/h_2$  such that Eq. (7.67) holds. One can then calculate  $h_2$  by inverting Eq. (7.69). This means that if  $h_1$ ,  $h_2$  are chosen appropriately, one can always map the solution of the scalar-field action onto the cosmological perturbations.

For an arbitrary choice of  $h_1$  and  $h_2$ , if  $z \propto a$ , the equations of motion are guaranteed to take the same form in general, therefore the mapping from the test scalar field to the primordial perturbations becomes the identity. (Note that for a specific choice of  $h_1$  and  $h_2$ , e.g.  $h_1 = 1/z$ ,  $h_2 = 1$ , the requirement for form-similarity may be less restrictive.)

Overall, the conclusion that one may use the solutions [Eq. \(7.59\)](#) as initial conditions for the associated perturbation mode in inflation is still valid for a different choice of (extended) canonical variables.

### Other vacuum choices

**Bunch–Davies vacuum:** A well-known vacuum prescription left out of our analysis is the Bunch–Davies vacuum [\[172\]](#), which is set by requiring that the positive-frequency mode function matches the Minkowski solution in the limit of all mode lengthscales being subhorizon ( $k/(aH) \rightarrow \infty$ ), in the infinite conformal past. The prescription only specifies the vacuum uniquely if inflation did not have a start, and otherwise results in ambiguities associated with the artificially imposed  $k \rightarrow \infty$  limit [\[198\]](#). This issue aside, the Bunch–Davies prescription relies on the choice of canonical variables being  $(\eta, v_k)$ , because it yields an oscillator equation of motion with a ‘bare’  $k^2$  term in the squared frequency, thus the frequency can be adiabatically constant. This is no longer the case when the equation of motion is rewritten in terms of a different time and field, i.e. under field redefinitions. This preference for variable choice is the reason why the Bunch–Davies vacuum was not considered.

**Adiabatic vacuum:** The adiabatic vacuum formalism [\[86, 198\]](#) is similar to the Bunch–Davies one in that it states that sometimes a unique vacuum can be defined when the frequency of the perturbation modes is adiabatically constant, and given by an  $n^{\text{th}}$  order Wentzel–Kramers–Brillouin (WKB) approximation. Since canonical transformations leave the equation of motion invariant, they will not affect the adiabatic vacuum if the expansion is carried out to arbitrarily high order, but otherwise the (truncated) WKB solution of an equation of motion rewritten in terms of different independent/dependent variables will generally differ. This is easily seen by observing that the lowest order adiabatic initial conditions match those obtained via Hamiltonian diagonalisation [Eq. \(7.85\)](#).

**$\alpha$ -vacua:** The behaviour of  $\alpha$ -vacua [\[177\]](#) under canonical transformations has been reviewed in [\[180\]](#). These vacua were found not to be canonically invariant.

### 7.4.2 Primordial power spectra

Primordial power spectra show the power in the Fourier modes of gauge-invariant curvature perturbations ( $\mathcal{R}_k$ ) after horizon exit – when their characteristic lengthscale  $k^{-1}$  first exceeds the size of the comoving Hubble horizon  $(aH)^{-1}$  during inflation. When this happens, the amplitude of fluctuations becomes constant, as shown in [Fig. 7.1](#) alongside definitions of quantities used in this section. Primordial power spectra are not directly observable<sup>3</sup>, but are the first physical quantities one can derive from the perturbations, and hence the first physical

<sup>3</sup>They are, however, reconstructable [\[169, 199, 200\]](#).




Fig. 7.1 Schematic evolution of the comoving Hubble horizon (black) and a ‘pivot’ perturbation  $\mathcal{R}_{k_*}$  (red), whose characteristic lengthscale is  $0.05 \text{ Mpc}^{-1}$ . Inflation starts as the comoving horizon begins to shrink, and continues until it reaches its minimum. The total number of e-folds by which the universe has grown during inflation is  $N_{\text{tot}} = N_* + N_{\dagger}$ , with  $N_*$  e-folds passing between the pivot perturbation exiting the horizon and the end of inflation.

objects affected by the existence of a range of quantum initial conditions. To compute them, first the ‘cosmological background’ quantities such as the scale factor and Hubble horizon are computed by solving the cosmological field equations [201], which are derived from the action Eq. (7.13) using Einstein’s field equations.

In this work, it is assumed that slow-roll inflation had a start and was preceded by a non-inflationary phase. An eternal inflation scenario would allow one to initialise all primordial perturbation modes deep inside the Hubble horizon, which would result in a power-law primordial power spectrum that is independent of the choice of primordial initial conditions and the choice of vacuum. Our result concerning the (non-)invariance of vacuum choices under field redefinitions and addition of surface terms would be theoretically important nonetheless, but the aim of this and the following sections is to demonstrate the observational impact of the choice of canonical variables, and therefore a physically well-motivated setup is chosen in which the choice of vacuum affects the primordial power spectrum. Generally it may also be attractive to choose a model without too long an inflationary phase in order to avoid having to account for trans-Planckian physics when initialising the primordial perturbations (see [173, 192, 196, 202] and references therein), due to lengthscales observed today having been smaller than the Planck length at some earlier epoch.

The numerical solution is initiated at the start of inflation, which is assumed to be preceded by a *kinetically dominated* [115, 117, 118] phase. In kinetic dominance the inflaton particle’s kinetic energy dominates over its potential energy, causing the comoving Hubble horizon to grow and reach a maximum as the universe enters slow-roll inflation.

For simplicity, a single-field inflationary model will be considered with a potential that is quadratic in the field,

$$V(\phi) = \frac{1}{2}\mu^2\phi^2. \quad (7.70)$$

For convenience, the number of e-foldings  $N \equiv \ln a$  is chosen to be the independent variable in the field equations. With this choice, they become

$$\frac{d \ln D}{dN} = 4 + D \left( 4K - 2e^{2N} V(\phi) \right), \quad (7.71)$$

$$\left( \frac{d\phi}{dN} \right)^2 = 6 + D \left( 6K - 2e^{2N} V(\phi) \right), \quad (7.72)$$

where  $D = (aH)^{-2}$  and  $K$  is the spatial curvature that can take values  $0, \pm 1$  for flat, closed, and open universes, respectively. We set  $K = 0$  when calculating the primordial power spectra and all subsequent computations. At the start of inflation, the field equations give

$$D = D_i = \frac{2e^{-2N_i}}{V(\phi_i)}, \quad (7.73)$$

leaving  $N_i$  and  $\phi_i$  as adjustable parameters. In a flat universe one can set  $N_i = 0$  without loss of generality, so the only free parameter left in the background initial conditions is the initial field strength,  $\phi_i$ . This determines  $N_{\text{tot}}$ , the total number of e-folds during inflation, and  $N_{\dagger}$ , the number of e-folds between the start of inflation and a ‘pivot’ perturbation mode with lengthscale  $k_* = 0.05 \text{ Mpc}^{-1}$  exiting the Hubble horizon. Observations constrain the value of  $N_*$ , the number of e-folds between the pivot scale exiting the Hubble horizon and the end of inflation, to be between 50 and 60 [203–205]. A value for  $\phi_i$  consistent with the ‘just enough inflation’ scenario [206, 207], where  $N_* \sim N_{\text{tot}}$ , is chosen. In the inflationary potential,  $\mu$  denotes the inflaton mass, which determines the overall ‘normalisation’ of the primordial power spectrum and can be calculated using the slow-roll approximation from  $N_*$  and  $A_s$  to yield the observed normalisation. The primordial parameters and the cosmological parameters used in the following sections are summarised in [Table 7.2](#).

Once the cosmological field equations have been solved, the perturbations  $\mathcal{R}_k$  are initialised simultaneously (independently of  $k$ ) at the start of inflation with two arbitrary sets of initial conditions (e.g.  $\{\mathcal{R}_k = 1, d\mathcal{R}_k/dN = 0\}$ ,  $\{\mathcal{R}_k = 0, d\mathcal{R}_k/dN = 1\}$ ) and for each  $k$ -mode the Mukhanov–Sasaki equation is solved to obtain their evolution in  $N$ . We use the numerical solver `oscode` [1] to carry out this computation efficiently. The amplitude of each perturbation is read off when they are well outside the Hubble horizon,  $k < 10^{-2} aH$ . The two solutions for

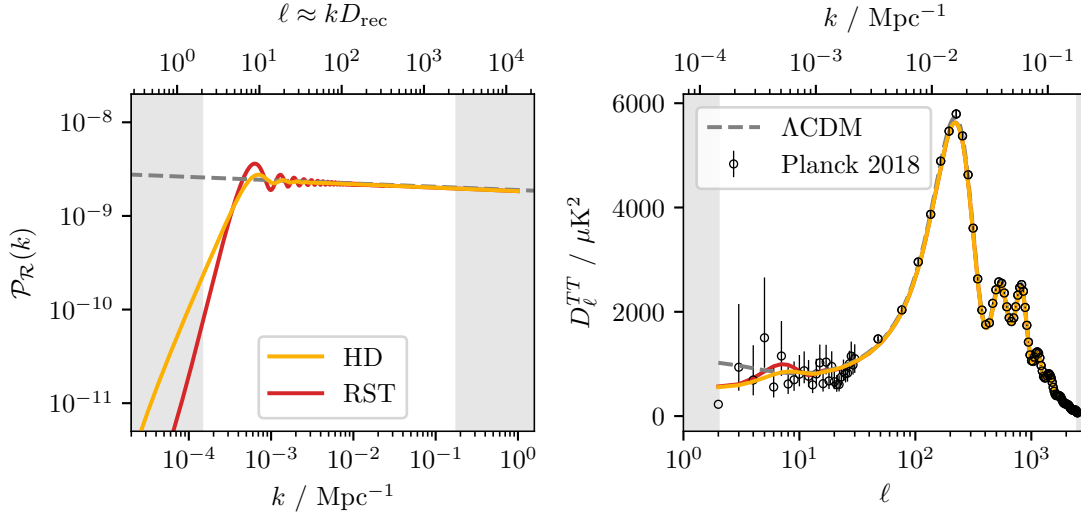


Fig. 7.2 Primordial power spectrum (top panel) and angular CMB power spectrum of temperature anisotropies (bottom panel) for the standard Hamiltonian Diagonalisation (HD) and Renormalised Stress-Energy-Tensor (RST) vacuum conditions.

$\mathcal{R}_k$  can then be linearly combined to satisfy any initial condition at the start of inflation from the sets derived, without having to re-compute the evolution of perturbations. Finally, the primordial power spectrum is constructed according to

$$\mathcal{P}_{\mathcal{R}}(k) = \frac{k^3}{2\pi^2} |\mathcal{R}_k|^2. \quad (7.74)$$

Fig. 7.2 shows example primordial power spectra generated using conventional Hamiltonian diagonalisation Eq. (7.28) and renormalised stress–energy tensor Eq. (7.41) initial conditions, and the corresponding CMB angular  $TT$  spectra. The details of how the latter kind of spectra are computed can be found in Section 7.4.3. Note the dual  $x$ -axis shows both the wavenumber  $k$  and the multipole  $\ell$ . The conversion between the two is performed using the Limber approximation,

$$\ell \approx k D_{\text{rec}} = k \frac{r_s}{\theta_s}, \quad (7.75)$$

where  $D_{\text{rec}}$  is the distance to the last scattering surface at recombination,  $r_s$  is the sound horizon at recombination, and  $\theta_s$  is the angular parameter. The latter two can be derived from the Planck baseline CMB parameters listed in Table 7.2.

To show the full range of primordial power spectra achievable using Hamiltonian diagonalisation initial conditions Eq. (7.50), one needs to consider the two-dimensional space spanned by the two initial condition parameters,  $h_2/h_0$  and  $h_1/h_0$ . Since it is difficult to visualise individual spectra in this two-dimensional space, contour plots of different features in the power spectra are shown instead. The features considered are listed below.

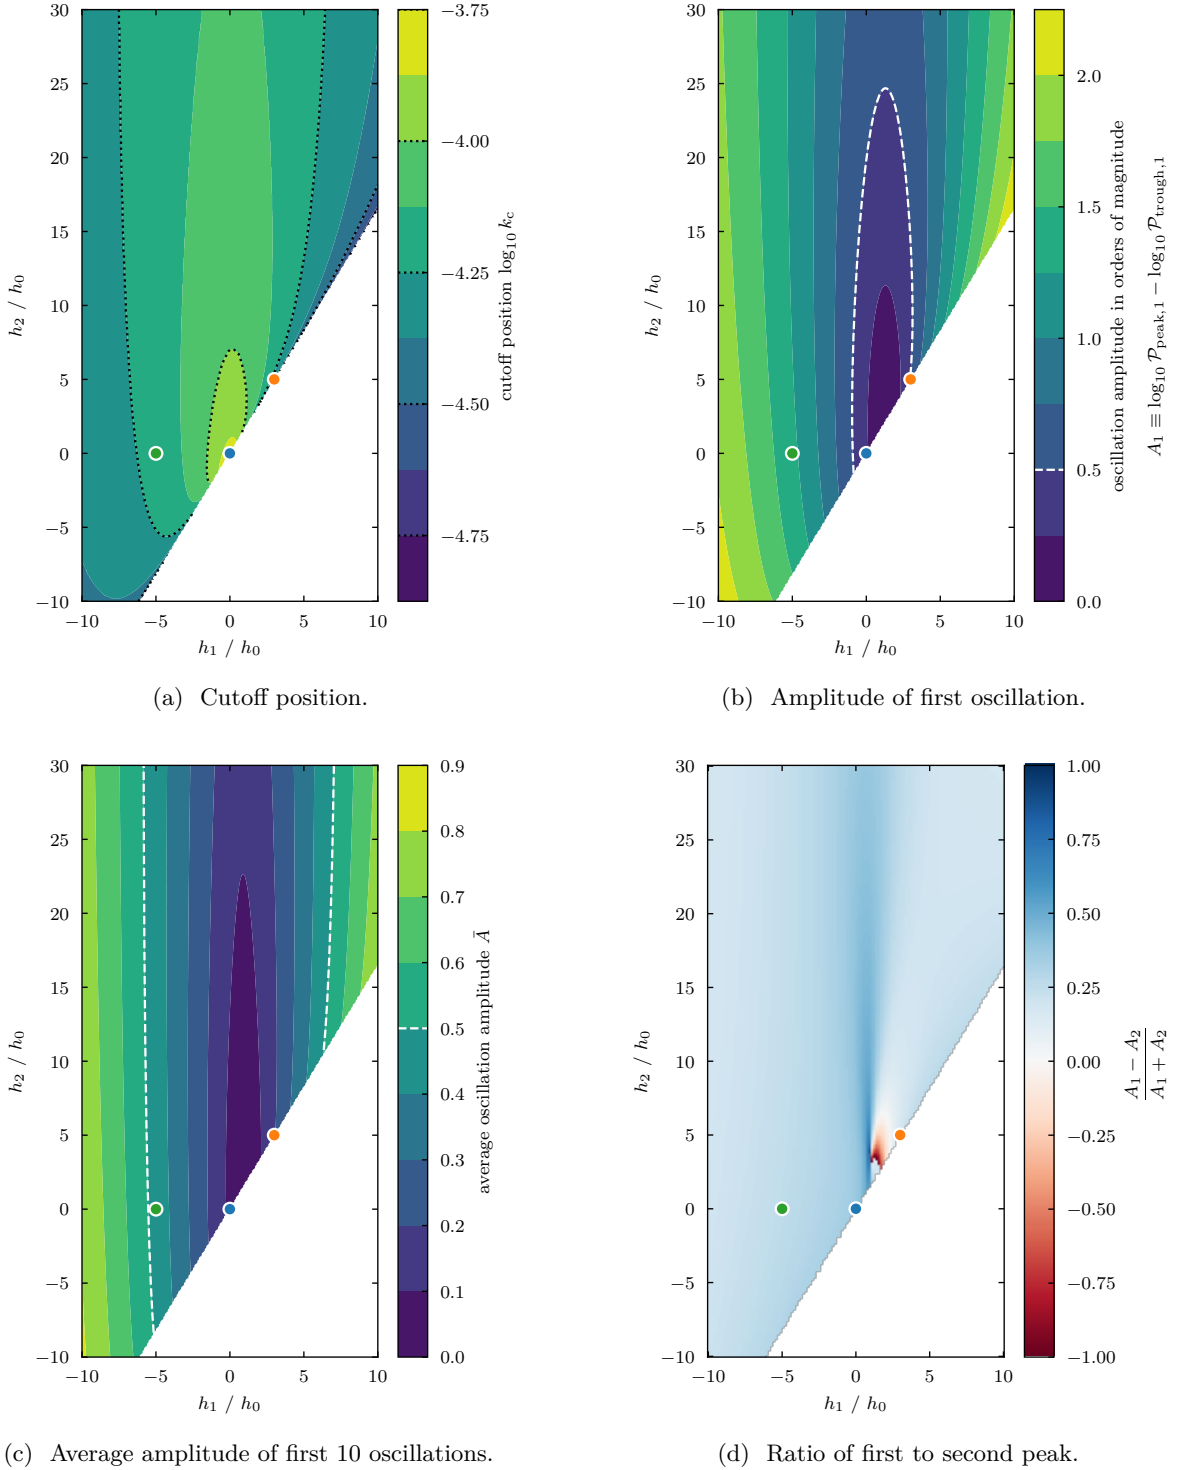


Fig. 7.3 Dependence of features of the primordial power spectrum derived from Hamiltonian Diagonalisation initial conditions Eq. (7.50) on the parameters  $h_1/h_0$  and  $h_2/h_0$ . The cutoff position is measured as the wavenumber  $k_c$  where the power spectrum drops below a value of  $10^{-10}$  towards large scales. The oscillation amplitude is measured as the log-difference between a peak and the following trough. The three coloured points correspond to parameter pairs considered in Section 7.4.3, their individual primordial power spectra are shown in Fig. 7.5.

Table 7.2 The fiducial values used for the cosmological parameters come from Planck 2018 TT,TE,EE+lowE [165].


parameter	fiducial value	parameter	fiducial value
$\Omega_b h^2$	0.02236	$A_s$	$2.101 \times 10^{-9}$
$\Omega_c h^2$	0.1202	$n_s$	0.9649
$h$	0.6727	$\phi_i$	16 $m_p$
$\tau_{\text{reio}}$	0.0544	$N_*$	55

1. **Cutoff position:** position of the low- $k$  cutoff present in all primordial power spectra considered. The cutoff is a result of considering kinetic dominance instead of eternal inflation, and is caused by perturbation modes that do not enter the Hubble horizon. The cutoff as the position where the amplitude of the power spectrum drops below  $10^{-10}$ .
2. **Amplitude of first oscillation:** calculated as the logarithmic difference between the first peak and the following trough.
3. **Average amplitude of the first 10 oscillations.**
4. **Ratio of the first to second peak:** defined as  $(A_1 - A_2)/((A_1 + A_2))$ , where  $A_{1,2}$  are the amplitudes of the first and second peaks, respectively.
5. **Frequency of oscillations:** the leading frequency in the Fourier transform of  $P_{\mathcal{R}}(k)$  was also considered as a feature initially. However, the frequency was not expected to change with  $(h_1/h_0, h_2/h_0)$ , which was indeed what has been found, and therefore the resulting contour plot is not shown.

Figs. 7.3a to 7.3d show contour plots of the features 1–4 in primordial power spectra in  $(h_1, h_2)$ -space, with the three coloured points corresponding to parameter pairs considered in Section 7.4.3. The empty ‘forbidden’ region in all plots is a result of  $h_2/h_0$  and  $h_1/h_0$  appearing under a square root in Eq. (7.50), causing there to be a lower limit to possible  $k$  values if


$$\frac{h_2}{h_0} + 2 \frac{h_1 z'}{h_0 z} < 0. \quad (7.76)$$

Figs. 7.3b and 7.3d exhibit an interesting feature around  $h_1/h_0 = 1.5$ ,  $h_2/h_0 = 3.5$ : the emergence of a new peak from the low- $k$  region of the primordial power spectrum. The phenomenon is shown in Figs. 7.4a to 7.4d: Fig. 7.4a and Fig. 7.4b mark the region of interest in parameter space and the extent of inversion, while Fig. 7.4c and Fig. 7.4d show the emergence of the new peak along orthogonal directions in parameter space.



(a) Amplitude of first oscillation. (zoomed in from Fig. 7.3b)

(b) Ratio of first to second peak. (zoomed in from Fig. 7.3d)



(c) Sample spectra along horizontal line in upper panels. (d) Sample spectra along vertical line in upper panels.

Fig. 7.4 Zoom-in of the centre region of the contour plots in Figs. 7.3b and 7.3d. The dotted horizontal and vertical lines correspond to the line collections of primordial power spectra in the lower panels. The sharp lines in the contour plot correspond to the (dis-)appearance of the first peak.

It is to be noted that one can arrive at the renormalised stress–energy tensor initial conditions Eq. (7.41) from Hamiltonian diagonalisation, by choosing the field to be quantised as  $\mathcal{R}$ , i.e. choosing  $h = 1$ . This choice translates to  $h_1/h_0 = 0$ ,  $h_2/h_0 = 0$  in Figs. 7.3, 7.4 and 7.8. The ‘standard’ Danielsson prescription leads to the same initial conditions as minimising the renormalised stress–energy tensor, and so also corresponds to  $h_1/h_0 = 0$ ,  $h_2/h_0 = 0$ . Conventional Hamiltonian diagonalisation is equivalent to choosing  $h = z^{-1}$ , or  $h_1/h_0 \approx 0.82$ ,  $h_2/h_0 \approx 1.37$  in Figs. 7.3, 7.4 and 7.8.

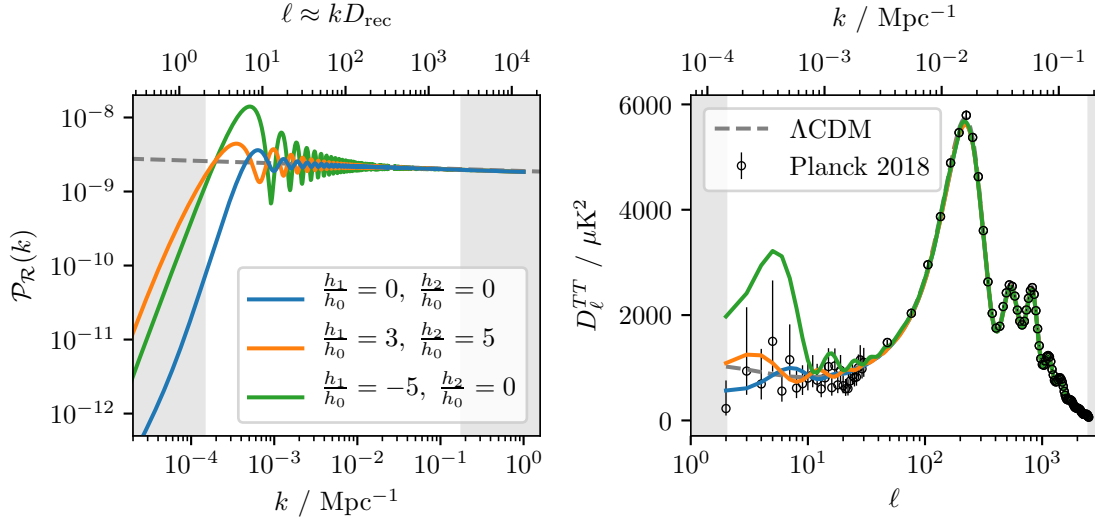


Fig. 7.5 Primordial power spectrum (top panel) and angular CMB power spectrum of temperature anisotropies (bottom panel) for three different combinations of  $h_1/h_0$  and  $h_2/h_0$  as used in Eq. (7.50).

### 7.4.3 Cosmic Microwave Background

The angular CMB  $TT$ ,  $TE$ , and  $EE$  power spectra are generated from primordial power spectra using the Boltzmann code CLASS [126–129], with Planck 2018 baseline [165] cosmological parameters as presented in Table 7.2.

There are no obvious choices of physical features in the range of angular spectra observable today that would allow their visualisation in contour plots, therefore without any claim to completeness, primordial power spectra and the corresponding  $TT$ ,  $TE$  and  $EE$  spectra are shown of the three points in  $(h_1/h_0, h_2/h_0)$ -space which were marked in colour in Figs. 7.3a to 7.3d and 7.4a to 7.4d. The primordial power spectra and the  $TT$  spectra are shown in Fig. 7.5, while Fig. 7.6 shows the  $TE$  and  $EE$  spectra. The low- $\ell$  region of all these CMB spectra are shown separately in Fig. 7.7.

Immediately it is clear that in the ‘just enough inflation’ case, the choice of the function  $h$  (through the values  $h_1/h_0$  and  $h_2/h_0$ ) can influence the CMB observed today, despite  $h$  stemming from the choice of canonical variables. Again it must be stressed that in the eternal inflation case, the different initial conditions discussed would yield the same observable spectra.

The question of whether observations could theoretically rule out certain regions of the  $(h_1/h_0, h_2/h_0)$  parameter space is worth investigating, even though the inflationary potential used in this work,  $V(\phi) \propto \phi^2$  has mostly been ruled out by the Planck 2015 results [200], and is only used here for the sake of simplicity. The posterior probability of  $(h_1/h_0, h_2/h_0)$  given the Planck 2018 low- $\ell$  ( $\ell \leq 30$ ) temperature-only likelihood [167] is shown in Fig. 7.8, where a uniform prior was used over the range  $-10 \leq h_1/h_0 \leq 10$ ,  $-10 \leq h_2/h_0 \leq 30$ , excluding the




Fig. 7.6 CMB angular power spectra of the TE cross-spectrum and the EE auto-spectrum for three different combinations of  $h_1/h_0$  and  $h_2/h_0$  as in Fig. 7.5.

forbidden region. The posterior has been explored using the PolyChord sampler [136, 137]. The lowT data was used because Figs. 7.5 and 7.6 suggest that it is the low- $\ell$  region the different initial conditions have the most effect on.

It is clear that observational data can be used to select a preferred region of the initial condition parameters, but the point this section aims to emphasise is the wide-range effects that the choice of canonical variables can have on observations when using Hamiltonian diagonalisation to define the vacuum.

## 7.5 Conclusions

In part, this chapter serves to highlight some points made by Fulling [171] regarding the lack of robustness of Hamiltonian diagonalisation under canonical transformations, in that it confirmed that this procedure yields physically distinct vacua, and different initial conditions for the scalar curvature perturbations. All sets of initial conditions derived this way have been shown to take the general form proposed by [196], and therefore could be argued to be physically sensible. On the other hand, this chapter has illustrated the effect that the choice of initial conditions has on the primordial power spectrum and the CMB angular power spectra under the ‘just enough inflation’ assumption. It has been shown that in such inflation models, different choices of canonical variables in a canonically non-invariant vacuum prescription such as Hamiltonian diagonalisation would be distinguishable by observation, noting that in a model where all perturbation modes spend a significant amount of time within the Hubble horizon, all initial conditions considered would yield the same observable results. Two other choices of vacuum have been subjected to the same transformations: one obtained by minimising the 00-component

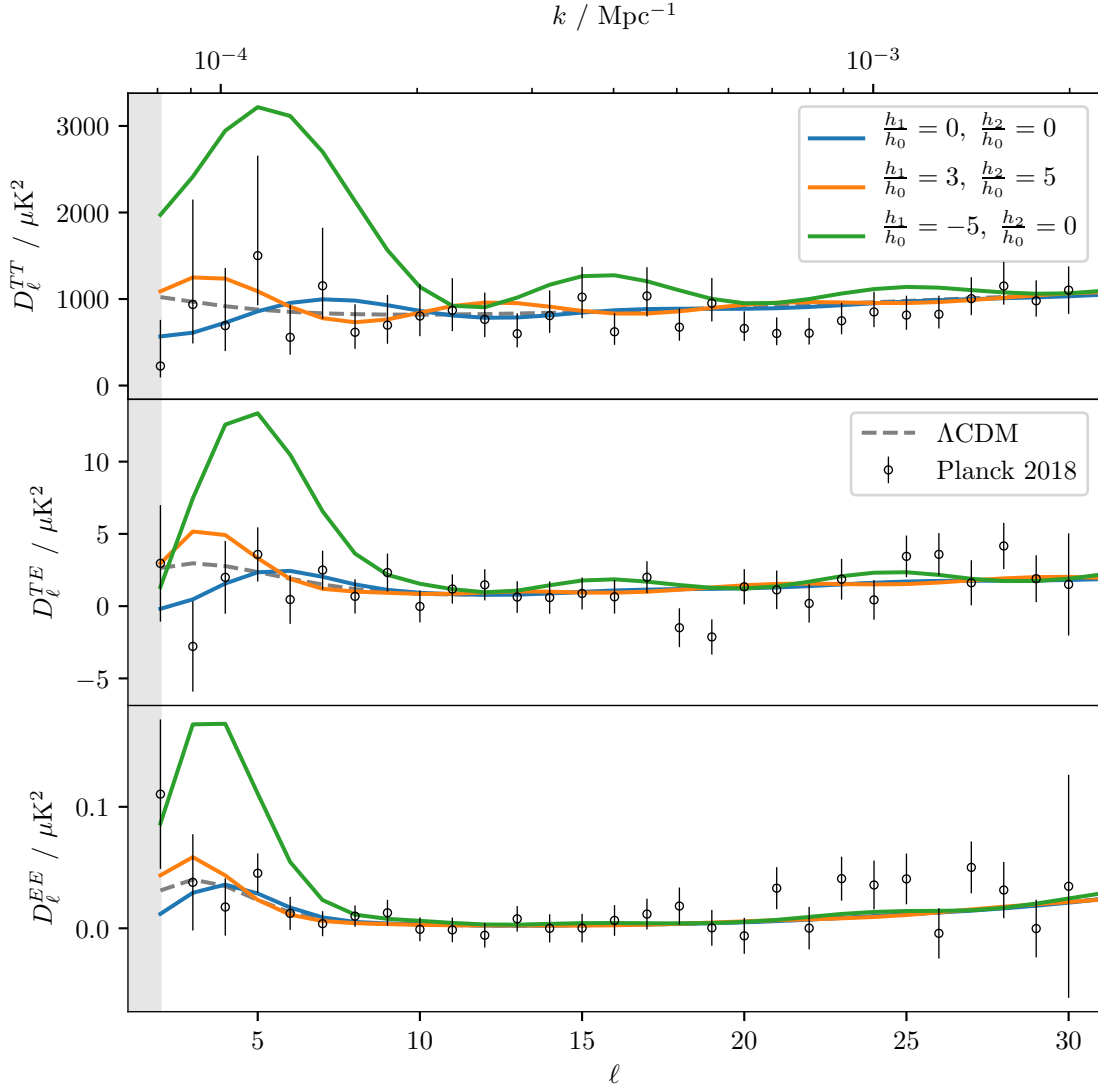


Fig. 7.7 Low- $\ell$  region of the CMB angular power spectra from Figs. 7.5 and 7.6.

of the renormalised stress–energy tensor [12], and the Danielsson vacuum [173]. In doing so, it has been found that the Danielsson vacuum is inherently sensitive to canonical field redefinitions and addition of surface terms, but the renormalised stress–energy tensor minimisation is not, despite not having been constructed to be invariant in this manner.

Many pieces of work since [171] acknowledge the existence of Hamiltonians related via canonical transformations, but suggest that there is a preferred Hamiltonian which leads to the conventionally used initial conditions. It has been noted that the initial conditions derived via the renormalised stress–energy tensor, which have thus been shown to be invariant under canonical transformations, differ from the initial conditions associated with this preferred Hamiltonian. Since the procedure of minimising the renormalised stress–energy tensor yields




Fig. 7.8 Prior samples and posterior contours showing the best-fit values of the  $(h_1/h_0, h_2/h_0)$  parameters, given Planck 2018 lowT data. The coloured dots refer to parameter pairs described in Section 7.4.3 and Fig. 7.5.

unambiguous initial conditions for the primordial perturbations and unambiguous quantum mechanical observables (expectation values of operators) under transformations that leave other physical traits (such as the equation of motion and commutator structure) invariant, it is recommended as the preferred prescription for setting the vacuum.

## Appendix 7.A Hamiltonian diagonalisation under field redefinitions and addition of surface terms

This appendix presents the details of deriving initial conditions via Hamiltonian diagonalisation for the variable pair  $(\tau, \chi)$ , related to cosmic time and the gauge-invariant curvature perturbation

$\mathcal{R}$  by Eq. (7.42). For the derivations to follow in general, our computational strategy can be summarised as:

1. Write down the action.
2. Perform the transformation in question, expressing the action in terms of the field of interest.
3. Find the equation of motion.
4. Find the condition for the equation of motion to be of an undamped harmonic oscillator (for the purposes of quantising the field).
5. Find the momentum conjugate to the field; quantise both as quantum oscillators.
6. Write down the commutators of the system, and derive conditions for these to be canonical (thus ensuring that the transformation itself is canonical).
7. Write down the definition of vacuum, i.e. the quantity to be minimised, and minimise it with respect to the Fourier modes of the quantised field, subject to the conditions derived.

### 7.A.1 Field redefinition

Changing variables in the perturbed action Eq. (7.15) leads to

$$S = \int d^3x d\tau \frac{C}{2} \left[ (\partial_\tau \chi)^2 + 2\chi (\partial_\tau \chi) \frac{\partial_\tau h}{h} + \chi^2 \left( \frac{\partial_\tau h}{h} \right)^2 - (\nabla \chi)^2 \left( \frac{h^2 z^2}{C} \right)^2 \right], \quad (7.77)$$

with  $C = h^2 \dot{\tau} a z^2$ . This will in general yield the equation of motion of an oscillator with a non-zero first-derivative term. If the field  $\chi$  is to be quantised as an oscillator, the first-derivative term needs to vanish, which holds if

$$C = C_0 = \text{const.} \quad (7.78)$$

Setting  $C$  to be constant, the equation of motion in terms of the Fourier modes of  $\chi$  becomes

$$0 = \partial_{\tau\tau} \chi_k + \left[ \left( \frac{kh^2 z^2}{C_0} \right)^2 + \frac{\partial_{\tau\tau} h}{h} - 2 \left( \frac{\partial_\tau h}{h} \right)^2 \right] \chi_k. \quad (7.79)$$

To follow the conventional procedure, we integrate Eq. (7.77) by parts to eliminate the term linear in  $\partial_\tau \chi$  to get

$$S = \int d^3x d\tau \frac{C_0}{2} \left[ (\partial_\tau \chi)^2 + \chi^2 \left( \frac{\partial_{\tau\tau} h}{h} - 2 \left( \frac{\partial_\tau h}{h} \right)^2 \right) - (\nabla \chi)^2 \left( \frac{h^2 z^2}{C_0} \right)^2 \right]. \quad (7.80)$$

The momentum conjugate to  $\chi$  is then

$$\pi_\chi = C_0 \partial_\tau \chi. \quad (7.81)$$

We now quantise  $\chi$  in the same way as  $v$  in Eq. (7.22). Imposing canonical commutation relations on the creation and annihilation operators results in the constraint

$$(\partial_\tau \chi_k) \chi_k^* - (\partial_\tau \chi_k^*) \chi_k = -\frac{i}{C_0}, \quad (7.82)$$

because the momentum conjugate to  $\chi$  was scaled by  $C_0$ . This is a necessary condition for the field redefinition to be a canonical transformation. The normal-ordered Hamiltonian then takes the general form

$$H = \frac{1}{2} \int \frac{d^3 k}{(2\pi)^3} \left[ \hat{a}_k \hat{a}_{-k} F_k(\tau) + \hat{a}_k^\dagger \hat{a}_{-k}^\dagger F_k^*(\tau) + \left( 2\hat{a}_k^\dagger \hat{a}_k + \delta^{(3)}(0) \right) E_k(\tau) \right], \quad (7.83)$$

with

$$\begin{aligned} E_k(\tau) &= |\partial_\tau \chi_k|^2 + \omega_k^2 |\chi_k|^2, \\ F_k(\tau) &= (\partial_\tau \chi_k)^2 + \omega_k^2 \chi_k^2, \\ \omega_k^2 &= \left( \frac{kh^2 z^2}{C_0} \right)^2 + \frac{\partial_{\tau\tau} h}{h} - 2 \left( \frac{\partial_\tau h}{h} \right)^2. \end{aligned} \quad (7.84)$$

Minimising the vacuum expectation value of the Hamiltonian for each  $k$ -mode separately, subject to the constraint Eq. (7.82) thus gives the generalised initial conditions

$$\begin{aligned} |\chi_k|^2 &= (2C_0 \omega_k)^{-1}, \\ \partial_\tau \chi_k &= -i\omega_k \chi_k. \end{aligned} \quad (7.85)$$

### 7.A.2 Surface terms

In order to get to the form of the Hamiltonian Eq. (7.83), we integrated by parts once and discarded a boundary term:

$$\begin{aligned} \mathcal{S} &\supset \int d^3 x d\tau \frac{C_0}{2} \left[ 2\chi(\partial_\tau \chi) \frac{\partial_\tau h}{h} + \chi^2 \left( \frac{\partial_\tau h}{h} \right)^2 \right] \\ &= \left[ \chi^2 \frac{\partial_\tau h}{h} \right] - \int d^3 x d\tau \left[ \frac{\partial_{\tau\tau} h}{h} \chi^2 - 2 \left( \frac{\partial_\tau h}{h} \right)^2 \chi^2 \right]. \end{aligned} \quad (7.86)$$

This is equivalent to adding the surface term

$$- \partial_\tau \left( \chi^2 \frac{\partial_\tau h}{h} \right). \quad (7.87)$$

If we instead kept the action in its original form, Eq. (7.77), the momentum conjugate to  $\chi$  would be

$$\pi = C_0 \left( \partial_\tau \chi + \frac{\partial_\tau h}{h} \chi \right), \quad (7.88)$$

and the normal-ordered Hamiltonian would take the same form as Eq. (7.83), but now with

$$\begin{aligned} E_k(\tau) &= |\partial_\tau \chi_k|^2 + \sigma_k^2 |\chi_k|^2, \\ F_k(\tau) &= (\partial_\tau \chi_k)^2 + \sigma_k^2 \chi_k^2, \\ \sigma_k^2 &= \left( \frac{kh^2 z^2}{C_0} \right)^2 - \left( \frac{\partial_\tau h}{h} \right)^2. \end{aligned} \quad (7.89)$$

As seen by symmetry, this yields initial conditions for  $(\tau, \chi)$  of the same form as Eq. (7.47), but with  $\omega_k$  replaced with  $\sigma_k$ . Switching back to  $(\eta, v)$ , they are

$$\begin{aligned} |v_k|^2 &= \frac{1}{2\sqrt{k^2 - \left(\frac{h'}{h}\right)^2}}, \\ v'_k &= \left( -i\sqrt{k^2 - \left(\frac{h'}{h}\right)^2} + \frac{h'}{h} + \frac{z'}{z} \right) v_k. \end{aligned} \quad (7.90)$$

## Appendix 7.B Minimising the renormalised stress–energy tensor under field redefinitions

We generalise the results Eqs. (7.40) and (7.41) in analogy with Section 7.A.1. We rewrite the action Eq. (7.19) in terms of the redefined time and field

$$t \rightarrow \tau, \quad \varphi \rightarrow \chi(x, \tau) \equiv \frac{\varphi}{h} = \frac{y}{ah}, \quad (7.91)$$

keeping in mind that the metric then changes to  $g_{\mu\nu} = \text{diag}(\dot{\tau}^{-2}, -a^2, -a^2, -a^2)$ . This yields

$$\begin{aligned} S &= \int d^3x d\tau \frac{C}{2} \left[ (\partial_\tau \chi)^2 + 2(\partial_\tau \chi) \chi \frac{\partial_\tau h}{h} + \chi^2 \left( \frac{\partial_\tau h}{h} \right)^2 \right. \\ &\quad \left. - (\nabla \chi)^2 \left( \frac{h^2 a^2}{C_0} \right)^2 - \chi^2 m^2 \left( \frac{h^2 a^3}{C_0} \right)^2 \right], \end{aligned} \quad (7.92)$$

with  $C = h^2 \dot{\tau} a^3$ . As before, the equation of motion resulting from the variation of this action can be made first-derivative-free by setting

$$C = C_0 = \text{const}, \quad (7.93)$$

in which case it becomes, in Fourier space,

$$\partial_{\tau\tau}\chi_k + \left[ \left( \frac{kh^2a^2}{C_0} \right)^2 + \frac{\partial_{\tau\tau}h}{h} - 2 \left( \frac{\partial_{\tau}h}{h} \right)^2 + \left( \frac{mh^2a^3}{C_0} \right)^2 \right] \chi_k = 0. \quad (7.94)$$

The field  $\chi$  is then quantised, which makes  $\varphi$  take the form

$$\varphi(x) = \int \frac{d^3k}{(2\pi)^3} h(\tau) \left[ \hat{a}_k \chi_k(\tau) e^{i\mathbf{k}\cdot\mathbf{x}} + \hat{a}_k^\dagger \chi_k^*(\tau) e^{-i\mathbf{k}\cdot\mathbf{x}} \right]. \quad (7.95)$$

The Hadamard Green function can then be written

$$G^{(1)}(x, x') = \int \frac{d^3k}{(2\pi)^3} h(\tau) h(\tau') \left[ \chi_k(\tau) \chi_k^*(\tau') e^{i\mathbf{k}\cdot(\mathbf{x}-\mathbf{x}')} + \chi_k^*(\tau) \chi_k(\tau') e^{-i\mathbf{k}\cdot(\mathbf{x}-\mathbf{x}')} \right]. \quad (7.96)$$

Acting on this with the bi-scalar derivative function and taking the 00 component thus gives

$$\begin{aligned} \langle 0 | T_{00} | 0 \rangle_{\text{ren}} = \tilde{T} + \frac{1}{2} \int \frac{d^3k}{(2\pi)^3} h^2 \left[ \left( \partial_{\tau} \chi_k + \frac{\partial_{\tau} h}{h} \chi_k \right) \left( \chi_k^* + \frac{\partial_{\tau} h}{h} \chi_k^* \right) \right. \\ \left. + \left( \frac{k^2}{a^2 \dot{\tau}^2} + \frac{m^2}{\dot{\tau}^2} \right) \chi_k \chi_k^* \right]. \end{aligned} \quad (7.97)$$

The normalisation constraint on the mode functions  $\chi_k$  takes the same form as for the generalised perturbed classical action in [Section 7.A.1](#), and is therefore given by [Eq. \(7.82\)](#). Minimising the expression [Eq. \(7.97\)](#) subject to this constraint gives for the mode functions  $\chi_k$

$$\begin{aligned} |\chi_k|^2 &= \frac{1}{2\sqrt{h^4 a^4 (k^2 + m^2 a^2)}}, \\ \partial_{\tau} \chi_k &= \left( -\frac{i}{C_0} \sqrt{h^4 a^4 (k^2 + m^2 a^2)} - \frac{\partial_{\tau} h}{h} \right) \chi_k. \end{aligned} \quad (7.98)$$



# Chapter 8

## Conclusions and future outlook

### 8.1 Computational tools

[Chapter 2](#) of this thesis discussed the contrast between the numerical analysis of non-oscillatory and highly oscillatory ODEs. Methods for solving the former kind of ODEs customarily rely on Taylor’s theorem or other polynomial expansions (in the stepsize  $h$ ) for forecasting the numerical solution at the end of each step. The reason conventional methods fail to integrate oscillatory ODEs efficiently is that in the presence of oscillations, these polynomial approximations only hold for  $h \lesssim \lambda$ , where  $\lambda$  is the approximate wavelength of oscillations. The number of steps standard (e.g. linear multistep) methods have to take while traversing an oscillatory ODE therefore scales with the frequency, which quickly raises the computational cost. Standard methods struggle to solve even  $\ddot{x}(t) + \omega^2 x(t) = 0$ , the simplest analytic example of an oscillatory ODE and oscillatory equivalent of the Dahlquist test equation. This hints at the possibility of a solution using an approximately sinusoidal approximation instead of Taylor’s theorem to forecast the solution. Oscillation can, however, take many forms and arise from many different types of ODE systems, which makes it difficult to find a fit-for-all approximation for all kinds of oscillatory ODEs. In [Chapter 2](#) I chose to focus on the class of oscillatory ODEs which can be written

$$\ddot{x}(t) + 2\gamma(t)\dot{x}(t) + \omega^2(t)x(t) = 0, \quad (8.1)$$

and reviewed efforts to construct efficient numerical methods to tackle them. The takeaway from this introduction is that the realm of oscillatory ODEs is far less unified than that of their non-oscillatory counterparts and is in need of more systematic review, and that the most successful numerical methods for oscillatory equations exploit our prior knowledge about the structure of the ODE system and the form of the solution.

Motivated by the abundance of oscillatory ODEs of the form [Eq. \(8.1\)](#) in physics and cosmology, I developed an efficient numerical routine and associated open-source software, `oscode`. Its core algorithm builds on the ideas of [\[68\]](#) and works as a generic stepper algorithm

that switches between an oscillatory and a slowly varying approximation for forecasting the solution at the end of a step of size  $h$ , depending on which kind gives the largest possible step while keeping the local error estimate within a user-specified tolerance. The oscillatory approximation is based on the Wentzel–Kramers–Brillouin expansion familiar from quantum mechanics, but generalised to include a potential damping term  $2\gamma(t)x(t)$  in Eq. (8.1). Due to the WKB expansion being a good approximation over distances  $h \gg \lambda$ , `oscode` is able to take huge steps in the oscillatory regions of the solution (if the frequency of oscillations is slowly changing) and is extremely efficient. In the slowly varying regions of the solution, `oscode` is able to switch to a 5th order Runge–Kutta method whose abscissas have been chosen so as to minimise the number of  $\omega$  and  $\gamma$  evaluations necessary per step. All critical components of the algorithm were discussed in Chapter 3, including local error estimates for the two kinds of steps, the switching mechanism, and the adaptive stepsize algorithm. The chapter also discussed some example applications of `oscode` through which its error properties and performance were assessed.

Chapter 4 then presented the method by which `oscode` is able to compute dense output, i.e. a continuous interpolation of the solution in the highly oscillatory (and non-oscillatory) regime.

After an introduction to the relevant cosmology in Chapter 5, Chapter 6 discussed `oscode`'s use in cosmology to quickly evolve primordial perturbations and compute their power spectrum. I showed results from [159], which builds heavily on `oscode` for quick forward modelling of closed universe models of varying primordial spatial curvature.

### 8.1.1 Future work

`Oscode`, which embeds the asymptotic WKB expansion in an iterative solver and pairs it with a Runge–Kutta method, proved to be extremely efficient at solving ODEs of the form  $\ddot{x}(t) + 2\gamma(t)\dot{x}(t) + \omega^2(t)x(t) = 0$  in both the highly oscillatory and the slowly varying regimes. Its implementation as a general numerical solver and high-quality peer-reviewed code make it easily applicable beyond cosmology. These also provide a good basis to expand from, and there are several possible extensions worthy of exploration.

#### Expanding to higher orders

The accuracy `oscode` can achieve is adequate for most applications in physics, where measurement error provides a natural lower limit for the tolerance. This applicability could be extended within and beyond the physical sciences by expanding the tolerance range within which `oscode` can meaningfully operate. For relative tolerances below  $\sim 10^{-6}$ , the 5th order Runge–Kutta method `oscode` uses is not recommended – in practice, a 7th–8th order pair is used in this tolerance regime, e.g. those computed by Dormand and Prince [208]. The use of higher-order solvers allows the stepsize to be *reduced less* for a larger gain in accuracy, and the decrease in the number of steps the methods thus needs to take outweighs the increased computational cost

per step. Gaining accuracy in the highly oscillatory regions is more challenging, as the main appeal of the WKB expansion (with a *fixed* number of terms) is the gentle scaling of local error with increasing stepsize – which also causes the stepsize to have to be reduced drastically for the same accuracy gain. One way to resolve this could be to adapt the number of terms included in the WKB series to the accuracy required. Since the WKB expansion is asymptotic, the last terms would need to be monitored throughout the integration in order to avoid a potential growth of terms, and the number of terms included would need to be updated dynamically. Another potential solution would be to allow for the algorithm to switch to the phase function method [56, 57] in regions of high oscillation. I have already started correspondence with James Bremer and discussed this promising possibility.

### Improved, dynamic error estimation

`oscode`'s adaptive stepsize algorithm is based on an asymptotic error estimate. As discussed in Section 2.3.2, the asymptotic error model is valid at small stepsizes, when the  $\mathcal{O}(h^{p+2})$  term in the local error estimate can safely be neglected, and this can always be achieved by restricting the stepsizes. There are cases when we may wish not to do this, an immediate example being when the tolerance requirements are not too strict and we can prioritise efficiency by taking large steps, as would be the case in some cosmological applications of `oscode`. Another important example is that of stiff equations being solved by explicit methods: the local error estimates of these methods will typically be significant underestimates, which results in the computation leaving its stability region (and to the unstable stepsize behaviour described in Section 2.3.1). As discussed there, these symptoms can be treated by choosing a more appropriate controller, but the root cause can and should also be eliminated by devising a better, *dynamic* error model. Such models have been designed for Runge–Kutta methods and are described in [209]. No dynamic error models have been created for iterative solvers based on the WKB approximation, nor have their stability properties been explored. Although challenging, this would be a huge step for the field. Adaptive stepsize algorithms based on control theory have been brought to my attention by Imre Fekete and Gustaf Söderlind with whom I will explore this problem further.

### Miscellaneous optimisations

I identified additional places in `oscode` where efficiency gains can be made: in the RK-WKB switching mechanism and in the choice of Runge–Kutta method used. At the moment, `oscode` will always attempt both a WKB and RK step of the same stepsize at every iteration, regardless of how oscillatory the solution is. This goes against the philosophy of exploiting information about the solution and is inefficient. The effort of computing an entire RK or WKB step can be saved by implementing a few preliminary checks at each step to check whether the solution is oscillatory or slowly changing, and only attempting the relevant step. For example, the

algorithm could compute the expected number of oscillations in the proposed step as  $\approx \omega h / (2\pi)$  or  $\approx h\dot{x}/x$ , then only attempt a WKB step if this value is above a threshold (and only attempt a RK step if it is below it). The Runge–Kutta solver `oscode` uses is built such that the two steps together use minimal evaluations of the right-hand-side of the ODE, but as such it has less-than-optimal stability properties and error coefficients. This is apparent from Fig. 6.4, where BINGO’s built-in Runge–Kutta method – a highly optimised (4,5) pair – takes larger strides in the slowly-varying region of the solution. `oscode` is still more efficient overall, but with the above checks implemented it would be less important for the WKB and RK steps to have as many function evaluations in common as possible as only one step would take place most of the time. Thus the priority would shift towards using a more optimal Runge–Kutta solver in regions of slow variation. Testing will be necessary to gauge how large a speedup can be achieved with this change.

### Generalising to coupled oscillators

Currently, a computationally limiting step in cosmological inference is the calculation of transfer functions via the Einstein–Boltzmann equations. This involves solving a set of coupled linear ODEs with time-dependent coefficients derived from the cosmological background. The coupled set of equations describes the evolution of perturbations (in matter species and the metric), and can be written as

$$\dot{\mathbf{x}}(t) = \mathbf{A}(t)\mathbf{x}(t), \quad (8.2)$$

where  $\mathbf{x} \in C^1(I, X)$  is a vector of dynamical variables in the  $N$ -dimensional vector space  $X$ , and  $\mathbf{A} : I \times X \mapsto X$  is an  $N$  by  $N$  matrix. The system presents a challenge to solve numerically: physical processes on multiple timescales are at play which makes the system stiff, with some components of  $\mathbf{x}$  being highly oscillatory and some varying slowly. It would be of great interest to generalise the RKWKB approach to deal with coupled oscillators of this form, the Einstein–Boltzmann equations being just one example in physics where more efficient solution of just one system of ODEs would lead to a dramatic reduction in computation time for the entire simulation.

The formal solution of Eq. (8.2) is the *Magnus expansion*, which forms the basis of a substantial portion of efforts to construct an efficient numerical solver for equations of this type (see, e.g. [210] and references therein). Recently, [79] investigated the performance of a new, Magnus expansion-based method (as a multi-oscillator generalisation of the RKWKB method) on the Einstein–Boltzmann equations. The expansion employed by the proposed Magnus–Jordan method does not reduce to the WKB in the case of a single oscillator, but has similar performance in this limit. In higher-dimensional cases however, neither this nor the other Magnus expansion-based methods yield considerable speedups.

Despite the WKB expansion being inherently one-dimensional<sup>1</sup>, it may be possible to use its underlying principles on a set of coupled oscillators if certain assumptions hold true. I have been in communication with Nils Schöneberg to devise a specialised method to solve the Einstein–Boltzmann equations. Under the assumption that all coupled components oscillate with the same frequency but are modulated by different time-dependent amplitudes, we propose the ansatz

$$\mathbf{x}(t) = \Re \left\{ \mathbf{x}_0(t) e^{\int_0^t a(t') dt'} \right\}, \quad (8.3)$$

and expand Eq. (8.2) in growing powers of the small oscillation period  $T$ . Denoting the term  $a \propto T^n$  as  $a^{(n)}$ , the proposed expansion is

$$\mathbf{A}(t) = \mathbf{B}^{(-1)} + \mathbf{B}^{(0)}(t), \quad (8.4)$$

$$\mathbf{x}_0(t) = \sum_{n=0}^{\infty} \mathbf{x}^{(n)}(t) T^n, \quad (8.5)$$

$$a(t) = \sum_{n=-1}^{\infty} a^{(n)}(t) T^n. \quad (8.6)$$

Note that the problems of interest in cosmology have  $\frac{d}{dt} \mathbf{B}^{(-1)}(t) = 0$ , i.e. the dominant,  $\propto 1/T$  part of  $\mathbf{A}$  that drives the oscillation is constant (e.g.  $\propto k$ , where  $k$  is a constant wavenumber). Substituting the above into Eq. (8.2) and matching terms of equal  $T$ -order yields a recursion relation whose initial value is the eigenvalue problem

$$a^{(-1)} \mathbf{x}^{(0)} = \mathbf{B}^{(-1)} \mathbf{x}^{(0)}. \quad (8.7)$$

Let the eigenvalues of  $\mathbf{B}^{(-1)}$  be  $\lambda_i$  and the associated eigenvectors  $\mathbf{v}_i$ , with  $i = 0..N-1$ , where  $N$  is the number of coupled differential equations in Eq. (8.2). The  $\mathbf{x}^{(n)}$  can then be expanded in the basis of eigenvectors. For  $n = 0$ , we have

$$\mathbf{x}^{(0)}(t) = \sum_{j=0}^{N-1} \mathbf{v}_j \alpha_j^{(0)}(t). \quad (8.8)$$

Each of the  $\lambda_i$  corresponds to a ‘fundamental’ frequency  $a^{(-1)} = \lambda_i$ , and the solutions computed for each  $\lambda_i$  are superimposed to form the general solution (due to the linearity of Eq. (8.2)). Focusing on a single solution denoted by the subscript 0, we set

$$\mathbf{x}^{(0)} = \mathbf{v}_0 \alpha_0^{(0)}, \quad (8.9)$$

$$a^{(-1)} = \lambda_0. \quad (8.10)$$

---

<sup>1</sup>In the sense that it applies to the equation of motion of a single oscillator.

Substituting the higher order terms  $\mathbf{x}^{(n)}$  (expanded in the basis of the eigenvectors) back into the recursion relation for  $n \geq 0$ , at each order  $n$  we obtain  $N$  equations for  $N + 1$  parameters. The phase of the complex vector  $\mathbf{x}^{(n)}$  is, however, degenerate with  $a^{(n)}(t)$  (which describes the oscillation frequency and damping at different  $T$ -orders). We can break this degeneracy by e.g. setting the first component of each  $\mathbf{x}^{(n)}(t)$  to one.<sup>2</sup> The  $\alpha_j^{(n)}$  and  $a^{(n)}(t)$  can then be found recursively for higher orders. The resulting expansion reduces to the WKB approximation in the limit of a single oscillator whose equation of motion is  $\ddot{x}(t) + 2\gamma(t)\dot{x}(t) + \omega^2(t)x(t) = 0$ .

Though this approach is theoretically promising, the feasibility of its implementation for large systems needs to be investigated. In particular, the analytic solution of the eigenvalue problem Eq. (8.7) proves challenging for computer algebra software for large  $N$ . This somewhat generalised WKB approximation will have to be implemented iteratively (i.e. re-initialised every so often) to maintain sufficient accuracy, and paired with an error estimate and adaptive stepsize control.

### Fast evaluation of highly oscillatory integrals

Rapid oscillation is ubiquitous in physics and another common form it can take is that of highly oscillatory integrals. It may be possible to use `oscode` to solve a highly oscillatory ODE and use the piecewise-continuous solution it computes for calculating integrals involving the solution of the ODE. When computing dense output, `oscode` essentially uses a piecewise polynomial to represent the slowly changing terms in the WKB expansion, which then appear in the exponent. If one were to integrate the result over a step from  $t$  to  $t + h$ , it would take the form

$$I = \int_t^{t+h} f(t)e^{i\phi(t)} dx, \quad (8.11)$$

where  $f(x)$  and  $\phi(x)$  are polynomials whose degree depends on the interpolation method used to construct dense output, and is up to 7 in `oscode`'s case. If Eq. (8.11) could be computed efficiently, the large stepsizes would ensure that the integral is carried out quickly over the entire interval. In the past decade, significant progress has been made in the development of efficient quadrature methods for integrals of the form

$$I[f] = \int_a^b f(x)e^{i\omega g(x)} dx, \quad (8.12)$$

where  $f, g \in C^\infty$  and  $\omega \gg 1$ . A review by Iserles et al. [48] summarises three new approaches in oscillatory quadrature: asymptotic, *Filon-*, and *Levin-type* methods. These have in common that their accuracy *improves* with increasing frequency  $\omega$ . Asymptotic methods expand the integral in growing inverse powers of  $\omega$ , but are prone to a growth of terms as they do not converge in the usual sense, and their accuracy is predetermined and difficult to tune. *Filon*

<sup>2</sup>We overparametrised the system by allowing  $a(t)$  to be expanded to higher orders.

methods interpolate  $f(x)$  with a polynomial in order to break the integral down to a sum of *moments*,

$$\mu_j(x) = \int x^j e^{i\omega g(x)} dt, \quad j \in \mathbb{Z}_+, \quad (8.13)$$

which have to be computed exactly and tabulated beforehand. This is not necessarily possible (and is not possible for the present problem) and traditional Filon-type methods are of little use for the problem at hand. Olver [49] developed an advanced, *moment-free* Filon method which, as the name suggests, does not rely on the moments in Eq. (8.13) being easy to compute, and is a promising candidate for computing the integral in Eq. (8.11). An alternative might be a Levin-type method, a type of quadrature based on the *collocation* of the entire integrand. Collocation is a type of approximation that becomes exact at a number of pre-determined points or nodes, and may refer to ODE solving or quadrature. In quadrature, the Gauss–Legendre, Gauss–Lobatto, Chebyshev, etc. methods are all examples of collocation. Levin methods exploit that  $I[f]$  is known exactly for functions that satisfy

$$f(x) = F'(x) + i\omega g'(x)F(x), \quad (8.14)$$

since then

$$\int f(x)e^{i\omega g(x)} dx = \int \frac{d}{dx} [F(x)e^{i\omega g(x)}] dx = F(x)e^{i\omega g(x)}. \quad (8.15)$$

Rewriting Eq. (8.14) as  $f = L[f]$ , Levin methods approximate  $F$  as a linear combination of  $v$  basis functions  $\psi_1, \psi_2, \dots, \psi_v$  by collocation with the operator  $L$ . That is, they solve

$$L[v](c_k) = f(c_k), \quad k = 1, 2, \dots, v, \quad (8.16)$$

where the  $c_k$  are appropriately chosen nodes with  $c_1, c_v$  being the endpoints of the integration interval. With some further generalisation, it may become possible to evaluate integrals such as

$$\int_{\eta_i}^{\eta_e} d\eta a^2(\eta) \epsilon_1^2(\eta) \left( f_{\mathbf{k}_1}^*(\eta) f_{\mathbf{k}_2}^*(\eta) f_{\mathbf{k}_3}^*(\eta) + \text{two permutations} \right), \quad (8.17)$$

where the  $f_{\mathbf{k}_i}$  are primordial perturbation modes obtained with `oscode`, and  $a(\eta)$  and  $\epsilon_1(\eta)$  are slowly varying functions that describe the homogeneous cosmological background. Eq. (8.17) describes one of the contributions to the bispectrum (Fourier transform of the three-point function) of scalar primordial perturbations, and the ability to compute it swiftly would be of great significance to cosmology, as discussed further in Section 8.2.

## 8.2 Theoretical results

In Chapter 7, the focus turned from `oscode`-driven numerical evolution of primordial perturbations to their quantum initial conditions. Initial conditions have to satisfy the equation of motion for the Fourier modes of the perturbations, but this in itself does not pin down the

mode functions completely. A choice of vacuum needs to be made, and the perturbations are assumed to start out from this vacuum state. I reviewed the conventional methods for defining the vacuum, including minimising the Hamiltonian density (Hamiltonian diagonalisation), the Danielsson vacuum,  $\alpha$ -vacua, the Bunch–Davies vacuum, and minimising the 00-component of the renormalised stress–energy tensor (RST). I investigated the behaviour of these vacuum choices under canonical transformations of the associated action, arguing that the transformations, like gauge transformations, should not change any of the associated physics including the vacuum state. I verified previous observations [171, 178–181] that Hamiltonian diagonalisation yields a vacuum state which changes under canonical transformations, and observed the same behaviour in all other vacuum choices *except* the one obtained via minimising the RST. I demonstrated what choosing a canonically non-invariant set of initial conditions means for the scalar primordial power spectrum and the CMB temperature and polarisation spectra in a toy inflationary model with reduced number of e-folds of inflation. The model was chosen because the initial conditions are washed out after a sufficiently long period of inflation and do not impact observables. Although a toy model was used for this demonstration, interest in models with fewer e-folds and featureful primordial power spectra has risen as these may give better fits to CMB data, an important example being models with a spatially curved unperturbed metric. The toy model has clearly shown that canonically non-invariant initial conditions result in a wide range of predictions being attainable with a given model, just by picking a different parametrisation of the associated action of primordial perturbations. I argued that quantum initial conditions derived via the minimisation of the RST should therefore be used wherever possible, especially in models which retain memory of the perturbations’ initial conditions.

### 8.2.1 Future work

#### **Renormalised stress–energy tensor initial conditions in spatially curved universes**

In [Chapter 7](#)’s analysis of quantum initial conditions for primordial perturbations, the set of initial conditions derived via minimising the renormalised stress–energy tensor proved to be unambiguous under canonical transformations, and their use was thus suggested in models with limited inflation where the initial conditions lead to features in the primordial power spectrum. The predictive power of these models is therefore tied to how certain we can be of the initial conditions. An obvious example for such models would be closed universes, but the form of the RST initial conditions was derived assuming spatial flatness ( $K = 0$ ). It is therefore only conjectured that the RST initial conditions will retain their canonical invariance on a spatially curved background. Deriving the curved form of these initial conditions and verifying their behaviour under canonical transformations would therefore be of great interest.

Once we can be more certain of initial conditions, the time when they are set, i.e. the moment when the Universe existed in that state could be constrained by observation. At the

very least, one could enquire how much current data can tell us about the time of setting initial conditions, and use this to guide future surveys.

### Reconstructing the initial conditions of primordial perturbations

With a highly efficient solver like `oscode` in hand, one could consider the more computationally challenging task of reconstructing functions of interest, e.g. the quantum initial conditions for primordial perturbations. If the Universe started from kinetic dominance, then [12] points out that the equation of motion for perturbation modes has the closed-form general solution

$$v_k = \frac{1}{2}\sqrt{\pi\eta} \left( A_k H_0^{(1)}(k\eta) + B_k H_0^{(2)}(k\eta) \right), \quad |B_k|^2 - |A_k|^2 = 1, \quad (8.18)$$

where the values of the constants  $A_k$  and  $B_k$  encode the choice of vacuum, and the  $H_i^{(1,2)}$  are the Hankel functions of the first and second kind, of order  $i$ . A node-based, free-form reconstruction [199, 211] of the initial conditions would start from representing the functions  $A_k(k\eta_0)$  and  $B_k(k\eta_0)$  (where the fixed parameter  $\eta_0$  represents the time of setting the initial conditions) as a set of values at  $N$  nodes. These nodes and function values are treated as parameters in a Bayesian inference loop, but will be marginalised over to obtain a *model-independent* reconstruction. The priors from which the node locations and values are drawn can be taken as independently uniform within what would give an observable range, e.g. the range of  $k$  values are related to the multipoles  $\ell$ , and the  $A_k$  and  $B_k$  determine the corresponding  $C_{\ell S}$ , if the reconstruction is based on CMB data. The evolution of the cosmological background and primordial perturbations then needs to be simulated from the (extremely early) time  $\eta_0$  until horizon crossing and freeze-out, after which a Boltzmann solver calculates the relevant observables via transfer functions from the primordial scalar and tensor curvature perturbations. Besides CMB intensity and polarisation maps, large-scale structure data from galaxy surveys and from the Lyman- $\alpha$  forest could also be added to further tighten the reconstruction.

### Axion production via vacuum realignment

In particle physics, an *axion* is a type of light particle represented by a scalar field. While this work exploits the fact that the gauge-invariant curvature perturbation  $\mathcal{R}$  is constant on super-horizon scales to *avoid* dealing with the unknown physics of reheating, in particle physics this phase is very much of interest. It therefore becomes necessary to consider the acceleration equation for the axion  $\theta$  with decay constant  $f_A$ ,

$$\ddot{\theta} + 3H\dot{\theta} + \frac{V'(\theta)}{f_A^2} = 0, \quad (8.19)$$

outside the slow-roll ( $\ddot{\theta} \approx 0$ ) limit. The region around an extremum of the potential  $V(\theta)$  can be always modelled as a parabola, and Eq. (8.19) becomes the equation of motion of a harmonic

oscillator with some time-dependent frequency and damping. This mechanism is called *vacuum realignment* and produces a population of cold axions, which may be candidates for dark matter. As a member of the GAMBIT collaboration, I was made aware that there would be interest in producing efficient reference code for the realignment calculation. The quantity of interest would be the average energy density associated with the axion for which one might also wish to average over initial conditions (for the field  $\theta$ ), therefore Eq. (8.19) will need to be solved repeatedly. Due to the slow numerical solution of Eq. (8.19) with conventional methods, analytic approximations are preferentially used, which only allow for specific potentials. With `oscode`, it would be possible to speed up the realignment calculation in the oscillatory regimes and thus allow the use of a much wider range of potentials, realistic axion mass models, and uncertainties.

### Investigating primordial non-Gaussianities

As introduced in Section 8.1.1, a fast solver for carrying out highly oscillatory integrals would have significant impact in cosmology. An example I bring here is primordial non-Gaussianities. The integral in Eq. (8.17) describes a contribution to the scalar primordial bispectrum, which is predicted to be non-zero if the initial distribution of primordial perturbations is non-Gaussian. As discussed in Section 5.5.1, most inflationary models predict some degree of non-Gaussianity, with predictions of models with featureful primordial power spectra being the largest. The observable, dimensionless parameter  $f_{\text{NL}}$  describing the amplitude of the bispectrum and thus the degree of non-Gaussianity is expected to be better constrained by future experiments [212] and is extremely powerful in constraining mechanisms of inflation. The measurement of non-Gaussianity is computationally and statistically complex, partly due to the need for evaluating highly oscillatory integrals such as Section 8.1.1 for all the different contributions [190, 213–216] to the scalar bispectrum. The development of an efficient quadrature routine built on `oscode` (which could simultaneously compute the  $f_{\mathbf{k}_i}(\eta)$  in Section 8.1.1) could be followed by its integration into the processing pipelines of future experiments (such as the Simons Observatory [217]) looking to better constrain  $f_{\text{NL}}$  and inflationary models.

# References

- [1] 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”. *Phys. Rev. Research* 2 (1 2020), p. 013030 (cit. on pp. [v](#), [2](#), [22](#), [56](#), [57](#), [63](#), [65](#), [131](#)).
- [2] F. J. Agocs. “(py)oscode: fast solutions of oscillatory ODEs”. *Journal of Open Source Software* 5.56 (2020), p. 2830 (cit. on pp. [v](#), [2](#)).
- [3] 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 pp. [v](#), [2](#)).
- [4] F. J. Agocs, L. T. Hergt, W. J. Handley, A. N. Lasenby, and M. P. Hobson. “Quantum initial conditions for inflation and canonical invariance”. *Phys. Rev. D* 102 (2 2020), p. 023507 (cit. on pp. [v](#), [2](#)).
- [5] D. Scott. “The Standard Model of Cosmology: A Skeptic’s Guide”. *arXiv e-prints* (2018). arXiv: [1804.01318](#) (cit. on pp. [2](#), [69](#), [95](#), [97](#)).
- [6] L. Perivolaropoulos and F. Skara. *Challenges for  $\Lambda$ CDM: An update*. 2021. arXiv: [2105.05208](#) (cit. on p. [2](#)).
- [7] N. Schöneberg, G. F. Abellán, A. P. Sánchez, S. J. Witte, c Vivian Poulin, and J. Lesgourgues. *The  $H_0$  Olympics: A fair ranking of proposed models*. 2021. arXiv: [2107.10291](#) (cit. on p. [2](#)).
- [8] E. Di Valentino, O. Mena, S. Pan, L. Visinelli, W. Yang, A. Melchiorri, D. F. Mota, A. G. Riess, and J. Silk. “In the realm of the Hubble tension - a review of solutions”. *Classical and Quantum Gravity* (2021) (cit. on pp. [2](#), [79](#), [97](#)).
- [9] E. Di Valentino, A. Melchiorri, and J. Silk. “Planck evidence for a closed Universe and a possible crisis for cosmology”. *Nature Astronomy* 4.2 (2019), pp. 196–203 (cit. on pp. [2](#), [95](#), [107](#)).
- [10] E. Di Valentino, A. Melchiorri, and J. Silk. “Investigating Cosmic Discordance”. *The Astrophysical Journal Letters* 908.1 (2021), p. L9 (cit. on pp. [2](#), [95](#), [107](#)).
- [11] W. Handley. “Curvature tension: Evidence for a closed universe”. *Phys. Rev. D* 103 (4 2021), p. L041301 (cit. on pp. [2](#), [107](#)).
- [12] W. J. Handley, A. N. Lasenby, and M. P. Hobson. “Novel quantum initial conditions for inflation”. *Physics Review D* 94.2 (2016). arXiv: [1607.04148](#) (cit. on pp. [2](#), [101](#), [104](#), [112](#), [121](#), [138](#), [153](#)).
- [13] L. Euler. “De integratione aequationum differentialium per approximationem”. *Opera Omnia* 11 (1913), p. 424 (cit. on p. [8](#)).
- [14] A. Iserles and G. Söderlind. “Global bounds on numerical error for ordinary differential equations”. *Journal of Complexity* 9.1 (1993), pp. 97–112 (cit. on p. [9](#)).

- [15] J. C. Butcher. “[Numerical Differential Equation Methods](#)”. John Wiley & Sons, Ltd, 2008. Chap. 2, pp. 51–135 (cit. on p. 11).
- [16] F. Bashforth and J. C. Adams. “An attempt to test the theories of capillary action by comparing the theoretical and measured forms of drops of fluid”. Cambridge University Press, 1883 (cit. on p. 12).
- [17] F. R. Moulton. “New methods in exterior ballistics”. Chicago, 1926 (cit. on p. 12).
- [18] E. Nyström. “[Ueber die numerische Integration von Differentialgleichungen](#)”. Acta Societatis Scientiarum Fennicae. 1925 (cit. on p. 12).
- [19] W. E. Milne. “[Numerical Integration of Ordinary Differential Equations](#)”. *The American Mathematical Monthly* 33.9 (1926), pp. 455–460 (cit. on p. 12).
- [20] W. E. Milne and W. Milne. “Numerical solution of differential equations”. Vol. 64. Wiley New York, 1953 (cit. on p. 12).
- [21] The Numerical Algorithms Group (NAG), Oxford, United Kingdom. *The NAG Library* (cit. on pp. 12, 14, 39).
- [22] *Scipy documentation* (cit. on pp. 12, 14).
- [23] H. H. Rosenbrock. “[Some general implicit processes for the numerical solution of differential equations](#)”. *The Computer Journal* 5.4 (1963), pp. 329–330 (cit. on p. 12).
- [24] C. Runge. “[Ueber die numerische Auflösung von Differentialgleichungen](#)”. *Mathematische Annalen* 46 (2 1895), pp. 167–178 (cit. on p. 12).
- [25] K. Heun. “Neue Methoden zur approximativen Integration der Differentialgleichungen einer unabhängigen Veränderlichen”. *Z. Math. Phys* 45 (1900), pp. 23–38 (cit. on p. 12).
- [26] W. Kutta. “Beitrag zur näherungsweise integration totaler differentialgleichungen”. *Z. Math. Phys.* 46 (1901), pp. 435–453 (cit. on p. 12).
- [27] A. Huta. “Une amélioration de la méthode de Runge-Kutta-Nyström pour la résolution numérique des équations différentielles du premier ordre”. *Acta Math. Univ. Comenian* 1 (1956), pp. 201–224 (cit. on p. 12).
- [28] A. Huta. “Contribution a la formule de sixieme ordre dans la methode de Runge-Kutta-Nystrom”. *Acta Math. Univ. Comenian* 2 (1957), pp. 21–24 (cit. on p. 12).
- [29] J. C. Butcher. “Coefficients for the study of Runge-Kutta integration processes”. *Journal of the Australian Mathematical Society* 3.2 (1963), pp. 185–201 (cit. on p. 12).
- [30] E. Hairer, S. Nørsett, and G. Wanner. “[Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems](#)”. Springer Series in Computational Mathematics. Springer, 1993 (cit. on p. 13).
- [31] J. C. Butcher. “[The numerical analysis of ordinary differential equations : Runge-Kutta and general linear methods](#)”. Chichester: Wiley, 1987 (cit. on pp. 14, 39).
- [32] L. F. Richardson and J. A. Gaunt. “VIII. The deferred approach to the limit”. *Philosophical Transactions of the Royal Society of London. Series A, containing papers of a mathematical or physical character* 226.636-646 (1927), pp. 299–361 (cit. on p. 14).
- [33] R. H. Merson. “An operational method for the study of integration processes”. *Proc. Symp. Data Processing*. Vol. 1. Weapons Research Establishment. 1957, p. 25 (cit. on p. 15).
- [34] E. Fehlberg. “Eine Methode zur Fehlerverkleinerung beim Runge-Kutta-Verfahren”. *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik* 38.11-12 (1958), pp. 421–426 (cit. on p. 15).

- [35] E. Fehlberg. “New high-order Runge-Kutta formulas with step size control for systems of first-and second-order differential equations”. *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik* 44.S1 S1 (1964), T17–T29 (cit. on p. 15).
- [36] E. Fehlberg. “Classical fifth-, sixth-, seventh-, and eighth-order Runge-Kutta formulas with stepsize”. *NASA Tech Rep R-287*, NASA (1968) (cit. on p. 15).
- [37] E. Fehlberg. “Low-order classical Runge-Kutta formulas with step size control and their application to some heat transfer problems.” *NASA Tech Rep R-315*, NASA (1969) (cit. on p. 15).
- [38] J. Dormand and P. Prince. “A family of embedded Runge-Kutta formulae”. *Journal of Computational and Applied Mathematics* 6.1 (1980), pp. 19–26 (cit. on p. 15).
- [39] L. F. Shampine. “Some Practical Runge-Kutta Formulas”. *Mathematics of Computation* 46.173 (1986), pp. 135–150 (cit. on pp. 15, 56, 62).
- [40] L. Stoller and D. Morrison. “A method for the numerical integration of ordinary differential equations”. *Mathematical Tables and Other Aids to Computation* 12.64 (1958), pp. 269–272 (cit. on p. 15).
- [41] M. Abramowitz. “Handbook of mathematical functions, with formulas, graphs, and mathematical tables.” New York: Dover, 1965, p. 888 (cit. on p. 16).
- [42] F. Ceschino and J. Kuntzmann. “Numerical solution of initial value problems”. Prentice-Hall, 1966 (cit. on p. 16).
- [43] “Highly Oscillatory Problems”. London Mathematical Society Lecture Note Series. Cambridge University Press, 2009 (cit. on p. 16).
- [44] L. R. Petzold, L. O. Jay, and J. Yen. “Numerical solution of highly oscillatory ordinary differential equations”. *Acta Numerica* 6 (1997), pp. 437–483 (cit. on p. 16).
- [45] A. Iserles. “Think globally, act locally: Solving highly-oscillatory ordinary differential equations”. *Applied Numerical Mathematics* 43.1 (2002). 19th Dundee Biennial Conference on Numerical Analysis, pp. 145–160 (cit. on pp. 16, 17).
- [46] A. Iserles. “Three stories of high oscillation”. *Bull. EMS* 87 (2013), pp. 18–23 (cit. on pp. 16, 17).
- [47] A. Arnold, N. B. Abdallah, and C. Negulescu. “WKB-Based Schemes for the Oscillatory 1D Schrödinger Equation in the Semiclassical Limit”. *SIAM Journal on Numerical Analysis* 49.4 (2011), pp. 1436–1460 (cit. on p. 18).
- [48] A. Iserles, S. Nørsett, and S. Olver. “Highly oscillatory quadrature: The story so far”. *Numerical mathematics and advanced applications*. Springer, 2006, pp. 97–118 (cit. on pp. 19, 150).
- [49] S. Olver. “Moment-free numerical integration of highly oscillatory functions”. *IMA Journal of Numerical Analysis* 26.2 (2006), pp. 213–227 (cit. on pp. 19, 151).
- [50] J. Körner, A. Arnold, and K. Döpfner. *WKB-based scheme with adaptive step size control for the Schrödinger equation in the highly oscillatory regime*. 2021. arXiv: 2102.03107 (cit. on pp. 19, 22).
- [51] D. Mace and L. Thomas. “An extrapolation formula for stepping the calculation of the orbit of an artificial satellite several revolutions time”. *The Astronomical Journal* 65 (1960), p. 300 (cit. on p. 19).
- [52] G. Taratynova. “Methods for the Numerical Solution of Finite Difference Equations and their Application to Computations of Satellite Orbits”. *Artificial Earth Satellites* 4 (1960), pp. 254–254 (cit. on p. 19).

- [53] O. Graf and D. Bettis. “Modified multirevolution integration methods for satellite orbit computation”. *Celestial Mechanics* 11.4 (1975), pp. 433–448 (cit. on p. 19).
- [54] L. R. Petzold. “An Efficient Numerical Method for Highly Oscillatory Ordinary Differential Equations”. *SIAM Journal on Numerical Analysis* 18.3 (1981), pp. 455–479 (cit. on pp. 19, 20, 22, 56).
- [55] E. Hairer, S. P. Norsett, and G. Wanner. “Solving Ordinary Differential Equations I”. Vol. 1. Springer Series in Computational Mathematics. Springer Berlin Heidelberg, 1993 (cit. on pp. 20, 24).
- [56] Z. Heitman, J. Bremer, and V. Rokhlin. “On the existence of nonoscillatory phase functions for second order ordinary differential equations in the high-frequency regime”. *Journal of Computational Physics* 290 (2015), pp. 1–27 (cit. on pp. 20, 147).
- [57] V. R. James Bremer. “Improved estimates for nonoscillatory phase functions”. *Discrete & Continuous Dynamical Systems* 36.8 (2016), pp. 4101–4131 (cit. on pp. 20, 147).
- [58] J. Bremer. “On the numerical solution of second order ordinary differential equations in the high-frequency regime”. *Applied and Computational Harmonic Analysis* 44.2 (2018), pp. 312–349 (cit. on pp. 20, 22, 53, 56).
- [59] A. Olde Daalhuis and F. Olver. “Hyperasymptotic solutions of second-order linear differential equations I”. *Methods and Applications of Analysis* 2.2 (1995), pp. 173–197 (cit. on p. 21).
- [60] A. Olde Daalhuis. “Hyperasymptotic solutions of second-order linear differential equations II”. *Methods and Applications of Analysis* 2.2 (1995), pp. 198–211 (cit. on p. 21).
- [61] A. Dutt, L. Greengard, and V. Rokhlin. “Spectral deferred correction methods for ordinary differential equations”. *BIT Numerical Mathematics* 40.2 (2000), pp. 241–266 (cit. on p. 21).
- [62] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. “Numerical Recipes 3rd Edition: The Art of Scientific Computing”. 3rd ed. USA: Cambridge University Press, 2007 (cit. on pp. 23, 24, 38, 56).
- [63] G. Söderlind, L. Jay, and M. Calvo. “Stiffness 1952–2012: Sixty years in search of a definition”. *BIT Numerical Mathematics* 55.2 (2015), pp. 531–558 (cit. on p. 22).
- [64] G. Söderlind. “Automatic control and adaptive time-stepping”. *Numerical Algorithms* 31.1 (2002), pp. 281–310 (cit. on pp. 25, 27).
- [65] G. Söderlind. “Digital filters in adaptive time-stepping”. *ACM Transactions on Mathematical Software (TOMS)* 29.1 (2003), pp. 1–26 (cit. on p. 25).
- [66] G. Söderlind. “Time-step selection algorithms: Adaptivity, control, and signal processing”. *Applied numerical mathematics* 56.3-4 (2006), pp. 488–502 (cit. on p. 25).
- [67] C. M. Bender and S. A. Orszag. “Advanced mathematical methods for scientists and engineers. I, Asymptotic methods and perturbation theory / Carl M. Bender, Steven A. Orszag.” New York, N.Y.: Springer, 1999 (cit. on pp. 34, 36, 37, 44, 56).
- [68] W. J. Handley, A. N. Lasenby, and M. P. Hobson. “The Runge-Kutta-Wentzel-Kramers-Brillouin Method”. *arXiv e-prints* (2016). arXiv: 1612.02288 (cit. on pp. 34–36, 145).
- [69] K. F. Riley, M. P. Hobson, and S. J. Bence. “Mathematical methods for physics and engineering”. 3rd ed. Cambridge: Cambridge University Press, 2006 (cit. on pp. 37, 56).
- [70] M. Abramowitz. “Handbook of mathematical functions, with formulas, graphs, and mathematical tables.” New York: Dover, 1965 (cit. on pp. 37, 56).
- [71] C. Jordan and K. Jordán. “Calculus of finite differences”. Vol. 33. American Mathematical Soc., 1965 (cit. on pp. 38, 57).

- [72] P. Bogacki and L. Shampine. “An efficient Runge-Kutta (4,5) pair”. *Computers & Mathematics with Applications* 32.6 (1996), pp. 15–28 (cit. on p. 39).
- [73] R. Brankin, I. Gladwell, and L. F. Shampine. “RKSUITE: A Suite of Explicit Runge-Kutta Codes”. in *Contributions in Numerical Mathematics*. Singapore: World, 1993, pp. 41–53 (cit. on p. 39).
- [74] W. Gander and J. Hřebíček. “Solving problems in scientific computing using Maple and MATLAB”. 4th, expanded and rev. ed. Berlin ; London: Springer, 2004 (cit. on p. 40).
- [75] B. V. Noumerov. “A Method of Extrapolation of Perturbations”. *Monthly Notices of the Royal Astronomical Society* 84.8 (1924), pp. 592–602 (cit. on p. 49).
- [76] J. Killingbeck. “Shooting methods for the Schrodinger equation”. *Journal of Physics A: Mathematical and General* 20.6 (1987), pp. 1411–1417 (cit. on p. 50).
- [77] K. Banerjee, S. P. Bhatnagar, V. Choudhry, and S. S. Kanwal. “The Anharmonic Oscillator”. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* 360.1703 (1978), pp. 575–586 (cit. on pp. 50, 52).
- [78] W. I. J. Haddadin and W. J. Handley. “Rapid numerical solutions for the Mukhanov-Sasaki equation”. *arXiv e-prints* (2018). arXiv: 1809.11095 (cit. on pp. 56, 105, 110).
- [79] J. Bamber and W. Handley. “Beyond the Runge-Kutta-Wentzel-Kramers-Brillouin method”. *Physical Review D* 101.4 (2020) (cit. on pp. 56, 148).
- [80] J. Stoer and B. Roland. “Introduction to numerical analysis”. New York: Springer-Verlag, 1983 (cit. on pp. 59, 60).
- [81] M. K. Horn. “Fourth- and Fifth-Order, Scaled Runge-Kutta Algorithms for Treating Dense Output”. *SIAM Journal on Numerical Analysis* 20.3 (1983), pp. 558–568 (cit. on p. 62).
- [82] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. Jarrod Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. Carey, Í. Polat, Y. Feng, E. W. Moore, J. Vand erPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and S. 1. 0. Contributors. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. *Nature Methods* 17 (2020), pp. 261–272 (cit. on p. 64).
- [83] V. Mukhanov and S. Winitzki. “Introduction to Quantum Effects in Gravity”. Cambridge University Press, 2007 (cit. on p. 69).
- [84] S. Dodelson. “Modern cosmology”. San Diego, Calif.: Academic Press, 2003 (cit. on pp. 69, 90, 99).
- [85] M. P. Hobson, G. P. Efstathiou, and A. N. Lasenby. “General Relativity: An Introduction for Physicists”. Cambridge University Press, 2006 (cit. on p. 69).
- [86] N. D. Birrell and P. C. W. Davies. “Quantum fields in curved space / N.D. Birrell and P.C.W. Davies.” Cambridge monographs on mathematical physics. Cambridge University Press, 1982 (cit. on pp. 69, 100, 121, 127, 129).
- [87] S. A. Fulling. “Aspects of Quantum Field Theory in Curved Spacetime”. London Mathematical Society Student Texts. Cambridge University Press, 1989 (cit. on p. 69).
- [88] V. Mukhanov. “Physical Foundations of Cosmology.” Cambridge University Press, 2005 (cit. on p. 69).
- [89] D. Baumann. “TASI Lectures on Inflation”. *arXiv e-prints* (2009). arXiv: 0907.5424 (cit. on pp. 69, 99, 116).

- [90] V. Mukhanov, H. Feldman, and R. Brandenberger. “Theory of cosmological perturbations”. *Physics Reports* 215.5 (1992), pp. 203–333 (cit. on pp. 69, 85, 87, 99, 116).
- [91] J. M. Stewart. “Perturbations of Friedmann-Robertson-Walker cosmological models”. *Classical and Quantum Gravity* 7.7 (1990), pp. 1169–1180 (cit. on pp. 69, 84, 85).
- [92] A. Friedmann. “Über die krümmung des raumes”. *Zeitschrift für Physik* 10 (1922), pp. 377–386 (cit. on p. 71).
- [93] G. Lemaître. “L’Univers en expansion”. *Annales de la Société Scientifique de Bruxelles* 53 (1933), p. 51 (cit. on p. 71).
- [94] H. P. Robertson. “Kinematics and World-Structure”. *The Astrophysical Journal* 82 (1935), p. 284 (cit. on p. 71).
- [95] H. P. Robertson. “Kinematics and World-Structure III.” *The Astrophysical Journal* 83 (1936), p. 257 (cit. on p. 71).
- [96] H. P. Robertson. “Kinematics and World-Structure III.” *The Astrophysical Journal* 83 (1936), p. 257 (cit. on p. 71).
- [97] A. A. Penzias and R. W. Wilson. “A measurement of excess antenna temperature at 4080 Mc/s.” *The Astrophysical Journal* 142 (1965), pp. 419–421 (cit. on p. 76).
- [98] R. H. Dicke, P. J. E. Peebles, P. G. Roll, and D. T. Wilkinson. “Cosmic black-body radiation.” *The Astrophysical Journal* 142 (1965), pp. 414–419 (cit. on p. 76).
- [99] G. F. Smoot, C. L. Bennett, A. Kogut, E. Wright, J. Aymon, N. Boggess, E. Cheng, G. De Amici, S. Gulkis, M. Hauser, et al. “Structure in the COBE differential microwave radiometer first-year maps”. *The Astrophysical Journal* 396 (1992), pp. L1–L5 (cit. on p. 76).
- [100] M. H. Abitbol, J. C. Hill, and J. Chluba. “Measuring the Hubble Constant from the Cooling of the CMB Monopole”. *The Astrophysical Journal* 893.1 (2020), p. 18 (cit. on p. 76).
- [101] R. J. Adler. “The geometry of random fields”. SIAM, 2010 (cit. on p. 77).
- [102] R. J. Thornton, P. A. R. Ade, S. Aiola, F. E. Angilè, M. Amiri, J. A. Beall, D. T. Becker, H.-M. Cho, S. K. Choi, P. Corlies, and et al. “The Atacama cosmology telescope: The polarization-sensitive ACTPol instrument”. *The Astrophysical Journal Supplement Series* 227.2 (2016), p. 21 (cit. on p. 78).
- [103] J. E. Austermann et al. “SPTpol: an instrument for CMB polarization measurements with the South Pole Telescope”. Ed. by W. S. Holland. Vol. 8452. International Society for Optics and Photonics. SPIE, 2012, pp. 393–410 (cit. on p. 78).
- [104] J. A. Peacock. “Implications of 2dFGRS results on cosmic structure”. *AIP Conf. Proc.* 666.1 (2003). Ed. by S. S. Holt and C. S. Reynolds, pp. 275–290. arXiv: astro-ph/0301042 (cit. on p. 78).
- [105] *The Sloan Digital Sky Survey: Mapping the Universe* (cit. on p. 78).
- [106] A. G. Riess, S. Casertano, W. Yuan, J. B. Bowers, L. Macri, J. C. Zinn, and D. Scolnic. “Cosmic Distances Calibrated to 1% Precision with Gaia EDR3 Parallaxes and Hubble Space Telescope Photometry of 75 Milky Way Cepheids Confirm Tension with  $\Lambda$ CDM”. *The Astrophysical Journal* 908.1 (2021), p. L6 (cit. on p. 79).
- [107] W. L. Freedman, B. F. Madore, D. Hatt, T. J. Hoyt, I. S. Jang, R. L. Beaton, C. R. Burns, M. G. Lee, A. J. Monson, J. R. Neeley, M. M. Phillips, J. A. Rich, and M. 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 (cit. on p. 79).

- [108] W. L. Freedman, B. F. Madore, T. Hoyt, I. S. Jang, R. Beaton, M. G. Lee, A. Monson, J. Neeley, and J. Rich. “Calibration of the Tip of the Red Giant Branch”. *The Astrophysical Journal* 891.1 (2020), p. 57 (cit. on p. 79).
- [109] 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 p. 80).
- [110] A. H. Guth. “Inflationary universe: A possible solution to the horizon and flatness problems”. *Phys. Rev. D* 23 (2 1981), pp. 347–356 (cit. on p. 80).
- [111] A. 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 p. 80).
- [112] A. Albrecht and P. J. Steinhardt. “Cosmology for Grand Unified Theories with Radiatively Induced Symmetry Breaking”. *Phys. Rev. Lett.* 48 (17 1982), pp. 1220–1223 (cit. on p. 80).
- [113] R. H. Brandenberger. “Inflationary cosmology: Progress and problems”. *IPM School on Cosmology 1999: Large Scale Structure Formation*. 1999. arXiv: [hep-ph/9910410](https://arxiv.org/abs/hep-ph/9910410) (cit. on p. 80).
- [114] R. H. Brandenberger. “Introduction to Early Universe Cosmology”. *Proceedings of 4th International Conference on Fundamental Interactions — PoS(ICFI 2010)*. Vol. 124. 2011, p. 001 (cit. on p. 80).
- [115] W. J. Handley, S. D. Brechet, A. N. Lasenby, and M. P. Hobson. “Kinetic initial conditions for inflation”. *Physical Review D* 89.6 (2014). arXiv: [1401.2253](https://arxiv.org/abs/1401.2253) (cit. on pp. 82, 100, 118, 131).
- [116] C. R. Contaldi, M. Peloso, L. Kofman, and A. Linde. “Suppressing the lower multipoles in the CMB anisotropies”. *Journal of Cosmology and Astro-Particle Physics* 2003 (2003). arXiv: [astro-ph/0303636](https://arxiv.org/abs/astro-ph/0303636) (cit. on pp. 82, 100).
- [117] L. T. Hergt, W. J. Handley, M. P. Hobson, and A. N. Lasenby. “Case for kinetically dominated initial conditions for inflation”. *Phys. Rev. D* 100 (2 2019), p. 023502 (cit. on pp. 82, 100, 104, 131).
- [118] L. T. Hergt, W. J. Handley, M. P. Hobson, and A. N. Lasenby. “Constraining the kinetically dominated universe”. *Phys. Rev. D* 100 (2 2019), p. 023501 (cit. on pp. 82, 100, 104, 131).
- [119] J. W. York Jr. “Covariant decompositions of symmetric tensors in the theory of gravitation”. *Annales de l’IHP Physique théorique*. Vol. 21. 4. 1974, pp. 319–332 (cit. on p. 85).
- [120] J. M. Bardeen. “Gauge-invariant cosmological perturbations”. *Phys. Rev. D* 22 (8 1980), pp. 1882–1905 (cit. on p. 86).
- [121] V. F. Mukhanov and G. V. Chibisov. “Quantum Fluctuations and a Nonsingular Universe”. *JETP Lett.* 33 (1981), pp. 532–535 (cit. on p. 87).
- [122] S. W. Hawking. “The development of irregularities in a single bubble inflationary universe”. *Physics Letters B* 115.4 (1982), pp. 295–297 (cit. on p. 87).
- [123] A. A. Starobinsky. “Dynamics of phase transition in the new inflationary universe scenario and generation of perturbations”. *Physics Letters B* 117.3 (1982), pp. 175–178 (cit. on p. 87).
- [124] A. H. Guth and S.-Y. Pi. “Fluctuations in the New Inflationary Universe”. *Phys. Rev. Lett.* 49 (15 1982), pp. 1110–1113 (cit. on p. 87).

- [125] V. F. Mukhanov, L. A. Kofman, and D. Y. Pogosian. “Cosmological Perturbations in the Inflationary Universe”. *Phys. Lett. B* 193 (1987), pp. 427–432 (cit. on p. 87).
- [126] J. Lesgourgues. “The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview”. *arXiv e-prints* (2011). arXiv: 1104.2932 (cit. on pp. 90, 136).
- [127] D. Blas, J. Lesgourgues, and T. Tram. “The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes”. *Journal of Cosmology and Astroparticle Physics* 2011.07 (2011), pp. 034–034 (cit. on pp. 90, 136).
- [128] J. Lesgourgues and T. Tram. “The Cosmic Linear Anisotropy Solving System (CLASS) IV: efficient implementation of non-cold relics”. *Journal of Cosmology and Astroparticle Physics* 2011.9 (2011). arXiv: 1104.2935 (cit. on pp. 90, 136).
- [129] T. Tram and J. Lesgourgues. “Optimal polarisation equations in FLRW universes”. *Journal of Cosmology and Astroparticle Physics* 2013.10 (2013). arXiv: 1305.3261 (cit. on pp. 90, 136).
- [130] A. Lewis, A. Challinor, and A. Lasenby. “Efficient computation of CMB anisotropies in closed FRW models”. *Astrophys. J.* 538 (2000), pp. 473–476. arXiv: astro-ph/9911177 (cit. on p. 90).
- [131] L. T. Hergt, W. J. Handley, M. P. Hobson, and A. N. Lasenby. “Bayesian evidence for the tensor-to-scalar ratio  $r$  and neutrino masses  $m_\mu$ : Effects of uniform versus logarithmic priors”. *Physical Review D* 103.12 (2021) (cit. on p. 93).
- [132] W. Handley and P. Lemos. “Quantifying dimensionality: Bayesian cosmological model complexities”. *Physical Review D* 100.2 (2019) (cit. on p. 93).
- [133] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. “emcee: The MCMC Hammer”. *Publications of the Astronomical Society of the Pacific* 125.925 (2013), p. 306. arXiv: 1202.3665 (cit. on p. 93).
- [134] A. Lewis and S. Bridle. “Cosmological parameters from CMB and other data: A Monte Carlo approach”. *Physical Review D* 66.10 (2002) (cit. on p. 93).
- [135] 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 (cit. on p. 93).
- [136] 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. 93, 137).
- [137] 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. 93, 137).
- [138] G. Efstathiou, W. J. Sutherland, and S. Maddox. “The cosmological constant and cold dark matter”. *Nature* 348.6303 (1990), pp. 705–707 (cit. on p. 94).
- [139] A. G. Kim, S. Gabi, G. Goldhaber, D. E. Groom, I. M. Hook, M. Y. Kim, J. C. Lee, C. R. Pennypacker, S. Perlmutter, I. A. Small, and et al. “Implications for the Hubble Constant from the First Seven Supernovae at  $z \geq 0.35$ ”. *The Astrophysical Journal* 476.2 (1997), pp. L63–L66 (cit. on p. 94).
- [140] W. Handley. “Curvature tension: Evidence for a closed universe”. *Physical Review D* 103.4 (2021) (cit. on p. 95).
- [141] S. Tsujikawa. “Quintessence: a review”. *Classical and Quantum Gravity* 30.21 (2013), p. 214003 (cit. on p. 95).

- [142] L. Husdal. “On Effective Degrees of Freedom in the Early Universe”. *Galaxies* 4.4 (2016) (cit. on p. 95).
- [143] R. G. Crittenden and N. Turok. “Doppler Peaks from Cosmic Texture”. *Phys. Rev. Lett.* 75 (14 1995), pp. 2642–2645 (cit. on p. 96).
- [144] W. Hu and M. White. “Acoustic Signatures in the Cosmic Microwave Background”. *The Astrophysical Journal* 471.1 (1996), pp. 30–51 (cit. on p. 96).
- [145] D. Langlois. “Isocurvature cosmological perturbations and the CMB”. *Comptes Rendus Physique* 4.8 (2003), pp. 953–959 (cit. on p. 96).
- [146] E. Di Valentino, L. A. Anchordoqui, Ö. Akarsu, Y. Ali-Haimoud, L. Amendola, N. Arendse, M. Asgari, M. Ballardini, S. Basilakos, E. Battistelli, and et al. “Snowmass2021 - Letter of interest cosmology intertwined I: Perspectives for the next decade”. *Astroparticle Physics* 131 (2021), p. 102606 (cit. on p. 97).
- [147] E. Di Valentino, L. A. Anchordoqui, Ö. Akarsu, Y. Ali-Haimoud, L. Amendola, N. Arendse, M. Asgari, M. Ballardini, S. Basilakos, E. Battistelli, and et al. “Snowmass2021 - Letter of interest cosmology intertwined I: Perspectives for the next decade”. *Astroparticle Physics* 131 (2021), p. 102606 (cit. on p. 97).
- [148] E. Di Valentino, L. A. Anchordoqui, Ö. Akarsu, Y. Ali-Haimoud, L. Amendola, N. Arendse, M. Asgari, M. Ballardini, S. Basilakos, E. Battistelli, and et al. “Cosmology intertwined III:  $f\sigma_8$  and  $S_8$ ”. *Astroparticle Physics* 131 (2021), p. 102604 (cit. on p. 97).
- [149] A. R. Liddle and D. H. Lyth. “Cosmological Inflation and Large-Scale Structure”. 2000, p. 414 (cit. on p. 99).
- [150] 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 (2003), pp. 213–231. arXiv: astro-ph/0302225 (cit. on pp. 99, 107).
- [151] D. K. Hazra, L. Sriramkumar, and J. Martin. “BINGO: a code for the efficient computation of the scalar bi-spectrum”. *Journal of Cosmology and Astro-Particle Physics* 2013 (2013). arXiv: 1201.0926 (cit. on pp. 99, 101).
- [152] D. S. Salopek, J. R. Bond, and J. M. Bardeen. “Designing density fluctuation spectra in inflation”. *Phys. Rev. D* 40 (6 1989), pp. 1753–1788 (cit. on pp. 99, 101).
- [153] M. J. Mortonson, H. V. Peiris, and R. Easther. “Bayesian Analysis of Inflation: Parameter Estimation for Single Field Models”. *Phys. Rev. D* 83 (2011), p. 043505. arXiv: 1007.4205 (cit. on p. 101).
- [154] R. Flauger, L. McAllister, E. Pajer, A. Westphal, and G. Xu. “Oscillations in the CMB from axion monodromy inflation”. *Journal of Cosmology and Astro-Particle Physics* 2010 (2010). arXiv: 0907.2916 (cit. on p. 104).
- [155] B. Allen. “Vacuum states in de Sitter space”. *Phys. Rev. D* 32 (12 1985), pp. 3136–3149 (cit. on pp. 104, 112).
- [156] R. P. Brent. *Algorithms for Minimization without Derivatives, chap. 4*. 1973 (cit. on p. 106).
- [157] W. Handley. “Primordial power spectra for curved inflating universes”. *Phys. Rev. D* 100 (12 2019), p. 123517 (cit. on pp. 106, 122).
- [158] A. Lasenby and C. Doran. “Closed universes, de Sitter space, and inflation”. *Physical Review D* 71.6 (2005). eprint: astro-ph/0307311 (cit. on pp. 106, 113).

- [159] L. T. Hergt, F. J. Agocs, W. J. Handley, M. P. Hobson, and A. N. Lasenby. “Finite inflation in curved space” (2021). in prep. (cit. on pp. 107, 109, 110, 146).
- [160] L. T. Hergt. “Constraining the kinetically dominated Universe: Bayesian methods and primordial cosmology”. pages 79–141. PhD thesis. 2020, pp. 79–141 (cit. on pp. 107, 109, 110).
- [161] 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”. *Astrophysical Journal, Supplement* 208 (2013). arXiv: 1212.5226 (cit. on p. 107).
- [162] J.-P. Uzan, U. Kirchner, and G. 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 pp. 107, 113).
- [163] Planck Collaboration. “Planck 2013 results. XVI. Cosmological parameters”. *Astronomy and Astrophysics* 571 (2014). arXiv: 1303.5076 (cit. on p. 107).
- [164] Planck Collaboration. “Planck 2015 results. XIII. Cosmological parameters”. *Astronomy & Astrophysics* 594 (2016), A13 (cit. on p. 107).
- [165] Planck Collaboration. “Planck 2018 results. VI. Cosmological parameters”. *arXiv e-prints* (2018). arXiv: 1807.06209 (cit. on pp. 107, 134, 136).
- [166] G. Efstathiou and S. 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. 107).
- [167] Planck Collaboration. “Planck 2018 results. V. CMB power spectra and likelihoods” (2019). arXiv: 1907.12875 (cit. on pp. 107, 136).
- [168] BICEP2 and Keck Array Collaborations. “BICEP2 / Keck Array X: 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 (cit. on p. 107).
- [169] Planck Collaboration. “Planck 2018 results. X. Constraints on inflation” (2018). arXiv: 1807.06211 (cit. on pp. 107, 129).
- [170] K. Freese, J. A. Frieman, and A. V. Olinto. “Natural inflation with pseudo Nambu-Goldstone bosons”. *Physical Review Letters* 65.26 (1990), pp. 3233–3236 (cit. on p. 107).
- [171] S. A. Fulling. “Remarks on positive frequency and Hamiltonians in expanding universes”. *General Relativity and Gravitation* 10.10 (1979), pp. 807–824 (cit. on pp. 112, 127, 137, 138, 152).
- [172] T. S. Bunch, P. C. W. Davies, and R. Penrose. “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 pp. 112, 129).
- [173] U. H. Danielsson. “Note on inflation and trans-Planckian physics”. *Phys. Rev. D* 66 (2002), p. 023511 (cit. on pp. 112, 119, 125, 130, 138).
- [174] N. A. Chernikov and E. A. Tagirov. “Quantum theory of scalar field in de Sitter space-time”. *Annales de l’I.H.P. Physique théorique* 9.2 (1968), pp. 109–141 (cit. on p. 112).

- [175] E. Tagirov. “Consequences of field quantization in de Sitter type cosmological models”. *Annals of Physics* 76.2 (1973), pp. 561–579 (cit. on p. 112).
- [176] E. Mottola. “Particle creation in de Sitter space”. *Phys. Rev. D* 31 (4 1985), pp. 754–766 (cit. on p. 112).
- [177] K. Goldstein and D. A. Lowe. “A note on  $\alpha$ -vacua and interacting field theory in de Sitter space”. *Nuclear Physics B* 669.1-2 (2003), pp. 325–340 (cit. on pp. 112, 129).
- [178] A. L. Matacz. “Coherent state representation of quantum fluctuations in the early Universe”. *Phys. Rev. D* 49 (2 1994), pp. 788–798 (cit. on pp. 112, 152).
- [179] V. Bozza, M. Giovannini, and G. Veneziano. “Cosmological perturbations from a new-physics hypersurface”. *Journal of Cosmology and Astroparticle Physics* 2003.05 (2003), pp. 001–001 (cit. on pp. 112, 152).
- [180] J. Grain and V. Venmin. “Squeezing formalism and canonical transformations in cosmology”. *arXiv e-prints* (2019). arXiv: 1910.01916 (cit. on pp. 112–114, 129, 152).
- [181] N. Weiss. “Consistency of Hamiltonian diagonalization for field theories in a Robertson-Walker background”. *Phys. Rev. D* 34 (6 1986), pp. 1768–1775 (cit. on pp. 113, 152).
- [182] B. Bonga, B. Gupta, and N. Yokomizo. “Inflation in the closed FLRW model and the CMB”. *Journal of Cosmology and Astroparticle Physics* 2016.10 (2016), pp. 031–031 (cit. on p. 113).
- [183] G. F. R. Ellis, W. Stoeger, P. McEwan, and P. Dunsby. “Dynamics of Inflationary Universes with Positive Spatial Curvature”. *General Relativity and Gravitation* 34.9 (2002), pp. 1445–1459 (cit. on p. 113).
- [184] G. F. R. Ellis, P. McEwan, W. Stoeger, and P. Dunsby. “Causality in Inflationary Universes with Positive Spatial Curvature”. *General Relativity and Gravitation* 34.9 (2002), pp. 1461–1481 (cit. on p. 113).
- [185] C. Lanczos. “The variational principles of mechanics”. Courier Corporation, 2012 (cit. on p. 115).
- [186] M. G. Calkin. “Lagrangian and Hamiltonian mechanics”. World Scientific, 1996 (cit. on p. 115).
- [187] A. V. Tsyganov. “Canonical transformations of the extended phase space and integrable systems”. *Theoretical and Mathematical Physics* 124.1 (2000), pp. 918–937 (cit. on p. 116).
- [188] W. Greiner. “Extended Hamilton–Lagrange Formalism”. *Classical Mechanics: Systems of Particles and Hamiltonian Dynamics*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 415–453 (cit. on p. 116).
- [189] J. Struckmeier and C. Riedel. “Canonical transformations and exact invariants for time-dependent Hamiltonian systems”. *Annalen der Physik* 11.1 (2002), pp. 15–38 (cit. on p. 116).
- [190] J. Maldacena. “Non-gaussian features of primordial fluctuations in single field inflationary models”. *Journal of High Energy Physics* 2003.05 (2003), pp. 013–013 (cit. on pp. 116, 154).
- [191] F. Lucchin and S. Matarrese. “Power-law inflation”. *Phys. Rev. D* 32 (6 1985), pp. 1316–1322 (cit. on p. 117).
- [192] J. Martin and R. H. Brandenberger. “Trans-Planckian problem of inflationary cosmology”. *Phys. Rev. D* 63 (12 2001), p. 123501 (cit. on pp. 117, 130).

- [193] R. Easther, B. R. Greene, W. H. Kinney, and G. Shiu. “Generic estimate of trans-Planckian modifications to the primordial power spectrum in inflation”. *Phys. Rev. D* 66 (2 2002), p. 023518 (cit. on p. 119).
- [194] D. Polarski and A. A. Starobinsky. “Semiclassicality and decoherence of cosmological perturbations”. *Classical and Quantum Gravity* 13.3 (1996), pp. 377–391 (cit. on p. 119).
- [195] G. Alberghi, R. Casadio, and A. Tronconi. “Trans-Planckian footprints in inflationary cosmology”. *Physics Letters B* 579.1 (2004), pp. 1–5 (cit. on p. 119).
- [196] C. Armendáriz-Picón and E. A. Lim. “Vacuum choices and the predictions of inflation”. *Journal of Cosmology and Astroparticle Physics* 12 (2003). arXiv: hep-th/0303103 (cit. on pp. 125, 130, 137).
- [197] Y. Yargic, L. Sberna, and A. Kempf. *Which part of the stress-energy tensor gravitates?* 2020. arXiv: 2001.04993 (cit. on p. 127).
- [198] D. J. H. Chung, A. Notari, and A. Riotto. “Minimal theoretical uncertainties in inflationary predictions”. *Journal of Cosmology and Astroparticle Physics* 2003.10 (2003), pp. 012–012 (cit. on p. 129).
- [199] W. J. Handley, A. N. Lasenby, H. V. Peiris, and M. P. Hobson. “Bayesian inflationary reconstructions from Planck 2018 data”. *Phys. Rev. D* 100 (10 2019), p. 103511 (cit. on pp. 129, 153).
- [200] Planck Collaboration. “Planck 2015 results. XX. Constraints on inflation”. *Astronomy and Astrophysics* 594 (2016). arXiv: 1502.02114 (cit. on pp. 129, 136).
- [201] A. Friedmann. “On the Curvature of Space”. *General Relativity and Gravitation* 31 (1999), p. 1991 (cit. on p. 130).
- [202] J. Martin and R. Brandenberger. “Dependence of the spectra of fluctuations in inflationary cosmology on trans-Planckian physics”. *Phys. Rev. D* 68 (6 2003), p. 063513 (cit. on p. 130).
- [203] A. R. Liddle and D. H. Lyth. “The cold dark matter density perturbation”. *Physics Reports* 231.1-2 (1993), pp. 1–105. arXiv: 9303019 (cit. on p. 131).
- [204] A. R. Liddle and S. M. Leach. “How long before the end of inflation were observable perturbations produced?” *Physical Review D* 68.10 (2003), p. 103503. arXiv: 0305263 (cit. on p. 131).
- [205] S. Dodelson and L. Hui. “Horizon Ratio Bound for Inflationary Fluctuations”. *Physical Review Letters* 91.13 (2003), p. 131301. arXiv: 0305113 (cit. on p. 131).
- [206] E. Ramirez and D. J. Schwarz. “Predictions of just-enough inflation”. *Phys. Rev. D* 85 (10 2012), p. 103516 (cit. on p. 131).
- [207] E. Ramirez. “Low power on large scales in just-enough inflation models”. *Phys. Rev. D* 85 (10 2012), p. 103517 (cit. on p. 131).
- [208] P. J. Prince and J. R. Dormand. “High order embedded Runge-Kutta formulae”. *Journal of Computational and Applied Mathematics* 7.1 (1981), pp. 67–75 (cit. on p. 146).
- [209] C. Arévalo, G. Söderlind, Y. Hadjimichael, and I. Fekete. “Local error estimation and step size control in adaptive linear multistep methods”. *Numerical Algorithms* 86 (2021), pp. 537–563 (cit. on p. 147).
- [210] S. Blanes, F. Casas, J.-A. Oteo, and J. Ros. “The Magnus expansion and some of its applications”. *Physics reports* 470.5-6 (2009), pp. 151–238 (cit. on p. 148).
- [211] Planck Collaboration. “Planck 2015 results. XX. Constraints on inflation”. *Astronomy and Astrophysics* 594 (2016). arXiv: 1502.02114 (cit. on p. 153).

- 
- [212] P. D. Meerburg, D. Green, R. Flauger, B. Wallisch, M. D. Marsh, E. Pajer, G. Goon, C. Dvorkin, A. M. Dizgah, D. Baumann, G. L. Pimentel, S. Foreman, E. Silverstein, E. Chisari, B. Wandelt, M. Loverde, and A. Slosar. “[Primordial Non-Gaussianity](#)”. *Bulletin of the AAS* 51.3 (2019) (cit. on p. 154).
- [213] X. Chen, R. Easther, and E. A. Lim. “Large non-Gaussianities in single-field inflation”. *Journal of Cosmology and Astroparticle Physics* 2007.06 (2007), p. 023 (cit. on p. 154).
- [214] X. Chen, R. Easther, and E. A. Lim. “Generation and characterization of large non-Gaussianities in single field inflation”. *Journal of Cosmology and Astroparticle Physics* 2008.04 (2008), p. 010 (cit. on p. 154).
- [215] J. Martin and L. Sriramkumar. “The scalar bi-spectrum in the Starobinsky model: The equilateral case”. *Journal of Cosmology and Astroparticle Physics* 2012.01 (2012), p. 008 (cit. on p. 154).
- [216] D. K. Hazra, J. Martin, and L. Sriramkumar. “[Scalar bispectrum during preheating in single field inflationary models](#)”. *Phys. Rev. D* 86 (6 2012), p. 063523 (cit. on p. 154).
- [217] The Simons Observatory Collaboration. “[The Simons Observatory: science goals and forecasts](#)”. *Journal of Cosmology and Astroparticle Physics* 2019.02 (2019), pp. 056–056 (cit. on p. 154).

