

# Investigating the Effect of Primordial Black Hole Hawking Radiation on the Cosmic Microwave Background



THE UNIVERSITY  
*of* ADELAIDE

Harry David Poulter

Principal Supervisor: Dr. Martin White

Co-Supervisor: Prof. Tony Williams

School of Physical Sciences  
University of Adelaide

This thesis is submitted towards the degree of  
*Master of Philosophy (Physics & Astronomy)*

November 2018



# Table of contents

<b>Abstract</b>	<b>vii</b>
<b>Declaration</b>	<b>ix</b>
<b>Acknowledgements</b>	<b>xi</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 Cosmology and the CMB</b>	<b>5</b>
2.1 Introduction . . . . .	5
2.2 Fundamentals of Cosmology . . . . .	5
2.2.1 Cosmological Principle . . . . .	6
2.2.2 Contents of the Universe . . . . .	6
2.3 The $\Lambda$ CDM Model . . . . .	10
2.3.1 Friedmann Equations . . . . .	10
2.3.2 Redshift-Scale Factor Relationship . . . . .	12
2.3.3 Evolution of the Hubble Parameter . . . . .	14
2.3.4 Parametrising the Early Universe . . . . .	17
2.4 The Cosmic Microwave Background . . . . .	18
2.4.1 Mathematical Description . . . . .	19
2.4.2 Temperature Anisotropies . . . . .	22
2.4.3 Polarisation Anisotropies . . . . .	24
2.4.4 Secondary Anisotropies . . . . .	25
<b>3 Recombination and Dark Matter</b>	<b>27</b>
3.1 Introduction . . . . .	27
3.2 The Recombination Epoch . . . . .	27
3.2.1 Imprints on CMB . . . . .	28
3.2.2 Three-Level Atom . . . . .	29
3.2.3 Effective Multi-Level Atom . . . . .	32
3.2.4 Recombination Codes . . . . .	32

3.3	Dark Matter as a Non-Standard Energy Source . . . . .	33
3.3.1	Energy Injection During Recombination . . . . .	34
3.3.2	Extension of TLA Model . . . . .	35
3.4	Primordial Black Holes . . . . .	36
3.4.1	Creation . . . . .	36
3.4.2	Evaporation . . . . .	37
3.4.3	Emission Spectra . . . . .	37
3.4.4	Energy Injection . . . . .	38
3.4.5	Deposition Efficiencies . . . . .	39
<b>4</b>	<b>Sampling <math>\Lambda</math>CDM Parameter Space</b> . . . . .	<b>41</b>
4.1	Introduction . . . . .	41
4.2	MULTINEST . . . . .	41
4.2.1	Bayesian Parameter Estimation . . . . .	41
4.2.2	Nested Sampling . . . . .	43
4.2.3	Prior Handling . . . . .	44
4.2.4	Convergence . . . . .	45
4.3	<i>Planck</i> Likelihood Code . . . . .	45
4.3.1	Map Making . . . . .	46
4.3.2	Likelihoods . . . . .	47
4.3.3	Parameters . . . . .	48
4.4	<code>pc_multinest</code> . . . . .	48
4.4.1	Parameter Handling . . . . .	49
4.4.2	Initialisation . . . . .	52
4.4.3	Likelihood Evaluation . . . . .	54
4.4.4	Comparison to <i>Planck</i> . . . . .	55
<b>5</b>	<b>Constraining Primordial Black Holes</b> . . . . .	<b>61</b>
5.1	Introduction . . . . .	61
5.2	PBH Deposition Efficiencies . . . . .	61
5.2.1	Effective Deposition Efficiencies . . . . .	62
5.2.2	Interpolation . . . . .	64
5.2.3	Cross-Validation . . . . .	66
5.2.4	Results . . . . .	67
5.3	PBH Mass Distributions . . . . .	67
5.3.1	Non-Monochromatic Masses . . . . .	67
5.3.2	Implementation . . . . .	69
5.4	Code Extensions . . . . .	71
5.4.1	CLASS . . . . .	71

---

5.4.2	RECFAST . . . . .	72
5.4.3	HYREC . . . . .	73
5.5	Results . . . . .	78
5.5.1	PBH Effects on Power Spectra . . . . .	79
5.5.2	PBH Effects on Full High- $\ell$ Likelihood . . . . .	81
5.5.3	Parameter Degeneracies . . . . .	82
5.5.4	Verification of Clark et al. Results . . . . .	84
5.5.5	Varying $\Lambda$ CDM Parameters . . . . .	85
5.5.6	Constraints on Distribution Widths . . . . .	86
5.5.7	Constraints on Distribution Types . . . . .	89
<b>6</b>	<b>Conclusions</b>	<b>93</b>
	<b>Appendix A Derivation of Friedmann Equations</b>	<b>95</b>
A.1	FLRW Metric . . . . .	95
A.2	Einstein Field Equations . . . . .	96
	<b>References</b>	<b>101</b>



## Abstract

Due in part to the recent gravitational wave events discovered by the LIGO collaboration, primordial black holes (PBHs) have made their way back in vogue as a possible dark matter candidate. Dynamically, PBHs are interesting objects as they have a mass-dependent energy injection mechanism. As such, their presence in the early universe would have interfered with the formation of the cosmic microwave background (CMB). This work examines the effect such injections would have on CMB formation through comparison with the latest *Planck* measurements. We consider light-mass PBHs in the range  $10^{15}$  g to  $10^{17}$  g which inject energy through Hawking radiation. In this work we present a review of the current state-of-the-art models of recombination and their extensions which incorporate energy injections from non-standard sources. We also give a basic introduction to Bayesian inference in the context of parameter inference, and explain in detail how nested sampling and the *Planck* likelihood code operate. With this knowledge, we create a new cosmological parameter estimation program, `pc_multinest`, which utilises the nested sampler MULTINEST, the Boltzmann code CLASS and the *Planck* likelihood code to place constraints on PBH parameter space. During the completion of the program, we also improve the existing HYREC recombination code to correctly account for large energy injections. We then reproduce the current limits on PBH fraction  $f_{\text{PBH}}$  presented by Clark et al. [1], before extending the analysis to examine the effects a set of free  $\Lambda$ CDM parameters has on the resulting bounds. By allowing the base  $\Lambda$ CDM parameters to vary, we show the exclusion bounds are alleviated by a full order of magnitude across the mass range. We also present results extending the typical monochromatic mass distribution to a log normal one. Fixing the PBH fraction and allowing both the mean and standard deviation of the distribution to vary indicates that delta mass distributions are more sensitive to freeing  $\Lambda$ CDM parameters than a mass distribution with finite spread. Results for different PBH mass distributions show that increasing the spread of the mass distribution results in tighter constraints at heavier masses and a slight relaxation for masses  $M_{\text{PBH}} \sim 10^{15}$  g. For a uniform mass distribution, the 95% exclusion limits are independent of PBH mass, giving  $f_{\text{PBH}} < 6.1 \times 10^{-7}$  when base  $\Lambda$ CDM parameters are fixed and  $f_{\text{PBH}} < 1.3 \times 10^{-5}$  when they are allowed to vary.



## Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

I acknowledge the support I have received for my research through the provision of an Australian Government Research Training Program Scholarship

---

Harry David Poulter

November 2018



## Acknowledgements

When I began this project, I was not entirely sure what I was getting into. What both my supervisors and I had originally thought would be a fairly straightforward task of replicating results became (as it seems to in research) a long and protracted effort on everyone's part. In retrospect, this project was no easy undertaking: there were many setbacks, bugs and technical issues along the way. To that end, I would like to acknowledge all those, man and machine, who have helped me push this project to its completion.

Beginning with the support of my two supervisors Martin and Tony, without which I would undoubtedly still be months off finishing. They have given me constructive advice and lifted me out of some rough patches on multiple occasions. Alongside them, I have been fortunate enough to have Jan Hamann supporting me from Sydney. Being the resident cosmologist, Jan has been indispensable on topics spanning CMB physics and its interpretation to cosmological perturbation theory, and for that I am very grateful. I would also like to thank my family, who although not as technically-minded as those above, have allowed on multiple occasions time for me to express my displeasures during times of drudgery. Alongside them, I also have my peers to thank for fruitful and engaging conversations: those with Violet Harvey straying at times quite far from the realm of physics (and reasonable lunch hours).

Outside my immediate support network, my correspondence with Yacine Ali-Haïmoud was instrumental in the completion of this work, allowing the extension of a computer program to better incorporate higher energy injections — a necessity for this work.

I would also like to acknowledge the Phoenix supercomputer at the University of Adelaide, the use of which made the high-intensity computations that formed the backbone of this work feasible. To a lesser extent, I would also like to acknowledge the CoEPP cloud computing cluster which allowed me to get my feet wet in the ocean that is job scheduling on high performance computing clusters.

This Masters degree was completed with the support of a Masters (No Honours) scholarship offered by the University of Adelaide.



# Chapter 1

## Introduction

The study of the universe on a rational, scientific level can be dated as far back as ancient Greek times, where philosophers of the pre-Socratic era speculated the origins of the universe and its fundamental constituents based on their own observations. It remained a philosophical topic for millennia until Nicolaus Copernicus revolutionised the current view of the universe with his heliocentric model of the solar system. The model was later revisited by early cosmologists such as Giordano Bruno in the 16th century leading to the idea of the mediocrity principle: that Earth's position in the universe is not a privileged place. This idea became the cornerstone of modern cosmology, nowadays manifesting itself in the form of the cosmological principle. It was not until the mid 20th century when Albert Einstein put forth his theory of general relativity — linking the matter and energy in the universe with its spatial curvature — that a model of the universe was able to be fully quantified. This led to the development of the Friedmann equations, which fully described the evolution of the universe and its homogeneous and isotropic constituents.

In recent decades, cosmology has evolved from a field based primarily on theoretical concepts to a high-precision science. This era of high-precision began with the serendipitous discovery of the cosmic microwave background (CMB) by Penzias and Wilson [2] in 1964, which was compelling evidence for a universe that began in a hot, dense state. Since then, balloon-borne experiments such as BOOMERanG and MAXIMA, and ground-based observations from observatories like the SPT have measured the small-scale variation in this signal. Similarly, space-based satellites like COBE, WMAP and the most recent *Planck* experiment have measured these fluctuations across the full sky. The study of these measurements prompted the formation of a new field to describe how quantum perturbations evolved into the large-scale structure of the universe visible today, known as cosmological perturbation theory. It wasn't until the observation of the accelerating expansion of the universe by Riess et al. [3] and Perlmutter et al. [4] that the  $\Lambda$ CDM model became the standard cosmological model. Its appeal lies in its simplicity: able to describe the complete evolution of the universe with only six parameters in the standard formulation.

The formation of the CMB is heavily governed by the manner in which the fully-ionised primordial plasma of hydrogen nuclei, helium nuclei and electrons became neutral during a period of the universe’s history known as recombination. The field of recombination physics was pioneered by the seminal works of Peebles [5] and Zeldovich, Kurt, and Syunyaev [6], who independently modelled recombination using an effective three-level hydrogen atom. More contemporary efforts have expanded on this model by incorporating higher level effects from the inclusion of more energy levels. This has led to computations of the ionisation history at an unprecedented level of precision. With such accurate models, the effects non-standard energy injections have on the ionisation history are able to be examined in closer detail. Whilst direct measurement of the CMB as a means to constrain the  $\Lambda$ CDM model is not the only way of doing so, this work will look at using these measurements to constrain the properties of dark matter, assuming a standard  $\Lambda$ CDM universe.

In its simplest form, the  $\Lambda$ CDM model only incorporates non-interacting dark matter, so a sensible extension to this theory would be to include species of dark matter that interact in some way. Assuming such dark matter was formed long before recombination, the CMB can be used to place constraints on any non-standard interactions that may have occurred during recombination due to this interacting dark matter. Recent works by Slatyer [7, 8, 9], Giesen et al. [10], and Galli et al. [11, 12, 13] have examined the effects annihilating and decaying particle dark matter has on recombination. Non-particle dark matter can also benefit from this work if their energy injection primarily constitutes of electrons, positrons and photons.

The latest discoveries of gravitational wave events by the LIGO collaboration [14] have prompted a burst of activity surrounding primordial black holes (PBHs). Due to the nature of some signals measured by the LIGO collaboration, some have suggested that the progenitor black holes may have been primordial in their nature [15, 16, 17, 18], forming during the inflationary epoch long before the first stars. Due to heavy exclusions from astrophysical observations, the currently viable mass range of PBHs has been split into two main regions: a high-mass range of  $10 M_{\odot}$  to  $10^2 M_{\odot}$ , and a low-mass range of  $10^{15}$  g to  $10^{17}$  g. The heavier mass region of the order of solar masses injects energy through the formation of an accretion disk. Recent papers by Ali-Haïmoud and Kamionkowski [19] and Poulin et al. [20] demonstrate the non-trivial nature of modelling and computing the energy output of an accretion disk around a black hole. By contrast, the lighter mass PBHs inject energy through Hawking radiation, a process that can be derived from purely thermodynamic arguments. Light mass PBHs have already been explored in the literature [1, 21], and typically consider the black holes to have a single fixed mass — a monochromatic mass distribution. Work done by Clesse and Garca-Bellido [22] shows that some models of inflation can lead to an extended distribution of masses for PBHs. Previous works have considered the effects of a non-monochromatic mass distribution [23, 24, 25, 26], but have not taken a ground-up approach for these light mass evaporating PBHs.

In this work, we aim to improve the current limits on PBHs in the early universe through examining their impact on the formation of the CMB. We aim to make use of the latest CMB data available from *Planck* to place constraints on the contribution PBHs make to recombination. To do this, we closely follow the method outlined in Clark et al. [1], which takes work done by Slatyer [8, 9] on computing the effects injected photons and electrons have on the recombination process to place constraints on evaporating PBHs. This work will extend the Clark et al. [1] analysis by examining the effects an extended mass distribution may have on the current bounds presented in that paper. In the process, a new program `pc_multinest` is written, combining the cosmological code `CLASS` with a parameter sampling program `MULTINEST` to evaluate the Bayesian posterior distributions of PBH parameters through use of the *Planck* likelihood code.

Chapter 2 outlines the basics of modern cosmology, and motivates the physics behind the  $\Lambda$ CDM model. It also covers the formation and subsequent description of the cosmic microwave background. Following this, Chapter 3 reviews the three-level and multi-level atom models in the context of recombination. The extension of the simple three-level atom model to incorporate effects of non-standard interacting dark matter is also presented. Chapter 4 covers how parameter estimation works in the context of the  $\Lambda$ CDM, and provides a review of Bayesian inference. Chapter 5 presents work towards extending the current analysis to incorporate the effects of mass distributions into recombination calculations, before citing the final results and exclusion regions for the set of PBH parameters under consideration. Finally, the work is concluded in Chapter 6, presenting an overview of the work undertaken, the results achieved and possible extensions to the analysis.



## Chapter 2

# Cosmology and the Cosmic Microwave Background

### 2.1 Introduction

To begin this work we present an introduction to the field of cosmology in Section 2.2, culminating in a full mathematical description of the current  $\Lambda$ CDM model in Section 2.3. We will then focus specifically on the cosmic microwave background, its physics and interpretation in Section 2.4.

### 2.2 Fundamentals of Cosmology

For a long time, the study of the universe was purely speculative in its nature, underpinned by philosophical principles and ideas. It was not until the theory of general relativity that cosmology became a physical science with testable predictions. In 1927, Georges Lemaître put forth the idea of an expanding universe based on Einstein’s proposal of a steady-state universe [27]. Previous observational work by Slipher [28, 29] supported this viewpoint, although it wasn’t until measurements by Hubble [30] that the idea of an expanding universe was fully embraced. This idea led to the formulation of the Big Bang theory of cosmology, where the universe began as an incredibly hot and dense primeval fireball, before expanding and cooling into the universe we see today.

The best confirmation for the Big Bang cosmology came with the discovery of the cosmic microwave background (CMB) by Penzias and Wilson [2]. The CMB is the first light from the Big Bang made of high energy photons that have since “cooled” due to the redshift associated with the expansion of the universe. As a result, they have moved from the high energy region to the microwave region of the electromagnetic spectrum. Since its discovery, purpose-made experiments such as the balloon-borne BOOMERanG [31] and MAXIMA [32] have further confirmed its existence, with satellite-based experiments being able to measure the CMB signal across the whole sky. The recent WMAP [33] and *Planck* [34] experiments have turned

cosmology into a precision science, allowing for much more precise tests of extensions to the model.

This section will review the principles upon which the current  $\Lambda$ CDM model of the universe rests, and summarise its components.

### 2.2.1 Cosmological Principle

In analogy with special relativity, measurement of quantities in the universe like velocity and pressure depend on the reference frame chosen. Cosmologists define a frame which is at rest with respect to the expanding universe, and call this the *fundamental observer's frame*. Two fundamental observers will seem to move relative to each other, but this motion is due to the expanding universe and these observers will still report the same measurements. Quantities they report are said to be *co-moving* with the expansion of the universe. Such observers can be seen as the cosmological equivalent to special relativity's rest frame.

Modern cosmological theories subscribe to the *cosmological principle*, which states [35]

*The universe as seen by fundamental observers is homogeneous and isotropic.*

This principle has been met with good agreement from galactic redshift surveys such as 2dFGRS [36], 6dFGS [37], wigglyZ [38] and BOSS [39, 40], which have shown that for sufficiently large distances, galaxies appear to be distributed homogeneously. Similarly, measurement of the CMB has shown deviations in photon temperature of just one part in  $10^5$  [34].

It is also worth considering the cosmological analogue to proper time. To do this we consider Weyl's postulate, which states that there exists a bundle of timelike geodesics which can be used to define perpendicular non-intersecting spacelike surfaces. These surfaces permit an equivalent definition of proper time, known as *cosmic time*.

For instance, consider a group of fundamental observers who synchronise their watches so at any given instant of time they all agree on the time they are measuring. The time that is being measured is in a sense co-moving with the expansion of the universe — as the fundamental observers move apart due to expansion, their watches will still agree at the same instant of time despite each observer seeing all other watches slow down relative to theirs.

### 2.2.2 Contents of the Universe

Cosmologists take a very boring interpretation of the universe's contents. Instead of dealing with the myriad of nebulae, satellite galaxies, stellar clusters and exoplanets, cosmologists categorise matter (and energy) into five main categories: baryonic matter, photons, neutrinos, dark matter and dark energy. These categories define particles and forces that have different properties and interactions. We shall see that for the standard cosmological model, these five groups are all that are needed to explain a large number of observations.

### Baryonic Matter, Photons and Neutrinos

For cosmologists, baryonic matter consists of protons, neutrons, electrons and other Standard Model particles. The name “baryonic” is a misnomer, as this group also includes leptons. However, this term is widely used in cosmological literature and so will be used throughout this work too. Baryonic matter does not usually encapsulate photons and neutrinos, as these particles behave differently. Quantities relating to baryonic matter will have a subscript  $b$ .

In the very early universe, baryonic matter is relativistic due to the large temperatures in excess of  $T > 1 \text{ GeV}$ . At later times, it becomes non-relativistic as the particles’ kinetic energies drop below their rest masses as the universe cools.

Photons on the other hand never become non-relativistic and so must be treated differently throughout the universe’s history. Neutrinos sit in the middle: having a small rest mass relative to those of baryonic particles, neutrinos behave relativistically for a much longer period but will eventually cool and become non-relativistic.

For times during the early universe such as recombination, baryonic matter behaves non-relativistically whilst neutrinos are relativistic. As such, neutrinos and photons can be considered effectively equivalent during this time period, and we will label them both as *radiation*. All quantities related to radiation will have a subscript  $\gamma$ .


### Dark Matter

Dark matter (DM) was first proposed by Zwicky [41] after he made measurements of galactic rotation curves that did not align with the standard Newtonian theory. Confirmation of these measurements was made much later by Rubin and Ford [42].

Rotation curves measure the angular velocity with which objects such as stars orbit the centre of their galaxy. Classically, one would expect that a conventional spiral galaxy behaves in much the same way as our Solar System: matter closer to the centre will orbit faster than matter further away. This can be seen by considering Newton’s universal law of gravitational attraction for a central mass  $M$  and orbiting mass  $m$  a distance  $r$  apart. The force acting between them is given by

$$\mathbf{F} = -\frac{GMm}{r^2} \hat{\mathbf{r}}.$$

However, observations of rotation curves showed that this doesn’t happen and instead galaxies behave more like a solid disk where angular velocity remains constant at large distances. Figure 2.1 shows the observed rotation curve of the spiral galaxy NGC 2403. We note the angular velocity of the atomic hydrogen gas does not decay as  $1/r^2$  and instead remains constant out to large  $r$ . Such observations suggest there is a hidden component of matter that acts to make spiral galaxies behave gravitationally as though they are solid disks. As this matter was not visible in any part of the electromagnetic spectrum, it was termed *dark matter*.



**Fig. 2.1** Rotation curve for galaxy NGC 2403, from de Blok et al. [43]. The solid black line is a fit to the atomic hydrogen gas velocities in the galaxy.

Strictly speaking, the proposal of DM is not the only explanation for these anomalous rotation curves. As these observations rely on Einstein’s general relativity for explanation, other alternative theories to DM involve modifying gravity itself to account for these phenomena. However, such theories will not be covered in this work as they do not form part of the standard cosmological model.

Astrophysical evidence found elsewhere in gravitationally-bound galactic clusters [44, 45], gravitational lensing [46] and the Bullet cluster [47] favours cold dark matter (CDM), which is non-relativistic, slow moving and able to form gravitationally-bound clumps with sizes of the order of galaxies.

Cold dark matter appears in the  $\Lambda$ CDM model as a substance that only interacts gravitationally with the other fluids. Potential  $\Lambda$ CDM extensions can consider specific DM models that contain extra interactions between baryonic matter and radiation, for instance.

## Dark Energy

The expansion of the universe has been known about since early measurements of galactic redshifts by Slipher [28, 29]. If a galaxy emits a photon of wavelength  $\lambda_e$ , and is received at Earth with wavelength  $\lambda_r$ , the redshift  $z$  is defined as

$$z := \frac{\lambda_r - \lambda_e}{\lambda_e}.$$

Further confirmation of these measurements came in the late 1920s when American astronomer Edwin Hubble was able to quantify the recessional velocities of galaxies [30]. By measuring the apparent magnitude of galaxies, Hubble was able to work out the distances to them. Combining this with the measurement of the recessional velocity resulted in an empirical




Fig. 2.2 Original plot from Hubble [30] showing Hubble's law for a set of observed galaxies.

law sharing his name relating the velocity  $v$  to the galaxy distance  $D$ :

$$v = H_0 D,$$

where  $H_0$  is known as the *Hubble constant*, currently measured by Planck Collaboration et al. [34] to be  $(67.31 \pm 0.96) \text{ km s}^{-1} \text{ Mpc}^{-1}$ . Figure 2.2 is the original plot from Hubble's paper showing the linear relationship of Hubble's law.

The explanation for this redshifting is that the light from galaxies is "stretched out" by the expanding universe, and as a result shifts the light towards the red end of the spectrum [48]. The mechanism is very different to that of a conventional Doppler shift, and as such is sometimes referred to as a cosmological redshift.

Redshift is used by cosmologists to refer to different times in the Universe unambiguously, as it can be measured without any knowledge of the evolution of  $a(t)$  and therefore any assumptions on the chosen model. Other measurements such as time and distance depend on the density of matter and energy in the universe and so do not make good standard distance measures. In the case of an expanding universe, there are also multiple ways with which to measure distances, each giving different results.

This picture of an expanding universe, although controversial at the time, eventually came to be accepted by the wider scientific community. Hubble's law stood until 1998 when teams led by Riess et al. [3] and Perlmutter et al. [4] decided to independently repeat the analysis using observations of type Ia supernovae. Due to their high luminosity, type Ia supernovae allowed probing of Hubble's law to much farther distances. The result both teams found not only confirmed that the universe was expanding, but also was increasing its rate of expansion.

This phenomenon is described by the presence of an extra component of energy, *dark energy*, which drives this acceleration. Although its properties and origin is still yet unknown, it is included in the  $\Lambda$ CDM model to explain this apparent accelerating expansion. This latent energy shares no interactions with matter or radiation, and as we will see exists purely as a vacuum energy associated with space itself.

## 2.3 The $\Lambda$ CDM Model

The  $\Lambda$ CDM model brings together all the observable elements of the universe in a cohesive and consistent way. It explains how the four main components of the universe (baryonic matter, cold dark matter, radiation and dark energy) work together and how they each influence the universe's evolution. The model also gives a quantitative description of the formation of the cosmic microwave background, and can predict the power spectrum of the actual CMB to very high accuracy.

To do this, the model uses the language of general relativity. Following this mathematical framework leads to the Friedmann equations which govern the dynamics and evolution of the universe in terms of the Hubble parameter  $H(t)$ .

### 2.3.1 Friedmann Equations

The  $\Lambda$ CDM model is based on the cosmological principle which is encapsulated by the Friedmann-Lemaître-Robertson-Walker (FLRW) metric, given in hyperspherical coordinates by:

$$ds^2 = -c^2 dt^2 + a^2(t) \left[ d\chi^2 + S_k^2(\chi) (d\theta^2 + \sin^2 \theta d\phi^2) \right].$$

Here, the scale factor  $a(t)$  is a multiplicative factor defining the expansion of the universe at some cosmic time  $t$ .  $k$  is the curvature constant, which is  $+1$  for a closed universe,  $0$  for a flat universe and  $-1$  for an open universe. The function  $S_k(\chi)$  changes form depending on the value of  $k$ :

$$S_k(\chi) = \begin{cases} \sin \chi, & k = +1, \\ \chi, & k = 0, \\ \sinh \chi, & k = -1. \end{cases}$$

This metric is the most general form that supports the cosmological principle. It is used in the Einstein field equations to derive the equations of motion for the universe, known as the Friedmann equations. Appendix A provides an explicit derivation of these equations from Einstein's field equations and the FLRW metric, noting the use of the  $(-, +, +, +)$  metric sign convention. The Friedmann equations are

$$\dot{H} + H^2 = -\frac{4\pi G}{3} \left( \rho + \frac{3p}{c^2} \right) + \frac{\Lambda c^2}{3}, \quad H^2 = \frac{8\pi G \rho}{3} - \frac{kc^2}{a^2} + \frac{\Lambda c^2}{3},$$

where  $H(t)$  is the Hubble parameter defined as  $H(t) = \dot{a}(t)/a(t)$ . The cosmological constant  $\Lambda$  is also included in the equations as an effective vacuum energy used to describe dark energy. Here we note the current value of the Hubble parameter at  $t = t_0$  is  $H(t_0) \equiv H_0$ , the Hubble constant as it appears in Hubble's law. The Hubble constant is usually rewritten in the literature as

$$H_0 = 100h \text{ km s}^{-1} \text{ Mpc}^{-1},$$

where  $h$  is effectively a dimensionless Hubble constant.

It is instructive to consider the simplest geometry for the Friedmann equations. Here, we take a flat universe ( $k = 0$ ) with no cosmological constant ( $\Lambda = 0$ ). In such a model, the Friedmann equations reduce to

$$\dot{H} + H^2 = -\frac{4\pi G}{3} \left( \rho + \frac{3p}{c^2} \right), \quad H^2 = \frac{8\pi G \rho}{3}.$$

At the present cosmic time  $t = t_0$ , the Hubble constant  $H_0$  is given by the right-hand equation:

$$H_0^2 = \frac{8\pi G \rho_c}{3}. \quad (2.3.1)$$

Rearranging, this gives

$$\rho_c = \frac{3H_0^2}{8\pi G},$$

which is known as the *critical density*. This is the density at which the universe needs to be today to appear spatially flat in the absence of any cosmological constant. This quantity is used to define the dimensionless *density parameters*, given by

$$\Omega_i(t) := \frac{\rho_i(t)}{\rho_c},$$

for  $i = \{\text{b, cdm}, \gamma, \Lambda\}$  being a label for each component of matter or energy considered: baryonic matter, cold dark matter, radiation or dark energy. This parametrisation is useful as it is defined independent of a given cosmological model, and effectively encapsulates the fraction of the universe made from each constituent compared to the critical density as measured today.

In a flat universe, we expect at the current time that

$$\Omega_{\text{b},0} + \Omega_{\text{cdm},0} + \Omega_{\gamma,0} + \Omega_{\Lambda,0} = 1. \quad (2.3.2)$$

However, for universes where  $k \neq 0$  and the total mass-energy density is above or below the critical density this equality will not hold. This can be rectified by defining an effective curvature “energy density”

$$\Omega_{k,0} := -\frac{kc^2}{H_0^2}, \quad (2.3.3)$$

which allows us to write

$$\Omega_{b,0} + \Omega_{\text{cdm},0} + \Omega_{\gamma,0} + \Omega_{\Lambda,0} = 1 - \Omega_{k,0},$$

which holds true for  $\Lambda$ CDM universes with arbitrary spatial curvature. This quantity will come to use later on when we consider the evolution of the Hubble parameter  $H(t)$  for a general  $\Lambda$ CDM universe with unspecified  $k$ .

### 2.3.2 Redshift-Scale Factor Relationship

Before considering how these density parameters evolve over cosmic time, it is worth examining the quantities of redshift  $z$  and scale factor  $a(t)$  more closely. As this relationship is used throughout the rest of the work (and across all other cosmology literature), it is worthwhile to follow the derivation explicitly.

First, consider light emitted from a galaxy with hyperspherical coordinates  $(\chi_1, \theta_1, \phi_1)$  at some time  $t_1$  with some wavelength  $\lambda_1$ . The light is later received at Earth at time  $t_0$ , the current time, with a cosmologically redshifted wavelength of  $\lambda_0$ .

Noting that light travels along null geodesics ( $ds^2 = 0$ ) with constant  $\theta = \theta_1$  and  $\phi = \phi_1$ , so from the FLRW metric:

$$\begin{aligned} ds^2 &= -c^2 dt^2 + a^2(t) \left( d\chi^2 + S_k^2(\chi) \left[ \sin^2 \theta d\phi^2 + d\theta^2 \right] \right) \\ \Rightarrow 0 &= -c^2 dt^2 + a^2(t) d\chi^2 \\ \Rightarrow \frac{c^2 dt^2}{a^2(t)} &= d\chi^2. \end{aligned}$$

Taking the square root of each side, we take the negative value corresponding to a light wave propagating towards the observer and integrate both sides to give

$$\int_{t_1}^{t_0} \frac{c dt}{a(t)} = \int_0^{\chi_1} d\chi.$$

Now, we consider the second crest of this light wave which is emitted a period  $T_1 = \lambda_1/c$  later, and received a period  $T_0 = \lambda_0/c$  later. As both periods are very small, we expect that there would be negligible apparent movement of the galaxy at  $\chi_1$  due to spacetime expansion, so  $\chi(t_1) = \chi(t_1 + T_1)$ . This allows us to write a similar equation for the second wave crest:

$$\int_{t_1+T_1}^{t_0+T_0} \frac{c dt}{a(t)} = \int_0^{\chi_1} d\chi.$$

Equating both expressions, we see

$$\int_{t_1+T_1}^{t_0+T_0} \frac{c dt}{a(t)} = \int_{t_1}^{t_0} \frac{c dt}{a(t)}.$$

We now split both integrals up in order to cancel off the overlap between the two ranges:

$$\begin{aligned}\int_{t_1+T_1}^{t_0+T_0} \frac{c dt}{a(t)} &\equiv \int_{t_1+T_1}^{t_0} \frac{c dt}{a(t)} + \int_{t_0}^{t_0+T_0} \frac{c dt}{a(t)}, \\ \int_{t_1}^{t_0} \frac{c dt}{a(t)} &\equiv \int_{t_1}^{t_1+T_1} \frac{c dt}{a(t)} + \int_{t_1+T_1}^{t_0} \frac{c dt}{a(t)}.\end{aligned}$$

Hence,

$$\int_{t_0}^{t_0+T_0} \frac{c dt}{a(t)} = \int_{t_1}^{t_1+T_1} \frac{c dt}{a(t)}. \quad (2.3.4)$$

To evaluate these integrals, we will show that both integrands are effectively constant between the bounds of integration. As this is applicable to both integrals, we shall drop the subscripts and consider an arbitrary time  $t$  with corresponding period  $T = \lambda/c$ . Taking the evaluation of the integrand at the end point  $t + \lambda/c$ , we Taylor expand to find

$$\frac{c}{a(t + \lambda/c)} = \frac{c}{a(t)} \left[ 1 - \lambda/c \left( \frac{\dot{a}(t)}{a(t)} \right) + \mathcal{O}((\lambda/c)^2) \right].$$

Considering the first term:

$$\lambda/c \left( \frac{\dot{a}(t)}{a(t)} \right) = (\lambda/c)H(t) \equiv \frac{\lambda/c}{t_{\text{H}}(t)},$$

where  $t_{\text{H}}$  is the *Hubble time*, defined as the reciprocal of the Hubble parameter:

$$t_{\text{H}}(t) := \frac{1}{H(t)} = \frac{a(t)}{\dot{a}(t)}.$$

In the case of a uniformly-expanding universe with  $a(t) = ct$ , the Hubble time reduces to cosmic time  $t_{\text{H}} = t$ .

As we are considering the ratio of an EM wave period to that of what is effectively cosmic time, then

$$\frac{\lambda/c}{t_{\text{H}}(t)} \ll 1,$$

and so we can truncate the Taylor expansion to the zeroth order term in  $\lambda/c$  to give

$$\frac{c}{a(t + \lambda/c)} \approx \frac{c}{a(t)},$$

to very good accuracy.

Hence, due to the many orders of magnitude difference in time scales, the integrand  $c/a(t)$  is constant and so can be directly evaluated:

$$\begin{aligned} \int_t^{t+\lambda/c} \frac{c dt'}{a(t')} &= \frac{c}{a(t)} \int_t^{t+\lambda/c} dt' \\ &= \frac{\lambda}{a(t)}. \end{aligned}$$

Substituting this back into the original equality in Eq. (2.3.4), we see

$$\frac{\lambda_0}{a(t_0)} = \frac{\lambda_1}{a(t_1)}.$$

From the definition of redshift, we have

$$1 + z = \frac{\lambda_0}{\lambda_1} = \frac{a(t_0)}{a(t_1)}.$$

Noting from the definition of the reduced scale factor  $a(t)$  that  $a(t_0) = 1$  this gives us

$$1 + z = \frac{1}{a(t)}.$$

This identity is known as the *redshift-scale factor relationship*. This relationship is widely used throughout cosmology as it links the theoretical scale factor  $a(t)$  from the FLRW metric with the observable redshift  $z$ .

Quantities from now onwards will be written in terms of redshift  $z$ . This is because redshift is independent of the model considered, as opposed to cosmic time  $t$  which is a function of constituent densities.

### 2.3.3 Evolution of the Hubble Parameter

Although we could solve the Friedmann equations explicitly for simple geometries, by taking a different, more physically-motivated approach we will find a general solution  $H(z)$  in terms of the relative fractions of the universe's constituents, the density parameters  $\Omega_i$ .

We will start by considering the redshift evolution of different components. For a cosmological fluid made of multiple components, its stress-energy tensor can be written as

$$T^{\mu\nu} = \sum_i (T^{\mu\nu})_i,$$

where each  $(T^{\mu\nu})_i$  is the stress-energy tensor corresponding to component  $i$ .

The components are modelled as anisotropic perfect fluids with pressures  $p_i(t)$  and densities  $\rho_i(t)$ , such that

$$(T^{\mu\nu})_i = \left( \rho_i + \frac{p_i}{c^2} \right) u^\mu u^\nu + p_i g^{\mu\nu}.$$

From the linearity of the covariant derivative, we expect that each component must satisfy the energy-momentum conservation equation separately, assuming each component is non-interacting on cosmological scales:

$$\nabla_\mu (T^{\mu\nu})_i = 0.$$

This equation leads to the relativistic conservation and continuity equations, the latter of which can be expressed as:

$$\frac{d\rho_i}{dt} + \left( \rho_i + \frac{p_i}{c^2} \right) \frac{3\dot{a}}{a} = 0,$$

where we have used that both pressure  $p_i$  and density  $\rho_i$  only have dependence on time in an isotropic, homogeneous universe.

Multiplying by  $a^3$ , the equation can be rewritten as:

$$\frac{d(\rho_i a^3)}{dt} = -\frac{3p_i \dot{a} a^2}{c^2},$$

leading to

$$\frac{d(\rho_i a^3)}{da} = -\frac{3p_i a^2}{c^2}.$$

All fluids have some relationship between their density and pressure. Relativistically, this may be written as

$$p_i = w_i \rho_i c^2,$$

where  $w_i$  is the *equation of state (parameter)* for component  $i$ . It is typically constant for most fluids, but some models of dark energy incorporate a running equation of state parameter. For this work, however, all  $w_i$ s are taken to be constant.

Using the density-pressure relationship, we can solve the differential equation for density giving

$$\rho_i(t) = \rho_{i,0} [a(t)]^{-3(1+w_i)},$$

where  $\rho_{i,0}$  is the component's density measured at the present time. From the scale factor-redshift relationship, this becomes <sup>1</sup>

$$\rho_i(z) = \rho_{i,0} (1+z)^{3(1+w_i)}.$$

---

<sup>1</sup>Note that density has been written as a function of redshift  $z$ , as opposed to cosmic time  $t$ . This is simply a reformulation that instead gives what redshift the measurements were made at, rather than at what time they were made.

$w_m = 0$  for matter (baryonic and CDM), as there is no pressure associated with sparsely-distributed grains. This gives

$$\rho_m(z) = \rho_{m,0}(1+z)^3,$$

which intuitively makes sense as we expect the density of matter to scale with the effective volume,  $a^{-3}(t)$ . This means matter is conserved with co-moving volume, which is precisely the result one would expect from using the continuity equation.

From electromagnetism, radiation pressure is given by  $p_\gamma = \rho_\gamma c^2/3$ , meaning  $w_\gamma = 1/3$ . This gives

$$\rho_\gamma(z) = \rho_{\gamma,0}(1+z)^4.$$

This extra factor of  $(1+z)$  for radiation can be seen as coming from the extra dilution of the radiation energy density due to cosmological redshifting.

In the  $\Lambda$ CDM model, dark energy takes the form of a cosmological constant  $\Lambda$  with an equation of state of  $w_\Lambda = -1$ . Hence,

$$\rho_\Lambda(z) = \rho_{\Lambda,0}.$$

Interestingly, this means that dark energy density is constant with the expanding universe. Physically, this may be interpreted as a vacuum energy with negative pressure.

Note here that we have decided to include dark energy  $\Lambda$  as part of the mass-energy content of the universe, as opposed to including it separately in the Friedmann equations as before. Hence, the second Friedmann equation may be equivalently written as

$$H^2 = \frac{8\pi G\rho}{3} - \frac{kc^2}{a^2},$$

where now the total density of the universe is

$$\rho(z) = \rho_m(z) + \rho_\gamma(z) + \rho_\Lambda(z).$$

The above can be equivalently written using the density parameter formulation as

$$\Omega(z) \equiv \frac{\rho(z)}{\rho_c} = \Omega_m(z) + \Omega_\gamma(z) + \Omega_\Lambda(z).$$

Hence, the Friedmann equation may be re-expressed as

$$\begin{aligned} H^2 &= \frac{8\pi G\rho_c}{3} \Omega(z) - \frac{kc^2}{a^2} \\ &= H_0^2 \left[ \Omega(z) - \frac{kc^2}{H_0^2 a^2} \right], \end{aligned}$$

where we have used the form for  $H_0^2$  in Eq. (2.3.1).

Using the definition of the curvature energy density  $\Omega_{k,0}$  in Eq. (2.3.3), we find

$$H^2(z) = H_0^2 \left[ \Omega_{m,0}(1+z)^3 + \Omega_{\gamma,0}(1+z)^4 + \Omega_{\Lambda,0} + \Omega_{k,0}(1+z)^2 \right],$$

where  $\Omega_{m,0}$  is the matter density, including both baryon and dark matter contributions:

$$\Omega_{m,0} = \Omega_{b,0} + \Omega_{\text{cdm},0}.$$

For the flat  $k = 0$   $\Lambda$ CDM universe, only three quantities are needed to fully specify the evolution of  $H(z)$ : the baryonic and dark matter quantities  $\Omega_{b,0}$  and  $\Omega_{m,0}$ , and the normalisation  $H_0$ . The radiation density parameter comes from the measurement of the average CMB temperature  $T_0$ . The remaining dark energy density is given by the completeness relationship in Eq. (2.3.2).

The final parameter needed in the  $\Lambda$ CDM model to describe this evolution is the normalisation constant,  $H_0$ . There are other parameters that can be used as a substitute for the Hubble constant, such as the age of the universe  $t_0$ . In such a case, from our previous discussion on the redshift-scale factor relationship we saw that  $1/H_0 \approx t_0$ .

In this work, the quantity  $\theta_s$  is used instead of  $H_0$  as their usage is interchangeable. This is the ratio of the sound horizon at decoupling to the angular diameter distance at decoupling:

$$\theta_s = \frac{r_s(z_*)}{D_A(z_*)},$$

where  $z_*$  is the redshift of recombination.

The sound horizon is the farthest distance a sound wave could have travelled in the time elapsed from the start of the universe to redshift  $z$ . Similarly, the angular diameter distance is the distance between an emission point and reception point at the time light was emitted for a flat universe. This angular size  $\theta_s$  acts as a fixed point to calibrate the evolution of the expansion  $H(z)$  in a manner similar to  $H_0$  and  $t_0$ , which is what allows us to consider it over a more traditional selection of  $H_0$ .

In order to remove some degeneracies between other parameters, the quantities  $\Omega_b h^2$  and  $\Omega_{\text{cdm}} h^2$  are sometimes used as parameters, where  $h$  is the dimensionless Hubble parameter. It is also more convenient to use  $100\theta_s$ , as  $\theta_s \sim \mathcal{O}(10^{-3})$ .

### 2.3.4 Parametrising the Early Universe

At this point, we are able to describe the macroscopic evolution of the universe and its constituents, and how this affects its expansion rate. However, we are yet to consider the early universe and in particular the conditions and physics relevant to the formation of the cosmic microwave background. A more thorough treatment of this subject would invoke cosmological perturbation theory (CPT). CPT examines perturbations  $h^{\mu\nu}$  to the spacetime

metric  $g^{\mu\nu} \rightarrow \tilde{g}^{\mu\nu} = g^{\mu\nu} + h^{\mu\nu}$ , and how such perturbations manifest themselves in various power spectra such as temperature and density. Although taking this approach lies outside the scope of the work, an important result to take from CPT is that the perturbation  $h^{\mu\nu}$  can be decomposed into two scalar, two vector and two tensor degrees of freedom. This decomposition leads to perturbations being categorised as being scalar, vector or tensor-like in their nature.

To begin the discussion, we shall consider the parameters that describe the initial perturbations in the matter density spectrum — a scalar perturbation field. In the very early universe, an inflationary epoch resulted in the exponential growth of the universe and turned quantum mechanical fluctuations of the inflation field  $\phi$  into physical perturbations to the matter density. By considering the slow-roll approximation for such a field  $\phi$ , the resulting power spectrum of these perturbations is very close to constant in wave number  $k$  [35]. This result motivates the following Fourier space form of the matter power spectrum:

$$P_s(k) = A_s k^{n_s - 1},$$

where  $k$  is the wave number of the perturbation under consideration. The scalar spectral index  $n_s \sim 1$  for the  $\Lambda$ CDM model, which suggests the primordial scalar power spectrum is close to being *scale invariant*, or that the spectrum contains equal power across all wave numbers  $k$ . Similarly,  $A_s$  governs the amplitude of the scalar perturbations. As it is usually a very small number, the value  $\ln(10^{10} A_s)$  is equivalently used in literature. These two parameters fully describe the contribution to the eventual large scale structure of the universe.

The final parameter needed describes the optical depth at reionisation,  $\tau$ . More generally, the optical depth  $\tau(z)$  is defined as

$$\tau(z) := \int_0^z \frac{n_e(z') c \sigma_T}{(1+z') H(z')} dz', \quad (2.3.5)$$

where  $n_e(z)$  is the electron number density and  $\sigma_T$  is the Thomson scattering cross-section. Hence, the optical depth as measured at the redshift of reionisation  $z_{\text{reio}}$  sets a limit on the amount of free electrons (or ionisation) happening due to stellar formation.


In total, only six parameters are needed to fully describe the base  $\Lambda$ CDM model of the universe:

$$(\Omega_b h^2, \Omega_{\text{cdm}} h^2, 100\theta_s, \tau, \ln(10^{10} A_s), n_s).$$

These parameters will be used later on when simulating  $\Lambda$ CDM universes as part of estimating the effects non-standard energy injection has on the early universe.

## 2.4 The Cosmic Microwave Background

During early times in the universe, light in the form of photons is trapped by electromagnetic interactions with the ionised baryon plasma. As the universe cools, the plasma becomes neutral



**Fig. 2.3** Geometry of CMB observations. Here,  $\hat{\mathbf{n}}$  is a normal vector in the direction of the observation point  $(\theta, \phi)$ , and  $\mathbf{k}$  is the wave vector of the CMB photon. Only a hemisphere is shown for brevity.

and it is during this time that photons are able to freely stream. The process, known as recombination, results in the formation of the cosmic microwave background (CMB): the first light visible from the Big Bang. The frequency spectrum of this light matches that of a black body with temperature  $T_0 = 2.725$  K to one part in  $10^5$ . The extremely small anisotropies are seeded from quantum fluctuations that were expanded to macroscopic scales during inflation, and form a powerful probe for early universe physics.

The CMB features both primary and secondary anisotropies. Primary anisotropies are those formed during recombination, whilst secondary anisotropies occur from the time of last scattering to when the photons arrive at the observer. Tertiary anisotropies due to foreground sources will be covered in Section 4.3. The ultimate goal is to decompose these anisotropies into spherical harmonics which allows us to create power spectra containing this information.

This section will outline the main sources of primary temperature and polarisation anisotropies and secondary anisotropies, as well as how they are turned into power spectra from the CMB signal measurement. Although the following is not integral to the understanding of this work, it is nonetheless important to have an appreciation for how quantities such as the E- and B-mode polarisations are computed and hence how variations in these quantities gives clues as to what was happening in the early universe.

### 2.4.1 Mathematical Description

As the CMB signal is present across the whole sky, it can be viewed as projected onto a sphere. The anisotropies are then described in terms of the amount of power contained within the spherical harmonics  $Y_{\ell m}(\hat{\mathbf{n}})$ .

The vector  $\hat{\mathbf{n}}$  is directed from the observer at the sphere's centre to a point on the sphere at  $(\theta, \phi)$ .  $\hat{\mathbf{n}}$  can also be taken as the direction that CMB photons are freely streaming from. Figure 2.3 shows the geometry being used in the following maths.

Electromagnetic radiation can be characterised by its wave vector  $\mathbf{k}$  and the electric field component (or polarisation vector)  $\mathbf{E}$ . Considering the geometry of the CMB, the wave vector is anti-parallel to  $\hat{\mathbf{n}}$ , *i.e.*  $\mathbf{k} \parallel \hat{\mathbf{n}}$ . Hence, the electric field lies in the plane perpendicular to  $\hat{\mathbf{n}}$ . This will become important later on when considering how polarisation anisotropies can be decomposed into their spherical harmonics.

The radiation has an associated intensity and polarisation, both of which can be evaluated using the electric field component  $\mathbf{E}(\hat{\mathbf{n}})$ . An efficient way of encapsulating this information is in the use of Stokes parameters, which are defined as the following time-averaged quantities:

$$\begin{aligned} I(\hat{\mathbf{n}}) &:= \langle E_x^2 \rangle + \langle E_y^2 \rangle, & Q(\hat{\mathbf{n}}) &:= \langle E_x^2 \rangle - \langle E_y^2 \rangle, \\ U(\hat{\mathbf{n}}) &:= \langle E_a^2 \rangle - \langle E_b^2 \rangle, & V(\hat{\mathbf{n}}) &:= \langle E_l^2 \rangle - \langle E_r^2 \rangle. \end{aligned}$$

The subscripts refer to the following three coordinate systems: the normal Cartesian basis  $(\hat{x}, \hat{y})$ , the same basis rotated by  $45^\circ$   $(\hat{a}, \hat{b})$  and the circular polar basis  $(\hat{l}, \hat{r})$ , where  $\hat{l} = \hat{x} + i\hat{y}$ .

Here,  $I(\hat{\mathbf{n}})$  and  $V(\hat{\mathbf{n}})$  correspond to the intensity and circular polarisation of the EM radiation. These quantities are both invariant under rotations around  $\hat{\mathbf{n}}$ , meaning they can both be decomposed in terms of spherical harmonics, allowing their associated power spectra to be calculated. The  $I$  parameter is associated with the ‘‘temperature’’ of the CMB photons, as their intensity is effectively a measure of the temperature of the corresponding black body emission. We will not be considering any spectra associated with the quantity  $V(\hat{\mathbf{n}})$ , as circular polarisation cannot be generated through Thomson scattering [49] and so we expect there to be no associated anisotropies in the  $V$  parameter.

$Q(\hat{\mathbf{n}})$  and  $U(\hat{\mathbf{n}})$  are orthogonal modes of linear polarisation. However, these quantities depend on the axes used to define them. As a consequence, rotating the coordinate basis by an angle  $\psi$  transforms  $Q(\hat{\mathbf{n}})$  and  $U(\hat{\mathbf{n}})$  as

$$\begin{pmatrix} Q \\ U \end{pmatrix} \rightarrow \begin{pmatrix} Q' \\ U' \end{pmatrix} = \begin{pmatrix} \cos 2\psi & \sin 2\psi \\ -\sin 2\psi & \cos 2\psi \end{pmatrix} \begin{pmatrix} Q \\ U \end{pmatrix}.$$

As  $Q(\hat{\mathbf{n}})$  and  $U(\hat{\mathbf{n}})$  are not invariant under rotations, they cannot be written in terms of spherical harmonics. However, the above property allows us to define the following quantities:

$${}_{\pm 2}S(\hat{\mathbf{n}}) := (Q \pm iU)(\hat{\mathbf{n}}),$$

which have the following transformation properties:

$${}_{\pm 2}S \rightarrow {}_{\pm 2}S' = e^{\mp 2\psi i} {}_{\pm 2}S.$$

The quantities  ${}_{\pm 2}S(\hat{\mathbf{n}})$  have a definite value of spin,  $\pm 2$ . This is defined in a purely mathematical sense, where more generally a function on a sphere  ${}_s f(\theta, \phi)$  is said to have spin  $s$

if it has the following transformation under a rotation of the coordinate basis by  $\psi$ :

$${}_s f \rightarrow {}_s f' = e^{-s\psi i} {}_s f.$$

By having definite spin values, the  ${}_{\pm 2}S(\hat{\mathbf{n}})$  functions can be decomposed in terms of *spin-weighted* spherical harmonics  ${}_{\pm 2}Y_{\ell m}(\hat{\mathbf{n}})$ . Hence, decomposing both  $I(\hat{\mathbf{n}})$  and  ${}_{\pm 2}S(\hat{\mathbf{n}})$ :

$$I(\hat{\mathbf{n}}) = \sum_{\ell m} a_{\ell m}^{\text{T}} Y_{\ell m}(\hat{\mathbf{n}}), \quad {}_{\pm 2}S(\hat{\mathbf{n}}) = \sum_{\ell m} a_{\ell m}^{\pm 2} Y_{\ell m}(\hat{\mathbf{n}}).$$

The coefficients are given by the inner product of their corresponding spherical harmonic with the original function:

$$\begin{aligned} a_{\ell m}^{\text{T}} &= \int Y_{\ell m}^*(\hat{\mathbf{n}}) I(\hat{\mathbf{n}}) \, \text{d}\Omega, \\ a_{\ell m}^{\pm 2} &= \int {}_{\pm 2}Y_{\ell m}^*(\hat{\mathbf{n}}) {}_{\pm 2}S(\hat{\mathbf{n}}) \, \text{d}\Omega, \end{aligned} \quad (2.4.1)$$

where the integral is over solid angle  $\text{d}\Omega$ .

For convenience, we define coefficients:

$$a_{\ell m}^{\text{E}} := -\frac{1}{2} (a_{\ell m}^{+2} + a_{\ell m}^{-2}), \quad a_{\ell m}^{\text{B}} := \frac{i}{2} (a_{\ell m}^{+2} - a_{\ell m}^{-2}),$$

which give the following relationship between the original Stokes parameters:

$$a_{\ell m}^{\text{E}} \pm i a_{\ell m}^{\text{B}} = - \int {}_{\pm 2}Y_{\ell m}^*(\hat{\mathbf{n}}) (Q + iU)(\hat{\mathbf{n}}) \, \text{d}\Omega.$$

These coefficients are known as the E- and B-modes of polarisation. As polarisation is a vector field  $\mathbf{E}(\hat{\mathbf{n}})$ , it can be represented as a sum of a curl- and divergence-free field. What we have done through spin-weighted spherical harmonics is effectively create these fields from the polarisation field. Here, the curl-free field is known as the E-mode, and the divergence-free field is known as the B-mode, mirroring that of electric and magnetic fields in electromagnetism.

The resulting CMB power spectra can be created by considering the auto- and cross-correlations of these coefficients. The B-mode has opposite parity to T and E, and so the corresponding TB and EB cross-correlations disappear. What we are left with are the following spectra, defined by:

$$\begin{aligned} C_{\ell}^{\text{TT}} &:= \frac{1}{2\ell + 1} \sum_m \langle a_{\ell m}^{\text{T}*} a_{\ell m}^{\text{T}} \rangle, & C_{\ell}^{\text{EE}} &:= \frac{1}{2\ell + 1} \sum_m \langle a_{\ell m}^{\text{E}*} a_{\ell m}^{\text{E}} \rangle, \\ C_{\ell}^{\text{BB}} &:= \frac{1}{2\ell + 1} \sum_m \langle a_{\ell m}^{\text{B}*} a_{\ell m}^{\text{B}} \rangle, & C_{\ell}^{\text{TE}} &:= \frac{1}{2\ell + 1} \sum_m \langle a_{\ell m}^{\text{T}*} a_{\ell m}^{\text{E}} \rangle. \end{aligned}$$

Here, we have a normalisation of  $2\ell + 1$  to satisfy the following relations:

$$\begin{aligned} \langle a_{\ell m}^{\text{T}*} a_{\ell' m'}^{\text{T}} \rangle &= \delta_{\ell\ell'} \delta_{mm'} C_{\ell}^{\text{TT}}, & \langle a_{\ell m}^{\text{E}*} a_{\ell' m'}^{\text{E}} \rangle &= \delta_{\ell\ell'} \delta_{mm'} C_{\ell}^{\text{EE}}, \\ \langle a_{\ell m}^{\text{B}*} a_{\ell' m'}^{\text{B}} \rangle &= \delta_{\ell\ell'} \delta_{mm'} C_{\ell}^{\text{BB}}, & \langle a_{\ell m}^{\text{T}*} a_{\ell' m'}^{\text{E}} \rangle &= \delta_{\ell\ell'} \delta_{mm'} C_{\ell}^{\text{TE}}. \end{aligned}$$

When measuring the power spectra, we must be aware that what is being measured is just one possible realisation of the CMB signal. In effect, there are many different realisations of the  $C_{\ell}$ s and as such we can only make an estimate on what their true values are. The measured  $\hat{C}_{\ell}$  are actually random variables drawn from a  $\chi^2$  distribution with a variance given by

$$\frac{\Delta \hat{C}_{\ell}}{\hat{C}_{\ell}} = \sqrt{\frac{2}{2\ell + 1}}.$$

This is known as the *cosmic variance*, and limits how well our measurements  $\hat{C}_{\ell}$  across one sky can estimate the true  $C_{\ell}$ s.


### 2.4.2 Temperature Anisotropies

Although the CMB temperature is homogeneous to one part in  $10^5$ , small deviations from the mean are still present. In the very early universe, density fluctuations left over after inflation through a variety of different processes gave rise to these temperature anisotropies. We will explore the various processes that form the characteristic CMB temperature signal.

Non-interacting cold dark matter forms gravitational potentials which attract baryons. As they are compressed, they may pull in more photons due to the increased gravitational potential, releasing excess energy as photons in the process. The radiation pressure from these photons acts as a restoring force, dissipating the baryons and the associated potential. As a consequence, acoustic waves propagate through the medium. Because recombination happens on a shorter timescale than the formation and propagation of these waves, the waves are effectively “frozen” into the CMB when photons finally decouple from the matter. This creates the characteristic acoustic peaks and troughs in the temperature spectrum as seen in Figure 2.4.

Due to the finite lifetime of the universe, sound waves can only travel a finite distance defined by the sound horizon. This horizon behaves like a boundary, meaning that sound waves create stationary modes. Odd peaks in the spectrum correspond to modes in their maximally compressed state, whilst the even peaks correspond to maximally rarefied states. Both compressions and rarefactions produce peaks in the spectrum as compressions heat baryons due to the increase in pressure which as a result blueshifts photons as they fall in to the potential well. Similarly, rarefactions result in a net cooling of photons due to redshifting from potential hills. As such, both result in temperature fluctuations about the mean temperature  $T_0$ .

Alongside the oscillatory behaviour, emitted photons can be red or blueshifted by the Doppler effect. This occurs if there is any bulk motion of matter prior to recombination.



**Fig. 2.4** TT power spectrum showing the acoustic peaks caused by sound waves propagating through the early universe.

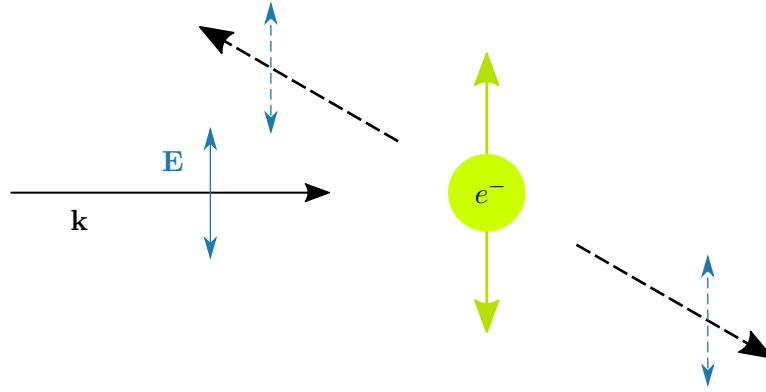
Photons that fall into gravitational potentials are blueshifted to higher energies. This energy gain is lost if it climbs out of the same, unchanged potential. However, due to the constantly changing environment, this gravitational potential can either increase or decrease, so by the time the photon leaves it has a net energy change. This is known as the Sachs-Wolfe effect and contributes to the total anisotropy.

### Spectrum Peaks

The form of Figure 2.4 can be interpreted to give information about the nature of the early universe.

The first peak corresponds to the largest acoustic wave. The size of this wave is constrained by the furthest distance sound can travel in the time it has to do so: the sound horizon. As such, the angular size of this peak corresponds to the angular diameter of the universe at recombination, and its location on the multipole axis puts constraints on the curvature of the universe.

The amount of baryonic matter present in the early universe governs the amount of compression in each wave. A larger quantity of baryons pulls more fluid into potential wells which enhance the compressional peaks. This effect is known as baryon drag, and is why the compression peaks in the temperature spectrum are larger than their consecutive rarefaction peaks.



**Fig. 2.5** Incoming light with  $\mathbf{E} \parallel \hat{\mathbf{z}}$  hitting an electron  $e^-$  will cause the electron to oscillate parallel to the electric field  $\mathbf{E}$ . The electron will emit radiation with the same polarisation preferential to the plane perpendicular to  $\mathbf{E}$ .

Anisotropies at high- $\ell$ , or small angular scales, are suppressed by the diffusion of photons during late recombination known as Silk damping. The resulting photon mixing exponentially dampens relative temperature differences on small scales.

### 2.4.3 Polarisation Anisotropies

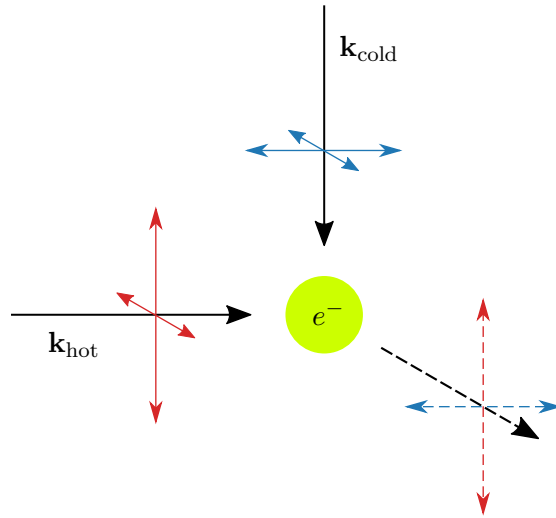
CMB light can also become linearly polarised during its formation. This is done through Thomson scattering, which is the low-energy limit of Compton scattering. In such a case, the electron's kinetic energy and the photon's frequency remain unchanged during the interaction. The differential cross-section for this process is given by

$$\frac{d\sigma}{d\Omega} = \frac{3}{8\pi} |\mathbf{E}' \cdot \mathbf{E}| \sigma_T,$$

where  $\mathbf{E}$ ,  $\mathbf{E}'$  are the polarisations of the incoming and scattered light respectively, and  $\sigma_T$  is the Thomson scattering cross-section. This process preferentially scatters photons perpendicular to the polarisation vector  $\mathbf{E}$ . Figure 2.5 shows how light incident on an electron may be scattered.

A net polarisation can be generated if the intensity of the radiation field exhibits a quadrupole anisotropy. Figure 2.6 shows how cold and hot photons hitting an electron at the same time produces an outgoing photon with component of linear polarisation due to Thomson scattering. Here,  $\mathbf{k}_{\text{hot}}$  corresponds to the hot photon which has increased intensity compared to the cold photon  $\mathbf{k}_{\text{cold}}$ . Due to this difference in intensities, the resulting scattered photon has a net polarisation parallel to  $\hat{\mathbf{z}}$ .

Polarisation can also be generated from vector and tensor effects. In this case, vector effects such as vorticity arising from hot and cold regions shearing can develop polarisation. Similarly, tensor effects such as primordial gravitational waves can also create polarisation.



**Fig. 2.6** An electron surrounded by a quadrupole of unpolarised photons (hot photons along  $x$ -axis, cold along  $z$ -axis) produces a net polarised photon along the  $y$ -axis.

#### 2.4.4 Secondary Anisotropies

After the surface of last scattering, light from the CMB still interacts with gravitational potentials and matter along its path to the observer. These interactions generate secondary anisotropies in the CMB spectrum.

The integrated Sachs-Wolfe (ISW) effect works similarly to the Sachs-Wolfe effect, except the gravitational potentials are encountered during the time taken to travel from the surface of last scattering to the observer. This results in a net integrated result as the CMB photon encounters evolving potentials. This effect is sometimes split into early- and late-time parts. The early-time ISW effect occurs directly after recombination where the radiation-dominated universe has evolving potentials due to the presence of large amounts of radiation. The late-time ISW effect happened much more recently during the dark energy-dominated universe, where gravitational potentials of galaxy clusters are dissipated due to the dark energy driving the galaxies apart.

Photon polarisation is also affected by gravitational potentials. B-mode anisotropies can be generated through weak gravitational lensing, which mixes E-modes with B-modes as photons propagate to the observer. At linear order in CPT, T-, B- and E-modes do not mix, but as gravitational lensing is a non-linear process we see mode mixing.

Alongside the evolution of gravitational potentials, CMB photons can be scattered by high-energy electrons, with the resulting anisotropy known as the Sunyaev-Zel'dovich (SZ) effect. This effect can be due to either thermal electrons (electrons with an emission spectrum similar to a blackbody), or kinetic electrons (those which have been accelerated by other astrophysical means). As such, the SZ effect is commonly referred to as either thermal or kinetic.



## Chapter 3

# Recombination and Dark Matter

### 3.1 Introduction

Observations of dark matter have led to a plethora of different theories. There is, however, an implicit assumption all of these theories make: that dark matter was created alongside the rest of “normal” baryonic matter and radiation during the Big Bang. Certain models of dark matter can introduce candidates that have additional non-gravitational interactions with baryonic matter and radiation. As an example, some WIMP theories of dark matter predict weakly interacting subatomic particles that can annihilate against each other or decay into SM particles.

For models of this interacting dark matter, the early universe could provide the right conditions for the dark matter to inject energy. This could be due to the large densities counteracting small annihilation cross-sections in WIMP models, or the formation of black holes which subsequently evaporate or accrete as in models of primordial black holes. The extra injection of energy during the early universe would upset the standard process of recombination, leaving an imprint in the CMB detectable by experiments such as *Planck*.

This chapter will explore the recombination epoch in further detail in Section 3.2, outlining the two main methods used to model this point in the universe’s past. Section 3.3 covers the extension of these models to incorporate non-standard energy injections from generic dark matter, before Section 3.4 details the specific model of dark matter considered throughout the rest of the work: primordial black holes.

### 3.2 The Recombination Epoch

During the early stages of the universe, hydrogen and helium nuclei were stripped of their electrons due to the extreme heat. It was only when the expanding universe cooled that the electrons combined with the hydrogen and helium nuclei to form neutral atoms. This part of the universe’s history is known as the recombination epoch. It is governed by electron transitions

between the various energy levels of hydrogen and helium atoms, and their interactions with the ambient photon field.

The problem for cosmologists remained how to accurately describe the evolution of this system from a hot plasma to a neutral gas. A full, mathematical description of these processes is vital to understanding the formation of the cosmic microwave background.

It was not until the late 60s that pioneering work by Peebles [5] and Zeldovich, Kurt, and Syunyaev [6] found a way of quantifying this problem. Both groups took a similar approach to solving the problem, which became known as the effective three-level atom. The three-level atom, although basic, provided the first insight into the recombination epoch, and allowed for the calculation of CMB power spectra from a purely theoretical standpoint.

An improvement to this method has recently been proposed, and treats the hydrogen atom in more depth whilst still remaining computationally viable. This is known as the effective multi-level atom.

This section will cover the effects recombination has on the formation of the CMB, and offer an explanation of the two predominant atomic models used to compute recombination.

### 3.2.1 Imprints on CMB

Understanding the recombination epoch is vital to understanding the CMB. The two quantities we are interested in when considering reionisation are the ionisation history and the matter temperature. Here, “matter” only refers to neutral atoms, protons, electrons and other nuclei, as we assume at this point any dark matter component has already dropped out of thermal equilibrium.

The ionisation history is given by the evolution of the free electron fraction  $x_e$ , defined by

$$x_e(z) := \frac{n_e(z)}{n_p(z) + n_H(z)},$$

where  $n_e(z)$ ,  $n_p(z)$ ,  $n_H(z)$  are the number densities of electrons, protons and atomic hydrogen respectively.

Changes to the ionisation history  $x_e(z)$  affect the CMB through changing the visibility function and the amount of Silk damping.

### Visibility Function

The visibility function is defined in terms of the optical depth  $\tau$ , given by Eq. (2.3.5):

$$\tau(z) := \int_0^z \frac{n_e(z') c \sigma_T}{(1+z') H(z')} dz'. \quad (2.3.5)$$

The visibility function  $g(z)$  gives the probability that a photon was last scattered in the interval  $[\tau, \tau + d\tau]$  per redshift:

$$g(z) := e^{-\tau(z)} \frac{d\tau}{dz} = \frac{n_e(z)c\sigma_T}{(1+z)H(z)} e^{-\tau(z)}.$$

$g(z)$  is dependent on the number density of free electrons  $n_e(z)$ , and governs the probability distribution for the last scattering of photons. If photon re-scattering increases through a change in  $n_e(z)$ , the resulting CMB anisotropy should be less than expected as scattering homogenises the temperature differences in the CMB photon field.

The effect  $g(z)$  has on the CMB is more important at low redshifts, as the visibility function is exponentially damped at high redshifts by the optical depth. Hence, increased ionisation leads to a larger tail in the visibility function at low redshifts therefore decreasing the overall amplitude of CMB anisotropies.

### Silk Damping

Before completely decoupling from the plasma at last scattering, photons undergo a random walk as they scatter off the remaining free electrons. The mean free path for photons at redshift  $z > z_{\text{rec}}$  is given by

$$L \approx \frac{1}{n_e(z)\sigma_T}.$$

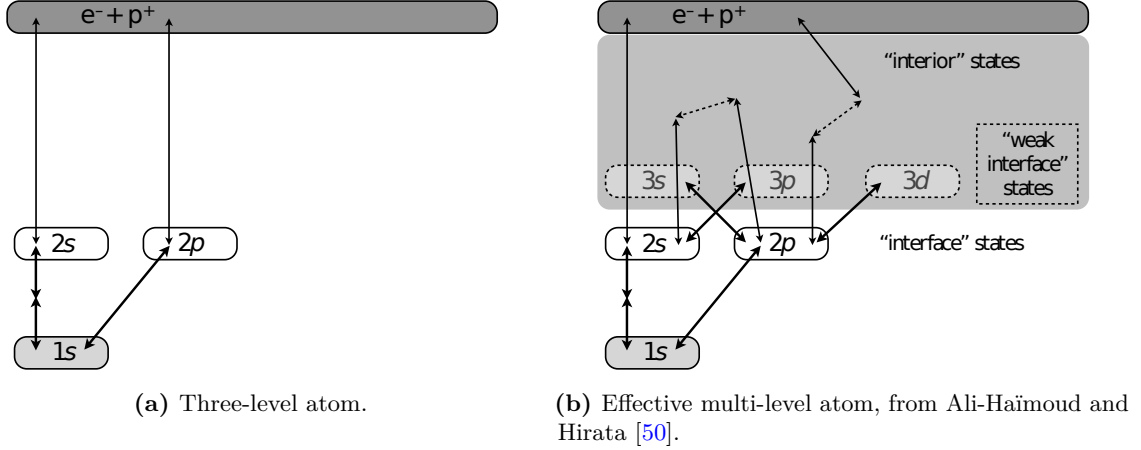
This random diffusion of photons prior to last scattering acts to dampen small scale fluctuations in the temperature power spectrum. This phenomena is known as Silk damping, and is dependent on the electron density and hence the ionisation history. A larger amount of ionisation will reduce the mean free path and decrease the amount of damping in the tail of the spectrum.

#### 3.2.2 Three-Level Atom

The three-level atom (TLA) compacts the infinite energy levels of the hydrogen atom into just three: the ground state, the first excited  $n = 1$  state, and the continuum state at  $n = \infty$ . The model assumes that all higher excited states are in Boltzmann equilibrium with the first excited state and so can be effectively ignored. Here, Boltzmann equilibrium is taken to mean the equilibrium reached when all energy states share the same occupation numbers and transition rates. Figure 3.1a gives a schematic diagram of the levels involved in calculations for the TLA.

Direct recombination to the ground state is inefficient, as doing so will emit a photon which will immediately photoionise another H atom in its ground state. Instead, recombination occurs at excited states, which then quickly cascade down to the 2s and 2p states [51, 52].

Atoms in the 2p state decay to the ground state by Ly- $\alpha$  emission. However, the emitted photon cannot escape easily as it will be re-absorbed by other hydrogen atoms in the ground



**Fig. 3.1** H atom models used to solve the recombination problem.

state not unlike the case of direct recombination. The photon can only escape if it is redshifted away from the resonant Ly- $\alpha$  line before it encounters another hydrogen atom.

Decay to the 2s state is classically forbidden due to the large lifetime of the 2s state compared to the 2p state. During recombination, the transition rate is comparable to the 2p decay to 1s, and results in a double photon decay to the ground state. As the difference in orbital angular momentum is zero between the 2s and 1s levels, two photons need to be released to conserve angular momentum. This process is  $10^8$  times slower than the Ly- $\alpha$  decay, but the photons escape immediately due to the lack of resonances at the photons' energy.

The net result is that both processes happen on comparable timescales, allowing hydrogen to fall into its ground state and become bound and neutral. Such a simplified model results in two coupled differential equations to solve:

$$\frac{dx_e}{dz} = \frac{1}{(1+z)H(z)} \{R_s(x_e, T_m; z) - I_s(x_e, T_m; z)\}, \quad (3.2.1)$$

$$\frac{dT_m}{dz} = \frac{1}{(1+z)H(z)} \left\{ \frac{8\sigma_T a_r T_\gamma^4(z)}{3m_e c} \frac{x_e}{1 + f_{\text{He}} + x_e} [T_m - T_\gamma(z)] + 2H(z)T_m \right\}. \quad (3.2.2)$$

Here,  $\sigma_T$  is the Thomson scattering cross-section,  $a_r = 4\sigma/c$  is the radiation constant,  $T_\gamma(z) = (1+z)T_0$  is the radiation temperature,  $m_e$  is the electron mass and  $f_{\text{He}}$  is the fraction of helium nuclei to hydrogen nuclei.

Eqs. (3.2.1) and (3.2.2) are those used when not considering any extra energy injections from external sources, and hence will be referred to as the standard equations. As there is a lot to unpack in these equations conceptually, we will start by considering Eq. (3.2.1).

### Free Electron Fraction

In (3.2.1),  $R_s$  and  $I_s$  are standard recombination and ionisation terms respectively, and take the form

$$\begin{aligned} R_s(x_e, T_m; z) &= C(x_e, T_m; z) \alpha_B(T_m, T_\gamma) x_e^2 n_H(z), \\ I_s(x_e, T_m; z) &= C(x_e, T_m; z) \beta_B(T_\gamma) (1 - x_e) e^{-h\nu_{2s}/k_B T_m}, \end{aligned}$$

where  $C(x_e, T_m; z)$  is Peebles' coefficient which will be defined later,  $\nu_{2s}$  is the frequency of the  $2s \rightarrow 1s$  transition, and  $k_B$  is the Boltzmann constant.

$\alpha_B(T_m, T_\gamma)$  and  $\beta_B(T_\gamma)$  are recombination and ionisation rates from fully-ionised to excited H states respectively. They both originated from theoretical descriptions of nebulae emission, and describe so-called ‘‘case B’’ emission where emitted Ly- $\alpha$  photons are quickly reabsorbed by surrounding neutral hydrogen atoms. It is quite easy to see why these coefficients have since found great use in recombination physics, as the same process happens in the early universe too.

### Matter Temperature

In Eq. (3.2.2), the first term is a Compton cooling term. It comes from considering the energy transferred from electrons to photons via Compton scattering, originally proposed by Weymann [53]. The rest of the particle content (protons, ionised and neutral hydrogen and helium states) are kept at the same temperature through collisions, so the net cooling effect extends Weymann's by considering the energy of all particles involved in recombination [54].

The second term is an adiabatic cooling term due to the expansion of the universe. Assuming matter forms an ideal gas,  $\gamma = 5/3$  for adiabatic expansion which implies  $T_m \propto (1 + z)^2$ . This implicitly leads to  $dT_m/dz = 2T_m/(1 + z)$ , the term that features in Eq. (3.2.2).

### Peebles' Coefficient

In order to account for the probability of a hydrogen atom being re-ionised by CMB photons before it reaches the ground state, Peebles introduced a coefficient  $C(x_e, T_m; z)$  defined by

$$C(x_e, T_m; z) := \frac{1 + K(z) \Lambda_{2s1s} n_H(z) (1 - x_e)}{1 + K(z) [\Lambda_{2s1s} + \beta_B(T_\gamma)] n_H(z) (1 - x_e)}.$$

In the above,  $\Lambda_{2s1s} = 8.22458 \text{ s}^{-1}$  is the decay rate of the 2s state, and  $K(z)$  is given by

$$K(z) := \frac{\lambda_\alpha^3}{8\pi H(z)},$$

and accounts for the cosmological redshifting of Ly- $\alpha$  photons with wavelength  $\lambda_\alpha$ .

The coefficient can be viewed as the probability that a hydrogen atom in the first excited state transitions to the ground state through either the 2s or 2p transition before being ionised by ambient CMB radiation.

### 3.2.3 Effective Multi-Level Atom

As the precision of various CMB anisotropy experiments increased, the accuracy of the TLA approach was brought into question [50, 51, 55]. In order to increase the accuracy of calculations to match experimental precision, higher order effects needed to be accounted for. This was achieved by extending the TLA model to incorporate such effects.

This extension is known as the multi-level atom (MLA), and improves upon the TLA model in two ways: by incorporating non-equilibrium effects of higher energy levels in the H atom, and by solving for the Ly- $\alpha$  escape probability explicitly and its effect on the 2s $\rightarrow$ 1s transition. The first approach effectively leads to a change in  $\alpha_B(T_m, T_\gamma)$ , whilst the second approach effectively changes Peebles' coefficient  $C(x_e, T_m; z)$ .

Although more accurate, the MLA model is also a lot slower, as the population for each extra excited state considered must be solved for explicitly. Instead, an effective approach was taken by Ali-Haimoud and Hirata [50], known as the effective multi-level atom (EMLA). Compared to the full MLA approach, the EMLA works by merging various higher-order effects into *effective* recombination and ionisation rates. Figure 3.1b gives a schematic diagram of the levels involved in calculations for the EMLA, highlighting the interface states which use effective transition rates between interior states to speed up calculations.

EMLA works by splitting the MLA problem into two parts: the calculation of bound-bound and bound-free transition rates between excited states of hydrogen in a thermal photon field, and the full evolution of the hydrogen atom accounting for these transitions. By integrating out these excited state transitions, the MLA method can be reduced to an effective one where only a handful of excited states need to be considered (*e.g.* 2s, 2p, 3s), and effective transition rates to higher energy levels are used in place of explicit calculations. As the effective rates can be calculated beforehand, the results can be interpolated on-the-fly during the EMLA calculation.

### 3.2.4 Recombination Codes

Many numerical codes have been written to solve for ionisation histories. For this work, we will consider the two codes packaged with the Boltzmann code CLASS: RECFAST and HYREC.

At the onset of this project, RECFAST was used for its apparent resemblance to TLA physics despite it being based on MLA calculations. However, due to its inability to correctly incorporate non-standard DM physics it was eventually discarded in favour of HYREC, a more robust code that accurately accounts for DM physics through the EMLA method.

## RECFAST

RECFAST [54, 55] was the first recombination code to implement an “effective” MLA approach. It computes the ionisation history quickly and accurately through the use of a so-called “fudge factor” that allows it to reproduce results from more accurate but slower MLA codes. It does this by using fudged TLA equations and various approximations [56] to calculate reionisation histories. As a consequence, it does not lend itself to exotic physics extensions. However, RECFAST presents a simple interface with which to examine the internals of TLA and naively extend it.

It was recently extended by Scott and Moss [56] to make use of a matter-radiation temperature equality approximation. By computing a characteristic time at which Compton interactions freeze out, the matter and radiation temperatures can be fixed equal until the matter temperature must be evolved in full.

## HyRec

HYREC [57] implements the true EMLA model, with support for a RECFAST-like model with fudge factors.

HYREC uses various approximation schemes to evolve the system through redshift  $z$ . For high redshifts, the program uses an analytic approximation based on the closeness of helium levels to Saha equilibrium, a so-called post-Saha expansion. It then fully evolves the hydrogen recombination, which incorporates many higher-order effects not included in previous analyses. This part involves solving the radiative transfer equation explicitly using a sparse matrix formulation to evolve the ambient photon field alongside state occupation numbers and the free electron fraction.

HYREC is based on a modular structure, allowing the user to select what model of recombination to use, such as the original Peebles model, a RECFAST-like model and the full EMLA model. HYREC also provides switches for various higher-order interactions included in the radiative transfer equation, so the user can remove such effects if execution speed is needed over accuracy.

## 3.3 Dark Matter as a Non-Standard Energy Source

As astrophysical observations convince us of the existence of dark matter, it seems reasonable to assume that it must have been created in the Big Bang along with all other matter. Depending on when the dark matter freezes out from potential interactions with baryonic matter and radiation, there could be some amount of it around the time of recombination, no longer interacting with the ordinary matter. Assuming there exists an avenue through which the dark matter can inject energy (*e.g.* for WIMP DM: annihilation or decay), this extra energy can

affect recombination. In this section, we will explore how the incorporation of dark matter can impact the ionisation history of the universe through non-standard energy injection.

### 3.3.1 Energy Injection During Recombination

The primary motive for considering recombination in such detail is to extend the processes to incorporate the effects of dark matter. Dark matter in the early universe, depending on what form it takes, can inject extra energy into the primordial photon-baryon plasma.

In order to include these extra effects during recombination, we must parametrise the different ways in which non-standard energy injections can affect the coupled photon-baryon plasma. The plasma can undergo ionisation or excitation at the atomic level in both hydrogen and helium. It can also be heated up through larger scale thermal excitations, and have extra continuum photons injected into the photon field with energies too low for them to interact ( $< 10.2$  eV). These are known as the various deposition channels for the plasma. We now need a way of converting the injected energy into the energy deposited into the plasma.

#### Injection vs Deposition

Energy injected in the early universe can come from a wide variety of sources, and consist of an array of different particles. The fine details of how this energy and matter interact with the primordial plasma can be abstracted away by considering the fraction of injected energy that is actually deposited into the plasma,  $f_{X,c}(z)$ , for each deposition channel  $c$ . This is known as an effective deposition efficiency, and is defined by the following:

$$\left(\frac{dE}{dV dt}\right)_{c,\text{dep}} := f_{X,c}(z) \left(\frac{dE}{dV dt}\right)_{\text{inj}},$$

where energy injection is given as a rate per unit co-moving volume per unit cosmic time. We also introduce the subscript  $X$  in this work to explicitly indicate that these efficiencies originate from a non-standard source  $X$ .

One could argue that this differential quantity should be written with a  $d^2E$  on the numerator, however throughout this work we use the convention of taking multiple differential quantities in the denominator to outline what the numerator is defined in terms of. Looking back at our definition of the energy injection, the inclusion of  $dV dt$  is simply there to indicate the quantity is defined per unit co-moving volume per unit cosmic time.

To simplify matters, the deposition efficiency was first approximated by assuming energy is deposited into the plasma at the same time it is injected [11, 58, 59]. This is known as *on-the-spot* energy injection. In such a case, deposited energy is given as a fixed fraction of the total injected energy and redshift-dependent effects are not taken into account. This approximation is only valid for high redshifts where cooling processes are efficient enough that energy deposition is effectively instantaneous. However, this assumption breaks down when

considering energy injections at  $z \sim 1000$ , where the photon cooling rate becomes comparable to the Hubble time [60].

More recently, work done by Slatyer [7, 8, 9] aims to circumvent these assumptions by simulating energy deposition for electrons, positrons and photons from first principles. This work considers the fine details of electron/positron and photon interactions in the early universe, and presents effective deposition efficiencies that can be used to compute ionisation histories. As the only difference between an electron and positron is the sign of their charges, any property of electrons in the context of particle radiation and energy deposition is equivalent for positrons, and as such we will not mention positrons explicitly in the interests of brevity.

In general, energy injections in the form of particle showers consist of more than just electrons and photons. However, other particles such as gauge bosons, baryons and mesons produce both electromagnetic and hadronic showers. The final states may include neutrinos, which are stable and weakly interacting so escape easily; and protons, which ineffectively transfer energy to the plasma and so can be ignored [58]. This work closely follows the analysis taken by Clark et al. [1], which neglects the effects of hadrons as primordial black holes with masses greater than  $10^{15}$  g primarily radiate electrons and photons.

### 3.3.2 Extension of TLA Model

For this work, we consider only three energy deposition channels: hydrogen ionisation and excitation, and heating. Helium exists in a much smaller fraction relative to hydrogen, and becomes neutral long before hydrogen recombines. As such, it does not greatly influence the reionisation history and can be effectively ignored when considering non-standard energy injections.

In the presence of non-standard energy injections, the effective TLA equations become [58]

$$\begin{aligned} \frac{dx_e}{dz} &= \left( \frac{dx_e}{dz} \right)_{\text{std}} - \frac{1}{(1+z)H(z)} [I_{X_i}(x_e, T_m; z) + I_{X_\alpha}(x_e, T_m; z)], \\ \frac{dT_m}{dz} &= \left( \frac{dT_m}{dz} \right)_{\text{std}} - \frac{2}{3k_B(1+z)H(z)} \frac{K_h(x_e, T_m; z)}{1 + f_{\text{He}} + x_e} \end{aligned}$$

The non-standard terms have been split up into their respective channels. They are given by

$$I_{X_i}(x_e, T_m; z) = f_{X,i}(z) \frac{(dE/dV dt)_{\text{inj}}}{n_{\text{H}}(z)E_i}, \quad (3.3.1)$$

$$I_{X_\alpha}(x_e, T_m; z) = [1 - C(x_e, T_m; z)] f_{X,\alpha}(z) \frac{(dE/dV dt)_{\text{inj}}}{n_{\text{H}}(z)E_\alpha}, \quad (3.3.2)$$

$$K_h(x_e, T_m; z) = f_{X,h}(z) \frac{(dE/dV dt)_{\text{inj}}}{n_{\text{H}}(z)}, \quad (3.3.3)$$

where we have included the energy deposition efficiencies  $f_{X,c}(z)$  for each channel  $c$ , and the total injected energy per unit volume per unit cosmic time  $(dE/dV dt)_{\text{inj}}$ . Here, H ionisation is given by the label i, H excitation by  $\alpha$  and heating by h. The energies used in the equations are the hydrogen ionisation energy  $E_i$  and the Lyman- $\alpha$  transition energy  $E_\alpha$ . The deposition efficiencies  $f_{X,c}(z)$  are taken from the Slatyer results. We will revisit this in Section 5.2, where we will explicitly follow the steps taken in using the Slatyer results.

### 3.4 Primordial Black Holes

Primordial black holes (PBHs) are black holes formed early in the universe's history from the gravitational collapse of large over-densities of matter. Originally proposed as a dark matter candidate in the mid '70s by Chapline [61], PBHs enjoyed a period of inattention before recently coming back in vogue through measurements by LIGO of a binary black hole merger [14, 15, 16, 17, 18]. PBHs are an interesting DM candidate to consider due to their heavily constrained parameter space. Astrophysical observations such as diffuse gamma ray background measurements [62, 63] and microlensing surveys [64] exclude some regions of PBH mass. However, there are two mass windows that have still not been fully excluded, both of which can potentially inject energy during recombination.

The first spans masses of the order of solar masses, from  $10 M_\odot$  to  $10^2 M_\odot$ . This high-mass region features black holes that inject energy through the formation of an accretion disk [19].

The second, lighter region contains PBHs with masses in the range  $10^{15}$  g to  $10^{17}$  g (roughly  $10^{-19} M_\odot$  to  $10^{-17} M_\odot$ ). Heavier PBHs are better constrained by femtolensing observations, whilst PBHs with  $M_{\text{PBH}} \lesssim 5 \times 10^{14}$  g would have evaporated by today, and so are not considered [65].

In order to maintain simplicity throughout this work, we will be only dealing with the low-mass range, and consider energy injection from black hole evaporation. We will not be considering the solar mass region due to the complexity of calculating radiation luminosity from accretion disks.

#### 3.4.1 Creation

Depending on the model of PBH considered, there are many different ways that overdense regions in the early universe could have formed. Models due to inflation [66, 67], pre-heating after inflation [68, 69], first-order phase transitions [70] and curvatures [71, 72] have all been put forth. For this work, we consider inflation-driven models, where an inflationary epoch expanded quantum fluctuations into macroscopic mass density fluctuations which later collapsed into black holes.

PBH analyses up until recently have implicitly assumed a monochromatic mass distribution, *i.e.* PBHs of a single fixed mass. However, it is not reasonable to expect that actual PBHs

generated from early universe mechanisms will all have the same mass. For instance, different inflation mechanisms can result in mass distributions with non-zero variance [22]. We will be considering the effects these extended mass distributions have on constraints for PBHs later on.

### 3.4.2 Evaporation

The idea of black hole evaporation was first proposed by Hawking [73, 74], and originated from purely thermodynamical arguments. The phenomenon became known as Hawking radiation, which for a black hole with no charge or angular momentum relates the mass of a black hole to its temperature:

$$k_{\text{B}}T_{\text{BH}} := \frac{\hbar c^3}{8\pi G M_{\text{BH}}}. \quad (3.4.1)$$

For PBHs, it is reasonable to assume the black holes themselves do not have charge or angular momentum, as these will be lost through particle emission quicker than the mass.

The temperature is that of a black body, which fixes the amount of energy released by the black hole according to the black body spectrum. Unlike a black body which traditionally only emits photons, black holes emit different particle species depending on how massive they are. Simply speaking, this is because higher black hole temperatures lead to larger ejection energies, and heavier particle species are ejected as a result as there is more energy with which to produce them. As a consequence, black holes lose mass.

The rate of change for black hole mass is given by [65]

$$\frac{dM_{\text{BH}}}{dt} = -5.34 \times 10^5 \times \left[ \sum_i f_i^{\text{em}}(M_{\text{BH}}) \right] \left( \frac{10^{10} \text{ g}}{M_{\text{BH}}} \right)^2 \text{ g s}^{-1}, \quad (3.4.2)$$

where the sum is over particle emission fractions  $f_i^{\text{em}}(M_{\text{BH}})$ , including relativistic particle degrees of freedom. For example, if electrons were being considered, a factor of four would be included alongside  $f_{e^\pm}^{\text{em}}(M_{\text{BH}})$  due to the two spin states and the matter/antimatter duality. The sum over  $f_i^{\text{em}}(M_{\text{BH}})$  is normalised to unity for a black hole with mass  $M_{\text{BH}} \gg 10^{17} \text{ g}$ .

### 3.4.3 Emission Spectra

Eq. (3.4.1) gives a good indication of what particle energies are ejected from a black hole when compared to Planck's law for black bodies. However, the specific energy spectrum for particle emission can also be calculated separately. For a general black hole, the ejection rate of particles per degree of particle freedom with total energies between  $E$  and  $E + dE$  is given by

$$\frac{dN}{dt} = \frac{\Gamma_s dE}{2\pi\hbar} \frac{1}{\exp\left(\frac{E}{k_{\text{B}}T_{\text{BH}}}\right) - (-1)^{2s}}, \quad (3.4.3)$$

where  $s$  is the particle's spin and  $\Gamma_s$  is its associated absorption coefficient.

The average energies of emitted electrons, photons and neutrinos for this spectrum are

$$\bar{E}_{e^\pm} = 4.18k_B T_{\text{BH}}, \quad \bar{E}_\gamma = 5.71k_B T_{\text{BH}}, \quad \bar{E}_\nu = 4.22k_B T_{\text{BH}}. \quad (3.4.4)$$

The peaks of the spectra are within 7% of these values [75], so using the average energies is a good approximation to the full spectra.

### Emission Fractions

In order to compute the mass loss in Eq. (3.4.2), we need to know the particle emission fractions  $f_i^{\text{em}}(M_{\text{BH}})$ . For the mass range under consideration, emitted particles are effectively massless and the emission fractions for each particle are given as follows [76]:

$$f_{e^\pm}^{\text{em}} = 0.142, \quad f_\gamma^{\text{em}} = 0.060, \quad f_\nu^{\text{em}} = 0.147.$$

#### 3.4.4 Energy Injection

The total Hawking radiation injected by PBHs is given by the energy injected per PBH multiplied by the number density of PBHs in the universe:

$$\left( \frac{dE}{dV dt} \right)_{\text{PBH}} := - \frac{dM_{\text{PBH}}}{dt} c^2 \times n_{\text{PBH}}(z).$$

As mass is being lost by the PBH, the energy injected into the medium must be positive, hence the negative sign in the above equation. Noting that the number density of PBHs can be rewritten in terms of its density parameter  $\Omega_{\text{PBH}}(z)$ :

$$n_{\text{PBH}}(z) = \frac{\rho_c \Omega_{\text{PBH}}(z)}{M_{\text{PBH}}},$$

where the density  $\Omega_{\text{PBH}}(z)$  evolves in the same way as matter:

$$\Omega_{\text{PBH}}(z) = \Omega_{\text{PBH},0}(1+z)^3,$$

we have

$$\left( \frac{dE}{dV dt} \right)_{\text{PBH}} = - \frac{\dot{M}_{\text{PBH}}}{M_{\text{PBH}}} c^2 \rho_c \Omega_{\text{PBH},0} (1+z)^3,$$

using the notation  $\dot{x} \equiv dx/dt$ .

Defining the fraction of dark matter that consists of PBHs as

$$f_{\text{PBH}} := \frac{\Omega_{\text{PBH},0}}{\Omega_{\text{cdm},0}},$$

and using Eq. (3.4.2) we can finally write

$$\left(\frac{dE}{dV dt}\right)_{\text{PBH}} = 5.34 \times 10^{-5} \times \left[\sum_i f_i^{\text{em}}(M_{\text{PBH}})\right] \left(\frac{10^{10} \text{ g}}{M_{\text{PBH}}}\right)^3 c^2 \rho_c \Omega_{\text{cdm},0} f_{\text{PBH}} (1+z)^3 \text{ s}^{-1}. \quad (3.4.5)$$

Given the PBH mass, the fraction of dark matter made of PBHs  $f_{\text{PBH}}$ , and the current density of dark matter  $\Omega_{\text{cdm},0}$ , the combined Hawking radiation of all PBHs can be calculated. In order to fully calculate the effect this extra radiation has on the CMB, deposition efficiencies also need to be calculated.

### 3.4.5 Deposition Efficiencies

We follow the simple approach outlined by Clark et al. [1] to calculate the deposition efficiencies for PBHs. By only considering injected electrons and photons from PBHs, the deposition efficiencies are calculated by weighting the particles' individual deposition efficiencies  $f_{\text{eff},c}^{e^\pm}(E_e, z)$  and  $f_{\text{eff},c}^\gamma(E_\gamma, z)$  by their particle degrees of freedom as mentioned previously in Eq. (3.4.2). The particle efficiencies are given by Slatyer [8], with the resulting deposition efficiencies for PBHs then given by

$$f_{\text{PBH},c}(M_{\text{PBH}}, z) = \frac{4f_{e^\pm}^{\text{em}} f_{\text{eff},c}^{e^\pm}(E_e, z) + 2f_\gamma^{\text{em}} f_{\text{eff},c}^\gamma(E_\gamma, z)}{4f_{e^\pm}^{\text{em}} + 2f_\gamma^{\text{em}}}. \quad (3.4.6)$$



## Chapter 4

# Sampling $\Lambda$ CDM Parameter Space

### 4.1 Introduction

When presented with a theoretical model describing a non-standard extension of a well-tested theory, it is well within reason to want to know how viable such a model is when compared to observations. In order to accomplish this, we turn to statistical inference. For this work we will be dealing exclusively with Bayesian inference.

To begin with, a summary of the MULTINEST sampling algorithm is presented in Section 4.2, with an outline of Bayesian inference. Section 4.3 describes the likelihoods used by *Planck*. Section 4.4 then outlines the creation and structure of `pc_multinest`, a driver program written as part of this work which links the *Planck* likelihood code together with the cosmological code CLASS and the sampling program MULTINEST to estimate parameter values.

### 4.2 MultiNest

MULTINEST is an implementation of the multimodal nested sampling algorithm [77, 78], written by Feroz, Hobson, and Bridges [79]. The nested sampling technique was introduced as an efficient way of evaluating the Bayesian evidence, but as a by-product also produces posterior samples. These concepts will be defined in the following section.

#### 4.2.1 Bayesian Parameter Estimation

In the field of statistics, parameter estimation encompasses many different methods, of which most are based in some way on the concept of a *likelihood*, or how well we expect the theory to fit the observed data. Parameter estimation aims to compute the probability that parameters of a given model have a given value. The interpretation of what this probability represents gives rise to the two schools of thought: Bayesian and frequentist statistics. The likelihood is defined independently of these differing interpretations, and is the probability of obtaining the data given a particular model and its set of parameters.

Bayesian inference requires knowledge of a model with parameters  $\boldsymbol{\theta}$ , and a set of data  $D$  to compare the model to. The machinery of Bayesian inference is governed by Bayes' theorem, written as:

$$P(\boldsymbol{\theta}|D) = \frac{\mathcal{L}(\boldsymbol{\theta}; D) \pi(\boldsymbol{\theta})}{Z(D)}. \quad (4.2.1)$$

Here,  $\mathcal{L}(\boldsymbol{\theta}; D)$  is the aforementioned likelihood,  $\pi(\boldsymbol{\theta})$  is the *prior* distribution and  $Z(D)$  is the *evidence*. The quantity on the left of Eq. (4.2.1),  $P(\boldsymbol{\theta}|D)$ , is known as the *posterior* distribution.

The likelihood can be loosely seen as the goodness for which the chosen parameters  $\boldsymbol{\theta}$  can replicate the data. In the case of this work, it describes how well the chosen model parameters fit a selection of power spectra as observed by *Planck*.

The prior is, strictly speaking, the initial probability that the parameters will have some value, defined independently of the data. In practice, this becomes a probability distribution that is used to draw samples from the posterior. It is the one quantity in Eq. (4.2.1) that can be changed by the user, and as such must be carefully chosen to not add extra information to the posterior which would bias the posterior. Such priors are known as *non-informative* priors.

The evidence has an analytical expression given by the law of total probability:

$$Z(D) := \int_{\boldsymbol{\theta}} P(D|\boldsymbol{\theta}) P(\boldsymbol{\theta}) d^n \boldsymbol{\theta} \equiv \int_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}; D) \pi(\boldsymbol{\theta}) d^n \boldsymbol{\theta},$$

where  $\int_{\boldsymbol{\theta}} d^n \boldsymbol{\theta}$  is the  $n$ -dimensional integral over all  $n$  free parameters  $\boldsymbol{\theta}$ . The form of the evidence as it appears in Bayes' theorem can also be seen as a normalisation condition, forcing

$$\int_{\boldsymbol{\theta}} P(\boldsymbol{\theta}|D) d^n \boldsymbol{\theta} = 1.$$

It is common to consider the natural logarithm of the likelihood  $\ln \mathcal{L}$  when computing posteriors. As the logarithm is a strictly increasing function of its argument, the logarithm of a function will have its maximum at the same value as the original function. As such, it is suitable to consider the log likelihood in place of the plain likelihood, especially since the likelihood can be very close to zero for large regions of parameter space. Taking the logarithm of this value makes it easier for numerical methods to converge on a maximum likelihood, as the difference across points in the local neighbourhood is no longer negligible.

Another benefit of considering the logarithm of the likelihood is that for a Gaussian likelihood, the logarithm is quadratic in its parameters. The corresponding Taylor expansion is then finite, and this makes it easier to find the maximum analytically for such problems where this is possible.

### 4.2.2 Nested Sampling

For problems where the likelihood  $\mathcal{L}(\boldsymbol{\theta}; D)$  does not take an analytic form, numerical approaches using sampling algorithms are used. The challenge for sampling algorithms is to sample the posterior  $P(\boldsymbol{\theta}|D)$  using the likelihood efficiently (which may be costly to evaluate).

In the field of cosmology where many problems are unimodal in nature, Markov chain Monte Carlo (MCMC) sampling methods are predominantly chosen as the sampling method. This is because MCMC sampling is able to efficiently compute the posterior for problems that are unimodal. However, it is non-trivial to define when a Markov chain has converged, and many convergence diagnostics can sometimes fail to detect convergence failure [80].

MULTINEST was instead chosen for this project for its concrete convergence criterion. It also performs better for multimodal problems which would otherwise be a lot harder for MCMC methods. MULTINEST achieves this by using nested sampling.

Nested sampling works by creating nested regions in parameter space defined by surfaces of constant likelihood (*iso-likelihood surfaces*), where each region iteratively becomes smaller before eventually converging on the maximum likelihood. The method continues to work when multiple local maxima (or modes) are present in the posterior function, due to the nested nature of the algorithm.

It is based on reducing the multidimensional integration of the evidence to one dimension. This is done by defining the *prior volume*  $X$ :

$$X(\lambda) := \int_{\mathcal{L}(\boldsymbol{\theta}) > \lambda} \pi(\boldsymbol{\theta}) \, d^n \boldsymbol{\theta}.$$

The volume takes on values in the range  $[0, 1]$ , as we assume that the prior  $\pi(\boldsymbol{\theta})$  is properly normalised. The volume  $X(\lambda)$  is one when the value of  $\lambda$  is low enough to encapsulate the full likelihood region of interest. It tends to zero as  $\lambda$  increases, until it finally goes to zero when  $\lambda$  reaches the maximum likelihood value.

The evidence is then given by

$$Z = \int_0^1 \mathcal{L}(X) \, dX,$$

where we have now written the likelihood as a function of the prior volume  $X$ .

By evaluating likelihoods  $\mathcal{L}_j = \mathcal{L}(X_j)$  such that consecutive prior volumes are decreasing:

$$1 = X_0 > X_1 > \dots > X_M > 0,$$

then the evidence can be evaluated through numerical integration. Generally, this corresponds to evaluating

$$Z = \sum_{j=1}^M \mathcal{L}_j w_j,$$

for some weights  $w_j$ . For instance, for integrals evaluated through the trapezoidal rule,  $w_j = (X_{j-1} - X_{j+1})/2$ .

### 4.2.3 Prior Handling

MULTINEST draws a sample of  $N$  so-called “live points”, which it uses to explore the posterior. The number of live points generated is an option given by the user.

The success of nested sampling relies on the assumption that each consecutive step draws samples from a decreasing prior volume. In practice, this sampling is done by imposing the condition  $\mathcal{L} > \mathcal{L}_j$  at each sampling step  $j$ . As the prior volume is defined to be the interior of an iso-likelihood surface, then consecutively smaller volumes correspond to larger likelihood values. After drawing each live point from the current prior volume, MULTINEST sets the next likelihood bound to the minimum likelihood of the current set of live points. This acts to decrease the prior volume over time.

MULTINEST works by sampling uniformly from a unit hypercube, *i.e.* a cube in  $n$  dimensions with each side spanning the range  $[0, 1]$ . The collection of live points that are sampled from the hypercube are related to the free parameters of the likelihood through a transformation. It is up to the user to transform these samples  $u_i$  from this unit space to the physical priors chosen by the user such that the likelihood can be properly evaluated. For a separable prior where each parameter’s prior is independent of other parameters, we may write the full prior as

$$\pi(\boldsymbol{\theta}) = \pi(\theta_1)\pi(\theta_2)\cdots\pi(\theta_i)\cdots\pi(\theta_n).$$

The physical value of each parameter  $\theta_i$  is then given by solving

$$u_i = \int_{-\infty}^{\theta_i} \pi(\theta'_i) d\theta'_i,$$

for each uniformly sampled  $u_i \sim U(0, 1)$ .

In the case of a uniform prior with  $\theta_i \sim U(a_i, b_i)$ , the above reduces to

$$\theta_i = a_i + (b_i - a_i)u_i. \quad (4.2.2)$$

More complex priors can also be transformed this way, but many, such as the Gaussian priors used by *Planck*, do not have analytic expressions. For such a prior  $\pi(\theta_i)$ , this pitfall is avoided by instead reformulating both the likelihood and prior:

$$\mathcal{L}(\boldsymbol{\theta}) \rightarrow \mathcal{L}'(\boldsymbol{\theta}) = \mathcal{L}(\boldsymbol{\theta})\pi(\theta_i), \quad \pi(\theta_i) \rightarrow \pi'(\theta_i) = \text{const.}$$

This reformulation correctly accounts for the extra information given by more complex priors. However, posterior points are instead sampled uniformly as opposed to being sampled proportional to their original prior.

In terms of the log likelihood, this transformation can be written as

$$\ln \mathcal{L}(\boldsymbol{\theta}) \rightarrow \ln \mathcal{L}(\boldsymbol{\theta}) + \ln \pi(\theta_i),$$

which for a Gaussian prior  $\theta_i \sim N(\mu, \sigma)$  gives

$$\ln \mathcal{L}(\boldsymbol{\theta}) \rightarrow \ln \mathcal{L}(\boldsymbol{\theta}) - \frac{(\theta_i - \mu)^2}{2\sigma^2}, \quad (4.2.3)$$

up to some additive constant which can be ignored.

### Ellipsoidal Sampling

MULTINEST also needs a way of imposing  $\mathcal{L} > \mathcal{L}_j$  when sampling the likelihood at a given sampling step  $j$ . To maintain computational efficiency, this is done by approximating the iso-likelihood surfaces by multidimensional ellipses. The ellipses are constructed from the covariance matrix computed from the current set of live points.

In general, a likelihood might not have perfectly ellipsoidal iso-surfaces. MULTINEST accounts for this by allowing the user to define a *sampling efficiency*  $e$  such that the prior volume encapsulated by the ellipse is increased by a factor of  $1/e$  after each step, *i.e.*  $X \rightarrow X' = X/e$ . The efficiency typically lies in the range of  $[0, 1]$ , but a value of  $e > 1$  may be chosen to quickly find best-fit parameters as a first estimate.

In order to account for overlapping ellipses, a point residing in  $n_e$  ellipses has a probability of  $1/n_e$  of being kept for that run.

#### 4.2.4 Convergence

As mentioned previously, MULTINEST has a concrete convergence criterion that arises naturally from the computation of the evidence. At any given step  $i$  in the iteration, the maximum change in the evidence possible from the remaining posterior volume is  $\Delta Z_i = \mathcal{L}_{\max} X_i$ , where  $\mathcal{L}_{\max}$  is the maximum likelihood of the current set of live points and  $X_i$  is the current prior volume of the step. Convergence is met when this value falls below some user-defined value.

## 4.3 *Planck* Likelihood Code

Now we have chosen a sampling algorithm, we also need a likelihood function to use in the posterior. This work uses the *Planck* likelihood code created by the *Planck* Collaboration, which uses data from the *Planck* satellite to compute the likelihood. An overview of the process is presented to provide context for the use of different *Planck* likelihoods later on.

The creation of the likelihood begins with all-sky measurements across four microwave bands by the *Planck* satellite. Taking these data and turning them into a likelihood function is a highly-involved and non-trivial task that has been already carried out by the *Planck*

Collaboration. Many foreground contributions need to be accounted for before a dataset (map) containing only the CMB signal + noise can be produced. From this, a method of evaluating the likelihood must be chosen that is still computationally viable. The result, the *Planck* likelihood code, is used in this work to place constraints on extensions of the  $\Lambda$ CDM model.

This section will cover in some detail how the *Planck* likelihood is computed, including how the raw data are turned into signal maps. It will also explore the various likelihoods available, and which ones were used in this work. This will involve taking a closer look at the parameters used by *Planck* in formulating the likelihoods with a quick introduction to nuisance parameters.

### 4.3.1 Map Making

Before any likelihood calculations can be made, the raw data from the *Planck* satellite need to be turned into a map of pixels. This process needs to remove any foreground emission from the galactic plane, interstellar dust and point sources. *Planck* does this by using the *Commander* code [81, 82], which uses a Bayesian approach to fit various astrophysical models to effectively separate the CMB signal from the astrophysical foreground emission. A low-resolution *Commander* map forms the basis for the low- $\ell$  likelihood, which will be discussed later.

In order to fit the data in this way, *Planck* uses the following parametric expression for the signal  $\mathbf{s}_\nu(\boldsymbol{\theta})$  at some frequency  $\nu$  with parameters  $\boldsymbol{\theta}$  [83]:

$$\mathbf{s}_\nu(\boldsymbol{\theta}) = g_\nu \sum_{i=1}^{N_{\text{comp}}} F_\nu^i \mathbf{a}_i + T_\nu \mathbf{m}_\nu,$$

where  $g_\nu$  is a calibration factor,  $\mathbf{a}_i$  are amplitude maps for component  $i$  and  $\mathbf{m}_\nu$  is a set of template correction amplitudes. The matrix  $T_\nu$  contains the templates themselves, and the other matrix  $F_\nu^i$  accounts for the effects of spectral changes in a given component  $i$ . The templates are simply theoretical distributions of different emission types, such as synchrotron and inverse Compton.

For the data vector  $\mathbf{d}_\nu$  given by the sum of the signal  $\mathbf{s}_\nu(\boldsymbol{\theta})$  and Gaussian-distributed noise  $\mathbf{n}_\nu$ , the likelihood is given by

$$\mathcal{L}(\mathbf{a}_i, \beta_i, g_\nu, \mathbf{m}_\nu, \Delta_\nu) \propto \exp\left(-\frac{1}{2}[\mathbf{d}_\nu - \mathbf{s}_\nu(\boldsymbol{\theta})]^\top N_\nu^{-1}[\mathbf{d}_\nu - \mathbf{s}_\nu(\boldsymbol{\theta})]\right),$$

where  $N_\nu$  is the noise covariance matrix.

With some choice of parameters, the fit is done using Gibbs sampling and returns the signal  $\mathbf{s}$  alongside the various astrophysical foregrounds. This is then used to create the maps which become part of the likelihood calculation.

### 4.3.2 Likelihoods

As outlined previously, *Planck* divides the sky up into pixels in order to quantify its measurements. Here, a pixel denotes a measurement of some quantity (temperature anisotropy, polarisation, *etc.*) corresponding to a point in the sky. In general, the data vector  $\mathbf{m}$  is given by

$$\mathbf{m}^X = \mathbf{s}^X + \mathbf{n}^X,$$

where  $\mathbf{s}^X$  is the signal vector containing pixels indexed by spherical harmonic indices  $\ell, m$  which are assumed to Gaussian-distributed. Similarly,  $\mathbf{n}^X$  is the instrument noise, also assumed to be Gaussian-distributed. The label  $X$  can be T, E, or B for temperature, and E- and B-mode polarisations respectively.

The signal vectors have corresponding auto- and cross-correlated power spectra  $C_\ell^{XY}$  and a covariance matrix given by

$$S(C_\ell) = \sum_{\ell=2}^{\ell_{\max}} \sum_{XY} C_\ell^{XY} P_\ell^{XY},$$

where the spectra under consideration are restricted to only TT, TE, EE and BB, and the quantity  $P_\ell^{XY}$  is a sum over Legendre polynomials weighted by the instrumental beam.

We again assume the corresponding noise covariance matrix  $N$  is Gaussian-distributed. This allows us to compute the full data covariance matrix given by  $M = S + N$ .

The full pixel-based likelihood is then [84]

$$\mathcal{L}(C_\ell) \equiv P(\mathbf{m}|C_\ell) = \frac{1}{2\pi\sqrt{\det M}} \exp\left(-\frac{1}{2} \mathbf{m}^\top M^{-1} \mathbf{m}\right). \quad (4.3.1)$$

The aim for pixel-based likelihoods is to invert this  $n \times n$  covariance matrix  $M$ , where  $n$  is the number of pixels in the experiment. In general, this approach is computationally demanding due to the large number of pixels used and the cubic scaling of matrix inversion computation. To maintain computational viability of the likelihood, *Planck* splits the likelihood evaluation across two regions: the high- $\ell$  region ( $\ell \geq 30$ ) and the low- $\ell$  region ( $\ell < 30$ ). We will briefly detail the approaches used in both the low- $\ell$  and high- $\ell$  cases to mitigate this problem.

#### Low- $\ell$ Likelihood

Low multipoles correspond to large angular scales. Because of this, the number of pixels needed for adequate resolution of features at these scales is markedly less than for high multipoles. For *Planck*, the total number of pixels across all TT, TE and EE spectra is  $n = 9216$  [84]. This allows the likelihood given in Eq. (4.3.1) to be computed by brute force.

The low- $\ell$  likelihood used in this work incorporated both temperature and polarisation maps. This allows both  $A_s$  and  $\tau$  to be estimated, as the temperature and polarisation maps alone cannot constrain them [84].

### High- $\ell$ Likelihood

For high- $\ell$ , the covariance matrix inversion is no longer tractable as on such scales the *Planck* maps contain of the order of  $5 \times 10^7$  pixels [85]. More pixels are needed in order to describe the smaller angular scales associated with higher multipole. Instead, a theoretical likelihood using Gaussian fields is used together with pseudo- $C_\ell$ s.

For an observed spectra matrix  $\hat{C}_\ell$  of  $n$  Gaussian fields, the log likelihood function may be written as [86]

$$-\ln \mathcal{L}(C_\ell | \hat{C}_\ell) = \frac{2\ell + 1}{2} \left\{ \text{Tr} [\hat{C}_\ell C_\ell^{-1}] - \ln \det(\hat{C}_\ell C_\ell^{-1}) - n \right\}. \quad (4.3.2)$$

The pseudo-spectra are computed such as to compress the full spectra with minimal information loss. As an example, the pseudo-spectrum for the temperature map is calculated using the following spherical coefficients [85]

$$\tilde{a}_{lm}^T = \sum_s \Delta T_s w_s^T \Omega_s Y_{lm}(\theta_s),$$

where the sum is over pixels  $s$  in the map. Here,  $\Delta T_s$  and  $w_s^T$  are respectively the temperature anisotropies and weighting functions for each pixel, and  $\Omega_s$  is the area of each pixel.

The resulting pseudo-spectra contain less points, and so the likelihood as given in Eq. (4.3.2) can be computed quicker. This method is known as the `plik` likelihood.

#### 4.3.3 Parameters

When fitting the data using `Commander`, *Planck* adds extra parameters to account for astrophysical foreground amplitudes and instrument uncertainties. These are termed “nuisance” parameters as they are necessary to the evaluation of the likelihood, but hold no influence over the physical parameters of interest.

When making PBH results, we use likelihoods where these nuisance parameters have already been marginalised out. In some settings these nuisance parameters may be of interest in constraining galactic dust emission, and the interested reader is referred to the relevant *Planck* paper [84].

## 4.4 `pc_multinest`

`pc_multinest` was written for this work to interface three programs: the *Planck* likelihood code (PLC), the cosmological code `CLASS` and the sampling code `MULTINEST`. Its name comes from the concatenation of its predecessor `plc_class`, and the parameter sampler `MULTINEST`. `plc_class` was written at the onset of this work as a proof-of-concept interface between both

```

enum param_t {
    free1,
    free2,
    ...
    UP_TO_FREE_PARAMS,
    fixed1 = UP_TO_FREE_PARAMS,
    fixed2,
    ...
    UP_TO_FIXED_PARAMS,
    derived1 = UP_TO_FIXED_PARAMS,
    derived2,
    ...
    TOTAL_PARAM_AMT,
    FREE_PARAM_AMT = UP_TO_FREE_PARAMS,
    FIXED_PARAM_AMT = UP_TO_FIXED_PARAMS - UP_TO_FREE_PARAMS,
    DERIVED_PARAM_AMT = TOTAL_PARAM_AMT - UP_TO_FIXED_PARAMS
};

```

**Listing 4.1** Schematic representation of parameter enumeration. Here, parameters are defined in blocks associated with their type: free, fixed or derived. The last block defines the number of parameters in each type.

PLC and CLASS. `pc_multinest` is written in C++ to utilise the CLASS wrappers, but otherwise does not make use of the object oriented paradigm in its implementation.

This section will explain in detail the design choices made in implementing `pc_multinest`, including how model parameters are handled and how the likelihood is calculated. Figure 4.1 summarises the caller hierarchy for each function in `pc_multinest`, and serves as an indication for what processes occur in which parts of the code.

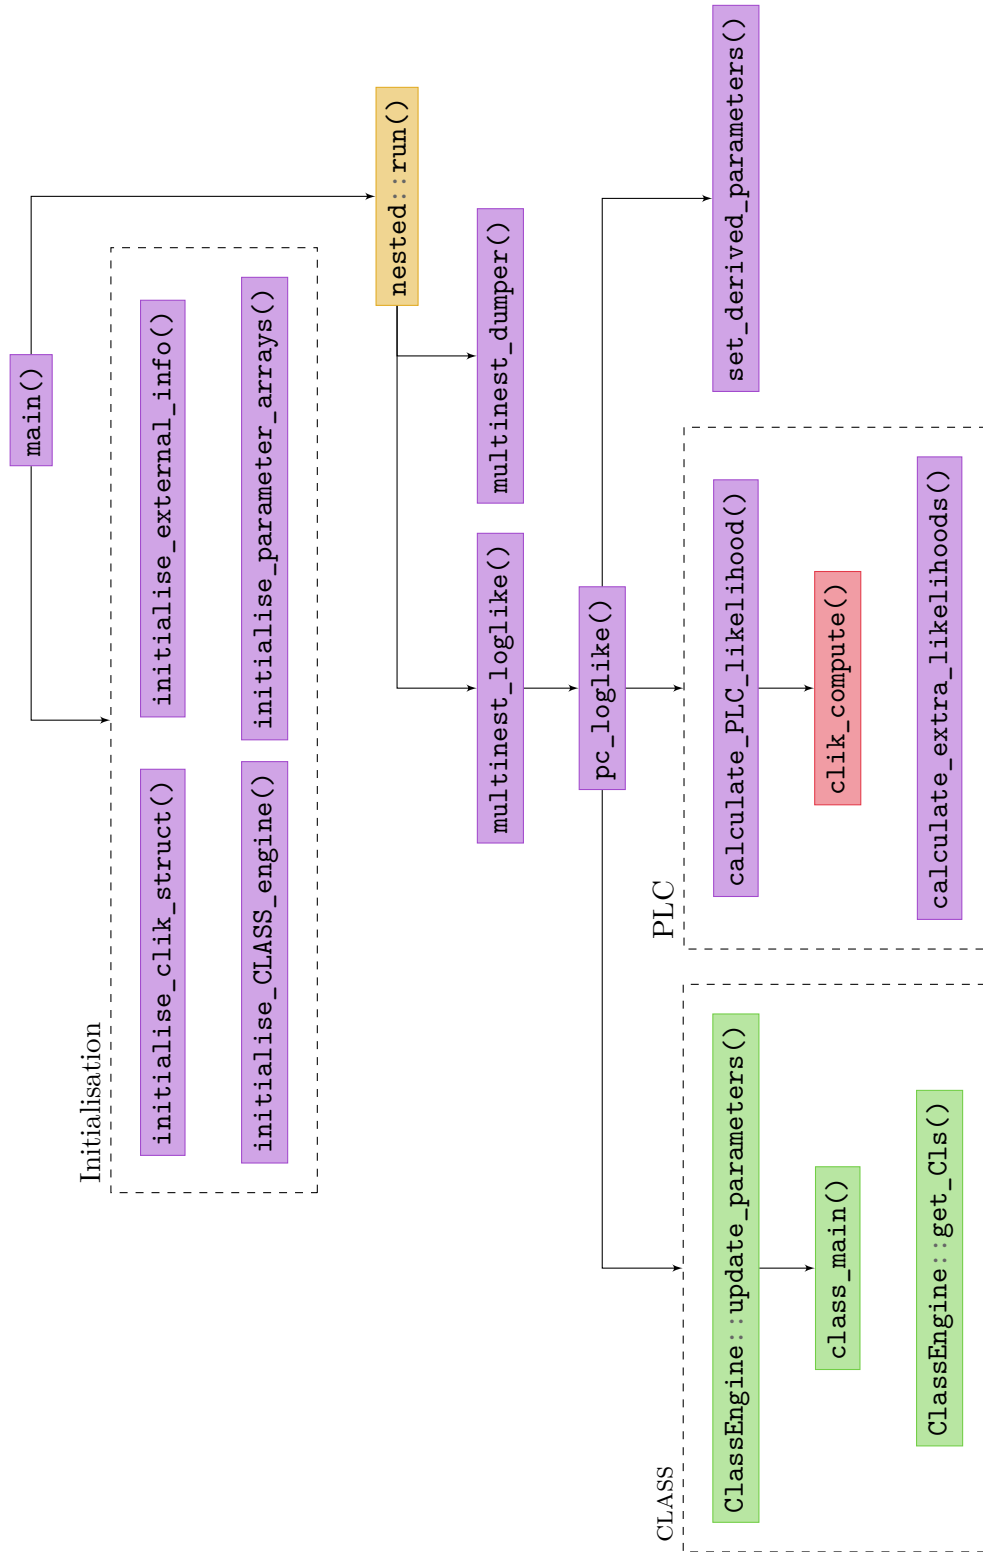
`pc_multinest` has been made available at [https://github.com/hdp1213/pc\\_multinest](https://github.com/hdp1213/pc_multinest) for public scrutiny.

#### 4.4.1 Parameter Handling

`pc_multinest` must handle a variety of different parameters during its execution. There are three types of parameters `pc_multinest` needs to consider: free, fixed and derived parameters. Free parameters are sampled from their priors by MULTINEST, and so need to have their priors stored somewhere. Fixed parameters are not used by MULTINEST at all, as they only have a single value given at compilation. Derived parameters are ones that can be seen as output from the model, such as parameters like  $H_0$  and  $t_0$ .

`pc_multinest` uses an enumeration to define each parameter used in the model, with a schematic representation of this given in Listing 4.1. Using enumerations also allows each parameter to be used as an index for arrays containing values associated with the parameter.

The `enum` environment is structured to keep track of which parameters are free, fixed or derived. The resulting structure enables the number of free, fixed and derived parameters to be



**Fig. 4.1** Calling hierarchy map for `pc.multinest`. Purple boxes correspond to functions defined in `pc.multinest`, the yellow box with `nested::run()` is defined in `MULTINEST`, the red box with `clik_compute()` is defined in `PLC` and green boxes are those defined in `CLASS`.

known at compile time, meaning the number does not need to be updated by the user when some parameters are changed.

### Value Arrays

Depending on the parameter's type, it can have associated values. These values are stored in global arrays. Although `pc_multinest` is a multi-threaded program, the choice of global arrays is well justified since the values contained within them are constants and are not rewritten or modified at any point during the program.

Free parameters each have a corresponding flat prior with bounds given by values in two arrays: `m_min` and `m_max`. As an example, if we consider a free parameter  $\theta_1$  with a flat prior of  $[a, b]$ , then this information is encoded in `pc_multinest` as

```
m_min[theta1] = a;   m_max[theta1] = b;
```

Similarly, fixed parameters only have a single value which is defined in a `m_value` array. Given a fixed parameter  $\theta_2 = c$ , this is encoded in `pc_multinest` as

```
m_value[theta2 - UP_TO_FREE_PARAMS] = c;
```

Note the inclusion of `UP_TO_FREE_PARAMS` which acts as an offset so that the first fixed parameter is indexed to the first element in the `m_value` array.

Derived parameters do not have an associated array, as their value is decided during the execution of `MULTINEST` and written to an internal array.

Some parameters in the *Planck* analysis also have extra Gaussian priors. To quantify which parameters actually have these Gaussian priors, the global array `m_has_gaussian_prior` contains a boolean flag. Similar to the flat prior case, Gaussian priors have their values set in the following arrays: `m_mean` and `m_stddev`.

Under some circumstances, a flat prior in linear space is not the most optimal prior. Instead, for parameters that act as scaling factors a flat prior in  $\log_{10}$  space is more optimal [87]. Other parameters may also benefit from the increased resolution at low values, such as the PBH fraction. For such parameters, the flat prior is specified in log space. `MULTINEST` then samples this parameter uniformly in log space before it is transformed back into linear space when calculating the likelihood.

There are two arrays that handle this transformation: `m_is_log10` and `m_transform`. The former contains flags for which parameters use  $\log_{10}$  space, and the latter contains the transformation functions. Before `MULTINEST` is run, the function array is initialised with the correct transformations. For regular parameters sampled in linear space, the transform is simply the identity.

## Parameter Files

In this work, a given scan uses parameters defined by a full set of enumerations with corresponding flat and Gaussian priors. As many different scans may need to be undertaken, this full set is placed into its own file. To change the parameters used in a scan, one could simply include a different parameter file during compilation.

The parameter files are split into a header file, which contains the enumeration definitions outlined earlier, and two source files which contain the uniform prior and Gaussian prior information respectively.

### 4.4.2 Initialisation

As mentioned before, `pc_multinest` uses the PLC and CLASS to evaluate the *Planck* likelihood. Before this is done, objects associated with PLC and CLASS need to be initialised. This allows the likelihood evaluation to only focus on the evaluation (rather than the initialisation) of these objects.

For PLC, the corresponding object in need of initialisation is a `clik_object` structure. Similarly, the CLASS code requires a `ClassEngine` object to be initialised. The results of each initialisation are added to a single object, `plc_bundle`, which is used in the likelihood evaluation.

#### `clik_object`

Initialisation of the `clik_object` requires the location of a likelihood file. For this work we make use of the high- $\ell$  TT,TE,EE `plik` likelihood with the Commander low- $\ell$  TEB likelihood (lowTEB), referred to as TT,TE,EE+lowTEB. The high- $\ell$  likelihood is evaluated using the TT, TE and EE cross spectra, and the low- $\ell$  uses the same with the addition of the BB spectrum. The corresponding files are `plik_lite_v18_TTEEE.clik` for the high- $\ell$  and `lowl_SMW_70_dx11d_2014_10_03_v5c_Ap.clik` for the low- $\ell$ .

After initialisation, each likelihood can be queried to access the number of nuisance parameters it needs, the maximum multipole  $\ell_{\max}$  it is computed to, and the types of spectra it contains. These values are read and stored alongside the `clik_object` so they can be used later on for sanity checks.

As each `clik_object` needs to know the values of their nuisance parameters to evaluate their likelihoods, each nuisance parameter enumeration is also stored alongside the `struct`. When storing these enumerations, care is taken in preserving the same order the `clik_object` expects them to appear in.

The resulting structure, `clik_struct`, contains the `clik_object`,  $\ell_{\max}$ , the size of the array used to evaluate the likelihood, the number of nuisance parameters and their enumerations, and a record of which spectra are contained in the `clik_object`.

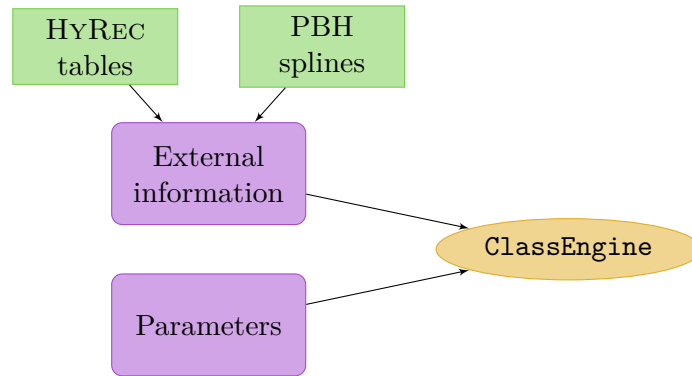


Fig. 4.2 Flowchart describing how `ClassEngine` is initialised.

### ClassEngine

`ClassEngine` is initialised in the way given by the flowchart shown in Figure 4.2, requiring both external information and CLASS parameters.

The external information reads in external files, such as HYREC effective rate tables and PBH spline coefficients. This is then passed on to the `ClassEngine` which stores the values for later use.

The `ClassEngine` also takes in a set of CLASS parameters. These parameters correspond to those given in an `.ini` file during a normal CLASS run. The following shows schematically how parameters are added to the `ClassEngine`:

```

default_params.add("model_class1", model1);
default_params.add("model_class2", model2);
...
default_params.add("internal_class1", internal1);
default_params.add("internal_class2", internal2);
...
  
```

Here, `default_params` is an array-like object used to initialise a `ClassEngine`. It adds parameters with their name and initial value. All `model_class*` parameters are physical parameters used in evaluating the given  $\Lambda$ CDM model. They are free parameters, and are sampled by MULTINEST later on. The `internal_class*` parameters are parameters used solely by CLASS, and do things like set the recombination model, output format, and the maximum multipole to compute spectra to. As such, they are not fixed parameters in the previous sense and hence do not have their own enumerations. Ordering all the free parameters first allows them to be updated during likelihood evaluation.

Note that each value given must be valid, as after this initialisation the `ClassEngine` does a validation run to confirm it can run with the given parameters.

### 4.4.3 Likelihood Evaluation

Likelihood evaluation begins with `MULTINEST`. `MULTINEST` takes as input the log likelihood function `multinest_loglike()`, `MULTINEST`-specific settings and an optional pointer containing extra data needed by the log likelihood function. In our work, this is the aforementioned `plc_bundle` structure which contains the initialised `PLC` and `CLASS` objects.

In the case of `clik_objects`, initialisation takes quite a long time, so having access to an already-initialised object in the likelihood function is essential to keeping evaluations quick. Similarly, updating an existing `ClassEngine` is much faster and more efficient than initialising one from scratch.

`MULTINEST` uses an array called `Cube` that stores the value of free and derived parameters of the current evaluation. At the start of the likelihood evaluation, the first `FREE_PARAM_AMT` entries of the array are initialised with  $u_i \sim U(0, 1)$ . The remaining `DERIVED_PARAM_AMT` entries are initialised to zero, and will eventually contain the value of any derived parameters. As each free parameter has a uniform prior, these initial  $u_i$  values are transformed to their physical values  $\theta_i$  using Eq. (4.2.2). The program also takes into account any parameter defined in  $\log_{10}$  space by transforming them back to linear space.

To get the likelihood, first the `ClassEngine` needs to be updated with these new parameters. This allows the spectra to be extracted, which is then used in the `clik_objects` to evaluate their likelihoods.

#### `ClassEngine`

The `ClassEngine` is updated with `CLASS` parameters in the same order they were declared when being initialised. Hence, for the vector of updated values `class_params`, we expect it takes the following form:

```
class_params.push_back(Cube[free_class1]);
class_params.push_back(Cube[free_class2]);
...
```

Here, `free_class*` are parameter enumerations for any free parameters needed by `CLASS`. The order in which these values are pushed into the vector is the same order that the `ClassEngine` was initialised with. We note that any fixed parameters do not need to have their value pushed into the vector, as these parameters can have their value fixed internally on initialisation without needing an enumeration to begin with.

After parameters have been updated, `ClassEngine` reruns `CLASS` with these new parameters. Then, the spectra are extracted from the engine and stored to be used by the `clik_objects`.

**clik\_object**

To evaluate the likelihood associated with a `clik_object`, a single array of values is needed, known as the `cl_and_pars` array. This array contains all of the  $C_\ell$ s for each spectrum included in the likelihood, as well as any nuisance parameters in the order the `clik_object` expects them.

The  $C_\ell$ s are filled from the `ClassEngine` output, whilst the nuisance parameter values are added using the `clik_object`'s nuisance parameter enumerations to access the corresponding values in the `Cube` array.

**Extra Likelihoods**

When all `clik_object` likelihoods are computed and summed up, any extra likelihoods are added onto the result. These include the Gaussian priors reformulated according to Eq. (4.2.3), as well as the flat [20, 100] prior placed on  $H_0$  and an extra Gaussian prior placed on a linear combination of two *Planck* nuisance parameters.

**4.4.4 Comparison to *Planck***

To check the validity of `pc_multinest`, we ran the baseline TT,TE,EE+lowTEB *Planck* analysis with a modified  $\Lambda$ CDM+PBH model implemented in `CLASS` and compared the results. The aims of this analysis was twofold: to show `pc_multinest` works as intended, and that the PBH modifications to `CLASS` behave like the  $\Lambda$ CDM model in the limit of no added PBHs. The analysis involved using the full high- $\ell$  likelihood with its 94 nuisance parameters. Using such a likelihood enabled more consistency checks to be performed due to the large number of parameters.

This baseline analysis required a scan over the standard six  $\Lambda$ CDM parameters and 27 nuisance parameters, with the remaining nuisance parameters set to values given in *Planck* Collaboration et al. [84]. The PBH modifications to `CLASS` will be discussed in the following chapter, but for now it is enough to know that the fraction of cold DM made of PBHs was set to  $10^{-90} \sim 0$  such that the modified model would behave like the vanilla  $\Lambda$ CDM model. The changes modify the behaviour of `RECFAST` and `HYREC`, both of which were used in separate runs as part of the analysis, with no difference between them seen in the final posterior distributions. This is as to be expected, as both models should agree when no extra energy is being injected during recombination. Here, we present the results computed with `HYREC`, which will be used later on in the full PBH physics computations.

Priors for the  $\Lambda$ CDM parameters were selected to fully encompass a region of  $8\sigma$  centred on the mean *Planck* value, with sigmas given by *Planck*'s 68% limits [88]. Similarly, the remaining 27 nuisance parameters use smaller priors compared to those used by the *Planck* analysis in order to speed up convergence, but care was chosen to not add unphysical bounds by truncating too much. The full set of priors used is given in Table 4.1. Some parameters also required an

extra Gaussian prior which is taken unchanged from Planck Collaboration et al. [84] and given in the form  $\mu \pm \sigma$  for  $\mu$  the mean and  $\sigma$  the standard deviation.

Alongside the priors cited in Table 4.1, two extra priors were included in the scan in alignment with the *Planck* analysis: one for the linear combination of the following two nuisance parameters

$$A^{kSZ} + 1.6A_{143}^{tSZ} = 9.5 \pm 3.0,$$

and a flat [20.0, 100.0] prior on  $H_0$  checked after it is computed by CLASS.

In order to best reproduce the posteriors of the *Planck* analysis, MULTINEST parameters were chosen to approximate the behaviour of a typical MCMC run which draws samples from a fixed prior volume. This was achieved by making the sampling efficiency small so that the ellipsoids that samples are drawn from encapsulate the full prior volume for a longer time. The values are given in Table 4.2. Due to their ability to reproduce the *Planck* analysis posteriors, these parameters were also used in the rest of this work.


Figure 4.3 shows the marginalised posteriors for the  $\Lambda$ CDM parameters plotted against the *Planck* results for the same analysis. Here, we note that all parameters agree well, with some small differences which will be discussed later on.

Similarly, Figure 4.4 shows the marginalised posteriors for the free nuisance parameters. The large discrepancy in  $A_{143 \times 217}^{\text{dust}TE}$  is because the *Planck* Collaboration released these chains when they were using a Gaussian prior of  $0.30 \pm 0.09$ . The final results updated this prior to  $0.60 \pm 0.18$ , which is the one used by `pc_multinest` in this run. The slight discrepancies in the remaining TE parameters are most likely due to this change in the Gaussian prior.

During the development of `pc_multinest`, the underlying CLASS source underwent version changes, bringing it from v2.5.0 to v2.6.3. As such updates brought overall improvements and optimisations to CLASS's operation, the decision was made to merge this update with the working branch of the PBH model extension. Due to some changes in how background quantities are calculated across versions v2.6.1 and v2.6.2 of CLASS, the values of some  $\Lambda$ CDM parameters shift slightly from the baseline *Planck* analysis. As these are only small deviations that arose due to changes in CLASS that are outside the scope of this work, we take this result as evidence that `pc_multinest` correctly calculates the *Planck* TT,TE,EE+lowTEB baseline result and as such works as intended.

Not only this, but this result also shows that in the limit of no added PBHs, we recover the physics of the  $\Lambda$ CDM, an outcome we would have expected but nonetheless serves as a sanity check on the modifications to CLASS.

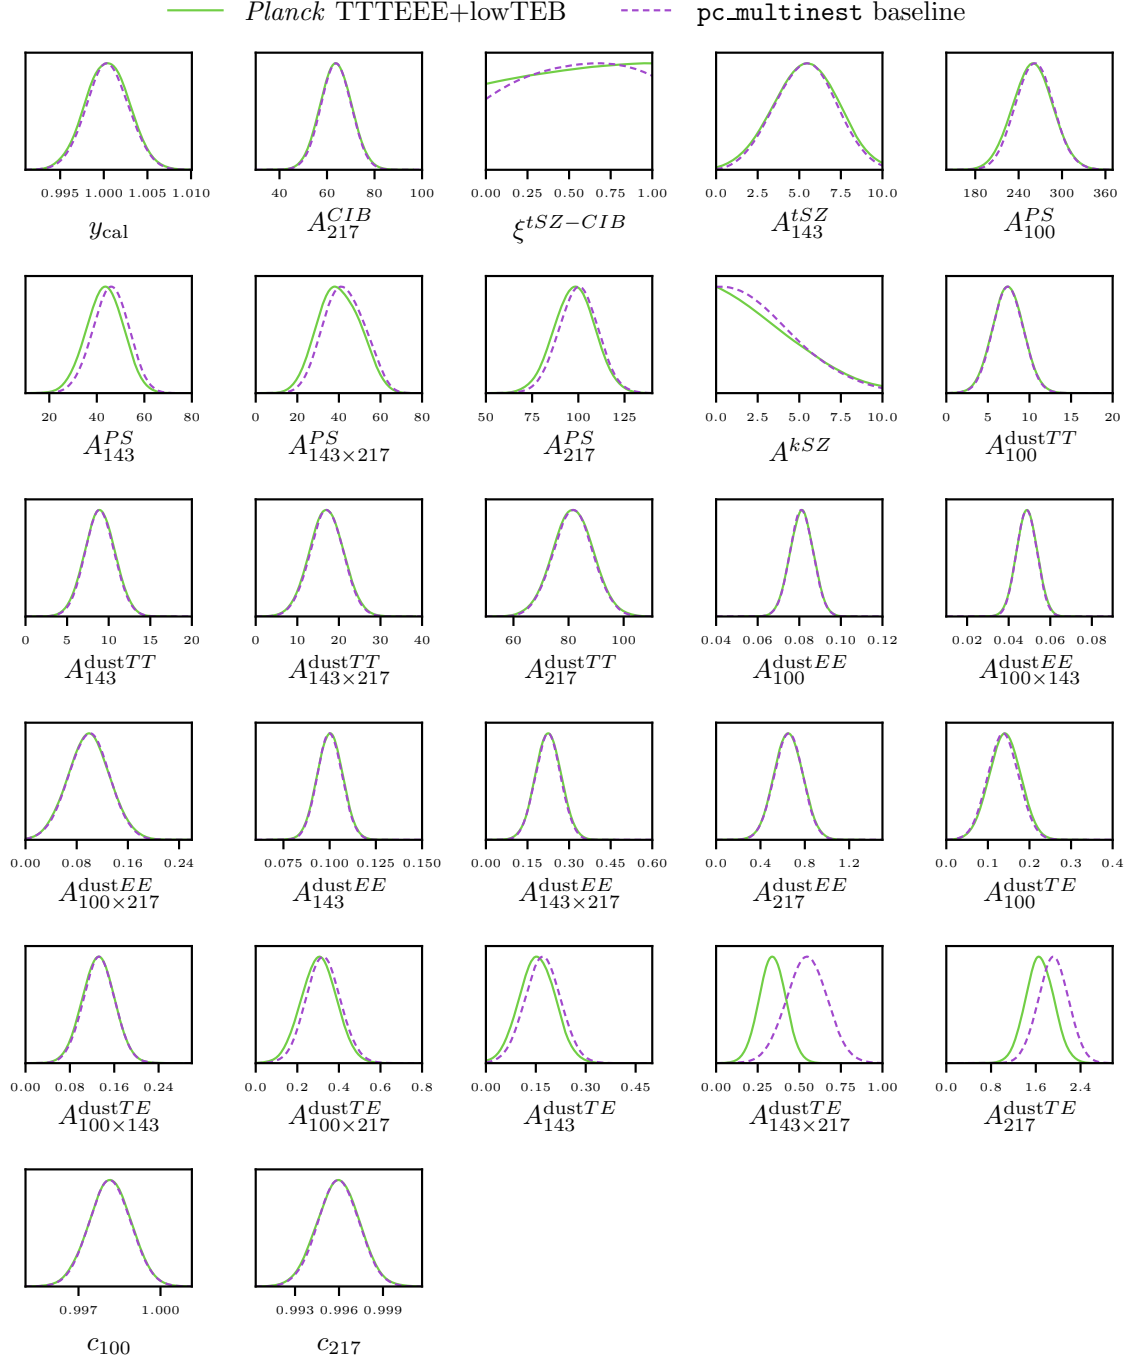
In order to mitigate any effect these slight discrepancies might have on the PBH limits calculated later on, we chose to extract a separate set of best-fit  $\Lambda$ CDM values, given in Table 4.3. The reasoning behind this is that if the *Planck* values were used when fixing  $\Lambda$ CDM parameters, they would not correspond exactly to the peak of the posterior as calculated by



**Fig. 4.3** Marginalised posteriors for  $\Lambda$ CDM parameters using the *Planck* TT,TE,EE+lowTEB and *Commander* likelihoods. The 2D contours show 68% (darker) and 95% (lighter) credible regions respectively.

*pc\_multinest*. Because of this, the corresponding results would favour regions of parameter space that would otherwise have been excluded if the  $\Lambda$ CDM parameters were more in line with their *Planck* best-fit values.

The values in Tab. 4.3 were computed using the ‘lite’ high- $\ell$  likelihood, which has all nuisance parameters marginalised out except for the total normalisation constant  $y_{\text{cal}}$ . From this point forth we choose to use this ‘lite’ likelihood over the full one due to the much smaller parameter space. The marginalisation of the nuisance parameters also places more focus on the  $\Lambda$ CDM parameters, which in contrast to the nuisance parameters are the model parameters of



**Fig. 4.4** Marginalised 1D posteriors for nuisance parameters using the *Planck* TT,TE,EE+lowTEB and Commander likelihoods.

most interest to us. We will see the effects that PBHs have on the full high- $\ell$  likelihood in the following chapter.

**Table 4.1** Uniform and Gaussian priors used for  $\Lambda$ CDM and nuisance parameters throughout this work. The top half of the table gives  $\Lambda$ CDM parameter priors, the bottom half gives the nuisance parameter priors. Note that nuisance parameter priors were only used in the *Planck* comparison whilst  $\Lambda$ CDM priors are used throughout the rest of this work when needed.

Parameter	Uniform prior	Gaussian prior
$\Omega_b h^2$	[0.016, 0.028]	—
$\Omega_{\text{cdm}} h^2$	[0.108, 0.130]	—
$100\theta_s$	[1.039, 1.043]	—
$\tau$	[0.01, 0.15]	—
$\ln(10^{10} A_s)$	[2.98, 3.20]	—
$n_s$	[0.92, 1.04]	—
$y_{\text{cal}}$	[0.9, 1.1]	$1.000 \pm 0.025$
$A_{217}^{CIB}$	[30.0, 100.0]	—
$\xi^{tSZ-CIB}$	[0.0, 1.0]	—
$A_{143}^{tSZ}$	[0.0, 10.0]	—
$A_{100}^{PS}$	[140.0, 370.0]	—
$A_{143}^{PS}$	[10.0, 80.0]	—
$A_{143 \times 217}^{PS}$	[0.0, 80.0]	—
$A_{217}^{PS}$	[50.0, 140.0]	—
$A^{kSZ}$	[0.0, 10.0]	—
$A_{100}^{\text{dust}TT}$	[0.0, 20.0]	$7.0 \pm 2.0$
$A_{143}^{\text{dust}TT}$	[0.0, 20.0]	$9.0 \pm 2.0$
$A_{143 \times 217}^{\text{dust}TT}$	[0.0, 40.0]	$21.0 \pm 8.5$
$A_{217}^{\text{dust}TT}$	[50.0, 110.0]	$80.0 \pm 20.0$
$A_{100}^{\text{dust}EE}$	[0.0, 2.0]	$0.060 \pm 0.012$
$A_{100 \times 143}^{\text{dust}EE}$	[0.0, 2.0]	$0.050 \pm 0.015$
$A_{100 \times 217}^{\text{dust}EE}$	[0.04, 0.12]	$0.110 \pm 0.033$
$A_{143}^{\text{dust}EE}$	[0.01, 0.09]	$0.10 \pm 0.02$
$A_{143 \times 217}^{\text{dust}EE}$	[0.00, 0.26]	$0.240 \pm 0.048$
$A_{217}^{\text{dust}EE}$	[0.060, 0.150]	$0.72 \pm 0.14$
$A_{100}^{\text{dust}TE}$	[0.0, 0.6]	$0.140 \pm 0.042$
$A_{100 \times 143}^{\text{dust}TE}$	[0.0, 1.5]	$0.120 \pm 0.036$
$A_{100 \times 217}^{\text{dust}TE}$	[0.0, 0.4]	$0.30 \pm 0.09$
$A_{143}^{\text{dust}TE}$	[0.0, 0.3]	$0.240 \pm 0.072$
$A_{143 \times 217}^{\text{dust}TE}$	[0.0, 0.8]	$0.60 \pm 0.18$
$A_{217}^{\text{dust}TE}$	[0.0, 0.5]	$1.80 \pm 0.54$
$c_{100}$	[0.0, 1.0]	$0.999\,000\,4 \pm 0.001\,000\,0$
$c_{217}$	[0.0, 4.5]	$0.995\,01 \pm 0.002\,00$

**Table 4.2** Values of MULTINEST parameters used throughout this work. Other parameters are excluded as they are either automatically set by the included parameter set at compile time or otherwise do not impact the resulting parameter scan.

MULTINEST parameter	Value	Description
IS	<code>false</code>	Switch for importance sampling
mmodal	<code>false</code>	Switch for multimodal search
ceff	<code>true</code>	Constant efficiency mode
nlive	1000	Number of live points
efr	0.05	Sampling efficiency $e$
tol	1E-1	Convergence tolerance $\Delta Z_i$

**Table 4.3** Best-fit values calculated from a lite high- $\ell$  likelihood run. These values are used later on when fixing  $\Lambda$ CDM and nuisance parameters to constant values. They are given to the same precision as those cited by Planck Collaboration [88].

$\Lambda$ CDM parameter	Best-fit value
$\Omega_b h^2$	0.022 220
$\Omega_{\text{cdm}} h^2$	0.119 83
$100\theta_s$	1.041 788
$\tau$	0.0803
$\ln(10^{10} A_s)$	3.0957
$n_s$	0.963 52
$y_{\text{cal}}$	1.000 47

# Chapter 5

## Constraining Primordial Black Holes

### 5.1 Introduction

Now that a theoretical model of primordial black holes injecting energy through evaporation has been laid out, and an adequate sampling regime has been chosen, the next step is to use both to place constraints on the model parameters. This chapter will outline the process that takes the PBH theory and implements it to constrain parameter space. This is done through the extension of the cosmological codes CLASS and HYREC.

Section 5.2 will cover the construction of the energy deposition efficiencies for primordial black holes, and Section 5.3 covers the extension of mass distributions to non-delta function distributions. The modifications made to the cosmological codes to account for injecting PBHs are outlined in Section 5.4. Finally, results and constraints are presented in Section 5.5.

### 5.2 PBH Deposition Efficiencies

As outlined in Section 3.4, primordial black holes inject energy into the primordial plasma non-trivially. These deposition efficiencies are dependent on both the deposition redshift  $z$  and PBH mass  $M_{\text{PBH}}$ , and are calculated by Eq. (3.4.6):

$$f_{\text{PBH},c}(M_{\text{PBH}}, z) = \frac{4f_{e^\pm}^{\text{em}} f_{\text{eff},c}^{e^\pm}(E_e, z) + 2f_\gamma^{\text{em}} f_{\text{eff},c}^\gamma(E_\gamma, z)}{4f_{e^\pm}^{\text{em}} + 2f_\gamma^{\text{em}}}. \quad (3.4.6)$$

The fractions also require effective energy depositions for electron-positron pairs (hereafter simply “electrons”) and photons. These are taken from Slatyer [9], who has provided tabulated forms for  $f_{\text{eff},c}^{e^\pm}(E_e, z)$  and  $f_{\text{eff},c}^\gamma(E_\gamma, z)$  hereafter referred to as the Slatyer tables.

This section details the steps taken in turning the effective electron and photon deposition efficiencies into usable PBH deposition efficiencies.

### 5.2.1 Effective Deposition Efficiencies

In Ref. [9], the mechanisms by which injected electrons and photons deposit their energy into the photon-baryon plasma are modelled explicitly. The results are tabulated based on the particle species considered and the resulting channel the energy is deposited into, given by

$$T_{c,ijk}^{\text{species}} \equiv T_c^{\text{species}}(z_{\text{dep}}^i, E^j, z_{\text{inj}}^k) d \ln(1 + z_{\text{dep}}^i).$$

This quantity is defined as the differential rate at which energy is absorbed into a channel  $c$  at redshift  $z_{\text{dep}}^i$  for a particle injected at redshift  $z_{\text{inj}}^k$  with energy  $E^j$ . Here, we note the inclusion of an injection redshift  $z_{\text{inj}}$ .

The rate  $T_c^{\text{species}}$  is used to compute the following particle deposition efficiency [7]:

$$f_{\text{eff},c}^{\text{species}}(z_{\text{dep}}) := \frac{\int E \int \frac{dN^{\text{species}}}{dE d \ln(1+z_{\text{inj}})} T_c^{\text{species}}(z_{\text{dep}}, E, z_{\text{inj}}) d \ln(1 + z_{\text{inj}}) dE}{\int E \frac{dN^{\text{species}}}{dE d \ln(1+z_{\text{dep}})} dE}. \quad (5.2.1)$$

This is the energy deposited into channel  $c$  at some redshift  $z$  by a particle species normalised by the total energy injected by the particle species at the same redshift  $z$ . Here,  $dN^{\text{species}}/dE d \ln(1+z)$  is the spectrum of injected particles as a function of energy and redshift, and can be species-dependent. All integrations over redshift are done with respect to  $d \ln(1+z)$ , as opposed to simply  $dz$ . This is because the tables give the redshift axes in log space, and so the integration must be performed over the same space when using the tabulated forms.

The above effective fraction is similar to the  $f_c(z)$  fractions defined in Section 3.3, but corresponds to the deposition efficiencies of the particles themselves. In order to simplify the expression, we consider the particle spectrum which may be rewritten as:

$$\begin{aligned} \frac{dN^{\text{species}}}{dE d \ln(1+z)} &= \left( \frac{dN^{\text{species}}}{dE dV dt} \right) dV \frac{dt}{d \ln(1+z)} \\ &\propto - \left( \frac{dN^{\text{species}}}{dE dV dt} \right) \frac{1}{H(z)(1+z)^3}, \end{aligned}$$

where we have used the following relations:

$$\frac{d \ln(1+z)}{dt} = -H(z), \quad dV(z) \propto \frac{1}{(1+z)^3}.$$

This allows us to write

$$f_{\text{eff},c}^{\text{species}}(z_{\text{dep}}) = \frac{H(z_{\text{dep}})(1+z_{\text{dep}})^3}{\int E \left( \frac{dN^{\text{species}}}{dE dV dt} \right) dE} \times \int E \int \left( \frac{dN^{\text{species}}}{dE dV dt} \right) \frac{T_c^{\text{species}}(z_{\text{dep}}, E, z_{\text{inj}})}{H(z_{\text{inj}})(1+z_{\text{inj}})^3} d \ln(1+z_{\text{inj}}) dE. \quad (5.2.2)$$

PBHs with mass  $> 10^{15}$  g have an ejection spectrum predominantly governed by electrons, photons and neutrinos. In a manner of speaking, these holes behave similarly to decaying WIMP particles. We will use this similarity to model the volume dependence of the resulting Hawking radiation. As particle decays are a function of the particle density (rather than the density squared as with annihilations), the associated energy spectrum is proportional to the expanding volume, or  $(1+z)^3$ . Using this, we may write

$$\left( \frac{dN^{\text{species}}}{dE dV dt} \right) = (1+z)^3 \frac{d\bar{N}^{\text{species}}}{dE},$$

where  $d\bar{N}^{\text{species}}/dE$  is the redshift-independent emission spectrum for the particles under consideration. In the case of PBHs, this is given by Eq. (3.4.3):

$$\frac{d\bar{N}^{\text{species}}}{dE} = \frac{\Gamma_s}{2\pi\hbar} \frac{1}{\exp\left(\frac{E}{k_B T_{\text{BH}}}\right) - (-1)^{2s}},$$

where in the above we have dropped the  $dt$  despite it still being a rate of injection in order to remain consistent with the notation of Liu, Slatyer, and Zavala [89].

Instead of using the computationally expensive emission spectra given above, we follow Clark et al. [1] and simplify matters by assuming that all particles of a given species ejected from the PBH have the same energy  $\bar{E}_{\text{species}}$  given by Eq. (3.4.4). Hence, we effectively set

$$\frac{d\bar{N}^{\text{species}}}{dE} = \delta(E - \bar{E}_{\text{species}}),$$

which reduces Eq. (5.2.2) to the following:

$$f_{\text{eff},c}^{\text{species}}(z_{\text{dep}}, \bar{E}_{\text{species}}) = H(z_{\text{dep}}) \int \frac{T_c^{\text{species}}(z_{\text{dep}}, \bar{E}_{\text{species}}, z_{\text{inj}})}{H(z_{\text{inj}})} d \ln(1+z_{\text{inj}}).$$

Following Clark et al. [1], we also assume that the ratio  $H(z_{\text{dep}})/H(z_{\text{inj}}) \approx 1$  over the range of injection redshifts considered. This allows us to write

$$f_{\text{eff},c}^{\text{species}}(z_{\text{dep}}, \bar{E}_{\text{species}}) \approx \int T_c^{\text{species}}(z_{\text{dep}}, \bar{E}_{\text{species}}, z_{\text{inj}}) d \ln(1+z_{\text{inj}}).$$

Using this approximation makes the deposition efficiencies independent of the cosmology considered, and allows these quantities to be computed beforehand and read in later on.

The deposition efficiency curves for  $f_{\text{eff},c}^{\text{species}}(z_{\text{dep}}, \bar{E}_{\text{species}})$  are given in Figure 5.1, and agree well with those presented in Figure 3 of Slatyer [9].

### 5.2.2 Interpolation

In order to convert the particle energies  $E_e$  and  $E_\gamma$  into a PBH mass, two steps need to be taken: first a conversion from energy to PBH temperature, and then from temperature to PBH mass.

To make the first conversion, we use the relations given in Eq. (3.4.4):

$$\bar{E}_{e\pm} = 4.18k_{\text{B}}T_{\text{BH}}, \quad \bar{E}_\gamma = 5.71k_{\text{B}}T_{\text{BH}}, \quad \bar{E}_\nu = 4.22k_{\text{B}}T_{\text{BH}}. \quad (3.4.4)$$

Here, we implicitly assume that electrons and photons ejected from PBHs have a fixed energy given by the above relations. Since the above relations are given in terms of the *average* particle energies ejected from a black hole, this is a reasonable enough assumption to make.


As the Slatyer tables are evaluated at fixed energies, and the corresponding PBH temperatures are different for the different species, the tables for each species must first be interpolated before being combined at the specific temperature value of interest. To do this, a bicubic b-spline was selected. B-splines are basis functions for more complex splines, and as such they describe curves in the least amount of parameters.

A bicubic b-spline is one that fits a cubic function between spline points (or *knots*) across two dimensions, ensuring the interface between two splines is continuous in its first and second derivatives. This selection was deemed appropriate, as it gives a better fit than a bilinear spline without introducing too many degrees of freedom and with it the risk of overfitting.

Whilst fitting the Slatyer tables, care was taken to ensure no edge effects crept in and that the interpolated values returned were actually positive. The region of energy space covered by the Slatyer tables is much larger than the range of black hole energy ejections. As a consequence, the spline only needed to be fit across a smaller region of energy space. In practice, the spline was fit across an energy region slightly larger than the one under consideration to avoid any edge effects from the spline having to extrapolate, rather than interpolate values lying close to the edge.

The deposition efficiencies being interpolated must be positive, as a negative fraction implies that the injected energy is deposited as a deficit, and that the surrounding plasma loses energy due to the injection rather than gains from it. However, during testing of the interpolation, some interpolated values returned were negative. This was because sharp drops in some regions of deposition efficiency space meant the interpolation occasionally dipped below zero.

In order to force positivity of the resulting spline, the spline was fit on the square root of the deposition efficiencies. The resulting interpolated value was then squared to return the



**Fig. 5.1** Energy deposition efficiencies for electrons and photons across different channels. The left column is electron depositions, the right column is photon depositions. From top to bottom, the channels are hydrogen ionisation, Lyman- $\alpha$  excitation and heating.

correct fraction. Fitting the spline on the square root of the data also improved the overall fit of the data, because the square root shifts all points in the range  $[0, 1]$  closer to one and away from zero. This effectively smoothed out “problem areas” the spline found hard to fit to.

### 5.2.3 Cross-Validation

In order to confirm the accuracy of using a bicubic b-spline, a cross-validation method was used. Because there was only one dataset per particle species per deposition channel, the data were resampled to create an ensemble of data, from which the goodness-of-fit for interpolated values was evaluated. This was done by resampling the data every  $n$ th point in both dimensions. Because the resulting subset is more coarse than the original dataset, an ensemble of independent subsets can be used to calculate an average with an associated error. Given a resampling using every  $n$ th point, this resulted in  $n^2$  subsets. The bicubic spline was then fit on each subset. In order to compare the two fits, the “fractional difference” was chosen:

$$\frac{|f_{\text{raw},c}(E, z) - f_{\text{eval},c}(E, z)|}{|f_{\text{raw},c}(E, z)| + |f_{\text{eval},c}(E, z)|}.$$


This quantity is zero when the spline-evaluated data matches the raw data, and is positive otherwise and gives the effective percentage difference between the raw data and the fitted. The reason for the sum of the two terms on the denominator is to prevent the metric blowing up at very small values of deposition efficiency. For example, if we consider a raw fraction value of  $10^{-20}$  and compare it with an interpolated value of  $10^{-19}$ , relatively speaking the difference is quite large. In practice, when considering the actual energy injections, the difference between the two is negligible.

When fitting a spline, we do not care about how well it fits the data it was fitted to. Instead, we care about how accurately it interpolates the data. With this in mind, the resulting fractional difference is set to zero at the points used in the fit. The total is then summed and divided by  $n^2 - 1$ , as we have effectively lost one point from each subset due to the fit. This creates an average best-fit with an estimate on the uncertainty given by  $\sigma/\sqrt{n^2 - 1}$ , where  $\sigma$  is the error in the ensemble.

The results of this are given in Table 5.1. Here we see that the average discrepancy across the ensemble is quite low. In the  $n = 2$  case, we note that two of the three subsets will interpolate a given point better than the third. Looking at Figure 5.2, the black points are those the spline is fit to. The upper-right diagonally striped green points are interpolated better than the upper-left diagonally striped purple point because the green points lie between points used in the fit. The purple point almost has a compounded uncertainty associated with it as

**Table 5.1** Average discrepancies for bicubic b-spline fits of the Slatyer tables.

Species	H ionisation	Ly- $\alpha$ excitation	Heating
Electron	$(4.45 \pm 3.86) \%$	$(3.35 \pm 2.74) \%$	$(4.21 \pm 3.87) \%$
Photon	$(2.95 \pm 3.27) \%$	$(2.88 \pm 3.22) \%$	$(2.98 \pm 3.35) \%$



**Fig. 5.2** Schematic representation of how coarsening the grid on which the spline is fit results in different accuracy interpolations. See text for interpretation.

the closest fitted points to it are diagonally adjacent. However, it receives information from four points rather than just two.

### 5.2.4 Results

With confirmation that spline interpolation will give accurate results, the actual construction of the deposition efficiencies was done. Figure 5.3 contains plots of the final interpolated PBH deposition efficiencies in PBH mass space. These resulting splines span 49 log-spaced bins in PBH mass space and 1989 log-spaced bins in redshift space. The mass range is from  $10^{14.8}$  g to  $10^{17.2}$  g, as the range needed to be extended to remove any potential edge effects. The redshift range runs from  $1 + z = 12$  to 2000 inclusive.

## 5.3 PBH Mass Distributions

Up until recently, all cosmological analyses of PBHs have assumed that PBHs share the same mass. As mentioned in Section 3.4, some models of inflation can result in non-delta mass distributions for PBHs. Hence, it makes sense to consider the effects a finite mass distribution has on constraints derived from measurements of the CMB. In this work, the previous analysis is extended to include a log-normal mass distribution.

This section will cover the theory and implementation of a log-normal mass distribution.


### 5.3.1 Non-Monochromatic Masses

Taking the non-standard terms defined in Eqs. (3.3.1)–(3.3.3) and replacing the injection fractions with the PBH fractions defined in the previous section, we get the following:

$$I_{X_i}(x_e, T_m; z) = f_{\text{PBH},i}(M_{\text{PBH}}, z) \frac{(dE/dV dt)_{\text{inj}}}{n_{\text{H}}(z)E_i}, \quad (5.3.1)$$

$$I_{X_\alpha}(x_e, T_m; z) = [1 - C(x_e, T_m; z)] f_{\text{PBH},\alpha}(M_{\text{PBH}}, z) \frac{(dE/dV dt)_{\text{inj}}}{n_{\text{H}}(z)E_\alpha}, \quad (5.3.2)$$

$$K_{\text{h}}(x_e, T_m; z) = f_{\text{PBH},\text{h}}(M_{\text{PBH}}, z) \frac{(dE/dV dt)_{\text{inj}}}{n_{\text{H}}(z)}, \quad (5.3.3)$$



**Fig. 5.3** Energy deposition efficiencies  $f_{\text{PBH},c}(M_{\text{PBH}}, z)$ . Clockwise, the channels are hydrogen ionisation, Lyman- $\alpha$  excitation and heating. The temperature axis in Figure 1 of Clark et al. [1] has been transformed into a mass axis in this work.

These terms form the basis for the typical PBH analysis with monochromatic masses.

What we are interested in now is incorporating some mass distribution into these terms. To do so, we define the following quantity:

$$\begin{aligned}
 F_{\text{PBH},c}(z; \boldsymbol{\theta}) &:= \int p(M'_{\text{PBH}}|\boldsymbol{\theta}) \left( \frac{dE}{dV dt} \right)_{c,\text{dep}} dM'_{\text{PBH}} \\
 &= \int p(M'_{\text{PBH}}|\boldsymbol{\theta}) f_{\text{PBH},c}(M'_{\text{PBH}}, z) \left( \frac{dE}{dV dt} \right)_{\text{inj}} dM'_{\text{PBH}}, \quad (5.3.4)
 \end{aligned}$$

where  $p(M'_{\text{PBH}}|\boldsymbol{\theta})$  is a PBH mass distribution dependent on some shape parameters  $\boldsymbol{\theta}$ .

This term is a weighted mean of the deposited energy across the PBH mass range considered. In the case of a delta function mass distribution where

$$p(M'_{\text{PBH}}|M_{\text{PBH}}) = \delta(M'_{\text{PBH}} - M_{\text{PBH}}),$$

we find

$$F_{\text{PBH},c}(z; M_{\text{PBH}}) = f_{\text{PBH},c}(M_{\text{PBH}}, z) \left( \frac{dE}{dV dt} \right)_{\text{inj}},$$

which are the energy deposition terms appearing in the non-standard terms given in Eqs. (5.3.1)–(5.3.3).

Hence, these equations can be rewritten as

$$\begin{aligned} I_{X_i}(x_e, T_m; z) &= \frac{F_{\text{PBH},i}(z; \boldsymbol{\theta})}{n_{\text{H}}(z) E_i}, \\ I_{X_\alpha}(x_e, T_m; z) &= [1 - C(x_e, T_m; z)] \frac{F_{\text{PBH},\alpha}(z; \boldsymbol{\theta})}{n_{\text{H}}(z) E_\alpha}, \\ K_{\text{h}}(x_e, T_m; z) &= \frac{F_{\text{PBH},\text{h}}(z; \boldsymbol{\theta})}{n_{\text{H}}(z)}. \end{aligned}$$

### Log-Normal Mass Distribution

As the PBH mass range under consideration spans two orders of magnitude, it makes more sense to consider  $\log_{10}$  space, where a range of  $10^{15}$  to  $10^{17}$  becomes a range of 15 to 17. We can formulate various mass distributions in such a space, but for this work we will be considering a normal distribution. This is in agreement with recent work by Garca-Bellido and Ruiz Morales [66], who provide a model of inflation leading to PBHs with normally-distributed masses. As this will be implemented in  $\log$  space, in linear space it becomes a log-normal distribution.

Defining the following dimensionless “reduced” mass term:

$$M'_{\text{red.}} := \log_{10} \left( \frac{M'_{\text{PBH}}}{10^{10} \text{ g}} \right),$$

in  $\log_{10}$  space the mass distribution takes the form

$$p(M'_{\text{red.}} | \mu_{10}, \sigma_{10}) := \frac{1}{\sigma_{10} \sqrt{2\pi}} \exp \left[ -\frac{(M'_{\text{red.}} - \mu_{10})^2}{2\sigma_{10}^2} \right],$$

where we have defined  $\mu_{10}$  to be the mean PBH mass, and  $\sigma_{10}$  to be the standard deviation (or effective width) of the distribution. Note that in such a “reduced” space, the original mass range becomes 5 to 7.

### 5.3.2 Implementation

In practice, as the deposition efficiencies  $f_{\text{PBH},c}(M_{\text{PBH}}, z)$  are non-analytic, the integral in Eq. (5.3.4) is done numerically. As it is infeasible to perform this integral from  $-\infty$  to  $\infty$ , the integral (and all subsequent ones involving the mass distribution) is performed numerically over the range  $[\mu_{10} - 4\sigma_{10}, \mu_{10} + 4\sigma_{10}]$ .

Considering this calculation must be called for each redshift step used to calculate the recombination history, where each calculated recombination history corresponds to one likelihood evaluation, the speed of evaluation was considered a priority. As such, a quick integration routine was needed for the feasibility of the calculation. Due to the simplicity of the normal distribution, a simple trapezoidal integration routine with a fixed number of points was used. Having a fixed number of points meant that the mesh the distribution was sampled over expanded and contracted with different distribution widths. Having this feature enabled a distribution of arbitrary thickness to be correctly sampled over.

The routine was evaluated against a more accurate quadrature algorithm to ensure the number of points the function is integrated over gives an accurate enough result. It was found that 41 points were sufficient to provide the desired accuracy.

Due to the finite PBH mass range under consideration, care had to be taken when implementing the mass distribution so that in the case of truncation, the resulting distribution compensates for it. Firstly, we note that the deposition efficiencies for PBHs are asymmetrical. Figure 5.3 shows how lighter PBHs deposit more of their energy into each channel compared to heavier PBHs. Similarly, the total injected energy is inversely proportional to the cube of the PBH mass.

Overall, we expect that Hawking radiation from lower mass PBHs are more important than those from higher mass PBHs. As a consequence, if truncation of the mass distribution occurs on the low mass limit of  $10^{15}$  g, the mass distribution should be renormalised to unity. For the distribution  $p(M'_{\text{red.}}|\mu_{10}, \sigma_{10})$  truncated at this level, the normalisation factor  $\lambda$  is given by

$$\lambda(\mu_{10}, \sigma_{10}) := \int_5^7 p(M'_{\text{red.}}|\mu_{10}, \sigma_{10}) dM'_{\text{red.}} \neq 1.$$

The final quantities  $F_{\text{PBH},c}(z; \mu_{10}, \sigma_{10})$  are then re-weighted appropriately:

$$F_{\text{PBH},c}(z; \mu_{10}, \sigma_{10}) \rightarrow F'_{\text{PBH},c}(z; \mu_{10}, \sigma_{10}) = \frac{F_{\text{PBH},c}(z; \mu_{10}, \sigma_{10})}{\lambda(\mu_{10}, \sigma_{10})}.$$

This correction is only evaluated if the following condition holds:

$$\mu_{10} - 4\sigma_{10} < 5,$$

remembering 5 is the lower mass bound in reduced space. The value of  $4\sigma_{10}$  was deemed suitable, as for a normal distribution this value encapsulates 99.99% of the total area of the distribution.

This truncation compensation is not done for the higher mass PBHs, as their Hawking radiation is comparatively negligible compared to the lower mass PBHs.

## 5.4 Code Extensions

The Boltzmann code CLASS creates CMB power spectra that can be compared to the *Planck* data to obtain a likelihood for a point in PBH parameter space. As such, it must be modified alongside the recombination codes it calls so it can incorporate the effects that evaporating primordial black holes would have had on the early universe. This section will detail the steps taken in extending these codes and verifying their accuracy. All extensions are freely available for public scrutiny at the following GitHub repository branch: [https://github.com/hdp1213/class\\_public/tree/phoenix](https://github.com/hdp1213/class_public/tree/phoenix).

### 5.4.1 CLASS

CLASS effectively becomes a wrapper for the modifications to RECFAST and HYREC. As a program, it must correctly handle the inputting of new PBH parameters, as well as read in the precomputed deposition efficiency tables correctly. As progress was made on CLASS, other problems and issues came up that were also addressed.

#### Input Parameters

In CLASS, all parameter inputs pass through the `input.c` module. For CLASS to work with a modified RECFAST and HYREC, it needed extra PBH parameters to be defined, which are given in Table 5.2.

#### B-Spline Interpolation

Since the PBH deposition efficiencies are precomputed tables, they can be read in externally and interpolated on the run. The choice was made to store the table as spline coefficients as opposed to the raw values. This enabled faster interpolation, as the coefficients could be read in instead of having to be generated each time a value was needed. As a consequence, CLASS was extended to read in and interpret these spline coefficients.

**Table 5.2** Parameters added to CLASS to extend the vanilla  $\Lambda$ CDM to include the effects of PBHs.

Variable	Symbol	Description
<code>Omega_pbh_ratio</code>	$f_{\text{PBH}}$	Fraction of dark matter made of evaporating PBHs
<code>pbh_mass_mean</code>	$M_{\text{PBH}}$	Mean mass of PBH mass distribution
<code>pbh_mass_width</code>	$\sigma_{10}$	Width/standard deviation of PBH mass distribution
<code>pbh_mass_dist</code>	–	PBH mass distribution. Either none, delta or log normal
<code>read_external_files</code>	–	Flag to prevent external files being read by CLASS during each evaluation when run in batch

The table was stored as coefficients for a bicubic b-spline across redshift and PBH mass space. Instead of rewriting a b-spline interpolator from the ground-up, an existing one was used. The chosen interpolator was an algorithm written in Fortran 77 taken from the `fitpack` library. It is currently used in the Python package `scipy` as part of its interpolation suite<sup>1</sup>. An interfacing function was written in C to call the Fortran routine, and the compiled Fortran code was included in the final executable.

### Miscellaneous Issues

Alongside the main work done on `CLASS` to prepare it for PBH physics, other peripheral tasks were completed to address various issues that arose. These included updating the C++ wrappers for `CLASS`, and fixing a substantial memory leak.

**C++ Wrappers** Noting that `pc_multinest` is written in C++, the associated C++ wrappers for `CLASS` were also updated to allow PBH information to be passed through. This primarily consisted of adding extra input parameters to functions to allow the passage of pointers containing information needed by `CLASS` to evaluate PBH physics.

**Memory Leaks** In its current state, `CLASS` is intended to be run as a standalone program. Whenever an error occurs in a function, `CLASS` stops itself from running any further by returning from each successive caller function until it returns an exit code from the top of the program. As such, it does not free any allocated memory. This is not an issue when `CLASS` is called as a standalone program, as all the memory allocated during execution of the program is then freed by the operating system when the `CLASS` process finishes.

However, when calling `CLASS` as part of a larger sampling code the process never finishes on a `CLASS` error. Instead, the process continues with a different choice of parameters. This implicit memory leak increases the memory footprint of the sampling program until it either finishes sampling or hits the memory limit for the computer. In order to fix this leak, any allocated memory must be explicitly deallocated in the event of an error.

In the interests of time, only leaks that arose from errors that were known to occur were patched. A more robust solution would be to keep track of all currently allocated memory such that on an error the current heap could be cleared, and no memory would leak as a result.

### 5.4.2 RECFAST

In `CLASS`, `RECFAST` is already extended to incorporate WIMP dark matter so its extension to include PBHs closely followed this. The explicit implementation of PBHs will not be covered for `RECFAST`, as this was not the final recombination code used in the results for the reasons outlined below.

<sup>1</sup><https://github.com/scipy/scipy/tree/master/scipy/interpolate/fitpack>

### Usage Issues

At the beginning of this work, RECFAST was used over HYREC due to its simplicity and apparent ability to incorporate non-standard energy injections as evidenced by its extension to include DM annihilation and decay. During some early runs, however, it was noted that RECFAST would periodically crash. The problem was eventually found to be due to the approximation schemes RECFAST used to increase its speed, namely the extension by Scott and Moss [56]. The approximation defined an effective measure of the duration of Compton heating: the Compton timescale  $t_C$  given by

$$t_C(T_\gamma, x_e; z) = \frac{3m_e c}{8\sigma_T a_T T_\gamma^4} \frac{1 + f_{\text{He}} + x_e}{x_e}.$$

In the absence of other heating, comparing this timescale to the Hubble time  $t_H(z)$  defined a part of recombination where heat exchange is achieved primarily through Compton processes and so the matter and radiation temperature were almost identical up to some difference  $\epsilon(z)$ . However, this assumption breaks down when non-standard heating is considered. An attempt to re-derive the limit on  $t_H/t_C$  was made incorporating the non-standard effects, however this did not work when implemented in RECFAST, suggesting that the code was simply not designed for non-standard physics. There was later confirmation by the author Douglas Scott that RECFAST was never intended to support cosmologies far from a vanilla  $\Lambda$ CDM model [90].

In the end, RECFAST was deemed inappropriate for the given task and so HYREC was chosen as its successor. Not all work on RECFAST was wasted, however. By extending RECFAST, we obtained a way of verifying that the subsequent HYREC extension was performing as expected for small energy injections.

#### 5.4.3 HyRec

To obtain our final results, we used a beta version of HYREC generously provided by Yacine Ali-Haïmoud which incorporates WIMP dark matter energy injections. This was subsequently extended to incorporate PBH energy injections, in a manner not dissimilar to RECFAST.

#### Implementation of PBHs

The following methods contain explicit recombination calculations and were flagged to be updated for PBHs: `rec_TLA_dxHIIIdlna`, `rec_HMLA_dxHIIIdlna`, `populateTS_2photon`, `rec_Tmss`, and `rec_dTmdlna`. Other methods were also updated, but only in their calling so they could pass the needed parameters through to the above methods. Due to the similarity between annihilating WIMPs and evaporating PBHs, all PBH terms added to HYREC followed the WIMP dark matter additions closely in their form.

**`rec_TLA_dxHIIIdlna`** This method implements the TLA approximation that HYREC can run as one of its models. As such, the extension followed the same logic covered in Section 3.3. The

form of the equations differs slightly compared to the theory as in practice, these equations have been reformulated in terms of differences from equilibrium in order to account for almost-cancelling terms at high- $z$ .

**rec\_HMLA\_dxHIIIdlna** This method implements the EMLA for hydrogen evolution. It has a similar form to **rec\_TLA\_dxHIIIdlna**, but interpolates the effective rates and uses a linear combination of the generalised 2s and 2p versions of Peebles' coefficient. The non-standard PBH terms are simply added to the result in the same fashion as for the TLA case, making note of the Peebles coefficient substitution:

$$C(x_e, T_m; z) \rightarrow 0.25C_{2s}(x_e, T_m; z) + 0.75C_{2p}(x_e, T_m; z).$$

The choice of these coefficients is made by considering that excitations predominantly come from Ly- $\alpha$  photons (exciting to the 2p channel), with some support from collisional excitation from fast electrons (exciting to the 2s channel).

**populateTS\_2photon** This method solves the radiative transfer equation by computing transition rates for two-photon processes. Here, the Hawking radiation must be added to the source vectors  $S_i$  for  $i = 2s, 2p$ . The contributions are weighted in the same 1:3 way as above:

$$S_{2s} = (S_{2s})_{\text{std}} + 0.25 \frac{F_{\text{PBH},\alpha}(z; \boldsymbol{\theta})}{n_{\text{H}}(z)E_{\alpha}}, \quad S_{2p} = (S_{2p})_{\text{std}} + 0.75 \frac{F_{\text{PBH},\alpha}(z; \boldsymbol{\theta})}{n_{\text{H}}(z)E_{\alpha}},$$

where the  $(S_i)_{\text{std}}$  are the standard source vectors given in Eq. (92) of Ali-Haïmoud and Hirata [57] and  $F_{\text{PBH},\alpha}(z; \boldsymbol{\theta})$  is the energy deposited into the excitation channel, from Eq (5.3.4).

**rec\_Tmss** This method uses a quasi-steady state approximation, where the matter temperature is approximately equal to the radiation temperature. In order to extend this, the technique in Appendix A of Giesen et al. [10] was used.

**rec\_dTmdlma** This method implements the full matter temperature evolution, whose extension has already been outlined in Section 3.3.

## Verification

In order to check the implementation of PBHs in HYREC, the resulting ionisation history  $x_e$  and matter temperature  $T_m$  were compared against the RECFast implementation. We aim to compare the behaviour of the two codes on the basis that for small amounts of Hawking radiation, the RECFast implementation is correct. We are allowed to assume this as its formulation closely followed the TLA approach, meaning sections of the code that needed to be changed had a form identical to the recombination equations in Section 3.3.2. Nonetheless, we

expect some discrepancies even without adding Hawking radiation due to the different ways both codes approach the recombination problem.

Figure 5.4 compares  $x_e(z)$  made using both RECFAST and HYREC, and similarly Figure 5.5 compares the matter temperature  $T_m(z)$ . The histories in both cases are calculated without any reionisation effects. A value of  $f_{\text{PBH}} = 10^{-7}$  was chosen for all PBH masses, which allows us to examine how the behaviour for both RECFAST and HYREC diverge as the injected energy increases with decreasing PBH mass. Both plots also show the discrepancy between RECFAST and HYREC in the case of no added PBHs (given by the  $M_{\text{PBH}} = 0$  g line) which serves to isolate the discrepancies purely due to PBH modelling that become visible at larger amounts of Hawking radiation.

Fig. 5.4 shows that for less Hawking radiation, HYREC agrees closely with RECFAST, with discrepancies at the sub-percent level. There is a small amount of disagreement between the no PBH case and the  $M_{\text{PBH}} = 10^{16}$  g case at low redshift, but otherwise this shows that in the small injected energy regime, both codes agree to well within their differences. For  $M_{\text{PBH}} = 10^{15}$  g we get a larger discrepancy of just over 11 % during recombination, which tells us that RECFAST seems to underestimate the amount of PBH ionisation at later times for high Hawking radiation.

Similarly, Fig. 5.5 shows good agreement in the low Hawking radiation limit. We see a more pronounced underestimation for  $M_{\text{PBH}} = 10^{16}$  g, which would be due to HYREC's improved modelling at later times. Just before these two lines diverge, the deviation of  $T_m(z)$  from the no PBH case is at 10 %, which suggests the prior agreement is not a symptom of the low Hawking radiation having little effect on the matter temperature, but is evidence for agreement across both HYREC and RECFAST.


Overall, both plots show the need for HYREC's improved modelling for high Hawking radiation, especially considering RECFAST cannot handle PBH fractions above  $10^{-7}$  for all masses in  $10^{15}$  g to  $10^{17}$  g.

### Effective Table Extension

In its current form, HYREC draws from a two-dimensional table of precomputed effective transition rates evaluated at differing matter and radiation temperatures. The axes the table is evaluated over are radiation temperature  $T_\gamma$ , and the ratio of matter to radiation temperature  $T_m/T_\gamma$ . Before our work, the table only contained values up to  $T_m/T_\gamma = 1$ .

In a standard recombination scenario, we would naively not expect the matter temperature to exceed that of the thermal photon field. However, when including non-standard energy injections this limitation becomes artificial and needs to be rectified.

In order to use HYREC with non-standard energy injections, it is necessary to extend the effective transition rate table past matter-radiation temperature equality. This was solely made possible thanks to the author of HYREC Yacine Ali-Haïmoud, who generously lent the code used to compute the effective rates for this purpose.




**Fig. 5.4** Comparison between HYREC and RECFAST for PBHs with differing masses. The  $f_{\text{PBH}}$  chosen for each curve was  $10^{-7}$ . All histories were evaluated with best-fit TT,TE,EE+lowTEB  $\Lambda$ CDM parameters given by [91] and no reionisation.

The inputs to the effective rates code are the bounds and number of points for both the  $T_\gamma$  and  $T_m/T_\gamma$  axes, and the maximum principle quantum number  $n_{\text{max}}$ . The outputs of the code are the effective recombination rates  $\mathcal{A}_{2s}(T_m, T_\gamma)$  and  $\mathcal{A}_{2p}(T_m, T_\gamma)$ , and the transition rate  $\mathcal{R}_{2s,2p}(T_\gamma)$ . The only quantities dependent on the matter temperature are the effective transition rates, and so these are the only ones that need to be extended.

The tables for the  $\mathcal{A}_i$  given in HYREC are effective rates extrapolated to  $n_{\text{max}} = \infty$ . In order to correct for this in our own extended tables, the method outlined in Appendix B of Ali-Haïmoud and Hirata [57] was followed. The code was run for multiple  $n_{\text{max}}$ , the results of which were then used to extrapolate to  $n_{\text{max}} = \infty$  using the following empirical equation:

$$\mathcal{A}_i(T_m, T_\gamma; n_{\text{max}}) = \mathcal{A}_i(T_m, T_\gamma; \infty) \left( 1 - \frac{\kappa}{n_{\text{max}}^\gamma} \right),$$

where  $\mathcal{A}_i(T_m, T_\gamma; \infty)$  is the coefficient being solved for and  $\kappa, \gamma$  are nuisance parameters taking whatever values the fit requires. As there are three unknowns, at least three points evaluated at different  $n_{\text{max}}$  are needed for a successful fit. For this work, the values of  $n_{\text{max}} = 200, 400, 600, 800$  were used, and the fit was completed using a standard numerical least-squares approach.




**Fig. 5.5** Comparison between HYREC and RECFAST for PBHs with differing masses. The  $f_{\text{PBH}}$  chosen for each curve was  $10^{-7}$ . All histories were evaluated with best-fit TT,TE,EE+lowTEB  $\Lambda$ CDM parameters given by [91] and no reionisation.

This equation is fit for each point in temperature space across both axes. Figure 5.6 shows the results of the extrapolation compared to the effective rates evaluated at  $n_{\text{max}} = 1000$ . As can be seen, the overall discrepancy between these two values is quite low, but at low matter/radiation temperatures the improved accuracy of the  $n_{\text{max}} = \infty$  extrapolation shows.

The final extended tables contain 200 log-spaced bins in  $T_{\text{m}}/T_{\gamma}$  space, spanning  $T_{\text{m}}/T_{\gamma} = 0.1$  to 165 482. Log-spacing was chosen over the original linear-spacing, as many orders of magnitude needed to be covered whilst keeping the size of the table small. With this in mind, the number of bins was also chosen to effectively keep the same spacing of the original table at low values of  $T_{\text{m}}/T_{\gamma}$ . The upper bound was chosen such that the resulting table would contain the value  $T_{\text{m}}/T_{\gamma} = 1$ , which is needed by HYREC to compute recombination rates at  $T_{\text{m}} = T_{\gamma}$ .

### Error Handling

During the execution of HYREC, checks are made on values to assert various physical conditions needed in order to successfully compute the ionisation history. In the event that these conditions are not met, HYREC throws an error and stops the program. Before HYREC could be used in CLASS for parameter sampling, these errors needed to be changed to not stop the instance,



**Fig. 5.6** Relative error between the  $\mathcal{A}_i(T_m, T_\gamma; n_{\max} = 1000)$  and  $\mathcal{A}_i(T_m, T_\gamma; \infty)$  tables. Left figure is for 2s, right is for 2p states.

and instead allow it to keep running. To do this, the error handling framework implemented by CLASS was adopted. This involved rewriting all error-throwing functions to return an exit code which is then checked by the `class_call` macro. In the event the exit code is not zero, the program then iteratively returns from each calling function, using macros to effectively build up a stack trace pointing to what part of the code failed. Of course, this sort of error handling does not account for purely programming-based errors such as memory errors leading to segmentation faults.

## 5.5 Results

This section presents the final results, including exclusion bounds for various PBH mass distributions. Firstly, as these results are made using Bayesian parameter estimation, a full disclosure of the priors used is needed. For  $\Lambda$ CDM parameters, the uniform priors are given in the top half of Table 4.1, with priors corresponding to PBH parameters given in Table 5.3. All results are also computed with the MULTINEST parameters given in Table 4.2.

All results given in this work present regions of parameter space that can be excluded from consideration as potentially viable up to some confidence “limit”. These regions are the inverse of the Bayesian 95% credible regions taken from the 2D marginalised posteriors of each pair of parameters plotted. The Bayesian credible region for some limit  $x\%$  is defined as the region which encloses  $x\%$  of the posterior probability. Hence, the associated exclusion regions encapsulate  $(100 - x)\%$  of the probability. It is calculated by performing an integration from the maximum of the probability distribution outwards in such a fashion that the integration boundary is coincident with a surface of constant probability at all times. In the case of our so-called 95% exclusion regions, the total integrated probability of parameters lying in such regions is only 5%.

### 5.5.1 PBH Effects on Power Spectra

To begin with, we can examine the effect Hawking radiation from PBHs has on the observable power spectra. Naively, we would expect that any viable point in the  $\Lambda$ CDM+PBH parameter space will affect the power spectra minimally so as to not be discounted by the accurate *Planck* measurements. As such, we will plot the relative difference between  $\Lambda$ CDM+PBH and a vanilla  $\Lambda$ CDM model run with the same base parameters, taken from Table 4.3. We use a monochromatic mass distribution with  $f_{\text{PBH}} = 10^{-7}$ , and vary the PBH mass to examine the effects this has on the power spectra. Later on, we use a likelihood based on the TT, TE and EE spectra. Consequently, we will examine the effects PBHs have on these spectra here, which will also be an indicator for how well-constrained these parameters might be.

The results are given in Figure 5.7, where the darker line of the same hue indicates the value was negative before its absolute value was taken. As expected, we see heavier PBHs lie much closer to the vanilla model across each spectra, with effects of  $M_{\text{PBH}} = 10^{17}$  g and  $10^{16}$  g PBHs on the spectra sitting below the sub-percent mark.

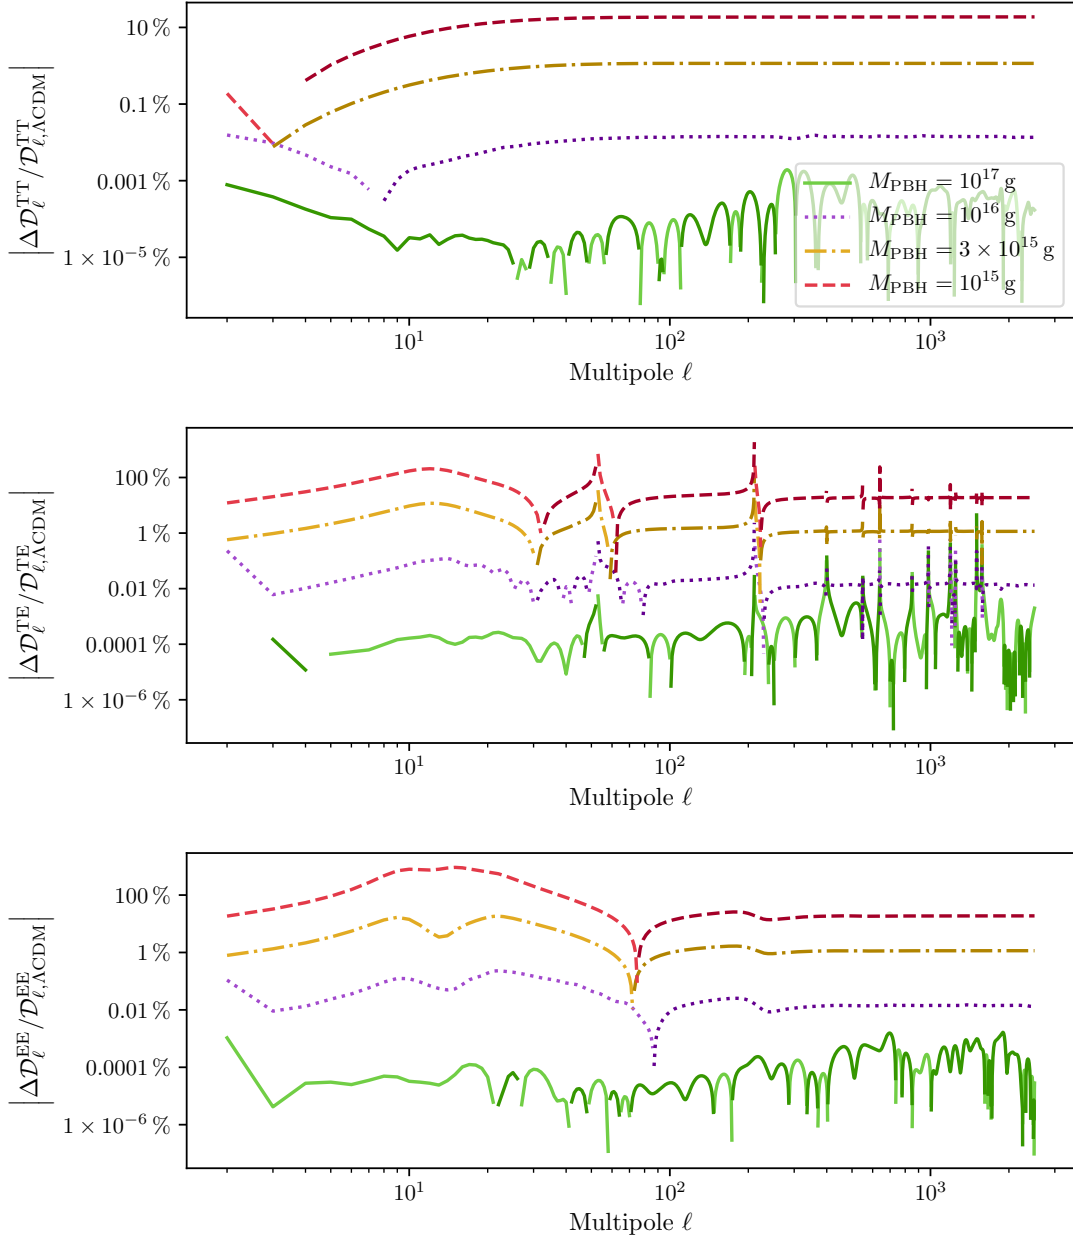
In the TT cross-spectrum, the lighter PBHs do not seem to affect the low multipole region  $\ell < 10$  as much as they do the higher region. As the heating increases, we see that small-scale anisotropies are washed out by a constant factor, which increases considerably in the transition from PBHs of mass  $10^{16}$  g to  $10^{15}$  g. This could be a result of the extra heating from Hawking radiation, which would move the plasma from its original thermal equilibrium resulting in a reduction of oscillations about this temperature at smaller scales.

The EE cross-spectrum displays the opposite behaviour, with the largest discrepancies between PBH and  $\Lambda$ CDM universes appearing over smaller multipoles. In this case, a universe containing PBHs of mass  $10^{15}$  g gives roughly 100 times more power in the polarisation signal at  $\ell \sim 10$  than that given by the vanilla  $\Lambda$ CDM model. This amplification is most likely caused by the Hawking radiation increasing the amount of quadrupole anisotropy that is responsible for polarisation.

We also see a dip in the relative difference where the PBH and  $\Lambda$ CDM spectra cross over. That is to say, for low multipoles we see that the EE spectrum is overestimated by a large amount, before crossing over at  $\ell \sim 70$ . As  $\ell$  increases, we see signal reductions that are approximately of the same order as those for the TT spectrum at high  $\ell$ .

**Table 5.3** Uniform priors used for each PBH parameter. Note all parameters have their priors defined in  $\log_{10}$  space. Parameter definitions can be found throughout Section 5.3.

PBH Parameter	Prior
$\log_{10}(f_{\text{PBH}})$	$[-8, 0]$
$M_{\text{red.}}$	$[5, 7]$
$\log_{10}(\sigma_{10})$	$[-2, 1]$



**Fig. 5.7** Absolute values of the relative differences for TT (top), TE (middle) and EE (bottom) PBH power spectra with a vanilla  $\Lambda$ CDM model. Here,  $f_{\text{PBH}} = 10^{-5}$ , and dark line segments correspond to a negative difference between the vanilla and PBH spectra.

The TE spectrum shows a similar story, with a large increase in signal seen at large scales, before turning to a reduction at smaller scales. The large spikes in the values are purely an artefact of the way we have compared the two spectra. As the TE spectrum has positive and negative values (correlations and anti-correlations), the points where the value changes sign gives rise to large differences if the spectra don't cross zero at the same multipole. If, for instance, at some multipole both spectra have the same small value close to zero, but with a

differing sign, the resulting relative difference between the two will be 200%. Hence, the spikes can be ignored in this case, as their only significance is to indicate where both spectra went to zero.

Overall, it is clear to see that a universe containing PBHs with a single mass of  $M_{\text{PBH}} = 10^{15}$  g at  $f_{\text{PBH}} = 10^{-5}$  would be excluded from *Planck* measurements due to the highly discrepant values of its power spectra compared to the pure  $\Lambda$ CDM model. However, we are unable to say with complete confidence what might happen to PBHs with  $M_{\text{PBH}} = 3 \times 10^{15}$  g, nor are we likely to find the minimum allowed PBH mass at this level of  $f_{\text{PBH}}$ . In order to answer these questions, we first need to establish which likelihood to use, which involves taking a look at how PBHs affect *Planck* nuisance parameters.

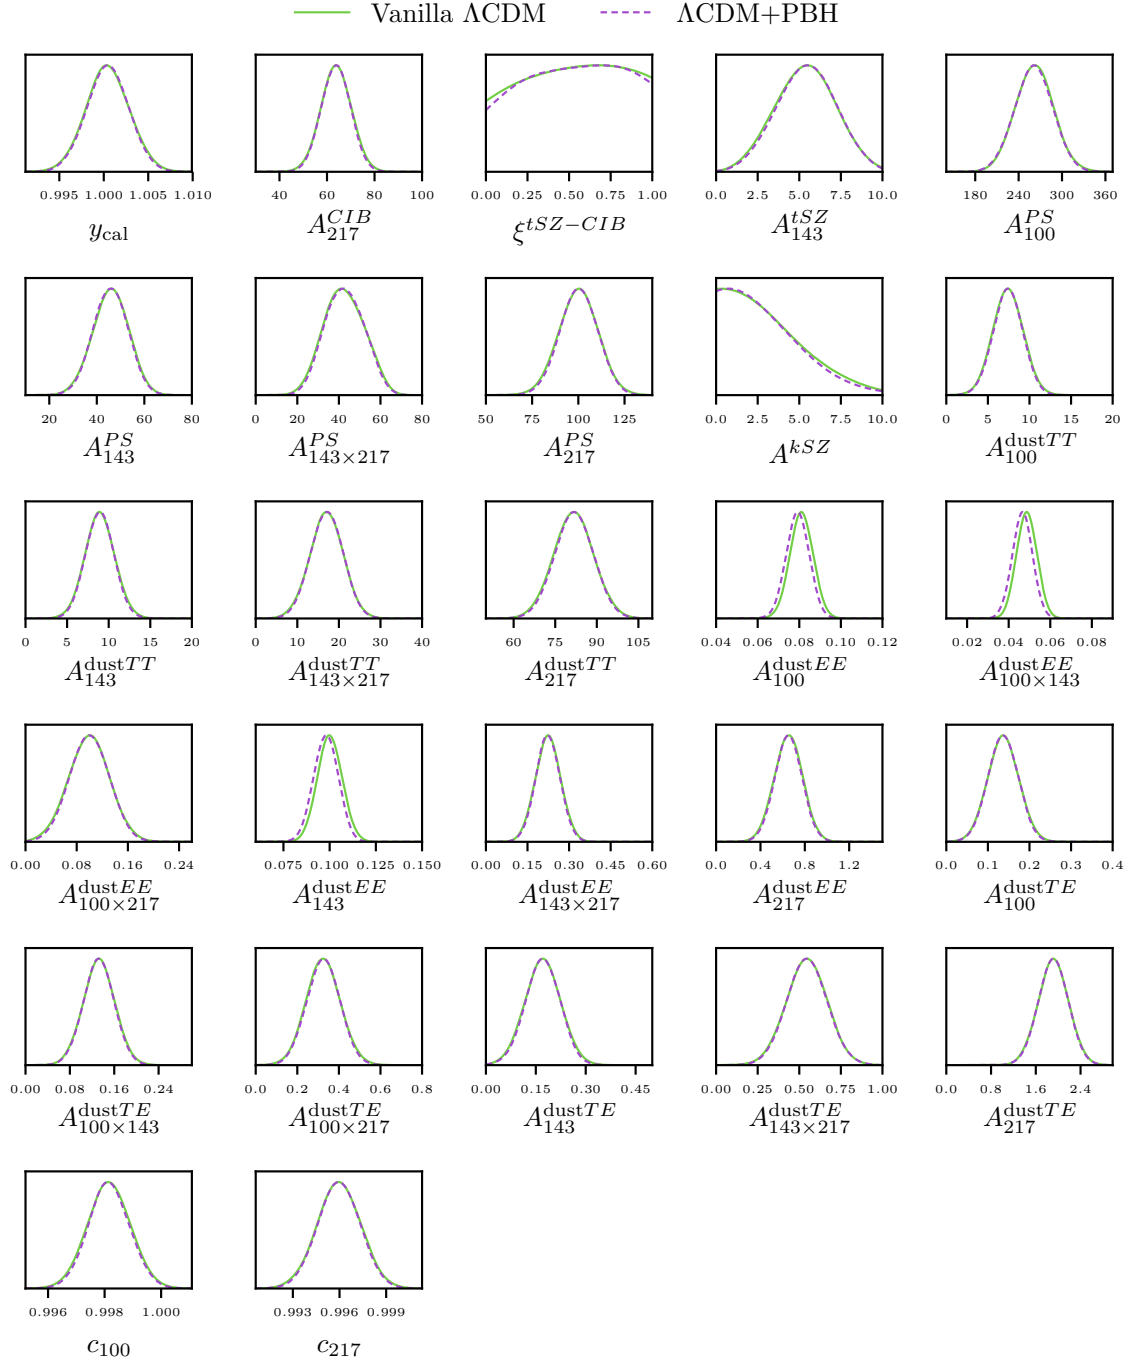
### 5.5.2 PBH Effects on Full High- $\ell$ Likelihood

As mentioned at the end of Section 4.4.4, we must examine the effects that PBHs have on the nuisance parameters to justify their exclusion later on. Here, we will be comparing the nuisance parameter values with added PBHs against the `pc_multinest` baseline result given in Section 4.4.4. This is purely so that any discrepancies present between the two runs is only due to the addition of PBHs, and not any other reason. We will include PBHs into the model through the PBH fraction of total DM,  $f_{\text{PBH}}$ .

To get a better understanding of how  $f_{\text{PBH}}$  might influence the nuisance parameters, we will fix the PBH mass to  $10^{16}$  g, as well as use a uniform prior of  $[0, 1]$  for  $f_{\text{PBH}}$ . The point of using this prior over the usual uniform prior in log space is that we effectively force MULTINEST to sample larger values of  $f_{\text{PBH}}$ , where we expect Hawking radiation will be higher and as a consequence will have more of an effect on the nuisance parameters.

Figure 5.8 shows the 1D marginalised posteriors for all 27 free *Planck* nuisance parameters. We can see that for the most part, there is excellent agreement with the vanilla  $\Lambda$ CDM model. However,  $A_{100}^{\text{dust}EE}$ ,  $A_{100 \times 143}^{\text{dust}EE}$  and  $A_{143}^{\text{dust}EE}$  deviate slightly from their vanilla values. These nuisance parameters determine the amount of residual galactic dust contamination present at  $\ell = 500$  in an EE spectrum made from the frequency bands listed as subscripts. For clarity, these are 100 GHz  $\times$  100 GHz, 100 GHz  $\times$  143 GHz, and 143 GHz  $\times$  143 GHz respectively. This suggests the presence of PBHs could appear as a slight lessening of residual galactic dust contamination in the satellite itself across the 100 GHz and 143 GHz EE cross-spectra.

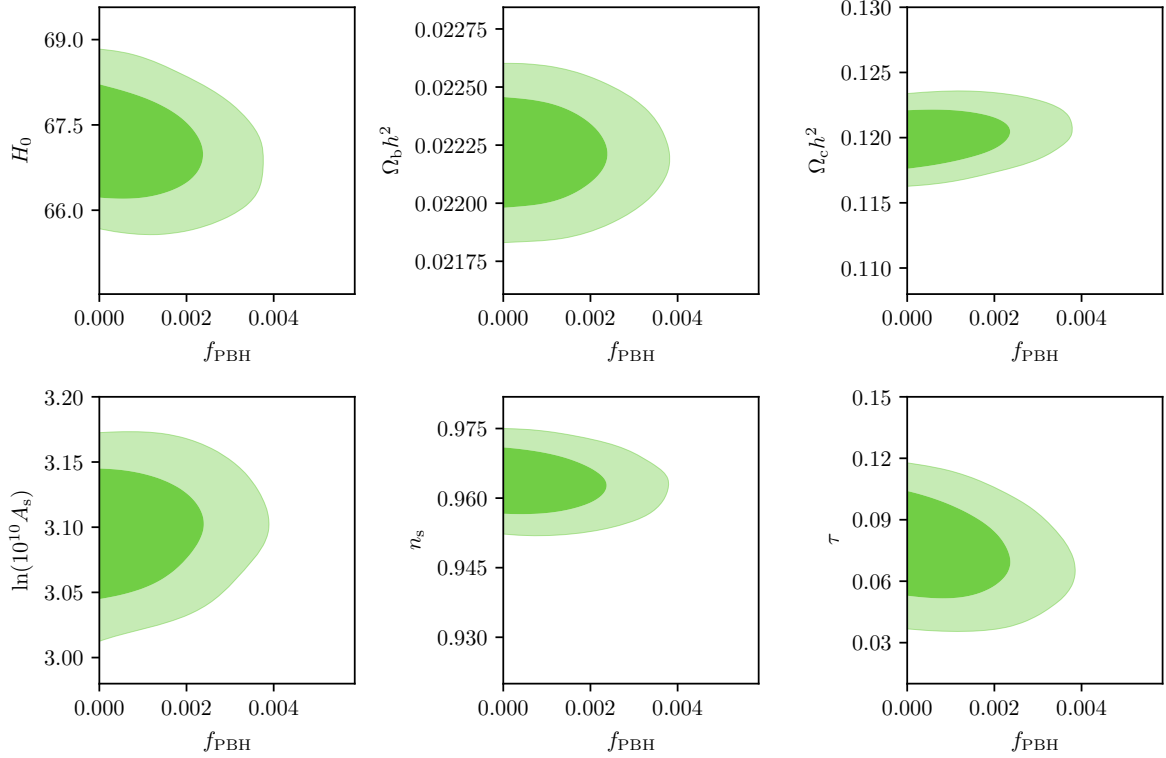
However, in this work we are not interested in the implications this conclusion may carry, and as these are slight deviations from the vanilla model from a set of nuisance parameters that otherwise agree perfectly, we will make use of the lite high- $\ell$  likelihood. This particular likelihood has all nuisance parameters bar the overall *Planck* normalisation  $y_{\text{cal}}$  marginalised out. As the effect of PBHs has only shifted, rather than skewed or otherwise deformed the posteriors of  $A_{100}^{\text{dust}EE}$ ,  $A_{100 \times 143}^{\text{dust}EE}$  and  $A_{143}^{\text{dust}EE}$ , the effect marginalising the vanilla  $\Lambda$ CDM values of these parameters will have on the results will be negligible.



**Fig. 5.8** Comparison of marginalised 1D posteriors for nuisance parameters in a vanilla  $\Lambda$ CDM model, and in a  $\Lambda$ CDM+PBH model with  $f_{\text{PBH}} \in [0, 1]$  a free parameter and  $M_{\text{PBH}} = 10^{16}$  g.

### 5.5.3 Parameter Degeneracies

When adding PBHs to the  $\Lambda$ CDM model, we might expect there to be some degeneracy between the two parameter sets: for example, a change in one PBH parameter might be effectively



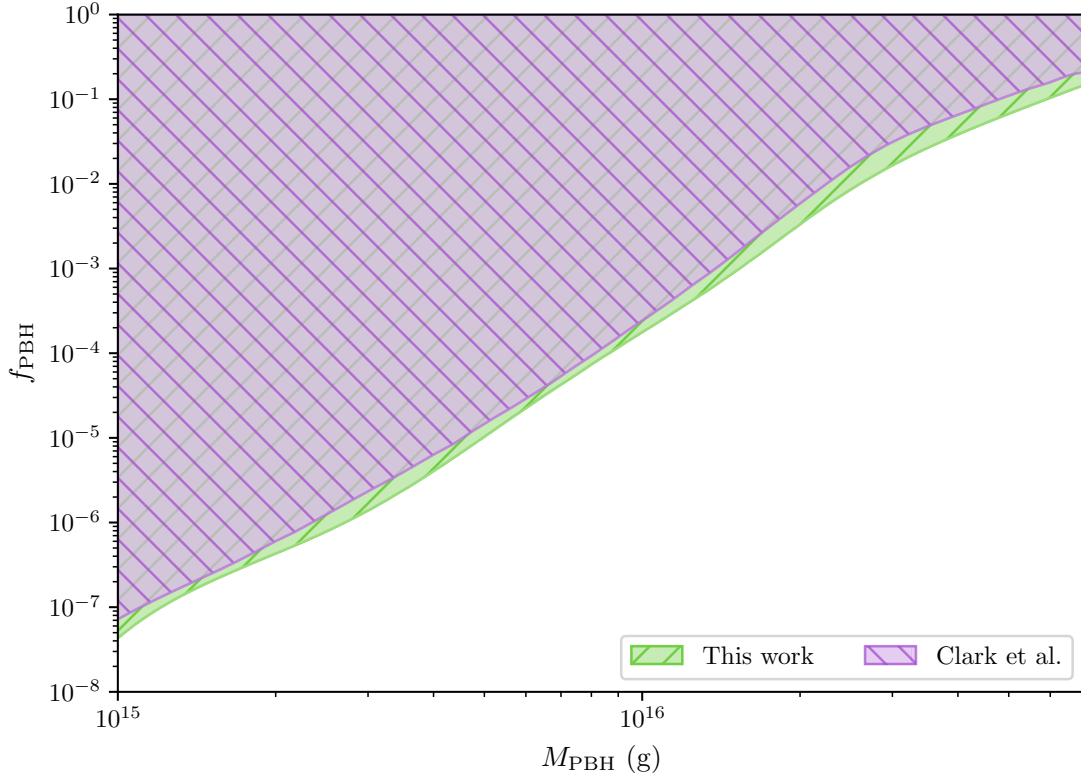
**Fig. 5.9** 2D marginalised posteriors showing the 68% (light) and 95% (dark) Bayesian credible regions for  $f_{\text{PBH}}$  versus each  $\Lambda\text{CDM}$  parameter.

equivalent to a change in an  $\Lambda\text{CDM}$  parameter. For this, we will be following the above analysis in considering the PBH fraction,  $f_{\text{PBH}}$ .

Again, we use a flat prior of  $[0, 1]$  for  $f_{\text{PBH}}$ , and fix  $M_{\text{PBH}}$  to  $10^{16}$  g. For this analysis, we'll also be fixing the CMB data normalisation nuisance parameter  $y_{\text{cal}}$ . In fixing this nuisance parameter, we can be certain that the only variation present in the results are coming directly from varying the  $\Lambda\text{CDM}$  parameters, and not by scaling the CMB data by a different amount. The value we fix  $y_{\text{cal}}$  to is in Table 4.3.

Figure 5.9 shows the 2D marginalised posteriors for  $f_{\text{PBH}}$  against each  $\Lambda\text{CDM}$  parameter. The parameters that show the most degeneracy with  $f_{\text{PBH}}$  are  $\tau$  and  $H_0$ , with slighter dependencies exhibited by  $\ln(10^{10} A_s)$  and  $\Omega_{\text{cdm}} h^2$ . We would expect that Hawking radiation from PBHs would leave the plasma more ionised during recombination than without. This means that by the time the universe goes through reionisation, less matter is able to be ionised. This may then lessen the effect of reionisation, leading to a lower optical depth  $\tau$ .

Similarly, we see the value of  $H_0$  decreases as  $f_{\text{PBH}}$  increases. We note the value of  $H_0$  derived from the vanilla  $\Lambda\text{CDM}$  model is already in tension with astrophysical observations. The effect of PBHs then is to push this value further from agreement, which suggests either PBHs could only be present in very small amounts, or that the *Planck* measurement of  $H_0$




**Fig. 5.10** Comparison of 95 % exclusion regions from this work and Clark et al. [1]. In Clark et al.,  $\Lambda$ CDM parameters are set to the *Planck* best-fit values for the TT,TE,EE+lowTEB likelihood. In this work,  $\Lambda$ CDM parameters are set to the best-fits given in Tab. 4.3. Results are shown up to  $M_{\text{PBH}} = 7 \times 10^{16}$  g.

must be reconciled with astrophysical observations another way. In a manner of speaking, this degeneracy is accidental: the reason we see this link is that the effects of changing  $f_{\text{PBH}}$  and  $H_0$  on the angular power spectra are similar enough that the addition of  $f_{\text{PBH}}$  affects the inference of  $H_0$  from the spectra.

#### 5.5.4 Verification of Clark et al. Results

As much of this work is based off the work done in Clark et al. [1], we begin by presenting our own results using the same analysis to demonstrate the approach taken in this work agrees with the limits given by Clark et al. These results were made by fixing the  $\Lambda$ CDM parameters to their best-fit values found in Table 4.3, and allowing both the reduced PBH mean mass  $M_{\text{red}}$  and PBH fraction  $f_{\text{PBH}}$  to vary with a delta mass distribution. Figure 5.10 shows the associated 95 % exclusion region.

We note that there is very good agreement across the mass range, with this work giving slightly more strict bounds. This may be a result of the confidence limit defined by Clark et al., which differs from the exclusion region used in this work as being the limit “... defined as the cumulative distribution centred around the median...”. The actual reason for this



**Fig. 5.11** 95 % exclusion regions for both free and fixed  $\Lambda$ CDM parameters. Both bounds use a delta mass distribution.

slight discrepancy is not that important, as this plot stands as a good indicator that the PBH recombination physics have been correctly incorporated into CLASS and is also independent verification of the Clark et al. results.

### 5.5.5 Varying $\Lambda$ CDM Parameters

One of the assumptions made in Clark et al. was that when varying PBH parameters, the variation in the base  $\Lambda$ CDM parameters was negligible enough to fix them. By relaxing this assumption, we hope to see the exclusion limits for PBH parameters to be lessened as any effects of Hawking radiation can be effectively nullified by a change in the  $\Lambda$ CDM parameters. Figure 5.11 shows the resulting 95 % exclusion regions. Allowing the base  $\Lambda$ CDM parameters to vary is akin to asking what a universe with PBHs would look like, rather than asking what sort of PBHs we would be able to see in a universe that does not contain them, as we were when fixing the  $\Lambda$ CDM parameters to their best-fit values.

What we see is that the exclusion bounds are lessened by approximately an order of magnitude across the entire mass range. This is in alignment with the prediction outlined by Clark et al. which suggested the constraints on  $f_{\text{PBH}}$  may be weakened by up to a factor of three if  $\Lambda$ CDM parameters are allowed to vary. More interestingly, the bounds seem to be

relaxed more for the extreme ends of the mass range rather than the middle of the range. For the low-mass, high-temperature end, this could be due to the base  $\Lambda$ CDM parameters needing to vary greatly in order to compensate for the large energy injections, whilst at the opposite end the low injected energy may be negligible enough that  $\Lambda$ CDM parameters are allowed to vary as though there is no external injection.

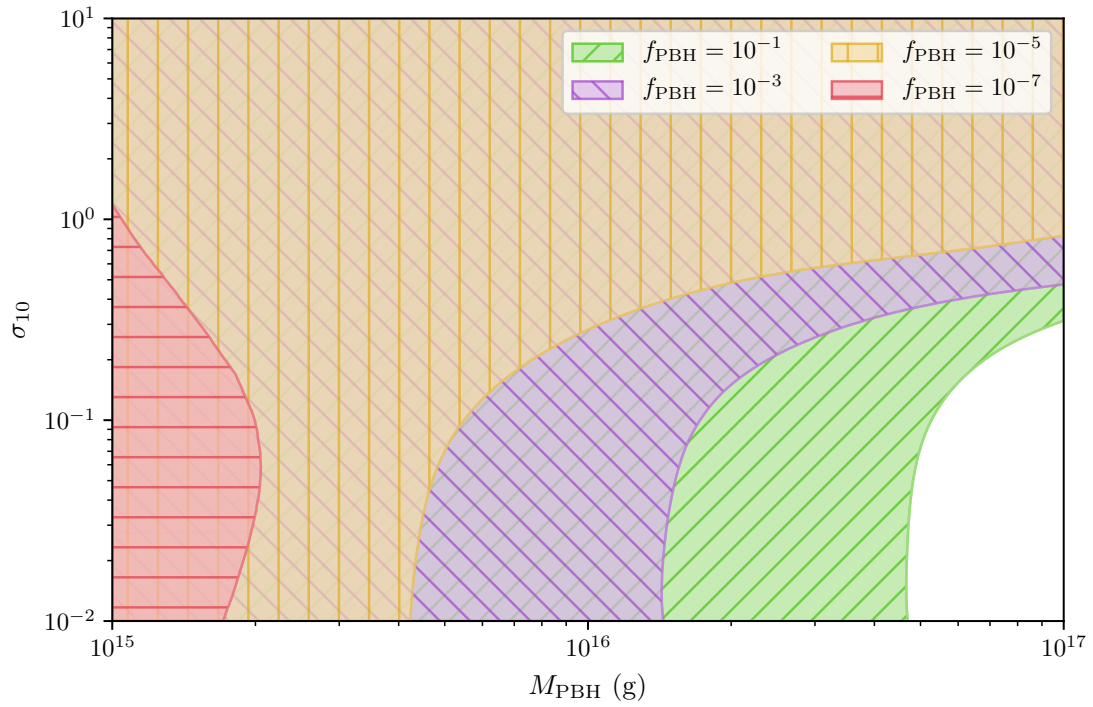
### 5.5.6 Constraints on Distribution Widths

Next, we consider the effects of introducing a finite log normal mass distribution for the PBHs. What might be of interest to examine is the allowed values of the shape parameters  $M_{\text{PBH}}$  and  $\sigma_{10}$  for some fraction of PBHs  $f_{\text{PBH}}$ . Figure 5.12 shows the 95% exclusion regions of such a scenario with  $\Lambda$ CDM values fixed to their best-fit values as given in Table 4.3, whilst Figure 5.13 shows the same regions but for  $\Lambda$ CDM values allowed to vary. In both figures, each exclusion region is overlapping, meaning for example that the  $f_{\text{PBH}} = 10^{-1}$  exclusion region extends beneath all other regions up to the graph's upper left boundaries.

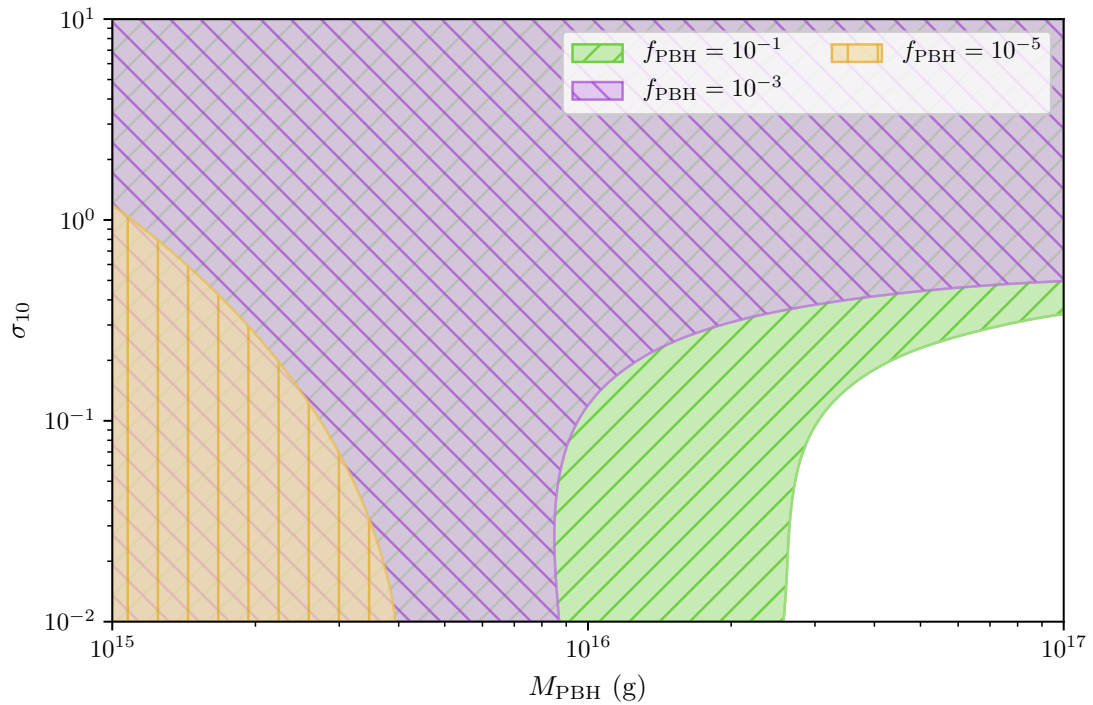
In Fig. 5.12, the exclusion regions for  $f_{\text{PBH}} = 10^{-1}$ ,  $10^{-3}$  and  $10^{-5}$  are such that they completely exclude the region from  $\sigma_{10} = 1$  upwards. At these high values of  $\sigma_{10}$ , the associated log normal distribution spans more orders of magnitude than present in the mass range of  $10^{15}$  g to  $10^{17}$  g, effectively making the distribution uniform and thus independent of  $M_{\text{PBH}}$ . However, for  $f_{\text{PBH}} = 10^{-7}$  we find this region of  $\sigma_{10}$  values is allowed. As this region does not depend on  $M_{\text{PBH}}$ , we expect that it can either be completely excluded or allowed for any given constraint, which is what we see in the results. More interesting is the idea that for some value of  $f_{\text{PBH}}$ , the uniform region switches from being prohibited to allowed. This effectively defines a ‘‘critical value’’ of  $f_{\text{PBH}}$  for which a uniform distribution is excluded, which we will encounter later on.

For those regions that do exclude the high- $\sigma_{10}$  region, their exclusion bounds appear to reach a limiting  $\sigma_{10}$  asymptotically as PBH mass increases. These ‘‘asymptotes’’ are interesting as they tell us that each PBH fraction has some maximum value of  $\sigma_{10}$  they constrain in the high-mass PBH limit. As the width of the PBH mass distribution increases, there must exist some maximum value for which the distribution is excluded no matter its mean value. This is because the distribution approaches a uniform distribution, and as a result incorporates more low-mass PBHs, leading to higher relative Hawking radiation from the  $M_{\text{PBH}}^{-3}$  dependence in Eq. 3.4.5. We see that the decreasing values of  $f_{\text{PBH}}$  correspond to increasing maximum values of  $\sigma_{10}$ , which is what we would expect given the above reasoning.

Similarly, the bounds appear to converge to some  $M_{\text{PBH}}$  value as  $\sigma_{10}$  decreases. We expect to see this behaviour as the width parameter  $\sigma_{10}$  spans several orders of magnitude in  $\log_{10}$  space, meaning that a thinly-peaked distribution with  $\sigma_{10} = 10^{-2}$  behaves almost the same as a delta mass distribution. Hence, as width decreases, the exclusion bounds should ‘‘converge’’ to the corresponding delta distribution bound at that value of  $f_{\text{PBH}}$ . As we have already



**Fig. 5.12** 95% exclusion regions for log normal PBH mass distributions with fixed  $f_{\text{PBH}}$  and  $\Lambda$ CDM parameters.



**Fig. 5.13** 95% exclusion regions for log normal PBH mass distributions with fixed  $f_{\text{PBH}}$  and free  $\Lambda$ CDM parameters. The final  $f_{\text{PBH}} = 10^{-7}$  limit is not shown as it imposes no constraint when  $\Lambda$ CDM parameters are allowed to vary.

computed these bounds for a delta distribution explicitly, we can check to see if these bounds match. Using Fig. 5.11, we found the values of  $M_{\text{PBH}}$  where the exclusion bounds begin for each value of  $f_{\text{PBH}}$  used. Table 5.4 shows the comparison between the two distributions. As can be seen, the agreement between the two is quite good, although not perfect. We must keep in mind, however, that due to the stochastic nature of the Monte Carlo integration used within MULTINEST, and the kernel density estimation employed by the plotting tool `GetDist`, we cannot expect perfect agreement between the two methods. We are also approximating the delta distribution with a thin log normal distribution, which although quite thin at  $\sigma_{10} = 10^{-2}$ , only reaches the delta distribution in the limit as  $\sigma_{10} \rightarrow 0$ .

Comparing Figs. 5.12 and 5.13, we again see that allowing  $\Lambda$ CDM parameters to vary lessens the resulting exclusion bounds. More specifically, for  $f_{\text{PBH}} = 10^{-1}$  and  $10^{-3}$  we see a larger relaxation of the exclusion bounds along the horizontal mean mass direction than the vertical mass width direction. This is similar to what we saw across differing  $f_{\text{PBH}}$  in the fixed  $\Lambda$ CDM case: that increasing the width of the distribution increases the total Hawking radiation much more than reducing the mean PBH mass. As the  $\Lambda$ CDM parameters are free to change, their variation may counteract high Hawking radiation, in turn making the resulting power spectrum seem closer to a vanilla  $\Lambda$ CDM one than it could have been if all  $\Lambda$ CDM parameters were fixed. This then leads to the enlargement of permitted space that we see in the results. For  $f_{\text{PBH}} = 10^{-5}$  we see instead that the freeing of  $\Lambda$ CDM parameters has “freed” the high- $\sigma_{10}$  region corresponding to a uniform distribution. Further, the region excluded by  $f_{\text{PBH}} = 10^{-7}$  disappears, meaning the entire parameter space is permitted when  $\Lambda$ CDM parameters vary.

The region of  $\sigma_{10} \sim 0.3$  is an interesting area. When we consider the maximum allowed values of  $\sigma_{10}$ , a value of 0.3 sits below these in a region that is still quite sensitive to changes in  $\sigma_{10}$ . We can see this by considering values of  $\sigma_{10}$  around this value, and examining the large shifts in allowed masses. On the contrary, varying  $M_{\text{PBH}}$  in this range does not influence the bounds on  $\sigma_{10}$  by as much. Having a width of 0.3 also means the distribution occupies a substantial amount of the mass range. As such, a distribution with this width is representative of a fairly standard log normal — the width is not too large for it to be considered uniform,

**Table 5.4** Comparison between different values of  $M_{\text{PBH}}$  for the delta distribution exclusion regions (Fig. 5.11) and the log normal exclusion regions (Figs. 5.12 and 5.13) at  $\sigma_{10} = 10^{-2}$ .

$f_{\text{PBH}}$	Fixed $\Lambda$ CDM		Free $\Lambda$ CDM	
	Delta $M_{\text{PBH}}$ (g)	$\sigma_{10} = 10^{-2}$ $M_{\text{PBH}}$ (g)	Delta $M_{\text{PBH}}$ (g)	$\sigma_{10} = 10^{-2}$ $M_{\text{PBH}}$ (g)
$10^{-7}$	$1.2 \times 10^{15}$	$1.7 \times 10^{15}$	—	—
$10^{-5}$	$5.0 \times 10^{15}$	$4.2 \times 10^{15}$	$2.9 \times 10^{15}$	$3.9 \times 10^{15}$
$10^{-3}$	$1.5 \times 10^{16}$	$1.4 \times 10^{16}$	$1.1 \times 10^{16}$	$8.7 \times 10^{15}$
$10^{-1}$	$5.9 \times 10^{16}$	$4.7 \times 10^{16}$	$3.5 \times 10^{16}$	$2.6 \times 10^{16}$

nor is it too small to be a delta distribution. We will be considering the properties of this distribution in the next section.


### 5.5.7 Constraints on Distribution Types

The final few plots considered in this work examine the effects changing the mass distribution has on the exclusion regions for  $f_{\text{PBH}}$ . The mass distributions under consideration are a uniform distribution, a log normal distribution with fixed width, a delta distribution and a log normal distribution with the width marginalised over. Although not strictly speaking an actual distribution, marginalising over  $\sigma_{10}$  effectively takes into account all shapes the log normal distribution can take, weighting them by their posterior probability. For the log normal distribution, a width of  $\sigma_{10} = 0.3$  was chosen as this corresponded to the interesting region in Figs. 5.12 and 5.13 where the exclusion bounds were changing the most rapidly in  $\sigma_{10}$ . The uniform distribution was modelled by selecting a width of  $\sigma_{10} = 10$ . As previously discussed, although a uniform distribution had not been implemented into CLASS explicitly, choosing this width results in a distribution that behaves effectively the same as a true uniform distribution.


Figure 5.14 presents the 95 % exclusion regions for these four distributions with base  $\Lambda$ CDM parameters set to their best-fit values in Table 4.3. Similarly, Figure 5.15 shows exclusion regions for the same distributions but with  $\Lambda$ CDM parameters allowed to vary. Both plots have exclusion regions with darker boundaries to more clearly show the bounds when they overlap.

As is to be expected, the uniform mass distribution in Fig. 5.14 has an exclusion region independent of its mean mass value  $M_{\text{PBH}}$ . However, in Fig. 5.15 we see that as mass decreases, the exclusion region recedes a small, but noticeable, amount. In the case of freeing the  $\Lambda$ CDM base parameters, the differing values would counteract larger amounts of Hawking radiation, leading to this slight curve. As a reminder, we expect there to be a slight difference in total energy injected for differing distribution means in this model due to the truncation correction that occurs when the distribution tail falls outside of the range  $10^{15}$  g to  $10^{17}$  g, as covered in Section 5.3.2. Examining the mean  $f_{\text{PBH}}$  values of each exclusion bound, the corresponding 95 % exclusion limits this distribution imposes are  $f_{\text{PBH}} < 6.1 \times 10^{-7}$  for fixed  $\Lambda$ CDM and  $f_{\text{PBH}} < 1.3 \times 10^{-5}$  for free  $\Lambda$ CDM.

The uniform distribution is also more heavily constrained than a delta distribution. We see that for much of the mass range ( $\gtrsim 2 \times 10^{15}$  g for fixed  $\Lambda$ CDM,  $\gtrsim 3 \times 10^{15}$  g for free  $\Lambda$ CDM) the uniform distribution has stronger constraints. Although the uniform distribution has its relative contribution equally spread across  $10^{15}$  g to  $10^{17}$  g, the inclusion of the low masses increases the combined Hawking radiation much more due to its  $M_{\text{PBH}}^{-3}$  proportionality. However, we still see that at low masses, a delta distribution is still more heavily constrained. In fact, the intersection of these two bounds gives us the mass at which a delta distribution behaves like a uniform distribution. For reference, we see that for fixed  $\Lambda$ CDM parameters a PBH mass



**Fig. 5.14** 95 % exclusion regions for varying PBH mass distributions with  $\Lambda$ CDM parameters fixed. The uniform distribution is a log normal distribution with  $\sigma_{10} = 10$ , effectively behaving as a uniform distribution. The log normal distribution has  $\sigma_{10} = 0.3$ , and the marginalised distribution is a log normal marginalised over  $\sigma_{10}$ .



**Fig. 5.15** 95 % exclusion regions for varying PBH mass distributions with  $\Lambda$ CDM parameters freed. The uniform distribution is a log normal distribution with  $\sigma_{10} = 10$ , effectively behaving as a uniform distribution. The log normal distribution has  $\sigma_{10} = 0.3$ , and the marginalised distribution is a log normal marginalised over  $\sigma_{10}$ .

of  $2.3 \times 10^{15}$  g is equivalent to a uniform distribution over  $10^{15}$  g to  $10^{17}$  g, and similarly when  $\Lambda$ CDM parameters are allowed to vary the equivalent mass is at  $3.1 \times 10^{15}$  g.

The remaining mass distributions follow the more characteristic triangle-shaped exclusion region already encountered in the delta distribution exclusion regions in Fig. 5.11. Here, we see that both the log normal and marginalised distributions are more excluded at higher mean mass PBHs than the delta distribution. Due to their finite sizes, both mass distributions incorporate lighter PBHs which inject more energy proportional to  $M_{\text{PBH}}^{-3}$ . As a consequence, for the log normal distribution contributions from the medium to high PBH mass range are effectively dominated by PBHs with masses roughly 0.3 orders of magnitude lighter. The marginalised distribution has comparatively weaker limits as it not only incorporates wider log normal distributions, but as the prior range for  $\sigma_{10}$  spans  $[10^{-2}, 10^1]$ , the marginalisation incorporates mostly thinner log normal distributions that behave closer to a delta distribution.

Interestingly, both distributions have weaker constraints at low PBH mass compared to the delta mass distribution. The log normal distribution has its probability smeared out compared to the delta distribution, meaning contributions at the mean PBH mass value  $M_{\text{PBH}}$  will be weaker in comparison. At the light end of the PBH mass region ( $M_{\text{PBH}} \sim 10^{15}$ ), a log normal distribution is truncated, and as covered in Section 5.3.2, is re-weighted to satisfy the normalisation condition which adjusts for the lost Hawking radiation. As such, a log normal distribution will become increasingly truncated at the light mass end and even though truncation correction will increase its probability (or weighting) around  $M_{\text{PBH}}$  it will always be less than that of a delta distribution. In a similar manner to the previous discussion, the inclusion of heavier PBHs in the non-truncated tail of log normal distributions contributes negligibly to the total Hawking radiation due to the inverse cube dependence on their mass. Consequently, the total Hawking radiation will be less than that of a pure delta distribution, which is why the bounds of the log normal and marginalised distributions cross with the delta distribution bounds at the low-mass end of the spectrum.

Comparing both the fixed and free  $\Lambda$ CDM plots, we see again that allowing the base  $\Lambda$ CDM parameters to vary lessens the exclusion bounds across all mass distributions. The reduction is roughly an order of magnitude across the PBH mass range, in alignment with the previously-examined delta mass distribution.

The net effect of spreading out the PBH mass distribution from a delta distribution to a log normal is to strengthen the exclusion bounds for masses  $M_{\text{PBH}} \gtrsim 2 \times 10^{15}$  g by a few orders of magnitude and weaken exclusion bounds for masses  $M_{\text{PBH}} \lesssim 2 \times 10^{15}$  g by considerably less.



# Chapter 6

## Conclusions

This work has set out to improve the current understanding of the effects primordial black holes might have had on the early universe. Beginning with an outline of our current understanding of the universe, manifesting in the  $\Lambda$ CDM model, the era of recombination was examined in more detail. The standard treatment of the primordial plasma was then extended to show how effects from non-standard contributions such as dark matter and primordial black holes (PBHs) might be incorporated.

Following this, the theory of light PBHs was explained, which presented a method of computing their combined Hawking radiation at some redshift  $z$  for some mass  $M_{\text{PBH}}$  and fraction of dark matter made from PBHs  $f_{\text{PBH}}$ .

In order to use any of this information to place constraints, a Bayesian approach was taken to turn prior probabilities into posterior probabilities through the use of the likelihood function. Nested sampling, and its implementation MULTINEST, was chosen over the more conventional MCMC sampling method due to its intuitive convergence criteria and general robustness.

It was also shown how injected energy from PBHs is split up into different deposition efficiency curves, which by using results given in Slatyer [9] was able to be incorporated into CLASS, allowing PBH energy injections to be calculated on-the-fly. This led to the creation of `pc_multinest`, a program utilising this PBH-extended version of CLASS, MULTINEST and the *Planck* likelihood code to help constrain these new PBH parameters. `pc_multinest` was then used to replicate the *Planck* baseline TT,TE,EE+lowTEB analysis, which it did to good precision.

The final results were computed using `pc_multinest` and the extended version of CLASS incorporating early universe PBH physics. Before calculating exclusion bounds, we saw that Hawking radiation from PBHs acts to dampen small scale fluctuations in the TT, TE and EE power spectra, whilst increasing power in the large to mid-scale region. When considering the effects PBHs have on the *Planck* nuisance parameters, we found a slight shift in only three of the 27 parameters, permitting us to use a lighter version of the TT,TE,EE+lowTEB likelihood which had significantly less parameters to sample. We also examined their effects on the base

$\Lambda$ CDM parameters, and found the Hubble parameter  $H_0$  and the optical depth to recombination  $\tau$  share a slight degeneracy with  $f_{\text{PBH}}$  for PBHs of mass  $10^{16}$  g.

Moving to computing constraints, we were able to reproduce the results presented in Clark et al. [1], before extending them to explicitly show the weakening of exclusion bounds when the base  $\Lambda$ CDM parameters are unfixed and allowed to vary.

We also presented results incorporating a non-monochromatic mass distribution for PBHs. These results showed that constraints on the width of log normal distributions are bound by some maximum width  $\sigma_{10}$  and some minimum mean mass  $M_{\text{PBH}}$ . The maximum width bound was due to distributions with  $\sigma_{10} \sim 1$  behaving increasingly like a uniform distribution, meaning this region could only be all excluded or all allowed at any given time as a uniform distribution is not affected by the choice of a mean mass value. The minimum mean mass bound occurred when the distribution width spanned 0.01 orders of magnitude, where the distribution behaved closer to a delta distribution and as such would reach the constraints of such a distribution. We compared the mass values at these points with their corresponding constrained masses for an actual delta mass distribution, and found good agreement across the two methods.

We also saw that a uniform mass distribution is more heavily constrained than a delta mass distribution at masses above  $2 \times 10^{15}$  g for fixed  $\Lambda$ CDM base parameters, and  $3 \times 10^{15}$  g for free parameters. Likewise, a log normal mass distribution with a width of 0.3 orders of magnitude gives stronger constraints than a uniform one. Overall, non-monochromatic distributions are more heavily constrained than a monochromatic one due to the inclusion of lower-mass PBHs in the tails of their distributions which effectively drown out any contributions from heavier holes.

In the future, this work may be extended by incorporating measurements of baryon acoustic oscillations (BAOs) into the full likelihood computation. These would be an extra independent constraint on early universe physics, and would most likely tighten PBH constraints. It is also well within reason to extend this analysis to more mass distributions, such as a power law  $p(M_{\text{PBH}}|\gamma) \sim M_{\text{PBH}}^{\gamma-1}$  for some selection of  $\gamma$ . As this work is built on many simplifying assumptions, it would also be useful to examine the effects that relaxing these assumptions would have on the final constraints. For example, one might use the full expression in Eq. (5.2.1), similar to the work of Stöcker et al. [21]. One could also use an event generator such as PYTHIA to calculate the PBH particle spectrum from first principles instead of using the analytic form of Eq. (3.4.3).

In conclusion, we have successfully expanded on the current PBH analysis by considering the effects extended mass distributions have on recombination, and also examining the effects produced by varying base  $\Lambda$ CDM parameters.

# Appendix A

## Derivation of Friedmann Equations

The Friedmann equations that govern the dynamics of the universe are derived using the Friedmann-Lemaître-Robertson-Walker metric in Einstein's field equations. We will follow the derivation explicitly here, using the metric sign convention of  $(-, +, +, +)$  throughout.

### A.1 Friedmann-Lemaître-Robertson-Walker Metric

The Friedmann-Lemaître-Robertson-Walker (FLRW) metric is the most general metric permissible by the cosmological principle [35]. It corresponds to a 4D space-time containing a maximally-symmetric hyperspherical 3D subspace. In terms of cosmic time  $t$  and the general hyperspherical coordinates  $(\chi, \theta, \phi)$ , the metric has an invariant interval of

$$ds^2 = -c^2 dt^2 + a^2(t) \left[ d\chi^2 + S_k^2(\chi) (d\theta^2 + \sin^2 \theta d\phi^2) \right], \quad (\text{A.1.1})$$

where the coordinate  $\chi$  has been rescaled such that the curvature takes values of  $k = -1$  for negative curvature,  $k = 0$  for no curvature and  $k = +1$  for positive curvature. The function  $S_k(\chi)$  is piecewise on these three  $k$  values and as a reminder is given by

$$S_k(\chi) = \begin{cases} \sin \chi, & k = +1, \\ \chi, & k = 0, \\ \sinh \chi, & k = -1. \end{cases}$$

Throughout this work we will use the *normalised* scale factor  $a(t)$  defined in terms of the full scale factor  $R(t)$  by

$$a(t) := \frac{R(t)}{R(t_0)}.$$

We will reduce the complexity of our subsequent calculations by changing variables in Eq. (A.1.1) from the hyperspherical  $\chi$  to the physically more meaningful radial coordinate  $r$ .

This results in a simplified invariant interval of the form

$$ds^2 = -c^2 dt^2 + a^2(t) \left[ \frac{dr^2}{1 - kr^2} + r^2 (d\theta^2 + \sin^2 \theta d\phi^2) \right], \quad (\text{A.1.2})$$

for  $k \in \{-1, 0, +1\}$  as before.

Looking at Eq. (A.1.2), we see the spatial part of the metric shares some similarities to a spherical metric in 3D, where  $r$  is the radial coordinate and  $\theta$  and  $\phi$  are the polar and azimuthal angles. In fact, for the case where  $k = 0$ , this part reduces to the 3D spherical polar metric, which is what we would expect for a flat space-time.

The scale factor  $a(t)$ , as its name suggests, acts as a scaling function of the three spatial coordinates as can be seen from observation of Eq. (A.1.2). The time dependence of the scale factor is a mathematical encapsulation of the expanding universe, which as we shall see depends on the components present in the universe. We will use the Einstein field equations to solve for its time evolution.

## A.2 Einstein Field Equations

Einstein's ten field equations in tensor component form are

$$R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}, \quad (\text{A.2.1})$$

where  $R_{\mu\nu}$  is the Ricci tensor,  $R := R^\mu{}_\mu$  is the Ricci scalar,  $\Lambda$  is the cosmological constant and  $T_{\mu\nu}$  is the stress-energy-momentum tensor.

The metric components  $g_{\mu\nu}$  are defined in terms of the invariant interval  $ds^2$ :

$$ds^2 := g_{\mu\nu} dx^\mu dx^\nu,$$

where  $dx^\mu := (dt, dr, d\theta, d\phi)$  and Greek spacetime indices take the coordinate values  $(t, r, \theta, \phi)$ .

By examination of Eq. (A.1.2), the only non-zero components of the metric are given by

$$\begin{aligned} g_{tt} &= -c^2, & g_{rr} &= \frac{a^2(t)}{1 - kr^2}, \\ g_{\theta\theta} &= a^2(t)r^2, & g_{\phi\phi} &= a^2(t)r^2 \sin^2 \theta. \end{aligned}$$

Before continuing, it is convenient to define a shorthand for the 4-divergence  $\partial_\mu$  acting on some tensor  $S_{\alpha\beta}$ :

$$\partial_\mu S_{\alpha\beta} \equiv S_{\alpha\beta,\mu}.$$

All non-zero partial derivatives of the metric components are given by

$$\begin{aligned} g_{rr,t} &= \frac{2a(t)\dot{a}(t)}{1-kr^2}, & g_{rr,r} &= \frac{2kra^2(t)}{(1-kr^2)^2}, \\ g_{\theta\theta,t} &= 2a(t)\dot{a}(t)r^2, & g_{\theta\theta,r} &= 2a^2(t)r, \\ g_{\phi\phi,t} &= 2a(t)\dot{a}(t)r^2 \sin^2 \theta, & g_{\phi\phi,r} &= 2a^2(t)r \sin^2 \theta, \\ g_{\phi\phi,\theta} &= 2a^2(t)r^2 \sin \theta \cos \theta, \end{aligned}$$

where overhead dots represent derivatives with respect to cosmic time.

The non-zero components of the inverse metric tensor are trivially given by the reciprocal of the corresponding component of the metric tensor, as the tensor is diagonal:

$$\begin{aligned} g^{tt} &= -\frac{1}{c^2}, & g^{rr} &= \frac{1-kr^2}{a^2(t)}, \\ g^{\theta\theta} &= \frac{1}{a^2(t)r^2}, & g^{\phi\phi} &= \frac{1}{a^2(t)r^2 \sin^2 \theta}. \end{aligned}$$

We now move towards evaluating the connection coefficients  $\Gamma_{\alpha\beta}^\mu$ . We assume a torsion-free universe with a metric-compatible connection such that

$$\Gamma_{\alpha\beta}^\mu = \frac{1}{2} g^{\mu\lambda} (g_{\alpha\lambda,\beta} + g_{\beta\lambda,\alpha} - g_{\alpha\beta,\lambda})$$

holds.

From the metric components, the non-zero components of the connection are given by

$$\begin{aligned} \Gamma_{rr}^t &= \frac{a(t)\dot{a}(t)}{c^2(1-kr^2)}, & \Gamma_{\theta\theta}^t &= \frac{a(t)\dot{a}(t)r^2}{c^2}, & \Gamma_{\phi\phi}^t &= \frac{a(t)\dot{a}(t)r^2 \sin^2 \theta}{c^2}, \\ \Gamma_{rt}^r &= \Gamma_{tr}^r = \frac{\dot{a}(t)}{a(t)}, & \Gamma_{\theta\theta}^r &= -r(1-kr^2), & \Gamma_{\phi\phi}^r &= -r(1-kr^2) \sin^2 \theta, \\ \Gamma_{\theta t}^\theta &= \Gamma_{t\theta}^\theta = \frac{\dot{a}(t)}{a(t)}, & \Gamma_{\theta r}^\theta &= \Gamma_{r\theta}^\theta = \frac{1}{r}, & \Gamma_{\phi\phi}^\theta &= -\sin \theta \cos \theta, \\ \Gamma_{\phi t}^\phi &= \Gamma_{t\phi}^\phi = \frac{\dot{a}(t)}{a(t)}, & \Gamma_{\phi r}^\phi &= \Gamma_{r\phi}^\phi = \frac{1}{r}, & \Gamma_{\phi\theta}^\phi &= \Gamma_{\theta\phi}^\phi = \cot \theta, \\ \Gamma_{rr}^r &= \frac{kr}{1-kr^2}. \end{aligned}$$

From here, the next step would be to compute the curvature tensor components given by

$$R^\alpha{}_{\beta\gamma\delta} = \partial_\gamma \Gamma_{\delta\beta}^\alpha - \partial_\delta \Gamma_{\gamma\beta}^\alpha + \Gamma_{\delta\beta}^\lambda \Gamma_{\gamma\lambda}^\alpha - \Gamma_{\gamma\beta}^\lambda \Gamma_{\delta\lambda}^\alpha.$$

However, we can save many steps of unnecessary algebra by instead considering the Ricci tensor, defined as a contraction of the curvature tensor across two indices:

$$\begin{aligned} R_{\mu\nu} &:= R^{\alpha}{}_{\mu\alpha\nu} = \partial_{\alpha}\Gamma_{\nu\mu}^{\alpha} - \partial_{\nu}\Gamma_{\alpha\mu}^{\alpha} + \Gamma_{\nu\mu}^{\lambda}\Gamma_{\alpha\lambda}^{\alpha} - \Gamma_{\alpha\mu}^{\lambda}\Gamma_{\nu\lambda}^{\alpha} \\ &= \partial_{\alpha}\Gamma_{\mu\nu}^{\alpha} - \partial_{\nu}\Gamma_{\alpha\mu}^{\alpha} + \Gamma_{\mu\nu}^{\beta}\Gamma_{\alpha\beta}^{\alpha} - \Gamma_{\alpha\mu}^{\beta}\Gamma_{\nu\beta}^{\alpha}, \end{aligned}$$

where the last line comes from the symmetry of the connection coefficients in the lower two indices and making the dummy substitution  $\lambda \leftrightarrow \beta$ .

Using the above form for the Ricci tensor, we find that all the non-zero components are diagonal:

$$\begin{aligned} R_{tt} &= -3\frac{\ddot{a}(t)}{a(t)}, \\ R_{rr} &= \frac{\ddot{a}(t)a(t) + 2\dot{a}^2(t) + 2kc^2}{c^2(1 - kr^2)}, \\ R_{\theta\theta} &= \frac{r^2[\ddot{a}(t)a(t) + 2\dot{a}^2(t) + 2kc^2]}{c^2}, \\ R_{\phi\phi} &= \frac{r^2[\ddot{a}(t)a(t) + 2\dot{a}^2(t) + 2kc^2]\sin^2\theta}{c^2}. \end{aligned}$$

The Ricci scalar is then given by

$$\begin{aligned} R &:= R^{\mu}{}_{\mu} \equiv R_{\mu\nu}g^{\mu\nu} \\ &= -3\frac{\ddot{a}(t)}{a(t)}\left(-\frac{1}{c^2}\right) + \frac{\ddot{a}(t)a(t) + 2\dot{a}^2(t) + 2kc^2}{c^2(1 - kr^2)}\left(\frac{1 - kr^2}{a^2(t)}\right) \\ &\quad + \frac{r^2[\ddot{a}(t)a(t) + 2\dot{a}^2(t) + 2kc^2]}{c^2}\left(\frac{1}{a^2(t)r^2}\right) \\ &\quad + \frac{r^2[\ddot{a}(t)a(t) + 2\dot{a}^2(t) + 2kc^2]\sin^2\theta}{c^2}\left(\frac{1}{a^2(t)r^2\sin^2\theta}\right) \\ &= 3\frac{\ddot{a}(t)}{c^2a(t)} + 3\frac{\ddot{a}(t)a(t) + 2\dot{a}^2(t) + 2kc^2}{c^2a^2(t)} \\ &= 6\frac{\ddot{a}(t)a(t) + \dot{a}^2(t) + kc^2}{c^2a^2(t)}. \end{aligned}$$

The final component we need in order to solve the Einstein field equations is the stress-energy tensor  $T_{\mu\nu}$ . We simplify matters by modelling the contents of the universe as a fluid with no viscous properties. This is known as a *perfect fluid*, and has a stress-energy tensor of the form

$$T_{\mu\nu} = \left(\rho + \frac{p}{c^2}\right)u_{\mu}u_{\nu} + pg_{\mu\nu},$$

where  $\rho$  is the proper density of the fluid and  $p$  is its pressure. Due to the fundamental assumptions of isotropy and homogeneity, neither of these quantities depend on position but they can depend on cosmic time  $t$ .

We will consider the solutions to Einstein's field equations in the frame where the universe appears isotropic and homogeneous — the *co-moving frame*. The 4-velocity of the fluid must take the form of

$$u^\mu = (1, 0, 0, 0) \equiv \delta_t^\mu,$$

as it is stationary in the comoving frame.

The covariant components are then simply

$$u_\mu := g_{\mu\nu}u^\nu = g_{\mu t} \equiv g_{tt}\delta_\mu^t = -c^2\delta_\mu^t,$$

by the diagonal form of the tensor, and with  $g_{tt} = -c^2$ .

Using this, the stress-energy tensor simplifies to

$$T_{\mu\nu} = (\rho c^2 + p)c^2\delta_\mu^t\delta_\nu^t + pg_{\mu\nu}.$$

As the metric, Ricci and stress-energy tensors are diagonal, there are only four Einstein field equations. From Eq. (A.2.1), the equation given by the  $tt$  components simplifies to

$$3 \frac{\dot{a}^2(t) + kc^2}{a^2(t)} - \Lambda c^2 = 8\pi G\rho. \quad (\text{A.2.2})$$

The remaining  $rr$ ,  $\theta\theta$  and  $\phi\phi$  spatial components of the Einstein field equations lead to the same equation:

$$-2\ddot{a}(t)a(t) - \dot{a}^2(t) - kc^2 + \Lambda c^2 a^2(t) = \frac{8\pi G}{c^2} pa^2(t).$$

Rearranging the above and substituting in Eq. (A.2.2):

$$\begin{aligned} \frac{\ddot{a}(t)}{a(t)} &= -4\pi G \frac{p}{c^2} - \frac{1}{2} \left( \frac{\dot{a}^2(t) + kc^2}{a^2(t)} \right) + \frac{1}{2} \Lambda c^2 \\ &= -4\pi G \frac{p}{c^2} - \frac{1}{6} (8\pi G\rho + \Lambda c^2) + \frac{1}{2} \Lambda c^2 \\ &= -\frac{4\pi G}{3} \left( \rho + \frac{3p}{c^2} \right) + \frac{1}{3} \Lambda c^2. \end{aligned}$$

The above equation, alongside a rearranged Eq. (A.2.2), form the Friedmann equations and govern the evolution of the scale parameter  $a(t)$  for given density  $\rho$ , pressure  $p$  and cosmological constant  $\Lambda$ :

$$\frac{\ddot{a}(t)}{a(t)} = -\frac{4\pi G}{3} \left( \rho + \frac{3p}{c^2} \right) + \frac{\Lambda c^2}{3}, \quad \left( \frac{\dot{a}(t)}{a(t)} \right)^2 = \frac{8\pi G\rho}{3} - \frac{kc^2}{a^2(t)} + \frac{\Lambda c^2}{3}.$$

These equations can be rewritten in a more succinct form by defining the Hubble parameter

$$H(t) := \frac{\dot{a}(t)}{a(t)},$$

giving

$$\dot{H} + H^2 = -\frac{4\pi G}{3} \left( \rho + \frac{3p}{c^2} \right) + \frac{\Lambda c^2}{3}, \quad H^2 = \frac{8\pi G \rho}{3} - \frac{kc^2}{a^2} + \frac{\Lambda c^2}{3}.$$

# References

- [1] S. J. Clark et al. “Planck constraint on relic primordial black holes”. In: *Phys. Rev. D* 95.8, 083006 (Apr. 2017), p. 083006. DOI: [10.1103/PhysRevD.95.083006](https://doi.org/10.1103/PhysRevD.95.083006). arXiv: [1612.07738](https://arxiv.org/abs/1612.07738).
- [2] A. A. Penzias and R. W. Wilson. “A Measurement of Excess Antenna Temperature at 4080 Mc/s.” In: *ApJ* 142 (July 1965), pp. 419–421. DOI: [10.1086/148307](https://doi.org/10.1086/148307).
- [3] A. G. Riess et al. “Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant”. In: *AJ* 116 (Sept. 1998), pp. 1009–1038. DOI: [10.1086/300499](https://doi.org/10.1086/300499). arXiv: [astro-ph/9805201](https://arxiv.org/abs/astro-ph/9805201).
- [4] S. Perlmutter et al. “Measurements of  $\Omega$  and  $\Lambda$  from 42 High-Redshift Supernovae”. In: *ApJ* 517 (June 1999), pp. 565–586. DOI: [10.1086/307221](https://doi.org/10.1086/307221). arXiv: [astro-ph/9812133](https://arxiv.org/abs/astro-ph/9812133).
- [5] P. J. E. Peebles. “Recombination of the Primeval Plasma”. In: *ApJ* 153 (July 1968), p. 1. DOI: [10.1086/149628](https://doi.org/10.1086/149628).
- [6] Y. B. Zeldovich, V. G. Kurt, and R. A. Syunyaev. “Recombination of Hydrogen in the Hot Model of the Universe”. In: *Zhurnal Eksperimentalnoi i Teoreticheskoi Fiziki* 55 (July 1968), pp. 278–286.
- [7] T. R. Slatyer. “Energy injection and absorption in the cosmic dark ages”. In: *Phys. Rev. D* 87.12, 123513 (June 2013), p. 123513. DOI: [10.1103/PhysRevD.87.123513](https://doi.org/10.1103/PhysRevD.87.123513). arXiv: [1211.0283](https://arxiv.org/abs/1211.0283) [[astro-ph](https://arxiv.org/abs/astro-ph)].
- [8] T. R. Slatyer. “Indirect dark matter signatures in the cosmic dark ages. I. Generalizing the bound on s -wave dark matter annihilation from Planck results”. In: *Phys. Rev. D* 93.2, 023527 (Jan. 2016), p. 023527. DOI: [10.1103/PhysRevD.93.023527](https://doi.org/10.1103/PhysRevD.93.023527). arXiv: [1506.03811](https://arxiv.org/abs/1506.03811) [[hep-ph](https://arxiv.org/abs/hep-ph)].
- [9] T. R. Slatyer. “Indirect dark matter signatures in the cosmic dark ages. II. Ionization, heating, and photon production from arbitrary energy injections”. In: *Phys. Rev. D* 93.2, 023521 (Jan. 2016), p. 023521. DOI: [10.1103/PhysRevD.93.023521](https://doi.org/10.1103/PhysRevD.93.023521). arXiv: [1506.03812](https://arxiv.org/abs/1506.03812).
- [10] G. Giesen et al. “CMB photons shedding light on dark matter”. In: *J. Cosmology Astropart. Phys.* 12, 008 (Dec. 2012), p. 008. DOI: [10.1088/1475-7516/2012/12/008](https://doi.org/10.1088/1475-7516/2012/12/008). arXiv: [1209.0247](https://arxiv.org/abs/1209.0247).
- [11] S. Galli et al. “CMB constraints on dark matter models with large annihilation cross section”. In: *Phys. Rev. D* 80.2, 023505 (July 2009), p. 023505. DOI: [10.1103/PhysRevD.80.023505](https://doi.org/10.1103/PhysRevD.80.023505). arXiv: [0905.0003](https://arxiv.org/abs/0905.0003) [[astro-ph](https://arxiv.org/abs/astro-ph)].
- [12] S. Galli et al. “Updated CMB constraints on dark matter annihilation cross sections”. In: *Phys. Rev. D* 84.2, 027302 (July 2011), p. 027302. DOI: [10.1103/PhysRevD.84.027302](https://doi.org/10.1103/PhysRevD.84.027302). arXiv: [1106.1528](https://arxiv.org/abs/1106.1528) [[astro-ph](https://arxiv.org/abs/astro-ph)].
- [13] S. Galli et al. “Systematic uncertainties in constraining dark matter annihilation from the cosmic microwave background”. In: *Phys. Rev. D* 88.6, 063502 (Sept. 2013), p. 063502. DOI: [10.1103/PhysRevD.88.063502](https://doi.org/10.1103/PhysRevD.88.063502). arXiv: [1306.0563](https://arxiv.org/abs/1306.0563).

- [14] B. P. Abbott et al. “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: *Physical Review Letters* 116.6, 061102 (Feb. 2016), p. 061102. DOI: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102). arXiv: [1602.03837 \[gr-qc\]](https://arxiv.org/abs/1602.03837).
- [15] M. Sasaki et al. “Primordial Black Hole Scenario for the Gravitational-Wave Event GW150914”. In: *Physical Review Letters* 117.6, 061101 (Aug. 2016), p. 061101. DOI: [10.1103/PhysRevLett.117.061101](https://doi.org/10.1103/PhysRevLett.117.061101). arXiv: [1603.08338](https://arxiv.org/abs/1603.08338).
- [16] I. Cholis et al. “Orbital eccentricities in primordial black hole binaries”. In: *Phys. Rev. D* 94.8, 084013 (Oct. 2016), p. 084013. DOI: [10.1103/PhysRevD.94.084013](https://doi.org/10.1103/PhysRevD.94.084013). arXiv: [1606.07437 \[astro-ph.HE\]](https://arxiv.org/abs/1606.07437).
- [17] S. Blinnikov et al. “Solving puzzles of GW150914 by primordial black holes”. In: *J. Cosmology Astropart. Phys.* 11, 036 (Nov. 2016), p. 036. DOI: [10.1088/1475-7516/2016/11/036](https://doi.org/10.1088/1475-7516/2016/11/036). arXiv: [1611.00541 \[astro-ph.HE\]](https://arxiv.org/abs/1611.00541).
- [18] S. Clesse and J. Garca-Bellido. “The clustering of massive Primordial Black Holes as Dark Matter: Measuring their mass distribution with advanced LIGO”. In: *Physics of the Dark Universe* 15 (Mar. 2017), pp. 142–147. DOI: [10.1016/j.dark.2016.10.002](https://doi.org/10.1016/j.dark.2016.10.002). arXiv: [1603.05234](https://arxiv.org/abs/1603.05234).
- [19] Y. Ali-Haïmoud and M. Kamionkowski. “Cosmic microwave background limits on accreting primordial black holes”. In: *Phys. Rev. D* 95.4, 043534 (Feb. 2017), p. 043534. DOI: [10.1103/PhysRevD.95.043534](https://doi.org/10.1103/PhysRevD.95.043534). arXiv: [1612.05644](https://arxiv.org/abs/1612.05644).
- [20] V. Poulin et al. “CMB bounds on disk-accreting massive primordial black holes”. In: *Phys. Rev. D* 96.8, 083524 (Oct. 2017), p. 083524. DOI: [10.1103/PhysRevD.96.083524](https://doi.org/10.1103/PhysRevD.96.083524). arXiv: [1707.04206](https://arxiv.org/abs/1707.04206).
- [21] P. Stöcker et al. “Exotic energy injection with ExoCLASS: application to the Higgs portal model and evaporating black holes”. In: *J. Cosmology Astropart. Phys.* 3, 018 (Mar. 2018), p. 018. DOI: [10.1088/1475-7516/2018/03/018](https://doi.org/10.1088/1475-7516/2018/03/018). arXiv: [1801.01871](https://arxiv.org/abs/1801.01871).
- [22] S. Clesse and J. Garca-Bellido. “Massive primordial black holes from hybrid inflation as dark matter and the seeds of galaxies”. In: *Phys. Rev. D* 92.2, 023524 (July 2015), p. 023524. DOI: [10.1103/PhysRevD.92.023524](https://doi.org/10.1103/PhysRevD.92.023524). arXiv: [1501.07565](https://arxiv.org/abs/1501.07565).
- [23] B. Carr et al. “Primordial black hole constraints for extended mass functions”. In: *Phys. Rev. D* 96.2, 023514 (July 2017), p. 023514. DOI: [10.1103/PhysRevD.96.023514](https://doi.org/10.1103/PhysRevD.96.023514). arXiv: [1705.05567](https://arxiv.org/abs/1705.05567).
- [24] F. Kühnel and K. Freese. “Constraints on primordial black holes with extended mass functions”. In: *Phys. Rev. D* 95.8, 083508 (Apr. 2017), p. 083508. DOI: [10.1103/PhysRevD.95.083508](https://doi.org/10.1103/PhysRevD.95.083508). arXiv: [1701.07223](https://arxiv.org/abs/1701.07223).
- [25] N. Bellomo et al. “Primordial black holes as dark matter: converting constraints from monochromatic to extended mass distributions”. In: *J. Cosmology Astropart. Phys.* 1, 004 (Jan. 2018), p. 004. DOI: [10.1088/1475-7516/2018/01/004](https://doi.org/10.1088/1475-7516/2018/01/004). arXiv: [1709.07467](https://arxiv.org/abs/1709.07467).
- [26] J. Calcino, J. Garca-Bellido, and T. M. Davis. “Updating the MACHO fraction of the Milky Way dark halowith improved mass models”. In: *MNRAS* 479 (Sept. 2018), pp. 2889–2905. DOI: [10.1093/mnras/sty1368](https://doi.org/10.1093/mnras/sty1368). arXiv: [1803.09205](https://arxiv.org/abs/1803.09205).
- [27] G. Lemaître. “Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques”. In: *Annales de la Société Scientifique de Bruxelles* 47 (1927), pp. 49–59.
- [28] V. M. Slipher. “The radial velocity of the Andromeda Nebula”. In: *Lowell Observatory Bulletin* 2 (1913), pp. 56–57.
- [29] V. M. Slipher. “Spectrographic Observations of Nebulae”. In: *Popular Astronomy* 23 (Jan. 1915), pp. 21–24.

- [30] E. Hubble. “A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae”. In: *Proceedings of the National Academy of Science* 15 (Mar. 1929), pp. 168–173. DOI: [10.1073/pnas.15.3.168](https://doi.org/10.1073/pnas.15.3.168).
- [31] B. P. Crill et al. “BOOMERANG: A Balloon-borne Millimeter-Wave Telescope and Total Power Receiver for Mapping Anisotropy in the Cosmic Microwave Background”. In: *ApJS* 148 (Oct. 2003), pp. 527–541. DOI: [10.1086/376894](https://doi.org/10.1086/376894). arXiv: [astro-ph/0206254](https://arxiv.org/abs/astro-ph/0206254).
- [32] S. Hanany et al. “MAXIMA-1: A Measurement of the Cosmic Microwave Background Anisotropy on Angular Scales of 10' to 5°”. In: *ApJ* 545 (Dec. 2000), pp. L5–L9. DOI: [10.1086/317322](https://doi.org/10.1086/317322). arXiv: [astro-ph/0005123](https://arxiv.org/abs/astro-ph/0005123).
- [33] G. Hütsi et al. “WMAP7 and future CMB constraints on annihilating dark matter: implications for GeV-scale WIMPs”. In: *A&A* 535, A26 (Nov. 2011), A26. DOI: [10.1051/0004-6361/201116914](https://doi.org/10.1051/0004-6361/201116914). arXiv: [1103.2766](https://arxiv.org/abs/1103.2766).
- [34] Planck Collaboration et al. “Planck 2013 results. I. Overview of products and scientific results”. In: *A&A* 571, A1 (Nov. 2014), A1. DOI: [10.1051/0004-6361/201321529](https://doi.org/10.1051/0004-6361/201321529). arXiv: [1303.5062](https://arxiv.org/abs/1303.5062).
- [35] M. P. Hobson, G. P. Efstathiou, and A. N. Lasenby. *General Relativity. An Introduction for Physicists*. 2007 reprint. Cambridge, United Kingdom: Cambridge University Press, 2006. ISBN: 9780521829519.
- [36] W. J. Percival et al. “The 2dF Galaxy Redshift Survey: the power spectrum and the matter content of the Universe”. In: *MNRAS* 327 (Nov. 2001), pp. 1297–1306. DOI: [10.1046/j.1365-8711.2001.04827.x](https://doi.org/10.1046/j.1365-8711.2001.04827.x). arXiv: [astro-ph/0105252](https://arxiv.org/abs/astro-ph/0105252).
- [37] D. H. Jones et al. “The 6dF Galaxy Survey: samples, observational techniques and the first data release”. In: *MNRAS* 355 (Dec. 2004), pp. 747–763. DOI: [10.1111/j.1365-2966.2004.08353.x](https://doi.org/10.1111/j.1365-2966.2004.08353.x). arXiv: [astro-ph/0403501](https://arxiv.org/abs/astro-ph/0403501).
- [38] Morag I. Scrimgeour et al. “The WiggleZ Dark Energy Survey: the transition to large-scale cosmic homogeneity”. In: *MNRAS* 425 (Sept. 2012), pp. 116–134. DOI: [10.1111/j.1365-2966.2012.21402.x](https://doi.org/10.1111/j.1365-2966.2012.21402.x). arXiv: [1205.6812](https://arxiv.org/abs/1205.6812) [[astro-ph](https://arxiv.org/abs/astro-ph).CO].
- [39] K. S. Dawson et al. “The Baryon Oscillation Spectroscopic Survey of SDSS-III”. In: *AJ* 145, 10 (Jan. 2013), p. 10. DOI: [10.1088/0004-6256/145/1/10](https://doi.org/10.1088/0004-6256/145/1/10). arXiv: [1208.0022](https://arxiv.org/abs/1208.0022).
- [40] L. Anderson et al. “The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: baryon acoustic oscillations in the Data Releases 10 and 11 Galaxy samples”. In: *MNRAS* 441 (June 2014), pp. 24–62. DOI: [10.1093/mnras/stu523](https://doi.org/10.1093/mnras/stu523). arXiv: [1312.4877](https://arxiv.org/abs/1312.4877).
- [41] F. Zwicky. “Die Rotverschiebung von extragalaktischen Nebeln”. In: *Helvetica Physica Acta* 6 (1933), pp. 110–127.
- [42] V. C. Rubin and W. K. Ford Jr. “Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions”. In: *ApJ* 159 (Feb. 1970), p. 379. DOI: [10.1086/150317](https://doi.org/10.1086/150317).
- [43] W. J. G. de Blok et al. “High-Resolution Rotation Curves and Galaxy Mass Models from THINGS”. In: *AJ* 136, 2648–2719 (Dec. 2008), pp. 2648–2719. DOI: [10.1088/0004-6256/136/6/2648](https://doi.org/10.1088/0004-6256/136/6/2648). arXiv: [0810.2100](https://arxiv.org/abs/0810.2100).
- [44] S. M. Kent and J. E. Gunn. “The dynamics of rich clusters of galaxies. I - The Coma cluster”. In: *AJ* 87 (July 1982), pp. 945–971. DOI: [10.1086/113178](https://doi.org/10.1086/113178).
- [45] G. des Forêts et al. “Simultaneous study of optical and X-ray properties of the Coma Cluster by multi-mass models”. In: *ApJ* 280 (May 1984), pp. 15–25. DOI: [10.1086/161963](https://doi.org/10.1086/161963).
- [46] J. A. Tyson et al. “Deep CCD images of 2345 + 007 - Lensing by dark matter”. In: *AJ* 91 (June 1986), pp. 1274–1278. DOI: [10.1086/114102](https://doi.org/10.1086/114102).

- [47] M. Markevitch et al. “Direct Constraints on the Dark Matter Self-Interaction Cross Section from the Merging Galaxy Cluster 1E 0657-56”. In: *ApJ* 606 (May 2004), pp. 819–824. DOI: [10.1086/383178](https://doi.org/10.1086/383178). arXiv: [astro-ph/0309303](https://arxiv.org/abs/astro-ph/0309303).
- [48] T. M. Davis and C. H. Lineweaver. “Expanding Confusion: Common Misconceptions of Cosmological Horizons and the Superluminal Expansion of the Universe”. In: *PASA* 21 (2004), pp. 97–109. DOI: [10.1071/AS03040](https://doi.org/10.1071/AS03040). arXiv: [astro-ph/0310808](https://arxiv.org/abs/astro-ph/0310808).
- [49] M. Zaldarriaga and U. Seljak. “All-sky analysis of polarization in the microwave background”. In: *Phys. Rev. D* 55 (Feb. 1997), pp. 1830–1840. DOI: [10.1103/PhysRevD.55.1830](https://doi.org/10.1103/PhysRevD.55.1830). arXiv: [astro-ph/9609170](https://arxiv.org/abs/astro-ph/9609170).
- [50] Y. Ali-Haïmoud and C. M. Hirata. “Ultrafast effective multilevel atom method for primordial hydrogen recombination”. In: *Phys. Rev. D* 82.6, 063521 (Sept. 2010), p. 063521. DOI: [10.1103/PhysRevD.82.063521](https://doi.org/10.1103/PhysRevD.82.063521). arXiv: [1006.1355](https://arxiv.org/abs/1006.1355).
- [51] C. M. Hirata. *Recent Developments in Cosmological Recombination*. Mar. 2009. URL: [http://cosmology.lbl.gov/talks/Hirata\\_09.pdf](http://cosmology.lbl.gov/talks/Hirata_09.pdf).
- [52] J. Chluba. *Recombination Physics and What this has to do with Cosmology and Particle Physics*. Nov. 2013. URL: [http://www.jb.man.ac.uk/~jchluba/Science/Recombination\\_Physics.pdf](http://www.jb.man.ac.uk/~jchluba/Science/Recombination_Physics.pdf).
- [53] R. Weymann. “Diffusion Approximation for a Photon Gas Interacting with a Plasma via the Compton Effect”. In: *The Physics of Fluids* 8.11 (1965), pp. 2112–2114. DOI: [10.1063/1.1761165](https://doi.org/10.1063/1.1761165). eprint: <http://aip.scitation.org/doi/pdf/10.1063/1.1761165>. URL: <http://aip.scitation.org/doi/abs/10.1063/1.1761165>.
- [54] S. Seager, D. D. Sasselov, and D. Scott. “How Exactly Did the Universe Become Neutral?”. In: *ApJS* 128 (June 2000), pp. 407–430. DOI: [10.1086/313388](https://doi.org/10.1086/313388). arXiv: [astro-ph/9912182](https://arxiv.org/abs/astro-ph/9912182).
- [55] S. Seager, D. D. Sasselov, and D. Scott. “A New Calculation of the Recombination Epoch”. In: *ApJ* 523 (Sept. 1999), pp. L1–L5. DOI: [10.1086/312250](https://doi.org/10.1086/312250). arXiv: [astro-ph/9909275](https://arxiv.org/abs/astro-ph/9909275).
- [56] D. Scott and A. Moss. “Matter temperature during cosmological recombination”. In: *MNRAS* 397 (July 2009), pp. 445–446. DOI: [10.1111/j.1365-2966.2009.14939.x](https://doi.org/10.1111/j.1365-2966.2009.14939.x). arXiv: [0902.3438](https://arxiv.org/abs/0902.3438).
- [57] Y. Ali-Haïmoud and C. M. Hirata. “HyRec: A fast and highly accurate primordial hydrogen and helium recombination code”. In: *Phys. Rev. D* 83.4, 043513 (Feb. 2011), p. 043513. DOI: [10.1103/PhysRevD.83.043513](https://doi.org/10.1103/PhysRevD.83.043513). arXiv: [1011.3758](https://arxiv.org/abs/1011.3758).
- [58] X. Chen and M. Kamionkowski. “Particle decays during the cosmic dark ages”. In: *Phys. Rev. D* 70.4, 043502 (Aug. 2004), p. 043502. DOI: [10.1103/PhysRevD.70.043502](https://doi.org/10.1103/PhysRevD.70.043502). arXiv: [astro-ph/0310473](https://arxiv.org/abs/astro-ph/0310473).
- [59] N. Padmanabhan and D. P. Finkbeiner. “Detecting dark matter annihilation with CMB polarization: Signatures and experimental prospects”. In: *Phys. Rev. D* 72.2, 023508 (July 2005), p. 023508. DOI: [10.1103/PhysRevD.72.023508](https://doi.org/10.1103/PhysRevD.72.023508). arXiv: [astro-ph/0503486](https://arxiv.org/abs/astro-ph/0503486).
- [60] T. R. Slatyer, N. Padmanabhan, and D. P. Finkbeiner. “CMB constraints on WIMP annihilation: Energy absorption during the recombination epoch”. In: *Phys. Rev. D* 80.4, 043526 (Aug. 2009), p. 043526. DOI: [10.1103/PhysRevD.80.043526](https://doi.org/10.1103/PhysRevD.80.043526). arXiv: [0906.1197](https://arxiv.org/abs/0906.1197) [[astro-ph.CO](https://arxiv.org/abs/astro-ph/0906.1197)].
- [61] G. F. Chapline. “Cosmological effects of primordial black holes”. In: *Nature* 253 (Jan. 1975), p. 251. DOI: [10.1038/253251a0](https://doi.org/10.1038/253251a0).
- [62] B. J. Carr et al. “Constraints on primordial black holes from the Galactic gamma-ray background”. In: *Phys. Rev. D* 94.4, 044029 (Aug. 2016), p. 044029. DOI: [10.1103/PhysRevD.94.044029](https://doi.org/10.1103/PhysRevD.94.044029). arXiv: [1604.05349](https://arxiv.org/abs/1604.05349).

- [63] The Fermi-LAT Collaboration. “Search for Gamma-Ray Emission from Local Primordial Black Holes with the Fermi Large Area Telescope”. In: *ArXiv e-prints* (Jan. 2018). arXiv: [1802.00100](https://arxiv.org/abs/1802.00100) [[astro-ph.HE](#)].
- [64] H. Niikura et al. “Microlensing constraints on primordial black holes with the Subaru/HSC Andromeda observation”. In: *ArXiv e-prints* (Jan. 2017). arXiv: [1701.02151](https://arxiv.org/abs/1701.02151).
- [65] B. J. Carr et al. “New cosmological constraints on primordial black holes”. In: *Phys. Rev. D* 81.10, 104019 (May 2010), p. 104019. DOI: [10.1103/PhysRevD.81.104019](https://doi.org/10.1103/PhysRevD.81.104019). arXiv: [0912.5297](https://arxiv.org/abs/0912.5297) [[astro-ph.CO](#)].
- [66] J. Garca-Bellido and E. Ruiz Morales. “Primordial black holes from single field models of inflation”. In: *Physics of the Dark Universe* 18 (Dec. 2017), pp. 47–54. DOI: [10.1016/j.dark.2017.09.007](https://doi.org/10.1016/j.dark.2017.09.007). arXiv: [1702.03901](https://arxiv.org/abs/1702.03901).
- [67] J. M. Ezquiaga, J. Garca-Bellido, and E. Ruiz Morales. “Primordial black hole production in Critical Higgs Inflation”. In: *Physics Letters B* 776 (Jan. 2018), pp. 345–349. DOI: [10.1016/j.physletb.2017.11.039](https://doi.org/10.1016/j.physletb.2017.11.039). arXiv: [1705.04861](https://arxiv.org/abs/1705.04861).
- [68] T. Suyama et al. “Are black holes overproduced during preheating?” In: *Phys. Rev. D* 71.6, 063507 (Mar. 2005), p. 063507. DOI: [10.1103/PhysRevD.71.063507](https://doi.org/10.1103/PhysRevD.71.063507). arXiv: [hep-ph/0410247](https://arxiv.org/abs/hep-ph/0410247).
- [69] T. Suyama et al. “Black hole production in tachyonic preheating”. In: *J. Cosmology Astropart. Phys.* 4, 001 (Apr. 2006), p. 001. DOI: [10.1088/1475-7516/2006/04/001](https://doi.org/10.1088/1475-7516/2006/04/001). arXiv: [hep-ph/0601108](https://arxiv.org/abs/hep-ph/0601108).
- [70] K. Jedamzik and J. C. Niemeyer. “Primordial black hole formation during first-order phase transitions”. In: *Phys. Rev. D* 59.12, 124014 (June 1999), p. 124014. DOI: [10.1103/PhysRevD.59.124014](https://doi.org/10.1103/PhysRevD.59.124014). arXiv: [astro-ph/9901293](https://arxiv.org/abs/astro-ph/9901293).
- [71] K. Kohri, C.-M. Lin, and T. Matsuda. “Primordial black holes from the inflating curvaton”. In: *Phys. Rev. D* 87.10, 103527 (May 2013), p. 103527. DOI: [10.1103/PhysRevD.87.103527](https://doi.org/10.1103/PhysRevD.87.103527). arXiv: [1211.2371](https://arxiv.org/abs/1211.2371) [[hep-ph](#)].
- [72] E. V. Bugaev and P. A. Klimai. “Primordial Black Hole Constraints for Curvaton Models with Predicted Large Non-Gaussianity”. In: *International Journal of Modern Physics D* 22, 1350034 (June 2013), p. 1350034. DOI: [10.1142/S021827181350034X](https://doi.org/10.1142/S021827181350034X). arXiv: [1303.3146](https://arxiv.org/abs/1303.3146) [[astro-ph.CO](#)].
- [73] S. W. Hawking. “Black hole explosions?” In: *Nature* 248 (Mar. 1974), pp. 30–31. DOI: [10.1038/248030a0](https://doi.org/10.1038/248030a0).
- [74] S. W. Hawking. “Particle creation by black holes”. In: *Communications in Mathematical Physics* 43 (Aug. 1975), pp. 199–220. DOI: [10.1007/BF02345020](https://doi.org/10.1007/BF02345020).
- [75] J. H. MacGibbon and B. R. Webber. “Quark- and gluon-jet emission from primordial black holes: The instantaneous spectra”. In: *Phys. Rev. D* 41 (10 May 1990), pp. 3052–3079. DOI: [10.1103/PhysRevD.41.3052](https://doi.org/10.1103/PhysRevD.41.3052). URL: <http://link.aps.org/doi/10.1103/PhysRevD.41.3052>.
- [76] J. H. MacGibbon. “Quark- and gluon-jet emission from primordial black holes. II. The emission over the black-hole lifetime”. In: *Phys. Rev. D* 44 (July 1991), pp. 376–392. DOI: [10.1103/PhysRevD.44.376](https://doi.org/10.1103/PhysRevD.44.376).
- [77] J. Skilling. “Nested Sampling”. In: *AIP Conference Proceedings* 735.1 (2004), pp. 395–405. DOI: [10.1063/1.1835238](https://doi.org/10.1063/1.1835238). eprint: <http://aip.scitation.org/doi/pdf/10.1063/1.1835238>. URL: <http://aip.scitation.org/doi/abs/10.1063/1.1835238>.
- [78] F. Feroz and M. P. Hobson. “Multimodal nested sampling: an efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analyses”. In: *MNRAS* 384 (Feb. 2008), pp. 449–463. DOI: [10.1111/j.1365-2966.2007.12353.x](https://doi.org/10.1111/j.1365-2966.2007.12353.x). arXiv: [0704.3704](https://arxiv.org/abs/0704.3704).

- [79] F. Feroz, M. P. Hobson, and M. Bridges. “MULTINEST: an efficient and robust Bayesian inference tool for cosmology and particle physics”. In: *MNRAS* 398 (Oct. 2009), pp. 1601–1614. DOI: [10.1111/j.1365-2966.2009.14548.x](https://doi.org/10.1111/j.1365-2966.2009.14548.x). arXiv: [0809.3437](https://arxiv.org/abs/0809.3437).
- [80] M. K. Cowles and B. P. Carlin. “Markov Chain Monte Carlo Convergence Diagnostics: A Comparative Review”. In: *Journal of the American Statistical Association* 91.434 (1996), pp. 883–904. DOI: [10.1080/01621459.1996.10476956](https://doi.org/10.1080/01621459.1996.10476956). URL: <https://www.tandfonline.com/doi/abs/10.1080/01621459.1996.10476956>.
- [81] H. K. Eriksen et al. “Cosmic Microwave Background Component Separation by Parameter Estimation”. In: *ApJ* 641 (Apr. 2006), pp. 665–682. DOI: [10.1086/500499](https://doi.org/10.1086/500499). arXiv: [astro-ph/0508268](https://arxiv.org/abs/astro-ph/0508268).
- [82] H. K. Eriksen et al. “Joint Bayesian Component Separation and CMB Power Spectrum Estimation”. In: *ApJ* 676, 10-32 (Mar. 2008), pp. 10–32. DOI: [10.1086/525277](https://doi.org/10.1086/525277). arXiv: [0709.1058](https://arxiv.org/abs/0709.1058).
- [83] Planck Collaboration et al. “Planck 2015 results. X. Diffuse component separation: Foreground maps”. In: *A&A* 594, A10 (Sept. 2016), A10. DOI: [10.1051/0004-6361/201525967](https://doi.org/10.1051/0004-6361/201525967). arXiv: [1502.01588](https://arxiv.org/abs/1502.01588).
- [84] Planck Collaboration et al. “Planck 2015 results. XI. CMB power spectra, likelihoods, and robustness of parameters”. In: *A&A* 594, A11 (Sept. 2016), A11. DOI: [10.1051/0004-6361/201526926](https://doi.org/10.1051/0004-6361/201526926). arXiv: [1507.02704](https://arxiv.org/abs/1507.02704).
- [85] Planck Collaboration et al. “Planck 2013 results. XV. CMB power spectra and likelihood”. In: *A&A* 571, A15 (Nov. 2014), A15. DOI: [10.1051/0004-6361/201321573](https://doi.org/10.1051/0004-6361/201321573). arXiv: [1303.5075](https://arxiv.org/abs/1303.5075).
- [86] S. Hamimeche and A. Lewis. “Likelihood analysis of CMB temperature and polarization power spectra”. In: *Phys. Rev. D* 77.10, 103013 (May 2008), p. 103013. DOI: [10.1103/PhysRevD.77.103013](https://doi.org/10.1103/PhysRevD.77.103013). arXiv: [0801.0554](https://arxiv.org/abs/0801.0554).
- [87] D. Sivia and J. Skilling. *Data Analysis: A Bayesian Tutorial*. OUP Oxford, 2006. ISBN: 9780191546709. URL: <https://books.google.com.au/books?id=Kxx8CwAAQBAJ>.
- [88] Planck Collaboration. *Planck 2015 results: Cosmological Parameter Tables*. In table base\_plikHM\_TTTEEE\_lowTEB. Jan. 2015. URL: [https://wiki.cosmos.esa.int/planckpla2015/images/f/f7/Baseline\\_params\\_table\\_2015\\_limit68.pdf](https://wiki.cosmos.esa.int/planckpla2015/images/f/f7/Baseline_params_table_2015_limit68.pdf).
- [89] H. Liu, T. R. Slatyer, and J. Zavala. “Contributions to cosmic reionization from dark matter annihilation and decay”. In: *Phys. Rev. D* 94.6, 063507 (Sept. 2016), p. 063507. DOI: [10.1103/PhysRevD.94.063507](https://doi.org/10.1103/PhysRevD.94.063507). arXiv: [1604.02457](https://arxiv.org/abs/1604.02457).
- [90] D. Scott. *Dartmouth-TRIUMF-University of Washington High-Energy Physics/Cosmology Tools Bootcamp: RECFAST*. Webinar. Oct. 2017.
- [91] Planck Collaboration et al. “Planck 2015 results. XIII. Cosmological parameters”. In: *A&A* 594, A13 (Sept. 2016), A13. DOI: [10.1051/0004-6361/201525830](https://doi.org/10.1051/0004-6361/201525830). arXiv: [1502.01589](https://arxiv.org/abs/1502.01589).