THE UNIVERSITY *of* EDINBURGH

The Impact of Baryonic Astrophysics and Intrinsic Alignments on Dark Energy and Neutrino Constraints in Cosmology

David Copeland

Doctor of Philosophy
University of Edinburgh
December 2019

Abstract

In the early 2020s, cosmology will enter an era of unprecedented precision when the next generation of large scale structure surveys begin receiving data. As a result, it is expected that stronger constraints on major features of cosmology like dark energy and massive neutrinos are forthcoming. The former could help to favour an explanation for the present accelerated expansion of the Universe, while the latter has the potential to enhance our understanding of a prominent intersection between particle physics and cosmology. However, there are many systematics that must be accounted for in parameter forecasts. One of the most prominent theoretical cases is the influence of baryonic astrophysics on large scale structure (e.g., AGN and supernova feedback, adiabatic contraction etc.) and the effect that marginalising over a limited theoretical understanding of the associated phenomena will have on forecasts. This question is one of the core concerns of this thesis. It will be applied respectively in Chapters 3, 4 and 5 to constraints on dynamical dark energy, the neutrino mass sum and a possible coupling between dark energy and dark matter. While forecasts are the primary focus here, much of this work has implications for parameter inference more broadly, and could be used to inform the direction that model building or simulation development should take in pursuit of the goal of more accurate parameter constraints.

The key statistics used here to probe the growth of structure derive from the power spectrum of matter overdensities. This permits the use of both weak gravitational lensing of light from background sources by foreground objects, and galaxy clustering. The former requires accounting for the intrinsic alignments of ellipticities and shears, a systematic examined in depth for its impact on forecasts. This thesis presents Fisher analyses using weak lensing and galaxy clustering probes for parameter forecasts for a Euclid-like survey.

The approach here to modelling the baryonic phenomena is to adopt a generic treatment of their global redistribution of the dark matter content in haloes, via energy transfer to their surroundings. Different baryonic effects are separated into three general but distinct categories: large scale adiabatic contraction caused by radiative cooling; high impact energy transfer from specific, localised sources; and small-scale effects that manifest as inner halo cores. I introduce the inner halo core through analytic modelling. A central tenet of this work is the use of analytic modelling, rather than numerical simulations, in capturing the relevant

physics so as to circumvent computational expenses and underlying systematics associated with the latter, while retaining the useful physical insight offered by the former.

Employing a maximum likelihood method, matter and weak lensing power spectra are varied around a fiducial cosmology given by Planck Collaboration et al. (2016b). For a Euclid-like survey covering 15000 sq. deg. of sky, measuring 10 redshift bins in the range $0 < z < 2$, the w_0-w_a dark energy Figure of Merit is shown to experience a 40% degradation due to the combination of baryon effects. This thesis presents a detailed analysis of the relative dark energy and baryon sensitivities over the range of available lensing modes. Ultimately, it is found that the application of cosmic microwave background (CMB) priors alleviate the baryon impact for individual errors on the dark energy parameters but the relative degradation to the Figure of Merit for the parameter space remains.

A similar approach is used to address the question of whether Stage IV surveys, whilst accounting for baryons and intrinsic alignments, can make a positive detection of the neutrino mass and whether they can distinguish between the normal and inverted hierarchy of mass eigenstates. Combined forecasts from weak lensing, the CMB and galaxy clustering preclude a meaningful distinction of hierarchies but do achieve a positive detection of the mass sum, overcoming the significant degradation by factors of $\sim$2 to that arise when marginalising over baryons for weak lensing alone. These results could be improved upon with future CMB priors on the spectral index and information from neutrinoless double beta decay to achieve a 2-σ distinction of the hierarchies. The effect of intrinsic alignments on forecasts is shown to be minimal, with constraints even experiencing mild improvements due to information from the intrinsic alignment signal.

Finally, this work explores the prospects for using large scale structure to constrain the strength of a possible coupling between dark matter and a dark energy scalar field. While the growth of structure in the linear regime in this model has been well-explored, the non-linear regime is more challenging. Forecasts have been made using polynomial fits to power spectra directly. However, this thesis presents a more physically motivated approach to the problem. I show that a single parameter, the halo virial density, is responsible for describing most of the impact of the coupling on small scales. By computing the changes to the virial density in this model, a fit can be found that allows for a full physically-motivated halo model. This allows Fisher forecasts to be made for the coupling strength, including an assessment of the impact of baryons. Degradations to the coupling strength constraint of $\sim$20% due to this systematic are found. While CMB and galaxy clustering priors notably improve the absolute errors without marginalising over baryons, when this systematic is accounted for these priors provide little improvement and the relative degradation increases.

Taken as a whole, this thesis provides a comprehensive analysis of the impact of baryons and intrinsic alignments on constraints for a wide range of cosmological phenomena responsible for the accelerated expansion, the neutrino mass hierarchy and beyond ΛCDM physics coupling dark matter and dark energy. Through modified and improved approaches to halo modelling, this work demonstrates which of these phenomena are subject to the most severe degeneracies with baryonic effects. This rigorous analysis, grounded in empirically motivated parameterisations, is designed to inform optimal mitigation strategies to minimise the impact on forecasts. In turn, this provides a wide scope for making future improvements to modelling and simulations that can advance efforts to constrain cosmology still further.

Lay Summary

Cosmology is rather unique among the sciences in that the observable Universe represents the results of a single 'experiment' to which we have access. The challenge of recent decades has been to improve our observational capabilities in order to tackle the underlying questions of the field with a level of precision that will allow us to draw meaningful scientific conclusions. These questions are among the boldest that are posed in science - what is the origin of the Universe? What are its constituents and how does it look like on the largest scales? What is its history and what is its ultimate future?

As the era of precision cosmology has come more sharply into focus, a broadly consistent model of the Universe has developed but there are a considerable number of problems that remain unsolved. One of the most striking is that observations indicate that only $\sim 4\%$ of the matter-energy content of the Universe is provided by the everyday 'baryonic' matter we are familiar with. The orbits of galaxies and stars indicate that $\sim 26\%$ of the cosmological budget is taken up by 'dark matter', which leaves an observable gravitational signature but does not interact with light. There are many proposed candidates for what dark matter may be - a population of familiar bodies like brown dwarfs with luminosities too low to observe; a manifestation of an incomplete understanding of gravity itself; or some new type of particle.

Our understanding of the large scale structure of the Universe emerges from modelling the gravitational dynamics of dark matter, giving rise to a complex framework of filamentary structures punctuated by collapsed 'halo' objects. Large scale structure bears the imprints of gravitational dynamics and the history of the expanding Universe in which it forms. It therefore provides an excellent probe into many of the questions cosmology is currently focused on. One of the main issues is explaining what the remaining 70% of the matter-energy content of the Universe is.

Observations of supernovae in 1998 famously showed that today the Universe is accelerating in its expansion, that galaxies recede from us at greater velocities the farther they are from the observer. Explaining this remarkable result is a significant challenge. Quantum field theory indicates that the vacuum of space itself should generate an accelerated expansion but one that is far greater (by 120

orders of magnitude) than observations suggest. In lieu of a solution, it is often assumed that some mechanism cancels the gravitation of the vacuum and that the remaining 70% is associated with a 'dark energy' component of the Universe. This could range from a new particle to an additional feature of gravitation beyond General Relativity.

Forthcoming large scale structure surveys aim to constrain the dark energy model space, and other interesting physics like the configuration of neutrino mass states, through precision measurements. In order for this to be successful, it is important that systematics are properly accounted for. A major example, with which this thesis is concerned, is the influence of baryonic astrophysics on large scale structure. Star formation, supernovae, and the activity generated by matter accreting around black holes all contribute to significantly distort the matter distribution. This in turn biases our statistical probes.

My work in this thesis explores this issue in detail, adopting and expanding models of structure with prescriptions for the effects of baryons. This allows me to quantify the impact of baryons on cosmological forecasts, and to identify potential routes for mitigating this issue. Chapters 3 and 4 performs this analysis for dark energy and neutrinos respectively. Chapter 5 addresses halo structure in the context of a coupling between dark matter and dark energy, and assesses the baryon impact on forecasts of the coupling strength.

Declaration

I declare that this thesis is my own work, subject only to helpful comments from my PhD supervisors, Andy Taylor and Alex Hall, and corrections made at the request of my viva examiners, Alan Heavens and Joe Zuntz. Where works by other authors are discussed these are appropriately referenced. I affirm that I have submitted no part of this thesis for any other degree or qualification at this or any other institution.

Chapter 3 is based on the 2018 paper "The impact of baryons on the sensitivity of dark energy measurements", D. Copeland, A. Taylor, A. Hall. This was published in the Monthly Notices of the Royal Astronomical Society (MNRAS), vol. 480, pp. 2247. It is also available on ArXiv (arxiv:1712.07112).

Chapter 4 is based on the 2019 paper "Towards determining the neutrino mass hierarchy: weak lensing and galaxy clustering forecasts with baryons and intrinsic alignments", D. Copeland, A. Taylor, A. Hall. This is available on ArXiv (arxiv:1905.08754) and has also been submitted to MNRAS.

(David Copeland, 05/12/2019)

Acknowledgements

This thesis has only been made possible by the dedicated efforts of my supervisors, Andy Taylor and Alex Hall, in guiding my work over the years. I would especially like to acknowledge your patience and encouragement during times of particular challenge when it seemed as if obstacles I was facing might prove insurmountable. The discussions in our weekly meetings and the insights provided by you both have helped me to improve my own abilities significantly, and helped this work to reach a much higher standard as a result.

I would also like to thank Alex Mead for numerous enlightening conversations throughout my PhD, and to acknowledge his role in providing the foundations (through his halo model code, HMCODE) which underpin much of this work. John Peacock, Jorge Peñarrubia and Sadegh Khochfar have also helped to shape the direction I took in a number of valuable discussions, for which I am grateful. I greatly appreciate the interesting and engaging conversations with Alan Heavens and Joe Zuntz during my viva, which I hope have added an additional robustness to this work. I would like to extend a wider thank you to all those at the Observatory and further afield whose friendship and insightful conversation have enriched this work and my experience of studying at Edinburgh.

Special thanks are also owed to the many friends outside of Astronomy during these years. Though too many to name, the Coldean Crew as a collective fuelled my interest in physics and were a source of grand camaraderie throughout my undergraduate days and beyond. My St. Leonard's friends are all owed a debt of loving gratitude for arriving in my life by chance during a critical year, for keeping me both sane and wildly entertained throughout it, and for being a well of true support. You are all superlative people and I will miss our Buffy nights. Kyle Lambe has earned high praise from me before and receives it here again for being my first friend in Edinburgh, a man of inestimable wit, and a longtime artistic collaborator when the day was long and it was time to unwind. They should pay him a hero's wage. Rachel Dowse (defender of the wild and warden of the woods) knows all the words I have to share are hers already, but my thanks goes out to her for her unfailing love and support during this journey, and for always inspiring the best from me.

Without the support of my family, my work would never have reached this stage.

Their daily motivation throughout this process has been invaluable. The loving upbringing of my parents, the lifelong friendship with my brother, Matthew, and the caring interest from my grandparents, aunt and uncle have been foundational to me. My grandparents are sadly not here to be able to see this work, my Grandad having passed away last year. I know he would have been excited to have a read of this, and so I hope that this work serves as some small tribute to him.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Cosmology is a very broad scientific field. It would be impractical to discuss all or even most of its subject matter here, so this Chapter will be confined to an explanation of the principal tenets of the standard model of cosmology, with a focus on those elements necessary to understand and interpret the methodology and conclusions of later Chapters.

The interpretation of gravity on large scales provided by General Relativity underpins a modern understanding of the cosmos, and so will be discussed first. This will lead into an examination of the various components of material permeating the Universe, and their role in defining different eras of cosmological evolution and expansion. Particular attention will be devoted to dark matter, as this thesis is grounded in the matter-dominated era. Dark energy and cosmological neutrinos will be discussed in detail, as later Chapters are devoted to efforts to constrain these features of cosmology.

1.1 General Relativity and Cosmology

An understanding of Einstein gravity and its role in cosmology begins with the acknowledgement that Newtonian gravity does not satisfy the special principle of relativity that physical laws should be invariant for all inertial frames (Einstein, 1905a). Such a frame is defined by a set of time and space coordinates, (t, x, y, z). An event observed in one frame will have a different set of coordinates to the same event observed in another, with a Lorentz transformation governing the relation between them. Under a Lorentz transformation of the Newtonian gravitational potential in the Poisson equation,

$$\nabla^2\Phi = 4\pi G\rho, \tag{1.1}$$

where G is the Universal gravitational constant and ρ is a source matter density, time derivatives will generally appear alongside spatial derivatives. There is a clear incompatibility here with the principle of relativity, although Newtonian gravity remains an accurate approximation provided that the speed of an object is much less than the speed of light, $v \ll c$, and that gravitational fields are correspondingly weak.

The expansion of Special Relativity to include gravity and non-inertial frames requires three equivalence principles:

1. **Weak equivalence principle:** All freely falling test particles of negligible active gravitational mass and with identical initial positions and velocities follow the same trajectory.

2. **Einstein equivalence principle:** In addition to the weak equivalence principle, Special Relativity is valid in the local region of a freely falling test particle of negligible active mass. Here the concept of 'locality' requires that variations of gravitational fields are vanishingly small.

3. **Strong equivalence principle:** For a freely falling massive body, a locally inertial frame always exists in which physical laws are identical in form to those in a frame with no gravitational field.

The (strong) equivalence principle is the basis for General Relativity (Einstein, 1905b), a theory that interprets the gravitational force on cosmological scales as a manifestation of the curvature of a four-dimensional 'spacetime' manifold. A review of the fundamental formalism of the theory is presented in Appendix A. The application of General Relativity to constructing the standard model of cosmology follows.

1.1.1 General Relativistic Cosmology and The FLRW Metric

A solution to Einstein's field equations (see equation A.9) relating the geometry of spacetime to the matter-energy content of the Universe is required that can describe the physical Universe. This in turn raises the question of what should be considered the foundational features of the cosmos that a descriptive model must accommodate. Here, the Cosmological Principle may be introduced. It states that on large scales, the content of the Universe is statistically homogeneous and isotropic (Milne, 1936). The homogeneity assumption requires that the statistical properties of the Universe are invariant under translations in space on sufficiently large scales. In other words, if one selects a number of large independent volumes they should have e.g., consistent mean matter densities, and likewise for other properties (Hogg et al., 2005). Isotropy requires that observations made in any direction are indistinguishable. This represents a culmination of a historical trend away from conceptions of the centrality of Earth and towards Copernican views that human observers occupy no privileged position in the Universe[1].

On sufficiently small scales, $\leq$ 100Mpc, there is neither homogeneity nor isotropy, with the Universe revealing a complex, irregular structure of connected dark matter filaments and haloes. On larger scales, the evidence for the Cosmological Principle is strong but not absolute. The angular distribution of temperature and density perturbations in cosmic microwave background (CMB) radiation (see § 1.4.1) provides the most incisive test, with the most recent accurate observations from *Planck* (Planck Collaboration et al., 2016a) finding that although the CMB is extremely isotropic there exist anisotropies at the level of $O\left(10^{-5}\right)$ which

[1] A further extension to the Perfect Cosmological Principle asserts homogeneity and isotropy in time as well across space, although this is not consistent with observations of the Universe.

are large enough to be statistically significant.

The Universe is assumed to be homogeneous on large scales. This is one of the central tenets of the standard cosmological model. The assumption has been explored by a number of studies examining the scale at which observations are homogeneous to within a certain statistical threshold (e.g., Hogg et al., 2005), although these implicitly assume that the geometry associated with standard cosmology is valid. Many works have used observations of galaxy distributions and how uniform their associated length scales are to probe for homogeneity indirectly (by absence of significant inhomogeneities). As Maartens (2011) notes one cannot directly test homogeneity. Observers on Earth have access to a lightcone to observe the Universe through but cannot study intersecting spatial surfaces. This places a limit on the capacity to probe the validity of the homogeneity assumption. Galaxy surveys can only provide information on the galaxy distribution for slices of the sky at constant redshifts.

The strongest basis to make observational claims about homogeneity is to do so through isotropy, for which the CMB is an effective probe. This requires that the Copernican Principle that observers do not occupy any special place in the Universe. Homogeneity then emerges as a consequence if all observers measure the CMB to be isotropic. For a detailed demonstration of this result, the reader is referred to Maartens (2011). However, as the CMB is observed to feature small anisotropies it does not immediately follow that there is near-homogeneity. This claim rests on a number of assumptions about the derivatives of CMB multipoles (Maartens, 2011).

The modern formulation of the Cosmological Principle is a consequence of Einstein's relativistic interpretation of cosmology. A Newtonian approach to an isotropic Universe requires a finite spatial extent, which leads to special locations for observation (Barrow, 1989). By contrast, the Einstein equations can be solved exactly by invoking the assumptions of isotropy and homogeneity as outlined above. Friedmann (1922) showed that a metric (known as the Friedmann-Lemaître-Robertson-Walker (FLRW) metric) satisfying these conditions takes the form

$$\mathrm{d}s^2 = -c^2\mathrm{d}\tau^2 = -c^2\mathrm{d}t^2 + R^2\left(t\right)\left[\mathrm{d}r^2 + f_K^2\left(r\right)\left(\mathrm{d}\theta^2 + \sin^2\theta\mathrm{d}\phi^2\right)\right], \qquad (1.2)$$

where $R(t)$ is the scale factor of cosmic expansion with dimensions of length, allowing r to be a dimensionless radial coordinate corresponding to the 'comoving distance' between a source and observer which is constant with respect to the expansion. The 'physical distance' is then given by $R(t)\,r$. It should be noted that in order to satisfy the Cosmological Principle, the scale factor depends only the 'cosmic time', t, rather than the most general possible form, $R(t, r, \theta, \phi)$, as a function of position as well. The cosmic time is defined to be the time experienced by a physical clock travelling with the homogeneous expansion, which will be measured by all observers with zero peculiar velocity. It can therefore be seen that the imposition of the Cosmological Principle reduces the contribution of Einstein's equations to determining the behaviour of $R(t)$, through the time-time component, g_{00}, of the metric tensor[2]. This is of course much more straightforward than contending with a fully general solution of the equations, and lays the basis for the Standard Model of cosmology.

For a polar decomposition of the spatial coordinates, the angular components are given by θ and ϕ in common with convention. If a coordinate transformation is performed such that $\sqrt{g_{11}}\mathrm{d}r \longrightarrow \mathrm{d}r$, then the remaining spatial curvature is described in hyperspherical coordinates by

$$f_K(r) = \begin{cases} \sinh(r), & k = -1 \\ r, & k = 0 \\ \sin(r), & k = +1 \end{cases} \tag{1.3}$$

with $k = \{-1, 0, +1\}$ being a 'curvature' constant corresponding to spatial hypersurfaces that are open, flat or closed respectively. These can be conceptualised by considering the appropriate 2D geometric surface for a hyperbola, a plane or a sphere. Note that, in order to preserve isotropy, $f_K(r)$ is a radial function, expressing no angular dependence in θ or ϕ.

It is useful to define a dimensionless form of the scale factor,

$$a(t) \equiv \frac{R(t)}{R_0}, \tag{1.4}$$

[2]The g_{0i} and g_{i0} components controlling any mixing of time and spatial coordinates are set to zero in order to satisfy the time invariant symmetry, $t \longleftrightarrow -t$, that ensures physical laws yield consistent results when equations are solved forwards or backwards in time. The FLRW metric contains only $\mathrm{d}t^2$ and $\left(\mathrm{d}r^2, \mathrm{d}\theta^2, \mathrm{d}\phi^2\right)$ terms as a result.

normalised to unity at the present day. This is the form that will be referred to as 'the scale factor' hereafter, with the time-dependence of the quantity assumed rather than denoted explicitly. Rewriting the metric in terms of a and rescaling $r \longrightarrow R_0\, r$ (so that the comoving distance is now defined as the physical distance today between an observer and source) gives

$$\mathrm{d}s^2 = -c^2\mathrm{d}t^2 + a^2\left(t\right)\left[\mathrm{d}r^2 + f_K^2\left(r\right)\left(\mathrm{d}\theta^2 + \sin^2\theta\mathrm{d}\phi^2\right)\right], \tag{1.5}$$

with the curvature described by

$$f_K\left(r\right) = \begin{cases} R_0 \sinh\left(\frac{r}{R_0}\right), & k = -1 \\ r, & k = 0 \\ R_0 \sin\left(\frac{r}{R_0}\right), & k = +1. \end{cases} \tag{1.6}$$

1.2 Cosmic Expansion

1.2.1 Hubble's Law and Redshift

The notion of cosmic expansion is best understood as the isotropic recession of galaxies from an observer at any and all positions. The relationship between the velocity of a galaxy and its physical separation from an observer measuring this velocity is given by Hubble's law as

$$\mathbf{v} = H\left(t\right)\,\mathbf{r}, \tag{1.7}$$

where $H\left(t\right)$ is the Hubble parameter. Its present value, H_0, is the constant of proportionality relating the speed of recession to the distance of a galaxy today. The most recent measurements from the *Planck* survey find $H_0 = (67.4 \pm 0.5)\ \mathrm{kms}^{-1}\mathrm{Mpc}^{-1}$ (Planck Collaboration et al., 2018). It is generally repackaged with a dimensionless parameter, h, as

$$H_0 = 100h\,\mathrm{kms}^{-1}\mathrm{Mpc}^{-1}. \tag{1.8}$$

The observational evidence for Hubble's law begins with measurements by Slipher (1917) of spiral nebulae[3]. The chemical composition of these objects results in

[3]In this context, 'nebula' is an outdated term for a spiral galaxy

a discrete set of wavelengths required to excite electrons within atoms of the different elements present. Absorption lines corresponding to these wavelengths appear in a measured spectrum of light from these sources. The key observation was that the spectral lines were shifted for observed spectra with respect to their position from spectra produced under laboratory conditions. When interpreted as Doppler shifts these observations could be related to the velocity of the spirals, which appeared to increase with distance. As has been discussed, Friedmann (1922) showed that there were theoretical grounds for an expanding Universe within the framework of Einstein's relativity. Lemaître (1927) and Hubble (1929) independently established the proportional relationship between velocity and distance from observations of receding objects that became Hubble's law, with Hubble (1929) providing an early value for the Hubble constant of $H_0 \approx 500\,\mathrm{kms}^{-1}\mathrm{Mpc}^{-1}$. The form of the Hubble parameter as a function of time but not the spatial coordinates (r, θ, ϕ) is a requirement of the Cosmological Principle. It can be seen that by expressing the physical separation in terms of the comoving separation (now denoted by $\mathbf{x}$) and the scale factor,

$$\mathbf{r} = a\,(t)\,\mathbf{x}, \tag{1.9}$$

according to the FLRW conventions previously discussed, that the Hubble parameter can be written as

$$H\,(t) \equiv \frac{\dot{a}}{a}, \tag{1.10}$$

where a dot indicates a derivative with respect to cosmic time. This is a simple consequence of writing the velocity in Hubble's law as

$$\mathbf{v} = \dot{\mathbf{r}} = \dot{a}\mathbf{x} = \frac{\dot{a}}{a}\mathbf{r}, \tag{1.11}$$

in which it is assumed that $\dot{\mathbf{x}} = 0$.

The Doppler shift experienced by light observed from a receding galaxy means that the initial wavelength, λ_i, of the emission is changed by a factor

$$z \equiv \frac{\lambda_0 - \lambda_i}{\lambda_i} \approx \frac{v}{c} \tag{1.12}$$

so that it reaches the observer with wavelength, λ_0. This factor is called the 'redshift' because light emitted by receding galaxies is observed with a longer

wavelength, while light received by a body with net incoming motion would be 'blueshifted' to shorter wavelengths.

A photon travels along a null geodesic with $\mathrm{d}s^2 = 0$ in the (FLRW) metric. For photons emitted from a source with zero peculiar velocity, a coordinate system can be defined such that the photon travel is entirely along the radial direction with the source located at the origin. This gives the relationship

$$\mathrm{d}t = \frac{a\left(t\right)}{c}\mathrm{d}r \tag{1.13}$$

between time and (radial) comoving distance intervals. A photon emitted at time t_{em} and observed at time t_{obs} at radial coordinate r_{obs} travels a comoving distance

$$r_1 = \int_0^{r_{\mathrm{obs}}} \mathrm{d}r = \int_{t_{\mathrm{em}}}^{t_{\mathrm{obs}}} \mathrm{d}t \frac{c}{a\left(t\right)} \tag{1.14}$$

between the events of emission and observation. A second photon, emitted a short time interval after the first, travels

$$r_2 = \int_0^{r_{\mathrm{obs}}} \mathrm{d}r = \int_{t_{\mathrm{em}}+\Delta\, t_{\mathrm{em}}}^{t_{\mathrm{obs}}+\Delta\, t_{\mathrm{obs}}} \mathrm{d}t \frac{c}{a\left(t\right)}. \tag{1.15}$$

It can be seen that $r_1 = r_2$. This is because the comoving distance between the source and observer is time-independent. All the time dependence is associated with the cosmic expansion and captured by $a\left(t\right)$. By equating the two expressions above and rewriting the integrals via

$$\int_{t_{\mathrm{em}}}^{t_{\mathrm{obs}}} \frac{\mathrm{d}t}{a\left(t\right)} = \int_{t_{\mathrm{em}}+\Delta\, t_{\mathrm{em}}}^{t_{\mathrm{obs}}+\Delta\, t_{\mathrm{obs}}} \frac{\mathrm{d}t}{a\left(t\right)} = \int_{t_{\mathrm{em}}}^{t_{\mathrm{obs}}} \frac{\mathrm{d}t}{a\left(t\right)} + \int_{t_{\mathrm{obs}}}^{t_{\mathrm{obs}}+\Delta\, t_{\mathrm{obs}}} \frac{\mathrm{d}t}{a\left(t\right)} - \int_{t_{\mathrm{em}}}^{t_{\mathrm{em}}+\Delta\, t_{\mathrm{em}}} \frac{\mathrm{d}t}{a\left(t\right)}, \tag{1.16}$$

the interval between two photons at emission can be related to the interval at observation such that

$$\int_{t_{\mathrm{obs}}}^{t_{\mathrm{obs}}+\Delta\, t_{\mathrm{obs}}} \frac{\mathrm{d}t}{a\left(t\right)} = \int_{t_{\mathrm{em}}}^{t_{\mathrm{em}}+\Delta\, t_{\mathrm{em}}} \frac{\mathrm{d}t}{a\left(t\right)}. \tag{1.17}$$

If these intervals are sufficiently small, this can be approximated as

$$\frac{\Delta\, t_{\mathrm{obs}}}{a_{\mathrm{obs}}} = \frac{\Delta\, t_{\mathrm{em}}}{a_{\mathrm{em}}}. \tag{1.18}$$

The relation between redshift and scale factor can be deduced from this by setting $\Delta\, t_{\rm em}$ equal to the period of the emitted photon. The ratio $\Delta\, t_{\rm obs}/\Delta\, t_{\rm em}$ is therefore equal to the ratio of observed and emitted wavelengths. It then follows from equation (1.12) that

$$1 + z = \frac{a_{\rm obs}}{a_{\rm em}}. \tag{1.19}$$

If the scale factor is normalised to unity today then the relations

$$z = \frac{1}{a} - 1 \longleftrightarrow a = \frac{1}{1+z}, \tag{1.20}$$

can be used with regard to present day observations. If the source has a peculiar velocity $v_{\rm pec}$ with respect to the expansion, an additional redshift, $z_{\rm pec}$, is induced. For a time-independent $v_{\rm pec}$, equation (1.15) can be re-expressed as

$$\int_0^{r_{\rm obs}} {\rm d}r + \int_{t_{\rm em}}^{t_{\rm em}+\Delta\, t_{\rm em}} {\rm d}t \frac{v_{\rm pec}}{a\,(t)} = \int_{t_{\rm em}+\Delta\, t_{\rm em}}^{t_{\rm obs}+\Delta\, t_{\rm obs}} {\rm d}t \frac{c}{a\,(t)}, \tag{1.21}$$

which leads, via the procedure outlined above, to the observed redshift,

$$1 + z_{\rm obs} = (1+z)\,(1 + v_{\rm pec}/c) = (1+z)\,(1+z_{\rm pec}) \tag{1.22}$$

(in the non-relativistic limit) as a combination of the cosmic redshift and the fedshift due to the peculiar velocity.

In equation (1.12) the approximate relationship between the recession velocity of a source and redshift holds well but on sufficiently large scales general relativistic effects cause Hubble's law to break down. Large-scale spacetime curvature can be approximated as by a flat Euclidean plane locally so galaxies within this environment obey $z \approx v/c$. On cosmological scales, however, it is more appropriate to determine redshifts using the scale factor as this encodes contributions from General Relativity directly. This approach will generally be used for objects at redshifts $z \gtrsim 1$.

1.2.2 The Friedmann Equation and Parameterising Density

Having established the FLRW metric that allows one to solve the Einstein field equations in adherence to the Cosmological Principle and the conventions of the scale factor and Hubble parameter, the Friedmann equations describing the evolution of the cosmic expansion can now be presented. The form of the metric in equation (1.5) is used. The 00 component of $G_{\mu\nu}$ (equation (A.9)) encapsulating the global geometry of spacetime can then be ultimately derived from the metric components required to construct the Christoffel symbols (equation (A.7)), Riemann tensor (equation (A.8)) and Ricci tensor. This can be equated to $T_{00} = \rho$, matter-energy source term of the field equations, to give (Friedmann, 1922)

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{kc^2}{a^2}. \tag{1.23}$$

This is the first Friedmann equation. It should be noted that the freedom to add a cosmological constant term is accounted for by incorporating its contribution to the total energy density alongside matter and radiation into ρ such that

$$\rho = \rho_m + \rho_r + \rho_\Lambda = \rho_m + \rho_r + \frac{\Lambda c^2}{8\pi G}. \tag{1.24}$$

The role played by the cosmological constant will be discussed in more detail in forthcoming sections but for now it will suffice to think of it as contributing to a vacuum energy density.

A version of this Friedmann equation can also be derived from Newtonian arguments by considering a homogeneous sphere of matter undergoing isotropic expansion and a test particle at the boundary experiencing the Newtonian gravitational force from the enclosed matter. Using the same definition of the scale factor relating comoving and proper distances, and then applying energy conservation arguments leads straightforwardly to the same dynamical equation presented above. The key difference is that the curvature term has an analogue in the form of a total energy term, the sign of which dictates whether the Universe is open, closed or flat in the same manner as the sign of k for the relativistic form. The robustness of the Newtonian approach is due to the fact that the shell theorem - which states that the test mass experiences no net force from matter

shells beyond the radius of the sphere (and that the enclosed matter has an indistinguishable effect from a point source of the same mass on the test particle) - has a relativistic analogue in the Jebsen-Birkhoff theorem (Jebsen, 2005; Birkhoff and Langer, 1923)[4] that states that any spherically symmetric solution of the field equations in a vacuum must be asymptotically flat. In other words, in this scenario, gravity reduces to the Newtonian limit at large distances where one would otherwise expect the curvature of spacetime to become relevant.

A flat Universe is associated with a particular time-varying value of the density called the 'critical density', given by the Friedmann equation when $k = 0$ as

$$\rho_{\text{crit}}(t) = \frac{3H^2(t)}{8\pi G}. \tag{1.25}$$

It is useful to define a dimensionless parameter,

$$\Omega(t) \equiv \frac{\rho(t)}{\rho_{\text{crit}}(t)} = \frac{8\pi G\rho(t)}{3H^2(t)}, \tag{1.26}$$

simply called the 'density parameter'. The value of Ω implicitly encodes the curvature. Corresponding to densities that require $k = +1$ ($k = -1$) to satisfy the value of H in the Friedmann equation, $\Omega < 1$ ($\Omega > 1$) describes a closed (open) Universe. Following the division of matter-energy components in equation (1.24), independent density parameters for each component can be written as

$$\Omega_m(a) = \frac{8\pi G\rho_m(a)}{3H^2(a)}, \quad \Omega_r(a) = \frac{8\pi G\rho_r(a)}{3H^2(a)}, \quad \Omega_\Lambda(a) = \frac{\Lambda\rho_\Lambda(a)c^2}{3H^2(a)}, \tag{1.27}$$

where the scale factor has replaced time in evolving functions as a (or, equivalently, z) is generally a more meaningful parameter to use. For a flat Universe, the sum of these Ω is unity. For an arbitrary Universe, a density parameter associated with curvature can be defined as

$$\Omega_k = -\frac{kc^2}{(aH)^2} \tag{1.28}$$

so that

$$\Omega_{\text{tot}} = \Omega_m(a) + \Omega_r(a) + \Omega_\Lambda(a) + \Omega_k(a) = 1 \tag{1.29}$$

[4]Jebsen (2005) is a posthumous republication in English of Jebsen's original 1921 article demonstrating this result.

is true generally.

The second Friedmann equation is derived by collecting the remaining independent contributions to the Einstein field equations by computing their trace. Combining this result with the first Friedmann equation leads to the acceleration equation,

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}\left(\rho + \frac{3p}{c^2}\right). \tag{1.30}$$

Differentiating equation (1.23) with respect to time and substituting this acceleration result gives the continuity equation,

$$\dot{\rho} + 3\frac{\dot{a}}{a}\left[\rho + \frac{p}{c^2}\right] = 0. \tag{1.31}$$

This is also known as the fluid equation and can be derived independently by thermodynamic arguments, assuming that the matter-energy 'fluid' is adiabatic. The importance of the pressure term alongside density in the acceleration and fluid equations becomes apparent when describing the expansion of a Universe with multiple matter-energy components, each with their own relationship between p and ρ governed by an equation-of-state,

$$w = \frac{p}{\rho c^2}. \tag{1.32}$$

Values of w will be quoted throughout this work with $c \equiv 1$ implicitly assumed. Different equations-of-state for different species can have significantly varying impacts on the above equations. These define distinct eras of cosmological evolution depending on which species dominates.

1.2.3 Matter Domination

At redshift $z \approx 3600$ an early radiation dominated era of the Universe (see § 1.2.4 and § 1.4) transitioned to one of matter domination. In this context, matter refers to all material travelling at non-relativistic speeds and behaving as a pressureless fluid. Matter therefore has an equation-of-state, $w = 0$. In forthcoming sections, the different effects of e.g., cold dark matter and baryons on large scale structure will be explored but are grouped together when discussing the evolution of the background expansion.

The continuity equation with $w = 0$ becomes

$$\dot{\rho} + 3\frac{\dot{a}}{a}\rho = 0. \tag{1.33}$$

This can be solved simply to show the expected result that matter density scales with volume, which is represented by a^3 for a cosmological context. Hence, substituting

$$\rho_m(a) = \rho_{m,0} a^{-3}, \tag{1.34}$$

into the Friedmann equation gives a differentiable equation,

$$\frac{\mathrm{d}a}{\mathrm{d}t} = \frac{8\pi G \rho_{m,0}}{3}\frac{1}{a} \tag{1.35}$$

that must be solved to discover how the expansion evolves with time. If one assumes a power-law solution, $a(t) \propto t^{\alpha}$ then $\alpha = 2/3$ satisfies the Friedmann equation. This leads to a Hubble parameter,

$$H(t) = \frac{\dot{a}}{a} = \frac{2}{3t}, \tag{1.36}$$

that decays with time as the matter content becomes increasingly diffuse. This result is asymptotic so the expansion continues forever but becomes infinitely slow. A flat Universe with no cosmological constant and containing only matter is known as an Einstein de-Sitter (EdS) Universe and shares this evolution (Einstein and de Sitter, 1932).

1.2.4 Radiation Domination

Radiation collectively refers to the photon content of the Universe and relativistic particles that have a non-zero pressure, such as neutrinos (under certain conditions; see § 1.5) or warm dark matter particles (see § 1.6.2). The equation-of-state for radiation is $w = 1/3$, which yields a continuity equation,

$$\dot{\rho} + 4\frac{\dot{a}}{a}\rho = 0, \tag{1.37}$$

and a density evolution,

$$\rho_r(a) = \rho_{r,0} a^{-4}, \tag{1.38}$$

with a dependence on an additional factor of a beyond the scaling with volume. This can be understood by considering that the fractional change to the wavelength of radiation with the cosmic expansion is proportional to a via equation (1.12). As the density of radiation is proportional to its energy which depends on wavelength as $E \propto \lambda^{-1}$, it is apparent that ρ_r decreases by a factor of a due to this effect.

Alternatively, one can show the same result through thermodynamic arguments. Adiabatic expansion, in which the total heat of the Universe is conserved, requires the transfer of internal energy, $\mathrm{d}U$, to work done, $p\mathrm{d}V$ to obey the first law of thermodynamics, $\mathrm{d}U + p\mathrm{d}V = 0$. The change in energy density, expressed as $u \equiv U/V$, can then be shown to be

$$\mathrm{d}u = -3\left(u + p\right)\frac{\mathrm{d}a}{a}. \tag{1.39}$$

Considering the radiation equation of state, $w = 1/3$, it is clear from this that an additional power of the scale factor will appear in the radiation density evolution,

$$\mathrm{d}u_r = -4u_r\frac{\mathrm{d}a}{a}, \tag{1.40}$$

compared to matter which has $w = 0$. Solving the Friedmann equation with the substitution, $\rho_r \propto a^{-4}$, with a power law assumption for the scale factor, gives $a \propto t^{1/2}$ and a Hubble parameter,

$$H\left(t\right) = \frac{1}{2t}. \tag{1.41}$$

It is worth noting that during radiation domination the decay of ρ_r with time is $\rho_r \propto t^{-2}$ while matter decays slower as $\rho_m \propto a^{-3} \propto t^{-3/2}$. However, applying the finding that $a \propto t^{2/3}$ during matter domination gives $\rho_m \propto t^{-2}$ in this era while radiation drops off faster $\rho_r \propto t^{-8/3}$. It is therefore clear that radiation domination is inherently temporary when any non-zero matter component is present and eventually transitions to a period of matter domination that continues indefinitely, provided that there are no additional contributions to the total density to consider (Liddle, 2003).

1.2.5 Accelerated Expansion

The cosmological constant is associated with the energy density of the vacuum. If thermodynamic arguments are applied for a volume of the vacuum, V, then $\mathrm{d}\left(\rho_\Lambda c^2 V\right) = -p_\Lambda \mathrm{d}V$ implies that $\Lambda > 0$ corresponds to a negative pressure that drives the expansion of the vacuum with equation of state $w_\Lambda = -1$. At sufficiently late times, the cosmological constant term dominates over other terms in the Friedmann equation. When the matter, radiation and curvature contributions are negligible, the scale factor can be simply derived as

$$a\left(t\right) = \exp\left[Ht\right] = \exp\left[\sqrt{\frac{\Lambda c^2}{3}}t\right]. \tag{1.42}$$

This is known as the de Sitter solution and is also relevant for describing a period of inflationary expansion in the early Universe (see § 1.4.2).

A generic dark energy component entering the Friedmann equation as another fluid term is defined by having $w_{DE} < 0$ giving rise to a negative pressure contribution. For an evolving equation of state, $w_{DE}\left(a\right)$, the continuity equation has a general solution,

$$\rho_{DE} = \rho_{DE,0} \exp\left[3\int_a^1 \mathrm{d}\ln a' \left(1 + w_{DE}\left(a'\right)\right)\right]. \tag{1.43}$$

As one would expect, substituting $w_{DE} = -1$ returns a constant energy density. It can be shown from the result above that $w_{DE} < -1/3$ is the condition required for dark energy to eventually come to dominate over kc^2/a^2 in the Friedmann equation and produce accelerated expansion. Values in the range $-1 < w_{DE} \leq -1/3$ provide stable accelerating solutions. By contrast, $w_{DE} < -1$ leads to phantom solutions that would give rise to a host of extreme particle physics phenomena, an increasing dark energy density and a Big Rip spacetime singularity in the future.

1.2.6 The Hubble Parameter

Substituting the relations for the evolution of each density component into the Friedmann equation allows it to be re-expressed as the Hubble parameter,

$$\frac{H^2(z)}{H_0^2} = \Omega_m (1+z)^3 + \Omega_r (1+z)^4 + \Omega_\Lambda \exp\left[3\int_0^z \mathrm{d}z' \, \frac{(1+w(z'))}{1+z'}\right] + \Omega_k (1+z)^2 . \tag{1.44}$$

This formulation replaces the present day values of the densities with the corresponding density parameter values. The label explicitly indicating a present day value is dropped for convenience, so hereafter it may be assumed, unless stated otherwise, that $\Omega \equiv \Omega(z=0)$[5]. Through equation (1.25) the factor of H_0 encapsulates the present day value of the critical density, which relates ρ to Ω. The time evolution has been expressed in terms of redshift but could equivalently use the scale factor. The final term on the right-hand side derives straightforwardly from equation (1.28). The Hubble parameter encodes cosmological distances; through this expression these can be computed for an object at a given redshift provided that one has knowledge of the relative contributions of the different matter-energy and curvature components to the total density today. As such, this form of the Friedmann equation is extremely useful for describing key cosmological observables.

1.3 Cosmological Distances

There are a number of specific distance measures made in cosmology. Observations are often based on the apparent angular size or the luminosity of an object, leading to the 'angular diameter distance' and the 'luminosity distance' being important measures that are useful to understand in terms of the previously discussed comoving distance. This can be expressed in terms of the Hubble parameter in equation (1.44) by changing the integration variable in equation (1.14) from t to z. This gives

$$\chi(z) = \int_0^z \mathrm{d}z' \, \frac{c}{H(z')}, \tag{1.45}$$

where χ is specifically the comoving distance to an object at redshift z. This notation will be used hereafter.

[5]Chapter 5 presents a general exception to this notation as it deals with growth histories in coupled dark energy cosmologies and so frequently refers to density parameters as fully redshift-dependent quantities for these purposes.

The angular diameter distance, $D_A(z)$, can be understood by considering two objects at redshift z, separated by an angular interval $\mathrm{d}\phi$, and separated from an observer at $z = 0$ by $a(z) f_K[\chi(z)]$. The arc length between the objects, or equivalently the physical size of a single extended object, can therefore be written

$$\mathrm{d}l = \frac{1}{(1+z)} f_K[\chi(z)]\,\mathrm{d}\phi. \tag{1.46}$$

The angular diameter distance is then defined through this relation as the distance required to compute this size or separation:

$$D_A(z) \equiv \frac{1}{(1+z)} f_K[\chi(z)]. \tag{1.47}$$

The luminosity distance, D_L, is defined as the distance measure in an expanding and curved spacetime that recovers the Euclidean form of the relationship,

$$L = 4\pi D_L^2 S, \tag{1.48}$$

between the luminosity, L, of a source and the observed flux, S. Several factors determine the relationship between D_L and other distance measures. The rate of energy flow for emitted photons decreases by a factor of $(1+z)$ due to the expansion. There is also a time dilation effect due to the decrease of frequency of observed photons by a factor of $(1+z)$ with respect to the frequency of emitted photons. As the flux depends on radiation energy transferred per unit time, this leads to S decreasing by a factor of $(1+z)^2$. The surface area of a sphere centred at the source increases by another factor of $(1+z)^2$ due to the expansion. Overall the flux decreases by a factor of $(1+z)^4$. To satisfy the above form of the luminosity-flux relationship, the luminosity distance is therefore defined as

$$D_L(z) \equiv (1+z)^2 D_A(z) = (1+z) f_K[\chi(z)]. \tag{1.49}$$

The importance of luminosity distance as a cosmological observable is perhaps most apparent when Riess et al. (1998) and Perlmutter et al. (1999) showed constraints on $D_L(z)$ derived from measurements of type-Ia supernovae provided the first compelling evidence of an accelerating Universe. This famous discovery

and its implications for dark energy are discussed in detail in § 1.7.1.

1.4 The Early Universe

1.4.1 Cosmic Microwave Background

The early universe provides one of the most powerful cosmological observables. A blackbody spectrum of microwave radiation characterised by a high degree of isotropy provides a signature of an early, high-temperature radiation-dominated era. Measurements of this cosmic microwave background (CMB) from the *Planck* survey mission are used to test the validity of the Cosmological Principle and also provide some of the tightest cosmological constraints (Planck Collaboration et al., 2018). The CMB is nearly a perfect blackbody and maintains this spectrum as photons propagate through the Universe[6] so the energy of radiation can be treated as depending on temperature according to the Stefan-Boltzmann law,

$$E_r = \rho_r c^2 \propto T^4. \tag{1.50}$$

As the density of radiation decreases with expansion via $\rho_r \propto (1+z)^4$, the cosmic temperature increases with redshift as

$$T(z) = (1+z)\, T_0, \tag{1.51}$$

where $T_0 = 2.718 \pm 0.021\,\mathrm{K}$ is the value measured today (Planck Collaboration et al., 2016b). This relationship suggests that the temperature of the Universe becomes indefinitely high at sufficiently early times. When the CMB was discovered by Penzias and Wilson (1965) it provided the first compelling observational evidence in support of ideas proposed by Alpher and Herman (1948) that a microwave spectrum observed today, arising from early Universe high-energy radiation redshifted by the expansion, would imply a Hot Big Bang in the past.

The origins of the CMB can be understood in the following simple terms. Before

[6]There are small deviations due to the Sunyaev-Zel'dovich effect (Sunyaev and Zeldovich, 1970, 1980), which describes how CMB photons interact with high-energy electrons in galaxy clusters. Inverse Compton scattering increases the energy of the photon enough to produce a small spectral distortion, $\lesssim 1\,\mathrm{mK}$ in the CMB.

the energy of radiation had decreased sufficiently to be comparable to the threshold, $k_B T(z) \simeq 13.6\,\mathrm{eV}$, at which neutral hydrogen forms, the Universe could be described as a hot ionised plasma of photons and free protons and electrons. Thomson scattering coupled the photons and electrons, with a short mean free path for the former due to the high density of the latter. The Universe is therefore referred to as 'opaque' during this epoch, which ended when the radiation temperature had decreased too far for photons to satisfy the ionisation potential of hydrogen. As a result, the 'decoupling' of photons and electrons occurred due to the latter combining with free protons to form neutral atoms. This marked the time (at $z_\star = 1089.80 \pm 0.21$, see Planck Collaboration et al. (2018)) of a last scattering surface for the photons before being free to radiate through space with a blackbody spectrum. This would also suggest that the CMB would be completely isotropic. This is nearly the case, with anisotropies only of order $O(10^{-5})$ observed. These arise due to quantum fluctuations in the dense early Universe producing temperature variations that are then amplified by the expansion history. Perturbations in the background can be treated as the progenitors for the large scale structure of collapsed overdensities observed today.

1.4.2 Inflationary Cosmology

The standard cosmological picture that has been sketched so far encounters a number of significant problems that can be resolved by introducing a period of exponential, quasi-de Sitter expansion in the early Universe referred to as 'inflation' (Guth, 1981; Linde, 1982).

The first obstacle inflation overcomes is the horizon problem. An observer receiving photons from opposite directions from two different regions is just coming into causal contact with each of them. These regions should therefore not be causally connected following a history of FLRW expansion. However, the extreme isotropy of the CMB contradicts this reasoning, implying that these distant regions have attained thermal equilibrium prior to decoupling. The comoving particle horizon, which defines the distance a photon can travel in time, t, is given by

$$\tau \equiv \int_0^t \frac{\mathrm{d}t'}{a(t')} = \int_0^a \frac{\mathrm{d}a'}{[a']^2 H}, \tag{1.52}$$

which can be shown for radiation and matter dominated eras to be $\tau \propto a$ and $\tau \propto a^{1/2}$ respectively. Therefore, as the particle horizon grows with time these regions should also have been causally disconnected at last scattering. Indeed, the particle horizon only spans about a degree on the sky at this time. Some additional mechanism is therefore required for attaining such uniform temperatures across the CMB.

The flatness problem becomes clear by examining equation (1.28) and noticing that Ω_k, the curvature parameter should diverge with time unless $k = 0$ exactly. It is therefore also clear that unless Ω_k was extremely close to zero at all times in the past then the Universe should not be flat today. As $\Omega_{\rm tot} \simeq 1$ is observed today this requires an extreme fine tuning of Ω_k in the past; at the Planck scale it is required to be of order $\Omega_k \leq O\left(10^{-61}\right)$ (Baumann, 2009).

There is also the monopole problem (Guth and Tye, 1980). A significantly simplified statement of the problem is that scalar fields associated with Grand Unified Theories in particle theory are predicted to have their vacuum energy values become aligned across different locations following a phase transition from an earlier chaotic state. Topological remnants of the pre-phase conditions are predicted to remain in the form of magnetic monopoles. These should have very large energy densities due to having central scalar field values of zero. As such they should that eventually come to dominate the Universe, in contradiction to what is observed (Zeldovich and Khlopov, 1978).

Inflation solves these problems by proposing a scalar field called the inflaton. The energy density of the inflaton field drops and freezes out in the very early Universe. Any inflaton occupying a central, local peak or 'false vacuum' of its potential at this time would exhibit the following behaviour. In the false vacuum the energy density would be constant, which in analogy to the cosmological constant, is associated with a negative pressure and the generation of accelerated de Sitter expansion. After a very short time interval, the inflaton would roll down to a stable minimum of its potential, precipitating its decay into other particles. After exiting inflation, the Universe would follow the FLRW behaviour described previously.

This accelerated phase allows for a particle horizon that was once larger than

it is today. In inflationary cosmology the observable regions of the Universe were initially encapsulated by a very small volume that was able to attain thermal equilibrium before inflation, which then expanded the volume to the scales now observed. This solves the horizon problem. The flatness problem is also solved with enough inflation. This can be seen by substituting a scale factor undergoing de Sitter expansion, in analogy to equation (1.42), into equation (1.28). The curvature parameter is exponentially driven to zero. With ~ 60 'e-folds' (an increase by of a by a factor of e^{60}) of inflationary expansion, Ω_k is small enough that the subsequent evolution of the Universe has not been sufficient to cause it to deviate strongly from zero. Finally, the monopole problem is resolved by the accelerated expansion driving the monopole density to be sufficiently small that no monopoles are present in the observable Universe.

1.5 Cosmological Neutrinos

Neutrinos have three available 'flavours' associated with the electron, muon and tauon charged leptons that give rise to them through weak interactions. Neutrinos are massive particles[7] (though very light) with three eigenstates that correspond to a specific flavour through superposition. This entails a massive neutrino oscillating between different flavours, a phenomenon which allowed the property of mass to be observationally confirmed. Determining the absolute masses remains a significant challenge, however, with solar, atmospheric and accelerator experiments restricted to placing constraints on the sum of mass squared differences (e.g., Maltoni et al., 2004; Fogli et al., 2006). This leads to competing hierarchies of mass eigenstates. The consequences of this for cosmological observables, and the capacity of forthcoming cosmological surveys to constrain the mass sum in different hierarchies will be the focus of Chapter 4. The current upper bound from cosmological observations for the mass sum is $\Sigma < 0.12\,\mathrm{eV}$ (Planck Collaboration et al., 2018). Here a brief introduction to neutrinos in cosmology will be presented to contextualise this work. For a more expansive discussion of the ideas in this section, see e.g., Lesgourgues and Pastor (2006); Gerbino (2018).

In the early Universe, neutrinos exist in thermal equilibrium with the primordial

[7] Although there is a possibility that one neutrino is massless.

plasma provided that their interaction rate surpasses the expansion rate, $\Gamma > H$. When the temperature of the Universe has decreased sufficiently to $T = T_{\nu,\mathrm{dec}}$, then $\Gamma = H$ and the neutrinos decouple. The decoupling temperature can be shown to be $\sim 1\,\mathrm{Mev}$ by comparing the relative dependences on temperatures. The Hubble parameter is influenced primarily by radiation during this epoch, so can be written $H \propto T^2$. The interaction rate is given by $\Gamma = n\sigma v$, where $n \propto T^3$ is the number density, $v \simeq c$ because the neutrinos are relativistic at this time. The cross-section is given by $\sigma \simeq G_F^2 T^2$ where $G_F \simeq 10^{-5}\,\mathrm{GeV}^{-2}$ is the Fermi constant. This leads to

$$\frac{\Gamma}{H} \simeq \left(\frac{T}{1\,\mathrm{Mev}}\right)^3. \tag{1.53}$$

Following decoupling, the temperature for neutrinos scales with the expansion as $T_\nu \propto a^{-1}$. The significance of the decoupling temperature is that it is greater than the temperature at which electron-positron creation ends in the plasma. Electron-positron annihilation therefore proceeds to increase the photon number density. This is not true of neutrinos, which have decoupled. The entropy density of the plasma, which is conserved, is given by

$$s \propto g_\star T^3, \tag{1.54}$$

where $g_\star$ is the number of degrees of freedom. Before annihilation, this is $g_\star = 11/2$, given by the sum of two photon polarisation states each contributing a factor of 1 and both the electron and positron contributing two spin states each contributing a factor of $7/8$ due to their fermionic nature. After annihilation, only the photon contribution remains, requiring a temperature increase to compensate. Having decoupled, neutrinos do not experience a transfer of entropy, so it can be shown that their temperature relates to the photon temperature as

$$T_\nu = \left(\frac{4}{11}\right)^{\frac{1}{3}} T_\gamma. \tag{1.55}$$

The present day neutrino temperature is then given through knowledge of the present day observations of the CMB temperature, $T_{\gamma,0} \approx 2.725\,K$, to be $T_{\nu,0} \approx 1.95\,K$. The (relativistic) neutrino density parameter can also now be written in terms of the photon radiation density parameter. As fermions, each of the three neutrino eigenstates contributes $7/8$ degrees of freedom. A correction is introduced such that the effective number of neutrinos is $N_{\mathrm{eff}} = 3.045$ to account

for the flavour oscillations and the time required for neutrinos to fully decouple. As energy density scales as $\rho_r \propto T^4$, the relation between the density parameters is therefore

$$\Omega_\nu = \frac{7}{8} N_{\rm eff} \left(\frac{4}{11}\right)^{\frac{4}{3}} \Omega_\gamma. \tag{1.56}$$

This is the relativistic limit of the most general expression for neutrino energy density that comes from integrating over the phase-space density (weighted by energy), which is given by the Fermi-Dirac distribution (Reif, 1965),

$$F_\nu = \frac{1}{\exp\left[\left(E_\nu - \mu_\nu\right)/T_\nu + 1\right]}, \tag{1.57}$$

for neutrino fermions in thermal equilibrium, where μ_ν is a negligible chemical potential. The non-relativistic limit gives,

$$\Omega_\nu \simeq \frac{\sum_i m_{\nu,i}}{94.1\, h^2\, \mathrm{eV}}. \tag{1.58}$$

which is the relevant expression for the massive neutrinos this work is concerned with.

Neutrinos exert an influence on cosmological observables. Knowledge of the density parameter, and its relation to the mass sum, allows one to predict the effect on the Hubble parameter and hence on cosmological distances. Observables like the sound horizon at recombination or baryon decoupling should therefore bear some signature of the neutrino effect. The redshift of matter-radiation equality will also be shifted, with both the total matter and total radiation density parameters depending on neutrino contributions. This in turn influences the integrated Sachs-Wolfe (ISW) effect - where CMB photons travelling along the line-of-sight are redshifted due to travelling through gravitational potential wells generated by matter perturbations. Time variation of the potentials occurs for the radiation dominated era and later for the dark energy dominated era, but not for the matter dominated era. A shift of $z_{\rm eq}$ therefore impacts measurements of the ISW effect.

As will be discussed in some detail in Chapter 4, massive neutrinos have a substantial influence on the evolution of large scale structure in the matter dominated Universe. The interplay between gravitational and pressure forces acting on neu-

trinos is such that there exists a 'free-streaming' scale, which depends on their mass according to (Ringwald and Wong, 2004)

$$k_{\rm fs} \sim \sqrt{\frac{\Omega_m h^2}{1+z}} \left(\frac{m_\nu}{\rm eV}\right) \ {\rm Mpc}^{-1}. \tag{1.59}$$

For scales smaller than $k_{\rm fs}$, neutrinos free-stream and damp perturbations, acting to inhibit the growth of structure. Indeed, a matter perturbation growing above the free-streaming scale with neutrinos contributing to clustering scales as $\delta \propto a$ during matter domination but below the free-streaming scale this becomes $\delta \propto a^{1-3f_\nu/5}$ on linear scales, where f_ν is the ratio of the neutrino density parameter to the total matter density parameter (see e.g., Lesgourgues and Pastor, 2006). The larger the neutrino mass sum, the more significant the effect that neutrinos have on small-scale matter clustering.

1.6 Dark Matter

The matter content of the Universe can be divided into two categories: baryonic and non-baryonic. The former category is composed of particle species described in the Standard Model of particle physics. In astrophysical terms direct observables such as stars and galaxies are composed of baryons. The non-baryonic category constitutes 'dark matter' which interacts only weakly with forces other than gravity and which can therefore be observed only indirectly through its influence on the distribution of matter throughout the Universe. The relative contribution to the total matter budget of the Universe is heavily weighted towards dark matter. The present day matter density parameter, $\Omega_m = 0.3156$, accounts for a baryonic density parameter of $\Omega_b = 0.049$. Understanding the nature of dark matter and how to properly model the large scale structure that is dominated by it are both among the most important pre-occupations of cosmology.

1.6.1 Observational Evidence For Dark Matter

There are several observations that require a substantial dark matter component to explain, some of which will be detailed here. The first is that the rotation curves of spiral galaxies tend towards the velocity of matter at large radii being nearly constant (Rubin et al., 1980). The tangential velocity, v, of an object at

radius r depends on the mass interior to the orbit via

$$v\left(r\right) = \left(\frac{GM\left(<r\right)}{r}\right)^{\frac{1}{2}}, \tag{1.60}$$

assuming that the gravitational field across the galaxy is sufficiently weak that relativistic effects do not need to be accounted for. Equation (1.60) only formally applies to spherically symmetric mass distributions but approximates the behaviour for other distributions such as those of spiral galaxies. The observed flattening of the rotation curves (such that velocities at large radii are greater by several factors than predicted by a $v \propto r^{-1/2}$ decline) at large r implies that $M \propto r$, indicating that there must be more matter present than the amount of luminous matter observed. If the galactic disk, composed of baryonic matter, is surrounded by a dark matter halo with a density, $\rho\left(r\right) \sim r^{-2}$, at large radii, the rotation curves can be explained. Alternatives that rely on additional baryonic matter in low mass stars or brown dwarfs are generally not favoured by observations.

On the larger scales of galaxy clusters, further evidence for a substantial dark matter component can be found. Indeed, in 1933 Zwicky (1933) was the first to propose that the idea after calculating mass-to-luminosity ratios for galaxies in the Coma cluster that had values of several hundred. Dominating over the contribution from stars, luminous matter is mostly present in the form of a diffuse gas called the intracluster medium (ICM). This gas has been heated due to infall into the gravitational potential well of the cluster, and is consequently observable through X-ray imaging. This allows the temperature and density of the ICM to be measured, which in turn provides an estimate of the pressure. This pressure acts in opposition to the self-gravitation of the gas. Observations indicate that the self-gravity is insufficient to keep the gas bound within the cluster without an additional matter component. Indeed the total cluster mass would need to be larger by approximately an order of magnitude to reconcile observations (e.g., Allen et al., 2011). Cluster measurements have been supported by findings from studies that indicate a dark matter presence is responsible for the strength of observed lensing of light signals from background objects by foreground galaxies and clusters (a consequence of the warping of spacetime by massive objects in General Relativity, discussed further in § 2.3). For a details of the implications

for dark matter from lensing (see e.g., Massey et al. (2010)).

Baryon acoustic oscillations are another signature of dark matter. The tight coupling between baryons and photons in the primordial plasma prior to decoupling provides a radiation pressure that acts on the baryons in opposition to the gravitational attraction. This leads to an oscillation of the radiation driving baryons outwards, collapse re-occurring as the pressure drops away, and the renewed increase of pressure due to the increased density via $p = \rho c^2/3$. In other words, the initial perturbation is subject to a sound wave that carries baryons and photons together from a central dark matter component experiencing a gravitational interaction only. Following decoupling the rapid decay of the radiation pressure as the photons diffuse deposits the baryonic matter at a fixed radius (the 'sound horizon') from the centre of the perturbation. This leads to a signature in the CMB power spectrum and the matter power spectrum, and serves as a 'standard ruler' by establishing the comoving size of a physical observable as a benchmark for other measurements of cosmological distance. As the oscillations require a non-baryonic component that does not interact with radiation to provide a suitably strong gravitational potential well, the observation of BAOs also corroborates the existence of dark matter.

1.6.2 Candidates

N-body simulations by Davis et al. (1985) are fit well by 'cold dark matter' (CDM), which has the properties of a weakly interacting (i.e., almost exclusively interacting gravitationally) fluid composed of collisionless particles travelling at non-relativistic velocities. Weakly Interacting Massive Particles (WIMPs) represent a class of candidates with natural explanations arising from particle physics. They typically have a mass in the range $m_X \sim 10\,\mathrm{GeV}$ - $1\,\mathrm{TeV}$ (the approximate scale of electroweak interactions) and could be attributed to a range of particles. An attractive feature of WIMPs (known as the 'WIMP miracle') is that supersymmetric counterparts to Standard Model particles (e.g., the neutralino) in this mass range could also have a self-annihilation cross-section that would allow the observed dark matter abundance to be explained by the following process. When the temperature of the early Universe decreased below the corresponding mass scale of the dark matter particle, there would have been insufficient energy for

pairs of the dark matter particle and its anti-particle to form. Annihilations between particles and anti-particles would then lead to the exponential suppression of the number density of dark matter particles in the Universe (see e.g., Kolb and Turner, 1990). This emerges from the Boltzmann equation,

$$\frac{\mathrm{d}n}{\mathrm{d}t} = -3Hn - \langle \sigma v \rangle \left(n^2 - n_{\mathrm{eq}}^2 \right), \tag{1.61}$$

where n is the number density (with n_{eq} being the number density during thermal equilibrium) and $\langle \sigma v \rangle$ is the cross-section. The cosmic expansion dilutes the dark matter content via $-3Hn$ while the final term on the right-hand side of the equation controls both the creation and annihilation of pairs. As a consequence of the expansion, the incidence of annihilations drops with the ultimate consequence of dark matter 'freezing out' and asymptotically approaching a constant number density. It can be shown at the time of freeze out, that the Boltzmann equation gives a number density of the form (see e.g., Feng, 2010),

$$n_f \sim (m_X T_f)^{\frac{3}{2}} \exp\left[-\frac{m_X}{T_f} \right] \sim \frac{T_f^2}{\langle \sigma v \rangle}, \tag{1.62}$$

where T_f is the temperature at freeze out. The observed abundance of dark matter emerges from an electroweak-scale particle forming in thermal equilibrium, a result termed the 'WIMP miracle'. It should be noted, however, that detections of WIMPs remain elusive.

An alternative approach is that dark matter is composed of axions, particles that were first proposed as a solution to problems in quantum chromodynamics (Peccei and Quinn, 1977). Specifically, a fixed axion field is used to minimise the vacuum energy, requiring an extremely light mass to satisfy cosmological constraints. Unlike WIMPs, a non-thermal mechanism allows for the possibility of axions composing dark matter. In the early Universe, the axion field undergoes a phase transition from a random value to the value required to minimise the vacuum in the present. This generates a relic density that could be associated with dark matter. An example of an axionic dark matter candidate is fuzzy dark matter. For a review of axion cosmology see Marsh (2016).

1.7 Dark Energy

In previous sections, the effect of the cosmological constant in the Friedmann equation generating an era of accelerated expansion was discussed. Here, the broader observational evidence for cosmic acceleration will be reviewed, the motivation for dark energy as a class of mechanisms for explaining it will be outlined, and a range of the most prominent candidate models will be discussed. It is important to state that, in the context of this Chapter, 'dark energy' will be used as an umbrella term that includes the effect of a cosmological constant, additional scalar fields and modifications to gravity.

1.7.1 Evidence

In the present, the observed flatness of the Universe, a matter density parameter of $\Omega_m \simeq 0.3$ (including baryons and dark matter) and a negligibly small contribution from radiation requires an additional matter-energy component. This is attributed to dark energy which, as $\Omega_{DE} \simeq 0.7$, dominates over the other components, driving an era of accelerated expansion.

Evidence for a Universe featuring a cosmological constant was mounting in the years preceding the famous type-Ia supernovae results (Krauss and Turner, 1995). For example, it was apparent from distance measurements to the Virgo Cluster that the Hubble constant required for a flat geometry predicted a Universe that was several billion years younger than the measured ages of the oldest globular clusters. This could be explained by a cosmological constant effectively generating a repulsive force. The Universe would have to have been expanding for longer than a matter-dominated Universe with the same present observed matter content and Hubble constant.

Observations of type-Ia supernovae (SNIa) (Riess et al., 1998; Perlmutter et al., 1999) were a historic result towards ruling out a decelerating expansion. SNIa are a result of a primary star in a binary system losing material to a secondary white dwarf. The accretion of matter onto the white dwarf will, if it is carbon-oxygen rich and has a low rotational velocity, cause the dwarf to approach the Chandrasekhar mass, $1.44 M_\odot$. Beyond this threshold, self-gravity overcomes electron degeneracy pressure which would lead to collapse into a neutron star. However,

SNIa are believed to occur due to carbon fusion being induced in the white dwarf core as its temperature rises significantly when approaching the limit. The final consequence of the ignition of the core is the high-energy release of a supernova (Wheeler, 2007).

These events can be considered 'standardisable candles'. A 'standard candle' is a class of objects that share an absolute magnitude (a measure of intrinsic luminosity). SNIa do not share have this property but their absolute magnitude is strongly correlated with the rate of decrease of their flux density. A correction can then be introduced to normalise the light curve (showing the evolution of luminosity) (Phillips, 1993) and allow these supernovae to be effectively used as standard candles. As SNIa are also significantly more luminous than galaxy populations, this makes them useful arbiters of relative distances in cosmology, with Hamuy et al. (1996) finding that they could be calibrated to 7% accuracy for this purpose.

Cepheid variable stars constitute a set of local standard candles through a relationship between their luminosity and their pulsation period (e.g., Macri et al., 2006). These allow for calibrations of SNIa that lead to measurements of the luminosity distance. Through this H_0 can be constrained. Riess et al. (1998) and Perlmutter et al. (1999) showed that these results from SNIa indicated an accelerating expansion.

Cosmic acceleration arises when the relative contributions to the matter-energy content of the Universe are such that $\ddot{a} > 0$ is satisfied. Historically, it was expected that the cosmic expansion would slow down so a dimensionless 'deceleration parameter',

$$q \equiv \frac{\ddot{R}R}{\dot{R}^2} = -\frac{\ddot{a}}{a\,H^2}, \tag{1.63}$$

was defined with the sign convention anticipating negative $\ddot{a}$. This can be expressed in terms of density parameters by appropriately substituting the Friedmann and acceleration equations to give the current deceleration parameter,

$$q_0 = \frac{1}{2}\sum_i (1 + 3w_i)\,\Omega_i \tag{1.64}$$

where $i = \{m, r, \Lambda\}$ in the standard model. The contribution from radiation can be neglected for much of the expansion. When $w_m = 0$ and $w_\Lambda = -1$ are substituted the condition for for current acceleration can be written as

$$\frac{\Omega_m}{2} - \Omega_\Lambda < 0. \tag{1.65}$$

Riess et al. (1998) found that measurements of the distances to SNIa from two different methods gave Ω_m and Ω_Λ consistent with $q \leq 0$, and hence a currently accelerating cosmos, with confidence in the range 2.8σ to 3.9σ (assuming a prior of $\Omega_m \geq 0$). This is illustrated in Figures 6 and 7 of Riess et al. (1998), to which the interested reader is referred.

This result is corroborated by a number of other cosmological observables. One example that was previously alluded to is the integrated Sachs-Wolfe effect. Photons from the CMB travelling through gravitational potential wells experience redshifting and blueshifting. These effects do not counter each other during accelerated expansion because the potential well is diminished in the presence of dark energy over the timescale of the passage of the photon. This leaves a residual redshift that acts as a dark energy signature. CMB anisotropies provide an additional repository of information that supports a flat Universe by constraining $\Omega_m = 0.315 \pm 0.007$ and $\Omega_K = 0.001 \pm 0.002$ (Planck Collaboration et al., 2018). This requires an additional matter-energy component that can be attributed to dark energy, $\Omega_{DE} = 0.6847 \pm 0.0073$. The constraint on the curvature density parameter also incorporates information from the BAO scale which, as has been discussed, acts as a standard ruler. This provides constraints on the angular diameter distance, which depends on dark energy through the Hubble parameter. The matter power spectrum also features a BAO signature. As will be seen through the rest of this work the matter power spectrum is a valuable probe for determining cosmological constraints, depending on Ω_m and $H(z)$.

1.7.2 Equation of State For Dark Energy

Efforts to use observational probes to determine the nature of the underlying dark energy mechanism for acceleration focus on constraining the equation of state. As previously discussed, the condition for acceleration is that $w < -1/3$. Observations (e.g., CMB+BAO+SNIa results from Planck Collaboration et al.

(2018)) are compatible with a cosmological model featuring cold dark matter and a cosmological constant (ΛCDM), so $w = -1$ is usually a good choice as an assumed (or 'fiducial') value for the equation of state.

In general, however, there is no requirement that the equation of state have a static value. (Chevallier and Polarski, 2001) and Linder (2003) introduced a parameterisation in terms of the scale factor,

$$w(a) = w_0 + w_a(1 - a), \tag{1.66}$$

where w_0 is the value of the equation of state in the present and $w_a \equiv -\mathrm{d}w/\mathrm{d}a$ controls how rapidly it varies over time. For ΛCDM, $w_a = 0$. This is also known as the Chevallier-Polarski-Linder (CPL) model and can in principle capture the behaviour that is relevant for observational probes for of the wide range of dark energy models that have a dynamic component. Implementing the (w_0, w_a) parameterisation in equation (1.44) gives the Hubble parameter,

$$\frac{H^2(a)}{H_0^2} = \Omega_m a^{-3} + \Omega_r a^{-4} + \Omega_{DE} a^{-3(1+w_0+w_a)} \exp\left[-3w_a(1-a)\right] + \Omega_k a^{-2}, \tag{1.67}$$

in terms of the scale factor. An equivalent formulation of the CPL model is

$$w(a) = w_0^p + w_a^p(a_p - a), \tag{1.68}$$

where $w_0^p = w_0 + w_a(1 - a_p)$, $w_a^p = w_a$ and $a_p = 1/(1 + z_p)$ corresponds to a 'pivot' redshift, z_p at which w is most tightly constrained. If $z_p = 0$ the standard CPL model is recovered. Alternatively, z_p can be treated as a free parameter with its true value determined by the choice of priors, probes and underlying model. This parameterisation is designed to maximise the constraints available for a dark energy parameter; in general, w_0^p can be more tightly constrained than w_0. Many works use the pivot parameterisation but in this thesis, the w_0-w_a parameter space is used. This is partly to retain a clearer interpretation of the physical meaning of the dark energy parameters but also because the focus of e.g., Chapter 3 is to constrain the overall $w(a)$ space, which is not determined by the choice of parameterisation.

Having detailed some of the observational evidence for dark energy and the pa-

rameters of interest for dark energy probes, a discussion of the different proposals for the nature of dark energy will be presented. This will begin with a deeper examination of the cosmological constant itself and its relationship to physical interpretations of the energy of the vacuum.

1.7.3 The Cosmological Constant

The original introduction of the cosmological constant in Einstein's field equations was to counteract an expanding spacetime and allow for a static Universe. The observations supporting Hubble expansion lead to the concept being abandoned. However, establishing a physical understanding of what the cosmological constant represented remained an engaging question. The term $\Lambda g^{\mu\nu}$ contributes a perfect fluid with a negative pressure ($w_\Lambda = -1$) to the energy-momentum tensor. Negative pressure is not a property of classical fluids, but can be arise in a system governed by quantum mechanics. Heisenberg's uncertainty principle (Heisenberg, 1927; Kennard, 1927),

$$\sigma_x \sigma_p \geq \frac{\hbar}{2} \tag{1.69}$$

where $\hbar = h/(2\pi)$ is the reduced Planck's constant, sets a fundamental boundary on the precision to which the position and momentum of a particle can be measured. This has an equivalent formulation in terms of energy and time, which has the consequence of allowing pairs of particles and antiparticles to manifest for brief time intervals. The lowest energy state of a vacuum is therefore non-zero. Variations in the zero-point energy generate forces via e.g., the Casimir effect that are important for laboratory experiments in condensed matter and particle physics, but the significance for cosmology is that through Einstein's field equations the absolute energy is what determines the evolution of the Universe.

If one attempts to introduce the energy density associated with the vacuum into the field equations, while imposing both Lorentz invariance for inertial observers of the vacuum energy properties and general covariance, an energy-momentum tensor must be written down in the form

$$T_V^{\mu\nu} = -\rho_v c^2 g^{\mu\nu}. \tag{1.70}$$

This has the same form as the cosmological constant term in the field equations, and yields a negative pressure with $w = -1$. However, it should be stressed that the vacuum energy density and the cosmological constant are separate concepts.

Determining a value for the vacuum energy density that can be satisfactorily applied to Λ while being consistent with observations of the expansion presents a severe challenge. Quantum field theory provides a prescription for calculating energy densities by quantising each harmonic mode of the field in question and then summing the energy contribution, $(\hbar\omega)/2$ from each frequency state. For the vacuum, a crude estimate of the energy density is given by

$$\rho_v c^2 = \int_0^{\Lambda_v} \frac{\mathrm{d}^3 k}{(2\pi)^3}\, \hbar\omega = \frac{\Lambda_v^4}{16\pi^2} + ... \tag{1.71}$$

where a series of terms have been neglected for either not being of leading order in Λ_v or for being unphysical by breaking Lorentz invariance[8] (Koksma and Prokopec, 2011). If the cut-off energy scale, Λ_v, is given by the Planck mass, $M_{\rm Pl} \equiv \sqrt{\hbar c/8\pi G}$, then the prediction for the vacuum energy density is $\approx 10^{76}\,\mathrm{GeV}^4$. However, cosmological observations indicate that the energy density associated with a cosmological constant is $\approx 10^{-47}\,\mathrm{GeV}^4$ in order to satisfy $\Omega_{DE} \approx 0.7$. This presents a fine-tuning problem (called the cosmological constant problem) for Λ in which a discrepancy of roughly 120 orders of magnitude must be explained (Weinberg, 1989; Adler et al., 1995). It should be noted that the problem as presented here is a surface-level explanation and neglects a number of nuances entering the assumptions regarding e.g., whether quantum field theory is a robust description up to the Planck mass.

Efforts to resolve the cosmological constant problem generally require a theoretical basis for cancelling the vacuum energy density predicted by quantum field theory, and invoking a separate explanation for the accelerated expansion and the observed value of Λ. Appeals have been made to super-symmetry, which would allow the contribution of zero-point energy from a fermion field to be negated by the zero-point energy from its bosonic superpartner's field (Bailin and Love, 1994). However, the need to accommodate the breaking of super-symmetry

[8] In fact, the remaining term also violates this condition but is generally retained to allow the cosmological constant to be renormalisable, i.e., to treat infinite terms arising from the calculations and produce a sensible result.

tends to require that a time-varying gravitational constant that is not compatible with constraints (Martin, 2012). Alternatively, quantum gravity approaches have sought to show that the Hartle-Hawking wavefunction of the Universe (Hartle and Hawking, 1983) naturally peaks at $\Lambda = 0$, but these efforts have not resolved a number of theoretical obstacles. With a satisfactory solution unavailable at present, several of the candidates for other dark energy prescriptions that can explain the acceleration will now be explored.

1.7.4 Dark Energy Models

The model space of dark energy candidates is rich and vast. Most of the proposals in the literature will not be discussed here. For a review of dark energy that focuses on the introduction of scalar fields see Copeland et al. (2006), while for modified gravity see Clifton et al. (2012). The presentation of models will also broadly follow discussions presented in these reviews. It is possible to construct models with a dark energy component interacting with dark matter. Coupled dark energy (in the context of large scale structure probes) is the focus of Chapter 5, and so the background to this class of theories is deferred to there rather than being discussed here.

Quintessence

Scalar fields are an attractive approach to generating the observed cosmic acceleration. The far more extreme, rapid acceleration during inflation is facilitated through a scalar field, so it is reasonable to suppose that dark energy may also be explained this way. They also open the possibility of a motivation rooted in particle physics. A leading scalar field candidate model is 'quintessence' (Wetterich, 1988). The action, which captures the dynamics of a system, for a quintessence cosmology is

$$S = \int \mathrm{d}^4x\sqrt{-g}\left[-\frac{1}{2}g^{\mu\nu}\partial_\mu\phi\partial_\nu\phi - V(\phi)\right]. \tag{1.72}$$

The quintessence scalar field, ϕ, has a potential energy density, $V(\phi)$, that is chosen to produce sensible results and perhaps to relate to an underlying physical motivation, for example by treating ϕ as arising from string theory. The first term in brackets in the integrand provides the kinetic part of the action. If conditions of spatial flatness and a homogeneous ϕ are imposed then varying the action with

respect to ϕ gives the equation of motion,

$$\ddot{\phi} + 3H\dot{\phi} + \frac{\mathrm{d}V}{\mathrm{d}\phi} = 0. \tag{1.73}$$

An alternative definition for the energy-momentum tensor that has not been presented previously is through varying the action with respect to the metric:

$$T_{\mu\nu} \equiv -\frac{2}{\sqrt{-g}}\frac{\delta S}{\delta g^{\mu\nu}} = \partial_\mu\phi\partial_\nu\phi - g_{\mu\nu}\left[\frac{1}{2}g^{\rho\sigma}\partial_\rho\phi\partial_\sigma\phi + V(\phi)\right]. \tag{1.74}$$

The energy density and pressure of the field are found by computing $-T^0_0$ and T^i_i respectively. If the spatial derivatives of the field can be neglected, the equation of state can be shown to be

$$w_\phi = \frac{\frac{\dot{\phi}^2}{2} - V(\phi)}{\frac{\dot{\phi}^2}{2} + V(\phi)}. \tag{1.75}$$

It can be seen that $-1 \leq w_\phi \leq 1$, with the trend towards ΛCDM behaviour, $w_\phi \longrightarrow -1$, emerging when $\dot{\phi}^2 \ll V(\phi)$. This is known as the 'slow-roll' limit. In order to generate accelerated expansion, $w_\phi < -1/3$, which corresponds to the energy density of the field scaling with the expansion as $\rho \propto a^{-k}$ for $0 < k < 2$.

K-essence

K-essence represents the opposite approach to quintessence in using scalar fields to generate accelerated expansion (Chiba et al., 2000). In the latter this result arises through the potential energy but in K-essence the kinetic term (which is condensed into the notation, X, here) in the action is the source of acceleration. This action is written in the form

$$S = \int \mathrm{d}^4x\sqrt{-g}f(\phi)\hat{p}(X), \tag{1.76}$$

which corresponds to the form of the action of a low-energy effective string theory. With the appropriate redefinition of the kinetic term and the field in terms of the coupling functions in such an action, it is possible to write

$$S = \int \mathrm{d}^4x\sqrt{-g}f(\phi)\left(-X + X^2\right). \tag{1.77}$$

This leads to the equation of state

$$w_\phi = \frac{1 - X}{1 - 3X}, \tag{1.78}$$

which corresponds to accelerated expansion for $-1/2 \leq X \leq 2/3$ and the cosmological constant for $X = 1/2$.

Modified Gravity

Accelerated expansion can also be achieved by extending or modifying the Einstein field equations. Generally, modified gravity models can be categorised as those that rely on i) higher derivatives; ii) extra dimensions; or iii) extra fields.

An example of the second category is $f(R)$ gravity (Buchdahl, 1970) which generalises the Ricci scalar in the Einstein-Hilbert action that gives rise to field equations. This leads to fourth order derivatives appearing whereas derivatives no higher than second order are found in General Relativity. Many modified gravity mechanisms generate a period of late-time accelerated expansion. However, these typically conflict with solar system tests that place stringent constraints on deviations from General Relativity. To reconcile these discrepancies screening mechanisms are introduced. An example is the 'chameleon mechanism' which entails non-minimally coupled scalar fields acquiring effective mass in high density environments, which allows for gravitation that resembles General Relativity locally and produces acceleration on larger scales for which the density is much lower (Khoury and Weltman, 2004).

The detection of gravitational waves from a binary neutron star merger (LIGO Scientific Collaboration and Virgo Collaboration, 2017) strongly constrains the low redshift speed of gravitational waves to be very close to the speed of light. This restricts the form that the Lagrangian describing gravitational dynamics can take, as first shown by Lombriser and Taylor (2016). As a result many modified gravity models have been disfavoured, particularly those that rely on quartic-order derivatives and quintic-order couplings. Quintessence, k-essence and $f(R)$ gravity (among other theories) do satisfy the gravitational wave constraints, however (Kase and Tsujikawa, 2019).

1.8 Summary

This Chapter has provided a foundational review of the standard cosmological model. Beginning with the essential precepts of General Relativity, the FLRW metric has been introduced and used to construct the Friedmann equation describing an expanding Universe. Particular attention has been paid to the importance of the Cosmological Principle and its conditions of homogeneity and isotropy. The observational evidence for expansion was also outlined. The Hubble parameter as a function of redshift was discussed, along with the key measures of cosmological distance.

The contribution of matter, radiation and a cosmological constant (or dark energy, more generally) to the evolution of the Universe has been detailed, establishing the different eras of domination by different components and expansion rate during these times. A brief explanation of the important features of the early Universe has been presented, in particular the CMB as one of the most powerful observables in cosmology. The paradigm of cosmic inflation in the early Universe has been introduced, primarily to emphasise its importance in solving several problems in standard cosmology and providing an explanation for large scale structure at later times.

Cosmological neutrinos, dark matter and dark energy have each been explored as these components of the matter-energy content of the Universe play important roles in this thesis. The formalism describing their effects on cosmological evolution, the observational evidence for their contribution to the total matter-energy budget, and major candidate models have been discussed in each case.

Having provided the basis for the background cosmology, it is now necessary to explore the formation and evolution of large scale structure within the Universe. This thesis focuses on large scale structure probes to explore constraints on dark energy and neutrinos, so it is important to devote some attention to outlining the most significant aspects of this field.

Chapter 2

Probes of Large Scale Structure

2.1 Linear Structure Formation and Evolution

2.1.1 Linear Density Fields

There are a number of thorough reviews of the familiar ideas presented in this section. Peacock (1999) is recommended in particular for a discussion of greater depth. The large scale structure observed today has emerged as a complex network of connected haloes and filaments. On scales, $\lesssim 100\,\mathrm{Mpc}$, where features of structure can be distinguished, the homogeneity of the Universe breaks down. This is the eventual result of small inhomogeneities, fluctuations in the large scale matter density field that have been seeded in the early Universe by a mechanism like inflation. Perturbations that increase the density of local regions sufficiently compared to the background are subject to collapse through self-gravitation, as more matter is accumulated by the region which enhances the local density contrast still further. An evolving, comoving spatial perturbation, $\delta(\mathbf{x}, t)$, about the mean cosmological matter density, $\bar{\rho}(t)$, is defined such that the matter density field is written as

$$\rho(\mathbf{x}, t) = \bar{\rho}(t)\left[1 + \delta(\mathbf{x}, t)\right]. \tag{2.1}$$

It is useful to interpret cosmological perturbations through the FLRW framework that has been laid out previously. The metric can accommodate perturbations most generally by introducing the two Bardeen scalar potentials, Ψ and Φ, such

that (Bardeen, 1980)

$$\mathrm{d}s^2 = -\left(1+\frac{2\Psi}{c^2}\right)c^2\mathrm{d}t^2 + \left(1-\frac{2\Phi}{c^2}\right)a^2(t)\left[\mathrm{d}r^2 + f_K^2(r)\left(\mathrm{d}\theta^2+\sin^2\theta\mathrm{d}\phi^2\right)\right]. \tag{2.2}$$

When deriving Einstein's field equations from this metric, it can be shown that in order for the energy-momentum tensor to represent a perfect fluid with no anisotropic stress that $\Psi = \Phi$ is a necessary condition. Otherwise non-zero off diagonal terms appear in $T_{\mu\nu}$. The potential, Φ, is analogous to the gravitational potential, Φ, in Newtonian gravitational dynamics that has already been introduced in the Poisson equation.

As has been discussed, within a spherical region of uniform density matter experiences no net gravitational force from external matter. A spherical region can then be roughly treated as a separate universe, embedded within the FLRW background. As the background is well described by $\Omega_{\mathrm{tot}} = 1$ then for an overdense spherical region, $\Omega < 1$. This is equivalent to the region behaving as a closed universe with $k = 1$ which will collapse following a period of expansion. An associated Friedmann equation using the physical radius of the perturbation, r, normalised at some early time to the scale factor is

$$\left(\frac{\dot{r}}{r}\right)^2 = \frac{8\pi G}{3}\rho - \frac{k}{r^2}. \tag{2.3}$$

The term on the left-hand side represents the effective Hubble parameter. By setting this equal to the Hubble parameter for the background at some early time and comparing this Friedmann equation to that of the background, the (initial) linear density contrast emerges directly as

$$\delta \equiv \frac{\rho}{\bar{\rho}} - 1 \propto \left(\bar{\rho}r^2\right)^{-1}. \tag{2.4}$$

Close to the initial time, when $r(t) \simeq a(t)$, the perturbation is small and scales as $\delta \propto a^2$ during the radiation dominated era and $\delta \propto a$ during the matter dominated era.

At this stage, it is useful to introduce the Fourier transform of $\delta(\mathbf{x})$. This allows the overdensity (and derived quantities like the power spectrum), to be expressed

in terms of wave modes:

$$\tilde{\delta}\left(\mathbf{k},t\right)=\int \mathrm{d}^3x\,\delta\left(\mathbf{x},t\right)\mathrm{e}^{i\mathbf{k}\cdot\mathbf{x}}. \tag{2.5}$$

This is a particularly useful tool in linear theory because in Fourier space, the modes evolve independently (in a homogeneous Universe). As will become apparent in later derivations, working in Fourier space also simplifies many computations, transforming convolutions of quantities into straightforward multiplications. The wavenumber, k, is related to a (comoving) wavelength via $k=(2\pi)/\lambda$. Growing modes that satisfy the condition, $\lambda<\lambda_H\left(a_{\mathrm{eq}}\right)$, where $\lambda_H\left(a_{\mathrm{eq}}\right)$ represents the comoving horizon scale (the distance a photon could travel since the origin of the Universe) at matter-radiation equality are said to have 'entered the horizon' during the radiation dominated era. As expansion is more rapid than during matter domination, there is a relative suppression of growth (Meszaros, 1974).

2.1.2 Linear Growth

The evolution of perturbations can be understood by applying the non-relativistic fluid equations. A fluid of density ρ with velocity $\mathbf{v}$ can be described by the continuity equation,

$$\frac{\partial\rho}{\partial t}+\nabla\cdot\left(\rho\mathbf{v}\right)=0, \tag{2.6}$$

the Euler equation,

$$\frac{\partial\mathbf{v}}{\partial t}+\left(\mathbf{v}\cdot\nabla\right)\mathbf{v}=-\frac{\nabla P}{\rho}-\nabla\Phi, \tag{2.7}$$

where the gravitational potential is related to the density via the Poisson equation,

$$\nabla^2\Phi=4\pi G\rho, \tag{2.8}$$

which is restated here for convenience. The continuity equation ensures that if the density of a region decays it is accompanied by an outward flux of matter through the surface of the region. The Euler equation describes the acceleration of fluid particles under gravitational and pressure forces. Note that the left-hand side is the form describing the rate of change of flow as perceived in a frame that comoves with the fluid. It therefore combines the time evolution perceived in a frame at a position, $\mathbf{x}$, that is stationary with respect to the flow, and the local

spatial gradient of the flow given by ∇.

One can derive the same set of equations for the matter perturbation, δ, the peculiar velocity perturbation, $\mathbf{u}$, given by $\mathbf{v} = H\mathbf{x} + \mathbf{u}$ and a perturbation, ϕ, in the gravitational potential. It is useful to change coordinates from $\mathbf{x}$ to a set of Lagrangian coordinates, $\mathbf{q} = a^{-1}\mathbf{x}$, that are comoving with the Hubble flow of the background fluid. This leads to

$$\dot{\delta} + \frac{1}{a}\nabla \cdot \mathbf{u} = 0 \tag{2.9}$$

$$\dot{\mathbf{u}} + H\mathbf{u} = -\frac{c_s^2}{a}\nabla\delta - \frac{1}{a}\nabla\phi \tag{2.10}$$

$$\nabla^2\phi = 4\pi G a^2 \bar{\rho}\delta, \tag{2.11}$$

at linear order in the perturbations. For CDM and radiation, the term $c_s^2 \equiv \mathrm{d}p/\mathrm{d}\rho = wc^2$ effectively defines the sound speed of a perfect fluid. The linear growth equation describing the evolution of the density perturbation results from a combination of these equations. In Fourier space the spatial second derivative, ∇^2, is equivalent to a factor of $-k^2$, which gives

$$\ddot{\delta} + 2H\dot{\delta} = \left(\frac{3}{2}H^2\Omega - \frac{c_s^2 k^2}{a^2}\right)\delta. \tag{2.12}$$

This can be applied to each individual matter-energy component. During the radiation dominated era, the sound speed of the photon-baryon fluid is given by $c_s \lesssim c/\sqrt{3}$ (due to the baryon component being pressure-free) and is therefore of the order of the speed of light. Therefore the 'sound horizon' is approximately given by the comoving horizon. The radiation pressure is sufficiently strong to prevent gravitational collapse for perturbations smaller than horizon scales. The above expression can therefore be written purely in terms of the matter perturbation consisting of coupled CDM and baryons. The final right-hand side term then disappears as $c_s = 0$ for matter. The second term on the left-hand side is known as the 'Hubble drag' and acts as an effective friction term, slowing the growth of structure.

On sufficiently large scales, the spatial gradients of the perturbation vanish. In linear theory one can write $\delta(\mathbf{q}, t) = \delta(\mathbf{q}, t_0)\, D(t)$, thereby isolating the time-evolution in the form of the linear growth factor, $D(t)$. When expressed in terms

of this quantity the growth equation yields two solutions such that

$$D(t) = c_1 t^{\frac{2}{3}} + c_2 t^{-1}, \tag{2.13}$$

where the constants, c_1 and c_2, control the respective contributions of a growing mode and a decaying mode to the growth factor.

2.1.3 Linear Matter Power

The statistical features of the density field at different scales describe structure in the Universe and underpin the probes of cosmology that this thesis focuses on. One of the most important statistical tools is the power spectrum of matter fluctuations. On large scales ($k \lesssim 0.1\,h$/Mpc today), this can be constructed from linear overdensities that evolve according to the perturbation equations that have been outlined. On smaller scales, the linear approximations break down, more complex structure develops and sophisticated models of the non-linear matter power spectrum are required. Here, only the linear power spectrum will be explored, while later an approach to the non-linear power spectrum through the halo model (used in this thesis) will be discussed.

Fourier modes of δ are treated as uncorrelated and evolving independently in a homogeneous Universe. Due to the further fact that the quantum fluctuations assumed to give rise to the initial density perturbations are modelled as an inherently random Gaussian quantum harmonic oscillator (in the ground state), the density field can be treated as a random Gaussian field. If the field is sampled, one will find randomly varying δ values at different positions, x. Furthermore, if one could generate multiple iterations of the Universe the value of δ would vary randomly. The mean value of the overdensity is necessarily $\langle\delta\rangle = 0$, so the variance of the field is $\langle\delta^2\rangle$. As there is a single iteration of the Universe to observe, one must rely on there being a sufficiently large volume of points at which the density field can be measured that averaged values are statistically equivalent to averaging over multiple iterations of the Universe. In other words, the density field is assumed to be ergodic. The two-point correlation function,

$$\xi(|\mathbf{x}|) = \langle\delta(\mathbf{x}+\mathbf{x}')\,\delta^{\star}(\mathbf{x}')\rangle \tag{2.14}$$

encodes all of the statistical properties of the density field. In Fourier space, it is given by performing the transform,

$$\langle \tilde{\delta}\left(\mathbf{k}\right) \tilde{\delta}^{\star}\left(\mathbf{k}'\right) \rangle = \int \mathrm{d}^3 x' \, \mathrm{e}^{i\mathbf{k}\cdot\mathbf{x}'} \int \mathrm{d}^3 x \, \mathrm{e}^{-i\mathbf{k}'\cdot\left(\mathbf{x}+\mathbf{x}'\right)} \xi\left(|\mathbf{x}|\right) = \left(2\pi\right)^3 \delta_D\left(\mathbf{k}-\mathbf{k}'\right) P\left(k\right), \tag{2.15}$$

where

$$P\left(k\right) \equiv \int \mathrm{d}^3 x \, \mathrm{e}^{-i\mathbf{k}\cdot\mathbf{x}} \xi\left(|\mathbf{x}|\right) \tag{2.16}$$

defines the matter power spectrum. Here, δ_D represents the delta function. It is often more convenient to use the dimensionless form,

$$\Delta^2\left(k\right) \equiv \frac{k^3}{2\pi^2} P\left(k\right), \tag{2.17}$$

which is equivalent to the fractional contribution to the variance of the matter distribution per logarithmic interval of k,

$$\sigma^2\left(R\right) = \int_0^{\infty} \mathrm{d} \ln k \; \Delta^2\left(k\right) W^2\left(kR\right), \tag{2.18}$$

where the field is often smoothed over some scale R using the window function of a spherical top-hat profile,

$$W\left(x\right) = \frac{3}{x^3}\left(\sin x - x\cos x\right). \tag{2.19}$$

The power spectrum is often normalised by fixing the value of σ at $R = 8\,\mathrm{Mpc/h}$.

On sufficiently large scales, the power spectrum takes the form $P\left(k\right) \propto k^{n_s}$. The spectral index, n_s, describing the slope of the power has a value close to 1 (which satisfies an expectation that the power spectrum is nearly scale-invariant on large scales) (see e.g., Peacock, 1999). Matter-radiation equality leads to a break scale in the linear power spectrum, with larger k giving power as $P\left(k\right) \propto k^{n_s-4}$. This is due to the Meszaros effect (Meszaros, 1974) that describes how the rapid cosmic expansion of the radiation dominated era freezes perturbations associated with modes that have entered the horizon. Their growth is prevented until matter-radiation equality, while superhorizon modes (associated with k values smaller than the break scale) are not suppressed. The suppression of δ for subhorizon modes scales as k^2, leading to the linear power being suppressed by 4 powers of k. This form is, however, the most generic conceptualisation. The power spectrum

is dependent on a number of cosmological parameters. For example, the density parameters control important features such as matter-radiation equality and the break scale. The equation of state of dark energy influences the Hubble expansion, which determines the linear growth rate (which the matter power spectrum is proportional to the square of) and hence the amplitude of the power.

2.2 The Halo Model

An accurate prescription for matter power on non-linear scales presents a steeper challenge than those of modelling the linear power. One could attempt to scale the linear power spectrum by some functional form that can be fit to simulations. A more sophisticated approach to relating non-linear and linear power is available through the halo model. In this framework it is assumed that all the matter in the Universe is contained in (spherical) haloes[1]. Modelling the distribution of these haloes as a function of their mass, along with a prescription for their internal density profiles provides a route to constructing a non-linear power spectrum.

The halo model allows the large-scale galaxy distribution to be well-approximated by treating the halo occupation number for galaxies as a function of halo mass, and positioning within each halo a central galaxy around which other galaxies trace the halo profile as satellites (Seljak, 2000; Peacock and Smith, 2000). In the halo model a distribution of spherically-collapsed halo structures randomly populate the linear density field, allowing for the effective separation of power into two distinct source terms,

$$\Delta^2(k) = \Delta^2_{1h}(k) + \Delta^2_{2h}(k). \tag{2.20}$$

The 2-halo term, $\Delta^2_{2h}(k)$, describes the correlations in the distribution of haloes themselves. As this occurs on large scales, an acceptable approximation is to equate this term to the linear matter power spectrum,

$$\Delta^2_{2h}(k) = \Delta^2_{\rm lin}(k). \tag{2.21}$$

[1]This is of course only partially accurate; observed and simulated large scale structure exhibits a complex network of haloes, connecting filaments and voids.

By contrast, the 1-halo term, $\Delta^2_{1h}(k)$, represents the internal halo structure on small scales. Computing this statistic requires averaging the self-convolutions of haloes over the full range of halo masses, weighted by the total number density of pairs of haloes of mass M. In Fourier space these convolutions become simple multiplication operations, leading to the integral

$$\Delta^2_{1h}(k) = \frac{k^3}{2\pi^2}\int_0^\infty \mathrm{d}M\,\frac{M^2 n(M)}{\bar{\rho}^2}u^2(k|M)\,, \tag{2.22}$$

where $n(M)$ is the comoving number density of halos per mass interval $\mathrm{d}M$, known as the halo mass function, and $u(k|M)$ is the halo density profile in Fourier space. Press and Schechter (1974) showed that the mass function is nearly independent of cosmology if expressed in terms of the density peak height, $\nu \equiv \delta_c/\sigma(M)$, given by the spherical collapse overdensity, δ_c, when the density field is smoothed at the mass scale, M. The spherical collapse overdensity marks the extrapolated linear density value at the time that the overdensity of a perturbation diverges according to the non-linear version of the growth equation[2]. Sheth and Tormen (1999) determined a form of the 'universal' mass function,

$$f(\nu) = \frac{M}{\bar{\rho}}n(M)\,\frac{\mathrm{d}M}{\mathrm{d}\nu} \tag{2.23}$$

that fit simulations to be

$$f(\nu) = 0.2162\left[1+\left(0.707\nu^2\right)^{-0.3}\right]\exp\left[-\frac{0.707\nu^2}{2}\right]. \tag{2.24}$$

Assuming spherical symmetry, the density profile in Fourier space can be written as the transform,

$$u(k|M) = \frac{4\pi}{M}\int_0^{r_v} r^2\mathrm{d}r\,\frac{\sin(kr)}{kr}\rho(r,M)\,, \tag{2.25}$$

with the prefactor normalising the profile by halo mass. The virial equilibrium of energy exchange between gravitationally interacting matter shells is a natural threshold at which to truncate the profile. A halo is therefore characterised by

[2]Further details on the derivation of this quantity, and non-linear growth in general are deferred to Chapter 5 which focuses on these topics in the context of coupled dark energy cosmologies.

its virial density, Δ_v, and radius, r_v, which are related by

$$r_v = \left(\frac{3M}{4\pi \bar{\rho} \Delta_v} \right)^{\frac{1}{3}}. \tag{2.26}$$

Spherical collapse calculations in the relevant cosmology inform Δ_v so the virial radius is fixed for a given halo mass. The process of virialisation (in the context of coupled dark energy cosmologies) forms a major focus of Chapter 5, to which a deeper discussion is deferred.

The form of the density profile is typically a matter of fitting to simulations of collisionless dark matter particles. The most common is the NFW profile (Navarro, Frenk, and White, 1997),

$$\rho\left(r, M\right) = \frac{\rho_s}{\left(\frac{r}{r_s}\right)\left[1 + \left(\frac{r}{r_s}\right)\right]^2}. \tag{2.27}$$

The scale radius, r_s, defines a break scale between the linear and cubic declines of density in the inner and outer regions of the halo respectively. This scale also dictates the normalisation factor, ρ_s.

2.3 Weak Gravitational Lensing

A principal drawback of the matter power spectrum as a statistic is that it cannot be directly measured. The advantage of weak gravitational lensing as an observational technique is its capture of the information contained in the three dimensional power spectrum across all scales, via a two-dimensional projection. Photon trajectories propagating from background source galaxies are subject to deflections when passing through the gravitational potential wells of high mass foreground objects. This distorts images of background galaxies, perpendicular to the line of sight, when observed from Earth. These shear effects are related to the distribution of matter on large scales, and so depend on the matter density field and power spectrum. As weak lensing measurements involve two-point real-space angular correlations of shear, the required statistics are the corresponding harmonic coefficients of the convergence power spectrum, itself a function of the 3D matter power spectrum.

This can be constructed (see e.g., Limber, 1953; Kaiser, 1992; Kilbinger, 2015) by first considering a light ray propagating through a perturbed FLRW spacetime. As $\mathrm{d}s^2 = 0$, the time coordinate can be expressed as an integral over the path taken, λ, in proper coordinates,

$$t = \int_{\lambda_A}^{\lambda_B} \mathrm{d}\lambda \, h\left[\mathbf{r}\right] \left| \frac{\mathrm{d}\mathbf{r}}{\mathrm{d}\lambda} \right|, \tag{2.28}$$

where

$$h\left[\mathbf{r}\right] = \frac{1}{c}\left(1 - \frac{2\Phi}{c^2}\right) \tag{2.29}$$

when setting $\Phi = \Psi$. Fermat's principle that $\delta t = 0$ leads to the construction of the Euler-Lagrange equations,

$$h\,\dot{\mathbf{e}} = -\mathbf{e}\left[\left(\nabla_{\mathbf{r}} h\right) \cdot \mathbf{e}\right] + \nabla_{\mathbf{r}} h. \tag{2.30}$$

For a suitable choice of λ, the unit vector tangent to the path is $\mathbf{e} \equiv \dot{\mathbf{r}}$, with $\cdot \equiv \mathrm{d}/\mathrm{d}\lambda$. The right-hand side is the perpendicular component of the gradient of h with respect to the path. We can use the approximation

$$\dot{\mathbf{e}} = \nabla_{\mathbf{r},\perp} \ln h \approx -\frac{2}{c^2} \nabla_{\mathbf{r},\perp} \Phi. \tag{2.31}$$

To first order, the deflection angle between emitted and received light rays is given by integrating this quantity over the trajectory as viewed from Earth:

$$\hat{\boldsymbol{\alpha}} = -\frac{2}{c^2} \int \mathrm{d}r \, \nabla_{\mathbf{r},\perp} \Phi. \tag{2.32}$$

The comoving arc length between two converging light rays separated by angle θ at comoving distance χ to their source,

$$\mathbf{x}_0\left(\chi\right) = f_K\left(\chi\right) \boldsymbol{\Theta}. \tag{2.33}$$

If a deflection is experienced by both rays A and B the separation between them is modified to

$$\mathbf{x}\left(\chi\right) = \mathbf{x}_0\left(\chi\right) - \frac{2}{c^2} \int_0^{\chi} \mathrm{d}\chi' \, f_K\left(\chi - \chi'\right) \left[\nabla_\perp \Phi_A\left(\mathbf{x}, \chi'\right) - \nabla_\perp \Phi_B\left(\chi'\right)\right], \tag{2.34}$$

where χ' is the comoving distance to the deflector. The Jacobian, $\boldsymbol{A} = \partial\boldsymbol{\beta}/\partial\boldsymbol{\theta}$,

maps between the angle measured in the absence of lensing, $\boldsymbol{\beta} = \boldsymbol{x}(\chi)/f_K(\chi)$, and the apparent angle via the lens equation, $\boldsymbol{\beta} = \boldsymbol{\theta} - \boldsymbol{\alpha}$. The deflection, $\boldsymbol{\alpha}$, is given by the final term in equation (2.34) with a factor of $f_K(\chi)$ scaled out. The resulting inverse amplification matrix

$$A = \begin{pmatrix} 1-\kappa-\gamma_1 & -\gamma_2 \\ -\gamma_2 & 1-\kappa+\gamma_1 \end{pmatrix} \tag{2.35}$$

is commonly decomposed into the traceless part as two-component, spin-two shear $\gamma = (\gamma_1, \gamma_2) = \gamma_1 + i\gamma_2 = |\gamma|\exp(2i\varphi)$ and the diagonal part as the scalar convergence κ. Varying the angle, φ, separating the shear components by 2π corresponds to a rotation by π of an elliptical image. The shear components are defined as second derivatives of the 2D Laplacian of the 'lensing potential',

$$\psi = \frac{2}{c^2}\int_0^{\chi} \mathrm{d}\chi' \, \frac{f_K(\chi-\chi')}{f_K(\chi)\, f_K(\chi')} \Phi\left(f_K(\chi')\,\boldsymbol{\theta}, \chi'\right). \tag{2.36}$$

such that

$$\gamma_1 = \frac{1}{2}\left(\partial_1\partial_1 - \partial_2\partial_2\right)\psi, \tag{2.37}$$

$$\gamma_2 = \frac{1}{2}\left(\partial_1\partial_1 - \partial_2\partial_2\right)\psi, \tag{2.38}$$

while as a measure of the isotropic magnification (with anisotropic distortion quantified by shear), the convergence is given by

$$\kappa = \frac{1}{2}\nabla^2\psi. \tag{2.39}$$

As the scattering angle is expected to almost always be small, the Born approximation from scattering theory (Born, 1926) has been applied here, permitting the integration to be taken along the original, undeflected direction so that the zeroth order approximation $\mathbf{x} \approx \mathbf{x}_0$ is used. The above relations, in concert with the Poisson equation, indicate that the convergence at χ is an integral over the large scale matter density perturbation, weighted by the geometry of the configuration of source and deflector with respect to Earth. The average convergence for an entire population of sources requires an additional weighting with respect to the number distribution of galaxies that is taken by e.g., the Euclid space survey

mission (Laureijs et al., 2011; Amendola et al., 2018) to be

$$n(z) \propto z^2 \exp\left[-\left(\frac{z}{z_0}\right)^{\frac{3}{2}}\right], \tag{2.40}$$

where the peak redshift of the distribution is given in terms of the median redshift by $z_0 = z_{\text{mean}}/1.412$. The probability of finding a galaxy within $\mathrm{d}\chi$ of position χ is $\mathrm{d}\chi n(\chi)$ if $n(\chi)$ is normalised. The convergence for a population extending out to some $\chi_{\max}$ defined by the limits of the survey can then be shown to be

$$\kappa(\boldsymbol{\theta}) = \frac{3H_0^2\Omega_m}{2c^2}\int_0^{\chi_{\max}} \mathrm{d}\chi \, \frac{g(\chi)}{a(\chi)} f_K(\chi)\, \delta_m\left(f_K(\chi)\,\boldsymbol{\theta}, \chi\right). \tag{2.41}$$

where

$$g(\chi) = \int_\chi^{\chi_{\max}} \mathrm{d}\chi' \, n(\chi') \, \frac{f_K(\chi - \chi')}{f_K(\chi')} \tag{2.42}$$

represents the efficiency of the lens at distance χ, a total weighting function over the distribution of sources and their relative distance from deflectors.

It is also necessary to address photometric redshift errors. For any particular redshift bin assumed to have an independent source distribution, the galaxy distribution when normalised over the survey is given by

$$n_i(z) = \frac{n(z)\int_{z_{i,-}}^{z_{i,+}} \mathrm{d}z_{ph}\, p_{ph}(z_{ph}|z)}{\int_{z_{\min}}^{z_{\max}} \mathrm{d}z' \, n(z') \int_{z_{i,-}}^{z_{i,+}} \mathrm{d}z_{ph}\, p_{ph}(z_{ph}|z')}, \tag{2.43}$$

where $p_{ph}(z_{ph}|z)$ is the probability of galaxies at z being measured at redshift z_{ph}. The integrals are equivalent to convolving the true distribution with a binning top-hat function, whilst accounting for the probability of galaxies being measured at z_{ph} outside the bin. The denominator is the result of normalising over the bin. This can be shown to be equivalent to the distribution used in Taylor et al. (2007). The frequently employed probability distribution (Ma et al., 2006; Taylor et al., 2007),

$$p_{ph}(z_{ph}|z) = \frac{1}{\sqrt{2\pi}\sigma_z(1+z)} \exp\left\{-\left[\frac{z - z_{ph}}{\sqrt{2}\sigma_z(1+z)}\right]^2\right\}. \tag{2.44}$$

is used throughout this work. For the photometric redshift error, I use the Laureijs et al. (2011) value of $\sigma_z = 0.05$ for Euclid-like surveys.

On large scales, a statistical description for weak lensing is useful. This invites using the two-point correlation function of convergence,

$$\xi_\kappa(\phi) = \langle \kappa(\boldsymbol{\theta})\, \kappa^\star(\boldsymbol{\theta} + \boldsymbol{\phi}) \rangle, \tag{2.45}$$

for which isotropy ensures only the magnitude of the angular separation between pairs, ϕ, determines the average over the sky. This in turn leads directly to the weak lensing angular power spectrum (assuming a flat sky) via a Fourier transform,

$$C_\kappa(\ell) = \int_\Omega \mathrm{d}^2\phi\, \xi_\kappa(\phi)\, \mathrm{e}^{i\boldsymbol{\ell}\cdot\boldsymbol{\phi}}. \tag{2.46}$$

Provided that the flat-sky limit is taken, this power spectrum, in terms of harmonic wavenumber ℓ, is related to the analogous 3D matter spectrum in terms of k by taking the inverse Fourier transform of the matter correlation function within ξ_κ such that

$$\begin{aligned}\xi_\kappa(\phi) = \frac{9}{4}\Omega_m^2 \left(\frac{H_0}{c}\right)^4 \int \mathrm{d}\chi \frac{g(\chi)}{a(\chi)} f_K(\chi) \int \mathrm{d}\tilde{\chi} \frac{g(\tilde{\chi})}{a(\tilde{\chi})} f_K(\tilde{\chi}) \\ \times \int \frac{\mathrm{d}^3 k}{(2\pi)^3} \int \frac{\mathrm{d}^3 \tilde{k}}{(2\pi)^3} \langle \delta_m(\mathbf{k})\, \delta_m^*\left(\tilde{\mathbf{k}}\right) \rangle \\ \times \exp\left[-i\left\{f_K(\chi)\, \boldsymbol{\theta}\cdot\mathbf{k}_\perp + f_K(\tilde{\chi})\,(\boldsymbol{\theta}+\boldsymbol{\phi})\cdot\tilde{\mathbf{k}}_\perp\right\}\right] \exp\left[-i\left\{k_\parallel \chi + \tilde{k}_\parallel \tilde{\chi}\right\}\right],\end{aligned} \tag{2.47}$$

where the wave vector $\mathbf{k}$ has been decomposed into parallel and perpendicular vectors. Evaluating these integrals under the Limber approximation that only modes perpendicular to the line of sight contribute (Limber, 1953), and similarly that additional source-source and source-lens clustering can be neglected, the result is substituted into the expression for the power spectrum to yield

$$C_{ij,\kappa}(\ell) = \frac{9}{4}\Omega_m^2 \left(\frac{H_0}{c}\right)^4 \int_0^{\chi_{\max}} \mathrm{d}\chi\, \frac{g_i(\chi)\, g_j(\chi)}{a^2(\chi)} P_\delta\left(k = \frac{\ell}{f_K(\chi)}, \chi\right), \tag{2.48}$$

with the remaining subscripts i, j denoting different tomographic redshift bins, which are treated as having independent source probability distributions. For surveys probing a finite number of bins, this distinction is useful and leads to an array of auto- and cross- convergence power spectra.

2.4 Weak Lensing: Observational Considerations and Biases

Accurate cosmological inference through weak lensing depends significantly on the observational approach to measuring galaxy shapes. The practical procedures for generating and processing this image data encounter a number of sources of bias that must be carefully accounted and calibrated for.

In the initial stage, images of galaxies are observed through by either ground-based telescopes such as LSST (LSST Science Collaboration et al., 2009) or space-based telescopes such as Euclid (Laureijs et al., 2011). Generally, multiple images are taken with CCDs for each area of sky covered by a weak lensing survey. This presents an opportunity to process either individual images or a single stacked image for each area. An immediate source of bias that must be addressed is that the images of galaxies can be sheared by the instrument as well as by the intervening matter distribution. It is also important to recognise that images may correspond to stars as well as galaxies. This typically requires a prescribed algorithm for identifying and selecting the galaxy images for data processing. Following the necessary corrections to the image, the values for the underlying shape parameters can be determined.

The main challenge facing observational weak lensing becomes apparent when one considers the range of sources of bias that affect galaxy images. One highly prominent example is the need to account for the point-spread function (PSF). This describes the degree to which an imaging system blurs a point source. On one level this is a problem simply because already faint galaxies will be blurred such that the lensing of their image is significantly obscure. If there are anisotropies within the PSF effect on shape images (as one would expect) then there is a substantial risk that the effect could be degenerate with the weak lensing signal itself by inducing its own coherent alignment of shapes (see e.g., Jee and Tyson, 2011). Calibrating for this bias requires a robust model for the PSF as it varies in amplitude and anisotropy over time and position on the observed sky. These changes can be complicated to model as they can depend sensitively on the configuration of the instrumentation. For example, as discussed in Jee and Tyson (2011), LSST has a small effective focal ratio (to accommodate a large field of

view) and is sensitive to focal-plane alignment and height deviations across the CCDs positioned on this plane will have significant effects on the PSF as it is partly a function of variations across and within CCDs.

There are various efforts that have been proposed to model the PSF. For example, a principal component analysis can be applied to the data itself, which provides the foundation for PSF modelling. This approach is used in e.g., Bertin (2011). An alternative option is to use the approach of 'shapelets', in which image objects are linearly decomposed into a series of basis functions for different shapes, as is the method used by KiDS (Refregier, 2003; Hildebrandt et al., 2017).

While the PSF has been discussed in some detail, there are further biases that must be calibrated for in the observing process. These include Poisson pixel noise, the pixelisation of images by detectors, the blending of source galaxies with other galaxies or stars via physical interaction or projection, and a range of other effects (see e.g., Huterer et al., 2006; Massey et al., 2013; Taylor and Kitching, 2016). A large reservoir of data will become available from Stage IV surveys which can mitigate the impact of large statistical errors for shear measurements but does not address the general problem of measurement bias as one of the main obstacles to cosmological parameter inference. Individually these biases present modelling challenges but it is also not obvious how bias effects may be correlated with each other, or indeed if one could expect measurements of the bias to itself be biased. It is therefore of increasingly pressing need for surveys like Euclid and LSST to develop generalised robust theoretical frameworks for assessing the impact of multiple sources of bias on parameter inference from cosmic shear.

Heymans et al. (2006) proposed to model bias generically with a linear model such that it is decomposed into multiplicative, m, and additive, c, parts via

$$\hat{g} = (1 + m)\, g + c + \delta_g, \tag{2.49}$$

where $\hat{g}$ represents the measured value of the 'true' reduced shear, g, and δ_g is a residual noise such that $\langle \delta_g \rangle = 0$. Massey et al. (2013) found that if the uncertainty on the multiplicative and additive biases exceeds the limits $\sigma_m \lesssim 2 \times 10^{-3}$ and $\sigma_c \lesssim 2 \times 10^{-4}$ then dark energy error requirements for Stage IV surveys cannot be met. Therefore, it is important to either devise minimal bias measurement

methods or develop schemes to accurately correct for bias. This will inevitably require careful calibration approach to simulated data. However, these calibration approaches will be limited to some degree; indeed Gillis and Taylor (2019) show that residual biases and increased errors are a typical feature of current calibration methods.

Given the range and diversity of bias effect that must be accounted for, and the lack of knowledge about their underlying relationships to the bias they induce, it will be necessary to develop accurate approximation schemes for an arbitrary bias. These should be optimised to accommodate e.g., the breadth of functional forms that could plausibly describe the relationship between the size of an effect and the resulting bias; different magnitudes of bias; different spacing and numbers of bins sampling the bias; and different weighting functions for the effect size.

Estimators of bias can in principle be used to correct shear measurements for bias and reproduce the true shear. The most straightforward approach is to replace the true biases in equation (2.49) with measured values, $\hat{m} = m + \delta_m$ and $\hat{c} + \delta_c$, and rearrange so that a correction to the shear is given by

$$\hat{g}' = \frac{\hat{g} - \hat{c}}{1 + \hat{m}}. \tag{2.50}$$

Note that the error, δ_g, intrinsic to the measurement in equation (2.49) is accounted for by the measurement in the bias. This correction is used by e.g., Fenech Conti et al. (2017). For $m \longrightarrow -1$, this approach is problematic because it involves dividing by a term that both contains noise and that approaches zero. This former will result in non-Gaussian distributions for the corrected shear and the latter will entail divergent behaviour. The benefits of any statistical analysis using this correction will therefore be limited in these cases. Gillis and Taylor (2019) apply a Taylor expansion to equation (2.50) that improves the statistical qualities of the correction at the expense of making it slightly more biased:

$$\hat{g}' = (\hat{g} - \hat{c}) \left(1 - \hat{m} + \hat{m}^2\right). \tag{2.51}$$

The authors show that the expected biases on the correction are

$$\langle m' \rangle = \sigma^2 [m] (1 + \hat{m}) + \hat{m}^3 \tag{2.52}$$

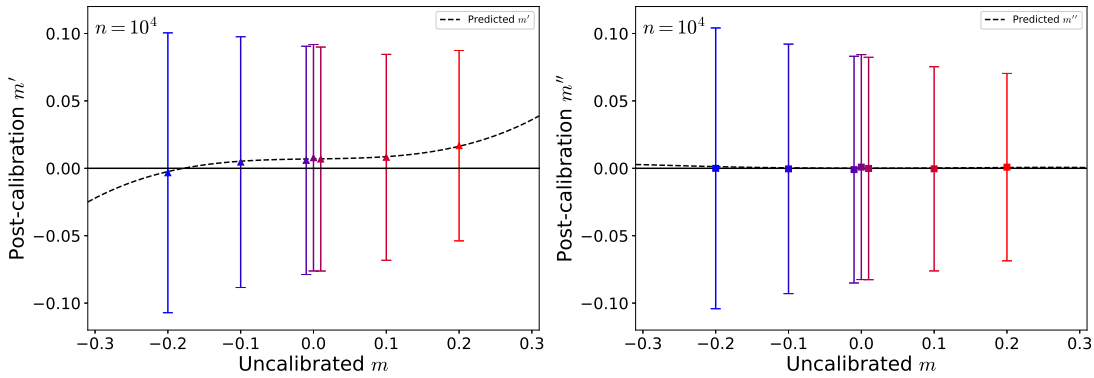

Figure 2.1: Post-calibration bias from first (left panel) and second (right panel) order corrections applied to 10^4 realisations of datasets containing 10^4 mock shear values. The uncalibrated bias values measured are $m = \{-0.2, -0.1, -0.01, 0., 0.01, 0.1, 0.2\}$ while c is set to zero. The markers denote values averaged over the set of realisations, and the error bars represent the scatter due to noise in the dataset used to determine the correction. The dashed curves show the predicted values for the correction, given by equation (2.52) and equation (2.55) for the first and second order corrections respectively.

and

$$\langle c' \rangle = 0. \tag{2.53}$$

They note that the presence of the $\hat{m}^3$ term fundamentally limits the bias on the correction, even for an infinite sample size of shears that could make the $\sigma^2[m]$ term negligible. One solution to this is to correct iteratively. To this end Gillis and Taylor (2019) provide a second order correction,

$$\hat{g}'' = \hat{g}' \left(1 - \sigma^2[m](1 - 2\hat{m}) - \hat{m}^3\right). \tag{2.54}$$

This is determined by applying another Taylor expansion and retaining terms to first order in $\hat{m}'$. The measured value of the corrected bias is incorporated by writing $\hat{m} = m + \delta_m$ in equation (2.52) and taking the expectation value. The bias on the second order correction is given by

$$\langle m'' \rangle \approx \sigma^2[m]\left(\sigma^2[m] + 3\hat{m}^2\right) - m^6. \tag{2.55}$$

One can follow the approach of Gillis and Taylor (2019) to numerically evaluating the post-calibration bias using the first and second order corrections on samples of 10^4 simulated shear values with intrinsic scatter and errors given above. A true bias is applied to this data set and the measured values of bias are determined by

performing a linear regression. The correction is then applied, and another linear regression performed to determine the measured values of the corrected bias and the error on the correction. The correction, using the measured bias values, is then applied to a data set with zero error to determine the post-calibration residual bias. This process is performed for $N = 10^4$ realisations which are averaged over to provide the final results. In Figure 2.1, these results are shown for several values of multiplicative bias (and setting $c = 0$ as it is far more straightforward to correct for additive bias) for the first and second order corrections. This is essentially a reproduction of Figure 1 in Gillis and Taylor (2019).

2.5 Spectroscopic galaxy clustering

The clustering of galaxies provides an additional probe through the anisotropic galaxy power spectrum, $P_{gg}(k, \mu; z)$. For the Euclid-like survey forecasts of later Chapters, the specific observable is Hα-emitting galaxies. The approach to constructing $P_{gg}(k, \mu; z)$ outlined here closely follows e.g., Ballinger et al. (1996) and Seo and Eisenstein (2003).

Sources of apparent anisotropy make it necessary to include the dependence on the cosine of the wave mode and radial vector, μ, in the power spectrum. These include redshift-space distortions arising from a peculiar velocity component in observed redshifts (Kaiser, 1987). As a result, the observed strength of galaxy clustering will be sensitive to the orientation of the pair-separation vector. To account for this, the Kaiser factor is incorporated into $P_{gg}(k, \mu; z)$. The underlying dark matter distribution is biased by the galaxies that trace it (see e.g., Desjacques et al., 2018), so the galaxy power spectrum can then be constructed from the matter power spectrum according to

$$P_{gg}(k, \mu; z) = \left(b\sigma_8(z) + f\sigma_8(z)\mu^2\right)^2 \frac{P_m(k; z)}{\sigma_8^2(z)}, \tag{2.56}$$

where $b(z) = \sqrt{1+z}$ is the effective bias of the sample of Hα-emitters used in Amendola et al. (2018) and $f(z)$ is the logarithmic growth rate of structure, which relates the peculiar velocity to the density. The redshift-dependence of the (linear) matter power spectrum is removed by factoring out $\sigma_8(z)$, which leads to $b\sigma_8(z)$ and $f\sigma_8(z)$ becoming free parameters in galaxy clustering forecasts

(White et al., 2009).

Converting measured angles and redshifts into radial and transverse distances requires assumptions about cosmology. If these are incorrect, distortions arise through the AP effect. The ratios,

$$\alpha_{\perp}(z) \equiv \frac{D_A(z)}{D_{A,\mathrm{fid}}(z)} \tag{2.57}$$

and

$$\alpha_{\parallel}(z) \equiv \frac{H_{\mathrm{fid}}(z)}{H(z)}, \tag{2.58}$$

of the assumed and true angular diameter distance and Hubble parameter (Alcock and Paczynski, 1979) change wavenumbers and angles from their 'fiducial' values to (Ballinger et al., 1996)

$$k(k_{\mathrm{fid}}, \mu_{\mathrm{fid}}; z) = \frac{k_{\mathrm{fid}}}{\alpha_{\perp}(z)} \sqrt{1 + \mu_{\mathrm{fid}}^2 \left(\frac{\alpha_{\perp}^2(z)}{\alpha_{\parallel}^2(z)} - 1 \right)} \tag{2.59}$$

and

$$\mu(\mu_{\mathrm{fid}}; z) = \mu_{\mathrm{fid}} \frac{\alpha_{\perp}(z)}{\alpha_{\parallel}(z)} \frac{1}{\sqrt{1 + \mu_{\mathrm{fid}}^2 \left(\frac{\alpha_{\perp}^2(z)}{\alpha_{\parallel}^2(z)} - 1 \right)}}. \tag{2.60}$$

Under these projections the matter power spectrum transforms as

$$P_m(k_{\mathrm{fid}}, \mu_{\mathrm{fid}}; z) \longrightarrow \frac{1}{\alpha_{\perp}^2 \alpha_{\parallel}} P_m(k, \mu; z). \tag{2.61}$$

Hereafter, unless explicitly stated, k and μ may be assumed to account for the AP effect.

Finally, redshift uncertainties along the line-of-sight are accounted for via the factor (e.g., Wang et al., 2013),

$$F_z(k; z) = \exp\left[-k^2 \mu^2 \sigma_r^2(z) \right], \tag{2.62}$$

where

$$\sigma_r(z) = \frac{c}{H(z)} (1 + z) \sigma_{z=0}, \tag{2.63}$$

with the spectroscopic redshift error given by $\sigma_{z=0} = 0.001$ (Laureijs et al., 2011). The full observed linear galaxy power spectrum is then

$$P_{gg}(k,\mu;z) = \frac{1}{\alpha_\perp^2 \alpha_\parallel} \left(b\sigma_8(z) + f\sigma_8(z)\mu^2\right)^2 \frac{P_m(k;z)}{\sigma_8^2(z)} F_z(k;z) + P_{\rm shot}(z)\,. \tag{2.64}$$

The final term represents any residual shot noise that is not accounted for by Poisson sampling the underlying CDM density field with Hα-emitting galaxies. A fiducial value of $P_{\rm shot,fid} = 0$ for all z is typically assigned when performing parameter forecasts for surveys. Non-linear contributions, such as the Finger-of-God effect (Hamilton, 1998) or the damping of the BAO component of $P(k;z)$ (Eisenstein et al., 2007; Seo and Eisenstein, 2007), are also important but the galaxy clustering presented in this thesis is restricted to linear scales. Discussions of these features in the references provided are therefore recommended for the interested reader.

2.6 Summary

In this Chapter the formalism underpinning the main probes of large scale structure has been introduced. The growth of structure in the linear regime has been discussed in detail, leading to an outline of the linear matter power spectrum. This statistic for the density field on large scales provides the foundation for the probes that are used in this thesis to forecast constraints on cosmology. Having defined the linear matter power, one of the most powerful frameworks for constructing a model of power on non-linear scales has been explored. This approach to describing smaller scales through a distribution of spherical halo structures is also central to one of the core focuses of this thesis. Namely, how modifications of the halo model to capture the effects of baryonic astrophysics can be used for large scale structure survey forecast analyses assessing the impact of such a systematic on constraining e.g., dark energy or the neutrino mass sum. The specific probes that this thesis uses for these analyses - weak gravitational lensing and spectroscopic galaxy clustering - have also been introduced. There is a wider and deeper literature available for both but the foundational information relevant to understanding the following Chapters has been provided here.

Chapter 3

The Impact of Baryons On The Sensitivity of Dark Energy Measurements

This Chapter provides an examination of the degradation to Stage IV dark energy forecasts that results from marginalising over parameters introduced to correct for the effects of baryonic phenomena on large-scale structure. Percent-level precision is required by next generation surveys for estimates of the dark energy equation of state parameters. The detailed physics of baryon effects is not yet well-understood and, as such, baryons represent a major systematic for surveys. Therefore, it is important this problem is addressed with an appropriate depth and thoroughness. The work presented here is an expansion of that provided in Copeland et al. (2018).

I detail analytic modifications to the halo model that have been adopted to capture the impact of adiabatic contraction of baryons and feedback on the matter power spectrum. A generalisation is also provided of the Navarro-Frenk-White profile to account for a possible inner halo core, potentially induced by baryons or, potentially, an exotic form of dark matter.

The Fisher information formalism underlying my statistical analysis will be outlined. This will provide the basis for understanding a full Fisher analysis quantifying the impact of marginalising over the baryon parameters on forecasts. A discussion of the capacity of Stage IV surveys to constrain the baryonic parame-

ters themselves is also included.

An exploration of the scales at which baryons and dark energy contribute most to the Fisher information is presented to determine whether probing deeper into the non-linear regime can effectively reduce baryon degradation. Other avenues for potential mitigation will be examined, including the improvement that can be expected from adding baryon information from external sources. The strength of priors required to limit degradation to the percent level will be evaluated. The alternative approach of adding information through the cosmological sector in the form of cosmic microwave background priors provided by the Planck 2015 data release will be subject to a detailed study.

The issue of model bias will also be addressed. The level of calibration of baryon modelling required to limit bias in dark energy forecasts to within an acceptable threshold for Stage IV surveys will be quantified.

The layout of this Chapter is as follows. Sections 3.1 and 3.2 will provide a primer on the range of astrophysical baryonic phenomenology influencing large scale structure, and previous studies of its impact on parameter forecasts. Sections 3.3 and 3.4 will introduce the formalism of the baryon halo model, and its extension to incorporate inner halo cores. The methodology underpinning Fisher forecasts will be detailed in Sections 3.5 and 3.6. The results of this process will be analysed in Sections 3.7 and 3.8 for dark energy, standard cosmological and baryon parameter forecasts. Section 3.9 is concerned with mitigation approaches and Section 3.10 explores the issue of baryon model bias in the results.

3.1 Baryonic Astrophysical Phenomena

Here I introduce the main baryonic phenomena dealt with in this Chapter. The purpose is to provide the reader with sufficient background to contextualise the approaches to modelling their effects on halo structure that are discussed in following sections.

3.1.1 Adiabatic Contraction and Radiative Cooling

Gas undergoing radiative cooling leads to clustering on small scales. This strengthens the local gravitational potential which in turn results in an infall of dark matter. The concentration of haloes consequently increases, leading to an enhancement of the matter power spectrum. This becomes the dominant baryonic effect on power for scales $k \gtrsim 10\,h\,\mathrm{Mpc}^{-1}$. The adiabatic contraction model has been used to partially explain the behaviour of the dark matter distribution since Eggen et al. (1962) used adiabatic invariant particle trajectories to predict the dynamics of a proto-galaxy undergoing contraction (Gnedin et al., 2004).

Assuming particles traversing circular orbits in a system with spherical symmetry, a constant fractional rate of contraction and the conservation of angular momentum (such that $M\,(r)\,r = \text{const.}$), the ratio of initial and final orbital radii, r_0 and r_1,

$$\frac{r_1}{r_0} = \frac{M_{DM}\,(r_0) + M_{\mathrm{gas}}\,(r_0)}{M_{DM}\,(r_1) + M_{\mathrm{gas}}\,(r_1)} \tag{3.1}$$

is given in terms of the mass of dark matter, M_{DM}, and gas, M_{gas}, enclosed by them (Blumenthal et al., 1986). It was shown by Gnedin et al. (2004) that this model was not sufficiently sophisticated to accurately match the results of hydrodynamic simulations, which indicated elliptical orbits. Gnedin et al. (2004) relaxed the circular orbit assumption and introduced orbit-averaged trajectories described by a power law, $\langle \bar{r} \rangle \sim r^n$. This has since been refined further, for example by Gustafsson (2006) who highlight the need to account for more complex 'baryonic pinching' effects like stellar feedback on contracting haloes.

Another interesting effect is the increased sphericity of the mass distribution. The gravitational potential well generated by the central baryon region is itself more spherical than for the matter distribution as a whole. This leads to the or-

bits of dark matter particles and condensed baryonic objects like stars becoming more circular in turn (Kazantzidis et al., 2004; Debattista et al., 2008).

3.1.2 AGN and Supernovae Feedback

Active galactic nuclei refer to extremely luminous, small scale regions at the centres of galaxies. The most widely accepted mechanism explaining them is the infall of gas from an accretion disk onto a central supermassive black hole. The cold matter that initially forms the disk is heated from friction as gravitational forces contract the disk and increase its angular momentum. Atoms in nearby cold material are excited by the radiation of the disk and re-radiate an emission line spectrum that provides a signature of the AGN. Two different modes of energy output characterise AGN: the kinetic mode is associated with highly directional relativistic jets of ionised matter from the accretion disk itself, while the radiative mode describes an isotropic transfer of thermal energy to the surrounding environment (Alexander and Hickox, 2012).

Supernovae exert a similar influence on the matter distribution by heating gas in their surroundings. The gravitational core collapse of very massive stars when their ability to sustain thermal pressure through nuclear fusion is exhausted leads to the explosive release of vast quantities of matter and radiation (Burrows, 2013). An alternative mechanism that could act as a supernova progenitor is a white dwarf accreting helium or generating runaway carbon fusion that disrupts the entire body (Nomoto, 1982). While AGN are single localised sources of thermal energy production, the effect of supernovae is a cumulative one with many sources within a single dark matter structure injecting energy into the environment.

The baryonic feedback effect associated with AGN and supernovae[1] has a significant impact on the matter distribution on megaparsec scales. Dark matter haloes can experience a bloating effect up to and beyond their virial radii. Consequently, baryonic feedback has been shown to have an effect of suppressing the matter power spectrum by $\sim 30\%$ in the nonlinear regime (Semboloni et al.,

[1]It is not uncommon for authors to refer to any and all baryonic phenomena impacting large scale structure as 'baryonic feedback' effects. However, this work will use the term strictly to refer to the types of thermal feedback effects associated with AGN and supernovae, drawing a clear distinction with e.g., adiabatic contraction.

2011). This becomes sub-dominant with respect to adiabatic contraction on the smallest scales but for $k \sim 5\,h/\mathrm{Mpc}$ (scales that large scale structure surveys are sensitive to), it is a highly prominent feature.

3.1.3 Star Formation Rates and Miscellaneous Phenomena

A number of additional astrophysical effects have been shown (e.g., van Daalen et al., 2011) to have $> 1\%$ level impacts an the matter power spectrum at relevant scales. Adiabatic contraction and baryonic feedback are the principal phenomena, but can themselves give rise to further secondary effects. For example, stellar distributions are dependent on local supernova activity, with stronger feedback reducing the star formation rate which can boost the dark matter power spectrum.

Other effects that should be accounted for include the relationship between the metallicity of cooling gas distributions and the adiabatic contraction induced by their radiative cooling. Careful modelling of stellar winds generated in and propagated through the environments surrounding supernovae is also important.

3.2 Introduction To Baryon Impact Studies

The impact of the phenomena described above on the large scale matter distribution has prompted a number of works recalibrating the halo model to match N-body and hydrodynamic simulations implementing baryonic effects (see e.g., Semboloni et al., 2011; Zentner et al., 2013; Semboloni et al., 2013; Mohammed et al., 2014; Mead et al., 2015; Schneider and Teyssier, 2015). A robust baryon-halo model in principle would allow one to examine the sensitivity of the matter power spectrum and its weak lensing projection to the parameters controlling the strength of the baryonic effects on halo structure. In turn this allows for large scale structure survey forecasts of cosmological parameters to be made that account for uncertainty in the baryon parameters.

Different approaches for implementing baryons have generated a considerable range of results for the impact on dark energy error forecasts. Among the most

alarming are those of Zentner et al. (2013) who find, by modifying the halo concentration relation, that 1-σ errors for w_0 and w_a increase by $\sim 50\%$. These degradations would compound to severely reduce the Figure of Merit, given by

$$\mathrm{FOM} = \frac{1}{\sqrt{(\sigma_{w_0}\sigma_{w_a})^2 - \sigma^2_{w_0 w_a}}}, \tag{3.2}$$

that constrains the w_0-w_a parameter space.

3.2.1 Simulations

Previous efforts to capture baryonic physics using N-body simulations have largely been replaced by more recent hydrodynamical simulations, which are able to incorporate more sophisticated prescriptions for the mechanical and fluid dynamics of the baryon sector. One particularly comprehensive project are the OverWhelmingly Large Simulations (OWLS; Schaye et al., 2010). OWLS is capable of implementing prescriptions for multiple sources of baryon physics. In Mead et al. (2015), a halo model variant called HMCODE was calibrated by making fits to a fiducial dark-matter only model (DMONLY), and then to three baryonic models. The first introduces radiative cooling and supernovae feedback along with prescriptions for stellar formation and evolution (REF). These features are also included in the model designed primarily to account for AGN feedback (simply referred to as AGN), and in a model that uses a stellar initial mass function to produce more supernovae as well as introducing a stronger effect from stellar winds (DBLIM).

Mead et al. (2015) also make a fit to non-linear power spectra generated by an 'emulator' code (COSMIC EMU) for the high resolution N-body simulations from the Coyote Universe project (Heitmann et al., 2009, 2010; Lawrence et al., 2010; Heitmann et al., 2014). These power spectra were used as the basis for the HMCODE claim of $\simeq 5$ percent accuracy for scales $k \leq 10\,h/\mathrm{Mpc}$ and redshifts $z \leq 2$ (improving by several factors over HALOFIT at non-linear scales). This motivates the decision taken in this work to use these fits to COSMIC EMU within a version of HMCODE that will be further modified as needed.

3.2.2 Modified Halo Models

In their halo model variant Semboloni et al. (2011) determine a bias as high as 40% in predictions of w_0 when neglecting baryons which reduces to $\sim 10\%$ when accounting for feedback by fitting mass fractions for separate profiles for dark matter, gas and stars. Semboloni et al. (2013) argue that this overcomes a shortcoming in the modelling of Zentner et al. (2013) that neglects differences between the distributions of dark matter and hot gas. Mohammed et al. (2014) adopt a similar approach, by modelling stellar contributions with a central galaxy, introducing a hot plasma in hydrostatic equilibrium and accounting for the baryon-induced adiabatic contraction of dark matter due to cooling. They find a degradation of $\sim 10\%$ and $\sim 30\%$ to the forecasted errors on w_0 and w_a respectively.

The approach of these works (see also Fedeli, 2014; Fedeli et al., 2014) of fitting for stellar and gas physics within the halo is different to that advocated by Mead et al. (2015) (hereafter M15). Their corrections to the halo model through HMCODE (extended in Mead et al., 2016) are designed specifically to calibrate the power spectrum accurately for the non-linear regime. This requires empirically motivated modifications to internal halo structure relations that can be directly attributed to the effects of adiabatic contraction and feedback. As the power spectrum is the statistic underpinning my study of forecast degradations, I adopt the model of M15 as the most suitable for the purposes of this work.

In this Chapter, HMCODE is extended to include inner halo cores. This is motivated by e.g., Martizzi et al. (2012) and Governato et al. (2012) who show that baryons can produce inner cores of the order of 10 kpc. Possible mechanisms range from dynamical friction effects in black hole orbit decays to AGN feedback removing dark matter from central halo regions by disturbing the gravitational potential. Alternatively, axions could be responsible in the form of solitons (Marsh and Pop, 2015). Whatever the underlying physics, my version of baryon-halo model will be the basis for a more comprehensive and robust analysis of the baryon impact on dark energy constraints than has been done previously.

This approach also allows for an exploration of how well the baryon parameters used in the modified HMCODE could be constrained by Stage IV surveys, given the uncertainty in cosmology. It should be stressed that this does not amount

to direct constraints on baryonic phenomena, but on the approximate redistribution of matter in haloes. This could provide opportunities for informing future hydrodynamic simulations, and insight into the possible nature of inner cores.

3.2.3 Impact Forecasts

To understand the scope of baryon impacts it is important to analyse the sensitivity of information from baryons and dark energy at different lensing scales. This goes beyond examining the effect of increasing the scale limit, $\ell_{\rm max}$, of an analysis (e.g., Semboloni et al., 2011; Zentner et al., 2013; Mohammed et al., 2014). Weak lensing power responses to varying w_0 and w_a exhibit subtle scale-dependencies due to competing influences on the growth of structure and the geometry governing distances. Understanding how the w_0-w_a degeneracy is broken in this interpretation, and how this is complicated by the inclusion of baryons, is essential for informing a mitigation strategy.

The prescription presented here is based on improvements offered by optimising $\ell_{\rm max}$, incorporating external baryon information, and adding *Planck* CMB priors. The last step makes use of strong constraints that have not been available for previous baryon impact studies. At the same time the advent of next generation surveys is fast approaching so now is an optimal moment to revisit the issue. The accuracy of power spectra provided by HMCODE and the scope of effects available through the inclusion of inner cores make this work uniquely placed to assess baryon degradation and target how to mitigate it.

3.3 Baryon-Halo Model

To incorporate baryons into the halo model I adopt and extend the treatment of M15. Three general but distinct baryonic effects are parameterised: large-scale adiabatic contraction caused by radiative cooling; high-impact energy transfer from localised sources; and the formation of inner halo cores with radius r_b due to small-scale physics. M15 capture the first two by varying internal halo structure relations through their parameters, A_B (referred to as A in M15) and η_0. An inner core is discussed in § 3.4.

Parameter	Fiducial value
Ω_m	0.3156
Ω_b	0.04917
h	0.6727
σ_8	0.831
n_s	0.9645

Table 3.1: Fiducial cosmological parameters used for the weak lensing survey forecasts presented in this work.

It is worth noting that M15 identify a degeneracy between A_B and η_0 from fitting to multiple OWLS simulations. A likelihood analysis by Hildebrandt et al. (2017) exploits this by fixing $\eta_0 = 1.03 - 0.11A_B$, where A_B becomes the single free baryon parameter. However, this work retains both parameters as independent terms. Having different baryon effects characterised by multiple parameters allows for the opportunity to explore how well surveys could constrain these particular phenomena, given the uncertainty in cosmology. It will also make it easier to identify whether degeneracies between baryon and cosmological parameters are primarily due to a specific baryonic effect.

Throughout this Chapter, the parameter values corresponding to the base ΛCDM *Planck* TT,TE,EE+lowP likelihood (see Table 4 in Planck Collaboration et al., 2016b) are chosen for the fiducial cosmology, along with the fiducial dark energy parameters, $w_{0,\text{fid}} = -1$, $w_{a,\text{fid}} = 0$. These values are provided in Table 3.1.

3.3.1 Modelling Adiabatic Contraction

Adiabatic contraction is the most straightforward baryon effect to model. Clustering of baryonic matter due to radiative cooling induces the gravitational infall of dark matter, so the total matter distribution undergoes contraction. Simulations have shown that the impact is at several percent for non-linear clustering (Duffy et al., 2010; Gnedin et al., 2011).

In M15 this is captured by modifying the concentration factor, $c\,(M, z)$, which relates the scale and virial radii via $r_s = cr_v$. The amplitude, A_B, in the concen-

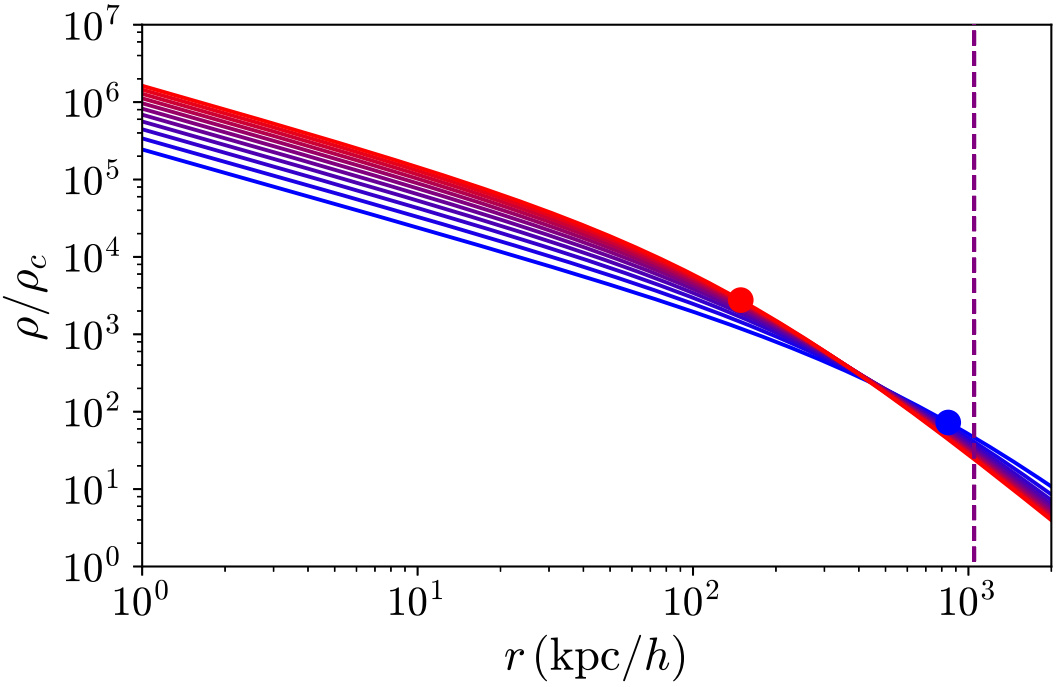

Figure 3.1: NFW halo density profiles (in units of the critical density, ρ_c) for a halo of mass $M = 3 \times 10^{14}\, M_\odot$ at $z = 0$ with different amplitudes, A_B, of the concentration factor. Blue (red) curves correspond to the lowest (highest) values in the range $0.5\, A_B \leq A_{B,\mathrm{fid}} \leq 1.5\, A_{B,\mathrm{fid}}$. Blue (red) dots indicate the scale radius, $r_s = 507$ (169) h^{-1}kpc, corresponding to the lowest (highest) A_B values. The purple dashed line marks the virial radius, $r_v = 1050\, h^{-1}$kpc, at which the profile must be truncated.

tration factor,

$$c\,(M, z) = A_B \frac{1 + z_f}{1 + z}, \tag{3.3}$$

is allowed to vary around a fiducial value, $A_{B,\mathrm{fid}} = 3.13$. This was chosen because M15 found it produced the best fit to COSMIC EMU power spectra. Fits by M15 to the different OWLS simulations satisfy the range $2 < A_B < 4$, which could be used to inform a prior. The dependence on halo mass enters the above expression via z_f, the formation redshift at which a fraction $f = 0.01$ of the total matter in a density fluctuation has collapsed. Figure 3.1 shows the impact of varying the concentration amplitude on the NFW density profile of a halo with fixed virial radius and mass. Adiabatic contraction then manifests through a reduced scale radius. This corresponds to suppressing the halo density on large scales, $r \gtrsim 10^2\, h^{-1}$kpc, while enhancing it at smaller scales, $r \lesssim 10\, h^{-1}$kpc.

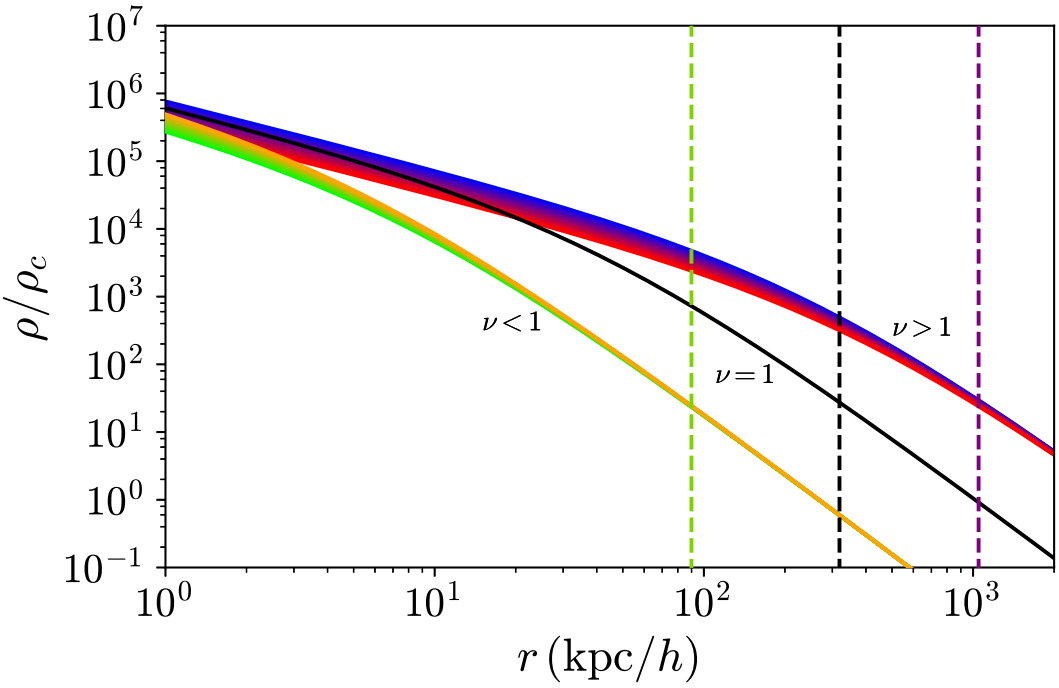

Figure 3.2: NFW halo density profiles (in units of the critical density, ρ_c) for haloes of mass $M = 2\times10^{11}\,M_\odot$ ($\nu < 1$) and $M = 3\times10^{14}\,M_\odot$ ($\nu > 1$) at $z = 0$. The feedback parameter is varied in the range $0.5\,\eta_{0,\mathrm{fid}} \leq \eta_0 \leq 1.5\,\eta_{0,\mathrm{fid}}$. For the low mass halo, low values of η are represented by green curves and high values by orange curves. For the high mass halo, low values of η are represented by blue curves and high values by red curves. The virial radii, $r_v = 89.9\,h^{-1}$kpc and $r_v = 1050\,h^{-1}$kpc, for the low and high mass halo are marked respectively by the green and purple dashed lines. I also include a reference profile (black) for a halo of mass $M = 10^{12}\,h^{-1}M_\odot$ corresponding to $\nu = 1$, and mark its virial radius $r_v = 319\,h^{-1}$kpc.

3.3.2 Modelling Feedback

Feedback mechanisms are more complex to model than adiabatic contraction. Objects like supernovae and AGN release large quantities of energy into their environments, heating gas which expands to virial radius scales (Pontzen and Governato, 2012; Lagos et al., 2013). A large range of scales from subparsecs up to megaparsecs are influenced, so the effect cannot be easily modelled analytically (Schaye et al., 2010; van Daalen et al., 2011; Martizzi et al., 2014). The impact on halo structure is also dependent on mass. Simulations have shown (e.g., Pontzen and Governato, 2012; Teyssier et al., 2013; Martizzi et al., 2013) that similar mechanisms can describe the expulsion of gas from the central regions of both lower and higher mass haloes by AGN. In the former case haloes are subjected to stronger expulsions, resulting in the loss of substantial baryonic matter. Larger

mass haloes are not so devastated by violent feedback, merely bloating outwards as heated gas expands through the structure.

M15 accounts for the scale and mass dependence of feedback by transforming the scale of the halo window function according to

$$u\left(k|M\right) \longrightarrow u\left(\nu^{\eta}k|M\right), \tag{3.4}$$

where

$$\nu \equiv \frac{\delta_c}{\sigma\left(M\right)} \tag{3.5}$$

is the ratio of the collapse overdensity to the standard deviation of the density field, smoothed over a mass scale, M. By performing the rescaling $r \longrightarrow \nu^{\eta}\, r$ in equation (2.25) that in real space the modified density profile takes the form

$$\rho\left(r, M\right) = \frac{\rho_s}{\nu^{3\eta}} \frac{1}{\left(\frac{r}{\nu^{\eta} r_s}\right)\left[1 + \left(\frac{r}{\nu^{\eta} r_s}\right)\right]^2}, \qquad r \leq \nu^{\eta} r_v. \tag{3.6}$$

As shown in Figure 3.2, more positive values of η increasingly bloat higher mass haloes (characterised by $\nu > 1$) while lower mass haloes ($\nu < 1$) are left relatively reduced by gas being fully expelled. A non-zero value of η was also required to make empirical corrections to the halo bloating to ensure accurate power spectra when fitting to dark-matter-only simulations.

When fitting power spectra to COSMIC EMU simulations it was found (see Table 2 in M15) that a number of parameters required redshift-dependent modifications. This includes η which is decomposed into a constant, η_0, that controls the degree of feedback, and a fixed dependence on $\sigma_8\left(z\right)$ such that

$$\eta = \eta_0 - 0.3\,\sigma_8\left(z\right). \tag{3.7}$$

A fiducial value of $\eta_{0,\mathrm{fid}} = 0.603$ is used here, as this is determined by M15 to best fit COSMIC EMU spectra. Fits to OWLS simulations lie within the range $0.5 \leq \eta_0 \leq 0.8$, which could potentially be used to define a prior.

3.4 Inner Cores

3.4.1 Inner Core Mechanisms

A long-running debate about the nature of the inner most region of the halo motivates incorporating inner cores into the baryon-halo model. The cusp-core problem arises from a discrepancy between N-body simulations that predict the divergent $\rho \propto r^{-1}$ NFW cusp in halo centres (see Dubinski and Carlberg, 1991; Navarro et al., 1997), and observations like dwarf galaxy rotation curves that indicate constant density cores of the order of a few kpc (Kuzio de Naray et al., 2008; Walker and Peñarrubia, 2011; Oh et al., 2011). Explanations for this fall into three categories:

1. Simulations are systematically neglecting some aspect of structure formation.

2. Replacing traditional CDM with e.g., fuzzy dark matter in the form of ultra-light axions (alternatively, self-interacting dark matter or warm dark matter) can generate cores of a few kpc (see e.g., Marsh and Pop, 2015; Zhang et al., 2016).

3. Baryonic processes flatten cusps into cores (de Blok et al., 2003; Pontzen and Governato, 2012; Peñarrubia et al., 2012).

A prominent proposal (Pontzen and Governato, 2012) motivated by baryon arguments is that supernovae feedback transfers sudden, repeated bursts of energy to surrounding gas, causing oscillations in the central gravitational potential. This in turn induces rapid orbits of dark matter particles that flatten the cusp (see also Read and Gilmore, 2005; Governato et al., 2012). However, the core physics is ultimately of limited relevance here compared to understanding the impact that the phenomenology has on large-scale structure probes.

3.4.2 Incorporating an Inner Core Into The Halo Model

There are various ways to introduce cores analytically within an NFW-like profile (see e.g., Einasto, 1965; Zhao, 1996; Navarro et al., 2004). The simplest extension, of the form

$$\rho\left(r\right)=\frac{\rho_{N}}{\left(\frac{r+r_{b}}{r_{s}}\right)\left(1+\frac{r}{r_{s}}\right)^{2}}, \tag{3.8}$$

is sufficient for this work, which is more focused on an exploratory study of generic core impacts than in rigorously implementing more sophisticated inner core modelling to haloes. This formalism has also been employed by Peñarrubia et al. (2012), although here the baryon-induced core radius is denoted as r_b. Setting r_b to zero reduces the profile to normal NFW. To retain the advantages of a semi-analytical halo model it is useful for a modified profile to have an analytic Fourier transform. Numerically computing the 1-halo term of the power spectrum given in equation (2.22) requires tabulated values of the integrand. This is much less computationally expensive if there is an analytic expression for the window function. Inevitably, condition restricts the number of viable profiles, with a general trend towards simpler extensions being more likely to meet this condition. The choice of profile here is therefore partly motivated by this practical concern.

The Fourier transform of equation (3.8) is derived as follows. In general,

$$u\left(k|M\right)=\frac{4\pi}{M}\int_{0}^{r_{v}}r^{2}\mathrm{d}r\,\frac{\sin\left(kr\right)}{kr}\rho\left(r,M\right). \tag{3.9}$$

The well-known transform of the NFW profile is given by Cooray and Sheth (2002) as

$$u_{NFW}\left(k|M_{NFW}\right)=\frac{4\pi\rho_{s}r_{s}^{3}}{M_{NFW}}\left\{F\left(k,c\right)\cos\left(kr_{s}\right)+G\left(k,c\right)\sin\left(kr_{s}\right)-\frac{\sin\left(ckr_{s}\right)}{kr_{s}\left(1+c\right)}\right\}, \tag{3.10}$$

where

$$\begin{aligned}F\left(k,c\right)&\equiv\mathrm{Ci}\left(\frac{kr_{v}}{c}\left(1+c\right)\right)-\mathrm{Ci}\left(\frac{kr_{v}}{c}\right)\\G\left(k,c\right)&\equiv\mathrm{Si}\left(\frac{kr_{v}}{c}\left(1+c\right)\right)-\mathrm{Si}\left(\frac{kr_{v}}{c}\right)\end{aligned} \tag{3.11}$$

and the scale radius, r_s, is defined in terms of the concentration, c, and virial radius, r_v, such that $r_v = c\, r_s$. Integrating the profile up to the virial radius defines the halo mass,

$$M \equiv 4\pi \int_0^{r_v} r^2 \rho\,(r, M)\ \mathrm{d}r, \tag{3.12}$$

which is evaluated and expressed as a function of the concentration factor,

$$M_{NFW}\,(c) = 4\pi\rho_s r_s^3 \left\{ \ln\,(1+c) - \frac{c}{1+c} \right\}. \tag{3.13}$$

Retaining an analytic Fourier transform when incorporating an inner core is part of the motivation for the simple modification made in equation (3.8). The resulting window function is

$$\begin{aligned} u\,(k|M) = \frac{4\pi\rho_s r_s^3}{M}\frac{b}{b-c}\Bigg\{ & \frac{M_{NFW}}{4\pi\rho_s r_s^3} u_{NFW}\,(k|M_{NFW}) \\ & + \frac{c}{b-c}\frac{1}{kr_s}\Big(\Big[\, G\,(k,c)\cos\,(kr_s) - F\,(k,c)\sin\,(kr_s) \Big] \\ & \qquad - \Big[G\,(k,b)\cos\,(kr_b) - F\,(k,b)\sin\,(kr_b)\ \Big] \Big) \Bigg\}, \end{aligned} \tag{3.14}$$

where $b \equiv r_v/r_b$ defines an effective 'baryon concentration factor'. The halo mass can then be determined as a function of both concentration factors, such that

$$M\,(b,c) = \frac{4\pi\rho_s r_s^3}{(b-c)^2}\left\{ b\,(b-2c)\,\frac{M_{NFW}}{4\pi\rho_s r_s^3} + c^2\left[\ln\,(1+b) - \frac{b}{1+c} \right] \right\}. \tag{3.15}$$

By taking the limit $r_b \to 0$ for both u and M, the NFW case is recovered. Despite the simplicity of the r_b modification, it generates a significantly more complex window function. More sophisticated formulations of an inner core typically require numerical Fourier transformations, so they are far less practical for the purposes of this work.

It is also possible to introduce a halo mass dependence, for example by allowing $r_b \propto r_s$. However, this would entail a dependence in turn on the halo concentration and therefore a possible degeneracy with A_B. This avenue could be further explored in future work but in this Chapter r_b is assumed to be determined by some combination of processes largely independent from specific halo properties. It should be emphasised that without introducing a more robust physical motivation in the modelling the accuracy of results will be limited. However, the main

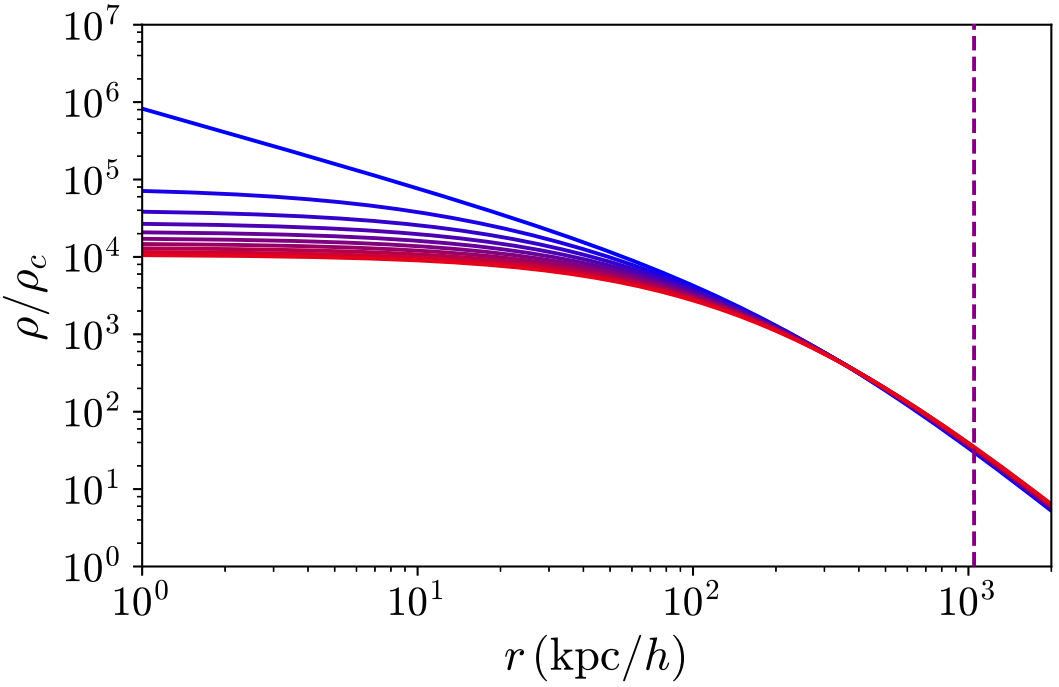

Figure 3.3: Modified NFW profile for a halo of mass $M = 3 \times 10^{14}\, M_\odot$ featuring an inner core radius, r_b, in the range $r_b = 0 - 100\, h^{-1}$kpc at $z = 0$. Blue curves represent profiles closer to standard NFW while red curves indicate prominent baryon cores. The purple dashed line marks the virial radius, $r_v = 1050\, h^{-1}$kpc.

concern of this work is including the generic feature of an inner core to explore its impact on the halo profile and matter power spectrum. This can provide useful insight into the potential consequences of marginalising over uncertainty in r_b, but should still be treated as a broad, first-pass implementation.

In Figure 3.3 the modified density profile is plotted for a range of cores up to $r_b = 100\, h^{-1}$kpc to emphasise the deviation from an NFW profile at small scales. Once the scale is reduced to r_b the density turns off from the NFW branch and becomes constant. The profile is therefore increasingly suppressed by larger inner cores.

3.5 Bayesian Statistics

A unique feature of cosmology as a scientific endeavour is that there is just one experiment available. Conclusions can be inferred only from observations made

of the single realisation of the evolution of the Universe that is apparent. Important properties of this data set can be described by cosmological models that are defined by a set of parameters. Determining the accuracy to which the values of these parameters can be known first requires an appropriate and useful interpretation of probability.

The frequentist interpretation is rooted in the notion that the probability of an event is given by the relative frequency of its occurrence during a large number of repetitions of an experiment. Underlying physical causes are treated as essentially random, either through some intrinsic property of the system or as deterministic phenomena requiring too much information to directly predict. If this approach were to be applied to cosmology the probability attached to a parameter value would have to be determined by the distribution of the instances of its measurement across an infinite number of experiments. This is clearly not viable, for the reasons stated above, but is also problematic on a more fundamental level. Even if a large but finite number of experiments are performed in place of an ill-defined infinite limit an additional uncertainty arises on the probability itself, which will be calculated at slightly different values for different sets of experiments. The objective probability that should supposedly emerge from this approach is instead compromised by a circular definition of frequency probability itself.

An alternative approach, the propensity interpretation, invokes objective probabilities that emerge directly from the physical properties of the system. This has the advantage of allowing probabilities to be assigned to single events. However, while this approach may prove useful when considering particle states in the Copehagen interpretation of quantum mechanics it is not viable for cosmology because intrinsic probabilities do not feature in reasonable cosmological models.

A preferable approach would be one that naturally accounts for the limitations of data. Bayesian probability provides this through a subjective approach that leads to assertions about the ‘degree of belief’ in a result, or the probability that a measurement is correct given the available data. The greatest advantage of this is that results using the same data (and applying the same prior confidence in a measurement) will be consistent. Concerns about the subjectivity of conclusions are mitigated by noting that in cosmology different data sets are ultimately subsets of a single data set provided by the Universe.

Bayes' theorem,

$$P\left(\Theta|D,M\right)=\frac{P\left(D|\Theta,M\right)P\left(\Theta|M\right)}{P\left(D|M\right)} \tag{3.16}$$

relates the 'posterior' probability, $P\left(\Theta|D,M\right)$, that a set of parameter values, Θ, are the true values given the data, D, and assuming a model, M, to the 'likelihood', $P\left(D|\Theta,M\right)$, of obtaining the data for that model with those parameter values. All previous information known about Θ is contained in the 'prior', $P\left(D|M\right)$. The 'evidence', $P\left(D|M\right)$, represents the probability of obtaining the data given the model. As this quantity is independent of the parameter values, it is more convenient, for a flat (i.e., constant) prior, to write the theorem as

$$P\left(\Theta|D,M\right)\propto L\left(D|\Theta,M\right), \tag{3.17}$$

when comparing likelihoods.

3.6 Likelihood Estimates and The Fisher Matrix Formalism

One can evaluate uncertainties for a set of parameter values $\Theta=(\theta_1,...,\theta_N)$ from a data set $\mathbf{x}$ by using the Fisher matrix formalism (Fisher, 1935; Tegmark et al., 1997). This provides a basis for quantifying the constraining power of an experiment, or in the case of this work, the capacity of a large scale structure survey to measure cosmological parameters to within a certain accuracy. As will be detailed in this section, there is a robust statistical underpinning for the Fisher formalism. It could be argued that it falls short of methods like generating mock galaxy catalogues or by running Monte Carlo experiments, which use random sampling algorithms to iteratively determine confidence forecasts for parameters. While the Fisher methodology lacks computational superiority this does allow it to retain an advantage of being quick and practical to execute. Moreover, it is far more plausible to extract the source of important features in its results. The effects of different parameters on probes, like the power spectrum, entering the Fisher matrix, can be studied and so it is often possible to trace similar effects through each step of the calculation to explain degenerate 2-parameter confidence

regions in the final forecasts. As the focus of this work is to perform these types of diagnosis in depth, the Fisher formalism is useful.

In the most general sense, the power of the Fisher formalism is that no data is required to produce forecasts. The essential requirement is that the expected data from an experiment satisfies a multivariate Gaussian distribution. In the case of correlated data, the Fisher formalism still applies. At this point the only information required is a knowledge of the sensitivity of the mean of the data and the data covariance to the underlying parameters. The Fisher matrix can then be constructed prior to the relevant experiment itself being performed. The forecast errors that become available through the Fisher matrix provides a quantitative assessment for how effective an experiment will be and therefore how best to optimise it. It is, however, important to note that the Fisher formalism cannot be applied in the case of data that is parameter-dependent. This can arise in certain circumstances where the data is treated as a construction of some underlying measurement, and hence there may be further parameters that need to be carefully considered. The appropriate course of action is then to assume a fiducial set of parameters that are not changed with the data.

It has been shown (Tegmark et al., 1997) that the inverse of the parameter covariance matrix is given by the expectation value of the curvature of the log-likelihood function $\ln L\left(\mathbf{x}|\Theta\right)$ around its maximum, occurring at the set of fiducial, most likely parameter values, Θ_{fid}. This quantity is the Fisher information matrix,

$$F_{\alpha\beta} \equiv \Big\langle \frac{\partial^2 \mathcal{L}}{\partial\theta_\alpha \partial\theta_\beta} \Big\rangle, \tag{3.18}$$

in which the definition $\mathcal{L} = -\ln L$ is a useful convention. The probability distribution of observables is treated as Gaussian such that

$$L = \frac{1}{\sqrt{\left(2\pi\right)^N \det \mathbf{C}}} \exp\left[-\frac{1}{2}\left(\mathbf{x}-\boldsymbol{\mu}\right)\mathbf{C}^{-1}\left(\mathbf{x}-\boldsymbol{\mu}\right)^T\right], \tag{3.19}$$

where $\mathbf{C}$ is the data covariance matrix and $\boldsymbol{\mu} = \langle \mathbf{x} \rangle$ is the mean data. Assuming equal priors on Θ and $\mathbf{x}$, the relation $P\left(\mathbf{x}|\Theta\right) = P\left(\Theta|\mathbf{x}\right)$ holds. A Taylor

expansion of the log-likelihood around $\Theta_{\rm fid}$ is

$$\mathcal{L} = \mathcal{L}|_{\Theta_{\rm fid}} + \frac{\partial \mathcal{L}}{\partial \theta_\alpha}\Big|_{\Theta_{\rm fid}} (\theta_\alpha - \theta_{\alpha,{\rm fid}}) + \frac{1}{2} (\theta_\alpha - \theta_{\alpha,{\rm fid}}) \frac{\partial^2 \mathcal{L}}{\partial \theta_\alpha \partial \theta_\beta}\Big|_{\Theta_{\rm fid}} (\theta_\beta - \theta_{\beta,{\rm fid}}) + O\left((\theta - \theta_{\rm fid})^3\right), \quad (3.20)$$

where there is an implicit sum over α and β. The first order term vanishes because it is evaluated at the peak of the log-likelihood. The expansion series is truncated after second order because modern cosmological experiments are expected to yield small errors on Θ which allow third order terms to be treated as negligible. The exponential of the remaining terms has the form of a Gaussian distribution. Hence, it is now possible to relate the data covariance matrix to the Hessian matrix,

$$T_{\alpha\beta} = \frac{\partial^2 \mathcal{L}}{\partial \theta_\alpha \partial \theta_\beta}, \quad (3.21)$$

the expectation value of which defines the Fisher information. An involved derivation invoking various matrix identities allows Tegmark et al. (1997) to show that the Fisher matrix can be written as

$$F_{\alpha\beta} = \frac{1}{2} {\rm Tr}\left[\mathbf{C}^{-1}\mathbf{C}_{,\alpha}\mathbf{C}^{-1}\mathbf{C}_{,\beta} + \mathbf{C}^{-1}M_{\alpha\beta}\right]. \quad (3.22)$$

The final term $M_{\alpha\beta} = \boldsymbol{\mu}_{,\alpha}\boldsymbol{\mu}_{,\beta}^T + \boldsymbol{\mu}_{,\beta}\boldsymbol{\mu}_{,\alpha}^T$ is the expectation value of the second derivative of the data matrix $(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T$ under Gaussian conditions. Derivatives with respect to parameters are denoted by $,\alpha \equiv \partial/\partial\alpha$. For weak lensing one can approach the Fisher matrix in one of two ways. One possibility is treating the shear field as the data. The underlying field can be approximated as a Gaussian here and the mean of the data can be approximated to be zero. This leads to the second term in the above equation disappearing. Alternatively, the mean can be defined as the lensing power spectrum directly. In this case, the covariance is effectively the four-point function of the underlying shear field. In this work, the power spectrum is treated as the data signal and is assumed to be Gaussian. The applicability of the Gaussian approximation requires more careful consideration here. The power spectrum is a quantity that is quadratic in the underlying field. While the shear field may be treated as Gaussian, this statistical property is lost when two random Gaussian variables are combined as a product. However, the central limit theorem can be applied here whereby the sum of a

sufficiently large number of independent random variables tends towards a Gaussian distribution even if the the variables are not themselves normally distributed.

In this scenario of treating the power as the data, the covariance should be parameter-independent and so the first term in the equation above vanishes. The corresponding Fisher matrix can then ultimately be expressed as (Tegmark et al., 1997; Takada and Jain, 2004)

$$F_{\alpha\beta} = \frac{1}{2} f_{\rm sky} \sum_{\ell} (2\ell + 1) \sum_{(ij)} \sum_{(pq)} C^{ij}_{\ell,\alpha} C^{pq}_{\ell,\beta} \left[{\rm Cov}^{-1} \right]_{\ell,(ij),(pq)}, \tag{3.23}$$

in which the spherical harmonic ℓ and m modes are summed over, and $f_{\rm sky}$ is the fraction of sky accessible to the survey. The auto- and cross-correlations of observed power in redshift bins $i, j, p, q = (1, ..., N_{\rm bin})$ are captured by the covariance matrix,

$${\rm Cov}_{\ell,(ij),(pq)} = \hat{C}^{ip}_{\ell} \hat{C}^{jq}_{\ell} + \hat{C}^{iq}_{\ell} \hat{C}^{jp}_{\ell}. \tag{3.24}$$

Contributions from different modes are treated as separable so that the matrix is block diagonal in ℓ. The full observed power spectrum is constructed by adding the shape noise, $\sigma_e = 0.3$, when averaged over the number of galaxies to the auto-correlations of power within each bin. Provided there is no intrinsic alignment to account for, this is given by

$$\hat{C}_{\ell,ij} = C_{\ell,ij} + \frac{\sigma_e^2}{n_i} \delta_{ij}, \tag{3.25}$$

where δ_{ij} is the Kronecker delta and n_i is the number density of galaxies in redshift bin i. This is computed by integrating over the redshift distribution of the galaxy number density, for which the common parameterisation,

$$n\left(z\right) \propto z^2 \exp\left[- \left(\frac{z}{z_0} \right)^{\frac{3}{2}} \right], \tag{3.26}$$

is used, where the peak density depends on the median redshift of the survey through $z_0 = z_{\rm med}/\sqrt{2}$.

A 2-parameter confidence region is finally determined by inverting the Fisher matrix to marginalise over the other parameters, and extracting the resulting 1-σ

errors, $\sigma_{\alpha\alpha} = \sqrt{[F^{-1}]_{\alpha\alpha}}$. The Cramer-Rao inequality,

$$\sigma_\alpha \geq \frac{1}{\sqrt{F_{\alpha\alpha}}}, \tag{3.27}$$

sets a lower limit, known as the 'conditional error' on a parameter. This is achievable only if all other parameters are known with absolute accuracy. Otherwise, the inversion of the Fisher matrix incorporates uncertainty on the other parameters into the error estimate to provide the 'marginal error'. The amount of degeneracy between parameters is related to the correlation coefficient (see e.g., Coe, 2009),

$$\rho \equiv \frac{[F^{-1}]_{\alpha\beta}}{\sqrt{[F^{-1}]_{\alpha\alpha}\,[F^{-1}]_{\beta\beta}}}. \tag{3.28}$$

It is clear that ρ lies within the range $-1 \leq \rho \leq 1$, with $\rho = -1$ representing total anti-correlation, $\rho = 1$ representing total correlation and $\rho = 0$ corresponding to completely independent posterior probability distributions. The parameters involved may or may not have a physical dependence on each other.

Simple geometric arguments can be used to show that the dimensions of a confidence ellipse derive from the errors and correlation coefficient via the following expressions. The semi-major and semi-minor axis lengths are given by

$$a^2 = \frac{\sigma_\alpha^2 + \sigma_\beta^2}{2} + \sqrt{\left(\frac{\sigma_\alpha^2 - \sigma_\beta^2}{2}\right)^2 + (\rho\sigma_\alpha\sigma_\beta)^2} \tag{3.29}$$

$$b^2 = \frac{\sigma_\alpha^2 + \sigma_\beta^2}{2} - \sqrt{\left(\frac{\sigma_\alpha^2 - \sigma_\beta^2}{2}\right)^2 + (\rho\sigma_\alpha\sigma_\beta)^2}, \tag{3.30}$$

while the orientation is set by the angle,

$$\phi = \frac{1}{2}\arctan\left[\frac{2\rho\sigma_\alpha\sigma_\beta}{\sigma_\alpha^2 - \sigma_\beta^2}\right]. \tag{3.31}$$

The confidence level required for the ellipse dictates a factor by which the lengths, a and b, are multiplied, with a 1-σ 68.3% confidence threshold corresponding to a factor of 1.52 (Lampton et al., 1976). The inverse of the area (Coe, 2009),

$$A = \pi\sigma_\alpha\sigma_\beta\sqrt{1-\rho^2} \tag{3.32}$$

determines the Figure of Merit statistic that quantifies the strength of constraints on a 2-parameter region like w_0-w_a (Albrecht et al., 2006).

3.7 Impacts of Baryons On Dark Energy Forecasts

3.7.1 Matter Power Spectra Responses

The sensitivity of the cosmological probe to baryons and cosmology can be examined by evaluating power spectra responses to varying each parameter with respect to its fiducial value while fixing the other parameters. Responses for $\Delta^2(k)$ for fixed σ_8 at redshift $z = 0$ are shown in Figure 3.4, and for C_ℓ in Figure 3.6.

The most subtle response to baryons is for η_0, where averaging over the scale- and mass-dependent influences produces a peak response. The bloating impact on higher mass haloes dominates over the reduction effect on lower mass ones, so lower values of η_0 overall enhance $\Delta^2(k)$. The peak is also a function of the evolution of halo populations, occurring for smaller scales at earlier times. For r_b, I plot the ratio of power for cores up to $100\,h^{-1}$kpc with respect to the case of a cusp, $r_b = 0\,h^{-1}$kpc. As expected, larger inner cores increasingly damp $\Delta^2(k)$ on small scales. Adiabatic contraction produces the opposite effect. Increasing A_B boosts $\Delta^2(k)$ on non-linear scales, corresponding to enhanced density profiles in this regime.

Figure 3.4 indicates that w_0 and w_a are degenerate with A_B and r_b at $z = 0$. For power spectra normalised to σ_8, only the non-linear influence of structure affects $P(k)$ at this redshift. At earlier times the degeneracy is broken because the A_B and r_b impacts are largely redshift-independent while increasing w_0 and w_a uniformly enhances $\Delta^2(k)$ over all scales for any $z > 0$ (e.g., see Figure 3.5 for $\Delta^2(k)$ responses at $z = 0.5$). This is because, for less negative values of w, dark energy becomes energetically relevant earlier. By $z = 0$ there has been more acceleration and therefore greater suppression of structure growth, so $\Delta^2(k)$ is boosted for fixed σ_8.

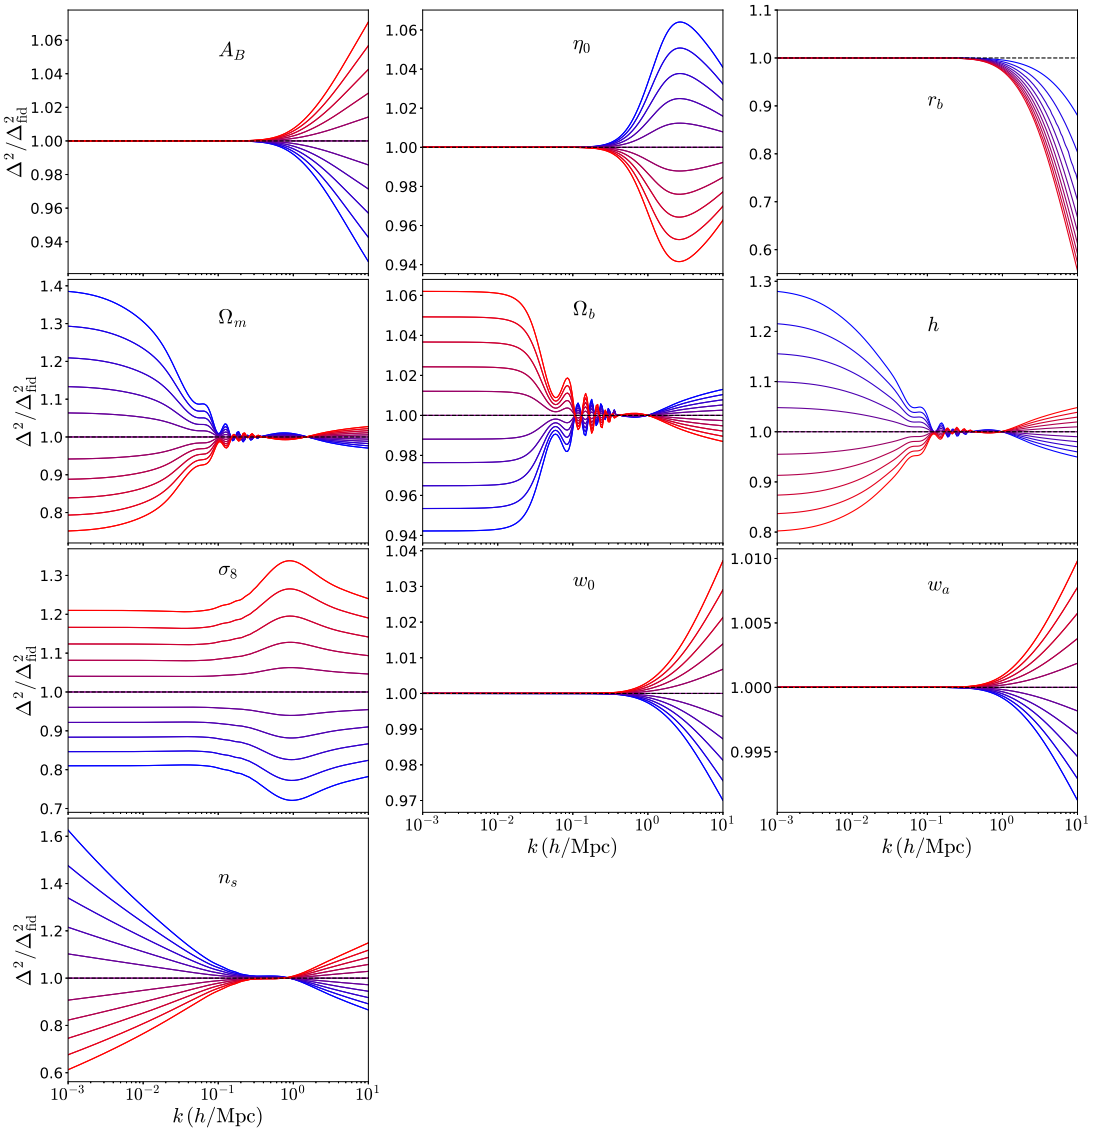

Figure 3.4: The ratio of matter power spectra at $z = 0$ for different iterations of parameters in $\Theta = (A_B, \eta_0, \Omega_m, \Omega_b, h, \sigma_8, n_s, w_0, w_a, r_b)$, with respect to a fiducial power spectrum computed with parameter values found by Planck Collaboration et al. (2016b). Bluer (redder) curves correspond to lower (higher) values for parameters in the range $0.9\,\Theta_{\rm fid} < \Theta < 1.1\,\Theta_{\rm fid}$, except in the case of the dynamic dark energy parameter which varies between $-0.1 < w_a < 0.1$, and r_b which is varied between core sizes of $r_b = 0 - 100\,h^{-1}$kpc and plotted with respect to the fiducial $r_b = 0\,h^{-1}$kpc.

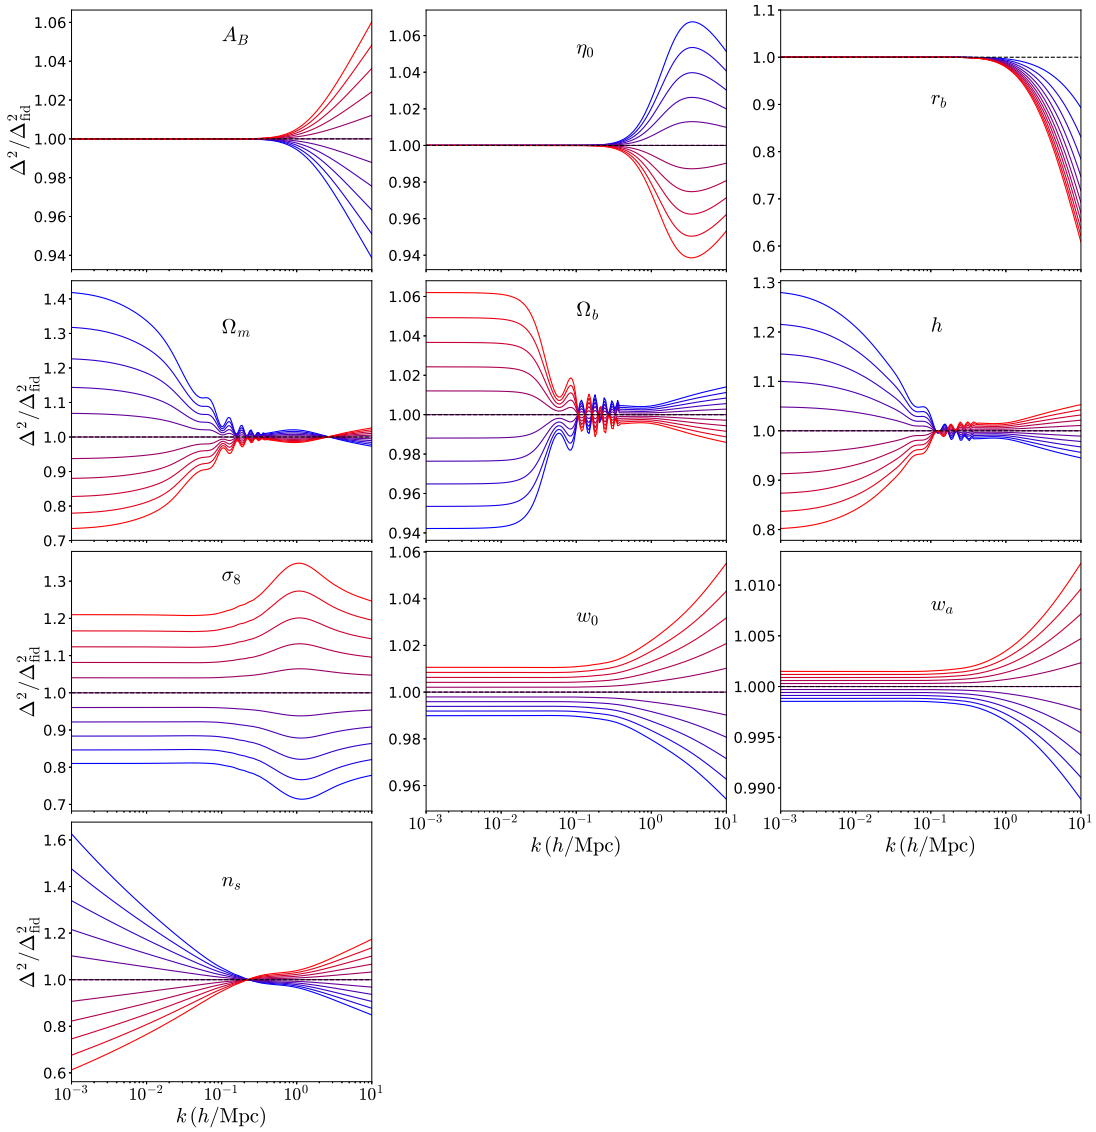

Figure 3.5: The ratio of matter power spectra at $z = 0.5$ for different iterations of parameters in $\Theta = (A_B, \eta_0, \Omega_m, \Omega_b, h, \sigma_8, n_s, w_0, w_a, r_b)$, with respect to a fiducial power spectrum given by *Planck* parameters. Bluer (redder) curves correspond to lower (higher) values for parameters in the range $0.9\,\Theta_{\rm fid} < \Theta < 1.1\,\Theta_{\rm fid}$, except in the case of the dynamic dark energy parameter which varies between $-0.1 < w_a < 0.1$, and r_b which is varied between core sizes of $r_b = 0 - 100\,h^{-1}$kpc and plotted with respect to the fiducial $r_b = 0\,h^{-1}$kpc.

3.7.2 Weak Lensing Convergence Power Spectra Responses

As weak lensing incorporates information from sources and foregrounds over a range of redshifts it is important to be aware of the changes in linear power over the growth history. Integrals of $\Delta^2(k)$ along the line of sight average over these redshift-dependent effects. Varying w_0 and w_a now induces the opposite response for C_ℓ (see Figure 3.6) than for $\Delta^2(k)$. Notably there is also a broad peak on large scales ($\ell \sim 100$). This is because dark energy influences cosmological distances and therefore rescales the lensing weight functions. More negative w increases this geometric contribution to the lensing signal (Huterer, 2002), boosting C_ℓ on all scales. On non-linear scales this is damped by the opposite influence from the growth of structure, which enters through $\Delta^2(k)$. For the smallest ℓ, linear k can only be accessed at larger distances and therefore earlier times, when more positive values of w boost $\Delta^2(k)$. This also has a mitigating effect on the influence of geometry, leading to the broad response peak. The results presented here are consistent with those of Zorrilla Matilla et al. (2017), who thoroughly examine the competing effects of geometry and growth on the sensitivity of lensing observables to Ω_m and w.

3.7.3 Forecast Results

Fisher matrices using the parameter set $(A_B, \eta_0, r_b, \Omega_m, \Omega_b, h, n_s, \sigma_8, w_0, w_a)$ are constructed in accordance with the specifications of a Euclid-like survey. In Table 3.2, the survey parameters given by the Euclid survey report (Laureijs et al., 2011) are summarised. $N_{\rm bin} = 10$ redshift bins are chosen in the range $0 < z < 2$ such that each bin contains an equal number density of galaxies,

$$\mathcal{N} = \frac{1}{N_{\rm bin}} \int_0^{z_{\rm max}} \mathrm{d}z\, n(z)\,, \tag{3.33}$$

where $n(z)$ is the redshift distribution of the number density. A large range of scales from $\ell_{\rm min} = 10$ to $\ell_{\rm max} = 5000$ are covered. Hence, in practice it is prudent to compute the summation in equation (3.23) at logarithmic intervals.

The full set of 2-parameter 1-σ confidence ellipses from this analysis for every parameter combination in $\Theta = (A_B, \eta_0, \Omega_m, \Omega_b, h, \sigma_8, n_s, w_0, w_a, r_b)$ is presented in

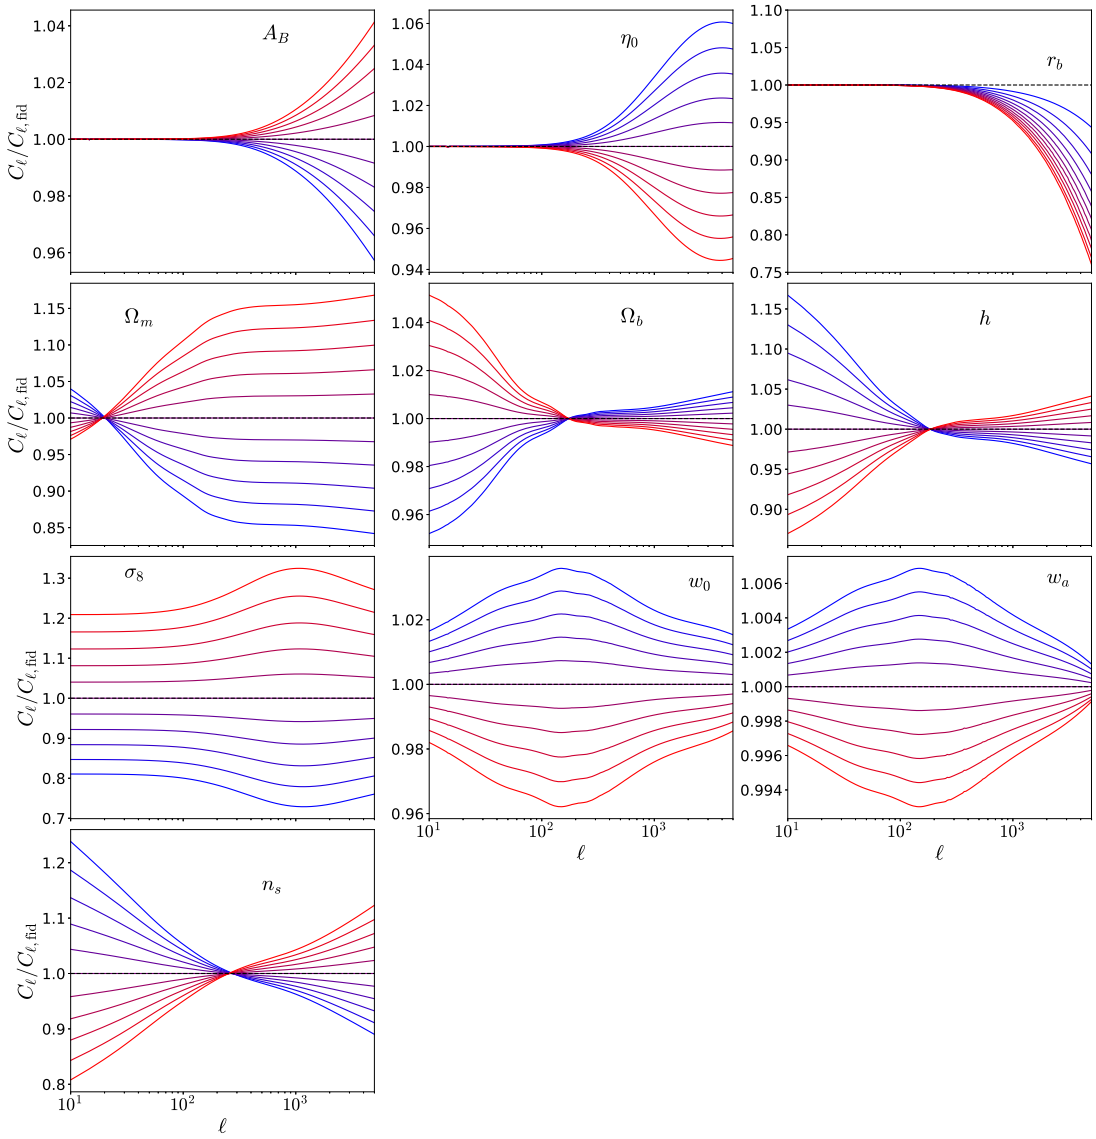

Figure 3.6: The ratio of weak lensing convergence power spectra in a $0.9 < z < 1.1$ redshift bin for different parameters in $\Theta = (A_B, \eta_0, \Omega_m, \Omega_b, h, \sigma_8, n_s, w_0, w_a, r_b)$, with respect to a fiducial power spectrum computed with Planck Collaboration et al. (2016b) parameter values. Bluer (redder) curves correspond to lower (higher) values for parameters in the range $0.9\,\Theta_{\rm fid} < \Theta < 1.1\,\Theta_{\rm fid}$, except in the case of the dynamic dark energy parameter which varies between $-0.1 < w_a < 0.1$, and r_b which is varied between core sizes of $r_b = 0 - 100\,h^{-1}$kpc and plotted with respect to the fiducial $r_b = 0\,h^{-1}$kpc.

Parameter	Euclid value
$A_{\rm sky}$	$15,000\,{\rm deg}^2$
$z_{\rm min}$	0.
$z_{\rm max}$	2.0
$z_{\rm med}$	0.9
$N_{\rm bin}$	10
$n_{\rm gal}$	30 gal/arcmin^2
σ_z	0.05
σ_e	0.3
$l_{\rm min}$	10
$l_{\rm max}$	5000

Table 3.2: Survey parameters for a Euclid-like space mission, including the area $A_{\rm sky}$ of sky probed, the redshift range and median redshift value $z_{\rm med}$, the number of redshift bins $N_{\rm bin}$, the number density of galaxies surveyed, $n_{\rm gal}$, the photometric error σ_z, the intrinsic ellipticity σ_e, and the range of accessible harmonic wavenumbers.

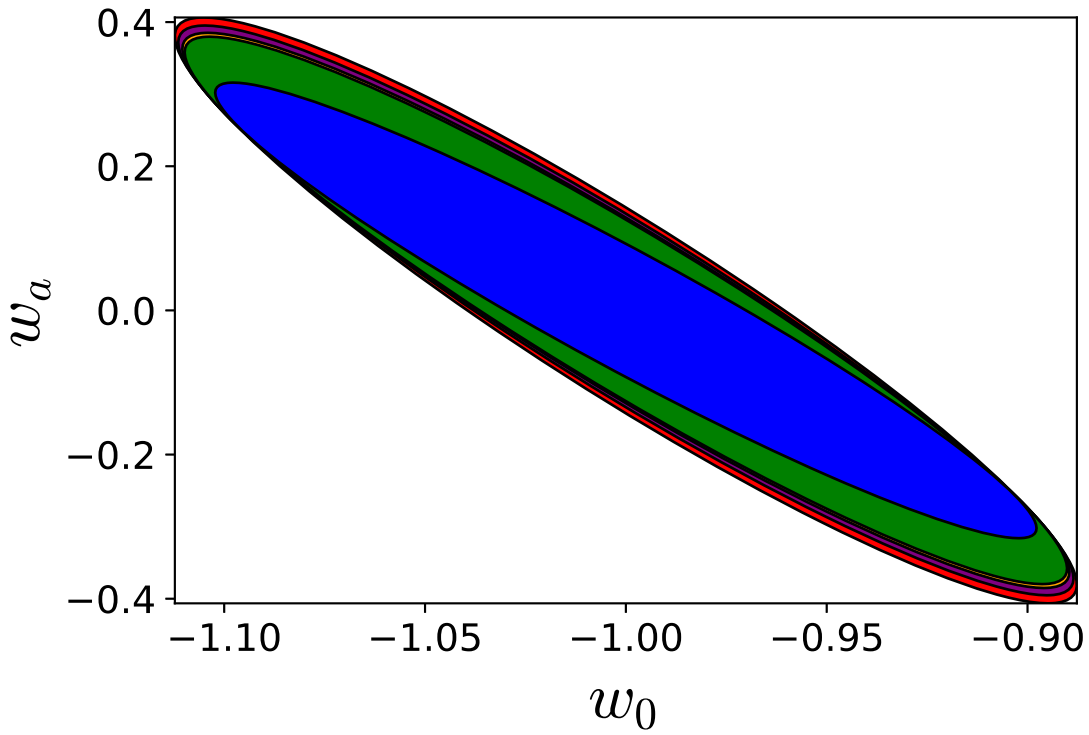

Figure 3.7: 1-σ 2-parameter confidence ellipses for w_0 and w_a. In each case, $\Omega_m, \Omega_b, h, n_s, \sigma_8$ have been marginalised over. Results are shown for all baryon parameters fixed to their fiducial values (blue); one baryon parameter fixed to their fiducial value, A_B (orange), η_0 (green), r_b (purple); and all baryon parameters marginalised over (red; just visible as the largest ellipse).

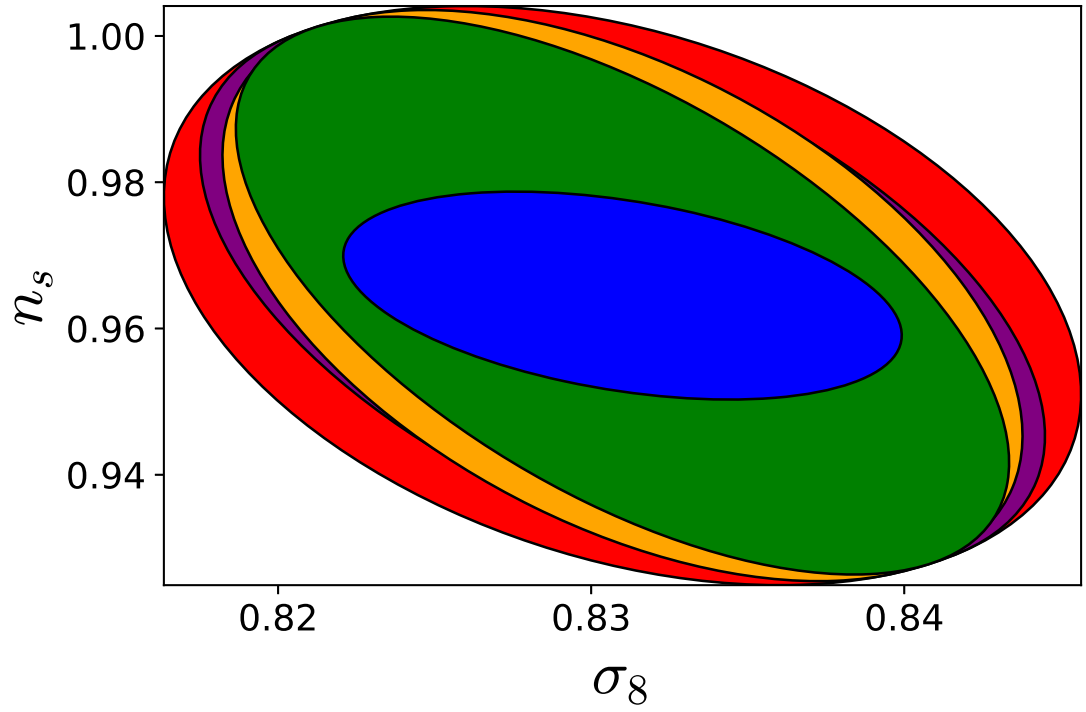

Figure 3.8: 1-σ 2-parameter confidence ellipses for σ_8 and n_s. In each case, $\Omega_m, \Omega_b, h, w_0, w_a$ have been marginalised over. Results are shown for all baryon parameters fixed to their fiducial values (blue); one baryon parameter fixed to their fiducial value, A_B (orange), η_0 (green), r_b (purple)); and all baryon parameters marginalised over (red).

	1-σ (no baryon marg.)	1-σ (inc. baryon marg.)	$f_{\rm deg}$
w_0	0.0673	0.0739	1.10
w_a	0.208	0.267	1.29
σ_8	0.00587	0.00964	1.64
n_s	0.00935	0.0260	2.79

Table 3.3: 1-σ error forecasts for dark energy and selected cosmological parameters for a Euclid-like survey, without and including marginalisation over baryon parameters. The ratio, $f_{\rm deg}$, of errors with baryons marginalised over to those without baryon marginalisation quantifies the degradation.

	FOM (no baryon marg.)	FOM (inc. baryon marg.)	$R_{\rm FOM}$
w_0-w_a	106	62.4	1.70
n_s-σ_8	8540	1830	4.65

Table 3.4: Figures of merit for w_0-w_a and n_s-σ_8 without and including marginalisation over baryonic physics. I include the reduction factor, $R_{\rm FOM}$, calculated as the ratio of the FOM without baryon marginalisation to the FOM with baryon parameters kept fixed.

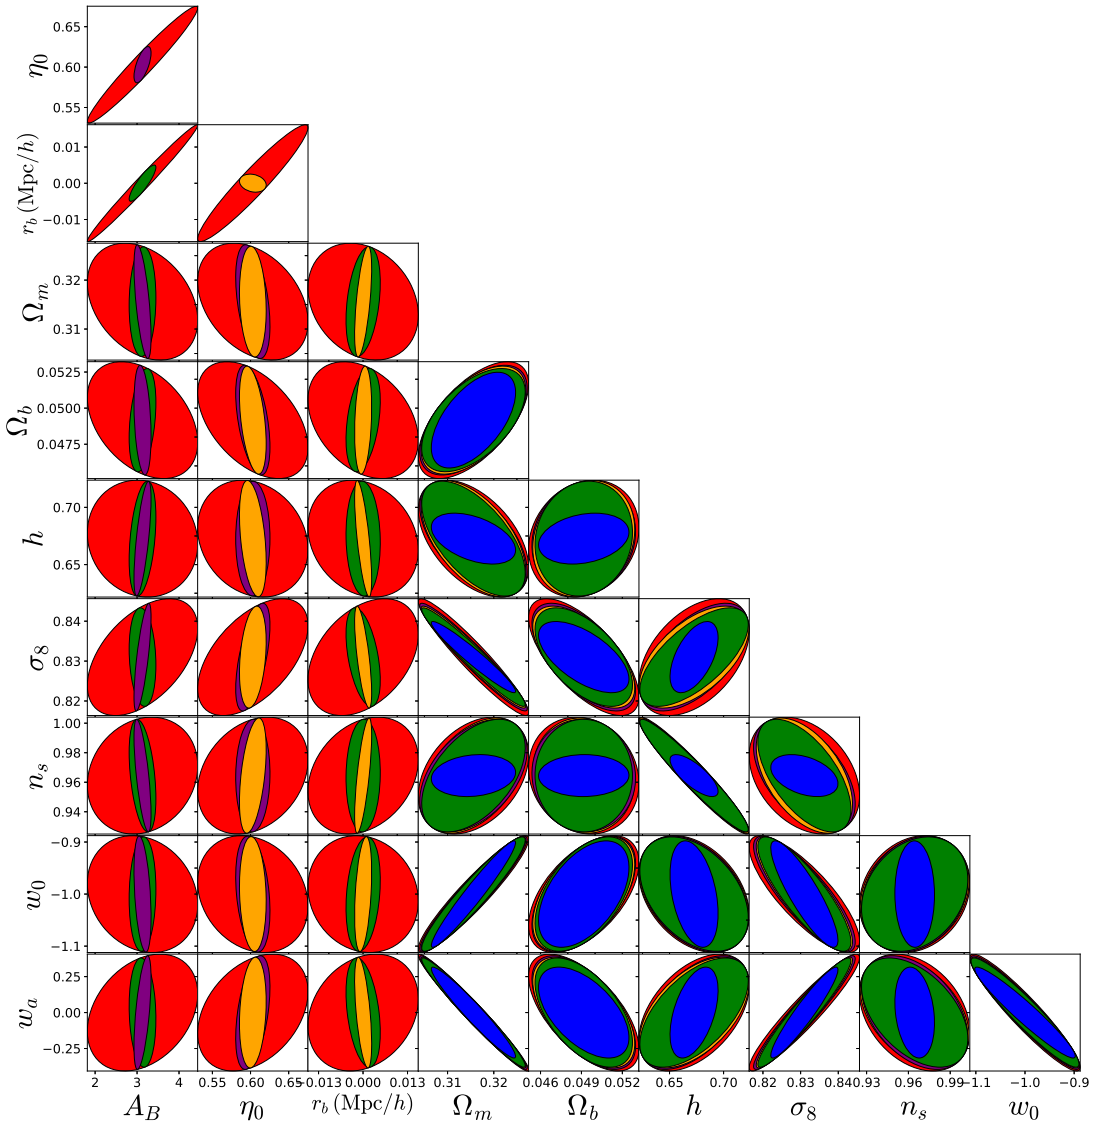

Figure 3.9: 1-σ 2-parameter confidence ellipses with different combinations of parameters marginalised over: all parameters in $\Theta = (A_B, \eta_0, r_b, \Omega_m, \Omega_b, h, \sigma_8, n_s, w_0, w_a)$ marginalised over (red); A_B fixed to its fiducial value (orange); η_0 fixed (green); r_b fixed (purple); A_B, η_0 and r_b fixed (blue).

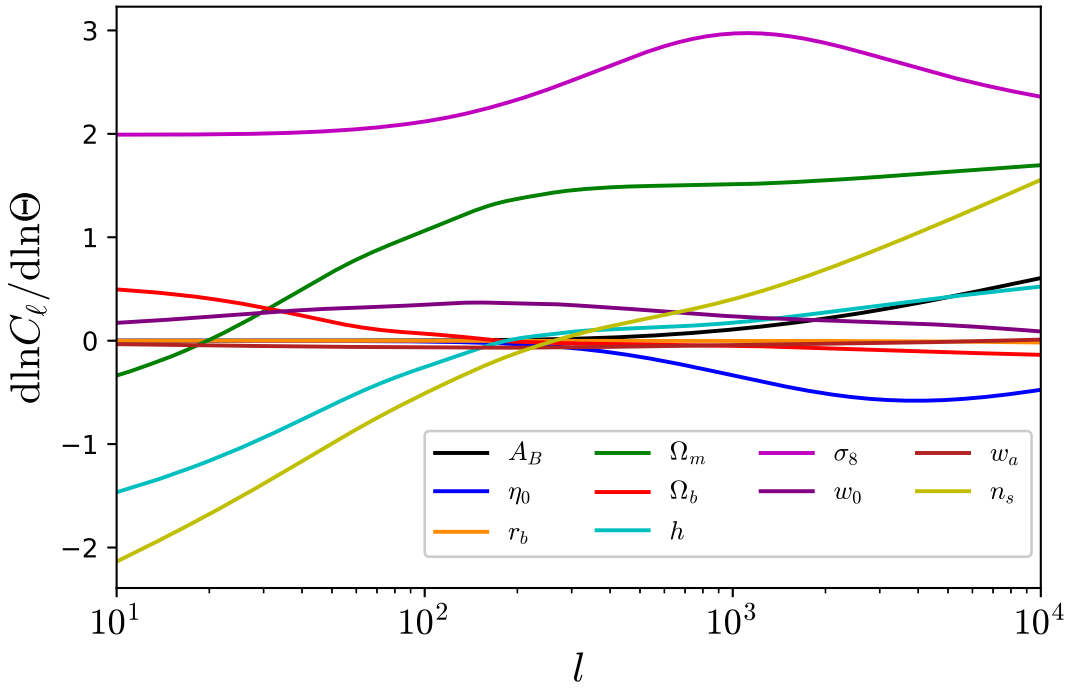

Figure 3.10: Logarithmic derivatives $\mathrm{d}\ln C_\ell/\mathrm{d}\theta$ of the weak lensing convergence power spectrum within the redshift bin $0.9 < z < 1.1$, with respect to baryon and cosmological parameters $\Theta = (A_B, \eta_0, \Omega_m, \Omega_b, h, \sigma_8, n_s, w_0, w_a, r_b)$.

Figure 3.9. This Figure includes results where different combinations of baryon parameters are systematically fixed, or marginalised over alongside the cosmological parameters.

The logarithmic derivatives of the weak lensing power spectrum with respect to each parameter are shown in Figure 3.10 as this is the essential contribution to the Fisher matrix. This is somewhat obscured by correlations of power between different redshift bins). I have chosen the redshift bin, $0.9 < z < 1.1$, in which to show these results. I plot the w_a derivative around $w_a = 1$ instead of $w_a = 0$. Similar to the power spectrum response plots, examining the derivatives is a useful tool for contextualising parameter degeneracies apparent in the forecast results. Parameters with associated derivatives of the same sign have correlated influences on the power. Similarly shaped derivatives are more likely to lead to degeneracies, although the effect of marginalising over other parameters adds a layer of complexity. The weighting of Fisher information per mode is greater for non-linear scales so the behaviour of the derivatives on these scales should provide the most insight into whether two parameters are likely to exhibit degeneracy in their confidence ellipse.

In Figure 3.7 the dark energy error forecasts are shown. There is a baryon degradation factor (computed as the ratio of the 1-σ errors when marginalising over baryons to those when baryon parameters are fixed), $f_{\rm deg} = 1.10$, on the w_0 error, and a more substantial degradation of $f_{\rm deg} = 1.29$ on the w_a error (see Table 3.3). These compound such that the dark energy FOM is reduced by a factor $R_{\rm FOM} = 1.70$.

When baryons are fixed to their fiducial values, the w_0-w_a parameter space can be constrained because the w_0-w_a degeneracy is broken by two sources. The first is differences in the scale dependence of the lensing power response to varying w_0 and w_a. Figure 3.6 illustrates that on non-linear scales, for a given redshift slice, competing influences from geometry and the matter power spectrum damp the net sensitivity of the lensing power to w_a. By contrast, the effect of varying w_0 on geometry remains dominant over its impact on the matter power deeper into the non-linear regime. The second source is contributions from multiple photometric redshift bins. The evolution of the lensing power spectrum depends on the growth rate and integrals over the line of sight, which respond differently to

w_0 and w_a. However, degeneracies between baryons and dark energy on non-linear scales obscure these distinctions in the Fisher analysis when baryons are marginalised over. Breaking the w_0-w_a degeneracy now depends on evolution alone and consequently the FOM experiences significant $\sim 40\%$ degradation.

Fixing any single baryon parameter does not significantly reduce the degradation on dark energy. Setting the inner core radius to zero, for example, will yield very limited improvement. However, Figure 3.8 exhibits more varied impacts on n_s-σ_8 constraints when selectively fixing different baryon parameters. This signals the importance of accounting for multiple baryon influences. Figure 3.11 shows there are correlations between baryon parameters in Stage IV forecasts but they are not sufficiently degenerate to motivate reduction to a single parameter, as in Hildebrandt et al. (2017).

The dependence of forecasts on the choice of baryon fiducial values should also be briefly noted. For example, selecting A_B and η_0 values that best fit the OWLS AGN simulation ($A_B = 2.32$, $\eta_0 = 0.76$) results in a $\sim$10% change in w_0 constraints, giving $\sigma_{w_0,\mathrm{AGN}} = 0.104$.

Figure 3.11 illustrates that other cosmological parameters are also vulnerable to baryon degradation, using n_s-σ_8 as an example. $P(k)$ and C_ℓ experience non-linear peak responses to σ_8 and η_0 (see Figures 3.4 and 3.6). The spectral index amplifies power with scale $k > 1\,h\,\mathrm{Mpc}^{-1}$ similarly to varying A_B and r_b. These combined degeneracies reduce the n_s-σ_8 FOM to $\sim 20\%$ of its value before baryons are marginalised over. Though this work is focused mainly on dark energy, I include this result to highlight the importance of understanding baryonic effects for constraining all cosmological parameters.

3.8 Constraints On Baryon Parameters

3.8.1 Results

The analysis framework presented in this Chapter also allows one to quickly compute confidence limits for prospective measurements of the different baryonic effects by future surveys. Such information could then be useful for modelling

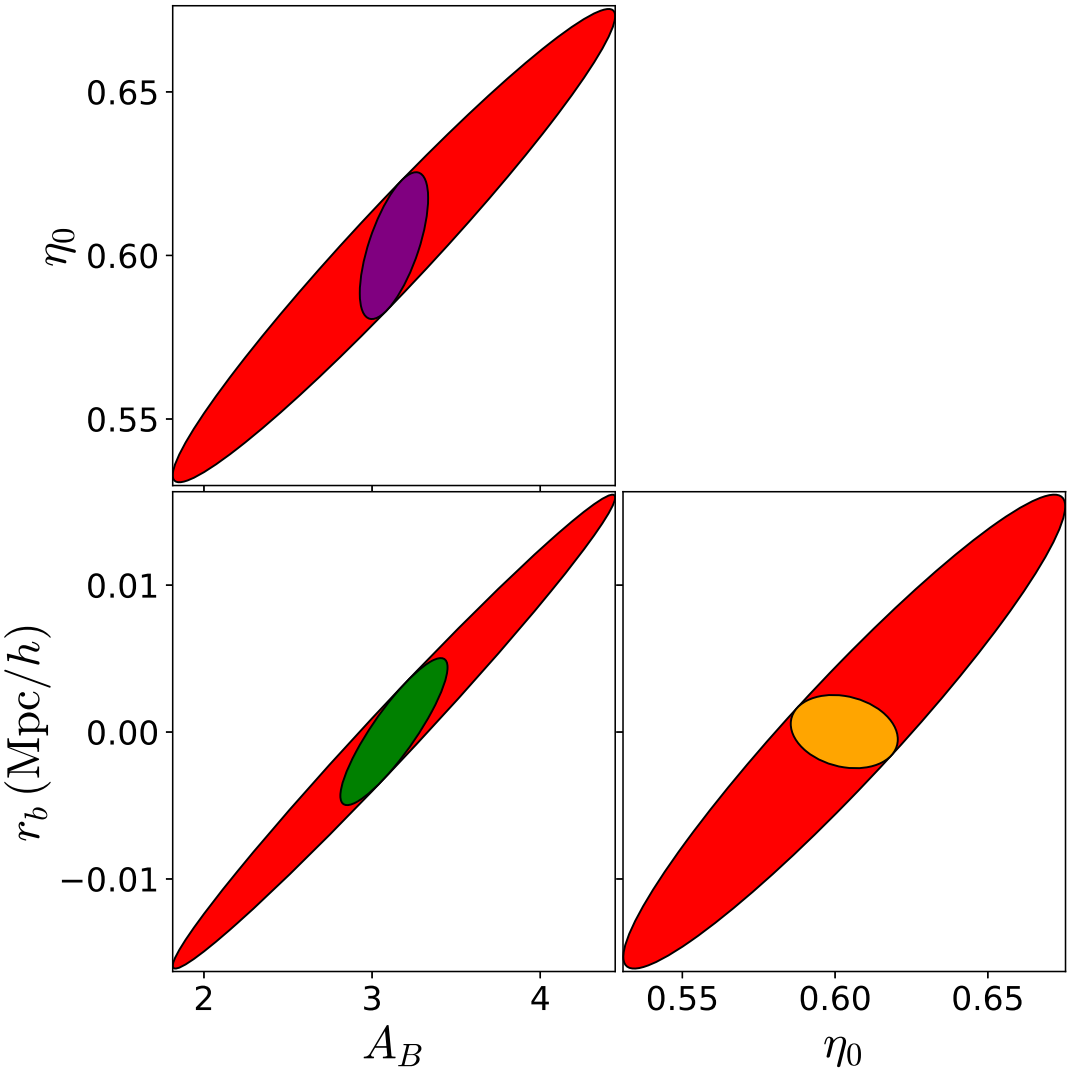

Figure 3.11: 1-σ 2-parameter confidence ellipses for (A_B, η_0, r_b): marginalised over all remaining parameters in $\Theta = (A_B, \eta_0, r_b, \Omega_m, \Omega_b, \ h, \sigma_8, n_s, w_0, w_a)$ (red); A_B fixed to its fiducial value (orange); η_0 fixed (green); and r_b fixed (purple).

baryons in future simulations. This represents an advantage of using multiple baryon parameters that are kept independent from each other, rather than imposing relationships between them that lead to further degeneracies.

In Figure 3.11, it is shown that A_B and η_0 could be constrained by a Euclid-like survey at the 50% and 10% level respectively, with 1-σ errors $\sigma_{A_B} = 0.866$ and $\sigma_{\eta_0} = 0.0476$. A significant improvement is made, particularly for constraining adiabatic contraction, by fixing the inner core radius to zero, as this breaks the degeneracy between A_B and r_b. This reduces the errors by several factors to $\sigma_{A_B,\mathrm{cusp}} = 0.134$ and $\sigma_{\eta_0,\mathrm{cusp}} = 0.0148$ for a cuspy halo model, while the A_B-η_0 FOM is greater by a factor of 6.7. The A_B-r_b and η_0-r_b constraints experience similar improvements when fixing the third baryon parameter in each case, further highlighting the degeneracies that occur between the different effects. It should also be noted that η_0 and r_b exhibit a positive correlation despite having opposite effects on C_ℓ. This is because there is also marginalisation over A_B, which has a dominant influence compared to r_b.

Surprisingly, the results presented here imply that a Euclid-like survey could be highly sensitive to cores on the smallest scales, within $0.02\,h^{-1}$Mpc. Various axionic mechanisms like fuzzy dark matter or solitonic field configurations of self-gravitating bosons that generate halo cores are posited to exist on kpc scales (e.g., Marsh and Pop, 2015). Such a preference in favour of the axion sector over baryon-induced mechanisms of core formation would be significant for the cusp-core debate. However, it is not realistic to expect the modelling underpinning these forecasts to be robust at these scales. As an example, one of the major sources of bias in lensing forecasts - degradation due to intrinsic alignments - would be expected to have a significant impact. This Chapter does not address that particular systematic, although work presented in the following Chapter discusses its significance in detail. I report the finding primarily as an indication that a more sophisticated approach to incorporating inner cores could draw substantial benefits from Stage IV forecasts.

As the scope of this work is primarily focused on dark energy constraints, the main result of this section is that only 59% of the dark energy FOM is retained when baryons are marginalised over (see Table 3.4). It is now necessary to explore possible mitigation strategies and identify with greater precision the scales

	1-σ (no baryon marg.)	1-σ (inc. baryon marg.)	$f_{\rm deg}$
w_0	0.0656	0.0725	1.10
w_a	0.197	0.258	1.31
σ_8	0.00531	0.00903	1.70
n_s	0.00740	0.0247	3.34

Table 3.5: 1-σ error forecasts for dark energy and selected cosmological parameters for a survey probing up to $\ell_{\rm max} = 10000$, without and including marginalisation over baryon parameters. The final column shows the degradation factor induced in the errors by baryonic physics.

$\ell_{\rm max}$	${\rm FOM}_{\rm w_0-w_a}$ (no baryon marg.)	${\rm FOM}_{\rm w_0-w_a}$ (inc. baryon marg.)	$R_{\rm FOM}$
5000	106	62.4	1.70
10000	121	66.5	1.82
$r_{10000/5000}$	1.15	1.07	1.07

Table 3.6: Figures of merit for w_0-w_a without and including marginalisation over baryon parameters, at $\ell_{\rm max} = 5000$ to $\ell_{\rm max} = 10000$. The final column and row respectively show the reduction fraction, $R_{\rm FOM}$, of the FOM when including marginalisation over baryonic physics and the ratio, $r_{10000/5000}$, of quantities computed using $\ell_{\rm max} = 10000$ to those computed using $\ell_{\rm max} = 5000$.

at which constraining information becomes compromised.

3.9 Mitigation

3.9.1 Fisher Information Sensitivity

An analysis that extends further into non-linear scales might be expected to improve upon the baryon impact. Instead it can be seen in Tables 3.5 and 3.6 that doubling the survey limit from $\ell_{\rm max} = 5000$ to $\ell_{\rm max} = 10000$ offers a relatively minor improvement of $\sim 7\%$ on the dark energy FOM. More than twice this gain is available when the baryon parameters are kept fixed to their fiducial values. This reflects the fact that baryon degradation actually worsens at higher $\ell_{\rm max}$, as shown in Figure 3.12.

I find degradation is minimised at $\sim 20\%$ for the optimal cutoff $\ell^{\star}_{\rm max} \approx 1700$. The corresponding marginalised errors on w_0 are $\sigma^{\star}_{w_0,b} = 0.080$ and $\sigma^{\star}_{w_0,nb} = 0.074$ for

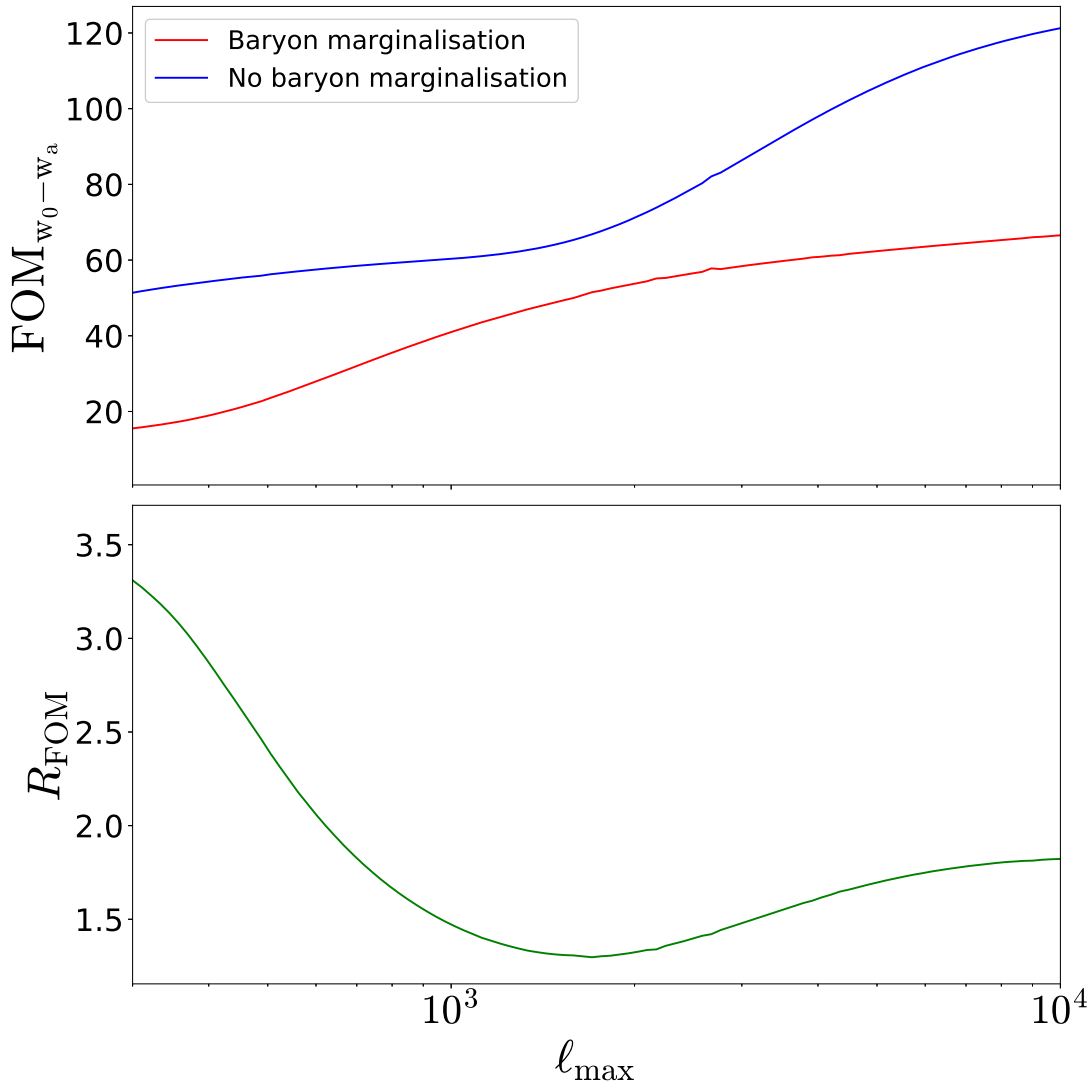

Figure 3.12: Top panel: FOM for w_0-w_a at different ℓ-mode cutoffs, with baryon parameters fixed to their fiducial values (blue) and marginalised over (red). Bottom panel: the reduction factor, $R_{\rm FOM}$ on the FOM due to baryon uncertainty.

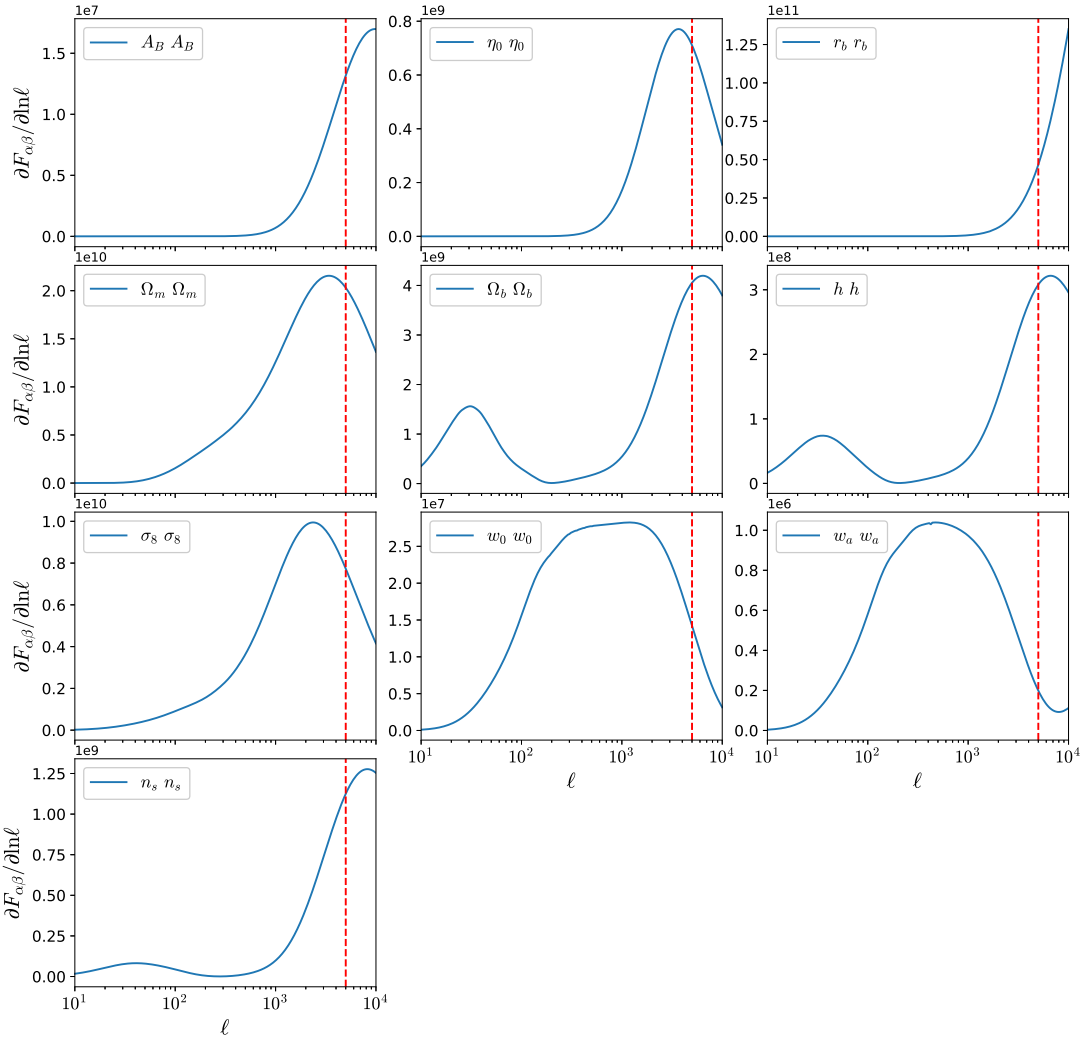

Figure 3.13: Diagonal contributions to the Fisher information at each ℓ-mode for each cosmological and baryon parameter. The vertical line (red, dashed) in each panel marks the conventional cutoff for Euclid-like surveys at $\ell_{\max} = 5000$. The derivative with respect to $\ln \ell$ is shown to more accurately convey the integration area covered by the logarithmically spaced ℓ-range.

marginalising over and fixing baryons respectively. For larger $\ell_{\max}$, the degradation worsens while the overall improvement in the FOM tails off. Despite these limitations, it should be emphasised that increasing $\ell_{\max}$ is still beneficial because the information gained at these scales reduces errors. However, it should not necessarily be assumed that simply increasing $\ell_{\max}$ is the most effective approach to mitigating the baryon impact.

The optimal $\ell_{\max}$ for reducing degradation can be understood by considering the scale dependence of Fisher information contributions from different parameters. Figure 3.13 shows this for the diagonal terms of the Fisher matrix. The peak contributions for the baryon parameters occur almost entirely over non-linear scales. However, the dark energy contribution is evenly distributed across a range of linear and quasi-nonlinear scales of a few hundred ℓ, quickly dropping off beyond $\ell \sim 1000$. This is due to the combination of decreasing sensitivity of lensing power to w_0 and w_a on these scales (see Figure 3.6) and the increasing influence of shape noise.

In Figure 3.12 the FOM branch without baryon marginalisation experiences an upturn at a greater $\ell_{\max}$ than the scale of the dark energy Fisher information peak. This is because the differences in lensing power sensitivity to w_0 and w_a on non-linear scales help to break the w_0-w_a degeneracy. Baryons impair this capacity, so the baryon marginalisation branch relies almost entirely on responses varying with evolution to break the degeneracy. Therefore, there is no upturn, resulting in the relative degradation increasing.

For low $\ell_{\max}$, though most of the dark energy information is available, baryons provide almost no information. The Fisher matrix is therefore close to singular at low $\ell_{\max}$ and the forecast FOM after marginalising over baryons is likely inaccurate, although the decrease in FOM with decreasing $\ell_{\max}$ is expected to be qualitatively correct. In practice, for very low values of $\ell_{\max}$ the baryon parameters would be fixed, allowing them to vary only when $\ell_{\max}$ is large enough that the data is sufficiently informative that fixing them does not induce significant biases.

There is a caveat to the analysis presented in this work. While calculating Fisher information using power spectra covariance matrices that are block diagonal in ℓ is a reasonable approximation, it is not strictly accurate. Particularly at low

redshifts, non-linear modes couple and further correlate power. The traditional Fisher framework (Tegmark et al., 1997) can be extended to account for this additional information. For instance, Kiessling et al. (2011) developed a method to model non-linear mode-coupling by treating the shear-power distribution as a multivariate Gaussian but with a covariance matrix derived from mock weak lensing surveys. It is beyond the scope of this work to include the necessary non-Gaussian corrections, here but the issue is raised to highlight that they should be considered for future studies.

3.9.2 External Baryon Priors

As little mitigation is offered by extending the minimum scale of an analysis, one can instead consider adding information from independent sources. I first explore the possibility of an external baryon prior to break degeneracies with dark energy and increase the total Fisher information available. For example, Hildebrandt et al. (2017) adopt a top-hat prior, $2 < A_B < 4$, based on the range of fits to different OWLS simulations M15 find for A_B and η_0. A stronger prior will have to come from future observations or simulations, so it is important to quantify the level of further information required. I provide a recommendation for limiting degradation of the error on w_0 to 1% by adding a diagonal baryon prior Fisher element, αF_{bb}, so that

$$F'_{bb} = F_{bb}\,(1 + \alpha)\,, \tag{3.34}$$

where α is the external improvement factor. For simplicity a single prior is imposed, assuming a prospective scenario with a broadly equivalent information increase on all baryonic phenomena. Figure 3.14 indicates that reducing degradation to 1% requires $\alpha = 0.47$. This corresponds to an external prior of $\sigma'_{b,\mathrm{con}} = 0.82\,\sigma_{b,\mathrm{con}}$ in terms of the conditional baryon errors.

At first, baryon degradation is highly sensitive to relatively small increases in information. This is cause for tentative optimism that if sufficient external data can be used to further inform Stage IV surveys, substantial improvements can be made comparatively easily. However, the rate of improvement with increasing information is soon damped, so significantly stronger priors are required to achieve negligible degradation.

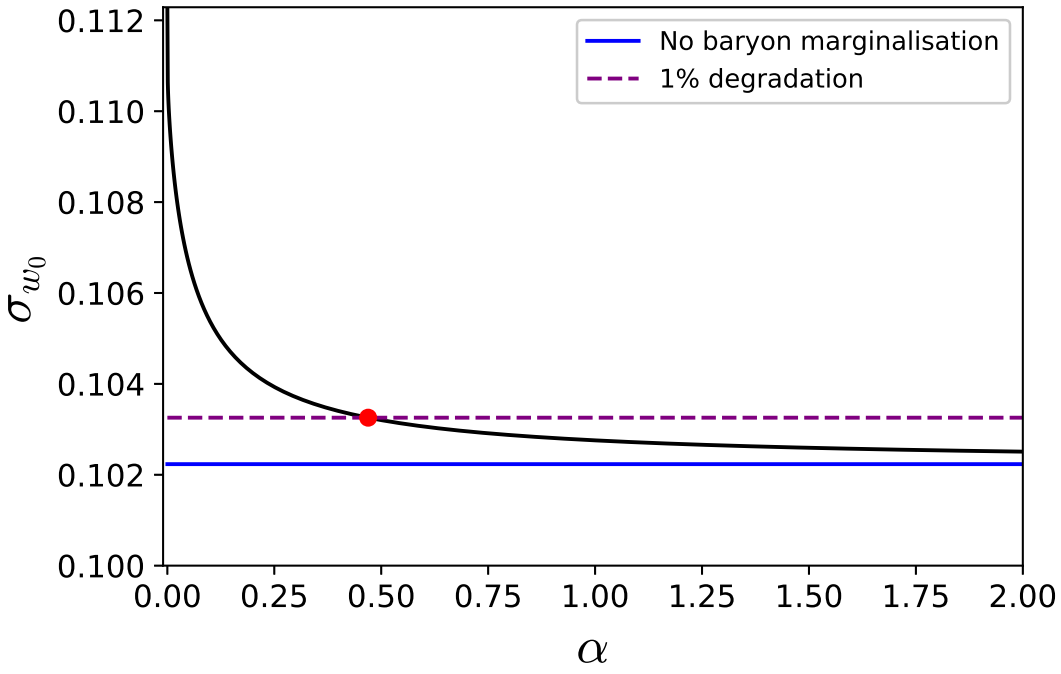

Figure 3.14: 1-σ marginalised errors for w_0 as a function of the enhanced baryon Fisher information parameter α. The blue line indicates the absolute best case scenario of no baryon uncertainty, while the purple dashed line represents an acceptable degradation threshold of 1%. The red dot marks the $\alpha \approx 0.47$ value required to achieve this improvement.

	w_0-w_a	n_s-σ_8
FOM_{WL} (no baryon marg.)	106	8540
FOM_{WL} (inc. baryon marg.)	62.4	1830
$R_{\text{FOM,WL}}$	1.70	4.65
$\text{FOM}_{\text{WL+CMB}}$ (no baryon marg.)	283	128000
$\text{FOM}_{\text{WL+CMB}}$ (inc. baryon marg.)	145	75800
$R_{\text{FOM,WL+CMB}}$	1.96	1.69

Table 3.7: Figures of merit for w_0-w_a and n_s-σ_8 without and including marginalisation over baryonic physics, and with and without the addition of priors on ΛCDM cosmological parameters from *Planck* CMB measurements. For the cases with and without priors, I include the reduction factor, R_{FOM}, of the FOM when including baryon marginalisation to the FOM when baryons are fixed.

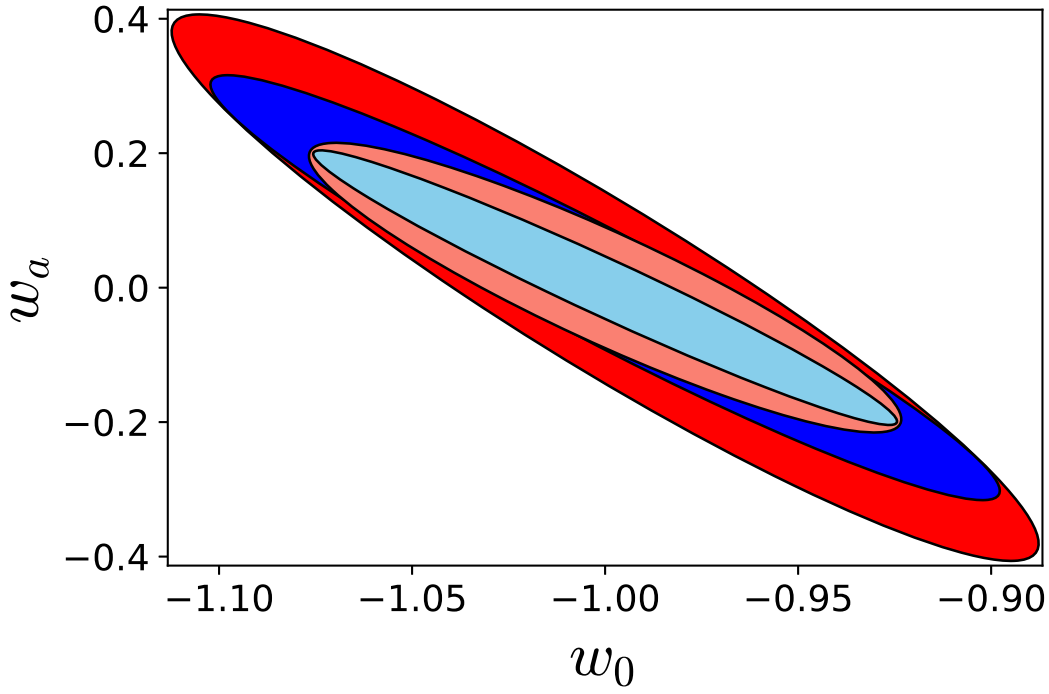

Figure 3.15: 1-σ 2-parameter confidence ellipses for w_0 and w_a. In each case, $\Omega_m, \Omega_b, h, n_s, \sigma_8$ have been marginalised over. Results are shown for when all baryon parameters are fixed to their fiducial values (without *Planck* CMB priors: blue; with priors: light blue) and when all baryon parameters are marginalised over (without *Planck* CMB priors: red; with priors: pink).

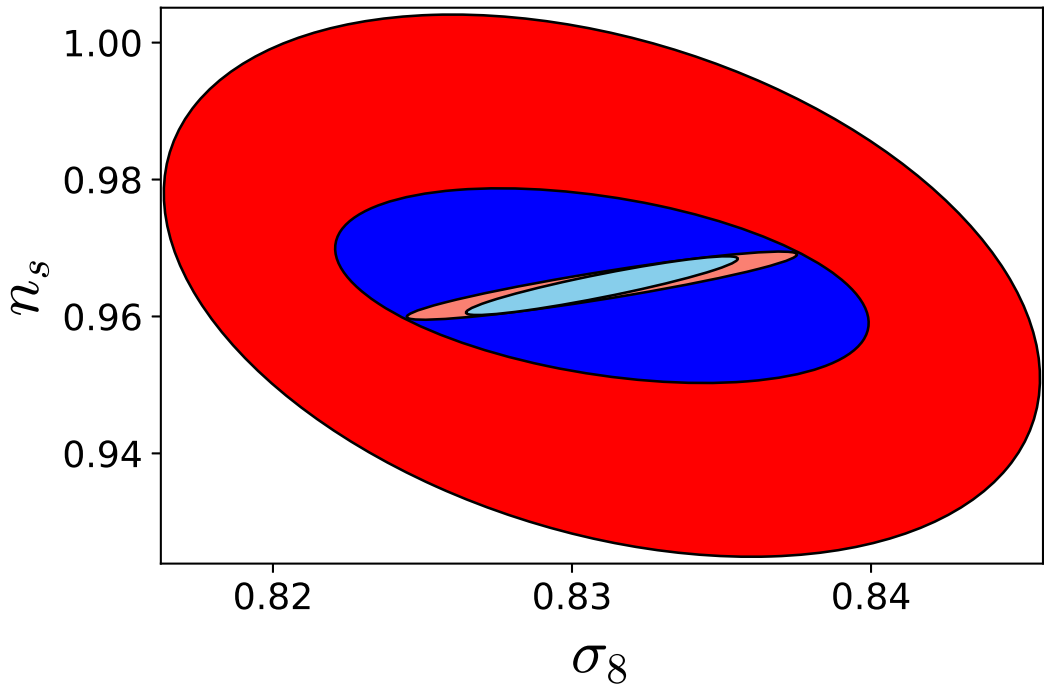

Figure 3.16: 1-σ 2-parameter confidence ellipses for σ_8 and n_s. In each case, $\Omega_m, \Omega_b, h, w_0, w_a$ have been marginalised over. Results are shown for when all baryon parameters are fixed to their fiducial values (without *Planck* CMB priors: blue; with priors: light blue) and when all baryon parameters are marginalised over (without *Planck* CMB priors: red; with priors: pink).

3.9.3 Cosmic Microwave Background Priors

If the price of external baryon information is too steep, priors can also be added on the cosmological parameters from sources like the early Universe that are independent of Stage IV large-scale structure survey constraints. Inverting the Fisher matrix propagates this information through to the dark energy errors, potentially mitigating baryon degradation.

Excellent information on the cosmic geometry and matter-energy density is provided by the most recent CMB anisotropy measurements from *Planck*. I use the publicly available MCMC chains for the base ΛCDM combined TT, TE and EE power spectra (see Table 4 in Planck Collaboration et al., 2016b). Constraints from the CMB on w_0 and w_a alone are extremely weak without adding information from weak lensing and external sources like BAOs (Planck Collaboration et al., 2016c). I therefore use ΛCDM constraints, which I derive by constructing a covariance matrix from the MCMC chains for $(\Omega_m, \Omega_b, h, n_s, \sigma_8)$. Inverting this incorporates uncertainties from the cosmological parameters into the resulting prior Fisher matrix, $F_{\rm CMB}$ [2]. The total Fisher information from weak lensing (WL) via a Euclid-like survey and from the CMB via *Planck* is then

$$F_{\rm tot} = F_{\rm WL} + F_{\rm CMB}. \tag{3.35}$$

Rows and columns of zeroes corresponding to baryon and dark energy parameters have been added to $F_{\rm CMB}$ to satisfy the parameter space dimensionality. The resulting improvements on dark energy constraints for w_0 and w_a are shown in Figure 3.15, and for n_s and σ_8 in Figure 3.16. Both sets of results are summarised in Table 5.1.

The CMB provides very strong constraints for n_s and σ_8, dramatically improving the forecast obtained from weak lensing alone, and removing much of the relative baryon degradation. There is an interesting comparison with the w_0-w_a constraints. These parameters are not themselves constrained by the CMB but adding the priors still more than doubles the FOM, including when baryons are marginalised over. This is mainly due to breaking degeneracies between dark energy and Ω_b, h, and n_s. The scale of the transition to non-linear power is affected

[2] Even though the parameter space is non-Gaussian, this approach is consistent within the Fisher approximation.

by Ω_b and h, while n_s tilts $P(k)$ around $k = 1\,h\,\mathrm{Mpc}^{-1}$. Hence, as illustrated in Figures 3.4 and 3.6, comparable boosts to the non-linear power spectrum occur from raising or lowering these parameters. Dark energy similarly amplifies non-linear power with scale. Adding CMB information on Ω_b, h and n_s alleviates these degeneracies, so the dark energy constraints improve substantially.

Figure 3.15 shows that one linear combination of (w_0, w_a) no longer suffers from baryon degradation when CMB priors are included. However there is comparatively limited alleviation of the relative degradation for other directions in the parameter space. This could be due to CMB data being unable to provide information about the relationship between w_0 and w_a, and baryons. This contributes to the relative degradation of the FOM being largely unchanged after the inclusion of priors. Therefore it should be emphasised that, despite the FOM doubling, the key statistic for constraining the dark energy parameter space is no less impacted by baryons.

3.10 Model Bias

The fiducial values of A_B and η_0 have been determined by fitting to simulations in M15. This does not account for systematic limitations of the simulations or incorrect physics. It is important to know how far from the fitted values the true values can lie before w_0 and w_a estimates are severely biased. Taylor et al. (2007) showed that a first-order approximation of the bias in a cosmological parameter, θ, can be related to the bias in a nuisance parameter, ψ, (in this case, baryons) through sub-blocks of the full Fisher matrix, such that

$$\delta\theta_i = -\left[F^{\theta\theta}\right]^{-1}_{ik} F^{\theta\psi}_{kj}\delta\psi_j, \tag{3.36}$$

in which k is implicitly summed over. Figure 3.17 shows the relative biases induced in w_0 and w_a when the 'true' values of the baryon parameters deviate from the fitted values. Biases of up to 25% can occur for w_0 if the true values lie at the edges of the ranges $2 < A_B < 4$ and $0.5 < \eta_0 < 0.8$ found by M15 fits to OWLS simulations. A line of minimal bias emerges for both w_0 and w_a due to the first-order cancellation of A_B and η_0 biases. This happens to be almost perpendicular to the minor axis of the marginalised baryon confidence ellipse.

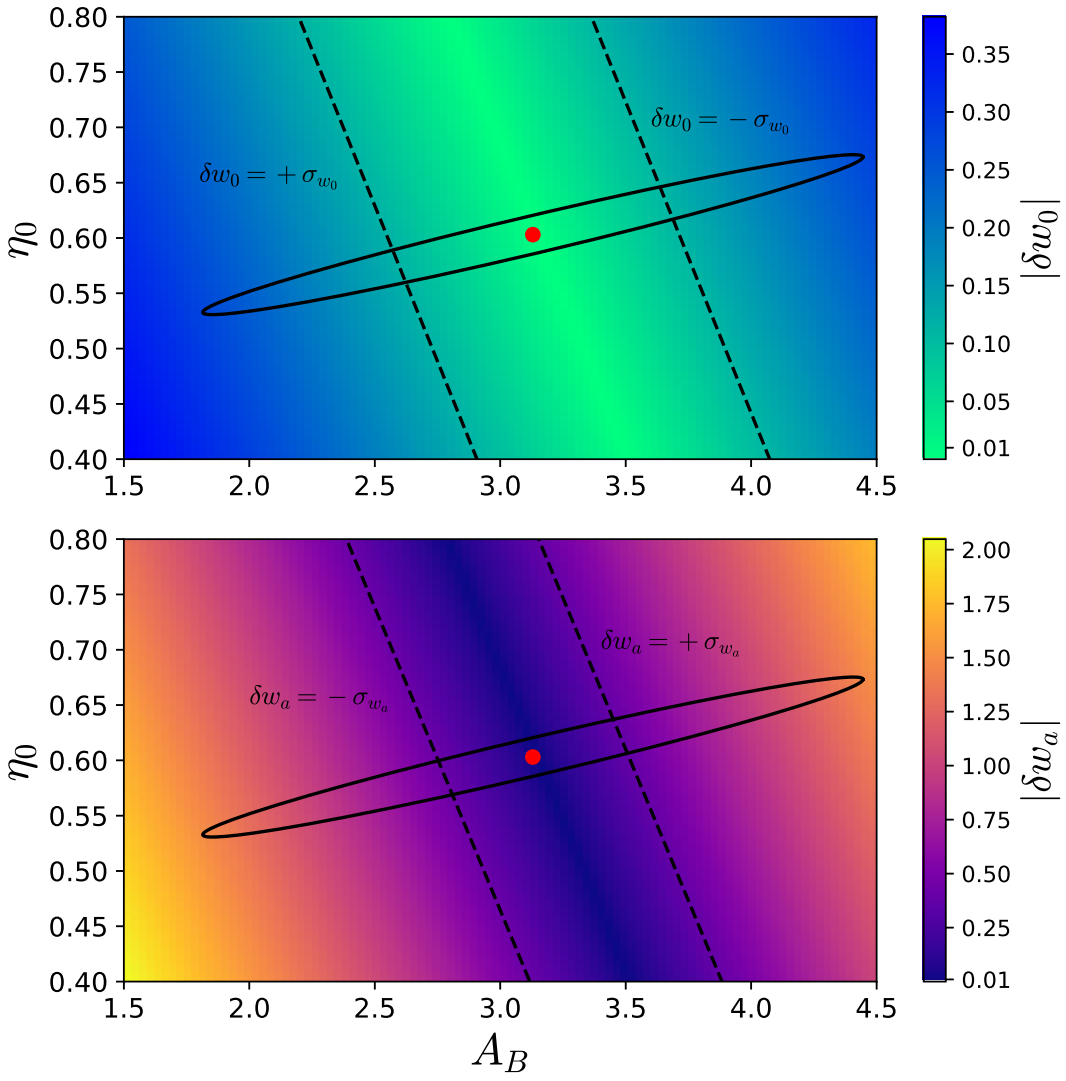

Figure 3.17: Absolute bias in w_0 (top) and w_a (bottom) due to model bias in A_B and η_0. The ellipse represents the marginalised 1-σ confidence region for the baryon parameters, with the red dot marking the fiducial point ($A_{B,\mathrm{fid}} = 3.13, \eta_{0,\mathrm{fid}} = 0.603$). The dashed lines mark the A_B-η_0 bias corresponding to the marginalised 1-σ errors for w_0 and w_a.

The A_B-η_0 regions in which the resulting bias to w_0 and w_a is within the 1-σ marginalised error limits is given by

$$\begin{cases} |\delta w_0| \leq \sigma_{w_0}, & -0.563A_B + 2.04 \leq \eta_0 \leq -0.563A_B + 2.69 \\ |\delta w_a| \leq \sigma_{w_a}, & -0.547A_B + 2.11 \leq \eta_0 \leq -0.547A_B + 2.53 \end{cases} . \tag{3.37}$$

A significant proportion of the baryon 1-σ confidence region generates an acceptable level of bias. However, positions in the parameter space that would generate biases approaching 35% for w_0, and even more severe effects for w_a, remain within the bounds of my forecasts.

Model bias is difficult to mitigate because it arises directly from subgrid limitations. A solution likely requires external data on baryon phenomenology, so it is beyond the scope of this work to do more than assess the potential impact of the issue.

3.11 Summary

The work presented in this Chapter has built upon previous analytic modifications of the halo model to account for the effect of baryonic astrophysical phenomena on the distribution and power spectrum of matter. I have used the baryon-halo model of M15 to incorporate the impact of adiabatic contraction on halo concentration, and the halo mass-dependent bloating effects of baryonic feedback from e.g., AGN and supernovae. The model of M15 was chosen because it provides accurate fits for the power spectrum to within a few percent by calibrating parameters to the COSMIC EMU (Heitmann et al., 2014) and OWLS (Schaye et al., 2010) simulations. Other approaches (e.g., Semboloni et al., 2011; Mohammed et al., 2014) focus on precisely modelling stellar, gas and dark matter distributions. This underlying physics is ultimately of less interest than the broad, empirically motivated corrections in M15 to the power spectrum. I have extended the model by incorporating an inner halo core, r_b, to account for small-scale structure. This was motivated by an array of baryonic feedback mechanisms, or the condensation of ultra-light axions instead of CDM in the inner halo.

I have examined the degradation that marginalising over the baryon parameters (A_B, η_0, r_b) has on constraints on the w_0-w_a dark energy parameter space forecast for a Euclid-like Stage IV cosmological survey. This was done by studying the impact that varying cosmological and baryon parameters has on $P(k)$ and C_ℓ at different scales, which informed my interpretation of a full Fisher analysis. The baryon degradation to the errors on w_0 and w_a is $\sim$10% and $\sim$30% respectively. However, as the FOM is quadratic in parameter uncertainty this translates to a $\sim$40% degradation in the capacity of a Stage IV survey to deliver accurate measurements. Though I applied my methodology to a Euclid-like survey, it could also be used for other next generation surveys like LSST.

This work has also highlighted that the effect of baryons is not limited to dark energy, showing the severe degradations on forecasts for n_s-σ_8 errors as an example. This illustrates the potential risk in making confident claims from these surveys even for cosmological parameters which are otherwise well-constrained from sources like the CMB.

I have shown that my framework can forecast constraints on baryons, marginalised over the uncertainty remaining in cosmology. This could potentially provide useful information for modelling baryons in simulations. Euclid-like surveys only constrain A_B and η_0 at the 50% and 10% level respectively, with $\sigma_{A_B} = 0.866$ and $\sigma_{\eta_0} = 0.0476$, although these improve significantly if the inner core is zero, reducing to $\sigma_{A_B,\mathrm{cusp}} = 0.134$ and $\sigma_{\eta_0,\mathrm{cusp}} = 0.0148$. The results presented here imply that the inner cores themselves could be constrained to a few kpc. The underlying modelling is unlikely to be robust on such scales but this remains an interesting indicator of the capacity of such surveys to potentially forecast inner cores in more sophisticated models. If this result were accurate it would have important implications for the cusp-core debate as it would address the question of cores arising from axion condensation on these scales.

The degradation found to dark energy forecasts is of a similar level to Mohammed et al. (2014), who also demonstrate a $\sim$10% baryon impact on w_0 constraints. The model used here has more freedom to vary individual baryonic effects, and has accurate power spectra fits to COSMIC EMU and OWLS. However, the consistent results should be seen as an encouraging sign that the magnitude of the baryonic impact on the w_0-w_a parameter space is well-understood. This should

temper concerns from the far more pessimistic predictions of Zentner et al. (2013) of 50% level degradations to w_0 and w_a. The larger impact could be attributed to inaccuracies in the baryon modelling by not accounting for distinct distributions of heated gas and cold dark matter (as noted by Semboloni et al., 2013), and calibrating to less accurate power spectra than have since become available (M15).

To inform a possible mitigation strategy I first explored the lensing scales on which Fisher information is most sensitive to both cosmological and baryon effects. It was determined that the region of maximum sensitivity for dark energy occurs at $\ell \sim 100$, i.e. on substantially larger scales than are typically assumed. This is due to the competing effects of geometry and growth broadening the impact of varying w_0 and w_a across a wide range of scales. It was illustrated that raising $\ell_{\rm max}$ has a limited improvement on the FOM and, in fact, suffers from an increasingly worse relative degradation.

A small amount of external baryon information from simulations or observations provided substantial improvements to the degradation on w_0 errors. However, the rate of improvement with information soon tails off so reaching a 1% degradation threshold requires priors of the order of the baryon conditional errors $\sigma_{\rm b,prior} = 0.82\sigma_{\rm b,con}$. This may prove challenging but as significant improvements are still possible, I consider this motivation to acquire stronger observational data for the influence of baryons on large-scale structure.

Constraints on dark energy greatly improved when including the strong cosmological priors offered by *Planck* CMB measurements. Particularly promising was the result that degradation on the errors for one linear combination of w_0 and w_a were almost completely removed. An important qualification is that due to dark energy itself being poorly constrained by the CMB there is no relative improvement to the degradation on the cross-covariance between w_0-w_a. Therefore, while the absolute improvements on the dark energy FOM are significant, the key statistic for constraining the parameter space remains as afflicted by baryons.

Finally, model bias emerging from incorrect calibrations of baryon parameters has been considered. I have calculated to first-order the bias that w_0 and w_a would experience due to the true values for A_B and η_0 deviating from fiducial values. Within the A_B-η_0 confidence ellipse there is a significant area correspond-

ing to w_0 and w_a bias within 1-σ error forecasts. I consider this region to be generally protected from inducing detrimental bias but close to half of the confidence ellipse overlaps with areas of larger biases. It is important to quantify these limitations, though it will ultimately require additional baryon information or improved simulations to mitigate the concern.

In summary, by incorporating inner cores into the baryon-halo model of M15, I have been able to encompass the full range of broad, empirically motivated baryonic effects on haloes. The framework presented here allows for quick and flexible predictions on both baryon and dark energy constraints. I anticipate that baryons will have a substantial but not catastrophic effect on the capacity of next generation surveys to constrain dark energy. Mitigation remains an issue. My thorough examination of the complex interplay of cosmological, baryon and dark energy effects on C_ℓ showed the limited value of enhancing the survey scope, or redirecting observing power to more linear scales. A combination of external baryon information and CMB priors offers significant improvements and reason for optimism, but there is still work to be done before degradation could be made negligible.

Chapter 4

Neutrino Mass Hierarchy Forecasts In The Baryon-Halo Model

The focus of this Chapter is an exploration of the capacity of Stage IV large scale structure surveys to measure the neutrino mass sum and differentiate between the normal and inverted neutrino mass hierarchies, when accounting for the impact of nuisance parameters controlling small-scale baryonic astrophysics and intrinsic alignments. The work presented here closely follows Copeland et al. (2019).

The basic framework outlined and discussed in Chapter 3 for capturing the impact of baryonic feedback and adiabatic contraction on power spectra generated from the halo model will also be applied here. Further modifications to non-linear halo power that account for the effects of massive neutrinos will also be discussed. The following Fisher analysis will be extended to galaxy clustering as well as the weak lensing probe of Chapter 3. In the case of the former, intrinsic alignments will be introduced to the lensing modelling and their impact as a major systematic assessed for the hierarchy forecasts. This work will therefore represent the first combined weak lensing and galaxy clustering Fisher analysis with baryons, intrinsic alignments, and massive neutrinos for both hierarchies.

The degradation from marginalising over baryonic effects to forecasts of the neutrino mass sum, Σ, will be assessed first for weak lensing probes. The prospects for reducing this degradation by adding cosmic microwave background *Planck*

priors will then be determined. Similarly, the improvements available from including galaxy clustering constraints from Euclid and BOSS (Alam et al., 2017) will be examined in detail. This analysis will be performed for both the normal and inverted hierarchies, with comparisons between the models and their responses to marginalising over nuisance parameters and introducing priors will be made throughout.

The forecasts combining these various probes and cosmological information sources will provide a comprehensive basis for assessing the confidence level that Σ can be expected to be measured by Stage IV surveys and whether a meaningful distinction between hierarchies is a realistic prospect. Potential improvements from future independent information sources, such as neutrinoless double beta decay will also be quantified.

The effect of intrinsic alignments (IA) on forecasts will be assessed with specific emphasis placed on the impact on forecasts when including IA in the weak lensing power spectrum. This is a separate issue from the important matter of bias in the predicted measured parameter values. However, model bias will also be investigated thoroughly in the context of miscalibrations of the baryon parameters from simulations, alongside intrinsic alignment bias.

The layout of this Chapter is as follows. In § 4.1, an introduction to the current state of neutrino mass sum forecasts is provided. In § 4.2 I outline the neutrino and baryon modifications to the halo model. The formalism of intrinsic alignments, and the model choices made for their implementation in this work, are discussed in § 4.3. The application of the Fisher formalism for galaxy clustering will be detailed in § 4.4. Results for the mass constraints on the normal and inverted hierarchies are presented in § 4.5. This will include an analysis of the degradation due to baryons and intrinsic alignments, and an exploration of the information gains available from combining weak lensing and galaxy clustering for a Euclid-like survey with priors from BOSS (Alam et al., 2017) and *Planck*. The significance of model bias will be assessed in § 4.6, with concluding points and a summary provided in § 4.7.

4.1 Introduction to neutrino mass sum forecasts

Flavour oscillations measured in solar, atmospheric, reactor and accelerator experiments have constrained the mass-squared differences, Δm^2_{12} and $|\Delta m^2_{23}|$, for the three neutrino mass eigenstates but cannot determine the absolute masses (e.g., Maltoni et al., 2004; Fogli et al., 2006). The sign of the largest splitting defines two possible orderings, known as the normal (NH, $\Delta m^2_{23} > 0$) and inverted (IH, $\Delta m^2_{23} < 0$) hierarchies. Lower bounds on the sum of the masses for each hierarchy, $\Sigma_{\rm NH} \gtrsim 0.06\,{\rm eV}$ and $\Sigma_{\rm IH} \gtrsim 0.1\,{\rm eV}$, can be arrived at numerically from the oscillation constraints (Lesgourgues and Pastor, 2006). Cosmology has been providing increasingly tight upper bounds. Data from the Lyman-α forest in combination with Planck Collaboration et al. (2016b) cosmic microwave background (CMB) constraints (Palanque-Delabrouille et al., 2015; Yèche et al., 2017) and the more recent Planck Collaboration et al. (2018) analysis of CMB and baryon acoustic oscillation (BAO) data both find $\Sigma \lesssim 0.12\,{\rm eV}$ (95% confidence level). Though this starts to confront the limits of the IH, it remains an open question as to whether future cosmological data will be able to distinguish between the hierarchies by constraining the mass sum. The degeneracy between hierarchies is partially broken by different sets of free-streaming scales associated with the mass eigenstates, but these differences are sufficiently small for a fixed mass sum that there is no realistic prospect of achieving a detection in the near future (Hall and Challinor, 2012).

Forthcoming Stage IV surveys, such as Euclid (Laureijs et al., 2011) and LSST (LSST Science Collaboration et al., 2009), are aiming to achieve Σ errors small enough to determine the hierarchy by using tomographic weak lensing and spectroscopic galaxy clustering. The formalism underpinning these probes has been discussed in previous Chapters. Understanding how galaxies trace the underlying matter distribution through clustering measurements provides a repository of information on the properties of the Universe (Hauser and Peebles, 1973; Davis and Peebles, 1983) that adds to that from weak lensing.

These probes are expected to observe the imprint of neutrinos through their suppression of matter clustering. Neutrinos are very light and so have large free-streaming lengths. For scales below these they do not cluster. For a fixed total matter density, this results in matter perturbations being slightly smoothed.

However, the cumulative effect on the matter power spectrum, $P(k)$, by $z = 0$ is a decrease on linear scales by a factor, $\sim 1-8f_\nu$, in linear theory, where $f_\nu \equiv \Omega_\nu/\Omega_m$ is the ratio of the neutrino and total matter density parameters (Hu et al., 1998). On non-linear scales, neutrinos are modelled through perturbative methods (e.g., Pietroni, 2008; Lesgourgues et al., 2009; Levi and Vlah, 2016) or accounted for in simulations. Generally, the most accurate approaches for the latter have implemented neutrinos as separate low-mass collisionless particle species alongside cold dark matter to fully capture their non-linear evolution (e.g. Bird et al., 2012; Massara et al., 2014; Liu et al., 2018). These indicate that the effect of neutrino suppression on haloes depends on their mass, leading to a greater power reduction than linear theory on intermediate scales of $k \sim 1h\,\mathrm{Mpc}^{-1}$ where larger haloes are significant. On deeper non-linear scales where smaller haloes predominate, the power, while suppressed overall, is boosted relative to the linear prediction.

Depending on the choice of probes, methodology and underlying model, 1-σ errors on the neutrino mass have been forecast broadly within the range $0.02 - 0.04\,\mathrm{eV}$ for Euclid-like surveys (e.g., Carbone et al., 2011; Audren et al., 2013; Boyle and Komatsu, 2018). A major systematic neglected in these forecasts is the redistributing of matter on halo scales by baryonic astrophysics. The previous Chapter discussed, at length, the impact of baryons on the matter distribution from a range of effects (e.g., adiabatic contraction and AGN and supernova feedback) across a range of scales. Overall, $P(k)$ experiences a $\sim$30% suppression in the mildly non-linear regime, which gives way to a net increase of power through adiabatic contraction on deeply non-linear scales of $k \sim 10h\,\mathrm{Mpc}^{-1}$ (e.g., Semboloni et al., 2011). As neutrinos and baryons both suppress matter power on weakly non-linear scales it is possible that a partial degeneracy exists between the two effects. It is therefore necessary to simultaneously constrain them both. Such degeneracies could potentially be broken by the scale-dependence of the baryonic effect and the redshift-dependence of neutrino suppression.

Marginalising over the uncertainty in baryon processes can significantly impact parameter forecasts. Significant biases in neutrino mass measurements were found by Natarajan et al. (2014) when failing to account for baryons, using simple arguments to model both components in $P(k)$, although the impact of marginalising over baryon parameters was not explored. More recently, Parimbelli et al. (2018) found that the neutrino mass is unlikely to be significantly biased by degen-

eracies with baryons, based on a Markov Chain Monte Carlo (MCMC) method using weak lensing and galaxy clustering that is limited by excluding freedom in the cosmological parameters. This work explores similar territory but addresses different key questions. I am primarily interested in the degradations to error forecasts from marginalising over baryonic effects, although the separate issue of bias in the neutrino mass itself is examined in the context of miscalibrations in baryon modelling. The specific focus in this Chapter is how degradation impacts the capacity to measure $\Sigma_{\rm NH}$ and $\Sigma_{\rm IH}$, and whether it will be feasible to distinguish between the mass hierarchies. To do this a Fisher analysis of the full range of cosmological parameters probed by Euclid-like weak lensing and galaxy clustering surveys will be performed.

The weak lensing analysis, addresses one further source of uncertainty than in the preceding Chapter. On large scales, gravitational fields cause tidal distortions in galaxies, which contribute to correlations between the intrinsic ellipticities and gravitational shear of galaxies. Failing to account for these intrinsic alignments can significantly bias lensing parameter estimates (Joachimi et al., 2011; Troxel and Ishak, 2015).

This work is the first to combine baryons, IA and massive neutrinos in a self-consistent Fisher framework in the context of distinguishing the NH and IH in Stage IV surveys. By applying this framework to weak lensing, spectroscopic galaxy clustering and CMB probes, this Chapter presents a uniquely comprehensive analysis of the current prospects for determining the hierarchy.

4.2 Modelling neutrino effects

I employ analytic modifications of the halo model to capture the effects of neutrinos and baryons on $P(k)$ and the weak lensing convergence power spectrum, C_ℓ. The Mead et al. (2015) halo model, HMCODE, is again used for its empirically driven baryon prescription of calibrating internal halo structure relations to match OWLS. The preceding Chapter provides a comprehensive discussion of the implementation of baryons in this halo model, which the reader is encouraged to refer to place this work in its proper context. Mead et al. (2016) (hereafter M16)

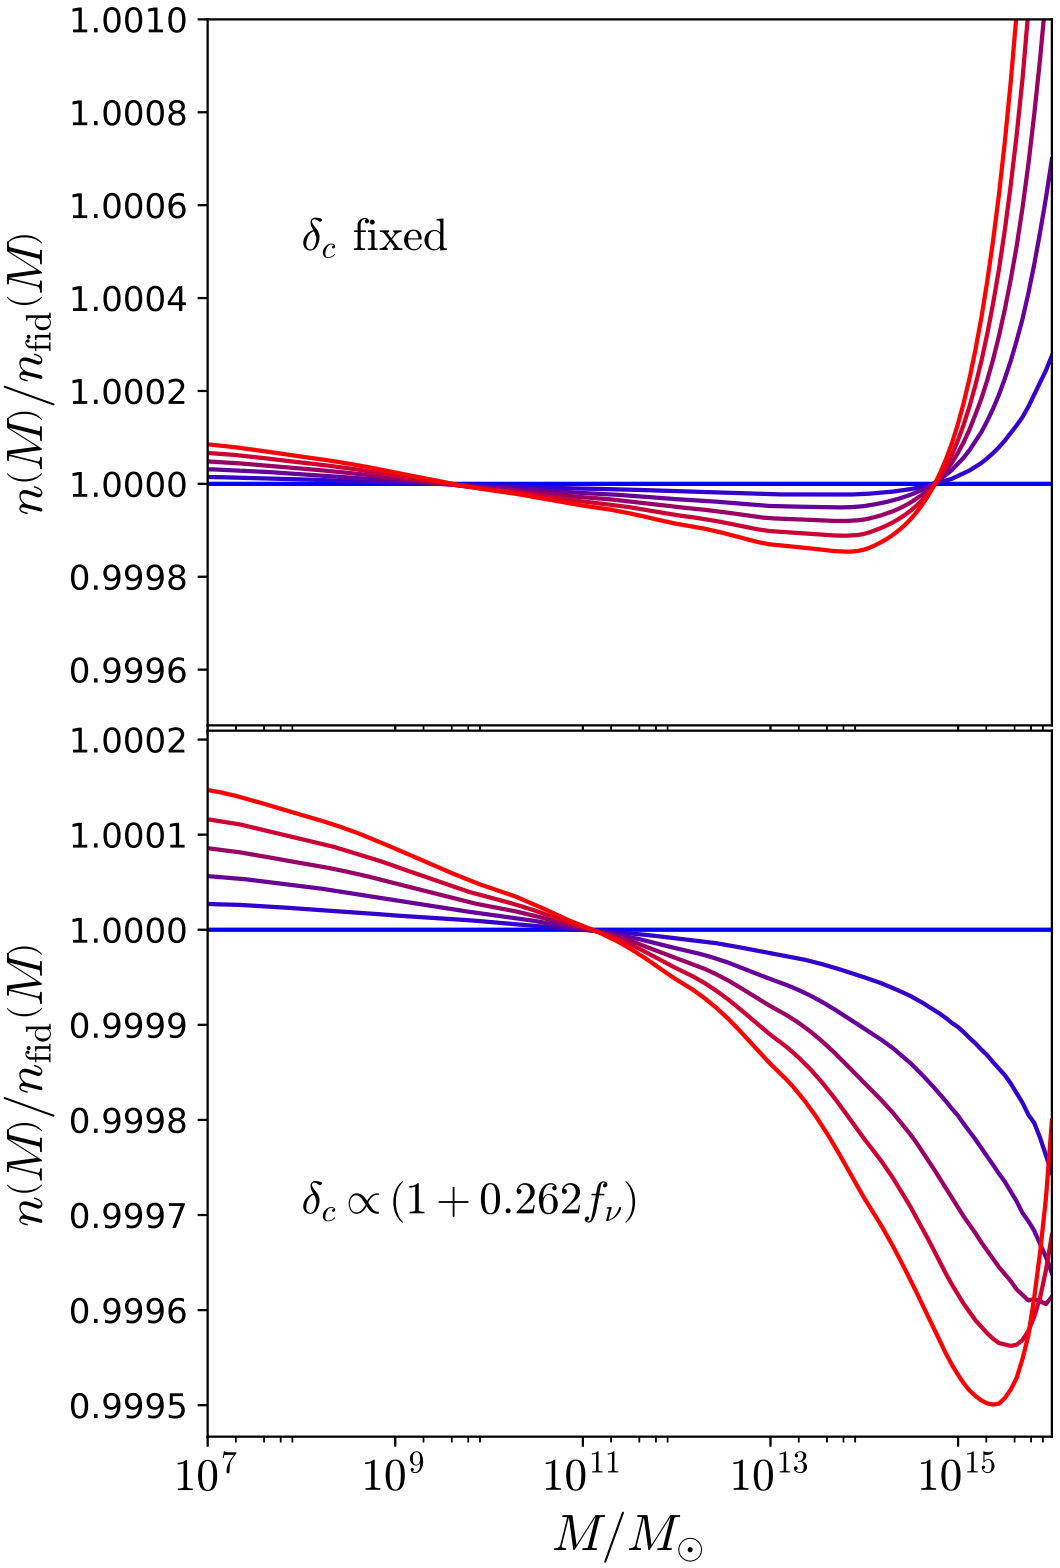

Figure 4.1: Response of the halo mass function at $z = 0$ to the neutrino mass sum for the normal hierarchy, with the horizontal blue line corresponding to the fiducial case using the minimal NH neutrino mass sum, $\Sigma_{\rm NH,min} = 0.06\,{\rm eV}$ and increasingly red curves corresponding to Σ in the range $\Sigma_{\rm NH,min} < \Sigma_{\rm NH} \leq 1.1\,\Sigma_{\rm NH,min}$. Top panel: the mass function is sensitive to the mass sum through $\sigma\,(M)$, which in turn depends on Σ through the linear matter power spectrum. There is no dependence of δ_c on Σ impacting the response. Bottom panel: the mass function depends on Σ through $\sigma\,(M)$ and δ_c, the latter being sensitive to f_ν via equation (4.1).

extend their model to account for non-linear neutrino influences by modifying the parameters governing spherical collapse, as this will be directly impacted by the reduction in matter clustering. Fitting to the simulations of Massara et al. (2014) reproduces $P(k)$ to within a few percent up to $k = 10\,h\,\mathrm{Mpc}^{-1}$ for multiple redshifts. This is a slightly better performance than the fitting formula provided by Bird et al. (2012). The CAMB Boltzmann code (Lewis et al., 2000) is used to generate linear $P(k)$ for the NH and IH.

Massive neutrinos impact the matter distribution primarily by damping matter perturbations as they free-stream through structure. For fixed Ω_m[1], the suppression of matter clustering is well-approximated as depending only on f_ν. This affects the spherical collapse overdensity, δ_c, and the virial density, Δ_v, which are characteristic descriptors of halo structure. The former affects the shape of the mass function and determines the halo concentration by defining the formation redshift at which a certain fraction of mass has undergone spherical collapse. The latter defines limits on the density profile. By introducing a linear correction in terms of f_ν, such that

$$\delta_c \propto 1 + 0.262 f_\nu \tag{4.1}$$

and

$$\Delta_v \propto 1 + 0.916 f_\nu, \tag{4.2}$$

M16 achieve a percent-level fit to N-body massive neutrino and CDM simulations by Massara et al. (2014) up to $k = 10\,h\,\mathrm{Mpc}^{-1}$. These relations can be interpreted as the damping of matter fluctuations due to neutrino free-streaming resulting in a smaller fraction of perturbations collapsing by a given redshift, equivalent to raising the density threshold for collapse. Potential improvements could be attained by implementing more sophisticated methods by Ichiki and Takada (2012) or LoVerde (2014) in which, for example, non-linear neutrino clustering effects are considered. However, the M16 fit is sufficiently accurate for my purposes and permits one to use f_ν as a free parameter in fast power spectrum calculations for forecasts. In practice, the mass sum, Σ, is treated as the actual free parameter, via the relation,

$$f_\nu = \frac{\Omega_\nu}{\Omega_m} \simeq \frac{\Sigma}{94.1\,\Omega_m h^2\,\mathrm{eV}}. \tag{4.3}$$

[1] I define the total matter density parameter, $\Omega_m = \Omega_c + \Omega_b + \Omega_\nu$, as the sum of contributions from cold dark matter, baryons and neutrinos.

The Massara et al. (2014) simulations are designed to test the halo model with neutrinos, specifically by providing accurate halo power spectrum responses to different neutrino masses. This makes the M16 fits to these simulations particularly relevant for the purposes of this Chapter, in which an accurate power spectrum is less important than the sensitivity of the power spectrum to changes in cosmology for making forecasts. Other implementations of neutrinos in the halo model (e.g., Takahashi et al., 2012; Bird et al., 2012) generate slightly different non-linear responses in the power spectrum. It should be noted that none of these treatments is entirely robust, being subject to the limited accuracy of the underlying halo model fit on non-linear scales and the challenge of properly accounting for neutrino non-linear clustering effects. There may be future scope to improve on these approaches by extending the methodology of Mead (2017) and Cataneo et al. (2018), who demonstrate the increased accuracy on non-linear scales available to the halo model by incorporating spherical collapse and growth results into power spectrum responses.

In Figure 4.1 I show the impact of varying Σ on the halo mass function with and without the modification to δ_c in equation (4.1). Neutrinos have a very small impact on the matter distribution on non-linear scales so I have illustrated it through the ratios of the mass function to the case with the minimal NH neutrino mass sum, $\Sigma_{\rm NH,min} = 0.06\,{\rm eV}$.

Considering first the case with δ_c unmodified, I find that for fixed Ω_m and σ_8, increasing the neutrino mass increases the mass function for high halo masses in the top panel of Figure 4.1. This reflects the fact that the suppression of power on small scales from neutrino free streaming requires that linear power is boosted on large scales to keep σ_8 fixed. As a result, the variance of the matter distribution is increased for the corresponding large mass scales. This is the regime where $\sigma\left(M\right)$ is smallest, and hence where ν is largest. At these scales the mass function is dominated by an exponential tail, as shown in equation (2.23). This leads to small changes in ν, from varying Σ, generating large changes in $n\left(M\right)$.

The non-linear regime begins on smaller mass scales where the clustering is damped by higher Σ, reducing the probability of matter fluctuations being dense enough to collapse and contribute to $n\left(M\right)$. Equation (2.23) is normalised to the mean density, so in order to satisfy the condition that integrating the mass

function over all scales returns $\bar{\rho}$, massive neutrinos amplify $n(M)$ for very small M. However, these are highly non-linear scales that are inaccessible to Stage IV surveys and so are of limited interest. The higher threshold for collapse in equation (4.1) suppresses the number density of high-mass haloes. This will propagate into a reduction of matter power on the weakly non-linear scales corresponding to such objects. The suppression of clustering due to Σ increases by a factor of ~ 5 with this change to δ_c, although it should be emphasised that the overall effect of massive neutrinos on the mass function remains very small.

The impact of massive neutrinos on the density profile has a less significant (but still non-trivial) contribution to the matter power spectrum compared to changes in the mass function. The collapse overdensity partially controls the formation redshift determining the concentration relation in equation (3.3). At late times, regions are more likely to collapse due to the amplitude of matter fluctuations being greater. Increasing δ_c makes halo formation rarer, favouring later times over earlier times, and therefore lowering z_f. The resulting reduction in halo concentration translates to an increased scale radius for fixed r_v. Consequently, for a fixed halo mass, the density is suppressed on small scales, with a milder boost on large scales. However, when the reduction of r_v through increasing Δ_v via Σ is also accounted for, this picture is complicated. If this effect is considered in isolation, for fixed mass and concentration, the density on small scales must increase to reflect the fact that haloes are smaller when meeting the increased density threshold for virialisation. Interestingly, when the sensitivity to both δ_c and Δ_v is included, the two contributions mostly cancel, leaving the halo density profile with minimal sensitivity to the neutrino mass on small scales. The relative insensitivity of the density profile, and more importantly the mass function, highlights, at the outset, the challenges of using the non-linear regime to constrain Σ.

4.3 Intrinsic alignments

In previous Chapters the formalism of gravitational lensing has been discussed and how the Fisher formalism is applied to that probe. In this Chapter, the weak lensing treatment is extended to include intrinsic alignments. This section pro-

vides a brief discussion of the background formalism and the choice of alignment model that is made when incorporating IA into the weak lensing convergence power spectrum. Although the observable is modified, the Fisher formalism itself used for weak lensing is the same as presented in Chapter 3.

The tidal forces that correlate ellipticities, ϵ, and shears, γ, are an intrinsic contribution to the observed signal. Accounting for this to linear order by writing the observed ellipticity as $\epsilon = \gamma + \epsilon^I$, the weak lensing power spectrum becomes

$$C_{\ell,ij} = C^{\gamma\gamma}_{\ell,ij} + C^{I\gamma}_{\ell,ij} + C^{II}_{\ell,ij}, \tag{4.4}$$

where $C^{\gamma\gamma}_{\ell,ij}$ is given by equation (2.3) and the term, $C^{I\gamma}_{\ell,ij}$, encompasses correlations of foreground ellipticities with background shear and of foreground shear with background ellipticities. The latter correlations should be very small under the assumption that redshift bins do not overlap. The IA-shear and IA-IA terms are given by

$$C^{I\gamma}_{\ell,ij} = \frac{3}{2}\Omega_m \left(\frac{H_0}{c}\right)^2 \int_0^{\chi_{\max}} \mathrm{d}\chi \, \frac{g_i(\chi)\, n_j(\chi) + g_j(\chi)\, n_i(\chi)}{a(\chi)\,\chi} \times P_{mI}\left(k = \frac{\ell}{f_K(\chi)}, \chi\right), \tag{4.5}$$

and

$$C^{II}_{\ell,ij} = \int_0^{\chi_{\max}} \mathrm{d}\chi \, \frac{n_i(\chi)\, n_j(\chi)}{\chi^2} P_{II}\left(k = \frac{\ell}{f_K(\chi)}, \chi\right), \tag{4.6}$$

where

$$P_{mI}(k,\chi) = -a_{\mathrm{IA}} \frac{c_{\mathrm{IA}} \rho_{\mathrm{crit}} \Omega_m}{D(z(\chi))} P_m(k,\chi) \tag{4.7}$$

and

$$P_{II}(k,\chi) = \left(a_{\mathrm{IA}} \frac{c_{\mathrm{IA}} \rho_{\mathrm{crit}} \Omega_m}{D(z(\chi))}\right)^2 P_m(k,\chi) \tag{4.8}$$

are the matter-IA and IA-IA source power spectra in the non-linear alignment model (Bridle and King, 2007) and $D(z)$ is the linear growth factor. The amplitude, a_{IA}, can be taken as a free parameter with a fiducial value of unity, while fits to SuperCOSMOS galaxy data determine the normalisation, $c_{\mathrm{IA}} \approx 0.0134/\rho_{\mathrm{crit}}$, in terms of the critical density (Brown et al., 2002; Hirata and Seljak, 2004). In Figure 4.2 I show comparisons of the shear-shear, IA-shear and IA-IA auto- and cross- lensing power spectra for several combinations of redshift bins. The

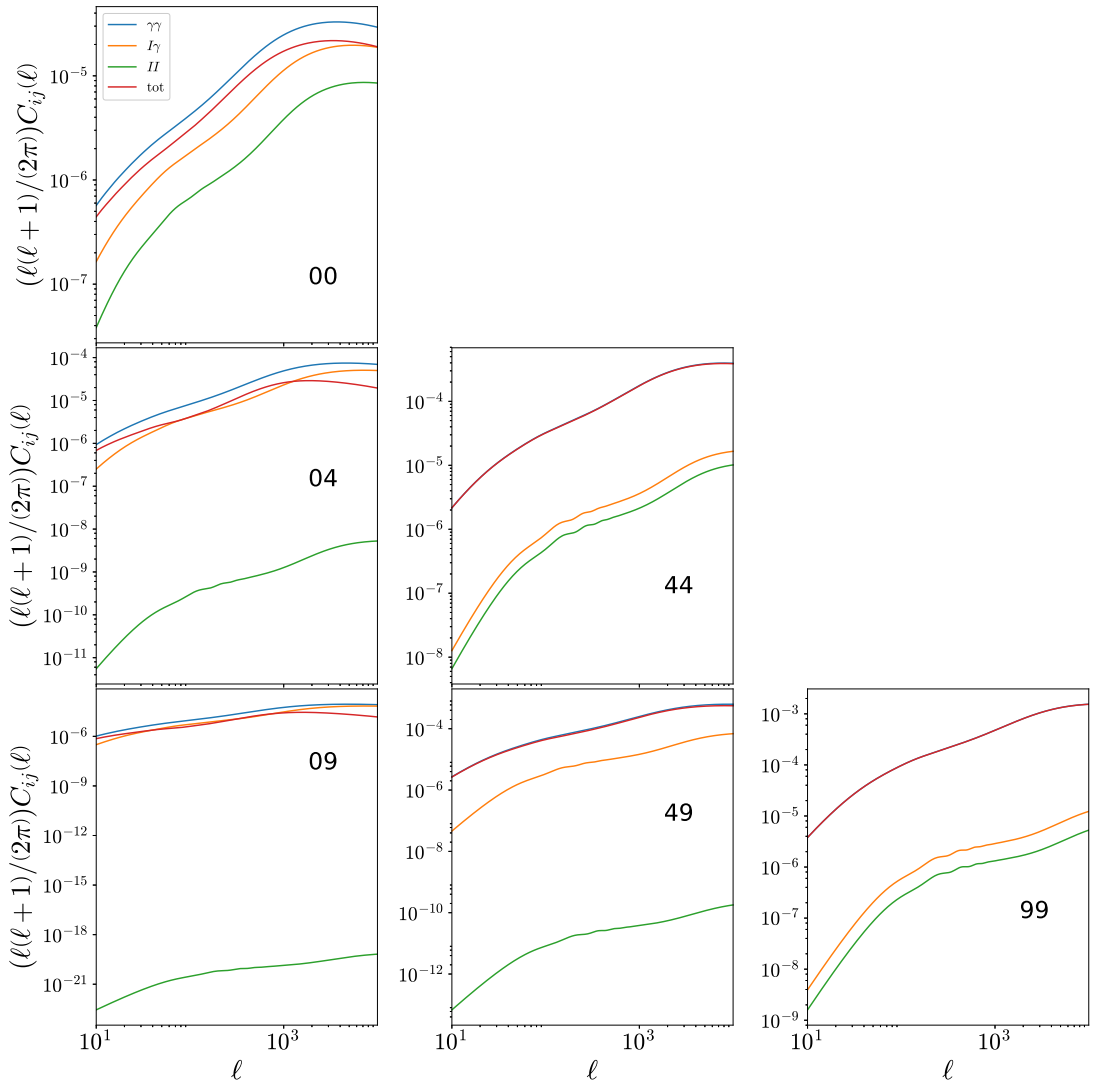

Figure 4.2: Auto- and cross- lensing power spectra in the 00, 44 and 99 redshift bins: $0 < z \leq 0.41$, $0.79 < z \leq 0.89$, $1.52 \leq z \leq 2.0$. Blue: shear-shear power spectrum; orange: absolute value of the IA-shear power spectrum; green: IA-IA power spectrum; red: total power spectrum.

IA-IA contributions to the signal are most significant for low redshift auto-power spectra, while the IA-shear term has its greatest impact on cross-power spectra of widely separated bins because it describes the correlations of foreground galaxy ellipticities with background shear.

4.4 Fisher formalism for galaxy clustering

The Fisher formalism in the context of weak lensing has been explored thoroughly in the preceding Chapter. However, for galaxy clustering there are some further substantial complexities and nuances to address, particularly pertaining to the treatment of nuisance parameters and projections between different parameter sets. Here a brief overview is provided.

The Fisher matrix for galaxy clustering (Seo and Eisenstein, 2003),

$$F_{\alpha\beta} = \frac{1}{8\pi^2} \sum_{i=1}^{N_{z,\mathrm{GC}}} \int_{-1}^{1} \mathrm{d}\mu \int_{k_{\min}}^{k_{\max}} k^2 \, \mathrm{d}k \times \frac{\partial \ln P_{gg}\left(k, \mu; z_i\right)}{\partial \theta_\alpha} V_{\mathrm{eff}}\left(k, \mu; z_i\right) \frac{\partial \ln P_{gg}\left(k, \mu; z_i\right)}{\partial \theta_\beta}, \tag{4.9}$$

is computed by integrating over the angular and radial contributions of the logarithmic galaxy power derivatives across $N_{z,\mathrm{GC}}$ redshift bins. In this respect the Fisher calculation is more straightforward than for weak lensing. For galaxy clustering a reasonable assumption can be made that cross-correlations between redshift bins will have a limited impact on the results, so they can be safely neglected. This allows information across the range of redshift bins to be accumulated through a simple summation of contributions. In principle, this make it easier to disentangle important features in the results compared to weak lensing which must account for covariance between bins.

An upper bound of $k_{\max} = 0.2\,h^{-1}\mathrm{Mpc}$ is chosen for the calculation in order to restrict the GC forecasts made in this Chapter to linear scales. It is beyond the scope of this work to address the impact of non-linear systematics like Finger-of-God effects. The choice of lower bound has very limited bearing on results so

I use the common value, $k_{\rm min} = 0.001\,h^{-1}$Mpc. The effective volume,

$$V_{\rm eff}\left(k,\mu;z_i\right) = V_{\rm survey}\left[\frac{1}{P_{gg}\left(k,\mu;z_i\right)+\frac{1}{n(z)}}\right]^2. \tag{4.10}$$

encodes the covariances and is calculated from the survey volume, V_s, and the number density, $n\left(z\right)$, of Hα-emitters, for which Euclid survey values provided by the reference case in Table 3 of Amendola et al. (2018) are used.

Note that all scales and angles refer to the 'fiducial' case, although this label has been dropped for convenience. Free parameters governing the AP effect, growth, bias and residual shot noise within each redshift bin independently, $\{D_A\left(z_i\right), H\left(z_i\right), f\sigma_8\left(z_i\right), b\sigma_8\left(z_i\right), P_{\rm shot}\left(z_i\right)|i=1,...,N_{z,\rm GC}\}$, can then be propagated in a Fisher analysis alongside redshift-independent parameters, $\{\omega_m,\omega_b,\omega_\nu,h,n_s\}$, controlling the shape of $P\left(k\right)$. The physical densities are related to the original density parameters via $\{\omega_i=\Omega_i\,h^2|i=m,b,\nu\}$. I marginalise over the nuisance parameters, $\{b\sigma_8\left(z_i\right), P_{\rm shot}\left(z_i\right)|i=1,...,N_{z,\rm GC}\}$, by inverting the Fisher matrix, removing rows and columns corresponding to these parameters, and then inverting the remaining sub-block. This follows a prescription first detailed by Seo and Eisenstein (2003). The resulting Fisher matrix, $\tilde{F}$, treats $\{D_A\left(z_i\right), H\left(z_i\right), f\sigma_8\left(z_i\right)|i=1,...,N_{z,\rm GC}\}$ as free parameters contributing information alongside $\{\omega_m,\omega_b,\omega_\nu,h,n_s\}$. Collecting all these parameters into the set, $\tilde{\Theta}$, a final Fisher matrix is derived for the wCDM parameter set, $\Theta = \{\Omega_m,\Omega_b,\Omega_\nu,h,n_s,\sigma_8,w\}$, by performing the transformation,

$$F = J^T\tilde{F}J, \tag{4.11}$$

where $J = \partial\tilde{\Theta}/\partial\Theta$ is the Jacobian matrix of parameter derivatives. Note that because $\tilde{\Theta}$ is evaluated in the wCDM model it can be completely determined by Θ.

I make the decision to evaluate the Fisher information available to a Euclid-like spectroscopic galaxy survey in the redshift range $0.75 < z \leq 2.05$, with bin width $\Delta z = 0.1$. For lower redshifts I refer to BOSS as it provides superior information at these z than Euclid-like surveys are designed to achieve. This entails using covariance information provided by Alam et al. (2017) on $\{D_A\left(z_i\right), H\left(z_i\right), f\sigma_8\left(z_i\right)\}$, where the i label runs over three partially overlap-

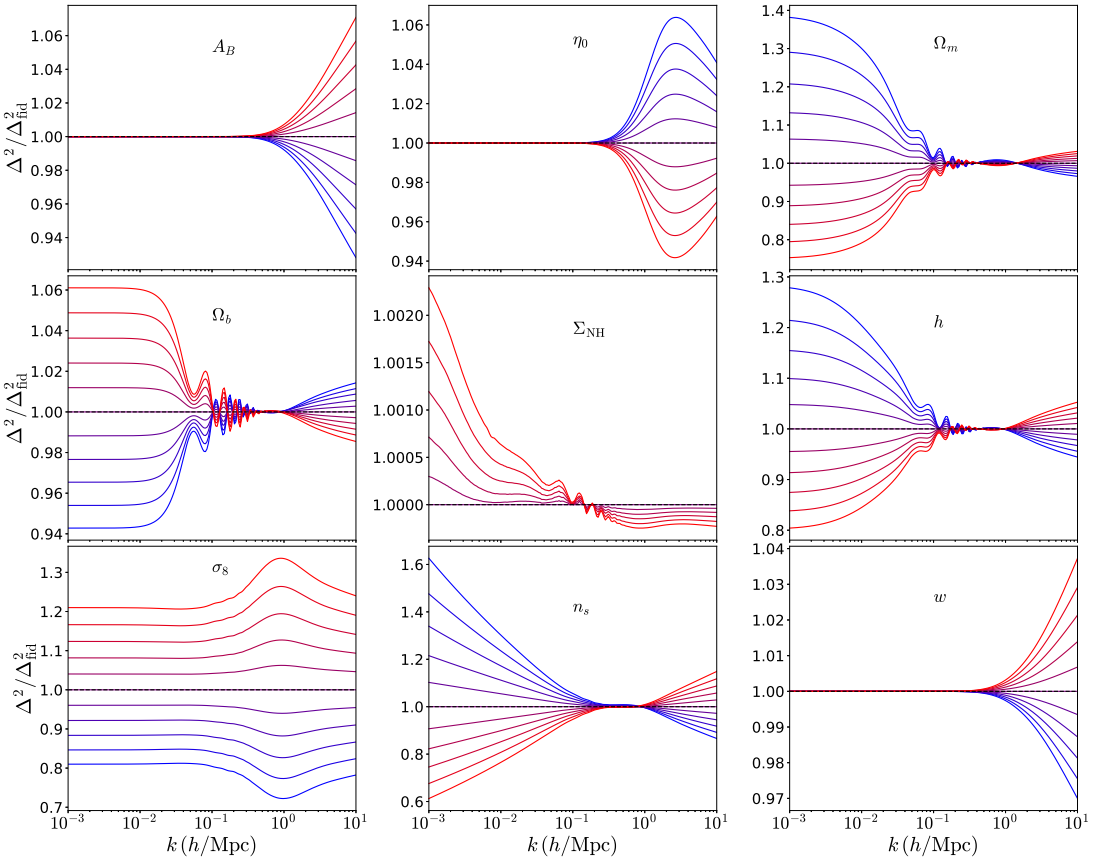

Figure 4.3: The ratio of matter power spectra at $z = 0$ using NH for different iterations of parameters in $\Theta = (A_B, \eta_0, \Omega_m, \Omega_b, \Sigma_{\rm NH}, h, \sigma_8, n_s, w)$, with respect to a fiducial power spectrum computed with parameter values found by Planck Collaboration et al. (2016b). Bluer (redder) curves correspond to lower (higher) values for parameters in the range $0.9\,\Theta_{\rm fid} \leq \Theta \leq 1.1\,\Theta_{\rm fid}$, except in the case of the neutrino mass parameter, which varies between $\Sigma_{\rm NH,min} \leq \Sigma_{\rm NH} \leq 1.1\,\Sigma_{\rm NH,min}$ with purple curves representing the minimal mass case, $\Sigma_{\rm NH,min} = 0.06\,{\rm eV}$. Chapter 3 showed similar $\Delta^2\,(k)$ response plots but did not include massive neutrinos.

ping redshift bins, spanning the total range $0.2 < z \leq 0.75$. A matrix inversion is performed to derive the corresponding Fisher information matrix. The BOSS bins overlap with the full z range available to the Euclid-like survey, so to ensure that covariance between the surveys is limited I only use redshift bins for the latter that lie beyond the range of BOSS. The final galaxy clustering Fisher matrix is then evaluated from the combination of Euclid-like high z forecasts and BOSS low z data.

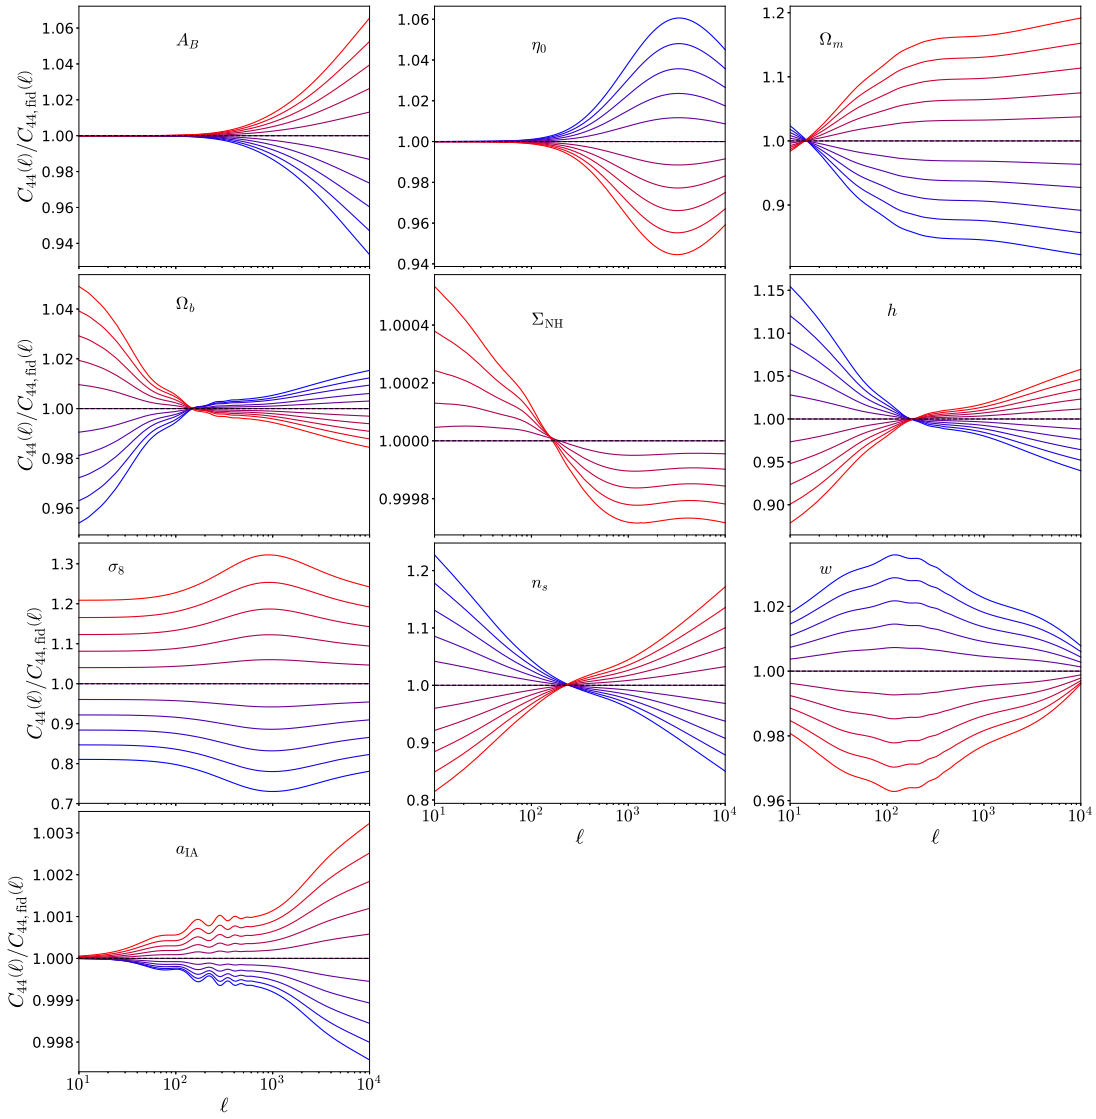

Figure 4.4: Lensing power spectrum responses using the NH in the $ij = 44$ redshift bin $0.79 \leq z \leq 0.89$. Blue (red) lines correspond to the lowest (highest) parameter values for $\Theta = (A_B, \eta_0, \Omega_m, \Omega_b, \Sigma_{\rm NH}, h, \sigma_8, n_s, w, a_{\rm IA})$ in the range $0.9\Theta_{\rm fid} \leq \Theta \leq 1.1\Theta_{\rm fid}$, except in the case of the neutrino mass parameter, which varies between $\Sigma_{\rm NH,min} \leq \Sigma_{\rm NH} \leq 1.1\,\Sigma_{\rm NH,min}$ with purple curves representing the minimal mass case, $\Sigma_{\rm NH,min} = 0.06\,{\rm eV}$. Chapter 3 showed similar C_ℓ response plots but did not include massive neutrinos.

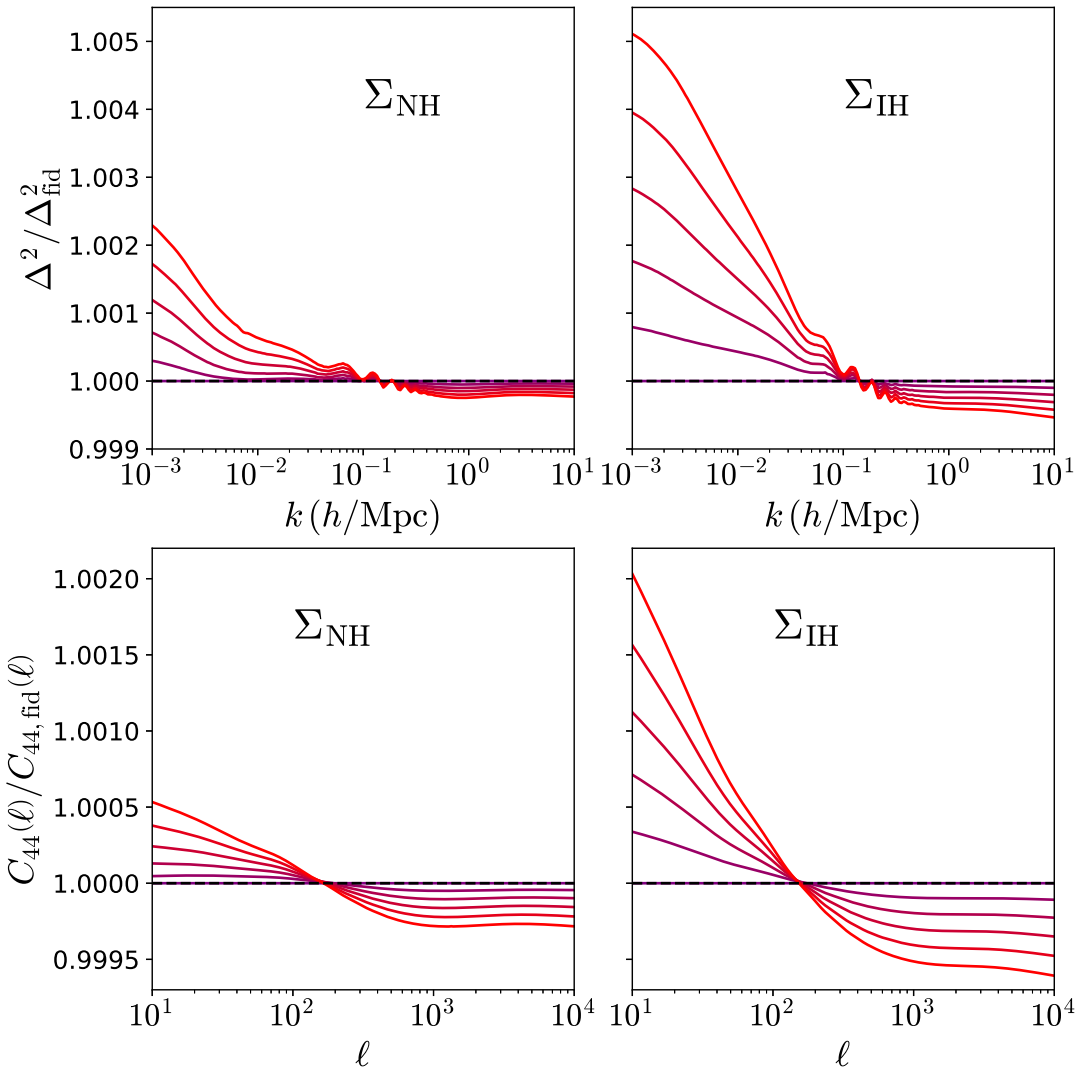

Figure 4.5: Top row: The ratio of matter power spectra at $z = 0$ for different iterations of Σ using the NH (left) and IH (right), with respect to a fiducial power spectrum computed with parameter values found by Planck Collaboration et al. (2016b). Bottom row: Lensing power spectrum responses for Σ using NH (left) and IH (right) in the $ij = 44$ redshift auto-bin $0.79 \leq z \leq 0.89$. The mass sum is varied in the range $\Sigma_{\rm min} \leq \Sigma \leq 1.1\,\Sigma_{\rm min}$ for both hierarchies with purple curves representing the minimal mass cases, $\Sigma_{\rm NH,min} = 0.06\,{\rm eV}$ and $\Sigma_{\rm IH,min} = 0.1\,{\rm eV}$.

4.5 Neutrino Mass Forecasts

Fisher matrices are constructed for the final parameter set, $(A_B, \eta_0, \Omega_m, \Omega_b, \Sigma, h, n_s, \sigma_8, w, a_{\rm IA})$, for Euclid-like weak lensing and spectroscopic galaxy surveys. The primary goals of this work are to assess the impacts of IA (a phenomenon exclusive to weak lensing) and baryonic effects on neutrino mass forecasts. As the influence of baryons degrades information significantly on scales $k > 1h\,\mathrm{Mpc}^{-1}$ this effect is of far greater concern to weak lensing forecasts than those from galaxy clustering. Hence, the analysis here takes the approach of first discussing weak lensing results and the extent of their degradation due to baryons and IA. Information from *Planck* and galaxy clustering is then added to form a more complete picture of the prospects of a Euclid-like survey to measure the neutrino mass with sufficient accuracy to distinguish the hierarchies.

4.5.1 Normal and inverted hierarchy results from weak lensing

In Table 3.2, I state the weak lensing survey parameters specified by the Euclid survey report (Laureijs et al., 2011). $N_{z,\rm WL} = 10$ redshift bins are chosen in the range $0 \leq z \leq 2$ such that each bin contains an equal number density of galaxies. A large range of scales from $\ell_{\rm min} = 10$ to $\ell_{\rm max} = 5000$ are covered so in practice I compute the summation in equation (3.23) at logarithmic intervals. I include full confidence ellipse plots for the NH in Figure 4.6.

The fiducial values of the cosmological parameters in this work are again given by the base ΛCDM Planck TT,TE,EE+lowP likelihood in Planck Collaboration et al. (2016b), summarised in Table 3.1. The Euclid-like weak lensing survey parameters have been given in Table 3.2

This Chapter applies the same philosophy as the preceding one of gaining intuition into Fisher forecasts by examining the response of the power spectra to varying parameters across relevant scales. This diagnostic can quickly indicate which scales most Fisher information comes from or is potentially lost due to degenerate power spectrum responses between parameters. In Figures 4.3 and 4.4 I show, for the normal hierarchy, responses of $\Delta^2(k)$ at $z = 0$ and C_ℓ respectively to varying each parameter with respect to its fiducial value while fixing the other

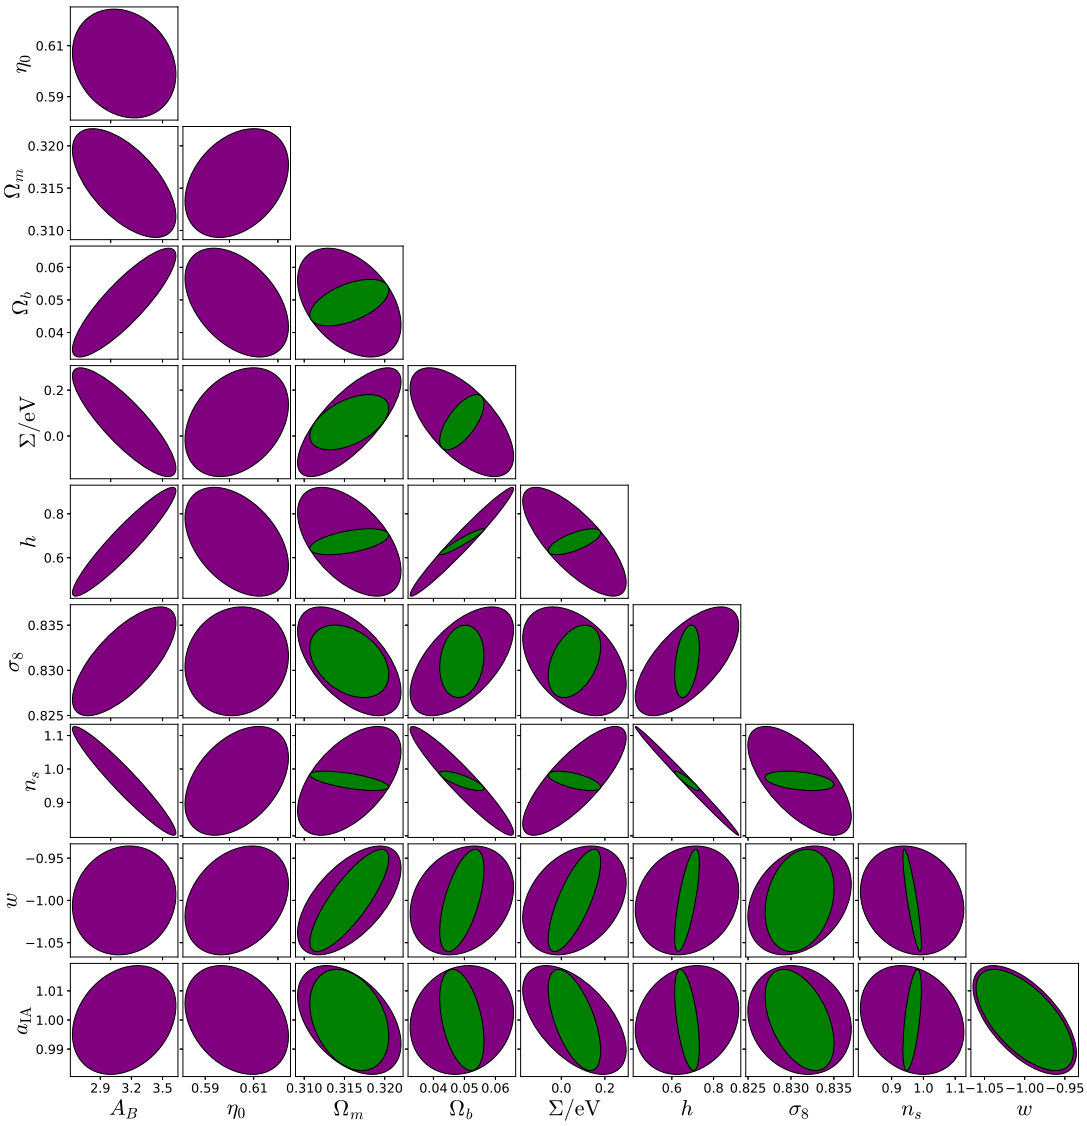

Figure 4.6: 1-σ 2-parameter confidence ellipses for the normal hierarchy with all parameters in $\Theta = (A_B, \eta_0, \Omega_m, \Omega_b, \Sigma_{\rm NH}, h, \sigma_8, n_s, w, a_{\rm IA})$ marginalised over (purple); and with the baryon parameters, A_B and η_0, fixed to their fiducial values (dark green).

parameters. In § 4.5.2 the responses to the baryon parameters are discussed in detail.

$\Delta^2(k)$ responses to most parameters exhibit nodes at the scale corresponding to σ_8, at which I normalise power. For example, the response to Σ is an enhancement of power on linear scales in order to satisfy power being damped in the non-linear regime due to free streaming. The impact of neutrinos on these scales is small, determined by the limited sensitivity of the mass function and the halo structure parameters, δ_c and Δ_v, to Σ (discussed in detail in § 4.2). Most of the C_ℓ responses reflect the corresponding $\Delta^2(k)$ responses, when accounting for redshift-dependent effects along the line of sight. However, parameters that affect the lensing weight function through their impact on cosmological distances can exhibit significantly different $\Delta^2(k)$ and C_ℓ responses. For example, the competing influences of geometry and growth are apparent in the w responses, which is examined in depth in Chapter 3.

I compare the Σ responses for $\Delta^2(k)$ and C_ℓ for the NH and IH directly in Figure 4.5. $\Sigma_{\rm IH,min} = 0.1\,{\rm eV}$ is less than double $\Sigma_{\rm NH,min} = 0.06\,{\rm eV}$ but the IH exhibits power responses that are more than twice as large as the NH responses. The leading order cause of this is likely a dependence of power on Σ that grows with the mass sum from a quadratic minimum which is reached when Σ approaches zero.

An additional effect to account for is the influence of the neutrino free-streaming scale, $k_{\rm fs}$, which increases with mass. There are different $k_{\rm fs}$ for each mass eigenstate in a hierarchy but, within the approximation of the NH as $\{m_1 \approx m_2 \approx 0\,{\rm eV}, {\rm m}_3 \approx 0.06\,{\rm eV}\}$ and the IH as $\{m_1 \approx m_2 \approx 0.05\,{\rm eV}, {\rm m}_3 \approx 0\,{\rm eV}\}$, the most relevant free-streaming scales satisfy $k_{\rm fs,IH} < k_{\rm fs,NH}$. Fixing σ_8 means that, for increases in Σ, large-scale power is enhanced relative to the fiducial case to satisfy the suppression of power by neutrinos on small scales. There is a similar contribution from the peak of $P(k)$ shifting with a decrease in the redshift of matter-radiation equality, which depends on Ω_c (which decreases to keep Ω_m constant for increasing Σ) due to the relativistic nature of neutrinos during this epoch. As the effect of free-streaming neutrinos arises at smaller k for the IH, the relative enhancement of the large-scale IH power is greater than for the NH. Similarly, the responses for the IH are more sensitive to the same fractional change

in Σ compared to the NH, which I see reflected in Figure 4.5.

When propagating $\Delta^2(k)$ into C_ℓ, the lowest ℓ for a Euclid-like survey correspond to k where the matter power response is weak. As a result, the strongest linear responses for C_ℓ in Figure 4.5 are several factors smaller than for $\Delta^2(k)$. Other parameter responses generally experience significantly more modest changes from one probe to another. It should be noted that the background geometry is not sensitive to increasing Σ when neutrinos are non-relativistic; the angular distances entering equation (2.3) depend on Ω_m, which is kept constant.

A consequence of the weak lensing probe only being sensitive in the linear regime to scales where the neutrino response is most limited is that the magnitude of the C_ℓ response is similar between the linear and non-linear regimes. The response for the latter is due mainly to the modest increases in δ_c and Δ_v. This highlights the importance of pursuing robust physical modelling for neutrino effects on structure at these scales because they are as sensitive to Σ as linear scales.

Weak lensing forecasts of Σ are particularly challenging due to this limited sensitivity and the degeneracies apparent between Σ and a range of other parameters on non-linear scales. A well-known example is the degeneracy with σ_8, in which both exhibit a 'spoon'-shaped response (see e.g., Massara et al., 2014). An increase in Σ most strongly affects high mass haloes, which dominate over small mass haloes on intermediate scales. Figure 4.1 shows that in this region higher Σ reduces the linear clustering of matter, and therefore the number of high mass haloes that can form, lowering the halo mass function and hence the non-linear matter power. Raising the collapse and virialisation density thresholds with Σ deepens the resulting dip in the power response and shifts it to larger scales associated with higher masses. At higher k the impact of Σ is somewhat decreased and becomes less scale-dependent.

The bump in the σ_8 response is qualitatively similar. By altering the clustering amplitude, a halo mass-dependent and hence scale-dependent power response is induced through the change in the peak height, $\nu = \delta_c/\sigma(M, z)$, a quantity that is also sensitive to Σ. The response becomes roughly scale-independent on smaller scales. This is similar to the Σ response but with the opposite sign, enhancing rather than suppressing power. As most information comes from small

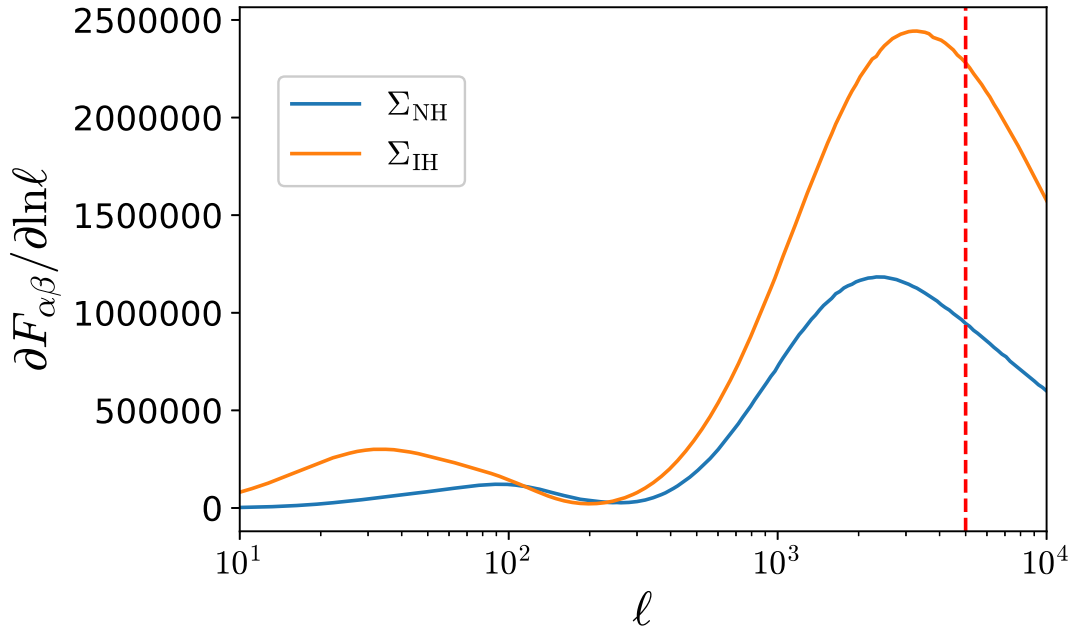

Figure 4.7: Contributions to the ΣΣ Fisher information element for the NH (blue) and IH (orange) at each ℓ-mode, weighted by ℓ. The vertical line (red, dashed) corresponds to the conventional cutoff for Euclid-like surveys at $\ell_{\rm max} = 5000$. The Fisher derivative with respect to $\ln \ell$ is used in order to convey the integration area for the logarithmically spaced ℓ-range.

scales this degeneracy is significant.

In Figure 4.7 the contribution per $\ln \ell$ to the Fisher matrix for the mass sum is shown for the NH and IH. This effectively weights the Fisher information by the signal-to-noise in each ℓ mode, which increases with ℓ. The greater number of independent modes in the non-linear regime, and consequently lower noise, results in most of the available information being drawn from these scales. The greater sensitivity of information at high ℓ highlights the limitations on achieving strong constraints, given the limited neutrino influence on non-linear power. The IH contributes almost twice as much information as the NH at the peak sensitivity.

The main difference between the hierarchies in their information contribution is from different logarithmic derivatives of C_ℓ. These are partially predicted by the power responses. On non-linear scales the C_ℓ response for the IH is larger than double the NH for the same fractional change in Σ. As $\Sigma_{\rm IH,min} = 0.1\,{\rm eV}$ is

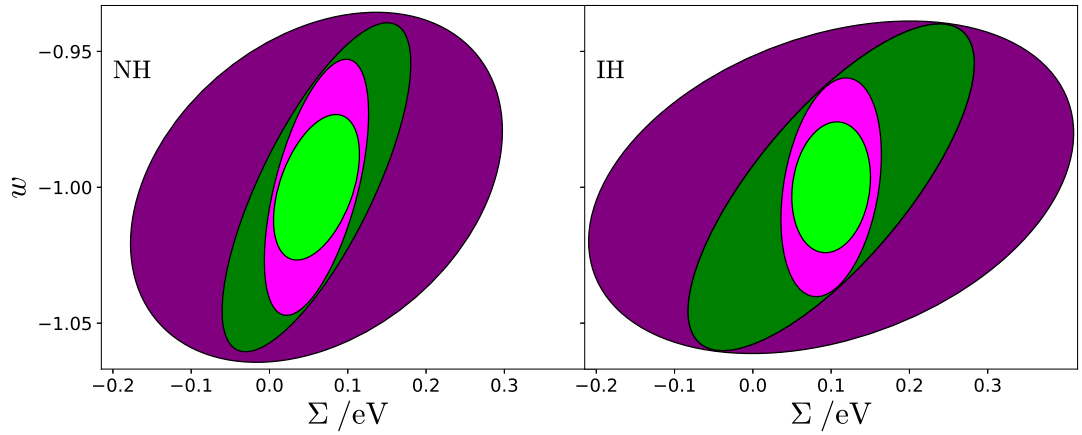

Figure 4.8: 1-σ w-Σ confidence ellipses for the normal (left; $\Sigma_{\rm fid} = 0.06\,{\rm eV}$) and inverted hierarchy (right; $\Sigma_{\rm fid} = 0.1\,{\rm eV}$) for a Euclid-like weak lensing survey with all parameters in $\Theta = (A_B, \eta_0, \Omega_m, \Omega_b, h, \sigma_8, n_s, w, a_{\rm IA})$ marginalised over (purple); and with the baryon parameters, A_B and η_0, fixed to their fiducial values (dark green). *Planck* CMB priors are included for the magenta (light green) ellipses with (without) baryon marginalisation.

	$\sigma_{\rm WL}$/eV (ex bar.)	$\sigma_{\rm WL}$/eV (inc. bar.)	$R_{B,{\rm WL}}$	$\sigma_{\rm WL+CMB}$/eV (ex bar.)	$\sigma_{\rm WL+CMB}$/eV (inc. bar.)	$R_{B,{\rm WL+CMB}}$
$\Sigma_{\rm NH}$	0.079	0.156	1.97	0.036	0.044	1.21
$\Sigma_{\rm IH}$	0.120	0.204	1.70	0.033	0.042	1.28

Table 4.1: 1-σ error forecasts for a Euclid-like survey of the neutrino mass sum in the normal and inverted hierarchies, without and including marginalisation over baryon parameters and the addition of priors on the ΛCDM cosmological parameters from *Planck* CMB measurements. I also include the response factors, R_B, to including baryon marginalisation.

less than double $\Sigma_{\rm NH,min} = 0.06\,{\rm eV}$, this leads to $\partial \ln C_\ell / \partial \Sigma_{\rm IH}$ being sufficiently larger than $\partial \ln C_\ell / \partial \Sigma_{\rm NH}$ that the Fisher sensitivity, which depends quadratically on the logarithmic derivative, is up to twice as large.

Figure 4.8 shows NH and IH confidence ellipses for w-Σ for a Euclid-like weak lensing survey. It is not expected that Σ can be well constrained from WL alone, but I am interested in using it as a baseline probe, through which I can examine both the impact of systematics like baryons and IA and the improvements available from adding CMB and galaxy clustering information. Before marginalising over baryons, WL errors of $\sigma_{\Sigma_{\rm NH}} = 0.079\,{\rm eV}$ and $\sigma_{\Sigma_{\rm IH}} = 0.120\,{\rm eV}$ (see by Table 5.1) are found. These are larger than the threshold, $|\Delta\Sigma_{\rm min}|=$

$|\Sigma_{\rm NH,min} - \Sigma_{\rm IH,min}| = 0.04\,{\rm eV}$, below which hierarchies can be distinguished, and are indeed too large to permit any positive detection of neutrinos.

The latter issue can be alleviated by adding priors on cosmology from sources like galaxy clustering or the early Universe. I focus here on the latter through *Planck* CMB anisotropy measurements, while impact of galaxy clustering is discussed in detail in S 4.5.4. The CMB provides a wealth of information on the matter energy-density and the geometry of the Universe, which is one of the main contributions to LSS signals. I use the publicly available MCMC chains for the base $\nu\Lambda$CDM combined TT, TE and EE power spectra (see Planck Collaboration et al., 2016b) to construct a covariance matrix which is inverted into a Fisher matrix, $F_{\rm CMB}$, that encapsulates the information available from the CMB.

As can be seen in Figure 4.8 there is a substantial improvement to Σ forecasts for the NH and IH, with errors reduced below the distinction threshold to $\sigma_{\Sigma_{\rm NH}} = 0.036\,{\rm eV}$ and $\sigma_{\Sigma_{\rm IH}} = 0.033\,{\rm eV}$. This is partly a result of parameter degeneracies being broken, for example the CMB provides strong constraints on σ_8 which propagate through to my final Σ forecasts. It should be noted that this requires CMB experiments to accurately measure the reionisation optical depth, τ, as it is degenerate with the primordial power amplitude. I marginalise over τ, along with a range of other parameters, when constructing $F_{\rm CMB}$. Cosmological neutrino measurements stand to benefit significantly from future 21 cm experiments that will probe the reionisation epoch with high precision (Allison et al., 2015; Liu et al., 2016). However, this work does not extend to examine this in detail.

The CMB prior provides a significant contribution to the forecasts by breaking parameter degeneracies and, as it is the same for both hierarchies, it leads to less difference between the NH and IH for the WL+CMB forecasts. This is reflected in the weaker WL constraint, $\sigma_{\Sigma_{\rm IH}}$, which improves by 72%, benefiting significantly more than $\sigma_{\Sigma_{\rm NH}}$, which improves by 54% (see Table 5.1). My results are comparable to findings from Audren et al. (2013) who use an MCMC approach to explore the degenerate hierarchy case. However, even these improvements would only achieve a close to 2σ detection of Σ in the case of the NH (although the IH case fares better with a 3σ detection) and a roughly 1σ determination of the hierarchy.

4.5.2 Impact of baryons

When the baryon systematic is accounted for, the prospects of using weak lensing to accurately distinguish hierarchies recede. Figures 4.3 and 4.4 include the $\Delta^2(k)$ and C_ℓ responses to the adiabatic contraction parameter, A_B, and the baryonic feedback parameter, η_0. The preceding Chapter discusses these in detail but they are briefly highlighted again here. The boost to non-linear power from increasing A_B is a reflection of the enhanced halo density profiles in this regime. The response to η_0 exhibits a peak, which is a more subtle consequence of capturing the effects of baryonic feedback over a range of mass and spatial scales. The bloating of higher-mass haloes is the dominant influence compared to the reduction effect experienced by lower-mass haloes, so decreasing η_0 generates a net increase of $\Delta^2(k)$. The peak occurs at deeper non-linear scales for higher redshifts, reflecting the evolution of halo populations.

Baryons and neutrinos both impact the amplitude of the matter power spectrum on small scales, from which most information is drawn, leading to significant degeneracies between both baryon parameters and Σ. For example, the mass-dependent response for η_0 emerges from its effect on the peak height which, as discussed above, determines non-linear halo properties in a similar manner to the neutrino mass sum. In the M16 model, the peak height depends on Σ directly. The total baryon impact on the Σ forecast is shown in Figure 4.8 to be severe. The errors for NH and IH double for the former and increase by 70% for the latter. The greater sensitivity of $\Sigma\Sigma$ Fisher information for the IH relative to the NH on non-linear scales where the baryons limit information results in this comparatively lesser degradation.

In the preceding Chapter, several strategies for mitigating baryon impacts on dark energy forecasts were explored. These included increasing the Euclid-like survey limit from $\ell_{\rm max} = 5000$ to $\ell_{\rm max} = 10000$ to access information from modes deeper in the non-linear regime. When this is done for the Σ forecasts, I find that baryon degradation is reduced, to 84% and 61% for the NH and IH respectively. Figure 4.7 shows that there is a repository of potential Fisher information on Σ at higher ℓ. The response plots suggest that the degeneracies between A_B, η_0 and Σ may be less severe for very high ℓ where the neutrino impact on very small haloes approaches scale-independence (Massara et al., 2014) in contrast to the baryons.

This may explain why the degradation is reduced, despite the increasing influence of baryons on non-linear scales overall, but ultimately this improvement is limited.

I have also explored the gains from adding Fisher information from an external baryon source (e.g., new simulations or improved observations) which in general does mitigate the baryon impact on Σ, but to remove it as a statistically significant systematic requires more baryon information than could realistically be accessed at present. As such these results are not shown.

Instead the focus here applied to the most promising source of baryon mitigation in Chapter 3, the constraints on cosmology from the CMB (Planck Collaboration et al., 2016b). These propagate through the Fisher analysis, contributing to the breaking of parameter degeneracies. This limits the baryon degradation, and in the case of neutrinos I find a significant improvement with the impacts of marginalising over baryons reduced to 21% and 28% for the NH and IH. Although the IH still has a smaller Σ error it now experiences greater relative baryon degradation than the NH when CMB priors are included. Despite the improvements, the baryon impact is detrimental for the weak lensing probe, raising the WL+CMB errors beyond the 1σ hierarchy distinction threshold, $|\Delta\Sigma_{\rm min}|$, to $\sigma_{\Sigma,\rm NH} = 0.044$ and $\sigma_{\Sigma,\rm IH} = 0.042$. These would constitute 1σ and 2σ detections for $\Sigma_{\rm NH}$ and $\Sigma_{\rm IH}$ respectively. It will therefore be important to include additional information from other sources, such as galaxy clustering.

4.5.3 Including intrinsic alignments

Biased parameter estimations from failing to account for intrinsic alignments have been extensively studied but it is also important to consider the impact on forecasts of including IA in the C_ℓ signal and then marginalising over $a_{\rm IA}$. The results of this Chapter are broadly consistent with the predictions of Krause et al. (2016) for a Euclid-like survey, with the caveat that I do not extend my analysis to marginalising over nuisance parameters governing a potential luminosity scaling of the IA amplitude, and so slightly less significant impacts are found. I show, in Figure 4.9, a direct comparison between hierarchies for the w-Σ forecasts with and without IA.

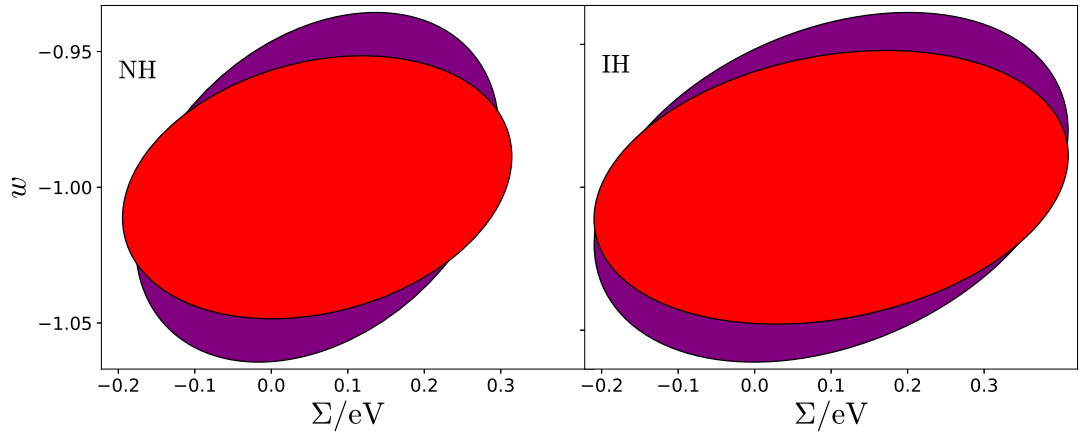

Figure 4.9: 1-σ w-Σ confidence ellipses for the normal (left; $\Sigma_{\rm fid} = 0.06\,{\rm eV}$) and inverted hierarchy (right; $\Sigma_{\rm fid} = 0.1\,{\rm eV}$) for a Euclid-like weak lensing survey. Purple: intrinsic alignments included in the modelling via the amplitude parameter, $a_{\rm IA}$; red: intrinsic alignments not included. All other cosmological parameters in $\Theta = (A_B, \eta_0, \Omega_m, \Omega_b, h, \sigma_8, n_s, w)$ are marginalised over in both cases.

These results indicate that the impact is small for most parameters, including Σ. This can be understood by considering the C_ℓ responses with and without IA for different redshift bin combinations. The $C_\ell^{\rm tot}$ responses for the $ij = 44$ z-bin ($0.79 \leq z \leq 0.89$) autocorrelation in Figure 4.4 are sufficiently similar to the corresponding $C_\ell^{\gamma\gamma}$ responses that it is not useful to distinguish between them by explicitly including the latter. This domination of the shear-shear signal is most extreme for auto-correlations in the highest redshift bins, as Figure 4.2 shows. However, for cross-bin correlations the IA-shear signal becomes important as it arises from the alignment of foreground galaxy ellipticities with background lens potentials. This term is negative and large enough to provide significant cancellation to the shear-shear signal in these cases.

It should be noted that there is a small improvement in the $\Sigma_{\rm NH}$ constraint when including IA in the modelling. This would seem to contradict the intuition that marginalising over an additional source of uncertainty should increase the error. However, it has been ascertained that the determinant of the Fisher matrix, which effectively quantifies the total available information, is reduced by the inclusion of IA. When the inversion to a parameter covariance matrix is performed this manifests as some individual errors (and parameter correlations) decreasing while others increase. In this case, a very slight decrease is seen for the $\Sigma_{\rm NH}$ error but this is compensated for by increased errors for other parameters.

Figure 4.10, illustrates the impact on the $C_\ell^{\rm tot}$ and $C_\ell^{\gamma\gamma}$ responses for the $ij = 09$ cross-correlation. To display this clearly, results are shown for a redshift bin combination, $ij = 09$, at which the differences between the two cases are most pronounced. As the IA only significantly affect the cross-power between widely separated bins or the auto-power for low redshift bins, and most of the lensing signal comes from the minimally impacted high redshift bin auto-spectra, the Σ Fisher information is similar between cases with and without IA. Any degradation of forecasts would therefore come from degeneracies between Σ and $a_{\rm IA}$. Figures 4.4 and 4.10 show significantly different power responses to these parameters. When combined with the fact that only a small subset of auto- and cross-power spectra are significantly sensitive to the IA, I would therefore expect the impact on constraints to be small.

For both the NH and IH cases, the IA signal adds limited information on all parameters other than w and Ω_m. A possible reason for the improvements due to IA seen to constraints in this analysis (as in the case of the slight decrease in the error for $\Sigma_{\rm NH}$) could be that the NLA model that is being used here is insufficiently sophisticated to capture the full impact. However, for the purposes of this work a relatively simplistic model is sufficient to gain a broad level of insight. Another finding to note is that neutrinos have a stronger effect on the power spectrum for the IH compared to the NH, and I find an even smaller, almost negligible, impact from IA on the $\Sigma_{\rm IH}$ forecast than for $\Sigma_{\rm NH}$. This suggests there is some sensitivity to the fiducial values of the parameters chosen. In turn this implies non-Gaussianity in the posterior, which would limit the validity of the Fisher approximation. However, the impact of IA for neutrino mass forecasts is overall of little concern, with only small improvements and degradations. The main systematic that must be mitigated remains the baryonic phenomena.

4.5.4 Galaxy clustering

I combine my WL+CMB forecasts with information from galaxy clustering (GC), which is derived from BOSS data at low redshifts ($0.2 < z \leq 0.75$) and from a Euclid-like spectroscopic galaxy survey at higher redshifts ($0.75 < z \leq 2.05$). Figure 4.11 shows confidence ellipses comparing the NH and IH with and with-

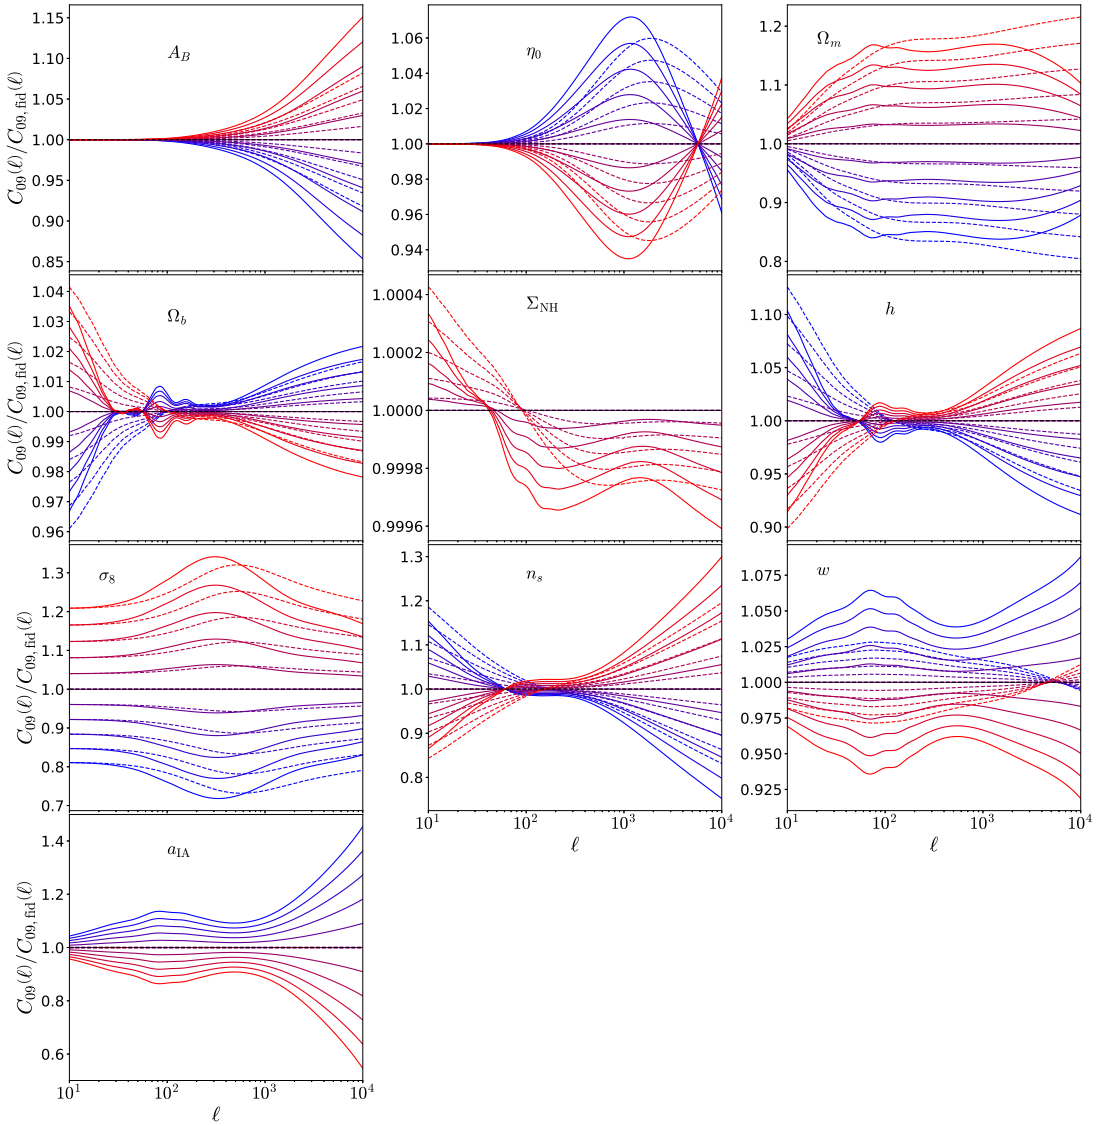

Figure 4.10: Lensing power spectrum responses using NH in the $ij = 09$ redshift bin. Blue (red) lines correspond to the lowest (highest) parameter values for $\Theta = (A_B, \eta_0, \Omega_m, \Omega_b, \Sigma_{\rm NH}, h, \sigma_8, n_s, w, a_{\rm IA})$ in the range $0.9\Theta_{\rm fid} \leq \Theta \leq 1.1\Theta_{\rm fid}$, except in the case of the neutrino mass parameter, which varies between $\Sigma_{\rm NH,min} \leq \Sigma_{\rm NH} \leq 1.1\,\Sigma_{\rm NH,min}$ with purple curves representing the minimal mass case, $\Sigma_{\rm NH,min} = 0.06\,{\rm eV}$. Solid lines correspond to the power spectrum including IA and dashed lines correspond to the $\gamma\gamma$ power spectrum where IA are not included.

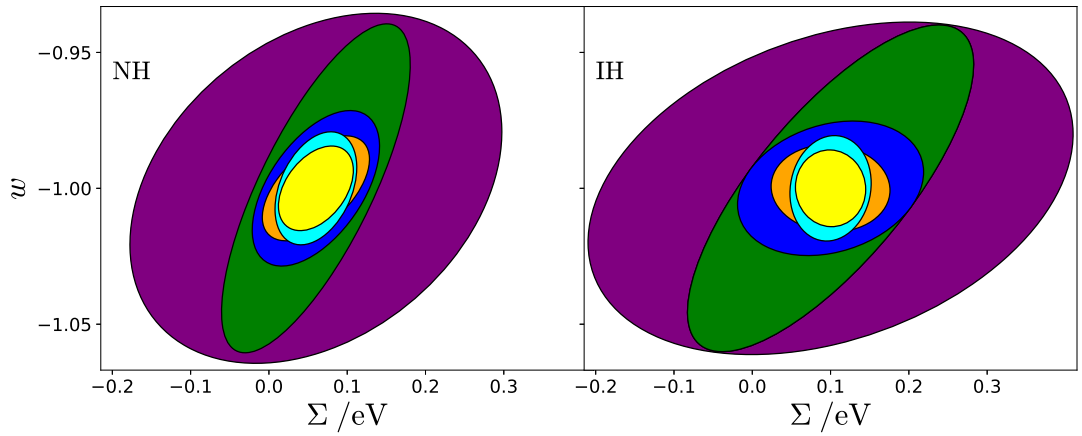

Figure 4.11: 1-σ w-Σ confidence ellipses for the normal (left; $\Sigma_{\rm min} = 0.06\,{\rm eV}$) and inverted hierarchy (right; $\Sigma_{\rm min} = 0.1\,{\rm eV}$) with all parameters in $\Theta = (A_B, \eta_0, \Omega_m, \Omega_b, h, \sigma_8, n_s, w, a_{\rm IA})$ marginalised over (purple: weak lensing only; blue: weak lensing and galaxy clustering; cyan: weak lensing, galaxy clustering and CMB priors combined); and with the baryon parameters, A_B and η_0, fixed to their fiducial values (dark green: WL only; orange: WL+GC; yellow: WL+GC+CMB).

	$\Sigma_{\rm NH}$	$\Sigma_{\rm IH}$
$\sigma_{\rm WL}$/eV (ex bar.)	0.079	0.120
$\sigma_{\rm WL}$/eV (inc. bar.)	0.156	0.204
$R_{B,{\rm WL}}$	1.97	1.70
$\sigma_{\rm WL+GC}$/eV (ex bar.)	0.045	0.050
$\sigma_{\rm WL+GC}$/eV (inc. bar.)	0.053	0.078
$R_{B,{\rm WL+GC}}$	1.19	1.57
$\sigma_{\rm WL+GC+CMB}$/eV (ex bar.)	0.031	0.030
$\sigma_{\rm WL+GC+CMB}$/eV (inc. bar.)	0.034	0.034
$R_{B,{\rm WL+GC+CMB}}$	1.09	1.16

Table 4.2: 1-σ error forecasts for a Euclid-like survey of the neutrino mass sum in the normal and inverted hierarchies, without and including marginalisation over baryon parameters and the addition of priors from galaxy clustering and *Planck* CMB measurements. I also include the response factors, R_B, to including baryon marginalisation.

out baryon marginalisation for three cases: weak lensing only; and weak lensing and galaxy clustering; weak lensing, galaxy clustering and CMB priors combined. Table 4.2 contains the corresponding 1-σ errors and the baryon degradation factors in each case. Neutrinos mainly impact the galaxy power spectrum through the matter power spectrum. The power sensitivity to Σ is greater for the IH than the NH, leading to more substantial improvements to forecasts for the former. However, both hierarchies experience a significant reduction in errors, with combined WL+GC+CMB probes achieving $\sigma_{\Sigma_{\rm NH}} = 0.034\,{\rm eV}$ and $\sigma_{\Sigma_{\rm IH}} = 0.034\,{\rm eV}$, when accounting for baryons. The degradation due to baryons is also approximately halved in the combined case. Again, even though the IH experiences a far greater improvement overall, the relative degradation is almost double that of the NH at 16% and 9% respectively. Interestingly, when examining the impact of adding GC to WL without the inclusion of CMB priors, it is apparent the NH degradation is reduced by a greater fraction than that of the IH.

The inclusion of BOSS data on low redshift clustering was found to have a significant impact on galaxy clustering constraints compared to the Euclid-like only case. However, the combined WL+GC errors are only mildly improved by $\sim$ 5-10% while there is no appreciable change for the WL+GC+CMB errors. Therefore, in my final analysis the dominant information contribution from galaxy clustering is provided by the Euclid-like forecasts rather than BOSS data. Ultimately, the combined forecast constraints represent a nearly 2σ (3σ) detection of $\Sigma_{\rm NH}$ ($\Sigma_{\rm IH}$) but are not significantly lower than the minimum threshold required to achieve a distinction between hierarchies.

My results for a GC+CMB combination are broadly comparable with recent Fisher forecasts by Boyle and Komatsu (2018), although, as Boyle (2018) notes, making comparisons between results across the literature is difficult. For example, differences in the implementation of CMB priors, the use of Fisher versus MCMC methodology, the choice of parameters to be varied, and their fiducial values leads to a range of results. Indeed, Audren et al. (2013) and Sprenger et al. (2018) find stronger constraints in studies restricted to the degenerate mass hierarchy. For galaxy clustering, where neutrino information is coming primarily from the shape of the matter power spectrum on linear and mildly non-linear scales, this choice can have a significant impact. I find that the different free-streaming lengths characterising the NH and IH (discussed in detail in § 4.5.1) can have a non-negligible

effect on the power sensitivity to Σ. It follows that addressing whether Stage IV surveys will be able to measure Σ or distinguish between hierarchies should not simply rely on predicting σ_Σ for an arbitrary model, but should robustly account for the effects of different hierarchy mass splittings. As this represents a sensitivity to fiducial parameters, which implies a potential breakdown of the Fisher approximation, beyond-Fisher methods such as MCMC may need to be considered necessary for future analyses (see e.g., Hall and Challinor, 2012).

4.5.5 Required improvements

An extension to the analysis of Chapter 3 of the effects of adding arbitrary priors by increasing the diagonal baryon elements of the Fisher matrix is provided here. In this case, the motivation is determine which parameters are contributing most to the size of the forecast Σ errors. Therefore, to quantify the information required from an arbitrary external source to reduce the errors on $\Sigma_{\rm NH}$ and $\Sigma_{\rm IH}$ sufficiently for a Euclid-like survey to distinguish the hierarchies beyond the 1σ level, each diagonal element of the (combined WL+GC+CMB) Fisher matrix is systematically increased by a factor, $(1+\alpha)$,

$$F'_{ii} \longrightarrow F_{ii}\,(1+\alpha)\,. \tag{4.12}$$

In Figure 4.12 I plot the change in NH and IH errors for increasing α for the parameters Σ and n_s. These provide the strongest improvements, as the impact of marginalising over other parameters on σ_Σ is already close to fully mitigated by the WL+GC+CMB combination of probes.

For Σ, while the WL+GC+CMB constraints alone constitute close to a 2σ detection for the NH (with the IH being constrained to close to 3σ), this can be improved to between 3σ and 4σ by adding priors in the range $0.7\,\sigma_{\Sigma,\rm con} \lesssim \sigma_{\Sigma,\rm prior} \lesssim 2.0\,\sigma_{\Sigma,\rm con}$, where $\sigma_{\Sigma,\rm con}$ is the conditional error on the mass sum. It may also be possible to achieve a distinction between hierarchies at 2σ or even 3σ with priors in the range $2.0\,\sigma_{\Sigma,\rm con} \lesssim \sigma_{\Sigma,\rm prior} \lesssim 2.3\,\sigma_{\Sigma,\rm con}$. This required information could be contributed to by e.g., improved data from the Lyman-α forest or possibly further constraints from particle physics. An example of such a particle experiment is measuring the decay rate of neutrinoless double beta decay, as this depends on the sum of mass eigenstates, $m_{\beta\beta}$, when weighted by Majorana

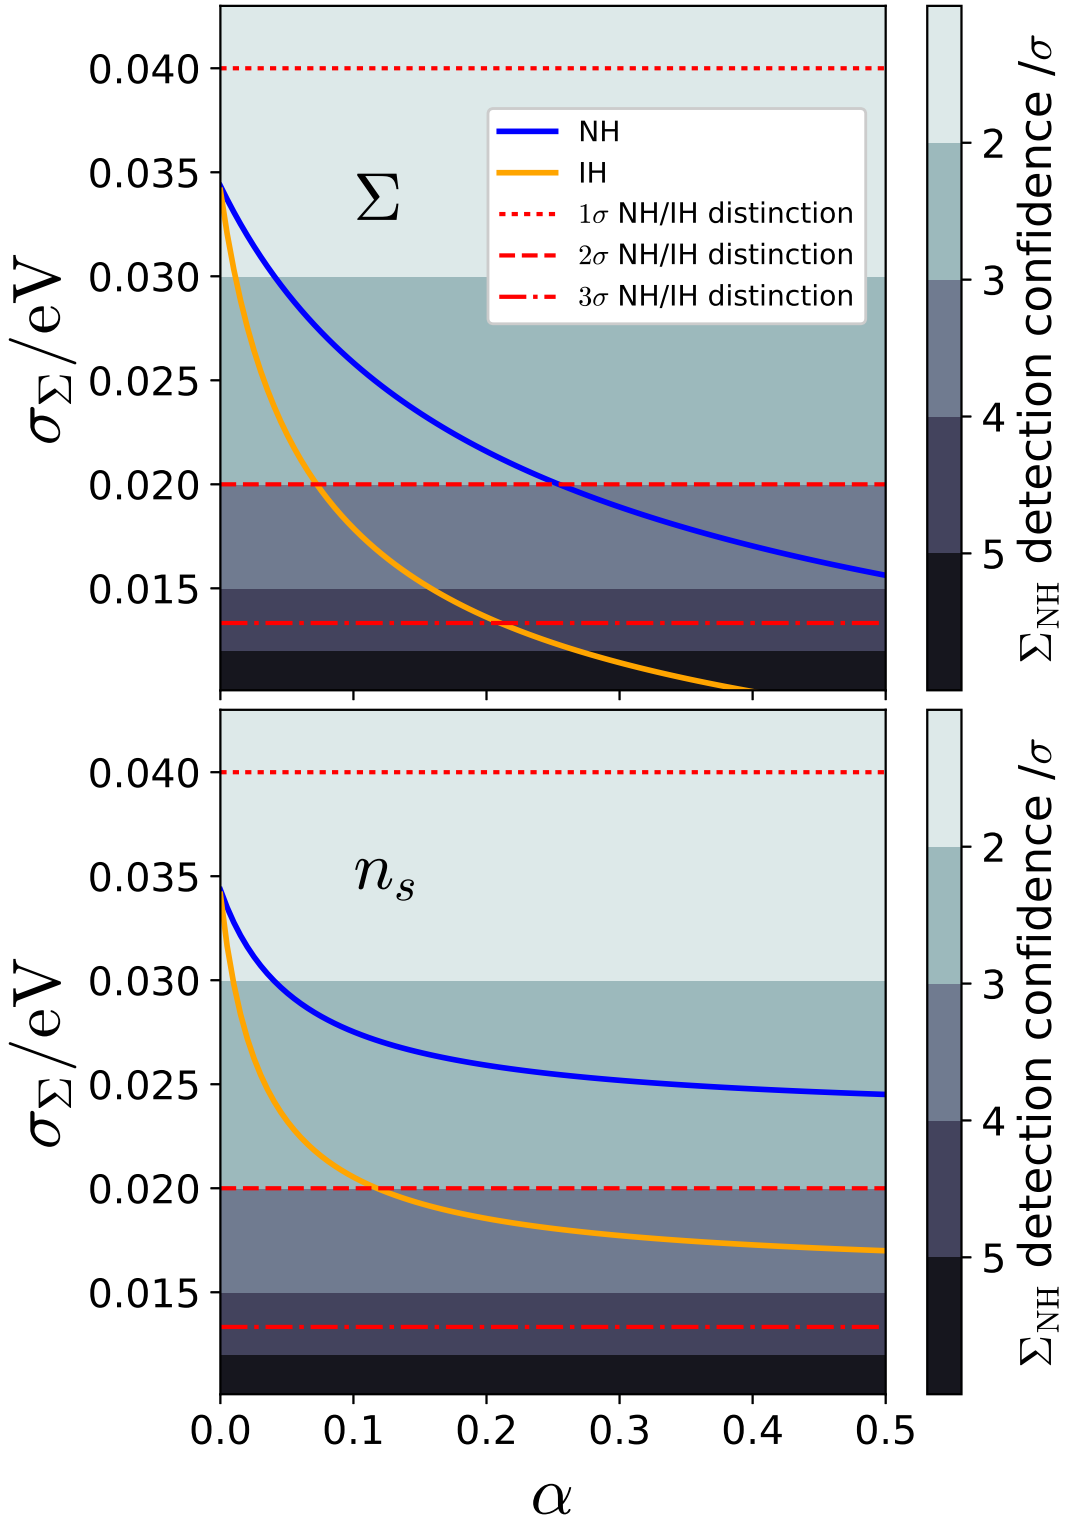

Figure 4.12: 1-σ WL+GC+CMB errors for Σ for the NH (blue) and IH (orange), as a function of arbitrary increases, α, to the Σ (top panel) and n_s (bottom panel) Fisher information, as defined in equation (4.12). The shaded regions correspond to the confidence level that Σ can be measured at in the case of the NH (the equivalent regions for the IH are not included because $\Sigma_{\rm IH}$ is already constrained to $\sim 3\sigma$ when $\alpha = 0$). The red lines mark the thresholds required to distinguish the NH and IH at 1σ (dotted), 2σ (dashed) and 3σ (dash-dotted). The difference between the NH and IH fiducial values, $|\Delta\Sigma_{\rm min}| = |\Sigma_{\rm NH,min} - \Sigma_{\rm IH,min}| = 0.04\,{\rm eV}$, represents the 1σ limit.

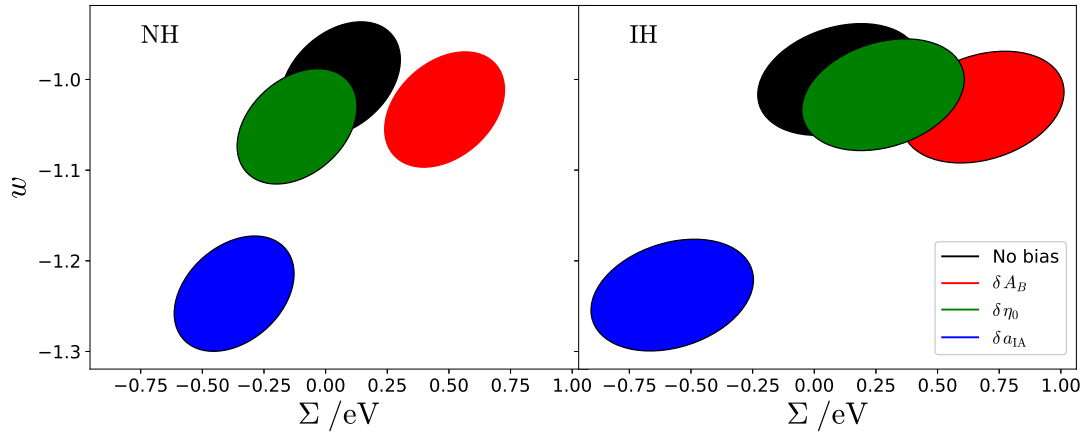

Figure 4.13: Biased 1-σ w-Σ confidence ellipses for the NH and IH with all parameters in $\Theta = (A_B, \eta_0, \Omega_m, \Omega_b, h, \sigma_8, n_s, w, a_{\rm IA})$ marginalised over. Bias in baryon parameters corresponds to values located in the extremes of the ranges identified in M15, $\delta\, A_B = -1.13$ (red) and $\delta\, \eta_0 = -0.203$ (green). The IA amplitude is biased by 10% such that $\delta\, a_{\rm IA} = 0.1$ in the blue ellipses. The black ellipses represent the unbiased forecasts.

phases[2] (see e.g., Capozzi et al., 2017). Although neutrinoless double beta decay has not yet been observed, the KamLAND-Zen experiment derives a lower bound of $m_{\beta\beta} \gtrsim 0.061{-}0.165\,\text{eV}$ (KamLAND-Zen Collaboration, 2016).

There is also considerable scope for improving the Σ constraints with additional priors on the spectral index which, if known with absolute certainty, would allow a 2σ (5σ) detection of Σ for the NH (IH) and a 2σ distinction between the hierarchies. The next generation CMB surveys CMB-S4 (Abazajian et al., 2016) predicts constraints on n_s tightening by a factor of 2-3 over *Planck*, which could contribute to the improvements that appear to be required for σ_Σ in Figure 4.12. It should be noted that the greater rate of decrease of $\sigma_{\Sigma_{\rm IH}}$ with α compared to $\sigma_{\Sigma_{\rm IH}}$ is a consequence of stronger degeneracies, e.g. between n_s-Σ, being broken for the IH compared to the NH.

4.6 Model Bias

The parameters controlling the baryon and IA systematics explored in this work are assigned their fiducial values from fits to simulations (Schaye et al., 2010;

[2]Neutrinoless double beta decay is subject to the assumption that at least one neutrino is a Majorana fermion, i.e. that $\nu \leftrightarrow \bar{\nu}$.

Mead et al., 2015) and shear observations (see e.g., Hirata and Seljak, 2004; Bridle and King, 2007) respectively. If the underlying physical models used in simulations are incorrect or if there are inherent limitations in the construction of the simulations, then these fiducial values are biased from the 'true' values. Similarly to Chapter 3, it is useful to consider the results of this work in this context by calculating the degree of calibration bias required to severely bias Σ forecasts, for example beyond their 1-σ errors.

Again, the Taylor et al. (2007) formalism for bias, captured by equation (3.36), is used with Figure 4.13 showing the bias in the weak lensing forecasts for w-Σ when the fiducial values of the baryon and IA parameters have been miscalibrated. I display results for model bias corresponding to the extreme values of the regions $2 < A_B < 4$ and $0.4 < \eta_0 < 0.8$ that broadly encompass the range of fits to different OWLS simulations made by M15, and for bias in $a_{\rm IA}$ of 10%.

Figure 4.13 shows that bias in Σ from miscalibrating the strength of the IA effect is significant, easily invalidating a measurement for either hierarchy with a 10% $a_{\rm IA}$ bias. This is in contrast to the very small impact on Σ forecasts from marginalising over $a_{\rm IA}$. However, as I find percent-level constraints on $a_{\rm IA}$ it might be expected that a 10% miscalibration would indeed produce a very large bias. The problem of IA bias for other parameters such as w is well-documented; this work does not address it further here beyond reporting the impact on Σ. The bias due to baryons is less severe, although an extreme miscalibration of the adiabatic contraction parameter would bias Σ beyond its 1-σ confidence region. However, biases of η_0, even lying at the extremes of the region of values identified by M15, do not displace the Σ estimate beyond the 1-σ limit. This difference between A_B and η_0 may also be reflected in the fact that marginalising over A_B alone has a significantly greater impact on Σ forecasts than marginalising over η_0 alone, with almost all of the baryon degradation being due to A_B. This highlights the importance of achieving tight calibrations of the adiabatic contraction parameter and the intrinsic alignment amplitude, while a similarly robust understanding of the strength of the baryonic feedback parameter is less urgent.

A subtlety worth noting is that a bias of $\delta\,\eta_0 = -0.2$ produces biases of opposite sign in Σ of different hierarchies. Because the sensitivity to η_0 is low, it may be that transitioning from one model to the other induces a small change to

the bias in Σ that is still sufficiently large to change its overall sign. This could complicate attempts to distinguish between hierarchies, but overall the η_0 bias remains subdominant to other effects.

As model bias is a product of subgrid limitations, resolving the issue is challenging without resorting to incorporating external data on baryon phenomenology into future simulations. In the case of IA, there has been substantial exploration into self-calibration methods of cleaning IA signals with independent information on IA correlations from deriving scaling relations that predict (and then subtract out) IA-shear signals from shear observables (Zhang, 2010; Troxel and Ishak, 2012; Yao et al., 2017, 2019). The IA-IA signal is strongest between close galaxies and so is generally assumed to be more straightforward to treat by purposely using sufficiently high z redshift bins to limit its contribution (Schneider and Bridle, 2010; Zhang, 2010). A future avenue of research into applying these methods to mitigate the bias impact on parameters like the neutrino mass sum would be worth pursuing.

4.7 Conclusions

This Chapter has presented detailed Fisher forecast analyses to assess the potential of a Euclid-like Stage IV survey to measure the neutrino mass sum within sufficient accuracy to distinguish between the normal and inverted mass hierarchies. Particular focus has been applied to the risk of forecast degradation presented by the need to marginalise over baryonic astrophysics, and further explored the impact of intrinsic alignments. To evaluate the widest possible scope of constraints available at present, forecasts for weak lensing were combined with those of spectroscopic galaxy clustering from Euclid and added prior information from *Planck* CMB constraints and BOSS measurements of low redshift clustering.

To explore the behaviour of baryons and neutrinos on non-linear scales, the halo model of Mead et al. (2015) and Mead et al. (2016) has been used. The impact of neutrinos is modelled by relating the mass sum to the spherical collapse overdensity and the virial density, which exert a significant influence on halo structure by changing the concentration function and the limits of the density profile re-

spectively. Fits of this model to simulations by Massara et al. (2014) reproduce $P(k)$ to within a few percent up to $k = 10\,h\,\mathrm{Mpc}^{-1}$ over different redshifts, which represents an improvement over a widely used fitting formula presented by Bird et al. (2012).

This work has examined both the impact of the Σ dependence of the power sensitivity to the mass sum and the effect of different free-streaming scales between hierarchies on the relative sensitivity of the IH compared to the NH. It was illustrated that for both hierarchies, there is low sensitivity on scales relevant to the cosmological probes, $\Delta^2(k)$ and C_ℓ. The resulting limited Fisher information available for a Euclid-like weak lensing survey lead to forecasts of $\sigma_{\Sigma_{\mathrm{NH}}} = 0.079\,\mathrm{eV}$ and $\sigma_{\Sigma_{\mathrm{IH}}} = 0.120\,\mathrm{eV}$. These are large enough to rule out making any definitive measurement for either hierarchy.

By incorporating prior CMB information from Planck Collaboration et al. (2016b), these constraints were tightened significantly to $\sigma_{\Sigma_{\mathrm{NH}}} = 0.036\,\mathrm{eV}$ and $\sigma_{\Sigma_{\mathrm{IH}}} = 0.033\,\mathrm{eV}$. Discussion was presented of the gains provided by this source of cosmological information in breaking degeneracies between Σ and other parameters, such as σ_8, which has a qualitatively similar effect as Σ on non-linear power. It was also illustrated that the relative advantages in constraining the IH over the NH were 72% and 54% improvements respectively when including CMB information.

This distinction was found to be more nuanced when the analysis was expanded to include marginalisation over the baryon parameters. Studying power responses illustrated specific degeneracies, e.g., between Σ and η_0 on non-linear scales, that contributed to the approximate doubling of mass sum error forecasts. While stronger non-linear sensitivity to Σ for the IH was found to result in a less severe impact than for the NH, when CMB priors were accounted for the degradation for the NH (21%) was appreciably less than that for the IH (28%). Despite this, the IH once again exhibited a much greater overall improvement than for the NH.

My weak lensing analysis has also included an investigation into the impact of including (and marginalising over) intrinsic alignments in my model on parameter forecasts. Unlike the significant bias induced in parameter estimations by neglecting IA, it was found that there is a relatively minimal effect for the separate

issue of forecasts because C_ℓ responses, which determine Fisher information, are very similar in both scenarios (especially for correlations between high redshift bins).

As the WL+CMB errors including baryons represent only a 1σ (2σ) confidence of measuring $\Sigma_{\rm NH}$ ($\Sigma_{\rm IH}$) and are greater than $|\Delta\Sigma_{\rm min}|= 0.04\,{\rm eV}$, I have included independent forecasts from a Euclid-like spectroscopic galaxy clustering survey. I have shown that this halved the baryon degradation factor for both hierarchies, with the NH still experiencing approximately half the impact of the IH. The galaxy clustering analysis has highlighted the importance of specifying the mass hierarchy when making neutrino forecasts as most information is derived from linear matter power responses. The configuration of free-streaming scales corresponding to each neutrino mass played a non-negligible role on these scales. Ultimately, the combination of WL+GC+CMB reduced NH (IH) errors to close to 2σ (3σ) and below the $|\Delta\Sigma_{\rm min}|$ threshold. The challenge of navigating the variation in forecast methodologies, and therefore results, made across the literature was noted(e.g., Carbone et al., 2011; Audren et al., 2013; Boyle and Komatsu, 2018; Sprenger et al., 2018). I argue that my approach of simultaneously incorporating the effects of neutrinos and baryons in the weak lensing power spectrum on non-linear scales in a physically well-motivated manner makes the forecasts presented here reliable and robust.

A potentially encouraging finding has been that additional information from e.g., neutrinoless double beta decay on the mass sum, or improved constraints on the spectral index from forthcoming Stage IV CMB experiments, could substantially reduce the errors on Σ. If these priors prove sufficient to increase the Fisher information on their corresponding parameters by 10-50%, it would achieve a 2σ (5σ) detection of Σ for the NH (IH) and distinguish between the hierarchies at the 2σ or even 3σ level.

Finally, the issue of model bias has been explored. I used first order approximations (Taylor et al., 2007) to determine the bias in Σ from miscalibrations of the baryon and intrinsic alignment parameters. I have found that large biases in A_B or $a_{\rm IA}$ induce a bias in Σ larger than σ_Σ. By contrast, I have found that even significant biases in the baryon feedback parameter, η_0, did not change the Σ estimate by more than σ_Σ. Though an understanding of model bias is important to

place the analysis in its proper context, it should be noted that it will ultimately require additional baryon information, improved simulations or self-calibrating methods for IA to successfully mitigate the issue.

An important limitation of my analysis to recognise is that the Fisher formalism underpinning it is not ideal for making forecasts in the case of distinct, competing models, as is the case with the normal and inverted mass hierarchies. Fisher information is determined by derivatives of likelihood functions around most-likely values. In cosmology these usually take the form of multi-variable unimodal Gaussians. In this work I have made assumptions that this is the case when analyzing each hierarchy separately. However, to make a robust forecast of Σ, given uncertainty over which hierarchy is correct, my analysis should incorporate the resulting bimodal Gaussian (with peaks at $\Sigma_{\rm NH,min} = 0.06\,{\rm eV}$ and $\Sigma_{\rm IH,min} = 0.1\,{\rm eV}$). This would require a non-trivial extension of the Fisher formalism, which is beyond the scope of this work but that I intend to pursue in future studies.

In summary, by using the M15 and M16 prescriptions for baryonic and neutrino effects on haloes, I have been able to perform a full forecast analysis of the neutrino mass sum in the normal and inverted hierarchies. By combining multiple large scale structure probes through Euclid-like weak lensing and spectroscopic galaxy clustering surveys, *Planck* CMB constraints and BOSS low redshift galaxy clustering data, I have shown that degradation due to baryon feedback can be reduced to between 9-16%. Using these sources of information, and accounting for baryons and intrinsic alignments, Stage IV surveys could be expected to measure Σ with 1-σ errors of $\sigma_{\Sigma_{\rm NH}} = 0.034\,{\rm eV}$ and $\sigma_{\Sigma_{\rm IH}} = 0.034\,{\rm eV}$ for the NH and IH. These approach the confidence level required to meaningfully distinguish the hierarchies, but with additional future priors on Σ and n_s there is tentative optimism for achieving more definitive results.

Chapter 5

Coupled Dark Energy In The Halo Model Using Virial Arguments

This Chapter returns to dark energy as its focus. However, rather than the standard $w\,(a)$CDM dynamical dark energy model that Chapter 3 addressed, this work explores a dark energy model defined by a non-minimal coupling to dark matter. Modelling the effects of such an interaction on large scale structure encounters significant challenges in the non-linear regime. Efforts to fit the matter power spectrum to results from coupled dark energy simulations have generally relied on complicated fitting functions with many free parameters to capture these scales.

This Chapter will present an alternative approach by using collapse arguments to compute the halo virial density in the presence of a coupling. I will show that varying this parameter can capture the changes to the non-linear power up to $k = 10\,h/\mathrm{Mpc}$ with percent level accuracy for coupling strengths and redshifts of interest. A fit will be made to simulations of the matter power spectrum in coupled dark energy models, using the baryon-halo model of previous Chapters. This allows for weak lensing forecasts for the coupling strength to be made, while presenting for the first time an analysis of the impact of marginalising over baryonic effects. An attempt at mitigating the baryon impact will be explored in detail.

An overview of the range of coupled dark energy theories, the simulations that currently exist to capture the response of the matter distribution to the coupling, and the state of forecasts for these models are discussed in an introduction in § 5.1. Following this, in § 5.2 the effects of coupled dark energy on dark matter haloes are explored. This involves an outline first of how the cosmological background can be computed in these models, before a detailed examination of the linear and non-linear growth equations in coupled dark energy. These inform computations of important halo parameters like the virial density threshold using spherical collapse arguments. These will be applied in an effort to fit to simulated power spectra in § 5.3. The Fisher forecast analysis used in previous Chapters will then be applied using the modified halo model. A discussion of how power spectrum responses inform the results of the forecasts, including the impact of baryons will be presented in § 5.4. The capacity of galaxy clustering and CMB priors to improve on these results will be assesssed before a conclusion is presented in § 5.5.

5.1 Introduction To Coupled Dark Energy

The model space of beyond ΛCDM cosmologies includes a range of scenarios featuring interactions between dark energy and other species. Among them are scalar-tensor theories of modified gravity that introduce a scalar field, ϕ, that is non-minimally coupled to the metric (see e.g., Clifton et al., 2012). Alternatively, models can be constructed in which dark energy interacts with matter-energy fields. These are generally confined to couplings to cold dark matter or neutrinos. A hypothetical interaction with baryons should be negligibly weak in order to satisfy a lack of observation of additional 'fifth forces' acting on this sector (Peebles and Ratra, 2003). Similarly, any coupling to radiation would perturb photons from geodesic trajectories, in violation of local observations which preclude this behaviour (Wang et al., 2016). There are a number of interesting propositions for relic neutrinos acquiring a growing mass by coupling to a scalar field. These allow for the possibility of late time dark energy domination being a consequence of neutrinos becoming non-relativistic and effectively freezing the evolution of the scalar field. I recommend e.g., Amendola et al. (2008), Wetterich (2007) and Casas et al. (2016) for the interested reader but, do not consider such models in this work, which I restrict to interactions between dark matter and dark energy.

These coupled dark energy (CDE) cosmologies are generally conceptualised as a cold dark matter component exchanging energy-momentum with a scalar field responsible for the observed accelerating expansion of the Universe (see e.g. Wetterich, 1995; Amendola, 2000; Mangano et al., 2003). As a result, the friction effect of the Hubble expansion on structure growth is modified, and a fifth force arises that acts alongside gravity on matter perturbations undergoing collapse. Both effects would be expected to play a significant role in the distribution of matter on large scales and have a subtle influence on structure in the non-linear regime (Macciò et al., 2004; Baldi et al., 2010; Li and Barrow, 2011; Baldi, 2011a).

As dynamical dark energy theories, CDE models circumvent the fine-tuning problem of the cosmological constant in ΛCDM (Weinberg, 1989). They are more specifically appealing to the extent that they resolve discrepancies between observations and predictions from the concordance model of cosmology. ΛCDM accounts for most cosmological data very well but encounters increasing chal-

lenges on smaller scales. For example, the rotation curves of dwarf galaxies indicate inner halo cores in place of predicted density profile cusps (e.g., Dubinski and Carlberg, 1991; Simon et al., 2003), there is an apparent deficit of observable satellite galaxies in haloes (Navarro et al., 1997), and high infall velocities characterise Bullet Cluster-like objects far in excess of ΛCDM simulation results (e.g., Lee and Komatsu, 2010). Complex astrophysical processes like baryonic phenomena may provide explanations for these and related problems. However, the alternative that the impacts on non-linear structure resulting from underlying interactions between matter and dark energy may be responsible is also well-motivated (see e.g., Baldi et al., 2010; Baldi and Viel, 2010; Penzo et al., 2016).

While the background dynamics and the large-scale evolution of perturbations in CDE are well-understood (Amendola, 2004), the challenge of interpreting the non-linear growth of structure for these cosmologies is considerable. The CoDECS project (Baldi, 2012) is arguably the most comprehensive set of simulations constructed for this purpose. It combines small-scale adiabatic hydrodynamical simulations with a set of N-body simulations using a box of 1 Gpc/h comoving length and 2×1024^3 CDM and baryon particles. CoDECS incorporates several effects of CDE to capture non-linear perturbations including the mass variation of cold dark matter particles due to the dark energy interaction; a modified expansion rate caused by an effective early dark energy component; an extra velocity-dependent acceleration of particles generated by the coupling in the Euler equation; and the fifth force, manifesting as a modification to the CDM-CDM gravitational interaction only (Baldi et al., 2010). The latter is achieved through a hybridisation of Tree and Particle Mesh algorithms. While the effects incorporated in CoDECS are extensive there are some limitations. For example, a significant branch of coupled dark energy theories is neglected, namely those models where dark energy couples to massive neutrinos. Growing neutrino quintessence, for example, may provide an explanation for why dark energy dominates only relatively recently in the history of the Universe (for details see e.g., Baldi et al., 2011). However, the complexities associated with including this additional coupling within simulations are not insignificant. Nevertheless, a future iteration of CoDECS implementing this physics could have significant benefits for expanding the understanding of the effects dark energy couplings on structure.

CoDECS provides non-linear matter power spectra across a range of redshifts for several CDE models where energy is transferred from dark matter to dark energy. This is valuable in part because the power spectrum is the underlying statistical probe for forthcoming Stage IV large scale structure surveys like Euclid (Laureijs et al., 2011) and LSST (LSST Science Collaboration et al., 2009). As has been discussed in previous Chapters, the matter power contributes to the weak gravitational lensing signal and the galaxy clustering power spectrum. Multiple photometric bins allow tomographic weak lensing surveys to constrain cosmological parameters through their influence on the cosmic expansion history and the growth of structure (Bartelmann and Schneider, 2001; Albrecht et al., 2006; Peacock et al., 2006). Spectroscopic galaxy clustering measurements provide an additional source of information through the relationship between the underlying dark matter distribution and the galaxies that populate it (Hauser and Peebles, 1973; Davis and Peebles, 1983).

In order to use lensing and clustering probes to constrain CDE parameters like the coupling strength (and also to perform full cosmological forecasts assuming an underlying CDE model), it is necessary to first have a model for the power spectrum that fits simulation predictions. Casas et al. (2016) determine a polynomial fit to the ratio of CDE to ΛCDM power in CoDECS, leading to a combined weak lensing and galaxy clustering forecast of 3.4% for the square of the coupling strength when power is normalised at early times. Cosmic microwave background (CMB) priors from *Planck* place on the coupling strength, β, an upper bound of $\beta \lesssim 0.06$.

The Casas et al. (2016) fit achieves percent level accuracy but requires 40 free parameters characterising polynomial and sigmoidal functions of the coupling strength, scale and redshift. The forms of these functions are motivated by the pursuit of accurate fits but have no physical justification in themselves. In this work I pursue an alternative avenue, inspired by the approach of Mead et al. (2015) and Mead et al. (2016) to constructing the HMCODE halo model power spectrum fit through empirically motivated modifications to internal halo structure properties. A prominent example was the underpinning for Chapter 4. The effects of free-streaming massive neutrinos on non-linear structure are captured with percent level accuracy to simulations by Massara et al. (2014) in HMCODE by modifying the spherical collapse overdensity, δ_c, and virial density threshold,

Δ_v. These quantities are characteristic descriptors of the halo model, with their modified forms predicted by solving the equations for spherical collapse in a massive neutrino cosmology (Ichiki and Takada, 2012; LoVerde, 2014).

Evidence that a similar strategy may be applicable for CDE is apparent in M16, who show that lowering the value of a parameter describing a linear combination of the concentration amplitude and the amount of bloating due to baryonic feedback provides a reasonable fit to the CoDECS non-linear power. However, they do not provide a basis for predicting the required changes to this parameter for different coupling strengths and at different redshifts. It was shown in equation (5.34) that the M15 modification for halo bloating effectively changes the virial radius to which a halo of a given mass must extend to satisfy the virial density. This suggests that Δ_v and δ_c, which affect the concentration amplitude, are the important underlying parameters. Moreover, the values of these parameters in CDE can, in principle, be predicted from physical arguments. Wintergerst and Pettorino (2010) derive the CDE non-linear growth equation and use it to compute the corresponding δ_c. This could also be extended to calculate Δ_v. One would expect the conservation of momentum in a CDE scenario where the mass of dark matter particles is reduced to lead to an enhanced acceleration of particles undergoing collapse into halo structures. As the process of gravitational collapse is accelerated, this would lower the density threshold at which virialisation occurs. Haloes would effectively be more bloated, thereby suppressing matter power on non-linear scales as seen in the CoDECS results.

This Chapter explores how well these parameters can be used to fit HMCODE predictions for the matter power (which is accurate to within $\simeq 5$ percent for scales $k \leq 10\,h\,\mathrm{Mpc}^{-1}$ and redshifts $z \leq 2$ for wCDM) to CoDECS, and the extent to which the required modifications can be predicted by spherical collapse calculations. This has the potential to improve on the Casas et al. (2016) fit by greatly reducing the number of free parameters and providing a physically grounded description of non-linear CDE power. These benefits are especially important for performing parameter inference from data and potentially being able to constrain CDE physics. More immediate for this Chapter are the benefits for a Fisher forecast analysis, which I perform for weak lensing and galaxy clustering probes. Using such a physically motivated fit would allow the source of potential degeneracies between parameters to be traced and understood. This in turn can

inform strategies for how the degeneracies may be mitigated by forthcoming surveys.

Incorporating CDE effects into a variant of HMCODE through the collapse and virial densities presents an excellent opportunity to also evaluate the impact of a major systematic like baryons on the capacity of surveys to constrain the coupling strength. Chapters 3 and 4 performed in-depth examinations into the baryon impact of $w(a)$ and massive neutrinos. Chapter 3 in particular discusses the details of the baryonic astrophysical phenomena that are also relevant for Fisher forecasts on the coupling strength in this Chapter. An early indication that there may be strong degeneracies between CDE and baryonic effects is that the parameter describing bloating due to feedback affects the virial limit of haloes, which also seems to plausibly be influenced by the coupling strength. A halo model parameterising both sets of effects therefore allows for an impact assessment of baryons on coupled dark energy constraints for the first time.

5.2 Coupled dark energy halo model

5.2.1 Coupled dark energy background

Dark energy models featuring scalar fields can be constructed to generate late-time accelerated expansion, for example by ensuring that the density of the field scales similarly to radiation at early times but after matter-radiation equality resembles a cosmological constant with its energy density dominating over that of matter. This partially addresses the cosmological constant problem by replicating the required present acceleration while maintaining a naturally low energy scale for dark energy[1] (Ratra and Peebles, 1988; Caldwell et al., 1998; Steinhardt et al., 1999). From the Lagrangian (Amendola, 2000),

$$\mathcal{L} = -\frac{1}{2}\partial^{\mu}\phi\partial_{\mu}\phi - V(\phi) - m_{\rm CDM}(\phi)\,\psi\bar{\psi} + \mathcal{L}_{\rm kin}[\psi]\,, \tag{5.1}$$

[1] Although of course the issue of providing a theoretical explanation for why the vacuum does not gravitate in this interpretation is not resolved.

for a spatially uniform scalar field, ϕ, coupled to a cold dark matter field, ψ, of mass, $m_{\rm CDM}(\phi)$, the evolution equation can be derived as

$$\ddot{\phi} + 3H\dot{\phi} + \frac{{\rm d}V(\phi)}{{\rm d}\phi} = \sqrt{\frac{2}{3}}\beta\rho_{\rm cdm}. \tag{5.2}$$

The term on the right hand side describes an exchange of energy between the dark matter and dark energy sectors. Here I use the convention of an overdot for derivatives with respect to time and set the reduced Planck mass $M_{pl} = 1/\sqrt{8\pi G} \equiv 1$ throughout this work. The coupling strength, β, is given by

$$\beta \equiv -\frac{{\rm d}\ln m_{\rm CDM}(\phi)}{{\rm d}\phi}. \tag{5.3}$$

The sign convention ensures that, for $\beta > 0$, a flow of energy from dark matter to dark energy (accompanied by decreasing $m_{\rm CDM}$) occurs for $\dot{\phi} > 0$ and in the other direction for $\dot{\phi} < 0$. The form of the scalar field potential,

$$V(\phi) = A\exp\left[-\sqrt{\frac{2}{3}}\alpha\phi\right] \tag{5.4}$$

is chosen here and normalised such that $\phi(t_0) = 0$ (where t_0 denotes the present time elapsed since the initial singularity) and becomes increasingly negative in the past. This also sets $A = V(t_0)$. For the cosmological models examined in this work, the remaining free parameter assumes a value, $\alpha \approx 0.1$, in order to satisfy the normalisation condition. I also restrict positive $\beta > 0$, so the coupling configuration in this work is always energy transfer from dark matter to dark energy. The energy density and pressure of the scalar field are then given respectively by

$$\rho_\phi = \frac{1}{2}\dot{\phi} + V(\phi), \tag{5.5}$$

$$P_\phi = \frac{1}{2}\dot{\phi} - V(\phi). \tag{5.6}$$

These define the equation of state parameter, $w_\phi \equiv P_\phi/\rho_\phi$. The background equations are completed with the density continuity equations for cold dark matter,

baryons and radiation,

$$\begin{aligned} \dot{\rho}_{\rm cdm} + 3H\rho_{\rm cdm} &= -\sqrt{\tfrac{2}{3}}\beta\rho_{\rm cdm}\dot{\phi} \\ \dot{\rho}_b + 3H\rho_b &= 0 \\ \dot{\rho}_r + 4H\rho_r &= 0 \end{aligned} \tag{5.7}$$

and the Hubble parameter, $H \equiv \dot{a}/a$, provided through the Friedmann equation,

$$H^2 = \sum_i \frac{\rho_i}{3}, \tag{5.8}$$

for $i = \{{\rm cdm}, b, r, \phi\}$. Note that in order to conserve energy, the right hand side of the cold dark matter equation has an equivalent but opposite in sign contribution as the right hand side of equation (5.2). This describes the loss of energy from dark matter to the dark energy sector. Indeed, the cold dark matter density is modified from its usual scaling with a^{-3} such that it is now given by

$$\rho_{\rm cdm}(a) = \rho_{\rm cdm}(a_0)\, a^{-3} \exp\left[-\sqrt{\frac{2}{3}}\beta\phi\right]. \tag{5.9}$$

The full set of background equations can be recast in terms of derivatives with respect to $\ln a$ (denoted by a prime) to give the more elegant, less computationally expensive form (Copeland et al., 1998; Amendola, 2000),

$$x' = \frac{x}{2}\left[3x^2 + 3y^2 + r^2 - 3\right] + \alpha y^2 + \beta\left[1 - x^2 - y^2 - r^2 - v^2\right] \tag{5.10}$$

$$y' = \frac{y}{2}\left[3x^2 + 3y^2 + r^2 - 3\right] - \alpha x\, y \tag{5.11}$$

$$v' = \frac{v}{2}\left[3x^2 + 3y^2 + r^2\right] \tag{5.12}$$

$$r' = \frac{r}{2}\left[3x^2 + 3y^2 + r^2 - 1\right] \tag{5.13}$$

$$H' = -\frac{H}{2}\left[3x^2 + 3y^2 + r^2 + 3\right], \tag{5.14}$$

$$\tag{5.15}$$

where

$$x \equiv \sqrt{\frac{\Omega_\phi(1+w_\phi)}{2}}, \quad y \equiv \sqrt{\frac{\Omega_\phi(1-w_\phi)}{2}}, \quad v \equiv \sqrt{\Omega_b}, \quad r \equiv \sqrt{\Omega_r}, \tag{5.16}$$

in which the density parameters are defined as $\Omega_i = \rho_i/3H^2$ and, as I impose flatness, the dark energy density parameter is given by $\Omega_\phi = 1 - \Omega_{\rm cdm} - \Omega_b - \Omega_r$.

As Amendola (2000) discusses in depth, this system of differential equations has a range of critical points to which solutions are attracted. However, as Baldi (2011b) notes, under a time transformation, $t \longrightarrow -t$, this behaviour is inverted such that the solutions to the equations diverge from the critical points towards non-sensible results in the past. This is relevant because I desire to solve the background to match a set of values of cosmological quantities at $z = 0$ (not at some initial z_i) which become my effective initial conditions in a backwards integration. Baldi (2011b) demonstrates that it is necessary to tune the present value of w_ϕ to find a stable solution in these circumstances. This work also follows this approach, setting the same condition for stability that $\Omega_r \sim 1$ at very early times, deep in the radiation dominated era.

5.2.2 Linear growth equation

The linear growth equations for a cosmology with cold dark matter coupled to dark energy have been derived by solving, in the Newtonian gauge, the linearised conservation equations for matter and scalar field perturbations, δ and $\delta\phi$, and then taking the Newtonian limit so that only modes, $k \gg aH$, much smaller than the horizon are considered (Amendola, 2004; Pettorino and Baccigalupi, 2008). Redefining the prime notation used in § 5.2.1 such that it now labels derivatives with respect to a rather than $\ln a$, the resulting growth equations can be written as

$$\begin{aligned} \delta''_{\rm cdm} &= -\left[\frac{3}{a} - \sqrt{\frac{2}{3}}\beta\phi' H + \frac{H'}{H}\right]\delta'_{\rm cdm} + \frac{3}{2a^2}\left[\Omega_b\delta_b + \Omega_{\rm cdm}\delta_{\rm cdm}\left(1 + \frac{4}{3}\beta^2\right)\right] \\ \delta''_b &= -\left[\frac{3}{a} + \frac{H'}{H}\right]\delta'_b + \frac{3}{2a^2}\left[\Omega_b\delta_b + \Omega_{\rm cdm}\delta_{\rm cdm}\right]. \end{aligned} \tag{5.17}$$

One may note that the effective changes due to the coupling are an 'anti'-friction term, $-\sqrt{2/3}\beta\phi' H$, and an amplification of the effective gravitational constant. The negative sign for the former is justified by the specific definition of the coupling strength such that for $\dot{\phi} > 0$, the coupling strength satisfies $\beta > 0$. This

represents a transference of energy from dark matter to dark energy, thereby reducing the friction term in the growth equation. Solving this system of equations allows one to evaluate the growth function for all matter following the convention of Baldi (2012),

$$G\left(z\right)=\frac{\Omega_{\rm cdm}\left(z\right)\delta_{\rm cdm}\left(z\right)+\Omega_{b}\left(z\right)\delta_{b}\left(z\right)}{\Omega_{\rm cdm}\left(z\right)+\Omega_{b}\left(z\right)}. \tag{5.18}$$

5.2.3 Non-linear growth equation

An aim of this Chapter is computing quantities like δ_c and Δ_v in a coupled dark energy cosmology. These emerge from the non-linear growth equation in the spherical collapse model. However, incorporating CDE into an equation for non-linear growth is not immediately straightforward. A typical approach would be to start with distinct matter density equations of the form in equation (5.9) for the background and for a density perturbation undergoing spherical collapse with a gravitational potential perturbed by the coupling to the scalar field. However, Wintergerst and Pettorino (2010) showed that such a result is incompatible with the form of the relativistic perturbation equations touched upon in § 5.2.2. Specifically, the spherical collapse argument yields no fifth force term associated with the coupled scalar field contributing to the effective gravitational Poisson term,

$$\nabla^{2}\Phi_{\rm eff}=\frac{a^{2}}{2}\rho_{m}\delta_{m}\left[1+\frac{4}{3}\beta^{2}\right], \tag{5.19}$$

where $\Phi_{\rm eff}=\Phi+\sqrt{2/3}\beta\delta\phi$. Instead Wintergerst and Pettorino (2010) derive the non-linear evolution equations directly from the non-relativistic Navier-Stokes equations, which gives

$$\begin{gathered}\dot{\delta}_{m}=-\mathbf{v}_{m}\nabla\delta_{m}-\left[1+\delta_{m}\right]\nabla\cdot\mathbf{v}_{m},\\ \dot{\mathbf{v}}_{m}=-\left[2H-\sqrt{\tfrac{2}{3}}\beta\dot{\phi}\right]\mathbf{v}_{m}-\mathbf{v}_{m}\cdot\nabla\mathbf{v}_{m}-\tfrac{1}{a^{2}}\nabla\left[\Phi+\sqrt{\tfrac{2}{3}}\beta\delta\phi\right],\\ \nabla^{2}\delta\phi=\sqrt{\tfrac{2}{3}}\beta a^{2}\rho_{m}\delta_{m},\\ \nabla^{2}\Phi=\tfrac{a^{2}}{2}\rho_{m}\delta_{m}.\end{gathered} \tag{5.20}$$

In this work only the cold dark matter, but not the baryons, is treated as coupled to the scalar field here. One may again note the effective 'anti'-friction term, $-\sqrt{2/3}\beta\dot{\phi}$ with the negative sign arising from the definition of β that for $\dot{\phi}>0$, one has $\beta>0$. This represents a flow of energy from dark matter to dark

energy, thereby decreasing the effective friction. Wintergerst and Pettorino (2010) propagate the effective gravitational source term through a successful derivation of the non-linear growth equation for CDE. For a cosmology with couplings only to cold dark matter, one finds

$$\begin{aligned}\delta''_{\rm cdm} = &-\left[\frac{3}{a} - \sqrt{\frac{2}{3}}\beta\phi' H \,+\, \frac{H'}{H}\right]\delta'_{\rm cdm} \\ &+ \frac{4}{3}\frac{\left[\delta'_{\rm cdm}\right]^2}{\left(1+\delta_{\rm cdm}\right)} + \frac{3}{2a^2}\left[\Omega_b\delta_b\left(1+\delta_b\right) + \Omega_{\rm cdm}\delta_{\rm cdm}\left(1+\delta_{\rm cdm}\right)\left(1+\frac{4}{3}\beta^2\right)\right] \\ \delta''_b = &-\left[\frac{3}{a} + \frac{H'}{H}\right]\delta'_b + \frac{4}{3}\frac{\left[\delta'_b\right]^2}{\left(1+\delta_b\right)} + \frac{3}{2a^2}\left[\Omega_b\delta_b\left(1+\delta_b\right) + \Omega_{\rm cdm}\delta_{\rm cdm}\left(1+\delta_{\rm cdm}\right)\right].\end{aligned} \tag{5.21}$$

For $\beta = 0$ this system of equations reduces to the familiar case of a non-coupled cosmology. If the system is linearised, then (5.17) is recovered.

5.2.4 Spherical collapse density threshold

The linear collapse overdensity threshold, δ_c, is computed from equation (5.21) by the following procedure. Initial matter overdensities are chosen with the caveat that $\delta_{\rm cdm,in} < 10^{-3}$, so that they reside within the Einstein-de Sitter (EdS) regime where $\delta \propto a$. This also sets the initial growth rate of the perturbation to be $\delta'_{\rm cdm} = 1$. I determine the initial size of the baryon perturbation using the same condition determined by Baldi (2012) that $\delta_b \sim 3.0\times10^{-3}\,\delta_{\rm cdm}$ at $z_{\rm CMB}$. Evolving both the non-linear growth equation from these initial conditions eventually leads to δ diverging due to the presence of the $[\delta']^2$ term. In the spherical collapse approximation, the radius of the density perturbation, $r \sim (1+\delta)^{-1/3}$ tends rapidly to zero when this happens. In my calculations I set the threshold for divergence to be $\delta > 10^8$, following a spherical collapse code provided by Mead (2017). This is an arbitrary choice, with results being insensitive to the divergence threshold provided that a very large value, $\delta \gg 1$, is used.

The linear growth equation does not exhibit this behaviour so δ_c is defined as the solution to linear growth extrapolated to the time that the non-linear growth equation diverges. This represents the time of collapse. By varying the initial density perturbation within the range $2.0\times10^{-7} \leq \delta_{\rm cdm,in} \leq 10^{-3}$, the time re-

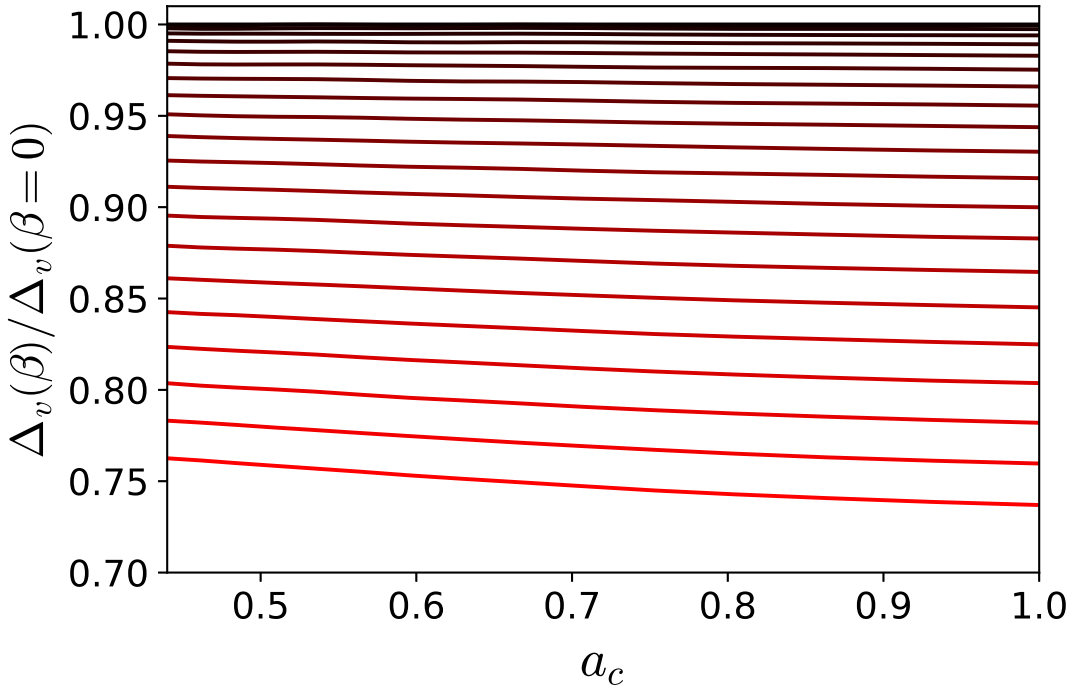

Figure 5.1: Virial density ratios to the ΛCDM case for linearly increasing (redder lines) coupling strengths in the range $0 \leq \beta \leq 0.5$.

quired to reach collapse also changes. In an EdS cosmology, the linear collapse threshold is insensitive to this initial condition. A larger initial density perturbation is compensated for by an earlier collapse time so that δ_c does not change. The threshold can be evaluated analytically as (Padmanabhan, 1993)

$$\delta_c = \frac{3}{20}(12\pi)^{2/3} \approx 1.686, \tag{5.22}$$

for all a_c. In ΛCDM or dark energy cosmologies, the era of matter domination gives way to a regime of acceleration which damps the growth of structure. In these models, the collapse threshold must be calculated numerically. For perturbations with early collapse times, the EdS result holds but δ_c decreases to $\delta_c \approx 1.676$ for those with $a_c = 1$ and $(\Omega_m = 0.3, \Omega_\Lambda = 0.7)$.

Wintergerst and Pettorino (2010) numerically compute the variation in $\delta_c(a_c)$ for different β in a CDE model. They find that for $\beta \lesssim 0.4$ and $z_c \geq 5$, the threshold increases by a factor quadratic in β, given by

$$\delta_c \propto \left(1 + 0.556\beta^2\right). \tag{5.23}$$

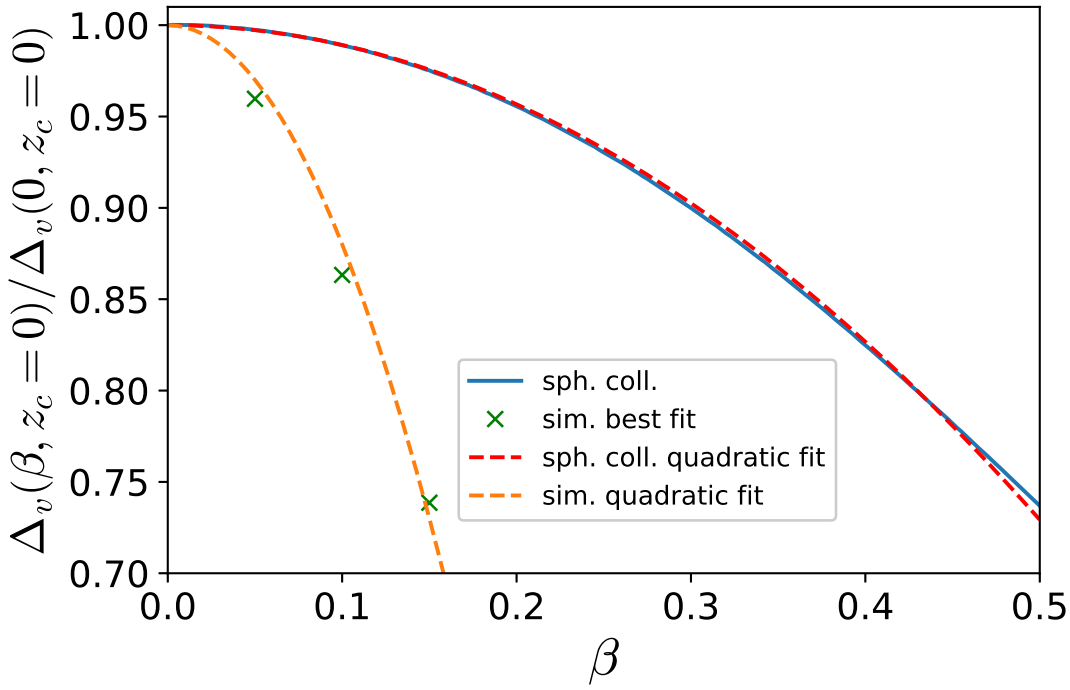

Figure 5.2: Ratio of virial densities computed at a collapse redshift, $z_c = 0$, for a CDE model with coupling strength, β, to the ΛCDM case. The solid blue and dashed red curves correspond respectively to the results from the spherical collapse calculation and a quadratic fit to those results. The green crosses represent the best fit values to the CoDECS simulations. The dashed orange curve is a quadratic fit to those values. The calculation of Δ_v involved baryonic and radiation components.

This functional form also holds for later collapse times up to $z_c = 0$ for smaller coupling strengths, such as $\beta \sim 0.05$, which are within observational limits. Wintergerst and Pettorino (2010) interpret their result by noting that linear and non-linear growth are enhanced in equations (5.17) and (5.21) by terms that enhance the gravitational source (in effect, modifying the effective gravitational strength) by a factor quadratic in β and reduce the Hubble drag term that damps structure growth by a factor linear in β. As a result the density perturbation diverges earlier in the non-linear growth equation, at a time when the linear density value in linear theory is lower. However, the linear growth has also been boosted. In the non-linear growth equation, the contribution from the reduction to the Hubble drag is made negligible during matter domination as it is a linear term. By contrast, in the linear growth equation this has a comparable effect to the enhanced gravitational potential. Therefore, the linear growth responds more than the non-linear growth to increased β, leading to the boost in equation (5.23).

5.2.5 Virial density threshold

In their work, Wintergerst and Pettorino (2010) do not calculate Δ_v for a CDE cosmology. Here, this calculation is performed so that there are spherical collapse predictions for both $\delta_c(\beta)$ and $\Delta_v(\beta)$ functions that can be used in the halo model to compute the CDE matter power. Before discussing the general prescription for calculating the virialisation density, it is necessary to introduce the underlying virial theorem. For a stably bound system of N particles, Clausius (1870) first showed through mechanical arguments that the total kinetic energy averaged over time can be expressed in terms of the separation between the particles and the forces acting on them via

$$\langle T \rangle = \frac{1}{2} \sum_{i=1}^{N} \langle \mathbf{F}_i \cdot \mathbf{r}_i \rangle. \tag{5.24}$$

In cosmology, where composite structures bound under gravity are the principal objects of interest, this theorem is extremely useful. If the force results from a Newtonian gravitational potential energy, $U_{\rm grav} \propto r^{-1}$ the virial theorem simplifies to

$$T = -\frac{1}{2} U, \tag{5.25}$$

where time-averaging is now implicit. The radius of a perturbation grows until it reaches a maximum turn-around radius, $r_{\rm ta}$, at which time the gravitational force

overcomes the background expansion and the perturbation should decrease in size until it completely collapses according to the growth equations. In physical structures, anisotropic collapse prevents this from happening, with virial equilibrium being obtained as the structure reaches a final bound state. Energy conservation, and the fact that $T_{\rm ta} = 0$, leads to the virial radius of an object being given by $r_v = r_{\rm ta}/2$. The turn-around radius can be found in a numerical computation by locating the maximum value of $r \sim (1+\delta)^{-1/3}$ when evolving the non-linear density. The virial density, or specifically $\Delta_v - 1 = (\rho/\bar{\rho})_v - 1$, at collapse is then determined by evaluating the value of the density perturbation at r_v via $r_{\rm ta}$, and multiplying by the factor $(a_c/a_v)^3$. For an EdS cosmology, $\Delta_v \approx 178$, but at $a_c = 1$ for a ΛCDM cosmology this increases to $\Delta_v \approx 310$ (Mead, 2017).

The approximations used here are strictly limited to an EdS cosmology. If there is a dark energy component, it is no longer true that energy is conserved between turn-around and collapse. The self-gravitation of matter is opposed by the effective force of accelerated expansion. Lahav et al. (1991) show that by adding a potential energy term for ΛCDM the relationship between the virial and turn-around radius is given by

$$\frac{r_v}{r_{\rm ta}} = \frac{1-\eta/2}{2-\eta/2} \tag{5.26}$$

where η is determined by the ratio of Λ to the turn-around matter density. This reduces the virial radius. The required modification of the virial theorem for CDE is likely to be less straightforward as the coupling between matter and dark energy complicates the effective gravitational potential and introduces a dependence on β. The focus of this work is not to reformulate the virial theorem itself for CDE, so I retain the $r_v = r_{\rm ta}/2$ approximation for the computation of the virial density. Mead (2017) notes that physical collapsed structures have a much more complex relationship with external forces in dark energy models than the spherical collapse model suggests, and that the shortcomings the EdS virial approximation may be less significant than these limitations. This work takes this position as well, but notes that differences between my spherical collapse predictions and the values of Δ_v indicated by simulations may be substantially different for these reasons. This approach therefore represents a first pass at including β-dependence in Δ_v.

In Figure 5.1, the virial density for different β, relative to the ΛCDM case, is

computed using the approach that has been discussed. Increasing β decreases the virial density, a reflection of the fact that perturbations grow to larger radial sizes at turn-around due to the energy transfer from dark matter to dark energy. It is found that the dependence on the time of collapse is limited but non-zero. Large coupling strengths exhibit greater dependence than smaller coupling strengths. For β that are small enough to be relevant for large scale structure survey forecasts, it is sufficiently minimal that it can be ignored.

The effect of β on the ratio of the CDE virial density at collapse redshift, $z_c = 0$, to the ΛCDM case is shown in Figure 5.2. I find that the quadratic fit,

$$\Delta_v \propto \left(1 - 1.08\,\beta^2\right) \tag{5.27}$$

matches the spherical collapse predictions to within 0.4% for $\beta \lesssim 0.45$. It is interesting to note that the dependence on the coupling strength is quadratic, as is the case for the linear collapse overdensity found by Wintergerst and Pettorino (2010). It may be that it is the β^2 modification to the effective gravitational potential in the non-linear growth equation that drives the changes to Δ_v. The fact that the $\beta\phi'$ modification to the Hubble drag becomes less relevant at the time of collapse in the non-linear growth equation further supports this possibility.

5.3 Fits to simulations

To determine values of collapse and virial quantities in CDE that will permit accurate fits of the HMCODE matter power spectrum to the CoDECS matter power spectra (Baldi, 2012), a least squares fit for the ratio, $R\left(k, \beta; z\right) \equiv P_m\left(k, \beta; z\right) / P_m\left(k, \beta = 0; z\right)$, is performed. This entails computing

$$\chi^2 = \sum_i^n \left(\frac{R_{\rm model}\left(k_i, \beta; z\right)}{R_{\rm sim}\left(k_i, \beta; z\right)} - 1 \right)^2 \tag{5.28}$$

for the (constant) coupling strengths, $\beta = \{0.05, 0.10, 0.15\}$, using the power spectrum provided by CoDECS at $n \sim 60$ different k values. I generate halo model power spectra with the linear power provided by CoDECS results directly. The non-linear 1-halo power depends on the clustering strength $\sigma\left(M\right)$, which in

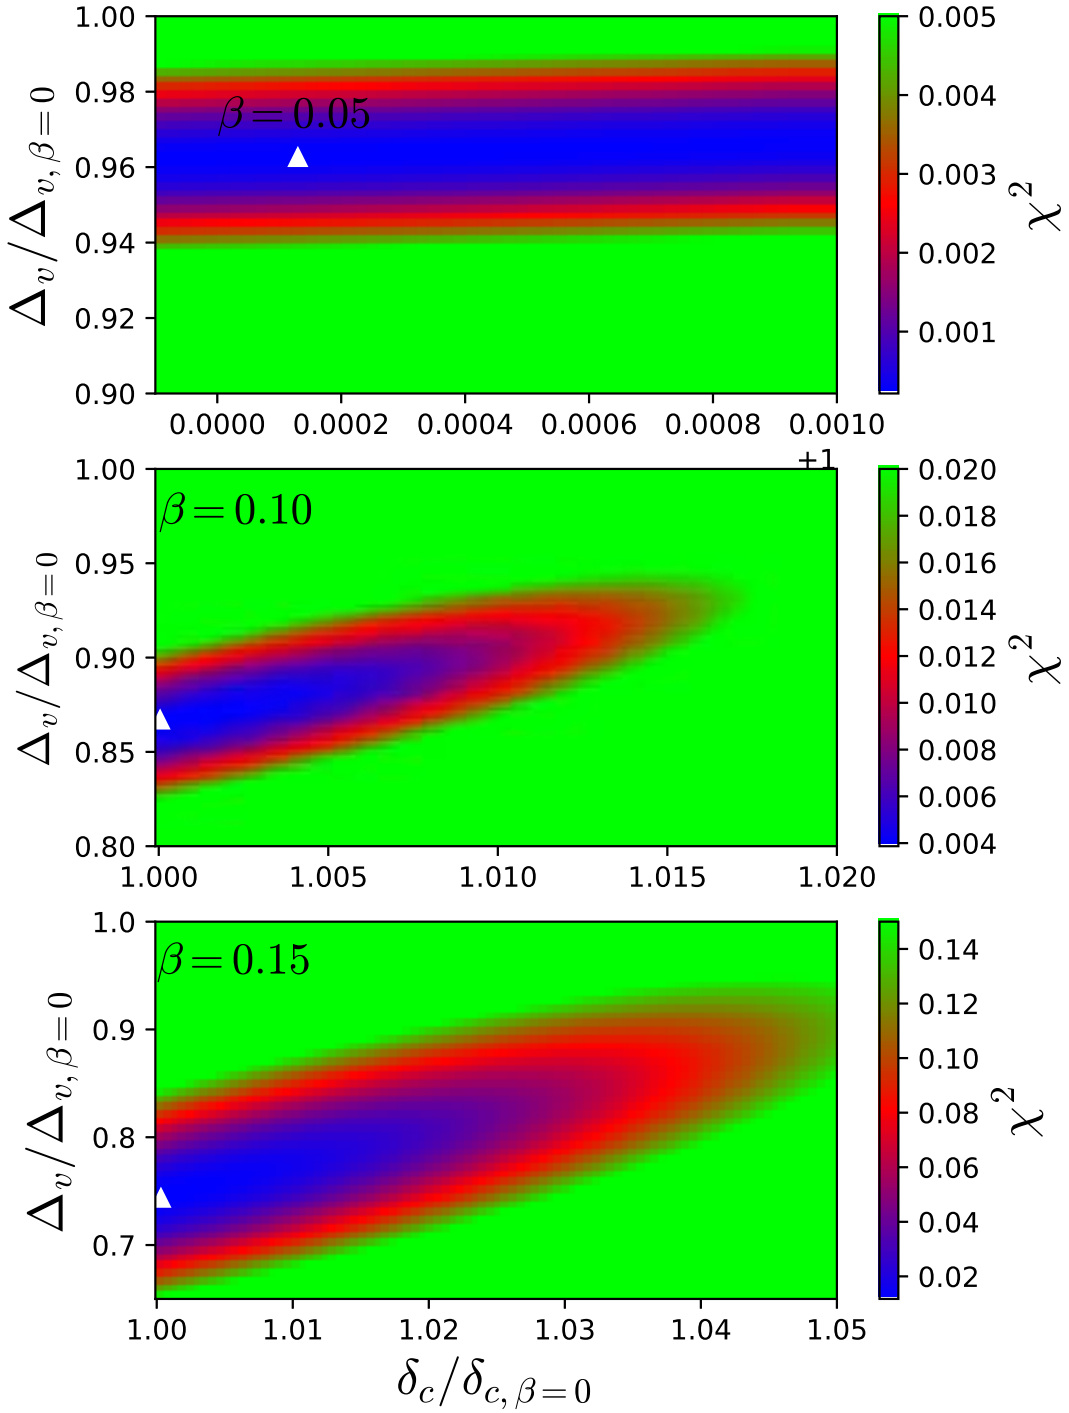

Figure 5.3: χ^2 fits of δ_c and Δ_v to the CoDECS ratio of CDE matter power to ΛCDM power, for different coupling strengths, β. The best fit values are denoted by the white markers. The Mead et al. (2016) power spectrum is used, including their ΛCDM fits for δ_c and Δ_v.

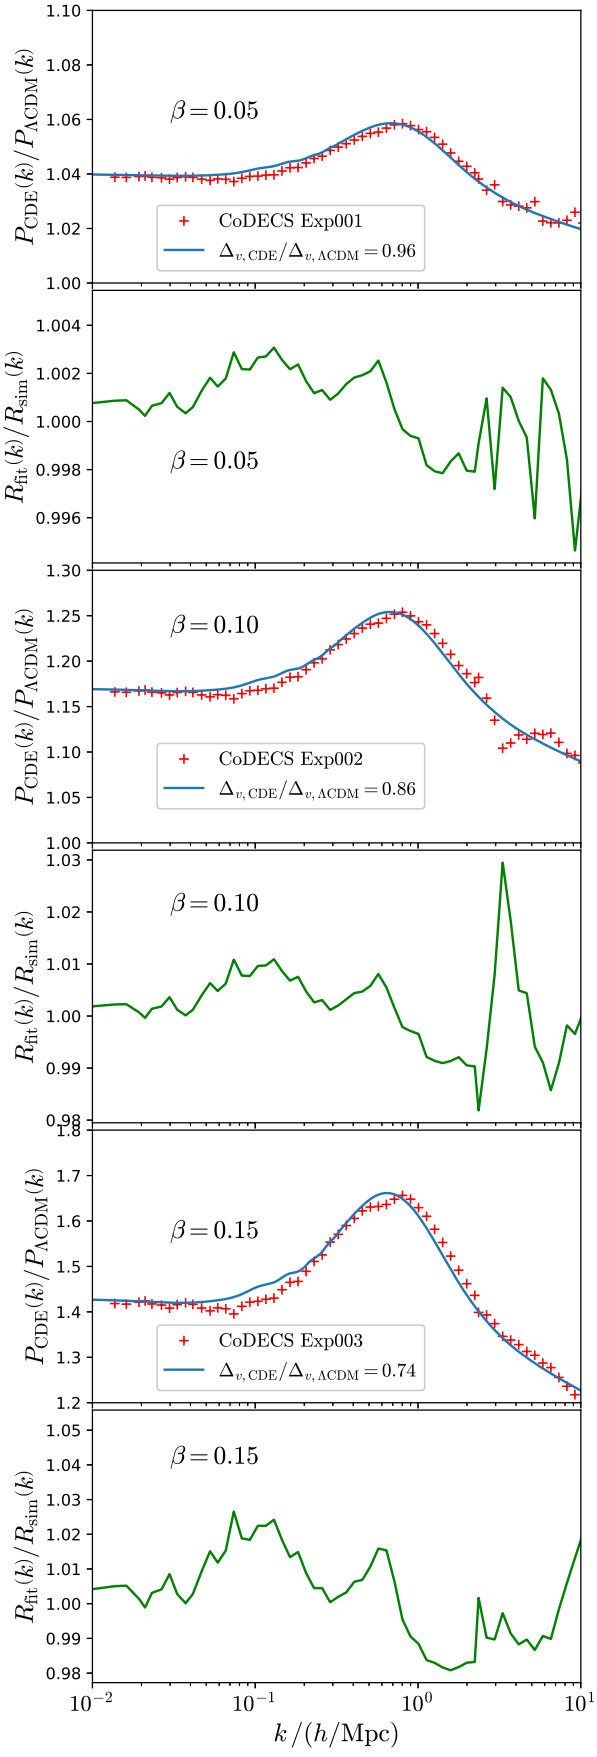

Figure 5.4: Ratios of matter power spectra at $z = 0$ for CDE to ΛCDM for $\beta = \{0.05, 0.10, 0.15\}$. Red crosses correspond to CoDECS simulation results (Baldi, 2012), while blue curves represent results from the halo model. The virial density has been modified to achieve the best fits to simulations, with the accuracy shown in green below the relevant panel.

turn depends on the linear power. As Baldi (2012) evolve their simulations from fixed initial conditions, the growth function and linear power varies between CDE models at $z = 0$. This leads to different $\sigma_8 = \{0.825, 0.875, 0.967\}$ for different CDE models compared to $\sigma_8 = 0.809$ for ΛCDM. I adjust the value of σ_8 appropriately for each $\beta > 0$ model before fitting for Δ_v and δ_c. The dark energy equation of state at $z = 0$ is also slightly varied as $w = \{-0.997, -0.995, -0.992\}$ to allow for a stable computation of the background. The remaining cosmological parameters are given by the WMAP7 results (Komatsu et al., 2011).

In Figure 5.3, I show my χ^2 results for each of the CDE models at $z = 0$ for a range of fractional changes to δ_c and Δ_v from their $\beta = 0$ values in HMCODE. I find that the best fit values are those with essentially no change to δ_c, but with significant changes for Δ_v. Linearly increasing the δ_c change requires the Δ_v change to be linearly reduced to achieve the best fit. However, this is not a line of degeneracy and the quality of the fit still degrades along it away from the best-fit value in the region. If, for example, spherical collapse results for the δ_c change were imposed, the χ^2 values increase notably from the best-fit case regardless of the Δ_v used.

At $z = 0$ the best fit Δ_v values for each β can be described by the fitting function,

$$\Delta_v \propto \left(1 - 12.00\,\beta^2\right), \quad \text{sim. fit} \tag{5.29}$$

to 1.2% accuracy (see Figure 5.2). This shares the quadratic form predicted by the spherical collapse calculations but the β^2 modification requires an amplitude greater by an order of magnitude. My results indicate that spherical collapse is not a sufficient framework in which to predict the impact of CDE on halo density parameters. One explanation for this could be that the approximations that have been used for the virial theorem are insufficient. As has been discussed, the change from the EdS relation between virial and turn-around radii to one correct for CDE would depend on β.

Another possibility becomes apparent by examining the work of Baldi et al. (2010) and Baldi (2011a) that preceded the final development of CoDECS. A limitation of the spherical collapse calculations is that they do not account for the fact that for non-linear virialised structures, particle trajectories include non-

negligible tangential components alongside the radial contribution. Compared to the linear regime, velocity dependent acceleration can have significantly different effects. This directionality does not feature in the spherical collapse model but is accounted for in the CoDECS simulations. Baldi et al. (2010) show that deriving the Euler equation from the perturbed stress-energy tensor in CDE leads to a modified expression for the acceleration,

$$\dot{\mathbf{v}} = -H\left(1 - \sqrt{\frac{2}{3}}\frac{\beta\dot{\phi}}{H}\right)\mathbf{v} - \frac{G\left(1 + 4/3\beta^2\right)M}{r}, \tag{5.30}$$

of a particle at position r with velocity $\mathbf{v}$. Here M is the mass of a spherical mass distribution at $r = 0$. The physical coordinates, $\mathbf{r} = a\left(t\right)\mathbf{x}$ of the particle can be written as

$$\dot{\mathbf{r}} = H\mathbf{r} + a\left(t\right)\dot{\mathbf{x}} \tag{5.31}$$

where the final term is a peculiar velocity. By recasting this equation in terms of the variable, $\mathbf{p} \equiv a^2\left(t\right)\dot{\mathbf{x}}$, the modified acceleration of the j-th particle in a system of N particles can be expressed in the form (Baldi et al., 2010),

$$\dot{\mathbf{p}}_i = \sqrt{\frac{2}{3}}\beta\,\dot{\phi}\,\mathbf{p} + \frac{1}{a}\sum_{i\neq j}\frac{G\left(1 + 4/3\,\beta^2\right)M_j\mathbf{x}_{ij}}{\left|\mathbf{x}_{ij}\right|^3}. \tag{5.32}$$

It is clear that in this formulation the acceleration experienced by particles in CDE is due to a (β modified) Newtonian gravitation and an additional velocity-dependent term that is entirely separate from the physics of the matter distribution. This again represents the effective anti-friction, reducing the Hubble drag and accelerating particles in the direction of their travel. This is attributable to the fact that for $\beta > 0$ and $\dot{\phi} > 0$ the transference of energy is from dark matter to dark energy. There is an analogous term reducing the drag in the growth equations that comes from taking the divergence of the above equation. The velocity divergence determines the flow of material, and hence how the density grows. However, an important distinction here is the effect of the orientation of particle velocities. In the non-linear regime, tangentially aligned velocities will contribute to deviations from the predictions of spherical collapse, which treats the velocity as purely radial. Indeed, Baldi (2011a) find that the velocity-dependent acceleration is the most significant contribution in simulations to changes to the matter power spectrum and the halo density profile in the non-linear regime. By con-

trast, effects like the effective fifth force arising from the modified gravitational potential or the loss of mass from cold dark matter particles contribute relatively little.

Beyond this discussion, this work does not resolve the discrepancy between the predictions of $\Delta_v(\beta)$ from spherical collapse and the best fits found to CoDECS power. A primary interest is whether modifying density thresholds that characterise the halo model can provide sufficiently accurate matches to simulated power. If this is achievable then the required fitting function allows β to be introduced to the forecast analysis performed in previous Chapters. In Figure 5.4 the best fits of HMCODE power spectra, with modified Δ_v, to CoDECS simulations at $z = 0$ are shown for $\beta = \{0.05, 0.10, 0.15\}$. I have appropriately increased σ_8 from its ΛCDM value to the value determined by changes to the growth function for each β. This increase boosts both linear and non-linear power. Δ_v is then artificially decreased (while the concentration stays fixed) to suppress the 1-halo term. The ensuing fits to the CoDECS results are accurate to at the sub-percent level for small $\beta \sim 0.05$, and to within a few percent for $\beta \sim 0.15$ across scales $0.01\,h/\mathrm{Mpc} \leq \mathrm{k} \leq 10\,\mathrm{h/Mpc}$. The spherical collapse overdensity has also been slightly modified but the required changes are less than percent level ($\delta_{c,\beta}/\delta_{c,\Lambda\mathrm{CDM}} = \{1.0001, 1.0002, 1.0012\}$) and so have not been marked on the plot.

For redshifts $z > 0$, I find that the best fits require the fractional virial density response to β to decrease as the cumulative effect of CDE on non-linear structure increases over time. Fits were performed for $\beta = \{0.05, 0.10, 0.5\}$ for 12 CoDECS simulations at different redshifts in the range $0 \leq z \leq 1.25$. Simulations at larger z do not provide sufficient data in the non-linear regime, extending to $k < 5\,h/\mathrm{Mpc}$, so they were not included. It was found that the quadratic dependence on β for Δ_v works well for these redshifts. I capture the diminishing CDE impact on Δ_v at earlier times with a function of $\Omega_m(z;\beta)$, to maintain the HMCODE approach of using physical parameters to make empirical fits. The best functional form was found to be

$$\Delta_v \propto \left(1 - c_1 |\ln \Omega_m(z;\beta)|^{c_2} \beta^2\right), \tag{5.33}$$

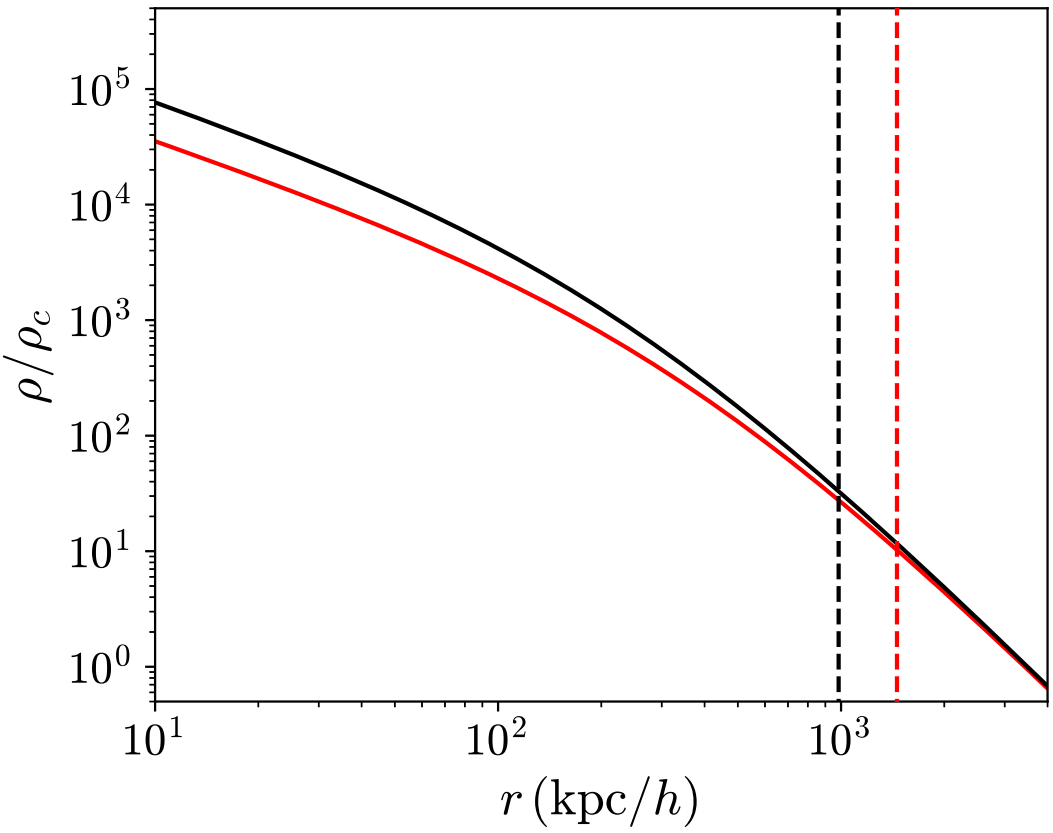

Figure 5.5: Halo density profiles for a halo of mass $2 \times 10^{14}\,h^1\,M_\odot$ for ΛCDM (black) and a coupled dark energy model with $\beta = 0.24$ (red). The corresponding virial limits are marked by the vertical dashed lines at $r_{v,\Lambda\mathrm{CDM}} = 983h^{-1}\,\mathrm{kpc}$ and $r_{v,\mathrm{CDE}} = 1466h^{-1}\,\mathrm{kpc}$.

where $c_1 = 11.829$ and $c_2 = 0.576$. At $z = 0$ this fit is accurate to $\sim 0.5\%$ for $\beta = 0.05$ on all scales of interest. For larger $\beta = 0.10$ and $\beta = 0.15$ the accuracy degrades to $\sim 1\%$ and $\sim 5\%$ respectively. The fitting function of Casas et al. (2016) achieves percent level or better fits for these β, but requires 40 free parameters compared to my two, with c_1 and c_2. My fit also has an physical interpretation within the halo model and a form motivated by collapse arguments. Moreover, the focus of this work is to perform forecasts using a fiducial coupling strength of $\beta = 0.05$, for which the accuracy is well below subpercent level at $z = 0$ and within $\sim 1.25\%$ for $z = 1.25$.

Figure 5.5 shows how the density profile responds to β through the virial density. In this case the chosen value is $\beta = 0.24$, far greater than the bounds placed by observations on the coupling strength. This serves to visibly illustrate the effect

of an increased virial radius (via a decreased Δ_v) on the halo density profile. For a fixed mass, haloes are bloated outwards, which suppresses the density profile on small scales as the concentration is also fixed. On large scales the CDE and ΛCDM profiles converge. Due to the increased virial radius, integrating over the profile returns the same mass in both models. This behaviour bears some similarity to the density profile response to changing the baryonic feedback parameter (see Chapter 3). This is to be expected as it has been shown that

$$\rho\left(r, M\right) = \frac{\rho_s}{\nu^{3\eta}} \frac{1}{\left(\frac{r}{\nu^{\eta} r_s}\right)\left[1 + \left(\frac{r}{\nu^{\eta} r_s}\right)\right]^2}, \qquad r \leq \nu^{\eta} r_v \tag{5.34}$$

for the baryon-halo model, with the virial radius increasing with η for $\nu > 1$. This indicates at the outset that CDE and baryonic phenomena may have degenerate effects on halo structure. The choice of $\beta = 0.24$ also allows for comparisons to Figure 6 of Baldi (2011a) which uses this value to compute the halo density profile through simulations. I find that the shape of my profile, the convergence to ΛCDM on large scales, and the factor ~ 2 suppression of the density on scales of $r = 10\,\mathrm{kpc/h}$ is consistent with their results. This lends further support for the ansatz that non-linear matter power in CDE can mostly be accounted for by an increased halo virial radius or, equivalently, a decreased virial density.

5.4 Forecasts

5.4.1 Power spectra responses

A Fisher analysis encounters challenges in determining the sources of degeneracies between parameters that may be apparent in forecast results. As has been discussed in Chapters 3 and 4, one way to gain an intuitive understanding is to study the response of the underlying probes to varying parameters across scales of interest. In this case, the relevant statistics are the matter power spectrum and the weak lensing convergence spectrum. I show response plots for the former at $z = 0$ in Figure 5.6 and the latter in the redshift bin $0.79 \leq z \leq 0.89$ in Figure 5.7. When computing the lensing power, it was important to incorporate the effects of β on the background (which changes e.g., the comoving distance, χ, via the Hubble parameter) and the growth factor (which is required to determine

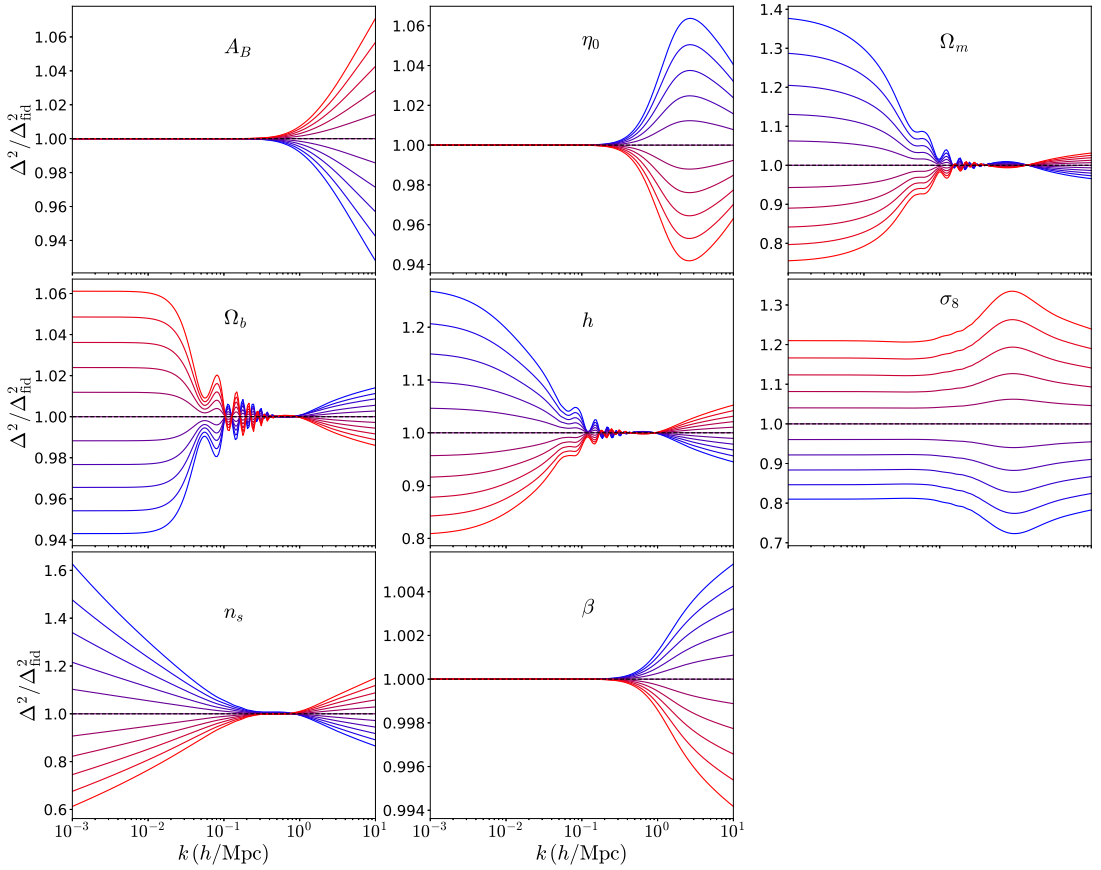

Figure 5.6: Matter power spectra responses at $z = 0$ with respect to a $\beta = 0.05$ CDE fiducial cosmological model. Parameters in $\Theta_{\rm fid} = (A_B, \eta_0, \Omega_m, \Omega_b, h, \sigma_8, n_s, \beta)$ are varied in the range $0.9\,\Theta_{\rm fid} \leq \Theta \leq 1.1\,\Theta_{\rm fid}$ with bluer (redder) lines corresponding to parameter values lower (higher) than the respective fiducial value.

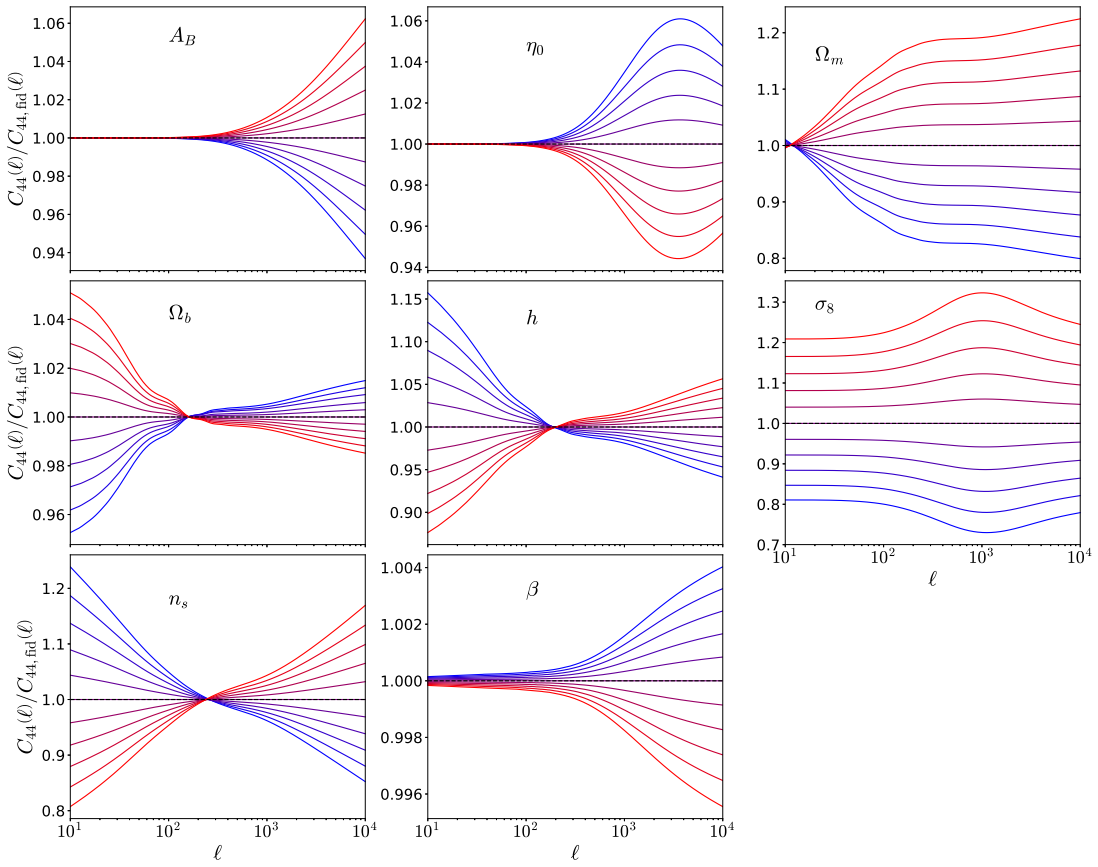

Figure 5.7: Lensing power spectrum responses for the $ij = 44$ redshift bin $0.79 \leq z \leq 0.89$. Parameters in $\Theta_{\mathrm{fid}} = (A_B, \eta_0, \Omega_m, \Omega_b, h, \sigma_8, n_s, \beta)$ are varied in the range $0.9\,\Theta_{\mathrm{fid}} \leq \Theta \leq 1.1\,\Theta_{\mathrm{fid}}$ with bluer (redder) lines corresponding to parameter values lower (higher) than the respective fiducial value.

the contributions of matter power for $z > 0$).

A feature shared by several of the response plots are nodes at the scales at which clustering is fixed to compute σ_8. As power is normalised at these scales, this has a significant impact on the shape of the power response. The response is closely related to the logarithmic derivative of power with respect to a parameter that control the Fisher information, so the normalisation exerts an important influence on the final forecasts.

The results presented here differ from those of e.g., Casas et al. (2016), who normalise power at the initial amplitude, A_s instead in their forecasts. As a result, varying the coupling strength can have a substantial effect on power at $z = 0$ as differences over the entire growth history contribute. By contrast, my normalisation ensures that only the non-linear effects on the virial density impact the power, with the large scale power given by constant $P_{\rm lin}(k)$. A fiducial value of $\beta_{\rm fid} = 0.05$ is used for the forecasts presented here, as this is small enough to be consistent with *Planck* CMB observations that restrict $\beta \lesssim 0.06$. Varying the coupling strength by 10% around this value changes the power by 0.4% at most, as only non-linear scales are influenced. As such, far less stringent constraints would be expected than those reported for β^2 by Casas et al. (2016). It should be emphasised that the amplitude of power on linear scales does respond to β for $z > 0$ via the growth factor. This is relevant for lensing forecasts, although the responses for non-linear modes provide the dominant contribution to the total Fisher information.

The σ_8 normalisation also explains the differences between the shape of the response to β in Figure 5.6 compared to the shape of the ratios of CDE power to ΛCDM power in Figure 5.4. The linear power spectrum at $z = 0$ is fixed in the response plot, with the decrease of power on increasingly non-linear scales being equivalent to the high-k tails in Figure 5.4 having been suppressed relative to a case where the virial density is not modified. The reduction in Δ_v with increasing β in equation (5.33) results in an increased virial radius. For a fixed mass, haloes are therefore more bloated, suppressing the density profile on small scales, as illustrated in Figure 5.5. Therefore, the matter power decreases in this regime.

The $\Delta^2(k)$ responses to varying parameters contribute to the corresponding C_ℓ

responses, and so the former are generally reflected in the latter. However, differences arise from the latter accounting for redshift-dependent influences on structure along the line of sight. The lensing weight function introduces a dependence on background geometry that can have a significant effect, especially if it competes with the contribution from the growth of structure provided by the matter power. This was explored in Chapter 3 for responses of power to varying the parameters controlling the dark energy equation of state. In this case, the $\Delta^2(k)$ and C_ℓ responses are different in sign and shape. For the coupling strength, increasing β suppresses the growth of structure through the transfer of energy from dark matter to dark energy, as discussed previously. This has the consequence of lowering the lensing power. Increasing β also further drives the background expansion, increasing $H(z)$. The comoving distances, χ, that characterise the lensing kernel must decrease to compensate for this. As a result, the lensing power is reduced. Growth and geometry both damp power, so the sign of the lensing power response is the same as the matter power response. However, the influence of β on the background quantities is significantly smaller than the influence from non-linear matter power or from the growth function, which becomes relevant for large scale modes due to integrating over the line of sight. Therefore, there is no substantial change to the shape of the response.

In the response plots, the effects on power of the baryon parameters, A_B and η_0, controlling adiabatic contraction and baryonic feedback are included. As has been seen in preceding Chapters focused on dynamic dark energy and massive neutrinos, marginalising over baryonic effects is one of the major systematics for Stage IV surveys, so it is necessary to explore this in order to make a robust assessment of the prospects for constraining β. The relevant details are once again highlighted here that A_B straightforwardly amplifies the power on non-linear scales, due to it enhancing the small-scale density profile. The feedback parameter is both scale- and mass-dependent such that high-mass haloes are bloated while lower mass ones are reduced. The former has the dominant effect on matter power, and so larger η_0 suppress $\Delta^2(k)$. It also determines the location of a peak in the response, where the largest impact on power occurs for mildly non-linear scales associated with large, high-mass haloes. At earlier times, the peak is translated to deeper non-linear scales as the predominant haloes are smaller.

There are potential degeneracies apparent between β and both A_B and η_0. The coupling strength uniformly suppresses non-linear power while A_B uniformly amplifies it. However, it is clear that the response to β does not grow exponentially at all non-linear scales like the A_B response. This is because it increases the virial density and decreases the virial radius. This changes the effective scale because, for a fixed concentration, it is the quantity $k\,r_v$ that is relevant in the Fourier transform of the NFW density profile (Cooray and Sheth, 2002),

$$u_{NFW}\left(k|M_{NFW}\right) = \frac{1}{\ln\left(1+c\right) - \frac{c}{1+c}} \times \left\{ F\left(k, r_v\right)\cos\left(\frac{kr_v}{c}\right) + G\left(k, r_v\right)\sin\left(\frac{kr_v}{c}\right) - \frac{\sin\left(kr_v\right)}{kr_v\frac{(1+c)}{c}} \right\}, \quad (5.35)$$

where

$$F\left(k, r_v\right) \equiv \mathrm{Ci}\left(\frac{kr_v}{c}\left(1+c\right)\right) - \mathrm{Ci}\left(\frac{kr_v}{c}\right)$$
$$G\left(k, r_v\right) \equiv \mathrm{Si}\left(\frac{kr_v}{c}\left(1+c\right)\right) - \mathrm{Si}\left(\frac{kr_v}{c}\right). \quad (5.36)$$

For highly non-linear scales, the 1-halo matter power is re-scaled by β according to

$$P_{1\mathrm{h}} \sim \left[k\left(\frac{\Delta_v\left(\beta\right)}{\Delta_v\left(\beta = 0\right)} \right)^{-1/3} \right]^{-n} \quad (5.37)$$

For mildly non-linear scales, there is a more complex interplay between trigonometric functions determining $u_{NFW}\left(k|M_{NFW}\right)$, representing the transitional regime between the cubic and linear behaviours of the profile on small and large scales. This manifests as a more gradual power response on intermediate scales, featuring a shape more comparable to the peaked response of power to η_0 in this regime. Some degeneracy may indeed be expected between β and η_0 as the latter also modifies the virial limit, as seen in equation (5.34).

5.4.2 Results

For a Euclid-like weak lensing survey, $(A_B, \eta_0, \Omega_m, \Omega_b, h, n_s, \beta, a_{\mathrm{IA}})$ Fisher matrices are computed. By including the intrinsic alignment amplitude as a free parameter, I account for this significant lensing systematic. However, it should

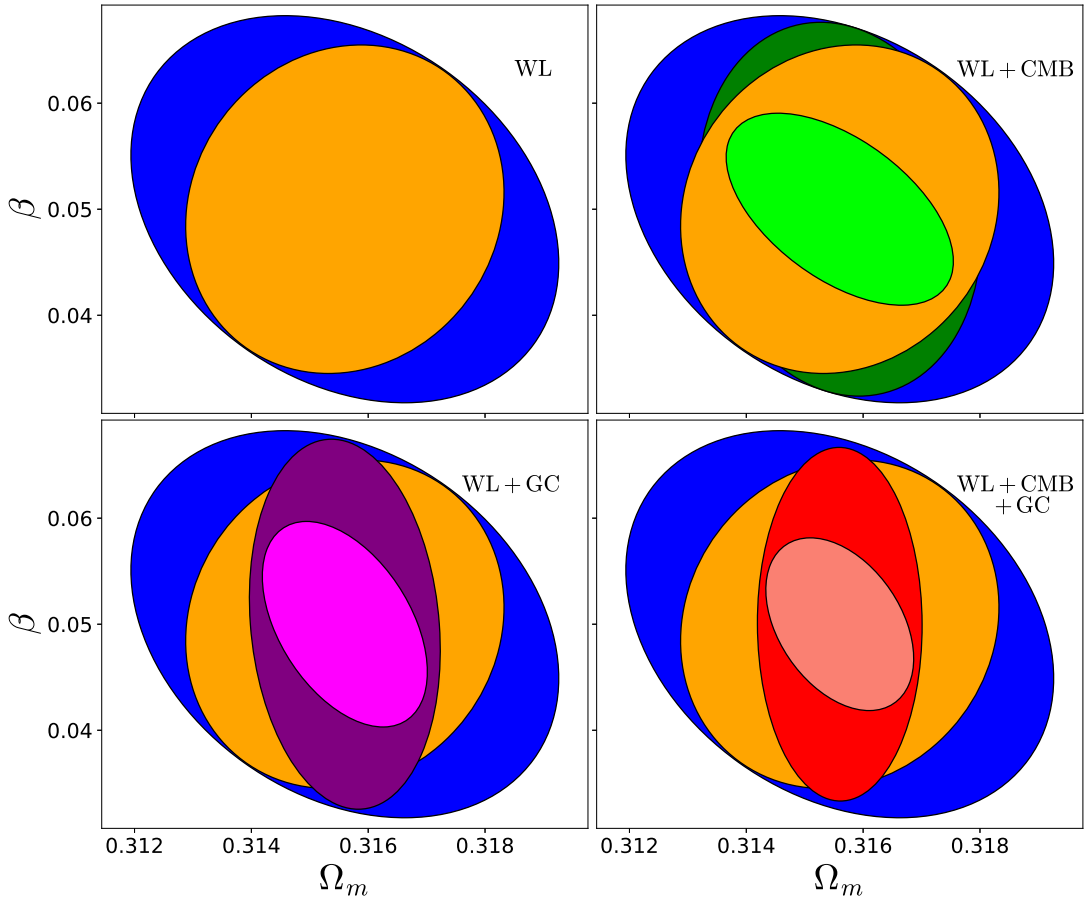

Figure 5.8: 1-σ 2-parameter weak lensing confidence ellipses for Ω_m-β using weak lensing only (top left), with baryon parameters fixed (marginalised over) shown as orange (blue) ellipses; weak lensing and CMB priors (top right), with baryons fixed (marginalised over) shown in light green (dark green); weak lensing and galaxy clustering (bottom left), with baryons fixed (marginalised over) shown in magenta (purple); and weak lensing, CMB priors and galaxy clustering (bottom right), with baryons fixed (marginalised over) shown in pink (red). All other parameters are marginalised over.

	σ	$\sigma_{\rm barfix}$	$R_{\rm bar}$
WL	0.0120	0.0102	1.178
WL + CMB	0.0116	0.0060	1.947
WL + GC	0.0115	0.0064	1.800
WL + CMB + GC	0.0110	0.0054	2.047

Table 5.1: 1-σ error forecasts for a Euclid-like survey of the coupling strength between dark matter and dark energy. I present the errors σ ($\sigma_{\rm barfix}$) for the cases where the baryon parameters are marginalised over (kept fixed), and the ratios of these errors, $R_{\rm bar} = \sigma/\sigma_{\rm barfix}$. Results are included for weak lensing only; weak lensing and CMB priors; weak lensing and galaxy clustering; and weak lensing, CMB priors and galaxy clustering.

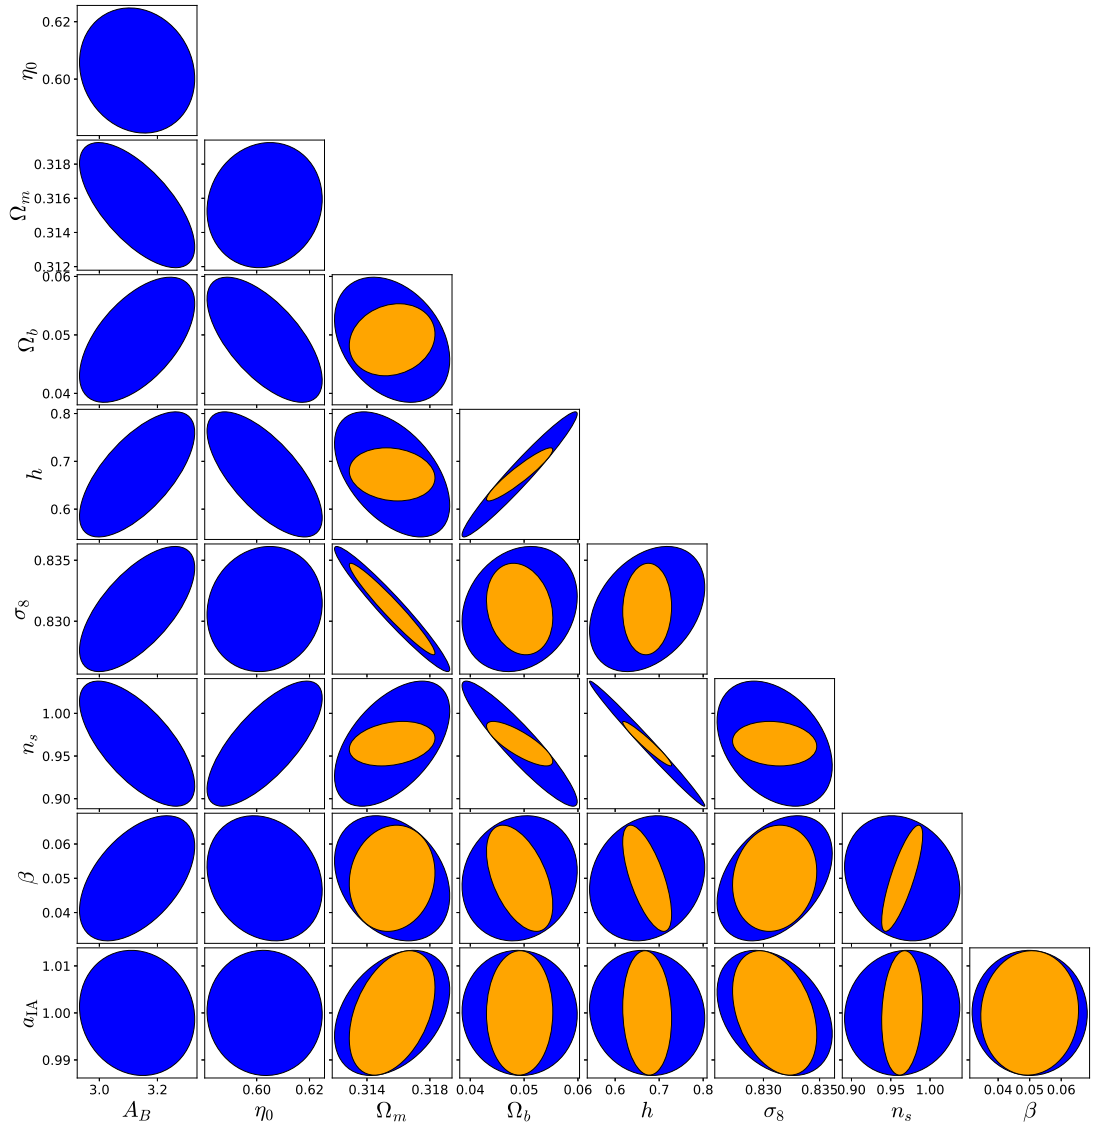

Figure 5.9: 1-σ 2-parameter weak lensing confidence ellipses for $\Theta_{\rm fid} = (A_B, \eta_0, \Omega_m, \Omega_b, h, \sigma_8, n_s, \beta, a_{\rm IA})$. Blue ellipses correspond to marginalisation over all other parameters. Orange ellipses have baryon parameters fixed to their fiducial values so that they are not marginalised over.

be noted that the main impact of intrinsic alignments is to induce bias in measurements. The error forecasts are generally not as severely impacted, as was demonstrated in Chapter 4. Intrinsic alignments are not dealt with in depth in this Chapter so a full analysis of their impact on coupling constraints will be deferred to future work.

The final set of parameters for the spectroscopic galaxy clustering forecast is $(\Omega_m, \Omega_b, h, n_s, \sigma_8, \beta)$, which excludes the baryon parameters as they are not relevant for the linear scales probed. Following the same prescription for galaxy clustering Fisher forecasts in Chapter 4, free parameters controlling the AP effect, growth, galaxy bias and residual shot noise within each redshift bin independently, $\{D_A(z_i), H(z_i), f\sigma_8(z_i), b\sigma_8(z_i), P_{\rm shot}(z_i) | i = 1, ..., N_{z,\rm GC}\}$, are propagated alongside redshift-independent parameters, $\{\omega_m, \omega_b, \omega_\nu, h, n_s\}$, controlling the shape of $P(k)$. Once again, the bias and shot noise parameters are marginalised over before the remaining parameters are projected via a Jacobian transform into the final parameter set.

The greatest constraints on the coupling strength are expected to come from the non-linear regime, as indicated by my power response plots. Therefore, I treat the weak lensing forecast as my baseline set of results to analyse. In this context, information from galaxy clustering is then added in the form of a prior. To acquire the most robust forecasts from a comprehensive range of cosmological information sources, CMB *Planck* priors are also included.

Table 3.2 provides the survey parameters used for a Euclid-like survey (Laureijs et al., 2011). Following the same Fisher methodology as previous Chapters, I use $N_{z,\rm WL} = 10$ redshift bins spanning the range $0 \leq z \leq 2$. The bin widths are chosen such that each bin contains an equal number density of galaxies. I compute the summation over lensing modes in (3.23) using logarithmic intervals, as multiple orders of magnitude are included in the range $\ell_{\rm min} = 10$ to $\ell_{\rm max} = 5000$. In Figure 5.9 the full weak lensing confidence plots are presented from my Fisher analysis for each pair of parameters in $\Theta = (A_B, \eta_0, \Omega_m, \Omega_b, h, \sigma_8, n_s, \beta, a_{\rm IA})$. Wherever relevant, confidence regions are shown for cases with and without marginalisation over baryonic effects. The fiducial values of the cosmological parameters are those given by Table 3.1.

Figure 5.8, shows results for Ω_m-β parameter confidence regions for different combinations of weak lensing, galaxy clustering and CMB priors, in each case with and without marginalising over baryons. The GC forecast methodology and treatment of CMB priors is the same as for work in previous Chapters, which provide the necessary detail. In particular, it should be noted that my GC results incorporate low redshift data from BOSS. I provide the corresponding 1-σ error forecasts for β in Table 5.1. Using the power spectra generated by the halo model with the$\Delta_v(\beta)$ fitting function on non-linear scales, it is found that a $\sigma_{\rm WL} = 0.0120$ constraint can be placed on β using weak lensing alone. However, if one neglects to marginalise over the baryon parameters (or, equivalently, assigns Dirac delta functions centred on their fiducial values as priors), it is found that the constraint improves by 18% to $\sigma = 0.0102$. This represents a significant degree of degradation due to baryon uncertainty. The weak lensing power spectrum response plot in Figure 5.7 foreshadows this. For $\ell \gtrsim 300$, there are similarly shaped C_ℓ responses for β and η_0 in particular. In both cases, their suppression of power is broadened by averaging over effects along the line of sight. Although I do not explore the issue in any depth here, I found that the forecasts for β improve only very slightly when neglecting to include (and marginalise over) intrinsic alignment effects in the lensing model. Chapter 4 discusses the wider implications for cosmological forecasts between cases with and without intrinsic alignments.

By adding CMB and GC priors, one can achieve notable improvements of 40% and 35% to give $\sigma_{\rm WL+CMB} = 0.0060$ and $\sigma_{\rm WL+GC} = 0.0064$ respectively on the β constraint in a scenario where baryon parameters are fixed. In each case this is largely a consequence of benefiting from a repository of information on the other cosmological parameters being marginalised over. For example, the CMB provides significant constraining power for parameters like the spectral index, while GC narrows the errors on parameters like Ω_m as seen in Figure 5.8. For β itself one would also expect some improvements from the GC probe measuring the changes to the linear growth of structure, but as has been discussed previously this is a subdominant effect compared to the non-linear impacts captured by the WL probe. Adding both sets of priors further improves the constraint by another $\sim$10% to $\sigma_{\rm WL+CMB+GC} = 0.0054$. This is not a dramatic improvement on the impact of either set individually, potentially indicating that some parameters are constrained well enough by WL+CMB and by WL+GC that the

impact of marginalising over them is minimal. Hence, further constraints in the WL+CMB+GC case have a limited contribution to the β forecast.

However, in the case where baryons are marginalised over it is found that CMB and GC priors lead to far less substantial improvements of $\sim$10%. As a result, the relative baryon degradation increases from $\sim$18% to a factor $\sim$2. This indicates that the main degeneracies with β are with the baryon parameters, and that these remain a dominant influence on forecasts even when degeneracies with cosmological parameters are significantly mitigated. This is reflected by the power spectra response plots which clearly indicate that the behaviour of β in changing matter and lensing power is mimicked fairly closely by A_B and η_0 but not by the cosmological parameters. An important conclusion, therefore, is that while the baryon degradation is not completely detrimental to coupling strength constraints, it cannot be ignored and, moreover, that multiple probes of large scale structure are unable to meaningfully mitigate its influence. To substantially improve constraints, it will be necessary to enhance the available information on the baryon sector itself, whether from future sets of observations or more powerful simulations. Chapter 3 explores this route for the purposes of mitigating baryon degradation of w_0-w_a constraints. However, the scope of this work is limited to reporting the initial results without attempting to assess the prospects for enhanced baryon priors. This is deferred to future work.

It should be noted that as the errors forecast for β are comparable to or smaller than the accuracy of my halo fitting function with respect to CoDECS simulations, there are grounds to claim that further improvements to the CDE halo model presented here will be required for Euclid-like surveys to accurately constrain the coupling strength.

5.5 Summary and Discussion

This Chapter has explored the influence of a coupling between dark matter and dark energy on non-linear structure in the halo model. I have expressed the virial density threshold, Δ_v, as a function of the coupling strength, β, allowing me to generate matter power spectra in HMCODE that fit the power spectrum

results of the CoDECS simulation suite to percent level accuracy or better at all relevant scales and redshifts for coupling strengths of interest to survey forecasts. I have found that computing the virial density using the spherical collapse formalism, modified by Wintergerst and Pettorino (2010), provided the correct quadratic form of the effects of β on Δ_v, but under predicted the amplitude by an order of magnitude. There has been a discussion of several possible shortcomings of the spherical collapse model that might contribute to an explanation of this discrepancy. These include the use of an Einstein-de Sitter approximation of the relationship between virialisation radius and turn-around radius of a density perturbation undergoing collapse. This approximation incorrectly assumes a conservation of energy between the turn-around and virialisation which is counter to a cosmology in which matter is transferring energy to dark energy. Another possible limitation is that spherical collapse calculations do not account for the complex non-linear behaviour of particles following trajectories in a CDE cosmology on non-linear scales, in particular the importance of a velocity-dependent acceleration which Baldi (2011a) show has the main impact on matter power and the halo density profile.

The $\Delta_v(\beta)$ fit that has been presented here is sufficiently accurate for the purposes of forecasts while using only two free parameters in contrast to the more accurate fit of Casas et al. (2016) which requires 40 free parameters. My fit has also been constructed using physical arguments rather than arbitrary fitting functions of β, k and z. This has provided a basis for interpreting changes to power with β through the underlying structural changes experienced by haloes. In turn this has conferred a significant advantage in extracting the important information from a Fisher forecast analysis, and has allowed for a more thorough diagnosis of any degeneracies limiting the constraints on β.

In the Fisher analysis of this Chapter it has been found that weak lensing forecasts give errors of $\sigma = 0.0102$ for β if one neglects to marginalise over the effects of baryonic astrophysics. However, I have shown through plots comparing the response of matter and weak lensing power to varying different cosmological parameters that the β response is similar to the response to adiabatic contraction and baryonic feedback. As a result, the error degrades by $\sim$20% to $\sigma = 0.0120$. I have applied priors from the CMB and galaxy clustering to assess the scope for mitigation. It has been found that the improvements to cosmological con-

straints from these probes propagate through the Fisher analysis to almost halve the errors on β provided that baryons are not marginalised over. However, when including baryons there is very little overall improvement. This could be understood by noting from the power spectrum response plots and the confidence ellipse plot that only mild degeneracies exist between β and the cosmological parameters. Improving upon cosmology constraints therefore has a limited impact on the main degeneracy with baryons.

Future work in this area would be well-motivated to explore the prospects of reducing the degradation to coupling strength constraints from baryons by using external information on the baryon sector from future data or simulations. As the forecast errors on β are comparable to the accuracy of the CDE power fitting function presented here, it would be desirable to make further improvements to the model in order for a Euclid-like survey to be able accurately constrain CDE parameters. A possible avenue to explore in this regard would be to modify the form of the halo concentration function as the Bullock et al. (2001) prescription is not designed to accommodate beyond ΛCDM physics. From a theoretical point of view it would also be interesting to better understand the reason for the large effect that β has on the virial density and whether spherical collapse arguments can be sufficiently improved upon to predict the values that best fit simulations.

Chapter 6

Conclusions and Future Work

This thesis has explored how modifications to the halo model that account for phenomena not included in the standard ΛCDM formulation can be used for the purposes of constraining these effects and analysing their impact on parameter forecasts for weak lensing and galaxy clustering surveys. A core focus of this work has been to perform Fisher analyses to determine the effects on dark energy and massive neutrino constraints when marginalising over parameters in the halo model that control baryonic astrophysical effects like adiabatic contraction and AGN and supernovae feedback.

I have extended the baryon-halo model of Mead et al. (2015) to include a prescription for inner halo cores (which may be caused by small-scale baryonic effects, or possibly exotic dark matter) alongside parameterisations of adiabatic contraction and baryonic feedback. The motivation behind this approach has been to use empirically motivated modifications directly to halo structure relaions that reproduce accurate power spectra. This has been more useful for the purposes of this work, which relies on accurate power spectra as probes for forecasts, than a precise implementation of specific stellar and gas physics (e.g., Semboloni et al., 2011; Mohammed et al., 2014). The inclusion of the inner cores into my baryon-halo model represents a positive step towards being able to encompass a wider range of baryonic effects on haloes.

However, there is naturally a great deal of work in the future that can be done to improve baryon modelling. An alternative approach to the one in this work is the 'baryonificiation' method. This entails modifying the displacements of particles

within haloes in N-body simulations based on baryon effects in the profile (Schneider and Teyssier, 2015). This has the benefit of having free baryon parameters that can be calibrated, while producing simulations that can accurately capture the matter distribution. Compared to the $\simeq 5$ percent accuracy of HMCODE, power spectra generated through this methodology are accurate to within $\simeq 2$ percent (Schneider et al., 2019). By modifying non-linear density fields directly, it is possible to move beyond the two-point statistics of this work. One of the reasons that baryonification holds promise in the future is that it could benefit from new deep learning techniques that could be applied to the complex behaviours of baryon dynamics in a large dark matter distribution (Chisari et al., 2019). Exploiting the recent advances in machine learning in this way could help to bridge the gap between gravity-only N-body simulations and hydrodynamic simulations.

In Chapter 3, I have shown that forecasts of the w_0-w_a Figure of Merit for a Euclid-like survey experiences a degradation of $\sim$40% due to marginalising over baryon parameters in a Fisher analysis, with $\sim$10% and $\sim$30% degradations for the individual w_0 and w_a 1-σ errors. This is a substantial impact that has significant implications for the next generation of surveys. A valuable direction for future work would be to perform the same analysis for other surveys, for example preliminary results showing comparable baryon impacts for an LSST survey that have not been included in this thesis. In the next few years it will be important to be able to interpret results from multiple Stage IV surveys launching in a similar timeframe. With regards to baryon impacts comparing forecasts between surveys may help to identify survey specific systematics that compound the issue.

I have also shown the advantages that my framework offers in providing constraints on baryon parameters themselves. While forecasting is an important part of this work, it should not be understated that these results are also important for being able to inform the future modelling of baryons in simulations. I have shown that Euclid-like surveys only constrain A_B and η_0 at the 50% and 10% level. Perhaps most significantly, my work indicates that if the inner core can be confidently assumed to vanish then constraints on the adiabatic contraction and feedback parameters reduce from $\sigma_{A_B} = 0.866$ and $\sigma_{\eta_0} = 0.0476$ to $\sigma_{A_B,\mathrm{cusp}} = 0.134$ and $\sigma_{\eta_0,\mathrm{cusp}} = 0.0148$. While constraints on the inner core itself need to be regarded with the suitable context that a Euclid-like survey is unlikely to be reliable probing such small scales, there are interesting implications for the

cusp-core debate if future surveys are able to measure any inner core to within a few kpc, as my results suggest.

In Chapter 4, I have shown that there is an even more severe baryon impact on neutrino mass sum constraints using a Euclid-like weak lensing survey than there is for dark energy. Forecasts of $\sigma_{\Sigma_{\rm NH}} = 0.079\,{\rm eV}$ and $\sigma_{\Sigma_{\rm IH}} = 0.120\,{\rm eV}$ were found, which roughly doubled to $\sigma_{\Sigma_{\rm NH}} = 0.156\,{\rm eV}$ and $\sigma_{\Sigma_{\rm IH}} = 0.204\,{\rm eV}$ when accounting for baryons. This is enough to easily rule out Stage IV surveys being able to distinguish between the normal and inverted hierarchy. By contrast, the effects of marginalising over intrinsic alignments on σ_Σ is fairly minimal. Indeed there are only small differences between forecasts where intrinsic alignments are included in the modelling and when they are excluded.

The baryon impacts for the final physical phenomenon explored in this thesis - a coupling between dark matter and dark energy - are of a similar order to those found for w_0 and w_a, with a degradation to the 1-σ error on β of $\sim$20%. This is not surprising as all these parameters are connected in a similar way to the growth of non-linear structure through their effect on acceleration.

The question of how best to mitigate baryon degradation in each case of dark energy, neutrino mass and coupling strength constraints is one that this thesis has addressed in detail. Chapter 3 found that, by exploring the relative sensitivity of w_0 and w_a Fisher information on non-linear scales with and without baryons, that simply extending the limits of the survey would have restricted improvements. On an intuitive level this can be understood as the baryonic effects being most prominent on non-linear scales and having obscured the substantial information gain otherwise afforded to w_0-w_a constraints on these very scales. While this has been useful to evaluate the additional priors required from the baryon sector to reduce the baryon impact on dark energy constraints, implementing this in a satisfactory way is likely to require future simulations with more sophisticated baryon prescriptions. The greatest gains for dark energy were found through CMB priors, relying on additional information on cosmological parameters propagating through the Fisher analysis. This has shown to essentially completely mitigate the individual impacts on σ_{w_0} and σ_{w_a}, although the the improvements to the Figure of Merit are a slightly more nuanced matter.

The reliance on cosmological priors was again applied to the neutrino forecasts and the coupling strength forecasts, with different degrees of success. In the case of the former, a combination of weak lensing, galaxy clustering and CMB priors was shown to reduce the baryon impact to the $\sim$10% level. I demonstrated how the mass forecasts were still not sufficient for a meaningful distinction of the hierarchies. However, the prospect of future improvements was discussed in detail. I identified that stronger CMB priors on n_s and information from neutrinoless double beta decay were the optimal way to eventually achieve a 2σ distinction. Exploring precisely how these priors may be attained, perhaps by Stage IV CMB surveys, is therefore a valuable avenue of future work.

The case of the coupling strength in CDE proved interesting, as it was found that galaxy clustering and CMB priors could make meaningful improvements to forecasts when baryons were not accounted for. However, this approach was largely ineffective at mitigating the baryon degradation.

Another feature of this work has been to discuss the limitations of the methodologies and modelling that have been employed here. In a sense, this is what much of Chapter 5 has been concerned with. The limitations of modelling coupled dark energy on non-linear scales through the halo model, by relying on fits to simulations with many parameters with no physical basis, have been tackled by identifying the virial density as the key parameter within the halo model that affects the non-linear power as it varies with the coupling strength. This has lead to a far simpler CDE halo model implementation, with a philosophy of interpreting the effects on structure through physical arguments. However, it must be emphasised that this approach is not complete, the work presented here serving more as a first pass effort. Discrepancies between predictions from spherical collapse predictions and values of the virial density required to fit simulations highlight a gap in understanding of how robust a halo model relying on the theory of spherical collapse is in a coupled dark energy cosmology. There are many ways that one could attempt to bridge this gap in the future. One way may be to explore modifying the virial theorem to account for this type of beyond ΛCDM physics, perhaps following similar approaches to those developed for modified gravity or dark energy by previous authors.

The underlying Fisher formalism itself is also subject to the fair critique that,

while a fast and computationally inexpensive tool, it falls short of the more powerful methods like MCMC algorithms that are available. However, the thrust of this work has consistently been to retain an understanding of the physical effects of different parameters on power spectrum probes and then trace these through a Fisher analysis constructed from the derivatives of the power with respect to the parameters. In this sense the Fisher formalism has provided key insights into the diagnostic process of identifying the source of parameter degeneracies and why some methods of mitigation of baryon degradation, for example, are superior to others. A specific shortcoming of the Fisher formalism in this work has been its fundamental limitation when attempting to place constraints on the neutrino mass sum given an underlying uncertainty about which of two models (the hierarchies) is correct. Ideally, a bimodal likelihood would be incorporated into the formalism (which is constructed from likelihood curvature around a single peak of a mulit-variate Gaussian) to correctly account for this. The challenges facing this are certainly considerable. While an analytical approach may fall short, numerical methods may provide a route forward if this project were to be further undertaken in future.

In summary, this thesis has provided a comprehensive assessment of the impact that properly accounting for baryonic astrophysical effects (and to a lesser extent, intrinsic alignments) on large scale structure probes will have on Stage IV parameter constraints. Specific focus has been directed towards dark energy and neutrino masses, two areas of particular interest to the cosmological paradigm. The former is central to our interpretation of the accelerating Universe (and 70% of its matter-energy content) and the latter ties directly into our understanding of fundamental particles that exert a significant influence on cosmology. This work has built on past research by modifying the halo model to account for beyond ΛCDM physics like inner cores and couplings between dark matter and dark energy. This has allowed this thesis to come to wide ranging conclusions on the overall significance of baryon impacts to parameter constraints, finding that for many cosmological parameters the effects are severe for Stage IV surveys. There remain prospects to achieve moderate to significant mitigation in certain cases, although much of this is dependent on the development of improved baryon modelling or more accurate simulations in future, or by achieving stronger cosmological priors from sources beyond weak lensing and galaxy clustering. The overall outlook for Stage IV surveys should therefore be one of measured op-

timism in cases but with due attention directed to overcoming the substantial issues presented by baryonic astrophysics.

Bibliography

Abazajian, K. N., P. Adshead, Z. Ahmed, S. W. Allen, D. Alonso, K. S. Arnold, C. Baccigalupi, J. G. Bartlett, N. Battaglia, B. A. Benson, C. A. Bischoff, J. Borrill, V. Buza, E. Calabrese, R. Caldwell, J. E. Carlstrom, C. L. Chang, T. M. Crawford, F.-Y. Cyr-Racine, F. De Bernardis, T. de Haan, S. di Serego Alighieri, J. Dunkley, C. Dvorkin, J. Errard, G. Fabbian, S. Feeney, S. Ferraro, J. P. Filippini, R. Flauger, G. M. Fuller, V. Gluscevic, D. Green, D. Grin, E. Grohs, J. W. Henning, J. C. Hill, R. Hlozek, G. Holder, W. Holzapfel, W. Hu, K. M. Huffenberger, R. Keskitalo, L. Knox, A. Kosowsky, J. Kovac, E. D. Kovetz, C.-L. Kuo, A. Kusaka, M. Le Jeune, A. T. Lee, M. Lilley, M. Loverde, M. S. Madhavacheril, A. Mantz, D. J. E. Marsh, J. McMahon, P. D. Meerburg, J. Meyers, A. D. Miller, J. B. Munoz, H. N. Nguyen, M. D. Niemack, M. Peloso, J. Peloton, L. Pogosian, C. Pryke, M. Raveri, C. L. Reichardt, G. Rocha, A. Rotti, E. Schaan, M. M. Schmittfull, D. Scott, N. Sehgal, S. Shandera, B. D. Sherwin, T. L. Smith, L. Sorbo, G. D. Starkman, K. T. Story, A. van Engelen, J. D. Vieira, S. Watson, N. Whitehorn, and W. L. Kimmy Wu (2016). CMB-S4 Science Book, First Edition. *arXiv e-prints*, arXiv:1610.02743.

Adler, R. J., B. Casey, and O. C. Jacob (1995). Vacuum catastrophe: An elementary exposition of the cosmological constant problem. *American Journal of Physics*, **63**, pp. 620–626. doi:10.1119/1.17850.

Alam, S., M. Ata, S. Bailey, F. Beutler, D. Bizyaev, J. A. Blazek, A. S. Bolton, J. R. Brownstein, A. Burden, C.-H. Chuang, J. Comparat, A. J. Cuesta, K. S. Dawson, D. J. Eisenstein, S. Escoffier, H. Gil-Marín, J. N. Grieb, N. Hand, S. Ho, K. Kinemuchi, D. Kirkby, F. Kitaura, E. Malanushenko, V. Malanushenko, C. Maraston, C. K. McBride, R. C. Nichol, M. D. Olmstead, D. Oravetz, N. Padmanabhan, N. Palanque-Delabrouille, K. Pan, M. Pellejero-Ibanez, W. J. Percival, P. Petitjean, F. Prada, A. M. Price- Whelan, B. A. Reid, S. A. Rodríguez- Torres, N. A. Roe, A. J. Ross, N. P. Ross, G. Rossi, J. A.

Rubiño-Martín, S. Saito, S. Salazar-Albornoz, L. Samushia, A. G. Sánchez, S. Satpathy, D. J. Schlegel, D. P. Schneider, C. G. Scóccola, H.-J. Seo, E. S. Sheldon, A. Simmons, A. Slosar, M. A. Strauss, M. E. C. Swanson, D. Thomas, J. L. Tinker, R. Tojeiro, M. V. Magaña, J. A. Vazquez, L. Verde, D. A. Wake, Y. Wang, D. H. Weinberg, M. White, W. M. Wood-Vasey, C. Yèche, I. Zehavi, Z. Zhai, and G.-B. Zhao (2017). The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample. *MNRAS*, **470**, pp. 2617–2652. doi:10.1093/mnras/stx721.

Albrecht, A., G. Bernstein, R. Cahn, W. L. Freedman, J. Hewitt, W. Hu, J. Huth, M. Kamionkowski, E. W. Kolb, L. Knox, J. C. Mather, S. Staggs, and N. B. Suntzeff (2006). Report of the Dark Energy Task Force. *ArXiv Astrophysics e-prints*.

Alcock, C. and B. Paczynski (1979). An evolution free test for non-zero cosmological constant. *Nature*, **281**, p. 358. doi:10.1038/281358a0.

Alexander, D. M. and R. C. Hickox (2012). What drives the growth of black holes? *New Astron. Reviews*, **56**(4), pp. 93–121. doi:10.1016/j.newar.2011.11.003.

Allen, S. W., A. E. Evrard, and A. B. Mantz (2011). Cosmological Parameters from Observations of Galaxy Clusters. *ARA&A*, **49**(1), pp. 409–470. doi:10.1146/annurev-astro-081710-102514.

Allison, R., P. Caucal, E. Calabrese, J. Dunkley, and T. Louis (2015). Towards a cosmological neutrino mass detection. *Phys. Rev. D*, **92**, 123535. doi:10.1103/PhysRevD.92.123535.

Alpher, R. A. and R. Herman (1948). Evolution of the Universe. *Nature*, **162**, pp. 774–775. doi:10.1038/162774b0.

Amendola, L. (2000). Coupled quintessence. *Phys. Rev. D*, **62**, 043511. doi:10.1103/PhysRevD.62.043511.

Amendola, L. (2004). Linear and nonlinear perturbations in dark energy models. *Phys. Rev. D*, **69**, 103524. doi:10.1103/PhysRevD.69.103524.

Amendola, L., S. Appleby, A. Avgoustidis, D. Bacon, T. Baker, M. Baldi, N. Bartolo, A. Blanchard, C. Bonvin, S. Borgani, E. Branchini, C. Burrage, S. Camera, C. Carbone, L. Casarini, M. Cropper, C. de Rham, J. P. Dietrich, C. Di

Porto, R. Durrer, A. Ealet, P. G. Ferreira, F. Finelli, J. García-Bellido, T. Giannantonio, L. Guzzo, A. Heavens, L. Heisenberg, C. Heymans, H. Hoekstra, L. Hollenstein, R. Holmes, Z. Hwang, K. Jahnke, T. D. Kitching, T. Koivisto, M. Kunz, G. La Vacca, E. Linder, M. March, V. Marra, C. Martins, E. Majerotto, D. Markovic, D. Marsh, F. Marulli, R. Massey, Y. Mellier, F. Montanari, D. F. Mota, N. J. Nunes, W. Percival, V. Pettorino, C. Porciani, C. Quercellini, J. Read, M. Rinaldi, D. Sapone, I. Sawicki, R. Scaramella, C. Skordis, F. Simpson, A. Taylor, S. Thomas, R. Trotta, L. Verde, F. Vernizzi, A. Vollmer, Y. Wang, J. Weller, and T. Zlosnik (2018). Cosmology and fundamental physics with the Euclid satellite. *Living Reviews in Relativity*, **21**, 2. doi:10.1007/s41114-017-0010-3.

Amendola, L., M. Baldi, and C. Wetterich (2008). Quintessence cosmologies with a growing matter component. *Phys. Rev. D*, **78**(2), 023015. doi:10.1103/PhysRevD.78.023015.

Audren, B., J. Lesgourgues, S. Bird, M. G. Haehnelt, and M. Viel (2013). Neutrino masses and cosmological parameters from a Euclid-like survey: Markov Chain Monte Carlo forecasts including theoretical errors. *Journal of Cosmology and Astro-Particle Physics*, **2013**, 026. doi:10.1088/1475-7516/2013/01/026.

Bailin, D. and A. Love (1994). *Supersymmetric gauge field theory and string theory*.

Baldi, M. (2011a). Clarifying the effects of interacting dark energy on linear and non-linear structure formation processes. *MNRAS*, **414**(1), pp. 116–128. doi:10.1111/j.1365-2966.2011.18263.x.

Baldi, M. (2011b). Time-dependent couplings in the dark sector: from background evolution to non-linear structure formation. *MNRAS*, **411**, pp. 1077–1103. doi:10.1111/j.1365-2966.2010.17758.x.

Baldi, M. (2012). The CoDECS project: a publicly available suite of cosmological N-body simulations for interacting dark energy models. *MNRAS*, **422**, pp. 1028–1044. doi:10.1111/j.1365-2966.2012.20675.x.

Baldi, M., V. Pettorino, L. Amendola, and C. Wetterich (2011). Oscillating nonlinear large-scale structures in growing neutrino quintessence. *MNRAS*, **418**(1), pp. 214–229. doi:10.1111/j.1365-2966.2011.19477.x.

Baldi, M., V. Pettorino, G. Robbers, and V. Springel (2010). Hydrodynamical N-body simulations of coupled dark energy cosmologies. *MNRAS*, **403**(4), pp. 1684–1702. doi:10.1111/j.1365-2966.2009.15987.x.

Baldi, M. and M. Viel (2010). The impact of coupled dark energy cosmologies on the high-redshift intergalactic medium. *MNRAS*, **409**(1), pp. L89–L93. doi:10.1111/j.1745-3933.2010.00954.x.

Ballinger, W. E., J. A. Peacock, and A. F. Heavens (1996). Measuring the cosmological constant with redshift surveys. *MNRAS*, **282**, p. 877. doi:10.1093/mnras/282.3.877.

Bardeen, J. M. (1980). Gauge-invariant cosmological perturbations. *Phys. Rev. D*, **22**, pp. 1882–1905. doi:10.1103/PhysRevD.22.1882.

Barrow, J. D. (1989). What is the principal evidence for the cosmological principle? *QJRAS*, **30**, pp. 163–167.

Bartelmann, M. and P. Schneider (2001). Weak gravitational lensing. *Phys. Rep.*, **340**, pp. 291–472. doi:10.1016/S0370-1573(00)00082-X.

Baumann, D. (2009). TASI Lectures on Inflation. *arXiv e-prints*, arXiv:0907.5424.

Bertin, E. (2011). *Automated Morphometry with SExtractor and PSFEx*, volume 442 of *Astronomical Society of the Pacific Conference Series*, p. 435.

Bird, S., M. Viel, and M. G. Haehnelt (2012). Massive neutrinos and the non-linear matter power spectrum. *MNRAS*, **420**, pp. 2551–2561. doi:10.1111/j.1365-2966.2011.20222.x.

Birkhoff, G. D. and R. E. Langer (1923). *Relativity and modern physics*.

Blumenthal, G. R., S. M. Faber, R. Flores, and J. R. Primack (1986). Contraction of dark matter galactic halos due to baryonic infall. *ApJ*, **301**, pp. 27–34. doi:10.1086/163867.

Born, M. (1926). Quantenmechanik der Stoßvorgänge. *Zeitschrift fur Physik*, **38**, pp. 803–827. doi:10.1007/BF01397184.

Boyle, A. (2018). Understanding the neutrino mass constraints achievable by combining CMB lensing and spectroscopic galaxy surveys. *arXiv e-prints*, arXiv:1811.07636.

Boyle, A. and E. Komatsu (2018). Deconstructing the neutrino mass constraint from galaxy redshift surveys. *Journal of Cosmology and Astro-Particle Physics*, **2018**, 035. doi:10.1088/1475-7516/2018/03/035.

Bridle, S. and L. King (2007). Dark energy constraints from cosmic shear power spectra: impact of intrinsic alignments on photometric redshift requirements. *New Journal of Physics*, **9**, p. 444. doi:10.1088/1367-2630/9/12/444.

Brown, M. L., A. N. Taylor, N. C. Hambly, and S. Dye (2002). Measurement of intrinsic alignments in galaxy ellipticities. *MNRAS*, **333**, pp. 501–509. doi:10.1046/j.1365-8711.2002.05354.x.

Buchdahl, H. A. (1970). Non-linear Lagrangians and cosmological theory. *MNRAS*, **150**, p. 1. doi:10.1093/mnras/150.1.1.

Bullock, J. S., T. S. Kolatt, Y. Sigad, R. S. Somerville, A. V. Kravtsov, A. A. Klypin, J. R. Primack, and A. Dekel (2001). Profiles of dark haloes: evolution, scatter and environment. *MNRAS*, **321**, pp. 559–575. doi:10.1046/j.1365-8711.2001.04068.x.

Burrows, A. (2013). Colloquium: Perspectives on core-collapse supernova theory. *Reviews of Modern Physics*, **85**(1), pp. 245–261. doi:10.1103/RevModPhys.85.245.

Caldwell, R. R., R. Dave, and P. J. Steinhardt (1998). Cosmological Imprint of an Energy Component with General Equation of State. *Physical Review Letters*, **80**, pp. 1582–1585. doi:10.1103/PhysRevLett.80.1582.

Capozzi, F., E. Di Valentino, E. Lisi, A. Marrone, A. Melchiorri, and A. Palazzo (2017). Global constraints on absolute neutrino masses and their ordering. *Phys. Rev. D*, **95**, 096014. doi:10.1103/PhysRevD.95.096014.

Carbone, C., L. Verde, Y. Wang, and A. Cimatti (2011). Neutrino constraints from future nearly all-sky spectroscopic galaxy surveys. *Journal of Cosmology and Astro-Particle Physics*, **2011**, 030. doi:10.1088/1475-7516/2011/03/030.

Casas, S., V. Pettorino, and C. Wetterich (2016). Dynamics of neutrino lumps in growing neutrino quintessence. *Phys. Rev. D*, **94**(10), 103518. doi:10.1103/PhysRevD.94.103518.

Cataneo, M., L. Lombriser, C. Heymans, A. Mead, A. Barreira, S. Bose, and B. Li (2018). On the road to per-cent accuracy: nonlinear reaction of the matter power spectrum to dark energy and modified gravity. *arXiv e-prints*, arXiv:1812.05594.

Chevallier, M. and D. Polarski (2001). Accelerating Universes with Scaling Dark Matter. *International Journal of Modern Physics D*, **10**, pp. 213–223. doi: 10.1142/S0218271801000822.

Chiba, T., T. Okabe, and M. Yamaguchi (2000). Kinetically driven quintessence. *Phys. Rev. D*, **62**, p. 023511. doi:10.1103/PhysRevD.62.023511.

Chisari, N. E., A. J. Mead, S. Joudaki, P. Ferreira, A. Schneider, J. Mohr, T. Tröster, D. Alonso, I. G. McCarthy, S. Martin-Alvarez, J. Devriendt, A. Slyz, and M. P. van Daalen (2019). Modelling baryonic feedback for survey cosmology. *arXiv e-prints*, arXiv:1905.06082.

Clausius, R. (1870). XVI. On a mechanical theorem applicable to heat. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, **40**(265), pp. 122–127. doi:10.1080/14786447008640370.

Clifton, T., P. G. Ferreira, A. Padilla, and C. Skordis (2012). Modified gravity and cosmology. *Phys. Rep.*, **513**, pp. 1–189. doi:10.1016/j.physrep.2012.01.001.

Coe, D. (2009). Fisher Matrices and Confidence Ellipses: A Quick-Start Guide and Software. *arXiv e-prints*, arXiv:0906.4123.

Cooray, A. and R. Sheth (2002). Halo models of large scale structure. *Phys. Rep.*, **372**, pp. 1–129. doi:10.1016/S0370-1573(02)00276-4.

Copeland, D., A. Taylor, and A. Hall (2018). The impact of baryons on the sensitivity of dark energy measurements. *MNRAS*, **480**, pp. 2247–2265. doi: 10.1093/mnras/sty2001.

Copeland, D., A. Taylor, and A. Hall (2019). Towards determining the neutrino mass hierarchy: weak lensing and galaxy clustering forecasts with baryons and intrinsic alignments. *arXiv e-prints*, arXiv:1905.08754.

Copeland, E. J., A. R. Liddle, and D. Wands (1998). Exponential potentials and cosmological scaling solutions. *Phys. Rev. D*, **57**, pp. 4686–4690. doi: 10.1103/PhysRevD.57.4686.

Copeland, E. J., M. Sami, and S. Tsujikawa (2006). Dynamics of Dark Energy. *International Journal of Modern Physics D*, **15**, pp. 1753–1935. doi: 10.1142/S021827180600942X.

Davis, M., G. Efstathiou, C. S. Frenk, and S. D. M. White (1985). The evolution of large-scale structure in a universe dominated by cold dark matter. *ApJ*, **292**, pp. 371–394. doi:10.1086/163168.

Davis, M. and P. J. E. Peebles (1983). A survey of galaxy redshifts. V - The two-point position and velocity correlations. *ApJ*, **267**, pp. 465–482. doi: 10.1086/160884.

de Blok, W. J. G., A. Bosma, and S. McGaugh (2003). Simulating observations of dark matter dominated galaxies: towards the optimal halo profile. *MNRAS*, **340**, pp. 657–678. doi:10.1046/j.1365-8711.2003.06330.x.

Debattista, V. P., B. Moore, T. Quinn, S. Kazantzidis, R. Maas, L. Mayer, J. Read, and J. Stadel (2008). The Causes of Halo Shape Changes Induced by Cooling Baryons: Disks versus Substructures. *ApJ*, **681**(2), pp. 1076–1088. doi:10.1086/587977.

Desjacques, V., D. Jeong, and F. Schmidt (2018). Large-scale galaxy bias. *Phys. Rep.*, **733**, pp. 1–193. doi:10.1016/j.physrep.2017.12.002.

Dubinski, J. and R. G. Carlberg (1991). The structure of cold dark matter halos. *ApJ*, **378**, pp. 496–503. doi:10.1086/170451.

Duffy, A. R., J. Schaye, S. T. Kay, C. Dalla Vecchia, R. A. Battye, and C. M. Booth (2010). Impact of baryon physics on dark matter structures: a detailed simulation study of halo density profiles. *MNRAS*, **405**, pp. 2161–2178. doi: 10.1111/j.1365-2966.2010.16613.x.

Eggen, O. J., D. Lynden-Bell, and A. R. Sandage (1962). Evidence from the motions of old stars that the Galaxy collapsed. *ApJ*, **136**, p. 748. doi: 10.1086/147433.

Einasto, J. (1965). On the Construction of a Composite Model for the Galaxy and on the Determination of the System of Galactic Parameters. *Trudy Astrofizicheskogo Instituta Alma-Ata*, **5**, pp. 87–100.

Einstein, A. (1905a). On the Electrodynamics of Moving Bodies. *Annalen der Physik*, **322**, p. 891.

Einstein, A. (1905b). The Foundation of the General Theory of Relativity. *Annalen der Physik*, **354**, p. 769.

Einstein, A. and W. de Sitter (1932). On the Relation between the Expansion and the Mean Density of the Universe. *Proceedings of the National Academy of Science*, **18**, pp. 213–214. doi:10.1073/pnas.18.3.213.

Eisenstein, D. J., H.-J. Seo, and M. White (2007). On the Robustness of the Acoustic Scale in the Low-Redshift Clustering of Matter. *ApJ*, **664**, pp. 660–674. doi:10.1086/518755.

Fedeli, C. (2014). The clustering of baryonic matter. I: a halo-model approach. *J. Cosmology Astropart. Phys.*, **4**, 028. doi:10.1088/1475-7516/2014/04/028.

Fedeli, C., E. Semboloni, M. Velliscig, M. Van Daalen, J. Schaye, and H. Hoekstra (2014). The clustering of baryonic matter. II: halo model and hydrodynamic simulations. *J. Cosmology Astropart. Phys.*, **8**, 028. doi:10.1088/1475-7516/2014/08/028.

Fenech Conti, I., R. Herbonnet, H. Hoekstra, J. Merten, L. Miller, and M. Viola (2017). Calibration of weak-lensing shear in the Kilo-Degree Survey. *MNRAS*, **467**(2), pp. 1627–1651. doi:10.1093/mnras/stx200.

Feng, J. L. (2010). Dark Matter Candidates from Particle Physics and Methods of Detection. *ARA&A*, **48**, pp. 495–545. doi:10.1146/annurev-astro-082708-101659.

Fisher, R. A. (1935). The Logic of Inductive Inference. *Journal of the Royal Statistical Society*, **98**(1), pp. 39–82. ISSN 09528385.

Fogli, G. L., E. Lisi, A. Marrone, and A. Palazzo (2006). Global analysis of three-flavor neutrino masses and mixings. *Progress in Particle and Nuclear Physics*, **57**, pp. 742–795. doi:10.1016/j.ppnp.2005.08.002.

Friedmann, A. (1922). Über die Krümmung des Raumes. *Zeitschrift fur Physik*, **10**, pp. 377–386. doi:10.1007/BF01332580.

Gerbino, M. (2018). Neutrino properties from cosmology. *arXiv e-prints*, arXiv:1803.11545.

Gillis, B. R. and A. N. Taylor (2019). The effects of calibration on the bias of shear measurements. *MNRAS*, **482**(1), pp. 402–415. doi:10.1093/mnras/sty2637.

Gnedin, O. Y., D. Ceverino, N. Y. Gnedin, A. A. Klypin, A. V. Kravtsov, R. Levine, D. Nagai, and G. Yepes (2011). Halo Contraction Effect in Hydrodynamic Simulations of Galaxy Formation. *ArXiv e-prints*.

Gnedin, O. Y., A. V. Kravtsov, A. A. Klypin, and D. Nagai (2004). Response of Dark Matter Halos to Condensation of Baryons: Cosmological Simulations and Improved Adiabatic Contraction Model. *ApJ*, **616**, pp. 16–26. doi:10.1086/424914.

Governato, F., A. Zolotov, A. Pontzen, C. Christensen, S. H. Oh, A. M. Brooks, T. Quinn, S. Shen, and J. Wadsley (2012). Cuspy no more: how outflows affect the central dark matter and baryon distribution in Λ cold dark matter galaxies. *MNRAS*, **422**, pp. 1231–1240. doi:10.1111/j.1365-2966.2012.20696.x.

Guth, A. H. (1981). Inflationary universe: A possible solution to the horizon and flatness problems. *Phys. Rev. D*, **23**, pp. 347–356. doi:10.1103/PhysRevD.23.347.

Guth, A. H. and S.-H. H. Tye (1980). Phase transitions and magnetic monopole production in the very early universe. *Physical Review Letters*, **44**, pp. 631–635. doi:10.1103/PhysRevLett.44.631.

Hall, A. C. and A. Challinor (2012). Probing the neutrino mass hierarchy with cosmic microwave background weak lensing. *MNRAS*, **425**, pp. 1170–1184. doi:10.1111/j.1365-2966.2012.21493.x.

Hamilton, A. J. S. (1998). Linear Redshift Distortions: a Review. In Hamilton, D. (ed.) *The Evolving Universe*, volume 231, p. 185. doi:10.1007/978-94-011-4960-0˙17.

Hamuy, M., M. M. Phillips, N. B. Suntzeff, R. A. Schommer, J. Maza, and R. Aviles (1996). The Absolute Luminosities of the Calan/Tololo Type IA Supernovae. *AJ*, **112**, p. 2391. doi:10.1086/118190.

Hartle, J. B. and S. W. Hawking (1983). Wave function of the Universe. *Phys. Rev. D*, **28**, pp. 2960–2975. doi:10.1103/PhysRevD.28.2960.

Hauser, M. G. and P. J. E. Peebles (1973). Statistical Analysis of Catalogs of Extragalactic Objects. II. the Abell Catalog of Rich Clusters. *ApJ*, **185**, pp. 757–786. doi:10.1086/152453.

Heisenberg, W. (1927). Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik. *Zeitschrift fur Physik*, **43**, pp. 172–198. doi:10.1007/BF01397280.

Heitmann, K., D. Higdon, M. White, S. Habib, B. J. Williams, E. Lawrence, and C. Wagner (2009). The Coyote Universe. II. Cosmological Models and Precision Emulation of the Nonlinear Matter Power Spectrum. *ApJ*, **705**, pp. 156–174. doi:10.1088/0004-637X/705/1/156.

Heitmann, K., E. Lawrence, J. Kwan, S. Habib, and D. Higdon (2014). The Coyote Universe Extended: Precision Emulation of the Matter Power Spectrum. *ApJ*, **780**, 111. doi:10.1088/0004-637X/780/1/111.

Heitmann, K., M. White, C. Wagner, S. Habib, and D. Higdon (2010). The Coyote Universe. I. Precision Determination of the Nonlinear Matter Power Spectrum. *ApJ*, **715**, pp. 104–121. doi:10.1088/0004-637X/715/1/104.

Heymans, C., L. Van Waerbeke, D. Bacon, J. Berge, G. Bernstein, E. Bertin, S. Bridle, M. L. Brown, D. Clowe, H. Dahle, T. Erben, M. Gray, M. Hetterscheidt, H. Hoekstra, P. Hudelot, M. Jarvis, K. Kuijken, V. Margoniner, R. Massey, Y. Mellier, R. Nakajima, A. Refregier, J. Rhodes, T. Schrabback, and D. Wittman (2006). The Shear Testing Programme - I. Weak lensing analysis of simulated ground-based observations. *MNRAS*, **368**(3), pp. 1323–1339. doi:10.1111/j.1365-2966.2006.10198.x.

Hildebrandt, H., M. Viola, C. Heymans, S. Joudaki, K. Kuijken, C. Blake, T. Erben, B. Joachimi, D. Klaes, L. Miller, C. B. Morrison, R. Nakajima, G. Verdoes Kleijn, A. Amon, A. Choi, G. Covone, J. T. A. de Jong, A. Dvornik, I. Fenech Conti, A. Grado, J. Harnois-Déraps, R. Herbonnet, H. Hoekstra, F. Köhlinger, J. McFarland, A. Mead, J. Merten, N. Napolitano, J. A. Peacock, M. Radovich, P. Schneider, P. Simon, E. A. Valentijn, J. L. van den Busch, E. van Uitert, and L. Van Waerbeke (2017). KiDS-450: cosmological parameter constraints

from tomographic weak gravitational lensing. *MNRAS*, **465**, pp. 1454–1498. doi:10.1093/mnras/stw2805.

Hirata, C. M. and U. Seljak (2004). Intrinsic alignment-lensing interference as a contaminant of cosmic shear. *Phys. Rev. D*, **70**, 063526. doi:10.1103/PhysRevD.70.063526.

Hogg, D. W., D. J. Eisenstein, M. R. Blanton, N. A. Bahcall, J. Brinkmann, J. E. Gunn, and D. P. Schneider (2005). Cosmic Homogeneity Demonstrated with Luminous Red Galaxies. *ApJ*, **624**(1), pp. 54–58. doi:10.1086/429084.

Hu, W., D. J. Eisenstein, and M. Tegmark (1998). Weighing Neutrinos with Galaxy Surveys. *Phys. Rev. E*, **80**, pp. 5255–5258. doi:10.1103/PhysRevLett.80.5255.

Hubble, E. (1929). A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae. *Proceedings of the National Academy of Science*, **15**, pp. 168–173. doi:10.1073/pnas.15.3.168.

Huterer, D. (2002). Weak lensing and dark energy. *Phys. Rev. D*, **65**(6), 063001. doi:10.1103/PhysRevD.65.063001.

Huterer, D., M. Takada, G. Bernstein, and B. Jain (2006). Systematic errors in future weak-lensing surveys: requirements and prospects for self-calibration. *MNRAS*, **366**(1), pp. 101–114. doi:10.1111/j.1365-2966.2005.09782.x.

Ichiki, K. and M. Takada (2012). Impact of massive neutrinos on the abundance of massive clusters. *Phys. Rev. D*, **85**, 063521. doi:10.1103/PhysRevD.85.063521.

Jebsen, J. T. (2005). On the general spherically symmetric solutions of Einstein's gravitational equations in vacuo. *General Relativity and Gravitation*, **37**(12), pp. 2253–2259. ISSN 1572-9532. doi:10.1007/s10714-005-0168-y.

Jee, M. J. and J. A. Tyson (2011). Toward Precision LSST Weak-Lensing Measurement. I. Impacts of Atmospheric Turbulence and Optical Aberration. *PASP*, **123**(903), p. 596. doi:10.1086/660137.

Joachimi, B., R. Mandelbaum, F. B. Abdalla, and S. L. Bridle (2011). Constraints on intrinsic alignment contamination of weak lensing surveys using the MegaZ-LRG sample. *A&A*, **527**, A26. doi:10.1051/0004-6361/201015621.

Kaiser, N. (1987). Clustering in real space and in redshift space. *MNRAS*, **227**, pp. 1–21. doi:10.1093/mnras/227.1.1.

Kaiser, N. (1992). Weak gravitational lensing of distant galaxies. *ApJ*, **388**, pp. 272–286. doi:10.1086/171151.

KamLAND-Zen Collaboration (2016). Search for Majorana Neutrinos near the Inverted Mass Hierarchy Region with KamLAND-Zen. *arXiv e-prints*, arXiv:1605.02889.

Kase, R. and S. Tsujikawa (2019). Dark energy in Horndeski theories after GW170817: A review. *International Journal of Modern Physics D*, **28**(5), 1942005. doi:10.1142/S0218271819420057.

Kazantzidis, S., A. V. Kravtsov, A. R. Zentner, B. Allgood, D. Nagai, and B. Moore (2004). The Effect of Gas Cooling on the Shapes of Dark Matter Halos. *ApJ*, **611**, pp. L73–L76. doi:10.1086/423992.

Kennard, E. H. (1927). Zur Quantenmechanik einfacher Bewegungstypen. *Zeitschrift fur Physik*, **44**, pp. 326–352. doi:10.1007/BF01391200.

Khoury, J. and A. Weltman (2004). Chameleon cosmology. *Phys. Rev. D*, **69**(4), 044026. doi:10.1103/PhysRevD.69.044026.

Kiessling, A., A. N. Taylor, and A. F. Heavens (2011). Simulating the effect of non-linear mode coupling in cosmological parameter estimation. *MNRAS*, **416**, pp. 1045–1055. doi:10.1111/j.1365-2966.2011.19108.x.

Kilbinger, M. (2015). Cosmology with cosmic shear observations: a review. *Reports on Progress in Physics*, **78**(8), 086901. doi:10.1088/0034-4885/78/8/086901.

Koksma, J. F. and T. Prokopec (2011). The Cosmological Constant and Lorentz Invariance of the Vacuum State.

Kolb, E. W. and M. S. Turner (1990). *The early universe.*

Komatsu, E., K. M. Smith, J. Dunkley, C. L. Bennett, B. Gold, G. Hinshaw, N. Jarosik, D. Larson, M. R. Nolta, L. Page, D. N. Spergel, M. Halpern, R. S. Hill, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright (2011). Seven-year Wilkinson Microwave

Anisotropy Probe (WMAP) Observations: Cosmological Interpretation. *ApJS*, **192**(2), 18. doi:10.1088/0067-0049/192/2/18.

Krause, E., T. Eifler, and J. Blazek (2016). The impact of intrinsic alignment on current and future cosmic shear surveys. *MNRAS*, **456**, pp. 207–222. doi: 10.1093/mnras/stv2615.

Krauss, L. M. and M. S. Turner (1995). The cosmological constant is back. *General Relativity and Gravitation*, **27**(11), pp. 1137–1144. doi: 10.1007/BF02108229.

Kuzio de Naray, R., S. S. McGaugh, and W. J. G. de Blok (2008). Mass Models for Low Surface Brightness Galaxies with High-Resolution Optical Velocity Fields. *ApJ*, **676**, 920-943. doi:10.1086/527543.

Lagos, C. d. P., C. G. Lacey, and C. M. Baugh (2013). A dynamical model of supernova feedback: gas outflows from the interstellar medium. *MNRAS*, **436**, pp. 1787–1817. doi:10.1093/mnras/stt1696.

Lahav, O., P. B. Lilje, J. R. Primack, and M. J. Rees (1991). Dynamical effects of the cosmological constant. *MNRAS*, **251**, pp. 128–136. doi: 10.1093/mnras/251.1.128.

Lampton, M., B. Margon, and S. Bowyer (1976). Parameter estimation in X-ray astronomy. *ApJ*, **208**, pp. 177–190. doi:10.1086/154592.

Laureijs, R., J. Amiaux, S. Arduini, J. . Augueres, J. Brinchmann, R. Cole, M. Cropper, C. Dabin, L. Duvet, A. Ealet, and et al. (2011). Euclid Definition Study Report. *ArXiv e-prints*.

Lawrence, E., K. Heitmann, M. White, D. Higdon, C. Wagner, S. Habib, and B. Williams (2010). The Coyote Universe. III. Simulation Suite and Precision Emulator for the Nonlinear Matter Power Spectrum. *ApJ*, **713**, pp. 1322–1331. doi:10.1088/0004-637X/713/2/1322.

Lee, J. and E. Komatsu (2010). Bullet Cluster: A Challenge to ΛCDM Cosmology. *ApJ*, **718**(1), pp. 60–65. doi:10.1088/0004-637X/718/1/60.

Lemaître, G. (1927). Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques. *Annales de la Société Scientifique de Bruxelles*, **47**, pp. 49–59.

Lesgourgues, J., S. Matarrese, M. Pietroni, and A. Riotto (2009). Non-linear power spectrum including massive neutrinos: the time-RG flow approach. *Journal of Cosmology and Astro-Particle Physics*, **2009**, 017. doi:10.1088/1475-7516/2009/06/017.

Lesgourgues, J. and S. Pastor (2006). Massive neutrinos and cosmology. *Phys. Rep.*, **429**, pp. 307–379. doi:10.1016/j.physrep.2006.04.001.

Levi, M. and Z. Vlah (2016). Massive neutrinos in nonlinear large scale structure: A consistent perturbation theory. *ArXiv e-prints*, arXiv:1605.09417.

Lewis, A., A. Challinor, and A. Lasenby (2000). Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models. *ApJ*, **538**, pp. 473–476. doi:10.1086/309179.

Li, B. and J. D. Barrow (2011). N-body simulations for coupled scalar-field cosmology. *Phys. Rev. D*, **83**(2), 024007. doi:10.1103/PhysRevD.83.024007.

Liddle, A. (2003). *An introduction to modern cosmology; 2nd ed.* Wiley, Chichester.

LIGO Scientific Collaboration and Virgo Collaboration (2017). GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. E*, **119**(16), 161101. doi:10.1103/PhysRevLett.119.161101.

Limber, D. N. (1953). The Analysis of Counts of the Extragalactic Nebulae in Terms of a Fluctuating Density Field. *ApJ*, **117**, p. 134. doi:10.1086/145672.

Linde, A. D. (1982). A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, **108**, pp. 389–393. doi:10.1016/0370-2693(82)91219-9.

Linder, E. V. (2003). Exploring the Expansion History of the Universe. *Physical Review Letters*, **90**(9), 091301. doi:10.1103/PhysRevLett.90.091301.

Liu, A., J. R. Pritchard, R. Allison, A. R. Parsons, U. Seljak, and B. D. Sherwin (2016). Eliminating the optical depth nuisance from the CMB with 21 cm cosmology. *Phys. Rev. D*, **93**, 043013. doi:10.1103/PhysRevD.93.043013.

Liu, J., S. Bird, J. M. Zorrilla Matilla, J. C. Hill, Z. Haiman, M. S. Madhavacheril, A. Petri, and D. N. Spergel (2018). MassiveNuS: cosmological massive neutrino

simulations. *Journal of Cosmology and Astro-Particle Physics*, **2018**, 049. doi: 10.1088/1475-7516/2018/03/049.

Lombriser, L. and A. Taylor (2016). Breaking a dark degeneracy with gravitational waves. *Journal of Cosmology and Astro-Particle Physics*, **2016**(3), 031. doi:10.1088/1475-7516/2016/03/031.

LoVerde, M. (2014). Spherical collapse in $\nu\Lambda$CDM. *Phys. Rev. D*, **90**, 083518. doi:10.1103/PhysRevD.90.083518.

LSST Science Collaboration, P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew, J. R. P. Angel, L. Armus, D. Arnett, S. J. Asztalos, T. S. Axelrod, and et al. (2009). LSST Science Book, Version 2.0. *ArXiv e-prints*.

Ma, Z., W. Hu, and D. Huterer (2006). Effects of Photometric Redshift Uncertainties on Weak-Lensing Tomography. *ApJ*, **636**, pp. 21–29. doi:10.1086/497068.

Maartens, R. (2011). Is the Universe homogeneous? *Philosophical Transactions of the Royal Society of London Series A*, **369**(1957), pp. 5115–5137. doi: 10.1098/rsta.2011.0289.

Macciò, A. V., C. Quercellini, R. Mainini, L. Amendola, and S. A. Bonometto (2004). Coupled dark energy: Parameter constraints from N-body simulations. *Phys. Rev. D*, **69**(12), 123516. doi:10.1103/PhysRevD.69.123516.

Macri, L. M., K. Z. Stanek, D. Bersier, L. J. Greenhill, and M. J. Reid (2006). A New Cepheid Distance to the Maser-Host Galaxy NGC 4258 and Its Implications for the Hubble Constant. *ApJ*, **652**, pp. 1133–1149. doi:10.1086/508530.

Maltoni, M., T. Schwetz, M. Tórtola, and J. W. F. Valle (2004). Status of global fits to neutrino oscillations. *New Journal of Physics*, **6**, p. 122. doi: 10.1088/1367-2630/6/1/122.

Mangano, G., G. Miele, and V. Pettorino (2003). Coupled Quintessence and the Coincidence Problem. *Modern Physics Letters A*, **18**(12), pp. 831–842. doi:10.1142/S0217732303009940.

Marsh, D. J. E. (2016). Axion cosmology. *Phys. Rep.*, **643**, pp. 1–79. doi: 10.1016/j.physrep.2016.06.005.

Marsh, D. J. E. and A.-R. Pop (2015). Axion dark matter, solitons and the cusp-core problem. *MNRAS*, **451**, pp. 2479–2492. doi:10.1093/mnras/stv1050.

Martin, J. (2012). Everything you always wanted to know about the cosmological constant problem (but were afraid to ask). *Comptes Rendus Physique*, **13**(6), pp. 566 – 665. ISSN 1631-0705. doi:https://doi.org/10.1016/j.crhy.2012.04.008. Understanding the Dark Universe.

Martizzi, D., I. Mohammed, R. Teyssier, and B. Moore (2014). The biasing of baryons on the cluster mass function and cosmological parameter estimation. *MNRAS*, **440**, pp. 2290–2299. doi:10.1093/mnras/stu440.

Martizzi, D., R. Teyssier, and B. Moore (2013). Cusp-core transformations induced by AGN feedback in the progenitors of cluster galaxies. *MNRAS*, **432**, pp. 1947–1954. doi:10.1093/mnras/stt297.

Martizzi, D., R. Teyssier, B. Moore, and T. Wentz (2012). The effects of baryon physics, black holes and active galactic nucleus feedback on the mass distribution in clusters of galaxies. *MNRAS*, **422**, pp. 3081–3091. doi:10.1111/j.1365-2966.2012.20879.x.

Massara, E., F. Villaescusa-Navarro, and M. Viel (2014). The halo model in a massive neutrino cosmology. *J. Cosmology Astropart. Phys.*, **12**, 053. doi:10.1088/1475-7516/2014/12/053.

Massey, R., H. Hoekstra, T. Kitching, J. Rhodes, M. Cropper, J. Amiaux, D. Harvey, Y. Mellier, M. Meneghetti, L. Miller, S. Paulin-Henriksson, S. r. Pires, R. Scaramella, and T. Schrabback (2013). Origins of weak lensing systematics, and requirements on future instrumentation (or knowledge of instrumentation). *MNRAS*, **429**(1), pp. 661–678. doi:10.1093/mnras/sts371.

Massey, R., T. Kitching, and J. Richard (2010). The dark matter of gravitational lensing. *Reports on Progress in Physics*, **73**(8), 086901. doi:10.1088/0034-4885/73/8/086901.

Mead, A. J. (2017). Spherical collapse, formation hysteresis and the deeply non-linear cosmological power spectrum. *MNRAS*, **464**, pp. 1282–1293. doi:10.1093/mnras/stw2312.

Mead, A. J., C. Heymans, L. Lombriser, J. A. Peacock, O. I. Steele, and H. A. Winther (2016). Accurate halo-model matter power spectra with dark energy, massive neutrinos and modified gravitational forces. *MNRAS*, **459**, pp. 1468–1488. doi:10.1093/mnras/stw681.

Mead, A. J., J. A. Peacock, C. Heymans, S. Joudaki, and A. F. Heavens (2015). An accurate halo model for fitting non-linear cosmological power spectra and baryonic feedback models. *MNRAS*, **454**, pp. 1958–1975. doi:10.1093/mnras/stv2036.

Meszaros, P. (1974). The behaviour of point masses in an expanding cosmological substratum. *A&A*, **37**(2), pp. 225–228.

Milne, E. A. (1936). Relativity, Gravitation, and World-Structure. *Philosophy*, **11**(41), pp. 95–97.

Mohammed, I., D. Martizzi, R. Teyssier, and A. Amara (2014). Baryonic effects on weak-lensing two-point statistics and its cosmological implications. *ArXiv e-prints*.

Natarajan, A., A. R. Zentner, N. Battaglia, and H. Trac (2014). Systematic errors in the measurement of neutrino masses due to baryonic feedback processes: Prospects for stage IV lensing surveys. *Phys. Rev. D*, **90**, 063516. doi:10.1103/PhysRevD.90.063516.

Navarro, J. F., C. S. Frenk, and S. D. M. White (1997). A Universal Density Profile from Hierarchical Clustering. *ApJ*, **490**, pp. 493–508. doi:10.1086/304888.

Navarro, J. F., E. Hayashi, C. Power, A. R. Jenkins, C. S. Frenk, S. D. M. White, V. Springel, J. Stadel, and T. R. Quinn (2004). The inner structure of ΛCDM haloes - III. Universality and asymptotic slopes. *MNRAS*, **349**, pp. 1039–1051. doi:10.1111/j.1365-2966.2004.07586.x.

Nomoto, K. (1982). Accreting white dwarf models for type I supernovae. I - Presupernova evolution and triggering mechanisms. *ApJ*, **253**, pp. 798–810. doi:10.1086/159682.

Oh, S.-H., C. Brook, F. Governato, E. Brinks, L. Mayer, W. J. G. de Blok, A. Brooks, and F. Walter (2011). The Central Slope of Dark Matter Cores in

Dwarf Galaxies: Simulations versus THINGS. *AJ*, **142**, 24. doi:10.1088/0004-6256/142/1/24.

Padmanabhan, T. (1993). *Structure Formation in the Universe*.

Palanque-Delabrouille, N., C. Yèche, J. Baur, C. Magneville, G. Rossi, J. Lesgourgues, A. Borde, E. Burtin, J.-M. LeGoff, J. Rich, M. Viel, and D. Weinberg (2015). Neutrino masses and cosmology with Lyman-alpha forest power spectrum. *J. Cosmology Astropart. Phys.*, **11**, 011. doi:10.1088/1475-7516/2015/11/011.

Parimbelli, G., M. Viel, and E. Sefusatti (2018). On the degeneracy between baryon feedback and massive neutrinos as probed by matter clustering and weak lensing. *ArXiv e-prints*, arXiv:1809.06634.

Peñarrubia, J., A. Pontzen, M. G. Walker, and S. E. Koposov (2012). The Coupling between the Core/Cusp and Missing Satellite Problems. *ApJ*, **759**, L42. doi:10.1088/2041-8205/759/2/L42.

Peacock, J. A. (1999). *Cosmological Physics*.

Peacock, J. A., P. Schneider, G. Efstathiou, J. R. Ellis, B. Leibundgut, S. J. Lilly, and Y. Mellier (2006). ESA-ESO Working Group on "Fundamental Cosmology". Technical report.

Peacock, J. A. and R. E. Smith (2000). Halo occupation numbers and galaxy bias. *MNRAS*, **318**, pp. 1144–1156. doi:10.1046/j.1365-8711.2000.03779.x.

Peccei, R. D. and H. R. Quinn (1977). CP Conservation in the Presence of Pseudoparticles. *Phys. Rev. Lett.*, **38**, pp. 1440–1443. doi:10.1103/PhysRevLett.38.1440.

Peebles, P. J. and B. Ratra (2003). The cosmological constant and dark energy. *Reviews of Modern Physics*, **75**(2), pp. 559–606. doi:10.1103/RevModPhys.75.559.

Penzias, A. A. and R. W. Wilson (1965). A Measurement of Excess Antenna Temperature at 4080 Mc/s. *ApJ*, **142**, pp. 419–421. doi:10.1086/148307.

Penzo, C., A. V. Macciò, M. Baldi, L. Casarini, J. Oñorbe, and A. A. Dutton (2016). Effects of coupled dark energy on the Milky Way and its satellites. *MNRAS*, **461**(3), pp. 2490–2501. doi:10.1093/mnras/stw1502.

Perlmutter, S., G. Aldering, G. Goldhaber, R. A. Knop, P. Nugent, P. G. Castro, S. Deustua, S. Fabbro, A. Goobar, D. E. Groom, I. M. Hook, A. G. Kim, M. Y. Kim, J. C. Lee, N. J. Nunes, R. Pain, C. R. Pennypacker, R. Quimby, C. Lidman, R. S. Ellis, M. Irwin, R. G. McMahon, P. Ruiz-Lapuente, N. Walton, B. Schaefer, B. J. Boyle, A. V. Filippenko, T. Matheson, A. S. Fruchter, N. Panagia, H. J. M. Newberg, W. J. Couch, and T. S. C. Project (1999). Measurements of Ω and Λ from 42 High-Redshift Supernovae. *ApJ*, **517**, pp. 565–586. doi:10.1086/307221.

Pettorino, V. and C. Baccigalupi (2008). Coupled and extended quintessence: Theoretical differences and structure formation. *Phys. Rev. D*, **77**, 103003. doi:10.1103/PhysRevD.77.103003.

Phillips, M. M. (1993). The absolute magnitudes of Type IA supernovae. *ApJ*, **413**, pp. L105–L108. doi:10.1086/186970.

Pietroni, M. (2008). Flowing with time: a new approach to non-linear cosmological perturbations. *Journal of Cosmology and Astro-Particle Physics*, **2008**, 036. doi:10.1088/1475-7516/2008/10/036.

Planck Collaboration, P. A. R. Ade, N. Aghanim, Y. Akrami, P. K. Aluri, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak, E. Battaner, K. Benabed, A. Benoît, A. Benoit-Lévy, J. P. Bernard, M. Bersanelli, P. Bielewicz, J. J. Bock, A. Bonaldi, L. Bonavera, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J. F. Cardoso, B. Casaponsa, A. Catalano, A. Challinor, A. Chamballu, H. C. Chiang, P. R. Christensen, S. Church, D. L. Clements, S. Colombi, L. P. L. Colombo, C. Combet, D. Contreras, F. Couchot, A. Coulais, B. P. Crill, M. Cruz, A. Curto, F. Cuttaia, L. Danese, R. D. Davies, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, F. X. Désert, J. M. Diego, H. Dole, S. Donzelli, O. Doré, M. Douspis, A. Ducout, X. Dupac, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, Y. Fantaye, J. Fergusson, R. Fernandez-Cobos, F. Finelli, O. Forni, M. Frailis, A. A. Fraisse, E. Franceschi, A. Frejsel, A. Frolov, S. Galeotta, S. Galli, K. Ganga, C. Gauthier, T. Ghosh, M. Giard, Y. Giraud-Héraud, E. Gjerløw, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gregorio, A. Gruppuso, J. E. Gudmundsson, F. K. Hansen, D. Hanson, D. L. Harrison, S. Henrot-Versillé, C. Hernández-Monteagudo, D. Herranz, S. R. Hildebrandt, E. Hivon, M. Hobson, W. A.

Holmes, A. Hornstrup, W. Hovest, Z. Huang, K. M. Huffenberger, G. Hurier, A. H. Jaffe, T. R. Jaffe, W. C. Jones, M. Juvela, E. Keihänen, R. Keskitalo, J. Kim, T. S. Kisner, J. Knoche, M. Kunz, H. Kurki-Suonio, G. Lagache, A. Lähteenmäki, J. M. Lamarre, A. Lasenby, M. Lattanzi, C. R. Lawrence, R. Leonardi, J. Lesgourgues, F. Levrier, M. Liguori, P. B. Lilje, M. Linden-Vørnle, H. Liu, M. López-Caniego, P. M. Lubin, J. F. Macías-Pérez, G. Maggio, D. Maino, N. Mand olesi, A. Mangilli, D. Marinucci, M. Maris, P. G. Martin, E. Martínez-González, S. Masi, S. Matarrese, P. McGehee, P. R. Meinhold, A. Melchiorri, L. Mendes, A. Mennella, M. Migliaccio, K. Mikkelsen, S. Mitra, M. A. Miville-Deschênes, D. Molinari, A. Moneti, L. Montier, G. Morgante, D. Mortlock, A. Moss, D. Munshi, J. A. Murphy, P. Naselsky, F. Nati, P. Natoli, C. B. Netterfield, H. U. Nørgaard-Nielsen, F. Noviello, D. Novikov, I. Novikov, C. A. Oxborrow, F. Paci, L. Pagano, F. Pajot, N. Pant, D. Paoletti, F. Pasian, G. Patanchon, T. J. Pearson, O. Perdereau, L. Perotto, F. Perrotta, V. Pettorino, F. Piacentini, M. Piat, E. Pierpaoli, D. Pietrobon, S. Plaszczynski, E. Pointecouteau, G. Polenta, L. Popa, G. W. Pratt, G. Prézeau, S. Prunet, J. L. Puget, J. P. Rachen, R. Rebolo, M. Reinecke, M. Remazeilles, C. Renault, A. Renzi, I. Ristorcelli, G. Rocha, C. Rosset, M. Rossetti, A. Rotti, G. Roudier, J. A. Rubiño-Martín, B. Rusholme, M. Sandri, D. Santos, M. Savelainen, G. Savini, D. Scott, M. D. Seiffert, E. P. S. Shellard, T. Souradeep, L. D. Spencer, V. Stolyarov, R. Stompor, R. Sudiwala, R. Sunyaev, D. Sutton, A. S. Suur-Uski, J. F. Sygnet, J. A. Tauber, L. Terenzi, L. Toffolatti, M. Tomasi, M. Tristram, T. Trombetti, M. Tucci, J. Tuovinen, L. Valenziano, J. Valiviita, B. Van Tent, P. Vielva, F. Villa, L. A. Wade, B. D. Wandelt, I. K. Wehus, D. Yvon, A. Zacchei, J. P. Zibin, and A. Zonca (2016a). Planck 2015 results. XVI. Isotropy and statistics of the CMB. *A&A*, **594**, A16. doi:10.1051/0004-6361/201526681.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, and et al. (2016b). Planck 2015 results. XIII. Cosmological parameters. *A&A*, **594**, A13. doi:10.1051/0004-6361/201525830.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, N. Bartolo, and et al. (2016c). Planck 2015 results. XIV. Dark energy and modified gravity. *A&A*, **594**, A14. doi:10.1051/0004-6361/201525814.

Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak, R. Battye, K. Benabed, J.-P. Bernard, M. Bersanelli, P. Bielewicz, J. J. Bock, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J.-F. Cardoso, J. Carron, A. Challinor, H. C. Chiang, J. Chluba, L. P. L. Colombo, C. Combet, D. Contreras, B. P. Crill, F. Cuttaia, P. de Bernardis, G. de Zotti, J. Delabrouille, J.-M. Delouis, E. Di Valentino, J. M. Diego, O. Doré, M. Douspis, A. Ducout, X. Dupac, S. Dusini, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, Y. Fantaye, M. Farhang, J. Fergusson, R. Fernandez-Cobos, F. Finelli, F. Forastieri, M. Frailis, E. Franceschi, A. Frolov, S. Galeotta, S. Galli, K. Ganga, R. T. Génova-Santos, M. Gerbino, T. Ghosh, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gruppuso, J. E. Gudmundsson, J. Hamann, W. Handley, D. Herranz, E. Hivon, Z. Huang, A. H. Jaffe, W. C. Jones, A. Karakci, E. Keihänen, R. Keskitalo, K. Kiiveri, J. Kim, T. S. Kisner, L. Knox, N. Krachmalnicoff, M. Kunz, H. Kurki-Suonio, G. Lagache, J.-M. Lamarre, A. Lasenby, M. Lattanzi, C. R. Lawrence, M. Le Jeune, P. Lemos, J. Lesgourgues, F. Levrier, A. Lewis, M. Liguori, P. B. Lilje, M. Lilley, V. Lindholm, M. López-Caniego, P. M. Lubin, Y.-Z. Ma, J. F. Macías-Pérez, G. Maggio, D. Maino, N. Mandolesi, A. Mangilli, A. Marcos-Caballero, M. Maris, P. G. Martin, M. Martinelli, E. Martínez-González, S. Matarrese, N. Mauri, J. D. McEwen, P. R. Meinhold, A. Melchiorri, A. Mennella, M. Migliaccio, M. Millea, S. Mitra, M.-A. Miville-Deschênes, D. Molinari, L. Montier, G. Morgante, A. Moss, P. Natoli, H. U. Nørgaard-Nielsen, L. Pagano, D. Paoletti, B. Partridge, G. Patanchon, H. V. Peiris, F. Perrotta, V. Pettorino, F. Piacentini, L. Polastri, G. Polenta, J.-L. Puget, J. P. Rachen, M. Reinecke, M. Remazeilles, A. Renzi, G. Rocha, C. Rosset, G. Roudier, J. A. Rubiño-Martín, B. Ruiz-Granados, L. Salvati, M. Sandri, M. Savelainen, D. Scott, E. P. S. Shellard, C. Sirignano, G. Sirri, L. D. Spencer, R. Sunyaev, A.-S. Suur-Uski, J. A. Tauber, D. Tavagnacco, M. Tenti, L. Toffolatti, M. Tomasi, T. Trombetti, L. Valenziano, J. Valiviita, B. Van Tent, L. Vibert, P. Vielva, F. Villa, N. Vittorio, B. D. Wandelt, I. K. Wehus, M. White, S. D. M. White, A. Zacchei, and A. Zonca (2018). Planck 2018 results. VI. Cosmological parameters. *ArXiv e-prints*.

Pontzen, A. and F. Governato (2012). How supernova feedback turns dark matter cusps into cores. *MNRAS*, **421**, pp. 3464–3471. doi:10.1111/j.1365-

2966.2012.20571.x.

Press, W. H. and P. Schechter (1974). Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation. *ApJ*, **187**, pp. 425–438. doi:10.1086/152650.

Ratra, B. and P. J. E. Peebles (1988). Cosmological consequences of a rolling homogeneous scalar field. *Phys. Rev. D*, **37**, pp. 3406–3427. doi:10.1103/PhysRevD.37.3406.

Read, J. I. and G. Gilmore (2005). Mass loss from dwarf spheroidal galaxies: the origins of shallow dark matter cores and exponential surface brightness profiles. *MNRAS*, **356**, pp. 107–124. doi:10.1111/j.1365-2966.2004.08424.x.

Refregier, A. (2003). Shapelets - I. A method for image analysis. *MNRAS*, **338**(1), pp. 35–47. doi:10.1046/j.1365-8711.2003.05901.x.

Reif, F. (1965). *Fundamentals of Statistical and Thermal Physics*. Waveland Press.

Riess, A. G., B. Bob, C. Burns, and D. Dody (1998). Observational evidence from supernovae for an accelerating universe and a cosmological constant. *Astron. J.*, **116**, pp. 1009–1038.

Ringwald, A. and Y. Y. Y. Wong (2004). Gravitational clustering of relic neutrinos and implications for their detection. *Journal of Cosmology and Astro-Particle Physics*, **2004**(12), 005. doi:10.1088/1475-7516/2004/12/005.

Rubin, V. C., W. K. Ford, Jr., and N. Thonnard (1980). Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 /R = 4kpc/ to UGC 2885 /R = 122 kpc/. *ApJ*, **238**, pp. 471–487. doi:10.1086/158003.

Schaye, J., C. Dalla Vecchia, C. M. Booth, R. P. C. Wiersma, T. Theuns, M. R. Haas, S. Bertone, A. R. Duffy, I. G. McCarthy, and F. van de Voort (2010). The physics driving the cosmic star formation history. *MNRAS*, **402**, pp. 1536–1560. doi:10.1111/j.1365-2966.2009.16029.x.

Schneider, A. and R. Teyssier (2015). A new method to quantify the effects of baryons on the matter power spectrum. *J. Cosmology Astropart. Phys.*, **12**, 049. doi:10.1088/1475-7516/2015/12/049.

Schneider, A., R. Teyssier, J. Stadel, N. E. Chisari, A. M. C. Le Brun, A. Amara, and A. Refregier (2019). Quantifying baryon effects on the matter power spectrum and the weak lensing shear correlation. *Journal of Cosmology and Astro-Particle Physics*, **2019**(3), 020. doi:10.1088/1475-7516/2019/03/020.

Schneider, M. D. and S. Bridle (2010). A halo model for intrinsic alignments of galaxy ellipticities. *MNRAS*, **402**, pp. 2127–2139. doi:10.1111/j.1365-2966.2009.15956.x.

Seljak, U. (2000). Analytic model for galaxy and dark matter clustering. *MNRAS*, **318**, pp. 203–213. doi:10.1046/j.1365-8711.2000.03715.x.

Semboloni, E., H. Hoekstra, and J. Schaye (2013). Effect of baryonic feedback on two- and three-point shear statistics: prospects for detection and improved modelling. *MNRAS*, **434**, pp. 148–162. doi:10.1093/mnras/stt1013.

Semboloni, E., H. Hoekstra, J. Schaye, M. P. van Daalen, and I. G. McCarthy (2011). Quantifying the effect of baryon physics on weak lensing tomography. *MNRAS*, **417**, pp. 2020–2035. doi:10.1111/j.1365-2966.2011.19385.x.

Seo, H.-J. and D. J. Eisenstein (2003). Probing Dark Energy with Baryonic Acoustic Oscillations from Future Large Galaxy Redshift Surveys. *ApJ*, **598**, pp. 720–740. doi:10.1086/379122.

Seo, H.-J. and D. J. Eisenstein (2007). Improved Forecasts for the Baryon Acoustic Oscillations and Cosmological Distance Scale. *ApJ*, **665**, pp. 14–24. doi:10.1086/519549.

Sheth, R. K. and G. Tormen (1999). Large-scale bias and the peak background split. *MNRAS*, **308**, pp. 119–126. doi:10.1046/j.1365-8711.1999.02692.x.

Simon, J. D., A. D. Bolatto, A. Leroy, and L. Blitz (2003). High-Resolution Measurements of the Dark Matter Halo of NGC 2976: Evidence for a Shallow Density Profile. *ApJ*, **596**(2), pp. 957–981. doi:10.1086/378200.

Slipher, V. M. (1917). Radial velocity observations of spiral nebulae. *The Observatory*, **40**, pp. 304–306.

Sprenger, T., M. Archidiacono, T. Brinckmann, S. Clesse, and J. Lesgourgues (2018). Cosmology in the era of Euclid and the Square Kilometre Array. *arXiv e-prints*, arXiv:1801.08331.

Steinhardt, P. J., L. Wang, and I. Zlatev (1999). Cosmological tracking solutions. *Phys. Rev. D*, **59**(12), 123504. doi:10.1103/PhysRevD.59.123504.

Sunyaev, R. A. and I. B. Zeldovich (1980). Microwave background radiation as a probe of the contemporary structure and history of the universe. *ARA&A*, **18**, pp. 537–560. doi:10.1146/annurev.aa.18.090180.002541.

Sunyaev, R. A. and Y. B. Zeldovich (1970). Small-Scale Fluctuations of Relic Radiation. *Astroph. Sp. Sci.*, **7**, pp. 3–19. doi:10.1007/BF00653471.

Takada, M. and B. Jain (2004). Cosmological parameters from lensing power spectrum and bispectrum tomography. *MNRAS*, **348**, pp. 897–915. doi:10.1111/j.1365-2966.2004.07410.x.

Takahashi, R., M. Sato, T. Nishimichi, A. Taruya, and M. Oguri (2012). Revising the Halofit Model for the Nonlinear Matter Power Spectrum. *ApJ*, **761**, 152. doi:10.1088/0004-637X/761/2/152.

Taylor, A. N. and T. D. Kitching (2016). Cosmic Shear Bias and Calibration in Cosmic Shear Studies. *arXiv e-prints*, arXiv:1605.09130.

Taylor, A. N., T. D. Kitching, D. J. Bacon, and A. F. Heavens (2007). Probing dark energy with the shear-ratio geometric test. *MNRAS*, **374**, pp. 1377–1403. doi:10.1111/j.1365-2966.2006.11257.x.

Tegmark, M., A. N. Taylor, and A. F. Heavens (1997). Karhunen-Loève Eigenvalue Problems in Cosmology: How Should We Tackle Large Data Sets? *ApJ*, **480**, pp. 22–35. doi:10.1086/303939.

Teyssier, R., A. Pontzen, Y. Dubois, and J. I. Read (2013). Cusp-core transformations in dwarf galaxies: observational predictions. *MNRAS*, **429**, pp. 3068–3078. doi:10.1093/mnras/sts563.

Troxel, M. A. and M. Ishak (2012). Self-calibration for three-point intrinsic alignment autocorrelations in weak lensing surveys. *MNRAS*, **423**, pp. 1663–1673. doi:10.1111/j.1365-2966.2012.20987.x.

Troxel, M. A. and M. Ishak (2015). The intrinsic alignment of galaxies and its impact on weak gravitational lensing in an era of precision cosmology. *Phys. Rep.*, **558**, pp. 1–59. doi:10.1016/j.physrep.2014.11.001.

van Daalen, M. P., J. Schaye, C. M. Booth, and C. Dalla Vecchia (2011). The effects of galaxy formation on the matter power spectrum: a challenge for precision cosmology. *MNRAS*, **415**, pp. 3649–3665. doi:10.1111/j.1365-2966.2011.18981.x.

Walker, M. G. and J. Peñarrubia (2011). A Method for Measuring (Slopes of) the Mass Profiles of Dwarf Spheroidal Galaxies. *ApJ*, **742**, 20. doi:10.1088/0004-637X/742/1/20.

Wang, B., E. Abdalla, F. Atrio-Barandela, and D. Pavón (2016). Dark matter and dark energy interactions: theoretical challenges, cosmological implications and observational signatures. *Reports on Progress in Physics*, **79**(9), 096901. doi:10.1088/0034-4885/79/9/096901.

Wang, Y., C.-H. Chuang, and C. M. Hirata (2013). Towards more realistic forecasting of dark energy constraints from galaxy redshift surveys. *MNRAS*, **430**, pp. 2446–2453. doi:10.1093/mnras/stt068.

Weinberg, S. (1989). The cosmological constant problem. *Reviews of Modern Physics*, **61**, pp. 1–23. doi:10.1103/RevModPhys.61.1.

Wetterich, C. (1988). Cosmology and the fate of dilatation symmetry. *Nuclear Physics B*, **302**, pp. 668–696. doi:10.1016/0550-3213(88)90193-9.

Wetterich, C. (1995). An asymptotically vanishing time-dependent cosmological "constant". *A&A*, **301**, p. 321.

Wetterich, C. (2007). Growing neutrinos and cosmological selection. *Physics Letters B*, **655**(5-6), pp. 201–208. doi:10.1016/j.physletb.2007.08.060.

Wheeler, J. C. (2007). *Cosmic Catastrophes: Exploding Stars, Black Holes, and Mapping the Universe*. Cambridge University Press, 2 edition. doi:10.1017/CBO9780511536625.

White, M., Y.-S. Song, and W. J. Percival (2009). Forecasting cosmological constraints from redshift surveys. *MNRAS*, **397**, pp. 1348–1354. doi:10.1111/j.1365-2966.2008.14379.x.

Wintergerst, N. and V. Pettorino (2010). Clarifying spherical collapse in coupled dark energy cosmologies. *Phys. Rev. D*, **82**, 103516. doi:10.1103/PhysRevD.82.103516.

Yao, J., M. Ishak, W. Lin, and M. Troxel (2017). Effects of self-calibration of intrinsic alignment on cosmological parameter constraints from future cosmic shear surveys. *Journal of Cosmology and Astro-Particle Physics*, **2017**, 056. doi:10.1088/1475-7516/2017/10/056.

Yao, J., M. Ishak, and M. A. Troxel (2019). Self-calibration method for II and GI types of intrinsic alignments of galaxies. *MNRAS*, **483**, pp. 276–288. doi: 10.1093/mnras/sty3188.

Yèche, C., N. Palanque-Delabrouille, J. Baur, and H. du Mas des Bourboux (2017). Constraints on neutrino masses from Lyman-alpha forest power spectrum with BOSS and XQ-100. *J. Cosmology Astropart. Phys.*, **6**, 047. doi: 10.1088/1475-7516/2017/06/047.

Zeldovich, Y. B. and M. Y. Khlopov (1978). On the concentration of relic magnetic monopoles in the universe. *Physics Letters B*, **79**, pp. 239–241. doi: 10.1016/0370-2693(78)90232-0.

Zentner, A. R., E. Semboloni, S. Dodelson, T. Eifler, E. Krause, and A. P. Hearin (2013). Accounting for baryons in cosmological constraints from cosmic shear. *Phys. Rev. D*, **87**(4), 043509. doi:10.1103/PhysRevD.87.043509.

Zhang, J., Y.-L. Sming Tsai, K. Cheung, and M.-C. Chu (2016). Ultra-Light Axion Dark Matter and its impacts on dark halo structure in N-body simulation. *ArXiv e-prints*.

Zhang, P. (2010). Self-calibration of Gravitational Shear-Galaxy Intrinsic Ellipticity Correlation in Weak Lensing Surveys. *ApJ*, **720**, pp. 1090–1101. doi: 10.1088/0004-637X/720/2/1090.

Zhao, H. (1996). Analytical models for galactic nuclei. *MNRAS*, **278**, pp. 488–496. doi:10.1093/mnras/278.2.488.

Zorrilla Matilla, J. M., Z. Haiman, A. Petri, and T. Namikawa (2017). Geometry and growth contributions to cosmic shear observables. *ArXiv e-prints*.

Zwicky, F. (1933). Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, **6**, pp. 110–127.

Appendix A

Formalism of General Relativity

The foundational material presented here is available in a number of reviews. For a more detailed exposition of these ideas, Peacock (1999) is recommended. Distances in curved spacetime can be encapsulated by a metric tensor, g, which generalises the concept of the scalar product of vectors in Euclidean space. The geometry of spacetime can be described by a differentiable manifold, M, upon which the length of a smooth curve, $\mathbf{x}(t) : (t_0, t_1) \longrightarrow M$, defined for $t_0 < t < t_1$, is given by

$$\int_{t_0}^{t_1} \mathrm{d}t \sqrt{g(X, X)|_{\mathbf{x}(t)}}. \tag{A.1}$$

Here, X is a tangent vector of the curve such that

$$X \equiv \frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t}. \tag{A.2}$$

The tensor properties of the metric allow it to be written as the outer product

$$g = g_{\mu\nu}\mathrm{d}x^{\mu} \otimes \mathrm{d}x^{\nu} \longrightarrow \mathrm{d}s^2 = g_{\mu\nu}\mathrm{d}x^{\mu}\mathrm{d}x^{\nu} = c^2\mathrm{d}\tau^2, \tag{A.3}$$

in which the right-hand side of the expression is a simplified notation that uses the cosmological convention of denoting $\mathrm{d}s$ as the infinitesimal separation of two points and τ as the quantity referred to as the 'proper time'. The components of the four-vector, $\mathrm{d}x$, may be evaluated as those given by $\mathrm{d}x_{\mu}$ by an observer in one frame, but observers in all frames will evaluate $\mathrm{d}s^2$ equally as this is the invariant quantity. Index notation, $\mu, \nu = \{0, 1, 2, 3\}$, will be used to denote four-vectors hereafter.

With a suitably defined metric, the General Relativity equation of motion accommodates the equivalence principle by allowing general transformations from freely falling frames to accelerating frames. There are a number of approaches to deriving the result, known as the geodesic equation. Arguably the most elegant is through applying the standard Euler-Lagrange formalism for extremising curves to timelike trajectories (defined by $\mathrm{d}s^2 < 0$)[1]. The extremal case is one in which a small perturbation to the curve has no effect on its length to first order. The interval between two points along such a trajectory, $\lambda(u)$, is the proper time, which can be expressed as the functional

$$\tau[\lambda(u)] = \int_0^1 \mathrm{d}u \sqrt{-g_{\mu\nu}(x(\lambda(u)))\,\dot{x}^\mu(\lambda(u))\,\dot{x}^\nu(\lambda(u))}, \tag{A.4}$$

where u is such that $\lambda(0) < \lambda(u) < \lambda(1)$ and a dot represents derivatives with respect to u. The curve of extremal proper time is that which satisfies the Euler-Lagrange equation,

$$\frac{\mathrm{d}}{\mathrm{d}u}\left(\frac{\partial G}{\partial \dot{x}^\mu}\right) - \frac{\partial G}{\partial x^\mu} = 0, \tag{A.5}$$

where G is the integrand of equation (A.4). Using the symmetric property of the metric and replacing u with the proper time (a physical parameter that is more convenient to use while solving the Euler-Lagrange equations), the geodesic equation is eventually derived as

$$\frac{\mathrm{d}^2 x^\mu}{\mathrm{d}\tau^2} + \Gamma^\mu_{\nu\rho}\frac{\mathrm{d}x^\nu}{\mathrm{d}\tau}\frac{\mathrm{d}x^\rho}{\mathrm{d}\tau} = 0, \tag{A.6}$$

where the Christoffel symbols that determine the metric connection are constructed entirely from the metric and derivatives of the metric (which are written using the convention, $g_{\mu\nu,\rho} \equiv \partial g_{\mu\nu}/\partial x^\rho$) such that

$$\Gamma^\mu_{\nu\rho} = \frac{1}{2}g^{\mu\sigma}\left(g_{\sigma\nu,\rho} + g_{\sigma\rho,\nu} - g_{\nu\rho,\sigma}\right). \tag{A.7}$$

[1]Timelike curves are followed by massive particles, which can be accelerated asymptotically towards the speed of light while never exceeding it. Photons, however, travel along null geodesics, curves for which $\mathrm{d}s^2 = 0$. In Special Relativity, the world lines of a Minkowski diagram associated with photons arriving at and emerging from some location define the causal regions that an observer at this location can (or will ever be able to) interact with. The resulting light cone is no longer necessarily defined by 45 degree lines in the warped spacetime of General Relativity. Beyond the light cone are regions defined by spacelike geodesics which are inaccessible to massive particles and photons and defined by $\mathrm{d}s^2 > 0$.

In flat (Minkowski) space, the geodesic reduces to a straight line between its start and end points. In curved spacetime the geodesic represents the shortest trajectory available between these points. An analogy can be drawn between the metric and the Christoffel symbols to gravitational potentials and gravitational forces respectively. This will become more apparent when (mass-)energy and momentum are introduced to this picture of geodesics in curved spacetime, but can already be seen by noting the construction of the Christoffel symbols as derivatives of the metric.

It is worth briefly discussing the term $g^{\mu\sigma}$ appearing in equation (A.7). This is defined in terms of what the metric must be contracted with in order to produce the invariant quantity, $\delta^{\mu}_{\nu} = g_{\nu\rho}g^{\mu\rho}$. Quantities with lowered indices, A_{μ}, and raised indices, A^{μ}, are called 'covariant' and 'contravariant' respectively. They are related by the metric such that $A_{\mu} = g_{\mu\nu}A^{\nu}$, which serves to effectively lower and raise indices. It can be seen that General Relativity deals in 'generally covariant' tensor equations, those which hold the same configuration of raised or lowered indices for all observers.

The curvature of spacetime is captured fully by the Riemann tensor. It describes the parallel transport of a four-vector around a small closed loop, along which there exist different freely falling frames due to the curvature of the manifold. This effectively characterises the tidal force experienced by a massive body traversing a geodesic, and so is an invaluable quantity in General Relativity. It can be constructed from the metric connection and its derivatives such that

$$R^{\mu}_{\alpha\beta\gamma} = \Gamma^{\mu}_{\alpha\gamma,\beta} - \Gamma^{\mu}_{\alpha\beta,\gamma} + \Gamma^{\mu}_{\rho\beta}\Gamma^{\rho}_{\gamma\alpha} + \Gamma^{\mu}_{\rho\gamma}\Gamma^{\rho}_{\beta\alpha}. \tag{A.8}$$

The Ricci tensor, $R_{\alpha\beta}$, follows by performing a contraction such that $\mu \longleftrightarrow \gamma$, and thence the curvature scalar, $R = g^{\alpha\beta}R_{\alpha\beta}$. These two quantities, along with the metric, are sufficient to determine the geometrical part of the Einstein field equation via the Einstein tensor,

$$G^{\mu\nu} \equiv R^{\mu\nu} - \frac{1}{2}g^{\mu\nu}R. \tag{A.9}$$

Having introduced this underlying formalism for General Relativity, it is now possible to relate geometrical perturbations to the spacetime manifold to the

large scale matter-energy distribution that sources these distortions and characterises the gravitational dynamics of the Universe. This will provide the energy-momentum part of the Einstein field equation. The relevant quantity to be introduced is the energy-momentum tensor, $T^{\mu\nu}$, of spacetime. This is defined through the flux, P^{μ}, of the μ component of the momentum 4-vector over a surface with fixed x^{ν} coordinate. This is related (via the invariant mass, m) to the corresponding 4-velocity, $U \equiv \gamma\,(c, \mathbf{v})$, in which γ is the Lorentz factor and $\mathbf{v}$ is the 3-velocity. On sufficiently large scales, such as those characterising cosmology, it is permissible to approximate the matter-energy distribution as a perfect fluid, in effect one that can be described entirely by macroscopic quantities like its total density, ρ, and (isotropic) pressure, p. The energy-momentum tensor then takes the form

$$T^{\mu\nu} = \left(\rho + \frac{p}{c^2}\right) U^{\mu}U^{\nu} - p\, g^{\mu\nu}. \tag{A.10}$$

Energy-momentum conservation is encapsulated by the statement

$$\nabla_{\mu} T^{\mu\nu} = 0, \tag{A.11}$$

where ∇_{μ} is the covariant derivative, which generalises the partial derivative, ∂, to curved manifolds.

The energy-momentum tensor is the source term for the Einstein tensor that describes the physical Universe. The two are related by

$$G^{\mu\nu} = -\frac{8\pi\, G}{c^4} T^{\mu\nu}. \tag{A.12}$$

When this is combined with the curvature described in equation (A.9) the Einstein field equations can be written in the form

$$R^{\mu\nu} - \frac{1}{2} g^{\mu\nu} R + \Lambda g^{\mu\nu} = -\frac{8\pi\, G}{c^4} T^{\mu\nu}. \tag{A.13}$$

The new addition characterises the contribution of a 'cosmological constant', Λ, that Einstein originally introduced to allow for a steady-state Universe that did not expand (or contract) as the field equations predicted. However, the field equations are free to accommodate additional linear terms, provided they are proportional to the metric, so the cosmological constant is self-consistent with General Relativity. Indeed, its resurrection in the modern era of cosmology with

a small positive value has been necessary to explain the observed acceleration of the Universe.

Einstein's field equations serve much the same purpose in curved spacetime as Poisson's equation. The weak field limit of the field equations, where Newtonian gravity is approximately true and the field varies slowly across (space-)time, does in fact resemble the Poisson equation:

$$\nabla^2\Phi \equiv \frac{c^2}{2}\nabla^2 g_{00} = \frac{4\pi G}{c^2}\left(\rho c^2 + 3p\right). \tag{A.14}$$

Note that the metric is perturbed from the Minkowski metric, $\eta_{\mu\nu}$ by a small quantity, $h_{\mu\nu}$ such that $|h_{\mu\nu}| \ll 1$. The effective source term differs by the dependence on the pressure term, which contributes to the 'active mass density' in General Relativity.

The formalism of General Relativity, which has been presented only through a brief introduction here, provides the framework through which the large-scale properties and dynamics of the Universe can be modelled in an entirely mathematically consistent way. Within this apparatus, modern cosmology begins to take shape.