Understanding neutrino mass constraints from galaxy clustering and CMB lensing

Aoife Boyle

München 2019

Understanding neutrino mass constraints from galaxy clustering and CMB lensing

Aoife Boyle

Dissertation
an der Fakultät für Physik
der Ludwig–Maximilians–Universität
München

vorgelegt von
Aoife Boyle
aus Cheshire, Vereinigtes Königreich

München, den 11. Juni 2019

Erstgutachter: Prof. Dr. Eiichiro Komatsu
Zweitgutachter: Prof. Dr. Jochen Weller
Tag der mündlichen Prüfung: 23. September 2019

Contents

List of Figures

List of Tables

Zusammenfassung

Diese Arbeit ist das Ergebnis eines langen Unterfangens, neue Perspektiven und Erkenntnisse für die Herausforderung der Extraktion starker und robuster Neutrinomassenbeschränkungen aus aktuellen und zukünftigen kosmologischen Messungen zu gewinnen. Obwohl die Kosmologie heute die stärksten Einschränkungen für die Neutrinomasse bietet und wahrscheinlich auch weiterhin wird, muss eine entscheidende Priorität sein sicherzustellen, dass diese Beschränkungen für absehbarer Zeit robust genug sind, um auch Wissenschaftler außerhalb der Kosmologie zu überzeugen.

Die überwiegende Mehrheit der aktuellen oder prognostizierten kosmologischen Neutrinomassenbeschränkungen in der Literatur gehen bei der Erlangung ihrer Ergebnisse vom konkordanten flachen ΛCDM-Modell aus. Trotz der vielen Erfolge ist das ΛCDM-Modell nicht ohne Herausforderungen, sowohl theoretischer als auch beobachtbarer Natur. Es ist wichtig, dass alle änderungen, die in Zukunft mit diesem Modell erlangt werden, unsere Neutrinomassenmessungen nicht wesentlich ändern. Das Ziel unserer Arbeit war es, die Bedeutung dieser kosmologischen Annahme zu verstehen, mögliche Entartungen zwischen der Neutrinomasse und anderen kosmologischen Parametern zu verstehen und zu untersuchen, wie viel wir wirklich über die Fähigkeit kosmologischer Messungen zur Messung der Neutrinomasse aussagen können. Wir verfolgen einen semi-analytischen Ansatz unter Verwendung des Fisher-Matrix-Formalismus.

In Kapitel 2 dekonstruieren wir gründlich die Beschränkungen der Neutrinomasse, die bei zukünftigen Galaxie-Rotverschiebungsdurchmusterungen zu erlangen sind. Galaxie-Clustering-Messungen reagieren in vielerlei Hinsicht empfindlich auf massive Neutrinos. Wir bieten isolierte Schranken aus Expansionsratenmessungen durch baryonakustische Schwingungen (BAOs) und den Alcock-Paczynski-Test. Wir zeigen auch die einschränkende Wirkung von Redshift-Raumverzerrungen (RSD). Wir entwickeln eine neue Methode zur Isolierung der charakteristischen skalenabhängigen Unterdrückung im Leistungsspektrum als Sonde der Neutrinomasse und zeigen, dass die damit verbundenen Schranken sowohl kosmologieunabhängig als auch relativ stark sind. Wir zeigen, dass alle anderen Einzelproben der Neutrinomasse und der kombinierten Randbedingungen stark kosmologieabhängig sind. Obwohl bereits bekannt war, dass es in kosmologischen Messungen eine Entartung zwischen der Neutrinomasse und der optischen Tiefe zur Reionisation (τ) gibt, erklären wir erstmals deutlich den Ursprung dieser Entartung.

In Kapitel 3 erweitern wir unsere Studie um prognostizierte Informationen aus zukünftigen CMB-Experimenten (Cosmic Microwave Background, kosmis-

cher Mikrowellenhintergrund), einschließlich verbesserter E-Moden-Polarisations- und CMB-Gravitationlinsenmessungen. Wir zeigen, dass das Hinzufügen optimaler primärer Anisotropieinformationen zu unseren Freiströmungsbeschränkungen die Neutrinomassenbeschränkungen signifikant verstärkt und gleichzeitig kosmologieunabhängig hält. Wir zeigen auch, dass Freiströmungssignale in den Leistungsspektren der CMB-Gravitationlinsenmessungen und der Galaxie-CMB-Gravitationlinsenmessungen auch isoliert werden können, um die Einschränkungen der Freiströmung weiter zu verbessern. Schließlich zeigen wir, dass kombinierte und reine BAO-Neutrinomassenbeschränkungen auch in Kombination mit zukünftigen modernsten CMB-Messungen stark kosmologieabhängig bleiben.

In den Kapiteln 2 und 3 arbeiten wir ausschließlich mit linearen Leistungsspektren, die auf einigen der von uns betrachteten Skalen wohl ungültig werden könnten. In Kapitel 4 erweitern wir unsere Analyse sowohl für Galaxienhaufen als auch für CMB-Linsen, um die Leistungsspektren der nächst höheren Ordnung zu nutzen. Die Einbeziehung von Beiträgen der zweitwichtigsten Ordnung verbessert die Freiströmungssignatur im Leistungsspektrum. Durch die Vielzahl zusätzlicher Störparameter im Leistungsspektrum sind die schlussendlichen Schranken jedoch noch etwas schwächer als im linearen Fall. Aber die meisten qualitativen Ergebnisse der vorangegangenen Kapitel bleiben weiterhin gültig, und die Freiströmungsschranken bleiben der einzige kosmologieunabhängige Test der Neutrino-Masse.

This thesis is the result of a long endeavour to bring new perspectives and insights to the challenge of extracting strong and robust neutrino mass constraints from current and future cosmological measurements. Although cosmology provides the strongest constraints on the neutrino mass today and will likely continue to into the foreseeable future, ensuring that these constraints are robust enough to also convince scientists outside of the cosmology community must be a crucial priority.

The vast majority of current or forecasted cosmological neutrino mass constraints in the literature assume the concordance flat ΛCDM model when extracting their results. Despite its many successes, the ΛCDM model is not without challenges, both theoretical and observational. It is important that any changes that are made to this model in future do not significantly alter our neutrino mass measurements. The goal of our work has been to understand the significance of this cosmological assumption, to understand possible degeneracies between the neutrino mass and other cosmological parameters, and to examine how much we can really claim about the capacity of cosmological measurements to measure the neutrino mass. We take a semi-analytical approach, using the Fisher matrix formalism.

In Chapter 2, we thoroughly deconstruct the constraints on the neutrino mass available from future galaxy redshift surveys. Galaxy clustering measurements are sensitive to massive neutrinos in many ways. We provide isolated constraints from expansion rate measurements provided by baryon-acoustic oscillations (BAOs) and the Alcock-Paczyński test. We also demonstrate the constraining power of redshift-space distortions (RSD). We develop a new method of isolating the characteristic scale-dependent suppression in the power spectrum as a probe of the neutrino mass, and show that the constraints it provides are both cosmology-independent and relatively powerful. We show that all of the other individual probes of the neutrino mass and the combined constraints are strongly cosmology-dependent. In addition, although it was already known that there is a degeneracy between the neutrino mass and the optical depth to reionisation (τ) in cosmological measurements, we explain clearly the origin of this degeneracy clearly for the first time.

In Chapter 3, we extend our study to include forecasted information from future cosmic microwave background (CMB) experiments, including improved E-mode polarisation and CMB lensing measurements. We show that adding optimal primary anisotropy information to our free-streaming constraints strengthens them significantly while still keeping them cosmology-independent. We also show that free-streaming signals in the CMB lensing and galaxy-CMB lensing power spectra can also be isolated to improve free-streaming constraints further. Finally, we show that combined and BAO-only neutrino mass constraints remain strongly cosmology-dependent even when combined with future state-of-the-art CMB measurements.

In Chapters 2 and 3, we work exclusively with linear power spectra, which could arguably become invalid on some of the scales we consider. In Chapter 4, we extend our analysis for both galaxy clustering and CMB lensing to make use of next-to-leading-order power spectra. Including next-to-leading-order contributions enhances the free-streaming signature in the power spectrum. However, the large number of additional nuisance parameters in the galaxy power spectrum mean that the final constraints are still somewhat weaker than in the linear case. But most of the qualitative results of the previous chap-

ters hold true, and the free-streaming constraints remain the only cosmology-independent probe of the neutrino mass.

Chapter 1

Introduction

The 2015 Nobel prize in physics was shared by Takaaki Kajita and Arthur B. McDonald *"for the discovery of neutrino oscillations, which shows that neutrinos have mass"*[1]. This represented the first experimentally proven clear deviation from the Standard Model of particle physics, which originally predicted neutrinos to be massless particles. Despite incredible efforts to measure neutrino masses from both the particle physics and cosmological angles, meaningful constraints on the masses of the three known neutrino mass states have yet to be achieved.

There are three major avenues of exploration underway in the quest to measure the neutrino mass scale more precisely, which are helpfully complementary. The first of these is neutrino oscillation experiments, which can now strongly constrain the squared mass differences between two pairs of mass values. Solar oscillation experiments constrain Δm_{21}^2 and atmospheric oscillation experiments Δm_{31}^2. This leaves us with two possible neutrino mass hierarchies depending on the sign of Δm_{31} - the so-called normal and inverted hierarchies (see Figure 1.1). Today, the 1σ constraints stand at $\Delta m_{21}^2 = 7.55 \pm 0.2 \ 10^{-5} \text{eV}^2$ and $\Delta m_{31}^2 = 2.5 \pm 0.03 \ 10^{-3} \text{eV}^2$ (the latter assuming normal ordering), and the normal hierarchy is currently favoured at $3.4\sigma^2$. This provides enough information to derive a lower bound on the total neutrino mass, $M_\nu = \sum_{i=1}^3 m_{\nu,i}$, of just under 0.06 eV for normal ordering and a little over 0.1 eV for inverted ordering. Achieving an upper bound on M_ν of below 0.1 eV could therefore exclude the inverted neutrino mass hierarchy.

A second set of laboratory-based attempts to measure the neutrino mass come from β decay and neutrinoless double-β decay experiments. In theory, precise measurements of the endpoint of the β decay electron spectrum should provide the cleanest measurements of the neutrino mass. With sufficient resolution, one could see the impact of the individual mass states on the electron energy spectrum. Despite the robustness of this approach, the projected uncertainties are relatively large. The KATRIN experiment aims for a 90% C.L. constraint of 0.2 eV on the electron neutrino mass [3]. Neutrinoless double-β decay experiments could soon reach a similar sensitivity [4], but are limited by the fact that the constraints only apply if neutrinos are Majorana particles.

---

[1]https://www.nobelprize.org/prizes/physics/2015/summary/
[2]https://globalfit.astroparticles.es/

Cosmological measurements complete the set of competitive neutrino mass measurements available currently or in the near future. Assuming $M_\nu = 0.06$ eV, massive neutrinos should constitute about 0.15% of the energy density of the universe today. This small component has several subtle but theoretically measurable effects on cosmological observables, which will be discussed in detail in both this and later chapters. In complement to oscillation experiments, cosmological large-scale structure (LSS) and cosmic microwave background (CMB) lensing experiments are primarily sensitive to the total neutrino mass M_ν. The strongest upper bounds on M_ν come from cosmology, with Planck offering a 2σ constraint of $M_\nu < 0.12$ eV last year, by combining CMB temperature and polarisation measurements with CMB lensing and baryon acoustic oscillation (BAO) measurements from galaxy redshift surveys [5], in the framework of the standard ΛCDM model with the neutrino mass as a single parameter extension. This result is intriguingly close to the inverted hierarchy lower limit. But despite these impressive results, cosmological measurements of the neutrino mass suffer from significant issues that laboratory experiments do not because of the complexity of the system being examined. The great success of modern cosmology is the concordance ΛCDM model, which has proven remarkably capable of explaining the observed properties of our universe using only six basic parameters. However, it is not without its problems, and there are a great number of theoretical extensions that could be added (see [6] for a review). Cosmological neutrino mass constraints are usually extracted using global fits to combinations of datasets, almost always varying the six base ΛCDM parameters and M_ν, and with many other possible parameters being assumed and kept fixed. Because the effects of the neutrino mass on cosmological observables are so subtle, this raises the possibility of very large relative errors because of possible degeneracies with other cosmological parameters. One would therefore expect neutrino mass constraints from cosmology to depend heavily on the cosmological model assumed. This is exactly the issue addressed by this thesis.

The rest of this chapter is intended to give the reader sufficient background to understand the results derived in Chapters 2-4. We begin by briefly reviewing the history of cosmological neutrinos (Section 1.1) and the free-streaming behaviour on small scales that distinguishes their effects on the evolution of cosmic structure from other forms of matter (Section 1.2). We then discuss how this behaviour impacts the evolution of the matter power spectrum (Section 1.3). In Section 1.4, we discuss the physical origins of important elements of the galaxy power spectrum (the observable in galaxy redshift surveys) that can be used to constrain cosmological parameters, including baryon acoustic oscillations (BAOs), the Alcock-Paczyński test and redshift-space distortions (RSD). In Section 1.5, we discuss how the lensing of the CMB by large-scale structure is sensitive to the neutrino mass. In Section 1.6, we briefly review how one can extend galaxy redshift survey and CMB lensing measurements into the regime where linear approximations of structure formation fail. Finally, in Section 1.7, we provide more background information on how we forecast neutrino mass constrains in this thesis in practise. We give an introduction to the Fisher matrix formalism, motivate the importance of exploring the parameter space beyond ΛCDM when extracting constraints, and explain why we focus on particular cosmological extensions and not others.
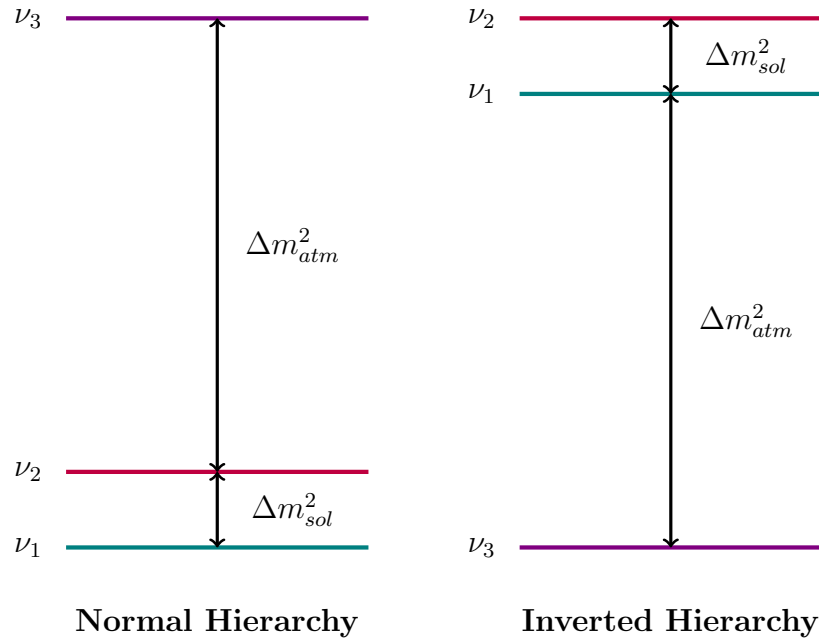
Figure 1.1: A graphical representation of the two possible neutrino mass hierarchies.

The results of this thesis are contained in Chapters 2-4. Chapter 2 focuses on disentangling the constraints on M_ν available from galaxy redshift surveys, using minimal CMB information. Chapter 3 extends this analysis to include full CMB anisotropy information and focuses on the contribution of CMB lensing. Chapter 4 extends the results of the two previous chapters into the mildly non-linear regime by making use of the next-to-leading-order or one-loop power spectra.

1.1 Cosmological Neutrinos

The cosmic neutrino background (CNB) was generated in the very early universe at high temperatures and held in equilibrium with the primordial plasma by weak interactions. Originally ultra-relativistic, the neutrinos gradually become non-relativistic over the course of the history of the universe. Their contribution as radiation at early times and as matter at late times is responsible for their distinctive effects on the evolution of the universe.

Here we use some simple approximations to describe the basic properties of this neutrino background. Originally coupled to other particles by weak interactions, the neutrinos hold a Fermi-Dirac momentum distribution

$$f(p) = \frac{1}{\exp[(p - \mu)/T] + 1}, \tag{1.1}$$

with p the momentum, T the temperature and μ the chemical potential (which should be negligible).

As the universe expands and cools, neutrinos decouple from the rest of the plasma. However, the distribution function in Equation 1.1 is maintained because neutrinos are non-interacting and the temperature and neutrino momentum redshift in the same way as the universe expands. This fact is very useful, as it allows us to derive many interesting parameters related to the neutrino distribution as a function of time.

The temperature of neutrino decoupling can be estimated by comparing how the weak interaction rate and the expansion rate evolve as a function of temperature, and turns out to be at around 2.5 MeV. Directly after decoupling, neutrinos and photons maintain the same temperature. However, electron pair production rates drop quickly once the temperature falls below the electron mass. If instantaneous decoupling of neutrinos is assumed, the excess energy is transferred to photons, but not to the decoupled neutrinos. One can use entropy conservation arguments to relate the temperatures of the cosmic microwave background (CMB) and the CNB

$$\frac{T_\nu}{T_\gamma} = \left(\frac{4}{11}\right)^{1/3}. \tag{1.2}$$

This gives a CNB temperature today of approximately 1.95 K.

One can also derive the temperature at which a particular neutrino mass state becomes non-relativistic by calculating the mean neutrino energy from Equation 1.1 as a function of temperature, to find that

$$1 + z_{nr,i} = 1890 \left(\frac{m_{\nu,i}}{1 \text{ eV}}\right). \tag{1.3}$$

We can therefore see that at least two of the neutrino mass states must be non-relativistic today.

Finally, another important parameter related to neutrino properties is N_{eff}, the effective number of relativistic species. It quantifies the contribution of neutrinos to the radiation density in the early universe.

$$\rho_R = \rho_\gamma \left(1 + \frac{7}{8} \left(\frac{4}{11}\right)^{4/3} N_{\text{eff}}\right) \tag{1.4}$$

N_{eff} is expected to have a value of 3.046 for three neutrino species, with the small deviation from 3 accounting for some non-thermal distortion of the spectra caused by electron-positron annihilations. N_{eff} can be constrained using Big Bang nucleosynthesis because ρ_γ determines the expansion rate during that epoch. N_{eff} is currently constrained to 2.99±0.34 at 95% C.L. with combinations of Planck and BAO data [5]. Any significant deviation from the expected value would imply the existence of additional unknown relativistic species in the early universe.

1.2 Neutrino Free-Streaming

We can infer from the current neutrino mass bounds that the three neutrino mass states were relativistic at decoupling (see Equation 1.3 and the bounds given at the beginning of this chapter). Relativistic massive neutrinos are classified as *hot* dark matter. Their behaviour when relativistic is distinguished from that of cold dark matter and baryons by the fact that they free-stream rather than clustering gravitationally. The transition of massive neutrinos from free-streaming particles to gravitationally responsive matter imprints a signature on large-scale structure observables that distinguishes them from other forms of matter, as we will see in the next section.

The scale above which neutrinos cluster in gravitational potential wells is called the free-streaming scale

$$k_{FS} = \left(\frac{4\pi G \bar{\rho} a^2}{v_{th}^2} \right)^{1/2}, \tag{1.5}$$

sometimes understood more intuitively as a free-streaming length

$$\lambda_{FS} = 2\pi \frac{a}{k_{FS}} = 2\pi \sqrt{\frac{2}{3}} \frac{v_{th}}{H}. \tag{1.6}$$

This is analogous to the Jeans length, the scale below which pressure prevents the gravitational collapse of a fluid. At early times, when neutrinos are relativistic, their free-streaming length is equal to the Hubble radius. This changes after the non-relativistic transition, when v_{th} falls as [7]

$$v_{th} = \frac{\langle p \rangle}{m} = \frac{3.15 T_\nu}{m} = \frac{3.15 T_\nu^0}{m} \left(\frac{a_0}{a} \right) \approx 158(1+z) \left(\frac{1 \text{ eV}}{m} \right) \text{ km s}^{-1}. \tag{1.7}$$

Note that $\langle p \rangle = 3.15 T_\nu$ holds true for a relativistic Fermi-Dirac distribution with negligible chemical potential, and we used the relation in Equation 1.2 for $T_\nu^0 = 1.95$ K. The free-streaming length therefore evolves after the non-relativistic transition as

$$\lambda_{FS} = 8.1(1+z) \left(\frac{H_0}{H(z)} \right) \left(\frac{1 \text{ eV}}{m} \right) h^{-1} \text{ Mpc}. \tag{1.8}$$

The free-streaming length grows as $(aH)^{-1}$ after transition, and so is at its minimum value at transition. Importantly, we see that the free-streaming scale is a function of the neutrino mass.

1.3 The Matter Power Spectrum

The matter power spectrum underlies all of the most important observables used to constrain the mass of neutrinos within the scope of this thesis. In the large-scale/linear regime, the galaxy power spectrum measured by galaxy redshift surveys (Chapter 2) is a simple

Figure 1.2: The relative change in the matter power spectrum $P_m(k)$ (dashed lines) and the baryon-CDM power spectrum $P_{bc}(k)$ (solid lines) with the addition of a given neutrino mass, holding Ω_m fixed. The lower mass bound of 0.06 eV uses the approximation of one massive neutrino and two massless neutrinos, while for the upper mass bound of 0.27 eV three degenerate neutrino mass states are assumed. One can see that the neutrinos contribute like additional cold dark matter on large scales, but suppress the power spectrum on small scales, and that they also suppress the growth of baryon-CDM perturbations.

function of the matter power spectrum modified by bias and a redshift-space distortion term and with a shot noise term. The CMB lensing power spectrum (Chapter 3) is derived from an integral over redshift of the matter power spectrum. Understanding how neutrinos affect the evolution of the matter power spectrum is therefore crucial to understanding how the neutrino mass can be constrained with cosmological surveys.

The matter power spectrum is defined as the Fourier transform of the two-point correlation function of matter perturbations

$$P(z, k) = \langle |\delta_m(z, k)|^2 \rangle, \tag{1.9}$$

where

$$\delta_m = \delta\rho_m / \bar{\rho}_m = \frac{\sum_i \bar{\rho}_i \delta_i}{\sum_i \bar{\rho}_i}. \tag{1.10}$$

and cold dark matter (CDM), baryons and massive neutrinos are the species of interest.

Assuming a Gaussian field, the power spectrum contains all of the information required to describe the field. The standard cosmological model describes the evolution of the universe in terms of the gravitational amplification of initially very small primordial fluctuations over time. These fluctuations are believed to be small quantum fluctuations enhanced to macroscopic scales by inflation. The primordial power spectrum is assumed to be almost scale invariant

$$P(k) = A_s k^{n_s}. \tag{1.11}$$

where n_s is close to 1. Radiation provides the dominant energy component of the early universe, and fluctuations that cross the horizon during radiation domination are prevented from collapsing under gravity by radiation pressure. Once the time of matter-radiation equality is reached, perturbations inside the horizon begin to grow as a result of gravitational collapse, and the power spectrum shape then deviates from its power law shape on small scales, with the matter-radiation equality scale being imprinted on the shape of the matter power spectrum. For cold dark matter and baryons (after the baryon drag epoch), the continuity and Euler equations can be combined to provide an evolution equation

$$\delta_m'' + \frac{a'}{a}\delta_m' - \frac{3}{2}\frac{a'^2}{a}\Omega_m(z)\delta_m = 0. \tag{1.12}$$

This equation corresponds to Equation C.8 in Appendix C, in which we provide an overview of the governing principles of linear and non-linear perturbation theory, and is derived in more detail in that appendix. On large scales, neutrino perturbations evolve in the same way, so the power spectrum is the same as it would be without massive neutrinos but with the same total Ω_m. But on small scales, neutrinos free-stream and do not cluster, so matter perturbations are suppressed automatically by a factor of $1 - f_\nu$, where $f_\nu = \Omega_\nu/\Omega_m$. For a neutrino mass of 0.06 eV, this corresponds to a suppression in the power spectrum of around 1% on small scales.

However, the strongest impact of the effect of massive neutrinos on the power spectrum comes from their effect on the growth of cold dark matter and baryon perturbations. If we consider Equation 1.12 for the growth of baryon and cold dark matter perturbations, the change in Ω_{bc} as a result of allowing for the extra neutrino component (while keeping Ω_m fixed) will suppress the growth of δ_{bc} on small scales. Dark matter and baryon perturbations normally grow proportionally to the scale factor a in the matter-dominated era. With massive neutrinos, they grow proportionally to $a^{1-\frac{3}{5}f_\nu}$ [8].

Accounting for this effect, the power spectrum in the small-scale limit is actually suppressed in total by a factor of approximately $(1 - 8f_\nu)$, and the baryon-CDM power spectrum by a factor of $(1 - 6f_\nu)$, with the addition of massive neutrinos. This result is derived analytically using linear perturbation theory by [7], and is a rough approximation. The precise effect of massive neutrinos on the power spectrum is best measured numerically using Boltzmann codes such as `CLASS` [9], as is done in Figure 1.2.

Figure 1.3: Source: [10]. The galaxy power spectra data and best-fit models from for the CMASS galaxy sample in the SDSS data release 9. The subplot shows the isolated BAO wiggles.

1.4   The Galaxy Power Spectrum

Naturally, the matter power spectrum cannot be directly observed. Galaxy redshift surveys instead measure the galaxy power spectrum (see Figure 1.3). Galaxies are not perfect tracers of the underlying matter field, and bias relates the galaxy over-density δ_g to the matter over-density δ_{bc}. Note that the galaxy power spectrum is a function of the baryon and cold dark matter power spectra only, because neutrino perturbations do not form galaxy haloes. On linear scales, the bias takes the form of a simple linear relation, $b = \delta_g / \delta_{bc}$. This relates the galaxy power spectrum to the matter power spectrum as

$$P_g(k) = b^2 P_{bc}(k) + n_g^{-1}, \tag{1.13}$$

where the final term accounts for shot noise because the galaxies are considered as discrete point-like objects. A real galaxy power spectrum from the Sloan Digital Sky Survey (SDSS) is shown in Figure 1.3. Equation 1.13 neglects the angular dependence of the measured galaxy power spectrum that arises due to redshift-space distortions and the Alcock-Paczyński effect. These will be dealt with separately in the following sections.

(a)

(b)

(c)
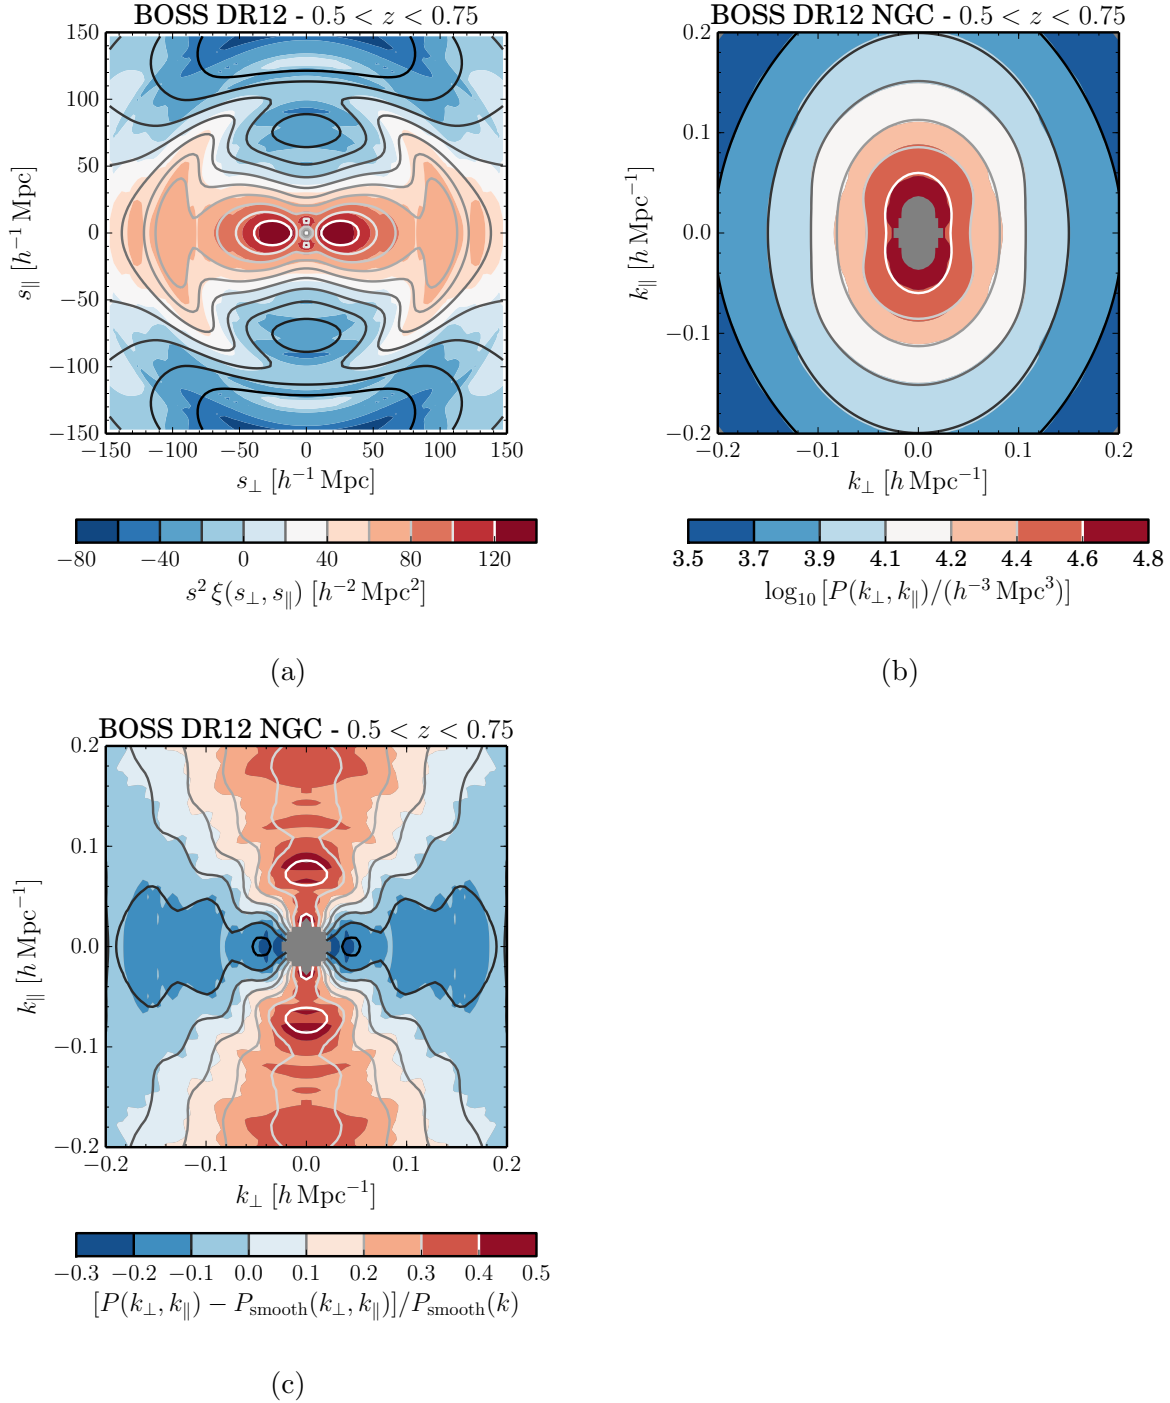
Figure 1.4: Source: [11]. Panels (a) and (b) show the unreconstructed correlation function and power spectrum from an SDSS-III BOSS DR12 galaxy sample in two dimensions (parallel and perpendicular to the line of sight). The clear anisotropy in both panels arises as a result of redshift-space distortions and the Alcock-Paczyński effect. Panel (c) demonstrates the presence of BAO wiggles by showing the power spectrum divided by the best-fit smooth power spectrum.

1.4.1   Baryon Acoustic Oscillations

Sound waves in the primordial baryon-photon plasma generated by opposing radiation pressure and gravity forces were frozen into the baryon distribution after decoupling. The radius reached by these waves by the time of decoupling is usually referred to as the sound horizon. This scale is now imprinted in the matter and galaxy correlation functions as a preferred clustering scale. The single peak corresponding to this scale in the correlation function becomes the characteristic wiggles seen in the power spectrum when Fourier transformed (see the subplot in Figure 1.3).

The BAO scale is much larger than the scale of virialised structures, making it relatively robust against non-linear effects. The angular size of the sound horizon is well constrained from CMB data, so the apparent size of the sound horizon as a function of redshift in large-scale structure can be used to infer cosmological distances. The BAO scale can therefore be used as a standard ruler to measure the expansion rate history of the universe. Since it was first measured [12, 13], BAOs have become one of the primary ways of constraining modern cosmology. A thorough review of all aspects of this topic is provided in [14]. Figure 1.4 (c) shows the BAO signature in real data in two dimensions.

The use of BAOs to map out the cosmological expansion history at low redshifts has been crucial to deriving the bounds on cosmological parameters we have today. From the Friedmann equations we have

$$\frac{H^2(z)}{H_0^2} = \Omega_R a^{-4} + \Omega_m a^{-3} + \Omega_k a^{-2} + \Omega_\Lambda. \tag{1.14}$$

We can therefore see that the expansion rate scales differently with different energy components of the universe. Precise measurements of the BAO scale as a function of redshift can therefore be fitted to extract constraints on these components.

1.4.2   Redshift-Space Distortions

When measuring the position of a galaxy, its position along the line of sight is determined from its redshift. However, the redshift that is measured is a function of two components - redshifting due to the expansion of the Universe (called the Hubble flow) and the unique peculiar velocity of a particular object along the line of sight.

Although the magnitude of the peculiar velocity component is usually small compared to that of the Hubble flow, it can have a significant effect on the measured power spectrum along the line of sight, introducing anisotropy into constructed galaxy clustering maps.

There are two regimes of interest. On large scales, objects will coherently fall towards large overdensities. Galaxies on the far side of such an overdensity will therefore appear closer to us, and those on the near side further away when their redshifts are measured. In redshift-space, these galaxies will then appear to be more clustered than they actually are. This is called the Kaiser effect [15]. On small scales, virialisation will induce significant random velocities. This will have the opposite effect, reducing the apparent strength of

clustering along the line of sight. This is called the Finger-of-God effect, so called because the effect manifests in clustering maps as long 'fingers' of galaxies pointing in our direction.

Although redshift-space distortions (RSD) make it more difficult to measure line-of-sight galaxy positions, they also add independent information that can also be used to constrain cosmology, by providing us with direct information on the velocity field and therefore on the growth of structure.

Accounting for the Kaiser effect, the power spectrum in Equation 1.13 is modified to

$$P_g(k) = (b + f\mu^2)^2 P_m(k) + n_g^{-1}. \tag{1.15}$$

The Finger-of-God effect is often approximated using an exponential function

$$P_{g,\text{FoG}}(k) = P_g(k) \exp\left(-k^2 \mu^2 f^2 \sigma_v^2\right), \tag{1.16}$$

where σ_v^2 is the velocity dispersion. When working in the linear regime in Chapters 2 and 3, we account only for the Kaiser effect.

An excellent (unpublished, but publicly available) review of the topic of redshift-space distortions has been provided by Shun Saito in the form of lecture notes[3].

1.4.3 Alcock-Paczyński Test

Another source of anisotropy in the redshift-space power spectrum arises from the Alcock-Paczyński (AP) effect [16]. In real galaxy surveys, three-dimensional galaxy positions are inferred by measuring angular positions on the sky and redshifts. In order to convert these measurements into three-dimensional maps in real space, values of the expansion rate $H(z)$ and angular diameter distance $D_A(z)$ must be assumed. If the values assumed are not in proportion to each other, the resulting distribution will be distorted, violating the requirement for isotropy given by the cosmological principle. The AP effect therefore provides an additional method of constraining the expansion rate history beyond standard ruler methods like BAOs (Section 1.4.1).

1.5 CMB Lensing

Large scale structure between us and the last scattering surface distorts the primary temperature and polarisation anisotropies in the CMB through gravitational lensing. This results in many significant effects, including a smearing of the acoustic peaks, the production of B-mode polarisation and the introduction of non-Gaussianity. A thorough review of the theory of the topic is provided in [17].

The CMB lensing power spectrum therefore provides a useful additional probe of large scale structure. It is a function of the projected matter distribution between us and the last scattering surface. This is of interest because it means that the relative suppression

---

[3]https://wwwmpa.mpa-garching.mpg.de/~komatsu/lecturenotes/Shun_Saito_on_RSD.pdf

Figure 1.5: The relative change in the CMB lensing and galaxy-CMB lensing power spectra $C_l^{\kappa\kappa}$ and $C_l^{g\kappa}$ (the latter at $z = 1.35$) with the addition of a given neutrino mass, holding Ω_m fixed. The masses are distributed between the three neutrino species as in Figure 1.2. Once again, the neutrinos suppress the power on small scales.

in matter clustering on small scales discussed in the previous section is also imprinted on the CMB lensing power spectrum (see Figure 1.5).

The CMB lensing potential is defined as

$$\phi(\hat{n}) = -2 \int_0^{\chi^\star} d\chi \left(\frac{d_A(\chi, \chi^\star)}{d_A(\chi^\star)d_A(\chi)} \right) \Psi(\boldsymbol{x}, \chi), \tag{1.17}$$

where Ψ is the gravitational potential, $\chi^\star$ is the the comoving distance last scattering, $\hat{n}$ is the direction of observation and d_A is the comoving angular distance, equal to the comoving distance in a flat universe, and otherwise defined in Equations 3.5 and 3.6 of Chapter 3.

To derive the angular power spectrum (see [18] for more detail), we take the two-point function of ϕ in Fourier space and expand it in spherical harmonics. This gives, in terms of the gravitational potential power spectrum $P_\Psi(k, z, z')$

$$C_l^{\phi\phi} = \frac{8}{\pi} \int k^2 dk \int_0^{\chi^\star} d\chi \int_0^{\chi^\star} d\chi' P_\Psi(k, d_A(\chi), d_A(\chi')) j_l(k, d_A(\chi)) j_l(k, d_A(\chi'))$$
$$\left(\frac{d_A(\chi, \chi^\star)}{d_A(\chi^\star)d_A(\chi)} \right) \left(\frac{d_A(\chi', \chi^\star)}{d_A(\chi^\star)d_A(\chi')} \right). \tag{1.18}$$

P_Ψ can be easily related to P_m using the Poisson equation

$$P_\Psi(k, \chi) = \frac{9}{4} \frac{H_0^4 \Omega_m^2 (1+z)^2}{k^4} P_m(k, z). \tag{1.19}$$

At high l, P_Ψ varies slowly compared to the Bessel functions and we can use the Limber approximation

$$\int k^2 dk j_l(k d_A(\chi)) j_l(k d_A(\chi')) = (\pi/d_A(\chi)^2) \delta(\chi - \chi'). \tag{1.20}$$

We can also replace k with $l/d_A(\chi)$ on small scales and change the integration variable to redshift to get

$$C_l^{\phi\phi} = l^{-4} \int_0^{z\star} dz W(z) P_m \left(\frac{l}{d_A(z)}, z \right). \tag{1.21}$$

$W(z)$ is usually called the lensing kernel

$$W(z) = 9 \frac{c^{-3} H_0^4}{H(z)} d_A^2(z) \left(\frac{d_A(z, z\star)}{d_A(z\star) d_A(z)} \right)^2 \Omega_m^2 (1+z)^2. \tag{1.22}$$

In this thesis, we work in terms of the lensing convergence (which is the observable) instead of the lensing potential when generating our results. The convergence is related to the lensing potential by

$$\kappa = \frac{1}{2} \nabla^2 \phi. \tag{1.23}$$

The convergence power spectrum is then given by

$$C_l^{\kappa\kappa} = \frac{9}{4} \int_0^{z\star} \frac{dz}{H(z)} (1+z)^2 H_0^4 \Omega_m^2 d_A^2(z) \left(\frac{d_A(z, z\star)}{d_A(z\star) d_A(z)} \right)^2 P_m(k). \tag{1.24}$$

This can be related to the form given in Chapter 3 using $\Omega_m = \rho_m/\rho_{\rm crit} = 8\pi G \rho_m/3H_0^2$. CMB lensing and two-dimensional galaxy positions can also be cross-correlated to give the galaxy-CMB lensing power spectrum, which also shows a relative suppression on small scales induced by the presence of massive neutrinos (see Figure 1.5).

1.6 The Next-to-Leading-Order Power Spectrum

The linear matter power spectrum is derived by assuming that fluctuations on different scales evolve independently of each other, through the linearisation of the governing equations (the Poisson, continuity and Euler equations). This works well on large scales and at early times (when $\delta_m \ll 1$). However, later gravitational dynamics leads to the coupling of fluctuations on different scales in Fourier space, and this approach no longer suffices. Moving beyond the linear regime requires higher order perturbation theory, taking as a first

step the next-to-leading-order or one-loop power spectrum. The derivation is summarised in Appendix C.

In Chapter 4, we extend our results for the galaxy power spectrum in Chapter 2 into the mildly non-linear regime, following the formalism for the calculation of the NLO galaxy power spectrum provided by [19]. The authors derive a complete description for the redshift-space NLO galaxy power spectrum, including both non-linear bias terms and selection effects.

We will briefly discuss the bias parameters required for the calculation of the next-to-leading-order galaxy power spectrum in Chapter 4. On small scales, the simple bias relation in Equation 1.13 is no longer sufficient (see [20] for an up-to-date review). But on scales where perturbation theory still holds, the growth of structure is still dominated by gravity, and the relationship of the galaxy density and the underlying matter density field can be encapsulated at a given order in a small number of bias parameters related to the matter density and tidal fields (and their derivatives)

$$\delta_g(\boldsymbol{x}, \tau) = \sum_O b_O(\tau) O(\boldsymbol{x}, \tau). \tag{1.25}$$

In the simplest case, one can treat the galaxy density field as an expansion in powers of the matter density field with associated bias parameters. This is sometimes called the local bias expansion.

$$\delta_g(\boldsymbol{x}, \tau) = \sum_n b_n(\tau) \delta_{bc}^n(\boldsymbol{x}, \tau) \tag{1.26}$$

Taking this only to $n = 1$ gives $\delta_g = b_1 \delta_{bc}$, as used in the linear galaxy power spectrum (Section 1.4). We include both b_1 and b_2 in the non-linear galaxy power spectrum in Chapter 4. Bias parameters can also be related to the tidal field

$$K_{ij} = D_{ij}\delta_m = \left(\frac{\partial_i \partial_j}{\nabla^2} - \frac{1}{3}\delta_{ij}\right)\delta_m. \tag{1.27}$$

and b_{K^2} in Chapter 4 is paired with the operator $K^2 = K_{ij}K^{ij}$. There is also a bias parameter associated with another tidal operator, O_{td}

$$O_{td} = \frac{8}{21}K_{ij}D^{ij}\left(\delta_m^2 - \frac{3}{2}K^2\right). \tag{1.28}$$

The set of bias parameters of interest in Chapter 4 is completed by those associated with a higher derivative operator of the density field $\nabla^2\delta$ and the velocity field $\nabla^2 v$, with the latter being required to account for the conversion from real space to redshift space.

The final galaxy power spectrum in Fourier space also contains three stochastic terms. The first of these, denoted $P_\epsilon^{\{0\}}$, corresponds to that in Equation 1.13 and represents shot noise due to the discreteness of the field. Another, $P_{\epsilon\varepsilon_\eta}^{\{2\}}$ corresponds to the Finger-of-God effect (Section 1.4.2).

1.7 Forecasting Constraints on Cosmological Parameters

1.7.1 The Fisher Matrix

Throughout this thesis, we use the Fisher matrix formalism to quantify our forecasted constraints. The Fisher matrix is a simple tool used to determine how much constraining information an observable x can contain about the parameters of an underlying model, θ. It is defined as

$$F_{\alpha\beta} = \left\langle \frac{\partial^2[-\ln L]}{\partial\theta_\alpha \partial\theta_\beta} \right\rangle. \tag{1.29}$$

where L is the likelihood. Inverting the Fisher matrix provides a covariance matrix, with the square roots of the diagonal elements giving the forecasted errors on the individual model parameters. The Cramér-Rao bound requires that the Fisher matrix bounds are the best achievable constraints.

Given the covariance matrix C of an observable P (for example, the galaxy power spectrum or a CMB power spectrum), a Fisher matrix can be easily constructed for the parameters of a cosmological model using

$$F_{\alpha\beta} = \frac{\partial P}{\partial\theta_\alpha} C^{-1} \frac{\partial P}{\partial\theta_\beta}. \tag{1.30}$$

The Fisher matrix has some limitations. It requires the assumption of a Gaussian likelihood (usually a reasonable assumption when CMB and low-redshift large scale structure information are combined, as the parameters are relatively well constrained). It also does not account for systematic errors. However, it is ideal for the purposes of this thesis for a number of reasons. Our goal is not to develop the most realistic forecasts, but to compare the constraining power of different sources of cosmological information, and the degeneracies inherent therein. The Fisher matrix also has some advantages over Markov chain Monte Carlo (MCMC) forecasts, which are the other popular choice in the literature. Besides its obvious efficiency, the Fisher matrix is ideal for deconstructing the constraining power of cosmological data, because the input derivatives can be modified to vary only specific variables that we want to contribute to constraints.

1.7.2 Extensions beyond flat ΛCDM

Given that the consensus ΛCDM model has so successfully connected cosmological observations from different epochs in the history of the universe, and the fact that many of its predictions are being further validated by ever-improving measurements, one might wonder why extensions to ΛCDM need to be considered at all. Besides philosophical issues with the fact that the two largest components of the model - cold dark matter and dark energy - are barely understood, there are increasingly concerning tensions in the measurements of

certain parameters between datasets, most importantly, in the present-day matter fluctuation amplitude σ_8 and the Hubble constant. In the latter case, tensions between Planck predictions and local measurements of the Hubble constant have now reached a significance of 3.6σ.

It would be impossible for us to investigate all possible extensions to the model in doing our analysis. A flat ΛCDM cosmology is generally defined by specifying six free parameters, though one has a degree of freedom in choosing some parameters because of their interdependence. In our work we follow the example of the Planck Collaboration, taking our free parameters to be the angular size of the sound horizon θ_s, the physical baryon and cold dark matter (CDM) density parameters $\omega_b = \Omega_b h^2$ and $\omega_{cdm} = \Omega_{cdm} h^2$, the curvature fluctuation amplitude A_s, the scalar spectral index n_s and the optical depth to reionisation τ. Of course, we also add a free M_ν and from Chapter 3 onwards, a free N_{eff}.

In our analysis in this thesis, we choose to analyse the cosmology-dependence of our forecasts by focusing on two particular extensions to this model - free curvature (Ω_k) and a free dark energy equation of state (w_0), which is also allowed to vary in time in some cases with the addition of a w_a parameter, so that

$$w = w_0 + w_a(1 - a). \tag{1.31}$$

We chose to focus on these two extensions for a reason. The curvature parameter would be expected to be degenerate with the neutrino mass in CMB because of the geometric degeneracy between the Ω_k and the matter density [21][22]. The dark energy equation of state provides a significant degree of freedom when allowed to also be time-dependent. In other words, these extensions were chosen because they seemed the *most* likely to significantly weaken constraints, and therefore provide conservative results. Although there are already very strong constraints from Planck on Ω_k ($= 0.001 \pm 0.002$) and w_0 (-1.03 ± 0.03), we will see in the course of this thesis that even this allows enough room for significant effects on neutrino mass constraints.

1.7.3   Future Surveys

The aim of this work is to qualitatively understand the origins of cosmological neutrino mass constraints and their robustness when permitting the possibility of more complex cosmological models than flat ΛCDM. Although the papers published from this work have included constraint forecasts for a number of planned and theoretical galaxy and CMB surveys, these have been a secondary outcome and intended to demonstrate principles rather than to be taken as either optimal or realistic constraints.

Because of this, the list of future experiments for which we have performed forecasts is not exhaustive. Throughout the main chapters of this thesis, we take Euclid[4] as our sample galaxy survey and Simons Observatory[5] as our sample CMB experiment. Both

---

[4]euclid-ec.org
[5]simonsobservatory.org

Figure 1.6: Source: [1]. The intended sky coverage of the Euclid mission over the course of six years. The different colours represent different years.
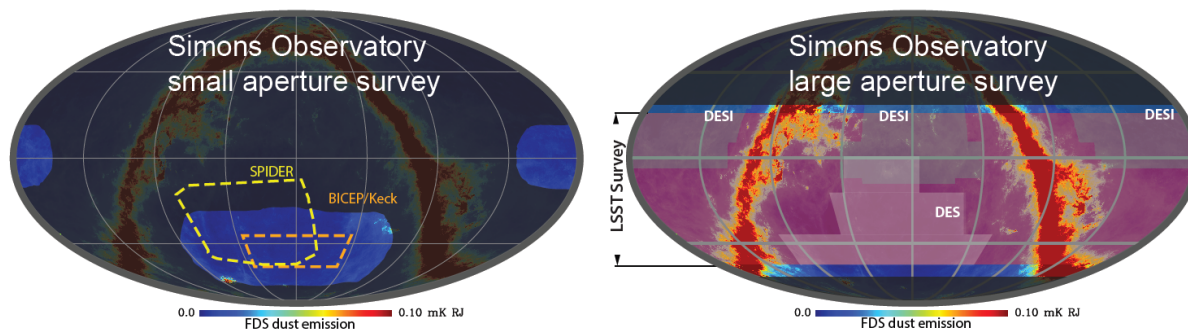
Figure 1.7: Source: [2]. The intended sky coverage (light regions) of the Simons Observatory small aperture telescope (SAT) and large aperture telescope (LAT), showing also the galactic dust emission distribution. The LAT survey aims to overlap as much as possible with the DESI and LSST surveys.

list improving neutrino mass constraints as a major motivation for their experiments. In Appendix B, we present limited forecasts for HETDEX[6], PFS[7], DESI[8] and WFIRST[9].

For Euclid [23, 24] we consider the 'Euclid Wide Survey', which aims to cover 15,000 deg^2 of the sky (see Figure 1.6) out to a redshift of 2.1, observing 2 billion galaxies in total, and obtaining spectra for about 50 million of them. It will carry out ambitious

---

[6]hetdex.org

[7]pfs.ipmu.jp

[8]desi.lbl.gov

[9]wfirst.gsfc.nasa.gov

measurements of both galaxy clustering and weak lensing.

The Simons Observatory project [2] is a combination of two surveys with two telescopes - the small aperture telescope (SAT), which will cover about 10% of the sky, and the large aperture telescope (LAT), which will cover about 40% (see Figure 1.7). We consider both in combination in this thesis, although the LAT is much more significant for our purposes.

Chapter 2

Galaxy Clustering

Much of the content of this chapter has been published in [25]. Some of the results may vary slightly as a result of a small correction made after the original submission (detailed in an Appendix of [25] and also in Section 2.4.4 of this chapter.)

2.1 Introduction

The galaxy power spectrum measured by galaxy redshift surveys provides multiple probes that can be used to constrain cosmological parameters. Massive neutrinos alter cosmological distances through their modification of the expansion rate, and cosmological distance scales can be constrained using standard rulers such as the baryon acoustic oscillation (BAO) scale (see Section 1.4.1), and through the Alcock-Paczyński (AP) test (see Section 1.4.3), which requires that isotropy is conserved when models are converted from redshift space. BAO information, in particular, is a popular probe of cosmological parameters because it is easily understood using linear theory and is easily measured, and does not depend heavily on an understanding of galaxy bias. For this reason, many previous studies have focused on extracting neutrino mass constraints from this source alone [26, 21]. Redshift-space distortions (RSD) (see Section 1.4.2) are used to constrain the growth of structure and are also affected by the additional matter provided by non-relativistic neutrinos.

However, there is nothing unique about the qualitative effect of massive neutrinos on the expansion rate of the universe or the large-scale structure growth rate. Either of these effects could be mimicked by the addition of other kinds of matter, or by changes in the nature of dark energy. The aim of this chapter is to deconstruct the information used to constrain the neutrino mass from galaxy redshift surveys, and to isolate the information available from those signals that are uniquely identifiable as the effects of massive neutrinos.

It is well known that neutrino free-streaming suppresses the growth of structure on small scales relative to that on large scales to an extent that is proportional to their mass, as outlined in Section 1.2. This results in small but distinctive signatures in the matter power spectrum $P_{\rm m}(k, z)$ (see Section 1.3) and in the structure growth rate $f(k, z)$. The possibility of constraining the neutrino mass through a scale-dependent measurement of $f(k, z)$ from

RSD was recently explored by [27]. The magnitude of the relative suppression also changes over time, leading to a redshift-dependence that also contributes to the uniqueness of the signal.

In this chapter, we aim to deconstruct the constraints on the neutrino mass offered by galaxy surveys, in order to understand how much constraining information is offered by each of the components listed above. In particular, we want to determine how sensitive the different constraints are to extensions to the cosmological model. We use minimal CMB information (we expand this in Chapter 3) and work only in the linear regime (the non-linear regime is investigated in Chapter 4).

The rest of this chapter is organised as follows. Section 2.2 provides a breakdown of our calculation method. Section 2.3 provides our findings, with some details expanded upon in Section 2.4. Section 2.5 comprises our conclusions.

2.2 Methodology

2.2.1 Overview

Model Parameters

Our simplest fiducial model consists of the six standard ΛCDM parameters and an additional total neutrino mass parameter. Fiducial values for the ΛCDM parameters were extracted from [28] based on the results from the `TT,TE,EE+lowP` dataset (see table 4 of [29]). At certain points we expand this model to free the curvature parameter Ω_k and the dark energy equation of state parameter w (which in some cases is allowed to be time-dependent). We assume a fiducial neutrino mass of 0.06 eV, which is close to the minimum limit implied by current neutrino oscillation experiments [see, for example, 30, for a relatively recent review of neutrino mass hierarchy measurements]. We make the approximation of one massive neutrino and two massless neutrinos. Our Fisher matrix parameters and their fiducial values are summarised in Table 2.1. Note that linear galaxy bias is also marginalised over as a nuisance parameter in all of our calculations, and its fiducial value varies depending on the survey and redshift bin. We carry out all of our forecasts with Euclid as an example survey (see Section 1.7.3). The survey parameters are given in Appendix A and results for an extended range of surveys in Appendix B.

Priors

All of our calculations are built upon a CMB prior. In the simplest case, we generate a minimalistic CMB 'compressed likelihood' prior from MCMC chains selected from the Planck Legacy Archive (the `base_mnu_plikHM_TT_lowTEB` dataset). The compressed likelihood prior compresses the information available from the CMB into four parameters that are effectively observables - the shift parameter $R = \sqrt{\Omega_m H_0^2} D_A(z_*)/c$ (where D_A is the angular diameter distance to the surface of last scattering), the angular scale of the sound

Parameter	Definition	Fiducial value
ω_b	Baryon density $\Omega_b h^2$	0.02225
ω_c	Cold dark matter density $\Omega_c h^2$	0.1198
$100\theta_s$	θ_s: Sound horizon size at last scattering (rad)	1.04077
τ	Optical depth to last scattering	0.079
$\ln(10^{10} A_s)$	A_s: Amplitude of the primordial power spectrum	3.094
n_s	Spectral index of the primordial power spectrum	0.9645
$M_\nu(\text{eV})$	Total neutrino mass	0.06
w_0	Time-independent dark energy equation of state parameter	-1
w_a	Time-dependent dark energy equation of state factor	0
Ω_k	Curvature parameter	0

Table 2.1: Summary of the model parameters. The first seven parameters are always free, and the final three are free in some cases. We also marginalise over a free linear bias parameter in each redshift bin.

horizon at last scattering $l_A = \pi/\theta_s$, ω_b and n_s (see section 5.1.6 of [29] for more information). We use the Fisher matrix mechanism to propagate these constraints into constraints on our cosmological parameter set. The advantage of the compressed likelihood prior is that the constraints it provides are relatively insensitive to variation in the curvature or dark energy equation of state. We also add very broad Gaussian priors on the parameters not constrained by the CMB prior to keep them within sensible ranges (see Section 2.3.1 for more information). In some cases, signficant improvements on the constraints can be achieved by including in the prior information on parameters related to the amplitude of CMB fluctuations, $A_s \exp(-2\tau)$ and τ, because of a significant degeneracy (see Section 2.4.3) and we highlight those cases.

The Full Galaxy Power Spectrum Fisher Matrix

To forecast the maximum amount of cosmological information available from a galaxy survey, constraints on the observed galaxy power spectrum $P_g(k, \mu)$ (where μ is the cosine of the angle with respect to the line of sight) are propagated directly into constraints on the cosmological parameters. The covariance of $P_g(k, \mu)$ can be most simply expressed for a given (k, μ) increment as [see, e.g. 31, 32]

$$\langle \Delta P_g(k, \mu)^2 \rangle = \frac{2\pi^2}{V k^2 \Delta k \Delta \mu} 2 P_g(k, \mu)^2, \qquad (2.1)$$

where V is the volume of the redshift bin being observed, and Δk and $\Delta \mu$ are the bin sizes for the wavenumber and angle with respect to the line of sight, respectively. Equation 2.1 applies in the case in which only one galaxy tracer population is assumed, with a single value assumed for the galaxy bias in each redshift bin. Here $P_g(k, \mu)$ is the full observed galaxy power spectrum including shot noise. Equation 2.1 can be appropriately generalised into a multi-dimensional band power matrix in cases in which multiple tracer populations

(with different biases) are used, which also accounts for their cross-correlation. For a single galaxy tracer population, we calculate the galaxy power spectrum in a particular redshift bin (including linear RSD) and shot noise as

$$P_g(k, \mu) = \left[b + f(k)\mu^2 \right]^2 P_{bc}(k) + \bar{n}_g^{-1}. \tag{2.2}$$

Here b is the fiducial bias of the galaxy sample, f is the growth function with $f = \frac{\mathrm{d}\ln D}{\mathrm{d}\ln a}$ (where D is the linear growth rate of perturbations) and P_{bc} is the real-space baryon-CDM power spectrum. $\bar{n}_g$ is the galaxy number density and the final term accounts for shot noise. In the linear regime, f is often taken as independent of scale, but massive neutrinos reduce the relative value of f on small scales by a small amount, so we include this effect here.

To convert observational measurements into a galaxy clustering model, fiducial values of $H(z)$ and $D_A(z)$ must be assumed. If the product of $H(z)$ and $D_A(z)$ assumed is incorrect, the three-dimensional model will be distorted. This is the AP test, and it provides another source of constraints on our cosmological parameters. Therefore, as a final step, we convert our k values into observable units and re-write the power spectra accordingly:

$$P(k_\parallel^{\mathrm{obs}}, k_\perp^{\mathrm{obs}}) = \frac{H(z)}{H_{fid}(z)} \left(\frac{D_{A,fid}(z)}{D_A(z)} \right)^2 P(k_\parallel^{\mathrm{com}}, k_\perp^{\mathrm{com}}), \tag{2.3}$$

where $k_\parallel^{\mathrm{obs}} = k_\parallel^{\mathrm{com}}(H_{fid}(z)/H(z))$ and $k_\perp^{\mathrm{obs}} = k_\perp^{\mathrm{com}}(D_A(z)/D_{A,fid}(z))$.

The linear matter power spectra used in our calculations were all generated using CLASS [9]. To generate the fiducial $P_{bc}(k)$ and $f(k)$ values as well as the numerical derivatives $\partial P_{bc}/\partial\theta_\alpha$ and $\partial f/\partial\theta_\alpha$, we generated matter power spectra for a very dense sample of z values, and stored the results in a two-dimensional table of k and z values. This table could then be interpolated to provide $P_{bc}(z, k)$ values. Values of $D(k)$ could be extracted by dividing the power spectra, and $f(z, k) = \mathrm{d}\ln D(z, k)/\mathrm{d}\ln a$ could then be calculated.

Care was taken with derivatives to ensure that they were not very sensitive to the increments by which the parameters were varied in their calculation. Increments that are too small can result in numerical scattering, while those that are too large lose finer elements of the structure. Derivatives were generally calculated as $(P[\theta + \epsilon] - P[\theta - \epsilon])/2\epsilon$, with P being either $P_{bc}(k)$ or $f(k)$, and with ϵ taking the values outlined in Table 2.2.

All of our constraints in this chapter are calculated with marginalisation over the linear galaxy bias b. Derivatives with respect to the bias parameter(s) can be calculated analytically using Equation 2.2. For a single tracer population:

$$\frac{\partial P_g(k, \mu)}{\partial b} = 2 \left[b + f(k)\mu^2 \right] P_{bc}(k). \tag{2.4}$$

It is important to define maximum and minimum usable k values in each redshift bin of a survey. k_{min} is calculated based on the dimensions of a particular redshift bin. k_{max} is a scale beyond which non-linear effects are too strong for linear approximations to be accurate. We choose $k_{\mathrm{max}} = 0.2 \ h \ \mathrm{Mpc}^{-1}$ here. In the case of BAO-only projections (the BAO signal is particularly robust against non-linear effects), it is common practice

Parameter	Increment (ϵ)
ω_b	0.001
ω_c	0.0025
$100\theta_s$	0.005
τ	0.025
$\ln(10^{10}A_s)$	0.05
n_s	0.01
$M_\nu(\text{eV})$	0.02
w_0	0.01
w_a	0.01
Ω_k	0.01

Table 2.2: List of the increment sizes used to calculate the derivatives numerically for each parameter.

to replace the sharp k-cutoff with an exponential degradation factor in the signal, which replicates a gradual smearing effect on the BAO peaks (see [33, 34]). In [35], it was suggested that the degraded BAO information could be recovered by reconstructing the original linear density field for a particular galaxy survey by using knowledge gleaned from the galaxy distribution to reverse the displacements of galaxies due to bulk flows and cluster formation. In the paper related to this chapter [25], for our BAO-only calculations, we replaced the sharp k cut-off with an exponential damping factor given by

$$P_{\text{BAO, damped}}(k, \mu) = P_{\text{BAO}} \exp\left[-\frac{1}{2}\left(k_\parallel^2 \Sigma_\parallel^2 + k_\perp^2 \Sigma_\perp^2\right)\right]. \tag{2.5}$$

The damping scales $\Sigma_\parallel$ and $\Sigma_\perp$ were calculated as a function of the structure growth rate $f(z)$ and the amplitude of the power spectrum $\sigma_8(z)$ as described by [34]. We followed the example of [32] to account for the possibility of improving constraints with reconstruction. We multiplied the damping scales for a given redshift bin by a reconstruction factor r calculated using the value of $n_g P_g(k = 0.14\ h\ \text{Mpc}^{-1},\ \mu = 0.6)$ in that bin. For high-density bins, r reached a maximum of 0.5, while in low-density bins it was just 1. For intermediate values, we interpolated over the same table of values given by [32]. In all other cases (beyond BAO), we used a sharp cut-off at $0.2\ h\ \text{Mpc}^{-1}$.

We abandoned this approach for our follow-up paper [36], switching instead to a hard cut-off k_{max}. This was done because we found that the results from the two approaches were not that different, and for ease of comparison with the other sets of results. In this chapter, for consistency with the next two chapters, we use once again a simple k_{max}. One can see how little difference this detail makes by referring to the results in the published paper for comparison.

Figure 2.1: The fiducial matter power spectrum for z=0.65 (the first Euclid redshift bin) with and without the BAO wiggles smoothed out.
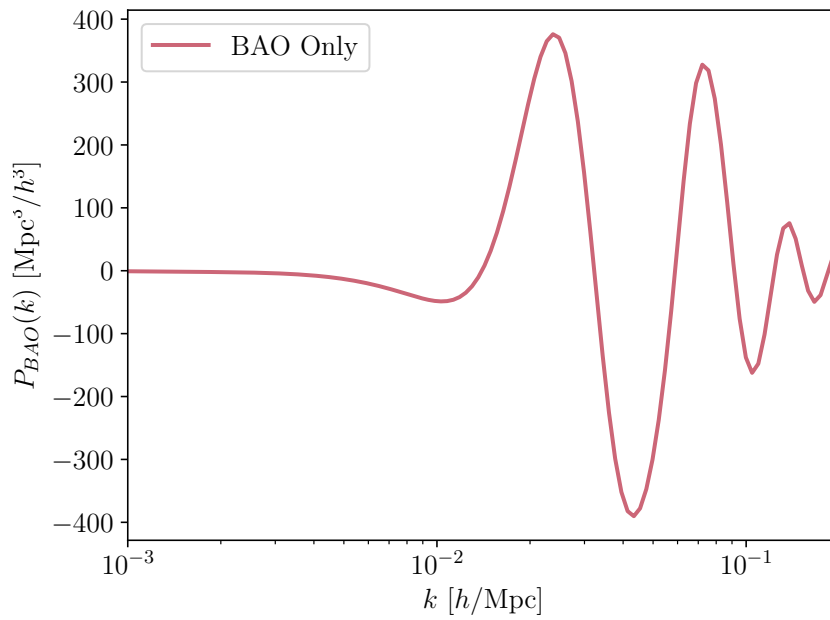
Figure 2.2: The fitted BAO signal removed from the power spectrum in Figure 2.1.

Removing Baryonic Oscillations from the Matter Power Spectrum

In the following sections, we attempt to determine the constraints that can be placed on the sum of the neutrino masses using different elements of the observed galaxy power spectrum. The sinusoidal BAO signal varies in both its phase and amplitude with many of our cosmological parameters. In some cases we need to remove the BAO signal from our Fisher derivatives of the matter power spectrum to isolate other effects, or to isolate the effects on the BAO signal alone.

We can consider the power spectrum to consist of two components, a BAO component and a smooth component (S): $P_{bc}(k) = P_{\mathrm{S}}(k) + P_{\mathrm{BAO}}(k)$. There are several common methods of extracting P_{S} (which can then be subtracted to obtain P_{BAO} alone), including fitting a spline to $P_{\mathrm{bc}}(k)$ that passes through the zero-points of the BAO oscillation, or using a formula for calculating the smooth power spectrum like that provided by [37]. These methods are unsuitable in our case, however, as we require the derivatives of P_{BAO} or P_{S} for insertion into the Fisher matrix, and small inaccuracies in the fitting of the matter power spectra can lead to artificially large or distorted derivatives. Therefore, we first calculate the derivatives of the full $P_{bc}(k)$, and then apply a smoothing function to the derivatives themselves to extract the smooth part of the derivative, which can be subtracted from the full derivative to obtain the derivative of the oscillatory part.

As a smoothing function we use a Savitzky-Golay filter [38]. The Savitzky-Golay filter sees the BAOs as noise and because of its averaging technique provides more consistent results than spline-fitting, which depends on manual selection of zero-point k values by sight. This application of the Savitzky-Golay method can be validated by applying it to a fiducial power spectrum (rather than a derivative) and then subtracting the fit from the original data to show a very regular and smooth BAO signal. In the case of derivative fitting, the smoothing is done in $\mathrm{d}\ln P(k)/\mathrm{d}\theta$ - log k space, and then both the original and smoothed spectra are plotted with a linear $P(k)$ scale to ensure that the fit remains reasonable. The derivative of the BAO component can then be obtained via subtraction and inspected. Figures 2.1 and 2.2 show the smoothed power spectrum and the isolated BAO wiggles for $z = 0.65$.

2.2.2   Distance Information

The most popular distance probe used to constrain cosmological parameters is the BAO feature. However, the broadband galaxy power spectrum also provides other means of constraining the cosmological distance parameters $H(z)$ and $D_A(z)$ [39]. The AP test requires $H(z)$ and $D_A(z)$ to scale appropriately with each other so that cosmological isotropy is preserved in real space. Other characteristic scales in the matter power spectrum, including the matter-radiation equality peak and the Silk damping scale, also provide distance constraints.

Full Distance Constraints

To extract the maximum amount of distance information from the galaxy power spectrum, we first use Equation 2.1 to extract constraints on $P_g(k, \mu)$. We then propagate these constraints into constraints on $\ln H(z)$ and $\ln D_A(z)$ (marginalised over bias b, the growth factor f and the underlying matter power spectrum P_{bc}) using the following derivatives (see e.g. [39]):

$$\frac{\partial P_g(k, \mu)}{\partial \ln D_A} = \frac{\partial P_g(k, \mu)}{\partial k} \frac{\partial k}{\partial \ln D_A} + \frac{\partial P_g(k, \mu)}{\partial \mu^2} \frac{\partial \mu^2}{\partial \ln D_A}, \tag{2.6}$$

$$\frac{\partial P_g(k, \mu)}{\partial \ln H} = \frac{\partial P_g(k, \mu)}{\partial k} \frac{\partial k}{\partial \ln H} + \frac{\partial P_g(k, \mu)}{\partial \mu^2} \frac{\partial \mu^2}{\partial \ln H}. \tag{2.7}$$

The derivatives of $P_g(k, \mu)$ with respect to k and μ^2 can be obtained directly from the calculated fiducial $P_g(k, \mu)$. The other terms are easily derived analytically:

$$\frac{\partial k}{\partial \ln D_A} = k(1 - \mu^2); \quad \frac{\partial k}{\partial \ln H} = -k(\mu^2); \quad \frac{\partial \mu^2}{\partial \ln D_A} = \frac{\partial \mu^2}{\partial \ln H} = -2\mu^2(1 - \mu^2). \tag{2.8}$$

BAOs

In the case that we want the information from the BAO signal alone, we must apply the method outlined in Section 2.2.1 to replace the full derivative of the matter power spectrum with just the oscillatory part when calculating the derivatives of $P_g(k, \mu)$ as above.

Seo and Eisenstein [34] provided a useful fitting function for forecasting H and D_A constraints from the BAO signal alone. We use our own fitting method here because it was most compatible in the context of our code, but our results agree well with published forecasts that use the Seo and Eisenstein method to predict constraints on $H(z)$ and $D_A(z)$, including a consistent $\ln H$-$\ln D_A$ correlation factor of 0.4.

AP Test

The AP test provides constraints on $H(z)$ and $D_A(z)$ by requiring that these values scale appropriately to preserve isotropy when the observed galaxy power spectrum is converted into real space coordinates. If the assumed product $H(z)D_A(z)$ is wrong, anisotropies will appear in the model. For AP information to be at its strongest, the redshift-space distortion effect must be well constrained so the two effects can be distinguished. Here, we extract tightest constraints that would be available from the AP test alone by holding the redshift space distortions fixed.

The AP test provides its constraints through changes in the observed galaxy power spectrum with the observation angle. The derivatives used are therefore as follows:

$$\frac{\partial P_g(k, \mu)}{\partial \ln D_A} = \frac{\partial P_g(k, \mu)}{\partial \mu^2} \frac{\partial \mu^2}{\partial \ln D_A}, \tag{2.9}$$

$$\frac{\partial P_g(k, \mu)}{\partial \ln H} = \frac{\partial P_g(k, \mu)}{\partial \mu^2} \frac{\partial \mu^2}{\partial \ln H}. \tag{2.10}$$

When the AP test is used alone, it provides a correlation coefficient between H and D_A of -1 [39]. Adding further distance information, such as standard rulers in the shape of the matter power spectrum, constrains H and D_A individually and allows this degeneracy to be broken.

2.2.3 Structure Growth

RSD

The RSD is an anisotropy that arises in the observed redshift-space power spectrum because the measured redshift of a particular galaxy is a function not just of the Hubble flow but also its peculiar velocity. The $[b + f(k)\mu^2]^2$ factor of Equation 2.2 is used to account for the Kaiser effect [15] resulting from structure formation, which causes an apparent strengthening of the clustering amplitude along the line of sight as objects fall into high-density regions. Galaxy survey measurements can be used to constrain $f\sigma_8$, where f is the structure growth rate and σ_8 the normalisation of the power spectrum amplitude, through analysis of this anisotropic signal.

White *et al.* [40] previously provided a method for isolating the information available from RSD in galaxy surveys, isolating f and keeping σ_8 fixed. We follow their example here, but include the scale-dependence of $f(k)$, and marginalise over both the bias and the entire matter power spectrum (which includes σ_8), using the following derivatives:

$$\frac{\partial P_g(k, \mu)}{\partial \theta} = \frac{\partial P_g(k, \mu)}{\partial f(k)} \frac{\partial f(k)}{\partial \theta} = \left[\frac{2\mu^2}{b + f(k)\mu^2}\right] P_g(k, \mu) \frac{\partial f(k)}{\partial \theta}, \tag{2.11}$$

$$\frac{\partial P_g(k, \mu)}{\partial P_{bc}(k)} = \left[b + f(k)\mu^2\right]^2. \tag{2.12}$$

We also marginalise over the distance parameters $H(z)$ and $D_A(z)$ using the derivatives given in Equations 2.6 and 2.7. We can also choose to extract constraints using the product $f\sigma_8$ if we do not wish to isolate the effect on the structure growth rate. σ_8 is calculated as an integral over the matter power spectrum and therefore also theoretically contains information on the suppression of the small-scale matter power by massive neutrinos.

The Small-Scale Suppression of the Structure Growth Rate

The treatment described in the previous section uses information from both the (constant) large-scale value of $f(k)$ and the small-scale suppression of $f(k)$ to derive constraints. We can now isolate the information available from the scale-dependent component of $f(k)$ alone. We re-write $f(k)$ as

$$f(k) = f_1.f_2(k). \tag{2.13}$$

Figure 2.3: The derivative of the matter power spectrum with respect to the neutrino mass, holding the amplitude of the power spectrum fixed.

Figure 2.4: The derivative of the structure growth rate with respect to the neutrino mass, holding the amplitude of the power spectrum fixed.

Here f_1 is the value of f on large scales ($k = 10^{-3}\ h$ Mpc^{-1} to be specific, where it is still constant, although the exact scale chosen should not matter) and $f_2(k)$ is a scale-dependent correction factor (equal to one on large scales). To isolate the information from the neutrino suppression of f on small scales, we replace f in our derivatives with f_2 multiplied by the fiducial value of f_1 and marginalise over f_1 as an additional Fisher matrix parameter alongside the other cosmological parameters. In other words, f_1 is taken as a constant outside the derivative and $df/d\theta_\alpha$ becomes $f_1.(df_2/d\theta_\alpha)$. This method removes any assumptions about the large-scale value of f and therefore can allow for the possibility of alternative gravity models. Figure 2.4 shows the derivatives of $f(k)$ with respect to the neutrino mass with f_1 held constant at a range of redshifts.

The Small-Scale Suppression of the Matter Power Spectrum

In order to obtain the information contained in the power spectrum shape alone, we must exclude the information from the BAO feature, which we have already accounted for. We extract the derivatives of P_s with respect to the cosmological parameters following the method outlined in Section 2.2.1. In the rest of this section, P_{bc} refers to the smoothed matter power spectrum.

As in the case of the structure growth factor f, we want to extract the scale-dependent component of P_{bc} alone as an information source, neglecting the amplitude of the power spectrum (see Figure 2.3). As before, we introduce a new parameterisation

$$P_{bc}(k) = P_1.P_2(k), \tag{2.14}$$

where P_1 is the matter power spectrum value at $k = 10^{-3}\ h$ Mpc^{-1} (the point of normalisation should not matter, and we try to avoid the edges where interpolation effects are stronger). The derivative of our Fisher matrix becomes:

$$\frac{\partial P(k,\mu)}{\partial \theta_\alpha} = \frac{\partial P(k,\mu)}{\partial P_{bc}}\frac{\partial P_{bc}}{\partial \theta_\alpha} = \left[b + f(k)\mu^2\right]^2 P_1\frac{\partial P_2}{\partial \theta_\alpha} \tag{2.15}$$

The derivatives of P_{bc} with respect to M_ν obtained from `CLASS` are provided in Figure 2.3 with P_1 held fixed before smoothing. We marginalise over P_1, bias, the structure growth rate f, and the distance parameters $H(z)$ and $D_A(z)$.

2.3 Results

2.3.1 Priors

We should begin by understanding how parameters are correlated with each other in the CMB prior. The prior only really constrains θ_s, n_s and ω_b. In some cases, we demonstrate the effect of adding information on parameters related to the amplitude of the fluctuations of the CMB: $A_s \exp(-2\tau)$ and τ (see Sections 2.3.3 and 2.4.3). The remaining parameters

Parameter	1σ error
$\ln(10^{10}A_s)$	1.0
τ	0.5
M_ν (eV)	1.0
ω_{cdm}	0.2
Ω_k	0.1
w_0	1.0
w_a	3.0

Table 2.3: 1σ uncertainties imposed on the cosmological parameters as initial priors. θ_s, n_s and ω_b are constrained by the CMB prior.

- M_ν, ω_{cdm}, Ω_k, w_0 and w_a - all modify R, but as the other prior parameters are held fixed, their effects are completely degenerate and they are unconstrained from the prior alone without additional information. A useful study of the degeneracies between M_ν and parameters in CMB data is provided in [41]. We add wide Gaussian priors on these unconstrained parameters to keep their values sensible. The general intention is that the results that follow should be relatively independent of the exact priors chosen. Table 2.3 summarises these prior values.

2.3.2 Distance Information: BAO and AP

Constraints on the distance parameters $H(z)$ and $D_A(z)$ are derived from two main sources in galaxy surveys. The AP test constrains the product of $H(z)$ and $D_A(z)$ by requiring that the galaxy clustering pattern derived from observations be isotropic. The BAO scale imprinted on the galaxy clustering pattern breaks the degeneracy between H and D_A and allows them to be measured individually. Roughly speaking, the AP test constrains $D_A(z)H(z)$ and the BAO signal constraints $D_A(z)^2/H(z)$ [e.g. 39].

H and D_A change with θ_s and the matter density. The strength of distance information in constraining the neutrino mass lies in its ability to break the correlation between ω_{cdm} and M_ν in the CMB prior. The effects of increasing either ω_{cdm} or M_ν on H and D_A are strongly degenerate, leading to strong anti-correlation between the two parameters that breaks the degeneracy created by the CMB constraints. The associated disadvantage, however, is that distance information is in fact sensitive to the sum of M_ν and ω_{cdm}, of which M_ν makes up a tiny fraction. Understanding the relationship between M_ν and ω_{cdm} is the key to understanding the constraints on M_ν provided by distance probes. Figures 2.5 and 2.6 show the effect of changing these parameters on H and D_A as a function of redshift.

Increasing ω_{cdm} or M_ν in the ΛCDM context requires that Ω_Λ is decreased to maintain the critical energy density. At higher redshifts, the increase in Ω_m is the dominant effect on $H(z)$, which is then increased relative to the fiducial model. In the later, dark-energy-dominated regime, the decrease in Ω_Λ dominates the change in $H(z)$, which is now reduced

compared to in the fiducial model. The effects are not completely degenerate as the crossover occurs at a higher redshift with additional massive neutrinos than with additional cold dark matter, as a result of the historical relativistic nature of the massive neutrinos. These changes in $H(z)$ mean that $D_A(z)$ is increased relative to the fiducial model at the redshifts covered by our surveys. The similarity in the effects of increasing ω_{cdm} and M_ν on the distance parameters (i.e. both increase H at high z and decrease it at low z, and both increase D_A, particularly at late times) makes the two parameters highly anti-correlated.

Figures 2.5 and 2.6 also show the impact of varying w_0, w_a and Ω_k on the distance parameters, helping us understand how freeing these parameters can affect the constraint on M_ν. Increasing w_0 or w_a reduces $H(z)$ in the later, dark-energy-dominated regime, but increases it at higher redshifts. This effect is qualitatively similar to that of increasing M_ν and ω_{cdm}, and the corresponding effects on $D_A(z)$ are also similar. Freeing w therefore weakens the constraint on ω_{cdm} considerably. However, the constraints on M_ν are much less affected because M_ν starts to suppress $H(z)$ at much higher redshifts, and the slope of the derivative is much less steep than for the other parameters, allowing this effect to be distinguished.

On the other hand, Figures 2.5 and 2.6 also show that the shape of the derivative of Ω_k is very similar to that of M_ν, though inverted. The effects of Ω_k and M_ν on both distance parameters are quite clearly degenerate, and the effect of an increase in Ω_k could quite clearly be compensated by a reduction in M_ν, and vice versa. These effects are reflected in the results we obtain.

Figure 2.7 shows a breakdown of the distance constraints on M_ν for a series of cosmological models. It is clear that the vast majority of the distance information comes from the BAO signal. It is also clear that in all cases, the constraints are significantly degraded when Ω_k is allowed to vary. If BAO information is included, the constraints are not very sensitive to assumptions about the dark energy equation of state.

With the BAO feature smoothed out, $H(z)$ and $D_A(z)$ become strongly anti-correlated, due to the AP test dominating the remaining information. Using the BAO feature alone, there is a consistent correlation coefficient of approximately 0.4. With all of the distance information, the correlation is approximately -0.4 to -0.5.

Table 2.4 shows the intermediate constraints on $H(z)$ and $D_A(z)$ for Euclid in each bin, which are often quoted in other works and therefore useful for comparisons.

2.3.3 Structure Growth Information

Sources of information on the growth of structure include the shape and amplitude of the matter power spectrum, and the observed anisotropies created by RSD.

RSD

We can use RSD to extract constraints on the product $f\sigma_8$ or on f marginalised over σ_8 if we want to isolate constraints from the structure growth rate alone. It is important to note that σ_8 is calculated as an integral over the matter power spectrum, and therefore also

Figure 2.5: Derivatives of $H(z)$ with respect to various cosmological parameters as a function of redshift.

Figure 2.6: Derivatives of $D_A(z)$ with respect to various cosmological parameters as a function of redshift.

Figure 2.7: A breakdown of the forecasted distance information constraints on M_ν for Euclid (including a CMB prior on θ_s, n_s and ω_b), for a variety of models. It is clear that the primary source of constraining information is the BAO signal. All of the constraints are weakened considerably if Ω_k is allowed to vary, and in the AP only case this results in effectively no constraint.

z	$\sigma(H(z))(\%)$	$\sigma(D_A(z))(\%)$
0.65	1.69	1.41
0.75	1.34	1.02
0.85	1.23	0.93
0.95	1.18	0.9
1.05	1.15	0.89
1.15	1.14	0.88
1.25	1.14	0.89
1.35	1.16	0.92
1.45	1.21	0.98
1.55	1.29	1.08
1.65	1.44	1.24
1.75	1.67	1.49
1.85	2.08	1.93
1.95	2.82	2.73
2.05	5.94	6.12

Table 2.4: Constraints on $H(z)$ and $D_A(z)$ using the BAO wiggles only.

contains information on the overall shape of the matter power spectrum, which is altered by massive neutrinos. The relationship between σ_8 and M_ν is therefore quite complex.

We begin by considering the information available from constraints on f alone, marginalised over σ_8. Figure 2.8 provides the key results for Euclid in two different cases. The upper panel demonstrates the constraining power of $f(k)$ values at all scales (including both the large-scale values and the relative change on small scales characteristic of massive neutrinos). We see that the constraints vary significantly depending on the cosmological model adopted. The lower panel shows the constraints if only the scale-dependence of f is considered, with the large-scale value of f being marginalised over. Here we see that the results are completely independent of the choice of cosmological model. This is because the scale-dependence of the structure-growth rate probed is a unique indicator of massive neutrinos, and is not replicated by the additional parameters we can include.

We briefly deconstruct the constraints we see in the upper panel of Figure 2.8. Figure 2.9 demonstrates the effect of changing the most relevant cosmological parameters on the large-scale value of f. Adding information on the large-scale structure growth rate to the CMB prior provides an improvement in the neutrino mass constraints by inverting the correlation between ω_{cdm} and M_ν in the CMB prior, as both parameters increase f by adding additional matter.

Increasing w_0 or w_a reduces f at higher redshifts and increases it at lower redshifts, so these two parameters become correlated with M_ν at the redshifts covered by our surveys, weakening the neutrino mass constraint. The effect of freeing Ω_k is quite complex. It can be seen from Figure 2.9 that the effects of ω_{cdm} and Ω_k on f are strongly degenerate. We may expect M_ν and Ω_k to be correlated because the former increases f and the latter

reduces it. However, ω_{cdm} is strongly anti-correlated with M_ν, as discussed previously, and is much more strongly correlated with Ω_k than M_ν is. So the net effect, including the distance prior, is an anti-correlation between M_ν and Ω_k. This results in the weakening of constraints with free curvature seen in Figure 2.8. In general, we marginalise over P_{bc}, b, $H(z)$ and $D_A(z)$ here. Fixing $H(z)$ and $D_A(z)$ in this case can actually improve constraints quite significantly (from 0.24 eV to 0.17 eV for Euclid in the ΛCDM case). This demonstrates the complementarity of BAO and RSD information.

We can next examine the constraints achievable from the combination $f(k)\sigma_8$. In this case, the choice of whether to include information on $A_s\exp(-2\tau)$ and τ in the CMB prior plays a crucial role in the constraints on M_ν obtained.

Figure 2.10, which shows a breakdown of the RSD constraints for Euclid with three possible choices of CMB prior, distinguished by different fill patterns. The largest error bars correspond to the four-parameter CMB prior we have used so far (constraining θ_s, n_s and ω_b). Constraining σ_8 requires a meaningful prior on A_s, which itself requires reasonable constraints on τ as $A_s\exp(-2\tau)$ is the parameter measured from the CMB. In the case in which $A_s\exp(-2\tau)$ and τ are not included in the CMB prior, the constraints on M_ν ultimately become controlled by the uncertainty on these two parameters, and freeing other parameters makes little difference, giving very uniform constraints across the cosmological models we consider. Including the constraining power of the CMB on $A_s\exp(-2\tau)$ and τ can improve the constraints significantly, as seen from the difference in the bars in Figure 2.10. When A_s is reasonably constrained, the constraint on M_ν becomes dominated by how well τ is known. The middle bars show the constraints when prior information from Planck on τ and $A_s\exp(-2\tau)$ is included. This corresponds to constraints of $\sigma(\ln 10^{10}A_s) \approx 0.04$ and $\sigma(\tau) \approx 0.02$ from the distance prior alone. The shortest bars include the prior information from Planck 2015 [28] on $A_s\exp(-2\tau)$ but assume that τ is known perfectly. Further discussion on this topic is provided in Section 2.4.3.

The Shape of the Matter Power Spectrum

The cumulative effect of neutrinos suppressing the structure growth rate f over time is a corresponding small-scale suppression in the matter power that increases in magnitude with time. We can isolate this effect by taking the power spectrum amplitude outside the derivative and smoothing out the BAO signal. Figure 2.11 shows the constraints obtained from this signal alone. We see that the constraints are relatively robust against variations in the model, as we would expect. We also note that these constraints are tighter than those from the scale-dependence of $f(k)$, as the fractional change in the matter power spectrum caused by massive neutrinos at low redshifts is larger than the fractional change in the growth factor.

2.3.4 Combining the Suppression Signals

The suppression of $P_{bc}(k)$ and $f(k)$ caused by massive neutrinos on small scales can be combined to maximise the constraint from this effect. The result is dominated by the shape
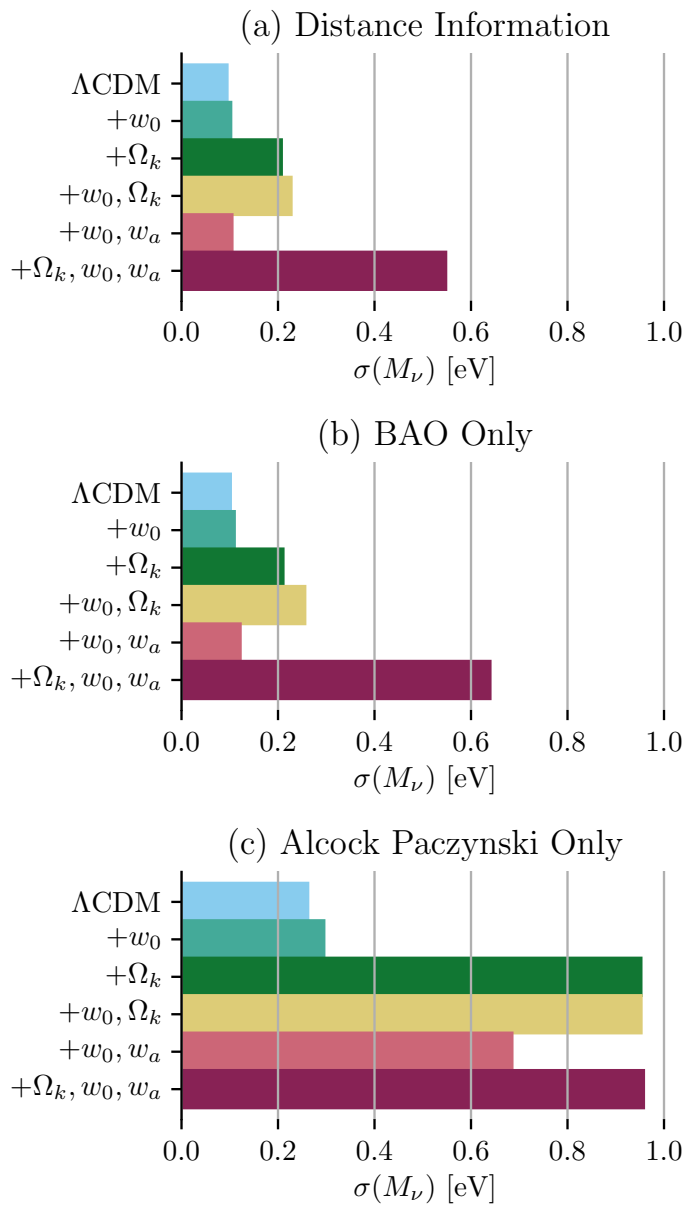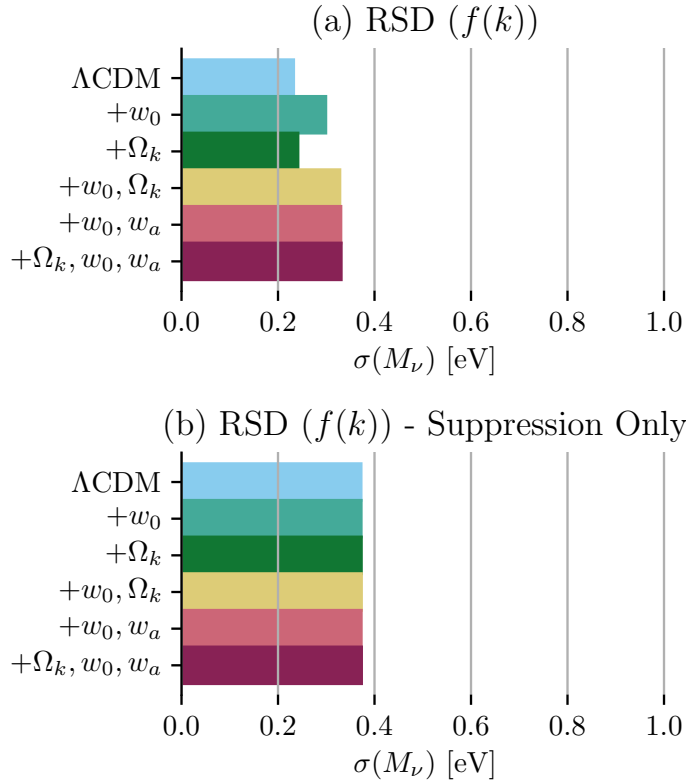
Figure 2.8: A breakdown of the forecasted constraints on M_ν from RSD for Euclid (including a CMB prior on θ_s, n_s and ω_b), for a variety of models. In both cases, RSD are used to constrain the structure growth rate f marginalised over the matter power spectrum (which includes σ_8). The lower panel gives the constraints from the scale-dependence of $f(k)$ alone (the large-scale amplitude of $f(k)$ is taken outside the derivatives and marginalised over). The upper panel gives the constraints available from both the large-scale amplitude and scale-dependence of $f(k)$.

of P_{bc}, and information from the effect on f is a much more minor contribution, but does provide some improvement on the constraints from P_{bc} alone. Figure 2.12 demonstrates this. This combination is the most robust probe of the neutrino mass that we identify in this work. As Figure 2.12 demonstrates, the constraints are not dependent on basic assumptions about the dark energy equation of state or curvature. The constraints provided by this combination are also competitive with constraints from distance probes and RSD.

2.3.5   Combined Information

Ultimately, the constraints from the total galaxy power spectrum can be broken down into two categories - distance constraints (BAO, AP, etc.) and constraints from the growth of structure (the shape and amplitude of the matter power spectrum, RSD, etc.). Combining

Figure 2.9: The derivatives of $f(z)$ with respect to the parameters M_ν, $\omega_{\rm cdm}$, w_0, w_a and Ω_k. The effects of adding extensions to the ΛCDM model on the neutrino mass constraints can be understood by comparing these derivatives. The values of f used are those for $k = 10^{-3} h$ Mpc^{-1}. The derivatives with respect to ω_{cdm} and Ω_k have been re-scaled by a factor of 0.1 for plotting purposes.

these two probes to extract the maximum amount of information is powerful. Figure 2.13 shows a breakdown in the combined constraints for Euclid, without any CMB prior being included for $A_s \exp(-2\tau)$ and τ (the effect of adding these priors is demonstrated in Section 2.4.3). We see that the constraints suffer considerably if more model parameters are allowed to vary. For example, assuming ΛCDM with Euclid gives a constraint of 0.037 eV. This weakens to 0.07 eV for our most complex model. This emphasises the inherent weakness of taking constraints from the entire observed galaxy power spectrum without closer analysis. We emphasise that the constraints derived here are somewhat larger than other published results. This is because of the choice of the compressed likelihood prior (see Section 2.2.1), which constrains only θ_s, n_s and ω_b, and does not include the constraining power of CMB lensing (see Section 2.4.2 for discussion).
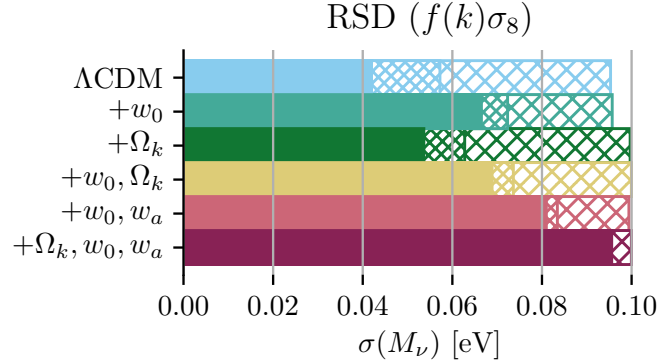
Figure 2.10: A breakdown of the forecasted constraints on M_ν from RSD for various priors and cosmological models. In this case, RSD is used to constrain $f\sigma_8$ (the shape of the matter power spectrum is marginalised over). The different fill patterns represent different prior conditions on $A_s \exp(-2\tau)$ and τ (from longest to shortest: no CMB information on $A_s \exp(-2\tau)$ and τ, Planck priors on $A_s \exp(-2\tau)$ and τ, and Planck 2015 priors on $A_s \exp(-2\tau)$ with perfectly-known τ). Note that the length of the x-axis has been altered here for greater precision.

Figure 2.11: A breakdown of the forecasted constraints on M_ν for Euclid from the scale-dependent suppression of the matter power spectrum characteristic of massive neutrino free-streaming. The amplitude of the matter power spectrum at $k = 10^{-3} \ h \ \mathrm{Mpc}^{-1}$ is marginalised over.

Combined Free-Streaming Signals
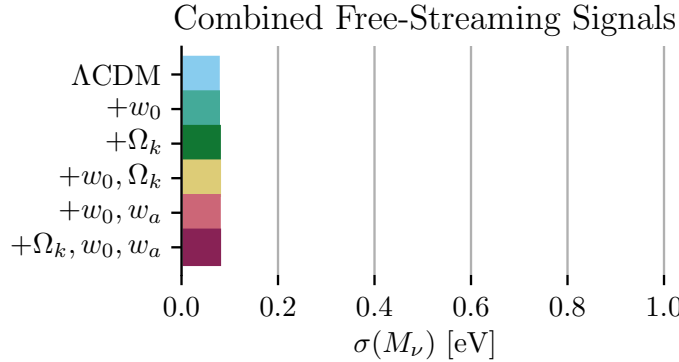
Figure 2.12: This figure combines the constraints from the free-streaming signals in both $f(k)$ and $P_{bc}(k)$ for Euclid. We see through comparison with Figure 2.11 that the constraint is dominated by information from the shape of the matter power spectrum.

Combined

Figure 2.13: The total combined neutrino mass constraint forecasts for Euclid. No CMB prior on $A_s \exp(-2\tau)$ or τ is included.

2.4 Discussion

2.4.1 Significance

In the previous section, we isolated the neutrino mass information available from the distinctive effects of neutrino free-streaming on $P_{bc}(k)$ and $f(k)$ in galaxy surveys. The constraints from this signal alone are weaker than the combined constraints that are usually quoted. However, the probes we have emphasised are directly relateable to neutrino physics and are relatively insensitive to the assumed cosmological model, unlike in the combined case. In a time when the ΛCDM model still leaves many open questions, an upper limit on the neutrino mass can only be asserted with confidence if the constraint is reasonably independent of small changes in the assumptions about the underlying cosmology. The unique effects of massive neutrinos on the growth of structure provide this. Our calculations

therefore provide conservative constraints that should be convincing to particle physicists and those working outside the cosmology community.

There are effects beyond the cosmological extensions that we have considered here that also induce a scale-dependence in the structure growth rate, for example, modified gravity models such as $f(R)$ theories (a study of the degeneracies between $f(R)$ gravity and massive neutrinos was carried out in [42]). However, we consider it unlikely that the characteristic scale and magnitude of the suppression caused by massive neutrinos would be exactly replicated by another effect, particularly if the redshift-dependence of the effect could be measured.

2.4.2 Choice of CMB prior

Our forecasted constraints are somewhat weaker than those in other published work because of our choice of a very conservative CMB prior, which neglects some important effects that can help constrain M_ν. The most significant of these is CMB lensing, which probes the shape and amplitude of the matter power spectrum directly, and is considered in [43, 44, 21, 26, 41]. As our aim in this chapter has been to disentangle neutrino mass constraints, using a very simple prior made sense, and for testing the cosmological dependence of different probes, we also needed a cosmology-independent prior. In Chapter 3, we expand on this work to analyse the effects of including full information from primary CMB anisotropies and CMB lensing in particular.

2.4.3 Sensitivity to priors on A_s and τ

The small-scale CMB temperature power spectrum is sensitive to the parameter combination $A_s \exp(-2\tau)$ [45], making A_s and τ very strongly correlated when using CMB temperature anisotropy data alone. Strengthening the constraint on τ is often specified as a recommended route towards improving constraints on M_ν [e.g. 21, 46]. The constraints on the neutrino mass from structure growth information (parameterised, for example, by $f\sigma_8$) are sensitive to the primordial amplitude A_s, but not sensitive to τ. However, because A_s and τ are so strongly correlated in the CMB prior, adding the prior makes the constraint on M_ν strongly dependent on the weak constraint on τ provided by CMB polarisation. Ultimately, when constraints are strong enough, the τ constraint becomes the limiting factor when trying to strengthen the constraint on M_ν. This is why we can obtain significant improvements in some of our results by extending our compressed likelihood CMB prior to also include constraints on $A_s \exp(-2\tau)$ and τ. We can expect the weak constraints on τ from Planck to be significantly improved by 21 cm emission measurements used to probe the epoch of reionisation [46], and we provide results for fixed τ to give forecasts in the most optimistic cases.

In Figure 2.14, we show how the neutrino mass constraint is limited by the constraint on τ when τ is included in the CMB prior. In Figure 2.15, we demonstrate the impact on the neutrino mass constraints from Euclid when τ is perfectly known. As previously noted, distance measurements have no dependence on A_s, and are therefore unaffected by

our choice to improve constraints on these parameters. But in the case of structure growth probes there is a significant improvement, which ultimately leads to an improvement in the combined constraints, particularly for the simpler models.

2.4.4 Choice of Power Spectrum

One subtlety that was pointed out to us after the submission of our article [25] is that the galaxy power spectrum (Eq. 2.2) should be calculated not as

$$P_g(k, \mu) = \left[b + f_m(k)\mu^2\right]^2 P_m(k) + \bar{n}_g^{-1},$$ (2.16)

but as

$$P_g(k, \mu) = \left[b + f_{bc}(k)\mu^2\right]^2 P_{bc}(k) + \bar{n}_g^{-1}.$$ (2.17)

While P_m is the power spectrum including the contribution of massive neutrinos, P_{bc} is the power spectrum of baryons and cold dark matter only, and f_{bc} is the growth rate corresponding to this power spectrum. The motivation for this is that neutrino perturbations do not contribute to the formation of galaxy haloes. In the published article [25], the results were calculated using Equation 2.16 with a brief discussion of the effects of using Equation 2.17 provided in an appendix. The effects can now be more fully understood by comparing the results provided in this chapter with those in that article. A comparison shows that the changes are minor and all of the qualitative results hold.

2.4.5 Non-Linear Effects

The minimum scale chosen for most of our calculations ($k_{\max} = 0.2\ h\ \mathrm{Mpc}^{-1}$) is a scale at which non-linear effects may play a small role. However, the primary aim of the present work is to examine the relative strength of the different probes we describe, and one could expect that all of these constraints would be modified similarly by the inclusion of non-linearities. We have also run our results for $k_{\max} = 0.15\ h\ \mathrm{Mpc}^{-1}$ and find that there is not a significant qualitative change in our conclusions. In Chapter 4, we extend the work of this chapter into the mildly non-linear regime.

2.5 Conclusions

In this chapter, we have analysed the various components of the observed linear galaxy power spectrum as will be measured by future galaxy redshift surveys to determine how these different components can contribute to a determination of the total neutrino mass, M_ν. Adding massive neutrinos to a cosmological model alters the expansion rate (an effect that can be identified through measurements of $H(z)$ and $D_A(z)$ derived from the AP test and standard rulers such as the BAO signal). Massive neutrinos also modify the structure growth rate (an effect that can be measured using RSD). The most unique

Figure 2.14: Contour plot demonstrating how the constraints on M_ν from the CMB prior and the combined galaxy redshift survey information become dominated by the constraints on τ when Planck priors on $A_s \exp(-2\tau)$ and τ are included. Note that in the CMB prior only case, an 'extra M_ν constraint' is included to keep the 1σ error on M_ν below 0.1 eV. This is purely for demonstrative purposes - without this the CMB prior would only impose an error of 1.0 eV (see Table 2.3), which would inflate the axis scales and make the contour for the CMB prior combined with Euclid ($\sigma M_\nu \approx 0.03$ eV) difficult to make out. We see that when the CMB prior alone is used, there is little correlation between M_ν and either A_s or τ, but the latter two are strongly correlated with each other. When the combined information from the Euclid survey is added, M_ν becomes strongly correlated with A_s, making it also strongly correlated with τ. τ is currently only weakly constrained by CMB polarisation information, so the neutrino mass constraint becomes limited by our knowledge of the value of τ.
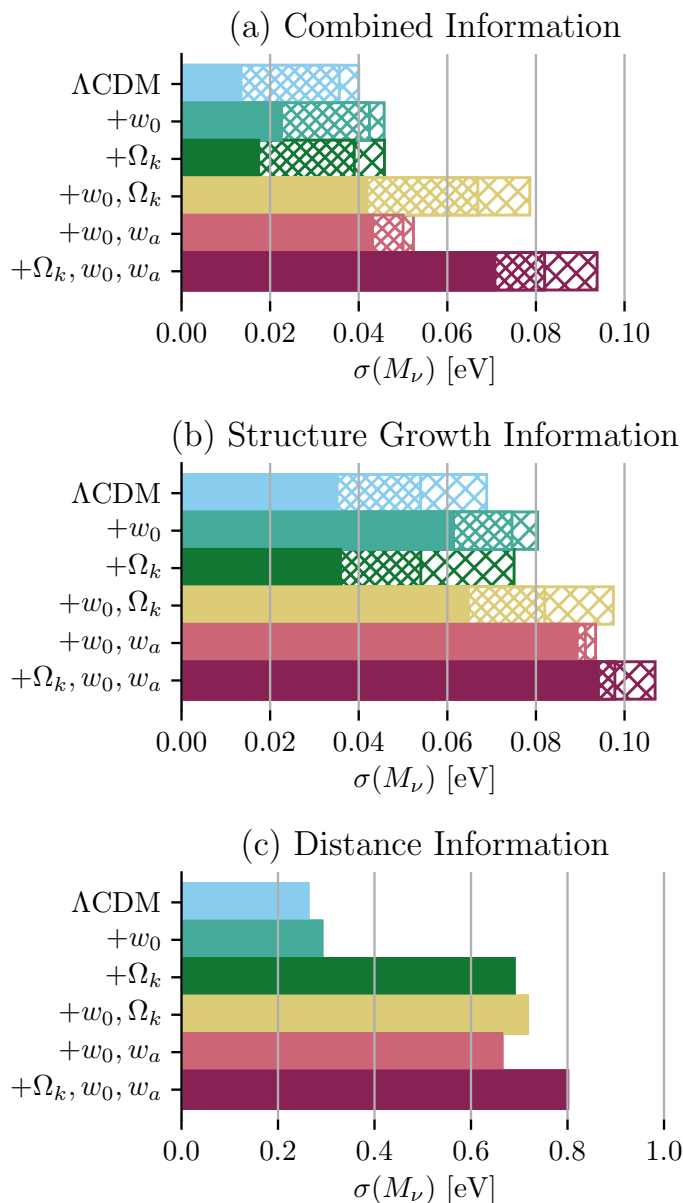
Figure 2.15: The forecasted constraints from Euclid, with the combined constraints broken down into their broad components, for various choices of prior for $A_s \exp(-2\tau)$ and τ. The different fill patterns represent these different priors (from longest to shortest: no CMB information on $A_s \exp(-2\tau)$ and τ, Planck 2015 priors on $A_s \exp(-2\tau)$ and τ, and Planck 2015 priors on $A_s \exp(-2\tau)$ with perfectly-known τ). The x-axis of the distance information panel is extended to keep roughly the same scale as the other two panels for comparison.

identifier of massive neutrinos, however, is a small-scale relative suppression in the matter power spectrum and in the structure growth rate, as the free-streaming of massive neutrinos reduces clustering on small scales. This is a probe of the neutrino mass that provides constraints that are independent of simple extensions to the assumed cosmological model. In this chapter, we have disentangled all of these measurement tools, and demonstrated the sensitivity of each to the assumed cosmological model. We have used a minimalistic CMB prior (constraining R, l_A, ω_b and n_s), which provides constraints that are relatively insensitive to the assumed curvature or dark energy equation of state.

We have also provided forecasts of the neutrino mass constraint using all of the information available in the observed galaxy power spectrum combined. These constraints can be weakened by more than a factor of two when the curvature or dark energy parameters are allowed to vary. We have confirmed that the combined constraints on M_ν are limited by the accuracy to which τ is known [21] and clarified the reason for this.

To extract the most robust constraints possible on M_ν, future surveys should focus on constraining $P_{\rm bc}(k)$ and $f(k)$ as precisely as possible, which requires solid measurements of the broadband shape of the galaxy power spectrum and redshift-space distortion analysis that considers the scale-dependence of f. Large-scale and small-scale values of $P_{\rm bc}$ and f can be compared to attempt to identify the relative suppression of structure growth caused by massive neutrino free-streaming. As the magnitudes of these two effects varies somewhat with redshift, surveys with deeper redshift surveys may provide more definitive results by showing a contrast between the signals measured at different redshifts.

Chapter 3

CMB Anisotropies and CMB Lensing

Much of the content of this chapter has been published in [36].

3.1 Introduction

In Chapter 2, we thoroughly deconstructed the constraints on the sum of neutrino masses, M_ν, achievable with upcoming spectroscopic galaxy surveys. Our focus was to understand how sensitive forecasts were to cosmological assumptions (specifically about curvature and the dark energy equation of state), and to understand where the majority of the constraining power came from. We found that constraints could degrade significantly when moving beyond ΛCDM depending on the data that was used. For example, we showed that constraints derived from distance probes, such as baryon acoustic oscillations (BAOs), could become several factors weaker if curvature was allowed to be non-zero.

We determined that the most reliable probe of M_ν, due to its robustness against changes in the underlying cosmology, was the distinctive scale-dependent free-streaming signature that massive neutrinos imprint on the underlying matter power spectrum and the structure growth rate, and provided a method of extracting isolated constraint forecasts from these effects alone. We demonstrated that, if measured, these signals could provide neutrino mass constraints that are insensitive to the assumed cosmology. Significantly, these constraints are also independent of τ, the optical depth to the cosmic microwave background (CMB).

In this work, we extend our analysis considerably to include forecasts from future CMB experiments. In Chapter 2, we used a very conservative Planck prior, the compressed likelihood prior [29], which aims to provide constraints from effective observables only, and is therefore insensitive to assumptions about curvature and dark energy. Here we want to expand our calculations to include the full benefits provided by CMB surveys, including the temperature power spectrum from Planck [47] and forecasted polarisation anisotropy measurements and CMB lensing.

We perform calculations for three combinations. First, we examine the most optimistic case, analysing the advantages of combining CMB lensing measurements with full broadband galaxy power spectra from spectroscopic surveys. Second, we look at the com-

bination of BAO measurements from galaxy surveys with CMB lensing, a combination that has been very popular in existing forecasts [48, 26, 21]. Finally, we demonstrate that our cosmology-independent free-streaming constraints from [25] can be improved by combining with powerful CMB measurements, while still remaining cosmology-independent. We consider these constraints the most robust forecasts.

The primary motivation for combining the free-streaming constraints with CMB forecasts is that the CMB lensing power spectrum also contains a relative suppression on small scales caused by the effects of neutrino free-streaming, analogous to that in the matter power spectrum. While the galaxy power spectrum is a biased measurement of the matter power spectrum, CMB lensing probes the matter power spectrum directly. Additionally, while the galaxy power spectrum contains contributions from the baryon and cold dark matter transfer functions only, the matter power spectrum probed by CMB lensing includes all matter, including massive neutrinos. It seems reasonable the CMB lensing could provide a useful complement to these measurements.

Additionally, forecasted constraints on M_ν from upcoming galaxy surveys based on the full galaxy power spectrum will primarily be limited by weak constraints on τ (see [25] and also [21]). The correlation between the two parameters arises from both being strongly correlated with A_s, as described in Chapter 2. As shown here and also recently by [49], CMB lensing could help overcome this obstacle somewhat through its potential for constraining A_s.

We work exclusively in the linear regime in this work. Analysis of the effects of implementing the next-to-leading-order (NLO) power spectrum is performed in Chapter 4.

This chapter is organised as follows. We outline our methodology in Section 3.2, and in Section 3.3 we provide a detailed breakdown of the effects of combining CMB lensing information with various types of galaxy survey forecasts, with some discussion. We conclude in Section 3.4.

3.2   Methodology

Our calculations in this work focus on the combination of CMB and spectroscopic galaxy surveys. Our Fisher matrix implementation for galaxy surveys and our fiducial cosmology are the same as in Chapter 2.

There are some minor changes to our method. The list of cosmological parameters we use and their fiducial values remain consistent with those in Chapter 2. However, we also now add to the list $N_{\rm eff}$ to account for the degeneracy between $N_{\rm eff}$ and M_ν in CMB observables. This was not possible when using the Planck compressed likelihood, as the compressed likelihood priors are provided for a certain combination of free parameters, and we could not obtain one that also kept $N_{\rm eff}$ free. Constraints in this paper are therefore marginalised over a total set of parameters: θ_s, A_s, $N_{\rm eff}$, n_s, ω_{cdm}, ω_b and τ in all cases. For the extended models, the list may be extended to include some or all of Ω_k, w_0 and w_a.

In this section, we discuss our Fisher matrix implementation for forecasting CMB constraints. The covariance matrix for a particular set of angular power spectra is given by

$$\langle \Delta C_l^{xy} \Delta C_l^{mn} \rangle = \frac{1}{(2l+1)f_{\text{sky}}dl}(C_l^{xm}C_l^{yn} + C_l^{xn}C_l^{ym}). \tag{3.1}$$

Here, f_{sky} is the fraction of the sky observed. The C_l values on the right-hand side of Equation 3.1 must include appropriate noise terms for auto-correlation power spectra. We propagate the forecasted C_l measurement accuracies into constraints on our cosmological parameters using the Fisher matrix formalism. We use the temperature anisotropy power spectrum and noise from the Planck Legacy Archive (2018 data release) [5]. At low l, the noise values are not symmetric. We take the larger values in each case. We use the temperature power spectrum in the range $2 \leq l \leq 2500$.

We forecast polarisation constraints for future surveys. We generate theoretical unlensed auto-correlation spectra C_l^{EE} using the Boltzmann code CLASS [9]. The noise term for the polarisation auto-correlation spectra is calculated as:

$$N_l^{-1} = \sum_i \Delta P_i^{-2} \exp\left[-l(l+1)\theta_i^2/8\ln 2\right], \tag{3.2}$$

where i indexes the frequency band, ΔP_i is measured in μK-arcmin and θ_i is the FWHM beam size in arcmin. We calculate the polarisation power spectra for l values of 30-2500.

Finally, we include the forecasted cross-correlation between the Planck temperature power spectrum and future polarisation power spectrum measurements (C_l^{TE}). We calculate the covariance using the existing temperature power spectrum and noise and the theoretical E-mode polarisation power spectrum and noise outlined above. We extract the noise for Planck for insertion into Equation 3.1 from the published variance values in the Planck Legacy Archive assuming $f_{\text{sky}} = 0.5$. In our calculations, we also include a prior of $\sigma(\tau) = 0.008$ (as quoted for TT,TE,EE+lowE in [5]).

When taking the derivatives of the theoretical CMB temperature and polarisation power spectrum with respect to the cosmological parameters, we do *not* include the effects of CMB lensing on the power spectrum. It is important to note that we take this approach only because we want to separate clearly the contributions of primary anisotropies and lensing to the constraints. For other purposes, it would make more sense to use the lensed TT, TE and EE spectra, as these are what are measured by real experiments, and delensing is very difficult in practice. It has been shown that cross-correlation terms between lensed temperature and polarisation spectra and the lensing power spectrum itself contribute negligibly to the covariance matrix [50, 51], so one would not need to worry about double-counting information when taking this approach.

For lensing forecasts, we require the CMB lensing convergence power spectrum $C_l^{\kappa\kappa}$. When two surveys cover a shared area of the sky, we can also use the cross-correlation between galaxy positions and the convergence map as an additional information source. This requires the angular galaxy clustering power spectrum $C_l^{g_i g_i}$ and the cross power spectrum $C_l^{g_i \kappa}$ (i indexes a specific redshift bin of the galaxy survey). All of these spectra can be derived from matter power spectra $P_m(k, z)$ (see e.g. [52]) generated using CLASS.

We calculate the galaxy power spectra for different redshift bins i making the approximation that all galaxies are at the mean redshift z_i (i.e. assuming a Dirac delta-function distribution). The Limber approximation fails for thin redshift distributions, so we use the exact equation

$$C_l^{g_i g_i} = \frac{2}{\pi} \int dk k^2 b_i^2 P_{bc}(z_i, k) j_l^2 [k d_A(z_i)]. \tag{3.3}$$

In the previous article, we assumed a maximum wavenumber in each redshift bin of $k_{\mathrm{max}} = 0.2 \ h \ \mathrm{Mpc}^{-1}$. For consistency, we calculate the angular galaxy power spectra in these calculations up to the corresponding appropriate l_{max} value, by converting k_{max} into units of Mpc^{-1} and then multiplying by the comoving angular diameter distance at the given mean redshift, $d_A(z) = D_A(z) \times (1 + z)$. We calculate the appropriate l_{min} value likewise, basing the value on the survey area.

The factor $b_i^2 P_{bc}(z_i, k)$ corresponds to the three-dimensional galaxy power spectrum, $P_{gg}(z_i, k)$. As in our previous work, we assume linear galaxy bias. The subscript bc emphasises that massive neutrinos do not contribute to the galaxy power spectrum but only cold dark matter and baryons. b_i is the fiducial linear galaxy bias and j_l is the spherical Bessel function.

For the convergence power spectrum, we use the Limber approximation [17]:

$$C_l^{\kappa\kappa} = \left(\frac{4\pi G \rho_{m,0}}{c^2} \right)^2 \int_0^{z\star} dz (1+z)^2 \left(\frac{d_A(z, z\star)}{d_A(z\star)} \right)^2 \frac{P_{mm} \left[k = \frac{l+1/2}{d_A(z)}, z \right]}{H(z)}. \tag{3.4}$$

$\rho_{m,0}$ is the comoving matter density, $z\star$ is the redshift of last scattering, and $d_A(z, z\star)$ represents the comoving angular diameter distance between the two redshifts. Accounting for curvature, this is calculated as [53]:

$$d_A(z_1, z_2) = f_{k,2}(\chi) \sqrt{1 + \Omega_k \left(\frac{f_{k,1}(\chi)}{D_H} \right)^2} + f_{k,1}(\chi) \sqrt{1 + \Omega_k \left(\frac{f_{k,2}(\chi)}{D_H} \right)^2}, \tag{3.5}$$

$$
\begin{aligned}
f_k(\chi) &= \frac{1}{\sqrt{k}} \sinh\left(\chi \sqrt{k} \right) & \Omega_k &> 0, \\
&= \chi & \Omega_k &= 0, \\
&= \frac{1}{\sqrt{k}} \sin\left(\chi \sqrt{k} \right) & \Omega_k &< 0,
\end{aligned}
\tag{3.6}
$$

where $k = -\Omega_k (H_0/c)^2$ and $D_H = (c/H_0)$ is the Hubble distance.

In this case, the requirement to evaluate $P_m(z)$ for such a large number of redshifts presents somewhat of an inconvenience. As massive neutrinos change the shape of the matter power spectrum over time, it is not sufficient to simply multiply $P_m(0)$ by a scale-independent growth factor $D(z)^2$ at each instance. We output $P_m(k)$ at a large number of redshifts using `CLASS`, and interpolate the table at the necessary redshifts.

Figure 3.1: Example convergence, galaxy and cross-power spectra for Euclid and the Simon's Observatory Large Aperture Telescope (LAT) with corresponding noise. The galaxy power spectrum corresponds to the central redshift bin of the Euclid survey (z=1.35) with a $b = 1.51$.

We also use the Limber approximation to calculate the galaxy-convergence cross-power spectrum [52]:

$$C_l^{g_i\kappa} = \left(\frac{4\pi G\rho_{m,0}}{c^2}\right)(1+z_i)\frac{d_A(z_i, z_\star)}{d_A(z_i)d_A(z_\star)}b_i P_{cb,m}\left[k = \frac{l+1/2}{d_A(z_i)}, z_i\right]. \tag{3.7}$$

Here the factor $b_i P_{bc,m}$ corresponds to the matter-galaxy cross-power spectrum.

In the case of the CMB lensing auto power spectrum, the corresponding noise is that associated with the reconstruction of the CMB lensing potential from CMB observations, which is calculated using the algorithm provided by Okamoto & Hu [54] by interfacing the Fortran module `FUTURCMB` (provided by [55]) with our code. This must be rescaled by a factor of $\frac{1}{4}[(l+2)!/(l-2)!]$ for use with the convergence κ (as opposed to the lensing potential ϕ) power spectrum. For the galaxy power spectra, the shot noise term is given by the inverse of the surface density of the galaxies in the particular redshift bin in steradians, n_g^{-1}.

We do not include cross-spectra between redshift bins. Following [56], we assume that the covariance between P_{gg} and $C_l^{g\kappa}$ can be neglected, and we can therefore do the Fisher matrix calculations for the two-dimensional and three-dimensional power spectra

calculations separately and simply add the output Fisher matrices, i.e.:

$$F = F(C_l^{TT}, C_l^{TE}, C_l^{EE}) + F(P_{gg}(k,z)) + F(C_l^{\kappa\kappa}) + F(C_l^{g\kappa}). \tag{3.8}$$

Because we treat CMB lensing with a separate Fisher matrix for analysis purposes, the cross-correlation between the temperature and polarisation power spectra and the lensing power spectrum is not included. We also ran tests in which $C_l^{T\kappa}$ and $C_l^{E\kappa}$ were included, and found the change in the constraints to be less than 1% in all cases.

3.3 Results and Discussion

3.3.1 Survey Data

In our galaxy clustering paper, we focused on constraints from Euclid. Here we focus on forecasts from the combination of Euclid and Simons Observatory (see Section 1.7.3) with existing information from Planck, in the forms of the CMB temperature power spectrum and a prior on τ of $\sigma(\tau) = 0.008$. For the galaxy-CMB lensing cross-correlation, we assume maximum overlap between Euclid and Simons Observatory. We present constraints for a wider range of survey combinations in Appendix B.

3.3.2 Results from the CMB Alone

Figure 3.2 shows the forecasted constraints on M_ν from Planck and Simons Observatory alone for various cosmologies. The constraints from temperature and polarisation alone are relatively weak but also insensitive to changes in curvature or the dark energy equation of state. Most of the information on M_ν comes from unlensed temperature anisotropy information, with the unlensed E-mode polarisation mostly improving constraints through tightening the constraints on other parameters rather than being directly sensitive to M_ν.

Adding CMB lensing improves the constraints significantly. To determine how much the free-streaming effect on $C_l^{\kappa\kappa}$ contributes to the constraints, panel (b) shows the constraints from the shape of $C_l^{\kappa\kappa}$ alone (changes in the overall amplitude of the power spectrum in the derivatives are neglected). Panel (c) shows the constraints when the full $C_l^{\kappa\kappa}$ is used. Both the dark energy equation of state and curvature parameters add a scale-dependent effect that is somewhat degenerate with the neutrino effect in panel (b). However, because Ω_k is quite well constrained from temperature and polarisation data, and the dark energy equation of state parameters are not, freeing w has a much more significant effect on the M_ν constraint.

When full lensing information is used in panel (c), the constraints on M_ν are much improved. However, here the curvature parameter degrades M_ν more so than w. This is because a very small change in Ω_k produces a much larger change in the amplitude of $C_l^{\kappa\kappa}$ than w. The results are now strongly cosmology-dependent.

Adding CMB lensing makes M_ν correlated with A_s, and therefore with τ (as A_s and τ are strongly degenerate in CMB measurements) unless we have a cosmic-variance-limited
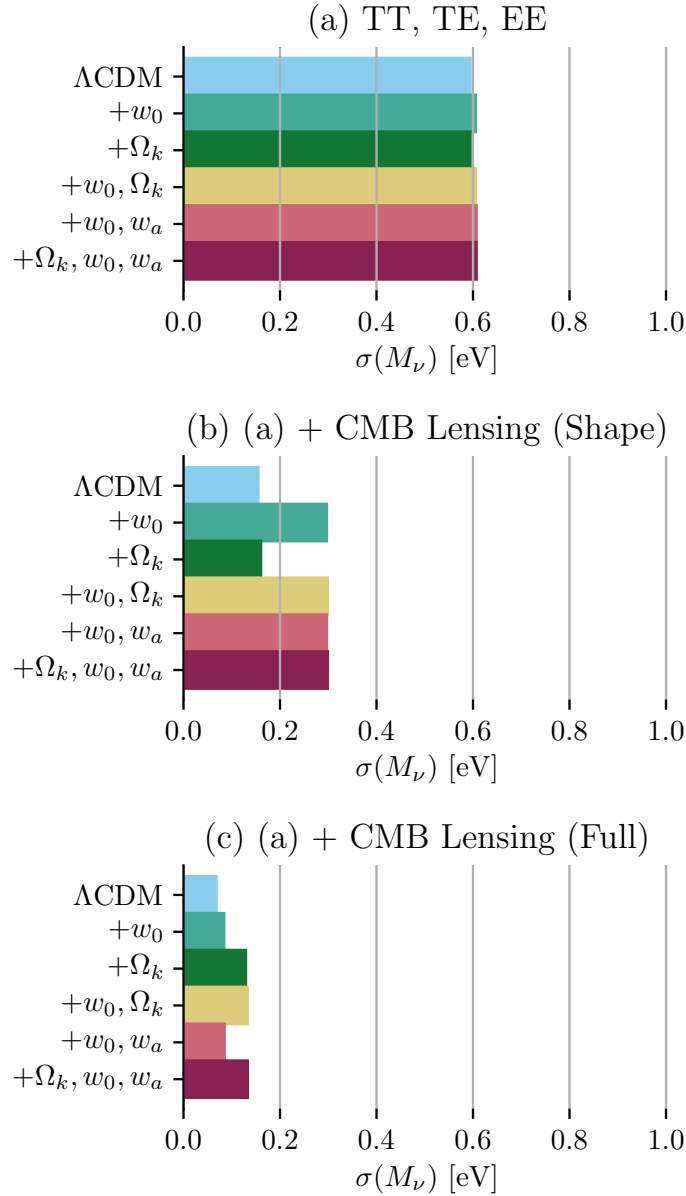
Figure 3.2: A breakdown of forecasted constraints from the CMB alone. 1σ constraints on M_ν are shown along the x-axes for the cosmological models shown on the y-axes. Panel (a) shows the constraints on M_ν that we derive using the current unlensed Planck temperature power spectrum ($2 \leq l \leq 2500$), forecasted E-mode polarisation information from Simons Observatory ($30 \leq l \leq 2500$) and their cross-correlation. Panel (b) adds constraints from the shape of the convergence power spectrum only (to capture the free-streaming effect). Panel (c) replaces this with full forecasted CMB lensing constraints.

measurement of E-mode polarisation at $l \leq 30$. This is analogous to what happened when the CMB prior was combined with galaxy power spectrum measurements in our previous work [25]. However, in the case of Figure 3.2, our constraint on τ from Planck is already quite strong. Fixing τ causes only a minor improvement in the ΛCDM case in panel (c). For the other models, the constraints on the additional extension parameters remain the limiting factor.

3.3.3 Full Galaxy Power Spectra

Figure 3.3 presents forecasted constraints from the Euclid full galaxy power spectrum and Planck/Simons Observatory both with and without CMB lensing information included for a range of cosmologies. Note that the x-axis scale is reduced by a factor of 10 compared to that in Figure 3.2. The cross-hatched bars are the results for a 1σ prior on τ of 0.008, and the solid bars are the results if τ is fixed.

We first examine the ΛCDM case, with the weaker constraint on τ. As emphasised in previous work [25], for a powerful galaxy survey like Euclid, if the full galaxy power spectrum is used with CMB information, the constraints on M_ν come to be limited by the weak constraints on τ. The cause of this is that the effects of M_ν and A_s on the galaxy power spectrum are strongly degenerate, and A_s and τ are measured in combination from the CMB. CMB lensing does not provide any additional direct information on τ, but does help constrain A_s better. This leads to the very modest improvement in $\sigma(M_\nu)$ seen in Figure 3.3 when comparing panels (a) and (b). Because of our strong τ prior, the relative gain from adding CMB lensing is relatively small (about 5%). The scale-dependent effect of M_ν on the CMB lensing power spectrum is sub-dominant here. Panel (c) shows that the gain when adding cross-correlation information between galaxy positions and the lensing map is more significant, particularly for the more complex models. This gain seems to come primarily from the scale-dependence in the galaxy-CMB cross-power spectrum imprinted by M_ν, and therefore reduces the cosmological dependence somewhat as well.

The solid bars show what can be achieved if τ is perfectly constrained. Adding CMB lensing when τ is already well constrained provides very little benefit, because A_s is already well constrained, though there is some mild improvement for the more complex models. Galaxy-CMB lensing helps very little here because the constraints are already quite strong.

It is significant to note that in all cases, the constraints on M_ν still depend quite heavily on the cosmological model assumed. CMB lensing does contribute to tightening constraints on curvature [57], but the dark energy equation of state can degrade constraints on M_ν considerably, particularly when the constraint on τ reaches its limit.

Neglecting galaxy-CMB lensing, the relationship between the ΛCDM constraints in the panels in Figure 3.3 can be understood completely in terms of the degeneracy between M_ν, A_s and τ. Imposing a particular relative improvement in the constraints on one of these parameters leads to an almost equal relative improvement in the constraints on the other two. Figure 3.4 shows this relationship in contour form. While CMB lensing is useful for improving the constraints on these three parameters, it is less powerful than the Planck τ prior we include.
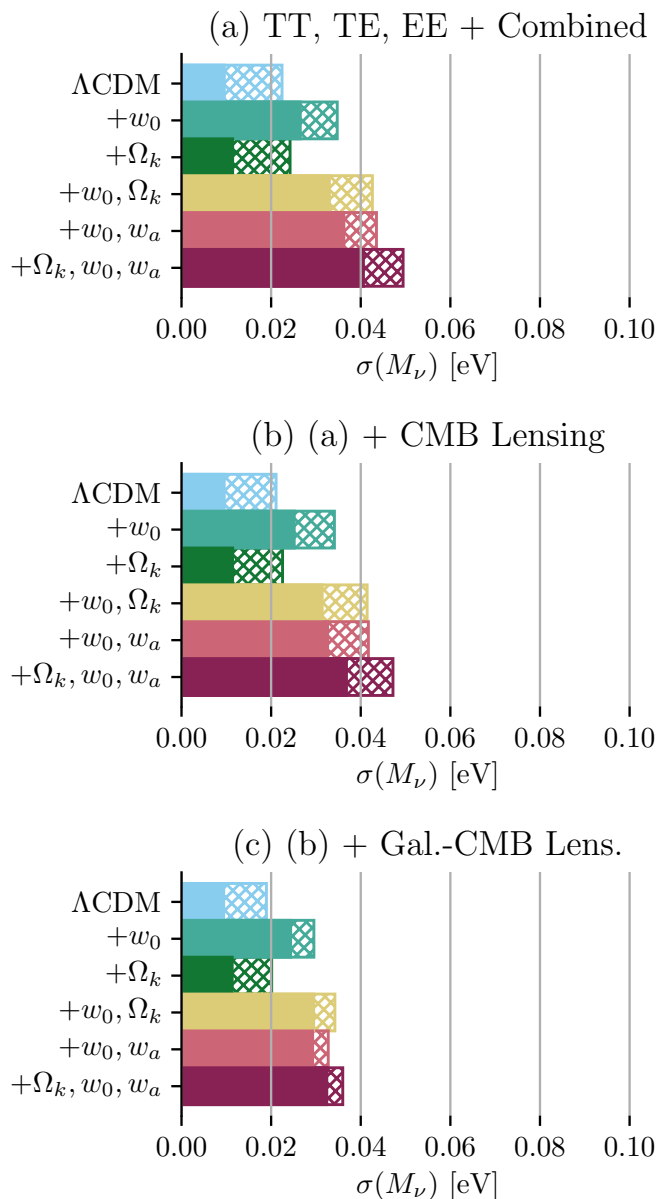
Figure 3.3: Forecasted constraints on M_ν for a combination of Simons Observatory/Planck and Euclid. The full (broadband) galaxy power spectra are used to generate the Euclid Fisher matrix. Panel (a) shows the constraints without any lensing information, panel (b) full CMB lensing information and panel (c) further adds the cross-correlation between galaxy positions in Euclid and the CMB lensing map (assuming maximum overlap between the two surveys). In each case, the cross-hatched bars represent the results with a 1σ prior on τ of 0.008, and the solid bars represent the results when τ is fixed.
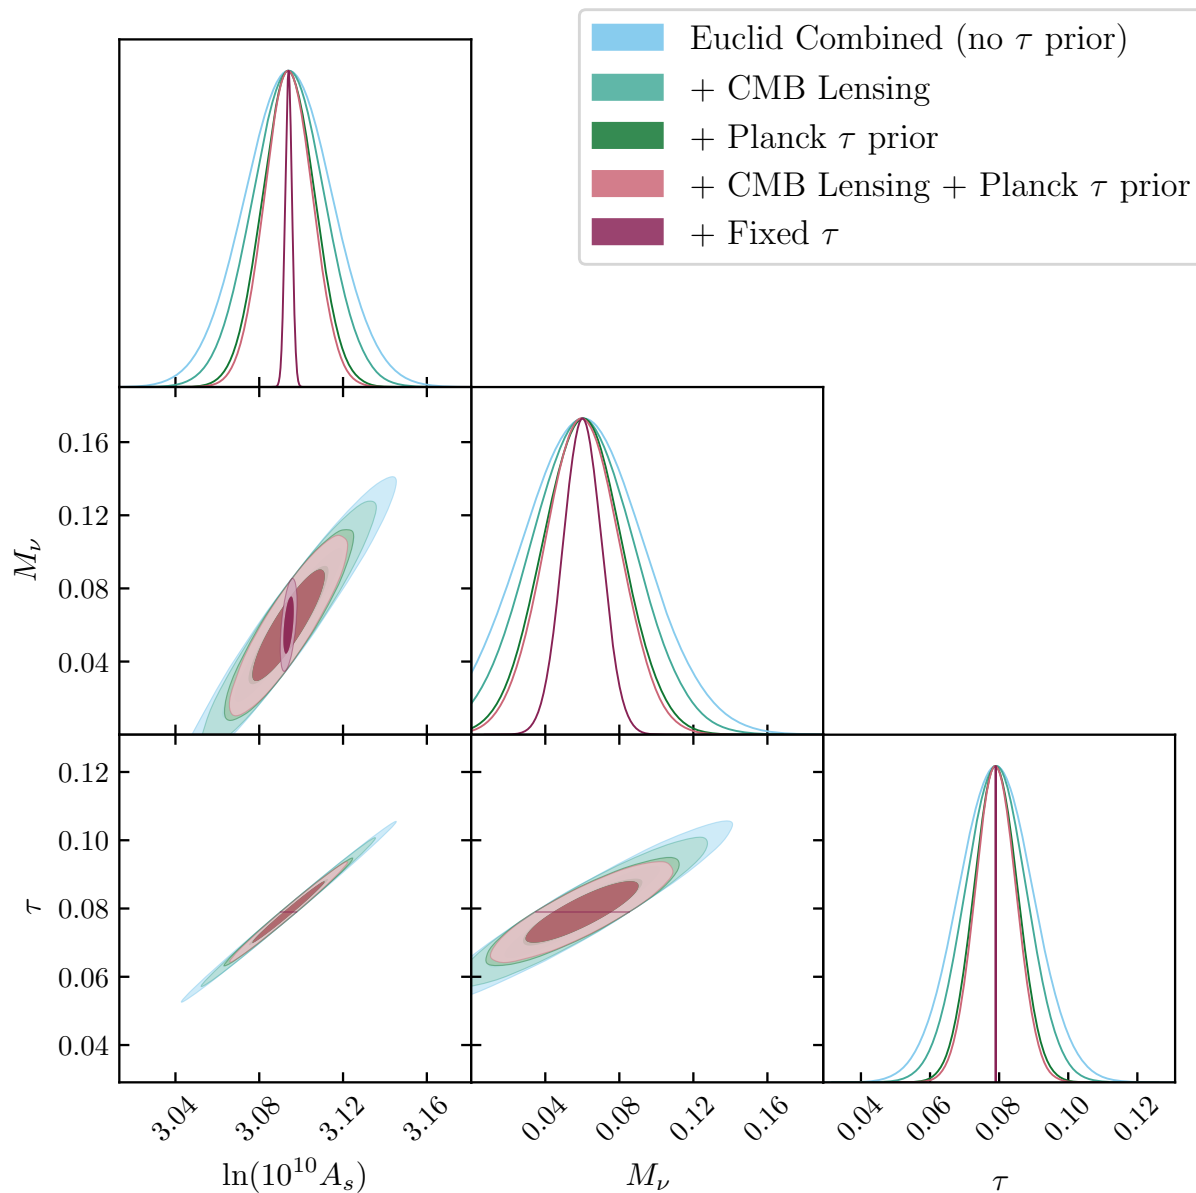
Figure 3.4: The relationship between the constraints on M_ν, A_s and τ from Euclid (marginalised over the other parameters in the text) with and without CMB lensing for various τ priors. The errors on these three parameters are very strongly correlated, to the extent that a change in the error on one of these parameters leads to an equal relative error change on the other two parameters. CMB lensing can help improve these constraints when τ is weakly constrained, but has much less impact when a τ prior from Planck is included.

3.3.4   BAO-Only Information

The combination of BAO and CMB lensing data is a common focus in forecasts [48, 26, 21]. As can be seen in Figure 3.5, the constraints from Euclid from BAOs alone are much weaker than those from the full power spectrum (note the x axis scale has been increased by a factor of 5 relative to Figure 3.3), but a better relative improvement is obtained by combining with CMB lensing information, particularly for the more complex models. As was highlighted in Chapter 2, there is a strong degeneracy between M_ν and Ω_k in their effects on cosmological distance parameters, as is clear from panel (a). In Chapter 2, there was very little interaction between M_ν and the dark energy equation of state parameters in the BAO-only case. However, we can see from panel (a) that there is a significant degeneracy between w_0/w_a and M_ν here. This is a result of using the Planck temperature power spectrum and Simons Observatory polarisation forecasts instead of the compressed likelihood prior used in Chapter 2.

We first examine the case without CMB lensing. BAO information does not constrain A_s or τ, but CMB lensing does. Because of the lack of information on A_s, the degeneracy between M_ν and A_s (and therefore τ) that arises in the combined case does not arise here. Therefore, looking at panel (a) of Figure 3.5, one can see that fixing τ makes no difference to the neutrino mass constraint.

Once CMB lensing is added, the τ degeneracy is re-established to a degree (see panels (b) and (c)). CMB lensing contributes in multiple ways here, primarily by adding structure growth information, including the scale-dependent suppression of the CMB lensing power spectrum by the neutrinos. Because the BAO-only constraints are sufficiently weak, CMB lensing can contribute significantly to improving constraints.

CMB lensing also tightens the constraints on curvature significantly (by a factor of 2 in the ΛCDM + Ω_k case), breaking the M_ν-Ω_k degeneracy in panel (a). CMB lensing does not contribute significantly to constraining the dark energy equation of state parameters, but the suppression in the power spectrum caused by massive neutrinos helps reduce the degeneracy between the two.

Galaxy-CMB lensing also helps improve constraints significantly, as can be seen in panel (c). However, although the scales in the plot make it difficult to see, the constraints are still relatively cosmology-dependent, degrading by about 40% for a free dark energy equation of state.

3.3.5   Free-Streaming Information

Figure 3.6 shows how CMB lensing affects our 'free-streaming' constraints. Massive neutrinos suppress the growth of structure on small scales to a degree that is primarily dependent on the total neutrino mass. This creates a small but distinctive scale-dependent signature in both the matter power spectrum and in the structure growth rate f (which can be measured independently using redshift-space distortions). If measured, the magnitude of the suppression can be used to obtain a cosmology-independent and τ-independent probe of the neutrino mass, as discussed extensively in Chapter 2. Our 'free-streaming' constraint

Figure 3.5: As for Figure 3.3, but with the constraining power from Euclid provided by constraints on $H(z)$ and $D_A(z)$ derived from the BAO wiggles. Note that the x axis here has been extended by a factor of 5 compared to Figure 3.3.
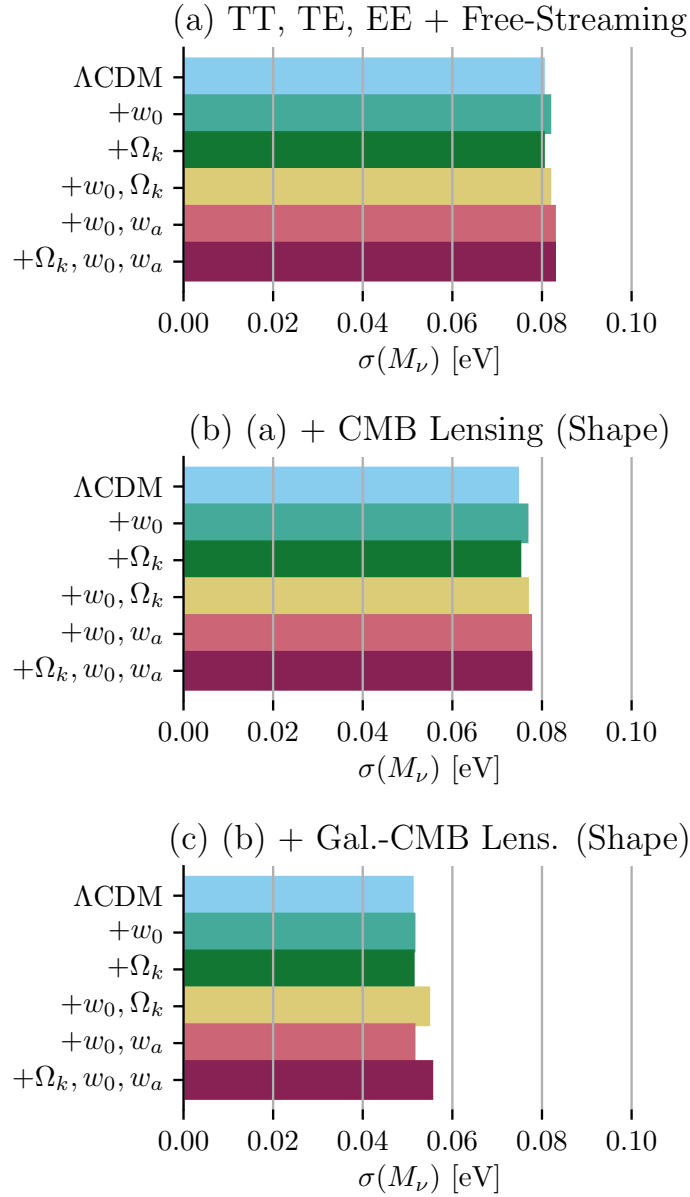
Figure 3.6: As for Figure 3.3, but with the constraining power from Euclid only deriving from the scale-dependence of the power spectrum and structure growth rate, using the method developed in [25]. In this case, the CMB lensing and galaxy-CMB lensing constraints are also derived only from the scale-dependence of the derivatives.

forecasts are calculated by isolating the scale-dependence in the matter power spectrum as the observable in our Fisher matrix, and marginalising over the overall amplitude, and doing likewise for the structure growth rate.

This is the first time that we present the free-streaming constraints in combination with a full CMB temperature and polarisation forecast. We see that the gains are significant over the Planck compressed-likelihood prior used previously, and that the final constraints remain effectively cosmology-independent (see panel (a)).

In panel (b), we add CMB lensing. Here we isolate the suppression in the convergence power spectrum alone and neglecting all amplitude information, and this improves constraints by about 7% in all cases. Although we showed in Section 3.3.2 that the free-streaming signature in the CMB lensing power spectrum can provide meaningful constraints, the constraints from the galaxy power spectrum are much stronger, so the relative improvement when adding the shape of the CMB lensing power spectrum is somewhat small.

However, in panel (c) we add galaxy-CMB lensing (excluding amplitude information as before, to extract the scale-dependence of the measurement only), and find that the constraints are much improved, while still remaining quite cosmology-independent. We can therefore conclude that by isolating the scale-dependent effect of free-streaming in the galaxy power spectrum, CMB lensing power spectrum, and galaxy-CMB lensing cross-power spectrum, one can extract robust, cosmology-independent (and τ-independent) constraints on the neutrino mass of approximately 0.05 eV when combining Simons Observatory and Euclid. We consider this measurement approach to be the most robust.

3.3.6 Comparisons with Previous Work

It is difficult to do direct comparisons between various forecasts of this type in the literature because of the many different assumptions that can be made, from survey choice to error management to cut-off scales. However, comparisons with some recent works can provide some reinforcement for the results provided here, particularly to manage scepticism about the Fisher matrix methodology, and can also provide some interesting insights. Here we examine how our full galaxy clustering constraints compare to those in the literature.

The authors of [58] provide some MCMC forecasts for Euclid galaxy clustering assuming ΛCDM+$M_\nu + N_{\text{eff}}$. They take Planck as their CMB survey and also include cosmic shear, but their results are very close to ours (21 meV in our case vs. 24 meV or 27 meV in their realistic and conservative cases, respectively). This could further support our conclusion here that galaxy clustering information really is dominant over CMB and lensing information (in this case, cosmic shear).

In [59], table 4 gives an uncertainty on M_ν of 17 meV for CMB-S4 and Euclid in the ΛCDM+$M_\nu + N_{\text{eff}}$ case. Our corresponding value is 21 meV. Their final constraint on τ is also stronger than ours, so correcting for that leaves our constraint at 17 meV. The fact that these values are identical despite the fact that we assume different CMB surveys highlights that future CMB surveys may not offer a huge improvement over Planck to M_ν constraints when combined with powerful measurements of the full redshift-space power

spectrum, other than by improving the constraint on τ.

3.4 Conclusions

This chapter represents a continuation of our work in Chapter 2. While there we focused on spectroscopic galaxy surveys alone, here we have expanded our analysis to include the full power of planned CMB experiments. We consider three possible methods of constraining M_ν from galaxy surveys and the effects of CMB lensing on each: the use of the full redshift-space galaxy power spectrum, the use of BAOs alone to infer distance constraints, and the use of the signatures of neutrino free-streaming only (a newly-developed method from Chapter 2).

Overall, we have shown that CMB lensing measurements are a much less powerful probe of the neutrino mass than large-scale structure surveys for the scales considered. When combined with the full galaxy power spectrum from a spectroscopic galaxy survey like Euclid, CMB lensing contributes to constraints on M_ν primarily by tightening constraints on A_s, as these two parameters are very strongly correlated in the measured galaxy power spectrum. This correlation is also the source of the M_ν-τ degeneracy that limits M_ν constraints when CMB and galaxy clustering measurements are combined. However, if τ is already well constrained, the primary gain from adding CMB lensing becomes redundant. Adding galaxy-CMB lensing can help improve neutrino mass constraints further, but requires significant overlap between the galaxy clustering and CMB surveys being considered.

CMB lensing primarily contributes to BAO-only constraints on M_ν by improving constraints on curvature and by adding information on structure growth. As we emphasised previously, M_ν and Ω_k are highly degenerate in their effects on distance parameters at low redshifts. This means allowing for a very small non-zero curvature can degrade the constraints on M_ν by several factors. CMB lensing can help mitigate this effect.

Finally, we look at the cosmology-independent free-streaming-only constraining method we developed in [25]. Combining the free-streaming constraints from a survey like Euclid with a full CMB temperature/polarisation forecast improves the constraints significantly while still keeping them cosmology-independent. The gains from including the shape of the CMB lensing power spectrum are small (see Figure 3.6). Although the free-streaming signal in the CMB lensing power spectrum can significantly improve CMB-only constraints, it is much weaker than the corresponding signal in the galaxy power spectrum. The constraining power of galaxy-CMB lensing is much more useful.

In combinations of galaxy clustering and CMB lensing measurements, the galaxy power spectrum is a much more powerful probe of M_ν. As τ is better constrained, the information provided by CMB lensing will become redundant in neutrino mass constraints. The constraints provided by the full galaxy power spectrum will also become increasingly cosmology-dependent as τ becomes better known. BAO-only constraints become more robust when combined with CMB lensing but waste a lot of valuable information. The constraints extracted through the effects of free-streaming on the power spectrum, on the

other hand, are cosmology-independent and independent of τ. They are also stronger than those from BAO measurements and CMB lensing combined for almost all models. Using the current Planck temperature power spectrum, the forecasted Simons Observatory E-mode polarisation spectrum and the free-streaming signals extracted from the Euclid galaxy power spectrum, reliable 1σ constraints on M_ν of approximately 0.08 eV can be achieved. This can be improved to 0.075 eV if CMB lensing is included, or 0.05 eV if galaxy-CMB lensing is included.

Chapter 4

Beyond Linearity

4.1 Introduction

A consistent caveat to our results in Chapters 2 and 3 was that we worked only with the linear galaxy and CMB lensing power spectra, a treatment which becomes increasingly invalid on small scales and at late times. We included scales no smaller than $k_{\mathrm{max}} = 0.2\ h\ \mathrm{Mpc}^{-1}$, as is a popular choice in the literature, but non-linear effects can become significant on larger scales than this (see Figure 4.1).

Major neutrino mass constraint forecast papers in the literature have focused on the linear power spectrum for galaxy redshift surveys (e.g. [59] [32]). They often make use of a gradual removal of information on small scales instead of a sharp cut-off as we do, but the results are still reasonably consistent. In this chapter, we take the next step and extend to the next-to-leading-order or one-loop power spectrum. The aim is to see whether the broad qualitative principles extracted in the previous chapters still hold true, and to develop an impression of how much constraints can be expected to relax once non-linear effects are consistently accounted for. Understanding the impact of these effects will be vital to allow us to make full use of the wealth of information that will be provided by upcoming surveys.

4.2 Methodology

Our basic approach follows that of Chapters 2 and 3, and the reader can refer to those chapters for the fiducial model. We forecast constraints for spectroscopic galaxy clustering measurements from Euclid and CMB lensing measurements from Simons Observatory. We use the same prior for CMB anisotropy information as that used in Chapter 3 - a combination of the existing Planck temperature data, forecasted Simons Observatory E-mode polarisation data and the cross-correlation between the two. We also include a prior on τ in all cases of $\sigma(\tau) = 0.008$. CMB lensing is analysed separately in Section 4.3.4. We work once again with $k_{\mathrm{max}} = 0.2\ h\ \mathrm{Mpc}^{-1}$.

Figure 4.1: A comparison of the fiducial linear and NLO galaxy power spectra at $z = 1.35$ (the central redshift bin for Euclid) with $\mu = 0$. The galaxy power spectrum is ultimately enhanced on small scales when moving to next-to-leading order.

4.2.1   Implementation of the NLO Power Spectra

In a recent paper, the authors of [19] presented the first complete expression for the perturbative next-to-leading-order (NLO) or 1-loop power spectrum in redshift space. Using the approach developed in [60], they were able to order the contributing terms in a consistent way and keep the contributions appropriately general. They also include the leading higher-derivative bias and velocity terms in their formalism. Their rigorous approach leads to a complex but complete final expression for the redshift-space power spectrum that makes minimal assumptions. We implement their formalism here. The authors emphasise the importance of accounting for line-of-sight-dependent selection effects in the calculation of the galaxy power spectrum unless physical arguments can be made for neglecting them. We neglect selection effects for most of this article, but justify this and provide some extended calculations in Section 4.4.2.

The NLO redshift-space power spectrum provided by [19] is calculated using a total of 28 independent loop integrals, and requires specification of 14 free bias parameters, 2 velocity bias parameters and 3 stochastic parameters (or 5 bias, 1 velocity bias and 3 stochastic parameters, if selection effects are neglected). The latter choice applies in most of this chapter. The next-to-leading-order galaxy power spectrum can be summarised in

m	0	1	2	3	4	5	6	7	8
p	0,2,4,6	1,3,5	0,2,4,6	1,3,5	2,4,6	3,5	4,6	5	6

Table 4.1: The (m, p) pairs summed over in Equation 4.5.

the following form in redshift space:

$$P_{gg,s}(k, \mu) = P_{gg,s}^{l+hd}(k, \mu) + P_{gg,s}^{2-2}(k, \mu) + 2P_{gg,s}^{1-3}(k, \mu), \tag{4.1}$$

where

$$\begin{aligned} P_{gg,s}^{l+hd}(k, \mu) = {}& [b_1 - b_\eta f(k)\mu^2]^2 P_L(k) + P_\epsilon^{\{0\}} \\ & - 2\{b_1 b_{\nabla^2\delta} - \mu^2 f(k) b_\eta [b_{\nabla^2\delta} + b_1 \beta_{\nabla^2 v} + b_1 \beta_{\delta_\parallel^2 v}\mu^2] \\ & + \mu^4 f(k)^2 b_\eta^2 [\beta_{\nabla^2 v} + \beta_{\delta_\parallel^2 v}\mu^2]\}k^2 P_L(k) \\ & + k^2 P_\epsilon^{\{2\}} - \mu^2 k^2 P_{\epsilon\varepsilon_\eta}^{\{2\}}, \end{aligned} \tag{4.2}$$

$$P_{gg,s}^{2-2}(k, \mu) = \sum_{n=0}^{4} \sum_{(m,p)} A_{n(m,p)}(f(k), b_\mathcal{O})\mathcal{I}_{mp}(k)\mu^{2n}, \tag{4.3}$$

$$P_{gg,s}^{1-3}(k, \mu) = \sum_{l=0}^{3} \sum_{n=1}^{5} C_n^{1-3,l}(f(k), b_\mathcal{O})\mathcal{I}_n(k)(\mu)P_L(k)\mathcal{L}_{2l}. \tag{4.4}$$

The coefficients $A_{n(m,p)}$ and $C_n^{1-3,l}$ and their input bias parameters $b_\mathcal{O}$ can be obtained from a `MATHEMATICA` notebook provided as supplementary material associated with [19][1]. These coefficients significantly improve the calculation time by combining contributions to give the minimum number of integrals. We highlight again that we include the scale-dependence in $f(k)$ when performing our calculations.

The loop integrals $\mathcal{I}_{mp}$ are calculated as

$$\mathcal{I}_{mp}(k) = 2 \int_{\boldsymbol{p}} \frac{p^{p-2}k^{6-p}}{|\boldsymbol{k} - \boldsymbol{p}|^4}(\hat{\boldsymbol{k}} \cdot \hat{\boldsymbol{p}})^m P_L(p)P_L(|\boldsymbol{k} - \boldsymbol{p}|) - \frac{2\delta_{p6}}{m+1}\left(\int_{\boldsymbol{p}}[P_L(p)]^2\right). \tag{4.5}$$

The second term simply subtracts the constant low-k limit when $p = 6$. The 23 (m, p) pairs are given in Table 4.1.

Equation 4.5 presents difficulties when $\boldsymbol{k} \approx \boldsymbol{p}$. We handle this using a method provided by [61], where it is noted that equal contributions are made to P_{22} as $\boldsymbol{p} \to 0$ and $\boldsymbol{p} \to \boldsymbol{k}$.

---

[1]`https://github.com/djeong98/pkgs_supplement`

We can take advantage of this fact by dividing up the integral into ranges of $\boldsymbol{p}$, excluding the range where $\boldsymbol{p} \to \boldsymbol{k}$ and doubling the term for $\boldsymbol{p} \to 0$.

$$\int_{\boldsymbol{p}} = 2 \int_0^\epsilon dp \int_{-1}^1 d\mu + \int_\epsilon^{k-\epsilon} dp \int_{-1}^1 d\mu + \int_{k-\epsilon}^{k+\epsilon} dp \int_{-1}^{(k^2+p^2-\epsilon)/2kq} d\mu +$$
$$\int_{k+\epsilon}^{k_{max}} dp \int_{(k^2+p^2-k_{max}^2)/2kq}^1 d\mu \tag{4.6}$$

The five loop integrals in Equation 4.4 are more straightforward to calculate and take the form

$$\mathcal{I}_1(k) = k^2 \int_{\boldsymbol{p}} \frac{\boldsymbol{k} \cdot \boldsymbol{p}}{p^2 |\boldsymbol{k} - \boldsymbol{p}|^2} (1 - \mu_{\boldsymbol{k},\boldsymbol{p}}^2) P_L(p),$$

$$\mathcal{I}_2(k) = k^2 \int_{\boldsymbol{p}} \frac{p^2}{p^2 |\boldsymbol{k} - \boldsymbol{p}|^2} (1 - \mu_{\boldsymbol{k},\boldsymbol{p}}^2) P_L(p),$$

$$\mathcal{I}_3(k) = \int_{\boldsymbol{p}} \left[\frac{p^2}{|\boldsymbol{k} - \boldsymbol{p}|^2} (1 - \mu_{\boldsymbol{k},\boldsymbol{p}}^2) - \frac{2}{3} \right] P_L(p), \tag{4.7}$$

$$\mathcal{I}_4(k) = \int_{\boldsymbol{p}} \left[\frac{(\hat{\boldsymbol{k}} \cdot \hat{\boldsymbol{p}})^2}{|\boldsymbol{k} - \boldsymbol{p}|^2} (1 - \mu_{\boldsymbol{k},\boldsymbol{p}}^2) - \frac{2}{15} \right] P_L(p),$$

$$\mathcal{I}_5(k) = \int_{\boldsymbol{p}} \left[\frac{(\hat{\boldsymbol{k}} \cdot \hat{\boldsymbol{p}})^4}{p^2 |\boldsymbol{k} - \boldsymbol{p}|^2} (1 - \mu_{\boldsymbol{k},\boldsymbol{p}}^2) - \frac{2}{35} \right] P_L(p).$$

We note that in calculating the galaxy power spectrum, $P_L(k)$ always means $P_{L,bc}(k)$ and f is f_{bc}, for the reasons outlined in Section 2.4.4.

Even with the simplification given in Equation 4.6, the requirement to calculate all of the integrals needed for the Fisher matrix calculation in a reasonable timeframe presented some issues. Calculating numerical derivatives requires two calculations per cosmological parameter (we have 11) plus one fiducial calculation per redshift bin (of which we have 15), with a total of 28 loop integrals calculated in each case. For the $I_n(k)$ integrals required for $P_{gg,s}^{1-3}$, the python `scipy.integrate.quad` function was sufficiently efficient and accurate with `limit=10`. But for the $I_{mp}(k)$ integrals required for $P_{gg,s}^{2-2}$, a `FORTRAN` implementation had to be adopted, using a Romberg integration function `qromb` from Numerical Recipes [2].

Our calculations clearly now have a far greater number of nuisance parameters. We choose to exclude selection effects for most of this chapter (see our discussion of this in Section 4.4.2). $P_{gg,s}^{l+hd}$ then contains three free bias parameters (b_1, $b_{\nabla^2\delta}$ and $b_{\nabla^2 v}$) and two new stochastic terms ($P_\epsilon^{\{2\}}$ and $P_{\epsilon\varepsilon_\eta}^{\{0\}}$). $P_{gg,s}^{2-2}$ adds two additional free bias parameters, b_2 and b_{K^2}. Adding $P_{gg,s}^{1-3}$ requires $b_{\rm td}$. In total, therefore, we require values for 5 new bias parameters and 3 stochastic amplitudes.

---

[2]Available at `https://wwwmpa.mpa-garching.mpg.de/~komatsu/crl/list-of-routines.html`.

Parameter	Fiducial value
b_1	$b_1(z)D(z) = 0.76$
$b_{\nabla^2\delta}$	R_L^2
$b_{\nabla^2 v}$	R_L^2
b_2	$0.412 - 2.143b_1 + 0.929(b_1)^2 + 0.008(b_1)^3$
b_{K^2}	$-\frac{2}{7}(b_1 - 1)$
$b_{\rm td}$	$\frac{23}{42}(b_1 - 1)$
$P_\epsilon^{\{0\}}$	n_g^{-1}
$P_\epsilon^{\{2\}}$	0
$P_{\epsilon\varepsilon_\eta}^{\{0\}}$	0

Table 4.2: Fiducial values for the NLO bias and stochastic parameters required for calculation of the NLO galaxy power spectrum.

The fiducial values for all the free bias and stochastic parameters are given in Table 4.2. R_L represents the Lagrangian radii of haloes, calculated using typical halo masses at each redshift derived from the halo mass function of [62]. n_g is the galaxy number density. The choice for b_1 is consistent with the previous chapters and [25, 36]. The fitting formula for b_2 was taken from N-body simulations and provided by [63]. The equations for b_{K^2} and b_{td} were both derived from Lagrangian LIMD calculations [19].

We want to avoid placing any significant priors on the parameters in Table 4.2. However, because of the large number of weakly constrained parameters in some cases, some initial prior must be placed on these parameters to keep the Fisher matrix invertible. We choose these priors to be large enough that increasing them by an order of magnitude has no significant impact on the final neutrino mass constraints. We settled on a 1σ error of 10 for each of the bias parameters and 100,000 for each of the stochastic parameters. The choice not to place meaningful priors on these parameters is a conservative approach.

The CMB lensing power spectrum is calculated as in Equation 3.4, but this time using the NLO matter power spectrum, calculated as in Appendix C (Equations C.17-C.21) using a publicly available code[3]. Calculating the matter power spectrum at the large number of redshifts required by Equation 3.4 would be very time-consuming, but it is also improper to scale between different redshifts using the growth factor $D(z)$ because of the redshift-dependent shape of the power spectrum when massive neutrinos are included. To account for this, we calculate the NLO matter power spectra independently for $0 < z < 3$, but scale backwards using the growth factor beyond $z = 3$, where the suppressin due to free-streaming is less substantial.

4.2.2   Deconstructing Constraints on M_ν

Following our work in Chapters 2 and 3, we seek not only to calculate the best possible neutrino mass constraints achievable from future surveys, but also to deconstruct the ori-

---

[3]https://wwwmpa-garching.mpg.de/komatsu/CRL/powerspectrum/density3pt/

gins of these constraints. In Chapter 3 we focused on comparing the 'combined' constraints from the full galaxy power spectrum with the constraints from BAOs only and those from the scale-dependent free-streaming signature in the power spectrum induced by massive neutrinos. Here we repeat this approach, dividing our results into three classes - combined, BAO-only, and free-streaming.

The combined constraints are the most straightforward. As in Chapter 2, in all cases we create a covariance matrix for the galaxy power spectrum using Equation 2.1. The derivatives of $P_{gg,s}(k, \mu)$ (Equation 4.1) with respect to the cosmological parameters are then used to propagate the constraints on the power spectrum into constraints on the cosmological parameters. To include full distance information, the scales in the power spectrum must be re-scaled to account for the dependence of the observed scales on $H(z)$ and $D_A(z)$, as in Equation 2.3.

In the BAO-only case, as in Chapter 2, the galaxy power spectrum is rewritten as the sum of two components, $P_{\mathrm{gg,s}}(k) = P_{\mathrm{S}}(k) + P_{\mathrm{BAO}}(k)$, with the fitting being done using the method outlined in Section 2.2.1. $P_{\mathrm{BAO}}(k)$ is used as the observable, and a 2×2 $H(z)$-$D_A(z)$ Fisher matrix is then extracted and used to generate constraints on the fiducial cosmology.

The free-streaming case is now more complicated. We want to isolate the information contained in the scale-dependence of the matter power spectrum, $P_L(k)$, and the structure growth rate, $f(k)$, to isolate the suppression in the small scale power caused by massive neutrinos. Before calculating the components of the power spectrum (Equations 4.2 - 4.4) to calculate our derivatives, we renormalise $P_m(k)$ and $f(k)$ to have the same amplitude as in the fiducial case at a scale of $0.005~h~\mathrm{Mpc}^{-1}$. Therefore, only the scale-dependence is varied.

4.3   Results

4.3.1   Full Galaxy Power Spectra

Figure 4.3 shows the forecasted constraints on M_ν for Euclid using the full galaxy power spectrum as a constraining tool. The cross-hatched bars show the constraints with $\sigma(\tau) = 0.008$, and the solid bars show the constraints if τ is held fixed. Looking at the former case first, we see as before that the constraints are weakened on extending to the NLO power spectrum. As emphasised in previous chapters, these combined constraints are strongly cosmology-dependent, and constraints assuming flatness or a simple cosmological constant should therefore be presented cautiously both in forecasts and when derived from real data. Here we see that the dependence of the constraints on the cosmology is even more stark when using the NLO power spectrum.

Figures 4.4 and 4.5 show the log derivatives of $P_{gg,s}$ with respect to the 9 nuisance parameters. We see that quite a number of these parameters change the power on small scales, and would therefore lessen the information on M_ν available from its characteristic suppression of the power spectrum on these scales. Extending to the NLO power spectrum
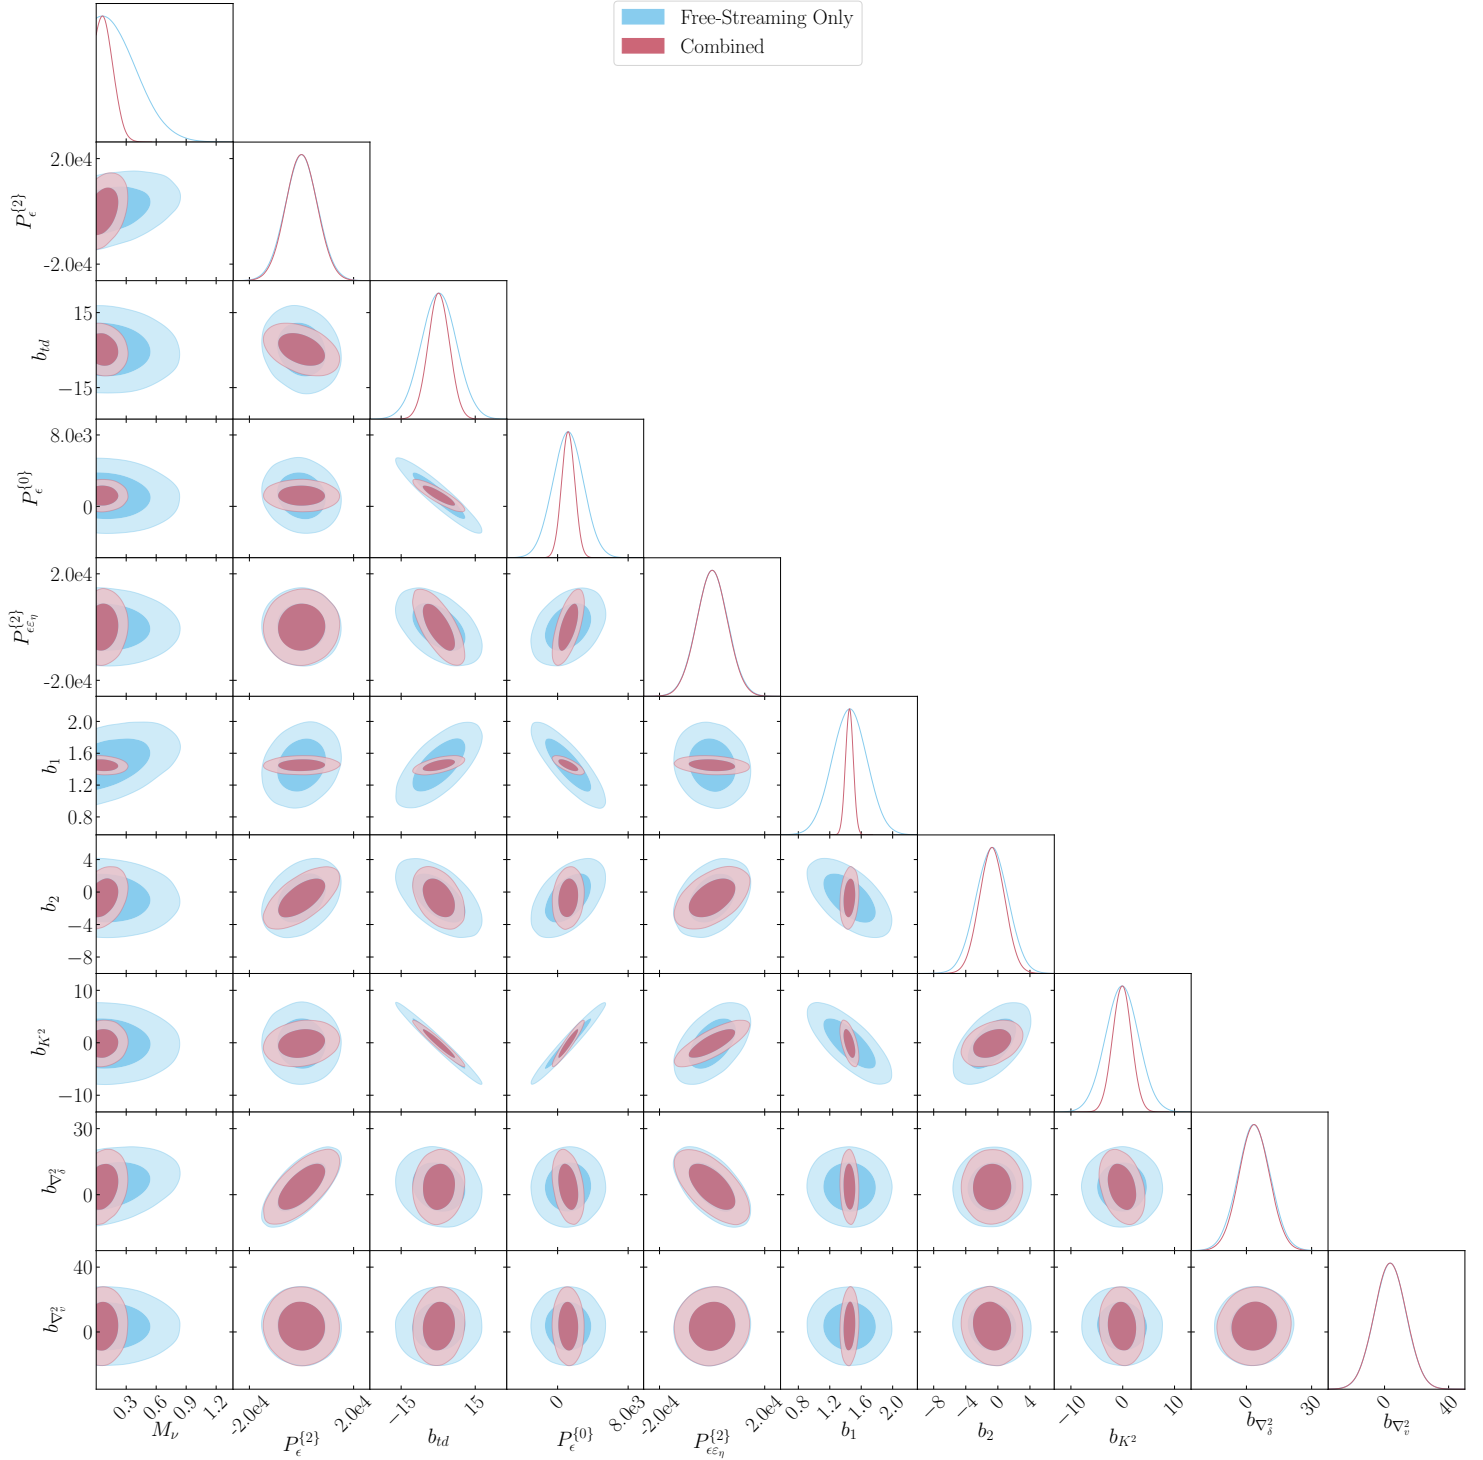
Figure 4.2: Contours showing the relationship between the different nuisance parameters and M_ν for the central redshift bin of our Euclid model plus a CMB prior. The pink contours represent the combined case, in which the full redshift-space galaxy power spectrum is used, and the blue contours represent the constraints when only the free-streaming signature in the power spectrum is used.

(a) TT, TE, EE + Combined (LO)

(b) TT, TE, EE + Combined (NLO)
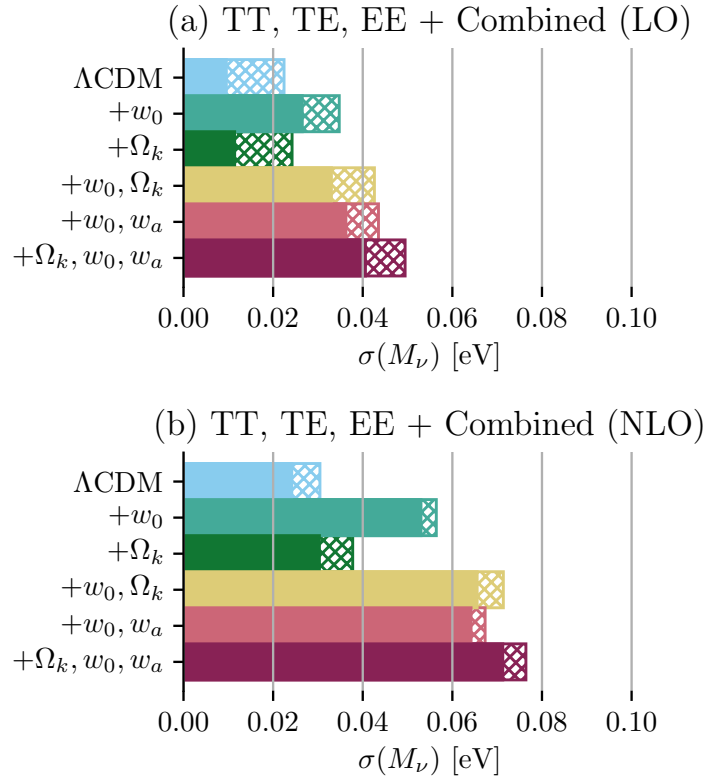
Figure 4.3: Constraint forecasts on M_ν for Euclid using the full galaxy power spectrum. The solid bars shows the constraints with τ fixed and the cross-hatched bars the constraints with $\sigma(\tau)=0.008$.

but fixing the new parameters recovers almost the same constraints as in the linear power spectrum case (0.023 eV vs. 0.022 eV) for ΛCDM, and when the new parameters are left free to vary the constraint weakens to 0.03 eV. However, on examining how the different parameters interact to produce this change, it becomes clear that there is no particular parameter that can specifically account for it. As might be expected from Figures 4.4 and 4.5, many of the nuisance parameters are quite strongly correlated with each other, and many of the new parameters show some degree of degeneracy with M_ν (see the pink contours in Figure 4.2). Nevertheless, we highlight the following interesting findings:

- In the linear case, imposing a prior on b_1 has the potential to significantly improve constraints on M_ν - if b_1 is fixed in each redshift bin, the M_ν constraint jumps from 0.022 eV to 0.16 eV. In the non-linear case, fixing b_1 improves the constraint only from 0.030 eV to 0.029 eV. The constraint on b_1 is weakened by more than a factor of 3 in a given redshift bin compared to in the linear case, but this still means it is relatively quite well-constrained compared to the other bias parameters. Its effect on the constraint on M_ν is therefore minimal. Fixing b_2 instead, for example, improves

Figure 4.4: The derivatives of the galaxy power spectrum with respect to the six non-linear bias parameters. One can see that there is some degeneracy in the effects of the different bias parameters as a function of scale. The derivative with respect to b_1 has had a factor of 1 subtracted for the readability of the plot.

the constraint to 0.028 eV.

- The nuisance parameter most strongly correlated with M_ν in this case is the stochastic parameter $P_\epsilon^{\{2\}}$. We note that adopting a model for the Finger-of-God effect could lead to a decent prior on $P_\epsilon^{\{2\}}$. The higher derivative bias parameter $b_{\nabla^2\delta}$ is also correlated with M_ν, and one can see from Figure 4.2 that $P_\epsilon^{\{2\}}$ and $b_{\nabla^2\delta}$ are also quite strongly correlated with each other, likely because both scale as k^2.

- The two bias parameters associated with operators derived from the tidal field, b_{K^2} and b_{td}, are very strongly anti-correlated with each other, as is clear from Figure 4.2. We discuss the reason for this strong degeneracy in Section 4.4.1. M_ν shows some correlation with both parameters. However, because of the strong degeneracy between the parameters, fixing one without the other has very little impact on the neutrino mass constraint, as the other can compensate for its effect. In this way, these two bias parameters behave like a single effective parameter in this case. Although constraining either does not improve the constraints on M_ν at all, imposing a prior on both leads to some improvement. $P_\epsilon^{\{0\}}$ is also strongly correlated with b_{K^2}.

In all, it is not unpromising that leaving all of the bias and stochastic parameters free only weakens the constraints on M_ν from 0.023 eV to 0.03 eV, particularly considering that

Figure 4.5: The derivatives of the galaxy power spectrum with respect to the three non-linear stochastic parameters. The derivative with respect to $P_\epsilon^{\{0\}}$ has been multiplied by a factor of 0.1 for the readability of the plot.

most of the new parameters show some degree of degeneracy with M_ν.

Finally, it is interesting to note that the M_ν-τ degeneracy, a significant focus of much recent forecast work on neutrino mass constraints from large-scale structure [21, 25, 36, 64, 49, 59] becomes much less significant when considering the NLO galaxy power spectrum. Instead the neutrino mass constraint is supported primarily by other degeneracies.

4.3.2   Free-Streaming Information

Figure 4.6 shows the results achievable using only the information contained in the scale-dependence of the matter power spectrum and structure growth rate, $f(k)$. This should isolate the effects of free-streaming as the constraining quantity. It can be seen that although the constraints are weakened somewhat by extending to the NLO power spectrum, they remain independent of changes in curvature and the dark energy equation of state. The relative change is also surprisingly small, considering the addition of 5 new bias parameters and 2 new stochastic parameters, all with a scale-dependent effect. The constraints change from about 0.08 eV in the linear case to about 0.12 eV in the NLO case.

Interestingly, fixing all of the new nuisance parameters in the NLO case improves the constraint only to 0.09 eV, which is weaker than the constraint of 0.082 eV in the linear case. From Figure 4.7, we can see that the free-streaming signal is actually enhanced in the NLO case, and on investigation of the Fisher matrix, the information content on M_ν

(a) TT, TE, EE + Free-Str. (LO)

(b) TT, TE, EE + Free-Str. (NLO)

Figure 4.6: Constraint forecasts on M_ν for Euclid allowing only the scale-dependence of the underlying matter power spectrum and structure growth rate to vary.

is actually increased. So the cause of this difference is not a change in the strength of the signal in the power spectrum itself. The issue lies in the new degeneracy with b_1. Note that in the case of the linear power spectrum, extracting only the scale-dependence means treating the amplitude of the power spectrum as a free parameter, and b_1 does not scale with k, so has no effect on the final constraint on M_ν. Fixing b_1 plus all the new nuisance parameters gives a slightly better constraint in the NLO case than in the linear case, as expected.

Now we can investigate, as before, which of the new parameters contribute in particular to the weakening of the constraints on M_ν seen in Figure 4.6. As can be seen from Figure 4.2, the story is generally similar to that in the combined case, except for b_1 and $P_\epsilon^{\{0\}}$. This makes sense, as the other nuisance parameters do not change the large-scale power spectrum, so their effects when only the scale-dependence is considered are identical to those in the combined case. The most significant difference is the impact of the constraint on b_1. Fixing b_1 improves the constraint on M_ν by about 10% (while this had a negligible effect in the combined case). We note that it would not be hard to place a prior on b_1 derived from analytical calculations or simulations, or to better constrain it by measuring the bispectrum.

Figure 4.7: The suppression in the linear and NLO power spectra caused by adding massive neutrinos, holding the overall matter density fixed. We can see that the suppression of the power spectrum is somewhat stronger in the NLO case.

Ultimately, considering that seven new nuisance parameters have been added to the calculation, all with scale-dependent effects, a degradation in the overall constraints of only 50% is quite remarkable. The cause of this is that the scale-dependence caused by neutrino free-streaming begins to take effect on much larger scales than those on which the scale-dependence of the bias parameters starts to be significant. Increasing the total neutrino mass shifts the free-streaming scale to higher k, and could therefore lead to greater degeneracy with the new parameters. This possibility is explored in Section 4.4.3.

4.3.3   BAO-Only Information

Figure 4.8 shows a comparison of the forecasted constraints on M_ν using constraints on $H(z)$ and $D_A(z)$ from the BAO feature in the power spectrum only. As is well-known, BAOs are favoured for their robustness against non-linear effects, and that is reflected here. We see that the constraints are degraded much less than in the free-streaming-only and combined cases on moving from the linear to NLO power spectrum.

Table 4.3 shows the intermediate constraints on $H(z)$ and $D_A(z)$ for these two calculations, and we can see that there is some moderate relaxation of the constraints in the NLO case. This has only a moderate impact on the final M_ν constraints, and this might suggest that the M_ν constraint from distance constraints has a limit, and does not improve persistently with decreasing error bars on $H(z)$ and $D_A(z)$.
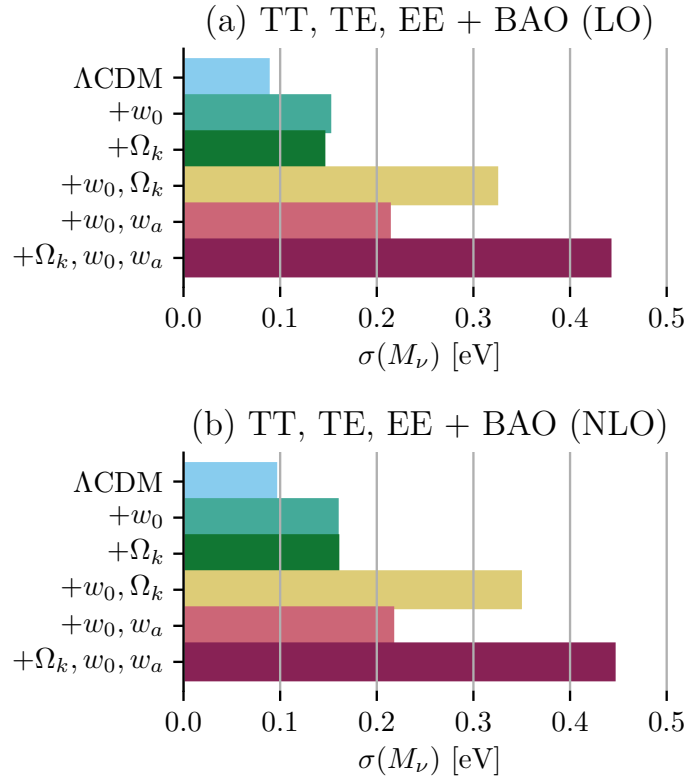
Figure 4.8: Constraint forecasts on M_ν for Euclid using distance information from the BAO feature only.

We can still seek to understand the small difference in the constraints. Fixing the bias and stochastic parameters in the NLO calculation does not improve the constraints on the distance parameters to the level of the LO case. The source of the degradation in the constraints is therefore not any degeneracy with the new bias parameters, but the weakening of the BAO feature in the NLO power spectrum itself. We have ignored the possibility of BAO reconstruction here, so it is a reasonable approximation to conclude that extending into the next-to-leading-order regime has no effect on BAO constraints on the neutrino mass. Extending to full distance information constraints (including, for example, the Alcock-Paczyński test) could result in a greater difference between the LO and NLO cases. However, BAOs dominate the overall distance constraints, as was seen in Chapter 2.

4.3.4 CMB Lensing

In Chapter 3, we showed that adding CMB lensing to galaxy clustering information only made a very significant difference to constraints on M_ν when only BAO information was used. The free-streaming-only and combined constraints were too strong to benefit much
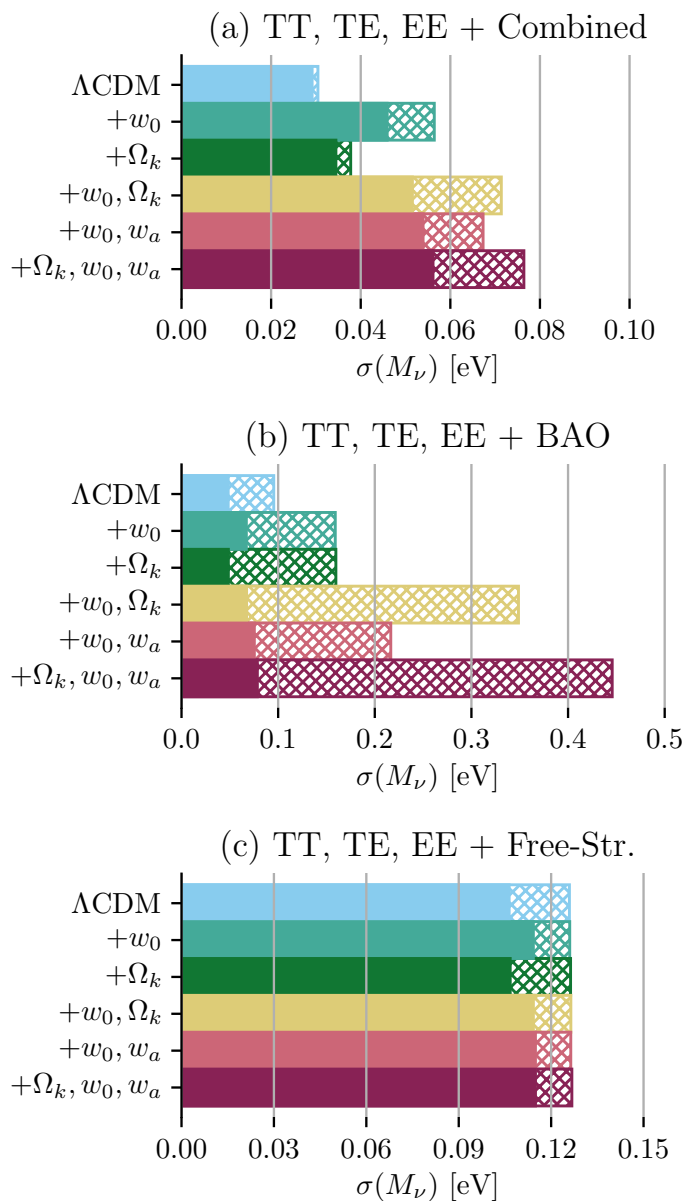
Figure 4.9: The NLO constraints on M_ν with (solid) and without (cross-hatched) CMB lensing, for the three measurement methods considered. Note the different x-axis scales.

z	$\sigma(H(z))_{LO}(\%)$	$\sigma(D_A(z))_{LO}(\%)$	$\sigma(H(z))_{NLO}(\%)$	$\sigma(D_A(z))_{NLO}(\%)$
0.65	1.69	1.41	2.24	1.88
0.75	1.34	1.02	1.78	1.35
0.85	1.23	0.93	1.62	1.22
0.95	1.18	0.9	1.53	1.16
1.05	1.15	0.89	1.47	1.12
1.15	1.14	0.88	1.44	1.11
1.25	1.14	0.89	1.42	1.11
1.35	1.16	0.92	1.42	1.13
1.45	1.21	0.98	1.46	1.18
1.55	1.29	1.08	1.54	1.28
1.65	1.44	1.24	1.69	1.46
1.75	1.67	1.49	1.93	1.74
1.85	2.08	1.93	2.36	2.22
1.95	2.82	2.73	3.17	3.13
2.05	5.94	6.12	6.63	7.04

Table 4.3: A comparison of the constraints on $H(z)$ and $D_A(z)$ achievable using the BAO feature in the linear and NLO galaxy power spectra, marginalised over nuisance parameters (bias and stochastic parameters).

from this addition. As the constraints are weaker when using the NLO power spectrum, it is worth returning to this analysis.

Figure 4.9 shows the effect of adding CMB lensing to the three galaxy clustering measurements considered thusfar. Note the different scales on the x-axes. Because the constraints from BAOs are roughly the same at linear and next-to-leading order, the results of adding CMB lensing are more or less the same as in Chapter 3. The combined constraints are weakened sufficiently that CMB lensing can now make a meaningful contribution to the constraints, at least for the more complex cosmological models, particularly those with a free dark energy equation of state.

The free-streaming constraints are also improved by CMB lensing. In this case we take only the shape of the CMB lensing power spectrum as the observable and marginalise over its amplitude so that the constraints come from the scale-dependence of the power spectrum only. The improvement is more significant than in the linear case and the results remain cosmology-independent. We have not implemented NLO galaxy-CMB lensing yet, but based on the results of Chapter 3 it is likely that that would also significantly strengthen the constraints. For now we leave this for future work.

4.4 Discussion

4.4.1 Degeneracy between Bias Parameters

We highlighted in Section 4.3.1 that there was a significant anti-correlation between the two parameters associated with the tidal field, b_{K^2} and b_{td}. An examination of the contributions of these terms to the power spectrum reveals why.

Both b_{td} and b_{K^2} contribute to the P_{13} term. In [19], the authors show that the contributions of the associated operators to P_{13} are linearly proportional. Their contributions are identical apart from a constant factor of $5/2$, making the effects of varying the two bias parameters fully degenerate.

b_{K^2} also contributes to the P_{22} term, while b_{td} does not, preventing the two parameters from being perfectly anti-correlated. However, on inspection, we see that the relative change in P_{22} when varying b_{K^2} is significantly smaller than that in P_{13}, so changing P_{13} is the dominant effect of b_{K^2}. b_{K^2} and b_{td} therefore remain quite strongly anti-correlated, even after marginalisation over other parameters.

4.4.2 Selection Effects

The observed galaxy power spectrum will always have a dependence on galaxy properties that alter the probability of a given galaxy being observed. Biases that arise in the observed power spectrum as a result of the specific subset of galaxies observed are called selection effects. Many selection effects are included by default in the bias expansion used here, but we have neglected selection effects that depend on the line-of-sight in this analysis. We will now justify this.

One particularly crucial selection effect arises from the dependence of the apparent brightness of a galaxy on its orientation with respect to us, as brighter galaxies are more likely to be detected. This selection effect is degenerate with the Kaiser effect [65], contributing the b_η in Equation 4.2, which we fix to -1 when neglecting selection effects. We can briefly examine the impact of freeing b_η to get an impression of its impact.

Without any prior on b_η, in the combined case for ΛCDM, the NLO constraint on M_ν degrades from 0.032 eV to 0.045 eV. Imposing a 1σ prior on b_η of 0.05 improves the constraint to 0.036 eV. Freeing b_η does not have any significant effect on the BAO-only or free-streaming constraints.

We finally note that Euclid will also image the galaxies that they observe [23]. This should allow for the intentional selection of a galaxy sample in such a way as to minimise any selection effects.

4.4.3 Varying M_ν

In Section 4.3.2 we found that the scale-dependent suppression in the power spectrum caused by neutrinos is not particularly degenerate with any of the new parameters. The suppression of the power spectrum by neutrino free-streaming starts at larger scales than

those at which the scale-dependent effects of the new bias parameters become significant. However, as we saw in Section 1.2, the free-streaming scale k_{FS} increases with the neutrino mass. It is therefore worth investigating whether the constraints for a higher neutrino mass would be further weakened.

We ran the results again for an increased M_ν of 0.27 eV assuming 3 degenerate mass states, and the constraints were better by about 10% than for our usual fiducial model. Therefore, any weakening of the constraints due to greater degeneracy with the non-linear bias parameters is more than compensated for by the increased amplitude of the suppression that comes with higher mass neutrinos.

4.4.4 Neutrino-Induced Bias

Something we have neglected in our treatment in this chapter is the effects of massive neutrinos themselves on the bias parameters. As we have seen, the free-streaming of massive neutrinos introduces a scale-dependence in the growth of perturbations, not just for neutrinos, but also cold dark matter and baryons. This should then result in some scale-dependence in the bias parameters. This was studied in detail by [66] for b_1, who modified the spherical collapse and peak-background split derivations of the large-scale bias to account for this. Including this scale-dependent bias should work counter to the suppression caused by massive neutrinos, theoretically weakening the signal. The same effect should also cause a scale-dependence in the higher-order bias parameters.

This is a complicated topic and a full model of the next-to-leading-order power spectrum that accounts for this effect has yet to be developed. In addition, there is some disagreement about whether the effect should be measurable at all, as it has not been identifiable in some N-body simulations [67][68][69]. We therefore leave this at present for future work.

4.4.5 Comparisons with Previous Work

Despite the wide array of cosmological neutrino mass constraint forecasts available in the literature, very few have been performed beyond the linear regime, and direct comparisons are often complicated by different initial assumptions.

A recent article [70] performed an MCMC forecast for Euclid to 1-loop order. Combining the current Planck likelihood with their Euclid forecast, they quote a constraint of 24 meV, which exactly matches our constraint for ΛCDM if N_{eff} is not allowed to vary (as in [70]).

This consistency may be a fluke, but we also showed in [36] that our Fisher matrix constraints were remarkably consistent with the MCMC forecasts of [59]. In both cases, the decision to use an existing Planck likelihood instead of a forecast CMB prior did not seem to significantly impact the results, emphasising the dominance of the galaxy clustering information. The choice to control errors on small scales using a gradually increasing uncertainty instead of a sharp k cut-off at $k = 0.2 \, h$ Mpc also does not seem to generate incongruencies between the results.

The authors of [70] also examined the specific contribution of BAOs to their final constraints, and found it to be modest. This emphasises our point in our previous work and here that neutrino mass constraints from BAOs alone waste a great deal of usable information.

4.5 Conclusions

In this chapter, we have dealt with an important caveat of our conclusions in the previous two chapters by extending our calculations to account for next-to-leading-order corrections to the galaxy and CMB lensing power spectra. Using the formalism of [19], we extended our Fisher matrix calculations to include an additional five higher-order galaxy bias parameters and two new stochastic parameters. Given that many of these new parameters suppress the galaxy power spectrum on small scales (see Figures 4.4 and 4.5), and since we identified the suppression of the small-scale power spectrum by neutrinos to be a crucial constraining component in previous chapters, it could be expected that M_ν would become significantly more difficult to measure robustly.

As in the previous chapter, we take Euclid as an example future survey and focus on three types of constraint - 'BAO-only'constraints, which use only constraints on the expansion rate derived from measurements of the BAO scale in the power spectrum; cosmology-independent 'free-streaming' constraints, which isolate the constraining information in the matter power spectrum and growth-factor only; and 'combined' constraints, which use the full observed galaxy power spectrum, naturally including all expansion information, redshift-space distortions, and the shape of the power spectrum. The BAO-only constraints were barely weakened at all by extending to the non-linear power spectrum, but therefore remained very sensitive to minimal extensions to the underlying cosmology, including non-zero curvature or $w \neq -1$, even when combined with full CMB information. The free-streaming constraints remained independent of these extensions, but the 1σ error expected from Euclid was degraded from about 0.08 eV to 0.12 eV. This was despite the small-scale suppression in power by massive neutrinos actually being enhanced in the NLO case (see Figure 4.7). Adding the shape of the CMB lensing power spectrum provided a small improvement. We did not consider galaxy-CMB lensing here but based on the results from Chapter 3, we can assume it was likely that this would also improve the constraints significantly. The forecasted combined constraints are weakened from 0.022 eV to 0.032 eV on extending to NLO, a comparable amount to the free-streaming constraints. These results also remain very sensitive to assumptions about cosmology, degrading by several factors if curvature and the dark energy equation of state are allowed to vary (see Figure 4.3), although CMB lensing measurements from a future experiment like Simons Observatory can help break some of the degeneracies to a degree (see Figure 4.9). The relative impact of improving constraints on the optical depth to the CMB, τ, is also significantly reduced due to the large number of new weakly-constrained parameters.

We chose not to include all possible selection effects in the general analysis, presenting only a summary analysis in Section 4.4.2. Including a full treatment of selection effects

would probably have the greatest effect on the combined constraints because of their degeneracy with redshift-space distortions. We also have not accounted for possible scale-dependence in the bias parameters caused by massive neutrinos (see Section 4.4.4), but leave this as future work. However, despite these caveats, our results are also quite conservative, considering we do not impose any priors on any of the bias parameters, although theoretically motivated priors or constraints from N-body simulations could reasonably be applied.

Chapter 5

Conclusions

Traditional experiments have so far failed to provide us with meaningful measurements of the neutrino mass scale. As an alternative, science has turned to precision cosmology. The strongest bounds on the total neutrino mass come from cosmological measurements, and this is likely to remain the case for the foreseeable future. The promise of an explicit neutrino mass detection has helped motivate a large number of currently planned cosmological surveys.

Cosmic microwave background and large-scale structure measurements provide complementary measurements of the early and late universe that we connect using the concordance ΛCDM model. This model has been incredibly successful in helping us to fit our observations of the universe using a surprisingly small number of parameters. However, it is not without its challenges. The goal of this thesis has been to provide insight into how confident we can be in constraints on M_ν that are extracted from fits to cosmological data, given our significant ignorance about many elements of our universe, and several issues that still plague the ΛCDM model despite ever-increasing precision.

Existing and forecasted neutrino mass constraints in the literature often combine several datasets and extract constraints using global fits to a very small number of parameters. In this thesis and the publications contributing to it, we have for the first time thoroughly deconstructed these constraints, seeking to understand in a more analytical way the degeneracies between different parameters and how robust different probes of the neutrino mass are against changes in the underlying cosmology. We believe this new approach could and should be applied more broadly in the field of precision cosmology, to a wide array of problems. As we have shown in this thesis, it can quite easily lead to important insights that are easily missed with more automated analysis approaches.

This work has resulted in many interesting conclusions. In Chapter 2, we analysed in detail the constraints on the neutrino mass available from galaxy redshift surveys, assuming minimal CMB information. The influence of massive neutrinos on the galaxy power spectrum can be divided into two broad categories. First, they modify the expansion rate of the universe. This is measured at low redshifts by galaxy surveys using standard rulers, like baryon acoustic oscillations (BAOs), and the Alcock-Paczyński test. Secondly, massive neutrinos change the growth of structure, modifying both the shape and amplitude

of the power spectrum and structure growth rate, f, with the latter being constrained using redshift-space distortions. Most importantly, massive neutrinos uniquely create a relative suppression in the growth of structure below a characteristic free-streaming scale. This introduces a scale-dependent effect in the power spectrum and structure growth rate. While we have shown that neutrino mass constraints based on BAO measurements only or on the full redshift-space galaxy power spectrum are heavily dependent on cosmological assumptions, if the scale-dependent signal is isolated and used alone as a probe ('free-streaming' constraints), the resulting constraints are cosmology-independent and still relatively powerful. We also clarified the origin of the degeneracy between the neutrino mass and the optical depth to the CMB in cosmological measurements for the first time. Our free-streaming constraints are independent of this degeneracy.

In Chapter 3, we extended our calculations to include more CMB information, including both primary anisotropies and CMB lensing. We showed that information from the temperature and E-mode polarisation spectra could be added to the free-streaming constraints while keeping them cosmology-independent, improving constraints significantly. We also showed that an analogous method to that used in Chapter 2 could be applied to the CMB lensing and galaxy-CMB lensing power spectra to extract only the free-streaming effect as a probe, and that this information could be combined with the galaxy survey free-streaming information to improve constraints even further while maintaining cosmology independence. We showed that the combined constraints on the neutrino mass from a survey like Euclid will not gain much from adding CMB lensing, particularly as τ is better constrained, but that BAO-only constraints improve significantly, largely because they are originally so weak. Both the combined constraints and BAO-only constraints remain strongly cosmology-dependent even when CMB lensing is added.

In Chapter 4, we analysed how the constraints from the previous chapters change when we extend our calculations to the next-to-leading-order power spectrum. Including next-to-leading-order contributions enhances the power spectrum on small scales. We discover that although the suppression of the power spectrum by massive neutrinos on small scales is actually enhanced in the NLO case, the total constraining power is still reduced because of degeneracies with the large number of additional nuisance parameters. Therefore, our free-streaming constraints are weakened somewhat, but remain cosmology independent. Our combined constraints are weakened by a greater fraction. BAO-only constraints are not altered much at all by non-linearities, but still waste a lot of useful information.

All in all, our free-streaming constraints remain the only cosmology-independent constraints on the neutrino mass, and we therefore consider this method the most robust approach to extracting constraints. Even though they are weakened when extending to the non-linear regime, they remain competitive with BAO and CMB lensing constraints. Constraints from all other information sources (including BAO-only and combined constraints) should be approached with caution.

What does all of this mean for the future? There are a great number of conclusions that can be drawn. Cosmological surveys intended to measure the neutrino mass robustly should focus on designing their surveys to measure the broadband shape of the power spectrum as precisely as possible, in order to allow extraction of the free-streaming signal.

When considering combined constraints, additional constraints on the τ, such as from 21-cm surveys or improved polarisation measurements, could be a useful complement. The constraining power of distance measurements like BAOs when it comes to the neutrino mass is very limited relative to what can be achieved using the full power spectrum.

In the future, this work could be extended in many different directions. We have already mentioned that the impact of neutrino free-streaming on the galaxy power spectrum is redshift-dependent. It would be very interesting to determine how measurable this redshift-dependence will be in future surveys, and to determine how surveys could be optimised to detect this.

Of course, the application of the results of this thesis to actual data, when the time comes, in order to provide robust cosmology-independent neutrino mass constraints, is the most exciting prospect.

Appendix A

Survey Parameters

Here we present the parameters for the planned galaxy clustering and CMB surveys we consider in this thesis. Euclid is the only galaxy survey considered in the main text, but we present forecasts for a further four galaxy surveys in Appendix B.

The additional surveys we account for are HETDEX (Table A.2), PFS (Table A.3), DESI (Table A.4) and WFIRST (Table A.5).

Our HETDEX [71] survey model consists of two redshift bins, with a total redshift range of $1.9 \leq z \leq 3.5$ and an area of 425 deg^2. The model comprises a total volume of 2 Gpc$^3/h^3$ and a total of 0.8 million galaxies [72]. We assume a constant bias of $b(z) = 1.5$.

For PFS, we use the survey parameters and bias values specified in table 2 of [73]. Under these specifications, PFS will have a redshift range of $0.6 \leq z \leq 2.4$ and an area of 1464 deg^2, providing a total volume of 9.91 Gpc$^3/h^3$ and 4.18 million galaxies across all redshift bins.

For DESI, we refer to table 2.3 of [74]. The authors provide a range of survey plans. The main survey covers a redshift range of $0.6 \leq z \leq 1.9$ with an area of 14000 deg^2. Redshift bin volumes are provided in units of Gpc$^3/h$ so we recalculated them in units of Gpc$^3/h^3$ to comply with our code. The total volume was then calculated to be 57.36 Gpc$^3/h^3$ with 22.35 million galaxies in total for three individual galaxy populations - emission line galaxies (ELGs), luminous red galaxies (LRGs) and quasars (QSOs). Here we use data from only the ELGs for a more direct comparison with other surveys. We use the fiducial bias formulae provided in section 2.4.2 of [74].

There is much less specific survey information available for Euclid [23, 24] and WFIRST [75, 76]. We follow tables 6 and 7 of [32], respectively, for these two surveys. In the case of Euclid, this assumes a survey area of 15000 deg^2 and a redshift range of $0.6 \leq z \leq 2.1$, corresponding to a survey volume of 72 Gpc$^3/h^3$ with a total galaxy count of 50 million galaxies. We assume a bias scaling of $b(z)D(z) = 0.76$. For WFIRST, the survey area is 2000 deg^2 and the redshift range is $1 \leq z \leq 2.8$. The survey volume is 13.55 Gpc$^3/h^3$, containing 26.5 million galaxies. We calculate the bias in this case as $b(z) = 1.5 + 0.4(z - 1.5)$.

The survey parameters assumed for Simons Observatory are taken from Table 1 of [2], with both the Small Aperture and Large Aperture Telescopes being included.

	Euclid		
z	Volume (h^{-3} Gpc^3)	N_g (millions)	b_1
0.65	2.59	1.65	1.07
0.75	3.07	4.425	1.12
0.85	3.52	5.7	1.17
0.95	3.93	5.85	1.23
1.05	4.29	5.6625	1.28
1.15	4.62	5.2875	1.34
1.25	4.9	4.875	1.39
1.35	5.14	4.275	1.45
1.45	5.35	3.525	1.51
1.55	5.52	2.775	1.57
1.65	5.66	2.0625	1.62
1.75	5.78	1.4625	1.68
1.85	5.88	0.975	1.74
1.95	5.95	0.6	1.8
2.05	6.01	0.225	1.86

Table A.1: Survey parameters for Euclid.

	HETDEX		
z	Volume (h^{-3} Gpc^3)	N_g (millions)	b_1
2.2	0.80971	0.445	1.5
3.0	1.18553	0.3468	1.5

Table A.2: Survey parameters for HETDEX.

	PFS		
z	Volume (h^{-3} Gpc^3)	N_g (millions)	b_1
0.7	0.59	0.1121	1.18
0.9	0.79	0.474	1.26
1.1	0.96	0.5568	1.34
1.3	1.09	0.8502	1.42
1.5	1.19	0.6545	1.5
1.8	2.58	0.7998	1.62
2.2	2.71	0.7317	1.78

Table A.3: Survey parameters for PFS.

DESI (ELG)

z	Volume (h^{-3} Gpc^3)	N_g (millions)	b_1
0.65	2.8	0.433	1.18
0.75	3.2	3.18	1.24
0.85	3.56	2.69	1.3
0.95	3.89	2.93	1.36
1.05	4.17	2.02	1.42
1.15	4.42	1.89	1.48
1.25	4.63	1.87	1.54
1.35	4.81	0.732	1.6
1.45	4.97	0.652	1.67
1.55	5.09	0.461	1.73
1.65	5.19	0.176	1.79

Table A.4: Surve parameters for DESI. Note that we only include the emission-line galaxy population here (see [74]).

WFIRST

z	Volume (h^{-3} Gpc^3)	N_g (millions)	b_1
1.05	0.57	2.1246	1.32
1.15	0.62	2.3552	1.36
1.25	0.65	2.7754	1.4
1.35	0.69	3.1054	1.44
1.45	0.71	3.378	1.48
1.55	0.74	3.1518	1.52
1.65	0.76	2.661	1.56
1.75	0.77	2.1836	1.6
1.85	0.78	1.7394	1.64
1.95	0.79	1.3436	1.68
2.05	0.8	0.322	1.72
2.15	0.81	0.3018	1.76
2.25	0.81	0.2736	1.8
2.35	0.81	0.2312	1.84
2.45	0.81	0.192	1.88
2.55	0.81	0.1562	1.92
2.65	0.81	0.1252	1.96
2.75	0.81	0.098	2.0

Table A.5: Survey parameters for WFIRST.

Simons Observatory LAT ($f_{\mathbf{sky}} = 0.4$)

Freq. (GHz)	T Noise (μK arcmin)	P Noise (μK arcmin)	FWHM (arcmin)
27.0	71.0	100.5	7.4
39.0	36.0	50.9	5.1
93.0	8.0	11.3	2.2
145.0	10.0	14.1	1.4
225.0	22.0	31.1	1.0

Table A.6: Survey parameters for Simons Observatory: Large Aperture Telescope.

Simons Observatory SAT ($f_{\mathbf{sky}} = 0.1$)

Freq. (GHz)	T Noise (μK arcmin)	P Noise (μK arcmin)	FWHM (arcmin)
27.0	35.0	49.5	91.0
39.0	21.0	29.7	63.0
93.0	2.6	3.7	30.0
145.0	3.3	4.7	17.0
225.0	6.3	8.9	11.0

Table A.7: Survey parameters for Simons Observatory: Small Aperture Telescope.

Appendix B

Extended Results

We hope the previous chapters leave the reader with little doubt that this thesis represents one of the broadest and deepest ranges of cosmological neutrino mass constraint forecasts ever conducted. Here we will try to provide a summary of results for other planned galaxy surveys for the curious reader. Using the broad conclusions reached, we try to focus on the most significant numbers, from which it should be possible to intuitively infer related results using the examples provided in this thesis.

We do not include here galaxy-CMB lensing from Chapter 3 or next-to-leading-order results from Chapter 4, because both are computationally time consuming to generate. It should be possible to extrapolate results for both by comparing with the examples for Euclid in those chapters.

We divide our results for galaxy surveys into two sections: results combined with existing Planck data and results combined with eventual Simons Observatory data. We extract the constraints from Planck using the same method to extract constraints from the temperature power spectrum alone in Chapter 3 Section 3.2, using the Planck 2018 data release, but also using the EE and TE power spectra, and using the lensed power spectra when CMB lensing is included. We do not include the combination of free-streaming clustering information and CMB lensing in the Planck case, as we have not extracted constraints from the shape of the CMB lensing power spectrum alone from the Planck data.

Euclid + Planck	ΛCDM	$+w_0$	$+\Omega_k$	$+w_0, \Omega_k$	$+w_0, w_a$	$+\Omega_k, w_0, w_a$
Combined	0.023	0.036	0.025	0.052	0.045	0.058
Combined + Fixed τ	0.01	0.028	0.013	0.044	0.038	0.051
Combined + CMB Lensing	0.023	0.036	0.025	0.052	0.045	0.058
Combined + CMB Lensing + Fixed τ	0.01	0.028	0.013	0.045	0.038	0.051
BAO	0.12	0.23	0.2	0.39	0.37	0.64
BAO + CMB Lensing	0.093	0.15	0.15	0.2	0.22	0.25
BAO + CMB Lensing + Fixed τ	0.09	0.15	0.14	0.19	0.21	0.24
Free-Streaming	0.084	0.087	0.085	0.087	0.1	0.1

Table B.1: Neutrino mass constraint forecast summary for Planck and Euclid.

WFIRST + Planck	ΛCDM	$+w_0$	$+\Omega_k$	$+w_0, \Omega_k$	$+w_0, w_a$	$+\Omega_k, w_0, w_a$
Combined	0.032	0.058	0.035	0.072	0.079	0.09
Combined + Fixed τ	0.02	0.052	0.024	0.064	0.074	0.083
Combined + CMB Lensing	0.032	0.058	0.035	0.072	0.079	0.09
Combined + CMB Lensing + Fixed τ	0.02	0.052	0.024	0.065	0.073	0.083
BAO	0.18	0.3	0.27	0.51	0.42	0.67
BAO + CMB Lensing	0.13	0.19	0.17	0.22	0.25	0.27
BAO + CMB Lensing + Fixed τ	0.12	0.18	0.16	0.22	0.24	0.26
Free-Streaming	0.18	0.18	0.18	0.18	0.18	0.18

Table B.2: Neutrino mass constraint forecast summary for Planck and WFIRST.

DESI (ELG) + Planck	ΛCDM	$+w_0$	$+\Omega_k$	$+w_0, \Omega_k$	$+w_0, w_a$	$+\Omega_k, w_0, w_a$
Combined	0.026	0.047	0.028	0.064	0.055	0.069
Combined + Fixed τ	0.014	0.041	0.017	0.058	0.049	0.062
Combined + CMB Lensing	0.026	0.047	0.028	0.065	0.056	0.069
Combined + CMB Lensing + Fixed τ	0.014	0.041	0.017	0.059	0.049	0.063
BAO	0.15	0.27	0.24	0.54	0.38	0.68
BAO + CMB Lensing	0.11	0.17	0.16	0.21	0.23	0.25
BAO + CMB Lensing + Fixed τ	0.11	0.16	0.15	0.2	0.23	0.25
Free-Streaming	0.11	0.12	0.11	0.12	0.12	0.12

Table B.3: Neutrino mass constraint forecast summary for Planck and DESI (ELG only).

PFS + Planck	ΛCDM	$+w_0$	$+\Omega_k$	$+w_0, \Omega_k$	$+w_0, w_a$	$+\Omega_k, w_0, w_a$
Combined	0.038	0.08	0.04	0.087	0.099	0.1
Combined + Fixed τ	0.029	0.074	0.032	0.081	0.094	0.097
Combined + CMB Lensing	0.038	0.08	0.04	0.087	0.099	0.1
Combined + CMB Lensing + Fixed τ	0.029	0.074	0.032	0.081	0.093	0.097
BAO	0.22	0.32	0.34	0.56	0.43	0.68
BAO + CMB Lensing	0.14	0.19	0.18	0.23	0.26	0.27
BAO + CMB Lensing + Fixed τ	0.14	0.19	0.17	0.22	0.25	0.27
Free-Streaming	0.22	0.22	0.22	0.22	0.22	0.22

Table B.4: Neutrino mass constraint forecast summary for Planck and PFS.

HETDEX + Planck	ΛCDM	$+w_0$	$+\Omega_k$	$+w_0, \Omega_k$	$+w_0, w_a$	$+\Omega_k, w_0, w_a$
Combined	0.09	0.17	0.09	0.17	0.22	0.23
Combined + Fixed τ	0.085	0.16	0.086	0.17	0.21	0.23
Combined + CMB Lensing	0.087	0.16	0.088	0.17	0.2	0.21
Combined + CMB Lensing + Fixed τ	0.081	0.16	0.082	0.16	0.2	0.21
BAO	0.55	0.57	0.61	0.67	0.62	0.71
BAO + CMB Lensing	0.25	0.31	0.25	0.32	0.33	0.34
BAO + CMB Lensing + Fixed τ	0.24	0.31	0.24	0.31	0.33	0.33
Free-Streaming	0.47	0.5	0.48	0.5	0.5	0.5

Table B.5: Neutrino mass constraint forecast summary for Planck and HETDEX.

Euclid + S.O.	ΛCDM	$+w_0$	$+\Omega_k$	$+w_0, \Omega_k$	$+w_0, w_a$	$+\Omega_k, w_0, w_a$
Combined	0.022	0.035	0.024	0.043	0.044	0.049
Combined + Fixed τ	0.0096	0.027	0.011	0.033	0.036	0.041
Combined + CMB Lensing	0.021	0.034	0.023	0.041	0.042	0.046
Combined + CMB Lensing + Fixed τ	0.0095	0.025	0.011	0.03	0.033	0.035
BAO	0.088	0.15	0.15	0.32	0.21	0.44
BAO + CMB Lensing	0.046	0.066	0.046	0.067	0.075	0.082
BAO + CMB Lensing + Fixed τ	0.034	0.052	0.035	0.057	0.061	0.073
Free-Streaming	0.078	0.08	0.078	0.08	0.081	0.081
Free-Streaming + CMB Lensing (Shape)	0.073	0.075	0.074	0.075	0.076	0.076

Table B.6: Neutrino mass constraint forecast summary for Simons Observatory and Euclid.

WFIRST + S.O.	ΛCDM	$+w_0$	$+\Omega_k$	$+w_0, \Omega_k$	$+w_0, w_a$	$+\Omega_k, w_0, w_a$
Combined	0.031	0.055	0.034	0.06	0.073	0.077
Combined + Fixed τ	0.02	0.048	0.022	0.051	0.067	0.07
Combined + CMB Lensing	0.031	0.051	0.033	0.053	0.06	0.061
Combined + CMB Lensing + Fixed τ	0.019	0.041	0.019	0.042	0.049	0.049
BAO	0.14	0.2	0.22	0.4	0.23	0.44
BAO + CMB Lensing	0.058	0.072	0.061	0.078	0.077	0.086
BAO + CMB Lensing + Fixed τ	0.043	0.059	0.051	0.069	0.064	0.079
Free-Streaming	0.16	0.16	0.16	0.16	0.16	0.16
Free-Streaming + CMB Lensing (Shape)	0.12	0.14	0.12	0.14	0.14	0.14

Table B.7: Neutrino mass constraint forecast summary for Simons Observatory and WFIRST.

DESI (ELG) + S.O.	ΛCDM	$+w_0$	$+\Omega_k$	$+w_0, \Omega_k$	$+w_0, w_a$	$+\Omega_k, w_0, w_a$
Combined	0.025	0.045	0.027	0.051	0.053	0.056
Combined + Fixed τ	0.014	0.038	0.015	0.043	0.045	0.048
Combined + CMB Lensing	0.024	0.043	0.026	0.047	0.049	0.052
Combined + CMB Lensing + Fixed τ	0.013	0.033	0.015	0.035	0.039	0.04
BAO	0.11	0.16	0.2	0.41	0.23	0.46
BAO + CMB Lensing	0.052	0.067	0.052	0.069	0.076	0.086
BAO + CMB Lensing + Fixed τ	0.039	0.053	0.041	0.058	0.063	0.079
Free-Streaming	0.1	0.1	0.1	0.1	0.1	0.1
Free-Streaming + CMB Lensing (Shape)	0.091	0.097	0.092	0.097	0.097	0.097

Table B.8: Neutrino mass constraint forecast summary for Simons Observatory and DESI (ELG only).

PFS + S.O.	ΛCDM	$+w_0$	$+\Omega_k$	$+w_0, \Omega_k$	$+w_0, w_a$	$+\Omega_k, w_0, w_a$
Combined	0.037	0.072	0.039	0.075	0.086	0.087
Combined + Fixed τ	0.028	0.066	0.029	0.067	0.079	0.08
Combined + CMB Lensing	0.036	0.061	0.037	0.061	0.068	0.068
Combined + CMB Lensing + Fixed τ	0.025	0.049	0.025	0.05	0.056	0.059
BAO	0.18	0.22	0.29	0.43	0.26	0.46
BAO + CMB Lensing	0.062	0.073	0.068	0.081	0.078	0.09
BAO + CMB Lensing + Fixed τ	0.046	0.061	0.058	0.074	0.066	0.085
Free-Streaming	0.18	0.18	0.18	0.18	0.18	0.18
Free-Streaming + CMB Lensing (Shape)	0.13	0.16	0.13	0.16	0.16	0.16

Table B.9: Neutrino mass constraint forecast summary for Simons Observatory and PFS.

HETDEX + S.O.	ΛCDM	$+w_0$	$+\Omega_k$	$+w_0, \Omega_k$	$+w_0, w_a$	$+\Omega_k, w_0, w_a$
Combined	0.087	0.14	0.088	0.16	0.16	0.18
Combined + Fixed τ	0.082	0.14	0.083	0.15	0.15	0.17
Combined + CMB Lensing	0.06	0.078	0.068	0.087	0.08	0.089
Combined + CMB Lensing + Fixed τ	0.045	0.069	0.058	0.081	0.071	0.084
BAO	0.46	0.48	0.48	0.51	0.49	0.52
BAO + CMB Lensing	0.067	0.083	0.13	0.13	0.084	0.13
BAO + CMB Lensing + Fixed τ	0.05	0.076	0.13	0.13	0.077	0.13
Free-Streaming	0.39	0.39	0.39	0.39	0.39	0.39
Free-Streaming + CMB Lensing (Shape)	0.15	0.26	0.16	0.26	0.26	0.26

Table B.10: Neutrino mass constraint forecast summary for Simons Observatory and HET-DEX.

Appendix C

Cosmological Perturbation Theory

The Vlasov or collisionless Boltzmann equation governs the evolution of the phase space distribution of a nonrelativistic, collisionless fluid

$$\frac{df}{d\tau} = \frac{\partial f}{\partial \tau} + \frac{\boldsymbol{p}}{ma} \cdot \nabla f - am\nabla\Phi \cdot \frac{\partial f}{\partial \boldsymbol{p}} = 0. \tag{C.1}$$

where $f = f(\boldsymbol{x}, \boldsymbol{p}, \tau)$ is the number density in phase space, $\boldsymbol{p}$ is the comoving momentum and Φ is the gravitational potential, given by the Poisson equation

$$\nabla^2 \Phi(\boldsymbol{x}, \tau) = \frac{3}{2}\mathcal{H}^2 \Omega_m(\tau)\delta(\boldsymbol{x}, \tau). \tag{C.2}$$

One can derive the continuity and Euler equations by taking moments of the Vlasov equation

$$\frac{\partial \delta(\boldsymbol{x}, \tau)}{\partial \tau} + \nabla \cdot [1 + \delta(\boldsymbol{x}, \tau)]\boldsymbol{v}(\boldsymbol{x}, \tau) = 0, \tag{C.3}$$

$$\frac{\partial \boldsymbol{v}(\boldsymbol{x}, \tau)}{\partial \tau} + \mathcal{H}(\tau)\boldsymbol{v}(\boldsymbol{x}, \tau) + \boldsymbol{v}(\boldsymbol{x}, \tau) \cdot \nabla\boldsymbol{v}(\boldsymbol{x}, \tau) = -\nabla\Phi(\boldsymbol{x}, \tau), \tag{C.4}$$

where $\mathcal{H} = Ha$ is the conformal expansion rate, $\boldsymbol{\delta}(\boldsymbol{x}, \tau)$ is the matter density field, $\boldsymbol{v}(\boldsymbol{x}, \tau)$ is the velocity field. We assume here that there is no anisotropic stress ($\sigma_{ij} = 0$).

Together, the continuity, Euler and Poisson equations determine the evolution of density and velocity perturbations over time.

C.1 Linear Perturbation Theory

On large scales and at early times, while the fluctuations in these fields are still small, the equations above can be linearised, giving

$$\frac{\partial \delta(\boldsymbol{x}, \tau)}{\partial \tau} + \theta(\boldsymbol{x}, \tau) = 0, \tag{C.5}$$

$$\frac{\partial \boldsymbol{v}(\boldsymbol{x},\tau)}{\partial \tau} + \mathcal{H}(\tau)\boldsymbol{v}(\boldsymbol{x},\tau) = -\nabla\Phi(\boldsymbol{x},\tau). \tag{C.6}$$

Here $\theta(\boldsymbol{x},\tau) = \nabla \cdot \boldsymbol{v}(\boldsymbol{x},\tau)$ is the divergence of the velocity field, and therefore satisfies

$$\frac{\partial \theta(\boldsymbol{x},\tau)}{\partial \tau} + \mathcal{H}(\tau)\theta(\boldsymbol{x},\tau) + \frac{3}{2}\Omega_m(\tau)\mathcal{H}^2(\tau)\delta(\boldsymbol{x},\tau) = 0. \tag{C.7}$$

Equation C.5 can now be solved using Equation C.7 to give the evolution of the matter density perturbations

$$\frac{d^2\delta(\boldsymbol{x},\tau)}{d\tau^2} + \mathcal{H}(\tau)\frac{d\delta(\boldsymbol{x},\tau)}{d\tau} - \frac{3}{2}\Omega_m(\tau)\mathcal{H}^2(\tau)\delta(\boldsymbol{x},\tau) = 0. \tag{C.8}$$

Equation C.8 always has two solutions that determine the evolution of the density perturbations as a function of the scale factor, expansion rate and matter density, usually called fast-growing and slow-growing modes. For example, one can show that during matter domination, density perturbations grow proportionally to the scale factor. In the radiation-dominated era, sub-horizon perturbations grow as $\ln(a)$. During the current era, with a significant contribution from dark energy, the rate of growth of perturbations with respect to the scale factor is somewhat suppressed compared to that during matter domination.

A final step usually taken is to rewrite Equation C.8 in Fourier space in terms of a scale-independent linear growth factor $D(\tau) = \delta(\boldsymbol{x},\tau)/\delta(\boldsymbol{x},0)$. Crucially, however, the growth factor is actually *not* scale-independent in the case of a cosmology with massive neutrinos, because of the scale-dependent contribution of neutrinos to the growth of perturbations (see Section 1.2).

C.2   Non-Linear Perturbation Theory

In linear perturbation theory, different modes evolve independently of each other. Moving beyond this regime, the coupling of different modes must be accounted for. We begin by returning to the original master equations (Equations C.2-C.4). We take the divergence of the Euler equation (Equation C.4) and substitute for Φ to get the following results in Fourier space

$$\frac{\partial \delta(\boldsymbol{k},\tau)}{\partial \tau} + \theta(\boldsymbol{k},\tau) = -\int d^3\boldsymbol{k_1}d^3\boldsymbol{k_2}\delta_D(\boldsymbol{k}-\boldsymbol{k_{12}})\alpha(\boldsymbol{k_1},\boldsymbol{k_2})\theta(\boldsymbol{k_1},\tau)\delta(\boldsymbol{k_2},\tau), \tag{C.9}$$

$$\frac{\partial \theta(\boldsymbol{k},\tau)}{\partial \tau} + \mathcal{H}(\tau)\theta(\boldsymbol{k},\tau) + \frac{3}{2}\Omega_m\mathcal{H}^2(\tau)\delta(\boldsymbol{k},\tau) = -\int d^3\boldsymbol{k_1}d^3\boldsymbol{k_2}\delta_D(\boldsymbol{k}-\boldsymbol{k_{12}})\beta(\boldsymbol{k_1}\boldsymbol{k_2})\theta(\boldsymbol{k_1},\tau)\theta(\boldsymbol{k_2},\tau), \tag{C.10}$$

where $\boldsymbol{k_{12}} = \boldsymbol{k_1} + \boldsymbol{k_2}$ and

$$\alpha(\boldsymbol{k_1},\boldsymbol{k_2}) = \frac{\boldsymbol{k_{12}} \cdot \boldsymbol{k_1}}{k_1^2}, \tag{C.11}$$

Figure C.1: Source: [77]. Feynman diagrams representing the one-loop contributions to the power spectrum.

$$\beta(\boldsymbol{k_1}, \boldsymbol{k_2}) = \frac{k_{12}^2(\boldsymbol{k_1} \cdot \boldsymbol{k_2})}{2k_1^2 k_2^2}. \tag{C.12}$$

Perturbation theory solves such equations with perturbative expansions around the linear solutions up to a given order

$$\delta(\boldsymbol{k}, \tau) = \sum_{n=1}^{\infty} \delta^n(\boldsymbol{k}, \tau), \tag{C.13}$$

$$\theta(\boldsymbol{k}, \tau) = \sum_{n=1}^{\infty} \theta^n(\boldsymbol{k}, \tau). \tag{C.14}$$

We can use Equations C.9 and C.10 to write δ and θ up to nth order as

$$\delta_n(\boldsymbol{k}) = \int d^3\boldsymbol{q_1}... \int d^3\boldsymbol{q_n} \delta_D(\boldsymbol{k} - \boldsymbol{q}_{1...n}) F_n(\boldsymbol{q_1}, ..., \boldsymbol{q_n}) \delta_1(\boldsymbol{q_1})...\delta_1(\boldsymbol{q_n}), \tag{C.15}$$

$$\theta_n(\boldsymbol{k}) = \int d^3\boldsymbol{q_1}... \int d^3\boldsymbol{q_n} \delta_D(\boldsymbol{k} - \boldsymbol{q}_{1...n}) G_n(\boldsymbol{q_1}, ..., \boldsymbol{q_n}) \delta_1(\boldsymbol{q_1})...\delta_1(\boldsymbol{q_n}). \tag{C.16}$$

To first order, $F_1 = G_1 = 1$ and we recover the linear case. F_n and G_n for other orders can be extracted by inserting these equations into Equations C.9 and C.10. A useful way to examine the contributions to the evolution at different orders is using Feynman diagrams. Examining the contribution to the power spectrum from second order perturbation theory, one discovers that there is a contribution of the same order that arises from third-order perturbation theory (see Figure C.1). These terms are combined to give the 'next-to-leading-order' or 'one-loop' power spectrum.

To summarise, at next-to-leading order, the power spectrum is equal to the linear power spectrum with two additional contributions

$$P_{NLO}(k) = P_L(k) + P^{(22)}(k) + 2P^{(13)}(k), \tag{C.17}$$

$$P_{mm}^{(22)}(k) = \langle \delta^{(2)}(\boldsymbol{k})\delta^{(2)}(\boldsymbol{k'}) \rangle = 2 \int_{\boldsymbol{p}} [F_2(\boldsymbol{p}, \boldsymbol{k} - \boldsymbol{p})]^2 P_L(p) P_L(|\boldsymbol{k} - \boldsymbol{p}|), \tag{C.18}$$

$$P_{mm}^{(13)}(k) = \langle \delta^{(1)}(\boldsymbol{k})\delta^{(3)}(\boldsymbol{k'})\rangle = 3P_L(k) \int_{\boldsymbol{p}} F_3(\boldsymbol{p}, \boldsymbol{k} - \boldsymbol{p})P_L(p). \tag{C.19}$$

The relevant kernels are given by [78]:

$$F_2(\boldsymbol{k_1}, \boldsymbol{k_2}) = \frac{5}{7} + \frac{2}{7}\frac{(\boldsymbol{k_1} \cdot \boldsymbol{k_2})^2}{k_1^2 k_2^2} + \frac{\boldsymbol{k_1} \cdot \boldsymbol{k_2}}{2k_1 k_2}\left(\frac{k_1}{k_2} + \frac{k_2}{k_1}\right), \tag{C.20}$$

$$
\begin{aligned}
F_3(\boldsymbol{k_1}, \boldsymbol{k_2}) = \frac{k_1^2}{252}&\left[12\left(\frac{k_1^2}{k_2^2}\right) - 158 + 100\left(\frac{k_2}{k_1}\right)^2 - 42\left(\frac{k_2}{k_1}\right)^4 + \right.\\
&\left. \frac{3}{k_1^5 k_2^3}(k_2^2 - k_1^2)^3(2k_1^2 + 7k_2^2)\ln\left(\frac{k_1 + k_2}{k_1 - k_2}\right)\right].
\end{aligned} \tag{C.21}
$$

Bibliography

[1] I. Tereno, C. S. Carvalho, J. Dinis, R. Scaramella, J. Amiaux, C. Burigana et al., *Euclid Space Mission: building the sky survey*, Proc. IAU **10** (May, 2014) 379–381.

[2] T. S. O. Collaboration, P. Ade, J. Aguirre, Z. Ahmed, S. Aiola, A. Ali et al., *The Simons Observatory: Science goals and forecasts*, J. Cosmol. Astropart. Phys. **2019** (Feb., 2019) 056–056.

[3] T. Thmmler, *Introduction to direct neutrino mass measurements and KATRIN*, Nuclear Physics B - Proceedings Supplements **229-232** (Aug., 2012) 146–151.

[4] M. J. Dolinski, A. W. P. Poon and W. Rodejohann, *Neutrinoless Double-Beta Decay: Status and Prospects*, arXiv:1902.04097 [hep-ex, physics:hep-ph, physics:nucl-ex, physics:nucl-th] (Feb., 2019) .

[5] P. Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi et al., *Planck 2018 results. VI. Cosmological parameters*, arXiv:1807.06209 [astro-ph] (July, 2018) .

[6] P. Bull, Y. Akrami, J. Adamek, T. Baker, E. Bellini, J. B. Jimnez et al., *Beyond ΛCDM: Problems, solutions, and the road ahead*, Physics of the Dark Universe **12** (June, 2016) 56–99.

[7] J. Lesgourgues, G. Mangano, G. Miele and S. Pastor, *Neutrino Cosmology*. Cambridge University Press, 2013.

[8] J. Lesgourgues and S. Pastor, *Massive neutrinos and cosmology*, Physics Reports **429** (July, 2006) 307–379.

[9] D. Blas, J. Lesgourgues and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS) II: Approximation schemes*, J. Cosmol. Astropart. Phys. **2011** (July, 2011) 034–034.

[10] L. Anderson, E. Aubourg, S. Bailey, D. Bizyaev, M. Blanton, A. S. Bolton et al., *The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations in the Data Release 9 Spectroscopic Galaxy Sample*, Monthly Notices of the Royal Astronomical Society **427** (Dec., 2012) 3435–3467.

[11] S. Alam, M. Ata, S. Bailey, F. Beutler, D. Bizyaev, J. A. Blazek et al., *The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample*, *Monthly Notices of the Royal Astronomical Society* **470** (Sept., 2017) 2617–2652.

[12] D. J. Eisenstein, I. Zehavi, D. W. Hogg, R. Scoccimarro, M. R. Blanton, R. C. Nichol et al., *Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies*, *ApJ* **633** (Nov., 2005) 560–574.

[13] S. Cole, W. J. Percival, J. A. Peacock, P. Norberg, C. M. Baugh, C. S. Frenk et al., *The 2df Galaxy Redshift Survey: Power-spectrum analysis of the final dataset and cosmological implications*, *Monthly Notices of the Royal Astronomical Society* **362** (Sept., 2005) 505–534.

[14] B. A. Bassett and R. Hlozek, *Baryon Acoustic Oscillations*, *arXiv:0910.5224 [astro-ph, physics:gr-qc, physics:hep-ex, physics:hep-ph]* (Oct., 2009) .

[15] N. Kaiser, *Clustering in real space and in redshift space*, *Mon.Not.Roy.Astron.Soc.* **227** (1987) 1–27.

[16] C. Alcock and B. Paczynski, *An evolution free test for non-zero cosmological constant*, *Nature* **281** (1979) 358–359.

[17] A. Lewis and A. Challinor, *Weak Gravitational Lensing of the CMB*, *Physics Reports* **429** (June, 2006) 1–65.

[18] F. Hassani, S. Baghram and H. Firouzjahi, *Lensing as a Probe of Early Universe: from CMB to Galaxies*, *J. Cosmol. Astropart. Phys.* **2016** (May, 2016) 044–044.

[19] V. Desjacques, D. Jeong and F. Schmidt, *The Galaxy Power Spectrum and Bispectrum in Redshift Space*, *J. Cosmol. Astropart. Phys.* **2018** (Dec., 2018) 035–035.

[20] V. Desjacques, D. Jeong and F. Schmidt, *Large-Scale Galaxy Bias*, *Physics Reports* **733** (Feb., 2018) 1–193.

[21] R. Allison, P. Caucal, E. Calabrese, J. Dunkley and T. Louis, *Towards a cosmological neutrino mass detection*, *Phys. Rev. D* **92** (Dec., 2015) 123535.

[22] J. R. Bond, G. Efstathiou and M. Tegmark, *Forecasting Cosmic Parameter Errors from Microwave Background Anisotropy Experiments*, *arXiv:astro-ph/9702100* (Feb., 1997) .

[23] L. Amendola, S. Appleby, A. Avgoustidis, D. Bacon, T. Baker, M. Baldi et al., *Cosmology and Fundamental Physics with the Euclid Satellite*, *Living Rev Relativ* **21** (Dec., 2018) 2.

[24] R. Laureijs, J. Amiaux, S. Arduini, J.-L. Augures, J. Brinchmann, R. Cole et al., *Euclid Definition Study Report*, *arXiv:1110.3193 [astro-ph]* (Oct., 2011) .

[25] A. Boyle and E. Komatsu, *Deconstructing the neutrino mass constraint from galaxy redshift surveys*, *J. Cosmol. Astropart. Phys.* **2018** (Mar., 2018) 035–035.

[26] Z. Pan and L. Knox, *Constraints on neutrino mass from Cosmic Microwave Background and Large Scale Structure*, *Mon. Not. R. Astron. Soc.* **454** (Dec., 2015) 3200–3206.

[27] O. F. Hernndez, *Neutrino Masses, Scale-Dependent Growth, and Redshift-Space Distortions*, *J. Cosmol. Astropart. Phys.* **2017** (June, 2017) 018–018.

[28] P. Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont et al., *Planck 2015 results. XIV. Dark energy and modified gravity*, *A&A* **594** (Oct., 2016) A14.

[29] P. Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont et al., *Planck 2015 results. XIII. Cosmological parameters*, *A&A* **594** (Oct., 2016) A13.

[30] X. Qian and P. Vogel, *Neutrino Mass Hierarchy*, *Progress in Particle and Nuclear Physics* **83** (July, 2015) 1–30.

[31] W. J. Percival, *Large Scale Structure Observations*, *arXiv:1312.5490 [astro-ph]* (Dec., 2013) .

[32] A. Font-Ribera, P. McDonald, N. Mostek, B. A. Reid, H.-J. Seo and A. Slosar, *DESI and other dark energy experiments in the era of neutrino mass measurements*, *J. Cosmol. Astropart. Phys.* **2014** (May, 2014) 023–023.

[33] D. J. Eisenstein, H.-j. Seo and M. White, *On the Robustness of the Acoustic Scale in the Low-Redshift Clustering of Matter*, *ApJ* **664** (Aug., 2007) 660–674.

[34] H.-J. Seo and D. J. Eisenstein, *Improved forecasts for the baryon acoustic oscillations and cosmological distance scale*, *ApJ* **665** (Aug., 2007) 14–24.

[35] D. J. Eisenstein, H.-j. Seo, E. Sirko and D. Spergel, *Improving Cosmological Distance Measurements by Reconstruction of the Baryon Acoustic Peak*, *ApJ* **664** (Aug., 2007) 675–679.

[36] A. Boyle, *Understanding the neutrino mass constraints achievable by combining CMB lensing and spectroscopic galaxy surveys*, *J. Cosmol. Astropart. Phys.* **2019** (Apr., 2019) 038–038.

[37] D. J. Eisenstein and W. Hu, *Baryonic Features in the Matter Transfer Function*, *ApJ* **496** (Apr., 1998) 605–614.

[38] A. Savitzky and M. J. E. Golay, *Smoothing and Differentiation of Data by Simplified Least Squares Procedures.*, *Anal. Chem.* **36** (July, 1964) 1627–1639.

[39] M. Shoji, D. Jeong and E. Komatsu, *Extracting Angular Diameter Distance and Expansion Rate of the Universe from Two-dimensional Galaxy Power Spectrum at High Redshifts: Baryon Acoustic Oscillation Fitting versus Full Modeling*, *ApJ* **693** (Mar., 2009) 1404–1416.

[40] M. White, Y.-S. Song and W. J. Percival, *Forecasting Cosmological Constraints from Redshift Surveys*, *Monthly Notices of the Royal Astronomical Society* **397** (Aug., 2009) 1348–1354.

[41] M. Archidiacono, T. Brinckmann, J. Lesgourgues and V. Poulin, *Physical effects involved in the measurements of neutrino masses with future cosmological data*, *J. Cosmol. Astropart. Phys.* **2017** (Feb., 2017) 052–052.

[42] M. Baldi, F. Villaescusa-Navarro, M. Viel, E. Puchwein, V. Springel and L. Moscardini, *Cosmic Degeneracies I: Joint N-body Simulations of Modified Gravity and Massive Neutrinos*, *Monthly Notices of the Royal Astronomical Society* **440** (May, 2014) 75–88.

[43] W. L. K. Wu, J. Errard, C. Dvorkin, C. L. Kuo, A. T. Lee, P. McDonald et al., *A Guide to Designing Future Ground-based CMB Experiments*, *ApJ* **788** (June, 2014) 138.

[44] K. N. Abazajian, K. Arnold, J. Austermann, B. A. Benson, C. Bischoff, J. Bock et al., *Neutrino Physics from the Cosmic Microwave Background and Large Scale Structure*, *Astroparticle Physics* **63** (Mar., 2015) 66–80.

[45] P. Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown et al., *Planck 2013 results. XVI. Cosmological parameters*, *A&A* **571** (Nov., 2014) A16.

[46] A. Liu, J. R. Pritchard, R. Allison, A. R. Parsons, U. Seljak and B. D. Sherwin, *Eliminating the optical depth nuisance from the CMB with 21 cm cosmology*, *Phys. Rev. D* **93** (Feb., 2016) 043013.

[47] P. Collaboration, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi et al., *Planck 2015 results. XI. CMB power spectra, likelihoods, and robustness of parameters*, *A&A* **594** (Oct., 2016) A11.

[48] M. Gerbino, M. Lattanzi, O. Mena and K. Freese, *A novel approach to quantifying the sensitivity of current and future cosmological datasets to the neutrino mass ordering through Bayesian hierarchical modeling*, *Physics Letters B* **775** (Dec., 2017) 239–250.

[49] B. Yu, R. Z. Knight, B. D. Sherwin, S. Ferraro, L. Knox and M. Schmittfull, *Towards Neutrino Mass from Cosmology without Optical Depth Information*, *arXiv:1809.02120 [astro-ph]* (Sept., 2018) .

[50] J. Peloton, M. Schmittfull, A. Lewis, J. Carron and O. Zahn, *Full covariance of CMB and lensing reconstruction power spectra*, *Phys. Rev. D* **95** (Feb., 2017) 043508.

[51] J. Rocher, K. Benabed and F. Bouchet, *Probing inflation with CMB polarization : weak lensing effect on the covariance of CMB spectra*, *J. Cosmol. Astropart. Phys.* **2007** (May, 2007) 013–013.

[52] D. Jeong, E. Komatsu and B. Jain, *Galaxy-CMB and galaxy-galaxy lensing on large scales: sensitivity to primordial non-Gaussianity*, *Phys. Rev. D* **80** (Dec., 2009) 123527.

[53] D. W. Hogg, *Distance measures in cosmology*, *arXiv:astro-ph/9905116* (May, 1999) .

[54] T. Okamoto and W. Hu, *CMB Lensing Reconstruction on the Full Sky*, *Phys. Rev. D* **67** (Apr., 2003) 083002.

[55] L. Perotto, J. Lesgourgues, S. Hannestad, H. Tu and Y. Y. Y. Wong, *Probing cosmological parameters with the CMB: Forecasts from full Monte Carlo simulations*, *J. Cosmol. Astropart. Phys.* **2006** (Oct., 2006) 013–013.

[56] E. Giusarma, S. Vagnozzi, S. Ho, S. Ferraro, K. Freese, R. Kamen-Rubio et al., *Scale-dependent galaxy bias, CMB lensing-galaxy cross-correlation, and neutrino masses*, *Phys. Rev. D* **98** (Dec., 2018) 123526.

[57] B. D. Sherwin, J. Dunkley, S. Das, J. W. Appel, J. R. Bond, C. S. Carvalho et al., *Evidence for dark energy from the cosmic microwave background alone using the Atacama Cosmology Telescope lensing measurements*, *Phys. Rev. Lett.* **107** (July, 2011) 021302.

[58] T. Sprenger, M. Archidiacono, T. Brinckmann, S. Clesse and J. Lesgourgues, *Cosmology in the era of Euclid and the Square Kilometre Array*, *J. Cosmol. Astropart. Phys.* **2019** (Feb., 2019) 047–047.

[59] T. Brinckmann, D. C. Hooper, M. Archidiacono, J. Lesgourgues and T. Sprenger, *The promising future of a robust cosmological neutrino mass measurement*, *J. Cosmol. Astropart. Phys.* **2019** (Jan., 2019) 059–059.

[60] M. Mirbabayi, F. Schmidt and M. Zaldarriaga, *Biased Tracers and Time Evolution*, *J. Cosmol. Astropart. Phys.* **2015** (July, 2015) 030–030.

[61] B. Jain and E. Bertschinger, *Second Order Power Spectrum and Nonlinear Evolution at High Redshift*, *ApJ* **431** (Aug., 1994) 495.

[62] J. L. Tinker, A. V. Kravtsov, A. Klypin, K. Abazajian, M. S. Warren, G. Yepes et al., *Toward a halo mass function for precision cosmology: the limits of universality*, Astrophysical Journal **688** (Dec., 2008) 709–728.

[63] T. Lazeyras, C. Wagner, T. Baldauf and F. Schmidt, *Precision measurement of the local bias of dark matter halos*, J. Cosmol. Astropart. Phys. **2016** (Feb., 2016) 018–018.

[64] S. Mishra-Sharma, D. Alonso and J. Dunkley, *Neutrino masses and beyond-ΛCDM cosmology with LSST and future CMB experiments*, Phys. Rev. D **97** (June, 2018) 123544.

[65] C. M. Hirata, *Tidal alignments as a contaminant of redshift space distortions*, Monthly Notices of the Royal Astronomical Society **399** (Oct., 2009) 1074–1087.

[66] M. LoVerde, *Halo bias in mixed dark matter cosmologies*, Phys. Rev. D **90** (Oct., 2014) 083530.

[67] E. Castorina, E. Sefusatti, R. K. Sheth, F. Villaescusa-Navarro and M. Viel, *Cosmology with massive neutrinos II: on the universality of the halo mass function and bias*, J. Cosmol. Astropart. Phys. **2014** (Feb., 2014) 049–049.

[68] E. Castorina, C. Carbone, J. Bel, E. Sefusatti and K. Dolag, *DEMNUni: The clustering of large-scale structures in the presence of massive neutrinos*, J. Cosmol. Astropart. Phys. **2015** (July, 2015) 043–043.

[69] F. Villaescusa-Navarro, A. Banerjee, N. Dalal, E. Castorina, R. Scoccimarro, R. Angulo et al., *The imprint of neutrinos on clustering in redshift-space*, ApJ **861** (July, 2018) 53.

[70] A. Chudaykin and M. M. Ivanov, *Measuring neutrino masses with large-scale structure: Euclid forecast with controlled theoretical error*, arXiv:1907.06666 [astro-ph, physics:hep-ph] (July, 2019) .

[71] G. J. Hill, K. Gebhardt, E. Komatsu, N. Drory, P. J. MacQueen, J. Adams et al., *The Hobby-Eberly Telescope Dark Energy Experiment (HETDEX): Description and Early Pilot Survey Results*, arXiv:0806.0183 [astro-ph] (June, 2008) .

[72] A. S. Leung, V. Acquaviva, E. Gawiser, R. Ciardullo, E. Komatsu, A. I. Malz et al., *Bayesian Redshift Classification of Emission-line Galaxies with Photometric Equivalent Widths*, ApJ **843** (July, 2017) 130.

[73] M. Takada, R. Ellis, M. Chiba, J. E. Greene, H. Aihara, N. Arimoto et al., *Extragalactic Science, Cosmology and Galactic Archaeology with the Subaru Prime Focus Spectrograph (PFS)*, Publ. Astron. Soc. Jpn **66** (Feb., 2014) R1.

[74] D. Collaboration, A. Aghamousa, J. Aguilar, S. Ahlen, S. Alam, L. E. Allen et al., *The DESI Experiment Part I: Science,Targeting, and Survey Design*, *arXiv:1611.00036 [astro-ph]* (Oct., 2016) .

[75] D. Spergel, N. Gehrels, J. Breckinridge, M. Donahue, A. Dressler, B. S. Gaudi et al., *Wide-Field InfraRed Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA Final Report, arXiv:1305.5422 [astro-ph]* (May, 2013) .

[76] J. Green, P. Schechter, C. Baltay, R. Bean, D. Bennett, R. Brown et al., *Wide-Field InfraRed Survey Telescope (WFIRST) Final Report, arXiv:1208.4012 [astro-ph]* (Aug., 2012) .

[77] V. Assassi, D. Baumann, E. Pajer, Y. Welling and D. van der Woude, *Effective Theory of Large-Scale Structure with Primordial Non-Gaussianity, J. Cosmol. Astropart. Phys.* **2015** (Nov., 2015) 024–024.

[78] N. Makino, M. Sasaki and Y. Suto, *Analytic approach to the perturbative expansion of nonlinear gravitational fluctuations in cosmological density and velocity fields, Phys. Rev. D* **46** (July, 1992) 585–602.

Acknowledgements

It is hard to do justice in such a small space to all of the people who have provided indispensable support to me in the last four years. Firstly I would like to thank my supervisor Eiichiro Komatsu, not only for his outstanding scientific guidance, but also for his faith in me, his patience and his reassurance. And equally importantly, a lot of fun!

I would also like to especially thank Fabian Schmidt, with whom I worked on the final chapter of this thesis. Without him always making time to help with my questions, and without his always insightful suggestions, this thesis would have been much less than it is.

I am also very grateful to Stuart Sim and Wolfgang Kerzendorf, both previous supervisors and collaborators who have now become wonderful mentors.

Philipp Busch and Tadeja Veršič deserve a very special mention for generally bringing light to my life, and especially for their indispensable, patient moral support in the difficult last six months of my PhD. Thank you for having faith in me when I sometimes did not.

I have been blessed with many wonderful friends in my time at MPA. I would like to thank Matteo Frigo, Inh Jee, Aniket Agrawal, Christina Kreisch, Leila Mirzagholi, Anabele Pardi, Vlas Sokolov, Linda Baronchelli, Philipp Plewa, Jens Stücker, Marius Berge Eide, Timo Halbesma, Malin Renneby, Max Eisenreich, Miranda Jarvis, Raffaella Capasso and Ricard Ardevol for some wonderful memories. I also have some long-distance friends to thank, who somehow still never feel so far away - Rachael McNeill, Robert McNeill and Bronagh Magee.

Last but not least, I have to thank my family, not just for the past few years but for my whole life. I would like to thank my parents and my nana for their unconditional love and unending moral support. I would also like to thank my siblings, Aine, Padraig and Bridgin. Although I am the eldest, you have taught me much more than I could ever teach you.