How Heavy Neutral Leptons Influence the Neutrino Mass Bound

Ploeg, Marco van der

How Heavy Neutral Leptons Influence the Neutrino Mass Bound

THESIS

submitted in partial fulfillment of the
requirements for the degree of

BACHELOR OF SCIENCE
in
PHYSICS

Author :	Marco van der Ploeg
Student ID :	2629283
Supervisor :	Alexey Boyarsky
Second corrector :	Subodh Patil

Leiden, The Netherlands, June 17, 2022

How Heavy Neutral Leptons Influence the Neutrino Mass Bound

Marco van der Ploeg

Instituut-Lorentz, Leiden University
P.O. Box 9500, 2300 RA Leiden, The Netherlands

June 17, 2022

Abstract

The ever-increasing accuracy of telescopes like *Planck* has allowed for the measurements of anisotropies in the CMB. For this project, the influence of heavy neutral leptons (HNLs) on the upper bound of the neutrino mass sum, as determined from the CMB anisotropies and additional datasets, has been studied. Specifically, an HNL with mass $m_N \sim 500$ MeV and lifetime $\tau_N \sim 0.05$ s is considered, which would cause a change of the number of relativistic species to $N_{\text{eff}} = 2.45$ and of the primordial helium fraction to $Y_P = 0.26$. With these conditions, the mass bound is found to be $\sum m_\nu < 0.12$ at 95% CL, while the bound in standard ΛCDM + $\sum m_\nu$ resulted in $\sum m_\nu < 0.13$ eV at 95% CL. The bound thus decreases slightly. This can be explained if N_{eff} and $\sum m_\nu$ have distinct effects, and other parameters' effects are able to counter changes due to the change in N_{eff}, so that the mass sum does not need to change much. HNLs thus have a tiny effect on the mass bound, but there may also be no effect, as the results could be slightly off if we have not run the fitting program long enough.

Contents

Chapter 1

Introduction

In the last 2-3 decades, we have entered the realm of precision cosmology, through measurements of the Cosmic Microwave Background (CMB). Due to more and more precise measurements by telescopes like WMAP and *Planck*, it has become possible to study the CMB in greater and greater detail. These measurements have revealed that at temperatures of $\sim 10^{-4}$ K, the CMB spectrum deviates from the Planck spectrum. These anisotropies reveal a great many things about the Universe. They allow one to determine cosmological parameters, for instance the density fraction $\Omega_x = \rho_x/\rho_{crit}$ for particle species x, like baryons, cold dark matter, or neutrinos. For this project, we use the influence of neutrinos on the anisotropies to study an unsolved problem.

This problem fits in both cosmology and particle physics, and concerns the masses of the neutrinos. Although it is known that their masses must be very small, they have not yet been determined. We know they must have masses in order to explain neutrino oscillation experiments. These experiments can determine the difference between the different neutrino masses, $\Delta m^2_{21} = m^2_2 - m^2_1 \approx 7.62 \cdot 10^{-5}$ eV^2 and $|\Delta m^2_{31}| = |m^2_3 - m^2_1| \approx 2.55 \cdot 10^{-3}$ eV^2 [17] (note the absolute value signs: we do not yet know the sign of Δm^2_{31}!). Additionally, this shows a bigger problem: that the Standard Model (SM) of particle physics is not complete, as the SM requires neutrinos to be massless.

Now cosmology allows us to also determine an upper bound on the sum of the neutrino masses $\sum m_\nu = m_1 + m_2 + m_3$. This may be done by considering $\sum m_\nu$ as one of the cosmological parameters in our model of the Universe. It then has its own influence on the CMB, from which it can be determined. Recently the Planck Collaboration determined an upper bound on the neutrino masses of $\sum m_\nu < 0.12$ eV at 95% confidence [18].

In this thesis, we will consider how the upper bound of the neutrino mass sum is affected by adding a new type of particle, namely the Heavy Neutral Lepton (HNL), to the Standard Model. An HNL is basically a heavier version of the neutrino, and is also referred to as 'sterile neutrino' or 'right-handed neutrino'. It has a right-handed spin, in contrast to the known neutrinos which always have a left-handed spin. Adding HNLs gives us better insight into an extended SM and into the neutrino masses. This follows a recent trend of considering the effects of adding HNLs to the SM to solve its incompleteness; HNLs could also explain why neutrinos have mass, why neutrino oscillations occur, why there is baryon asymmetry in the Universe, and what dark matter is [7]. By considering their effect on the neutrino mass bound, we may add to the total knowledge of HNLs, and determine if their effect is present in the cosmological data, as a way to prove if they actually exist. For this project, we specifically look at an HNL with mass $m_N \sim 500$ MeV and lifetime $\tau_N \sim 0.05$ s.

We will start by reviewing the necessary background theory of the nature of the CMB anisotropies, the characteristics of the matter power spectrum and the ΛCDM model (Chapters 2, 3 and 4). Then we study the influence of neutrinos on the CMB anisotropies, and we look at the effect of HNLs on neutrinos (Chapter 5 and 6). Next the methods are described that are used for this project to determine the parameters of interest (Chapter 7). The results of this are shown in Chapter 8, and discussed in Chapter 9. Finally we conclude in Chapter 10.

Chapter 2

The CMB and its Anisotropies

The CMB spectrum we observe now is well described by a blackbody at a mean temperature of $T_0 = 2.7255$ K (or 0.234 meV) [18], so that it peaks at wavelengths of about 1.06 mm. It was cooled due to gravitational redshift. The mean temperature allows us to determine the photon number density $n_\gamma = 2\zeta(3)T_0^3/\pi^2$ and the energy density $\rho_\gamma = (\pi^2/15)T_0^4$, from which the density fraction $\Omega_\gamma = \rho_\gamma/\rho_{crit}$ can be determined. The high homogeneity/isotropy of the CMB tells us that the Universe is also highly homogeneous/isotropic. It however also contains small anisotropies.

The anisotropies are described by a function expansion

$$\frac{\Delta T(\theta,\phi)}{T_0} = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\theta,\phi), \tag{2.1}$$

where the $Y_{\ell m}(\theta,\phi)$ are the spherical harmonics, $\Delta T(\theta,\phi) = T(\theta,\phi) - T_0$ are (just) the anisotropies and $T(\theta,\phi)$ is the measured CMB temperature, including anisotropies.

This expansion is useful, since the CMB is measured on a sphere in the sky (as it comes from the surface of last scattering), and can be seen as the sum of many independent oscillations. Spherical harmonics are like the 2D versions of such oscillations, and well-defined on a spherical surface. No direct relation exists between ℓ and θ, but the spherical harmonic $Y_{\ell m}(\theta,\phi)$ describes angular variations of a size $\theta \sim \pi/\ell$.

The value of ℓ gives the multipole: $\ell = 0$ is the monopole, which can be considered as the mean temperature of the CMB. For this reason, $\ell = 0$ is omitted in the sum in formula 2.1, as the mean has been subtracted already. (If we did start the sum at $\ell = 0$, we would have to set $a_{00} = 0$, so we might as well omit it.) $\ell = 1$ is the dipole, which corresponds to the

Doppler shift in the spectrum and tells us about the velocity of the Solar System relative to the CMB. This term could also be omitted in formula 2.1, if the dipole term is subtracted from $T(\theta, \phi)$ and the sum starts at $\ell = 2$. For multipoles $\ell > 1$, the effects are of cosmological origin, and therefore of great interest to us.

But first, let us consider the anisotropy data itself. The usual way of looking at the anisotropies is by using the angular power spectrum, given by $\ell(\ell+1)C_\ell/2\pi$, where $C_\ell = \frac{1}{2\ell+1}\sum_m |a_{\ell m}|^2$. The power spectrum, as measured by multiple telescopes, is plotted in figure 2.1. One way to think about the power is that it counts how many hotter/colder spots there are (relative to the average temperature T_0) within a certain angular scale. The power spectrum is then like a histogram of these counts. Here we look at the TT (T from Temperature) mode; there are also the TE and EE modes, E coming from an analogy with electric fields. The TT mode describes anisotropies in the temperature, and is the strongest signal. The EE mode describe anisotropies in the polarisation of the CMB due to Thomson scattering [18]. The TE mode describes correlations between temperature and polarisation.

2.1 Sachs-Wolfe effect

The anisotropies in the spectrum for multipoles $\ell \lesssim 100$ are caused by the so-called Sachs-Wolfe effect. There are two versions of this; firstly, the non-integrated Sachs-Wolfe effect. This is caused by gravitational redshift occurring at the surface of last scattering. As the energy distribution at last scattering is not homogeneous, the redshift varies too. For a perturbation $\delta\phi$ of the potential, the resulting temperature fluctuation is $\delta T/T_0 \approx (1/3)\delta\phi/c^2$ [18].

Secondly, there is the integrated Sachs-Wolfe (ISW) effect. This is again due to redshift, but now the gravity causing this comes from matter that the photons pass by during their journey from the last scattering surface to our telescopes. As the photons move through the potential well, they are red- or blue-shifted. However, if the potential is not constant while the photon crosses it, the amount of blue- and red-shifting while the photon enters and leaves does not cancel each other out, causing a net shift [14]. Since the distribution of energy is not homogeneous for each path a photon could travel, this causes anisotropies.

The ISW effect can again be divided in the early and late ISW effect [13]. The early ISW effect occurs if the Universe is not completely matter dominated at last scattering; then radiation pressure causes the gravitational

Version of June 17, 2022– Created June 17, 2022 - 11:38

potentials at the surface to decay. This decay is what changes the potential, resulting in the ISW effect. The late ISW effect again happens due to the decay of potentials, but now the decay is caused by dark energy. Once the Universe becomes dark energy dominated, this decay starts.

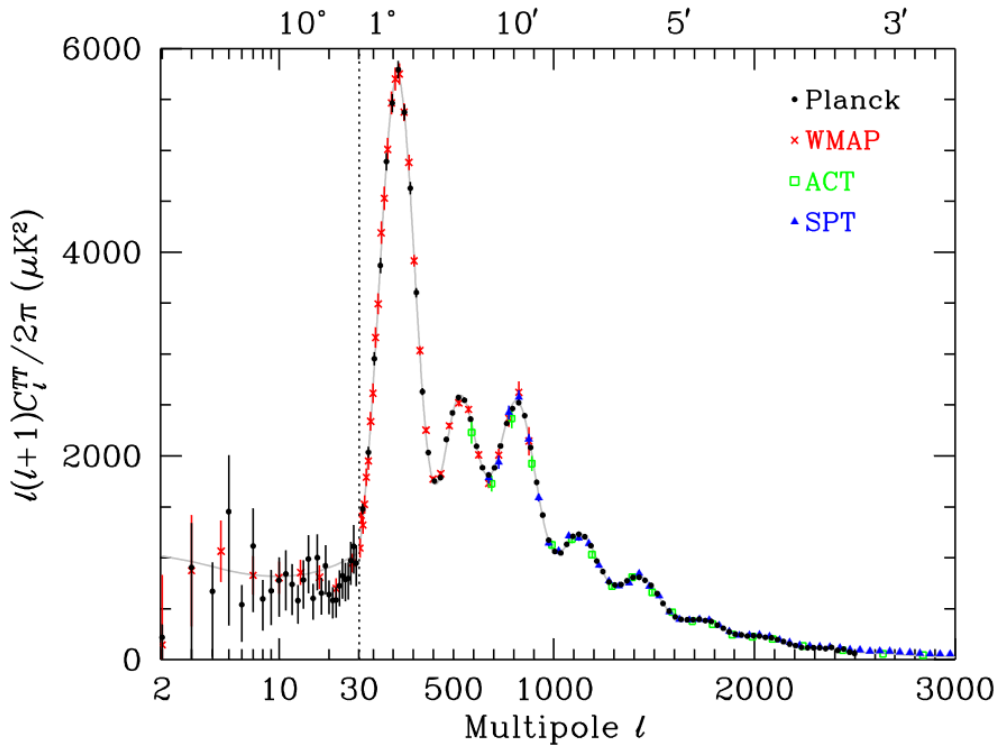

Figure 2.1: *The perturbations in the CMB power spectrum as a function of multipole ℓ. A fit is plotted over data from 4 telescopes. Here we look at the TT mode. Figure taken from [18].*

2.2 Baryon Acoustic Oscillations

For $100 \lesssim \ell \lesssim 1000$, one can see the effect of baryon acoustic oscillations (BAO). This effect is caused by the interplay of photon pressure and gravitational attraction on baryons, before recombination. The gravitational attraction pulls baryons together, but this increases the photon pressure on them. Once this pressure gets strong enough, it will drive the baryons apart again. Then they will be pulled together again by gravity, and so on, so oscillations occur. The system can be thought of as a mass on a spring, with the mass being the effective mass of the baryon-photon fluid and the spring being the photon pressure. Then when recombination occurs, the

photons are released from the oscillations. Their density at the moment of recombination is thus frozen. Since the photon's temperature is related to the number density, we can see these oscillations as anisotropies in the CMB. The odd-numbered peaks in the spectrum represent baryons that were at maximal compression and the even peaks represent baryons that were at maximal rarefaction at the moment of recombination. Generally the odd peaks are higher than the even ones because the rarefaction also has to counter the inertia the baryons built up [24].

2.3 Damping tail

For $\ell \gtrsim 1000$, there is a damping tail which shows that recombination was not instantaneous. This is because non-instantaneous recombination means that the surface of last scattering is not really a surface, but has a certain thickness. The photons must then escape through this thickness, which causes some of them to be scattered. This transfers energy between hot and cold regions, smoothing out the anisotropies. Thus the fluctuations damp, with the damping increasing exponentially as the photon diffusion length gets larger than the wavelength. The anisotropies disappear at scales smaller than the last scattering surface's thickness, corresponding to high ℓ.

The damping tail also has an angular scale θ_D associated with it. This is the angular distance that photons could travel during recombination. Before recombination, the photons couldn't get anywhere far due to the strong baryon-photon coupling. But while recombination was happening, the photons were able to travel further and further. θ_D corresponds to the average mean free path of the photons in this period.

It can be calculated from $\theta_D = r_D/D_M$ [24], where D_M is the comoving angular diameter distance to recombination, $D_M = (1+z)D_A = D_A/a(z)$ for D_A the usual angular distance (which, in case of a far away object and a flat Universe, is the real distance to the object from Earth). D_M is the distance to the object if the Universe were not expanding; comoving coordinates take out this expansion. We can also write $D_M = \int_0^{z_{dec}} cdz/H(z)$ [24]. And r_D is the damping scale, given by

$$r_D^2 = \pi^2 \int_0^{a_{dec}} \frac{da}{a^3 \sigma_T n_e(a) H(a)} \frac{R^2 + \frac{16}{15}(1+R)}{6(1+R^2)}, \tag{2.2}$$

with $R = \frac{p_b+\rho_b}{p_\gamma+\rho_\gamma} = 3\rho_b/4\rho_\gamma$ the baryon-to-photon momentum density ratio. Then in the spectrum, the damping tail is described by a term $e^{-(\theta_D \ell)^2}$.

Version of June 17, 2022– Created June 17, 2022 - 11:38

In figure 2.2, we see the power spectrum again, as predicted by the ΛCDM model (which will be introduced in chapter 4), with the described effects noted.

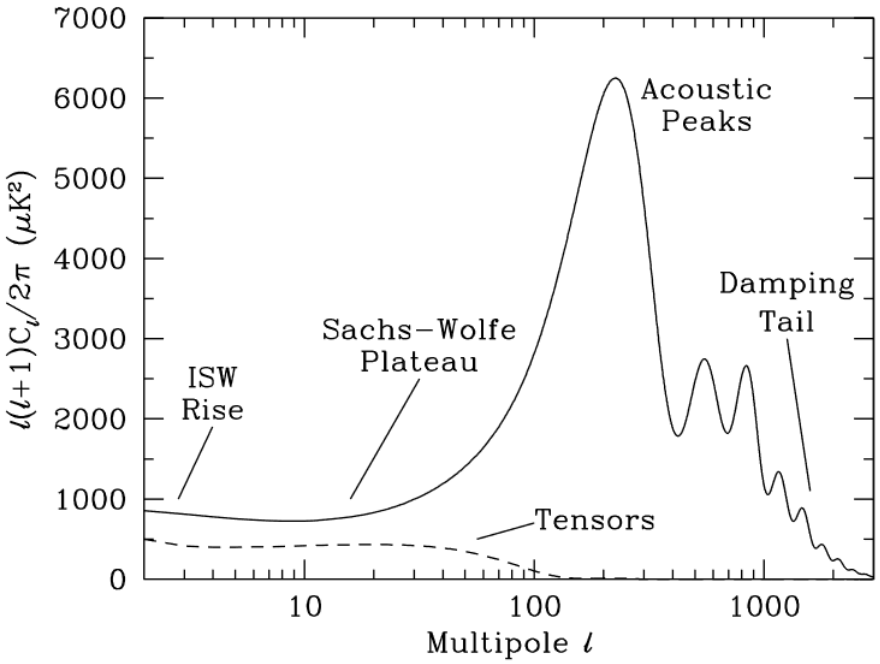

Figure 2.2: *The power spectrum of the CMB anisotropies, with each effect noted. The ISW rise is the same as the ISW effect, and the Sachs-Wolfe plateau shows where the ISW effect occurs. The 'tensor' contribution, accounting for influence of gravitational waves (mathematically described with tensors), is neglected for our purposes, as it is much smaller than the other (scalar) modes [18]. Image taken from [23].*

Chapter 3

The Matter Power Spectrum

Similar to the CMB anisotropies, measurements of the matter power spectrum $P(k, z)$ (also sometimes $P_m(k, z)$) can also help to constrain the cosmological parameters. The matter power spectrum describes the variations of matter densities in the Universe as a function of scale k (in Fourier space). Where the CMB angular power spectrum characterised temperature differences on an angular scale, the matter power spectrum thus characterises density differences on a distance scale. Now since we can observe galaxies at different redshift, while the CMB is fixed to one redshift at last scattering, P is also a function of z.

3.1 Measurements

In order to determine the matter power spectrum, we first need the overdensity $\delta_g(\vec{x}) = \frac{\rho_g(\vec{x}) - \overline{\rho_g}}{\overline{\rho_g}}$, where $\rho_g(\vec{x})$ is the number density of galaxies and $\overline{\rho_g}$ is the mean galaxy density [24].

The overdensity can be calculated using Large Scale Structure (LSS) measurements. These are measurements of galaxies' redshift, together with their angular position in the sky. This allows us to determine the position of the galaxies, using redshift as a distance measurement. From these positions, the densities can be found. Then we use that the Fourier transform of the overdensity is proportional to the galaxy power spectrum $P_g(k, z)$. This power spectrum describes variations in density of galaxies instead of matter. It is related to the matter power spectrum through $P_g(k, z) = b(z)^2 \cdot P(k, z)$, for some bias function $b(z)$ which depends on the specific LSS measurements [24].

Now these LSS measurements only constrain a certain k-range. For

Version of June 17, 2022– Created June 17, 2022 - 11:38

other ranges, we may use Lyα, X-ray, or CMB (weak lensing and SZ effect) measurements. Firstly, it is possible to determine the spectrum using data from the Lyman-α forest (Lyα). These are measurements of the absorption line of the Lyman-α transition in hydrogen; there are multiple lines because the spectra from different sources are redshifted differently. By Fourier expanding these spectra, the flux power spectrum $P_F(z,k)$ can be obtained. This can then be related to the matter power spectrum $P(k,z)$; again with a bias, in a more complicated form $P_F(k) = b^2[k,P]P(k)$ [11]. However, P_F only gives information about mildly non-linear scales, whereas P is for linear scales. This makes relating the two very complicated [17], decreasing constraining strength. Lyα can typically determine the matter power spectrum in ranges $0.3 < k < 3\ h/\text{Mpc}$ and $2 < z < 5$ [17].

Other observations that helps constrain the spectrum are those of X-rays, emitted by hot gas in galaxy clusters, measured by e.g. the ROSAT All-Sky Survey. Additionally, the CMB measurements include the Sunyaev-Zeldovitch (SZ) effect. The SZ effect is the increase in energy of the CMB photons after scattering off of hot electrons in galaxy clusters. X-ray and SZ effect measurements can be used to count galaxy clusters. From these, like LSS measurements, the overdensity δ_g can be determined to find the matter power spectrum. Additionally, the CMB measurements also include weak lensing effects. Large matter densities deform the CMB radiation through gravitational lensing, affecting the C_ℓ. This can be transformed into $P(k,z)$. The transformation however depends on which model is used.

In figure 3.1 we see a plot showing measurements of the matter power spectrum by *Planck* and other probes.

3.2 Characteristics

Let us now discuss some of the characteristic aspects of the matter power spectrum. The maximum of the spectrum, where it goes from increasing to decreasing, is at the so-called pivot point wavenumber $k_0 \approx 0.05\ \text{Mpc}^{-1}$. Then in formula, the primordial spectrum is $P_R(k) = A_s(k/k_0)^{n_s-1}$ [21]. This primordial spectrum is only a part of the full spectrum; we may write $P(k) \propto P_R(k)$. The primordial matter power spectrum is described by the same parameters A_s and n_s as the primordial CMB spectrum, which we will see in more detail in section 4.3. So we see that $P(k_0) = A_s$, regardless of n_s: at this scale the amplitude can be measured. The choice of $k_0 \approx 0.05\ \text{Mpc}^{-1}$ is convention, often chosen as the scale that is best constrained by

observations [10]. For $k > k_0$, the effect of BAO on this spectrum is visible as oscillations of the line.

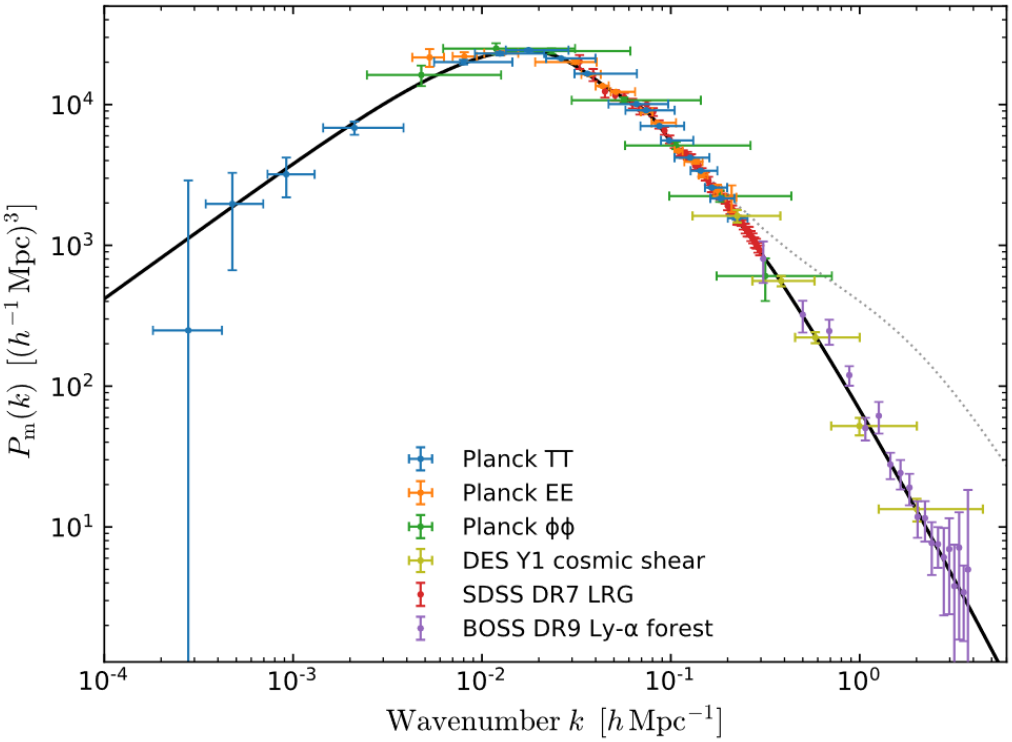

Figure 3.1: *The matter power spectrum, showing perturbations in the matter density of the Universe as a function of scale k. A ΛCDM fit is plotted through the data. Figure taken from [19].*

Now for the influence of neutrinos on the matter power spectrum, as we will look into more thoroughly in chapter 5, we may define a (physical) free-streaming wavelength $\lambda_{FS}(t) = 2\pi\sqrt{2/3}v_{th}(t)/H(t)$, with thermal velocity v_{th} (c when UR, p/m when NR) and a corresponding comoving free-streaming wavenumber $k_{FS} = 2\pi a(t)/\lambda_{FS}$, as in [17] (the wavenumber is always the transformation of the comoving system!). This wavelength is the maximal physical distance a particle can travel between a set time t_i and t. Now λ_{FS} increases with time, but after the neutrino transition to NR, it increases more slowly. The comoving wavelength $\lambda_{FS}/a(t)$ even decreases. This means that there is a maximum comoving wavelength, and so a minimal (comoving) wavenumber, given by [17]

$$k_{nr} \approx 0.018\Omega_m^{1/2}(m_\nu/1\text{eV})^{1/2}h\,\text{Mpc}^{-1}. \tag{3.1}$$

Version of June 17, 2022– Created June 17, 2022 - 11:38

Then the effect of neutrinos on the spectrum will become visible on scales $k > k_{nr}$ [17]. This is because the neutrinos are only able to free-stream over distances $\lambda < \lambda_{FS}$, so over wavenumbers $k > k_{FS} \geq k_{nr} \Rightarrow k > k_{nr}$. The physical effect of free-streaming is that it dampens density fluctuations, reducing structure growth, and so decreases power in the matter power spectrum.

Version of June 17, 2022– Created June 17, 2022 - 11:38

Chapter 4

The ΛCDM Model

The shape of the CMB power spectrum is affected by the cosmological parameters. These are defined within the ΛCDM model. In this chapter, we look at this model and the effect of its parameters.

The ΛCDM model assumes the Universe consists of ordinary matter, radiation (which contains both photons and relativistic neutrinos), cold dark matter (CDM) and dark energy in the form of the cosmological constant (Λ). Here CDM is a type of dark matter that is non-relativistic, hence calling it 'cold'. Additionally, ΛCDM assumes that General Relativity can be used to describe gravity in cosmology, and that the Universe is homogeneous and isotropic on a large enough scale. Then the Universe was once much hotter and denser, and has been expanding. The Universe must also be (close to) flat.

The 6 cosmological parameters in ΛCDM are: (1) the baryon density $\omega_b = \Omega_b h^2$, where $h = H_0/(100\ \mathrm{km/s/Mpc})$, (2) the cold dark matter density $\omega_c = \Omega_c h^2$, (3) the spectral index n_s, (4) the scalar perturbation amplitude A_s, (5) the angular scale of the acoustic oscillations θ_*, and (6) the optical depth to reionization τ. We now look at each parameter and its physical meaning more closely, and consider its effect on the CMB anisotropies. (The detailed effects of these parameters on the matter power spectrum will not be discussed here, but see e.g. [24] figure 4.6 and related text.)

4.1 Baryon density ω_b

We use the baryon density $\omega_b = \Omega_b h^2$ instead of just Ω_b, because it is directly proportional to only the physical density of baryons, while Ω_b is

Version of June 17, 2022– Created June 17, 2022 - 11:38

not: $\Omega_b = \rho_b/\rho_{crit}$ and $\rho_{crit} \propto H_0^2 \propto h^2$ so that $\Omega_b \propto \rho_b h^{-2}$, giving $\omega_b \propto \rho_b$. Now ρ_b is certainly a parameter of interest because the density of baryons affects a great many things in the primordial baryon fluid, from the speed of sound waves to the redshift of radiation-matter equality.

The effect of larger ω_b on the CMB is that it causes that odd peaks are enhanced and even peaks are suppressed in the power spectrum [12]. As we saw in the BAO discussion, this is due to the inertia of the baryons. Considering the mass on a spring model, where a larger ω_b means a larger effective mass, one can see that a larger mass would mean that the compressed state is more easily realised and the rarefied state less easily, explaining the enhancement/suppression of the peaks.

Additionally, increasing ω_b shifts the damping tail to higher multipoles. This is because increasing ω_b means more baryons and therefore also more electrons are around, which increases scattering of photons during decoupling. This increases damping, shifting the damping tail to higher multipoles. On top of this, increasing the baryon density increases the Hubble constant and through that decreases the age of the Universe, which in turn reduces how much time photons have for their random walk during decoupling. This decreases damping, shifting the damping tail to lower multipoles. The increasing effect is stronger than the decreasing one, giving a net increase of the damping tail.

4.2 Cold dark matter density ω_c

Again, as for ω_b, we use ω_c because it is directly proportional to the dark matter density ρ_c. *

The effect of ω_c is that the larger it is, the more it dampens the oscillations, reducing the amplitude. This happens because the dark matter causes a part of the gravitational pull on the baryons, and as can be understood from the spring model, it will be harder for the mass to oscillate if the potential well is deeper.

On top of this, increasing dark matter enhances the higher peaks (as in, those at higher multipoles). This is because at some point, the Universe goes from being radiation dominated to dark matter dominated. When this happens the baryons feel a weaker gravitational force, since the gravity before that point due to photons disappears. This causes the oscillations to become much larger. Only higher peaks are affected because

*One may think we are now leaving out the dark energy density ω_Λ, but we're not: since we assume a flat Universe, $\Omega_m + \Omega_r + \Omega_\Lambda = 1$. Now $\Omega_m = \Omega_b + \Omega_c$, and Ω_r is known from the CMB, so we can calculate Ω_Λ once we have Ω_b and Ω_c.

Version of June 17, 2022– Created June 17, 2022 - 11:38

these correspond to modes with higher frequency/smaller wavelength, and the smaller the wavelength, the earlier the oscillation began. Thus these modes started during radiation domination, and are affected by the end of it. The moment at which the domination ends depends on the ratio of Ω_c and Ω_γ, and can be seen in the spectrum by an increase of higher peaks compared to lower ones. This allows us to determine ω_c [15].

Lastly, like ω_b, ω_c influences the Hubble constant and so the age of the Universe, reducing damping during photon decoupling. So increasing ω_c shifts the damping tail to lower multipoles.

4.3 Spectral index n_s and scalar perturbation amplitude A_s

The spectral index describes how density fluctuations vary with scale. The effect of this parameter is that it determines the overall tilt in the power spectrum [12]. This is because the power depends on it as

$$C_\ell \propto A_s(\ell/\ell_0)^{n_s-1}, \tag{4.1}$$

for some pivot point ℓ_0 where A_s is fixed, so increasing n_s increases power at $\ell > \ell_0$ while decreasing power at $\ell < \ell_0$. For $n_s = 1$, the spectrum is scale invariant.

The scalar perturbation amplitude determines the overall magnitude of the scalar perturbations in the power spectrum [12], as can be seen from formula 4.1 for the power. Note that this formula only includes the contribution of the primordial perturbations, described by A_s and n_s, and not the contributions leading to the other effects discussed in chapter 2.

4.4 Angular scale of the acoustic oscillations θ_*

The angular scale of the acoustic oscillations, also the angular scale of the sound horizon, or simply the angular scale, denoted θ_*, is the angular distance that sound waves could travel between the Big Bang and recombination (also accounting for the expansion of space). It is then the position of the first peak in the spectrum, about 1 degree in the sky.

The angular scale is given by [21]

$$\theta_* = r_*/D_M, \tag{4.2}$$

where r_* is the comoving (that is, it stays constant despite the Universe's expansion) sound horizon at recombination, given by

$$r_* = \int_{z_{dec}}^{\infty} \frac{c_s(z)}{H(z)} dz, \tag{4.3}$$

with sound speed $c_s(z) = \frac{c}{\sqrt{3}} \left(1 + \frac{3}{4}\frac{\rho_b(z)}{\rho_\gamma(z)}\right)^{-1/2}$ and z_{dec} being the redshift at photon decoupling [2]. D_M is again the comoving angular diameter distance (see section 2.3).

The effect of θ_* is that it determines peak and through positions; the greater the angle, the more the peaks shift to smaller ℓ [12], since angle and multipole are inversely related.

4.5 Optical depth to reionization τ

The optical depth to reionization τ accounts for scattering of CMB photons after reionization. Reionization started around 400 million years after the Big Bang and was caused by newly formed objects which were radiating energetically enough photons to destroy the atoms that were formed during recombination, bringing free electrons back to the Universe. As the Universe was now much larger and thus the free electron density much smaller, the CMB photons were still able to stream more or less freely, only occasionally getting scattered by the free electrons. This is quantified by [19]

$$\tau = n_H(0)\sigma_T c \int_0^{z_{max}} \frac{(1+z)^2 x_e(z) dz}{H(z)}, \tag{4.4}$$

where $n_H(0)$ is the number density of hydrogen nuclei today, σ_T is the Thomson scattering cross-section, and $x_e(z) = n_e(z)/n_H(z)$ is the free electron fraction. z_{max} is the redshift at which reionization started.

The effect of τ on the CMB is that it decreases the power spectrum for $\ell > 10$ by a factor $e^{-2\tau}$ [19]. This is because the photon density decreases due to the scattering after reionization, decreasing the power. This only affects photons on scales smaller than the Hubble scale at reionization ($cH(t_{rei})^{-1}$), which corresponds to $\ell > 10$.

Version of June 17, 2022– Created June 17, 2022 - 11:38

Chapter 5

The Influence of Neutrinos

We now look at the effect that neutrinos have on the CMB spectrum and the matter power spectrum. This influence is parameterised through two parameters: the effective number of relativistic species, in this case neutrinos, N_{eff} and the sum of the neutrino masses $\sum m_\nu$. The first is useful for the period that neutrinos were ultra-relativistic (UR), the second when they were non-relativistic (NR).

5.1 Thermal history

First, we must discuss (a part of) the thermal history of the Universe and neutrinos. While the temperature of the plasma of particles that formed the Universe was well above 1 MeV, the plasma consisted of neutrinos, electrons, positrons, protons, neutrons and photons, all in thermal equilibrium with each other. This equilibrium existed due to the strong coupling between particles. In the case of neutrinos, coupling to other particles was established through weak reactions like $e^- + e^+ \leftrightarrow \nu_e + \overline{\nu}_e$. Then, at a temperature of about 2 MeV, these weak reactions no longer occurred because their reaction rates dropped below the expansion rate of the Universe, and so neutrinos gradually decoupled. After decoupling was complete, the then free-streaming neutrinos formed the Cosmic Neutrino Background (CνB), analogous to the CMB (but formed much earlier). Shortly after, the temperature became low enough that electron-positron annihilation ($e^- + e^+ \rightarrow \gamma + \gamma$) started. This caused more energy to be added to the photon gas, but not to the neutrino gas, as it had decoupled already. By then considering the change in the effective number of degrees of freedom in entropy, photon and neutrino properties can be equated [4].

Version of June 17, 2022– Created June 17, 2022 - 11:38

Assuming that neutrino decoupling took place instantaneously, this results in the following: $T_\nu/T_\gamma = (4/11)^{1/3}$, so that the energy densities are related through $\frac{\rho_\nu}{\rho_\gamma} = \frac{7}{8}\left(\frac{T_\nu}{T_\gamma}\right)^4 N_{\text{eff}} = 7/8(4/11)^{4/3}N_{\text{eff}}$, where N_{eff} is the effective amount of relativistic species (3 in the SM, for the 3 neutrino flavours). If small corrections due to flavour oscillation and non-instantaneous decoupling are included, this formula can be used with $N_{\text{eff}} = 3.045$ [18].

N_{eff} as a parameter is defined through the density of photons when neutrinos were UR: $\rho_\nu = \rho_{\text{UR}} = 7/8(4/11)^{4/3}N_{\text{eff}}\,\rho_\gamma$, so that the total radiation density becomes $\rho_{\text{rad}} = \rho_\gamma + \rho_{\text{UR}} = \rho_\gamma(1 + 7/8(4/11)^{4/3}N_{\text{eff}})$. N_{eff} thus increases the effect photon-only radiation would have.

The sum of the masses is used at later times when neutrinos became NR. We know neutrinos are NR today, because the current neutrino background temperature is $T_\nu^0 = (4/11)^{1/3}T_\gamma^0 = 1.7 \cdot 10^{-4}$ eV [18]. From the neutrino oscillation experiments we obtain the differences between the masses, giving a lower bound (if we set the lightest neutrino to be massless) of $\sum m_\nu \geq 0.06$ eV. So at least two neutrino species are NR today, since the transition occurs when the average momentum $\langle p \rangle$ is below the mass of the particle, and for Fermi-Dirac distributed particles like neutrinos $\langle p \rangle = 3.15\,T$. We may then write their energy density as $\rho_\nu = \rho_{\text{NR}} = n_\nu m_\nu$ (per species). If we assume the neutrino number density stays the same as when they were UR, $n_\nu = \rho_{\text{UR}}/\langle p_\nu \rangle$. If this is also a thermal spectrum, we know $T_\nu \sim N_{\text{eff}}^{1/4}$ so $\langle p_\nu \rangle \sim N_{\text{eff}}^{1/4}$ so that $\rho_{\text{NR}} \sim N_{\text{eff}}^{3/4} m_\nu$. Note that N_{eff} thus still has an effect when neutrinos are NR!

For the effects of neutrinos, we use that they were all still UR during photon decoupling. This is valid, as for Fermi-Dirac distributed particles, the transition in terms of redshift happens when $z = m_\nu/(3.15T_\nu^0) - 1$ [18]. Then $m_\nu = (0.53\text{ meV}) \cdot (z+1)$, so since photon decoupling occurred around $z \approx 1100$, the neutrino mass for NR transition at decoupling is $m_\nu \approx 0.6$ eV. Now even *Planck*'s weakest constraint puts $\sum m_\nu < 0.54$ eV [18], so we can conclude that all neutrino species were still relativistic at photon decoupling.

5.2 Effects of the effective number of relativistic species

We start with the influence on the CMB spectrum. To figure out the effect of the effective number of relativistic species, like neutrinos in early times,

we need to make sure that the resulting effect is only due to N_{eff} and no other parameter. Then something else needs to be fixed, but the choice of fixing certain parameters influences the effect of N_{eff}. To obtain a characteristic effect of N_{eff}, one needs to keep the radiation-to-matter equality z_{eq} and matter-to-vacuum equality z_Λ fixed [18]. This can be achieved by increasing the total radiation density, total matter density and total vacuum density by a factor

$$\alpha = \frac{1 + 7/8(4/11)^{4/3} N_{\text{eff}}}{1 + 7/8(4/11)^{4/3} \cdot 3.045}. \tag{5.1}$$

One can do this and also keep the parameters $\omega_b, \Omega_\Lambda, A_s, n_s, \theta_*$ and τ constant by varying the Hubble parameter h with N_{eff}, transforming h and ω_c as follows: $h \to h\sqrt{\alpha}$ and $\omega_c \to \omega_c + (\alpha - 1)\omega_m$ [16, 24]. Then all effects of the other parameters that are degenerate with the effect of N_{eff} are absorbed.

The first effect of increasing N_{eff} on the CMB spectrum is an increase in the diffusion scale at the time of decoupling, so that the power spectrum decreases faster at large multipoles. This shows that θ_D is increased: fixing θ_D in the ΛCDM + N_{eff} model while keeping the other parameters fixed is impossible [24]. One can however keep θ_D fixed if one varies the primordial helium abundance Y_P [24]. This is because increasing Y_P decreases the free electron density n_e in equation 2.2 (since helium has a much larger binding energy than hydrogen, so more helium means that more electrons are already bound before recombination for hydrogen begins), increasing θ_D. So we see there is a degeneracy between N_{eff} and Y_P in affecting the damping scale. Additionally, a second effect is that the acoustic peaks are decreased and shifted to smaller multipoles [18]. The effects can be seen in figure 5.1.

The resulting effect of increasing N_{eff} on the matter power spectrum is more power on small scales (which means larger k), and a damping of BAO [24]. This is because ω_c increases due to the transformation, since $\alpha \geq 1$ if $N_{\text{eff}} \geq 3.045$. We have seen before how increasing ω_c dampens the oscillations. Additionally, more dark matter will cause it to cluster more due to the extra gravity it generates, which will in turn also attract more ordinary matter. More clustering of matter on smaller scales then causes the increase in power. The effects can be seen in figure 5.2.

Version of June 17, 2022– Created June 17, 2022 - 11:38

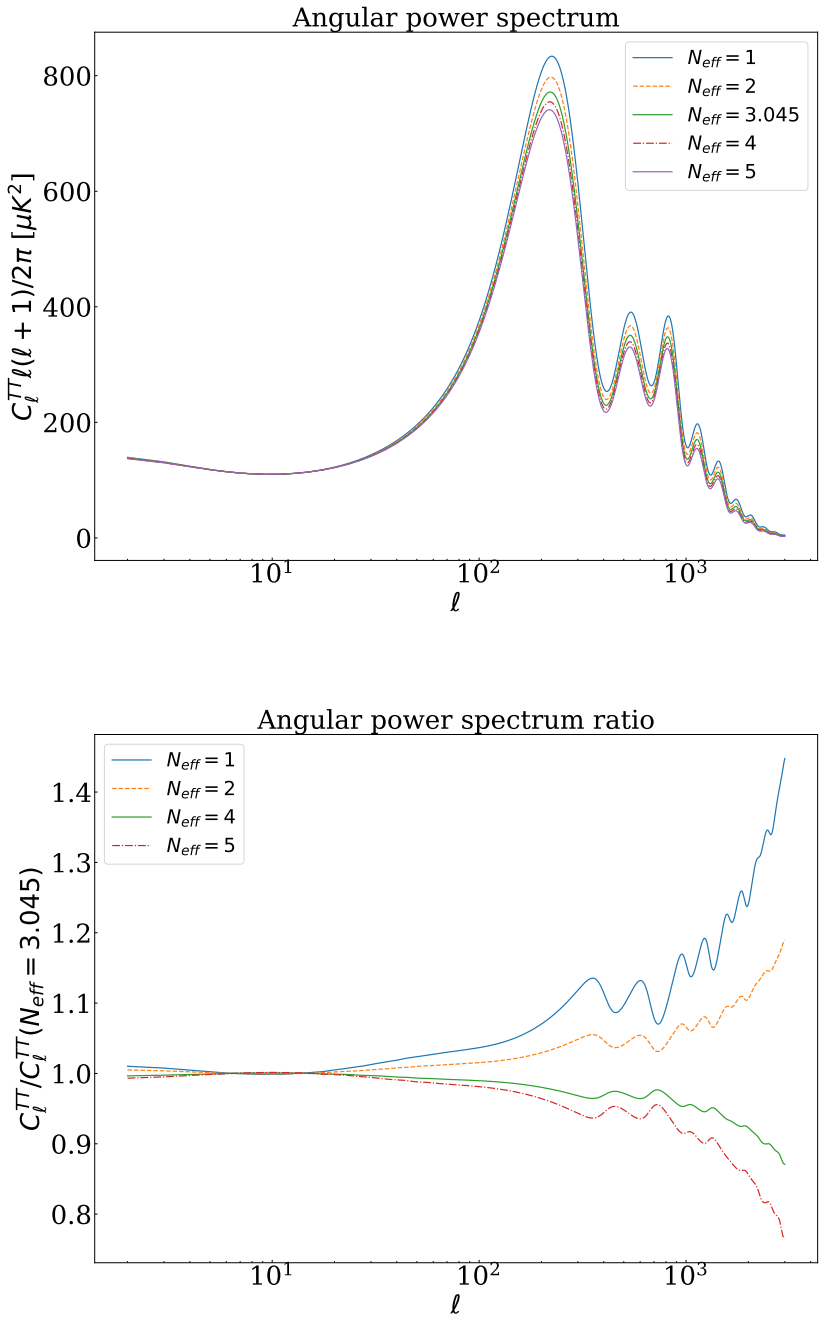

Figure 5.1: *The effect of varying N_{eff} on the CMB power spectrum. The effect is the greatest at high multipoles, as can be seen in the ratio plot, where θ_D is changed. We also see the peaks decrease in amplitude as the effective number increases. The plot with $N_{\text{eff}} = 3.045$ is shown as a reference to the standard ΛCDM case.*

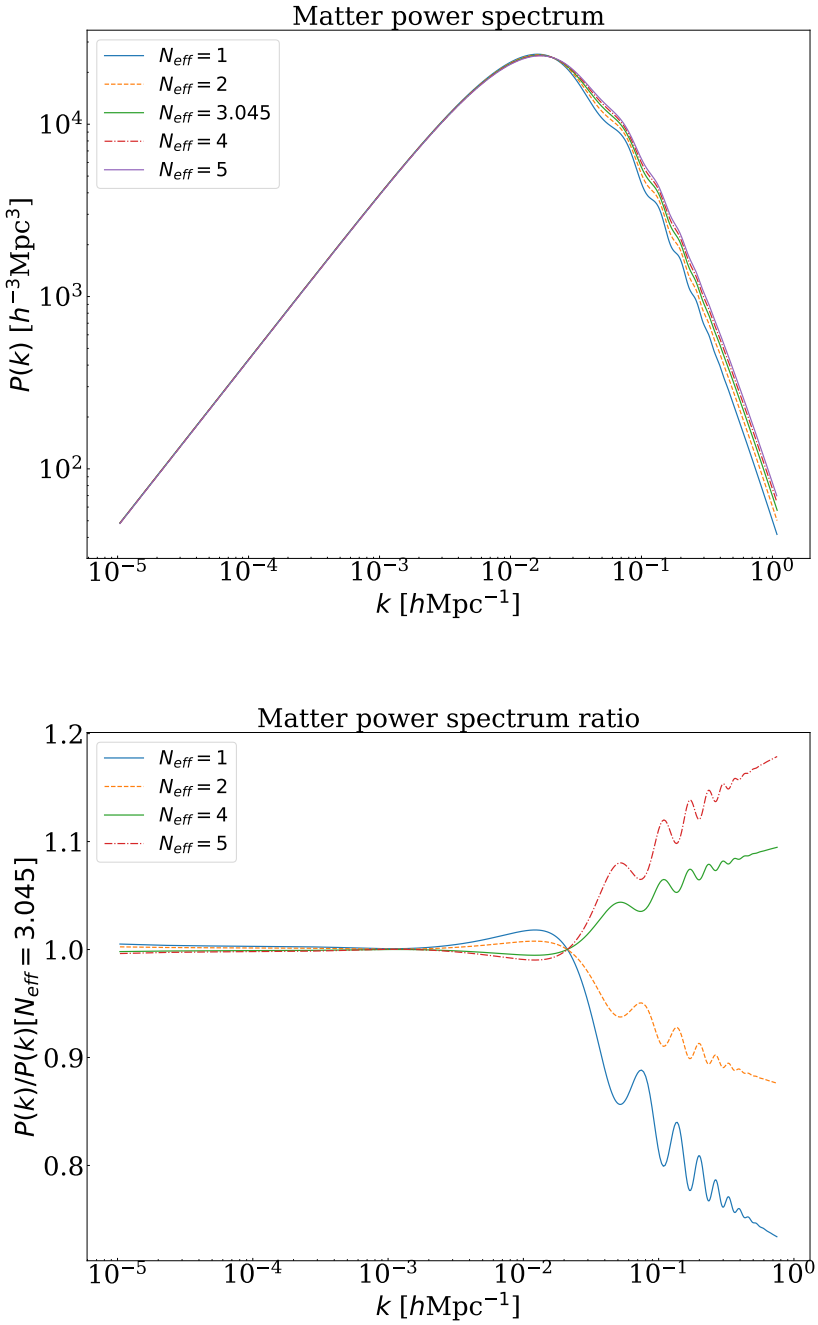

Figure 5.2: *The effect of varying N_{eff} on the matter power spectrum. The effect mostly occurs on small scales, as visible in the ratio plot. The power stays constant at the pivot point $k_0 \approx 0.05\,\text{Mpc}^{-1}$ (as it should). The damping of BAO is better visible in the upper plot, where we see the oscillations in the line flatten out.*

5.3 Effects of the neutrino mass sum

For the effect of the neutrino mass sum, we again need to fix other parameters to obtain a characteristic effect of $\sum m_\nu$. In this case we want to keep ω_m and ω_b constant, so we vary ω_c like $\omega_c \to \omega_c - \omega_\nu$ to account for the increase of ω_m due to neutrinos. To calculate the required neutrino density, we can use $\omega_\nu = \sum m_\nu / (93.14 \text{ eV})$ [18]. This choice is made because it shows the effect of neutrino free-streaming well [18]; its effect can mostly be seen in the matter power spectrum, where keeping Ω_m constant leads to k_{nr} (equation 3.1) being constant.

The effect of $\sum m_\nu$ on the CMB spectrum is that the peaks (mostly the first and second) have an increased amplitude. This happens because of the decrease in ω_c, as discussed before. There is also the shift of the damping tail due to decreasing ω_c. In figure 5.3 we see a plot showing the effect of increasing the neutrino mass sum on the CMB power spectrum.

The effect of $\sum m_\nu$ on the matter power spectrum only occurs on small scales $k > k_{nr}$, where the free-streaming neutrinos do add to the background density, but not to density fluctuations. Then structure growth is reduced, decreasing power on these small scales. The effect is shown in figure 5.4. Note that theoretically, there is no degeneracy between the effect of $\sum m_\nu$ and other parameters on the matter power spectrum, so that one can determine it from the spectrum directly. But to do this, the spectrum must be measured over a large enough interval, which we have not been able to do so far. Within the limited region where we do have measurements, the effect of the neutrino masses is degenerate, so it cannot be determined from LSS data alone [17]. If e.g. the Lyα data becomes more accurate, it would be possible to constrain the mass from only the matter power spectrum. This bound would probably be better than the CMB-only bound, because of the characteristic step-like effect that neutrinos have in the matter power spectrum [17], as can be seen in the ratio plot of figure 5.4.

Note that to produce these plots, `CLASS` code (see section 7.2.1) is used. `CLASS` does not work with three neutrinos with different masses with the right mass differences, which we know is how it should be, but assumes all three neutrinos have the same mass: implementing different masses in the code is much harder. However, this simplification is reasonable as cosmology at its current precision is only sensitive to the sum of the masses, not the individual masses (assuming the three have the same number densities) [17]. So one would not be able to determine individual masses anyway; the net effect of e.g. $m_1 = 0.1, m_2 = 0.2, m_3 = 0.3$ eV is the same as that of $m_1 = m_2 = m_3 = 0.2$ eV, since their sum is the same.

Version of June 17, 2022– Created June 17, 2022 - 11:38

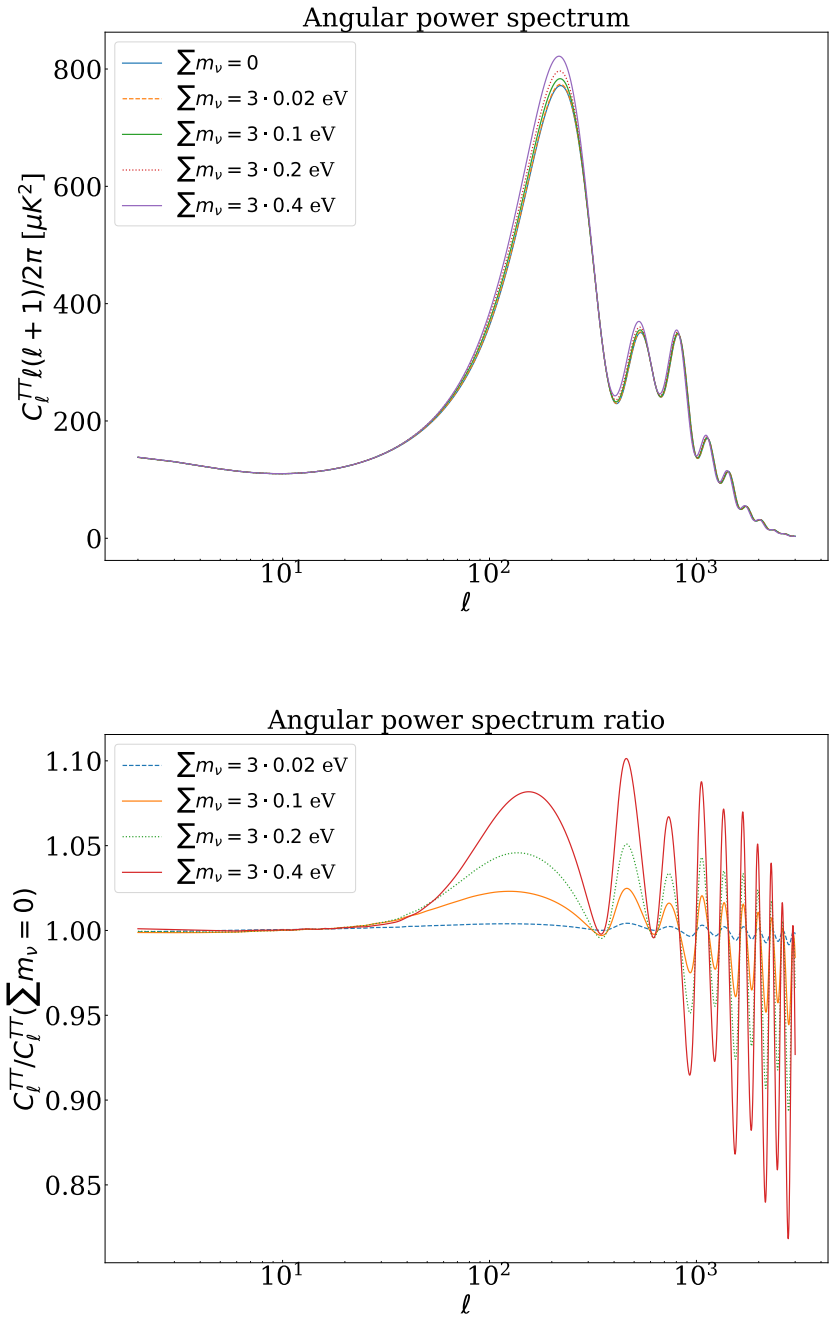

Figure 5.3: *The effect of varying $\sum m_\nu$ on the CMB power spectrum. The increase of the peaks is clearly visible. The oscillatory pattern in the ratio plot shows the shift of the damping tail. The plot with $\sum m_\nu = 0$ is shown as a reference to the standard ΛCDM case. The plot with $\sum m_\nu = 3 \cdot 0.02$ eV shows the minimum of the masses as dictated by oscillation experiments.*

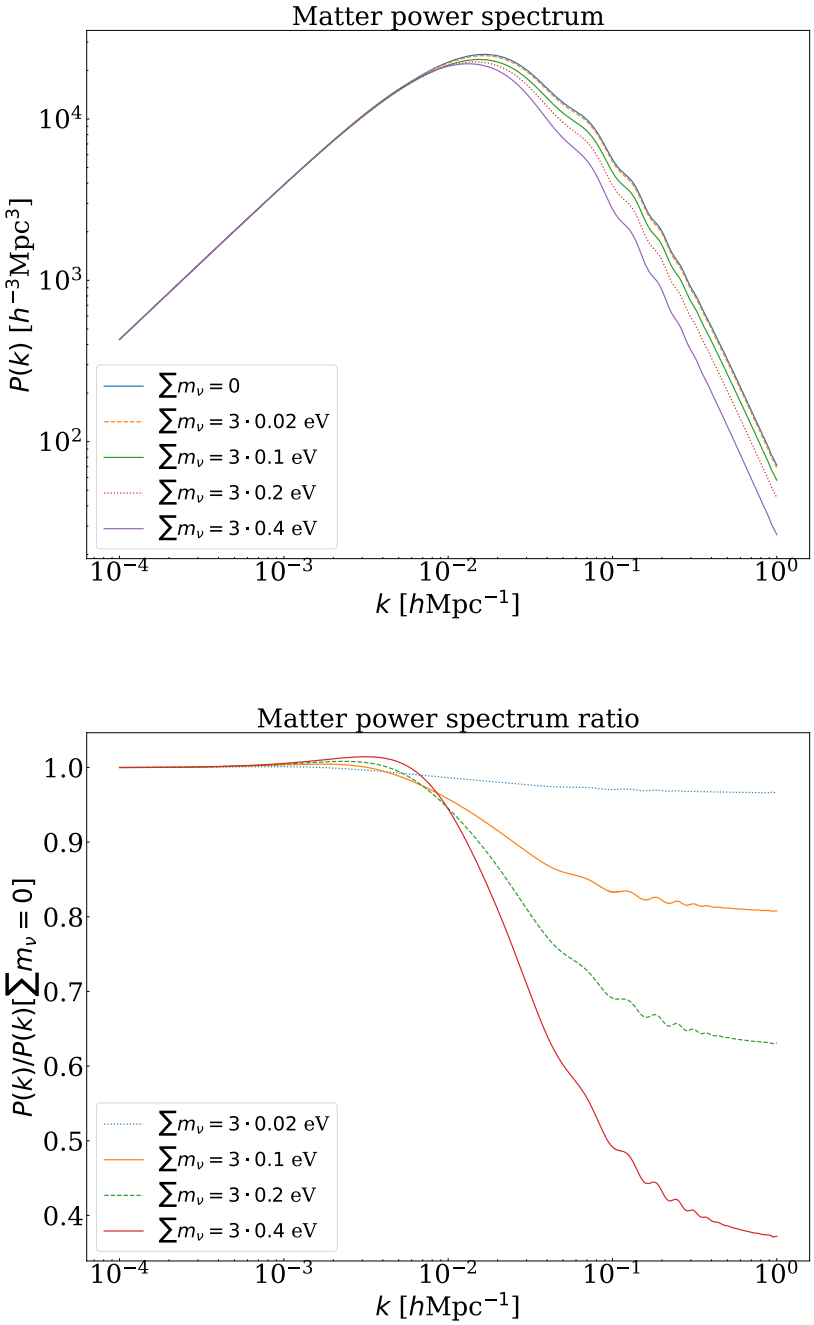

Figure 5.4: *The effect of varying $\sum m_\nu$ on the matter power spectrum. The slight increase followed by a large decrease in the ratio plot shows the step-like effect of increasing the masses.*

Chapter 6

The Influence of HNLs

In this chapter, we will discuss how Heavy Neutral Leptons influence the thermal history, and so the CMB. We add three HNLs to the SM, giving rise to the Neutrino Minimal Standard Model (νMSM). Their effect mostly comes through influencing N_{eff}. This may in turn change the CMB spectrum in such a way that the best fit value for the neutrino mass sum also changes.

6.1 Big Bang Nucleosynthesis

Firstly, HNLs affect Big Bang Nucleosynthesis (BBN) because heavy enough HNLs can decay into mesons (particles made of one quark and one anti-quark) $h = \pi, K$, etc., which may then drive the proton-neutron reactions (instead of the neutrino), for instance for pions:

$$\pi^- + p \to n + \pi^0/\gamma, \tag{6.1}$$

$$\pi^+ + n \to p + \pi^0, \tag{6.2}$$

because these have a much larger cross section than the weak interactions [7]. This brings the proton and neutron densities in equilibrium. If the HNLs have a long enough lifetime, the mesons can be around until after the weak reactions have stopped, which allows neutrons to still form while this was not possible in standard BBN. Then there will be more neutrons around, so that eventually the primordial helium abundance Y_P increases.

6.2 Neutrino population

Secondly, HNLs affect the population of neutrinos. They could either increase or decrease N_{eff} (by most of the injected energy ending up in either the neutrino or photon plasma, respectively) depending on the mass and lifetime of the HNL and on the temperature at which a neutrino from HNL decay is injected into the plasma [8]. HNLs can decay in neutrinos or heavier particles, depending on their own mass. The lifetime is relevant, as HNLs that decay before neutrino decoupling influence thermal history through interactions in the plasma. If they live longer, they automatically increase N_{eff} by simply adding more energetic neutrinos to the plasma. The mass is relevant because this determines what fraction of the energy is injected into which plasma. This happens because the mass determines into which particles the HNL can decay (and different particles will add energy to the neutrino or photon plasma differently), but also because mass determines the energy that a neutrino from the decay has, which in turn is added to either the neutrino or photon plasma (this is before decoupling, so neutrinos are still interacting with electrons). The temperature of the decay-neutrino is relevant as this determines whether the neutrino is in thermal equilibrium with the other neutrinos or not, affecting how much energy is transferred. For a plot of the fraction of energy that is injected into the neutrino plasma as a function of HNL mass (which decay before neutrino decoupling), see figure 6.1. For a plot of how N_{eff} then changes as a function of lifetime for different HNL masses, see figure 6.2.

Additionally, HNLs change the neutrino spectrum by adding the high-energy neutrinos after decaying [8]. For a thermal spectrum, we had $T_\nu \sim \rho_\nu^{1/4} \sim N_{\text{eff}}^{1/4}$. Adding HNLs distorts this as, due to their decay to neutrinos, there are now more neutrinos with $E \gtrsim T_\nu$ and less with $E \lesssim T_\nu$ (after interacting with high-energy neutrino). These distortions cause that now $\langle p_\nu \rangle \nsim N_{\text{eff}}^{1/4}$ so that $\rho_{\text{NR}} \nsim N_{\text{eff}}^{3/4} m_\nu$. Distortions also cause more energy to be added to the photon than neutrino plasma, changing N_{eff}. Now since $\rho_{\text{NR}} \sim m_\nu N_{\text{eff}} / \langle p_\nu \rangle$, a decrease in N_{eff} due to HNLs weakens the effect of NR neutrinos on observables. However, $N_{\text{eff}} = 3.045$ from ΛCDM is able to fit the data quite well, and so it is undesirable to change it.

Version of June 17, 2022– Created June 17, 2022 - 11:38

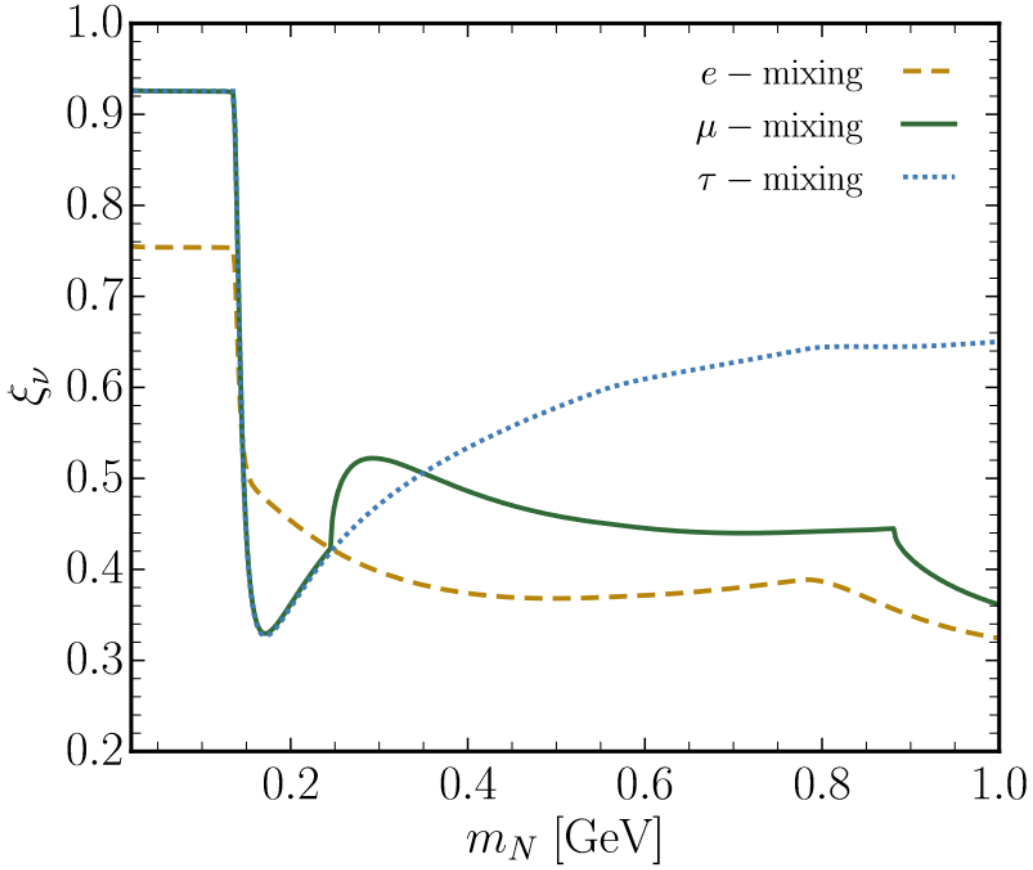

Figure 6.1: *The fraction of the HNL mass ξ_ν that is injected into the neutrino plasma as a function of the HNL mass. The quick drop around $m_N = 135$ GeV happens because here HNLs become heavy enough to decay into pions, which decay in photons, adding energy to photon instead of neutrino plasma. The lines for different flavour-mixing show what neutrino the HNL can mix with, so that neutrino oscillations between the two can take place. Figure taken from [8].*

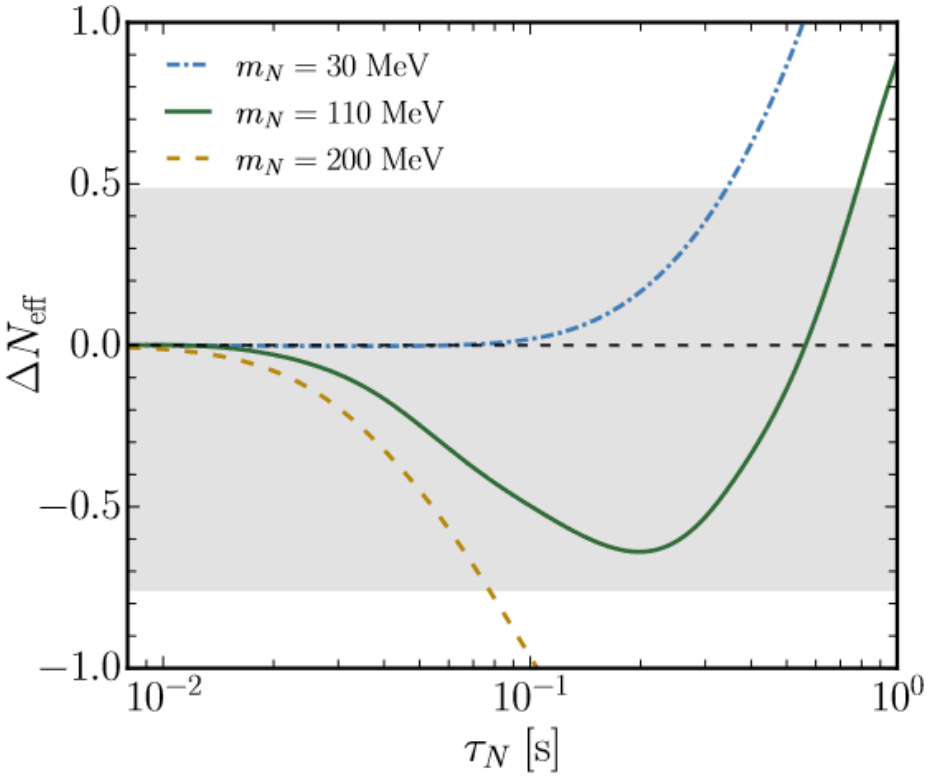

Figure 6.2: *How $N_{\rm eff}$ changes due to HNLs ($\Delta N_{\rm eff} = N_{\rm eff} - N_{\rm eff}^{\rm SM}$, here $N_{\rm eff}^{\rm SM} \simeq 3.044$) as a function of their lifetime, for three different HNL masses. The smallest mass HNL decays mostly into neutrinos, and so only increases $N_{\rm eff}$. The intermediate mass HNL can also mostly decays into neutrinos, but through interactions with photons their energy can still end up in the photon plasma, so that $N_{\rm eff}$ decreases. If the lifetime is long enough however, neutrinos are decoupled and all energy ends up in the neutrino plasma, increasing $N_{\rm eff}$. The heaviest mass HNL can also decay into heavier particles, which decay into photons only, decreasing $N_{\rm eff}$. The grey band shows the current sensitivity of the Planck data. Figure taken from [8].*

Chapter 7

Methods

In this chapter, the data and code that were used during the project are presented, as well as the basic ideas on how the code works, and the way in which the results were obtained.

7.1 Data

In order to determine the neutrino mass, CMB, BAO and LSS data will be used. The CMB data are measurements of the TT, TE and EE mode of the CMB anisotropies at high ℓ and the TT and EE mode at low ℓ, as measured by the *Planck* telescope* [20, 21].

The BAO data [1] are measurements of redshifts of galaxies, measured by the Baryon Oscillation Spectroscopic Survey (BOSS) from the Sloan Digital Sky Surveys (SDSS). Now these are related to BAO, because they can be used to calculate $D_M/r_* = 1/\theta_*$ [2].

The LSS (Large Scale Structure) data [22] are measurements of the redshifts in combination with angular positions of galaxies, also measured by the SDSS. Here was looked at Luminous Red Galaxies (LRGs). From these measurements, the density of galaxies can be determined [24]. This can be used to find the matter power spectrum, as discussed before. Furthermore, the *Planck* lensing data is included; these are again measurements of the CMB power spectrum, specifically how it is influenced by gravitational lensing due to large matter densities that the photons pass by while travelling through space. The information that the lensed spectra contain can be used to find the matter power spectrum, as also discussed before.

*Data can be downloaded at `http://pla.esac.esa.int/pla/`.

7.2 Code

The theoretical predictions of the extended ΛCDM model can be used to fit to the described data and determine the cosmological parameters, including neutrino parameters.

7.2.1 CLASS

To calculate the theoretical spectra, `CLASS` code is used[†] [6], the Cosmic Linear Anisotropy Solving System. `CLASS` is a Boltzmann code, that is, given a set of cosmological parameters, it calculates the power spectra in ΛCDM or an extension of it. To do this, it uses the Boltzmann equation. In its most abstract form, this is $\frac{df}{dt} = C[f]$, where C is the collision operator (describing encounters between particles) and $f(\vec{x}, \vec{p}, t)$ is the particle's distribution function, defined through the number of particles in a phase space volume as $N(\vec{x}, \vec{p}, t) = f(\vec{x}, \vec{p}, t)(\Delta x)^3 \frac{(\Delta p)^3}{(2\pi)^3}$ [10]. It describes the dynamics of a rarified gas, like photons or baryons in the Universe. It can thus be used to calculate the evolution of the CMB photon gas, and so the CMB spectra.

7.2.2 MontePython

To actually determine the cosmological parameters, `MontePython` code is used[‡] [3, 9]. This uses a Monte Carlo Markov Chain (MCMC) to determine the best-fit parameters. This is necessary, as simply trying a fit for an array of values will take far too long: 6/7 base parameters plus order 10 nuisance parameters (parameters with no physical usefulness, but necessary to fix calibration, systematics, etc. of the experiment), for a considerable amount of different values, all to be evaluated one by one, quickly grows too large to handle.

The MCMC program starts of with a certain set of parameters. Then it chooses a new set of parameters from a proposal distribution (one for each parameter). This proposal distribution is usually a Gaussian, with the current value as its mean and (optimally) $2.4/\sqrt{D}$ (D: dimensionality of parameter space) as its variance [24]. So random parameters close to the current ones are picked with a probability as described by these Gaussians.

[†]Can be downloaded via `http://class-code.net/`.

[‡]Can be downloaded via `https://github.com/brinckmann/montepython_public`.

Then the program determines if these parameters are better than the previous ones through the Metropolis-Hasting algorithm [24]. This is done by calculating the power spectrum resulting from the parameters with `CLASS`, from which the ratio of posterior distributions $p(\vec{\theta}|d)$ can be calculated, for parameters $\vec{\theta}$ (which is the chance that these parameters $\vec{\theta}$ describe the data). Denoting the proposal distribution as $q(\vec{\theta}^*|\vec{\theta})$, for newly chosen parameters $\vec{\theta}^*$, the probability of accepting the new parameters is then given by

$$\alpha = \min\left(1, \frac{p(\vec{\theta}^*|d)}{p(\vec{\theta}|d)} \frac{q(\vec{\theta}^*|\vec{\theta})}{q(\vec{\theta}|\vec{\theta}^*)}\right). \tag{7.1}$$

Note that only the previous parameter set is considered at each step, which is what makes this a Markov Chain.

So with each parameter having its own chain, the newly accepted parameter will be added to the chain and the process starts over. The amount of steps used as input determines how many times the program checks if a new parameter is accepted. The actual length of the chain will be about $1/4$ of this, as optimally, about $1/4$ of steps is accepted. The reason this method works, is because we only use the ratio of the posterior distributions. This can be calculated since $p(\vec{\theta}|d) \propto p(d|\vec{\theta})P(\vec{\theta})$ [24], where $p(d|\vec{\theta})$ is the chance certain data is described by those parameters (calculated with `CLASS`) and $P(\vec{\theta})$ is the prior distribution (specified by our input). Knowing the ratio of those two, gives the ratio of posterior distributions. Now the posterior distribution grows larger as the parameters fit the data better, and so the probability of accepting better parameters also increases. This leads to ending up with the best fitting parameters.

7.2.3 Specific settings of the code

In the code, as discussed before, we work with three neutrinos that have the same mass, instead of accounting for the mass differences. This single mass is defined as m_{ncdm} (ncdm: non-cold dark matter, which neutrinos may be counted as), so that for us $\sum m_\nu = 3m_{ncdm}$.

On top of this, the code is made to work with the parameter T_{ncdm} instead of N_{eff}. For our purposes, keeping other parameters constant, these are related through

$$\frac{N_{\text{eff}}}{3.045} = \left(\frac{T_{ncdm}}{0.71611}\right)^4, \tag{7.2}$$

where we have set their standard values as $N_{\text{eff}} = 3.045$ and $T_{ncdm} = 0.71611$. This ncdm temperature is a replacement of the $T_\nu/T_\gamma = (4/11)^{1/3}$

Version of June 17, 2022– Created June 17, 2022 - 11:38

term in the definition of the effective number of relativistic species. It is therefore not really a temperature, and its value slightly differs from $(4/11)^{1/3} = 0.71377$ to ensure that $\omega_\nu = \sum m_\nu / (93.14 \text{ eV})$.

Lastly, we fix the primordial helium fraction to $Y_P = 0.247$ instead of letting the program calculate it from BBN, so that this well-determined parameters stays constant. (Note that this is for the case that we determine the neutrino mass bound without the influence of HNLs.)

7.2.4 Adding HNLs

We will consider HNLs with a short lifetime, which themselves do not affect the CMB directly. They do influence the neutrino population, which in turn affects the CMB. HNLs change the neutrino population, as discussed in chapter 6, by changing N_{eff} and by adding spectral distortions. The influence of the spectral distortions is important because these stay present when neutrinos become non-relativistic and thus influence the late-time effects of the mass sum on the CMB and matter power spectrum (and the influence of N_{eff} we have discussed in chapter 5). Additionally, we have seen that HNLs affect Y_P, which in turn affects the CMB as well, as also discussed in chapter 5.

Now we may characterise the distortions caused by HNLs by a ratio $\xi = \frac{\langle p_\nu \rangle}{3.15 T_\nu} - 1$, where we know $T_\nu \sim N_{\text{eff}}^{1/4}$ and remember $\langle p_\nu \rangle = 3.15 T_\nu$ in case of a Fermi-Dirac distribution. Then for our purposes, as given to me by my supervisor, $\xi < 0.1$ so we may neglect the spectral distortions. Note that we then still have $\rho_{NR} \sim N_{\text{eff}}^{3/4} m_\nu$. Thus the influence of HNLs only occurs through changing N_{eff} and Y_P. We then add HNLs to the code by simply changing these parameters to new fixed values.

We will look at the specific case of an HNL with mass $m_N \sim 500$ MeV and lifetime $\tau_N \sim 0.05$ s. As given to me by my supervisor, this results in a change of parameters to $N_{\text{eff}} = 2.45$ and $Y_P = 0.26$. These are used as standard values when determining the neutrino mass bound.

7.3 Determining parameters

Now, using these programs, determining parameters works as follows: in the input file we define the 6 cosmological parameters plus N_{eff} or $\sum m_\nu$. Additionally, we define the nuisance parameters. These depend on which data is used, as each experiment has their own nuisance parameters. For each parameter, we give a guess of the mean and standard deviation and

an upper and lower bound, fixing the parameter space in which the best-fit can be searched for (you can also pass no bounds, but for e.g. the neutrino mass sum, you want the sum to be at least 0.06 eV, as required by the mass splittings). Then a run can be started. Each run takes a certain amount of steps N. At first, we use $N \sim 1000$ and use the calculated chain to obtain best fits for the parameters. The first run won't have good convergence, that is, one cannot determine the standard deviation of the calculated best fit values. But these best fits can then be used as a new input with $N \sim 5000$, which will have better convergence, and from which we can find a covariance matrix. This matrix has the variance of the parameters as its diagonal elements, and the covariance between two parameters as off-diagonal elements. Adding this as input in the next run will increase convergence even more. So for the final run we give both best fits and a covariance matrix, and run for a long time, with at least $N \sim 100\,000$. This should give a good estimate of the best fits with standard deviation. Doing 100 000 steps however takes very long, so in practice (having only my own laptop available) we do this in parts of 10 000 to 40 000 steps, taking about 3 to 12 hours, and add these chains together.

If the result is not good enough after $N \sim 100\,000$ steps, we keep going. To see if the result is good, we look at the single parameter plots of the posterior distribution from the program's output. If these look like Gaussians and their mean and standard deviation can be calculated, the result should be good. The contour plots of two parameters should also be smooth and elliptical of shape.

We will determine three parameters: N_{eff} (as an extension of standard ΛCDM only) and $\sum m_\nu$ both with and without the influence of the HNL. We determine N_{eff} to cross-check if the method works properly, as we can compare this to e.g. the *Planck* team's result. We determine $\sum m_\nu$ without the HNL to have a reference point from this same method to compare the result with the HNL to.

Chapter 8

Results

In this chapter, we look at the results of fitting the cosmological parameters plus N_{eff} or $\sum m_\nu$, both with and without the influence of HNLs present, to the data with `MontePython`.

8.1 Effective number of relativistic species

We start with fitting the ΛCDM $+N_{\text{eff}}$ model, to determine if the data (here *Planck* TT,TE,EE+lowT+lowE+BAO was used) allows for additional relativistic species (like neutrinos). The result for the effective number of relativistic species from `MontePython` is, at 95 % CL, $N_{\text{eff}} = 2.98^{+0.36}_{-0.33}$, with a standard deviation of 0.16. This agrees with the standard ΛCDM value of 3.045. For this result, 97 764 steps were taken, resulting in chains of length 26 515. The minimum χ^2 is 2779. The values found for each parameter are listed in table 8.1, and plots showing the posterior distributions for each parameter and for two parameters together as a contour plot can be seen in figure 8.1. These simply show the probability that each parameter has a certain value, given the data. In both the table and figure, the nuisance parameters are left out.

8.2 Neutrino mass sum

For the ΛCDM $+\sum m_\nu$ model, we use *Planck* TT,TE,EE + lowT + lowE + lensing + BAO + LSS data to determine the neutrino mass sum. In the prior we have set $\sum m_\nu > 0.056$ eV, or $m_{ncdm} > 0.0187$ eV, so that the minimum from oscillation experiments is considered. The result from `MontePython` is then an upper bound of $\sum m_\nu < 0.13$ eV at 95% CL, with a standard

deviation of 0.02 eV. Here 297 892 steps were taken, resulting in chains of length 78 651. The minimum χ^2 is 2835. The values found for each parameter are listed in table 8.2, and the plots of the posterior distributions in figure 8.2.

Table 8.1: *The results of the* `MontePython` *program for the six ΛCDM parameters plus N_{eff} (as T_{ncdm}). The mean value is determined from a probability distribution of good fit values, of which the best-fit value is the most probable. The last two columns show the 95% CL intervals.*

Parameter	best-fit	mean$\pm\sigma$	95% lower	95% upper
$100 \cdot \omega_b$	2.244	$2.239^{+0.015}_{-0.015}$	2.208	2.268
ω_c	0.1183	$0.1181^{+0.0034}_{-0.0028}$	0.112	0.124
$100 \cdot \theta_*$	1.042	$1.042^{+0.0005}_{-0.00056}$	1.041	1.043
$\ln(10^{10} A_s)$	3.047	$3.041^{+0.017}_{-0.017}$	3.009	3.072
n_s	0.9656	$0.9648^{+0.006}_{-0.0058}$	0.953	0.9762
$100 \cdot \tau$	5.506	$5.485^{+0.72}_{-0.73}$	4.024	6.948
T_{ncdm}	0.7134	$0.7121^{+0.011}_{-0.01}$	0.6917	0.7326

Table 8.2: *The results of the* `MontePython` *program for the six ΛCDM parameters plus $\sum m_\nu$ (as m_{ncdm}). The mean value is determined from a probability distribution of good fit values, of which the best-fit value is the most probable. The last two columns show the 95% CL intervals.*

Parameter	best-fit	mean$\pm\sigma$	95% lower	95% upper
$100 \cdot \omega_b$	2.244	$2.248^{+0.013}_{-0.014}$	2.221	2.275
ω_c	0.1184	$0.1185^{+0.00086}_{-0.00087}$	0.1168	0.1202
$100 \cdot \theta_*$	1.042	$1.042^{+0.00027}_{-0.00029}$	1.041	1.043
$\ln(10^{10} A_s)$	3.048	$3.05^{+0.014}_{-0.016}$	3.021	3.08
n_s	0.9688	$0.9681^{+0.0036}_{-0.0037}$	0.9611	0.9754
$100 \cdot \tau$	5.964	$5.849^{+0.7}_{-0.78}$	4.394	7.364
m_{ncdm} (eV)	0.01946	$0.02773^{+0.0019}_{-0.009}$	0.0187	0.04367

Version of June 17, 2022– Created June 17, 2022 - 11:38

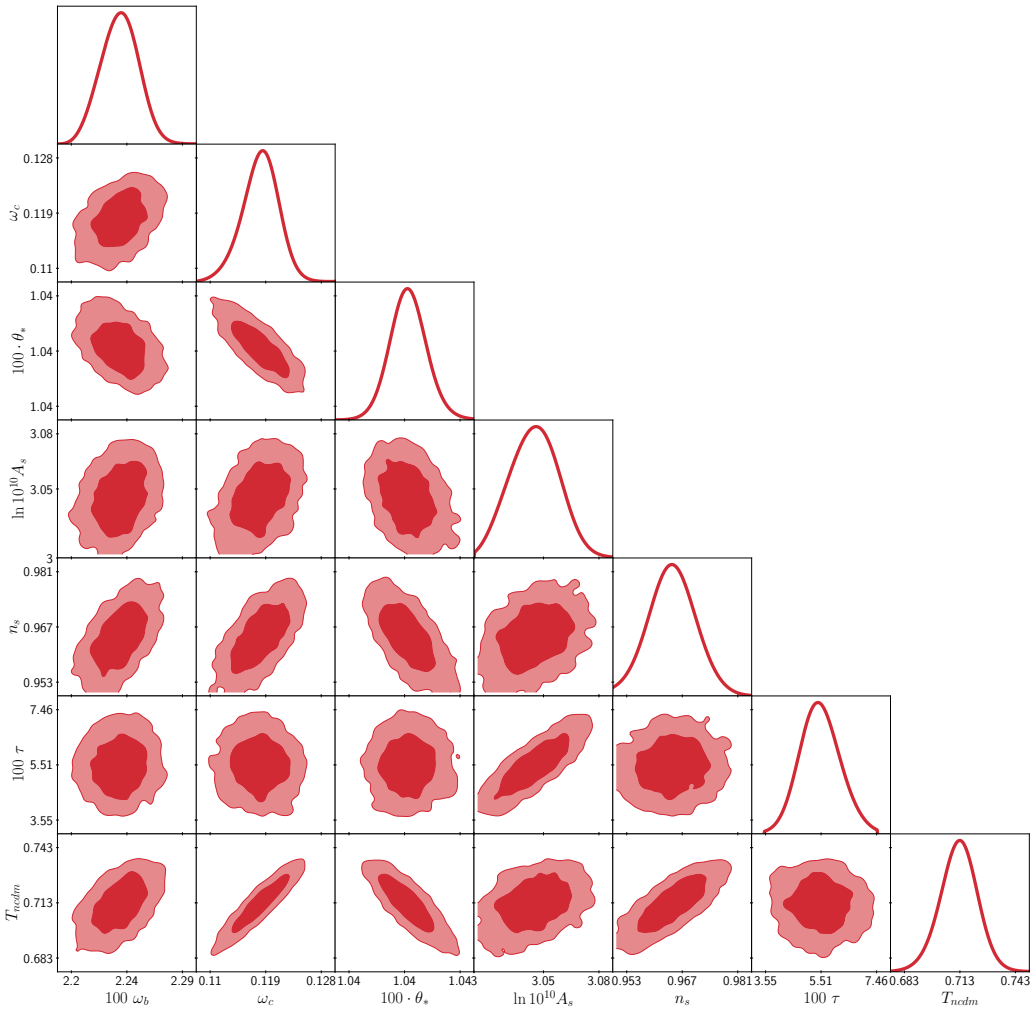

Figure 8.1: *The constraints on the parameters from* `MontePython` *when varying* N_{eff} *(as* T_{ncdm}*). The contour plots for two parameters show 68 and 95 % confidence intervals. The single parameter plots (implicitly) have probability on their vertical axis.*

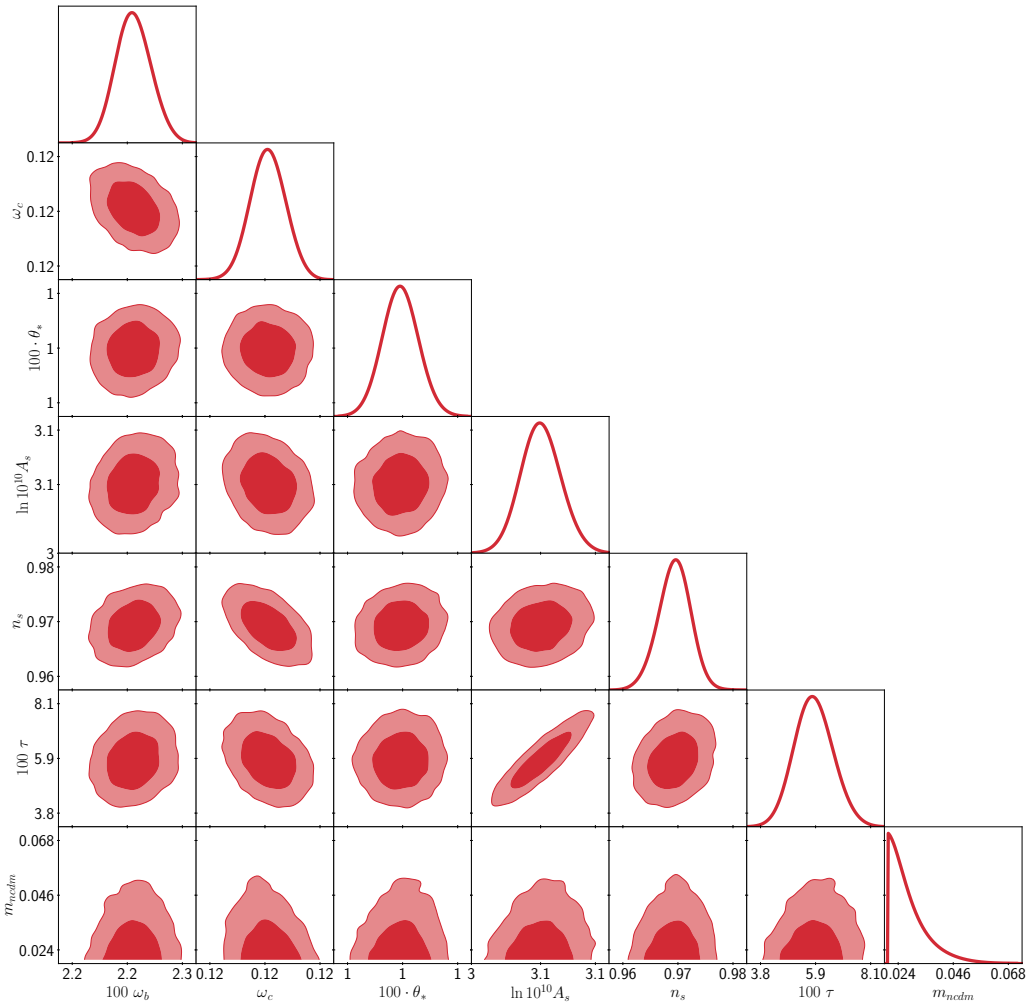

Figure 8.2: *The constraints on the parameters from* `MontePython` *when varying* $\sum m_\nu$ *(as* m_{ncdm}*). The contour plots for two parameters show 68 and 95 % confidence intervals. The single parameter plots (implicitly) have probability on their vertical axis. We can see how the neutrino mass wants to be smaller than the minimum value, even smaller than zero, because the contours would overlap those values if we had allowed them to.*

8.3 Neutrino mass sum with HNLs

For this result we consider the ΛCDM model extended by one HNL with mass $m_N \sim 500$ MeV and lifetime $\tau_N \sim 0.05$ s, added to the code by setting $N_{\text{eff}} = 2.45$ and $Y_P = 0.26$. We used the same data as for the neutrino mass sum without the HNL: *Planck* TT,TE,EE + lowT + lowE + lensing + BAO + LSS data. In the prior we again set $\sum m_\nu > 0.056$ eV ($m_{ncdm} > 0.0187$ eV) to account for the minimum dictated by oscillation experiments.

The result is then an upper bound of $\sum m_\nu < 0.12$ eV at 95% CL, with a standard deviation of 0.01 eV. This is slightly lower than the bound without HNLs, by 0.01 eV. For this result 406 325 steps were taken, resulting in chains of length 106 646. The minimum χ^2 is 2840. The values found for each parameter are listed in table 8.3, and the plots of the posterior distribution in figure 8.3.

Table 8.3: *The results of the* `MontePython` *program for the six ΛCDM parameters plus $\sum m_\nu$ (as m_{ncdm}), when the influence of HNLs is present. The mean value is determined from a probability distribution of good fit values, of which the best-fit value is the most probable. The last two columns show the 95% CL intervals.*

Parameter	best-fit	mean$\pm\sigma$	95% lower	95% upper
$100 \cdot \omega_b$	2.224	$2.225^{+0.013}_{-0.014}$	2.199	2.251
ω_c	0.1096	$0.1091^{+0.00078}_{-0.00074}$	0.1076	0.1106
$100 \cdot \theta_*$	1.044	$1.044^{+0.00029}_{-0.00028}$	1.043	1.044
$\ln 10^{10} A_s$	3.029	$3.031^{+0.014}_{-0.015}$	3.003	3.059
n_s	0.9561	$0.9566^{+0.0033}_{-0.0037}$	0.9501	0.9632
$100 \cdot \tau$	5.72	$5.935^{+0.66}_{-0.78}$	4.527	7.357
m_{ncdm}(eV)	0.02316	$0.02563^{+0.0014}_{-0.0069}$	0.0187	0.03856

Version of June 17, 2022– Created June 17, 2022 - 11:38

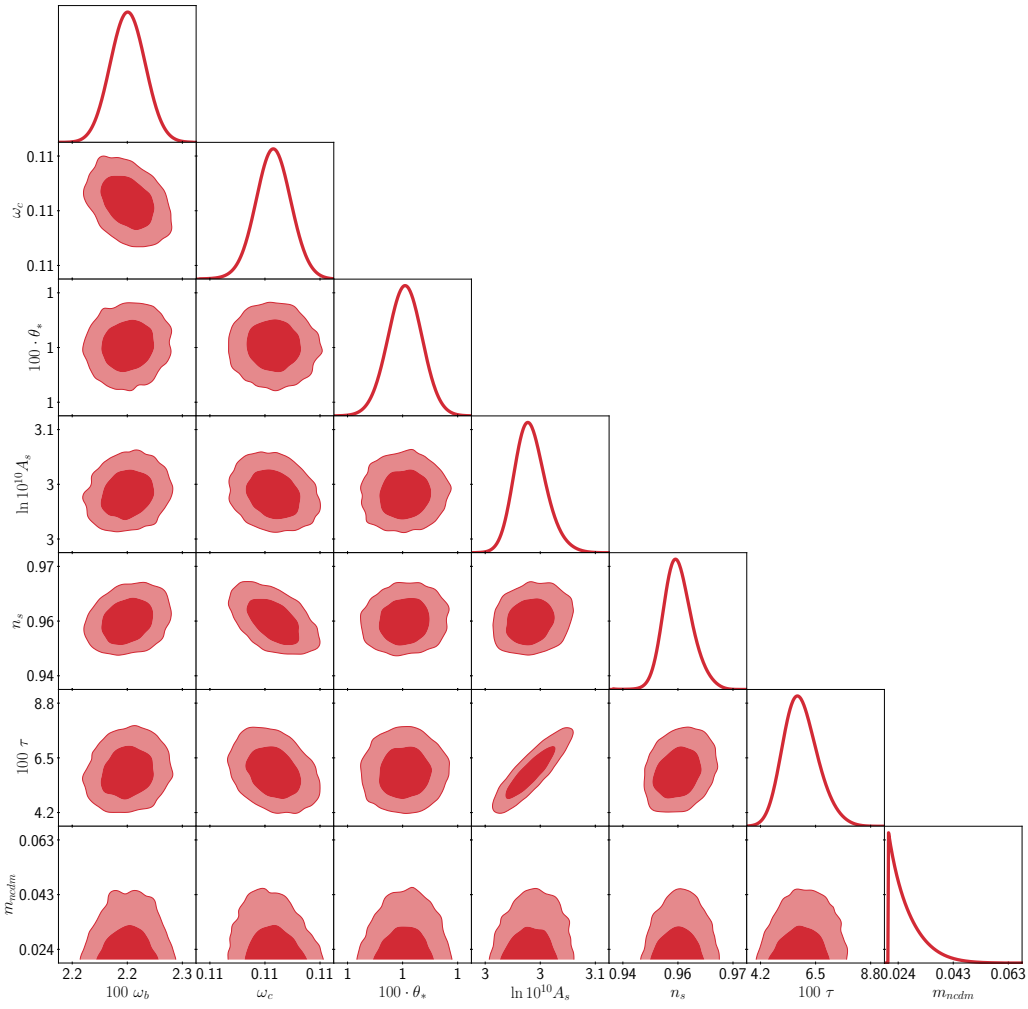

Figure 8.3: *The constraints on the parameters from* `MontePython` *when varying* $\sum m_\nu$ *(as* m_{ncdm}*) in the presence of HNL influence. The contour plots for two parameters show 68 and 95 % confidence intervals. The single parameter plots (implicitly) have probability on their vertical axis. It can be seen that the mass again wants to be zero, but is cut off by the minimum we set.*

Chapter 9

Discussion

We now discuss the results obtained and how they relate to previous research. The limitations of the method are also discussed, and at the end an outlook is given.

9.1 Effective number of relativistic species

The outcome of $N_{\text{eff}} = 2.98^{+0.36}_{-0.33}$ agrees with the standard ΛCDM value of 3.045, and so this shows no proof of additional particle species. This result is also in line with the results of the *Planck* team [21] (eq. 67b), who found $N_{\text{eff}} = 2.99^{+0.34}_{-0.33}$ (95% CL). They did also include lensing data, which I didn't, and I included low ℓ TT data, which they didn't, so that may cause the small discrepancy.

This discrepancy may also be due to a limitation of this project, which has to do with running time. Working on my own laptop, as discussed in the Methods chapter, runs quickly take a very long time to complete. Due to time constraints, I therefore have performed around 100 000 steps to reach this result, but this is a minimum amount of steps one would have to take, and more steps are preferable. One can see in figure 8.1 that the contours are not as smooth as they should be, and that some are slightly outside of the bounds for the parameter; doing more steps would likely fix this. Additionally, the program works better when all steps are taken within the same run. Since I was restricted to 40 000 steps at most, this reduced performance of the program, so that it took longer to find the best-fit values. Having a more powerful computer available would therefore have helped greatly.

Comparing the other parameters in table 8.1 to the *Planck* findings for

standard ΛCDM [21], we see that all six parameters stay the same within about a 1σ error margin. Thus changing N_{eff} does not significantly change the other parameters, which is expected as N_{eff} does not really differ from its standard value either.

Since N_{eff} is allowed to be the standard value, this shows no proof that additional relativistic species are needed. At the same time, the influence of HNLs could vary N_{eff} within the error margin, so their existence also cannot be ruled out.

9.2 Neutrino mass sum

The outcome of $\sum m_\nu < 0.13$ eV without HNLs is also in line with the *Planck* results [21] (eq. 63b), who found $\sum m_\nu < 0.12$ eV (95% CL). Here the difference of 0.01 eV may be caused by the SDSS data I included, which they did not.

The difference may again also be caused by the limited running time I had. Indeed, looking at the results of consecutive runs of the neutrino mass sum, the upper bound decreased with order 0.0001 eV each run. It is possible the bound would have decreased down to 0.12 eV eventually, but I did not have enough time to run the program long enough to see if it did. On top of this, we see the minimum $\chi^2 = 2835$, while it should be around 2780 for the ΛCDM + $\sum m_\nu$ case. This shows that the fit values are not as good as they should be, which would also be fixed by taking more steps. By taking more steps, we may have also ended up with the *Planck* result.

Comparing the results of the other 6 parameters in table 8.2 to the *Planck* standard ΛCDM results [21], we see all parameters stay the same within the 95% CL intervals. Thus they do not change significantly when adding massive neutrinos.

9.3 Neutrino mass sum with HNLs

The result of $\sum m_\nu < 0.12$ eV with the influence of HNLs present is still in line with the *Planck* results [21] (eq. 63b), who found $\sum m_\nu < 0.12$ eV (95% CL); in this case the bound is even the same, unlike the result that was found without HNLs. To look at how HNLs change the mass sum though, it is better to look at the result obtained here with the same dataset and method, so here we do see a slight decrease of 0.01 eV due to HNLs, but the result is still very close. Interestingly though, while the upper bound

Version of June 17, 2022– Created June 17, 2022 - 11:38

and mean value have decreased with HNLs, the best-fit value has actually increased. So at least the fit moves away from the imposed lower bound on the masses a bit more.

Now, how can we explain the upper bound staying more or less the same? Naively, we may remember that the neutrino influence occurs through their energy density, for which at late times we have $\rho_{NR} \sim N_{\text{eff}}^{3/4} m_\nu$. This influence on the data is fixed, so if we decrease N_{eff}, we must increase m_ν. This will relax the upper bound on the masses. However, this is clearly incorrect, as the bound tightened a little.

So instead, we can conclude that although the influence of neutrinos manifests through their energy density, the two parameters N_{eff} and $\sum m_\nu$ we use to describe this influence do still have their own, distinct effects on the spectra. This is also found by the *Planck* team, when they vary both parameters at the same time and find the results to remain close to those of varying only one parameter [21]. (One may think this is obvious from the discussion about neutrino influence in chapter 5, where the effects were certainly different, but remember that the influences there were only for a certain set of fixed parameters and transformations. Here the parameters are not fixed, so the net effect of changing N_{eff} and $\sum m_\nu$ may be completely different.)

This does mean that the effect of decreasing N_{eff} needs to be countered by the other parameters. Looking at the 6 base parameters in table 8.3, we see that compared to the *Planck* standard ΛCDM results [21], some now do change significantly. Only τ stays within 1σ of the standard result, while ω_b and A_s only stay within 2σ, and ω_c, θ_* and n_s deviate more than 2σ from the standard results. Now looking at how large the change is in percentages, ω_b, n_s and A_s only decrease by less than 1%, while θ_* increases a tiny amount, τ increases about 6% and ω_c decreases the most, with about 8%. Now if we were to apply the transformation for N_{eff} from chapter 5 using the factor α to the standard results, we would get $\omega_c = 0.108$ (and an increased h). This is quite close to the value of ω_c with HNLs present; additionally, ω_b is almost the same, so we may expect that the influence of N_{eff} is similar to that discussed before. The decrease of N_{eff} would then cause that power decreases slower at high multipoles (decreasing θ_D), and the acoustic peaks have an increased amplitude and are shifted to higher multipoles.

Now the decrease of θ_D can already be fixed by the increase in Y_P that also comes from HNLs, since increasing Y_P increases θ_D too, as discussed before. The remaining effects indeed seem to be countered by the remaining 4 parameters: a lower A_s and higher τ decrease the amplitude again,

while a higher θ_* and lower n_s shifts the peaks to lower multipoles again. These parameters are not able to counter the effect of N_{eff} on the matter power spectrum of decreasing power at higher multipoles though; here the slight decrease of neutrino mass, which increases power again, may explain why the neutrino mass does not stay exactly the same.

In figure 9.1, we can see this. Here the ratio is plotted of the spectra resulting from the parameters we determined *but* with zero neutrino mass, and from the parameters we determined including the value for the neutrino mass. This way we can see what the influence of the neutrino mass is. Note that in the case where the neutrino mass is zero, the dark energy budget is increased to keep $\Omega_m + \Omega_r + \Omega_\Lambda = 1$. In the figure, we see that the CMB spectrum only slightly deviates from the final result when the neutrino mass is zero, showing the other parameters countered most changes due to $N_{\text{eff}} = 2.45$ already. The difference occurs through the Sachs-Wolfe effect, which makes sense as massive neutrinos add to this effect. The matter power spectrum deviates more from the final result, so there the neutrino mass has a more significant impact, which may explain the slight decrease.

We see that the influence of the specific HNL that was used on the mass bound is quite small. By extension, other HNLs probably also have a small influence, since HNLs with different masses/lifetimes just affect N_{eff} and Y_P differently, but then their effects on the spectra are still distinct from those of the neutrino mass sum and our discussion still applies. This may not be the case if the change in these two parameters becomes much larger, but since both are quite well constrained this would be an unphysical situation anyway.

Moreover, there may also be no difference between the mass bounds with and without HNLs, if one of the results is slightly off. If more steps were run by the program, a difference as small as 0.01 eV may have disappeared. This would also correspond with the 0.12 eV bound from the *Planck* team without HNLs [21], which is the same as the bound found here with HNLs (assuming this bound would not change if more steps were taken, which it actually might). So the result that the mass bound slightly decreases due to HNLs may be wrong, if not enough steps were taken. We were not able to do this due to time constraints.

We can however say for sure that the naive expectation of the mass bound increasing when N_{eff} decreases, due to $\rho_{NR} \sim N_{\text{eff}}^{3/4} m_\nu$, does not hold, as there seems to be a tendency of strengthening the mass bound instead, and the bound certainly does not relax.

Version of June 17, 2022– Created June 17, 2022 - 11:38

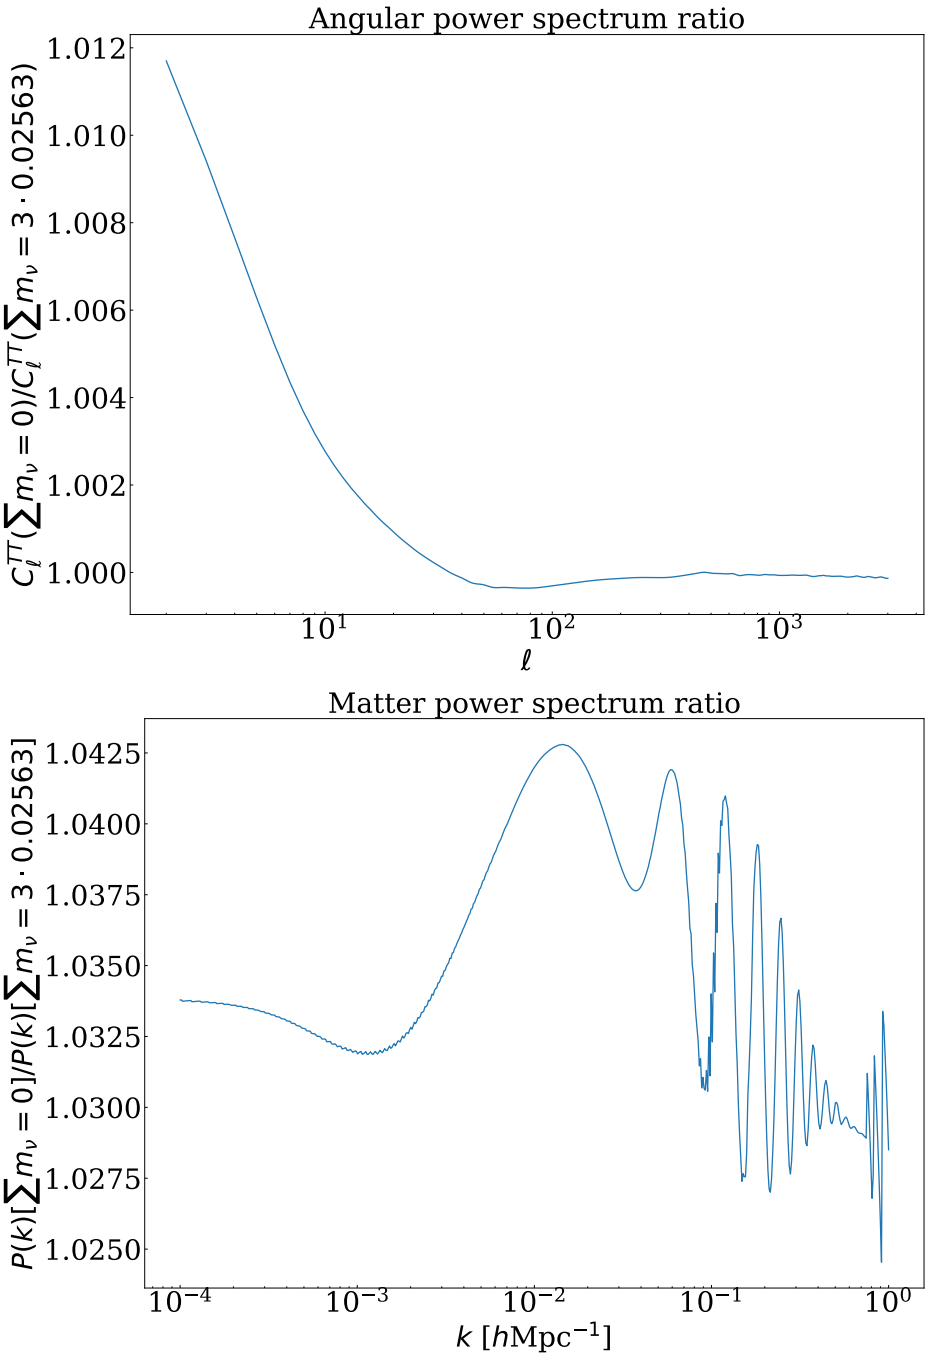

Figure 9.1: *The ratio of spectra with determined parameters and zero neutrino mass over nonzero neutrino mass. The relatively small effect of the masses on the CMB spectrum shows the other parameters already made up for the changes there, while the larger effect of the masses on the matter power spectrum may explain the slight decrease of the mass sum.*

Version of June 17, 2022– Created June 17, 2022 - 11:38

9.4 Outlook

One questions that remains, is why the standard ΛCDM model works so well. It is a (relatively) simple cosmological model, needing only 6 parameters to describe the Universe. Yet it describes the CMB spectrum splendidly; adding extra parameters just returns their previously fixed value. No neutrino mass needed, nor further complications, like spatial curvature, tensor modes, or running of the spectral index. It is surprising that the simplest model describes the data this well. At the same time, we do know that neutrinos must have mass, so unfortunately the simplest, best-fitting model does not seem the correct one.

In order to really figure out if the theoretical value of $N_{\text{eff}} = 3.045$ is correct, or if it is changed due to HNLs, and to find stronger constraints on the mass sum, more accurate measurements are required. For a list of upcoming experiments which aim to achieve this, see [18] section 26.4. Another experiment that is being planned is to measure the CνB directly, by the PTOMELY experiment [5]. This would certainly be useful data to constrain neutrino properties, although it is very difficult to measure the CνB at this moment.

Version of June 17, 2022– Created June 17, 2022 - 11:38

Chapter 10

Conclusions

Adding the effective number of relativistic species as a parameter to the ΛCDM model, it was found that $N_{\text{eff}} = 2.98^{+0.36}_{-0.33}$ (95% CL), in agreement with the standard value of 3.045. Other parameters also did not change significantly, as expected when N_{eff} itself does not change significantly either.

Adding the neutrino mass sum as a parameter to the ΛCDM model, it was found that $\sum m_\nu < 0.13$ eV (95% CL). Again, other parameters did not change significantly.

Adding both the neutrino mass sum as a parameter and the influence of HNLs to the ΛCDM model, it was found that $\sum m_\nu < 0.12$ eV (95% CL), very close to the bound without HNLs. This time the other parameters also changed significantly, which countered the effect of the lower N_{eff} due to HNLs. Because of this, the mass sum could stay close to its value without HNLs. Thus the influence of HNLs on the mass bound is tiny, but there may also be no effect at all, depending on whether the results would change if there had been more time to run `MontePython`. We do note that the naive expectation that the mass bound would increase due to decreasing N_{eff} certainly does not hold.

Version of June 17, 2022– Created June 17, 2022 - 11:38

Acknowledgements

I would like to thank Maksym Ovchynnikov for guiding me during the project and answering any questions I had. I would also like to thank Alexey Boyarsky and Subodh Patil for being my supervisor and second corrector, respectively.

This project obtained results based on observations obtained with *Planck* (`http://www.esa.int/Planck`), an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA, and Canada.

Bibliography

[1] S. Alam *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(3):2617–2652, 03 2017.

[2] E. Aubourg *et al.* Cosmological implications of baryon acoustic oscillation measurements. *Physical Review D*, 92(12), Dec 2015.

[3] B. Audren, J. Lesgourgues, K. Benabed, and S. Prunet. Conservative Constraints on Early Cosmology: an illustration of the Monte Python cosmological parameter inference code. *JCAP*, 1302:001, 2013.

[4] D. Baumann. Cosmology (lecture notes). Available online at `http://cosmology.amsterdam/education/cosmology/`, last accessed on May 23, 2022.

[5] S. Betts *et al.* Development of a Relic Neutrino Detection Experiment at PTOLEMY: Princeton Tritium Observatory for Light, Early-Universe, Massive-Neutrino Yield, 2013.

[6] D. Blas, J. Lesgourgues, and T. Tram. The cosmic linear anisotropy solving system (CLASS). part II: Approximation schemes. *Journal of Cosmology and Astroparticle Physics*, 2011(07):034–034, jul 2011.

[7] A. Boyarsky, M. Ovchynnikov, O. Ruchayskiy, and V. Syvolap. Improved big bang nucleosynthesis constraints on heavy neutral leptons. *Physical Review D*, 104(2):023517, July 2021.

[8] A. Boyarsky, M. Ovchynnikov, N. Sabti, and V. Syvolap. When FIMPs Decay into Neutrinos: The N_{eff} Story. *arXiv e-prints*, March 2021.

[9] T. Brinckmann and J. Lesgourgues. MontePython 3: boosted MCMC sampler and other features. *Phys. Dark Univ.*, 24:100260, 2019.

[10] S. Dodelson and F. Schmidt. *Modern Cosmology*. Academic Press, 2020.

[11] N. Y. Gnedin and A. J. S. Hamilton. Matter power spectrum from the Lyman-alpha forest: myth or reality? *Monthly Notices of the Royal Astronomical Society*, 334(1):107–116, 2002.

[12] D. S. Gorbunov and V. A. Rubakov. *Introduction to the theory of the early universe: cosmological perturbations and inflationary theory*. World Scientific Publishing Co. Pte. Ltd, Singapore, 2011.

[13] W. Hu. Early and Late ISW. Available online at `http://background.uchicago.edu/~whu/physics/aux/isw.html`, last accessed on May 25, 2022.

[14] W. Hu. ISW Effect. Available online at `http://background.uchicago.edu/~whu/physics/isw.html`, last accessed on May 25, 2022.

[15] W. Hu. Radiation Driving Force. Available online at `http://background.uchicago.edu/~whu/intermediate/driving.html`, last accessed on May 23, 2022.

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

[17] J. Lesgourgues and S. Pastor. Neutrino mass from cosmology. *Advances in High Energy Physics*, 2012, 2012.

[18] Particle Data Group. Review of Particle Physics. *Progress of Theoretical and Experimental Physics*, 2020(8), 08 2020.

[19] Planck Collaboration. Planck 2018 results. I. Overview and the cosmological legacy of Planck. *Astronomy & Astrophysics*, 641:A1, Sep 2020.

[20] Planck Collaboration. Planck 2018 results. V. CMB power spectra and likelihoods. *Astronomy & Astrophysics*, 641:A5, 2020.

[21] Planck Collaboration. Planck 2018 results. VI. Cosmological parameters. *Astronomy & Astrophysics*, 641:A6, Sep 2020.

[22] B. A. Reid *et al.* Cosmological constraints from the clustering of the Sloan Digital Sky Survey DR7 luminous red galaxies. *Monthly Notices of the Royal Astronomical Society*, 404(1):60–85, 04 2010.

[23] D. Scott and G. Smoot. Cosmic background radiation mini-review, 2004.

[24] S. Vagnozzi. *Weigh them all! - Cosmological searches for the neutrino mass scale and mass ordering*. PhD thesis, Stockholm University, June 2019.