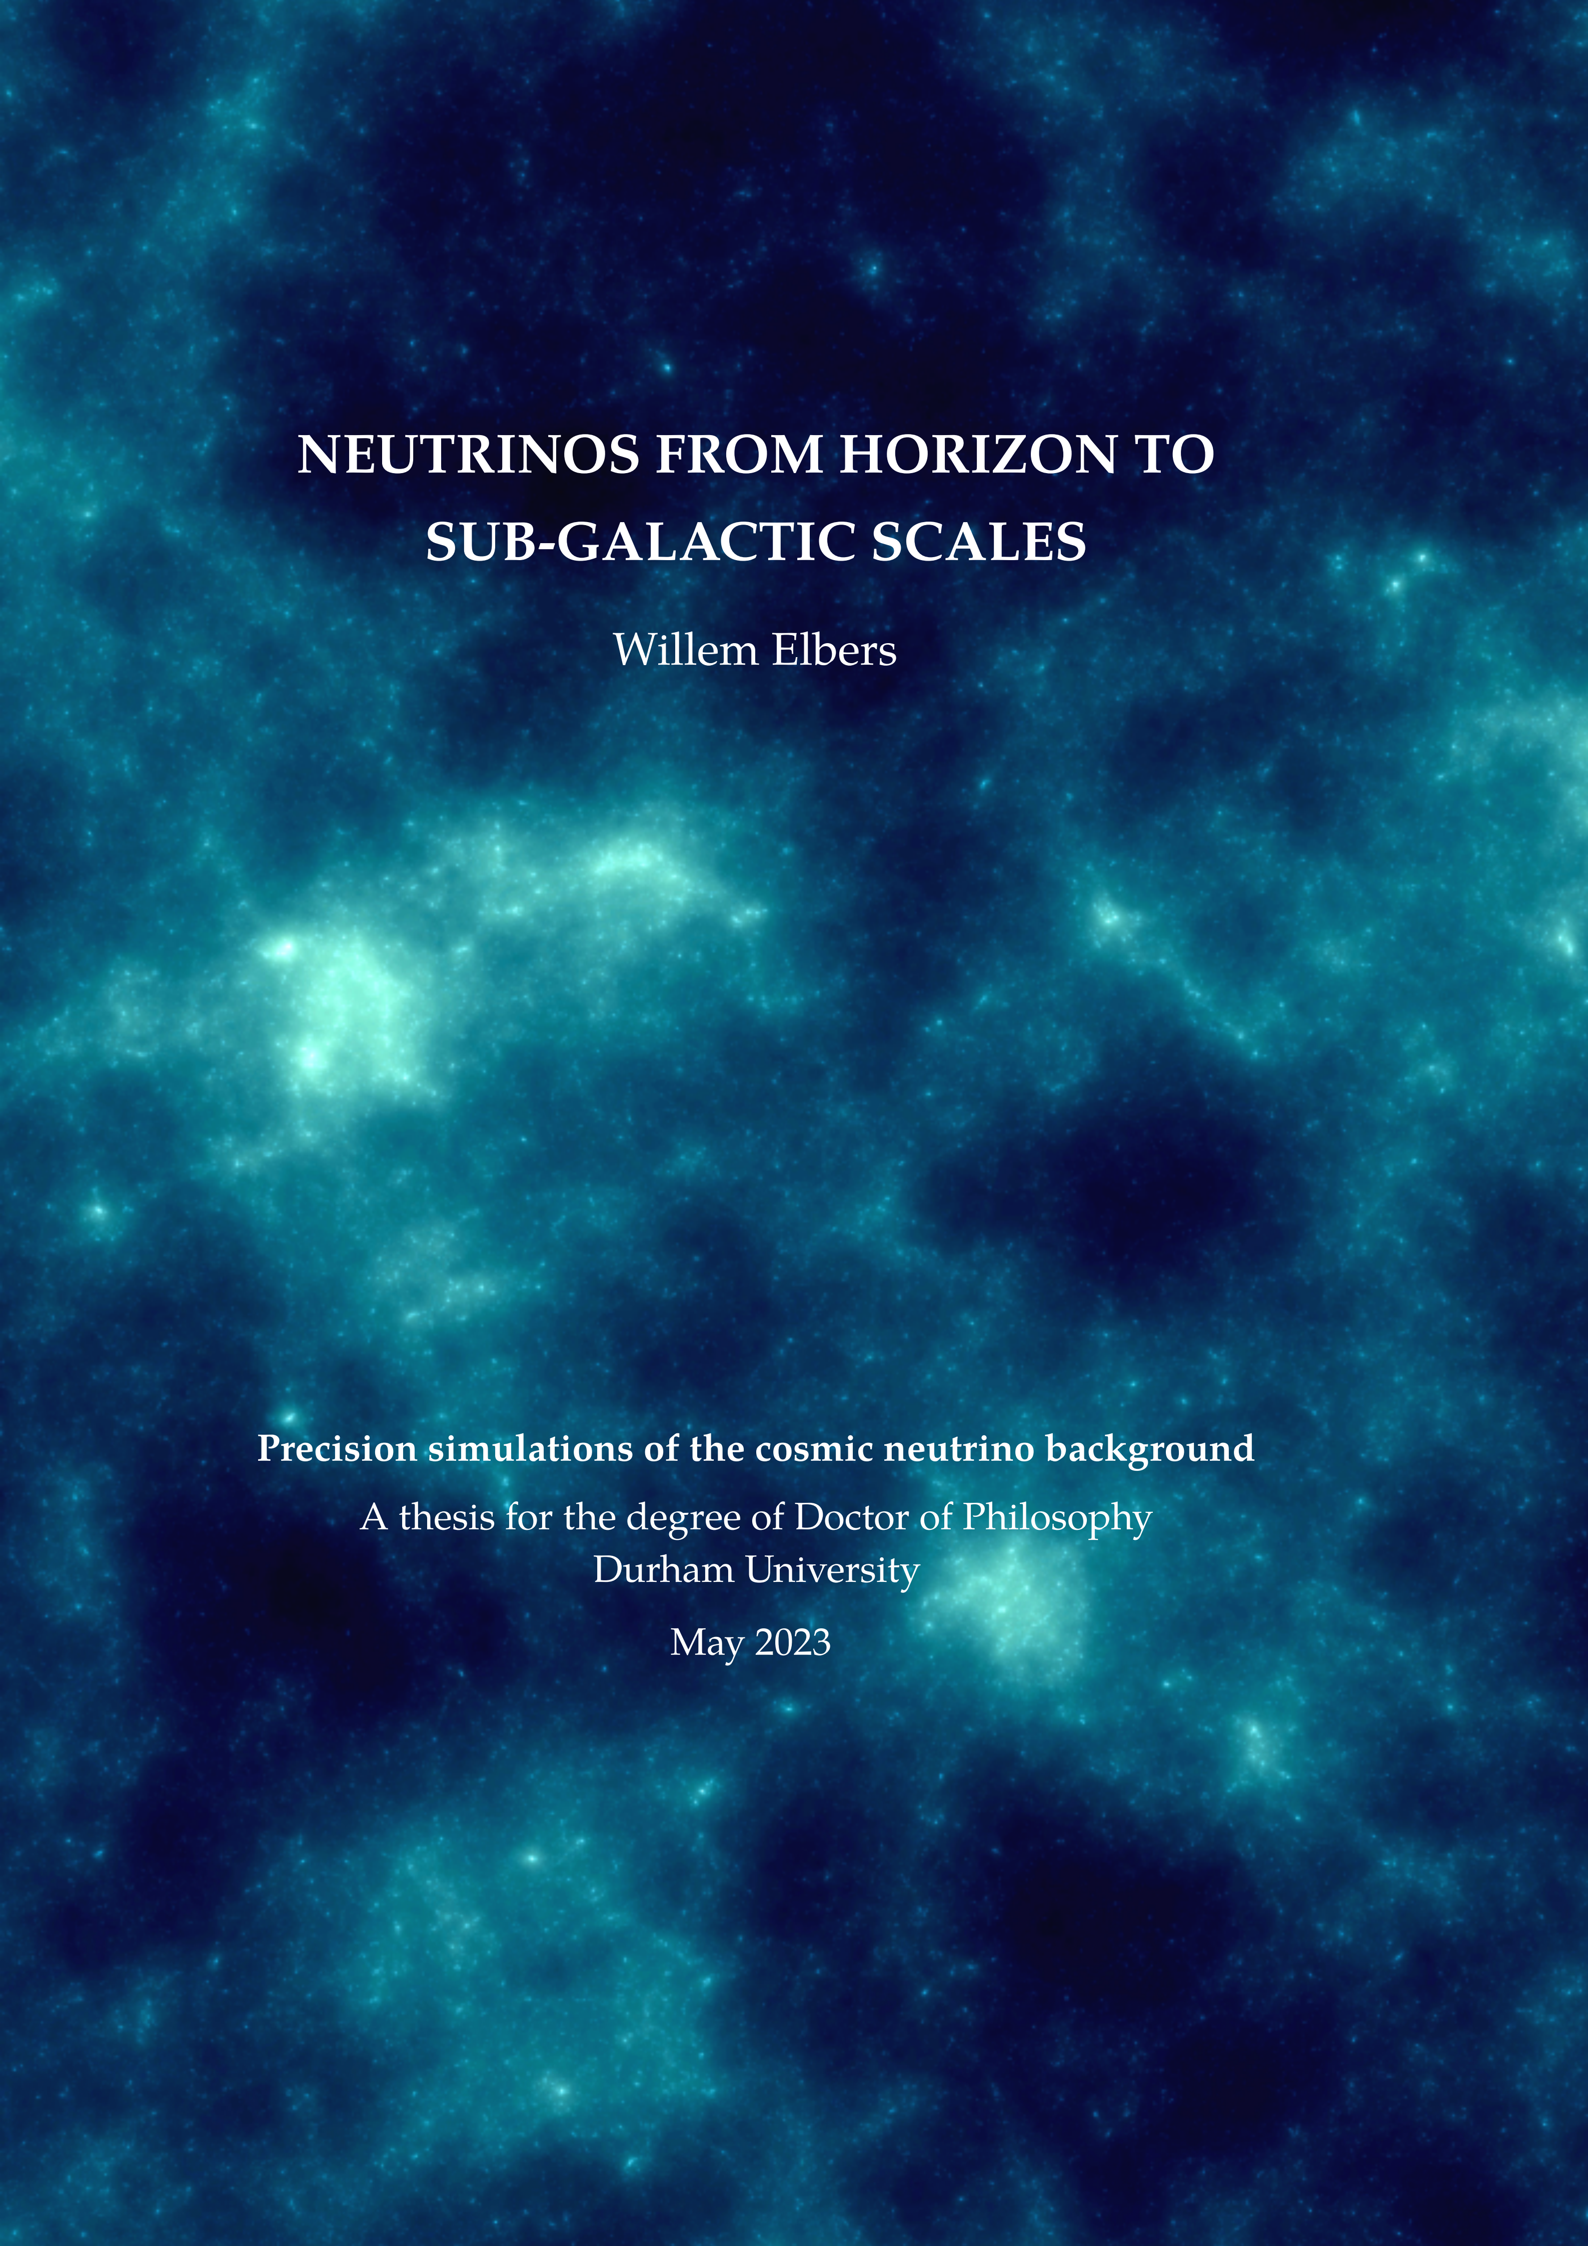

NEUTRINOS FROM HORIZON TO SUB-GALACTIC SCALES

Willem Elbers

Precision simulations of the cosmic neutrino background

A thesis for the degree of Doctor of Philosophy
Durham University

May 2023

Advisors:
Carlos Frenk
Adrian Jenkins
Baojiu Li
Silvia Pascoli

Institute:
Institute for Computational Cosmology
Department of Physics
Durham University

Location:
Durham, United Kingdom

Illustration:
Illustration based on a simulation carried out on the cosma supercomputer in Durham. The image represents the projected mass density in a simulated region of the Universe. The blue patches correspond to neutrino clouds that stretch over tens of millions of light-years. The tiny bright spots correspond to dark matter halos: dense regions where galaxies are forged.

Abstract

A first determination of the mass scale set by the lightest neutrino remains a crucial outstanding challenge for cosmology and particle physics, with profound implications for the history of the Universe and physics beyond the Standard Model. In this thesis, we present the results from three methodological papers and two applications that contribute to our understanding of the cosmic neutrino background.

First, we introduce a new method for the noise-suppressed evaluation of neutrino phase-space statistics. Its primary application is in cosmological N-body simulations, where it reduces the computational cost of simulating neutrinos by orders of magnitude without neglecting their nonlinear evolution. Second, using a recursive formulation of Lagrangian perturbation theory, we derive higher-order neutrino corrections and show that these can be used for the accurate and consistent initialisation of cosmological neutrino simulations. Third, we present a new code for the initialisation of neutrino particles, accounting both for relativistic effects and the full Boltzmann hierarchy. Taken together, these papers demonstrate that with the combination of the methods described therein, we can accurately simulate the evolution of the neutrino background over 13.8 Gyr from the linear and ultra-relativistic régime at $z = 10^9$ down to the non-relativistic yet nonlinear régime at $z = 0$. Moreover, they show that the accuracy of large-scale structure predictions can be controlled at the sub-percent level needed for a neutrino mass determination.

In a first application of these methods, we present a forecast for direct detection of the neutrino background, taking into account the gravitational enhancement (or indeed suppression) of the local density due to the Milky Way and the observed large-scale structure within $200h^{-1}$ Mpc. We determine that the large-scale structure is more important than the Milky Way for neutrino masses below 0.1 eV, predict the orientation of the neutrino dipole, and study small-scale anisotropies. We predict that the angular distribution of neutrinos is anti-correlated with the projected matter density, due to the capture or deflection of neutrinos by massive objects along the line of sight.

Finally, we present the first results from a new suite of hydrodynamical simulations, which includes the largest ever simulation with neutrinos and galaxies. We study the extent to which variations in neutrino mass can be treated independently of astrophysical processes, such as feedback from supernovae and black holes. Our findings show that baryonic feedback is weakly dependent on neutrino mass, with feedback being *stronger* for models with larger neutrino masses. By studying individual dark matter halos, we attribute this effect to the increased baryon density relative to cold dark matter and a reduction in the binding energies of halos. We show that percent-level accurate modelling of the matter power spectrum in a cosmologically interesting parameter range is only possible if the cosmology-dependence of feedback is taken into account.

Contents

Acknowledgement

This thesis is dedicated to everyone who joined me on this journey and to all the people I met along the way. ■ To my parents, Kees and Jeanet, who were there for me every step of the way. ■ To my family and friends back home, for making it seem as if no time had passed. ■ To my mentors over the years and in particular to Rien van de Weijgaert, for his invaluable advice. ■ To my incredible team of advisors in Durham: Carlos Frenk, Adrian Jenkins, Baojiu Li, and Silvia Pascoli, for lending me their expertise, ideas, patience, support, and much more. ■ To Daniel Eisenstein and Carlton Baugh for generously agreeing to examine this thesis. ■ To David Creed, for his advice and for supporting our dream to advance neutrino physics through cosmology. ■ To my friends at the department, my office mates present and past, and my collaborators further afield for the times we shared, from the boardgame nights at James' and Factorio games with Victor to evenings on the football pitch and in the pub. ■ To the COSMA team for the smooth operation of the supercomputer in difficult times. ■ Writing this thesis would not have been possible, nor half as much fun without you all!

Declaration

The work in this thesis is based on research carried out at the Institute for Computational Cosmology and Department of Physics in the University of Durham. No part of this thesis has been submitted elsewhere for any other degree or qualification, and it is the sole work of the author unless referenced to the contrary in the text.

Durham, United Kingdom, May 2023

The following chapters were published as:

Chapter 4 Willem Elbers, Carlos S. Frenk, Adrian Jenkins, Baojiu Li, and Silvia Pascoli (2021), *An optimal nonlinear method for simulating relic neutrinos*, MNRAS, 507 (2), 2614-2631.

Chapter 5 Willem Elbers, Carlos S. Frenk, Adrian Jenkins, Baojiu Li, and Silvia Pascoli (2022), *Higher order initial conditions with massive neutrinos*, MNRAS, 516 (3), 3821–3836.

Chapter 6 Willem Elbers (2022), *Geodesic motion and phase-space evolution of massive neutrinos*, JCAP, 11/058.

Chapter 7 Willem Elbers, Carlos S. Frenk, Adrian Jenkins, Baojiu Li, Silvia Pascoli, Jens Jasche, Guilhem Lavaux, and Volker Springel (2023), *Where shadows lie: reconstruction of anisotropies in the neutrino sky*, JCAP, 10/010.

The author also contributed to the following papers related to the PhD work:

I. Euclid Collaboration (2023), *Euclid: Modelling massive neutrinos in cosmology – a code comparison*, JCAP, 06/035.

II. Matthieu Scaller et al. (2023), SWIFT*: A modern highly-parallel gravity and smoothed particle hydrodynamics solver for astrophysical and cosmological applications*, arXiv:2305.13380.

III. Joop Schaye et al. (2023), *The FLAMINGO project: cosmological hydrodynamical simulations for large-scale structure and galaxy cluster surveys*, MNRAS, 526 (4), 4978-5020.

IV. Roi Kugel et al. (2023), *FLAMINGO: calibrating large cosmological hydrodynamical simulations with machine learning*, MNRAS, 526 (4), 6103-6127.

V. Ian G. McCarthy et al. (2023), *The FLAMINGO project: revisiting the S_8 tension and the role of baryonic physics*, MNRAS, 526 (4), 5494-5519.

VI. Amol Upadhye et al. (2023), *Non-linear CMB lensing with neutrinos and baryons: FLAMINGO simulations vs. fast approximations*, arXiv:2308.09755.

Part I

Background

Introduction

1

The first cosmologists lived in Miletus, once a Greek metropolis on the Meander looking out over the Icarian Sea. As early as the sixth century BC, they speculated about the Universe and relied on observation and reason to devise naturalistic theories of the cosmos [1–4]. They held that all of nature was made of one substance, taking celestial bodies out of the realm of the mystical and placing people and stars on equal footing. In so doing, they anticipated the existence of universal laws of nature. The first real demonstration of this premise occurred in the 17th century, when Newton discovered his law of universal gravitation. However, it was not until the 20th century, following the developments of quantum mechanics and nuclear physics, that hydrogen was revealed as the primary building block of stars and nuclear fusion as the energy source that fuels them [5, 6].

A Universe governed by laws becomes itself a laboratory. A modern example is the solar neutrino problem. Neutrinos are subatomic particles produced in nuclear reactions. In the 1960s, when the Homestake experiment first detected neutrinos produced in the sun, physicists discovered a tension between the prevailing model of the sun and the detected number of neutrinos, which was smaller than expected by a factor of three [7, 8]. One possible explanation was that the solar model was flawed, but no modification could accommodate all observations. An alternative solution was that neutrinos were not the massless particles predicted by the Standard Model of particle physics, but rather mixtures of particles with different masses [9, 10]. This allows the identity of neutrinos to fluctuate over time, a phenomenon known as neutrino oscillations. The problem was finally settled in the years between 1998 and 2002 when the Super-Kamiokande, Sudbury Neutrino Observatory, and KamLAND experiments found evidence of neutrino oscillations and reconciled the Homestake findings with the solar model prediction [11–13].

Neutrinos are the last Standard Model particles for which the masses have not been measured. Our knowledge of neutrino oscillations indicates that at least two of the three known neutrinos[1] have masses greater than zero, but the absolute mass scale is unknown. This problem is now under siege from two opposing directions. The first approach involves KATRIN, a laboratory experiment the size of a small factory that measures the energy spectrum of electrons produced in radioactive decay [14]. Since neutrinos are massive, they carry away a small amount of mass energy, which shows up

---

[1] The electron neutrino, muon neutrino, and tau neutrino; see Chapter 3.

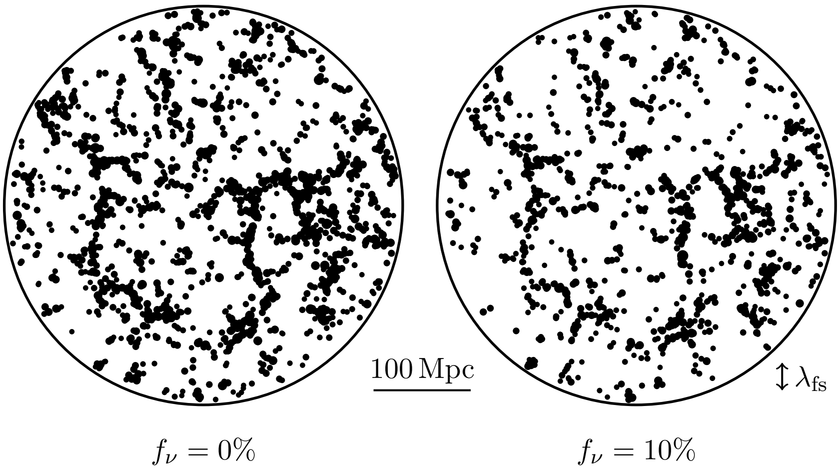

$f_\nu = 0\%$ $f_\nu = 10\%$

Figure 1.1: The distribution of galaxies depends on the neutrino mass fraction f_ν. The plots show the distribution of dark matter halos (where galaxies are born) in simulations without and with neutrinos. The arrow on the right shows the free-streaming length λ_{fs}, corresponding to the size of collapsing regions that neutrinos can escape. This was much larger in the past. The radius of the circles is about 650 million light-years. The size of the dots indicates halo mass.

in the energy spectrum of electrons produced in the same interaction. While KATRIN is ongoing, future experiments are planned with even greater sensitivity [15]. The second approach involves a multitude of astronomical observatories [16–18] that search for the gravitational effects of neutrinos. Although individual neutrinos must be very light, they were created in large numbers in the Big Bang (hundreds per cm^3 today), making them the most abundant particles in the Universe, second only to photons. This means that their collective mass has a gravitational effect of cosmological consequence [19–22].

One of the most promising sources of information on neutrinos is the large-scale distribution of galaxies in the Universe [23–25]. Because their masses are small, neutrinos move much faster than other particles given the same kinetic energy. This allows them to escape regions that are collapsing under the influence of gravity, which inhibits the growth and clustering of galaxies. Figure 1.1, which is based on computer simulations, shows that galaxies become less clustered as the neutrino mass fraction f_ν increases from 0% to 10%. At this level, the effect of neutrinos can even be seen by eye, but the mass fraction could be as small as 0.45%. To detect neutrino masses, we must therefore be able to distinguish between 0% and 0.45% with statistical significance. This requires not only large numbers of observations, but also theoretical calculations that are accurate at the sub-percent level. In Part II of this thesis, we will develop new tools for accurate neutrino calculations and show that this requirement can be met.

Completing the laboratory and cosmological measurements of the neutrino mass is one of the major challenges for physics in the coming years. If the two measurements agree, it would demonstrate that the additional radiation seen in the early Universe[2] is really due to neutrinos, confirming a basic prediction of the Big Bang model: the existence of a Cosmic Neutrino Background. It would also confirm that we understand the formation of structure in the Universe at the sub-percent level and constrain the properties of the dark sector [26, 27]. At the same time, as with the solar neutrino problem, any tension could indicate a flaw in our models, suggesting non-standard neutrino properties or cosmology. Even absent a laboratory measurement, cosmological data can be tested against neutrino oscillations and establish the way neutrino masses are ordered [28, 29]. The measurement is obviously important for particle physics, constraining the mechanism that generates neutrino masses, which must involve physics beyond the Standard Model [30–32].

The thesis is divided into three parts:

- **Part I. Background**: The first part introduces some background material on cosmology and neutrinos, but is intended to be brief.
- **Part II. Technology**: This part develops methods for doing highly accurate computer simulations of the evolution of neutrinos and the formation of structure in the Universe. We begin by introducing a new method for evaluating statistics of the neutrino distribution, which can be used on-the-fly in simulations and significantly reduces the amount of noise without making any approximations. We then develop a perturbation theory for dark matter in the presence of neutrinos and use this to accurately determine the initial conditions of the simulations. Finally, we consider the initial conditions of the neutrinos themselves, taking into account the fact that neutrinos move at nearly the speed of light throughout much of cosmic history.
- **Part III. Prediction**: In the third part, we present two applications of the new methods. First, we perform a comprehensive analysis of the local neutrino background. Among other things, we forecast event rates for direct detection experiments and show that the largest structures in the Universe cast shadows on the neutrino background. In the last chapter, we present the first results from a new suite of galaxy formation simulations, called FLAMINGO. We demonstrate that the outflows of black holes are more powerful in models with neutrinos and discuss the implications for future observations.

However, we begin in the next section with an accessible introduction to one of the first astronomical indicators of the neutrino mass, by way of analogy to moonlight.

---

[2] The detections of $N_{\text{eff}} \approx 3$ from the Cosmic Microwave Background and light elemental abundance.

Figure 1.2: This painting by Claude-Joseph Vernet (1771) depicts a Mediterranean seaport at night. I'd like to imagine that this is Miletus and that the philosophers by the fire are engaged in a debate about moonlight, perhaps putting their theory to test.

1.1. A moonlight parable

Could you use a magnifying glass to set a piece of paper on fire with moonlight? The answer is counter-intuitive to many, as reflected by online discussions of this topic [33–36]. The apparent magnitude of a full moon is $m_\circ = -12.74$, compared with $m_\odot = -26.74$ for the sun. This means that the moon is only $1/400\,000$ times as bright as the sun. Could we compensate with a really big magnifying glass or a clever system of lenses? We will use geometry to show that this is not possible. Bear with me; the same reasoning will also tell us something about the neutrino mass and the nature of dark matter.

There is a limit to the extent to which light can be magnified. If light is concentrated onto a smaller area, then its angular spread must increase. In the following box, we motivate the law of conservation of étendue, which says that the quantity

$$U = A\sin^2\theta, \tag{1.1.1}$$

is constant for a bundle of light rays. In this expression, A is the area covered by the light and θ is the spread in directions.

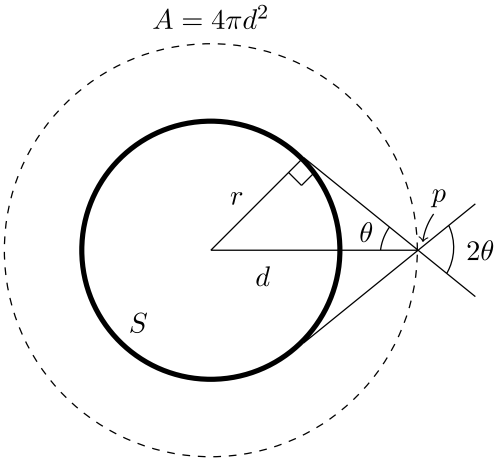

Figure 1.3: Geometric illustration of conservation of étendue, based on a similar figure in [37].

Consider the set-up depicted in Fig. 1.3. A spherical source S of radius r emits light that propagates outward until it hits a larger surface A of radius d. At a point p on A, the light from S is spread over an angle θ determined by the tangents to S that intersect in p. The area of the surface is $A = 4\pi d^2$. Trigonometry tells us that $A \sin^2\theta = 4\pi r^2 = A_S$, where A_S is the surface area of the source. As this holds for an arbitrary surface at any distance d, we have in fact shown that the product $A \sin^2\theta$ is conserved as the light spreads out.

If the surface A were the interface between two media with refractive indices n_1 and n_2, then the incoming angle θ_1 and outgoing angle θ_2 would be related by Snell's law: $n_1 \sin\theta_1 = n_2 \sin\theta_2$. Hence, in that case, the conserved quantity would be $An^2 \sin^2\theta$. This holds more generally[a] and is known as the law of conservation of étendue [37]. For simplicity, we set $n = 1$ above, but this does not affect our conclusion.

---

[a] Technically, the conserved quantity is $\mathrm{d}U = n^2 \mathrm{d}A \cos\theta \mathrm{d}\Omega$, where n is the refractive index of the medium, $\mathrm{d}A$ is the area element, and $\mathrm{d}\Omega$ is the solid angle.

What does this mean for our moonlit fire? The best we could hope to do is concentrate the light from an incoming area A_{in} onto an area A_{out}, related by

$$C_{\max} = \frac{A_{\text{in}}}{A_{\text{out}}} = \frac{\sin(\theta_{\text{out}})^2}{\sin(\theta_{\text{in}})^2} \approx 50\,000, \tag{1.1.2}$$

where we used the fact that the maximum angular spread for the output is $\theta_{\text{out}} = 90°$ and that the angular size of the moon is $\theta_{\text{in}} = 0.26°$. It follows that the maximum flux is still only 1/8 that of the unmagnified sun. Since pieces of paper do not spontaneously ignite when left out in the sun, no magnifying glass will do the trick with moonlight. The magnification makes it appear as if the moon fills the sky and, at some point, we simply run out of sky.

Notice that the étendue of a bundle of light is the product of the spread in location (A) and the spread in direction ($\sin^2\theta$) of the rays. The equivalent for a group of particles would be the product of the spread in position (Δx) and spread in momentum (Δp). In fact, the two are entirely equivalent: étendue can be seen as a volume in phase space. Phase space is the 6-dimensional space in which particles are simultaneously assigned a position (x, y, z) and momentum (p_x, p_y, p_z). Both étendue and phase space volume are related to entropy. They are all measures of spread and disorder. It is worth noting that while étendue is conserved for idealized optics, it may increase in practice but can never decrease. The equivalent statement for particles is that their volume in phase space can only grow or stay the same. This may call to mind the second law of thermodynamics.

In the 1930s, astronomers found that stars and galaxies were moving faster than could be explained with Newtonian mechanics, given the amount of matter that was visible [38–40]. This eventually led to the notion that galaxies could be immersed in halos of invisible dark matter [41–44]. A simple model for the distribution of mass in a galaxy is the singular isothermal sphere. At radius r, the mass density in this model is given by

$$\rho(r) = \frac{\sigma^2}{2\pi G r^2}, \tag{1.1.3}$$

where G is Newton's constant and σ is the spread in velocities of the particles.

During the 1970s, neutrinos were thought to be perfect candidates for dark matter [45–47]. They could be produced in abundance during the Big Bang, do not emit any light, and definitely exist. If neutrinos are responsible for the additional mass in galaxies, they must clump together sufficiently strongly. However, just as we saw for moonlight, neutrinos cannot be compressed beyond a certain limit. In a classic paper, Tremaine & Gunn [48] used the conservation of phase space volume to make a statement about the neutrino mass, under the assumption that neutrinos dominate the mass in galaxies.

In the early Universe, neutrinos were not clumped together but distributed nearly homogeneously. If the initial spread in positions was $(\Delta x_{\text{in}})^3$ and if the maximum concentration is $C_{\max}$, then the mass density of the neutrinos satisfies

$$\rho(r) \leq C_{\max}\,\frac{m_\nu}{(\Delta x_{\text{in}})^3}, \tag{1.1.4}$$

where m_ν is the neutrino mass.

Neutrinos belong to the class of fermion particles, which satisfy the Pauli exclusion principle. This principle states that no two particles can occupy the same quantum state. In phase space, quantum states can be thought of as discrete points separated by multiples of Planck's constant h. The Pauli exclusion principle then imparts the following inequality on the phase space volume:

$$\Delta x^3 \Delta p^3 \geq h^3. \tag{1.1.5}$$

Demanding that the initial phase space volume (in the early Universe) is less than or equal to the final phase space volume (when they are part of the halo), we find that the neutrinos cannot be compressed by more than

$$C_{\max} = \frac{(\Delta x_{\text{in}})^3}{(\Delta x_{\text{out}})^3} = \frac{(\Delta p_{\text{out}})^3}{(\Delta p_{\text{in}})^3} = \frac{m_\nu^3 \sigma^3}{h^3}\,(\Delta x_{\text{in}})^3, \tag{1.1.6}$$

where we used $(\Delta p_{\text{in}})^3 \geq h^3/(\Delta x_{\text{in}})^3$ and $\Delta p_{\text{out}} = m_\nu \sigma$. Note that it is the conservation of phase space volume and not just the exclusion principle that gives rise to the limit. The same argument would work for classical particles following a Maxwell-Botzmann distribution[3]. Combining equations (1.1.3), (1.1.4), and (1.1.6), we conclude that

$$m_\nu \geq \left[\frac{h^3}{2\pi G \sigma r^2}\right]^{1/4}. \tag{1.1.7}$$

This says that the neutrino mass m_ν must be larger than some quantity that depends on the radius r and velocity dispersion σ of galaxies, connecting the modest neutrino with objects comprising billions of suns.

In the original paper [48], this argument led to a value of $m_\nu \geq 24\,\text{eV}$. Not long after, a laboratory measurement of $m_\nu \approx 30\,\text{eV}$ for the electron neutrino was reported [49]. This was spectacular, because a mass of this magnitude would correspond to a flat Universe dominated by neutrinos. However, some of the first cosmological simulations showed that the distribution of galaxies and the properties of galaxy clusters in such a model would be incompatible with observations [24, 50, 51], suggesting that the reported value was wrong. Indeed, a recent laboratory constraint from KATRIN indicates that $m_\nu \leq 0.8\,\text{eV}$ [14], which on its own shows that neutrinos cannot be responsible for more than $\sim 20\%$ of the dark matter. In this case, the bound (1.1.7) does not apply and we must resort to alternative calculations. Nevertheless, this short history reveals that the simultaneous pursuit of cosmological and laboratory strategies is of tremendous value.

---

[3] But not for particles with a Bose-Einstein distribution, which is unbounded at low energies.

Cosmology

2

This chapter introduces the ΛCDM model, the Eulerian and Lagrangian approaches to the perturbation theory of structure formation, and the role of N-body simulations. We establish conventions and point out ways in which neutrinos are special.

2.1. The ΛCDM model

The ΛCDM model, fashionably called the standard model of cosmology, ties together a range of observations of the Universe at different epochs with an economical theory and a limited number of parameters [52–54]. Let us be equally parsimonious in our recitation of its ingredients. In its simplest form, the ΛCDM model assumes a spatially homogeneous and isotropic geometry given by the Friedmann-Lemaître-Robertson-Walker metric[1]

$$\mathrm{d}s^2 = -\mathrm{d}t^2 + a(t)^2\mathrm{d}\mathbf{x}^2, \tag{2.1.1}$$

with a scale factor $a(t)$, whose present value is normalized to 1. In an expanding Universe, the wavelengths of photons emitted at time t are redshifted by a factor $1+z = a(t)^{-1}$. We will frequently work with conformal time τ, defined such that $\mathrm{d}s^2 = -a^2(\mathrm{d}\tau^2 - \mathrm{d}\mathbf{x}^2)$.

According to the model, the main categories of matter are dark energy, cold dark matter, baryons, neutrinos, and photons. The assumptions of homogeneity and isotropy imply that matter behaves as a perfect fluid, described in terms of an energy density $\rho(t)$ and pressure $P(t)$. In the rest frame of the fluid, we can write the conservation of energy as

$$\dot{\rho} + 3\frac{\dot{a}}{a}(\rho + P) = 0, \tag{2.1.2}$$

where overdots denote time derivatives. For a constant equation of state $w = P/\rho$, the solution is

$$\rho(t) \propto a(t)^{-3(1+w)}. \tag{2.1.3}$$

---

[1] In this chapter and the next, we set $c = 1$.

A gas of non-relativistic particles has negligible pressure ($w = 0$), whereas a photon gas has radiation pressure ($w = 1/3$), and dark energy has negative pressure (in its simplest form, $w = -1$). Neutrinos are rather special, having transitioned from relativistic to non-relativistic velocities in recent cosmic history ($0 < w(t) < 1/3$). Given the total energy density ρ of all these species, the Friedmann equation gives the Hubble rate as

$$H^2 \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho. \tag{2.1.4}$$

Energy densities are conveniently expressed as fractions of a critical value $\rho_{\text{crit}} = 3H_0^2/(8\pi G)$, where $H_0 \equiv 100h\,\text{km}\,\text{s}^{-1}\,\text{Mpc}^{-1}$ is the value of the Hubble constant:

$$\Omega_i = \frac{\rho_i}{\rho_{\text{crit}}}, \quad \text{with} \quad i = \text{c}, \text{b}, \nu, \gamma, \Lambda \tag{2.1.5}$$

for cold dark matter (CDM), baryons, neutrinos, photons, and dark energy. The Hubble rate can then be expressed as

$$H^2 = H_0^2 \left[\frac{\Omega_{\text{c},0} + \Omega_{\text{b},0}}{a^3} + \Omega_\nu(a) + \frac{\Omega_{\gamma,0}}{a^4} + \Omega_\Lambda\right], \tag{2.1.6}$$

where subscript 0 indicates present values. Approximately, $\Omega_{\text{c},0} = 0.26$, $\Omega_{\text{b},0} = 0.05$, $\Omega_{\gamma,0} = 5 \times 10^{-5}$, $\Omega_\Lambda = 0.68$. The evolution of the neutrino density $\Omega_\nu(a)$ depends non-trivially on the neutrino temperature T_ν and masses $m_{\nu,i}$ (Chapter 3). Its present value is related to the sum of neutrino masses:

$$\Omega_{\nu,0} \cong \frac{\sum m_\nu}{93.14h^2}, \tag{2.1.7}$$

where the masses are in eV and the value is somewhere in the range $\Omega_{\nu,0} \lesssim 0.007$. Closure implies that $\sum_i \Omega_i = 1$, eliminating one free parameter. To describe the geometry, we therefore need the following parameters: $h, \Omega_{\text{cb}}, \Omega_\gamma, T_\nu, m_{\nu,i}$, where $\Omega_{\text{cb}} = \Omega_{\text{c}} + \Omega_{\text{b}}$.

2.2. Structure formation

At the time of CMB decoupling, $z \approx 1100$, density perturbations in the Universe are of the order 10^{-5} and nearly Gaussian, as shown in Fig. 2.1. Perhaps its greatest achievement is that the ΛCDM model can explain how these perturbations grow into the observed large-scale structure (depicted in simulated form in Fig. 1.1), through the processes of gravitational collapse and hierarchical structure formation.

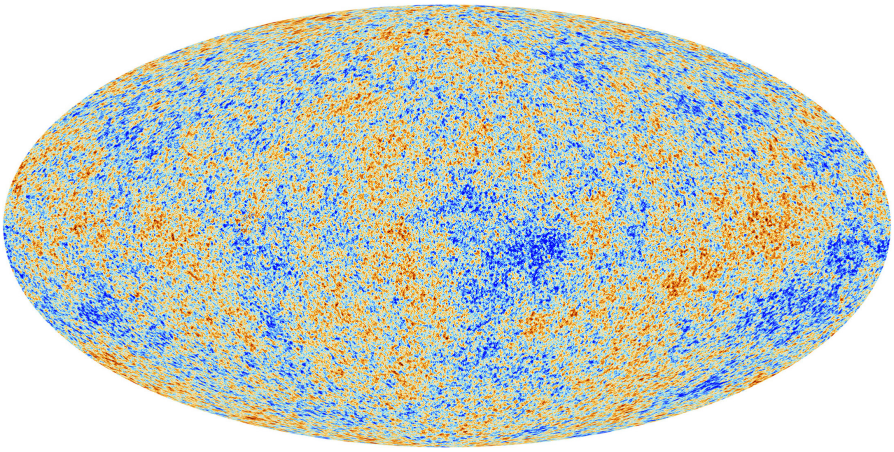

Figure 2.1: Temperature anisotropies of the Cosmic Microwave Background as measured by the Planck satellite [55]. Variations are $\mathcal{O}(10^{-5})$ after subtracting the dipole perturbation. In Chapter 7, we will make predictions for the analogous map of the Cosmic Neutrino Background.

The density field can be decomposed as

$$\rho(\mathbf{x},\tau) = \bar{\rho}(\tau)\left[1+\delta(\mathbf{x},\tau)\right], \tag{2.2.1}$$

where $\bar{\rho}(\tau)$ is the mean density and $\delta(\mathbf{x},\tau)$ the density contrast. The statistical properties of a Gaussian random field with mean zero are completely determined by its power spectrum $P(k)$, defined as

$$\langle\delta(\mathbf{k})\delta(\mathbf{k}')^*\rangle = (2\pi)^3\delta^{(3)}(\mathbf{k}+\mathbf{k}')P(k), \tag{2.2.2}$$

where $\delta(\mathbf{k})$ is the Fourier transform of the density contrast, $\delta^{(3)}$ the Dirac delta function, and we dropped the time dependence for a moment. In its simplest form, the model calls for two additional parameters (A_s, n_s) to describe the power spectrum:

$$P(k,\tau) = \frac{2\pi^2}{k^3}A_s\left(\frac{k}{k_p}\right)^{n_s-1}T(k,\tau)^2, \tag{2.2.3}$$

where A_s is the scalar amplitude, n_s the spectral index, $k_p = 0.05\,\mathrm{Mpc}^{-1}$ a pivot scale, and $T(k,\tau)$ a transfer function that relates the density perturbations at some later time τ to the primordial perturbations encoded by A_s and n_s. The normalization A_s can also be expressed in terms of the amplitude of matter fluctuations on $8h^{-1}\,\mathrm{Mpc}$ scales σ_8.

To set the stage, let us first consider the basics of structure formation without neutrinos, mostly following [56]. We will point out where common assumptions break down for neutrinos. Allowing for perturbations in the metric (2.1.1), assuming only scalar perturbations are present and working in Newtonian gauge, we have

$$\mathrm{d}s^2 = a(\tau)^2 \left[-(1+2\psi)\mathrm{d}\tau^2 + (1-2\phi)\mathrm{d}\mathbf{x}^2\right], \tag{2.2.4}$$

where ϕ and ψ are the perturbations, ϕ being equal to the Bardeen potential Φ. In the non-relativistic limit, anisotropic stress vanishes and $(\phi - \psi) = 0$. This is not the case in the presence of neutrinos, so long as the relativistic tail of the distribution function is important. Ignoring neutrinos, the Poisson equation reads

$$\nabla^2\Phi(\mathbf{x},\tau) = \frac{3}{2}\Omega_\mathrm{m}(\tau)(aH)^2\delta(\mathbf{x},\tau), \tag{2.2.5}$$

where $\Omega_\mathrm{m} = \Omega_\mathrm{cb} = \Omega_\mathrm{c} + \Omega_\mathrm{b}$ and $\delta(\mathbf{x},\tau)$ is the mass-weighted density contrast of CDM and baryons. For collisionless non-relativistic particles (i.e. generally not neutrinos) with mass m, moving in a gravitational potential Φ, the equations of motion are

$$\frac{\mathrm{d}\mathbf{x}}{\mathrm{d}\tau} = \frac{\mathbf{p}}{ma}, \qquad \frac{\mathrm{d}\mathbf{p}}{\mathrm{d}\tau} = -ma\nabla\Phi. \tag{2.2.6}$$

Liouville's theorem that phase-space density is conserved gives rise to the collisionless Boltzmann equation, also known as the Vlasov equation,

$$\frac{\mathrm{d}f}{\mathrm{d}\tau} \equiv Lf \equiv \left[\frac{\partial}{\partial\tau} + \frac{\mathrm{d}\mathbf{x}}{\mathrm{d}\tau}\cdot\nabla + \frac{\mathrm{d}\mathbf{p}}{\mathrm{d}\tau}\cdot\nabla_\mathbf{p}\right] f = 0, \tag{2.2.7}$$

where L is the Liouvillian and $f(\mathbf{x},\mathbf{p},\tau)$ the phase-space distribution function.

2.2.1. Eulerian approach

In Eulerian perturbation theory, also known as standard perturbation theory or SPT, the particle description is related to a fluid description by taking moments of the distribution function. For non-relativistic particles, we define

$$\rho(\mathbf{x},\tau) \equiv a^{-3}\int \mathrm{d}^3\mathbf{p}\; m\, f(\mathbf{x},\mathbf{p},\tau), \tag{2.2.8}$$

$$\rho(\mathbf{x},\tau)\mathbf{u}(\mathbf{x},\tau) \equiv a^{-3}\int \mathrm{d}^3\mathbf{p}\; \frac{\mathbf{p}}{a}\, f(\mathbf{x},\mathbf{p},\tau), \tag{2.2.9}$$

$$\rho(\mathbf{x},\tau)u_i(\mathbf{x},\tau)u_j(\mathbf{x},\tau) + \sigma_{ij}(\mathbf{x},\tau) \equiv a^{-3}\int \mathrm{d}^3\mathbf{p}\; \frac{p_i p_j}{ma^2}\, f(\mathbf{x},\mathbf{p},\tau), \tag{2.2.10}$$

where $\boldsymbol{\sigma}$ is the stress tensor. Taking moments of the Boltzmann equation (2.2.7) gives

$$\frac{\partial\delta}{\partial\tau} + \nabla\cdot[(1+\delta)\mathbf{u}] = 0, \tag{2.2.11}$$

$$\frac{\partial\mathbf{u}}{\partial\tau} + aH\mathbf{u} + \mathbf{u}\cdot\nabla\mathbf{u} = -\nabla\Phi - \frac{\nabla\boldsymbol{\sigma}}{\rho}. \tag{2.2.12}$$

Notice that the lowest moment of the Boltzmann equation gives the continuity equation (2.2.11), which depends on the first moment $\mathbf{u}$. Meanwhile the first moment of the Boltzmann equation gives the Euler equation (2.2.12), which depends on the second moment $\boldsymbol{\sigma}$. We can carry out this expansion to successively higher moments, which is a common strategy used to compute the neutrino perturbations (see Section 3.3). Unlike for neutrinos near or below the free-streaming scale k_{fs}, the stress tensor can often be neglected for CDM and baryons: $\boldsymbol{\sigma}\approx 0$. In this case, we obtain at first order:

$$\frac{\partial\delta}{\partial\tau} + \nabla\cdot\mathbf{u} = 0, \tag{2.2.13}$$

$$\frac{\partial\mathbf{u}}{\partial\tau} + aH\mathbf{u} = -\nabla\Phi. \tag{2.2.14}$$

Separating $\delta(\mathbf{x},\tau) = D(\tau)\delta(\mathbf{x})$ into temporal and spatial factors, these two equations can be combined with the Poisson equation (2.2.5) to give

$$\frac{\mathrm{d}^2D}{\mathrm{d}\tau^2} + aH\frac{\mathrm{d}D}{\mathrm{d}\tau} = \frac{3}{2}\Omega_{\text{m}}(\tau)(aH)^2D(\tau). \tag{2.2.15}$$

Picking out the fastest growing solution, one obtains the first-order growth factor $D_+(\tau)$. If $\Omega_{\text{m}} = 1$, one finds $D_+(\tau) = a(\tau)$. Analytic expressions can also be given in the cases that $\Omega_{\text{m}} < 1; \Omega_\Lambda = 0$ [57], $\Omega_{\text{m}} + \Omega_\Lambda = 1$ [58, 59], and relevant in the presence of neutrinos: $\Omega_{\text{cb}} + \Omega_\Lambda < 1$ (Chapter 5). The first-order solution can then be used in (2.2.11–2.2.12) to obtain equations for the next higher order and so on. Indeed, the SPT equations can be conveniently written and solved in recursive form [60, 61].

2.2.2. Lagrangian approach

An alternative philosophy is the Lagrangian approach. In Lagrangian perturbation theory (LPT), one follows fluid trajectories

$$\mathbf{x}(\mathbf{q},\tau) = \mathbf{q} + \boldsymbol{\psi}(\mathbf{q},\tau), \tag{2.2.16}$$

where $\boldsymbol{\psi}$ is a displacement vector that maps initial particle positions $\mathbf{q}$ to final positions $\mathbf{x}$. In the single-stream limit, the Eulerian density contrast can be related to the displacement via the conservation of mass equation

$$\delta(\mathbf{x},\tau) = \frac{1}{J(\mathbf{q},\tau)} - 1, \tag{2.2.17}$$

where $J(\mathbf{q},\tau)$ is the determinant of the Jacobian of the coordinate transformation $\mathbf{q} \to \mathbf{x}$. The single-stream limit is not a good approximation for neutrinos near or below the free-streaming scale k_{fs}, precluding the use of this equation. The Lagrangian version of the Euler equation (2.2.12) becomes

$$\frac{\partial^2 \mathbf{x}}{\partial \tau^2} + aH\frac{\partial \mathbf{x}}{\partial \tau} = -\nabla_{\mathbf{x}}\Phi. \tag{2.2.18}$$

To first order, Eq. (2.2.17) is simply $\delta = -\nabla_{\mathbf{q}}\boldsymbol{\psi}$. Using this in the equation of motion above, one obtains the Zel'dovich solution [62]

$$\nabla \cdot \boldsymbol{\psi}(\tau) = -D(\tau)\delta(\mathbf{q},\tau), \tag{2.2.19}$$

where $D(\tau)$ satisfies the same equation (2.2.15) as the first-order Eulerian solution. This perturbation theory can similarly be carried out to higher orders. As in the case of SPT, LPT can be formulated and solved recursively [59, 63–66].

2.3. Simulations

In addition to the basic Eulerian and Lagrangian schemes described above, many alternative perturbation theories have been developed over the years [e.g. 67–72]. Nevertheless, all perturbation theories eventually break down and their range of applicability is limited to linear and mildly nonlinear scales. This can be understood in many different ways. The spherical collapse model shows that self-gravitating objects form when the linear density contrast $\delta \approx 1.68$, at which point an expansion in δ^n cannot be expected to converge. In practice, SPT breaks down long before this. Gravitational collapse also generates virial motions, such that $\sigma_{ij} \neq 0$ inside collapsed regions. From the point of view of effective field theory, integrating out short wavelength modes generates an effective pressure term [73]. Either way, at this point the single stream limit is no longer a good approximation. Although formally LPT remains consistent until after first shell crossing [74], the use of (2.2.17) implies that it is no longer physical.

While perturbation theory eventually breaks down, the underlying phase-space description in terms of characteristic equations (2.2.6) and Boltzmann equation (2.2.7) remains valid. These equations can be solved even in the deeply nonlinear régime with numerical techniques. The essence of N-body simulations is to solve the Boltzmann equation by the method of characteristics, i.e. by sampling the phase-space distribution $f(\mathbf{x}, \mathbf{p}, \tau)$ with discrete particles and tracing their evolution along characteristic curves. Conceptually, the simplest approach is to compute the gravitational forces by direct summation over the $N(N-1)/2$ particle pairs. In practice, far more efficient algorithms are used.

Many modern N-body codes solve at least part of the problem with a spectral method, known as the particle mesh (PM) approach [75]. Pure PM codes are fast but have a force resolution determined by the size of the mesh, which limits their usefulness to low resolution applications. To resolve highly nonlinear structures, adaptive methods are called for. Two of the most common methods are the Tree-PM and FMM-PM algorithms, which use the PM method for periodic large-scale gravitational forces and octree- or multipole-based methods for the efficient evaluation of small-scale forces with greatly improved force resolutions. The two main codes used in this thesis, SWIFT [76] and GADGET-4 [77] employ these two algorithms. Other popular codes include[2] GLAM [78], FASTPM [79], ABACUS [80], PKDGRAV3 [81], RAMSES [82], AREPO [83], CUBEP^{3}M [84], GEVOLUTION [85], CONCEPT [86], and GRAMSES [87].

Of particular interest are applications at the interface between cosmological perturbation theory and simulations, a field that is now blossoming. The obvious example is the set-up of initial conditions for cosmological simulations, a major theme in this thesis. While in the single-stream limit (an assumption we cannot generally make for neutrinos; Chapter 6), the initial displacements and velocities of particles are uniquely determined by the LPT solutions $\boldsymbol{\psi}^{(n)}$ and $\mathbf{v}^{(n)} = \mathrm{d}\boldsymbol{\psi}^{(n)}/\mathrm{d}\tau$. Going to higher n brings with it substantial benefits in terms of accuracy and computational expense [88, 89] (see Chapter 5). But there are many other exciting applications. LPT has been to great success to speed up the convergence of PM simulations [90] and can also be used as a cheap control variate to improve the statistical properties of ensembles of simulations [91–93]. In similar vein, the basic method for neutrino simulations introduced in Chapter 4 can be extended to higher-order versions with perturbation theory. Combining the power of N-body simulations to model the collisionless dynamics of dark matter particles with a Lagrangian bias expansion allows the clustering of biased tracers to be modelled accurately and efficiently [94]. On the Eulerian side, the SPT recursion relations can be formulated on a grid, enabling generative analyses that are directly analogous to simulations [95].

Finally, N-body simulations provide the ideal platform for studying galaxy formation

---

[2] We will soon be adding our own to the mix: a scalable PM code called SEDULUS.

in the cosmological context. There exist two broad philosophies in this field [96–98]. The first is to develop semi-analytical galaxy formation models based on the merger histories and properties of dark matter halos in collisionless N-body simulations [e.g. 96]. The second is to integrate self-consistent galaxy formation models with hydrodynamical simulations [e.g. 98], with methods such as smoothed particle hydrodynamics [99]. Since the resolution of cosmological simulations is typically insufficient to resolve all relevant processes, such as the evolution of stellar populations or the physics of black hole accretion disks, effective "subgrid" models are used with free parameters that can be matched to observations or chosen based on theoretical considerations. We will consider a first application of the new FLAMINGO suite of hydrodynamical simulations in Chapter 8.

Neutrinos

3

This chapter introduces some background material on neutrinos in the standard model, neutrino oscillations, current mass bounds, the role of neutrinos in cosmology, and cosmological perturbation theory for massive neutrinos. The structure mirrors that of the previous chapter, introducing basic concepts first before dealing with linear perturbations.

3.1. Neutrinos in a nutshell

In the Standard Model of particle physics, neutrinos are massless fermions that only participate in weak interactions [100]. There are three left-handed particles accompanying charged leptons with different flavours: the electron neutrino ν_e, the muon neutrino ν_μ, and the tau neutrinos ν_τ, and three corresponding anti-neutrinos $\bar{\nu}_e, \bar{\nu}_\mu, \bar{\nu}_\tau$. The part of the Standard Model Lagrangian involving neutrinos is

$$\mathcal{L} \supset \frac{-g}{2\cos\theta_W} \sum_{\alpha=e,\mu,\tau} \bar{\nu}_{\alpha L}\gamma^\mu \nu_{\alpha L} Z^0_\mu - \frac{g}{\sqrt{2}} \sum_{\alpha=e,\mu,\tau} \bar{\nu}_{\alpha L}\gamma^\mu \ell^-_{\alpha L} W^+_\mu + \text{h.c.}, \tag{3.1.1}$$

where g is the coupling constant, θ_W the Weinberg angle, γ^μ are Dirac matrices, $\ell_{\alpha L} \in \{e_L, \mu_L, \tau_L\}$ are the left-handed components of the charged leptons, and W_μ and Z_μ are the weak bosons. The first term gives rise to neutral current interactions amongst neutrinos and the second term to charged current interactions between neutrinos and charged leptons. Neutrinos can be produced through processes like β-decay:

$$^3\text{H} \to {}^3\text{He}^+ + e^- + \bar{\nu}_e. \tag{3.1.2}$$

Interactions such as these produce neutrinos of definite flavour. However, the observation of neutrino oscillations implies that the flavour states ν_α (with $\alpha = e, \mu, \tau$) are superpositions of mass eigenstates ν_i (with $i = 1, 2, 3$). The two are related by a unitary matrix, known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [9, 10]:

$$\nu_\alpha = \sum_{i=1}^{3} U^*_{\alpha i}\nu_i. \tag{3.1.3}$$

The mass states ν_i are eigenstates of the Hamiltonian with eigenvalue[1] $E_i = (p^2 + m_i^2)^{1/2}$. We can derive the expression for the neutrino oscillation probability by making some simplifying assumptions, which ultimately gives the same result as a more complete treatment based on wave packets [101–103]. Since the time evolution is described by the Hamiltonian, it is the mass eigenstates that propagate: $|\nu_i(t)\rangle = e^{-iE_it}|\nu_i\rangle$. Now suppose that a neutrino is produced with a definite flavour α and definite momentum p. After propagating some distance, the flavour β is picked out with probability

$$P(\alpha \to \beta) = |\langle \nu_\beta | \nu_\alpha(t)\rangle|^2 = \left| \sum_i^3 U_{\beta i} U_{\alpha i}^* e^{-iE_it} \right|^2, \tag{3.1.4}$$

where it was used that the mass states form an orthonormal basis $\langle \nu_i | \nu_j \rangle = \delta_{ij}$. For relativistic neutrinos with equal momentum p, we can expand

$$E_i - E_j \cong \frac{\Delta m_{ij}^2}{2p}, \tag{3.1.5}$$

where $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$. Pulling out a common factor e^{2iE_jt} for any j, we then obtain

$$P(\alpha \to \beta) = \left| \sum_i^3 U_{\beta i} U_{\alpha i}^* \exp\left(\frac{-i\Delta m_{ij}^2 t}{2p} \right) \right|^2. \tag{3.1.6}$$

This equation demonstrates two key properties of neutrino oscillations. First, it shows that neutrino oscillations are only possible if the mixing matrix is non-trivial, $U \neq I$, and in the presence of non-zero neutrino masses, $\Delta m_{ij}^2 \neq 0$. It also shows that neutrino oscillations are sensitive to the squared mass differences Δm_{ij}^2, but not to the absolute mass scale. Cosmological observables are therefore complementary to neutrino oscillation experiments by being sensitive to the sum of neutrino masses $\sum m_\nu \equiv \sum_i m_i$.

The parameter space describing neutrino oscillations is quite large, consisting of the mass squared differences Δm_{ij}^2 and the matrix elements $U_{\alpha i}$, which can be parametrized in terms of 3 mixing angles θ_{ij} and up to 3 phases. Our knowledge of these parameters was accumulated over more than two decades by a large number of experiments using neutrinos from various sources [103]. Solar neutrinos were first observed by the Homestake experiment [7], while their oscillations were established by the Sudbury Neutrino Observatory [12]. Solar neutrino oscillations provide information about Δm_{21}^2, which is therefore known as the solar mass splitting. Atmospheric neutrinos are produced when cosmic rays hit the atmosphere, producing pions and kaons which subsequently decay, leading to

---

[1] We set the scale factor $a = 1$ in this subsection for simplicity.

predominantly muon and electron neutrinos. They were first detected at the Kolar Gold Fields [104], while their oscillations were discovered by the Super-Kamiokande experiment [11] and contribute information about $|\Delta m^2_{31}|$, which is known as the atmospheric mass splitting. In addition, particle accelerators and nuclear reactors provide controlled sources of neutrinos used by many oscillation experiments starting with KamLAND [13]. The current state of knowledge of the mass splittings, derived from global fits to the entire experimental landscape, can be summarized as [29]

$$\Delta m^2_{21} = m^2_2 - m^2_1 = 7.42^{+0.21}_{-0.20} \times 10^{-5}\ \mathrm{eV}^2, \tag{3.1.7}$$

$$\Delta m^2_{3\ell} = m^2_3 - m^2_\ell = \begin{cases} +2.514^{+0.028}_{-0.027} \times 10^{-3}\ \mathrm{eV}^2\ (\mathrm{NO}), \\ -2.497^{+0.028}_{-0.028} \times 10^{-3}\ \mathrm{eV}^2\ (\mathrm{IO}). \end{cases} \tag{3.1.8}$$

The sign of Δm^2_{21} is known to be positive, but the sign of Δm^2_{31} is unknown, which leaves two possible mass orderings: $m_1 < m_2 < m_3$ with $\ell = 1$ (normal ordering; NO) or $m_3 < m_1 < m_2$ with $\ell = 2$ (inverted ordering; IO). There is currently a slight preference for the normal ordering from oscillation data [28, 29]. Two crucial numbers can be derived from these mass splittings. If the lightest neutrino is massless, we find that the minimum possible value for the sum of neutrino masses is

$$\sum m_\nu \geq \begin{cases} 0.06\,\mathrm{eV} & (\mathrm{NO}), \\ 0.10\,\mathrm{eV} & (\mathrm{IO}). \end{cases} \tag{3.1.9}$$

This is illustrated by the red and black curves in Fig. 3.1. As mentioned above, cosmological observables are sensitive to the mass sum, but much less so to the individual masses [105, 106]. Even so, cosmology can rule out the inverted ordering by establishing that $\sum m_\nu < 0.1\,\mathrm{eV}$. The strongest cosmological bounds are now hinting in this direction [107–110], but the inverted ordering cannot yet be ruled out [111].

The best terrestrial constraint on the absolute mass scale comes from the KATRIN detector, which places a bound of $m_\beta < 0.8$ eV at the 90% C.L. [14] on an effective mass for neutrinos involved in the β-decay of tritium (3.1.2), given by

$$m_\beta = \sqrt{c^2_{12}c^2_{13}m^2_1 + s^2_{12}c^2_{13}m^2_2 + s^2_{13}m^2_3}, \tag{3.1.10}$$

where $c_{ij} = \cos\theta_{ij}$ and $s_{ij} = \sin\theta_{ij}$. Assuming a degenerate mass spectrum, $m_1 \approx m_2 \approx m_3$, this corresponds to a neutrino mass sum of $\sum m_\nu < 2.4\,\mathrm{eV}$. Recent cosmological limits are much stronger, as shown in Fig. 3.1, and are quoted below at the 95% C.L.

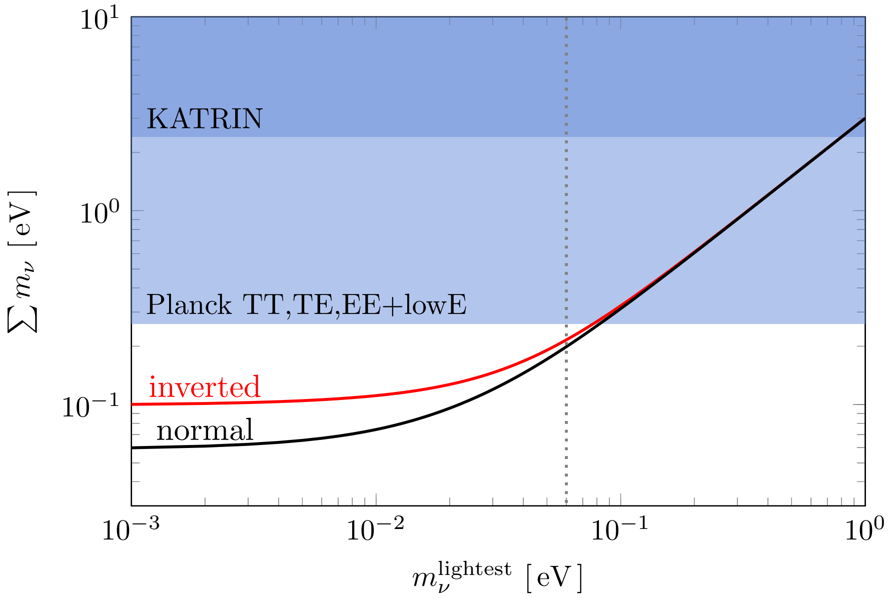

Figure 3.1: Constraints on the sum of neutrino masses, $\sum m_\nu$, from KATRIN [14] and Planck [112], showing the sum as a function of the lightest mass for the two possible mass orderings (normal and inverted). The dotted line indicates where the masses become degenerate: $m_1 \approx m_2 \approx m_2$.

Assuming a degenerate mass spectrum, the Planck temperature, polarization, and lensing likelihoods give a constraint of $\sum m_\nu < 0.24\,\text{eV}$ or $\sum m_\nu < 0.26\,\text{eV}$, depending on the details of the high-ℓ polarization analysis [112].

By complementing the CMB measurements from Planck with probes of the large-scale structure, even stronger constraints are possible. Adding BAO data from BOSS DR12, SDSS MGS, and 6dFGS leads to $\sum m_\nu < 0.12\,\text{eV}$ (degenerate), $\sum m_\nu < 0.15\,\text{eV}$ (normal), and $\sum m_\nu < 0.17\,\text{eV}$ (inverted) [113]. A full-shape analysis of the BOSS DR11 redshift-space power spectrum, combined with Planck temperature and polarization data and JLA of Type Ia supernovae leads to $\sum m_\nu < 0.18\,\text{eV}$ [114]. Combining instead SDSS DR14 BOSS and eBOSS Lyman-α forest data with Planck temperature, polarization, and lensing data and 6dFGS, SDSS, BOSS-LOWZ, CMASS DR12 measurements of BAO leads to the strongest constraint yet: $\sum m_\nu < 0.09\,\text{eV}$ [108]. A similarly strong bound was set through a combination of Planck temperature, polarization, and lensing data, Pantheon Type Ia supernovae, SDSS BAO and RSD measurements of the growth rate: $\sum m_\nu < 0.09\,\text{eV}$ by [109].

It is noteworthy that a recent analysis of the final data release of Planck (PR4) [115] found a significantly relaxed upper bound of $\sum m_\nu < 0.4\,\text{eV}$. This is primarily attributed to the disappearance of an anomalous preference for additional gravitational lensing ($A_{\mathrm{L}} = 1.036 \pm 0.051$ is now consistent with no additional lensing-induced smoothing), which is known to have a profound effect on the inferred neutrino mass [113, 115–117].

Based on this incomplete survey of results, all assuming the most basic ΛCDM + $\sum m_\nu$ model, we make the following observations:

1. The assumed mass ordering becomes important as we near the hierarchical régime.
2. Full-shape power spectrum analyses are not necessarily the most constraining.
3. Geometry (BAO, supernovae, etc.) contributes a lot of constraining power.
4. The most constraining mass bounds stack many different cosmological observables from different epochs.
5. Addressing the lensing anomaly in the CMB is important for a reliable neutrino mass constraint.

Neutrinos are exactly massless in the Standard Model [100]. Therefore, we know that the model is incomplete and must be extended to account for neutrino masses. This can be done in many different ways. One possibility is to add one or more right-handed neutrinos, ν_R, which allows Dirac mass terms of the form

$$\mathcal{L} = m\bar{\nu}_L\nu_R + \text{h.c.}, \tag{3.1.11}$$

to be generated, similar to the way that mass terms are generated for the charged leptons through the Higgs mechanism. If neutrinos are Majorana particles, masses can also be generated without introducing additional particles. However, this comes at a price. The neutrino would be its own anti-particle and conservation of lepton number would be broken. Hence, determining the way that neutrino masses are realized in nature would tell us about the symmetries and particles of the theory that lies beyond.

One way to determine whether neutrinos are Dirac or Majorana is to search for processes that violate lepton number, such as neutrinoless double β-decay [118]. Neither neutrino oscillations nor the usual cosmological probes can distinguish between Dirac and Majorana neutrinos. However, the distinction does become relevant for direct detection of the Cosmic Neutrino Background, as we will see in Chapter 7.

3.2. Thermal history

In the early Universe, neutrinos are in thermal equilibrium with a hot plasma of electrons, positrons, and photons through frequent weak interactions [53, 119–121]. As space expands and the plasma cools down, the interaction rate $\Gamma \sim T^5$ decreases faster than the expansion rate $H \sim T^2$, such that the interactions become rare when $\Gamma \sim H$. This happens around $z \sim 10^{10}$ when $T_{\rm dec} \sim 1\,{\rm MeV}$. While the neutrinos are in thermal equilibrium with the plasma, their energies follow a Fermi-Dirac distribution

$$f(\mathbf{x},\mathbf{p},\tau) = \frac{g_s}{(2\pi)^3}\left[1+\exp\left(\frac{p}{k_b T_\nu}\right)\right]^{-1}, \tag{3.2.1}$$

where g_s is the number of degrees of freedom and T_ν is the neutrino temperature. Because the masses $m_\nu \ll 1\,{\rm MeV}$, the neutrinos are relativistic at this stage and the energy $\epsilon = \sqrt{p^2 + a^2 m_\nu^2} \approx p$. Therefore, the energy distribution can be expressed in terms of the neutrino momentum p. After the neutrinos decouple, the neutrino temperature and momentum both scale with redshift $\propto (1+z)^{-1}$. Therefore, by Liouville's theorem, the momentum distribution (3.2.1) is preserved at zeroth order even after the neutrinos become non-relativistic.

The electron mass, $m_e = 0.51\,{\rm MeV}$, is not much smaller than the temperature at which neutrinos decouple. Hence, not long after neutrino decoupling, it becomes favourable for electrons and positrons to annihilate. Assuming that the neutrinos have completely decoupled at this point, the entropy from e^- and e^+ is entirely transferred to the photons, heating their temperature relative to the neutrino temperature by a factor

$$\frac{T_\gamma}{T_\nu} = \left(\frac{g_+}{g_-}\right)^{1/3} = \left(\frac{11}{4}\right)^{1/3}, \tag{3.2.2}$$

where $g_+ = 11/2$ and $g_- = 2$ are the effective numbers of degrees of freedom before and after electron-positron annihilation. Since the photon temperature is very well known from the CMB, $T_\gamma = 2.7255\,{\rm K}$, it is convenient to express neutrino-related quantities relative to the corresponding photon quantities. The number and energy densities of a relativistic particle species with phase-space distribution (3.2.1) are given by

$$n(\mathbf{x},\tau) = a^{-3}\int {\rm d}^3p\, f(\mathbf{x},\mathbf{p},\tau) = \frac{3}{4}\frac{\zeta(3)}{\pi^2}T^3, \tag{3.2.3}$$

$$\rho(\mathbf{x},\tau) = a^{-4}\int {\rm d}^3p\, \epsilon\, f(\mathbf{x},\mathbf{p},\tau) = \frac{7}{8}\frac{\pi^2}{30}T^4. \tag{3.2.4}$$

For bosons with a Bose-Einstein distribution, similar to (3.2.1) but with a minus sign, we obtain the same expression without the factors 3/4 and 7/8. In the relativistic limit, we therefore obtain

$$n_\nu = \frac{3}{11} N_{\text{eff}}\, n_\gamma, \qquad \rho_\nu = \frac{7}{8}\left(\frac{4}{11}\right)^{4/3} N_{\text{eff}}\, \rho_\gamma, \tag{3.2.5}$$

where we have introduced the effective number of neutrino species $N_{\text{eff}} \approx 3$. In reality, neutrino decoupling does not happen instantaneously and the process partially overlaps with electron-positron annihilation. This introduces small perturbations in the energy distribution (3.2.1) and also means that the temperatures are not exactly related by (3.2.2). A full treatment of these processes, also accounting for flavour differences and neutrino oscillations, yields a slightly greater value of $N_{\text{eff}} = 3.046$ [122, 123]. We can absorb the leading effect of the spectral distortions to (3.2.1) by slightly modifying the neutrino temperature whilst keeping the Fermi-Dirac form [124], which is the strategy followed in this thesis.

In the non-relativistic limit, $\epsilon = am_\nu$, we instead obtain $\rho_\nu = \sum m_{\nu,i}\, n_{\nu,i}$, which gives Eq. (2.1.7). In general, we need to evaluate the integral (3.2.4) numerically and sum over the different neutrino species to obtain $\Omega_\nu(a) = \rho_\nu(a)/\rho_{\text{crit}}$.

3.3. Linear perturbations

Although more limited in the range of scales where it is accurate, an Eulerian fluid approximation can also be given for massive neutrinos. This is sometimes useful for speeding up calculations or to gain physical insight [124, 125]. If we treat the neutrinos as a perfect fluid, the stress tensor becomes $\sigma_{ij} = -P\delta_{ij}$. We can then write the continuity equation (2.2.11) and Euler equation (2.2.12) to first order as

$$\frac{\partial \delta}{\partial \tau} + \nabla \cdot \mathbf{u} = 0, \tag{3.3.1}$$

$$\frac{\partial \mathbf{u}}{\partial \tau} + aH\mathbf{u} = -\nabla\Phi + c_s^2 \nabla^2 \delta, \tag{3.3.2}$$

where we introduced the sound speed $c_s^2 = \delta P/\delta\rho$. Although obviously incorrect, it is useful to consider the case where neutrinos are the dominant matter species with $\Omega_{\text{m}} = \Omega_\nu = 1$. In that case, the combination of the Poisson equation (2.2.5) with the continuity and Euler equations above yields

$$\frac{\partial^2 \delta}{\partial \tau^2} + aH\frac{\partial \delta}{\partial \tau} = \left[\frac{3}{2}(aH)^2 - k^2 c_s^2\right]\delta, \tag{3.3.3}$$

which is analogous to the Jeans equation. Defining the free-streaming scale as [120]

$$k_{\rm fs} = \sqrt{\frac{3}{2}}\frac{aH}{c_s}, \tag{3.3.4}$$

we see that on small scales, $k \gg k_{\rm fs}$, the growth of density perturbations is suppressed and δ oscillates with frequency kc_s. By contrast, on very large scales, $k \ll k_{\rm fs}$, neutrinos cluster like cold dark matter. If we approximate the sound speed by the velocity dispersion of the unperturbed Fermi-Dirac distribution (3.2.1), we find that

$$c_s^2 \approx \sigma_\nu^2 = \frac{15\zeta(5)}{\zeta(3)}\frac{T_\nu^2}{m_\nu^2}. \tag{3.3.5}$$

This additionally shows that the free-streaming wavenumber scales as $k_{\rm fs} \propto m_\nu$.

A complete treatment of linear neutrino perturbations is also possible [126]. We briefly review the derivation here. To start with, the neutrino phase-space distribution is decomposed into a background part and a perturbation part via

$$f(\mathbf{x},\mathbf{p},\tau) = \bar{f}(p)\left[1+\Psi(\mathbf{x},p,\hat{n},\tau)\right], \tag{3.3.6}$$

where $\bar{f}(p)$ is the homogeneous Fermi-Dirac distribution (3.2.1) and the momentum $\mathbf{p} = p\hat{n}$ has been decomposed into a magnitude p and a unit vector $\hat{n}$. With the metric in Newtonian gauge (2.2.4), the acceleration can be written to first order as

$$\frac{\mathrm{d}p}{\mathrm{d}\tau} = p\frac{\partial\phi}{\partial\tau} - \frac{\epsilon}{p}\mathbf{p}\cdot\nabla\psi. \tag{3.3.7}$$

Substituting this into the Boltzmann equation (2.2.7) and switching to Fourier space, we find that the perturbation Ψ evolves as

$$\frac{\partial\Psi}{\partial\tau} + i\frac{p}{\epsilon}(\mathbf{k}\cdot\hat{n})\Psi + \frac{\mathrm{d}\ln\bar{f}}{\mathrm{d}\ln p}\left[\frac{\partial\phi}{\partial\tau} - i\frac{\epsilon}{p}(\mathbf{k}\cdot\hat{n})\psi\right] = 0.$$

To solve this equation, Ψ is decomposed into a Legendre series[2]

$$\Psi(\mathbf{k},\hat{n},p,\tau) = \sum_{\ell=0}^{\infty}(-i)^\ell(2\ell+1)\Psi_\ell(\mathbf{k},p,\tau)P_\ell(\hat{k}\cdot\hat{n}). \tag{3.3.8}$$

---

[2] Note that we define the Ψ_ℓ slightly differently than [126].

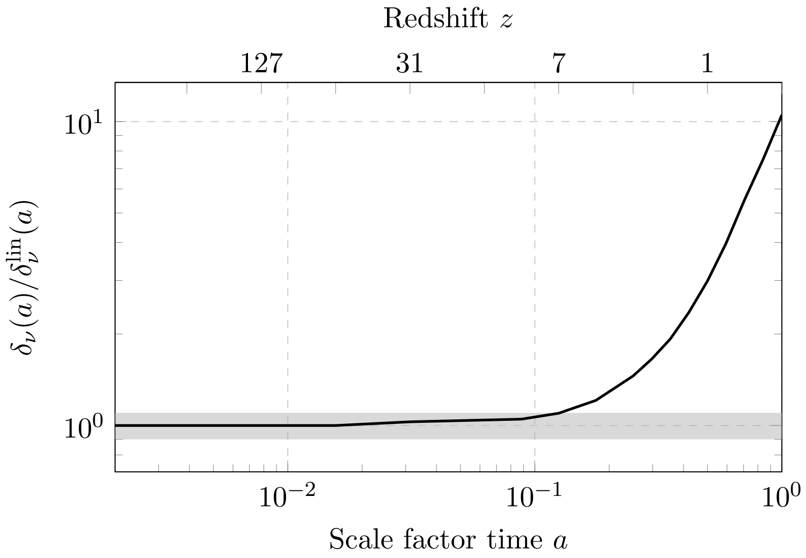

Figure 3.2: Non-linear density perturbations for a neutrino species with $m_\nu = 0.1\,\mathrm{eV}$, computed with the δf method (Chapter 4), relative to the linear theory prediction at $k = 0.6\,\mathrm{Mpc}^{-1}$. The grey bar represents a ±10% error.

The Boltzmann equation then becomes an infinite tower of equations:

$$\frac{\partial \Psi_0}{\partial \tau} = -\frac{pk}{\epsilon}\Psi_1 + \frac{\partial \phi}{\partial \tau}\frac{\mathrm{d}\ln \bar{f}}{\mathrm{d}\ln p}, \tag{3.3.9}$$

$$\frac{\partial \Psi_1}{\partial \tau} = \frac{pk}{3\epsilon}(\Psi_0 - 2\Psi_2) - \frac{\epsilon k}{3p}\psi\frac{\mathrm{d}\ln \bar{f}}{\mathrm{d}\ln p}, \tag{3.3.10}$$

$$\frac{\partial \Psi_\ell}{\partial \tau} = \frac{pk}{(2\ell+1)\epsilon}\left[\ell\Psi_{\ell-1} - (\ell+1)\Psi_{\ell+1}\right], \qquad \text{for} \quad \ell \geq 2. \tag{3.3.11}$$

This system of equations can be solved numerically if we truncate the hierarchy at some $\ell = \ell_{\rm max}+1$. In Einstein-Boltzmann codes like CAMB [127] and CLASS [128], this is often done with an Ansatz such as

$$\Psi_{\ell_{\rm max}+1} = \frac{(2\ell_{\rm max}+1)\epsilon}{pk\tau}\Psi_{\ell_{\rm max}} - \Psi_{\ell_{\rm max}-1}.$$

The momentum dependence sets these equations apart from the equivalent for photons. The neutrino perturbations Ψ need to be integrated for pairs of (k, p) instead of being functions of k alone. Without further approximations, this is often the slowest part of Einstein-Boltzmann runs. Although the growth of neutrino perturbations is suppressed below the free-streaming scale, they nevertheless become nonlinear at late times (Fig. 3.2), such that eventually the perturbation theory described above breaks down. Therefore, it is natural to use cosmological simulations (Section 2.3) to self-consistently solve for the gravitational evolution of cold dark matter, baryons, and massive neutrinos. Various methods exist to treat neutrinos in N-body simulations (see [129, 130] for reviews). In the next chapter, we will present a novel method and compare it with several alternatives from the literature.

Part II

Technology

Optimal nonlinear treatment

4

This chapter deals with the treatment of massive neutrinos in cosmological N-body simulations. It introduces a novel simulation method that is unbiased and consistent, converging to the exact solution in the large particle limit. We prove optimality, in terms of noise, within a broad family of hybdrid methods that encompasses the commonly-used particle and linear mesh methods as special cases.

An optimal nonlinear method for simulating relic neutrinos

Abstract: Cosmology places the strongest current limits on the sum of neutrino masses. Future observations will further improve the sensitivity and this will require accurate cosmological simulations to quantify possible systematic uncertainties and to make predictions for nonlinear scales, where much information resides. However, shot noise arising from neutrino thermal motions limits the accuracy of simulations. In this paper, we introduce a new method for simulating large-scale structure formation with neutrinos that accurately resolves the neutrinos down to small scales and significantly reduces the shot noise. The method works by tracking perturbations to the neutrino phase-space distribution with particles and reduces shot noise in the power spectrum by a factor of $\mathcal{O}\left(10^2\right)$ at $z = 0$ for minimal neutrino masses and significantly more at higher redshifts, without neglecting the back-reaction caused by neutrino clustering. We prove that the method is part of a family of optimal methods that minimize shot noise subject to a maximum deviation from the nonlinear solution. Compared to other methods we find permille level agreement in the matter power spectrum and percent level agreement in the large-scale neutrino bias, but large differences in the neutrino component on small scales. A basic version of the method can easily be implemented in existing N-body codes and allows neutrino simulations with significantly reduced particle load. Further gains are possible by constructing background models based on perturbation theory. A major advantage of this technique is that it works well for all masses, enabling a consistent exploration of the full neutrino parameter space.

4.1. Introduction

The discovery of neutrino masses [11–13] calls for extensions of the Standard Model of particle physics and provides the only known form of dark matter. Measuring the masses is crucial for understanding their origin and for constraining cosmological parameters. While the neutrino mass squared differences are known to a few percent, the absolute masses are unknown and there remain two possible mass orderings: normal and inverted. A rich experimental programme is aimed at determining the mass ordering, measuring the mass scale set by the lightest neutrino and completing the overall picture of neutrino properties. Cosmology plays a vital role in this programme due its ability to provide an independent and complementary constraint on the sum of neutrino masses, $\sum m_\nu$ [23, 25] with a potential sensitivity below 0.02 eV [131–133].

Ongoing and planned neutrino experiments will establish the mass ordering with a discovery expected by the end of the decade. Although oscillation data have shown persistent hints of normal ordering, this preference has decreased to 1.6σ over the past year [29]. The mass ordering can be established by exploiting matter effects in long baseline neutrino oscillation experiments, as in DUNE [134], and in the Earth for atmospheric neutrinos, as in ORCA [135] and Hyper-K [136], as well as vacuum oscillations in medium baseline reactor neutrino experiments, specifically Juno [137]. Each approach is challenging, so information from multiple sources is essential. Single β-decay is the experimental strategy of choice for direct mass searches and provides a model-independent determination of neutrino masses, in particular the effective electron-neutrino mass. The KATRIN experiment is ongoing and has put a bound of $m_\beta < 0.8$ eV , assuming quasi-degenerate neutrino masses, with the aim of reaching $m_\beta < 0.2$ eV in the near future [14]. Project 8 will have the potential to set a limit of $m_\beta < 0.04$ eV [15]. Neutrinoless double β-decay can also provide information on neutrino masses [138–140], albeit entangled with the value of the Majorana CP-violating phases and affected by uncertainty in the nuclear matrix elements [141]. For a recent review see e.g. [142].

The complementarity between these different strategies is of great interest. A cosmological measurement of $\sum m_\nu$ would provide a target for direct mass searches [143, 144]. An incompatibility between the two would indicate a non-standard cosmological evolution or new neutrino properties. A cosmological bound of $\sum m_\nu < 0.1\,\text{eV}$ would suggest a normal mass ordering, which should be confronted with evidence from neutrino experiments. Finally, there is a strong synergy with neutrinoless double β-decay. Knowing the mass ordering and the sum of neutrino masses would narrow down the range of values for the effective Majorana mass parameter, providing a clear target for future experiments.

Measuring the mass scale, and potentially ruling out the inverted mass ordering, is therefore a major target of near-term cosmological surveys, including DESI [16], EUCLID [17], and LSST [18]. In order to analyse these surveys and to extract a mass measurement, there has been a substantial effort to model precisely the effects of massive neutrinos on structure formation. From the analytical side, a swathe of new techniques such as time RG perturbation theory [114] and effective field theories [145, 146], promise to push the validity of perturbation theory into the quasi-linear régime. In the nonlinear régime, N-body simulations offer the most accurate picture of structure formation. Yet incorporating neutrinos into N-body simulations has proved to be a challenge and some doubts remain about the validity of neutrino simulations on small scales.

The main obstacle to simulating neutrinos is that, in contrast to cold dark matter and baryons, neutrinos have a significant velocity dispersion. This effectively turns the 3-dimensional problem of structure formation, for which N-body simulations are well suited, into a 6-dimensional phase-space problem. If no provisions are made, a far greater number of simulation particles is needed to sample properly the phase-space manifold. A further complication arises from the fact that neutrinos are relativistic at high redshifts, such that simulations need to handle both the régime where neutrinos are best described as radiation and the régime where neutrinos are better described as massive particles.

The first 3-dimensional cosmological neutrino simulations were carried out by [147] and [50], when neutrinos were thought to be much more massive and the velocity dispersion not as problematic. Modern simulations with sub-electronvolt neutrinos were pioneered by [148, 149]. Neutrinos are most commonly included in simulations as particles whose initial velocity is the sum of a peculiar gravitational component and a random component sampled from a Fermi-Dirac distribution [148–160]. The main difficulty with particle simulations is shot noise caused by the velocity dispersion. This problem is more severe for the smallest neutrino masses, which are observationally most relevant. Because neutrinos are a subdominant component, the error in the total matter distribution is relatively small. However, shot noise obscures the small-scale behaviour of the neutrinos and is clearly undesirable if one is interested in the neutrino component and its effect on structure formation.

To overcome the problems with particle simulations, grid simulations evolve the neutrino distribution using a system of fluid equations, which requires a scheme to close the moment hierarchy at some low order [149, 161–168], or as a linear response to the non-relativistic matter density [116, 169–171]. Even more efficiently, but in the same spirit of treating neutrinos perturbatively, the total effect of neutrinos has been included as a post-processing step in the form of a gauge transformation [172]. While these approaches do not suffer from shot noise, they are not able to capture the full nonlinear evolution of the neutrinos at late times. This problem becomes more severe for more massive neutrinos, but is present even for minimal neutrino masses. A number of hybrid simulations have therefore combined grid and particle methods [164, 173, 174], typically transitioning from a fluid method to a particle method at some redshift when the neutrinos become nonlinear. Another interesting alternative is to integrate the Poisson-Boltzmann equations directly on the grid [175].

The method proposed in this paper can be considered as a type of hybrid method that integrates neutrino particles but only uses the information contained in the particles to the extent that it is necessary. This is accomplished by dynamically transitioning from a smooth background model to a nonlinear model at the individual particle level. It relies on the noiseless (but approximate) background model as much as possible, thereby producing the smallest amount of shot noise possible whilst solving the full nonlinear system. The main idea is to decompose the phase-space distribution function $f(x,p,t)$ into a background model $\bar{f}(x,p,t)$ which can be solved without noise, and a perturbation which is carried by the simulation particles:

$$f(x,p,t) = \bar{f}(x,p,t) + \delta f(x,p,t). \tag{4.1.1}$$

The choice of background model is arbitrary, but the method performs best whenever $\bar{f}(x,p,t)$ is strongly correlated with $f(x,p,t)$, in a way that will be made precise below. If the choice of background model is poor, the method performs no worse than an ordinary N-body simulation, except for the small amount of overhead associated with evaluating $\bar{f}(x,p,t)$. Note that the background model is just an approximation of f and can itself be a perturbed Fermi-Dirac distribution.

This type of method has a long history in other fields and is variably known as the method of 'perturbation particles' or more commonly as the 'δf method', which is the name we shall adopt. [176] and [177] discussed the method of perturbation particles in stellar dynamics. Around the same time, the δf method arose in plasma physics [178–181]. While the method of perturbation particles is not widely known today in astronomy, the δf method is standard fare in plasma physics. A major difficulty in astronomical applications is the absence of a background model that captures enough of the dynamics to be useful. In contrast, plasma physicists are often interested in turbulent phenomena arising in an otherwise stable system, with a natural candidate for a background model $\bar{f}$ at hand. Our work is motivated by the fact that there is also a natural background model for cosmic neutrinos, namely the phase-space density predicted by perturbation theory. There is a major synergy between δf N-body simulations proposed here and work on improved perturbation theory methods. A better background model means a smaller dependence on the particles and therefore further reduced shot noise. We will show however that even the 0^{th} order approximation, which is just a homogeneous redshifted Fermi-Dirac distribution, provides a significant improvement over ordinary N-body methods.

The remainder of the paper is structured as follows. In Section 4.2, we derive the δf method and describe its use as a variance reduction method for N-body simulations. We also show that the method is part of a family of optimal hybrid methods. In Section 4.3, we illustrate the method with a one-dimensional test problem. In Section 4.4, we discuss how the method can be embedded in relativistic simulations. Our suite of simulations is then described in Section 4.5. The method is compared with commonly-used alternatives in Section 4.6. We consider higher-order background models based on perturbation theory in Section 4.7. Finally, we conclude in Section 4.8.

4.2. Derivation

The phase-space evolution of self-gravitating collisionless particles is described by the Poisson-Boltzmann equations, which in the single-fluid case read

$$Lf \equiv \left[\frac{\partial}{\partial t} + p\cdot\nabla - \nabla\Phi\cdot\nabla_p\right] f = 0, \tag{4.2.1}$$

$$\nabla^2\Phi = 4\pi G\rho = 4\pi G\int \mathrm{d}^3p\sqrt{m^2+p^2}f(x,p,t). \tag{4.2.2}$$

Here, Φ is the gravitational potential, ρ the energy density, and f the phase-space density. In general, the Liouville operator, L, acts on each fluid separately and the potential should be summed over all fluid components. In relativistic perturbation theory, this system can be written as a hierarchy of moment equations for the neutrinos, which is solved to first order with Boltzmann codes such as CLASS [128] or CAMB [127]. To extend our predictions to the nonlinear régime, we can use N-body codes, which solve the Poisson-Boltzmann system by the method of characteristics. Characteristic curves satisfy

$$\frac{\mathrm{d}x}{\mathrm{d}t} = p \qquad \text{and} \qquad \frac{\mathrm{d}p}{\mathrm{d}t} = -\nabla\Phi. \tag{4.2.3}$$

By construction, one finds that $\mathrm{d}f/\mathrm{d}t = Lf = 0$ along these curves. To infer statistics of the phase-space distribution, we simulate N of these trajectories using marker particles. We can freely choose the phase-space distribution, g, of our simulation particles at the initial time. For instance, assuming an initially homogeneous spatial distribution and momenta from the Fermi-Dirac distribution, we would have $g \propto (\exp\{p/T_\nu\} + 1)^{-1}$. Typically, one chooses $g(x, p, t_0) = f(x, p, t_0)$. Since $Lg = 0$, this equality then holds for all $t \geq t_0$. In general, a phase-space statistic is given by

$$\begin{aligned} A(x,t) = \langle A \rangle_p &= \int \mathrm{d}^3p \; f(x,p,t) A(x,p,t) \\ &\cong \frac{1}{N} \sum_{i=1}^{N} \frac{f(x_i,p_i,t)}{g(x_i,p_i,t)} A(x_i,p_i,t). \end{aligned} \tag{4.2.4}$$

Following the usual choice of setting $g(x, p, t_0) = f(x, p, t_0)$, the sum reduces to a simple average over marker particles. The error in our estimate of A is then $\sigma_A/\sqrt{N}$. Hence, if the distribution, $f(x, p, t)$, has a large intrinsic scatter, we need a large N to beat down the noise. Alternatively, we might construct an estimator with a smaller error. Let us therefore write the phase-space distribution function, f, as a background model, $\bar{f}$, together with some perturbation, δf:

$$f(x,p,t) = \bar{f}(x,p,t) + \delta f(x,p,t). \tag{4.2.5}$$

We can reduce the error by only using the particles to estimate the perturbed distribution, δf. We replace (4.2.4) with

$$A(x,t) = \int \mathrm{d}^3p \; \left[\bar{f}(x,p,t) + \delta f(x,p,t)\right] A(x,p,t) \tag{4.2.6}$$

$$\cong \bar{A}(x,t) + \frac{1}{N} \sum_{i=1}^{N} \frac{\delta f(x_i,p_i,t)}{g(x_i,p_i,t)} A(x_i,p_i,t). \tag{4.2.7}$$

This is useful if

$$\bar{A}(x,t) = \int \mathrm{d}^3p \; \bar{f}(x,p,t) A(x,p,t) \tag{4.2.8}$$

can be computed efficiently and if f and $\bar{f}$ are strongly correlated, so that the second term is small. The simplest choice of background model is a homogeneous Fermi-Dirac distribution

$$\bar{f}(x,p,t) = \frac{g_s}{(2\pi)^3} \frac{1}{e^{ap/(k_b T_\nu)} + 1}, \tag{4.2.9}$$

with g_s internal degrees of freedom. Here, $a = a(t)$ is the scale factor, $T_\nu = 1.95$ K the present-day neutrino temperature, and ap the present-day momentum. Since the noise reduction scales with the correlation between $\bar{f}$ and f, we can achieve further gains by adding more information to the background model. The obvious next step is to use perturbation theory to improve on (4.2.9). This option is considered in Section 4.7.

4.2.1. Implementation

In Appendix 4.E, we outline the practical steps needed to implement the method in cosmological N-body simulations. In essence, whenever we sum over neutrino particles, such as when calculating the gravitational force on a test particle, we replace the particle mass with a statistically weighted mass:

$$m \to m w_i = m \left[\frac{\delta f(x_i, p_i, t)}{g(x_i, p_i, t)} \right] = m \left[\frac{f(x_i, p_i, t) - \bar{f}(x_i, p_i, t)}{g(x_i, p_i, t)} \right]. \tag{4.2.10}$$

The weights are computed by comparing the true phase-space density with the background model. We know the background model density, because we can evaluate (4.2.9) at any time. We also know the true density for each particle, because $Lf = Lg = 0$ along characteristic curves. It is therefore sufficient to record the two numbers f and g at the initial sampled location of each particle in phase space. We note that any sampling distribution g is valid provided that $g \neq 0$ almost everywhere $f \neq 0$. We will continue to use the common choice, $g = f$, where f is the Fermi-Dirac distribution. In general, the optimal choice of g will depend on the phase-space statistic of interest. Choosing a distribution g that oversamples slower particles can provide an additional reduction in shot noise. Given the homogeneous Fermi-Dirac background model (4.2.9), the neutrino density becomes

$$\rho_\nu(x, t) = \bar{\rho}_\nu(t) + \sum_{i=1}^{N} m w_i \, \delta^{(3)}(x - x_i). \tag{4.2.11}$$

Cosmological N-body simulations only compute the perturbed potential, since the background density $\bar{\rho}$ is accounted for in the background equations. The only change affecting the force calculation is therefore the weighting of the particles. The mean squared weight, $I = \frac{1}{2}\langle w^2 \rangle$, is a convenient statistic to quantify the importance of including the neutrino particles. We show the evolution of I for a $\sum m_\nu = 0.1\,\mathrm{eV}$ simulation with the homogeneous background model (4.2.9) in Fig. 4.1. At early times, particles deviate very little from their initial trajectory and the weights are negligible. We find that $I = 4 \times 10^{-7}$ at $z = 20$, $I = 3 \times 10^{-6}$ at $z = 10$, and $I = 2 \times 10^{-5}$ at $z = 5$. This early reduction is important as shot noise at high redshifts inhibits the growth of physical structure and can seed additional fluctuations that grow by gravitational instability. At late times, when nonlinear effects become important, the weights increase to $I = 2 \times 10^{-4}$ at $z = 2$,

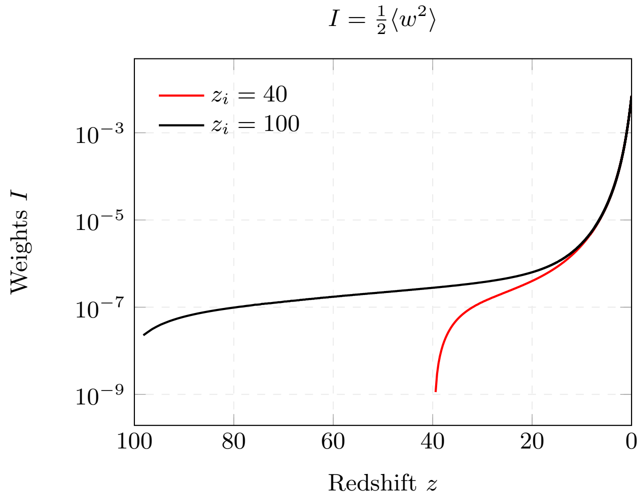

Figure 4.1: Evolution of particle weights for a $\sum m_\nu = 0.1\,\mathrm{eV}$ cosmology, starting at different redshifts z_i. The mean squared particle weight $\langle w^2 \rangle$ represents the effective reduction in shot noise.

$I = 1 \times 10^{-3}$ at $z = 1$, and $I = 6.7 \times 10^{-3}$ at $z = 0$, independently of the starting redshift of the simulation. This translates to a reduction in shot noise, $\sigma = 2VI/N$, or an effective increase in particle number at $z = 0$ by a factor $(2I)^{-1} = 75$. Finally, we note that one can save computational resources by integrating only a fraction of the neutrino particles as long as I remains small. We do not consider this possibility here.

4.2.2. Variance reduction

The δf method is an application of the much more general control variates method [181, 182]. This is a variance reduction technique commonly used in Monte Carlo simulations. See [91] for another recent application in cosmology. We briefly review the method here. Let A be a random variable with an unknown expectation $\mathrm{E}[A] = \mathcal{A}$. Given independent random samples A_i, the standard estimator is given by

$$\hat{\mathcal{A}} = \frac{1}{N}\sum_{i=1}^{N} A_i. \tag{4.2.12}$$

The error in $\hat{\mathcal{A}}$ is

$$\sigma^2_{\hat{\mathcal{A}}} = \mathrm{E}\big[(A - \hat{\mathcal{A}})^2\big] = \frac{\sigma^2_A}{N}. \tag{4.2.13}$$

Let B be another random variable for which the expected value $\mathrm{E}[B] = \mathcal{B}$ is known. By adding and subtracting, we can construct a control variate estimator for A:

$$\hat{\mathcal{A}}_{\text{cv}} = \frac{1}{N}\sum_{i=1}^{N}\left[A_i - \alpha B_i\right] + \alpha\mathcal{B}, \tag{4.2.14}$$

for any constant α. Like $\hat{\mathcal{A}}$, this is an unbiased and consistent estimator of $\mathrm{E}[A]$. However, the error in $\hat{\mathcal{A}}_{\text{cv}}$ is given by

$$\sigma^2_{\hat{\mathcal{A}}_{\text{cv}}} = \frac{1}{N}\left(\sigma_A^2 + \alpha^2\sigma_B^2 - 2\alpha\,\mathrm{cov}(A,B)\right). \tag{4.2.15}$$

Therefore, the error can be reduced if A and B are correlated. Differentiating, we see that the optimal value of α is given by

$$\alpha^* = \frac{\mathrm{cov}(A,B)}{\sigma_B^2}. \tag{4.2.16}$$

For the Fermi-Dirac model considered above, α^* is very close to unity and we simply set $\alpha = 1$ at all times. In general, the value of α^* could be estimated at runtime. This is useful if we add more information about the unknown variable and extend the method to a linear combination of multiple control variates (see Section 4.7). Furthermore, the method can still be useful when a control variate is not exactly known but can be estimated more efficiently than A.

4.2.3. Optimality

Let us consider how the δf method compares to other methods. To allow for the broadest possible comparison, we will write down an arbitrary hybrid method that involves some background model, $\bar{f}(x,p,t)$, such as a fluid description or linear response, and a discrete sampling of the distribution with arbitrary particle weights, $w_i(t)$:

$$f_{\text{hyb}}(x,p,t) = \alpha(t)\bar{f}(x,p,t) + \sum_i w_i(t)\delta^{(3)}(x-x_i)\delta^{(3)}(p-p_i), \tag{4.2.17}$$

where $\alpha(t)$ is a weight function for the background. This parametrization captures virtually all existing methods. The ordinary N-body particle method corresponds to $(\alpha, w_i) = (0,1)$ at all times. Pure grid-based methods have $(\alpha, w_i) = (1,0)$. Existing hybrid methods switch over from a grid method to a particle method after some time t_s, which corresponds to $(\alpha, w_i) = (1-q, q)$ with $q(t) = I[t \geq t_s]$ a step function. For simplicity, we consider only the case where all particles are switched on at the same time, but the argument extends readily to the more practical case where only some particles are switched on. Given a choice of weight function, $\alpha(t)$, for the background, what choice of particle weights is optimal?

Let $f(x,p,t)$ be the nonlinear distribution and $g(x,p,t)$ the sampling distribution of the markers. In the continuous limit, the expected error in the number density is given by

$$\begin{aligned}\langle\epsilon\rangle &= \int \mathrm{d}^3p\; f_{\mathrm{hyb}}(x,p,t) - \int \mathrm{d}^3p\; f(x,p,t) && (4.2.18)\\ &= \int \mathrm{d}^3p\Big(w(x,p,t)g(x,p,t) + \alpha(t)\bar{f}(x,p,t) - f(x,p,t)\Big). && (4.2.19)\end{aligned}$$

Meanwhile, the shot noise term in the power spectrum grows as the square of the particle weights, so we want to minimize

$$\tfrac{1}{2}\langle w^2\rangle = \int \mathrm{d}^3p\; \tfrac{1}{2}w(x,p,t)^2 g(x,p,t), \tag{4.2.20}$$

subject to the constraint $\langle\epsilon\rangle \leq \eta$ for some maximum error η. Assume that the bound is saturated. First, let us look for solutions that extremize the integral constraint. We find the unique solution

$$w = \frac{\delta f}{g} \quad \text{with} \quad \delta f = f - \alpha\bar{f}. \tag{4.2.21}$$

This is the δf method introduced above, with optimal α given by (4.2.16). Any further solution should extremize the Lagrangian,

$$\mathcal{L} = \tfrac{1}{2}w(x,p,t)^2 g(x,p,t) + \lambda\Big(w(x,p,t)g(x,p,t) + \alpha(t)\bar{f}(x,p,t) - f(x,p,t)\Big). \tag{4.2.22}$$

Writing down the Euler-Lagrange equations

$$[w+\lambda]\, g\nabla_p w + \tfrac{1}{2}w^2\nabla_p g + \lambda w\nabla_p g = \lambda\nabla_p f - \alpha\lambda\nabla_p\bar{f}, \tag{4.2.23}$$

one finds a family of quadratic solutions

$$w = -\lambda \pm \sqrt{\lambda^2 + 2\lambda\frac{\delta f}{g}} \quad \text{with} \quad \delta f = f - \alpha\bar{f}. \tag{4.2.24}$$

The case $\lambda = 0$ corresponds to the trivial solution $w = 0$. For $\lambda \neq 0$, we obtain the minima

$$w = \frac{\delta f}{g} - \frac{1}{2\lambda}\left(\frac{\delta f}{g}\right)^2 + \mathcal{O}\left(\frac{1}{\lambda^{2/3}}\frac{\delta f}{g}\right)^3. \tag{4.2.25}$$

These solutions correspond to small perturbations around the δf method that trade some accuracy for a possible reduction in shot noise. However, since the leading correction is $\propto (\delta f)^2$, this is only possible if the background model is skewed with respect to the nonlinear solution. Typically, the skewness and the additional reduction in shot noise is negligible. In fact, since the next-to-leading correction is positive, shot noise increases if the skewness is small.

We have shown that within the broad class of hybrid methods described by equation (4.2.17), δf-type methods of the form (4.2.25) minimize the amount of shot noise, subject to the constraint that the error in the number density remains below a certain bound. The δf method given by (4.2.21), recovered from (4.2.25) in the limit $\lambda \to \infty$, is the unique solution for which the expected error $\langle \epsilon \rangle = 0$. The optimal value of α is given by (4.2.16), but will be close to 1 if $\bar{f} \approx f$. This is the method we will use exclusively, with the choice $\alpha = 1$.

4.3. One-dimensional example

We now illustrate the method by applying it to a one-dimensional test problem with a known solution. Readers that are satisfied with the mathematical derivation may skip ahead to Section 4.4.

4.3.1. The elliptical sine wave

Consider the 1-dimensional collisionless Boltzmann equation

$$\frac{\partial f}{\partial t} + p\frac{\partial f}{\partial x} - \frac{\partial \Phi}{\partial x}\frac{\partial f}{\partial p} = 0, \tag{4.3.1}$$

where the particles move under a conservative force $F(x) = -\Phi'(x)$. Let us assume a periodic potential given by

$$\Phi(x) = \sin^2(x/2). \tag{4.3.2}$$

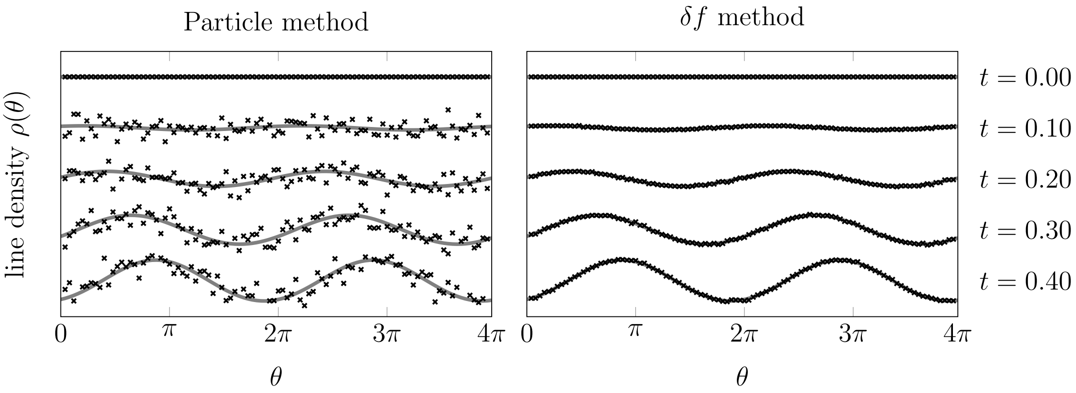

Figure 4.2: Density profiles for the 1-dimensional elliptical sine wave test problem. We counted particles in 100 bins of width $\Delta\theta = 4\pi/100$ to create the empirical density profiles. On the left, an ordinary N-body simulation with $N = 10^6$ particles was used. On the right, the N-body simulation was extended with a δf step.

The steady-state solution can be found to be:

$$f(x,p) = \frac{\bar{\rho}}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{p^2}{2\sigma^2} + \frac{\cos(x)}{2\sigma^2}\right), \tag{4.3.3}$$

in terms of the background density $\bar{\rho}$ and velocity dispersion σ. The corresponding density profile $\rho(x)$ is given by

$$\rho(x) = \int_{-\infty}^{\infty} f(x,p)\mathrm{d}p = \bar{\rho}\exp\left(\frac{\cos(x)}{2\sigma^2}\right). \tag{4.3.4}$$

To find the general time-dependent solution, we use the method of characteristics. The characteristic equations are

$$\frac{\mathrm{d}x}{\mathrm{d}t} = p, \qquad \frac{\mathrm{d}p}{\mathrm{d}t} = -\frac{1}{2}\sin(x). \tag{4.3.5}$$

These equations of motion can be solved in terms of the energy $E = \frac{1}{2}p^2 + \sin^2(x/2)$, which gives

$$\sin(x/2) = \mathrm{sn}\left(\pm\sqrt{E/2}(t-\tau)\right), \tag{4.3.6}$$

where τ is an integration constant and $\mathrm{sn}(x)$ is the Jacobi elliptic sine function with elliptic modulus $k = 1/\sqrt{E}$ [183][1]. Assuming a homogeneous Gaussian distribution with mean $\bar{p}$ for the initial momenta p at time $t = 0$,

$$f(x, p, 0) = \frac{\bar{\rho}}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(p-\bar{p})^2}{2\sigma^2}\right), \tag{4.3.7}$$

the general solution, $f(x, p, t)$, at later times is a complicated expression involving elliptic sines and arcsines. The details are given in Appendix 4.A. We replicate the problem using N-body methods. A large number of particles are initialized on the interval $x \in [0, 4\pi]$ with momenta drawn from the initial distribution (4.3.7). The particles are then integrated using

$$\Delta x = p\Delta t, \qquad \Delta p = -\frac{1}{2}\sin(x)\Delta t. \tag{4.3.8}$$

In addition to the ordinary N-body method, we use a δf method, where the background model is given by

$$\bar{f}(x, p, t) = \frac{\bar{\rho}}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(p-\bar{p})^2}{2\sigma^2}\right), \tag{4.3.9}$$

and the weights are updated during each step via $w = \delta f/f$. The corresponding density profiles are shown in Fig. 4.2. The plots were created using $N = 10^6$ particles and the model parameters are $\bar{\rho} = \sigma = 1$ and $\bar{p} = 10$. The results show that both the ordinary N-body simulation and the simulation with a δf step can reproduce the exact solution. However, the ordinary method is very noisy, whereas the δf method reproduces the expected profiles with remarkable accuracy. The reason for this is that while the distribution itself has a large dispersion, resulting in noisy results for the ordinary method, the perturbations from the steady solution are small, which allows the δf method to work. This is exactly analogous to the cosmic neutrino background.

4.4. Relativistic effects

Neutrinos are relativistic at early times, which introduces some subtleties when evolving such a species with a Newtonian code. Including relativistic effects is not necessary for the δf method, but we include them in our simulations to allow a consistent comparison with several recent works [154, 166, 172]. Furthermore, the higher-order δf methods

---

[1] For $E \to \infty$, we have $\operatorname{sn} x \to \sin x$, meaning that $x \propto t$. The particle 'ignores' the potential. For $E = k = 1$, $\operatorname{sn} x = \tanh x$, meaning the particle asymptotically approaches a potential peak. For $E < 1$, the particle is bounded and oscillates between peaks.

discussed in Section 4.7 provide a natural setting for including these effects without neglecting the nonlinear evolution of the neutrinos.

There are broadly speaking two philosophies when it comes to the inclusion of relativistic effects in Newtonian N-body codes, which one might call the "active" and "back-scaled" approaches. In the former, relativistic effects are actively included, while in the latter, the initial conditions are modified by rescaling the linear power spectrum in such a way that the desired outcome is recovered at $z = 0$. In this chapter, we take an active approach. In Chapter 5, we instead take the back-scaled approach. By construction, both methods agree at the linear level at $z = 0$, but predictions at high redshift may differ. Generally, however, these effects are quite small. In both cases, it is convenient to choose a gauge in which relativistic corrections are small to begin with.

We will work in the Newtonian motion framework of [184] and make modifications to the initial conditions, long-range force calculation, and particle equations of motion as outlined below.

4.4.1. Initial Conditions

To generate initial conditions for massive neutrinos and to set up the higher-order background models (Section 4.7), accurate calculation of the linear theory neutrino distribution function is indispensable. This can be done with the Boltzmann codes CAMB [127] and CLASS [128]. At their default settings, these codes produce accurate total matter and radiation power spectra (their intended purpose), but the neutrino related transfer functions (e.g. density and velocity) are not converged and can be very inaccurate [165]. To obtain converged results, we post-process perturbation vectors from CLASS by integrating source functions up to multipole $\ell_{\max} = 2000$. This prevents the artificial reflection that can happen for low $\ell_{\max}$. See Appendix 4.B for more details.

Initial conditions are then created using the post-processed transfer functions from CLASS in N-body gauge at $z = 100$. We do not follow the usual approach of back-scaling the present-day power spectrum, but use the so-called forward Newtonian motion approach [184, 185]. To our knowledge, forward Newtonian motion initial conditions have always been set up with the Zel'dovich approximation. However, this approximation is known to be inadequate for precision simulations [88]. To go beyond Zel'dovich initial conditions, we determine the Lagrangian displacement vectors $\boldsymbol{\psi} = \mathbf{x} - \mathbf{q}$ by solving the Monge-Ampère equation

$$\rho(\mathbf{x}) = \bar{\rho}(1+\delta(\mathbf{x})) = \det\left[1+\frac{\partial\psi_i(\mathbf{x})}{\partial q_j}\right]^{-1}. \tag{4.4.1}$$

This equation is solved numerically with a fixed-point iterative algorithm that exploits the fact that the density perturbation δ is small. We note that this approach is not equivalent to Lagrangian perturbation theory, but merely provides a more accurate map from the Eulerian initial density field to a Lagrangian displacement field compared to the Zel'dovich approximation. A detailed analysis of this method will be presented elsewhere. Velocities were determined independently using the transfer function for the energy flux $\theta = ik \cdot v$.

We used two different methods to generate initial conditions for the neutrino particles. In the first method, neutrino particles were displaced randomly in phase space according to the perturbed phase-space density function, $f_{\mathrm{PT}}(x,p,t)$, including terms up to $\ell = 5$. This method was used for our (256 Mpc)^3 simulations. However, we encountered some problems with this method and later switched to the method presented in Chapter 6. In this second method, which accounts also for higher multipoles, neutrino particles are integrated in linear theory from $z = 10^9$ to the starting redshift $z = 10^2$ [150, 154, 186]. Both methods agree on small scales, but the latter method was found to be more accurate on large scales, $k < 10^{-2}$ Mpc^{-1}. For this reason, we used the second method for our (1 Gpc)^3 simulations.

4.4.2. Long-range forces

In a relativistic setting, the gravitational evolution is governed by the Einstein-Boltzmann equations. We will approximate this system using a hybrid approach [187], in which dark matter and massive neutrinos are evolved using a Newtonian N-body code complemented with relativistic corrections to the fluid equations that are pre-solved in linear theory. We will work in N-body gauge, which allows the fluid equations for dark matter to be written in a particularly convenient form resembling the Newtonian equations solved by conventional N-body codes. The continuity and Euler equations can then be written as [185, 188]:

$$\dot{\delta} + \nabla\cdot\mathbf{v} = 0, \tag{4.4.2}$$

$$\dot{\mathbf{v}} + aH\mathbf{v} = -\nabla\phi + \nabla\gamma^{\mathrm{Nb}}, \tag{4.4.3}$$

where overdots denote conformal time derivatives, δ is the density contrast, $\mathbf{v}$ the peculiar velocity, and $H = \dot{a}/a^2$. All relativistic corrections are captured by the N-body gauge term, $\nabla\gamma^{\mathrm{Nb}}$, which arises from the anisotropic stress of relativistic species. In addition, the scalar potential ϕ receives contributions from all fluid components:

$$\nabla^2\phi = 4\pi G a^2 \sum_i \delta\rho_i, \tag{4.4.4}$$

where the sum runs over cold dark matter, baryons, neutrinos, and photons. Density perturbations are actively calculated for all species. In the case of massive neutrinos and the cold dark matter and baryon fluid, this is done with particles in the usual way. For photons and massless neutrinos (and for some runs, the massive neutrinos[2]), this is done by realizing the corresponding transfer functions from CLASS on a grid as part of the long-range force calculation in our N-body code SWIFT.

In the absence of relativistic species, the N-body gauge term, $\nabla\gamma^{\mathrm{Nb}}$, vanishes and the continuity and Euler equations agree with the Newtonian equations solved in conventional N-body codes. This makes N-body gauge useful as it allows one to set up initial conditions in N-body gauge, evolve them in a Newtonian simulation, and give the results a relativistic interpretation. The relativistic corrections become relevant at the 0.5% level on the largest scales in our Gpc simulations. Ordinarily, these corrections are accounted for in the initial conditions by back-scaling the present-day linear power spectrum, ensuring that the linear power spectrum is recovered on large scales at $z = 0$. Here, instead, we actively include these contributions with the aim of recovering the linear power spectrum at earlier times as well.

4.4.3. Particle content

When simulating light neutrinos from high redshifts, we are evolving relativistic particles in a Newtonian simulation. Such particles can reach superluminal speeds when evolved using the ordinary equations of motion. Following [85], we initially addressed this issue by replacing the equations of motion with special relativistic equations that are valid to

---

[2] Specifically, the linear theory runs and the runs with higher-order δf methods, as discussed in Sections 4.6 and 4.7, respectively.

all orders in u:

$$\dot{\mathbf{u}} = -\frac{2u^2 + a^2}{\sqrt{u^2 + a^2}} \nabla \left(\phi - \gamma^{\mathrm{Nb}}\right), \tag{4.4.5}$$

$$\dot{\mathbf{x}} = \frac{\mathbf{u}}{\sqrt{u^2 + a^2}}. \tag{4.4.6}$$

Here, a is the scale factor and $\mathbf{u}$ the comoving 3-velocity. However, we encountered two problems with our original approach. The first is that a direct leapfrog implementation of these equations is not symplectic due to the fact that the right-hand side of (4.4.5) depends on $\mathbf{u}$. As a consequence, phase-space density is not exactly conserved by the discrete Hamiltonian, as assumed by the δf method. We investigate this issue further in Appendix 4.D, where we offer an alternative solution and conclude that this is not an issue in practice. The second problem is that using these equations, we did not exactly reproduce linear theory on large scales. We ultimately traced this to the use of a gauge-dependent definition of momentum in the derivation of Eq. (4.4.5), as will be discussed in Chapter 6. For now, we simply note that a solution is to use the special relativistic equation (4.4.6) together with the non-relativistic version of (4.4.5):

$$\dot{\mathbf{u}} = -a \nabla \left(\phi - \gamma^{\mathrm{Nb}}\right). \tag{4.4.7}$$

This choice not only ensures that the neutrinos move the correct subluminal distance and that the integrator is symplectic, but also avoids the problem on large scales. We used this second approach for our Gigaparsec simulations. On small scales, the evolution of the neutrinos at late times is insensitive to their early evolution, making the difference between (4.4.5) and (4.4.7) immaterial. For this reason, we present the original results for our 256 Mpc simulations, which used the first approach with (4.4.5) and (4.4.6).

A separate matter from the equations of motion is that the neutrinos have a relativistic energy at early times. Using only the mass-energy, as is done for non-relativistic matter, leads to underestimation of the matter power spectrum on large scales. To rectify this, we replace the weighted mass of the particles with a weighted energy $\epsilon = m\sqrt{a^2 + u^2}$. Here, again the issue of the symplectic integrator plays a role. To ensure that $\dot{\mathbf{u}}$ does not depend on current particle velocities, we used the alternative form $\epsilon = m\sqrt{a^2 + u_0^2}$ for our Gigaparsec simulations. Here, u_0 is the initial particle velocity at $z = 100$. This approximation is extremely accurate, as substantial deviations from u_0 only occur for slow particles at late times in which case the mass term dominates (see Appendix 4.D). For the 256 Mpc simulations, we used the original form $\epsilon = m\sqrt{a^2 + u^2}$.

Table 4.1: Description of the simulations. The listed particle mass, m_p, refers to the combined cold dark matter and baryon particles. The neutrino fraction is listed as $f_\nu = \Omega_\nu/(\Omega_{\rm cb} + \Omega_\nu)$.

Side Length	N_c	$m_p\ [M_\odot]$	N_ν	$\sum m_\nu$	f_ν
1024 Mpc	1024^3	3.96×10^{10}	0	0.0 eV	0
1024 Mpc	1024^3	3.93×10^{10}	1024^3	0.1 eV	0.0073
1024 Mpc	1024^3	3.81×10^{10}	1024^3	0.5 eV	0.0376
256 Mpc	512^3	4.95×10^{9}	0	0.0 eV	0
256 Mpc	512^3	4.92×10^{9}	1024^3	0.1 eV	0.0073
256 Mpc	512^3	4.77×10^{9}	1024^3	0.5 eV	0.0376

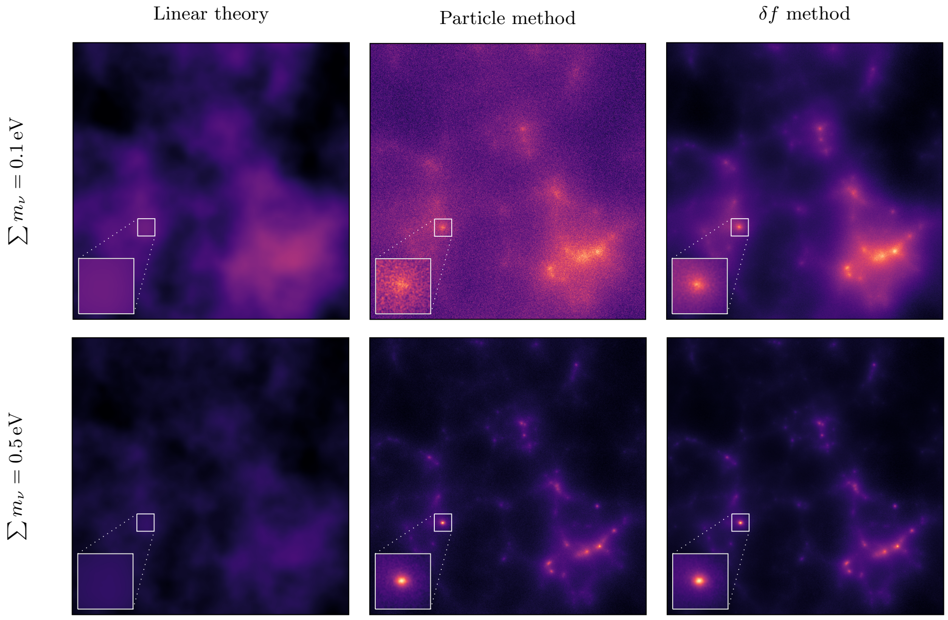

Figure 4.3: Neutrino density plots of (256 Mpc)^3 cubes at $z = 0$, simulated with two commonly-used methods and with the δf method. The particle and δf simulations used $N_\nu = 1024^3$ particles. Shot noise is clearly visible for the particle method, although noticeably less so for $\sum m_\nu = 0.5\,$eV. The linear theory model fails to reproduce the small-scale behaviour. The δf method solves both problems. The inset zooms in on a neutrino halo and should also be compared with the linear response prediction in Fig. 4.4, which uses the same colour scale. The neutrino fraction $f_\nu = \Omega_\nu/(\Omega_{\rm cb} + \Omega_\nu) = 0.0073$ for the top row and $f_\nu = 0.0376$ for the bottom row.

4.5. Simulations

We now describe our neutrino simulations, which were run on the COSMA6 computing facility in Durham. We have implemented the δf method in the cosmological hydrodynamics code SWIFT [189, 190]. SWIFT uses a combination of the Fast Multipole Method for short-range gravitational forces and the Particle Mesh method for long-range forces. It uses a 5^{th} degree polynomial kernel for the force softening with a single time-dependent softening length. The code uses a task-based parallelization paradigm to achieve strong scaling on large clusters and obtain significant speed-ups over competing N-body codes. The main simulations presented in this paper use the basic version of the δf method with a homogeneous Fermi-Dirac distribution as background model. Our choice of cosmological parameters, based on Planck 2018 [112], is $(h, \Omega_\mathrm{c} + \Omega_\nu, \Omega_\mathrm{b}, A_s, n_s) = (0.6737, 0.265, 0.0492, 2.097 \times 10^{-9}, 0.9652)$. We run two sets of simulations at different resolution to test the large-scale and small-scale behaviour of various methods. The cube sizes and particle numbers are listed in Table 8.1.

4.5.1. Choice of neutrino masses

Neutrino oscillations indicate that there are three neutrino mass eigenstates with unknown masses m_i. The mass splittings have been measured with good precision and are complemented by cosmological constraints on the sum of neutrino masses, as discussed in Section 3.1. Given the limits discussed in that section, we consider three values for $\sum m_\nu$, keeping the present-day value $\Omega_{\mathrm{m},0} = \Omega_{\mathrm{cb},0} + \Omega_{\nu,0}$ fixed. Scenario one contains three massless neutrinos, scenario two corresponds to the inverted mass ordering with $\sum m_\nu = 0.1\,\mathrm{eV}$[3], and scenario three to a degenerate spectrum with $\sum m_\nu = 0.5\,\mathrm{eV}$. The first two models bracket the most interesting range of values $0 < \sum m_\nu < 0.1\,\mathrm{eV}$. The last model has surely been ruled out, but is included for several reasons. First of all, the δf method reduces to the ordinary particle method in the large mass limit at late times. Hence, the $\sum m_\nu = 0.5\,\mathrm{eV}$ case provides a useful consistency check. Second, when simulations are used to emulate statistics for parameter extraction, we should allow for unlikely excursions in MCMC analyses without our simulation methods breaking down [172]. Finally, in the extended parameter space around ΛCDM, for example with a non-standard lensing amplitude, A_L, or curvature, or when varying the dark energy equation of state, the possibility of larger neutrino masses remains very relevant [113, 114, 116, 191].

The two massive scenarios considered in this paper have degenerate neutrino masses

---

[3] Specifically, two 0.0486 eV neutrinos and one massless neutrino.

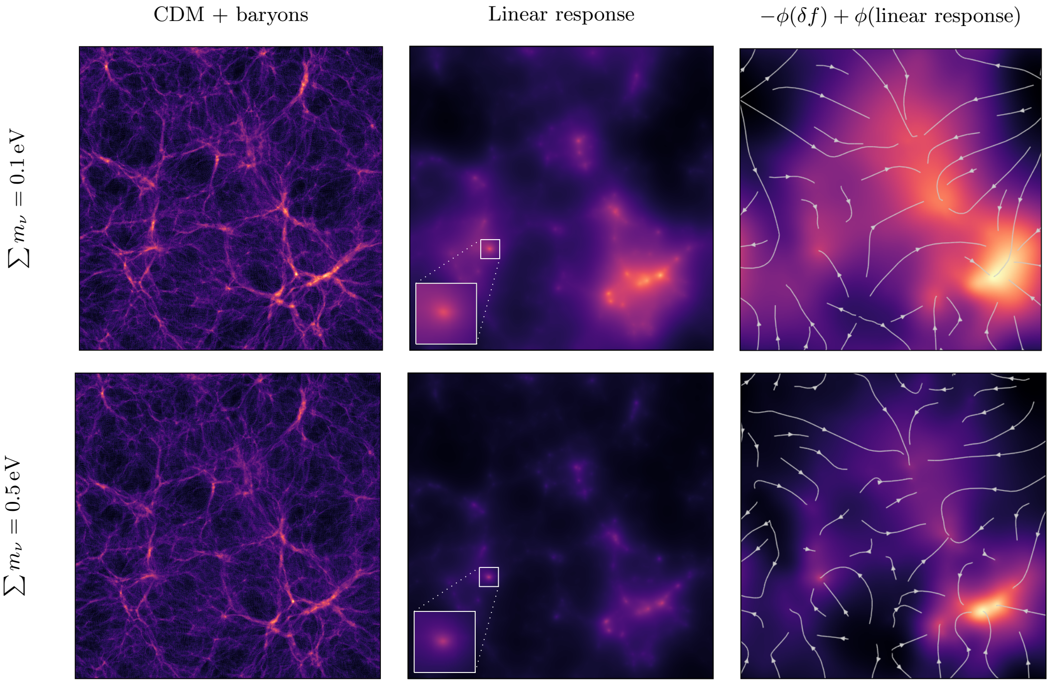

Figure 4.4: Density plots of (256 Mpc)^3 cubes at $z = 0$. The linear response method applies the linear theory ratio $\delta_\nu^{\mathrm{lin}}(k)/\delta_{\mathrm{cb}}^{\mathrm{lin}}(k)$ to the simulated CDM + baryon phases [169]. Compared to the linear theory prediction, it performs remarkably well, but the neutrino halos around clusters are significantly more diffuse compared to the particle and δf simulations (compare the zoomed in halo with the δf prediction in Fig. 4.3). The resulting potential difference is shown in the last column, with flowlines indicating the forces that are not present in the linear response model.

$(2 \times 0.05\,\mathrm{eV}$ and $3 \times 0.167\,\mathrm{eV})$. However, the δf method can easily be extended to account for mass splittings. In that case, particles would be labelled with a given mass state, i, and each state would have its own background model, $\bar{f}_i$. The reduction in shot noise is largest for the smallest neutrino masses, placing different masses on a level footing. This allows for better load balancing between different neutrino masses.

4.6. Results

We compare our neutrino δf method with three commonly-used alternatives. The most common alternative is the ordinary N-body particle method, which is the same in every respect as our method, but with the weighting step disabled. Next, we consider a linear theory method based on [166] that does not evolve neutrino particles but instead realizes

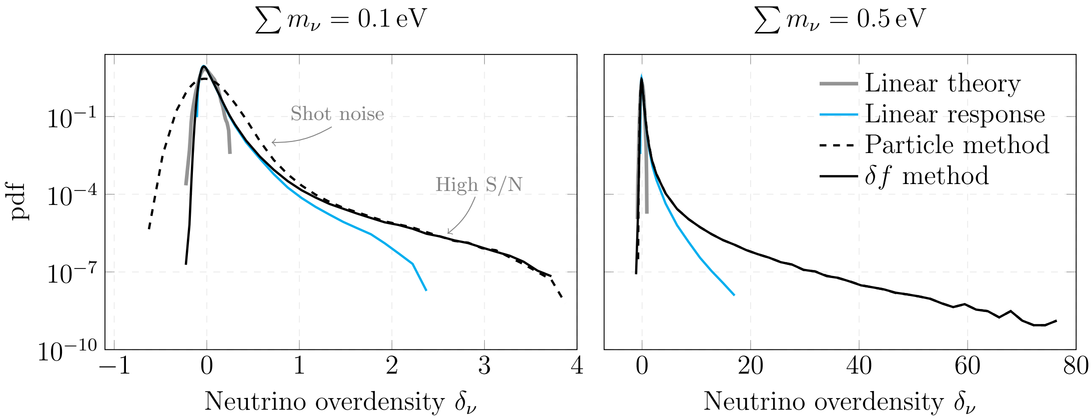

Figure 4.5: Neutrino probability density functions (pdf) at $z = 0$, computed on a 1024^3 grid from the 256 Mpc simulations, and smoothed with a Gaussian filter with radius $R = 256$ kpc. We compare the δf method with three commonly-used alternatives. The particle and δf methods agree in the high density tail, because the largest overdensities have enough particles to achieve a high signal-to-noise ratio. Shot noise plagues the particle method, particularly in underdense regions. The linear methods fail in the high density tail.

the linear theory neutrino perturbation in N-body gauge on a grid. The neutrinos are then fully accounted for in the long-range forces. Finally, we consider the linear response method of [169] in which the neutrino perturbation is calculated by applying the linear theory transfer function ratio, $\delta_\nu^{\mathrm{lin}}(k)/\delta_{\mathrm{cb}}^{\mathrm{lin}}(k)$, to the simulated CDM + baryon phases.

A visual inspection of the neutrino density plots shown in Figs. 4.3 and 4.4 reveals the strengths and weaknesses of the four methods. Broadly, we see that the linear theory method does not suffer from shot noise, but fails to reproduce the small-scale behaviour resolved by the particle and δf methods. At the same time, shot noise is clearly visible in the particle simulation with $\sum m_\nu = 0.1\,\mathrm{eV}$, despite using $N_\nu = 1024^3$ particles in a 256 Mpc cube. This is evidently cured in the δf plot. We also see that shot noise is much less of a problem for $\sum m_\nu = 0.5\,\mathrm{eV}$, but the δf plot is still less grainy than the corresponding particle plot. Finally, Fig. 4.4 shows that the linear response method greatly improves on the pure linear theory prediction, but still produces neutrino halos that are too diffuse compared to the particle and δf simulations.

4.6.1. Neutrino component

We start with an analysis of the probability density function of the neutrino density field, computed on a 1024^3 grid from the 256 Mpc simulations. Refer to the plots in Fig. 4.5,

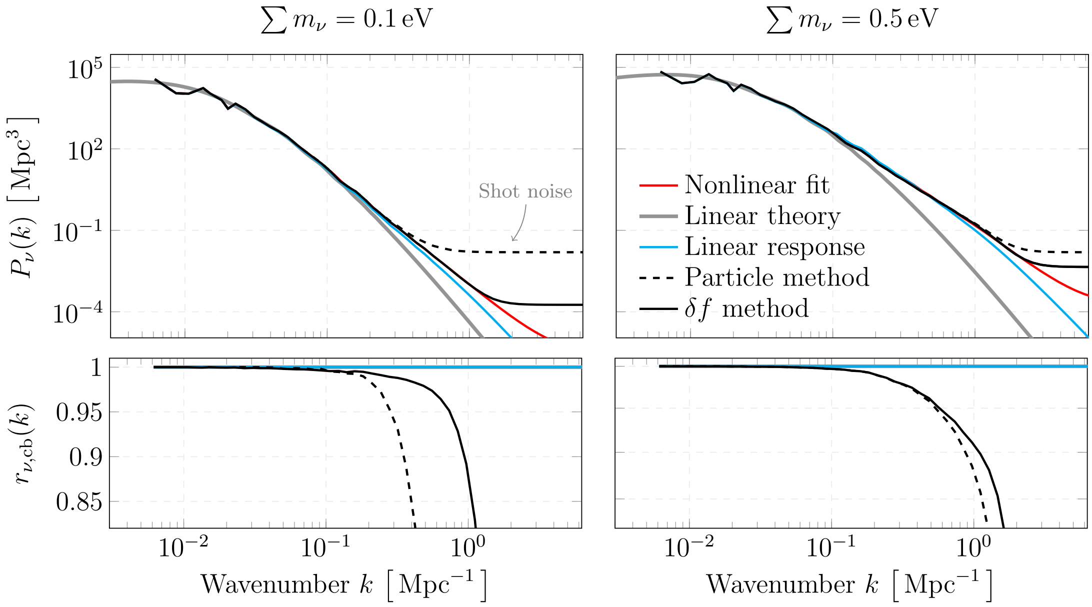

Figure 4.6: Neutrino power spectra at $z = 0$. We compare the δf method with three commonly-used alternatives. Shot noise enters the power spectrum at the constant level $V/N = 1/64$ for the particle method. We also show a fit to the δf power spectrum (red curves), given by eq. (4.6.1). The bottom panels show the cross-spectral coefficient $r_{\nu,\mathrm{cb}} = P_{\nu,\mathrm{cb}}/\sqrt{P_\nu P_\mathrm{cb}}$.

which bear out the basic picture sketched above. For the $\sum m_\nu = 0.1\,\mathrm{eV}$ neutrinos, the particle method is plagued by shot noise, but agrees with the δf method in the high density tail where the particle number is sufficient to obtain a good signal-to-noise ratio. The linear prediction fails in the high and low density tails. Finally, the linear response method, which applies the linear theory ratio $\delta_\nu(k)/\delta_\mathrm{cb}(k)$ to the CDM + baryon phases, is an intermediate case between the linear theory and δf methods. For the more massive scenario, the situation is much the same, except that shot noise is much less of a problem for the particle method on these scales.

Next, we consider two-point statistics and show the neutrino power spectrum at $z = 0$ in Fig. 4.6, combining the large and small simulations to show a wide range of scales. We use the Gpc simulations for $k < 0.1\,\mathrm{Mpc}^{-1}$ and the 256 Mpc simulations for $k \geq 0.1\,\mathrm{Mpc}^{-1}$. As expected, all methods agree on linear scales, $k \lesssim 0.1\,\mathrm{Mpc}^{-1}$, for both neutrino masses. On smaller scales, linear theory significantly underpredicts the amount of neutrino clustering. The linear response method also underpredicts the neutrino power spectrum, but not by as much. The relative difference between the nonlinear power spectrum and linear power spectrum is greater for neutrinos than for CDM and baryons. To account for this effect, we fit a nonlinear correction to the linear response power spectrum using the measured δf power spectrum up to $k = 1\,\mathrm{Mpc}^{-1}$:

$$P_\nu^{\mathrm{fit}}(k) = P_{\mathrm{cb}}(k)\left[\frac{\delta_\nu^{\mathrm{lin}}(k)}{\delta_{\mathrm{cb}}^{\mathrm{lin}}(k)}\right]^2 e^{\alpha+\beta k}, \tag{4.6.1}$$

and find $\alpha = 0.006 \pm 0.004$ and $\beta = 0.90 \pm 0.01$ ($\sum m_\nu = 0.1\,\mathrm{eV}$) and $\alpha = -0.06 \pm 0.03$ and $\beta = 0.34 \pm 0.09$ ($\sum m_\nu = 0.5\,\mathrm{eV}$). These are shown as the red curves in Fig. 4.6.

The particle simulations are clearly affected by shot noise, at the level of $V/N = 1/64$, obscuring the neutrino signal on scales smaller than $k = 0.2\,\mathrm{Mpc}^{-1}$ for the lightest scenario and on scales smaller than $k = 1\,\mathrm{Mpc}^{-1}$ for the more massive scenario. Using the δf method, shot noise is significantly reduced in the former case (factor of 87) and slightly reduced in the latter case (factor of 3.5), revealing a signal down to $k = 1 - 2\,\mathrm{Mpc}^{-1}$. Hence, δf simulations can achieve a similar resolution independently of mass without adjusting the particle number.

We also show the cross-spectral coefficient

$$r_{\nu,\mathrm{cb}}(k) = \frac{P_{\nu,\mathrm{cb}}(k)}{\sqrt{P_\nu(k)P_{\mathrm{cb}}(k)}}, \tag{4.6.2}$$

which captures phase differences between the dark matter and neutrinos. By definition, $r_{\nu,\mathrm{cb}} = 1$ according to the linear response method. However, this does not hold on small scales as can be seen in the bottom panels. Up to the point where shot noise becomes a problem, the particle and δf methods agree, demonstrating that $r_{\nu,\mathrm{cb}} < 1$ is a physical effect. This is particularly clear for $\sum m_\nu = 0.5\,\mathrm{eV}$.

Next, we consider how well the simulations can resolve the extended neutrino halos surrounding galaxies and clusters [192, 193]. In Fig. 4.7, we show stacked neutrino profiles for halos with virial CDM + baryon mass M_{cb} in the range $(5, 12) \times 10^{14} M_\odot$. The particle and δf methods agree almost perfectly, once again because of the high signal-to-noise ratio in the largest overdensities. In linear theory, the neutrino halos are completely absent as is evident also from the cross-sections in Fig. 4.3. Finally, the linear response method predicts neutrino halos that are too diffuse compared to the nonlinear simulations, and with too little dispersion from the mean profile. The larger dispersion found in the nonlinear simulations is not due to errors in individual profiles, but due to a stronger correlation between M_{cb} and the local neutrino density.

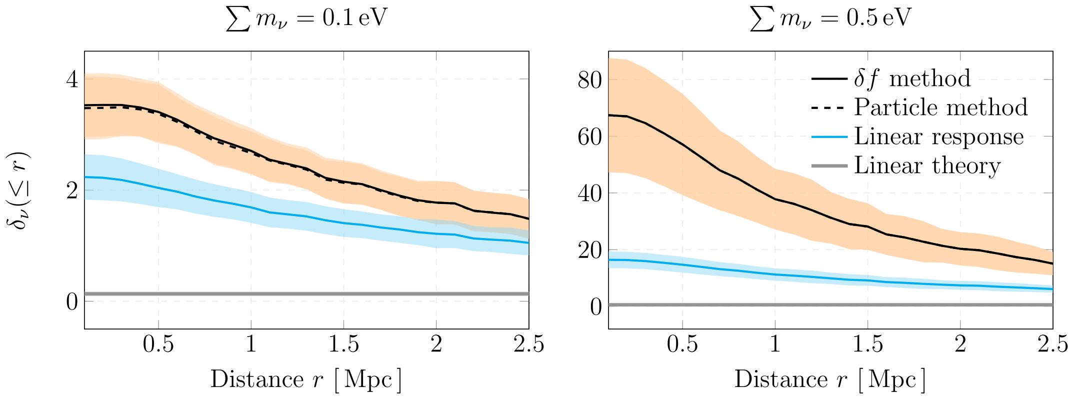

Figure 4.7: Stacked neutrino density profiles at $z = 0$ for halos with virial mass $M_{\rm cb}$ in the range $(5, 14) \times 10^{14}\,M_\odot$, computed with four different methods from the 256 Mpc simulations. The particle and δf curves overlap almost perfectly. The shaded area indicates the 1σ dispersion around the mean profile.

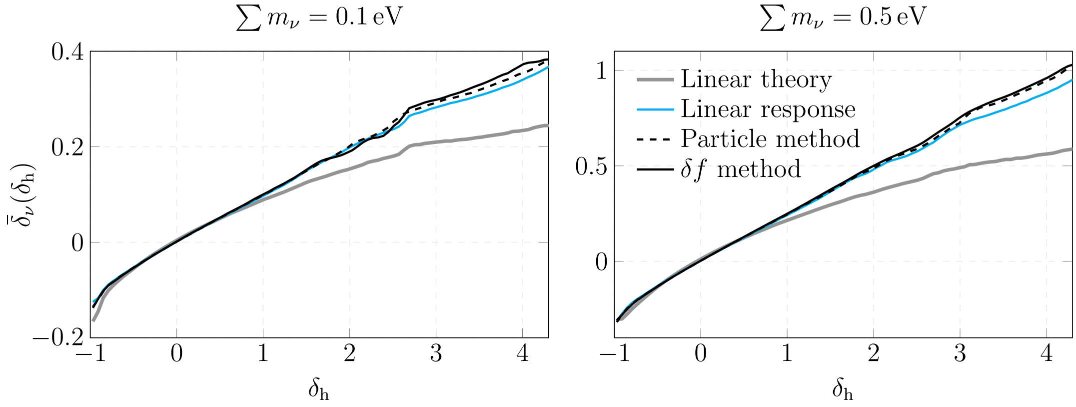

Figure 4.8: Neutrino bias relative to dark matter halos with virial mass, $M_{\rm cb} > 10^{12} M_\odot$, on scales, $R = 30h^{-1}$ Mpc, computed with four different methods from the Gpc simulations at $z = 0$. The $\bar{\delta}_\nu(\delta_{\rm h})$ relationship is approximately linear with slope equal to the neutrino bias b.

4.6.2. Neutrino bias

On larger scales, the neutrino density field can be reconstructed from the density of halos for a given neutrino mass spectrum [157]. We therefore construct the halo overdensity field,

$$\delta_{\mathrm{h}}(x) = \frac{n_{\mathrm{h}}(x) - \bar{n}_{\mathrm{h}}}{\bar{n}_{\mathrm{h}}}, \tag{4.6.3}$$

by calculating the number density, $n_{\mathrm{h}}(x)$, of halos and the mean density, $\bar{n}_{\mathrm{h}}$, at $z = 0$ in our Gpc simulations identified using the halo finder VELOCIRAPTOR [194]. We restrict attention to halos with virial mass, $M_{\mathrm{cb}} > 10^{12} M_{\odot}$, and smooth δ_{h} and δ_{ν} with a tophat filter of comoving radius $R = 30h^{-1}$ Mpc. Following [175], we study the mean neutrino density at constant halo density $\bar{\delta}_{\nu}(\delta_{\mathrm{h}})$, defined in terms of the joint probability density function $P(\delta_{\nu}, \delta_{\mathrm{h}})$ as

$$\bar{\delta}_{\nu}(\delta_{\mathrm{h}}) = \int \mathrm{d}\delta_{\nu} \delta_{\nu} P(\delta_{\nu}, \delta_{\mathrm{h}}). \tag{4.6.4}$$

This relationship is close to linear with slope equal to the neutrino bias, given by

$$b = \frac{\langle \delta_{\nu} \delta_{\mathrm{h}} \rangle}{\langle \delta_{\mathrm{h}}^2 \rangle}. \tag{4.6.5}$$

The degree of nonlinearity is captured by

$$\epsilon_{\mathrm{nl}}^2 = \frac{\langle \delta_{\mathrm{h}}^2 \rangle \langle \bar{\delta}_{\nu}^2 \rangle}{\langle \bar{\delta}_{\nu} \delta_{\mathrm{h}} \rangle^2} - 1, \tag{4.6.6}$$

which satisfies $\epsilon_{\mathrm{nl}} = 0$ if and only if the slope of $\bar{\delta}_{\nu}(\delta_{\mathrm{h}})$ is independent of δ_{h}. The scatter around the biasing relationship is characterized by the stochasticity,

$$\epsilon_{\mathrm{stoch}}^2 = \frac{\langle \delta_{\mathrm{h}}^2 \rangle \langle (\delta_{\nu} - \bar{\delta}_{\nu})^2 \rangle}{\langle \bar{\delta}_{\nu} \delta_{\mathrm{h}} \rangle^2}. \tag{4.6.7}$$

The nonlinearity and stochasticity are related to the correlation coefficient,

$$r_{\nu,\mathrm{h}} = \frac{\langle \delta_{\nu} \delta_{\mathrm{h}} \rangle}{\sqrt{\langle \delta_{\nu}^2 \rangle \langle \delta_{\mathrm{h}}^2 \rangle}}, \tag{4.6.8}$$

via $r_{\nu,\mathrm{h}} \simeq (1 + \epsilon_{\mathrm{nl}}^2 + \epsilon_{\mathrm{stoch}}^2)^{-1/2}$. This model is analogous to the nonlinear stochastic galaxy biasing model of [195, 196]. We compute the four quantities $(b, \epsilon_{\mathrm{nl}}^2, \epsilon_{\mathrm{stoch}}^2, r_{\nu,\mathrm{h}})$ for each of the methods under consideration. The results are listed in Table 4.2 and the biasing relationship is shown in Fig. 4.8. As expected on these large scales, we find good agreement with differences of a few percent in the bias. The greater the level of neutrino clustering resolved by a given method, the greater the bias b and correlation $r_{\nu,\mathrm{h}}$. The

Table 4.2: Neutrino bias relative to dark matter halos on scales $R = 30h^{-1}$ Mpc. Listed are the bias, b; nonlinearity, $\epsilon^2_{\rm nl}$; stochasticity, $\epsilon^2_{\rm stoch}$; and the correlation coefficient, $r_{\nu,\rm h}$.

	Method	b	$\epsilon^2_{\rm nl}$	$\epsilon^2_{\rm stoch}$	$r_{\nu,\rm h}$
0.1 eV	δf method	0.1032	0.0022	0.4883	0.8195
	Particle method	0.1028	0.0021	0.4955	0.8176
	Linear response	0.1015	0.0022	0.5065	0.8146
	Linear theory	0.0987	0.0206	0.5878	0.7889
0.5 eV	δf method	0.2556	0.0014	0.1969	0.9137
	Particle method	0.2546	0.0017	0.1927	0.9152
	Linear response	0.2502	0.0019	0.2031	0.9112
	Linear theory	0.2404	0.0257	0.2902	0.8719

stochasticity follows the opposite pattern. The nonlinearity follows no such pattern, but is very small in each case except (amusingly) for the linear theory runs. This is because linear theory does not resolve neutrino halos, causing the $\bar{\delta}_\nu(\delta_{\rm h})$ relation to level off in the high density tail.

The bias $b = 0.103$ for the 0.1 eV scenario is in excellent agreement with the bias $b = 0.071$ found by [175], when the difference in mass ordering is factored in using the approximately linear relationship between neutrino mass and bias in their results. [175] do not consider neutrino masses beyond 0.4 eV, but our finding of $b = 0.256$ for 0.5 eV is slightly lower than expected when extrapolating from their results. We also find a larger stochasticity and smaller correlation than might be expected, although the small nonlinearities agree. Given the mutual agreement between the different runs in Table 4.2, these differences are unlikely to be due to our choice of neutrino method. Differences in the N-body code or the identification of halos could also affect this comparison.

4.6.3. Matter power spectrum

The suppression of the total matter power spectrum at $z = 0$, relative to a massless neutrino cosmology, is shown in Fig. 4.9. We see that all methods are in excellent agreement and reproduce the famous spoon-like feature, which has recently been explained in terms of the halo model [197]. The differences between the methods are most pronounced around $k = 0.6\,\text{Mpc}^{-1}$, where the suppression is largest. The inset graphs zoom in on these scales. For both neutrino masses, the δf method predicts a smaller suppression than the particle and linear methods. This is in line with expectation, as the additional small-scale neutrino clustering, which is obscured by shot noise in the particle method and absent in linear theory, slightly offsets the suppression. Accordingly, the pure linear

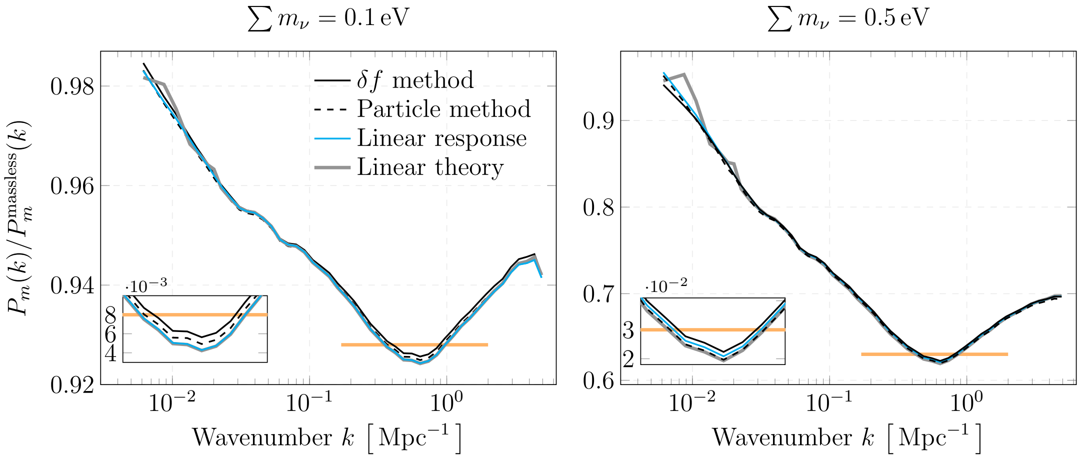

Figure 4.9: Total matter power spectra at $z = 0$, relative to a massless neutrino cosmology. The plots are based on (1024 Mpc)^3 simulations with $N_{\mathrm{cb}} = 1024^3$ and (for the particle and δf methods) $N_\nu = 1024^3$ simulation particles. The horizontal line is the empirical fitting formula, $\Delta P/P = -9.8 f_\nu$.

theory method predicts the least neutrino clustering and the largest suppression. It is interesting to see that the particle and δf methods do not agree for $\sum m_\nu = 0.5\,\mathrm{eV}$, despite having similar neutrino power spectra at $z = 0$. This is most likely due to shot noise at high redshift in the particle simulation. Compared to the cold dark matter and baryon fluctuations, the shot noise itself is negligible at $z = 0$, but it has two possible effects on structure formation at earlier times. It could seed non-physical density fluctuations or the random motions of neutrinos could obscure their real contribution to the growth of physical structure. Our results suggest that the latter effect dominates. In either case, these effects highlight the importance of using a hybrid method that eliminates shot noise at high redshift. The differences between the methods are at the permille level, corresponding to a shift in neutrino mass of several meV. In absolute terms, the differences are larger for $\sum m_\nu = 0.5\,\mathrm{eV}$, but less important overall.

The horizontal line corresponds to the empirical fitting formula, $\Delta P/P = -9.8 f_\nu$ [148]. Compared to this formula, we find a slightly greater suppression in each case, regardless of the method used to model the neutrinos. For the 0.1 eV simulations, this can be attributed to our use of the inverted mass ordering. The $\sum m_\nu = 0.5\,\mathrm{eV}$ case is perhaps more surprising, but seems to be in line with recent works. For example, [172] find increasingly larger differences with the fitting formula for increasing masses, although they do not consider models with $\sum m_\nu > 0.3\,\mathrm{eV}$.

Globally, the agreement between these very different methods is an encouraging sign and suggests that we have a good handle on the effects of massive neutrinos on the matter power spectrum. The differences, at most a few permille, may perhaps be relevant when trying to distinguish the effects of individual neutrino masses [105].

4.7. Higher-order δf methods

The performance of the δf method scales with the correlation between the nonlinear solution $f(x, p, t)$ and the background model $\bar{f}(x, p, t)$, so it is worth investigating other background models. We can go beyond the 0^{th} order Fermi-Dirac model by including the linear theory prediction. In that case, the distribution function can be written as

$$\bar{f}(x, p, t) = f_{\mathrm{FD}}(x, p, t)\left[1 + \Psi(x, p, t)\right], \tag{4.7.1}$$

where the perturbation is decomposed into multipole moments [126],

$$\Psi(k, p, t) = \sum_{\ell=0}^{\infty}(-1)^{\ell}(2\ell+1)\Psi_{\ell}(t, k, q)P_{\ell}(\hat{k}\cdot\hat{n}). \tag{4.7.2}$$

Here, $P_\ell(\cdot)$ are Legendre polynomials and the coefficients Ψ_ℓ satisfy an infinite hierarchy of moment equations. The Legendre representation yields simple expressions for the first few fluid moments, but is cumbersome for evaluating the distribution function itself. For our purposes, it is more convenient to use the following monomial representation

$$\Psi(\mathbf{k}, \hat{n}, q, \tau) = \sum_{\ell=0}^{\infty} i^{\ell}\Phi_{\ell}(\mathbf{k}, q, \tau)(\mathbf{k}\cdot\hat{n})^{\ell}, \tag{4.7.3}$$

where for a given $\ell_{\max}$, the odd (even) $\Phi_\ell(x, q, t)$ can be expressed in terms of all the odd (even) $\Psi_m(x, q, t)$ with $m \leq \ell$. See Appendix 4.C for details. With this choice of background model, the density integral becomes

$$\rho(x) = \bar{\rho}\left[1 + \delta_{\nu}(x)\right] + \sum_{i=1}^{N}\sqrt{m^2 + p_i^2}\, w_i\, \delta^{(3)}(x - x_i), \tag{4.7.4}$$

with particles weights $w_i = \delta f / f$ and $\bar{f}$ given by (4.7.1). Here, $\delta_\nu(x)$ is the linear neutrino overdensity, which is calculated using CLASS. The effect of the δ_ν perturbation should now be included in the long-range force calculation.

As shown in Fig. 4.10, adding the multipoles Φ_0 and Φ_1 significantly improves the

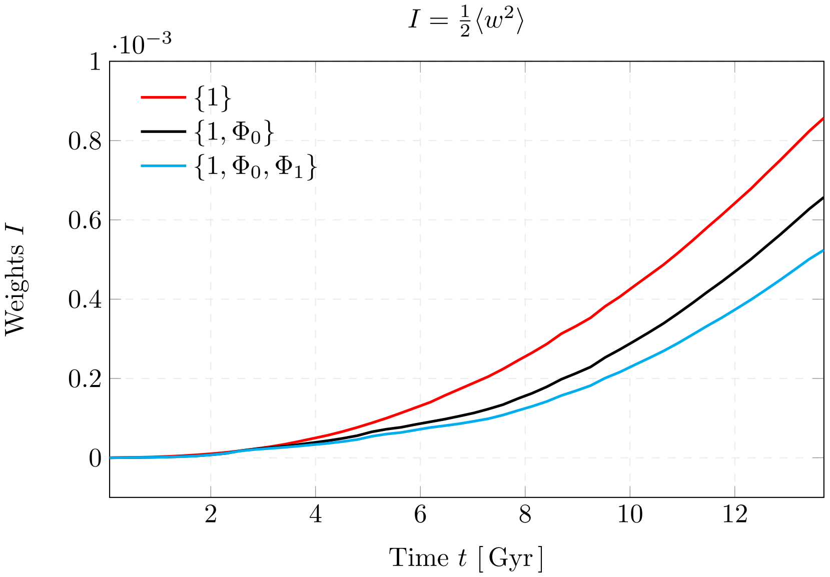

Figure 4.10: Evolution of the weights, or the effective reduction in shot noise, parametrized by the I-factor $I = \frac{1}{2}\langle w^2 \rangle$, when including higher-order perturbations: Φ_0 (density), and Φ_1 (energy flux).

correlation and therefore reduces the shot noise by almost 50%. It is likely that higher-order terms could contribute meaningfully too, as the multipole expansion converges only slowly. However, most of the gain is due to the 0^{th} order term, which on its own is much easier to implement.

4.8. Discussion and conclusions

Shot noise in N-body simulations is a major obstacle to modelling the nonlinear evolution of light relic neutrinos. In this paper, we demonstrate that the δf method, which decomposes the neutrino distribution into an analytically tractable background component, $\bar{f}$, and a nonlinear perturbation, δf, carried by the simulation particles, is an effective variance reduction technique. The reduction in shot noise scales with the dynamic particle weights, parametrized by $I = \frac{1}{2}\left\langle w^2 \right\rangle$. Because the weights are negligible until very late times, the simulation is mostly immunized against the effects of shot noise. Furthermore, shot noise is greatly reduced even at $z = 0$, which makes it possible to resolve neutrino clustering down to much smaller scales than is possible with conventional methods. Using higher-order versions of the δf method, which incorporate additional information from perturbation theory, shot noise can be reduced by another factor of $\mathcal{O}(2)$, and possibly more if moments $\ell > 2$ are included. Additional reduction in shot noise is possible by carefully tuning the sampling distribution of the marker particles.

The reduction in shot noise is more significant for smaller neutrino masses, because faster particles deviate less from their initial trajectory, resulting in smaller weights. This is fortunate as shot noise is most problematic for the fastest neutrinos. More generally, particles whose trajectories are not perturbed have negligible weights, whereas particles that are captured by halos have appreciable weights. This is again fortunate, because particles are needed in the vicinity of halos where grid methods tend to fail, while the unperturbed particles contain no information and contribute only noise. In between these extremes, particles will have intermediate weights. In this way, the δf method ensures an optimal combination of particles and background.

The method can in principle be combined with any grid or fluid background model to obtain an optimal hybrid method. Any simulation that evolves neutrino particles can be extended with a weighting step to minimize the shot noise as outlined in Section 4.2.3. It is not necessary, as was done here, to evolve the neutrino particles from the beginning. The I-statistic from a reference simulation can be used to gauge when the neutrinos become nonlinear and at what point they can safely be introduced (see Fig. 4.1).

We know from neutrino oscillations that at least one neutrino has a mass $m_\nu \gtrsim 0.05$ eV. Our results indicate that even for masses close to that bound, neutrinos are not particularly well modelled by linear approximations. For instance, the linear response neutrino power spectrum is off by 10% (60%) at $k = 0.1\,\mathrm{Mpc}^{-1}$ ($k = 1\,\mathrm{Mpc}^{-1}$) at $z = 0$, and the pure linear theory prediction is off by 14% (96%). Because the neutrinos make up only a small fraction of the total mass, the effect on the matter power spectrum is at most a few permille. This is the level at which the mass splittings are important [105]. Other statistics may be affected at a greater level, particularly if they are more sensitive to neutrino effects. For example, we have shown that the neutrino bias relative to dark matter halos is affected at the percent level on $30h^{-1}$ Mpc scales. In addition, some novel probes may require accurate modelling of the neutrino dynamics around halos, such as the neutrino-induced dynamical friction [198] and torque [199] on halos. By reducing shot noise without neglecting nonlinear terms, the δf method makes it feasible to calculate these effects even for the lightest neutrinos.

4.A. Elliptical sine wave solution

We define an integral of motion

$$E(x, p, t) = \tfrac{1}{2}p^2 + \sin^2(x/2), \tag{4.A.1}$$

which is interpreted as the energy of the particle. Hence,

$$p = \pm\sqrt{2E - 2\sin^2(x/2)}. \tag{4.A.2}$$

We have reduced the characteristic equations to

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \pm\sqrt{2E - 2\sin^2(x/2)}. \tag{4.A.3}$$

This equation is separable,

$$\int \frac{\mathrm{d}x}{\sqrt{2E - 2\sin^2(x/2)}} = \pm \int \mathrm{d}t. \tag{4.A.4}$$

Let τ be the time when $x(\tau) = 0$. Putting in the integration limits, setting $u = x/2$, and factoring out $2E$, we obtain

$$F(u) \equiv \int_0^u \frac{\mathrm{d}u'}{\sqrt{1 - \sin^2(u')/E}} = \pm\sqrt{E/2}\int_\tau^t \mathrm{d}t' \equiv \phi. \tag{4.A.5}$$

The elliptic sine function is defined such that $\mathrm{sn}(\phi) = \sin(u)$, where $u = F^{-1}(\phi)$. Hence,

$$\mathrm{sn}\left(\pm\sqrt{E/2}(t - \tau)\right) = \sin u. \tag{4.A.6}$$

There exist the following trigonometric identities [183]:

$$\mathrm{sn}(\phi)^2 + \mathrm{cn}(\phi)^2 = 1 \quad \text{and} \quad \frac{\mathrm{d}}{\mathrm{d}\phi}\mathrm{sn}(\phi) = \mathrm{cn}(\phi)\mathrm{dn}(\phi), \tag{4.A.7}$$

where $\mathrm{cn}(\phi) = \cos(u)$ and $\mathrm{dn}(\phi) = \sqrt{1 - \sin^2(u)/E}$ are the elliptic cosine and delta amplitude functions. Using these identities, one can confirm the solution (4.A.6). To find the phase-space distribution at time t, we use the fact that $f(x, p, t)$ is constant along its characteristic curves. At $t = 0$, let $x_0 = x(0)$ and $p_0 = p(0)$. Using the initial Gaussian distribution (4.3.7), we find

$$f(x,p,t) = f(x_0,p_0,0) \tag{4.A.8}$$
$$= \frac{\bar{\rho}}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(p_0-\bar{p})^2}{2\sigma^2}\right). \tag{4.A.9}$$

We need to express p_0 in terms of x, p, and t. First, we use conservation of energy to note that

$$p_0^2 = p^2 + 2\sin^2(x/2) - 2\sin^2(x_0/2). \tag{4.A.10}$$

What remains to show is how to express $\sin^2(x_0/2)$ in terms of x, p, and t. But this is simply,

$$\sin^2(x_0/2) = \mathrm{sn}^2\left(\mp\sqrt{\tfrac{1}{4}p^2 + \tfrac{1}{2}\sin^2(x/2)}\tau\right), \tag{4.A.11}$$

where the time, τ, is given by

$$\tau = \mp\sqrt{2/E}\,\mathrm{arcsn}\left(\sin(x/2)\right) + t. \tag{4.A.12}$$

Here, we used the inverse of the elliptic sine function, $\mathrm{arcsn}(x) = \phi$, with $x = \mathrm{sn}(\phi)$. It follows that

$$\sin^2(x_0/2) = \mathrm{sn}^2\left(\mathrm{arcsn}\left(\sin(x/2)\right) \mp \sqrt{\tfrac{1}{4}p^2 + \tfrac{1}{2}\sin^2(x/2)t}\right). \tag{4.A.13}$$

Therefore, the distribution function is

$$f(x,p,t) = \frac{\bar{\rho}}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2}\left[h(x,p,t) - 2\bar{p}\sqrt{h(x,p,t)} + \bar{p}^2\right]\right), \tag{4.A.14}$$
$$h(x,p,t) = k(x,p,1) - m(x,p,t) \tag{4.A.15}$$
$$k(x,p,t) = p^2 + 2\sin^2(x/2)t \tag{4.A.16}$$
$$m(x,p,t) = 2\mathrm{sn}^2\left(\mathrm{arcsn}\left(\sin(x/2)\right) \mp \sqrt{\tfrac{1}{4}k(x,p,t)}\right). \tag{4.A.17}$$

To find the density profile, $\rho(x,t)$, we integrate

$$\rho(x,t) = \int_{-\infty}^{\infty} f(x,p,t)\mathrm{d}p, \tag{4.A.18}$$

which can be done numerically. This gives the solution curves in Fig. 4.2.

4.B. Accurate calculation of neutrino moments

We reviewed the linear theory calculation of the neutrino distribution function in Newtonian gauge in Section 3.3. In this paper, we calculated the neutrino perturbations in synchronous gauge, with a calculation that proceeds along the same lines [126]. The precision of this calculation is set by two parameters: the maximum multipole $\ell_{\max}$ and the number of momentum bins N_q. By default, CLASS partly relies on a set of fluid equations and partly on integrating the hierarchy, using $\ell_{\max} = 50$ and $N_q = 28$ at the pre-set high precision settings [124]. The differences in the CMB anisotropies are at the permille level. However, the neutrino transfer functions have still not converged. To obtain converged results, [165] ran calculations with $N_q = 2000$ bins and $\ell_{\max} = 2000$, which each required hundreds of CPU hours. This is to be contrasted with a default CLASS run, which completes in seconds. To circumvent this computational cost, we use a different approach, which involves a post-processing step of CLASS tables.

To quickly integrate the Boltzmann hierarchy for high N_q and $\ell_{\max}$, we note that the source terms in the evolution equations depend on the matter content only through the scalar potential derivatives $\dot{h}$ and $\dot{\eta}$, which can be calculated accurately with much lower settings[4]. Therefore, we make the assumption that we can decouple the potential terms from most of the neutrino moments Ψ_ℓ. We first evolve all source functions in CLASS at a reasonable precision setting. This gives the metric perturbations $\dot{h}(k,\tau)$ and $\dot{\eta}(k,\tau)$, which we then take as given and use to integrate the multipole moments Ψ_ℓ at high precision where they are needed.

4.C. Monomial basis for the distribution function

Boltzmann codes can solve for the functions $\Psi_\ell(k,q,\tau)$. But evaluating the distribution function, $f(\mathbf{x},\mathbf{q},\tau)$, requires substituting these back into the definitions (4.7.1) and (4.7.2). This presents a challenge as the Ψ_ℓ are large discretely sampled arrays of amplitudes that need to be convolved with the random phases. It would be prohibitively expensive to do this repeatedly for each term in the Legendre expansion. We therefore adopt the following scheme. First, we use the following representation of the ℓth Legendre polynomial,

$$P_\ell(x) = 2^\ell \sum_{n=0}^{\ell} x^n \binom{\ell}{n} \binom{\frac{n+\ell-1}{2}}{\ell}, \tag{4.C.1}$$

---

[4] In the reference model with $N_q = 28$ and $\ell_{\max} = 50$, relative errors in $\dot{h}$ and $(\dot{h}/3 + 2\dot{\eta})$ are of order 10^{-4}. Although $\dot{\eta}$ still fluctuates at the several percent level, this term is much smaller than $\dot{h}$.

where the last factor is a generalized binomial coefficient. This allows us to expand Ψ and collect monomial terms in $\hat{k}\cdot\hat{n}$. We write

$$\Psi(\mathbf{k},\hat{n},q,\tau)=\sum_{\ell=0}^{\infty}i^{\ell}\Phi_{\ell}(\mathbf{k},q,\tau)(\mathbf{k}\cdot\hat{n})^{\ell}, \tag{4.C.2}$$

where the functions Φ_ℓ are defined by

$$\Phi_{\ell}(\mathbf{k},q,\tau)=\frac{1}{k^{\ell}}\sum_{n=0}^{\infty}(-2)^{n}\binom{n}{\ell}\binom{\frac{n+\ell-1}{2}}{n}(2n+1)\Psi_{n}(\mathbf{k},q,\tau). \tag{4.C.3}$$

Note that we factored out the magnitude of $\mathbf{k}=k\hat{k}$ and write the expansion in terms of $(\mathbf{k}\cdot\hat{n})^{\ell}$ and not $(\hat{k}\cdot\hat{n})^{\ell}$. This is to facilitate taking derivatives, as shown below. The Fourier transform of Ψ is

$$\Psi(\mathbf{x},\hat{n},q,\tau)=\int\frac{\mathrm{d}^{3}k}{(2\pi)^{3}}\Psi(\mathbf{k},\hat{n},q,\tau)e^{i\mathbf{x}\cdot\mathbf{k}}, \tag{4.C.4}$$

and similarly for the Ψ_ℓ and Φ_ℓ. We write the directional derivative along the unit vector $\hat{n}$ as $\mathcal{D}_{\hat{n}}=n^{i}\partial_{x_i}$. In other words,

$$\mathcal{F}\left\{\mathcal{D}_{\hat{n}}\Phi_{\ell}(\mathbf{x},\hat{n},q,\tau)\right\}\Longleftrightarrow i(\mathbf{k}\cdot\hat{n})\Phi_{\ell}(\mathbf{k},q,\tau). \tag{4.C.5}$$

Hence, we obtain

$$i^{\ell}\Phi_{\ell}(\mathbf{k},q,\tau)(\mathbf{k}\cdot\hat{n})^{\ell}\Longleftrightarrow\mathcal{F}\left\{\mathcal{D}_{\hat{n}}^{\ell}\Phi_{\ell}(\mathbf{x},q,\tau)\right\}. \tag{4.C.6}$$

And so, the overall perturbation, Ψ, is

$$\Psi(\mathbf{x},\hat{n},q,\tau)=\sum_{\ell=0}^{\infty}\mathcal{D}_{\hat{n}}^{\ell}\Phi_{\ell}(\mathbf{x},q,\tau). \tag{4.C.7}$$

A convenient numerical scheme is to store the Fourier transformed grids $\Phi_{\ell}(\mathbf{x},q,\tau)$, in which case we can evaluate the distribution function efficiently by taking finite differences.

4.D. Symplectic integrator

A point of concern is that the special relativistic equation of motion (4.4.5) may not be suited for the usual leapfrog integration in cosmological N-body simulations [200]. A straightforward substitution for the non-relativistic equation (4.4.7) produces a leapfrog integrator that is not symplectic and may therefore not explicitly conserve phase-space volume. The problem is that the equation for $\dot{\mathbf{u}}$ depends on $\mathbf{u}$. Here, we assess the impact of this error and provide a workaround. Fortunately, it is easy to construct a symplectic integrator that closely approximates the relativistic form. We simply replace equation (4.4.5) with:

$$\dot{\mathbf{u}} = -\frac{2u_0^2 + a^2}{\sqrt{u_0^2 + a^2}} \nabla \left(\phi - \gamma^{\mathrm{Nb}}\right), \tag{4.D.1}$$

where u_0 is the magnitude of $\mathbf{u}$ at the starting redshift of the simulation. Moreover, we use $\epsilon = m\sqrt{u_0^2 + a^2}$ when computing the energy density. As confirmed below, this is a good approximation due to the fact that $u_0 \ll a$ whenever u deviates much from u_0, i.e. for slow particles at late times. Equation (4.4.6) is unchanged. A leapfrog discretization of these equations is

$$\mathbf{x}_{k+1/2} = \mathbf{x}_k + \frac{1}{2}\Delta t \frac{\mathbf{u}_k}{\sqrt{u_k^2 + a^2}}, \tag{4.D.2}$$

$$\mathbf{u}_{k+1} = \mathbf{u}_k - \Delta t \frac{2u_0^2 + a^2}{\sqrt{u_0^2 + a^2}} \nabla \left(\phi(\mathbf{x}_{k+1/2}) - \gamma^{\mathrm{Nb}}(\mathbf{x}_{k+1/2})\right), \tag{4.D.3}$$

$$\mathbf{x}_{k+1} = \mathbf{x}_{k+1/2} + \frac{1}{2}\Delta t \frac{\mathbf{u}_{k+1}}{\sqrt{u_{k+1}^2 + a^2}}. \tag{4.D.4}$$

To determine whether this is symplectic, one considers the Jacobian, J, of the transformation $\psi\colon \mathbf{z}_k \mapsto \mathbf{z}_{k+1}$, where $\mathbf{z}_k = (\mathbf{x}_k, \mathbf{u}_k)$. One can confirm that

$$J^{\mathrm{T}}\Omega J = \Omega \equiv \begin{pmatrix} 0 & I_3 \\ -I_3 & 0 \end{pmatrix} \tag{4.D.5}$$

to show that ψ is a symplectomorphism. It follows that $\det(J) = 1$, which ensures that the leapfrog integrator is volume-preserving.

Since the validity of the δf algorithm depends on conservation of phase-space density along particle trajectories, we need to determine to what extent this is violated when using a nonsymplectic discretization of (4.4.5). Deviations will be of order $\mathcal{O}\left(u^2\right)$, which is small for any nontrivial neutrino orbit. Therefore, the difference should be negligible when the weights are large. To test this assertion, we evolve 256^3 neutrino and dark matter particles in a volume with sidelength 256 Mpc and assuming $\sum m_\nu = 3 \times 0.05\,\mathrm{eV}$, using both the special relativistic equations (4.4.5) and the alternative equation (4.D.1).

First, we confirm that (4.D.1) is a good approximation of (4.4.5), by checking that the ratios

$$r = \sqrt{\frac{u_0^2 + a^2}{u^2 + a^2}}, \qquad \text{and} \qquad s = \frac{1}{r}\frac{2u_0^2 + a^2}{2u^2 + a^2} \tag{4.D.6}$$

are close to unity. We find that $|r-1| < 0.02\%$ and $|s-1| < 0.03\%$ for 99% of particles at all times.

The evolution of the weights is shown in Fig. 4.11. At very high redshifts, when $I = \frac{1}{2}\langle w^2 \rangle < 10^{-6}$, the mean squared weight is about 60% larger for the nonsymplectic integrator. We interpret this as being due to small perturbations to neutrino trajectories that are absent in the symplectic case. The results converge after $z = 20$, when density perturbations approach nonlinearity, driving up the weights. The difference in I is 7% at $z = 10$, decreasing to 0.2% at $z = 2$, and 0.03% at $z = 0$. The difference in weights $|w_{\rm ns} - w_{\rm s}| < 0.025$ for 99% of particles at all times. The difference will be even smaller for larger neutrino masses.

As a result, we find that the use of the nonsymplectic integrator has a negligible effect on observables at late times. In particular, there is $< 0.1\%$ difference in the total matter and neutrino power spectra at $z = 0$. The difference in $P_\nu(k)$ grows to 0.3% at $z = 2$, and 2% at $z = 10$. For the total matter power spectrum, the difference is always below 0.1%. For future simulations, we recommend using a leapfrog discretization based on the modified expression (4.D.1) or using the non-relativistic version of (4.4.5) together with (4.4.6), as discussed in Section 4.4.3. See Chapter 6 for further details and recommendations.

4.E. Practical implementation

We briefly outline how to implement the δf method in a typical N-body code. First, a choice needs to be made for the background model. The simplest choice is the homogeneous Fermi-Dirac distribution, which we repeat here without pre-factor:

$$\bar{f}(x, p, t) = \bar{f}(p, a) = \left[\exp\left(\frac{ap}{k_b T_\nu}\right) + 1\right]^{-1}, \tag{4.E.1}$$

where $T_\nu = 1.95$ K is the present-day neutrino temperature, a the scale factor, and p the 3-momentum. The method can be implemented as follows:

1. Implement a function $\bar{f}(x, p, t)$ that returns the phase-space density at a particle's location according to the background model.

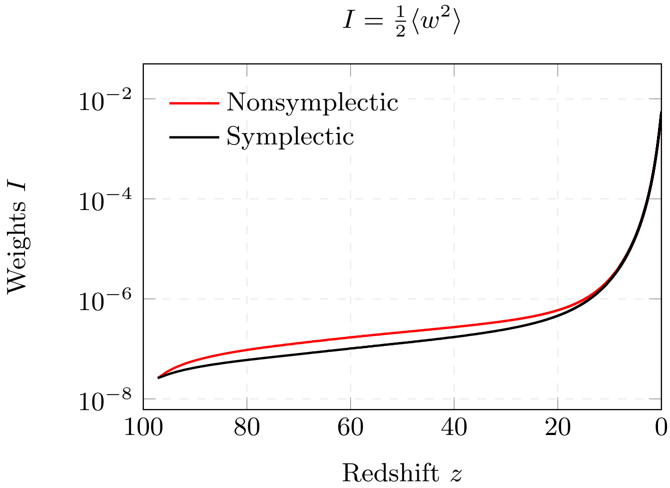

Figure 4.11: Evolution of particle weights for a $\sum m_\nu = 0.15\,\mathrm{eV}$ cosmology, using special relativistic equations of motion, evolved with symplectic and nonsymplectic leapfrog integrators.

2. Generate initial conditions with neutrino particles having a random initial momentum sampled from the background model. Perturbations can be applied afterwards.
3. For each particle, record the value of $\bar{f}(x, p, t) = f_0$ at the initial sampled position in phase space[5]. The numerical value of f_0 does not change if perturbations are applied to the initial conditions.
4. During subsequent time steps, for each particle:
    a) Compute $\bar{f}(x, p, t)$ using the new position and momentum.
    b) Compute the weight $w = (f_0 - \bar{f})/f_0$.
    c) For the purposes of the gravity force calculation, use the weighted mass mw.

The weights should be used when calculating statistics of the neutrino ensemble, such as the neutrino density, $\rho_\nu(x)$, that enters into the gravity force calculation. We emphasize that the weights should not be used in relationships such as $\mathbf{p} = m\mathbf{u}$.

For tree codes that perform a multipole expansion around the centre of mass, like SWIFT, there is a final point of consideration. Since neutrino particles can have negative weights, it is possible in rare circumstances that groups of particles have a nearly vanishing total mass. In that case, the centre of mass can be far removed from the particles and the

---

[5] Because the background model was used for sampling the initial momenta, we initially have $\bar{f} = f = g$. Conservation of phase-space density then ensures that $f = g = f_0$ for all particles at all times.

multipole expansion breaks down. There is a simple solution in such cases, which is to expand around any other point such as the geometric centre of the particles or the centre of absolute mass:

$$\mathbf{x}_{\mathrm{CoAM}} = \frac{\sum_i |m_i w_i| \mathbf{x}_i}{\sum_i |m_i w_i|}, \tag{4.E.2}$$

which has the advantage that it is very close to the ordinary centre of mass in most cases. When the background model agrees with the nonlinear solution, the weights are exactly zero. In that case, they can be set to a small value $w \ll 1$. This ensures that the centre of absolute mass is always well-defined.

Cold initial conditions

5

This chapter deals with the influence of massive neutrinos on the initial conditions of matter species that can be treated as cold on cosmological distance scales: baryons and cold dark matter. We show that the effects of neutrinos in Lagrangian perturbation theory are well described by scale-independent coefficients. We also find a new analytic solution for the first-order growth factor in ΛCDM with a non-clustering component.

Higher order initial conditions with massive neutrinos

Abstract: The discovery that neutrinos have mass has important consequences for cosmology. The main effect of massive neutrinos is to suppress the growth of cosmic structure on small scales. Such growth can be accurately modelled using cosmological N-body simulations, but doing so requires accurate initial conditions (ICs). There is a trade-off, especially with first-order ICs, between truncation errors for late starts and discreteness and relativistic errors for early starts. Errors can be minimized by starting simulations at late times using higher-order ICs. In this paper, we show that neutrino effects can be absorbed into scale-independent coefficients in higher-order Lagrangian perturbation theory (LPT). This clears the way for the use of higher-order ICs for massive neutrino simulations. We demonstrate that going to higher order substantially improves the accuracy of simulations. To match the sensitivity of surveys like DESI and Euclid, errors in the matter power spectrum should be well below 1%. However, we find that first-order Zel'dovich ICs lead to much larger errors, even when starting as early as $z = 127$, exceeding 1% at $z = 0$ for $k > 0.5$ Mpc^{-1} for the power spectrum and $k > 0.1$ Mpc^{-1} for the equilateral bispectrum in our simulations. Ratios of power spectra with different neutrino masses are more robust than absolute statistics, but still depend on the choice of ICs. For all statistics considered, we obtain 1% agreement between 2LPT and 3LPT at $z = 0$.

5.1. Introduction

The neutrino content of the Universe, $\Omega_\nu \cong \sum m_\nu/(93\text{ eV}\,h^2)$, becomes a powerful probe for cosmology once the implied neutrino masses are confronted with data from neutrino oscillations [29] and the kinematics of β-decay [14]. A non-zero detection of Ω_ν would be consequential for fundamental physics. It would confirm that a background of relic neutrinos survived until the epoch of structure formation, provide insight into the origin of neutrino mass, and constrain the search for dark matter and dark sectors. Oscillation experiments provide a lower bound of $\sum m_\nu > 0.058$ eV, while cosmology provides upper bounds of $\sum m_\nu < 0.15$ eV or better assuming ΛCDM [108, 109, 113, 201], with ongoing and future surveys promising significant further improvement. Planck and future cosmic microwave background experiments, together with large-scale structure surveys like DESI, Euclid, and Vera Rubin, could achieve sensitivities in the 0.01 - 0.02 eV range [132, 202–204]. Such small shifts in neutrino mass correspond to tiny 0.5% - 1.5% effects on the power spectrum of matter fluctuations on 0.1 Mpc^{-1} to 1 Mpc^{-1} scales, requiring theoretical predictions that are at least as accurate.

With this goal in mind, many groups have studied the effects of massive neutrinos on large-scale structure. At early times and on large enough scales, perturbation theory is the method of choice for this purpose. Cosmological perturbation theory [56] is essential for providing analytical insight and a necessary complement to more expensive numerical simulations. The effects of neutrinos on the nonlinear matter power spectrum were first calculated at one-loop by [205] and [206]. Subsequent work has dealt more realistically with the neutrino phase-space distribution [125, 168, 207–210], which parallels similar efforts on the numerical simulations side. Other advances were made by including neutrinos in the effective field theory of large-scale structure [145, 146] and using time renormalisation group perturbation theory [114, 211], which improved agreement with N-body simulations. More closely related to this work, [212] extended the hybrid COLA simulation method to cases with massive neutrinos using second-order Lagrangian perturbation theory (2LPT) and [213] incorporated nonlinear neutrino effects in Lagrangian perturbation theory up to third order (3LPT). On the numerical simulations side, where higher-order LPT has been used to great effect to produce accurate initial conditions (ICs) for conventional simulations without massive neutrinos [88, 214, 215], neutrino effects have not been included and higher-order LPT is therefore rarely used for neutrino simulations (but see [148, 216]). In this work, we propose a novel scheme for generating nLPT ICs for neutrino simulations based on all-order recursive solutions in the small-scale limit. We also generate ICs based on a full calculation of scale-dependent neutrino effects in 2LPT, dealing with frame-lagging terms following [213], and find near perfect agreement with our scheme in the final simulation product. This demonstrates that neutrino effects can be implemented beyond first order by working in the small-scale limit, paving the way for accurate neutrino simulation ICs.

N-body simulations are used to solve for the nonlinear gravitational dynamics of matter on small scales, where perturbation theory fails. Cosmological simulations with ICs based on LPT were pioneered by [50, 147] and [217]. Mixed dark matter simulations with sub-electronvolt neutrinos were first carried out by [148, 149, 161]. We refer the reader to [129] for a review of neutrino simulation methods. As with perturbation theory, the accuracy of modern surveys places stringent demands on simulations, popularly expressed as a requirement for 1% accurate calculations of the matter power spectrum [218]. A major source of uncertainty concerns the interface between perturbation theory and simulation, in the form of ICs, and associated transients [214]. We may distinguish two major sources of uncertainty related to the choice of ICs [89, 217]. The first arise from discrepancies between the ICs and the actual nonlinear solution at the initial time. When the solution is calculated perturbatively at order n, this uncertainty can be understood as the truncation error introduced by neglecting terms of order $n+1$ and greater. The second source of uncertainty relates to discreteness effects that build up over time as the continuous fluid equations are solved by means of a discrete particle representation [219, 220]. There is a tension between these two, as early starts minimize truncation errors but entail larger discreteness errors, while late starts do the opposite. For example, the first-order solution of [62] has the largest possible truncation error, driving practitioners to start simulations early when higher-order corrections are small. However, such simulations manifest a greater dependence on particle resolution due to discreteness errors. While such errors can be corrected [220], this reasoning provides strong motivation for using higher-order ICs at late times [89].

Neutrinos complicate this picture in two ways. First, neutrinos introduce an additional length scale into the problem. Due to their large thermal velocities, neutrinos free stream out of potential wells, otherwise stated in terms of a suppression of clustering on scales smaller than a typical free-streaming length [120]. This in turn causes a scale- and time-dependent suppression of dark matter and baryon clustering that must be accounted for in the initial conditions. [221] showed how to incorporate such scale-dependence in a first order back-scaling procedure, but a consistent framework for higher-order ICs has thus far been lacking[1]. The second complication is that late-time observables are more strongly correlated with the initial conditions and less determined by the internal structure of halos, when clustering is suppressed on small scales. This means that simulations with

---

[1] We note that after we submitted our paper to the journal, [222] presented a recipe for second-order neutrino ICs. Like us, they use a back-scaled transfer function for the cold dark matter and baryon species.

different neutrino masses are affected by errors to different degrees, contaminating ratios such as the suppression of the matter power spectrum. We will show that such ratios are more robust than absolute statistics, but still depend on the choice of initial conditions on small scales.

The paper is organized as follows. We begin by summarizing our recipe for generating higher-order ICs for neutrino simulations in Section 5.2. The second part of the paper is concerned with a derivation of the higher-order solutions necessary for ICs, starting with the set-up in Section 5.3, limiting solutions at all orders in Section 5.4.1, and the full second-order solution in Section 5.4.2. The final third of the paper contains a systematic analysis of higher-order ICs in Section 5.5. Finally, we conclude in Section 5.6. Throughout this paper, we use a default neutrino mass sum of $\sum m_\nu = 0.3$ eV to showcase our results, except where indicated otherwise.

5.2. N-body Initial conditions

We begin by outlining our approach for setting up for 3-fluid ICs with cold dark matter (c), baryons (b), and neutrinos (ν). Initially, we deal with a single cold fluid, described in terms of the the mass-weighted density contrast and velocity,

$$\delta_{\mathrm{cb}} = f_{\mathrm{c}}\delta_{\mathrm{c}} + f_{\mathrm{b}}\delta_{\mathrm{b}}, \tag{5.2.1}$$

$$\boldsymbol{v}_{\mathrm{cb}} = f_{\mathrm{c}}\boldsymbol{v}_{\mathrm{c}} + f_{\mathrm{b}}\boldsymbol{v}_{\mathrm{b}}, \tag{5.2.2}$$

where $f_{\mathrm{c}} = \Omega_{\mathrm{c}}/(\Omega_{\mathrm{c}} + \Omega_{\mathrm{b}})$ and $f_{\mathrm{b}} = 1 - f_{\mathrm{c}}$. In a final step, the cold fluid is separated into two components with distinct transfer functions. Our approach is based on a growing mode solution of the LPT equations in the small-scale limit, motivated by the hierarchy between the neutrino free-streaming scale and the nonlinear scale, $k_{\mathrm{fs}} \ll k_{\mathrm{nl}}$, at the redshifts relevant for ICs. In Section 5.5, we confirm that this is an excellent approximation suited for precision simulations. The recipe boils down to the following steps:

1. Compute a back-scaled transfer function $\delta_{\mathrm{cb}}(k)$
2. Compute particle displacements via Eqs. (5.2.3–5.2.11)
3. Compute particle velocities via Eqs. (5.2.12–5.2.14)
4. Perturb particle masses and velocities via Eqs. (5.2.15–5.2.19)

These steps can be performed using a modified version of the MONOFONIC code [89], which we have made publicly available[2]. We briefly discuss the steps in order and then deal with possible extensions in Section 5.2.5 and 5.2.6.

5.2.1. Transfer functions and back-scaling

In this paper, we follow the commonly used back-scaling approach. This approach begins by choosing a pivot redshift, typically $z = 0$, where the simulation should reproduce linear theory on the largest scales. This is necessary because conventional N-body codes solve Newtonian equations and therefore fail to capture the large-scale general relativistic dynamics in which matter and radiation are coupled through the Einstein-Boltzmann equations. We note that there exist alternative solutions to this problem [166, 172, 187, 188, 223] as well as fully relativistic N-body codes [87, 154, 224], which can avoid it altogether. In the back-scaling procedure, one uses a linear Einstein-Boltzmann code such as CLASS [128] or CAMB [127] to calculate the density transfer functions for each fluid species at z^{pivot}, which are then scaled back to the starting redshift of the simulation using the exact linear dynamics of the Newtonian code. For ΛCDM without neutrinos, this amounts to rescaling the dark matter transfer function by a constant growth factor ratio $D(z_i)/D(z^{\text{pivot}})$.

Adding massive neutrinos makes the linear solution scale-dependent, precluding a simple rescaling factor. Nevertheless, the same philosophy can be applied by solving the Newtonian dynamics of an N-body code with massive neutrinos at linear order. Following [221], we do this using a first-order Newtonian fluid approximation [125, 208], but see also [222] for a relativistic formulation. This back-scaling method for neutrino cosmologies was first implemented in the REPS code. To streamline the procedure for the end-user and to reduce the potential for human error, we built a lightweight back-scaling library ZWINDSTROOM that interfaces directly with CLASS and the initial conditions generator MONOFONIC. The final result of these steps is a rescaled density transfer function $\delta_{\text{cb}}(k) = D_{\text{cb}}(k, z_i)/D_{\text{cb}}(k, z^{\text{pivot}}) \cdot \delta_{\text{cb}}(k, z^{\text{pivot}})$ for a cold dark matter-baryon fluid (cb), where the growth factor ratio is computed with ZWINDSTROOM and the transfer function with CLASS.

---

[2] Up-to-date links to the software referenced in this paper are maintained at `https://www.willemelbers.com/neutrino_ic_codes/`.

5.2.2. Displacements

The displacement field, $\boldsymbol{\psi} = \boldsymbol{x} - \boldsymbol{q}$, relates the particle position $\boldsymbol{x}$ to the corresponding Lagrangian coordinate $\boldsymbol{q}$. To determine $\boldsymbol{\psi}$, we first obtain the linear potential by solving

$$\nabla^2 \varphi^{(1)}(\boldsymbol{q}) = \delta_{\mathrm{cb}}(\boldsymbol{q}). \tag{5.2.3}$$

Unless indicated otherwise, $\nabla = \nabla_{\boldsymbol{q}}$. We observe that $\varphi^{(1)}$ is not the gravitational potential, which also includes a neutrino contribution, but a notation that reflects the fact that we are solving for the displacements of cb fluid particles. Our fast approximate 3LPT [225–227] scheme for the displacement field in the presence of massive neutrinos has the simple form

$$\boldsymbol{\psi} = \boldsymbol{\psi}^{(1)} + C_2\boldsymbol{\psi}^{(2)} + C_3\boldsymbol{\psi}^{(3a)} + C_2C_3^1\boldsymbol{\psi}^{(3b)} + C_2\boldsymbol{\psi}^{(3c)}, \tag{5.2.4}$$

where C_n are scale-independent factors that capture the absence of neutrino perturbations in the small-scale limit, $C_n^i = C_n/C_i$, and $\boldsymbol{\psi}^{(n)}$ have the same form in terms of $\varphi^{(1)}$ as in ΛCDM. In the notation of [89], these are given by

$$\boldsymbol{\psi}^{(1)} = -\nabla\varphi^{(1)}, \quad \boldsymbol{\psi}^{(2)} = -\frac{3}{7}\nabla\varphi^{(2)}, \tag{5.2.5}$$

$$\boldsymbol{\psi}^{(3a)} = \frac{1}{3}\nabla\varphi^{(3a)}, \boldsymbol{\psi}^{(3b)} = -\frac{10}{21}\nabla\varphi^{(3b)}, \boldsymbol{\psi}^{(3c)} = \frac{1}{7}\nabla\times\boldsymbol{A}^{(3)}, \tag{5.2.6}$$

with higher-order potentials given by

$$\nabla^2\varphi^{(2)} = \frac{1}{2}\left[\varphi^{(1)}_{,ii}\varphi^{(1)}_{,jj} - \varphi^{(1)}_{,ij}\varphi^{(1)}_{,ij}\right], \tag{5.2.7}$$

$$\nabla^2\varphi^{(3a)} = \det \varphi^{(1)}_{,ij}, \tag{5.2.8}$$

$$\nabla^2\varphi^{(3b)} = \frac{1}{2}\left[\varphi^{(2)}_{,ii}\varphi^{(1)}_{,jj} - \varphi^{(2)}_{,ij}\varphi^{(1)}_{,ij}\right], \tag{5.2.9}$$

$$\nabla^2\boldsymbol{A}^{(3)} = \nabla\varphi^{(2)}_{,i} \times \nabla\varphi^{(1)}_{,i}, \tag{5.2.10}$$

where commas denote partial derivatives and we sum over repeated indices. In Section 5.4.1, we show that C_n can be expressed in terms of the neutrino fraction, $f_\nu = \Omega_\nu/\Omega_{\mathrm{m}}$. The correction, as it turns out, is small and approximately linear in f_ν:

$$C_n \cong 1 + \frac{2nf_\nu}{5(2n+3)}. \tag{5.2.11}$$

For a minimal neutrino mass sum of $\sum m_\nu = 0.06$ eV, one finds $C_2 - 1 = 5 \times 10^{-4}$. For our fiducial mass sum of $\sum m_\nu = 0.3$ eV, it is 0.3%. At $\sum m_\nu = 1$ eV, the effect is about one percent. The third-order correction C_3 is larger, but since $\boldsymbol{\psi}^{(3)}$ is suppressed by another power of the growth factor, the overall impact is smaller.

5.2.3. Velocities

The velocity field is $\boldsymbol{v}_{\text{cb}} = \mathrm{d}\boldsymbol{\psi}/\mathrm{d}t$. Given a satisfactory scheme for computing the displacement field, the time derivative can be evaluated numerically. This is our preferred method, since it requires no additional approximations. However, a faster method that avoids calculating higher order terms more than once is to use the asymptotic logarithmic growth rate

$$f_\infty = \lim_{k\to\infty} \frac{\mathrm{d}\log D_{\text{cb}}(k,a)}{\mathrm{d}\log a}, \tag{5.2.12}$$

to convert displacements to velocities, setting

$$\boldsymbol{v}_{\text{cb}} = aHf_\infty \left[\boldsymbol{\psi}^{(1)} + 2C_2\boldsymbol{\psi}^{(2)} + 3\left(C_3\boldsymbol{\psi}^{(3a)} + C_2C_3^1\boldsymbol{\psi}^{(3b)} + C_2\boldsymbol{\psi}^{(3c)}\right)\right]. \tag{5.2.13}$$

By construction, this gives the correct particle velocities on small scales. To recover also the correct behaviour on horizon scales, we add a large-scale correction $\boldsymbol{v}_{\text{cb}}^{(c)}$ given by

$$\boldsymbol{v}_{\text{cb}}^{(c)} = aHf_\infty \nabla^{-2}\nabla(\theta_{\text{cb}} - \delta_{\text{cb}}), \tag{5.2.14}$$

where θ_{cb} is the dimensionless energy flux transfer function computed with CLASS. We verified that the resulting simulated power spectrum agrees with linear theory to better than 0.1% at the pivot redshift of $z = 0$ on large scales. However, this approximation neglects possible nonlinear effects of scale-dependent growth on particle velocities. Another alternative is to rescale the velocities by the scale-dependent growth rate [221], which faces a similar problem beyond linear order.

5.2.4. Additional steps for 3-fluid ICs

The steps above are sufficient for simulations with neutrinos and a single cold fluid. To separate this cold fluid into baryon and CDM components with distinct transfer functions, we follow the approach of [228]. In short, the component densities are related to the mass-weighted average via[3]

$$\delta_\mathrm{c} = \delta_\mathrm{cb} - f_\mathrm{b}\delta_\mathrm{bc}, \tag{5.2.15}$$

$$\delta_\mathrm{b} = \delta_\mathrm{cb} + f_\mathrm{c}\delta_\mathrm{bc}, \tag{5.2.16}$$

where the difference variable, $\delta_\mathrm{bc} = \delta_\mathrm{b} - \delta_\mathrm{c}$, is constant at first order. The velocity difference too is conserved and vanishes at all orders: $\boldsymbol{v}_\mathrm{bc} = \boldsymbol{v}_\mathrm{b} - \boldsymbol{v}_\mathrm{c} = 0$. These results, derived for ΛCDM without massive neutrinos [229], carry over to the neutrino case, essentially due to the fact that the neutrino contribution cancels in the difference equations (Appendix 5.A). The transfer function difference, $\delta_\mathrm{bc}(k) = \delta_\mathrm{b}(k) - \delta_\mathrm{c}(k)$, is computed with CLASS at the pivot redshift and, since it is conserved, is not scaled back.

After assigning displacements and velocities to both particle species using the mass-weighted average fields, the density difference is implemented by setting the masses to

$$m_\lambda(\boldsymbol{q}) = \bar{m}_\lambda\big[1 + \delta_\lambda(\boldsymbol{q}) - \delta_\mathrm{cb}(\boldsymbol{q})\big], \tag{5.2.17}$$

with $\bar{m}_\lambda$ the mean particle mass for type $\lambda \in \{\mathrm{c},\mathrm{b}\}$. Perturbing the masses, rather than the displacements, was found by [228] to limit discreteness errors.

By construction, Newtonian simulations with initial conditions set up using the above procedure, reproduce the expected evolution of two cold fluids with a shared velocity field and a relative density contrast that is approximately conserved. However, like the large-scale velocity correction (5.2.14), a further modification is needed to bring the dynamics back into agreement with CLASS at first order:

$$m_\lambda(\boldsymbol{q}) \to m_\lambda(\boldsymbol{q}) + 2\bar{m}_\lambda\left[\left(\frac{D_\infty(z^\mathrm{pivot})}{D_\infty(z_i)}\right)^{1/2} - 1\right]\Theta_\lambda(\boldsymbol{q}), \tag{5.2.18}$$

$$\boldsymbol{v}_\lambda(\boldsymbol{q}) \to \boldsymbol{v}_\lambda(\boldsymbol{q}) + aHf_\infty\left(\frac{D_\infty(z^\mathrm{pivot})}{D_\infty(z_i)}\right)^{1/2}\nabla^{-2}\nabla\Theta_\lambda(\boldsymbol{q}), \tag{5.2.19}$$

where $D_\infty(z_i)$ is the small-scale growth factor at the starting redshift z_i and $\Theta_\mathrm{c} = -f_\mathrm{b}\theta_\mathrm{bc}$ and $\Theta_\mathrm{b} = f_\mathrm{c}\theta_\mathrm{bc}$. The difference, $\theta_\mathrm{bc}(k) = \theta_\mathrm{b}(k) - \theta_\mathrm{c}(k)$, of the dimensionless energy flux transfer functions is computed with CLASS at the pivot redshift.

---

[3] We remind the reader that $f_\lambda = \Omega_\lambda/\Omega_\mathrm{cb}$ for $\lambda \in \{\mathrm{c},\mathrm{b}\}$ even as $f_\nu = \Omega_\nu/\Omega_\mathrm{m} = \Omega_\nu/(\Omega_\mathrm{cb}+\Omega_\nu) = 1 - f_\mathrm{cb}$. Furthermore, $\delta_\mathrm{bc} \neq \delta_\mathrm{cb}$ and $\boldsymbol{v}_\mathrm{bc} \neq \boldsymbol{v}_\mathrm{cb}$.

5.2.5. Neutrino particles

Massive neutrinos can be included in N-body codes using a variety of methods. The most common approach is to solve for the neutrino perturbations self-consistently by including them as a separate N-body particle species [148, 149]. Initial conditions are then also needed for these neutrino particles. Capturing the full neutrino phase-space distribution is non-trivial even in linear theory and it is therefore not sufficient to compute only the first two moments, as is done for baryons and CDM. Accurate neutrino particle initial conditions can be generated by integrating geodesics from high redshift [150, 154], where the perturbed phase-space distribution can be expressed analytically [126], but care must be taken that the equations of motion remain valid in the ultra-relativistic régime [186]. This procedure can be carried out efficiently using our FASTDF code, introduced in Chapter 6. We stress that the focus of this paper is on dark matter and baryon ICs and the results apply regardless of whether the neutrino implementation uses particles.

5.2.6. Scale-dependent effects

Finally, we verified the approximations above by performing a full calculation of scale-dependent effects on the second-order displacement field. This is done by replacing (5.2.7) with a convolution of two copies of the first-order potential $\varphi^{(1)}(\boldsymbol{k})$, modulated by kernels $D_A^{(2)}(\boldsymbol{k}_1, \boldsymbol{k}_2)$ and $D_B^{(2)}(\boldsymbol{k}_1, \boldsymbol{k}_2)$, computed in Section 5.4.2. This numerical calculation is expensive, but we will show in Section 5.5 that simulations with ICs based on the full calculation agree extremely well with those based on the approximate scheme described above. The reason for this is the hierarchy of scales, $k_{\text{fs}} \ll k_{\text{nl}}$, which implies that higher-order corrections are important only on scales where neutrinos do not cluster, at least at redshifts that are relevant for ICs. Since the overall impact of the third-order correction factor, C_3, is smaller than that of C_2 and given the excellent agreement between the full and approximate solutions at second order, we expect the difference to be even smaller at third order. At the same time, the triple convolutions required for the third-order solution would be prohibitively expensive and would require a different approach. For this reason, we only consider 2LPT in Section 5.4.2.

5.3. Theoretical set-up

We now proceed with the set-up of a 3-fluid model, which is solved in Section 5.4. We consider three fluids indexed by $\lambda \in \{\mathrm{c}, \mathrm{b}, \nu\}$ for cold dark matter, baryons, and neutrinos. Throughout, we will treat baryons like dark matter particles and denote the mass-weighted CDM-baryon fluid by subscript cb. Let $\rho_\lambda(\boldsymbol{x})$ be the density, $\boldsymbol{u}_\lambda(\boldsymbol{x})$ the peculiar velocity flow, and $\boldsymbol{\sigma}_\lambda(\boldsymbol{x})$ the stress tensor. We also write $\delta_\lambda = \rho_\lambda/\bar{\rho}_\lambda - 1$ for the density contrast.

5.3.1. Euler equations

Taking moments of the Boltzmann equation yields the Euler fluid equations [56]

$$\partial_\tau \boldsymbol{u}_\lambda + \boldsymbol{u}_\lambda \cdot \nabla_{\boldsymbol{x}} \boldsymbol{u}_\lambda = -aH\boldsymbol{u}_\lambda - \nabla_{\boldsymbol{x}}\Phi - \frac{1}{\rho_\lambda}\nabla_{\boldsymbol{x}}(\rho_\lambda \boldsymbol{\sigma}_\lambda), \tag{5.3.1}$$

$$\partial_\tau \delta_\lambda + \nabla_{\boldsymbol{x}} \cdot [(1+\delta_\lambda)\boldsymbol{u}_\lambda] = 0, \qquad \text{for} \quad \lambda \in \{\mathrm{c}, \mathrm{b}, \nu\}, \tag{5.3.2}$$

where τ is conformal time, $H = \partial_\tau a/a^2$ is the Hubble constant (given explicitly below) and Φ the Newtonian potential. While the neutrino distribution function and its higher-order moments are complicated, the stress tensor can be neglected for the cold dark matter and baryon fluids on the scales of interest, $\boldsymbol{\sigma}_\mathrm{c} = \boldsymbol{\sigma}_\mathrm{b} = 0$, because we are restricting to scales much larger than the Jeans length and times before shell crossing. Taking the mass-weighted average of the cold dark matter and baryon equations, we obtain at all orders (see Appendix 5.A)

$$\partial_\tau \boldsymbol{u}_\mathrm{cb} + \boldsymbol{u}_\mathrm{cb} \cdot \nabla_{\boldsymbol{x}} \boldsymbol{u}_\mathrm{cb} = -aH\boldsymbol{u}_\mathrm{cb} - \nabla_{\boldsymbol{x}}\Phi, \tag{5.3.3}$$

$$\partial_\tau \delta_\mathrm{cb} + \nabla_{\boldsymbol{x}} \cdot [(1+\delta_\mathrm{cb})\boldsymbol{u}_\mathrm{cb}] = 0. \tag{5.3.4}$$

The potential is given by Poisson's equation,

$$\nabla^2_{\boldsymbol{x}}\Phi(\boldsymbol{x}) = \frac{3}{2}\frac{\Omega_\mathrm{m} H_0^2}{a}\delta_\mathrm{m}(\boldsymbol{x}), \tag{5.3.5}$$

in terms of the total matter density, $\delta_\mathrm{m} = f_\mathrm{cb}\delta_\mathrm{cb} + f_\nu \delta_\nu$, which includes a massive neutrino contribution. To complete the system, we assume the linear response approximation for the neutrino density:

$$\delta_\nu(\boldsymbol{k}) = \frac{\delta_\nu^\mathrm{lin}(k)}{\delta_\mathrm{cb}^\mathrm{lin}(k)}\delta_\mathrm{cb}(\boldsymbol{k}), \tag{5.3.6}$$

where $\delta_\lambda^\mathrm{lin}(k)$ refers to the density transfer function of $\lambda \in \{\nu, \mathrm{cb}\}$ computed in relativistic linear perturbation theory with CLASS. The total matter density contrast is then

$$\delta_\mathrm{m}(\boldsymbol{k}) = [1 + \alpha(k)]\, f_\mathrm{cb}\delta_\mathrm{cb}(\boldsymbol{k}), \tag{5.3.7}$$

where we have introduced the convenient notation $\alpha = f_\nu \delta_\nu^{\rm lin}/(f_{\rm cb}\delta_{\rm cb}^{\rm lin})$ for the linear theory ratio. The linear response approximation is accurate while neutrinos and dark matter remain in phase, which is a reasonable assumption at the early times considered here (see below). Inserting this in (5.3.5) yields

$$-k^2\Phi(\boldsymbol{k}) = \frac{B_0}{a}\left[1+\alpha(k)\right]\delta_{\rm cb}(\boldsymbol{k}), \tag{5.3.8}$$

where $B_0 = \frac{3}{2}(1-f_\nu)\Omega_{\rm m}H_0^2$ is written in terms of present-day values. We look for a growing solution of the form $\delta_{\rm cb}(\boldsymbol{k},\tau) = D_{\rm cb}(k,\tau)\delta_{\rm cb}(\boldsymbol{k},\tau_0)$. Linearising (5.3.3-5.3.5), we find

$$\partial_\tau^2 D_{\rm cb} + aH\partial_\tau D_{\rm cb} = \frac{B_0}{a}(1+\alpha)D_{\rm cb}. \tag{5.3.9}$$

In contrast to the ΛCDM case, this equation is scale-dependent due to the appearance of $\alpha(k)$. To proceed, we will take the limit $k\to\infty$. Since $\lim_{k\to\infty}\alpha(k) = 0$, we simply obtain

$$\partial_\tau^2 D_\infty + aH\partial_\tau D_\infty = \frac{B_0}{a}D_\infty \qquad (k\to\infty). \tag{5.3.10}$$

We denote the solution of (5.3.10) by D_∞ to indicate that this is the small-scale solution. At this point, an equally valid description could be given in the large-scale limit or indeed for an arbitrary pivot scale. We deliberately choose the small-scale limit for two reasons. First, most simulations are not large enough to realize the large-scale limit. Second, we are interested in nonlinear corrections to the initial conditions which are negligible on large scales.

5.3.2. Asymptotic form

We can find an analytic[4] solution to (5.3.10) if the contribution of radiation to the Hubble rate is neglected. We will return to this point further below. For now, let us assume that

$$H^2 = H_0^2\left[\Omega_\Lambda + \frac{\Omega_{\rm cb}+\Omega_\nu}{a^3}\right]. \tag{5.3.11}$$

---

[4] A function f is analytic at x if the Taylor series of f around x converges to f in a neighbourhood of x.

In this case, the growing mode can be expressed in terms of the hypergeometric function as (see Appendix 5.B)

$$D_\infty(a) = a^p\sqrt{1+\Lambda a^3}\,{}_2F_1\left(\frac{2p+7}{6}, \frac{2p+3}{6}, \frac{4p+7}{6}, -\Lambda a^3\right), \tag{5.3.12}$$

with $\Lambda = \Omega_\Lambda/\Omega_\mathrm{m}$ and $p = \sqrt{1+24(1-f_\nu)}/4 - 1/4$. This is normalized such that $\lim_{a\to 0} D_\infty/a^p = 1$. Taking $f_\nu = 0$, we recover the ΛCDM solution with $p = 1$ [59]. Taking instead $\Lambda \to 0$, we recover the solution during matter domination (MD)

$$D_\infty(a) = a^p = a^{\sqrt{1+24(1-f_\nu)}/4-1/4}, \tag{5.3.13}$$

which agrees with [23].

For ΛCDM without massive neutrinos, accurate nonlinear predictions can be made by substituting the growth factor for the scale factor, $a \to D$, in solutions obtained for the Einstein-de Sitter model. This is facilitated by using the growth factor as time variable (e.g. [59, 65, 229]). Here, we will pursue a similar strategy and make a change of time variables to D_∞. Defining the quantity

$$g_\infty = \frac{2}{3}\frac{B_0}{a}\left(\frac{D_\infty}{\partial_\tau D_\infty}\right)^2 \tag{5.3.14}$$

and the new velocity variable $\boldsymbol{v}_\mathrm{cb} = \partial_{D_\infty}\boldsymbol{x}$, the fluid equations can be rewritten as

$$\partial_{D_\infty}\boldsymbol{v}_\mathrm{cb} + \boldsymbol{v}_\mathrm{cb}\cdot\nabla_{\boldsymbol{x}}\boldsymbol{v}_\mathrm{cb} = -\frac{3g_\infty}{2D_\infty}(\boldsymbol{v}_\mathrm{cb} + \nabla_{\boldsymbol{x}}\varphi), \tag{5.3.15}$$

$$\partial_{D_\infty}\delta_\mathrm{cb} + \nabla_{\boldsymbol{x}}\cdot[(1+\delta_\mathrm{cb})\boldsymbol{v}_\mathrm{cb}] = 0, \tag{5.3.16}$$

$$\nabla^2_{\boldsymbol{x}}\varphi = \frac{\delta_\mathrm{cb}}{D_\infty} * (1+\alpha), \tag{5.3.17}$$

where the rescaled potential $\varphi = a\Phi/(B_0 D_\infty)$ is given in terms of a convolution, denoted by $*$, of δ_cb and the linear response $(1+\alpha)$. Although written in terms of D_∞, this is completely general.

Given suitable boundary conditions, Eqs. (5.3.15-5.3.17) are analytic at $D_\infty = 0$. In particular, we require that $\delta^\mathrm{ini}_\mathrm{m} = \delta^\mathrm{ini}_\mathrm{cb} = 0$. This agrees with our use of growing mode solutions for particle displacements, $\boldsymbol{q} \mapsto \boldsymbol{q} + \boldsymbol{\psi}$, where the unperturbed particle grid represents a uniform density field. The scaling, $H^2 \propto a^{-3}$, of the Hubble rate at early times ensures that such mass transport problems are well-posed [59, 230]. This scaling does not hold in the presence of radiation, a problem that already occurs in ΛCDM on

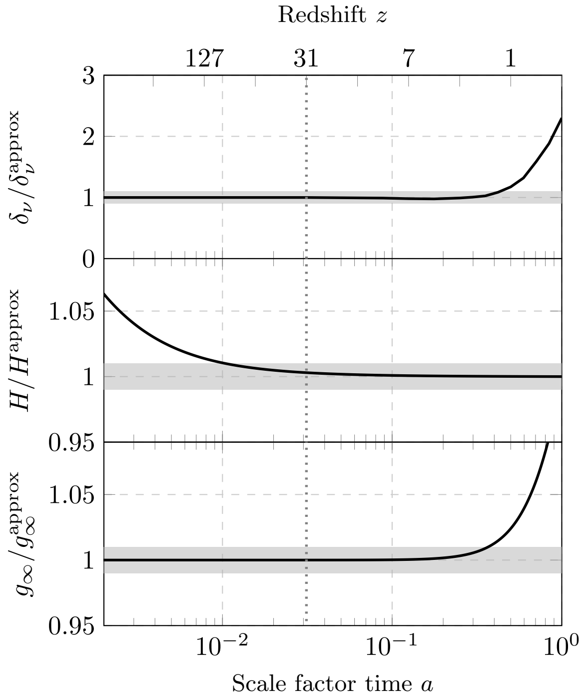

Figure 5.1: Accuracy of the linear response approximation (5.3.6) evaluated at $k = 0.60$ Mpc^{-1}, compared to a reference simulation (top), of the Hubble rate (5.3.11) when neglecting radiation (middle), and of (5.3.18) for the constant matter-dominated value for g_∞. The vertical dotted line indicates the fiducial starting redshift of $z_i = 31$. The neutrino mass sum is $\sum m_\nu = 0.3$ eV and the shaded region is 10% (top) and 1% (middle & bottom).

account of the cosmic microwave background radiation, but is certainly made worse by the inclusion of massive neutrinos, which scale like radiation in the relativistic régime. Therefore, we need to start the integration at a time when the relativistic contribution of neutrinos to the Hubble expansion can be neglected. Note that we make this assumption to ensure a consistent mathematical framework for the higher-order LPT solutions. However, it is not needed for the linear transfer functions, the back-scaling procedure or in the N-body code itself. In each of those cases, we do take the relativistic neutrino contribution into account.

Before proceeding, let us give the following convenient expression for g_∞ in the limit $\Lambda \to 0$:

$$g_\infty^{-1/2} = \frac{a^{3/2}H}{\sqrt{\frac{2}{3}B_0}}\frac{\mathrm{d}\log D_\infty}{\mathrm{d}\log a} = \frac{1}{4}\frac{\sqrt{1+24(1-f_\nu)}-1}{\sqrt{1-f_\nu}}. \tag{5.3.18}$$

Both numerator and denominator scale approximately as $(1 - f_\nu)^{1/2}$. The numerator is simply the exponent of the growing mode in (5.3.13), while the dependence of the denominator can be traced to the appearance of B_0 on the right-hand side of (5.3.10). The resulting smallness of $g_\infty - 1$ explains why neutrino corrections at nth order are small relative to D^n_∞: the lack of neutrino clustering is largely compensated by slower growth of the linear solution. In the next section, we will validate the assumptions made up to this point.

5.3.3. Validity of assumptions

Central to the approach of Section 5.4 is the linear response approximation (5.3.6) for the nonlinear neutrino density, $\delta_\nu(\boldsymbol{k})$. This approximation is very accurate at early times, but underestimates neutrino clustering on small scales and neglects the phase shift between neutrinos and dark matter that builds up at late times (see Fig. 4.6 in Chapter 4). The top panel of Fig. 5.1 shows the nonlinear neutrino density contrast, computed from a simulation with neutrino particles, relative to the linear neutrino response evaluated at $k = 0.60\text{ Mpc}^{-1}$. The neutrino mass is $\sum m_\nu = 0.3$ eV. The figure suggests that the approximation is valid at this scale up to $z \approx 1.5$, when perturbation theory has presumably already broken down. Hence, approximation (5.3.6) is well-suited for our application at much higher redshifts.

A second approximation is that we neglect the contribution of the relativistic tail of the neutrino distribution to the Hubble rate in (5.3.11). We reiterate that this approximation is only made for the calculation of the higher-order kernels and not in any of the calculations at first order. The middle panel of Fig. 5.1 shows that this approximation is accurate to better than 1% for $a > 0.01$, for our default neutrino mass of $\sum m_\nu = 0.3$ eV. In particular, at the fiducial starting redshift of $z_i = 31$, the error is 0.3%. We are helped in this regard by our preference for late starts.

Finally, we assume that g_∞ is constant in Section 5.4.1. The bottom panel of Fig. 5.1 shows that this is an excellent approximation, except at late times during Λ-domination. The figure suggests that there is a window where all assumptions are valid, potentially allowing us to push to even later starts, with the breakdown of LPT likely being the limiting factor.

5.4. Lagrangian approach

In the Lagrangian approach to gravitational instability [62, 226, 231–235], the objective is to describe fluid particle trajectories

$$\boldsymbol{x}(\boldsymbol{q}) = \boldsymbol{q} + \boldsymbol{\psi}(\boldsymbol{q}), \tag{5.4.1}$$

in terms of a displacement field $\boldsymbol{\psi}$. We use the Helmholtz decomposition, writing the Laplacian of a smooth vector field as

$$\nabla^2 \boldsymbol{\psi} = \nabla\left(\nabla \cdot \boldsymbol{\psi}\right) - \nabla \times \left(\nabla \times \boldsymbol{\psi}\right). \tag{5.4.2}$$

What remains is to solve for the longitudinal and transverse derivatives. The displacement is related to the Eulerian density, δ_{cb}, through the mass conservation equation

$$\delta_{\mathrm{cb}}(\boldsymbol{x}) = \frac{1}{J(\boldsymbol{q})} - 1, \tag{5.4.3}$$

where $J(\boldsymbol{q})$ is the determinant of the Jacobian of the coordinate transformation, $J_{ij} = \partial x_i / \partial q_j$, given by

$$J = \det J_{ij} = 1 + \psi_{i,i} + \frac{1}{2}\left[\psi_{i,i}\psi_{j,j} - \psi_{i,j}\psi_{j,i}\right] + \det \psi_{i,j}. \tag{5.4.4}$$

Let $(\partial/\partial D_\infty)_{\mathrm{L}} = (\partial_{D_\infty} + \boldsymbol{v}_{\mathrm{cb}} \cdot \nabla_{\boldsymbol{x}})$ be the Lagrangian derivative. The Lagrangian form of the Euler Eq. (5.3.15) can be written as

$$\mathcal{D}_\infty \boldsymbol{x} = -\frac{3g_\infty}{2D_\infty} \nabla_{\boldsymbol{x}} \varphi, \tag{5.4.5}$$

where we used $\boldsymbol{v}_{\mathrm{cb}} = (\partial \boldsymbol{x}/\partial D_\infty)_{\mathrm{L}}$ and introduced the linear operator

$$\mathcal{D}_\infty = \left(\frac{\partial}{\partial D_\infty}\right)^2_{\mathrm{L}} + \frac{3g_\infty}{2D_\infty}\left(\frac{\partial}{\partial D_\infty}\right)_{\mathrm{L}}. \tag{5.4.6}$$

Using (5.3.17) and taking the divergence and curl of (5.4.5), we find that the evolution of the displacement is governed by

$$\nabla_{\boldsymbol{x}} \cdot \mathcal{D}_\infty \boldsymbol{x}(\boldsymbol{q}) = -\frac{3g_\infty}{2D_\infty^2}\left[\delta_{\mathrm{cb}} * (1+\alpha)\right](\boldsymbol{x}), \tag{5.4.7}$$

$$\nabla_{\boldsymbol{x}} \times \mathcal{D}_\infty \boldsymbol{x}(\boldsymbol{q}) = 0. \tag{5.4.8}$$

To facilitate a fully Lagrangian description, we define the frame-lagging terms [212, 236]

$$F(\boldsymbol{q}) \equiv \left[(1/J-1) * \alpha\right](\boldsymbol{q}) - [\delta_{\rm cb} * \alpha](\boldsymbol{x}). \tag{5.4.9}$$

Frame-lagging terms arise from mapping the Eulerian neutrino response to Lagrangian coordinates. We give explicit expressions up to second order in Appendix 5.C. Transforming the derivatives on the left-hand side of (5.4.7) and (5.4.8) using $\partial_{x_i} = (\partial q_j/\partial x_i)\partial_{q_j} = J^{-1}_{ij}\partial_{q_j}$ and using the Monge-Ampère Eq. (5.4.3), we write these equations in Lagrangian coordinates as

$$J^{-1}_{ij}\mathcal{D}_\infty \psi_{i,j} = \frac{3g_\infty}{2D^2_\infty}\left[(1-1/J)*(1+\alpha)+F\right], \tag{5.4.10}$$

$$\epsilon_{ijk}J^{-1}_{jl}\mathcal{D}_\infty\psi_{k,l} = 0. \tag{5.4.11}$$

It will be the task in the following sections to find perturbative solutions for $\boldsymbol{\psi}$. We perform an expansion in displacements, writing

$$\boldsymbol{\psi} = \sum_{n=1}^{\infty}\boldsymbol{\psi}^{(n)}, \tag{5.4.12}$$

where $\boldsymbol{\psi}^{(n)}$ is of order $\left[\boldsymbol{\psi}^{(1)}\right]^n$.

5.4.1. Limiting solutions

Having set up the Lagrangian equations for the neutrino-cb fluid model, we are now in a position to look for approximate solutions. The aim is to find expressions for the displacement on large and small scales. In the small-scale limit, neutrinos do not cluster and only contribute to the background expansion as encoded by g_∞. Meanwhile, in the large-scale limit, neutrinos cluster like cold dark matter and one recovers behaviour analogous to ΛCDM. In both cases, we can find simple solutions in the form of LPT recursion relations [59, 63–66]. These limiting solutions will be used as initial conditions for the numerical integration of the general problem and provide the basis for the recipe of Section 5.2.

In this section, we assume that $g_\infty =$ constant, which is exact during matter domination (Eq. (5.3.18)), and a very good approximation in general (Fig. 5.1). On large scales, we also have $1+\alpha(k) = 1 + f_\nu/f_{\rm cb}$[5] and on small scales $1+\alpha(k)=1$. Hence, if all modes

---

[5] This is not strictly true, since $\delta_\nu > \delta_{\rm cb}$ on the largest scales due to the relativistic tail of the neutrino distribution. We ignore this small effect in the current section and in Fig. 5.2, but take it into account in Section 5.4.2.

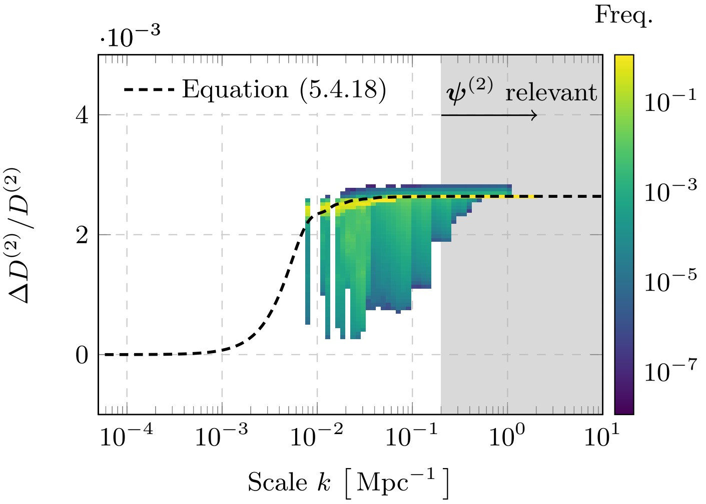

Figure 5.2: Correction to the ΛCDM prediction of $D^{(2)} = (3/7)D^2$ for the second-order growth factor, according to the approximate model of Eq. (5.4.18), for $\sum m_\nu = 0.3$ eV at $z = 31$ (dashed line). The colours represent a histogram of the full numerical solution, $D_B^{(2)}(\boldsymbol{k}_1, \boldsymbol{k}_2)$, evaluated on a 6D Fourier space lattice with physical dimension $L = 800$ Mpc (i.e. $\Delta k = 7.85 \times 10^{-3}$ Mpc^{-1}), projected onto the $k = \|\boldsymbol{k}_1 + \boldsymbol{k}_2\|$-axis and normalized per k-bin. For the large majority of configurations, the system attains the approximate value. The shaded region indicates the range of scales for which the power spectrum of $\boldsymbol{k} \cdot \boldsymbol{\psi}^{(2)}$ is at least 0.01% of that of $\boldsymbol{k} \cdot \boldsymbol{\psi}^{(1)}$.

involved in the problem are either large or small, we can approximate the convolution with the neutrino response as multiplication by a constant $\beta = 1 + \alpha(k)$. In such cases, the frame-lagging terms also vanish, as will be confirmed in Section 5.4.2. Given these assumptions, (5.4.10) reduces to

$$J_{ij}^{-1}\mathcal{D}_\infty \psi_{i,j} = \frac{3\beta g_\infty}{2D_\infty^2}(1 - 1/J). \tag{5.4.13}$$

Using the identities $JJ_{ij}^{-1} = (1/2)\epsilon_{jkp}\epsilon_{iqr}J_{kq}J_{pr}$ and $J = (1/6)\epsilon_{ijk}\epsilon_{pqr}J_{ip}J_{jq}J_{kr}$, we rewrite (5.4.13) and (5.4.11) as

$$\epsilon_{ijk}\epsilon_{pqr}J_{qj}J_{ip}\left[\mathcal{D}_\infty - \frac{\beta g_\infty}{2D_\infty^2}\right]J_{kr} + \frac{3\beta g_\infty}{D_\infty^2} = 0, \tag{5.4.14}$$

$$\epsilon_{lpq}J_{qk}\mathcal{D}_\infty \psi_{k,l} = 0. \tag{5.4.15}$$

Hence, using $J_{ij} = \delta_{ij} + \psi_{i,j}$ and substituting the expansion (5.4.12), we obtain equations for the longitudinal and transverse parts at order n in terms of perturbations of orders

$m_1+m_2=n$ (for $n\geq 2$) and $m_1+m_2+m_3=n$ (for $n\geq 3$):

$$\begin{aligned} \left[\mathcal{D}_\infty - \frac{3\beta g_\infty}{2D_\infty^2}\right]\nabla\cdot\boldsymbol{\psi}^{(n)} = \\ &- \sum_{m_1+m_2=n} \epsilon_{ijk}\epsilon_{ipq}\psi^{(m_1)}_{j,p}\left[\mathcal{D}_\infty - \frac{3\beta g_\infty}{4D_\infty^2}\right]\psi^{(m_2)}_{k,q} \\ &- \sum_{m_1+m_2+m_3=n} \epsilon_{ijk}\epsilon_{pqr}\frac{1}{2}\psi^{(m_1)}_{i,p}\psi^{(m_2)}_{j,q}\left[\mathcal{D}_\infty - \frac{\beta g_\infty}{2D_\infty^2}\right]\psi^{(m_3)}_{k,r}, \end{aligned} \tag{5.4.16}$$

$$\mathcal{D}_\infty\nabla\times\boldsymbol{\psi}^{(n)} = \sum_{m_1+m_2=n}\nabla\psi^{(m_1)}_i\times\mathcal{D}_\infty\nabla\psi^{(m_2)}_i. \tag{5.4.17}$$

The first-order equations separate. The longitudinal Eq. (5.4.16) has the particular time-dependent solution

$$D^{(1)} = D^q_\infty \quad \text{with} \quad q = \tfrac{1}{4}\sqrt{4+3g_\infty(8\beta+3g_\infty-4)} - \tfrac{3}{4}g_\infty + \tfrac{1}{2},$$

while the transverse Eq. (5.4.17) has constant and decaying solutions. Identifying the fastest growing solutions order by order, we find that $\boldsymbol{\psi}^{(n)}\propto D^{nq}_\infty$. In particular, we find that the fastest growing solution at second order grows as

$$\frac{D^{(2)}}{D^{2q}_\infty} = \frac{3g_\infty\beta}{4q(2q-1)+3g_\infty(2q-\beta)}. \tag{5.4.18}$$

Reinserting $\beta = 1+\alpha(k)$, we obtain a useful approximation of the magnitude of neutrino effects on the second-order coefficient, relative to the ΛCDM value of 3/7. This is shown by the dashed line in Fig. 5.2 for a model with $\sum m_\nu = 0.3$ eV at $z=31$. We stress that this approximation neglects the non-trivial coupling with the neutrino response in the general case. As we will see in the next section, the second-order solution can be described in full by two kernels, $D^{(2)}_A(\boldsymbol{k}_1,\boldsymbol{k}_2)$ and $D^{(2)}_B(\boldsymbol{k}_1,\boldsymbol{k}_2)$. For most configurations on the 6D Fourier space lattice that we use to generate N-body ICs, both k_1 and k_2 are large and the result is close to the estimate of Eq. (5.4.18). However, for cases with one mode large and one mode small or for squeezed configurations with $k = \lVert\boldsymbol{k}_1+\boldsymbol{k}_2\rVert \ll k_1\approx k_2$, the value may depart from this estimate, as shown by the histogram in Fig. 5.2. Nevertheless, the figure demonstrates that the large- and small-scale limits provide reasonable bounds on the effect at intermediate scales. Overall, the magnitude of the effect is $\mathcal{O}(10^{-3})$, in line with the estimate given in Section 5.2.2 for this mass. The figure also demonstrates that the ΛCDM value of 3/7 is only reached for $k<10^{-3}\text{ Mpc}^{-1}$, while the second-order potential is important for $k>10^{-1}\text{ Mpc}^{-1}$, reflecting the hierarchy between the neutrino free-streaming scale and the nonlinear scale, $k_\text{fs}\ll k_\text{nl}$, that motivates the approach of

Section 5.2.

Using $\boldsymbol{\psi}^{(n)} \propto D_\infty^{nq}$, we derive recursion relations for the fastest growing solution at order $n \geq 2$:

$$\begin{aligned} \nabla \cdot \boldsymbol{\psi}^{(n)} = &- \sum_{m_1+m_2=n} \frac{1}{2} \left[1 - \frac{4 m_1 m_2 q^2}{2nq(nq-1) + 3g_\infty (nq - \beta)} \right] \\ &\times \epsilon_{ijk} \epsilon_{ipq} \psi_{j,p}^{(m_1)} \psi_{k,q}^{(m_2)} \\ &- \sum_{m_1+m_2+m_3=n} \left[1 - \frac{4(m_1 m_2 + m_2 m_3 + m_3 m_1) q^2}{2nq(nq-1) + 3g_\infty (nq - \beta)} \right] \\ &\times \epsilon_{ijk} \epsilon_{pqr} \frac{1}{6} \psi_{i,p}^{(m_1)} \psi_{j,q}^{(m_2)} \psi_{k,r}^{(m_3)}, \end{aligned} \tag{5.4.19}$$

$$\nabla \times \boldsymbol{\psi}^{(n)} = \sum_{m_1+m_2=n} \frac{1}{2} \frac{m_2 - m_1}{n} \nabla \psi_i^{(m_1)} \times \nabla \psi_i^{(m_2)}. \tag{5.4.20}$$

For the purposes of higher-order ICs, we are primarily interested in deriving corrections to the ΛCDM coefficients in the small-scale limit with $\beta = q = 1$. Reading off coefficients from (5.4.19), we find that these can be conveniently expressed in terms of

$$C_n = \frac{(2n+3) g_\infty}{2n + 3g_\infty}. \tag{5.4.21}$$

Proceeding as in Appendix 5.D, we obtain the 3LPT form given in Section 5.2.2. Combining Eqs. (5.4.21) and (5.3.18) yields an accurate approximation of C_n in terms of f_ν:

$$C_n = \frac{8(1-f_\nu)(2n+3)}{n(S-1)^2 + (S^2-1)} \cong 1 + \frac{2 f_\nu n}{5(2n+3)}, \tag{5.4.22}$$

with $S = \sqrt{1 + 24(1-f_\nu)}$. For $n = 2$, the above expression agrees with that given by [212]. The next section is dedicated to relaxing the assumptions on g_∞ and $\alpha(k)$, finding the general solution at second order.

5.4.2. General solution

For the general solution, we need to deal with the frame-lagging terms $F(\boldsymbol{q})$. Here, we will follow the approach of [213]. We are interested in solutions at second order. The transverse Eq. (5.4.11) only has non-trivial solutions for $n \geq 3$. Therefore, we concentrate on the longitudinal part. We repeat (5.4.10) for convenience:

$$J^{-1}_{ij}\mathcal{D}_\infty\psi_{i,j} = \frac{3g_\infty}{2D^2_\infty}\left[(1-1/J)*(1+\alpha)+F\right]. \tag{5.4.23}$$

Using (5.4.4) and $J^{-1}_{ij} = \sum_{n=0}^{\infty}[(I-J)^n]_{ij}$, we can write this up to second order in the displacement:

$$\begin{aligned}\mathcal{D}_\infty\psi_{i,i} &= \psi_{i,j}\mathcal{D}_\infty\psi_{j,i} + \frac{3g_\infty}{2D^2_\infty}\psi_{i,i}*(1+\alpha)\\ &\quad - \frac{3g_\infty}{2D^2_\infty}\frac{1}{2}\left[\psi_{i,i}\psi_{j,j}+\psi_{i,j}\psi_{j,i}\right]*(1+\alpha) + \frac{3g_\infty}{2D^2_\infty}F^{(2)},\end{aligned} \tag{5.4.24}$$

where the second-order frame-lagging terms, $F^{(2)}$, are given in Appendix 5.C. At first order, the displacement admits a growing solution $\boldsymbol{\psi}^{(1)} \propto D^{(1)}$ with a growth factor that satisfies

$$\mathcal{D}_\infty D^{(1)} = \frac{3g_\infty}{2D^2_\infty}(1+\alpha)D^{(1)}. \tag{5.4.25}$$

This is simply a reformulation of the Eulerian equation for the first-order growth factor (5.3.9). Using the expansion (5.4.12) in (5.4.24) and collecting second-order terms then yields

$$\begin{aligned}\mathcal{D}_\infty\psi^{(2)}_{i,i} &= \frac{3g_\infty}{2D^2_\infty}\psi^{(2)}_{i,i}*(1+\alpha) + \psi^{(1)}_{i,j}\mathcal{D}_\infty\psi^{(1)}_{j,i}\\ &\quad - \frac{3g_\infty}{2D^2_\infty}\frac{1}{2}\left[\psi^{(1)}_{i,i}\psi^{(1)}_{j,j}+\psi^{(1)}_{i,j}\psi^{(1)}_{j,i}\right]*(1+\alpha) + \frac{3g_\infty}{2D^2_\infty}F^{(2)}.\end{aligned} \tag{5.4.26}$$

In Fourier space, each of the quadratic terms in (5.4.26), including the second-order frame-lagging term, is a convolution of derivatives of $\boldsymbol{\psi}^{(1)}(\boldsymbol{k}_1)$ and $\boldsymbol{\psi}^{(1)}(\boldsymbol{k}_2)$. Expressing the displacements in terms of potentials as

$$\boldsymbol{\psi}^{(1)} = -\nabla\varphi^{(1)}, \qquad \boldsymbol{\psi}^{(2)} = -\nabla\varphi^{(2)}, \tag{5.4.27}$$

and identifying terms, we thus obtain

$$\begin{aligned}\varphi^{(2)}(\boldsymbol{k}) = \frac{1}{2}\int_{\boldsymbol{k}_1,\boldsymbol{k}_2}&\frac{1}{(ik)^2}\frac{1}{D_1D_2}\varphi^{(1)}(\boldsymbol{k}_1)\varphi^{(1)}(\boldsymbol{k}_2)\\ &\times\left[D^{(2)}_A(\boldsymbol{k}_1,\boldsymbol{k}_2)k_1^2k_2^2 - D^{(2)}_B(\boldsymbol{k}_1,\boldsymbol{k}_2)k_{12}^2\right],\end{aligned} \tag{5.4.28}$$

where $\int_{\boldsymbol{k}_1,\boldsymbol{k}_2} = \int \mathrm{d}\boldsymbol{k}_1 \mathrm{d}\boldsymbol{k}_2 (2\pi)^{-6}\delta^{(3)}(\boldsymbol{k}_1+\boldsymbol{k}_2-\boldsymbol{k})$ and $k_{12}=\boldsymbol{k}_1\cdot\boldsymbol{k}_2$ and $D_i = D^{(1)}(k_i)$ for $i=1,2$. Notice the similarity of this equation with Eq. (5.2.7). The difference is that the two terms now have distinct scale- and time-dependent coefficients satisfying

$$\mathcal{D}_\infty D_A^{(2)} = \frac{3g_\infty}{2D_\infty^2}\left(1+\alpha(k)\right)D_A^{(2)} + \frac{3g_\infty}{2D_\infty^2}(1+A)D_1D_2, \tag{5.4.29}$$

$$\mathcal{D}_\infty D_B^{(2)} = \frac{3g_\infty}{2D_\infty^2}\left(1+\alpha(k)\right)D_B^{(2)} + \frac{3g_\infty}{2D_\infty^2}(1+B)D_1D_2, \tag{5.4.30}$$

where the functions A and B are given by

$$A(k,k_1,k_2) = \alpha(k) + \left[\frac{\alpha(k)-\alpha(k_2)}{k_1^2} + \frac{\alpha(k)-\alpha(k_1)}{k_2^2}\right]k_{12}, \tag{5.4.31}$$

$$B(k,k_1,k_2) = \alpha(k_1)+\alpha(k_2)-\alpha(k), \tag{5.4.32}$$

for $k = \|\boldsymbol{k}_1+\boldsymbol{k}_2\|$. The terms in square brackets correspond to the frame-lagging terms. In the small-scale limit with $k,k_1,k_2 \gg k_\mathrm{fs}$, we have $A=B=0$. Hence, $D_A^{(2)} = D_B^{(2)}$ and (5.4.28) factorizes as in Eq. (5.2.7). Similarly, in the large-scale limit with $k,k_1,k_2 \ll k_\mathrm{fs}$, we obtain again the approximate form described in Section 5.4.1 with $A=B\approx f_\nu/f_\mathrm{cb}$. In both limits, the frame-lagging terms drop out, as anticipated. Intermediate configurations will deviate from the asymptotic solutions, as was already discussed in Section 5.4.1 and shown in Fig. 5.2.

For the numerical solution, we begin the integration at a time when the non-relativistic neutrino fraction is 50%. For the fiducial neutrino mass, $\sum m_\nu = 0.3$ eV, this corresponds to $z=187$. We integrate Eqs. (5.4.25) for the first-order growth factor and (5.4.29-5.4.30) for the second-order kernels, using the approximate model of Eq. (5.4.18) as initial conditions. The results, projected onto the k-axis, are shown in Fig. 5.2. When generating 2LPT particle initial conditions, we begin by generating a realisation of the back-scaled first-order potential, $\varphi^{(1)}$. We then perform the convolution integral of Eq. (5.4.28) explicitly, interpolating from tables of $D_{A,B}^{(2)}(k,k_1,k_2)$. To ensure completion in a reasonable time frame, we impose cut-offs at $k_1 \leq k_\mathrm{cut}$ and $k_2 \leq k_\mathrm{cut}$. We performed convergence tests to ensure that the results are independent of the cut-off scale, finding that a cut-off at $k_\mathrm{cut} = 1\ \mathrm{Mpc}^{-1}$ was more than adequate for the resolutions considered in this paper.

5.5. Results

We will now discuss the power spectra, bispectra, and halo mass functions of massive neutrino simulations with different ICs. We introduce our simulation suite in Section 8.2. We then consider the impact of different approximation schemes for the second-order kernels in Section 5.5.2 and follow it up with a comparison of Zel'dovich (ZA), 2LPT, and 3LPT ICs at various starting redshifts in Section 5.5.3. Finally, we consider the impact of ICs on the suppression of the power spectrum as a function of neutrino mass in Section 5.5.4.

5.5.1. Simulations

We use the cosmological hydrodynamics code SWIFT [189, 190], which uses task-based parallelism, asynchronous communication, fast neighbour finding, and vectorised operations to achieve significant speed-ups. The code uses the Fast Multipole Method (FMM) for short-range gravitational forces and the Particle Mesh method for long-range forces. Neutrinos are modelled as a separate particle species. We employ the δf method to suppress the effects of shot noise [237] and generate neutrino particle initial conditions by integrating geodesics from high redshift using our FASTDF code. Additionally, we use fixed initial conditions to limit cosmic variance [238]. Apart from the neutrino mass, we use cosmological parameters based primarily on Year 3 results from the Dark Energy Survey [201] and Planck 2018 [112]. Our choice of parameters is $(h, \Omega_\mathrm{m}, \Omega_\mathrm{b}, A_s, n_s) = (0.681, 0.306, 0.0486, 2.09937 \times 10^{-9}, 0.967)$, with different choices for the neutrino density Ω_ν.

There is a subtle point regarding comparisons between simulations with and without massive neutrinos. Codes like SWIFT employ a multipole acceptance criterion to determine when the multipole approximation is sufficiently accurate to be used without further refinement. The adaptive criterion used for the runs in this paper is based on error analysis of forces on test particles. This means that the accuracy of the N-body calculation depends on the number of particles contained in any given volume. When comparing two runs with equal numbers of dark matter particles, one with neutrinos and the other without, all other things being equal, forces will be calculated more accurately in the run with neutrinos. To account for this difference, we included an equal number of massless 'spectator' neutrino particles in the $f_\nu = 0$ runs, with velocities corresponding to $m_\nu = 0.05$ eV neutrinos. These particles contribute no forces and only affect the N-body simulation through the multipole acceptance criterion, ensuring that the accuracy of the

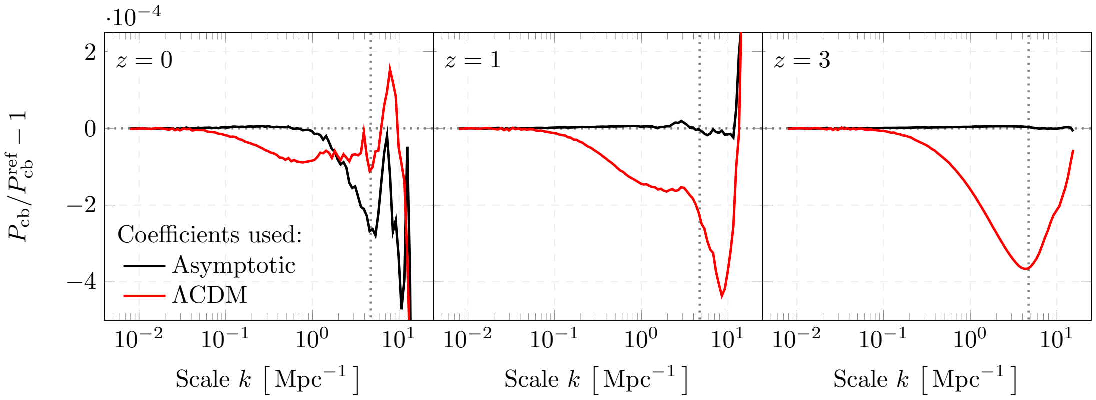

Figure 5.3: Impact of approximation schemes for the second-order potential on the CDM & baryon power spectrum. The reference run used initial conditions based on a numerical calculation of the scale-dependent 2LPT kernels. In the asymptotic approximation (black), we use Eqs. (5.2.4) and (5.2.13), but truncate third-order terms. In the ΛCDM approximation (red), we additionally set $C_2 = 1$. The vertical dotted line is the Nyquist frequency.

massless runs is comparable to that of the massive neutrino runs. Such massless runs are considered in Section 5.5.4.

5.5.2. Validation of approximate treatment

To validate our approach, we compare three different implementations of 2LPT, based on the following models:

1. The asymptotic model of Section 5.2
2. A model with ΛCDM coefficients
3. A reference model with scale-dependent effects

The first order displacements and velocities are identical in each of the approaches, obtained from the back-scaled linear power spectrum at $z = 0$. In the asymptotic scheme, we use Eqs. (5.2.4) and (5.2.13), but truncate the 3LPT terms. In the ΛCDM approximation, we additionally set $C_2 = 1$, which corresponds to neglecting neutrino effects at second order. Finally, we compare these two approximate methods with a reference run that relied on a numerical calculation of the scale-dependent 2LPT kernels, $D^{(2)}_A(\boldsymbol{k}_1, \boldsymbol{k}_2)$ and $D^{(2)}_B(\boldsymbol{k}_1, \boldsymbol{k}_2)$. With respect to Fig. 5.2, the asymptotic approximation corresponds to using the small-scale limit, the ΛCDM approximation corresponds to the large-scale limit, and the reference run corresponds to the underlying histogram. We use simulations with side length $L = 800$ Mpc and $N_{\rm cb} = 1200^3$ particles.

Fig. 5.3 shows the impact of these approximations on the power spectrum of the evolved CDM & baryon density field. The differences are most evident at $z = 3$ (right panel). On the largest scales, $k < 0.05$ Mpc^{-1}, nonlinear corrections are small and all simulations agree to machine precision. For $k > 0.05$ Mpc^{-1}, the ΛCDM simulation systematically underestimates clustering with a maximum error of 0.04% at $k = 4$ Mpc^{-1}. For the asymptotic run, the error is two orders of magnitude smaller over the same scales. Between $z = 31$ and $z = 3$, the evolution is virtually identical in the asymptotic and reference runs, but we begin to see some noise in the ratio on the smallest scales at $z = 1$ (middle panel). These perturbations continue to grow until $z = 0$ (left panel), where we find a scatter of 2×10^{-4} for $k > 1$ Mpc^{-1} in both the asymptotic/reference and ΛCDM/reference ratios. It is hard to attribute this noise to any particular run as the power spectrum on these scales is increasingly determined by the internal structure of poorly resolved halos. On larger scales, $k < 1$ Mpc^{-1}, the asymptotic run performs extremely well with errors below 10^{-5}, while the systematic deficit in the ΛCDM run persists.

These results demonstrate that, at second order, the effect of the suppressed neutrino perturbations can be absorbed into a scale-independent factor C_2 and that further scale-dependent neutrino effects are negligible as far as initial conditions are concerned. We expect that this continues to hold for third-order corrections, which are confined to even smaller scales. Including the correction factor C_2 is clearly superior to simply using the ΛCDM coefficient. However, we also observe that this higher-order neutrino effect is below 0.1%, and therefore beyond the sensitivity of current experiments. Hence, we conclude that for most purposes both the ΛCDM approximation and the asymptotic approximation are justified.

5.5.3. Choice of LPT order and starting time

We are now in a position to study the effects of LPT order and starting time on massive neutrino simulations, using the asymptotic approximation. Fig. 5.4 shows the late-time power spectrum for simulations with $L = 800$ Mpc and $N_{\text{cb}} = 1200^3$ particles, comparing in the first instance Zel'dovich (solid red) and 2LPT (solid black) with 3LPT (dotted gray) as a baseline. All three runs were started at $z_i = 31$. The most striking observation is that the differences are much larger than those shown in Fig. 5.3. This means that using higher-order LPT in some fashion is more important than getting the details right. Next, we find percent agreement between 2LPT and 3LPT over the entire range of scales probed for $z \leq 1$ and approximately a 1% error at $z = 3$ for $k > 2$ Mpc^{-1}. We also find that the Zel'dovich approximation performs very poorly with errors of (4, 7, 15)% for $k > 1$ Mpc^{-1} at $z = (0, 1, 3)$. This well-known fact [88] has motivated practitioners to

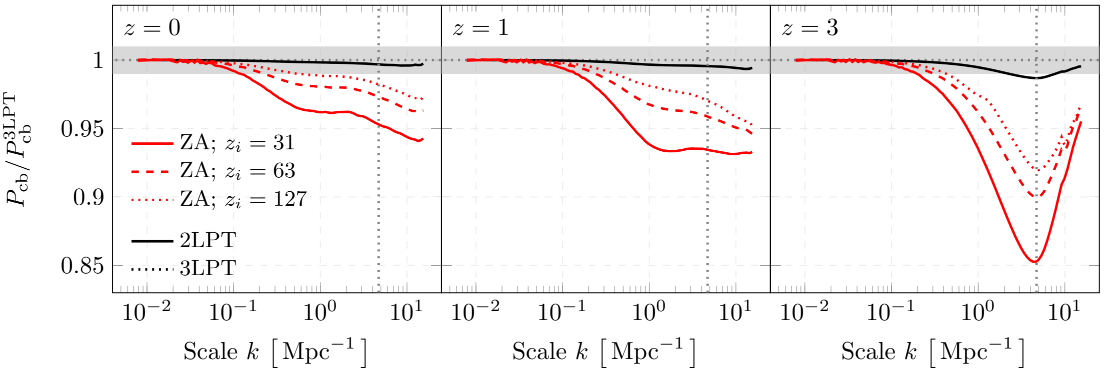

Figure 5.4: Impact of starting time and LPT order on the CDM & baryon power spectrum. The reference simulation used 3LPT and both it and the 2LPT simulation were started at $z_i = 31$. The shaded area is 1%.

start Zel'dovich simulations at higher redshifts, when truncation errors are smaller. We demonstrate this with Zel'dovich runs started at $z_i = 63$ (dashed, red) and $z_i = 127$ (dotted, red). While the agreement with the higher-order runs improves, we still find percent agreement only up to $k = 0.4\ \mathrm{Mpc}^{-1}$. Moreover, starting earlier introduces inaccuracies of a different sort. To see this, we repeat the exercise at a lower resolution with $N_{\mathrm{cb}} = 600^3$ particles. The resulting power spectra at $z = 0$ are shown in Fig. 5.5, with Zel'dovich runs compared against 3LPT in the left panel. We observe that for runs started at $z_i = 31$ (red), the error is almost independent of resolution. However, for earlier starts at $z = 63$ (black) and $z = 127$ (blue), the lower resolution runs increasingly underestimate the power spectrum on small scales. This shows that while truncation errors decrease, resolution effects increase as simulations are started earlier. The pattern reverses for 2LPT (right panel), with earlier starts performing worse than later starts. This can easily be explained by the fact that truncation errors are much smaller for 2LPT, such that the effect of increasing discreteness errors dominates. We confirm the finding of [89] that the size of discreteness errors is independent of LPT order. This demonstrates that, at fixed resolution and LPT order, starting earlier does not guarantee convergence onto the higher-order solution. As was the case for truncation errors, discreteness errors are much larger at $z = 1$, 3 (not shown).

We also consider three-point statistics, which are sensitive to transients from initial conditions [88] and an interesting probe of neutrino physics [239–241]. For the equilateral bispectrum, $B(k) = B(k_1, k_2, k_3)$ with $k = k_1 = k_2 = k_3$, shown in Fig. 5.6 at late times, the same pattern is broadly repeated as for the power spectrum. However, errors are approximately twice as large as for the power spectrum. In detail, we again find percent agreement between 2LPT and 3LPT for $z \leq 1$ with larger errors on small scales at $z = 3$.

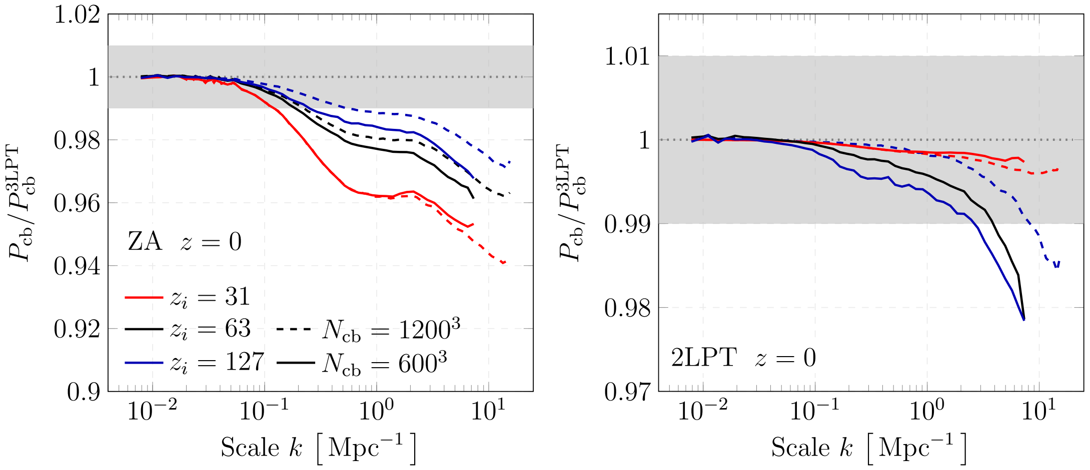

Figure 5.5: Impact of starting time and resolution on the CDM & baryon power spectrum. The simulations are compared against 3LPT runs with the same resolution ($N_{\rm cb} = 600^3$ or $N_{\rm cb} = 1200^3$), started at $z_i = 31$. The shaded area is 1%. Not all combinations were tested.

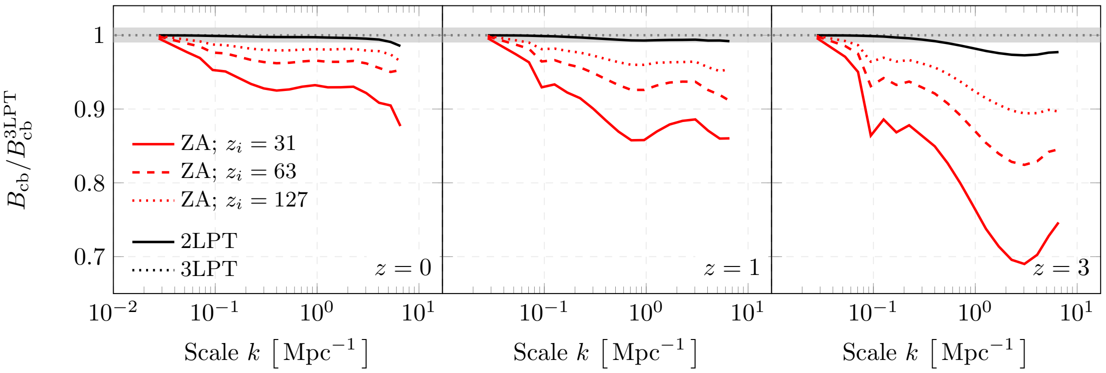

Figure 5.6: Impact of starting time and LPT order on the equilateral bispectrum of CDM & baryon density perturbations at late times. The reference simulation used 3LPT and both it and the 2LPT simulation were started at $z_i = 31$. All runs used $N_{\rm cb} = 1200^3$ particles. The shaded area is 1%.

For the Zel'dovich runs, we find significant errors compared to 3LPT, even when starting at $z = 127$, with percent agreement only up to $k = 0.1$ Mpc^{-1} at $z = 0$, and not even there for $z \geq 1$.

Finally, we compare the halo mass function at $z = 0$. Halos are identified with VELOCIRAPTOR [194] using a 6D friends-of-friends algorithm applied to the cb particles. Spherical overdensity masses are computed within spheres for which the density equals 200 times the mean CDM & baryon density $\bar{\rho}_{\rm cb}$. The reason for using $\bar{\rho}_{\rm cb}$ instead of the

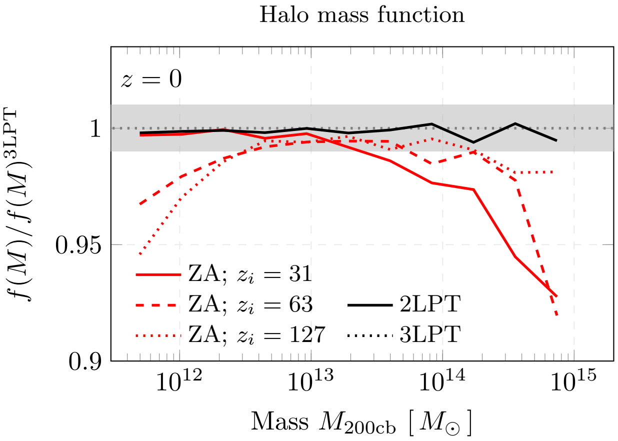

Figure 5.7: Impact of starting time and LPT order on the halo mass function, $f(M) = \mathrm{d}n/\mathrm{d}\log M$, at $z = 0$ for the $N_{\mathrm{cb}} = 1200^3$ runs. The reference simulation used 3LPT and was started at $z_i = 31$. The shaded area is 1%.

total mass density $\bar{\rho}_{\mathrm{m}}$ is that it is this cold density field that produces universal and unbiased results in halo model calculations [242–244]. The results are shown in Fig. 5.7. We once again find percent agreement between 2LPT and 3LPT over the entire mass range, but large errors for the Zel'dovich runs. There is an interesting pattern in the Zel'dovich error as the starting time is varied. For late starts (solid red), the simulation agrees well at the low-mass end but underestimates the number of very massive, $10^{15} M_\odot$, halos by more than 7%. This can be understood in terms of the deficit of power seen also in Fig. 5.4, resulting in a suppressed growth of large structures. Meanwhile, for early starts (dotted and dashed red), the agreement at the high-mass end improves like the small-scale power spectrum. However, the number of low-mass halos decreases by a similar factor, likely due to discreteness errors. This seems to be broadly consistent with the ΛCDM results of [89], but not with [245] who find little dependence on starting time at $z = 0$.

5.5.4. Dependence on neutrino mass

Thus far, we have focused on a single neutrino mass of $\sum m_\nu = 0.3$ eV. However, it is of great interest to determine the effect of initial conditions on the suppression of the power spectrum for different neutrino masses. We consider three cases:

1. massless neutrinos
2. degenerate $\sum m_\nu = 0.15$ eV neutrinos ($f_\nu = 0.011$),

3. degenerate $\sum m_\nu = 0.30$ eV neutrinos ($f_\nu = 0.023$).

In each case, we adjust $\Omega_{\rm cdm}$ to keep the total matter density $\Omega_{\rm m}$ fixed. To be able to carry out many variations, we primarily use lower resolution simulations with $N_{\rm cb} = 600^3$ particles in an $L = 800$ Mpc cube. This still enables us to study the impact of LPT order and starting time, as the following discussion reveals.

First, we consider the effect of LPT order. In Fig. 5.8, we show the suppression of the CDM & baryon power spectrum relative to the massless case, comparing ZA/ZA (solid), 2LPT/2LPT (dashed), and 3LPT/3LPT (shaded). Evidently, it is crucial to compare like with like simulation, keeping the LPT order and starting redshift the same. Not doing so introduces large errors in the ratio, as might be expected from the fixed neutrino mass results discussed above. We illustrate this by including a dotted line for the ZA/2LPT ratio, which is clearly off the mark. However, even when comparing like with like, we find a residual error that is proportional to the neutrino mass, rises with k, and peaks around the turn-over of the suppression. This feature is most clearly visible at $z = 1$ for ZA, with a maximum error of $0.05 f_\nu$. The effect is already present in the initial conditions and can be explained by a mass-dependent suppression of nonlinear terms. As virialized structures grow, both the turn-over of the suppression and the peak of the error move to larger scales. At $z = 0$, the error is $0.025 f_\nu$ around $k = 0.3$ Mpc^{-1}. On smaller scales, we see a scatter of order 0.5%, treading outside the scale-dependent error bars that correspond to a ±0.005 eV shift in $\sum m_\nu$. For 2LPT, both the systematic effect and the noise are greatly suppressed, resulting in 0.1%-level agreement with 3LPT even at early times.

Next, we consider the effect of the starting time of the simulation. In Fig. 5.9, we show the suppression for simulations with 2LPT ICs started at $z = 127$ (solid), $z = 63$ (dashed), $z = 31$ (shaded). Once again, we compare like with like simulations. Even so, we find a small residual effect with earlier starts overestimating the suppression. The differences between $z = 31$ and $z = 63$ are minimal for both neutrino masses. However, starting at $z = 127$ results in $(0.1, 0.2) f_\nu$ errors at $z = (0, 1)$ for $k > 1$ Mpc^{-1}. These errors once again exceed the threshold for a ±0.005 eV shift in $\sum m_\nu$. Based on the discussion above, and given that we are using 2LPT, we expect that truncation errors are small at both redshifts. This suggests that the differences are caused by resolution effects, which grow in importance with the starting redshift. To test this, we repeated some of the simulations at a higher resolution with $N_{\rm cb} = 1200^3$ particles, starting at $z = 127$ and $z = 31$. The ratio is shown as a dotted line in the bottom panels of Fig. 5.9. The agreement between the early and late starts improves to 0.1% up to $k = 10$ Mpc^{-1} at $z = 0$, comparable to the low-resolution $z = 63$ start. However, the suppression is still

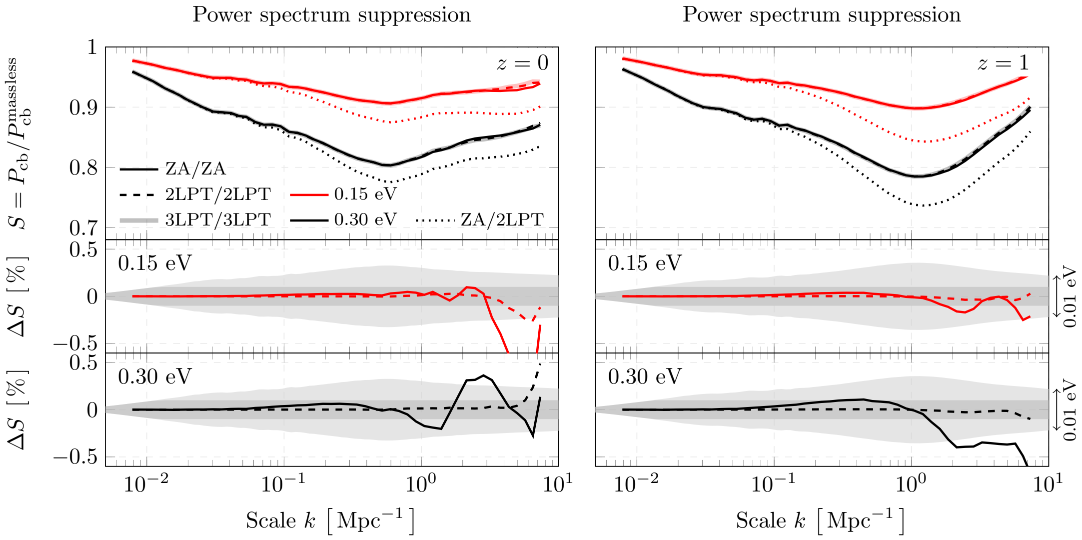

Figure 5.8: Impact of LPT order on the suppression of the CDM & baryon power spectrum for different choices of A/B, where A is the LPT order of the massive neutrino run and B the LPT order of the massless run. The neutrino masses are $\sum m_\nu = 0.15$ eV (red) and $\sum m_\nu = 0.30$ eV (black). The bottom panels show the suppression relative to 3LPT/3LPT, with shaded areas representing a ±0.005 eV shift (light) or a constant 0.1% error (dark) where this is smaller.

slightly overestimated at $z = 1$.

One possible alternative explanation is that errors could be introduced by the back-scaling procedure (Section 5.2.1). To test this hypothesis, we repeated some of the simulations with ICs that were not back-scaled, as in Chapter 4. We found nearly identical results for these runs, ruling out this explanation. Another possibility is that the errors could be the result of shot noise, since we use a particle-based implementation of neutrino perturbations. However, this is unlikely as the differences already appear at high redshift when shot noise is highly suppressed due to our use of the δf method. Finally, one might expect differences due to relativistic effects that are increasingly important for earlier starts. Once again, this is unlikely since relativistic effects would appear on the largest scales, where the differences shown in Fig. 5.9 are minimal. Since the error decreases for the higher resolution runs, discreteness effects likely account for the majority of the difference, with massive neutrino simulations being more sensitive to such errors, due to the suppressed growth of structure. Late starts can be utilized to minimize the effect of particle resolution, as shown in Fig. 5.5.

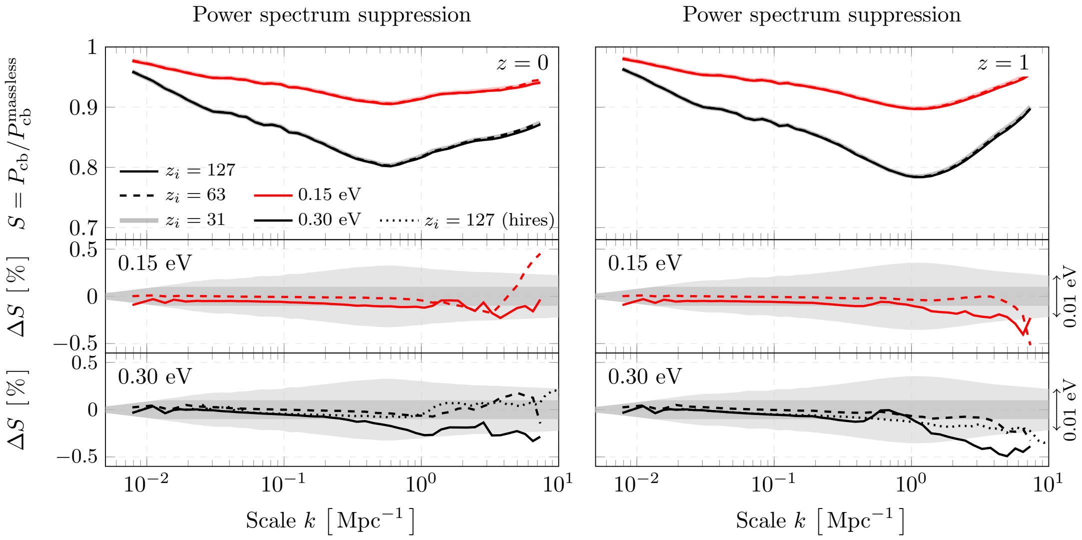

Figure 5.9: Impact of starting redshift and resolution on the suppression of the CDM & baryon power spectrum. The neutrino masses are $\sum m_\nu = 0.15$ eV (red) and $\sum m_\nu = 0.30$ eV (black). The bottom panels show the suppression relative to runs with the same resolution, but started at $z_i = 31$. The shaded areas represent a ± 0.005 eV shift (light) and a constant 0.1% error (dark) where this is smaller. All simulations used 2LPT initial conditions.

5.6. Discussion

We have investigated the use of higher-order Lagrangian initial conditions (ICs) for cosmological simulations with massive neutrinos. We solved the fluid equations for a neutrino-CDM-baryon model with approximate time-dependence in the large- and small-scale limits, finding that higher-order neutrino effects can be described by scale-independent coefficients that are easy to implement in existing IC codes. To validate our approach, we constructed ICs based on a rigorous treatment of the scale-dependent neutrino response in 2LPT, obtaining agreement with our scheme to better than one part in 10^5 up to $k = 1$ Mpc^{-1} in the power spectrum of the evolved CDM and baryon perturbations at late times.

Compared to these small differences, we find that the truncation error associated with using the first-order Zel'dovich approximation is much larger. For our fiducial model with $\sum m_\nu = 0.3$ eV and a starting redshift of $z_i = 31$, the error is 4% in the power spectrum and 7% in the equilateral bispectrum around $k = 0.5$ Mpc^{-1} at $z = 0$. Ratios of statistics from simulations with different neutrino masses can be calculated much more robustly, provided that the LPT order and starting redshift are the same. Nevertheless, even such ratios have a residual dependence on the ICs. For instance, Zel'dovich ICs introduce a mass-dependent error in the suppression of the power spectrum that grows with wavenumber k and redshift z, peaking around the turn-over of the suppression. We also find that the starting time of the simulation has an impact on the suppression over a wide range of scales and redshifts. Simulations started at $z_i = 127$ overestimate the suppression of the power spectrum on small scales, compared to later starts. While simulations can be started at higher redshifts to reduce truncation errors, this also increases the importance of particle resolution and relativistic effects. To minimize errors from initial conditions and particle resolution, simulations can be started at late times using higher-order ICs.

A major target of cosmological surveys is to measure the sum of neutrino masses. Assuming the minimum value allowed under the normal mass ordering, $\sum m_\nu = 0.06$ eV, cosmology could provide a 3σ detection and rule out the inverted mass ordering at 2σ by reaching a sensitivity of 0.02 eV, which is in reach of future cosmic microwave background and large-scale structure experiments [132, 202–204]. This corresponds to detecting 1% effects on the matter power spectrum on $0.1\text{ Mpc}^{-1} < k < 1\text{ Mpc}^{-1}$ scales. We should therefore aim for neutrino simulations with errors that are well below 1% on these scales. While Zel'dovich ICs fall short of this mark, our findings suggest that 2LPT is sufficiently accurate for most applications. Higher-order statistics at high redshift seem to be the notable exception, which could be relevant for Lyman-α forest simulations.

The accuracy of neutrino simulations depends on many factors: the accuracy of the linear transfer functions and back-scaling procedure [124, 221], the implementation of neutrino perturbations (e.g. [174], Chapter 4), neutrino initial conditions (Chapter 6), and dark matter and baryon initial conditions (this chapter). It has now been demonstrated that each of these factors can be controlled to within 1%. The remaining uncertainty is likely dominated by the choice of gravity solver. Achieving 1% agreement between different N-body codes is non-trivial even in the absence of neutrinos [218, 246, 247]. Fortunately, the accuracy of N-body codes should not in the first place be expected to deteriorate in the presence of neutrinos. In fact, the accuracy could even improve for particle-based implementations due to 'spectator' effects (Section 8.2). Indeed, since publication of this paper, a systematic comparison of neutrino simulations with different codes and identical initial conditions has shown that consistent results can be obtained if shot noise is satisfactorily addressed [130]. These explorations confirm our ability to simulate nonlinear clustering in Universes with massive neutrinos, allowing us to meet the demands of the next generation of surveys.

5.A. Difference and sum equations

As in (5.3.15–5.3.17), the component fluid Eqs. (5.3.1–5.3.2) can be rewritten using D_∞ as time variable and $\boldsymbol{v}_\lambda = \boldsymbol{u}_\lambda / \partial_\tau D_\infty$ as velocity:

$$\partial_{D_\infty} \boldsymbol{v}_\lambda + \boldsymbol{v}_\lambda \cdot \nabla_{\boldsymbol{x}} \boldsymbol{v}_\lambda = -\frac{3 g_\infty}{2 D_\infty} (\boldsymbol{v}_\lambda + \nabla_{\boldsymbol{x}} \varphi), \tag{5.A.1}$$

$$\partial_{D_\infty} \delta_\lambda + \nabla_{\boldsymbol{x}} \cdot [(1+\delta_\lambda) \boldsymbol{v}_\lambda] = 0, \tag{5.A.2}$$

for $\lambda \in \{\mathrm{c}, \mathrm{b}\}$ with $\varphi = a\Phi/(B_0 D_\infty)$ and g_∞ defined in (5.3.14). The initial conditions at $D_\infty = 0$ must be $\boldsymbol{v}_\mathrm{c} = \boldsymbol{v}_\mathrm{b} = -\nabla_{\boldsymbol{x}} \varphi$ for (5.A.1) not to diverge. Taking the difference of (5.A.1) for $\lambda = \mathrm{b}$ and $\lambda = \mathrm{c}$ gives

$$\partial_{D_\infty} \boldsymbol{v}_\mathrm{bc} + \boldsymbol{v}_\mathrm{b} \cdot \nabla_{\boldsymbol{x}} \boldsymbol{v}_\mathrm{bc} + \boldsymbol{v}_\mathrm{bc} \cdot \nabla_{\boldsymbol{x}} \boldsymbol{v}_\mathrm{c} = -\frac{3 g_\infty}{2 D_\infty} \boldsymbol{v}_\mathrm{bc}, \tag{5.A.3}$$

where $\boldsymbol{v}_\mathrm{bc} = \boldsymbol{v}_\mathrm{b} - \boldsymbol{v}_\mathrm{c}$. Notice that the neutrino contribution contained in $\nabla_{\boldsymbol{x}} \varphi$ has dropped out. Consequently, we obtain results analogous to the ΛCDM case without massive neutrinos [229]. Expand $\boldsymbol{v}_\lambda = \sum_{m=1}^\infty \boldsymbol{v}_\lambda^{(m)}$ for $\lambda \in \{\mathrm{c}, \mathrm{b}\}$ and $\boldsymbol{v}_\mathrm{bc} = \sum_{m=1}^\infty \boldsymbol{v}_\mathrm{bc}^{(m)}$. At first order, we find

$$\partial_{D_\infty} \boldsymbol{v}_\mathrm{bc}^{(1)} = -\frac{3 g_\infty}{2 D_\infty} \boldsymbol{v}_\mathrm{bc}^{(1)}. \tag{5.A.4}$$

Since g_∞ is strictly positive (see Fig. 5.1), the only non-decaying solution is $\boldsymbol{v}_\mathrm{bc}^{(1)} = 0$. As $\boldsymbol{v}_\mathrm{bc} = 0$ initially, this is the only solution. Suppose that $\boldsymbol{v}_\mathrm{bc}^{(m)} = 0$ for $m = 1, \ldots, n-1$. Then also

$$\partial_{D_\infty} \boldsymbol{v}_\mathrm{bc}^{(n)} = -\frac{3 g_\infty}{2 D_\infty} \boldsymbol{v}_\mathrm{bc}^{(n)}, \tag{5.A.5}$$

with the only solution being $\boldsymbol{v}_\mathrm{bc}^{(n)} = 0$. It follows that $\boldsymbol{v}_\mathrm{bc} = 0$ at all orders. Using this result and taking the mass-weighted average of the component equations yields at all orders:

$$\partial_{D_\infty} \boldsymbol{v}_\mathrm{cb} + \boldsymbol{v}_\mathrm{cb} \cdot \nabla_{\boldsymbol{x}} \boldsymbol{v}_\mathrm{cb} = -\frac{3 g_\infty}{2 D_\infty} (\boldsymbol{v}_\mathrm{cb} + \nabla_{\boldsymbol{x}} \varphi), \tag{5.A.6}$$

$$\partial_{D_\infty} \delta_\mathrm{cb} + \nabla_{\boldsymbol{x}} \cdot [(1+\delta_\mathrm{cb}) \boldsymbol{v}_\mathrm{cb}] = 0. \tag{5.A.7}$$

Converting back to τ-time gives (5.3.3–5.3.4). Letting $\delta_{\mathrm{bc}} = \delta_{\mathrm{b}} - \delta_{\mathrm{c}}$ and taking the difference of (5.A.2) for $\lambda = \mathrm{b}$ and $\lambda = \mathrm{c}$ also gives

$$\partial_{D_\infty} \delta_{\mathrm{bc}} + \nabla_{\boldsymbol{x}} \cdot [\delta_{\mathrm{bc}} \boldsymbol{v}_{\mathrm{cb}}] = 0. \tag{5.A.8}$$

Inserting $\delta_{\mathrm{bc}} = \sum_{m=1}^{\infty} \delta_{\mathrm{bc}}^{(m)}$, we find that $\delta_{\mathrm{bc}}^{(1)} =$ constant at first order, as in the case without neutrinos.

5.B. Analytic solution

We seek a solution to

$$\partial_\tau^2 D + aH\partial_\tau D = \frac{B_0}{a} D. \tag{5.B.1}$$

To express the solution as a function of the scale factor, $a(\tau)$, we switch time variables to $\log a$ and define the new velocity variable, $\tilde{\boldsymbol{u}}_{\mathrm{cb}} = \boldsymbol{u}_{\mathrm{cb}}/(aH)$. Eq. (5.B.1) is then written as

$$\frac{\mathrm{d}^2 D}{\mathrm{d}(\log a)^2} + \left[2 + \frac{\mathrm{d}\log H}{\mathrm{d}\log a}\right] \frac{\mathrm{d}D}{\mathrm{d}\log a} = \frac{B_0}{a^3 H^2} D. \tag{5.B.2}$$

The hypergeometric function ${}_2F_1(c, d, e, z)$ is a solution of the differential equation

$$z(1-z)\frac{\mathrm{d}^2 F}{\mathrm{d}z^2} + [e - (c+d+1)z]\frac{\mathrm{d}F}{\mathrm{d}z} - cdF = 0. \tag{5.B.3}$$

Given the Ansatz $D(a) = a^p \sqrt{1 + \Lambda a^3} F(z)$ with $z = -\Lambda a^3$ and $\Lambda = \Omega_\Lambda/\Omega_{\mathrm{m}}$, we obtain after some algebra

$$\begin{aligned} &(1-z)\frac{\mathrm{d}^2 F}{\mathrm{d}(\log a)^2} + \left[2(p+1)(1-z) - 3z - \frac{3}{2}\right]\frac{\mathrm{d}F}{\mathrm{d}\log a} = \\ &-\left[\left(p^2 + \frac{p}{2} - \frac{3}{2}(1 - f_\nu)\right) - \left(p^2 + 5p + \frac{21}{4}\right) z\right] F. \end{aligned} \tag{5.B.4}$$

To bring this in the form of (5.B.3), we require

$$p = \frac{1}{4}\left(\pm\sqrt{1 + 24(1 - f_\nu)} - 1\right), \tag{5.B.5}$$

where the positive sign picks the growing solution. Using this in (5.B.4), we obtain

$$\begin{gathered}z(1-z)\frac{\mathrm{d}^2F}{\mathrm{d}z^2}+\frac{1}{3}\left[2p+\frac{7}{2}-(2p+8)z\right]\frac{\mathrm{d}F}{\mathrm{d}z}= \\ \frac{1}{9}\left[p^2+5p+\frac{21}{4}\right]F.\end{gathered} \tag{5.B.6}$$

Identifying constants in (5.B.3) and (5.B.6), we derive the desired expression

$$D(a)=a^p\sqrt{1+\Lambda a^3}\,{}_2F_1\left(\frac{2p+7}{6},\frac{2p+3}{6},\frac{4p+7}{6},-\Lambda a^3\right), \tag{5.B.7}$$

with $p=\sqrt{1+24(1-f_\nu)}/4-1/4$.

5.C. Frame lagging

Let $S(\boldsymbol{x})=(\delta_{\mathrm{cb}}*\alpha)(\boldsymbol{x})$. Since S is itself first order, we have up to second order that

$$S(\boldsymbol{x})=S(\boldsymbol{q}+\boldsymbol{\psi})=S(\boldsymbol{q})+\left.\frac{\partial S}{\partial q_i}\right|_{\boldsymbol{q}}\psi_i(\boldsymbol{q}). \tag{5.C.1}$$

Denoting the Fourier transform of $S(\boldsymbol{x})$ as $\mathcal{F}\{S(\boldsymbol{x})\}$, we find that

$$\mathcal{F}\{S(\boldsymbol{x})\}=\mathcal{F}\{S(\boldsymbol{q})\}+\mathcal{F}\left\{\left.\frac{\partial S}{\partial q_i}\right|_{\boldsymbol{q}}\right\}*\mathcal{F}\{\psi_i(\boldsymbol{q})\}. \tag{5.C.2}$$

To be more explicit, we will denote the Fourier transform of $S(\boldsymbol{x})$ by $S^x(\boldsymbol{k})$ and the Fourier transform of $S(\boldsymbol{q})$ by $S^q(\boldsymbol{k})$. The above identity can then be written as

$$S^x(\boldsymbol{k})=S^q(\boldsymbol{k})+\int_{\boldsymbol{k}_1,\boldsymbol{k}_2}i\boldsymbol{k}_1^iS^q(\boldsymbol{k}_1)\psi_i(\boldsymbol{k}_2), \tag{5.C.3}$$

where $\int_{\boldsymbol{k}_1,\boldsymbol{k}_2}=\int\mathrm{d}\boldsymbol{k}_1\mathrm{d}\boldsymbol{k}_2(2\pi)^{-6}\delta^{(3)}(\boldsymbol{k}_1+\boldsymbol{k}_2-\boldsymbol{k})$. Similarly,

$$\delta^q_{\mathrm{cb}}(\boldsymbol{k})=\delta^x_{\mathrm{cb}}(\boldsymbol{k})-\int_{\boldsymbol{k}_1,\boldsymbol{k}_2}i\boldsymbol{k}_1^i\delta^q_{\mathrm{cb}}(\boldsymbol{k}_1)\psi_i(\boldsymbol{k}_2). \tag{5.C.4}$$

Combining the last two equations, we obtain

$$\alpha^x(\boldsymbol{k})\delta^x_{\mathrm{cb}}(\boldsymbol{k})=\alpha^q(\boldsymbol{k})\delta^x_{\mathrm{cb}}(\boldsymbol{k})-F(\boldsymbol{k}), \tag{5.C.5}$$

where we denote the so-called "frame-lagging" terms by

$$F(\boldsymbol{k}) = \int_{\boldsymbol{k}_1,\boldsymbol{k}_2} i\boldsymbol{k}_1^i \left[\alpha^q(\boldsymbol{k}) - \alpha^q(\boldsymbol{k}_1)\right] \delta_{\mathrm{cb}}^q(\boldsymbol{k}_1)\psi_i(\boldsymbol{k}_2). \tag{5.C.6}$$

Now, since $\delta_{\mathrm{cb}}^x = 1/J - 1$, we obtain the result used in Section 5.4.2:

$$[\delta_{\mathrm{cb}} * \alpha](\boldsymbol{x}) = [(1/J - 1) * \alpha](\boldsymbol{q}) - F(\boldsymbol{q}). \tag{5.C.7}$$

We now rewrite the second-order frame-lagging terms using the Monge-Ampère equation, obtaining

$$F^{(2)}(\boldsymbol{k}) = \int_{\boldsymbol{k}_1,\boldsymbol{k}_2} \left[\alpha(\boldsymbol{k}) - \alpha(\boldsymbol{k}_1)\right] \boldsymbol{k}_1^i \boldsymbol{k}_1^j \psi_i^{(1)}(\boldsymbol{k}_2)\psi_j^{(1)}(\boldsymbol{k}_1). \tag{5.C.8}$$

5.D. Terms up to third order

We give explicit expressions up to third order. For $n = 2$, both the cubic term on the right-hand side of (5.4.19) and the quadratic term on the right-hand side of (5.4.20) vanish. Hence, only the quadratic term in (5.4.19) contributes. Using $\epsilon_{ijk}\epsilon_{ipq} = \delta_{jp}\delta_{kq} - \delta_{jq}\delta_{kp}$, we find

$$\nabla \cdot \boldsymbol{\psi}^{(2)} = -\frac{3g_\infty}{4 + 3g_\infty}\frac{1}{2}\left[\psi_{i,i}^{(1)}\psi_{j,j}^{(1)} - \psi_{i,j}^{(1)}\psi_{i,j}^{(1)}\right]. \tag{5.D.1}$$

The corresponding ΛCDM coefficient $(3/7)$ is found by setting $g_\infty = 1$. Dividing these coefficients, one obtains the scale-independent factor $C_2 = 7g_\infty/(4 + 3g_\infty)$. For $n = 3$, we obtain two pieces from (5.4.19) and one piece from (5.4.20), giving $\boldsymbol{\psi}^{(3)} = \boldsymbol{\psi}^{(3a)} + \boldsymbol{\psi}^{(3b)} + \boldsymbol{\psi}^{(3c)}$. Using $\det A_{ij} = (1/6)\epsilon_{ijk}\epsilon_{pqr}A_{ip}A_{jq}A_{kr}$, we can write these as

$$\nabla \cdot \boldsymbol{\psi}^{(3a)} = -\frac{g_\infty}{2 + g_\infty}\det \psi_{i,j}^{(1)}, \tag{5.D.2}$$

$$\nabla \cdot \boldsymbol{\psi}^{(3b)} = -\frac{4 + 6g_\infty}{6 + 3g_\infty}\frac{1}{2}\left[\psi_{i,i}^{(1)}\psi_{j,j}^{(2)} - \psi_{i,j}^{(1)}\psi_{i,j}^{(2)}\right], \tag{5.D.3}$$

$$\nabla \times \boldsymbol{\psi}^{(3c)} = -\frac{1}{3}\nabla\psi_i^{(2)} \times \nabla\psi_i^{(1)}. \tag{5.D.4}$$

The corresponding ΛCDM terms are again found by setting $g_\infty = 1$. Expressing these in terms of potentials (5.2.7-5.2.10) and dividing the corresponding coefficients, we obtain the form given in Section 5.2.2 in terms of C_1, C_2, C_3.

Hot initial conditions

6

This chapter deals with the initial conditions of the hot matter species: neutrinos. The thermal distribution implies that another angle of attack is needed. Following Ma & Bertschinger (1994), we use the method of geodesic integration. Combined with the δf method of Chapter 4, we demonstrate that the phase-space distribution can be accurately described with particles even in the relativistic régime.

Geodesic motion and phase-space evolution of massive neutrinos

ABSTRACT: The non-trivial phase-space distribution of relic neutrinos is responsible for the erasure of primordial density perturbations on small scales, which is one of the main cosmological signatures of neutrino mass. In this paper, we present a new code, FASTDF, for generating 1%-accurate particle realisations of the neutrino phase-space distribution using relativistic perturbation theory. We use the geodesic equation to derive equations of motion for massive particles moving in a weakly perturbed spacetime and integrate particles accordingly. We demonstrate how to combine geodesic-based initial conditions with the δf method to minimise shot noise and clarify the definition of the neutrino momentum, finding that large errors result if the wrong parametrisation is used. Compared to standard Lagrangian methods with ad-hoc thermal motions, FASTDF achieves substantial improvements in accuracy. We outline the approximation schemes used to speed up the code and to ensure symplectic integration that preserves phase-space density. Finally, we discuss implications for neutrino particles in cosmological N-body simulations. In particular, we argue that particle methods can accurately describe the neutrino distribution from $z = 10^9$, when neutrinos are linear and ultra-relativistic, down to $z = 0$, when they are nonlinear and non-relativistic. FASTDF can be used to set up accurate initial conditions (ICs) for N-body simulations and has been integrated into the higher-order IC code MONOFONIC.

6.1. Introduction

It is expected that relic neutrinos of the early Universe outnumber the baryons by a factor of $n_\nu/n_b \approx 10^9$. The discovery of neutrino oscillations [11, 12] implies that at least two-thirds of these particles carry a mass, which though small, through sheer abundance should leave an imprint on the large-scale distribution of matter. Detecting this signature would provide a means of measuring the sum of neutrino masses $\sum m_\nu$ from cosmology [23, 25, 50], complementing an extensive programme of neutrino experiments on Earth. The imprint of massive neutrinos arises primarily from the fact that, during the era of structure formation, neutrinos are non-relativistic particles with a relativistic phase-space distribution. Neutrinos decouple from the primordial plasma at a temperature of $1\,\mathrm{MeV}$ and subsequently stream along geodesics, essentially without scattering, but maintaining a thermal phase-space distribution. After becoming non-relativistic, massive neutrinos have a thermal velocity $v_\text{th} \propto 1/m_\nu$ and cannot be contained effectively in regions smaller than v_th/H, where H is the Hubble rate. As a result, although neutrinos contribute like dust to the geometric expansion of the Universe, they cluster less effectively on small scales, slowing down the growth of matter perturbations. This effect has been used to put tight constraints on the sum of neutrino masses, with current limits of $\sum m_\nu < 0.15\,\mathrm{eV}$ or better [108, 109, 113, 248]. These constraints are an order of magnitude below the strongest laboratory constraint, $m_\nu < 0.8\,\mathrm{eV}$, from KATRIN [14], but come with the important assumption of ΛCDM cosmology, which highlights their complementarity.

Cosmological N-body simulations are widely used to make predictions for nonlinear structure formation in the presence of massive neutrinos and to study their effects on cosmological observables, which is needed to unlock the full potential of surveys like DESI and Euclid for neutrino science. Many approximate methods exist to incorporate neutrino effects in simulations, of which [168, 172, 222, 249, 250] are some recent examples. Methods that solve for the neutrino and dark matter perturbations self-consistently fall roughly into three categories: grid-based methods actively solve evolution equations on the grid [149, 161, 163, 164, 167, 175, 251, 252], linear methods use transfer functions computed with an Einstein-Boltzmann code [166, 171, 253], and particle-based methods sample the phase-space distribution with tracers [148–150, 155, 156, 159, 164, 237, 254–256]. While particle methods are uniquely suited to follow nonlinear neutrino clustering at late times, they typically disagree with linear theory in the neutrino component at early times, in part due to the way that initial conditions are handled and in part due to shot noise. The purpose of this paper is to address these shortcomings and to demonstrate how particle methods can be used to obtain accurate results at all times.

Particle initial conditions for N-body simulations are commonly set up with Lagrangian perturbation theory (LPT). This works very well for baryons and cold dark matter, even in the presence of neutrinos [212, 213, 257]. However, standard methods fail for the neutrino fluid itself. The free-streaming behaviour is usually implemented in an ad-hoc manner by drawing a random thermal velocity from the homogeneous Fermi-Dirac distribution and assigning it to the neutrino particles [254, 255]. This is typically combined with first-order Lagrangian perturbation theory (1LPT), more commonly known as the Zel'dovich approximation [62], in which particle displacements and velocities are proportional to one another: $\mathbf{v} = aHf\boldsymbol{\psi}$, where f is the linear growth rate and a the scale factor. It is easy to see that these steps are inconsistent. The result is illustrated in the top row of Fig. 7.8. Even though the displacement field, $\boldsymbol{\psi}$, can be chosen to reproduce the density field at the initial time, the imprinted density perturbations are wiped out by random motions after only a few steps. A better approach, already proposed by [150] and used recently by [154], is to integrate neutrinos along geodesics from high redshift, $z = 10^9$, down to the starting redshift of the simulation using metric perturbations obtained from an Einstein-Boltzmann code[1]. This, however, does nothing to address the issue of shot noise, which is particularly problematic at early times. We recently proposed the δf method as a way of minimising shot noise in neutrino simulations [237], inspired by similar efforts in plasma physics [179–181] and stellar dynamics [176, 177]. Here, we will show how these methods can be combined to produce accurate density fields from the very beginning of the simulation, as shown in the bottom panels of Fig. 7.8. To facilitate this approach for large simulations, we have made our FASTDF[2] code publicly available, and integrated it into the higher-order initial conditions generator MONOFONIC, along with other neutrino extensions [257, 260, 261].

The remainder of the paper is structured as follows. We will first describe our methods in Section 6.2. We then derive the required equations of motion directly from the geodesic equation in Section 6.3 and briefly remark on the Lagrangian derivation that was used previously. In Section 6.4, we present numerical results, comparing the proposed method with linear fluid calculations and standard methods, and evaluating the impact of the equations of motions. Finally, we discuss the implications for simulations in Section 6.5.

6.2. Methods

Throughout this paper, we work in Newtonian gauge with a metric given by

$$\mathrm{d}s^2 = a^2(\tau)\left[-(1+2\psi(\mathbf{x},\tau))\mathrm{d}\tau^2 + (1-2\phi(\mathbf{x},\tau))\delta_{ij}\mathrm{d}x^i\mathrm{d}x^j\right], \tag{6.2.1}$$

---

[1] Another solution could be to extend LPT to fluids with non-negligible velocity dispersion [258, 259].

[2] Fast Distribution Function; all codes available via `https://willemelbers.com/neutrino_ic_codes/`.

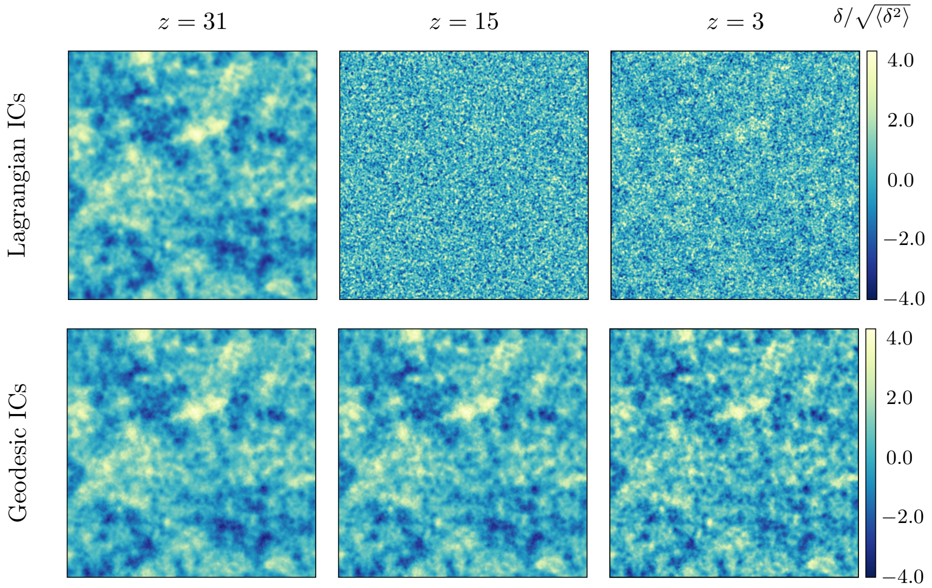

Figure 6.1: Neutrino density slices from a 3.2 Gpc cube for $z \in \{31, 15, 3\}$. In the top row, neutrino particles were set up at $z = 31$ using first-order Lagrangian perturbation theory (1LPT) and then integrated forward. The initial perturbations are immediately washed out and structure is only recovered over time. In the bottom row, the neutrino density field is faithfully reproduced at all times using geodesic integration together with the δf method [237].

where τ is conformal time and we consider only scalar metric perturbations: ϕ and ψ. Let $U_\mu = \mathrm{d}x_\mu/\sqrt{-\mathrm{d}s^2}$ be the 4-velocity and $P_\mu = mU_\mu$ the 4-momentum of a massive neutrino particle. The physical momentum measured by a cosmological observer is

$$p = \sqrt{g_{ij}P^iP^j}. \tag{6.2.2}$$

We define the comoving momentum as $q = ap$ and let $q^i = q_i = q\hat{n}_i$, where the unit vector $\hat{n}_i = P_i/P$ with $P^2 = \delta^{ij}P_iP_j$. Finally, we define the energy as $\epsilon = \sqrt{q^2 + m^2a^2}$. Our aim is to sample particles from the neutrino phase-space distribution,

$$f(\mathbf{x}, \mathbf{q}, \tau) = \bar{f}(q)\left[1 + \Psi(\mathbf{x}, \mathbf{q}, \tau)\right], \tag{6.2.3}$$

where $\bar{f}(q) = (1+\exp(q/T))^{-1}$ is the homogeneous Fermi-Dirac distribution and $T = 1.95\,\mathrm{K}$ the present-day neutrino temperature. In terms of f, the energy density is

$$\begin{aligned} \rho(\mathbf{x},\tau) &= a^{-4}\int \mathrm{d}^3q\,\epsilon\, f(\mathbf{x},\mathbf{q},\tau) && (6.2.4)\\ &= \bar{\rho}(\tau)\left[1+\delta(\mathbf{x},\tau)\right]. && (6.2.5)\end{aligned}$$

The evolution of f is governed by the collisionless Boltzmann equation (2.2.7). At linear order in the metric perturbations, solutions can be found by decomposing Ψ into a Legendre series in Fourier space [126][3]:

$$\Psi(\mathbf{k},\hat{n},q,\tau) = \sum_{\ell=0}^{\infty}(-i)^{\ell}(2\ell+1)\Psi_{\ell}(\mathbf{k},q,\tau)P_{\ell}(\hat{k}\cdot\hat{n}). \tag{6.2.6}$$

The Boltzmann equation (2.2.7) then becomes an infinite tower of equations in Ψ_ℓ, which is usually truncated at some high $\ell_{\max}$ using an algebraic Ansatz (see Section 3.3). We solve this system with CLASS [124, 128]. To obtain very accurate results, we turn off the default neutrino fluid approximation and use $N = 100$ momentum bins and an integration tolerance of 10^{-12}. In terms of Ψ_ℓ, the energy density and flux perturbations can then be written as[4]

$$\begin{aligned} \bar{\rho}(\tau)\delta(\mathbf{k},\tau) &= a^{-4}\int \mathrm{d}^3q\,\epsilon\,\bar{f}(q)\Psi_0(\mathbf{k},q,\tau), && (6.2.7)\\ \left(\bar{\rho}(\tau)+\bar{P}(\tau)\right)\theta(\mathbf{k},\tau) &= a^{-4}\int \mathrm{d}^3q\,qk\,\bar{f}(q)\Psi_1(\mathbf{k},q,\tau), && (6.2.8)\end{aligned}$$

where $\bar{\rho}$ and $\bar{P}$ are the background density and pressure.

6.2.1. Initial conditions

To sample particles from the full perturbed phase-space distribution (6.2.3), taking into account the non-trivial correlations between $\mathbf{x}$ and $\mathbf{q}$, we integrate particles along geodesics from high redshift. We begin shortly after decoupling at $z = 10^9$, when all modes of interest are outside the horizon and the neutrino phase-space distribution can be described in closed form, although in practice a slightly lower redshift would suffice. To recover the correct super-horizon evolution, we account for the initial monopole and dipole temperature perturbations. At early times on super-horizon scales, the first

---

[3] Note that our definition of Ψ_ℓ is slightly different.

[4] Recall that $\int_0^\pi P_\ell(\cos\theta)P_m(\cos\theta)\sin\theta\,\mathrm{d}\theta = \delta_{\ell m}2/(2\ell+1)$.

two moments of the distribution function are $\delta = -2\psi$ and $\theta = \frac{1}{2}k^2\tau\psi$ [126]. From (6.2.7–6.2.8), we find

$$\Psi_0 = -\frac{\delta}{\nu}\frac{\mathrm{d}\log\bar{f}}{\mathrm{d}\log q}, \qquad \Psi_1 = -\frac{\omega\epsilon\theta}{qk\nu}\frac{\mathrm{d}\log\bar{f}}{\mathrm{d}\log q}, \tag{6.2.9}$$

where $\nu \equiv \mathrm{d}\log\bar{\rho}/\mathrm{d}\log T = 4$ and $\omega \equiv 1 + w = (\bar{\rho} + \bar{P})/\bar{\rho} = 4/3$. It follows that

$$f(\mathbf{x}, \mathbf{q}, \tau) = \bar{f}\left(q\left[1 - \frac{\delta}{\nu} - \frac{\omega\epsilon\hat{q}}{\nu q}\cdot\nabla\left(\nabla^{-2}\theta\right)\right]\right). \tag{6.2.10}$$

Particle positions are sampled uniformly in the periodic simulation volume. We then apply the initial perturbations by sampling momenta from the unperturbed Fermi-Dirac distribution, $\bar{f}$, and rescaling the ith component of $\mathbf{q}$:

$$q_i \to q_i\left[1 + \frac{\delta}{\nu} + \frac{\omega\epsilon}{\nu q_i}\nabla_i\left(\nabla^{-2}\theta\right)\right]. \tag{6.2.11}$$

After setting up these “pre-initial” conditions, neutrinos are integrated using relativistic equations motion, derived in the next section. These depend on the scalar potentials, $\phi(\mathbf{x}, \tau)$ and $\psi(\mathbf{x}, \tau)$, whose transfer functions are computed with CLASS. The integration is done with the C-code FASTDF, which we make publicly available. Since the metric is computed in linear theory beforehand, each neutrino is completely independent, in principle allowing the code to be perfectly parallel. However, a large fraction of the computational expense is due to the potential grids, which can be shared if the particles are synchronised. To exploit this, FASTDF supports parallelisation through both OPENMP and MPI. The latter is also used to facilitate parallel data output through HDF5. Further gains in speed are made by realising that the metric perturbations are constant during pure radiation and pure matter domination. We therefore compute the potential fields only when the fractional change in the transfer functions exceeds 1% and linearly interpolate between these super-steps. This significantly reduces the required number of Fourier transforms and has a negligible impact on the accuracy.

6.2.2. The δf method

To handle particle shot noise, which is of particular concern at early times, we use the δf method [237]. This is a variance reduction technique in which the phase-space distribution is decomposed as $f = \bar{f} + \delta f$ in terms of an analytical background $\bar{f}$ and perturbation δf sampled by the particles. The density integral (6.2.4) is then decomposed into a smooth background, $\bar{\rho}(\mathbf{x},\tau)$, and a sum over simulation particles:

$$\rho(\mathbf{x},\tau) \cong \bar{\rho}(\mathbf{x},\tau) + \frac{M}{N}\sum_{k=1}^{N} w_k \epsilon_k \, W(\mathbf{x}-\mathbf{x}_k), \tag{6.2.12}$$

where M is a normalisation factor, $W(\mathbf{x})$ a smoothing kernel, ϵ_k the energy and w_k a statistical weight for particle k given by $w_k = \delta f_k / f_k$. The weights are simple to compute in practice. Conservation of phase-space density along geodesics implies that $f_k = \bar{f}(p_k)$ with p_k the initially sampled (unperturbed) value for particle k at $z = 10^9$. At any later point, we obtain $\delta f_k = \bar{f}(p_k) - \bar{f}(q_k)$. The method similarly extends to other phase-space statistics, such as the momentum density.

6.3. Equations of motion

We will derive the relativistic equations of motion starting directly from the geodesic equation and then comment on the differences with [150, 154].

6.3.1. Geodesic derivation

To derive equations of motion in terms of x^i and q_i, we begin with the geodesic equation $\nabla_P P = 0$. Its components read

$$P^\nu \frac{\mathrm{d}P^i}{\mathrm{d}x^\nu} = -\Gamma^i_{\mu\nu} P^\mu P^\nu. \tag{6.3.1}$$

To first order, the Christoffel symbols $\Gamma^i_{\mu\nu}$ are

$$\begin{aligned} \Gamma^i_{00} &= \partial_i \psi, \\ \Gamma^i_{j0} &= \delta^i_j (aH - \dot{\phi}), \\ \Gamma^i_{jk} &= -2\partial_{(j}\phi\delta_{k)i} + \partial_i\phi\delta_{jk}. \end{aligned} \tag{6.3.2}$$

Furthermore, using $q^2 = a^2 g_{ij} P^i P^j$ and $m^2 = -g_{\mu\nu} P^\mu P^\nu$, we express the momentum components in terms of the energy $\epsilon = \sqrt{q^2 + m^2 a^2}$ and the comoving 3-momentum q_i:

$$\begin{aligned} P^0 &= a^{-2}(1-\psi)\epsilon, \\ P^i &= a^{-2}(1+\phi)q^i. \end{aligned} \tag{6.3.3}$$

The left-hand side of (6.3.1) consists of two terms, the first being

$$P^0 \frac{\mathrm{d}P^i}{\mathrm{d}\tau} = a^{-4}(1-\psi)\epsilon\left(-2aHq^i(1+\phi) + \dot{\phi}q^i + (1+\phi)\frac{\mathrm{d}q^i}{\mathrm{d}\tau}\right), \tag{6.3.4}$$

whereas the second is simply

$$P^j \frac{\mathrm{d}P^i}{\mathrm{d}x^j} = a^{-4}(1+\phi)q^j\partial_j\phi q^i. \tag{6.3.5}$$

The right-hand side of (6.3.1) consists of three terms that can be written as

$$\Gamma^i_{\mu\nu}P^\mu P^\nu = \Gamma^i_{00}a^{-4}(1-2\psi)\epsilon^2 + 2\Gamma^i_{0j}a^{-4}(1-\psi+\phi)q^j\epsilon + \Gamma^i_{jk}a^{-4}(1+2\phi)q^jq^k \tag{6.3.6}$$

$$\begin{aligned} &= \partial_i\psi a^{-4}(1-2\psi)\epsilon^2 + \delta^i_j(2aH-2\dot{\phi})a^{-4}(1-\psi+\phi)q^j\epsilon \\ &\quad + \left(-2\partial_{(j}\phi\delta_{k)i} + \partial_i\phi\delta_{jk}\right)a^{-4}(1+2\phi)q^jq^k. \end{aligned} \tag{6.3.7}$$

Using (6.3.4),(6.3.5),(6.3.7) in the geodesic equation (6.3.1) and dividing by $a^{-4}\epsilon(1+\phi-\psi)$, we finally obtain the acceleration

$$\frac{\mathrm{d}q_i}{\mathrm{d}\tau} = -\epsilon\partial_i\psi - \frac{q^2}{\epsilon}\partial_i\phi + \frac{1}{\epsilon}q_iq^j\partial_j\phi + q_i\dot{\phi}. \tag{6.3.8}$$

From (6.3.3), we also obtain

$$\frac{\mathrm{d}x^i}{\mathrm{d}\tau} = \frac{q^i}{\epsilon}(1+\phi+\psi). \tag{6.3.9}$$

Eqs. (6.3.8) and (6.3.9) are the desired equations of motion. These have a different form from those used previously by [150, 154]. This is due to the choice of independent variables, as will be discussed in the next section.

6.3.2. Lagrangian derivation

The Lagrangian derivation[5] of [150] uses the same metric (6.2.1), while [154] also include vector and tensor perturbations. Rather than working directly with the geodesic equation (6.3.1), they start with the action:

$$S = \int \mathrm{d}\tau L = -m\int\sqrt{-\mathrm{d}s^2}. \tag{6.3.10}$$

---

[5] This use of 'Lagrangian' should not be confused with references to Lagrangian perturbation theory or Lagrangian methods elsewhere.

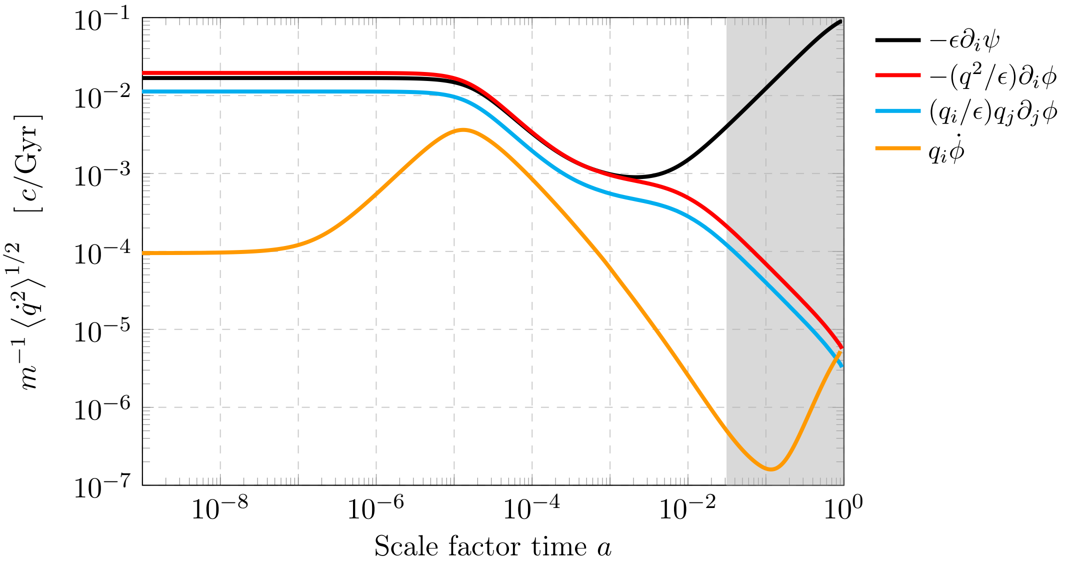

Figure 6.2: Contributions to the particle acceleration over time. The Newtonian acceleration, $-\epsilon\partial_i\psi$, dominates for $z \leq 31$ (shaded), but the relativistic terms are relevant at early times.

Expanding the Lagrangian L to first order in the metric perturbations yields

$$L = -ma\sqrt{1-u^2}\left[1 + \frac{\psi + u^2\phi}{1-u^2}\right], \tag{6.3.11}$$

where $u^i = \mathrm{d}x^i/\mathrm{d}\tau$ and $u^2 = \delta_{ij}u^iu^j$. Observe that the second term inside the square brackets of (6.3.11) diverges in the relativistic limit $u \to 1$, so attention must be paid to the radius of convergence for fast particles. Proceeding from (6.3.11), the conjugate momentum variable to x^i is found by differentiating the Lagrangian with respect to u^i:

$$P_i = \frac{\partial L}{\partial u^i} = \frac{mau_i}{\sqrt{1-u^2}}\left(1 - 2\phi - \frac{\psi + u^2\phi}{1-u^2}\right). \tag{6.3.12}$$

We note that [150] here use the symbol q_i for P_i, but stress that this conjugate momentum variable is in fact related to the comoving 3-momentum q_i by a factor of $q_i/P_i = (1+\phi)$. The Euler-Lagrange equation gives

$$\frac{\mathrm{d}P_i}{\mathrm{d}\tau} = \frac{\partial L}{\partial x^i} = -\epsilon\partial_i\psi - \frac{P^2}{\epsilon}\partial_i\phi, \tag{6.3.13}$$

where $P^2 = P_i P_j \delta^{ij}$ and we used that $\epsilon = ma/\sqrt{1-u^2}$ to zeroth order. Meanwhile, inserting $u^i \propto (1 + f\psi + g\phi)$ into (6.3.12) and solving for f and g gives

$$\frac{\mathrm{d}x^i}{\mathrm{d}\tau} = u^i = \frac{P_j \delta^{ij}}{\epsilon}\left[1 + \psi + \left(2 - \frac{P^2}{\epsilon^2}\right)\phi\right]. \tag{6.3.14}$$

The velocity corrections are small, so let us restrict attention to the acceleration equations (6.3.8) in terms of q and (6.3.13) in terms of P. Both equations contain the usual Newtonian acceleration $-\epsilon\partial_i\psi$ and a post-Newtonian term $-q^2/\epsilon\partial_i\phi$ or $-P^2/\epsilon\partial_i\phi$. However, the geodesic version (6.3.8) has two additional terms: another quadratic term $q_i q^j/\epsilon\partial_j\phi$ and a time-derivative or Sachs-Wolfe term $q_i\dot{\phi}$. These differences can be traced to the use of different momentum variables: the comoving 3-momentum q_i in Section 6.3.1 and the spatial part of the 4-momentum P_i in Section 6.3.2. The two quantities differ by a factor of $(1+\phi)$, which after insertion into (6.3.13) yields the time-derivative term $q_i\dot{\phi}$. The quadratic term, meanwhile, arises in the geodesic derivation from the $P^i P^\mu_{,i}$ term of $\nabla_P P = 0$. This quantity vanishes in the Lagrangian derivation, where the position x^i and its conjugate momentum variable P_i are independent. However, the term is generally non-zero when q_i and x^i are taken as independent instead. The question remains which choice of momentum variable is more suitable for neutrinos in N-body simulations. The advantage of q_i is that it is a physical quantity, eliminating the dependence on metric perturbations when evaluating $\bar{f}(q)$. Since this is a necessary step for neutrino simulations, particularly when using the δf method, we opt for the parametrisation in terms of q.

It is worth asking whether the relativistic corrections are needed in practice. In Fig. 6.2, we show the root mean square of the four terms of (6.3.8) between $z = 10^9$ and $z = 0$, for a 0.1 eV neutrino. As expected, the acceleration is dominated by the Newtonian term (black) at late times. However, the relativistic corrections are non-negligible for $z > 31$. Notably, the quadratic terms (red and blue) are always of the same order of magnitude and one should not be neglected if the other is included. Finally, the time-derivative term (yellow) is negligible during pure radiation or pure matter domination, but becomes relevant outside these régimes. While the relativistic terms are clearly needed for generating initial conditions, they are less relevant for N-body simulations that are started sufficiently late, as will be discussed in Section 6.5.

6.3.3. Symplectic integration

Symplectic integrators explicitly conserve phase-space density[6] and reduce the build-up of errors, which makes them suitable for N-body problems [200, 262]. For FASTDF, we follow the simple strategy proposed in Appendix 4.D and use separable equations of motion that closely approximate the relativistic form, yet admit a straightforward symplectic discretization, but see also Appendix A of [154] for a scheme involving a predictor-corrector step. Concretely, we approximate equations (6.3.8–6.3.9) with:

$$\frac{\mathrm{d}\mathbf{q}}{\mathrm{d}\tau} = -\epsilon_0 \nabla\psi - \frac{q_0^2}{\epsilon_0}\nabla\phi + \frac{1}{\epsilon_0}\mathbf{q}_0\left[\mathbf{q}_0\cdot\nabla\phi\right] + \mathbf{q}_0\dot{\phi}, \tag{6.3.15}$$

$$\frac{\mathrm{d}x^i}{\mathrm{d}\tau} = \frac{q^i}{\epsilon}, \tag{6.3.16}$$

where $\mathbf{q}_0 = \mathbf{q}(z = 10^9)$ and $\epsilon_0 = \sqrt{q_0^2 + m^2a^2}$. Eq. (6.3.15) is a good approximation because $q_0 \ll ma$ whenever q deviates much from q_0: for slow particles at late times, while (6.3.16) neglects the first-order term $|\phi + \psi| \ll 1$. A leapfrog discretization of these equations is

$$\mathbf{q}_{k+\frac{1}{2}} = \mathbf{q}_k + \Delta\tau_k^{k+\frac{1}{2}}\Big[-\epsilon_0\nabla\psi_k - \frac{q_0^2}{\epsilon_0}\nabla\phi_k + \frac{1}{\epsilon_0}\mathbf{q}_0\left[\mathbf{q}_0\cdot\nabla\phi_k\right] + \mathbf{q}_0\dot{\phi}_k\Big], \tag{6.3.17}$$

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \Delta\tau_k^{k+1}\frac{\mathbf{q}_{k+\frac{1}{2}}}{\sqrt{q_{k+\frac{1}{2}}^2 + m^2a^2}}, \tag{6.3.18}$$

$$\mathbf{q}_{k+1} = \mathbf{q}_{k+\frac{1}{2}} + \Delta\tau_{k+\frac{1}{2}}^{k+1}\Big[-\epsilon_0\nabla\psi_{k+1} - \frac{q_0^2}{\epsilon_0}\nabla\phi_{k+1} + \frac{1}{\epsilon_0}\mathbf{q}_0\left[\mathbf{q}_0\cdot\nabla\phi_{k+1}\right] + \mathbf{q}_0\dot{\phi}_{k+1}\Big], \tag{6.3.19}$$

where $\psi_k = \psi(\mathbf{x}_k, a_k)$ and similarly for ϕ. As is common in cosmological simulations, we use a constant step size $\Delta\log a$ and find the corresponding conformal time steps to be

$$\Delta\tau_k^\ell = \int_{\log a_k}^{\log a_\ell}\frac{\mathrm{d}\log a}{aH(a)}. \tag{6.3.20}$$

We observe that $(\partial\mathbf{x}_{k+1}/\partial\mathbf{x}_k)(\partial\mathbf{q}_{k+1}/\partial\mathbf{q}_k) = I_d + (\partial\mathbf{x}_{k+1}/\partial\mathbf{q}_k)(\partial\mathbf{q}_{k+1}/\partial\mathbf{x}_k)$, which ensures sympecticity. To verify the validity of (6.3.15–6.3.16), we also implemented a non-symplectic leapfrog scheme based directly on (6.3.8–6.3.9) and found relative differences in the resulting power spectra of order 10^{-5}, well below other sources of

---

[6] A linear map $J\colon \mathbb{R}^{2d} \to \mathbb{R}^{2d}$ is symplectic if $J^\mathrm{T}\Omega J = \Omega$ for $\Omega = \left(\begin{smallmatrix} 0 & I \\ -I & 0\end{smallmatrix}\right)$, with $I = I_d$ the $d \times d$ identity matrix. A differential map $f\colon U \to \mathbb{R}^{2d}$, with $U \subset \mathbb{R}^{2d}$ open, is symplectic if the Jacobian matrix J of f is everywhere symplectic. Conservation of phase-space density follows from $\det(J) = \det(\Omega) = 1$.

error.

6.4. Results

We set up 800^3 particles in a periodic volume with side length $L = 3.2\,\mathrm{Gpc}$, using (6.2.11) to generate pre-initial conditions at $z = 10^9$. For comparison, particles are also set up with first-order Lagrangian ICs at $z = 31$[7]. We consider two degenerate models with $\sum m_\nu = 0.15\,\mathrm{eV}$ ($f_\nu = 0.11$) and $\sum m_\nu = 0.3\,\mathrm{eV}$ ($f_\nu = 0.023$). Fixed initial conditions are used to facilitate comparison with linear theory on large scales [238] and the δf method is used in each case to suppress shot noise. First, we show the neutrino density power spectrum evaluated at various redshifts in Fig. 6.3. Power spectra are computed by dividing the neutrino ensemble in half and taking the cross-spectrum, which eliminates the constant shot noise plateau on small scales [157]. Note that we compute the power spectrum of the energy density, as expressed in (6.2.12), as opposed to the mass density. The results are compared with the linear fluid calculations from CLASS. We remind the reader that particles were integrated using linear metric perturbations, which should result in perfect agreement with CLASS. We see that this is indeed the case with the geodesic approach, while the power is significantly underestimated for the runs with Lagrangian ICs, recovering only over time. We also show the effect of using the alternative equations (6.3.13–6.3.14), essentially substituting the canonical momentum P for the comoving momentum q without accounting for the relative factor $(1 + \phi)$. In this case, the power spectrum is overestimated. In both cases, the errors are largest at early times, but persist on large scales down to $z = 0$. This can be seen more clearly in Fig. 6.4, where we show the ratios relative to CLASS for $\sum m_\nu = 0.3\,\mathrm{eV}$. Using the geodesic method, we obtain 1%-agreement independent of redshift, while the other methods result in significant errors on all scales. At $z = 0$, a $(-8\%, +5\%)$ error remains at $k = 2 \times 10^{-3}\,\mathrm{Mpc}^{-1}$ when using Lagrangian ICs or when substituting the canonical momentum P for q, respectively. In a full N-body simulation, this lack or excess of neutrino clustering would cause a back-reaction, resulting in still larger errors and contaminating the dark matter and baryon components.

In all cases, the power diminishes relative to CLASS beyond $k = 0.1\,\mathrm{Mpc}^{-1}$. This is due to the limited resolution of the runs. The precision and speed of FASTDF are mainly determined by two parameters: the step size $\Delta \log a$ and the size M of the mesh on which the potentials are calculated. A third parameter, the interpolation order used when computing forces, chosen from $r = 1$ or $r = 2$, has a small effect on the accuracy. We

---

[7] This is the fiducial starting redshift for neutrino ICs in [257]. Usually, $z = 31$ is too late for accurate first-order ICs, but this is not true for neutrinos. Moreover, all calculations are linear in this paper.

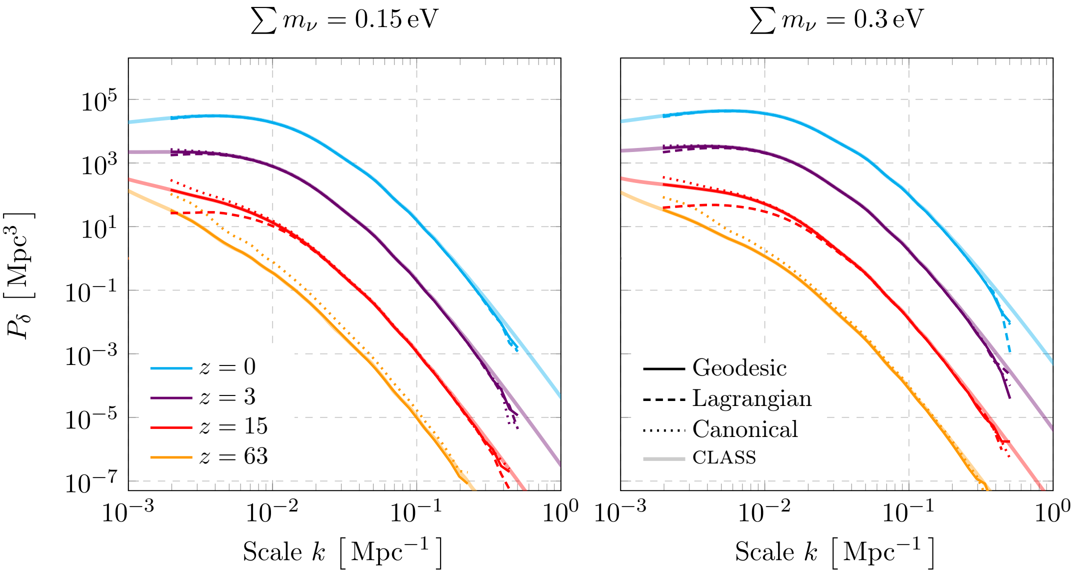

Figure 6.3: The linear neutrino density power spectrum at various redshifts computed from 800^3 particles in an $L = 3.2$ Gpc cube for $\sum m_\nu = 0.15\,\mathrm{eV}$ and $\sum m_\nu = 0.3\,\mathrm{eV}$. Particles were set up with Eq. (6.2.11) at $z = 10^9$ or with Lagrangian ICs at $z = 31$ and subsequently evolved forward using linear metric perturbations. We also show the effect of substituting the canonical momentum P for the comoving momentum q in the Fermi-Dirac function. The spectra are compared with the linear fluid prediction from CLASS. There is no line for the Lagrangian ICs at $z = 63$.

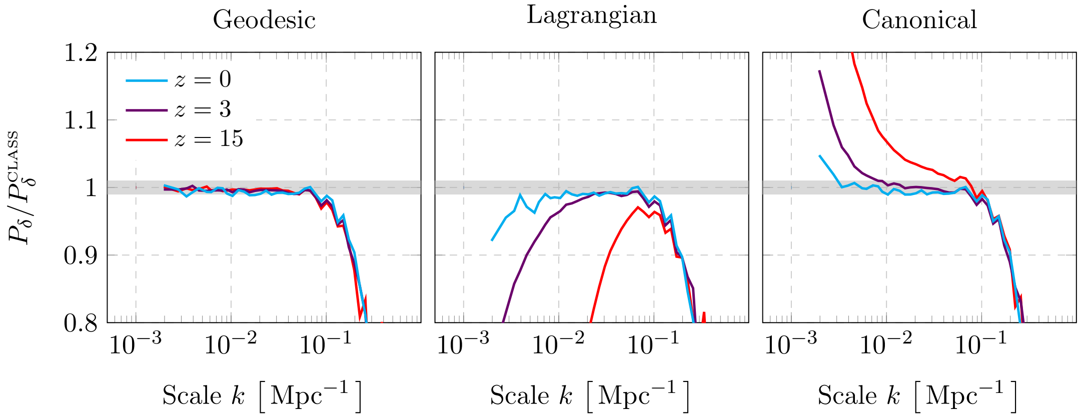

Figure 6.4: Ratios of the linear neutrino density power spectrum at various redshifts computed from 800^3 particles in an $L = 3.2$ Gpc cube for $\sum m_\nu = 0.3\,\mathrm{eV}$, relative to the linear fluid prediction from CLASS. Particles were set up with Eq. (6.2.11) at $z = 10^9$ (left) or with Lagrangian ICs at $z = 31$ (middle) and subsequently evolved forward using linear metric perturbations. We also show the effect of substituting the canonical momentum P for the comoving momentum q in the Fermi-Dirac function (right). The shaded area is 1%.

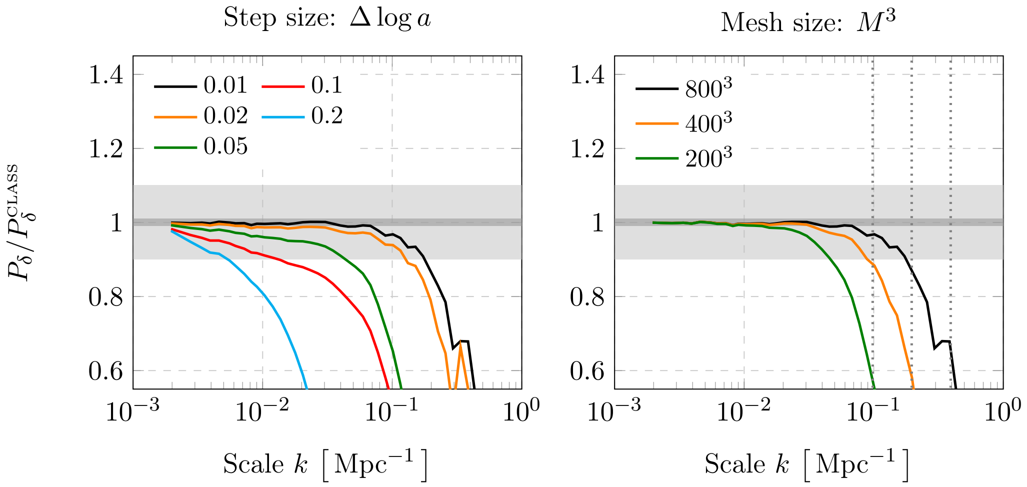

Figure 6.5: Impact of the step size (left) and the mesh size (right) on the neutrino density power spectrum at $z = 31$, computed from 800^3 particles in an $L = 3.2$ Gpc cube for $\sum m_\nu = 0.3$ eV. The spectra are compared with the linear fluid prediction from CLASS. The shaded areas are 1% (dark) and 10% (light). The vertical dotted lines on the right represent half the Nyquist frequency $k_N = \pi M/L$.

show the impact of the first two parameters on the neutrino power spectrum at $z = 31$ for $\sum m_\nu = 0.3\,\text{eV}$ in Fig. 6.5. For the main results in this paper, we used $\Delta \log a = 0.01$ together with $M = 800$, resulting in 1%-agreement with the fluid calculations up to $k = 0.07\,\text{Mpc}^{-1}$. However, errors decrease quickly on small scales in an N-body simulation once neutrinos become non-relativistic, so obtaining agreement on large scales is most important. For many applications, the parameters could therefore be relaxed to enable more rapid realisations of the neutrino distribution function.

To demonstrate that we can also reproduce higher-order moments of the distribution function, we show the power spectrum of the momentum perturbation, $(1 + \delta)\theta$, in Fig. 6.6. Despite the extreme precision settings, a small scatter can be seen at large k for the CLASS results at $z \geq 15$, reflecting the difficulty of solving the Boltzmann hierarchy numerically on small scales. We once again obtain excellent agreement between the geodesic results and CLASS, but find large errors at $z = 63$ when using the canonical momentum, especially for the lighter neutrinos. Large errors are also apparent for the Lagrangian ICs at $z = 15$. In contrast to the density power spectrum, however, these errors decrease quickly on large scales.

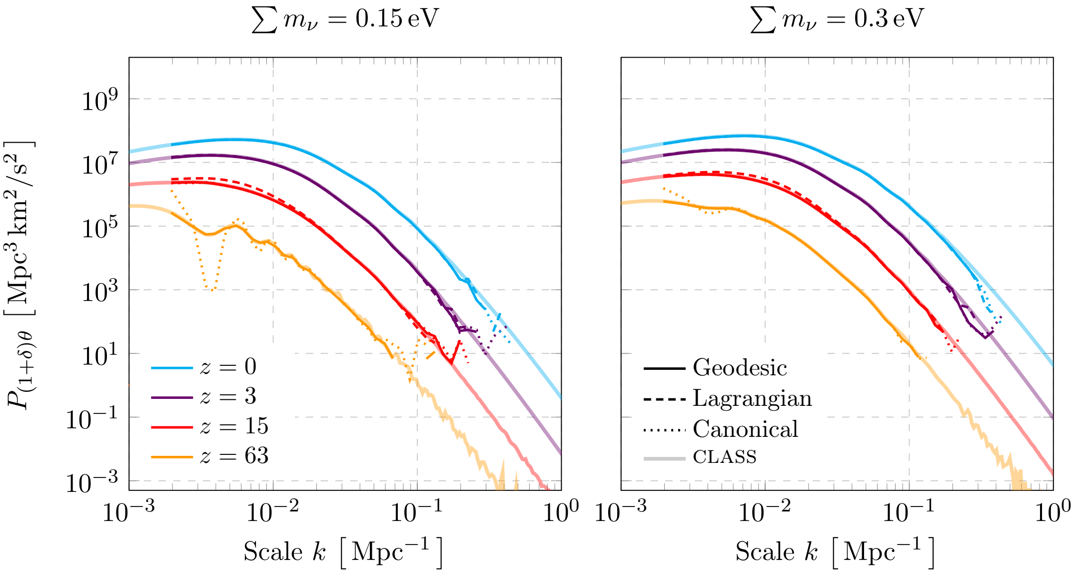

Figure 6.6: The linear neutrino momentum power spectrum at various redshifts computed from 800^3 particles in an $L = 3.2$ Gpc cube for $\sum m_\nu = 0.15\,\mathrm{eV}$ and $\sum m_\nu = 0.3\,\mathrm{eV}$. Particles were evolved in the linearly perturbed spacetime. The spectra are compared with the linear fluid prediction from CLASS, which shows some scatter on small scales at early times. There is no line for the Lagrangian ICs at $z = 63$.

6.5. Discussion

The accurate treatment of massive neutrinos in cosmological N-body simulations, consistent with the demand of surveys like DESI and Euclid for percent-level accurate modelling of large-scale structure observables, also calls for accurate neutrino initial conditions (ICs). In this paper, we have shown that by integrating neutrino particles from high redshift, it is possible to obtain 1%-agreement with linear fluid calculations, even at early times. To achieve this level of agreement, suitable pre-initial conditions must be generated at sufficiently early times, the equations of motion must remain valid in the relativistic limit, and shot noise must be significantly suppressed. We addressed these requirements by providing a closed form expression for the super-horizon perturbations of the Fermi-Dirac distribution $f(\mathbf{x}, \mathbf{q}, \tau)$, by expressing the geodesic equation in terms of $\mathbf{q}$, and by using the δf method to limit shot noise. We also used fixed ICs [238] to limit cosmic variance, which allowed a detailed comparison between linear particle and fluid methods.

When these conditions are not satisfied, significant errors in the neutrino component occur on large scales. For neutrino particles used in N-body simulations, this error causes a back-reaction in the dark matter and baryon components. Simulations that use Lagrangian or unperturbed ICs together with an ad-hoc momentum sampled from the homogeneous Fermi-Dirac distribution underestimate the clustering of neutrinos, leading to errors of a few percent on large scales at $z = 0$. These errors get progressively worse at higher redshifts. Neutrino clustering recovers over time, beginning on small scales where errors are less apparent. Some simulations use hybrid methods (e.g. [154, 173, 174]), transitioning from a linear or grid-based method at early times to a particle method at late times. This would mitigate the back-reaction arising from these errors. Nevertheless, we have demonstrated that a transition of this sort is not necessary if suitable ICs are used and shot noise is addressed.

These results have further implications for neutrino particles in N-body simulations. Aside from the ICs, some codes also use relativistic equations of motion for the neutrino particles in the simulation itself. For ordinary Newtonian simulations, [154] proposed using special relativistic equations of motion with Newtonian gravity. These can be obtained from (6.3.13–6.3.14) by assuming that $|\phi| \ll 1$ and $\phi \approx \psi$:

$$\frac{\mathrm{d}x^i}{\mathrm{d}\tau} = \frac{P^i}{\sqrt{P^2 + m^2a^2}}, \tag{6.5.1}$$

$$\frac{\mathrm{d}P_i}{\mathrm{d}\tau} = -\frac{2P^2 + m^2a^2}{\sqrt{P^2 + m^2a^2}}\partial_i\psi. \tag{6.5.2}$$

Based on (6.3.8–6.3.9) and Fig. 6.2, we instead propose the simpler form

$$\frac{\mathrm{d}x^i}{\mathrm{d}\tau} = \frac{q^i}{\sqrt{q^2 + m^2a^2}}, \tag{6.5.3}$$

$$\frac{\mathrm{d}q_i}{\mathrm{d}\tau} = -\sqrt{q^2 + m^2a^2}\,\partial_i\psi. \tag{6.5.4}$$

The velocities have the same form: including the Lorentz factor is crucial for sub-light neutrino speeds and physical free-streaming lengths. However, the accelerations are different due to the choice of momentum variable. By expressing the equations in terms of the physical quantity q, the potential ϕ need not be evaluated when computing the δf weights. Even so, the corrections to the acceleration matter less in the time frame where Newtonian simulations are used to best effect ($z \ll 10^2$). Simply using the Newtonian acceleration, $\dot{q}_i = -m\partial_i\psi$, together with the special relativistic velocity equation therefore seems to be a reasonable alternative with the benefit of having a straightforward symplectic discretization. Let us remark finally on the choice of gauge. While Newtonian gauge is convenient for the geodesic integration, recent years have seen the introduction of gauges more naturally suited for cosmological N-body simulations.

A popular choice is N-body gauge [185, 188], in which the spatial metric perturbation is constant and traceless, such that the relativistic dark matter density coincides with that of the Newtonian simulation. Using CLASS, it is possible to compute the shifts in density, $\Delta\delta(k)$, and energy flux, $\Delta\theta(k)$, from Newtonian to N-body gauge. Provided that the perturbations are small, the gauge transformation can then be applied to the neutrino ensemble in the same way as the pre-initial conditions, via (6.2.11), since the higher-order moments Ψ_ℓ are gauge-invariant. This feature is available in FASTDF.

The main application of the described method is to set up accurate and consistent neutrino particle initial conditions for simulations. Another interesting application would be to integrate particles back along the line of sight from Earth to analyse the angular dependence of the local neutrino flux. Sampling the full phase-space distribution with particles may be advantageous if, for instance, non-trivial selections are of interest (e.g. neutrinos with momenta in a given interval that passed through halos in a particular mass range). If the metric perturbations are treated in linear theory, as in this paper, the method could provide a cross-check of linear calculations [263, 264], while transitioning from an N-body simulation at late times would enable a fully nonlinear calculation. Another interesting extension would be to consider other massive thermal relics [124, 265].

Part III

Prediction

Local neutrino background

7

This chapter presents a Bayesian forecast of the imprint of the observed large-scale structure on the local neutrino background. It discusses the expected overdensity and bulk velocity of relic neutrinos in the Milky Way and deals with the distribution of angular anisotropies in the neutrino background.

Where shadows lie: reconstruction of anisotropies in the neutrino sky

Abstract: The Cosmic Neutrino Background (CNB) encodes a wealth of information, but has not yet been observed directly. To determine the prospects of detection and to study its information content, we reconstruct the phase-space distribution of local relic neutrinos from the three-dimensional distribution of matter within $200h^{-1}$ Mpc of the Milky Way. Our analysis relies on constrained realization simulations and forward modelling of the `2M++` galaxy catalogue. We find that the angular distribution of neutrinos is anti-correlated with the projected matter density, due to the capture and deflection of neutrinos by massive structures along the line of sight. Of relevance to tritium capture experiments, we find that the gravitational clustering effect of the large-scale structure on the local number density of neutrinos is more important than that of the Milky Way for neutrino masses less than 0.1 eV. Nevertheless, we predict that the density of relic neutrinos is close to the cosmic average, with a suppression or enhancement over the mean of (−0.3%, +7%, +27%) for masses of (0.01, 0.05, 0.1) eV. This implies no more than a marginal increase in the event rate for tritium capture experiments like PTOLEMY. We also predict that the CNB and CMB rest frames coincide for 0.01 eV neutrinos, but that neutrino velocities are significantly perturbed for masses larger than 0.05 eV. Regardless of mass, we find that the angle between the neutrino dipole and the ecliptic plane is small, implying a near-maximal annual modulation in the bulk velocity.

7.1. Introduction

Precise measurements of a near-perfect black-body energy spectrum and of a power-law spectrum of temperature fluctuations in the Cosmic Microwave Background (CMB) reveal detailed information about the state of the Universe at the time of decoupling around $t = 10^5\,\mathrm{yrs}$ [55, 266, 267]. There is strong but indirect evidence for another Big Bang fossil in the form of $N_\mathrm{eff} = 2.99^{+0.34}_{-0.33}$ species of fermionic particles that were relativistic when the radiation decoupled [112]. This is consistent with the prediction of $N_\mathrm{eff} = 3.045$ for the Cosmic Neutrino Background (CNB), consisting of three species that decoupled far earlier, at only $t = 1\,\mathrm{s}$ [122, 123, 268]. That these particles are indeed neutrinos could be confirmed if they were found to be non-relativistic today, given the standard prediction for the present-day neutrino temperature, $T_\nu = 1.68 \times 10^{-4}\,\mathrm{eV}$, and the minimum mass, $m_\nu \gtrsim 0.05\,\mathrm{eV}$, required by neutrino oscillations for the most massive species [11, 12, 29]. Although detecting the indirect cosmological effects of massive neutrinos is challenging, this target could soon be in reach, as suggested by improved constraints on the cosmic neutrino mass fraction [108, 109, 113, 248].

Direct detection of relic neutrinos will be more challenging still and is likely beyond our immediate capabilities. The Karlsruhe Tritium Neutrino Experiment (KATRIN) recently placed an upper bound of 9.7×10^{10} on the local neutrino overdensity relative to the cosmic mean [269], far greater than the density predicted in this paper and elsewhere. An experiment specifically designed for CNB detection has been proposed by the PTOLEMY collaboration [270–272]. Like KATRIN, the PTOLEMY proposal aims to capture neutrinos through the inverse β-decay of tritium [268, 273], but with targets bound to a graphene substrate to enable a larger target mass, which has its own challenges [272, 274]. Other detection proposals rely on the net momentum transfer from the neutrino wind to macroscopic test masses [275–278], absorption features in the cosmic ray spectrum [279, 280], blocking of neutrino emission from de-exciting atoms due to the Pauli exclusion principle [281] or the capture of neutrinos on high-energy ion beams [282]. We refer to [283] for a detailed review of the subject.

Like the CMB, the neutrino background carries both primordial or primary perturbations and secondary gravitational perturbations imprinted by the large-scale structure at late times [263, 264, 284–286]. Since neutrinos are massive particles, secondary perturbations are more significant and depend on the neutrino mass and momentum, giving the background additional structure compared to the CMB. In some cases, gravitational effects may lead to slight modifications of the expected signal and in others they open up entirely new ways of testing neutrino physics. For tritium capture experiments like PTOLEMY, the expected event rate is proportional to the local number density of neutrinos [271], given by the monopole moment of the phase-space distribution. If the tritium targets are polarized, PTOLEMY could measure the angular power spectrum by exploiting the dependence of the event rate on the angle between the polarization and neutrino momentum axes [287]. Some proposals depend on the velocity of neutrinos in the lab frame [275–278, 283], while the orientation of the dipole is important for methods that rely on periodic or angular modulation of the capture rate [287–289]. Pauli blocking could in principle probe the momentum distribution [281, 283]. Additionally, gravitational perturbations may change the flavour [290] and helicity [291–294] makeup of the neutrino background, affecting the ability of experiments like PTOLEMY to distinguish between Dirac and Majorana neutrinos.

To determine the prospects of current and future CNB detection proposals, we therefore need to model the phase-space distribution of relic neutrinos, including its higher-order directional perturbations. Previous studies have looked at the gravitational enhancement of the monopole moment due to the Milky Way [295–299] and nearby Andromeda and Virgo [298]. A very recent study also considered the gravitational influence of dark matter structures in a random $(25\,\mathrm{Mpc})^3$ region on the neutrino phase-space distribution [300]. Here, we expand on these works in several ways. First and foremost, we model the full six-dimensional phase-space distribution of relic neutrinos, taking into account perturbations imprinted on the neutrinos before they entered our galactic neighbourhood. Second, we use self-consistent cosmological simulations to accurately model the time evolution of the large-scale structure and the neutrino background. Third, we use an accurate nonlinear treatment of massive neutrinos [237], which includes the gravitational effects of the neutrinos themselves. Fourth, we model the large-scale distribution of matter within $200h^{-1}\,\mathrm{Mpc}$[1] over the full sky, using observations from the `2M++` galaxy redshift catalogue [301]. Fifth, we use a more recent estimate of the Milky Way mass from [302], which is significantly lower than the value used in previous studies, depressing the effect of the Milky Way.

Using our constrained phase-space simulations, we compute the expected density, velocity, and direction of relic neutrinos, as well as expected event rates for PTOLEMY. We also study the distribution of angular anisotropies, finding that local neutrino density perturbations are anti-correlated with the projected matter distribution, due to the capture and deflection of neutrinos by massive objects along the line of sight. The paper is organized as follows. We describe our simulation and calibration methods in Section 7.2. Our main results are presented in Section 7.3. We finally conclude in Section 7.4.

---

[1] In this expression, h is defined in terms of Hubble's constant as $h \equiv H_0/(100\,\mathrm{km/s/Mpc})$.

7.2. Methods

We now describe our simulation and analysis methods, starting with the details of the constrained simulations in Section 7.2.1, our calibration procedure for applying `2M++` constraints to different neutrino cosmologies in Section 7.2.2, and our treatment of nonlinear neutrino perturbations in Section 7.2.3.

7.2.1. Constrained simulations

Our analysis is based on constrained ΛCDM simulations of the local Universe. Whereas most cosmological simulations start from random initial conditions and only reproduce observations in a statistical sense, constrained simulations employ specialized initial conditions that give rise to an *in silico* facsimile of the observed large-scale structure. Within the precision of the constraints, objects appear in the right relative positions and with the right dimensions, enabling a one-to-one comparison with observations. The past few years have seen constrained simulations being used for a wide range of applications and employing a variety of methods to set up the initial conditions [303–308]. In this paper, we use a Bayesian forward modelling approach known as 'Bayesian Origin Reconstruction from Galaxies' (BORG) [309–311]. This approach uses a Hamiltonian Monte Carlo algorithm to draw samples from the posterior distribution of initial conditions, given a likelihood function that connects initial conditions with observations and a Gaussian prior. The forward model consists of a Comoving Lagrangian Acceleration (COLA) code [90] that approximates the process of structure formation in the ΛCDM paradigm and a nonlinear bias model that connects the final dark matter density field to observed galaxy positions. The Hamiltonian Monte Carlo algorithm is used to efficiently sample a high-dimensional parameter space, consisting of a grid of 256^3 initial phases, multiple bias parameters, and the observer velocity in the CMB frame.

The constraints used in this paper are based on galaxies from the `2M++` catalogue [301]. This is a catalogue of galaxy positions and redshifts, compiled from the 2MASS, 6dF, and SDSS redshift surveys, that covers the full sky out to a distance of $200h^{-1}$ Mpc. Previous simulations with initial conditions based on forward modelling of `2M++` galaxies include the `CSiBORG` suite [312–315] and the SIBELIUS-DARK simulation [307]. We refer the reader to [310, 311] for further details on the BORG analysis of the `2M++` catalogue. This analysis provides not only an accurate reconstruction of the three-dimensional density field in the local Universe, but also reproduces the masses of nearby clusters, with the notable exception of the Perseus-Pisces cluster for which the mass is biased low [311]. This is most likely due to a systematic error in the analysis, but could perhaps also indicate an observational issue [311]. Interestingly, the SIBELIUS-DARK simulation [307], which is based on a similar but older BORG reconstruction, found its most massive dark matter halo at the location of Perseus. However, SIBELIUS-DARK was less accurate in other respects, such as the motion of the Local Group, which is important for our purposes here. Our work is based on nine draws from an earlier version of the chain described in [311], which used ten COLA steps instead of twenty, but was identical in every other respect. We therefore expect the results to be broadly consistent. After discarding an initial burn-in portion, we selected every 432nd draw from the chain to minimize the serial correlation between consecutive draws. This sample of initial conditions allows us to estimate both the expected signal and the uncertainty in our predictions. To demonstrate the effectiveness of the constraints, we show slices of the dark matter and neutrino densities in a portion of the sky in Fig. 7.1, overlaid with `2M++` galaxies (white dots). All prominent structures present in the catalogue are reproduced by the simulations, revealing the underlying dark matter filaments and surrounding neutrino clouds.

Our simulations assume periodic boundary conditions in a $(1\,\mathrm{Gpc})^3$ cube, with the observer located at the centre. The `2M++` constraints mostly cover a central sphere of radius 200 Mpc and gradually taper off beyond that. This means that sufficiently far away from the centre, the initial conditions revert to purely random fluctuations. Given that the phases are provided in the form of 256^3 grids, the constraints only cover 4 Mpc scales and larger. Fluctuations on smaller scales are unconstrained and purely random. Dark matter initial conditions are generated with 3LPT at $z = 31$, using a modified version of MONOFONIC that adds corrections from massive neutrinos [257, 260, 261], while the neutrinos themselves are generated with FASTDF, using linear geodesic integration [186]. The transfer functions are computed with CLASS [124, 128].

The simulations were carried out with a version of GADGET-4 [77] that was modified to be bitwise reversible (see Appendix 7.A) and to add support for massive neutrinos and radiation. We use a 3rd-order Tree-PM algorithm for the gravity calculation. Neutrinos are followed with the δf method to minimize shot noise, boosting the effective particle number without neglecting their nonlinear evolution [237]. We use $N_{\mathrm{cb}} = 384^3$ dark matter and baryon particles[2] and $N_\nu = 384^3$ massive neutrino particles. In order to increase the sampling density of neutrinos locally, upon completion of a simulation, we isotropically inject an additional $N = 224^3 \sim 10^7$ 'spectator' neutrinos at the observer location and run the simulations backwards, allowing us to trace the neutrinos back in time through the evolving large-scale structure (see Section 7.2.3). To ensure that the

---

[2] We will treat cold dark matter and baryons as a single cold fluid and refer to it as dark matter on occasion.

accelerations are identical in the forwards and backwards directions, spectator neutrinos contribute no forces.

A final consideration is that Milky Way-sized perturbations have a characteristic length that is much smaller than 4 Mpc. Hence, our constraints are not sufficient to guarantee the formation of a Milky Way at the centre. Since we expect the Milky Way (MW) to have a considerable effect on the neutrino background, we run two backwards versions of each simulation. Initially, neutrinos are only traced back through the large-scale structure without accounting for MW effects. In the second version, we additionally apply forces from the MW dark matter halo. Following [302], we model the MW halo as an NFW profile [316] with a mass of $M_{200} = 0.82 \times 10^{12} M_\odot$ and a concentration of $c_{200} = 13.31$.[3] For computational simplicity, we use the uncontracted version of the model, since both versions fit the data nearly equally well. We place the centre of the NFW potential at a distance of 8 kpc from the centre of the simulation in the direction of Sag-A*. We also include the motion of the galactic centre in the CMB rest frame of the simulation, by letting the centre of the NFW potential move at a constant speed of 567 km/s in the direction of galactic coordinates $(l, b) = (267°, 29°)$ [55, 317, 318]. In Section 7.3.1, we additionally correct for the motion of the Sun relative to the CMB, $v_\odot = 369.8$ km/s towards $(l, b) = (264°, 48.3°)$ [55], which is otherwise unresolved by the simulations. Crucially, we note that we use a more recent and considerably smaller estimate of the MW mass than that used in previous related works [296, 298]. We therefore expect to find a smaller effect from the MW. Since we are mainly interested in the imprint of the large-scale structure, we do not include the various gaseous and stellar components of the MW, which are altogether less important than the dark matter halo itself.

7.2.2. Model selection

To derive constrained initial conditions with BORG, we have to assume a particular cosmological model. The constraints used in this paper were derived assuming a flat ΛCDM model with parameters $(\Omega_{\rm cdm}, \Omega_{\rm b}, h, A_s, n_s, \sum m_\nu) = (0.2621, 0.04897, 0.6766, 2.105 \times 10^{-9}, 0.9665, 0)$. Despite the fact that this model does not include massive neutrinos, we wish to run constrained simulations for different neutrino masses, without running an expensive MCMC analysis for each case. Doing this requires modifying the cosmological model slightly without altering the clustering on small scales, since otherwise the same phase information would give rise to structures that differ somewhat from the observations. We therefore take the following approach. When increasing $\sum m_\nu$, we

---

[3] Here, M_{200} is the mass contained in a spherical region of radius R_{200} with a density equal to 200 times the critical density and $c_{200} = R_{200}/R_s$, with R_s the scale radius of the NFW profile.

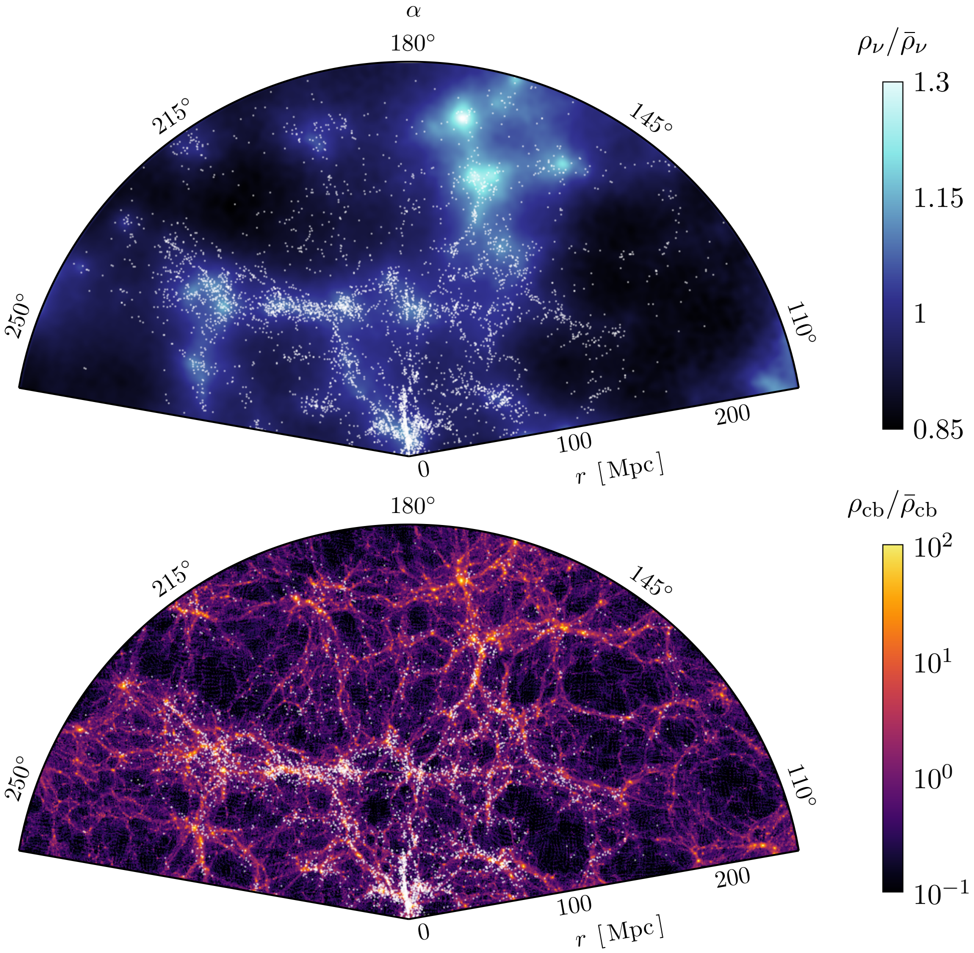

Figure 7.1: Slices of the expected neutrino (top) and dark matter (bottom) densities with right ascension $100° \leq \alpha \leq 260°$ within $r \leq 250\,\text{Mpc}$, assuming $\sum m_\nu = 0.06\,\text{eV}$. The white dots are galaxies from the `2M++` catalogue. From Earth, one would see a deficit in neutrino flux along lines of sight that intersect massive structures, due to the trapping of neutrinos in the surrounding neutrino clouds.

decrease $\Omega_{\rm cdm}$ such that $\Omega_{\rm m} = \Omega_{\rm cdm} + \Omega_{\rm b} + \Omega_\nu$ is fixed. In addition, we modify the primordial scalar amplitude A_s, such that the nonlinear power spectrum at $z = 0$ is fixed at the nonlinear scale $k_{\rm nl} = 1\,\text{Mpc}^{-1}$. Note that $P_{\rm cb}$, the power spectrum of cold dark matter and baryons, is the relevant power spectrum, given that halos are primarily biased with respect to the cold matter, as opposed to the total matter density [244, 319, 320]. To achieve this in practice, we perform a small number of calibration runs and iteratively select values of A_s that satisfy this condition.

Table 7.1: Cosmological parameters for our six neutrino models, which have been calibrated such that $\Omega_\nu + \Omega_{\rm cdm}$ and $P_{\rm cb}(k_{\rm nl})$ with $k_{\rm nl} = 1\,{\rm Mpc}^{-1}$ are fixed.

$\sum m_\nu$	Mass m_ν	N_ν	Ω_ν	$\Omega_{\rm cdm}$	A_s
0.01 eV	0.01 eV	1	2.353×10^{-4}	0.26189	2.107×10^{-9}
0.06 eV	0.06 eV	1	1.407×10^{-3}	0.26072	2.156×10^{-9}
0.15 eV	0.05 eV	3	3.518×10^{-3}	0.25861	2.243×10^{-9}
0.30 eV	0.10 eV	3	7.035×10^{-3}	0.25509	2.429×10^{-9}
0.45 eV	0.15 eV	3	1.055×10^{-2}	0.25157	2.641×10^{-9}
0.60 eV	0.20 eV	3	1.407×10^{-2}	0.24805	2.878×10^{-9}

As noted before, the `2M++` data mostly constrain scales larger than 4 Mpc within 200 Mpc of the observer. As shown in Fig. 7.2, this leaves enough flexibility on large scales to accommodate neutrino masses up to $\sum m_\nu \sim 0.6\,{\rm eV}$.[4] To see this, note that the left-hand panel shows total matter power spectra, $P_{\rm m}(k)$, for nine realizations assuming ΛCDM without massive neutrinos. Although the power spectrum is well-constrained on small scales, there is considerable variance on large scales ($k \lesssim 0.03\,{\rm Mpc}^{-1}$). The right-hand panel shows the power spectrum of dark matter and baryons, $P_{\rm cb}(k)$, for the calibrated models with different neutrino masses, relative to the massless case. For the largest mass considered, $\sum m_\nu = 0.6\,{\rm eV}$, the ratio is still within 1σ of the average. We also checked that the cross-correlation coefficients of the final density fields are within 1% for $k \leq k_{\rm nl}$ and $\sum m_\nu \leq 0.3\,{\rm eV}$ and within a few percent for $\sum m_\nu \leq 0.6\,{\rm eV}$, indicating that the phase information is the same on large scales. Finally, we performed a visual inspection to confirm that we recover the same large-scale structure for all neutrino masses. Hence, the outcome of this procedure is a plausible cosmological model with massive neutrinos that reproduces the `2M++` observations.

Although the resulting power spectra are compatible with the `2M++` constraints at the 1σ-level, one may wonder whether the 20% − 30% differences seen for $\sum m_\nu = 0.6\,{\rm eV}$ on the largest scales could still affect the results. We expect the impact of this offset to be small, because the distance travelled by neutrinos is inversely proportional to the mass, such that heavier neutrinos are less sensitive to large-scale density perturbations. Therefore, matching only the small-scale power spectrum for $\sum m_\nu = 0.6\,{\rm eV}$ is likely justified.

Using the above procedure, we calibrate six models with different neutrino masses: four

---

[4] We note that this breaks the agreement with CMB observations, which primarily constrain large scales. This is simply another way of stating that the combination of CMB and LSS data can rule out large neutrino masses in $\nu\Lambda$CDM, although we make no attempt to do this here.

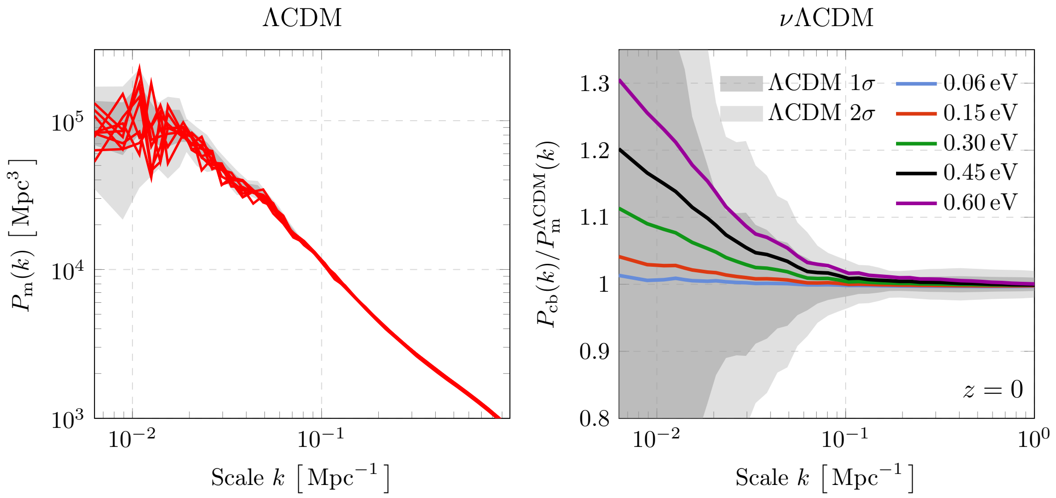

Figure 7.2: (Left) The red lines are nine nonlinear matter matter power spectra, $P_{\mathrm{m}}(k)$, drawn from the posterior distribution of the `2M++` reconstruction, assuming ΛCDM with massless neutrinos at $z = 0$. The shaded areas represent the 1σ and 2σ deviations from the mean. The spectra are well-constrained for $0.03\,\mathrm{Mpc}^{-1} \lesssim k \lesssim 1\,\mathrm{Mpc}^{-1}$, but the variance is considerable on large scales. (Right) Ratios of the nonlinear cold matter power spectrum, $P_{\mathrm{cb}}(k)$, for different neutrino masses relative to the massless ΛCDM case, calibrated to match the constraints in the small-scale limit. The shaded areas represent the 1σ and 2σ constraints.

models with three degenerate neutrino species, $\sum m_\nu \in \{0.15, 0.3, 0.45, 0.6\}$ eV[5], and two models with one species, $\sum m_\nu \in \{0.01, 0.06\}$ eV. The relevant model parameters are given in Table 7.1. Although not strictly allowed by oscillation data, the first four models assume a degenerate neutrino mass spectrum, neglecting the mass-squared differences $|\Delta m^2_{31}| = 2.5 \times 10^{-3}\,\mathrm{eV}^2$ and $\Delta m^2_{21} = 7.4 \times 10^{-5}\,\mathrm{eV}^2$ [29]. Of course, the last two models are also not allowed. The penultimate case is included to examine the behaviour of very light neutrinos. The last model is included as it approximates the cosmological effects of the minimal neutrino mass case under the normal mass ordering. In each case, the intent is only to recover the correct cosmological evolution for a given neutrino mass, m_ν, and for this purpose, the mass splittings have a negligible effect [106].

7.2.3. Neutrino treatment

Let us now discuss our treatment of neutrino perturbations. The evolution of the phase-space distribution, $f(\mathbf{x}, \mathbf{q}, \tau)$, is governed by the collisionless Boltzmann equation (2.2.7).

---

[5] Hence, the individual neutrinos have masses $m_\nu \in \{0.05, 0.1, 0.15, 0.2\}$ eV.

We solve this equation by generating particles from a sampling distribution $g(\mathbf{x},\mathbf{q})$ and tracing their evolution through the constrained volume using the relativistic equations of motion (6.5.3) and (6.5.4) from Chapter 6.

The sampling distribution g need not be the same as the physical distribution f and can be chosen arbitrarily, subject to being normalized and the set $\{g = 0 \wedge f \neq 0\}$ having measure zero. One of the main advantages of a particle-based approach is that we can follow particles into the nonlinear régime, which is particularly important for our purposes here, since we wish to describe neutrinos perturbed by the Milky Way halo. The main downside is the introduction of sampling noise, which we overcome with the δf method of Chapter 4. In this method, the phase-space distribution is decomposed as $f = \bar{f} + \delta f$ into an analytical background part $\bar{f}$ and a perturbation δf sampled by the particles. A noise-suppressed estimator of some phase-space quantity $A(\mathbf{x},\tau)$ is then given by (4.2.7). We can similarly define angular statistics. For example, the density of neutrinos at $\mathbf{x}$ with momenta oriented along the unit vector $\hat{n}$ is

$$n_\nu(\mathbf{x},\hat{n},\tau) = \int \mathrm{d}^3q\,\left[\bar{f}(\mathbf{x},\mathbf{q},\tau) + \delta f(\mathbf{x},\mathbf{q},\tau)\right]\delta^{(2)}(\mathbf{q}/q - \hat{n}) \tag{7.2.1}$$

$$\cong \frac{\bar{n}_\nu(\tau)}{4\pi} + \sum_{k=1}^{N}\frac{\delta f(\mathbf{x}_k,\mathbf{q}_k,\tau)}{g(\mathbf{x}_k,\mathbf{q}_k)}\delta^{(2)}(\mathbf{q}_k/q_k - \hat{n})\delta^{(3)}(\mathbf{x}-\mathbf{x}_k), \tag{7.2.2}$$

where $\bar{n}_\nu(\tau)$ is the mean number density and where $\delta^{(2)}(\hat{x}-\hat{y}) = \delta(\cos\theta - \cos\theta')\delta(\phi-\phi')$. Throughout, we use a standard Fermi-Dirac distribution, $\bar{f}(q) = (1+\exp(q/k_bT_\nu))^{-1}$, for the background model and we set $g = f$ when generating the initial conditions.

This approach is sufficient for describing the neutrino distribution on large scales, as illustrated in Fig. 7.1 for 0.06 eV neutrinos. However, given the $(1\,\mathrm{Gpc})^3$ ambient volume of our simulations, there is a more efficient way to estimate the properties of neutrinos incident on Earth. For this, we inject 'spectator' neutrinos at the location of Earth and run our simulations backwards. For these neutrinos, we adopt an isotropic Fermi-Dirac sampling distribution g. We then apply our δf logic in reverse: given the known sampling density g and the background density $\bar{f}(q)$ with the momentum q from the final $(z = 31)$ snapshot of the backwards simulation, we obtain the statistical weight $w = (\bar{f} - g)/g$. We again estimate phase-space statistics using Eq. (4.2.7). Note that in this case, the assumed sampling distribution g is not equal to the physical distribution f. In particular, we do not expect the distribution of local relic neutrinos to be exactly isotropic. However, the assumption of an isotropic and homogeneous Fermi-Dirac distribution at $z = 31$ still allows us to use Eq. (4.2.7) to obtain physical phase-space estimates. Finally, we note that running N-body simulations backwards is non-trivial and we refer the reader to

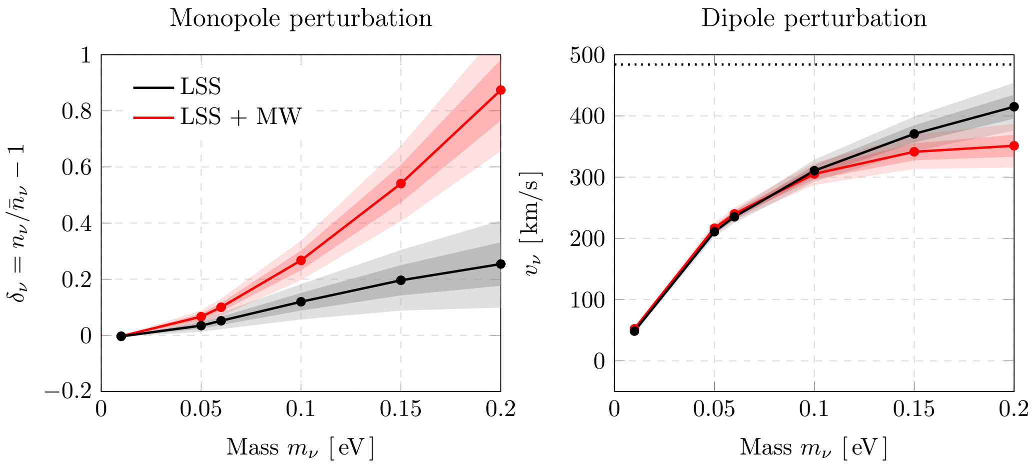

Figure 7.3: (Left) The predicted enhancement of the local neutrino density, δ_ν, as a result of the observed large-scale structure in the local Universe (LSS, black) and the combined effect of the large-scale structure and the Milky Way (LSS + MW, red). The mean and standard deviation are estimated from nine draws from the posterior distribution of the `2M++` reconstruction. (Right) The effect of the large-scale structure and Milky Way on the bulk neutrino velocity, v_ν, in the CMB rest frame. The horizontal dotted line indicates the bulk velocity of CDM and baryons within 10 Mpc of the observer.

Appendix 7.A for details on how this is accomplished.

7.3. Results

Having described our simulation methods, we are now in a position to discuss the results. In Section 7.3.1, we present the expected number density, bulk velocity, and deflection angles of relic neutrinos in the Milky Way. We also compute expected event rates for PTOLEMY. In Section 7.3.2, we turn to the angular distribution of neutrino anisotropies. In Section 7.3.3, we adopt a cosmographical perspective and look at maps of the large-scale distribution of neutrinos in the local Universe.

7.3.1. Local abundance and bulk motion

A crucial input for relic neutrino detection efforts is the expected gravitational enhancement of the local neutrino density. Using our constrained simulations, we are able for the first time to compute the total effect of the observed large-scale structure. The result is shown in the left-hand panel of Fig. 7.3. The black line (labelled LSS) shows the effect from the large-scale structure, excluding the Milky Way, on the neutrino overdensity, $\delta_\nu = n_\nu/\bar{n}_\nu - 1$, as a function of neutrino mass m_ν. The error bars indicate the dispersion among the nine constrained realizations. We see that the enhancement is negligible for $m_\nu \leq 0.05\,\text{eV}$. In fact, for the smallest mass of $0.01\,\text{eV}$, we find a small deficit of $\delta_\nu = -0.0038 \pm 0.0006$. From there, the density contrast increases approximately linearly with mass up to $\delta_\nu = 0.25 \pm 0.08$ for $0.2\,\text{eV}$.

The red line shows the combined effect of the large-scale structure and the Milky Way dark matter halo (LSS + MW). The importance of the MW increases with mass, relative to the LSS. For $m_\nu = 0.1\,\text{eV}$, they are approximately equally important. For $m_\nu = 0.2\,\text{eV}$, the MW is responsible for three-quarters of the effect. This is a result of the decrease in free-streaming length with mass: at average speed, an unperturbed $0.01\,\text{eV}$ neutrino has travelled $3.1\,\text{Gpc}$ since $z = 31$, while the number is only $200\,\text{Mpc}$ for $0.2\,\text{eV}$. As a result, lighter neutrinos are sensitive to more distant structures. We will confirm this explicitly in Section 7.3.2. Taking the difference between the results with and without the MW, we find that the galactic effect is well described by $\delta_\nu^{\text{MW}} = 27.6(m_\nu/1\,\text{eV})^{2.29}$. The near-quadratic scaling agrees with [297], who found $\delta_\nu^{\text{MW}} = 76.5(m_\nu/1\,\text{eV})^{2.21}$, but our amplitude is three times smaller. Similarly, we find significantly smaller overdensities compared to [296, 298, 299]. This may be partially due to the absence of gaseous and stellar Milky Way components in our simulations. However, the primary reason is most likely the more recent but lower estimate of the dark matter mass used in this work ($M_{200} = 0.82 \times 10^{12} M_\odot$ here compared to $M_{200} = 3.34 \times 10^{12} M_\odot$ in [296] and $M_{200} = 1.79 \times 10^{12} M_\odot$ in [298]).[6] To confirm this, we verified for one simulation that doubling the MW mass approximately restores agreement with [298]. On the other hand, both amplitude and scaling are in good agreement with the recent study [300], who also point to a difference in halo properties, rather than methodology, to explain the disagreement with [298].

Some detection proposals depend on the neutrino velocity in the lab frame [275–278, 283, 287]. From our simulations, we estimate the bulk neutrino velocity v_ν. Given that the simulation is carried out in the rest frame of the CMB, a value of $v_\nu = 0$ indicates that the neutrino dipole aligns with that of the CMB. We show the expected magnitude of the velocity perturbation in the right-hand panel of Fig. 7.3. As for δ_ν, the gravitational effect of the large-scale structure and Milky Way is negligible for $m_\nu = 0.01\,\text{eV}$. The velocity perturbation increases to $211\,\text{km/s}$ at $m_\nu = 0.05\,\text{eV}$ and trends towards $415\,\text{km/s}$ for $m_\nu = 0.2\,\text{eV}$. These neutrinos are approximately at rest with respect to the bulk flow

---

[6] In this comparison, we converted their virial masses to masses within a spherical region containing 200 times the critical density. We also note that [296] used a generalized NFW profile with an additional parameter, precluding an exact one-to-one comparison.

Table 7.2: Predictions for the neutrino dipole induced by the large-scale structure, compared with the measured CMB dipole from Planck [55]. The neutrino velocity, v_ν, is the mean velocity in the CMB rest frame. The difference, $v_\odot - v_\nu$, is the Sun's motion in the neutrino frame. The angles (l, b) correspond to the direction of the Sun's motion in the neutrino frame in galactic coordinates. The error is the standard deviation among nine realizations from the chain. The final row shows the velocity of CDM and baryons within 10 Mpc of the observer.

Mass m_ν	v_ν [km/s]	$v_\odot - v_\nu$ [km/s]	l [deg]	b [deg]
CMB	0	369.8	264.0	48.3
0.01 eV	48.5 ± 1.5	321.3 ± 1.5	263.6 ± 0.5	48.2 ± 0.1
0.05 eV	211.0 ± 4.3	193.8 ± 5.5	232.6 ± 2.4	37.9 ± 1.7
0.06 eV	235.3 ± 5.0	187.3 ± 7.6	225.7 ± 2.5	32.7 ± 1.2
0.10 eV	310.6 ± 9.0	193 ± 15	208.4 ± 2.6	14.1 ± 3.7
0.15 eV	371 ± 14	229 ± 22	199.6 ± 2.5	1.6 ± 5.6
0.20 eV	415 ± 20	265 ± 27	195.0 ± 2.7	-4.5 ± 6.7
Matter	484 ± 83	406 ± 67	206 ± 11	-10 ± 18

of matter in the inner 10 Mpc of the simulation (see Table 7.2). When we include the effect of the Milky Way, the velocity appears to converge for the largest neutrino masses. Combined with the increased density perturbation, this indicates that the simulated MW and the surrounding structure are capable of trapping 0.2 eV neutrinos in significant numbers.

In addition to the magnitude of the velocity perturbation, we can also predict its orientation. Table 7.2 shows the predicted direction of the neutrino dipole, for the runs without MW, in galactic coordinates and compares it with the measured values for the CMB dipole from Planck [55] and the direction of the simulated matter flow within 10 Mpc of the observer. For 0.01 eV, the predicted 1σ range of the neutrino dipole contains the measured CMB dipole. As m_ν increases to 0.2 eV, the values appear to converge towards the direction of the bulk flow of dark matter.[7] The results are broadly similar for the runs with MW. In the case of 0.01 eV, we find $(l, b) = (258.0^\circ \pm 0.5^\circ, 47.7^\circ \pm 0.1^\circ)$, which is still very close to the CMB dipole. For 0.2 eV, the direction changes somewhat more to $(l, b) = (203.2^\circ \pm 2.9^\circ, 7.2^\circ \pm 6.0^\circ)$. It is interesting to note that the ecliptic north pole is at $l = 97^\circ$, $b = 30^\circ$. This means that the neutrino dipole is close to the plane of Earth's orbit around the Sun, making an angle of $\phi \approx 10^\circ$. The Earth's orbital velocity is $v_\oplus \approx 30\,\mathrm{km/s}$, producing a $(2v_\oplus/v_\nu)\cos\phi \sim 20\%$ perturbation for a typical neutrino velocity of $v_\nu = 300\,\mathrm{km\,s^{-1}}$. Hence, for experiments that depend on

---

[7] Note that the uncertainties are larger for the bulk dark matter velocity, because it is computed from the forward simulations, which have a much lower sampling density near the observer.

Table 7.3: Predictions for the average deflection angle, $\cos\theta = (\mathbf{v}_\nu \cdot \mathbf{v}_\nu^{\mathrm{ini}})/(v_\nu v_\nu^{\mathrm{ini}})$, including the effects of the of the large-scale structure (LSS) and the Milky Way (LSS + MW). Using (4.2.7), we compute this from the backtraced particles with $\langle\cos\theta\rangle = (1 + \sum_i w_i \cos\theta_i)/(1 + \sum_i w_i)$, where w_i is the phase-space weight of particle i and $\cos\theta_i$ is its deflection angle.

Mass m_ν	(LSS) $\langle\cos\theta\rangle$	(LSS + MW) $\langle\cos\theta\rangle$
0.01 eV	0.999995 ± 0.000002	0.999987 ± 0.000003
0.05 eV	0.99806 ± 0.00072	0.99482 ± 0.00084
0.06 eV	0.9965 ± 0.0013	0.9905 ± 0.0015
0.10 eV	0.9847 ± 0.0058	0.9542 ± 0.0058
0.15 eV	0.958 ± 0.016	0.869 ± 0.013
0.20 eV	0.923 ± 0.029	0.754 ± 0.018

the neutrino velocity, an annual modulation may be detectable [288]. Finally, we note that the SIBELIUS-DARK simulation, which used similar techniques to set up the initial conditions, did not accurately reproduce the observed direction of the local matter flow [307]. We therefore caution that the theoretical uncertainty in the direction may be greater than the dispersion among the nine realizations given in Table 7.2.

A related quantity to the velocity perturbation is the deflection angle between the initial and final velocities, $\cos\theta = (\mathbf{v}_\nu \cdot \mathbf{v}_\nu^{\mathrm{ini}})/(v_\nu v_\nu^{\mathrm{ini}})$. For non-relativistic neutrinos, the gravitational effect on the spin is negligible, such that a deflection of the momentum vector by an angle θ implies a change in the helicity from ± 1 to $\pm\cos\theta$, with a probability $P = 1/2 - \cos\theta/2$ of observing a reversed spin [291]. It has recently been argued that the gravitational effect of the Virgo Supercluster might result in large deflection angles, significantly altering the helicity makeup of the neutrino background [294]. These authors compute deflection angles for neutrinos in halos of a similar mass to Virgo, $M = 1.48 \times 10^{15} M_\odot$, finding an average of $\langle\cos\theta\rangle = 0.54 - 0.60$ for $m_\nu = 0.05\,\mathrm{eV}$. Using our constrained simulations, which include Virgo, we can estimate directly the effect that the large-scale structure has on neutrinos that arrive on Earth. We give the average for different neutrino masses and for the cases with and without Milky Way in Table. 7.3. For 0.05 eV, we find $\langle\cos\theta\rangle = 0.99482 \pm 0.00084$, when including the Milky Way. Given that the deflection is even smaller for lighter neutrinos, we expect the effect of gravitational deflection to be negligible for the minimal neutrino mass case, $\sum m_\nu = 0.06\,\mathrm{eV}$.

Gravitational clustering also has the potential to alter the flavour composition of the local neutrino background [290]. The mass eigenstates ν_i considered so far are superpositions of flavour eigenstates ν_α, with $\alpha = e, \mu, \tau$, for electron, muon, and tau neutrinos. The two bases are related by the unitary Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix $U_{\alpha i}$

Table 7.4: Predicted number of events per year for PTOLEMY, including the effects from the large-scale structure (LSS) and the Milky Way (LSS + MW), for Dirac and Majorana neutrinos. We give the results for the individual mass states, with (7.3.1) giving the total rate. The uncertainty corresponds to the 1σ dispersion among nine realizations from the chain.

	(LSS)		(LSS + MW)	
Mass m_ν	$\Gamma^D_{i,\text{CNB}}$ $[\text{yr}^{-1}]$	$\Gamma^M_{i,\text{CNB}}$ $[\text{yr}^{-1}]$	$\Gamma^D_{i,\text{CNB}}$ $[\text{yr}^{-1}]$	$\Gamma^M_{i,\text{CNB}}$ $[\text{yr}^{-1}]$
0.01 eV	4.042 ± 0.002	8.075 ± 0.005	4.045 ± 0.002	8.080 ± 0.005
0.05 eV	4.20 ± 0.05	8.39 ± 0.09	4.33 ± 0.05	8.65 ± 0.09
0.06 eV	4.27 ± 0.06	8.53 ± 0.12	4.46 ± 0.06	8.92 ± 0.13
0.10 eV	4.54 ± 0.13	9.08 ± 0.26	5.14 ± 0.14	10.27 ± 0.29
0.15 eV	4.85 ± 0.22	9.70 ± 0.44	6.25 ± 0.27	12.49 ± 0.54
0.20 eV	5.09 ± 0.32	10.17 ± 0.63	7.60 ± 0.44	15.19 ± 0.88

[9, 10]. The flavour composition could be altered, since the degree of clustering depends on mass. For instance, assuming the mass ordering is normal, the contribution of ν_e to the heaviest mass state ν_3 is only $|U_{e3}|^2 = 2.3\%$. Therefore, if ν_3 is much more strongly clustered than ν_1 and ν_2, most relic neutrinos on Earth would be ν_μ or ν_τ. For this effect to be large, the masses must be hierarchical ($m_1 \ll m_3$ or $m_3 \ll m_1$), which requires $m_\nu \lesssim 0.1\,\text{eV}$. Fig. 7.3 shows that the differences in the density contrast δ_ν are then small, which implies that the fraction of ν_e is not significantly altered from its primordial value of 1/3. We nevertheless incorporate this effect in the calculation below.

We now have the necessary ingredients to compute the expected event rate for an experiment like PTOLEMY. The CNB capture rate,

$$\Gamma_\text{CNB} = \sum_{i=1}^{N_\nu} \Gamma_{i,\text{CNB}} |U_{ei}|^2, \tag{7.3.1}$$

is summed over all mass eigenstates that exceed the energy threshold of the experiment, weighted by the PMNS mixing elements, U_{ei}. The event rate for mass state ν_i is given by [321]

$$\Gamma_{i,\text{CNB}} = N\bar{\sigma}\left[n_i^+ A_i^+ + n_i^- A_i^-\right], \tag{7.3.2}$$

where N is the number of targets, $\bar{\sigma}$ is the average cross section, $n_i^\pm$ are the number densities for the two spin states, $A_i^\pm = 1 \mp v_i/c$ is a spin-dependent factor, and v_i is the velocity of the mass eigenstate. As discussed, gravitational deflection by an angle θ reverses the spin with probability $P = 1/2 - \cos\theta/2$. The number densities for both spin states are then given by

$$n_i^{\pm} = n_i \left[\frac{1}{2} \pm \frac{1}{2}\langle\cos\theta\rangle_i\right]. \tag{7.3.3}$$

In the absence of clustering and deflection, $\langle\cos\theta\rangle_i = 1$, such that $n_i^+ = n_i = \bar{n}$ and $n_i^- = 0$ for Dirac neutrinos. For Majorana neutrinos, the densities are both equal to the mean: $n_i^+ = n_i^- = \bar{n}$. Consequently, for non-relativistic neutrinos with $A_i^{\pm} = 1$, the expected signal is twice as large in the Majorana case. If we allow for gravitational effects, we instead obtain

$$\Gamma^D_{\text{CNB}} = N\bar{\sigma}\sum_{i=1}^{N_\nu}|U_{ei}|^2\left[1 + \langle\cos\theta\rangle_i\frac{v_i}{c}\right]n_i, \tag{7.3.4}$$

$$\Gamma^M_{\text{CNB}} = N\bar{\sigma}\sum_{i=1}^{N_\nu}|U_{ei}|^2 2n_i, \tag{7.3.5}$$

for the Dirac and Majorana cases, respectively. Plugging in the number $N = 100\,\text{g}/m_{^3\text{H}}$ of tritium atoms for PTOLEMY [270] and the average cross section $\bar{\sigma} = 3.834\times10^{-45}\,\text{cm}^2$ from [321], and a mean number density of $\bar{n} = 56\,\text{cm}^{-3}$ per degree of freedom, we obtain the event rates in Table 7.4. We report the values for the individual mass eigenstates. Comparing the most and least massive cases, we see that gravitational clustering only has a marginal effect, boosting the capture rate by less than a factor of two. For each mass, we predict a factor ~ 2 difference between the Dirac and Majorana cases. Let us now compute the total event rate for the minimal neutrino mass case, using $|U_{ei}|^2 = (0.678, 0.299, 0.023)$ [29]. We assume that only the heaviest neutrinos with $m_\nu = 0.05\,\text{eV}$ (ν_3 under the normal ordering or ν_1 and ν_2 under the inverted ordering) would produce peaks in the electron energy spectrum far enough beyond the β-decay endpoint to be detected by PTOLEMY with a reasonable energy resolution [271]. For the normal ordering, we then find $\Gamma_{\text{CNB}} \approx 0.1\,\text{yr}^{-1}$ (Dirac) or $0.2\,\text{yr}^{-1}$ (Majorana), while $\Gamma_{\text{CNB}} \approx 4\,\text{yr}^{-1}$ (Dirac) or $8\,\text{yr}^{-1}$ (Majorana) for the inverted ordering.

7.3.2. Angular anisotropies

Having presented our results for the monopole and dipole moments, we now turn to higher-order moments of the neutrino distribution. Fig. 7.4 presents maps of the predicted angular anisotropies in the number density, $\delta_\nu(\theta,\phi) = n_\nu(\theta,\phi)/(\bar{n}_\nu/4\pi)$, for four different masses, after subtracting the monopole and dipole perturbations. The maps show relative

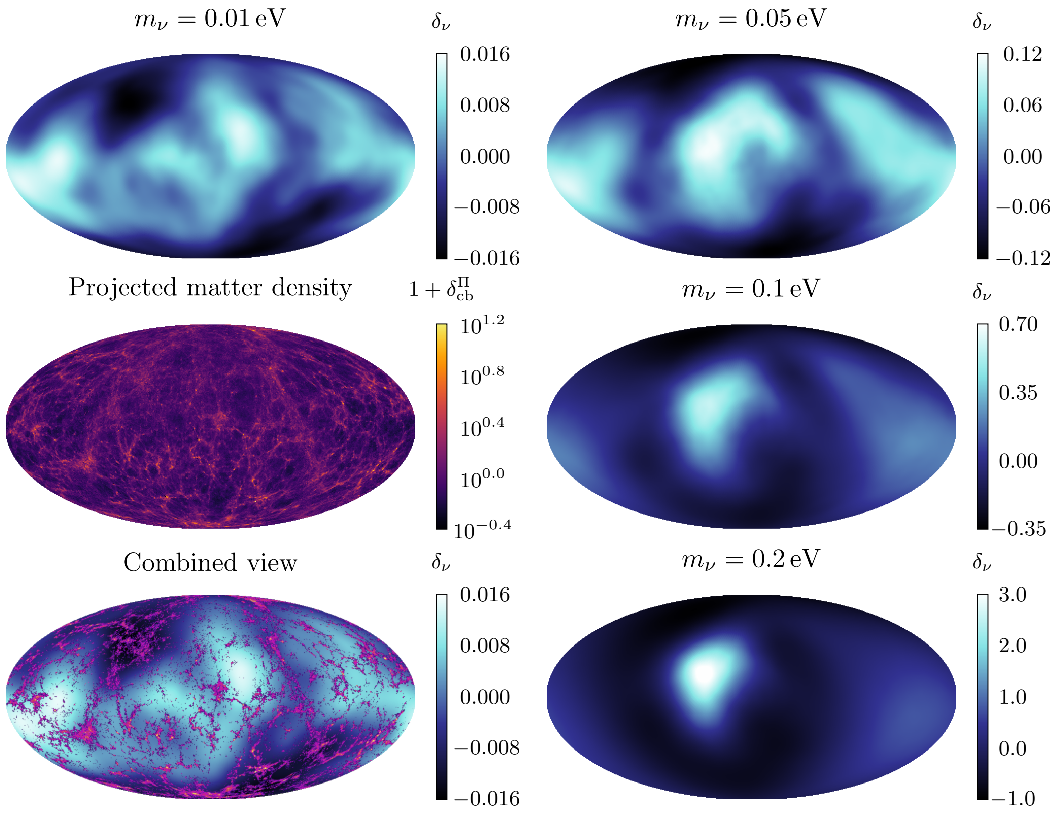

Figure 7.4: Angular anisotropies in the neutrino number density contrast, δ_ν, for $m_\nu = 0.01\,\mathrm{eV}$ (top left) and for $m_\nu \in \{0.05, 0.1, 0.2\}\,\mathrm{eV}$ (right). In all cases, we subtract the monopole and dipole moments and smooth over 3° scales. We also show the projected dark matter and baryon density, δ^Π_cb, within 200 Mpc of the observer, both separately (middle left) and overlaid on the top of the neutrino density for $m_\nu = 0.01\,\mathrm{eV}$ (bottom left). We observe that the projected dark matter density and the local neutrino density are anti-correlated on the sky. Except for the projected matter density, the maps are all based on backtraced particles.

variations in the neutrino density for individual mass eigenstates, computed via equation (7.2.2) by adding the weights of backtraced particles along each direction. As discussed in Section 7.2.3, these particles represent an ergodic ensemble of neutrino paths with weights that correct for the isotropic sampling distribution. In [186], it is shown that the statistical properties of such weighted particle ensembles are consistent with the transfer functions obtained from an Eulerian fluid calculation at the linear level. One advantage of the particle-based treatment, however, is its ability to describe the nonlinear growth of neutrino perturbations, which becomes important when the neutrino mass is large, as discussed below.

Each map is averaged over nine realizations from the `2M++` reconstruction. The top-left panel of Fig. 7.4 shows the map for $m_\nu = 0.01\,\mathrm{eV}$ and the right-hand panels show maps for $m_\nu \in \{0.05, 0.1, 0.2\}$ eV. First of all, we observe that the magnitude of the perturbations strongly depends on mass: they are $\mathcal{O}(10^{-2})$ for $m_\nu = 0.01\,\mathrm{eV}$ and $\mathcal{O}(1)$ for $m_\nu = 0.2\,\mathrm{eV}$. We also see that the largest neutrino mass maps have large-scale perturbations that are suppressed, relative to small-scale perturbations, for the smaller neutrino masses. The middle-left panel shows the projected density of dark matter and baryons,

$$1 + \delta^{\Pi}_{\mathrm{cb}}(\theta, \phi, R_{\max}) = \frac{\int_0^{R_{\max}} \rho_{\mathrm{cb}}(\mathbf{r})\,\mathrm{d}r}{\int_0^{R_{\max}} \bar{\rho}_{\mathrm{cb}}(\mathbf{r})\,\mathrm{d}r}, \tag{7.3.6}$$

up to a distance of $R_{\max} = 200\,\mathrm{Mpc}$ from the observer. Comparing this with the neutrino maps, we find that distant matter fluctuations are anti-correlated with local neutrino fluctuations. This can be seen more clearly in the bottom-left panel, in which the projected matter perturbations are overlaid on the neutrino perturbations for $m_\nu = 0.01\,\mathrm{eV}$. The anti-correlation is much more evident for smaller neutrino masses.

Next, we compute angular power spectra, C_ℓ^ν, from the neutrino overdensity maps. To compare our results with other works [263, 264, 300], we convert the spectra to temperature units by assuming that $\delta T_\nu/\bar{T}_\nu \sim \delta n_\nu/3\bar{n}_\nu$.[8] In Fig. 7.5, we show the results for five different masses, averaging over nine realizations from the chain. To uncover the perturbations imprinted by the large-scale structure, we fit smooth spectra of the form

$$C_\ell^{\mathrm{fit}} = \exp\left[c_1 + c_2 \log \ell + c_3 (\log \ell)^2\right], \tag{7.3.7}$$

to the simulation predictions, restricting to the multipoles with $1 \leq \ell \leq 15$, since higher-order multipoles are noisy and poorly constrained. The thick curves in Fig. 7.5 correspond to these fits, with the solid and dashed lines indicating the LSS-only and combined LSS + MW results, respectively. As expected from the previous section, the effect of the MW is most pronounced for the largest neutrino masses and the lowest-order multipoles. The difference between the dashed and solid curves is negligible for $m_\nu \leq 0.05\,\mathrm{eV}$, but clearly visible for $m_\nu = 0.2\,\mathrm{eV}$. We compute our maps in the rest frame of the simulations, without accounting for observer motion. Therefore, Fig. 7.5 shows the intrinsic dipole moment ($\ell = 1$) arising from large-scale matter fluctuations. The value is orders of magnitude larger than the intrinsic dipole expected for massless tracers like the CMB

---

[8] This follows from the idealized result, $n_\nu = 3\zeta(3)T_\nu^3/2\pi^2$, for the Fermi-Dirac distribution, even though the actual momentum distribution of clustered neutrinos is heavily perturbed.

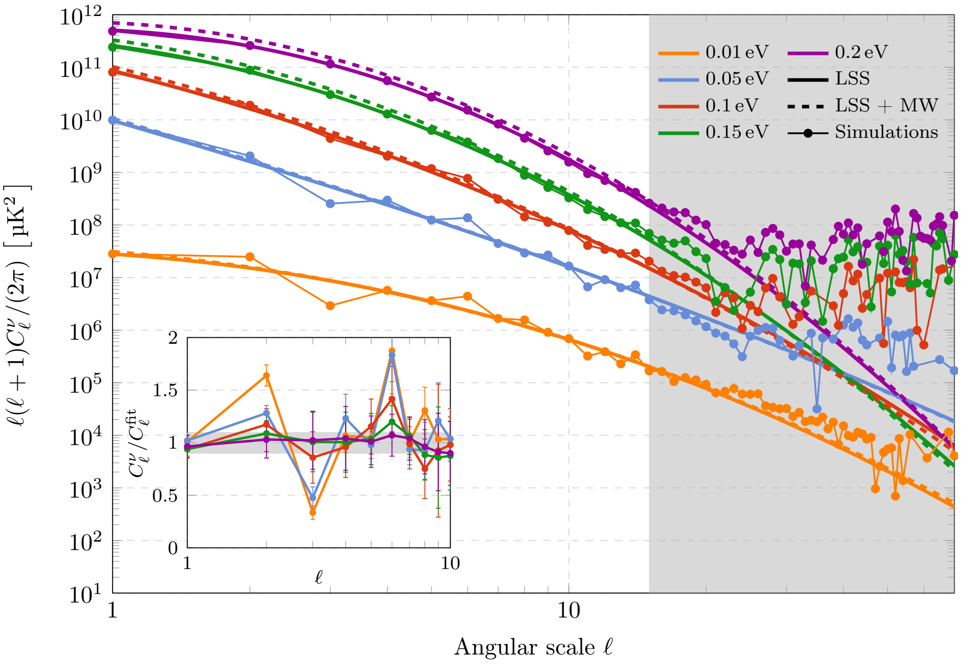

Figure 7.5: Angular neutrino temperature power spectra, C_ℓ^ν, for different masses. We fit a smooth spectrum, $C_\ell^{\text{fit}} = \exp(c_1 + c_2 \log \ell + c_3 \log^2 \ell)$, up to $\ell_{\text{max}} = 15$ for the simulations with and without a Milky Way (dashed and solid lines, respectively). To avoid clutter, we only show the data for the simulations without MW. The inset graph zooms in on the first ten multipoles, showing the data relative to the fit. The grey error bar represents ±10%. The oscillatory perturbations arise from the imprint of dark matter perturbations on the neutrino background and can ultimately be traced to cosmic variance in the matter distribution.

[322–324]. This is consistent with the behaviour seen in Fig. 7.5, which shows that low-multipole perturbations are strongly enhanced for heavier neutrinos.

Our results differ substantially from [300], who compute a range of temperature power spectra for $m_\nu = 0.1\,\text{eV}$ using different $(25\,\text{Mpc})^3$ simulations. We find a slope that is much steeper and an amplitude at low multipoles that is greater. This could be due to the absence of large-scale structure in their simulations, explaining the lack of power at low multipoles. Our results are in good agreement with the linear theory calculations of [264] for $m_\nu < 0.1\,\text{eV}$. For 0.1 eV, the normalization at low multipoles agrees, but we predict significantly more power beyond $\ell \geq 10$, where the linear calculation likely breaks down. Similarly, although our definition of the neutrino temperature power spectrum is somewhat different from [263], given that we do not define a power spectrum for each momentum bin separately but show the pointwise integrated result, we obtain at least qualitative agreement with their linear calculations for masses $m_\nu \leq 0.1\,\text{eV}$, the largest

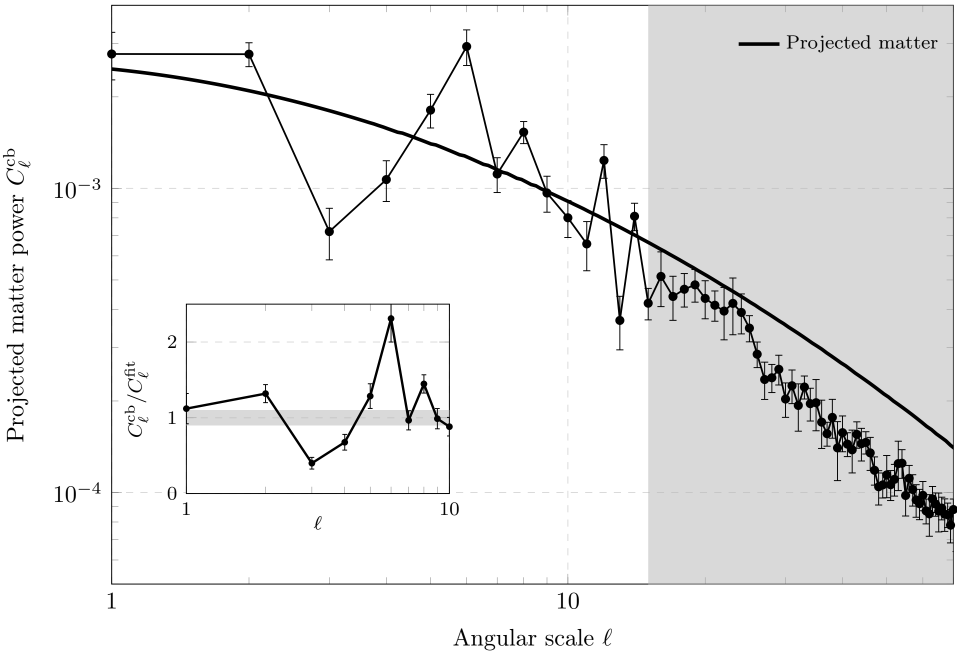

Figure 7.6: Dimensionless angular power spectrum of the projected CDM and baryon density contrast, $C_\ell^{\rm cb}$, out to 200 Mpc for 0.01 eV. By construction, the results are similar for different neutrino masses. We fit a smooth spectrum, $C_\ell^{\rm fit} = \exp(c_1 + c_2 \log \ell + c_3 \log^2 \ell)$, up to $\ell_{\rm max} = 15$ to the simulation spectrum (thick solid curve). The inset graph zooms in on the lowest-order multipoles, showing the data relative to the fit. The grey error bar represents ±10%.

mass considered amenable to their analysis. These authors model the gravitational deflection of neutrinos with a lensing potential, similar to what is done for the CMB [325]. A key difference between our results and the linear theory calculations [263, 264] is the presence of oscillatory perturbations around the smooth spectra in Fig. 7.5, which are much larger than the predicted lensing effect in [263]. This can be seen more clearly in the inset graph, which zooms in on the lowest-order multipoles ($\ell \leq 10$) and shows the simulation predictions relative to the smooth fits. The perturbations depend sensitively on mass, being most prominent for 0.01 eV and nearly absent for 0.2 eV.

The origin of these perturbations becomes clear when we plot the angular power spectrum, $C_\ell^{\rm cb}$, of the projected CDM and baryon density up to 200 Mpc, in Fig 7.6. In this case, we compute a dimensionless power spectrum directly from the maps of the projected density contrast, $\delta_{\rm cb}^{\Pi}(\theta, \phi)$, defined in Eq. (7.3.6). Fitting a smooth power spectrum (7.3.7) in the same way as for the neutrinos, reveals the same oscillatory perturbations. This suggests that cosmic variance in the matter density field is imprinted on the local neutrino background if the neutrino mass is sufficiently small. To confirm this explicitly,

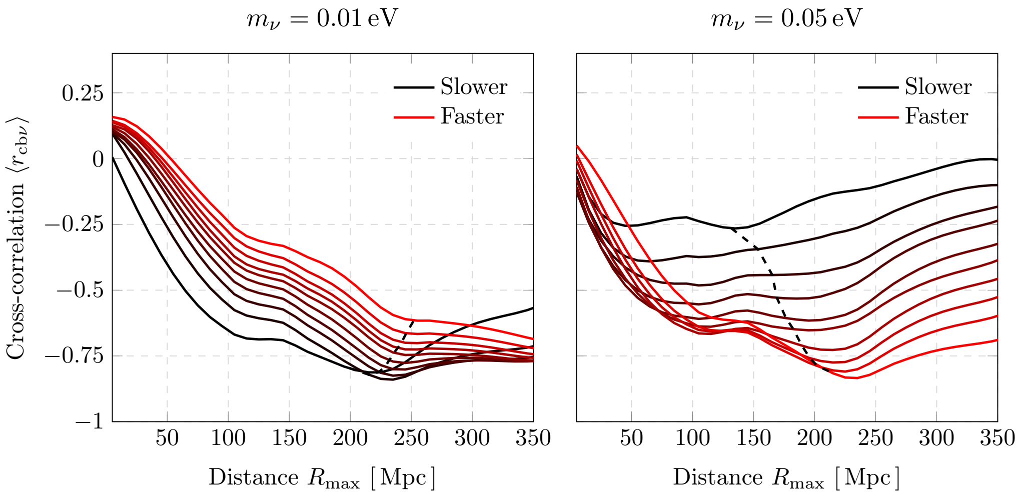

Figure 7.7: Cross-correlation coefficient, $r_{\mathrm{cb}\nu}(\ell) = C_\ell^{\mathrm{cb}\nu}/(C_\ell^\nu C_\ell^{\mathrm{cb}})^{1/2}$, between the local neutrino density and the projected CDM and baryon density, as a function of the maximum projected distance $R_{\max}$, for $m_\nu \in \{0.01, 0.05\}$ eV, split into ten equal-sized neutrino momentum bins. The coefficients are averaged over the multipoles $1 \leq \ell \leq 10$ and the curves are smoothed with a Savitzky-Golay filter. The dashed line indicates the locus of the barycentre of each curve, indicating that the sensitivity shifts to larger distances for faster neutrinos. Note that the overall momentum range is much wider for the 0.01 eV case.

we compute the cross-correlation coefficient, $r_{\mathrm{cb}\nu}(\ell) = C_\ell^{\mathrm{cb}\nu}/(C_\ell^\nu C_\ell^{\mathrm{cb}})^{1/2}$, between the local neutrino density and the projected dark matter and baryon density, as a function of the maximum projected distance $R_{\max}$. The results, averaged over the lowest-order multipoles, $1 \leq \ell \leq 10$, and smoothed with a Savitzky-Golay filter, are shown in Fig. 7.7. We additionally split the results into ten equal-sized neutrino momentum bins, with redder curves indicating faster neutrinos. For both neutrino masses shown, $m_\nu = 0.01\,\mathrm{eV}$ (left) and 0.05 eV (right), there is a strong anti-correlation that peaks around $r_{\mathrm{cb}\nu} = -0.8$. In both cases, faster neutrinos are sensitive to more distant matter fluctuations. To emphasize this point, we indicate the locus of the barycentre of each curve by a black dashed line.

Note that $r_{\mathrm{cb}\nu}$ trends upwards as $R_{\max}$ decreases, eventually becoming positive for the fastest neutrinos. This might be explained by the gravitational attraction of neutrinos to positive density perturbations close to the observer. In this case, a positive correlation should be expected. In line with expectation, the distance at which the correlation becomes positive increases with neutrino momentum. Interestingly, the anti-correlation becomes weaker with neutrino momentum for 0.01 eV and stronger with neutrino momentum for 0.05 eV. A simple explanation for this could be that the anti-correlation begins trending upwards earlier for faster neutrinos, causing a reversal in the trend, as can be seen for $R_{\max} < 100\,\text{Mpc}$ in the case of $m_\nu = 0.05\,\text{eV}$. For $m_\nu = 0.01\,\text{eV}$, this reversal may only happen at distances that are not constrained by the `2M++` data underlying our simulations.

Just before this paper was submitted, a related study appeared in which neutrino anisotropy maps are analysed for different random configurations of dark matter halos in a $(25\,\text{Mpc})^3$ volume [300]. For some configurations, they report positive or negative correlations between the neutrino and projected dark matter densities. Overall, the ensemble average of cross-power spectra is consistent with zero. Taking into account the smaller volume of the simulations, this can probably be understood in terms of the aforementioned transition from positive to negative correlations close to the observer.

7.3.3. Cosmography

In this section, we make a first attempt at neutrino cosmography. Given the limited resolution of our simulations, we focus on one illustrative example and run a higher-resolution constrained simulation with $N_\nu = N_{\text{cb}} = 1024^3$ particles for $\sum m_\nu = 0.06\,\text{eV}$. In Fig. 7.8, we present maps of the neutrino density (left) and dark matter and baryon density (right), in a slice of $500 \times 500 \times 60\,\text{Mpc}$ that includes the Local Group and several well-known clusters. A few striking observations can be made. First of all, the large-scale neutrino and dark matter densities are positively correlated. This explains the anti-correlation seen in the previous section. Relic neutrinos that are captured by massive objects form localized clouds. Hence, while they are visible from the hypothetical viewpoint[9] depicted in Fig. 7.8, they would not be seen from Earth along lines of sight that intersect those structures.

After plotting the locations of several famous galaxy clusters, we find massive dark matter structures associated with each of them. Surrounding most of these structures, we also identify neutrino clouds that stretch over 10 Mpc scales and reach central overdensities of 30%. Two interesting exceptions are the Perseus and Pisces clusters, which lie close to the Taurus void [326] and appear to inhabit a large region that is deficient in neutrinos (a ‘glade’ in the neutrino cloudscape). Although we see some collapsed dark matter structures at their locations, these are more dispersed compared to other clusters. This

---

[9] The viewpoint of a distant observer looking at the Milky Way in its cosmic environment. One might call this the Archimedean viewpoint, based on Archimedes’ claim that he could lift the Earth given only a fulcrum and a place to stand.

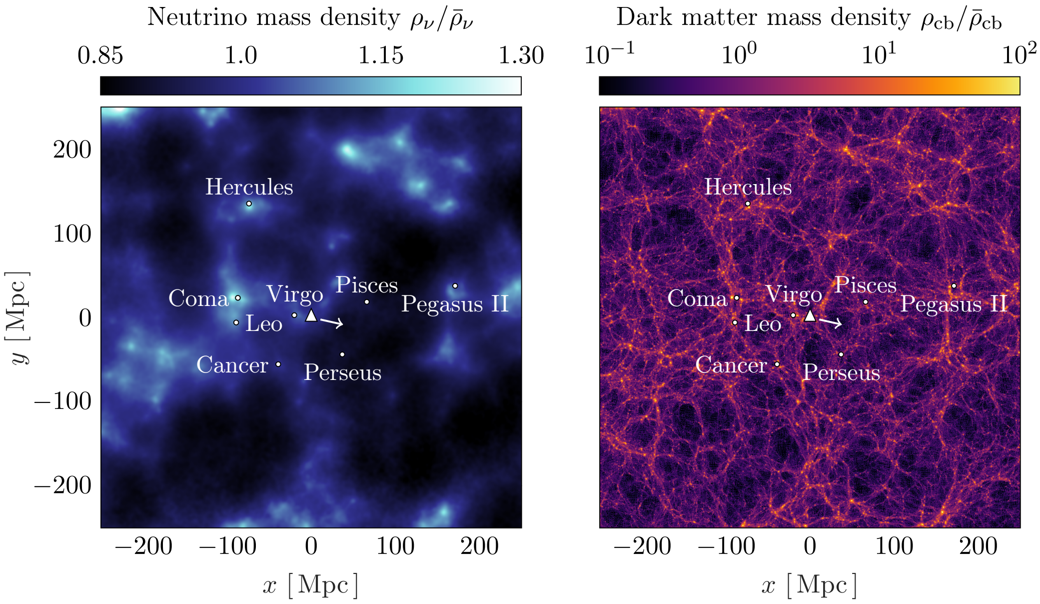

Figure 7.8: Slice of the projected neutrino (left) and dark matter (right) mass density, with thickness of 60 Mpc, containing the Local Group and nearby clusters, for a species with mass $m_\nu = 0.06\,\mathrm{eV}$. The location of the Milky Way is indicated by a white triangle. The arrow indicates the direction of the relative neutrino velocity. In terms of the comoving distance r, right ascension α, and declination δ, the coordinates are $(x, y) = r\cos\delta(\cos\alpha, \sin\alpha)$.

could be due to a failure of the constrained simulations to model the Perseus-Pisces wall accurately [311].

The Milky Way is marked by a white triangle, located along a filament that stretches towards the Virgo cluster. For this neutrino mass, $m_\nu = 0.06\,\mathrm{eV}$, we appear to inhabit a region with a large-scale neutrino overdensity that is not due to the Milky Way. It was this large-scale modulation of the neutrino density that originally motivated our investigation. Its effect was shown in Fig. 7.3 as a function of mass. For 0.01 eV, we predicted a small neutrino deficit. We now see that this could be due to our proximity to the Taurus/Perseus-Pisces glade. Hence, the local neutrino density depends on the interplay between the overdensities associated with Virgo and the Local Group and nearby underdensities. The direction of the neutrino dipole is indicated by a white arrow. It points away from the overdense region around the Coma cluster, which is consistent with our motion towards the Shapley Supercluster and the Great Attractor [327]. Correspondingly, it points towards an underdense region known as the Dipole Repeller, which causes an apparent repulsion [328]. In short, the behaviour of the CNB dipole is similar to that of the CMB when the neutrino mass is small, consistent with our findings in Section 7.3.1.

7.4. Conclusion

Direct detection of the Cosmic Neutrino Background (CNB) remains one of the great challenges in cosmology. In this paper, we have analysed the gravitational effects of the large-scale structure and the Milky Way on the local neutrino background. Through the use of the 'BORG' framework for Bayesian forward modelling of large-scale structure observations [309, 310], we have carried out constrained simulations of the local Universe for different neutrino cosmologies with masses between $\sum m_\nu = 0.01\,\mathrm{eV}$ and $\sum m_\nu = 0.6\,\mathrm{eV}$. The constraints are based on the `2M++` catalogue [301], which maps the local Universe out to a distance of $200h^{-1}\,\mathrm{Mpc}$. We account for the Milky Way dark matter halo, using an updated estimate of the mass from [302]. By tracing neutrinos back through the galaxy and large-scale structure with a bitwise reversible version of the N-body code Gadget-4 [77], keeping track of phase-space density perturbations, we compute statistics of the expected neutrino flux. Our results suggest that the gravitational clustering of neutrinos due to the large-scale structure is not negligible compared to the effect of the Milky Way, with both contributing about half of the total effect for 0.1 eV neutrinos.

Despite the inclusion of the large-scale structure, we find smaller overdensities compared to earlier studies [296–299]. We attribute this to a decrease in recent estimates of the Milky Way halo mass. We therefore predict only marginal increases in the event rates for tritium capture experiments like PTOLEMY [270–272]. Additionally, we also predict a smaller impact of gravitational deflection on the helicity distribution of the neutrino background compared to [294], due to our distance from the centre of the Virgo cluster. As a result, the difference between the event rates for Dirac and Majorana neutrinos is slightly smaller, though still close to 100% in most cases. Similarly, we also predict a smaller impact on the flavour composition compared to [290], with an electron-neutrino fraction that is close to 1/3 even for hierarchical masses. We also make predictions for the neutrino dipole. In the limit of very small neutrino masses, $m_\nu \leq 0.01\,\text{eV}$, we recover the CMB result with a dipole that corresponds to Solar motion towards $(l, b) = (264°, 48°)$ at a relative velocity of around 300 km/s. The velocities are significantly perturbed for larger masses and the dipole direction shifts, but remains nearly parallel to the ecliptic plane. This implies a near-maximal annual modulation in the neutrino velocity throughout Earth's orbit around the Sun.

Although perhaps unlikely, a future directional CNB detector might image the angular distribution of relic neutrinos. We have produced maps and power spectra of the nonlinear neutrino perturbations imprinted by the large-scale structure. Our findings are in qualitative agreement with the linear theory results of [263] for masses $m_\nu \leq 0.1\,\text{eV}$, but with a much larger gravitational effect that produces an oscillatory feature in the power spectrum. This feature is related to cosmic variance in the dark matter density field. Indeed, we find that local neutrino density perturbations, in principle detectable from Earth, are anti-correlated with the projected dark matter density up to at least 250 Mpc, the largest distance constrained by the `2M++` catalogue, although for very nearby structures and fast neutrinos, we instead predict a positive correlation. The distance at which neutrinos are most sensitive to the intervening cosmic structure increases with momentum and decreases with mass, potentially enabling a kind of neutrino tomography of the large-scale structure, which would be impervious to extinction by gas and dust. Finally, we presented maps of the forecasted neutrino distribution in the local Universe, identifying neutrino clouds associated with several well-known clusters, such as Coma and Hercules. We release our simulation data to the public, which we hope will be useful for future analyses of the neutrino background.

7.A. Reversible simulations

Running a cosmological N-body simulation backwards to recover the initial conditions is non-trivial (see [329–331] for related ideas). In principle, leapfrog integration is time-reversible [200]. However, in practice, small rounding errors inevitably accumulate in the backwards direction. This is problematic if one aims to recover a low entropy initial configuration (such as two merging galaxies that are initially well separated) from a final high entropy configuration (the merged galaxy). The root of the problem is the non-associativity of standard floating point arithmetic, causing different rounding errors in backwards integrations. Furthermore, floating point errors are not necessarily reproducible in parallel programs, because of the unpredictable execution order of threads.

We here briefly discuss the modifications necessary to make a cosmological code reversible, in anticipation that this may be useful for other applications.

To test the bitwise reversibility of GADGET-4, we periodically compute a hash of all particle data. The state of the simulation should be identical in the forwards and backwards directions at the beginning and end, respectively, of each corresponding step. Unsurprisingly, the code is not reversible by default. A first step towards achieving this is to store particle positions and velocities as integers. Implementing integer velocities is a natural step, because GADGET-4 already uses integer positions by default to achieve constant precision throughout the simulation domain [77]. However, this is by no means enough to guarantee reversibility, if only because the gravitational Tree-PM algorithm still relies on floating point operations.

To guarantee reversibility, we must therefore also ensure that different threads execute their tree calculations in the same order in both directions. Furthermore, there can be no time-asymmetric decision making. For instance, we use a basic geometric tree opening criterion [332], because the more adaptive opening criterion available in GADGET-4 depends on the particle accelerations from the previous step, which are different in the backwards direction. Similarly, the time step is usually chosen based on the maximum distance that particles can move or on the acceleration of particles in the previous time step, which again introduces an asymmetry. To address this problem without adopting a constant time step, we store a list of step sizes used in the forwards direction and feed this file back in the backwards direction. Special consideration is also needed for the neutrinos to ensure that the δf weighting is time-reversible. Special relativistic velocities (6.5.3) can be absorbed in the leapfrog integration scheme [186].

The domain decomposition is another point of concern. By default, GADGET-4 uses floating point arithmetic for load balancing, which can lead to differences between the forwards and backwards runs. These operations are therefore modified to use integers. As a final example, recall that we inject additional ‘spectator’ neutrinos at the start of the backwards runs. We take steps to ensure that their presence affects neither the domain decomposition of the original particles nor alters the gravity calculation. With these modifications, we exactly recover the initial conditions when running our constrained neutrino simulations backwards.

Baryonic effects

8

This chapter deals with the combined effect of cosmology and baryonic feedback on the large-scale structure. It shows that the effects are not exactly separable, but that the baryonic suppression of the power spectrum is stronger for models with massive neutrinos. This is explained in terms of changes in the baryon fraction and binding energy of halos. Using a halo concentration emulator, predictions are made for the non-factorizable corrections as a function of Ω_m and σ_8.

Cosmological dependence of baryonic feedback

Abstract: The large-scale distribution of matter is influenced not only by the choice of cosmology but also by astrophysical processes, such as feedback from active galactic nuclei and stellar winds. This has important implications for cosmological parameter estimation, particularly from weak-lensing surveys such as Euclid and LSST. In this study, we use the new FLAMINGO suite of large-volume hydrodynamical simulations to study the extent to which variations in cosmology can be treated independently of astrophysical processes. We find a weak dependence of baryonic effects on cosmology, leading to non-factorizable corrections to the power spectrum and halo mass function. In particular, we find that the baryonic suppression of the matter power spectrum is stronger for models in which structure formation is already suppressed, such as in models with massive neutrinos. To determine the physical origin of these corrections, we study an array of halo properties, including the halo concentration and environmental density, ranking each property along two axes: its sensitivity to cosmology and its importance in regulating feedback. Doing so, we determine that the baryon fraction and halo concentration are most important in the halo mass range relevant for the suppression of the power spectrum for $k \leq 10\,\mathrm{Mpc}^{-1}$. Our model predicts that non-factorizable corrections exceed 1% for cosmological parameter variations allowed by large-scale structure surveys and should therefore be taken into account.

8.1. Introduction

Two of the principal components invoked by the ΛCDM model of cosmology, dark energy and cold dark matter, cannot be explained by the Standard Model of particle physics and may only be effective descriptions of phenomena yet to be discovered. The success of the model in reproducing a wide array of observations indicates that signatures of new physics are likely to manifest either as small modifications to the ΛCDM prediction or in the relatively unexplored high-redshift régime. As such, the model is under intense scrutiny at both the high-precision and high-redshift frontiers. A 5σ tension between distance ladder measurements of the Hubble constant [333] and the ΛCDM prediction extrapolated from cosmic microwave background (CMB) observations [112] could be a first indication of new physics [e.g. 334, 335]. A similar discrepancy in the matter density and amplitude of fluctuations on $8h^{-1}$Mpc scales, $S_8 = \sigma_8\sqrt{\Omega_\mathrm{m}}$, comparing the CMB-extrapolated prediction with measurements from galaxy clustering and weak lensing observations [336–338], stands at $2 - 3\sigma$. Further motivating the work at the high-precision frontier is the possibility to measure the sum of neutrino masses. The imprint of massive neutrinos could be detected by galaxy surveys such as DESI [339], even for the minimum value allowed by oscillation data, $\sum m_\nu = 0.06\,\mathrm{eV}$ [29], but this requires percent-level accuracy in large-scale structure measurements and predictions [132, 204].

Complicating these efforts is the fact that astrophysical processes, such as feedback from supernovae and active galactic nuclei (AGN), can change the distribution of matter even on relatively large scales [340–342]. By heating and ejecting gas into the intergalactic medium, AGN feedback can suppress the power spectrum of matter fluctuations by $\mathcal{O}(10\%)$ on nonlinear scales, $1\,\mathrm{Mpc}^{-1} < k < 10\,\mathrm{Mpc}^{-1}$. On smaller scales, the power spectrum is boosted, first by star formation and then by gas cooling, both processes allowing matter to contract ([343]; but see [344]). A crucial question for the interpretation of large-scale structure observations concerns the dependence of these baryonic effects on cosmology. A number of previous studies [345–348] have shown that the effects of cosmology and baryonic physics are separable to a first approximation, but with residual effects of up to several percent. As these effects are approaching the statistical errors of upcoming galaxy surveys, we are interested in modelling the non-factorizable corrections that arise from variations in cosmology in the presence of baryonic physics. We will consider their origin and characterize their overall importance.

One could imagine different mechanisms through which baryonic processes, such as star formation and the growth of supermassive black holes, and hence baryonic feedback, could depend on cosmology. Although dark matter halos share a universal density profile [316], their concentrations depend on cosmology [349–353]. Cosmological model variations that slow the rate of structure formation (such as decreasing the matter density, Ω_m, or amplitude of clustering, σ_8) lead to less concentrated halos, lowering the gravitational binding energy and altering the balance between outflows and black hole accretion [354, 355]. Another potential channel is the formation history of dark matter halos. If halos assemble their mass more slowly, galaxy formation and feedback processes may be delayed. A third possibility is that a change in the large-scale distribution of matter affects the halo environment, which could affect halo properties indirectly through assembly bias [356–358] or affect feedback by changing the density of the halo outskirts. Finally, variations in the baryon density, Ω_b, alter the amount of gas that is available for star formation and black hole accretion.

Massive neutrinos could plausibly affect feedback through any of these channels. Neutrinos cluster less effectively on scales smaller than their free-streaming length [120], which results in less concentrated halos, delayed structure formation, and smoother halo environments. Moreover, neutrinos also affect the baryon fraction, $\Omega_\mathrm{b}/\Omega_\mathrm{c}$, relative to the cold dark matter density, Ω_c, given that a change in neutrino mass at fixed matter density, $\Omega_\mathrm{m} = \Omega_\mathrm{b} + \Omega_\mathrm{c} + \Omega_\nu$, and baryon density will alter the amount of gas that is available for a halo of a given cold dark matter mass. These channels are not necessarily mutually exclusive. To find out which, if any, play a role in regulating the amount of feedback, we will use the FLAMINGO suite of hydrodynamical simulations [359], which include several feedback and cosmology variations. We will use halo properties (such as the concentration and formation epoch) as proxies for the ways in which feedback could depend on cosmology. We will then formulate a model to predict the non-factorizable correction to the matter power spectrum.

The organization of the paper is as follows. In Section 8.2, we introduce the FLAMINGO suite of simulations. In Section 8.3, we present the main results, first studying the effects of baryons and neutrinos on global large-scale structure statistics and then considering the effect on individual halos. Finally, we discuss and conclude in Section 8.4.

8.2. Simulations

Our analysis is based on the new FLAMINGO suite of cosmological hydrodynamical simulations [359]. The FLAMINGO simulations use an updated version of the subgrid models used by the earlier Cosmo-OWLS [360] and BAHAMAS [361] projects. The FLAMINGO simulations also implement a number of technical improvements, related to the initial conditions and the treatment of massive neutrinos, aimed at improving

Table 8.1: An overview of the FLAMINGO simulations used in this paper. The number of baryon particles, N_b, is equal to the number of cold dark matter particles, N_c, for the simulations that have them. The number of neutrino particles is always $N_\nu = N_\mathrm{c}/1.8^3$. The bottom rows correspond to the gravity-only (DMO) counterparts of the top rows. The columns correspond to the side length, L, the number and the mass of cold dark matter particles, N_c and m_c, the initial mass of gas particles, m_g, and the cosmological parameters. All simulations assume a flat ($\Omega_k = 0$) ΛCDM Universe with massive neutrinos and with an amount of radiation corresponding to $T_\mathrm{CMB} = 2.7255\,\mathrm{K}$ and $N_\mathrm{eff} = 3.045$ effective relativistic neutrinos species at high redshift.

Identifier	L	N_c	m_c [$\mathrm{M}_\odot$]	m_g [$\mathrm{M}_\odot$]	h	Ω_m	Ω_c	Ω_b	$\sum m_\nu$	σ_8	$10^9 A_\mathrm{s}$	n_s
DES3_L2800	2.8 Gpc	5040^3	5.65×10^9	1.07×10^9	0.681	0.306	0.256	0.0486	0.06 eV	0.807	2.099	0.967
DES3_L1000	1.0 Gpc	1800^3	5.65×10^9	1.07×10^9	0.681	0.306	0.256	0.0486	0.06 eV	0.807	2.099	0.967
Planck	1.0 Gpc	1800^3	5.72×10^9	1.07×10^9	0.673	0.316	0.265	0.0494	0.06 eV	0.812	2.101	0.966
PlanckM240Var	1.0 Gpc	1800^3	5.67×10^9	1.06×10^9	0.662	0.328	0.271	0.0510	0.24 eV	0.772	2.109	0.968
PlanckM240Fix	1.0 Gpc	1800^3	5.62×10^9	1.07×10^9	0.673	0.316	0.261	0.0494	0.24 eV	0.769	2.101	0.966
LS8	1.0 Gpc	1800^3	5.65×10^9	1.07×10^9	0.682	0.305	0.256	0.0473	0.06 eV	0.760	1.836	0.965
DES3_L2800_DMO	2.8 Gpc	5040^3	6.72×10^9	–	0.681	0.306	0.256	0.0486	0.06 eV	0.807	2.099	0.967
DES3_L1000_DMO	1.0 Gpc	1800^3	6.72×10^9	–	0.681	0.306	0.256	0.0486	0.06 eV	0.807	2.099	0.967
Planck_DMO	1.0 Gpc	1800^3	6.78×10^9	–	0.673	0.316	0.265	0.0494	0.06 eV	0.812	2.101	0.966
PlanckM240Var_DMO	1.0 Gpc	1800^3	6.73×10^9	–	0.662	0.328	0.271	0.0510	0.24 eV	0.772	2.109	0.968
PlanckM240Fix_DMO	1.0 Gpc	1800^3	6.68×10^9	–	0.673	0.316	0.261	0.0494	0.24 eV	0.769	2.101	0.966
LS8_DMO	1.0 Gpc	1800^3	6.72×10^9	–	0.682	0.305	0.256	0.0473	0.06 eV	0.760	1.836	0.965

the accuracy of its large-scale structure predictions (Part II of this thesis). In another departure from its predecessors, the subgrid physics parameters were systematically calibrated by training emulators to predict key astrophysical quantities (the galaxy stellar mass function and cluster gas fractions) and comparing directly with observations [362]. Combined with the unprecedented volume of the simulations, these improvements make FLAMINGO ideal for precision cosmology applications.

The flagship simulation contains $N_\mathrm{c} = N_\mathrm{b} = 5040^3$ dark matter and baryon particles and $N_\nu = 2800^3$ massive neutrino particles in a periodic $(2.8\,\mathrm{Gpc})^3$ volume. In addition, there are many simulations with the same mass resolution in a $(1\,\mathrm{Gpc})^3$ volume. These simulations span a range of subgrid physics and cosmological parameter variations. We only use the cosmological variations in this paper and restrict to the fiducial subgrid model. For each hydrodynamical simulation, there is a gravity-only counterpart that treats dark matter and baryons as a single cold fluid, but still includes the effects of massive neutrinos (also called "dark matter only" or DMO). The hydrodynamical simulations use a new subgrid physics model, which includes improved prescriptions for gas cooling, star formation, black holes, and AGN feedback [359, 363–365]. The simulations were run with the SWIFT cosmological hydrodynamics code [76, 190] on the COSMA-8 facility in Durham, using the SPHENIX flavour of SPH [366], which is optimized for galaxy formation applications. The initial conditions were generated with third-order Lagrangian perturbation theory (3LPT) at $z = 31$ with separate transfer functions for dark matter, baryons, and neutrinos [228, 257], using the MONOFONIC [89, 261] and FASTDF [186] codes, with transfer functions computed with CLASS [128]. Neutrinos were implemented with the δf method [237], which minimizes shot noise without neglecting the nonlinear evolution of the neutrino phase-space distribution. The relevant parameters for the simulations used in this paper are listed in Table 8.1.

8.3. Results

We begin by studying the effects of neutrinos and baryons on global large-scale structure statistics in Section 8.3.1 and then study the origin of non-factorizable corrections through the lens of individual halo properties in Section 8.3.2. Finally, we construct a simple model for the non-factorizable correction to the power spectrum and consider its cosmological implications in Section 8.3.3.

8.3.1. Global statistics

We begin by looking at two fundamental global statistics relevant for large-scale structure studies: the matter power spectrum $P_\mathrm{m}(k)$ and the halo mass function $f(M) = \mathrm{d}n(M)/\mathrm{d}\log_{10}(M)$. To untangle the effects of neutrinos and baryons, we look at four simulations: the hydrodynamical Planck and PlanckM240Fix simulations, which differ only in their assumed massive neutrino contents of respectively $\sum m_\nu = 0.06\,\mathrm{eV}$ and $0.24\,\mathrm{eV}$, and their gravity-only counterparts. We show the corresponding power spectrum and halo mass function ratios in Fig. 8.1. In each case, we compare to one of the gravity-only Planck simulations. The left-hand panel shows the effects on the power spectrum and the right-hand panel shows the effects on the halo mass function.

Let us begin with the power spectrum. The basic effects of neutrinos and baryons have been studied before [e.g. 148, 152, 340]. The effect of increasing the neutrino mass, $F_\nu(k) = P_\mathrm{m}^{0.24\,\mathrm{eV}}(k)/P_\mathrm{m}^{0.06\,\mathrm{eV}}(k)$, is shown in red. On large scales, the ratio $F_\nu \to 1$. This is because on scales larger than the free-streaming length, neutrinos and dark matter cluster in the same way, barring relativistic effects. On scales $0.01\,\mathrm{Mpc}^{-1} \leq k \leq 0.6\,\mathrm{Mpc}^{-1}$, the ratio decreases to a minimum of $\Delta P/P \approx 10\Delta f_\nu \approx 14\%$ before turning over and increasing to around 10% on sub-Mpc scales. This happens because neutrinos suppress structure formation, which affects the 1-halo and 2-halo terms differently [197, 244].

The effect of baryonic feedback, $F_\mathrm{b}(k) = P_\mathrm{m}^\mathrm{hydro}(k)/P_\mathrm{m}^\mathrm{DMO}(k)$, obtained by dividing the hydrodynamical result with $\sum m_\nu = 0.06\,\mathrm{eV}$ by the corresponding gravity-only result, is shown in blue. This ratio equals unity on large scales up to about $k = 0.6\,\mathrm{Mpc}^{-1}$ and then decreases to a dramatic minimum at $k = 6\,\mathrm{Mpc}^{-1}$, before turning over. This happens primarily because feedback from AGN expels gas from $10^{14}M_\odot$ mass halos, which lowers their contribution at these scales [340, 342, 367]. On scales smaller than $k = 20\,\mathrm{Mpc}^{-1}$, the ratio is greater than 1 because gas cooling and star formation allows halos to contract, increasing the density on small scales [340, 367].

Now, let us consider the combined effect of massive neutrinos and baryonic feedback, $F_{\nu\mathrm{b}}(k) = P_\mathrm{m}^{\mathrm{hydro},0.24\,\mathrm{eV}}(k)/P_\mathrm{m}^{\mathrm{DMO},0.06\,\mathrm{eV}}(k)$. The true result is obtained by dividing the hydrodynamical PlanckM240Fix result by the gravity-only Planck result. This curve is shown in orange. The result is well approximated by the product of the individual effects: $F_{\nu\mathrm{b}} \approx F_\nu F_\mathrm{b}$, shown in purple. However, we clearly see that the true effect is stronger than would be expected from this approximation. The non-factorizable correction, $F_{\nu\mathrm{b}}/(F_\nu F_\mathrm{b})$ is shown in the bottom panel. It vanishes on large scales, where only neutrinos have an effect, but approaches -1% between $1\,\mathrm{Mpc}^{-1} \leq k \leq 20\,\mathrm{Mpc}^{-1}$ before increasing beyond $+1\%$ at very small scales. In the following subsections, we will try to understand the source of this effect. However, for now it is interesting to observe that the true suppression

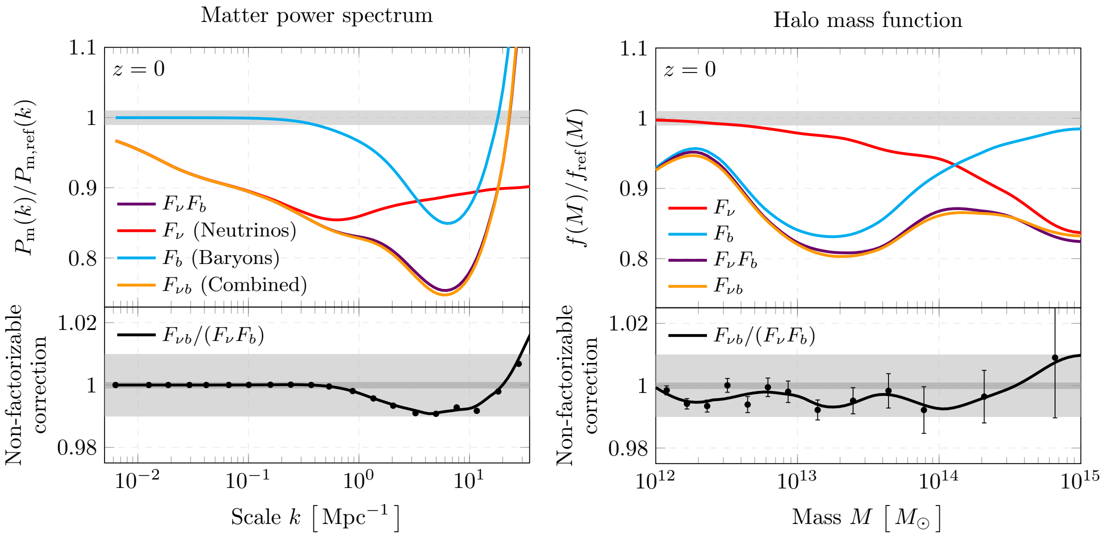

Figure 8.1: The effects of baryons (F_b, blue) and neutrinos (F_ν, red) on the matter power spectrum (left) and halo mass function (right). If the combined effect ($F_{\nu b}$, orange) were perfectly factorizable, it would equal the product $F_\nu F_b$ (purple). The bottom panels show the non-factorizable correction, $F_{\nu b}/(F_\nu F_b)$, to this approximation. It can be seen that baryonic effects are enhanced by increasing the neutrino mass. The simulated data were interpolated to a logarithmically spaced grid and smoothed with a 4th-order Savitzky-Golay filter. We show Poisson error bars for the halo mass function. Errors are much smaller for the power spectrum.

is stronger than would be expected from the factorizable approximation. This means that the baryonic effect is stronger for the model with the larger neutrino mass[1].

Next, let us consider the neutrino effect F_ν on the halo mass function [173, 243], defined in the same way as in the power spectrum case and shown in red in the right-hand panel of Fig. 8.1. The effect is very small for low mass halos, but increases towards larger halo masses. The number density of the most massive clusters is significantly suppressed by $\sim 20\%$. This happens because, while the normalization of density fluctuations $\sigma(M)$ decreases along with the power spectrum, the halo mass function remains approximately universal [242–244]. As a result, the effect is most pronounced around the exponential cut-off at high masses. Similar behaviour is seen when changing σ_8. In terms of the excursion set formalism [368, 369], the probability of rare peaks collapsing is very sensitive to a change in the normalization, but the number of small peaks is a power law that depends less sensitively on the normalization.

The baryonic effect, F_b, is more complex [370, 371]. The effect is very small for the

---

[1] Or equivalently that the neutrino effect is stronger for the hydrodynamical simulations.

most massive clusters, as these have a sufficiently deep gravitational potential well to prevent feedback from expelling a significant fraction of the gas. The effect increases as we move towards smaller halo masses, reaching a maximum suppression around a few times $10^{13} M_\odot$. For smaller halos than this, AGN feedback becomes less important and therefore the baryonic effect becomes smaller. The effect reaches another extremum at a few times $10^{12} M_\odot$, where supernova feedback starts to be able to overcome the gravitational potential well, driving the ratio back down for lighter halos [371].

When we consider the combined effect, $F_{\nu\mathrm{b}}$, and its factorizable approximation, $F_\nu F_\mathrm{b}$, we reach a similar conclusion as for the power spectrum, although the statistics are worse and the result is therefore not as clear cut. The non-factorizable correction appears to be largest at the extrema of the combined effect and works in the direction of making the feedback stronger for the model with the larger neutrino mass. The corrections to the halo mass function are slightly smaller than the corrections to the power spectrum, but still approach 1% at the maxima.

8.3.2. Halo properties

To gain insight into the origin of the non-factorizable corrections seen in the last section, we turn our attention to a simpler problem: the baryonic effect on individual halo masses. By matching the 10 most strongly bound particles of each field halo in the hydrodynamical and gravity-only versions of the same simulation, we determine the properties of the same halo with and without baryonic effects. The aim is to use halo properties as proxies for the different ways in which feedback might depend on cosmology. We do this without varying the cosmological model, instead relying on the inherent scatter in halo properties in the large DES3_2800 simulations.

We are interested in the baryonic effect on the halo mass, for which we define the ratio

$$\psi(M_\mathrm{DMO}) = \frac{M_\mathrm{hydro}}{M_\mathrm{DMO}}, \tag{8.3.1}$$

using M_200m masses. These correspond to the total mass contained in a spherical region with an average density equal to 200 times the mean matter density of the Universe. We obtain qualitatively similar results for different mass definitions, but opt for M_200m because of its relatively large aperture, producing smoother results. We exclusively consider field halos, ignoring satellites.

The baryonic effect, ψ, depends sensitively on the gravity-only mass. We therefore split the sample into bins of M_DMO. Within each bin, we rank the halos according to a secondary halo property and compute ψ for the five quintiles (i.e. equal-sized groups with rank between $[0, 0.2]$ up to $[0.8, 1]$). We considered a large array of possible halo properties, but selected the following four to characterize each of the channels discussed in the introduction:

1. *Concentration* c. We use a proxy for the halo concentration $c = V_{\max}/V$, based on the maximum circular velocity $V_{\max}$ and virial velocity $V = \sqrt{GM/R}$ where M and R are the mass and radius of the spherical overdensity [351].
2. *Formation epoch* a_x. The scale factor time a at which the halo first accreted $x\%$ of its present-day mass, computed by linearly interpolating between adjacent snapshots.
3. *Environmental density* δ_r. We characterize the environment by computing the total mass M_r enclosed by halos within a radius of $r\,$Mpc, excluding the halo itself. The environmental density δ_r is then defined as $\delta_r = M_r/\langle M_r \rangle - 1$, where $\langle M_r \rangle$ is the average value for all halos in the sample.
4. *Baryon fraction* f_{b}. This is the only property that is computed from the hydrodynamical simulation, rather than the gravity-only version. It is the mass fraction of baryons in the spherical overdensity.

The resulting $\psi(M_{\mathrm{DMO}})$ curves, split into quintiles of the secondary halo property, are shown in Fig. 8.2. First of all, we note that the general trend resembles that of the baryonic effect on the halo mass function seen in the right-hand panel of Fig. 8.1. It arises from the interplay between the depth of the gravitational potential well and the strength of stellar and AGN feedback at those masses.

Let us first consider the effect of the halo concentration c, shown in the top left panel of Fig. 8.2. There is a clear dependence of ψ on concentration. For halos with dark matter masses between $10^{12} M_\odot$ and $10^{13} M_\odot$, the suppression is smaller for less concentrated halos, but the trend reverses for $M > 10^{13} M_\odot$. This can be interpreted as follows. Halos with lower concentrations have smaller binding energies, which shifts the balance in favour of supernova and AGN feedback compared to gravitational infall. As a result, the curves shift to larger masses. Although concentration is anti-correlated with formation epoch, older halos being more concentrated on average [e.g. 372], the effect seen is not due to a dependence of feedback on formation history, at least for masses $M > 10^{13} M_\odot$.

To see this, consider the dependence on the formation epoch shown in the top right panel for a_{50}. In this case, we additionally fix the maximum circular velocity, $V_{\max}$, to the 40%-60% interpercentile range to reduce the correlation between c and a_{50}. For halos

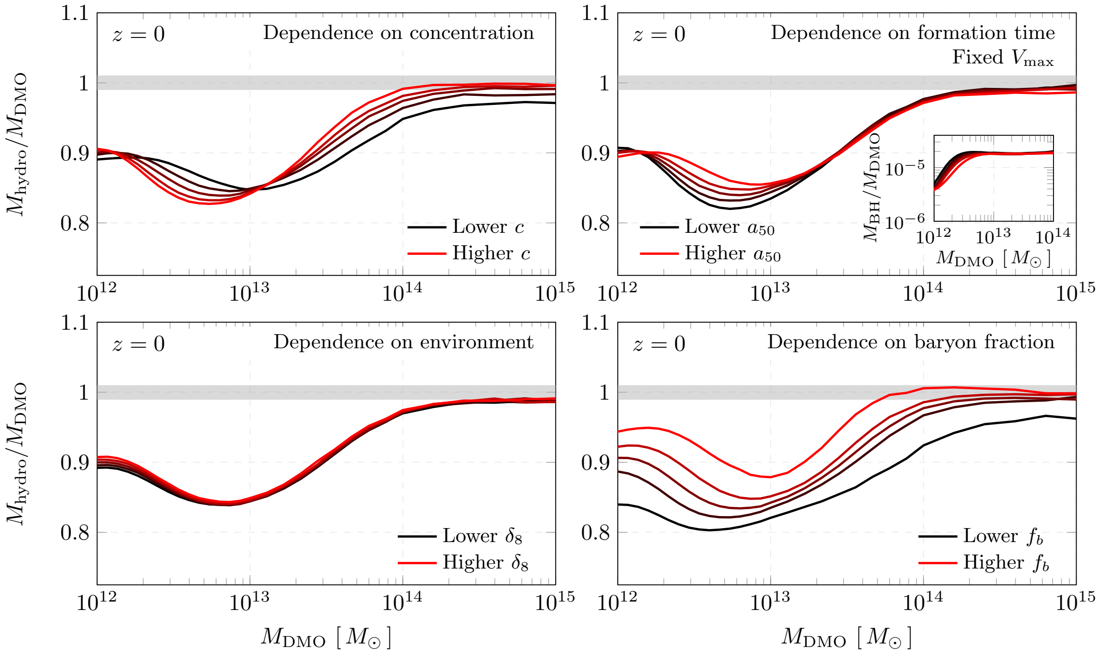

Figure 8.2: A closer look at the baryonic effect on halo mass $M_{\rm 200m}$, comparing the masses of matched field halos in the corresponding gravity-only and hydrodynamic simulations. We show the dependence of the mass ratio on four secondary halo properties: concentration (top left), formation time (top right), environment (bottom left), and baryon fraction (bottom right). The colours indicate quintiles of the secondary halo property, calculated within each mass bin, with black (red) indicating a lower (higher) value of that property.

$M < 10^{13} M_\odot$, the suppression does depend on formation epoch. In this mass range, the masses of older halos are more suppressed. However, there is no dependence on a_{50} for halos beyond $10^{13} M_\odot$. This shows that the dependence seen for the concentration is due to the binding energy and not the formation epoch. To understand this, we show the black hole mass relative to the gravity-only mass, $M_{\rm BH}/M_{\rm DMO}$, in the inset graph. We see that halos with masses between $10^{12} M_\odot$ and $10^{13} M_\odot$ fall into a critical transition range where nonlinear black hole growth is on the cusp of being triggered [355], depending on the mass and formation epoch. At the high mass end, all black holes are self-regulating and attain a nearly fixed fraction of the dark matter mass. From this point on, the dependence on formation epoch essentially disappears.

The third property is the environmental density, δ_8, defined in terms of the mass contained by halos within 8 Mpc. The bottom left panel of Fig. 8.2 shows that this property is barely correlated with the baryonic suppression of halo mass. The suppression is very slightly smaller for halos in denser regions, particularly towards the low mass end. As halos in dense regions tend to be more concentrated, a phenomenon known as assembly bias [356, 357], this could be due to the indirect effect of the concentration. Although we do not consider satellites, the behaviour at the small mass end may also be an artefact of overlapping spherical overdensities. In any case, the environment does not appear to play a major role in regulating feedback.

Finally, we consider the dependence on the baryon fraction f_{b}, computed directly from the hydrodynamical simulation. Over the entire mass range, halos with higher baryon fractions have a smaller suppression. This is consistent with the assumption that the suppression is primarily driven by baryonic outflows. Indeed, the suppression of the power spectrum can be related directly to the reduction in the baryon fraction in groups and clusters [347, 373]. The dependence is strongest at the low mass end, but even at the high mass end, the effect of the baryon fraction is larger than that of the concentration.

To quantify the importance of a given halo property θ for regulating the baryonic effect on the mass ψ, we compute the correlation coefficient

$$\rho_\theta = \frac{\langle \delta_\psi \delta_\theta \rangle}{\sqrt{\langle \delta_\psi^2 \rangle \langle \delta_\theta^2 \rangle}}, \tag{8.3.2}$$

where $\delta_X = X - \langle X \rangle$ are centred variables. A value of $\rho_\theta = 0.1$ implies that a 1σ change in θ results in a 0.1σ change in ψ. To determine the overall importance of this property when explaining the dependence on cosmology, we also require the change in that property when we change cosmological parameters $\boldsymbol{\alpha}$ to $\boldsymbol{\alpha} + \Delta\boldsymbol{\alpha}$:

$$\Delta\theta = \frac{\big\langle \theta(\alpha + \Delta\boldsymbol{\alpha}) \big\rangle - \big\langle \theta(\boldsymbol{\alpha}) \big\rangle}{\sqrt{\big\langle \delta_\theta^2(\boldsymbol{\alpha}) \big\rangle}}, \tag{8.3.3}$$

normalized by the standard deviation of θ at $\boldsymbol{\alpha}$. In this way, we rank each halo property along two dimensions: their importance for feedback and their sensitivity to cosmology, with the overall impact being the product of the two. We did this for a large array of halo properties, including the four properties studied above, for the shift from the Planck cosmology with $\sum m_\nu = 0.06\,\mathrm{eV}$ to the PlanckM240Fix model with $\sum m_\nu = 0.24\,\mathrm{eV}$. In all cases, the change in baryon fraction was found to be most important. At the low mass end, the early formation epochs a_{25} and a_{50} followed the baryon fraction in relevance. At the high mass end, the halo concentration c was the second most important property. Since the baryonic suppression of the matter power spectrum is mostly determined by halos with masses $10^{13} M_\odot < M < 10^{14} M_\odot$ on scales $k < 10 h\,\mathrm{Mpc}^{-1}$ [343], we conclude

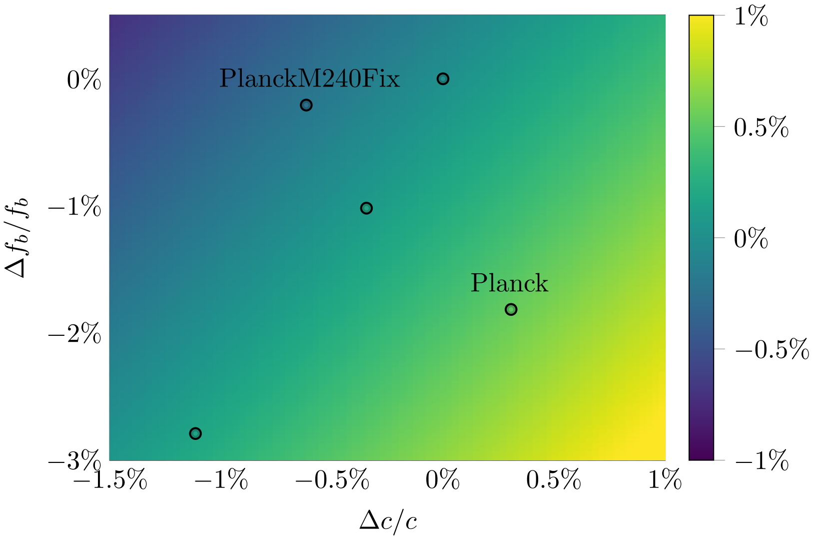

Figure 8.3: We can predict the non-factorizable correction to the power spectrum from the changes in baryon fraction and halo concentration. The points are based on the different $L = 1\,\mathrm{Gpc}$ simulations listed in Table 8.1, shown relative to the fiducial DES3_L1000 model. We evaluate the matter power spectra at $k = 10h\,\mathrm{Mpc}^{-1}$, where the corrections first become important.

that a model of the non-factorizable corrections should first account for the change in baryon fraction and halo concentration.

8.3.3. Predicting non-factorizable corrections

Based on the insights gathered from the DES3_2800 simulation, we now turn to the five cosmological variations in the FLAMINGO suite, listed in Table 8.1. We wish to explain the non-factorizable correction to the matter power spectrum in terms of the change in the universal baryon fraction, Δf_b, and the shift in the mass-concentration relation Δc:

$$\Delta\big(P_\mathrm{m}^\mathrm{hydro}/P_\mathrm{m}^\mathrm{DMO}\big) = a_b\frac{\Delta f_b}{f_b} + a_c\frac{\Delta c}{c}. \tag{8.3.4}$$

We evaluate the present-day power spectrum at $k = 10h\,\mathrm{Mpc}^{-1}$, where the non-factorizable corrections are prominently seen in Fig. 8.1. The mass-concentration relation is computed from the hydrodynamic simulations, using the velocity-based proxy $c = V_\mathrm{max}/V_{200}$. We take the median of the fractional change in $\Delta c/c(M)$ for all bins with $M > 10^{12}M_\odot$. The change in the baryon fraction, $\Delta f_b/f_b$, is simply taken from the cosmological parameters. The resulting fit is shown in Fig. 8.3. As expected, an increase

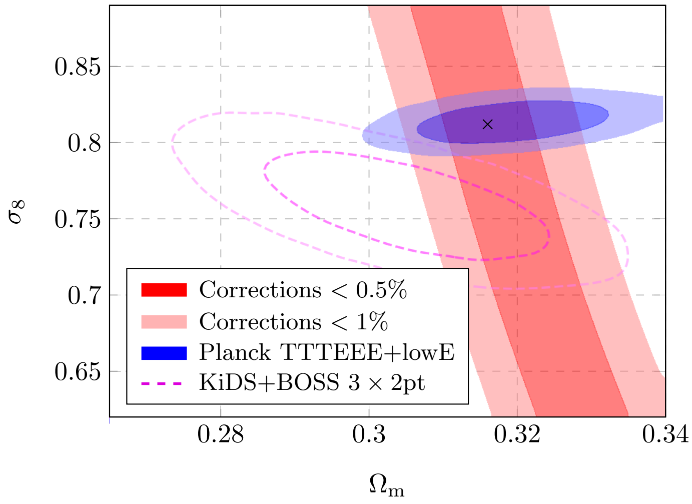

Figure 8.4: As a first application of the model, we show the cosmological parameter range around the nominal Planck-based cosmology (marked by a cross) for which non-factorizable corrections are within 0.5% and 1% (red), along with the Planck TTTEEE + lowℓ + lowE + lensing + BAO constraints (blue).

in the baryon fraction leads to a greater baryonic suppression ($a_b = -0.22 \pm 0.02$), as a larger gas reservoir is available for halos of a given mass. Also in line with expectation, an increase in the concentration leads to a smaller baryonic suppression ($a_c = 0.40 \pm 0.06$), which we attribute to the increased binding energy. Although it could not be included here, we have developed an analytical model in which the strength of AGN feedback is a function of f_b/c_v^2 (in prep.), consistent with the finding here that $a_c \approx -2a_b$. The two models compared in Fig. 8.1, Planck and PlanckM240Fix, are indicated in Fig. 8.3. We see that the non-factorizable correction between these two models is large for two reasons. First of all, halos are less concentrated in the large neutrino mass cosmology because the matter density is lower when halos collapse. Secondly, the baryon fraction is higher relative to the cold dark matter density, because Ω_m and Ω_b are fixed while Ω_ν is larger. Both changes increase the potency of baryonic feedback.

Although the model is statistically significant ($p = 0.00466$) and physically motivated, it is based on a limited number of simulations. Hence, the model should be tested on a wider sample before any definitive conclusions are drawn. This work is currently ongoing. For now, let us briefly consider the implications, assuming that the model holds more generally. As a proof of concept, we use (8.3.4) to predict the cosmological parameter range within which the non-factorizable corrections are smaller than 1%. To do this, we use an emulator for the mass-concentration relation [352, 374] and convert the output to our velocity-based proxy with the relation from [351]. Doing so, we predict the shift

$\Delta c/c$ in concentration as a function of Ω_m and σ_8. Since Ω_b is much better determined than Ω_m, we hold the former fixed at the fiducial Planck value [112], which gives $\Delta f_b/f_b$. Applying (8.3.4), we then obtain the red contours shown in Fig. 8.4. They indicate the parameter range where corrections are below 0.5% (dark) and 1% (light red). The contours are mostly parallel to lines of constant Ω_m, reflecting the importance of the baryon fraction, while the tilt is due to the dependence on halo concentration.

Also shown in Fig. 8.4 are the constraints from Planck temperature and polarization data (TTTEEE + lowE) in blue and the KiDS + BOSS 3 × 2pt galaxy clustering and weak lensing constraints are indicated by purple dashed lines. The model used by the Planck simulation in this paper is marked by a cross. The figure shows that the non-factorizable corrections are mostly below 1% if one restricts to the 68% Planck constraints. However, to sample a parameter space that covers both the constraints from Planck and large-scale structure surveys like KiDS and BOSS with 1% precision is not possible unless non-factorizable corrections are taken into account.

8.4. Discussion

In this paper, we presented early results from the new FLAMINGO hydrodynamical simulations project [359]. We studied the effects of massive neutrinos and baryons on global large-scale structure statistics, with a particular focus on non-factorizable corrections that arise from variations in cosmology in the presence of baryonic physics. We confirmed prior findings [345] that non-factorizable corrections to the matter power spectrum are at the percent level on nonlinear scales, $1\,\mathrm{Mpc}^{-1} \leq k \leq 10\,\mathrm{Mpc}^{-1}$. Turning to the baryonic effect on individual halo masses, we then studied the dependence of the baryonic suppression on secondary halo properties, such as the formation epoch, and determined that the baryon fraction and halo concentration or binding energy are most important in the halo mass range that is most relevant for the suppression of the power spectrum between $1\,\mathrm{Mpc}^{-1} \leq k \leq 10\,\mathrm{Mpc}^{-1}$.

Based on this insight, we constructed a simple bilinear model to explain the correction to the power spectrum, ΔP_m, in terms of the cosmological change in halo concentration, $\Delta c/c$, and baryon fraction, $\Delta f_b/f_b$. As a first application of the model, we determined the cosmological parameter range for which the results obtained assuming a nominal Planck cosmology can be applied to within 1% accuracy on nonlinear scales. The results indicate that non-factorizable corrections cannot be ignored, at this level of accuracy, if one wishes to sample a parameter space that covers both the results from large-scale structure surveys and CMB observations (Fig. 8.4). There are several promising ways forward. One possibility is to construct emulators that cover both changes in cosmology and feedback. However, our findings suggest that the dependence on cosmology may be modelled with a limited number of parameters. Non-factorizable corrections could therefore also be included in alternative analytic and semi-analytic approaches.

Finally, it is interesting to note that we find baryonic suppression to be stronger in models in which structure formation is already suppressed, such as in models with massive neutrinos. Although the effect we see is limited to scales that are too small to affect the S_8 parameter, this finding lends some credence to the idea that a combination of feedback and a mechanism to suppress structure formation could alleviate the S_8 tension. It is worth investigating whether alternative implementations of baryonic feedback combined with non-factorizable corrections from neutrinos can lead to a nonlinear resolution of the S_8 tension [375].

Outlook

9

9.1. Towards a neutrino mass detection

The large-scale structure of the Universe provides an ideal testing ground for low-energy neutrino physics. By detecting a gravitational signature of the cosmic neutrino background, galaxy surveys like DESI and Euclid could enable the first measurement of the sum of neutrino masses, $\sum m_\nu$, providing a unique window on the properties of neutrinos beyond the Standard Model, such as their lifetime [376–380] and the mechanism that generates their masses [30–32], as well as cosmology itself [26, 27].

Any cosmological constraint should be confronted with information from the laboratory. Neutrino oscillation experiments have measured the mass squared differences at the percent level [28, 29], while β-decay experiments have achieved constraints on the absolute mass scale with sub-eV precision [14, 15]. Together, these experiments imply that $0.058\,\mathrm{eV} < \sum m_\nu < 2.4\,\mathrm{eV}$. By comparison, some recent cosmological analyses yield bounds of $\sum m_\nu < 0.09\,\mathrm{eV}$ [108, 109], approaching the lower bound from neutrino oscillations. If the trend continues (see Fig. 9.1), aided by larger surveys and better techniques, we will soon either measure the sum of neutrino masses or find evidence for non-standard neutrino physics or cosmology.

This last possibility is worth further consideration. While terrestrial experiments are model-independent, the same is not true for the cosmological approach. The ΛCDM model has been successful in reproducing many observations, but a number of puzzling anomalies and tensions between observations persist, despite wide-ranging efforts to address them [334, 335, 381]. To obtain a robust measurement from cosmology, the need to combine disparate datasets for additional constraining power must be balanced against the increased dependence on cosmological modelling.

Cosmological simulations will play a key role in this balancing act. In Part II of this thesis, we developed new techniques and demonstrated that neutrino simulations can now be performed with the required level of precision. Building on these developments by exploiting a new generation of accurate simulations, we are continuing on the road towards a robust neutrino mass detection with work in several directions, from the development of novel probes to the treatment of uncertainties. We will discuss these in further detail below.

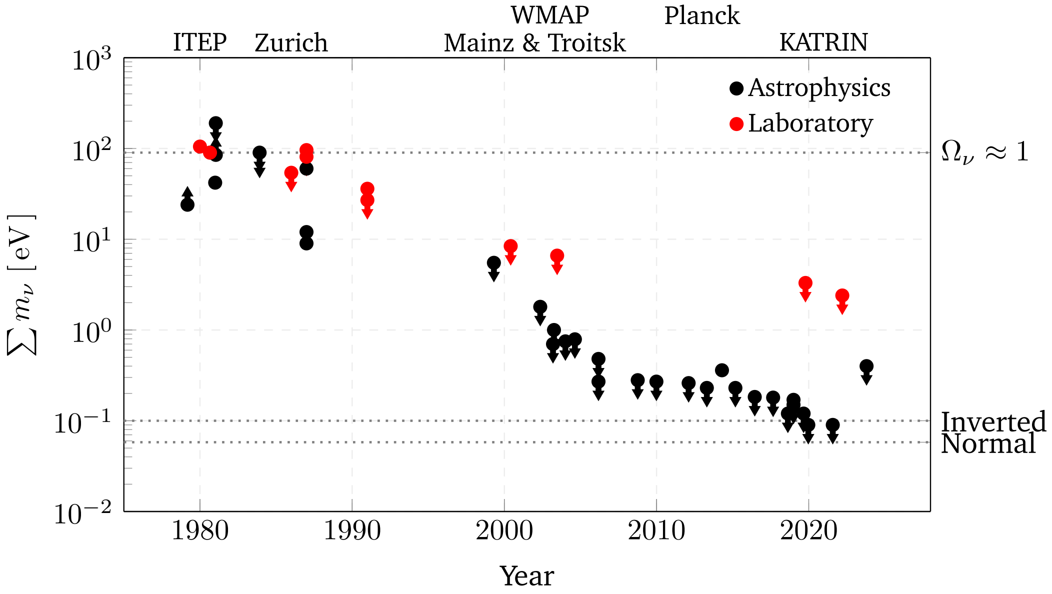

Figure 9.1: An informal census of published neutrino mass results, with early experimental findings based on [382] and extrapolated to three massive species. Upper bounds are indicated by downward arrows. The first upward arrow indicates the 1979 Tremaine & Gunn bound that $m_\nu \geq 24\,\mathrm{eV}$ if one species of neutrinos forms the dark matter in halos [48]. The double downward arrow indicates the finding that a neutrino-dominated Universe poses a challenge to the observed large-scale structure [24, 50]. The lower dotted lines indicate the lower bounds from neutrino oscillations [28, 29], assuming the inverted or normal mass ordering, respectively.

9.1.1. Novel neutrino probes

Simulations are ideal for developing novel neutrino probes. These are important not just for providing additional constraining power in the non-linear régime, where two-point statistics no longer capture all available information, but also to provide robustness. Two strategies can be distinguished. The first is to develop novel summary statistics that target large-scale structure in general, providing information about neutrinos primarily by breaking parameter degeneracies within the $\nu\Lambda$CDM model. Such statistics include the bispectrum [239–241], the marked power spectrum [383], density-split clustering [384, 385], the wavelet-scattering transform [386], and void statistics [387–389]. Analytical treatments are often untractable, which means that simulations are indispensable.

An alternative strategy is to search for signatures that are somehow unique to massive neutrinos. While more challenging, the pay-off would be immense if such signals were large enough, enabling an unambiguous detection. Work in this direction is also of general interest, as it improves our understanding of neutrino phenomenology. Some examples in this category include the possible effects of dynamical friction [198] and torque [199] from neutrinos on dark matter halos or the relative velocity between neutrinos and dark matter [390]. Analytical work that appears promising should be checked against simulations. Studies in this category would benefit especially from high-fidelity particle neutrino simulations.

Let us mention just a few more possibilities for future work. If neutrinos measurably affect dark matter halos, such a signal might show up in the kinematics of galaxies around clusters, given that effects like dynamical friction depend on the mass of the perturber. The velocity distribution of galaxies can be probed with the redshift space galaxy-cluster cross-correlation function [391], a statistic that should be simpler to model on small scales than galaxy-galaxy clustering [114, 392–394]. This probe would also be sensitive to neutrinos through the growth rate of structure and has already been used to constrain modified gravity [395]. Large-volume neutrino simulations, containing sufficient numbers of clusters, could be used to isolate the direct and indirect effects of massive neutrinos.

Another exciting possibility is to search for extended neutrino halos around clusters with weak lensing. Previous work, which assumed idealized spherical halo profiles, suggested that the effects may be detectable by Euclid at the 1%-level [193]. Now that Euclid has launched, more than ten years on, it is timely to revisit the problem. Using the same large-volume simulations as above, the signal could be studied for realistic stacked cluster profiles. Although we briefly looked at simulated neutrino halos in Chapter 4, a more detailed analysis of the weak lensing signal remains outstanding.

9.1.2. Uncertainty quantification

Simulations are also useful for the end-to-end modelling of large-scale structure observations in the form of mock catalogues that account for non-linear structure and galaxy formation, survey geometry, and systematic errors. Mock catalogues are crucial to forecast the performance of large-scale structure surveys, to analyse observational strategies, to identify flaws in analysis pipelines, to quantify uncertainties, and to provide confidence in any potential discovery. Realistic neutrino simulations are particularly important to study the impact of systematics on cosmological neutrino mass constraints.

Due to non-trivial correlations among cosmological and nuisance parameters, systematics can bias parameters in unexpected ways. One example is the CMB lensing anomaly, an oscillatory residual in the temperature power spectrum, present in many CMB analyses, that resembles the smoothing effect of gravitational lensing and has been parametrized by the nuisance parameter A_L [396–398]. Marginalising over A_L to account for possible systematics has a profound effect on the inferred neutrino mass [113, 115–117]. Similar correlations may exist for systematics in large-scale structure observations. Determining the impact on the inferred neutrino mass requires the end-to-end propagation of systematics from observation space to cosmological parameter space.

This mapping from observation space to cosmological parameter space, known as cosmological parameter estimation, amounts to the reduction of complex astronomical datasets to a number of point estimates that describe the cosmological model, together with a quantitative description of the uncertainty. This is traditionally accomplished by compressing observations to summary statistics and by assuming a Gaussian likelihood function with errors given by the covariance matrix. Under certain conditions, covariances can be estimated directly from the data using boostrap or jackknife methods [399–403]. Covariance matrices can also be computed analytically [404–408], but this is difficult for novel statistics in the non-linear régime. Alternatively, covariance matrices can be computed from ensembles of N-body simulations [409, 410].

Recently, simulation-based inference has emerged as an interesting alternative to the traditional approach based on explicit likelihoods [411–416]. These methods combine the power of simulations for the accurate forward modelling of observations with machine learning techniques to train models for the posterior distribution of the cosmological parameters. By relying on forward modelling, these methods are ideal for the propagation of systematic effects to cosmological parameter space. As such, they are a promising way forward to extract cosmological information from large-scale structure observations.

9.1.3. Next-generation simulations

To facilitate the applications described above, we have proposed an ambitious suite of next-generation neutrino simulations, using the methods and techniques developed in this thesis to ensure sub-percent accuracy in its large-scale structure predictions. In order to resolve the dark matter halos hosting all main tracer galaxies of the DESI survey [16] in a volume large enough to represent the DESI footprint and to capture large-scale neutrino effects, the volume and resolution of the simulations will significantly exceed that of existing neutrino simulation suites like Quijote [160] and DEMNUni [159]. A training data set, including a ladder of simulations varying the neutrino masses in small steps and a grid of cosmological parameter variations, will be used to develop and test novel neutrino probes and for emulation and machine learning applications. Some simulations will use modified initial conditions as part of a controlled experiment to identify effects unique to neutrinos. A further set will assume a blinded neutrino cosmology as part of a mock challenge to stress test analysis pipelines.

These high-fidelity simulations will be complemented with a large number of low-fidelity simulations for uncertainty quantification and to expand the effective volume of the suite with the use of control variates [91]. Generating the large numbers of realizations required for these applications implies the need for highly optimized codes [78, 79] or approximate methods [90, 417, 418]. To handle the case with massive neutrino, we developed a fast N-body code, called SEDULUS, that incorporates all techniques needed for robust neutrino simulations, as well as other advances for speed-ups compared to existing alternatives. Together, the high-fidelity and low-fidelity simulations will provide the essential theoretical infrastructure needed to measure the sum of neutrino masses with DESI and other Stage IV galaxy surveys.

9.2. Other directions

The work in Part III of this thesis also raises interesting questions for future work. In Chapter 7, we made detailed predictions for the rate of direct relic neutrino detections, assuming the specifications of the proposed PTOLEMY experiment [270–272]. We also analysed angular anisotropies in the neutrino background, but without reference to a specific detector. In principle, an experiment like PTOLEMY could measure CNB anisotropies using polarized tritium [287]. Hence, it is worth following this work up with a specific application to such an experiment. More generally, it is interesting to explore what could be learnt from a measurement of the neutrino power spectrum, be it through direct detection or indirect cosmological inference.

In Chapter 8, we studied the dependence of astrophysical feedback processes on cosmology, finding that this coupling gives rise to non-factorizable corrections to the matter power spectrum that are not negligible in the cosmological parameter space spanned by current observations. Our results show that baryonic feedback is stronger when combined with a mechanism that suppresses structure formation or removes dark matter from group-sized halos at late times, both increasing the baryon fraction and decreasing the halo concentration. This suggests that a combination of baryonic and non-baryonic suppression effects might alleviate the present tension between large-scale structure and CMB measurements of S_8 [336–338]. Hence, it is interesting to study extensions of ΛCDM that produce such effects. Besides replacing some fraction of cold dark matter with hot dark matter like neutrinos, a model in which dark matter decays into dark radiation [419–422] appears promising, but annihilating dark matter [423] or a Yukawa coupling between dark matter and quintessence [424] might also produce such an outcome.

9.3. Closing remarks

Confirming the existence of the cosmic neutrino background and measuring its fundamental properties has been a long-standing challenge for cosmology and particle physics. While their effects are subtle, neutrinos leave a variety of cosmological traces, which could be probed by observations of the Universe at different epochs and scales. The major challenge in obtaining a robust neutrino mass measurement is the need to combine and interpret observations consistently within a cosmological model. In this thesis, we have developed methods and techniques that enable the accurate modelling of neutrino effects on late-time cosmological observables. These will be used in the next generation of neutrino simulations whose aim is to facilitate the optimal use of galaxy observations in our quest to reveal these most elusive messengers from the early Universe.

Bibliography

[1] C. Rovelli, *The first scientist: Anaximander and his legacy*, Westholme Publishing (2011).

[2] G.E. Idang, *Thales, Anaximander and Anaximenes as Pathfinders of Modern Science*, *International Journal of Philosophy* **1** (2013) .

[3] K. Kalachanis, E. Panou, E. Theodossiou, I. Kostikas, V.N. Manimanis and M.S. Dimitrijević, *The cosmic system of the pre-Socratic philosopher Anaximenes and stars and their formation*, *Bulgarian Astronomical Journal* **23** (2015) 41.

[4] A. Kenny, *An Illustrated Brief History of Western Philosophy*, John Wiley & Sons (2018).

[5] A.S. Eddington, *The internal constitution of the stars*, *The Observatory* **43** (1920) 341.

[6] H.A. Bethe, *Energy Production in Stars*, *Physical Review* **55** (1939) 434.

[7] R. Davis, D.S. Harmer and K.C. Hoffman, *Search for Neutrinos from the Sun*, *Phys. Rev. Lett.* **20** (1968) 1205.

[8] J.N. Bahcall and J. Davis, Raymond, *Solar Neutrinos: A Scientific Puzzle*, *Science* **191** (1976) 264.

[9] Z. Maki, M. Nakagawa and S. Sakata, *Remarks on the Unified Model of Elementary Particles*, *Progress of Theoretical Physics* **28** (1962) 870.

[10] B. Pontecorvo, *Neutrino Experiments and the Problem of Conservation of Leptonic Charge*, *Soviet Journal of Experimental and Theoretical Physics* **26** (1968) 984.

[11] Y. Fukuda, T. Hayakawa, E. Ichihara, K. Inoue, K. Ishihara, H. Ishino et al., *Evidence for Oscillation of Atmospheric Neutrinos*, *Phys. Rev. Lett.* **81** (1998) 1562 [`hep-ex/9807003`].

[12] Q.R. Ahmad, R.C. Allen, T.C. Andersen, J.D. Anglin, J.C. Barton, E.W. Beier et al., *Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory*, *Phys. Rev. Lett.* **89** (2002) 011301 [`nucl-ex/0204008`].

[13] K. Eguchi, S. Enomoto, K. Furuno, J. Goldman, H. Hanada, H. Ikeda et al., *First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance*, *Phys. Rev. Lett.* **90** (2003) 021802 [`hep-ex/0212021`].

[14] M. Aker, A. Beglarian, J. Behrens, A. Berlev, U. Besserer, B. Bieringer et al., *First direct neutrino-mass measurement with sub-eV sensitivity*, *arXiv e-prints* (2021) arXiv:2105.08533 [`2105.08533`].

[15] A. Ashtari Esfahani, D.M. Asner, S. Böser, R. Cervantes, C. Claessens, L. de Viveiros et al., *Determining the neutrino mass with cyclotron radiation emission spectroscopy—Project 8*, *Journal of Physics G Nuclear Physics* **44** (2017) 054004 [`1703.02037`].

[16] M. Levi, C. Bebek, T. Beers, R. Blum, R. Cahn, D. Eisenstein et al., *The DESI*

Experiment, a whitepaper for Snowmass 2013, *arXiv e-prints* (2013) arXiv:1308.0847 [`1308.0847`].

[17] R. Laureijs, J. Amiaux, S. Arduini, J.L. Augueres, J. Brinchmann, R. Cole et al., *Euclid Definition Study Report*, *arXiv e-prints* (2011) arXiv:1110.3193 [`1110.3193`].

[18] Ž. Ivezić, S.M. Kahn, J.A. Tyson, B. Abel, E. Acosta, R. Allsman et al., *LSST: From Science Drivers to Reference Design and Anticipated Data Products*, vol. 873, p. 111, Mar., 2019, DOI [`0805.2366`].

[19] Y.Y.Y. Wong, *Neutrino Mass in Cosmology: Status and Prospects*, *Annual Review of Nuclear and Particle Science* **61** (2011) 69 [`1111.1436`].

[20] J. Lesgourgues and S. Pastor, *Neutrino mass from Cosmology*, *arXiv e-prints* (2012) arXiv:1212.6154 [`1212.6154`].

[21] S. Hannestad, *Neutrino physics and precision cosmology*, *arXiv e-prints* (2016) arXiv:1605.03829 [`1605.03829`].

[22] Z. Sakr, *A Short Review on the Latest Neutrinos Mass and Number Constraints from Cosmological Observables*, *Universe* **8** (2022) 284.

[23] J.R. Bond, G. Efstathiou and J. Silk, *Massive Neutrinos and the Large-Scale Structure of the Universe*, *Phys. Rev. Lett.* **45** (1980) 1980.

[24] S.D.M. White, C.S. Frenk and M. Davis, *Clustering in a neutrino-dominated universe*, *ApJ* **274** (1983) L1.

[25] W. Hu, D.J. Eisenstein and M. Tegmark, *Weighing Neutrinos with Galaxy Surveys*, *Phys. Rev. Lett.* **80** (1998) 5255 [`astro-ph/9712057`].

[26] S. Vagnozzi, S. Dhawan, M. Gerbino, K. Freese, A. Goobar and O. Mena, *Constraints on the sum of the neutrino masses in dynamical dark energy models with w (z)≥-1 are tighter than those obtained in Λ CDM*, *Phys. Rev. D* **98** (2018) 083501 [`1801.08553`].

[27] D. Green and J. Meyers, *Cosmological Implications of a Neutrino Mass Detection*, *arXiv e-prints* (2021) arXiv:2111.01096 [`2111.01096`].

[28] P.F. de Salas, D.V. Forero, S. Gariazzo, P. Martínez-Miravé, O. Mena, C.A. Ternes et al., *2020 global reassessment of the neutrino oscillation picture*, *Journal of High Energy Physics* **2021** (2021) 71 [`2006.11237`].

[29] I. Esteban, M.C. Gonzalez-Garcia, M. Maltoni, T. Schwetz and A. Zhou, *The fate of hints: updated global analysis of three-flavor neutrino oscillations*, *Journal of High Energy Physics* **2020** (2020) 178 [`2007.14792`].

[30] S.F. King, *Neutrino mass models*, *Reports on Progress in Physics* **67** (2004) 107 [`hep-ph/0310204`].

[31] R.N. Mohapatra and A.Y. Smirnov, *Neutrino Mass and New Physics*, *Annual Review of Nuclear and Particle Science* **56** (2006) 569 [`hep-ph/0603118`].

[32] S.F. King, A. Merle, S. Morisi, Y. Shimizu and M. Tanimoto, *Neutrino mass and mixing: from theory to experiment*, *New Journal of Physics* **16** (2014) 045018 [1402.4271].

[33] R. Munroe, *What If? 2: Additional Serious Scientific Answers to Absurd Hypothetical Questions*, Penguin (2022).

[34] Calmarius, "Is it possible to start fire using moonlight." Physics Stack Exchange, 2014. `https://physics.stackexchange.com/q/140927`.

[35] Spraff, "Why does conservation of étendue matter when showing one cannot focus light to arbitrary temperatures." Physics Stack Exchange. `https://physics.stackexchange.com/q/234996`.

[36] Anonymous, "What-if 145: Fire from moonlight." Reddit, 2014. `https://www.reddit.com/r/xkcd/comments/451qis/`.

[37] J. Chaves, *Introduction to nonimaging optics*, CRC press (2017).

[38] J.H. Oort, *The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems*, *Bull. Astron. Inst. Netherlands* **6** (1932) 249.

[39] F. Zwicky, *Die Rotverschiebung von extragalaktischen Nebeln*, *Helvetica Physica Acta* **6** (1933) 110.

[40] H.W. Babcock, *The rotation of the Andromeda Nebula*, *Lick Observatory Bulletin* **498** (1939) 41.

[41] J.P. Ostriker and P.J.E. Peebles, *A Numerical Study of the Stability of Flattened Galaxies: or, can Cold Galaxies Survive?*, *ApJ* **186** (1973) 467.

[42] J.P. Ostriker, P.J.E. Peebles and A. Yahil, *The Size and Mass of Galaxies, and the Mass of the Universe*, *ApJ* **193** (1974) L1.

[43] J. Einasto, E. Saar, A. Kaasik and A.D. Chernin, *Missing mass around galaxies - Morphological evidence*, *Nature* **252** (1974) 111.

[44] S.D.M. White and M.J. Rees, *Core condensation in heavy halos: a two-stage theory for galaxy formation and clustering.*, *MNRAS* **183** (1978) 341.

[45] R. Cowsik and J. McClelland, *An Upper Limit on the Neutrino Rest Mass*, *Phys. Rev. Lett.* **29** (1972) 669.

[46] R. Cowsik and J. McClelland, *Gravity of Neutrinos of Nonzero Mass in Astrophysics*, *ApJ* **180** (1973) 7.

[47] A.S. Szalay and G. Marx, *Neutrino rest mass from cosmology.*, *A&A* **49** (1976) 437.

[48] S. Tremaine and J.E. Gunn, *Dynamical role of light neutral leptons in cosmology*, *Phys. Rev. Lett.* **42** (1979) 407.

[49] V.A. Lubimov, E.G. Novikov, V.Z. Nozik, E.F. Tretyakov and V.S. Kosik, *An estimate of the ν_e mass from the β-spectrum of tritium in the valine molecule*, *Physics Letters B* **94** (1980) 266.

[50] C.S. Frenk, S.D.M. White and M. Davis, *Nonlinear evolution of large-scale structure in the universe*, *ApJ* **271** (1983) 417.

[51] S.D.M. White, M. Davis and C.S. Frenk, *The size of clusters in a neutrino-dominated universe*, *MNRAS* **209** (1984) 27P.

[52] O. Lahav and A.R. Liddle, *The Cosmological Parameters (2021)*, *arXiv e-prints* (2022) arXiv:2201.08666 [`2201.08666`].

[53] D. Baumann, *Cosmology*, Cambridge University Press (2022).

[54] E.D. Di Valentino, *Challenges of the Standard Cosmological Model*, *Universe* **8** (2022) 399.

[55] Planck collaboration, *Planck 2018 results. I. Overview and the cosmological legacy of Planck*, *Astron. Astrophys.* **641** (2020) A1 [`1807.06205`].

[56] F. Bernardeau, S. Colombi, E. Gaztañaga and R. Scoccimarro, *Large-scale structure of the Universe and cosmological perturbation theory*, *Phys. Rep.* **367** (2002) 1 [`astro-ph/0112551`].

[57] P.J.E. Peebles and E.J. Groth, *Statistical analysis of catalogs of extragalactic objects. V. Three-point correlation function for the galaxy distribution in the Zwicky catalog.*, *ApJ* **196** (1975) 1.

[58] D.J. Heath, *The growth of density perturbations in zero pressure Friedmann-Lemaître universes.*, *MNRAS* **179** (1977) 351.

[59] C. Rampf, B. Villone and U. Frisch, *How smooth are particle trajectories in a ΛCDM Universe?*, *MNRAS* **452** (2015) 1421 [`1504.00032`].

[60] M.H. Goroff, B. Grinstein, S.J. Rey and M.B. Wise, *Coupling of modes of cosmological mass density fluctuations*, *ApJ* **311** (1986) 6.

[61] B. Jain and E. Bertschinger, *Second-Order Power Spectrum and Nonlinear Evolution at High Redshift*, *ApJ* **431** (1994) 495 [`astro-ph/9311070`].

[62] Y.B. Zel'dovich, *Gravitational instability: An approximate theory for large density perturbations.*, *A&A* **5** (1970) 84.

[63] C. Rampf, *The recursion relation in Lagrangian perturbation theory*, *J. Cosmology Astropart. Phys.* **2012** (2012) 004 [`1205.5274`].

[64] V. Zheligovsky and U. Frisch, *Time-analyticity of Lagrangian particle trajectories in ideal fluid flow*, *Journal of Fluid Mechanics* **749** (2014) 404 [`1312.6320`].

[65] T. Matsubara, *Recursive solutions of Lagrangian perturbation theory*, *Phys. Rev. D* **92** (2015) 023534 [`1505.01481`].

[66] F. Schmidt, *An n-th order Lagrangian forward model for large-scale structure*, *J. Cosmology Astropart. Phys.* **2021** (2021) 033 [`2012.09837`].

[67] M. Crocce and R. Scoccimarro, *Renormalized cosmological perturbation theory*, *Phys. Rev. D* **73** (2006) 063519 [`astro-ph/0509418`].

[68] T. Matsubara, *Resumming cosmological perturbations via the Lagrangian picture: One-loop results in real space and in redshift space*, *Phys. Rev. D* **77** (2008) 063530 [`0711.2521`].

[69] M. Pietroni, *Flowing with time: a new approach to non-linear cosmological perturbations*, *J. Cosmology Astropart. Phys.* **2008** (2008) 036 [`0806.0971`].

[70] J. Carlson, M. White and N. Padmanabhan, *Critical look at cosmological perturbation theory techniques*, *Phys. Rev. D* **80** (2009) 043531 [`0905.0479`].

[71] J.J.M. Carrasco, M.P. Hertzberg and L. Senatore, *The effective field theory of cosmological large scale structures*, *Journal of High Energy Physics* **2012** (2012) 82 [`1206.2926`].

[72] R.A. Porto, L. Senatore and M. Zaldarriaga, *The Lagrangian-space Effective Field Theory of large scale structures*, *J. Cosmology Astropart. Phys.* **2014** (2014) 022 [`1311.2168`].

[73] D. Baumann, A. Nicolis, L. Senatore and M. Zaldarriaga, *Cosmological non-linearities as an effective fluid*, *J. Cosmology Astropart. Phys.* **2012** (2012) 051 [`1004.2488`].

[74] C. Rampf and O. Hahn, *Shell-crossing in a* Λ*CDM Universe*, *MNRAS* **501** (2021) L71 [`2010.12584`].

[75] R.W. Hockney and J.W. Eastwood, *Computer simulation using particles* (1988).

[76] M. Schaller, J. Borrow, P.W. Draper, M. Ivkovic, S. McAlpine, B. Vandenbroucke et al., *Swift: A modern highly-parallel gravity and smoothed particle hydrodynamics solver for astrophysical and cosmological applications*, *arXiv e-prints* (2023) arXiv:2305.13380 [`2305.13380`].

[77] V. Springel, R. Pakmor, O. Zier and M. Reinecke, *Simulating cosmic structure formation with the gadget-4 code*, *Mon. Not. Roy. Astron. Soc.* **506** (2021) 2871 [`2010.03567`].

[78] A. Klypin and F. Prada, *Cosmoligical n-body code glam.*, .

[79] Y. Feng, M.-Y. Chu, U. Seljak and P. McDonald, “FastPM: Scaling N-body Particle Mesh solver.” Astrophysics Source Code Library, record ascl:1905.010, May, 2019.

[80] L.H. Garrison, D.J. Eisenstein, D. Ferrer, N.A. Maksimova and P.A. Pinto, *The ABACUS cosmological N-body code*, *MNRAS* **508** (2021) 575 [`2110.11392`].

[81] D. Potter, J. Stadel and R. Teyssier, *PKDGRAV3: beyond trillion particle cosmological simulations for the next era of galaxy surveys*, *Computational Astrophysics and Cosmology* **4** (2017) 2 [`1609.08621`].

[82] R. Teyssier, *Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES*, *A&A* **385** (2002) 337 [`astro-ph/0111367`].

[83] V. Springel, *E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh*, *MNRAS* **401** (2010) 791 [`0901.4107`].

[84] J. Harnois-Déraps, U.-L. Pen, I.T. Iliev, H. Merz, J.D. Emberson and V. Desjacques, *High-performance* P^3M *N-body code: CUBEP*^{3}*M*, *MNRAS* **436** (2013) 540 [`1208.5098`].

[85] J. Adamek, D. Daverio, R. Durrer and M. Kunz, *gevolution: a cosmological N-body code based on General Relativity*, *JCAP* **07** (2016) 053 [`1604.06065`].

[86] J. Dakin, S. Hannestad and T. Tram, *The cosmological simulation code CONCEPT 1.0*, *MNRAS* **513** (2022) 991 [`2112.01508`].

[87] C. Barrera-Hinojosa and B. Li, *GRAMSES: a new route to general relativistic N-body simulations in cosmology. Part I. Methodology and code description*, *J. Cosmology Astropart. Phys.* **2020** (2020) 007 [`1905.08890`].

[88] M. Crocce, S. Pueblas and R. Scoccimarro, *Transients from initial conditions in cosmological simulations*, *MNRAS* **373** (2006) 369 [`astro-ph/0606505`].

[89] M. Michaux, O. Hahn, C. Rampf and R.E. Angulo, *Accurate initial conditions for cosmological N-body simulations: minimizing truncation and discreteness errors*, *MNRAS* **500** (2021) 663 [`2008.09588`].

[90] S. Tassev, M. Zaldarriaga and D. Eisenstein, *Solving Large Scale Structure in Ten Easy Steps with COLA*, *JCAP* **06** (2013) 036 [`1301.0322`].

[91] N. Chartier, B. Wandelt, Y. Akrami and F. Villaescusa-Navarro, *CARPool: fast, accurate computation of large-scale structure statistics by pairing costly and cheap cosmological simulations*, *MNRAS* **503** (2021) 1897 [`2009.08970`].

[92] N. Kokron, S.-F. Chen, M. White, J. DeRose and M. Maus, *Accurate predictions from small boxes: variance suppression via the Zel'dovich approximation*, *J. Cosmology Astropart. Phys.* **2022** (2022) 059 [`2205.15327`].

[93] J. DeRose, S.-F. Chen, N. Kokron and M. White, *Precision redshift-space galaxy power spectra using Zel'dovich control variates*, *J. Cosmology Astropart. Phys.* **2023** (2023) 008 [`2210.14239`].

[94] C. Modi, S.-F. Chen and M. White, *Simulations and symmetries*, *MNRAS* **492** (2020) 5754 [`1910.07097`].

[95] A. Taruya, T. Nishimichi and D. Jeong, *Grid-based calculation for perturbation theory of large-scale structure*, *Phys. Rev. D* **98** (2018) 103532 [`1807.04215`].

[96] C.M. Baugh, *A primer on hierarchical galaxy formation: the semi-analytical approach*, *Reports on Progress in Physics* **69** (2006) 3101 [`astro-ph/0610031`].

[97] R.S. Somerville and R. Davé, *Physical Models of Galaxy Formation in a Cosmological Framework*, *ARA&A* **53** (2015) 51 [`1412.2712`].

[98] M. Vogelsberger, F. Marinacci, P. Torrey and E. Puchwein, *Cosmological simulations of galaxy formation*, *Nature Reviews Physics* **2** (2020) 42 [`1909.07976`].

[99] V. Springel, *Smoothed Particle Hydrodynamics in Astrophysics*, *ARA&A* **48** (2010) 391 [`1109.2219`].

[100] R.L. Workman, V.D. Burkert, V. Crede, E. Klempt, U. Thoma, L. Tiator et al., *Review of Particle Physics*, *Progress of Theoretical and Experimental Physics* **2022** (2022) 083C01.

[101] B. Kayser, *Neutrino Oscillation Physics*, *arXiv e-prints* (2012) arXiv:1206.4325 [`1206.4325`].

[102] G. Fantini, A. Gallo Rosso, F. Vissani and V. Zema, *The formalism of neutrino oscillations: an introduction*, *arXiv e-prints* (2018) arXiv:1802.05781 [`1802.05781`].

[103] S. Pascoli, *Neutrino physics*, *CERN Yellow Rep. School Proc.* **6** (2019) 213.

[104] C.V. Achar, M.G.K. Menon, V.S. Narasimham, P.V. Ramana Murthy, B.V. Sreekantan, K. Hinotani et al., *Observation of a non-elastic cosmic ray neutrino interaction*, *Physics Letters* **19** (1965) 78.

[105] C. Wagner, L. Verde and R. Jimenez, *Effects of the Neutrino Mass Splitting on the Nonlinear Matter Power Spectrum*, *ApJ* **752** (2012) L31 [`1203.5342`].

[106] M. Archidiacono, S. Hannestad and J. Lesgourgues, *What will it take to measure individual neutrino mass states using cosmology?*, *J. Cosmology Astropart. Phys.* **2020** (2020) 021 [`2003.03354`].

[107] S. Vagnozzi, E. Giusarma, O. Mena, K. Freese, M. Gerbino, S. Ho et al., *Unveiling ν secrets with cosmological data: Neutrino masses and mass hierarchy*, *Phys. Rev. D* **96** (2017) 123503 [`1701.08172`].

[108] N. Palanque-Delabrouille, C. Yèche, N. Schöneberg, J. Lesgourgues, M. Walther, S. Chabanier et al., *Hints, neutrino bounds and WDM constraints from SDSS DR14 Lyman-α and Planck full-survey data*, *JCAP* **04** (2020) 038 [`1911.09073`].

[109] E. Di Valentino, S. Gariazzo and O. Mena, *Most constraining cosmological neutrino mass bounds*, *Phys. Rev. D* **104** (2021) 083504 [`2106.15267`].

[110] R. Jimenez, C. Pena-Garay, K. Short, F. Simpson and L. Verde, *Neutrino masses and mass hierarchy: evidence for the normal hierarchy*, *J. Cosmology Astropart. Phys.* **2022** (2022) 006 [`2203.14247`].

[111] S. Gariazzo, M. Gerbino, T. Brinckmann, M. Lattanzi, O. Mena, T. Schwetz et al., *Neutrino mass and mass ordering: no conclusive evidence for normal ordering*, *J. Cosmology Astropart. Phys.* **2022** (2022) 010 [`2205.02195`].

[112] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi et al., *Planck 2018 results. VI. Cosmological parameters*, *A&A* **641** (2020) A6 [`1807.06209`].

[113] S.R. Choudhury and S. Hannestad, *Updated results on neutrino mass and mass hierarchy from cosmology with Planck 2018 likelihoods*, *J. Cosmology Astropart. Phys.* **2020** (2020) 037 [`1907.12598`].

[114] A. Upadhye, *Neutrino mass and dark energy constraints from redshift-space distortions*, *J. Cosmology Astropart. Phys.* **2019** (2019) 041 [`1707.09354`].

[115] M. Tristram, A.J. Banday, M. Douspis, X. Garrido, K.M. Górski, S. Henrot-Versillé et al., *Cosmological parameters derived from the final (PR4) Planck data release*, *arXiv e-prints* (2023) arXiv:2309.10034 [`2309.10034`].

[116] I.G. McCarthy, S. Bird, J. Schaye, J. Harnois-Deraps, A.S. Font and L. van Waerbeke,

The BAHAMAS project: the CMB-large-scale structure tension and the roles of massive neutrinos and galaxy formation, *MNRAS* **476** (2018) 2999 [`1712.02411`].

[117] F. Beutler, S. Saito, J.R. Brownstein, C.-H. Chuang, A.J. Cuesta, W.J. Percival et al., *The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: signs of neutrino mass in current cosmological data sets*, *MNRAS* **444** (2014) 3501 [`1403.4599`].

[118] W.H. Furry, *On Transition Probabilities in Double Beta-Disintegration*, *Physical Review* **56** (1939) 1184.

[119] E.W. Kolb and M.S. Turner, *The early universe*, Westview Press (1990).

[120] J. Lesgourgues and S. Pastor, *Massive neutrinos and cosmology*, *Phys. Rep.* **429** (2006) 307 [`astro-ph/0603494`].

[121] S. Weinberg, *Cosmology*, Oxford University Press (2008).

[122] G. Mangano, G. Miele, S. Pastor, T. Pinto, O. Pisanti and P.D. Serpico, *Relic neutrino decoupling including flavor oscillations*, *Nucl. Phys. B* **729** (2005) 221 [`hep-ph/0506164`].

[123] P.F. de Salas and S. Pastor, *Relic neutrino decoupling with flavour oscillations revisited*, *JCAP* **07** (2016) 051 [`1606.06986`].

[124] J. Lesgourgues and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS) IV: efficient implementation of non-cold relics*, *J. Cosmology Astropart. Phys.* **2011** (2011) 032 [`1104.2935`].

[125] M. Shoji and E. Komatsu, *Erratum: Massive neutrinos in cosmology: Analytic solutions and fluid approximation [Phys. Rev. D 81, 123516 (2010)]*, *Phys. Rev. D* **82** (2010) 089901 [`1003.0942`].

[126] C.-P. Ma and E. Bertschinger, *Cosmological Perturbation Theory in the Synchronous vs. Conformal Newtonian Gauge*, *arXiv e-prints* (1994) astro [`astro-ph/9401007`].

[127] A. Lewis and A. Challinor, *CAMB: Code for Anisotropies in the Microwave Background*, `1102.026`.

[128] J. Lesgourgues, *The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview*, *arXiv e-prints* (2011) arXiv:1104.2932 [`1104.2932`].

[129] R.E. Angulo and O. Hahn, *Large-scale dark matter simulations*, *Living Reviews in Computational Astrophysics* **8** (2022) 1 [`2112.05165`].

[130] J. Adamek, R.E. Angulo, C. Arnold, M. Baldi, M. Biagetti, B. Bose et al., *Euclid: Modelling massive neutrinos in cosmology – a code comparison*, *arXiv e-prints* (2022) arXiv:2211.12457 [`2211.12457`].

[131] 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. Cosmology Astropart. Phys.* **2014** (2014) 023 [`1308.4164`].

[132] A. Chudaykin and M.M. Ivanov, *Measuring neutrino masses with large-scale structure:*

Euclid forecast with controlled theoretical error, *J. Cosmology Astropart. Phys.* **2019** (2019) 034 [1907.06666].

[133] T. Sprenger, M. Archidiacono, T. Brinckmann, S. Clesse and J. Lesgourgues, *Cosmology in the era of Euclid and the Square Kilometre Array*, *J. Cosmology Astropart. Phys.* **2019** (2019) 047 [1801.08331].

[134] R. Acciarri, M.A. Acero, M. Adamowski, C. Adams, P. Adamson, S. Adhikari et al., *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report Volume 1: The LBNF and DUNE Projects*, *arXiv e-prints* (2016) arXiv:1601.05471 [1601.05471].

[135] S. Adrián-Martínez, M. Ageron, F. Aharonian, S. Aiello, A. Albert, F. Ameli et al., *Letter of intent for KM3NeT 2.0*, *Journal of Physics G Nuclear Physics* **43** (2016) 084001 [1601.07459].

[136] K. Abe, T. Abe, H. Aihara, Y. Fukuda, Y. Hayato, K. Huang et al., *Letter of Intent: The Hyper-Kamiokande Experiment — Detector Design and Physics Potential —*, *arXiv e-prints* (2011) arXiv:1109.3262 [1109.3262].

[137] F. An, G. An, Q. An, V. Antonelli, E. Baussan, J. Beacom et al., *Neutrino physics with JUNO*, *Journal of Physics G Nuclear Physics* **43** (2016) 030401 [1507.05613].

[138] S.M. Bilenky, S. Pascoli and S. Petcov, *Majorana neutrinos, neutrino mass spectrum, CP violation and neutrinoless double beta decay. 1. The Three neutrino mixing case*, *Phys. Rev. D* **64** (2001) 053010 [hep-ph/0102265].

[139] S. Pascoli and S. Petcov, *The SNO solar neutrino data, neutrinoless double beta decay and neutrino mass spectrum*, *Phys. Lett. B* **544** (2002) 239 [hep-ph/0205022].

[140] H. Nunokawa, W. Teves and R. Zukanovich Funchal, *Constraining the absolute neutrino mass scale and Majorana CP violating phases by future 0 neutrino beta beta decay experiments*, *Phys. Rev. D* **66** (2002) 093010 [hep-ph/0206137].

[141] J.D. Vergados, H. Ejiri and F. Šimkovic, *Neutrinoless double beta decay and neutrino mass*, *International Journal of Modern Physics E* **25** (2016) 1630007.

[142] APPEC Committee collaboration, *Double Beta Decay APPEC Committee Report*, 1910.04688.

[143] G. Drexlin, V. Hannen, S. Mertens and C. Weinheimer, *Current Direct Neutrino Mass Experiments*, *arXiv e-prints* (2013) arXiv:1307.0101 [1307.0101].

[144] S. Mertens, *Direct Neutrino Mass Experiments*, in *Journal of Physics Conference Series*, vol. 718 of *Journal of Physics Conference Series*, p. 022013, May, 2016, DOI [1605.01579].

[145] L. Senatore and M. Zaldarriaga, *The Effective Field Theory of Large-Scale Structure in the presence of Massive Neutrinos*, *arXiv e-prints* (2017) arXiv:1707.04698 [1707.04698].

[146] T. Colas, G. d'Amico, L. Senatore, P. Zhang and F. Beutler, *Efficient cosmological analysis of the SDSS/BOSS data from the Effective Field Theory of Large-Scale Structure*, *J. Cosmology Astropart. Phys.* **2020** (2020) 001 [1909.07951].

[147] A.A. Klypin and S.F. Shandarin, *Three-dimensional numerical model of the formation of large-scale structure in the Universe.*, *MNRAS* **204** (1983) 891.

[148] J. Brandbyge, S. Hannestad, T. Haugbølle and B. Thomsen, *The Effect of Thermal Neutrino Motion on the Non-linear Cosmological Matter Power Spectrum*, *JCAP* **08** (2008) 020 [`0802.3700`].

[149] M. Viel, M.G. Haehnelt and V. Springel, *The effect of neutrinos on the matter distribution as probed by the Intergalactic Medium*, *JCAP* **06** (2010) 015 [`1003.2422`].

[150] C.-P. Ma and E. Bertschinger, *A Calculation of the Full Neutrino Phase Space in Cold + Hot Dark Matter Models*, *ApJ* **429** (1994) 22 [`astro-ph/9308006`].

[151] C.-P. Ma and E. Bertschinger, *Do Galactic Systems Form too Late in Cold + Hot Dark Matter Models?*, *ApJ* **434** (1994) L5 [`astro-ph/9407085`].

[152] S. Bird, M. Viel and M.G. Haehnelt, *Massive neutrinos and the non-linear matter power spectrum*, *MNRAS* **420** (2012) 2551 [`1109.4416`].

[153] F. Villaescusa-Navarro, F. Marulli, M. Viel, E. Branchini, E. Castorina, E. Sefusatti et al., *Cosmology with massive neutrinos I: towards a realistic modeling of the relation between matter, haloes and galaxies*, *J. Cosmology Astropart. Phys.* **2014** (2014) 011 [`1311.0866`].

[154] J. Adamek, R. Durrer and M. Kunz, *Relativistic N-body simulations with massive neutrinos*, *JCAP* **11** (2017) 004 [`1707.06938`].

[155] A. Banerjee, D. Powell, T. Abel and F. Villaescusa-Navarro, *Reducing Noise in Cosmological N-body Simulations with Neutrinos*, *JCAP* **09** (2018) 028 [`1801.03906`].

[156] J.D. Emberson et al., *Cosmological neutrino simulations at extreme scale*, *Res. Astron. Astrophys.* **17** (2017) 085 [`1611.01545`].

[157] D. Inman, J.D. Emberson, U.-L. Pen, A. Farchi, H.-R. Yu and J. Harnois-Déraps, *Precision reconstruction of the cold dark matter-neutrino relative velocity from N-body simulations*, *Phys. Rev. D* **92** (2015) 023502 [`1503.07480`].

[158] F. Villaescusa-Navarro, P. Bull and M. Viel, *Weighing Neutrinos with Cosmic Neutral Hydrogen*, *ApJ* **814** (2015) 146 [`1507.05102`].

[159] E. Castorina, C. Carbone, J. Bel, E. Sefusatti and K. Dolag, *DEMNUni: The clustering of large-scale structures in the presence of massive neutrinos*, *JCAP* **07** (2015) 043 [`1505.07148`].

[160] F. Villaescusa-Navarro, C. Hahn, E. Massara, A. Banerjee, A.M. Delgado, D.K. Ramanah et al., *The Quijote Simulations*, *ApJS* **250** (2020) 2 [`1909.05273`].

[161] J. Brandbyge and S. Hannestad, *Grid Based Linear Neutrino Perturbations in Cosmological N-body Simulations*, *JCAP* **05** (2009) 002 [`0812.3149`].

[162] S. Hannestad, T. Haugbølle and C. Schultz, *Neutrinos in non-linear structure formation — a simple SPH approach*, *J. Cosmology Astropart. Phys.* **2012** (2012) 045 [`1110.1257`].

[163] M. Archidiacono and S. Hannestad, *Efficient calculation of cosmological neutrino clustering in the non-linear regime*, *JCAP* **06** (2016) 018 [`1510.02907`].

[164] A. Banerjee and N. Dalal, *Simulating nonlinear cosmological structure formation with massive neutrinos*, *JCAP* **11** (2016) 015 [`1606.06167`].

[165] J. Dakin, J. Brandbyge, S. Hannestad, T. HaugbØlle and T. Tram, *νCONCEPT: cosmological neutrino simulations from the non-linear Boltzmann hierarchy*, *J. Cosmology Astropart. Phys.* **2019** (2019) 052 [`1712.03944`].

[166] T. Tram, J. Brandbyge, J. Dakin and S. Hannestad, *Fully relativistic treatment of light neutrinos in N-body simulations*, *JCAP* **03** (2019) 022 [`1811.00904`].

[167] D. Inman and H.-r. Yu, *Simulating the Cosmic Neutrino Background using Collisionless Hydrodynamics*, *Astrophys. J. Suppl.* **250** (2020) 21 [`2002.04601`].

[168] J.Z. Chen, A. Upadhye and Y.Y.Y. Wong, *The cosmic neutrino background as a collection of fluids in large-scale structure simulations*, *JCAP* **03** (2021) 065 [`2011.12503`].

[169] Y. Ali-Haïmoud and S. Bird, *An efficient implementation of massive neutrinos in non-linear structure formation simulations*, *MNRAS* **428** (2013) 3375 [`1209.0461`].

[170] J. Liu, S. Bird, J.M. Zorrilla Matilla, J.C. Hill, Z. Haiman, M.S. Madhavacheril et al., *MassiveNuS: cosmological massive neutrino simulations*, *J. Cosmology Astropart. Phys.* **2018** (2018) 049 [`1711.10524`].

[171] J.Z. Chen, A. Upadhye and Y.Y.Y. Wong, *One line to run them all: SuperEasy massive neutrino linear response in N-body simulations*, *JCAP* **04** (2021) [`2011.12504`].

[172] C. Partmann, C. Fidler, C. Rampf and O. Hahn, *Fast simulations of cosmic large-scale structure with massive neutrinos*, *J. Cosmology Astropart. Phys.* **2020** (2020) 018 [`2003.07387`].

[173] J. Brandbyge and S. Hannestad, *Resolving Cosmic Neutrino Structure: A Hybrid Neutrino N-body Scheme*, *JCAP* **01** (2010) 021 [`0908.1969`].

[174] S. Bird, Y. Ali-Haïmoud, Y. Feng and J. Liu, *An Efficient and Accurate Hybrid Method for Simulating Non-Linear Neutrino Structure*, *Mon. Not. Roy. Astron. Soc.* **481** (2018) 1486 [`1803.09854`].

[175] K. Yoshikawa, S. Tanaka, N. Yoshida and S. Saito, *Cosmological Vlasov–Poisson Simulations of Structure Formation with Relic Neutrinos: Nonlinear Clustering and the Neutrino Mass*, *Astrophys. J.* **904** (2020) 159 [`2010.00248`].

[176] D. Merritt, *Stability of elliptical galaxies. Numerical experiments.*, in *Structure and Dynamics of Elliptical Galaxies*, P.T. de Zeeuw, ed., vol. 127, pp. 315–329, Jan., 1987, DOI.

[177] F. Leeuwin, F. Combes and J. Binney, *N-body simulations with perturbation particles. I - Method and tests*, *MNRAS* **262** (1993) 1013.

[178] T. Tajima and F.W. Perkins, *Unpublished communication*, .

[179] S.E. Parker and W.W. Lee, *A fully nonlinear characteristic method for gyrokinetic simulation*, *Physics of Fluids B* **5** (1993) 77.

[180] A.M. Dimits and W.W. Lee, *Partially Linearized Algorithms in Gyrokinetic Particle Simulation*, *Journal of Computational Physics* **107** (1993) 309.

[181] A.Y. Aydemir, *A unified monte carlo interpretation of particle simulations and applications to non-neutral plasmas*, *Physics of Plasmas* **1** (1994) 822.

[182] S. Ross, *Simulation*, Knovel Library, pp. 147–154, Elsevier Science (2012).

[183] E.W. Weisstein, *Jacobi elliptic functions*, *From MathWorld–A Wolfram Web Resource. https://mathworld.wolfram.com/JacobiEllipticFunctions.html* (2002) .

[184] C. Fidler, T. Tram, C. Rampf, R. Crittenden, K. Koyama and D. Wands, *Relativistic initial conditions for N-body simulations*, *J. Cosmology Astropart. Phys.* **2017** (2017) 043 [`1702.03221`].

[185] C. Fidler, C. Rampf, T. Tram, R. Crittenden, K. Koyama and D. Wands, *General relativistic corrections to N-body simulations and the Zel'dovich approximation*, *Phys. Rev. D* **92** (2015) 123517 [`1505.04756`].

[186] W. Elbers, *Geodesic motion and phase-space evolution of massive neutrinos*, *J. Cosmology Astropart. Phys.* **2022** (2022) 058 [`2207.14256`].

[187] J. Brandbyge, C. Rampf, T. Tram, F. Leclercq, C. Fidler and S. Hannestad, *Cosmological N -body simulations including radiation perturbations*, *MNRAS* **466** (2017) L68 [`1610.04236`].

[188] C. Fidler, T. Tram, C. Rampf, R. Crittenden, K. Koyama and D. Wands, *General relativistic weak-field limit and Newtonian N-body simulations*, *JCAP* **12** (2017) 022 [`1708.07769`].

[189] M. Schaller, P. Gonnet, A.B.G. Chalk and P.W. Draper, *SWIFT: Using Task-Based Parallelism, Fully Asynchronous Communication, and Graph Partition-Based Domain Decomposition for Strong Scaling on more than 100,000 Cores*, in *Proceedings of the Platform for Advanced Scientific Computing Conference*, p. 2, June, 2016, DOI [`1606.02738`].

[190] M. Schaller, P. Gonnet, P.W. Draper, A.B.G. Chalk, R.G. Bower, J. Willis et al., *SWIFT: SPH With Inter-dependent Fine-grained Tasking*, `1805.020`.

[191] E. Di Valentino, A. Melchiorri and J. Silk, *Cosmological constraints in extended parameter space from the Planck 2018 Legacy release*, *J. Cosmology Astropart. Phys.* **2020** (2020) 013 [`1908.01391`].

[192] J. Brandbyge, S. Hannestad, T. Haugbølle and Y.Y.Y. Wong, *Neutrinos in non-linear structure formation — the effect on halo properties*, *J. Cosmology Astropart. Phys.* **2010** (2010) 014 [`1004.4105`].

[193] F. Villaescusa-Navarro, J. Miralda-Escudé, C. Peña-Garay and V. Quilis, *Neutrino halos*

in clusters of galaxies and their weak lensing signature, J. Cosmology Astropart. Phys. **2011** (2011) 027 [`1104.4770`].

[194] P.J. Elahi, R. Cañas, R.J.J. Poulton, R.J. Tobar, J.S. Willis, C.d.P. Lagos et al., *Hunting for galaxies and halos in simulations with velociraptor*, Publ. Astron. Soc. Australia **36** (2019) e021 [`1902.01010`].

[195] A. Taruya and Y. Suto, *Nonlinear Stochastic Biasing from the Formation Epoch Distribution of Dark Halos*, ApJ **542** (2000) 559 [`astro-ph/0004288`].

[196] K. Yoshikawa, A. Taruya, Y.P. Jing and Y. Suto, *Nonlinear Stochastic Biasing of Galaxies and Dark Halos in Cosmological Hydrodynamic Simulations*, ApJ **558** (2001) 520 [`astro-ph/0104361`].

[197] S. Hannestad, A. Upadhye and Y.Y.Y. Wong, *Spoon or slide? The non-linear matter power spectrum in the presence of massive neutrinos*, J. Cosmology Astropart. Phys. **2020** (2020) 062 [`2006.04995`].

[198] C. Okoli, M.I. Scrimgeour, N. Afshordi and M.J. Hudson, *Dynamical friction in the primordial neutrino sea*, MNRAS **468** (2017) 2164 [`1611.04589`].

[199] H.-R. Yu, U.-L. Pen and X. Wang, *Parity-odd neutrino torque detection*, Phys. Rev. D **99** (2019) 123532 [`1810.11784`].

[200] T.R. Quinn, N. Katz, J. Stadel and G. Lake, *Time stepping N body simulations*, `astro-ph/9710043`.

[201] A. Porredon, M. Crocce, J. Elvin-Poole, R. Cawthon, G. Giannini, J. De Vicente et al., *Dark Energy Survey Year 3 results: Cosmological constraints from galaxy clustering and galaxy-galaxy lensing using the MAGLIM lens sample*, Phys. Rev. D **106** (2022) 103530 [`2105.13546`].

[202] J. Hamann, S. Hannestad and Y.Y.Y. Wong, *Measuring neutrino masses with a future galaxy survey*, J. Cosmology Astropart. Phys. **2012** (2012) 052 [`1209.1043`].

[203] 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** (2015) 66 [`1309.5383`].

[204] T. Brinckmann, D.C. Hooper, M. Archidiacono, J. Lesgourgues and T. Sprenger, *The promising future of a robust cosmological neutrino mass measurement*, J. Cosmology Astropart. Phys. **2019** (2019) 059 [`1808.05955`].

[205] S. Saito, M. Takada and A. Taruya, *Impact of Massive Neutrinos on the Nonlinear Matter Power Spectrum*, Phys. Rev. Lett. **100** (2008) 191301 [`0801.0607`].

[206] Y.Y.Y. Wong, *Higher order corrections to the large scale matter power spectrum in the presence of massive neutrinos*, J. Cosmology Astropart. Phys. **2008** (2008) 035 [`0809.0693`].

[207] H. Dupuy and F. Bernardeau, *Describing massive neutrinos in cosmology as a collection of independent flows*, J. Cosmology Astropart. Phys. **2014** (2014) 030 [`1311.5487`].

[208] D. Blas, M. Garny, T. Konstandin and J. Lesgourgues, *Structure formation with massive neutrinos: going beyond linear theory*, *J. Cosmology Astropart. Phys.* **2014** (2014) 039 [`1408.2995`].

[209] F. Führer and Y.Y.Y. Wong, *Higher-order massive neutrino perturbations in large-scale structure*, *J. Cosmology Astropart. Phys.* **2015** (2015) 046 [`1412.2764`].

[210] M. Levi and Z. Vlah, *Massive neutrinos in nonlinear large scale structure: A consistent perturbation theory*, *arXiv e-prints* (2016) arXiv:1605.09417 [`1605.09417`].

[211] J. Lesgourgues, S. Matarrese, M. Pietroni and A. Riotto, *Non-linear power spectrum including massive neutrinos: the time-RG flow approach*, *J. Cosmology Astropart. Phys.* **2009** (2009) 017 [`0901.4550`].

[212] B.S. Wright, H.A. Winther and K. Koyama, *COLA with massive neutrinos*, *JCAP* **10** (2017) 054 [`1705.08165`].

[213] A. Aviles and A. Banerjee, *A Lagrangian Perturbation Theory in the presence of massive neutrinos*, *JCAP* **10** (2020) 034 [`2007.06508`].

[214] R. Scoccimarro, *Transients from initial conditions: a perturbative analysis*, *MNRAS* **299** (1998) 1097 [`astro-ph/9711187`].

[215] E. Sirko, *Initial Conditions to Cosmological N-Body Simulations, or, How to Run an Ensemble of Simulations*, *ApJ* **634** (2005) 728 [`astro-ph/0503106`].

[216] C. Yèche, N. Palanque-Delabrouille, J. Baur and H. du Mas des Bourboux, *Constraints on neutrino masses from Lyman-alpha forest power spectrum with BOSS and XQ-100*, *J. Cosmology Astropart. Phys.* **2017** (2017) 047 [`1702.03314`].

[217] G. Efstathiou, M. Davis, S.D.M. White and C.S. Frenk, *Numerical techniques for large cosmological N-body simulations*, *ApJS* **57** (1985) 241.

[218] A. Schneider, R. Teyssier, D. Potter, J. Stadel, J. Onions, D.S. Reed et al., *Matter power spectrum and the challenge of percent accuracy*, *J. Cosmology Astropart. Phys.* **2016** (2016) 047 [`1503.05920`].

[219] B. Marcos, T. Baertschiger, M. Joyce, A. Gabrielli and F. Sylos Labini, *Linear perturbative theory of the discrete cosmological N-body problem*, *Phys. Rev. D* **73** (2006) 103507 [`astro-ph/0601479`].

[220] L.H. Garrison, D.J. Eisenstein, D. Ferrer, M.V. Metchnik and P.A. Pinto, *Improving initial conditions for cosmological N-body simulations*, *MNRAS* **461** (2016) 4125 [`1605.02333`].

[221] M. Zennaro, J. Bel, F. Villaescusa-Navarro, C. Carbone, E. Sefusatti and L. Guzzo, *Initial conditions for accurate N-body simulations of massive neutrino cosmologies*, *MNRAS* **466** (2017) 3244 [`1605.05283`].

[222] P. Heuschling, C. Partmann and C. Fidler, *A minimal model for massive neutrinos in Newtonian N-body simulations*, *J. Cosmology Astropart. Phys.* **2022** (2022) 068 [`2201.13186`].

[223] C. Fidler and A. Kleinjohann, *Suitable initial conditions for Newtonian simulations with massive neutrinos*, *J. Cosmology Astropart. Phys.* **2019** (2019) 018 [`1810.12019`].

[224] C. Barrera-Hinojosa and B. Li, *GRAMSES: a new route to general relativistic N-body simulations in cosmology. Part II. Initial conditions*, *J. Cosmology Astropart. Phys.* **2020** (2020) 056 [`2001.07968`].

[225] T. Buchert, *Lagrangian Theory of Gravitational Instability of Friedman-Lemaitre Cosmologies - a Generic Third-Order Model for Nonlinear Clustering*, *MNRAS* **267** (1994) 811 [`astro-ph/9309055`].

[226] F.R. Bouchet, S. Colombi, E. Hivon and R. Juszkiewicz, *Perturbative Lagrangian approach to gravitational instability.*, *A&A* **296** (1995) 575 [`astro-ph/9406013`].

[227] A.L. Melott, T. Buchert and A.G. Weiss, *Testing higher-order Lagrangian perturbation theory against numerical simulations. II. Hierarchical models.*, *A&A* **294** (1995) 345 [`astro-ph/9404018`].

[228] O. Hahn, C. Rampf and C. Uhlemann, *Higher order initial conditions for mixed baryon-CDM simulations*, *MNRAS* **503** (2021) 426 [`2008.09124`].

[229] C. Rampf, C. Uhlemann and O. Hahn, *Cosmological perturbations for two cold fluids in ΛCDM*, *MNRAS* **503** (2021) 406 [`2008.09123`].

[230] Y. Brenier, U. Frisch, M. Hénon, G. Loeper, S. Matarrese, R. Mohayaee et al., *Reconstruction of the early Universe as a convex optimization problem*, *MNRAS* **346** (2003) 501 [`astro-ph/0304214`].

[231] T. Buchert, *A class of solutions in Newtonian cosmology and the pancake theory*, *A&A* **223** (1989) 9.

[232] F. Moutarde, J.M. Alimi, F.R. Bouchet, R. Pellat and A. Ramani, *Precollapse Scale Invariance in Gravitational Instability*, *ApJ* **382** (1991) 377.

[233] F.R. Bouchet, R. Juszkiewicz, S. Colombi and R. Pellat, *Weakly Nonlinear Gravitational Instability for Arbitrary Omega*, *ApJ* **394** (1992) L5.

[234] M. Gramann, *An Improved Reconstruction Method for Cosmological Density Fields*, *ApJ* **405** (1993) 449.

[235] T. Buchert and J. Ehlers, *Lagrangian theory of gravitational instability of Friedman-Lemaitre cosmologies – second-order approach: an improved model for non-linear clustering*, *MNRAS* **264** (1993) 375.

[236] A. Aviles and J.L. Cervantes-Cota, *Lagrangian perturbation theory for modified gravity*, *Phys. Rev. D* **96** (2017) 123526 [`1705.10719`].

[237] W. Elbers, C.S. Frenk, A. Jenkins, B. Li and S. Pascoli, *An optimal non-linear method for simulating relic neutrinos*, *MNRAS* **507** (2021) 2614 [`2010.07321`].

[238] R.E. Angulo and A. Pontzen, *Cosmological N-body simulations with suppressed variance*, *Mon. Not. Roy. Astron. Soc.* **462** (2016) L1 [`1603.05253`].

[239] C.-T. Chiang, W. Hu, Y. Li and M. LoVerde, *Scale-dependent bias and bispectrum in neutrino separate universe simulations*, *Phys. Rev. D* **97** (2018) 123526 [`1710.01310`].

[240] R. Ruggeri, E. Castorina, C. Carbone and E. Sefusatti, *DEMNUni: massive neutrinos and the bispectrum of large scale structures*, *J. Cosmology Astropart. Phys.* **2018** (2018) 003 [`1712.02334`].

[241] C. Hahn, F. Villaescusa-Navarro, E. Castorina and R. Scoccimarro, *Constraining* M_ν *with the bispectrum. Part I. Breaking parameter degeneracies*, *J. Cosmology Astropart. Phys.* **2020** (2020) 040 [`1909.11107`].

[242] K. Ichiki and M. Takada, *Impact of massive neutrinos on the abundance of massive clusters*, *Phys. Rev. D* **85** (2012) 063521 [`1108.4688`].

[243] 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. Cosmology Astropart. Phys.* **2014** (2014) 049 [`1311.1212`].

[244] E. Massara, F. Villaescusa-Navarro and M. Viel, *The halo model in a massive neutrino cosmology*, *JCAP* **12** (2014) 053 [`1410.6813`].

[245] T. Nishimichi, M. Takada, R. Takahashi, K. Osato, M. Shirasaki, T. Oogi et al., *Dark Quest. I. Fast and Accurate Emulation of Halo Clustering Statistics and Its Application to Galaxy Clustering*, *ApJ* **884** (2019) 29 [`1811.09504`].

[246] L.H. Garrison, D.J. Eisenstein and P.A. Pinto, *A high-fidelity realization of the Euclid code comparison N-body simulation with ABACUS*, *MNRAS* **485** (2019) 3370 [`1810.02916`].

[247] C. Grove, C.-H. Chuang, N.C. Devi, L. Garrison, B. L'Huillier, Y. Feng et al., *The DESI N-body simulation project - I. Testing the robustness of simulations for the DESI dark time survey*, *MNRAS* **515** (2022) 1854 [`2112.09138`].

[248] DES collaboration, *Dark Energy Survey Year 3 results: Cosmological constraints from galaxy clustering and weak lensing*, *Phys. Rev. D* **105** (2022) 023520 [`2105.13549`].

[249] C. Fidler, A. Kleinjohann, T. Tram, C. Rampf and K. Koyama, *A new approach to cosmological structure formation with massive neutrinos*, *JCAP* **01** (2019) 025 [`1807.03701`].

[250] M. Zennaro, R.E. Angulo, G. Aricò, S. Contreras and M. Pellejero-Ibáñez, *How to add massive neutrinos to your* Λ*CDM simulation - extending cosmology rescaling algorithms*, *Mon. Not. Roy. Astron. Soc.* **489** (2019) 5938 [`1905.08696`].

[251] S. Hannestad, T. Haugbolle and C. Schultz, *Neutrinos in Non-linear Structure Formation - a Simple SPH Approach*, *JCAP* **02** (2012) 045 [`1110.1257`].

[252] J. Dakin, J. Brandbyge, S. Hannestad, T. Haugbølle and T. Tram, ν*CONCEPT: Cosmological neutrino simulations from the non-linear Boltzmann hierarchy*, *JCAP* **02** (2019) 052 [`1712.03944`].

[253] Y. Ali-Haimoud and S. Bird, *An efficient implementation of massive neutrinos in*

non-linear structure formation simulations, Mon. Not. Roy. Astron. Soc. **428** (2012) 3375 [`1209.0461`].

[254] M. Davis, F.J. Summers and D. Schlegel, *Large scale structure in a universe with mixed hot and cold dark matter*, Nature **359** (1992) 393.

[255] A. Klypin, J. Holtzman, J. Primack and E. Regos, *Structure formation with cold plus hot dark matter*, Astrophys. J. **416** (1993) 1 [`astro-ph/9305011`].

[256] F. Villaescusa-Navarro et al., *The Quijote simulations*, Astrophys. J. Suppl. **250** (2020) 2 [`1909.05273`].

[257] W. Elbers, C.S. Frenk, A. Jenkins, B. Li and S. Pascoli, *Higher order initial conditions with massive neutrinos*, MNRAS **516** (2022) 3821 [`2202.00670`].

[258] M. Morita and T. Tatekawa, *Extending Lagrangian perturbation theory to a fluid with velocity dispersion*, Mon. Not. Roy. Astron. Soc. **328** (2001) 815 [`astro-ph/0108289`].

[259] T. Tatekawa, M. Suda, K.-i. Maeda, M. Morita and H. Anzai, *Perturbation theory in Lagrangian hydrodynamics for a cosmological fluid with velocity dispersion*, Phys. Rev. D **66** (2002) 064014 [`astro-ph/0205017`].

[260] M. Michaux, O. Hahn, C. Rampf and R.E. Angulo, *Accurate initial conditions for cosmological N-body simulations: Minimizing truncation and discreteness errors*, Mon. Not. Roy. Astron. Soc. **500** (2020) 663 [`2008.09588`].

[261] O. Hahn, M. Michaux, C. Rampf, C. Uhlemann and R.E. Angulo, *MUSIC2-monofonIC: 3LPT initial condition generator*, `2008.024`.

[262] P. Saha and S. Tremaine, *Symplectic integrators for solar system dynamics*, *The Astronomical Journal* **104** (1992) 1633.

[263] S. Hannestad and J. Brandbyge, *The Cosmic Neutrino Background Anisotropy - Linear Theory*, JCAP **03** (2010) 020 [`0910.4578`].

[264] C.G. Tully and G. Zhang, *Multi-messenger astrophysics with the cosmic neutrino background*, JCAP **06** (2021) 053 [`2103.01274`].

[265] A. Banerjee, S. Das, A. Maharana and R. Kumar Sharma, *Signatures of Light Massive Relics on nonlinear structure formation*, `2202.09840`.

[266] C.L. Bennett, A. Banday, K.M. Gorski, G. Hinshaw, P. Jackson, P. Keegstra et al., *Four year COBE DMR cosmic microwave background observations: Maps and basic results*, Astrophys. J. Lett. **464** (1996) L1 [`astro-ph/9601067`].

[267] WMAP collaboration, *Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation*, Astrophys. J. Suppl. **192** (2011) 18 [`1001.4538`].

[268] S. Weinberg, *Universal Neutrino Degeneracy*, Phys. Rev. **128** (1962) 1457.

[269] KATRIN collaboration, *New Constraint on the Local Relic Neutrino Background*

Overdensity with the First KATRIN Data Runs, Phys. Rev. Lett. **129** (2022) 011806 [2202.04587].

[270] PTOLEMY collaboration, *PTOLEMY: A Proposal for Thermal Relic Detection of Massive Neutrinos and Directional Detection of MeV Dark Matter*, 1808.01892.

[271] PTOLEMY collaboration, *Neutrino physics with the PTOLEMY project: active neutrino properties and the light sterile case*, JCAP **07** (2019) 047 [1902.05508].

[272] PTOLEMY collaboration, *Heisenberg's uncertainty principle in the PTOLEMY project: A theory update*, Phys. Rev. D **106** (2022) 053002 [2203.11228].

[273] A.G. Cocco, G. Mangano and M. Messina, *Probing low energy neutrino backgrounds with neutrino capture on beta decaying nuclei*, JCAP **06** (2007) 015 [hep-ph/0703075].

[274] Y. Cheipesh, V. Cheianov and A. Boyarsky, *Navigating the pitfalls of relic neutrino detection*, Phys. Rev. D **104** (2021) 116004 [2101.10069].

[275] R. Opher, *Coherent scattering of cosmic neutrinos*, *Astron. Astrophys.* **37** (1974) 135.

[276] L. Stodolsky, *Speculations on Detection of the Neutrino Sea*, Phys. Rev. Lett. **34** (1975) 110.

[277] V. Domcke and M. Spinrath, *Detection prospects for the Cosmic Neutrino Background using laser interferometers*, JCAP **06** (2017) 055 [1703.08629].

[278] J.D. Shergold, *Updated detection prospects for relic neutrinos using coherent scattering*, JCAP **11** (2021) 052 [2109.07482].

[279] B. Eberle, A. Ringwald, L. Song and T.J. Weiler, *Relic neutrino absorption spectroscopy*, Phys. Rev. D **70** (2004) 023007 [hep-ph/0401203].

[280] V. Brdar, P.S.B. Dev, R. Plestid and A. Soni, *A new probe of relic neutrino clustering using cosmogenic neutrinos*, Phys. Lett. B **833** (2022) 137358 [2207.02860].

[281] M. Yoshimura, N. Sasao and M. Tanaka, *Experimental method of detecting relic neutrino by atomic de-excitation*, Phys. Rev. D **91** (2015) 063516 [1409.3648].

[282] M. Bauer and J.D. Shergold, *Relic neutrinos at accelerator experiments*, Phys. Rev. D **104** (2021) 083039 [2104.12784].

[283] M. Bauer and J.D. Shergold, *Limits on the cosmic neutrino background*, JCAP **01** (2023) 003 [2207.12413].

[284] W. Hu, D. Scott, N. Sugiyama and M.J. White, *The Effect of physical assumptions on the calculation of microwave background anisotropies*, Phys. Rev. D **52** (1995) 5498 [astro-ph/9505043].

[285] R.J. Michney and R.R. Caldwell, *Anisotropy of the Cosmic Neutrino Background*, JCAP **01** (2007) 014 [astro-ph/0608303].

[286] J.Y.-Y. Lin and G. Holder, *Gravitational Lensing of the Cosmic Neutrino Background*, JCAP **04** (2020) 054 [1910.03550].

[287] M. Lisanti, B.R. Safdi and C.G. Tully, *Measuring anisotropies in the cosmic neutrino background*, *PRD* **90** (2014) 073006 [`1407.0393`].

[288] B.R. Safdi, M. Lisanti, J. Spitz and J.A. Formaggio, *Annual Modulation of Cosmic Relic Neutrinos*, *Phys. Rev. D* **90** (2014) 043001 [`1404.0680`].

[289] E. Akhmedov, *Relic neutrino detection through angular correlations in inverse β-decay*, *JCAP* **09** (2019) 031 [`1905.10207`].

[290] A.N. Baushev, *The Relic Neutrino Composition as Seen from Earth*, *Astron. Rep.* **64** (2020) 1005 [`2101.11405`].

[291] G. Baym and J.-C. Peng, *Evolution of primordial neutrino helicities in cosmic gravitational inhomogeneities*, *Phys. Rev. D* **103** (2021) 123019 [`2103.11209`].

[292] B. Hernandez-Molinero, R. Jimenez and C. Pena-Garay, *Distinguishing Dirac vs. Majorana neutrinos: a cosmological probe*, *JCAP* **08** (2022) 038 [`2205.00808`].

[293] J.-C. Peng and G. Baym, *Implication of Helicity Modifications of Primordial Neutrinos on Their Detection*, *JPS Conf. Proc.* **37** (2022) 020704.

[294] B. Hernández-Molinero, C. Carbone, R. Jimenez and C. Peña Garay, *Cosmic Background Neutrinos Deflected by Gravity: DEMNUni Simulation Analysis*, `2301.12430`.

[295] A. Ringwald and Y.Y.Y. Wong, *Gravitational clustering of relic neutrinos and implications for their detection*, *JCAP* **12** (2004) 005 [`hep-ph/0408241`].

[296] P.F. de Salas, S. Gariazzo, J. Lesgourgues and S. Pastor, *Calculation of the local density of relic neutrinos*, *JCAP* **09** (2017) 034 [`1706.09850`].

[297] J. Zhang and X. Zhang, *Gravitational clustering of cosmic relic neutrinos in the Milky Way*, *Nature Commun.* **9** (2018) 1833 [`1712.01153`].

[298] P. Mertsch, G. Parimbelli, P.F. de Salas, S. Gariazzo, J. Lesgourgues and S. Pastor, *Neutrino clustering in the Milky Way and beyond*, *JCAP* **01** (2020) 015 [`1910.13388`].

[299] E. Brinch Holm, I.M. Oldengott and S. Zentarra, *Local clustering of relic neutrinos with kinetic field theory*, *arXiv e-prints* (2023) arXiv:2305.13379 [`2305.13379`].

[300] F. Zimmer, C.A. Correa and S. Ando, *Influence of local structure on relic neutrino abundances and anisotropies*, `2306.16444`.

[301] G. Lavaux and M.J. Hudson, *The 2M++ galaxy redshift catalogue*, *Mon. Not. Roy. Astron. Soc.* **416** (2011) 2840 [`1105.6107`].

[302] M. Cautun, A. Benitez-Llambay, A.J. Deason, C.S. Frenk, A. Fattahi, F.A. Gómez et al., *The Milky Way total mass profile as inferred from Gaia DR2*, *Mon. Not. Roy. Astron. Soc.* **494** (2020) 4291 [`1911.04557`].

[303] G. Yepes, S. Gottlöber and Y. Hoffman, *Dark Matter in the Local Universe*, *New Astron. Rev.* **58** (2014) 1 [`1312.0105`].

[304] H. Wang, H.J. Mo, X. Yang, Y.P. Jing and W.P. Lin, *ELUCID - Exploring the Local*

Universe with reConstructed Initial Density field I: Hamiltonian Markov Chain Monte Carlo Method with Particle Mesh Dynamics, *Astrophys. J.* **794** (2014) 94 [1407.3451].

[305] N.I. Libeskind, E. Carlesi, R.J.J. Grand, A. Khalatyan, A. Knebe, R. Pakmor et al., *The HESTIA project: simulations of the Local Group*, *MNRAS* **498** (2020) 2968 [2008.04926].

[306] J.G. Sorce, Y. Dubois, J. Blaizot, S.L. McGee, G. Yepes and A. Knebe, *I – A hydrodynamical clone of the Virgo cluster of galaxies to confirm observationally driven formation scenarios*, *Mon. Not. Roy. Astron. Soc.* **504** (2021) 2998 [2104.13389].

[307] S. McAlpine, J.C. Helly, M. Schaller, T. Sawala, G. Lavaux, J. Jasche et al., *SIBELIUS-DARK: a galaxy catalogue of the local volume from a constrained realization simulation*, *Mon. Not. Roy. Astron. Soc.* **512** (2022) 5823 [2202.04099].

[308] M.S.S.L. Oei, R.J. van Weeren, F. Vazza, F. Leclercq, A. Gopinath and H.J.A. Röttgering, *Filamentary baryons and where to find them - A forecast of synchrotron radiation from merger and accretion shocks in the local Cosmic Web*, *Astron. Astrophys.* **662** (2022) A87 [2203.05365].

[309] J. Jasche and B.D. Wandelt, *Bayesian physical reconstruction of initial conditions from large scale structure surveys*, *Mon. Not. Roy. Astron. Soc.* **432** (2013) 894 [1203.3639].

[310] J. Jasche and G. Lavaux, *Physical Bayesian modelling of the non-linear matter distribution: new insights into the Nearby Universe*, *Astron. Astrophys.* **625** (2019) A64 [1806.11117].

[311] S. Stopyra, H.V. Peiris, A. Pontzen, J. Jasche and G. Lavaux, *Towards Accurate Field-Level Inference of Massive Cosmic Structures*, *arXiv e-prints* (2023) arXiv:2304.09193 [2304.09193].

[312] D.J. Bartlett, H. Desmond and P.G. Ferreira, *Constraints on Galileons from the positions of supermassive black holes*, *PRD* **103** (2021) 023523 [2010.05811].

[313] D.J. Bartlett, A. Kostić, H. Desmond, J. Jasche and G. Lavaux, *Constraints on dark matter annihilation and decay from the large-scale structure of the nearby Universe*, *PRD* **106** (2022) 103526 [2205.12916].

[314] H. Desmond, M.L. Hutt, J. Devriendt and A. Slyz, *Catalogues of voids as antihaloes in the local Universe*, *MNRAS* **511** (2022) L45 [2109.09439].

[315] M.L. Hutt, H. Desmond, J. Devriendt and A. Slyz, *The effect of local Universe constraints on halo abundance and clustering*, *MNRAS* **516** (2022) 3592 [2203.14724].

[316] J.F. Navarro, C.S. Frenk and S.D.M. White, *A Universal density profile from hierarchical clustering*, *Astrophys. J.* **490** (1997) 493 [astro-ph/9611107].

[317] R.B. Tully, E.J. Shaya, I.D. Karachentsev, H.M. Courtois, D.D. Kocevski, L. Rizzi et al., *Our Peculiar Motion Away from the Local Void*, *Astrophys. J.* **676** (2008) 184 [0705.4139].

[318] R.B. Tully, H.M. Courtois and J.G. Sorce, *Cosmicflows-3*, *Astron. J.* **152** (2016) 50 [1605.01765].

[319] K. Ichiki and M. Takada, *The impact of massive neutrinos on the abundance of massive clusters*, *Phys. Rev. D* **85** (2012) 063521 [`1108.4688`].

[320] 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*, *JCAP* **02** (2014) 049 [`1311.1212`].

[321] A.J. Long, C. Lunardini and E. Sabancilar, *Detecting non-relativistic cosmic neutrinos by capture on tritium: phenomenology and physics potential*, *JCAP* **08** (2014) 038 [`1405.7654`].

[322] O. Roldan, A. Notari and M. Quartin, *Interpreting the CMB aberration and Doppler measurements: boost or intrinsic dipole?*, *JCAP* **2016** (2016) 026 [`1603.02664`].

[323] P.D. Meerburg, J. Meyers and A. van Engelen, *Reconstructing the primary CMB dipole*, *PRD* **96** (2017) 083519 [`1704.00718`].

[324] P. da Silveira Ferreira and M. Quartin, *First constraints on the intrinsic CMB dipole and our velocity with Doppler and aberration*, *arXiv e-prints* (2020) arXiv:2011.08385 [`2011.08385`].

[325] A. Lewis and A. Challinor, *Weak gravitational lensing of the CMB*, *Phys. Rept.* **429** (2006) 1 [`astro-ph/0601594`].

[326] H. Meusinger, C. Rudolf, B. Stecklum, M. Hoeft, R. Mauersberger and D. Apai, *The galaxy population within the virial radius of the Perseus cluster*, *Astron. Astrophys.* **640** (2020) A30 [`2103.11209`].

[327] D. Lynden-Bell, S.M. Faber, D. Burstein, R.L. Davies, A. Dressler, R.J. Terlevich et al., *Spectroscopy and photometry of elliptical galaxies. V - Galaxy streaming toward the new supergalactic center*, *Astrophys. J.* **326** (1988) 19.

[328] Y. Hoffman, D. Pomarede, R. Brent Tully and H. Courtois, *The Dipole Repeller*, `1702.02483`.

[329] D. Levesque and L. Verlet, *Molecular dynamics and time reversibility*, *Journal of Statistical Physics* **72** (1993) 519.

[330] B. Li, L.J. King, G.-B. Zhao and H. Zhao, *A Semi-analytic Ray-tracing Algorithm for Weak Lensing*, *Mon. Not. Roy. Astron. Soc.* **415** (2011) 881 [`1012.1625`].

[331] H. Rein and D. Tamayo, *JANUS: A bit-wise reversible integrator for N-body dynamics*, *Mon. Not. Roy. Astron. Soc.* **473** (2018) 3351 [`1704.07715`].

[332] J. Barnes and P. Hut, *A Hierarchical O(N log N) Force Calculation Algorithm*, *Nature* **324** (1986) 446.

[333] A.G. Riess et al., *A Comprehensive Measurement of the Local Value of the Hubble Constant with 1 km* s^{-1} *Mpc*$^{-1}$ *Uncertainty from the Hubble Space Telescope and the SH0ES Team*, *Astrophys. J. Lett.* **934** (2022) L7 [`2112.04510`].

[334] E. Di Valentino, *A combined analysis of the H_0 late time direct measurements and the impact on the Dark Energy sector*, *MNRAS* **502** (2021) 2065 [2011.00246].

[335] E. Abdalla, G.F. Abellán, A. Aboubrahim, A. Agnello, Ö. Akarsu and Akrami, *Cosmology intertwined: A review of the particle physics, astrophysics, and cosmology associated with the cosmological tensions and anomalies*, *Journal of High Energy Astrophysics* **34** (2022) 49 [2203.06142].

[336] M. Asgari, C.-A. Lin, B. Joachimi, B. Giblin, C. Heymans, H. Hildebrandt et al., *KiDS-1000 cosmology: Cosmic shear constraints and comparison between two point statistics*, *A&A* **645** (2021) A104 [2007.15633].

[337] A. Amon, D. Gruen, M.A. Troxel, N. MacCrann, S. Dodelson, A. Choi et al., *Dark Energy Survey Year 3 results: Cosmology from cosmic shear and robustness to data calibration*, *Phys. Rev. D* **105** (2022) 023514 [2105.13543].

[338] L.F. Secco, S. Samuroff, E. Krause, B. Jain, J. Blazek, M. Raveri et al., *Dark Energy Survey Year 3 results: Cosmology from cosmic shear and robustness to modeling uncertainty*, *Phys. Rev. D* **105** (2022) 023515 [2105.13544].

[339] DESI collaboration, *Overview of the Instrumentation for the Dark Energy Spectroscopic Instrument*, *Astron. J.* **164** (2022) 207 [2205.10939].

[340] M.P. van Daalen, J. Schaye, C.M. Booth and C. Dalla Vecchia, *The effects of galaxy formation on the matter power spectrum: a challenge for precision cosmology*, *MNRAS* **415** (2011) 3649 [1104.1174].

[341] N.E. Chisari, A.J. Mead, S. Joudaki, P.G. Ferreira, A. Schneider, J. Mohr et al., *Modelling baryonic feedback for survey cosmology*, *The Open Journal of Astrophysics* **2** (2019) 4 [1905.06082].

[342] A. Schneider, R. Teyssier, J. Stadel, N.E. Chisari, A.M.C. Le Brun, A. Amara et al., *Quantifying baryon effects on the matter power spectrum and the weak lensing shear correlation*, *J. Cosmology Astropart. Phys.* **2019** (2019) 020 [1810.08629].

[343] S.N.B. Debackere, J. Schaye and H. Hoekstra, *The impact of the observed baryon distribution in haloes on the total matter power spectrum*, *MNRAS* **492** (2020) 2285 [1908.05765].

[344] V.J. Forouhar Moreno, A. Benítez-Llambay, S. Cole and C. Frenk, *Baryon-driven decontraction in Milky Way-mass haloes*, *MNRAS* **511** (2022) 3910 [2107.14245].

[345] B.O. Mummery, I.G. McCarthy, S. Bird and J. Schaye, *The separate and combined effects of baryon physics and neutrino free streaming on large-scale structure*, *MNRAS* **471** (2017) 227 [1702.02064].

[346] A. Schneider, N. Stoira, A. Refregier, A.J. Weiss, M. Knabenhans, J. Stadel et al., *Baryonic effects for weak lensing. Part I. Power spectrum and covariance matrix*, *J. Cosmology Astropart. Phys.* **2020** (2020) 019 [1910.11357].

[347] M.P. van Daalen, I.G. McCarthy and J. Schaye, *Exploring the effects of galaxy formation*

on matter clustering through a library of simulation power spectra, *MNRAS* **491** (2020) 2424 [`1906.00968`].

[348] G. Parimbelli, G. Scelfo, S.K. Giri, A. Schneider, M. Archidiacono, S. Camera et al., *Mixed dark matter: matter power spectrum and halo mass function*, *J. Cosmology Astropart. Phys.* **2021** (2021) 044 [`2106.04588`].

[349] V.R. Eke, J.F. Navarro and M. Steinmetz, *The Power Spectrum Dependence of Dark Matter Halo Concentrations*, *ApJ* **554** (2001) 114 [`astro-ph/0012337`].

[350] S.R. Knollmann, C. Power and A. Knebe, *Dark Matter Halo Profiles in Scale-Free Cosmologies*, in *Dark Matter in Astroparticle and Particle Physics*, H.V. Klapdor-Kleingrothaus and G.F. Lewis, eds., pp. 290–298, Apr., 2008, DOI.

[351] F. Prada, A.A. Klypin, A.J. Cuesta, J.E. Betancort-Rijo and J. Primack, *Halo concentrations in the standard LCDM cosmology*, *Mon. Not. Roy. Astron. Soc.* **423** (2012) 3018 [`1104.5130`].

[352] J. Kwan, S. Bhattacharya, K. Heitmann and S. Habib, *Cosmic Emulation: The Concentration-Mass Relation for wCDM Universes*, *ApJ* **768** (2013) 123 [`1210.1576`].

[353] A. Ragagnin, A. Saro, P. Singh and K. Dolag, *Cosmology dependence of halo masses and concentrations in hydrodynamic simulations*, *MNRAS* **500** (2021) 5056 [`2011.05345`].

[354] C.M. Booth and J. Schaye, *Dark matter haloes determine the masses of supermassive black holes*, *MNRAS* **405** (2010) L1 [`0911.0935`].

[355] R.G. Bower, J. Schaye, C.S. Frenk, T. Theuns, M. Schaller, R.A. Crain et al., *The dark nemesis of galaxy formation: why hot haloes trigger black hole growth and bring star formation to an end*, *MNRAS* **465** (2017) 32 [`1607.07445`].

[356] V. Avila-Reese, P. Colín, S. Gottlöber, C. Firmani and C. Maulbetsch, *The Dependence on Environment of Cold Dark Matter Halo Properties*, *ApJ* **634** (2005) 51 [`astro-ph/0508053`].

[357] R.H. Wechsler, A.R. Zentner, J.S. Bullock, A.V. Kravtsov and B. Allgood, *The Dependence of Halo Clustering on Halo Formation History, Concentration, and Occupation*, *ApJ* **652** (2006) 71 [`astro-ph/0512416`].

[358] L. Gao and S.D.M. White, *Assembly bias in the clustering of dark matter haloes*, *MNRAS* **377** (2007) L5 [`astro-ph/0611921`].

[359] J. Schaye, R. Kugel, M. Schaller, J.C. Helly, J. Braspenning, W. Elbers et al., *The FLAMINGO project: cosmological hydrodynamical simulations for large-scale structure and galaxy cluster surveys*, *MNRAS* (2023) .

[360] A.M.C. Le Brun, I.G. McCarthy, J. Schaye and T.J. Ponman, *Towards a realistic population of simulated galaxy groups and clusters*, *MNRAS* **441** (2014) 1270 [`1312.5462`].

[361] I.G. McCarthy, J. Schaye, S. Bird and A.M.C. Le Brun, *The BAHAMAS project: calibrated hydrodynamical simulations for large-scale structure cosmology*, *MNRAS* **465** (2017) 2936 [`1603.02702`].

[362] R. Kugel, J. Schaye, M. Schaller, J.C. Helly, J. Braspenning, W. Elbers et al., *FLAMINGO: Calibrating large cosmological hydrodynamical simulations with machine learning*, *MNRAS* (2023) [2306.05492].

[363] Y.M. Bahé, J. Schaye, M. Schaller, R.G. Bower, J. Borrow, E. Chaikin et al., *The importance of black hole repositioning for galaxy formation simulations*, *MNRAS* **516** (2022) 167 [2109.01489].

[364] E. Chaikin, J. Schaye, M. Schaller, A. Benítez-Llambay, F.S.J. Nobels and S. Ploeckinger, *A thermal-kinetic subgrid model for supernova feedback in simulations of galaxy formation*, *arXiv e-prints* (2022) arXiv:2211.04619 [2211.04619].

[365] F. Huško, C.G. Lacey, J. Schaye, M. Schaller and F.S.J. Nobels, *Spin-driven jet feedback in idealized simulations of galaxy groups and clusters*, *MNRAS* **516** (2022) 3750 [2206.06402].

[366] J. Borrow, M. Schaller, R.G. Bower and J. Schaye, *SPHENIX: smoothed particle hydrodynamics for the next generation of galaxy formation simulations*, *MNRAS* **511** (2022) 2367 [2012.03974].

[367] N.E. Chisari, M.L.A. Richardson, J. Devriendt, Y. Dubois, A. Schneider, A.M.C. Le Brun et al., *The impact of baryons on the matter power spectrum from the Horizon-AGN cosmological hydrodynamical simulation*, *MNRAS* **480** (2018) 3962 [1801.08559].

[368] W.H. Press and P. Schechter, *Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation*, *ApJ* **187** (1974) 425.

[369] A.R. Zentner, *The Excursion Set Theory of Halo Mass Functions, Halo Clustering, and Halo Growth*, *International Journal of Modern Physics D* **16** (2007) 763 [astro-ph/0611454].

[370] W. Cui, S. Borgani and G. Murante, *The effect of active galactic nuclei feedback on the halo mass function*, *MNRAS* **441** (2014) 1769 [1402.1493].

[371] M. Velliscig, M.P. van Daalen, J. Schaye, I.G. McCarthy, M. Cacciato, A.M.C. Le Brun et al., *The impact of galaxy formation on the total mass, mass profile and abundance of haloes*, *MNRAS* **442** (2014) 2641 [1402.4461].

[372] R.H. Wechsler, J.S. Bullock, J.R. Primack, A.V. Kravtsov and A. Dekel, *Concentrations of Dark Halos from Their Assembly Histories*, *ApJ* **568** (2002) 52 [astro-ph/0108151].

[373] E. Semboloni, H. Hoekstra, J. Schaye, M.P. van Daalen and I.G. McCarthy, *Quantifying the effect of baryon physics on weak lensing tomography*, *MNRAS* **417** (2011) 2020 [1105.1075].

[374] S. Bhattacharya, S. Habib, K. Heitmann and A. Vikhlinin, *Dark Matter Halo Profiles of Massive Clusters: Theory versus Observations*, *ApJ* **766** (2013) 32 [1112.5479].

[375] A. Amon and G. Efstathiou, *A non-linear solution to the* S_8 *tension?*, *MNRAS* **516** (2022) 5355 [2206.11794].

[376] J.F. Beacom, N.F. Bell and S. Dodelson, *Neutrinoless Universe*, *Phys. Rev. Lett.* **93** (2004) 121302 [`astro-ph/0404585`].

[377] P.D. Serpico, *Cosmological Neutrino Mass Detection: The Best Probe of Neutrino Lifetime*, *Phys. Rev. Lett.* **98** (2007) 171301 [`astro-ph/0701699`].

[378] P.D. Serpico, *Neutrinos and cosmology: A lifetime relationship*, in *Journal of Physics Conference Series*, vol. 173 of *Journal of Physics Conference Series*, p. 012018, June, 2009, DOI.

[379] Y. Farzan and S. Hannestad, *Neutrinos secretly converting to lighter particles to please both KATRIN and the cosmos*, *J. Cosmology Astropart. Phys.* **2016** (2016) 058 [`1510.02201`].

[380] G.F. Abellán, Z. Chacko, A. Dev, P. Du, V. Poulin and Y. Tsai, *Improved cosmological constraints on the neutrino mass and lifetime*, *Journal of High Energy Physics* **2022** (2022) 76 [`2112.13862`].

[381] P.J.E. Peebles, *Anomalies in physical cosmology*, *Annals of Physics* **447** (2022) 169159 [`2208.05018`].

[382] G. Taubes, *Are Neutrino Mass Hunters Pursuing a Chimera?*, *Science* **256** (1992) 731.

[383] E. Massara, F. Villaescusa-Navarro, S. Ho, N. Dalal and D.N. Spergel, *Using the Marked Power Spectrum to Detect the Signature of Neutrinos in Large-Scale Structure*, *Phys. Rev. Lett.* **126** (2021) 011301 [`2001.11024`].

[384] E. Paillas, C. Cuesta-Lazaro, P. Zarrouk, Y.-C. Cai, W.J. Percival, S. Nadathur et al., *Constraining* $\nu\Lambda$*CDM with density-split clustering*, *MNRAS* **522** (2023) 606 [`2209.04310`].

[385] C. Cuesta-Lazaro, E. Paillas, S. Yuan, Y.-C. Cai, S. Nadathur, W.J. Percival et al., *SUNBIRD: A simulation-based model for full-shape density-split clustering*, *arXiv e-prints* (2023) arXiv:2309.16539 [`2309.16539`].

[386] G. Valogiannis and C. Dvorkin, *Towards an optimal estimation of cosmological parameters with the wavelet scattering transform*, *Phys. Rev. D* **105** (2022) 103534 [`2108.07821`].

[387] E. Massara, F. Villaescusa-Navarro, M. Viel and P.M. Sutter, *Voids in massive neutrino cosmologies*, *J. Cosmology Astropart. Phys.* **2015** (2015) 018 [`1506.03088`].

[388] C.D. Kreisch, A. Pisani, C. Carbone, J. Liu, A.J. Hawken, E. Massara et al., *Massive neutrinos leave fingerprints on cosmic voids*, *MNRAS* **488** (2019) 4413 [`1808.07464`].

[389] G. Zhang, Z. Li, J. Liu, D.N. Spergel, C.D. Kreisch, A. Pisani et al., *Void halo mass function: A promising probe of neutrino mass*, *Phys. Rev. D* **102** (2020) 083537 [`1910.07553`].

[390] H.-M. Zhu and E. Castorina, *Measuring dark matter-neutrino relative velocity on cosmological scales*, *Phys. Rev. D* **101** (2020) 023525 [`1905.00361`].

[391] Y. Zu and D.H. Weinberg, *The redshift-space cluster-galaxy cross-correlation function - I. Modelling galaxy infall on to Millennium simulation clusters and SDSS groups*, *MNRAS* **431** (2013) 3319 [`1211.1379`].

[392] 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** (2018) 53 [`1708.01154`].

[393] J.E. García-Farieta, F. Marulli, A. Veropalumbo, L. Moscardini, R.A. Casas-Miranda, C. Giocoli et al., *Clustering and redshift-space distortions in modified gravity models with massive neutrinos*, *MNRAS* **488** (2019) 1987 [`1903.00154`].

[394] B.S. Wright, K. Koyama, H.A. Winther and G.-B. Zhao, *Investigating the degeneracy between modified gravity and massive neutrinos with redshift-space distortions*, *J. Cosmology Astropart. Phys.* **2019** (2019) 040 [`1902.10692`].

[395] Y. Zu, D.H. Weinberg, E. Jennings, B. Li and M. Wyman, *Galaxy infall kinematics as a test of modified gravity*, *MNRAS* **445** (2014) 1885 [`1310.6768`].

[396] E. Calabrese, A. Slosar, A. Melchiorri, G.F. Smoot and O. Zahn, *Cosmic Microwave Weak lensing data as a test for the dark universe*, *Physical Review D* **77** (2008) 123531.

[397] F. Couchot, S. Henrot-Versillé, O. Perdereau, S. Plaszczynski, B.R. d'Orfeuil, M. Spinelli et al., *Relieving tensions related to the lensing of the cosmic microwave background temperature power spectra*, *Astronomy & Astrophysics* **597** (2017) A126.

[398] G.E. Addison, C.L. Bennett, M. Halpern, G. Hinshaw and J.L. Weiland, *Revisiting the AL Lensing Anomaly in Planck 2018 Temperature Data* , *arXiv e-prints* (2023) arXiv:2310.03127 [`2310.03127`].

[399] P. Norberg, C.M. Baugh, E. Gaztañaga and D.J. Croton, *Statistical analysis of galaxy surveys - I. Robust error estimation for two-point clustering statistics*, *MNRAS* **396** (2009) 19 [`0810.1885`].

[400] A. Taylor, B. Joachimi and T. Kitching, *Putting the precision in precision cosmology: How accurate should your data covariance matrix be?*, *MNRAS* **432** (2013) 1928 [`1212.4359`].

[401] O. Friedrich, S. Seitz, T.F. Eifler and D. Gruen, *Performance of internal covariance estimators for cosmic shear correlation functions*, *MNRAS* **456** (2016) 2662 [`1508.00895`].

[402] F.G. Mohammad and W.J. Percival, *Creating jackknife and bootstrap estimates of the covariance matrix for the two-point correlation function*, *MNRAS* **514** (2022) 1289 [`2109.07071`].

[403] S. Trusov, P. Zarrouk, S. Cole, P. Norberg, C. Zhao, J.N. Aguilar et al., *2-point statistics covariance with fewer mocks*, *arXiv e-prints* (2023) arXiv:2306.16332 [`2306.16332`].

[404] H.A. Feldman, N. Kaiser and J.A. Peacock, *Power-Spectrum Analysis of Three-dimensional Redshift Surveys*, *ApJ* **426** (1994) 23 [`astro-ph/9304022`].

[405] Z. Slepian and D.J. Eisenstein, *Computing the three-point correlation function of galaxies in* $O(N\hat{2})$ *time*, *MNRAS* **454** (2015) 4142 [`1506.02040`].

[406] C. Howlett and W.J. Percival, *Galaxy two-point covariance matrix estimation for next generation surveys*, *MNRAS* **472** (2017) 4935 [`1709.03057`].

[407] C. Blake, P. Carter and J. Koda, *Power spectrum multipoles on the curved sky: an application to the 6-degree Field Galaxy Survey*, *MNRAS* **479** (2018) 5168 [1801.04969].

[408] D. Wadekar and R. Scoccimarro, *Galaxy power spectrum multipoles covariance in perturbation theory*, *Phys. Rev. D* **102** (2020) 123517 [1910.02914].

[409] C. Zhao, C.-H. Chuang, J. Bautista, A. de Mattia, A. Raichoor, A.J. Ross et al., *The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: 1000 multi-tracer mock catalogues with redshift evolution and systematics for galaxies and quasars of the final data release*, *MNRAS* **503** (2021) 1149 [2007.08997].

[410] M. Rashkovetskyi, D.J. Eisenstein, J.N. Aguilar, D. Brooks, T. Claybaugh, S. Cole et al., *Validation of semi-analytical, semi-empirical covariance matrices for two-point correlation function for early DESI data*, *MNRAS* **524** (2023) 3894 [2306.06320].

[411] J. Alsing, B. Wandelt and S. Feeney, *Massive optimal data compression and density estimation for scalable, likelihood-free inference in cosmology*, *MNRAS* **477** (2018) 2874 [1801.01497].

[412] P.L. Taylor, T.D. Kitching, J. Alsing, B.D. Wandelt, S.M. Feeney and J.D. McEwen, *Cosmic shear: Inference from forward models*, *Phys. Rev. D* **100** (2019) 023519 [1904.05364].

[413] N. Jeffrey, J. Alsing and F. Lanusse, *Likelihood-free inference with neural compression of DES SV weak lensing map statistics*, *MNRAS* **501** (2021) 954 [2009.08459].

[414] C. Hahn, M. Eickenberg, S. Ho, J. Hou, P. Lemos, E. Massara et al., *SIMBIG: A Forward Modeling Approach To Analyzing Galaxy Clustering*, *arXiv e-prints* (2022) arXiv:2211.00723 [2211.00723].

[415] C. Modi, S. Pandey, M. Ho, C. Hahn, B. R'egaldo-Saint Blancard and B. Wandelt, *Sensitivity Analysis of Simulation-Based Inference for Galaxy Clustering*, *arXiv e-prints* (2023) arXiv:2309.15071 [2309.15071].

[416] D. Prelogović and A. Mesinger, *Exploring the likelihood of the 21-cm power spectrum with simulation-based inference*, *MNRAS* **524** (2023) 4239 [2305.03074].

[417] P. Monaco, T. Theuns and G. Taffoni, *The pinocchio algorithm: pinpointing orbit-crossing collapsed hierarchical objects in a linear density field*, *MNRAS* **331** (2002) 587 [astro-ph/0109323].

[418] L.A. Rizzo, F. Villaescusa-Navarro, P. Monaco, E. Munari, S. Borgani, E. Castorina et al., *Simulating cosmologies beyond ΛCDM with PINOCCHIO*, *J. Cosmology Astropart. Phys.* **2017** (2017) 008 [1610.07624].

[419] B. Audren, J. Lesgourgues, G. Mangano, P.D. Serpico and T. Tram, *Strongest model-independent bound on the lifetime of Dark Matter*, *J. Cosmology Astropart. Phys.* **2014** (2014) 028 [1407.2418].

[420] K.L. Pandey, T. Karwal and S. Das, *Alleviating the H_0 and σ_8 anomalies with a decaying dark matter model*, *J. Cosmology Astropart. Phys.* **2020** (2020) 026 [1902.10636].

[421] J. Hubert, A. Schneider, D. Potter, J. Stadel and S.K. Giri, *Decaying dark matter: simulations and weak-lensing forecast*, *J. Cosmology Astropart. Phys.* **2021** (2021) 040 [`2104.07675`].

[422] S. Mau, E.O. Nadler, R.H. Wechsler, A. Drlica-Wagner, K. Bechtol, G. Green et al., *Milky Way Satellite Census. IV. Constraints on Decaying Dark Matter from Observations of Milky Way Satellite Galaxies*, *ApJ* **932** (2022) 128 [`2201.11740`].

[423] G. Arcadi, M. Dutra, P. Ghosh, M. Lindner, Y. Mambrini, M. Pierre et al., *The waning of the WIMP? A review of models, searches, and constraints*, *European Physical Journal C* **78** (2018) 203 [`1703.07364`].

[424] R. An, A.A. Costa, L. Xiao, J. Zhang and B. Wang, *Testing a quintessence model with Yukawa interaction from cosmological observations and N-body simulations*, *MNRAS* **489** (2019) 297 [`1809.03224`].