ALMA MATER STUDIORUM · UNIVERSITÀ DI BOLOGNA

School of Science
Department of Physics and Astronomy
Master Degree Programme in Astrophysics and Cosmology

Disentangling Massive Neutrino Cosmologies with Higher-Order Clustering Statistics

Graduation Thesis

Presented by:
Andrea Labate

Supervisor:
Prof. Michele Ennio Maria Moresco

Co-supervisor:
Dr. Massimo Guidi

Academic Year 2024/2025
Graduation date III

Abstract

Free-streaming of cosmic neutrinos on small scales affects the distribution and growth of cosmic structures, allowing constraints on the sum of their masses from clustering studies. In this context, it is of particular interest to explore the contribution of higher-order clustering statistics, and in particular of the 3-point correlation function (3PCF), since it is the first order capable of quantifying the non-Gaussian properties of the Large-Scale Structure. In this Thesis, we investigate for the first time the possibility of disentangling massive neutrino cosmologies with the 3PCF.

We measure the isotropic 3PCF ζ and the reduced 3PCF Q of halos in the QUIJOTE simulations, which include neutrinos as particles, with $\sum m_\nu = 0.1, 0.2$, and $0.4\,\mathrm{eV}$. We develop a framework to quantify the detectability of massive neutrinos for different triangle configurations and shapes, finding that, as an effect of free-streaming, the largest signal is found in isosceles triangles on scales smaller than $\sim 30\,h^{-1}\,\mathrm{Mpc}$, and in quasi-isosceles and squeezed triangles, increasing for decreasing redshift. For these configurations, we find that elongated triangles, tracing the filamentary structure of the cosmic web, are the most affected by neutrinos, with a 3PCF signal increasing with M_ν. A complementary source of information comes from perpendicular configurations in Q.

We also investigate whether the 3PCF can break the degeneracy between M_ν and σ_8 by measuring both ζ and Q for an additional set of QUIJOTE halo catalogues with $\sigma_8 = 0.819, 0.849$. We find that they respond to variations of M_ν and σ_8 differently as a function of triangle shapes and scales. In particular, the signature of neutrinos on perpendicular configurations in Q results significantly different in varying-σ_8 cosmologies, and hence can be used to effectively break their degeneracy.

ἔνθα καὶ ἠματίη μὲν ὑφαίνεσκεν μέγαν ἱστόν,
νύκτας δ᾽ ἀλλύεσκεν, ἐπὴν δαΐδας παραθεῖτο.

Homer, *Odyssey* 2.104-105

Contents

* * *

> *I have done a terrible thing: I have postulated a particle that cannot be detected!*
>
> Wolfgang E. Pauli, 1930

On 4th December 1930, the physicists convened at the Gauverein meeting in Tübingen received a rather peculiar letter from Zürich. The author himself was quite exceptional: the Austrian theoretical physicist Wolfgang E. Pauli (1900-1958), who was not in Tübingen because, as he had written, his presence was indispensable at a ball in Zürich.

The scientific community of "radioactive Ladies and Gentlemen" (in Pauli's own words) was discussing a rather puzzling phenomenon: electrons emitted during the β-decay of atomic nuclei showed a continuous energy spectrum. This contradicted the conservation of energy and momentum, assuming that the electron is the only particle emitted in the decay. The seriousness of the situation had been emphasized by a remark of Peter Debye (1884-1966), who had told Pauli: "Oh, it's better not to think about this at all, like new taxes!"

Pauli revealed that he had found a "desperate remedy": a particle not yet detected, electrically neutral, with very small mass, emitted during the decay along with the electron.

After the discovery of the neutron by James Chadwick (1891-1974) in 1932, the hypothetical particle was jokingly called "neutrino" in a conversation between the Italian physicists Edoardo Amaldi (1908-1989) and Enrico Fermi (1901-1954); this name entered into common use in the scientific community after the Solvay Conference in 1933.

Neutrinos were discovered in 1956 by Clyde Cowan (1919-1974) and Frederick Reines (1918-1998), who sent Pauli a telegram announcing the discovery. Pauli, who had waited for 26 years for his hypothesis to be confirmed, and who would die two years later, replied with the words: "Everything comes to him who knows how to wait".

Actually, this was just the beginning of the story.

* * *

Introduction

Scientific framework

The Standard Model (SM) of Particle Physics predicts that weak interactions create neutrinos in one of the three leptonic flavor eigenstates: *electron neutrino* (ν_e), *muon neutrino* (ν_μ), and *tau neutrino* (ν_τ). Within the SM, neutrinos are precisely massless. In 1998, the *Super-Kamiokande* experiment presented evidence for neutrino flavor oscillations (Fukuda et al. 1998), a phenomenon possible only if at least two neutrinos have mass. Indeed, according to the theory of neutrino oscillations, neutrinos are produced and detected through weak interactions as flavor eigenstates, but propagate as a quantum superposition of three mass eigenstates (labelled with ν_1, ν_2, and ν_3, with mass m_1, m_2 and m_3, respectively): since the three mass components travel at slightly different speeds, phase shift in their wave packets makes it possible for neutrinos produced with a given flavour to be detected in another flavour eigenstate (e.g., from electronic to muon neutrino). Vacuum oscillation experiments are only sensitive to the quantities Δm^2_{21} and $|\Delta m^2_{32}|$, with $\Delta m^2_{ij} \equiv m_i^2 - m_j^2$. Consequently, the neutrino mass hierarchy, i.e., how m_1, m_2 and m_3 are ordered, is still unknown. The most up-to-date best fit values of the mass splitting are (Particle Data Group: Navas et al. 2024):

$$\Delta m^2_{21}[10^{-5}\,\mathrm{eV}^2] = 7.50^{+0.19}_{-0.18} \qquad \Delta m^2_{32}[10^{-3}\,\mathrm{eV}^2] = 2.451^{+0.026}_{-0.026}\ (\mathrm{NO}),\ -2.527^{+0.034}_{-0.034}\ (\mathrm{IO}).$$

The acronyms NO and IO stand for, respectively, *normal ordering*, in which one assumes $m_1 < m_2 < m_3$, and *inverted ordering*, where the assumed ordering is $m_3 < m_1 < m_2$. By setting the lightest of the three masses to 0, one can determine the smallest possible total mass $\sum m_\nu$ allowed by either ordering:

$$\left(\sum m_\nu\right)^{\mathrm{NO}} > (59.980 \pm 0.304) \times 10^{-3}\,\mathrm{eV} \qquad \left(\sum m_\nu\right)^{\mathrm{IO}} > (99.824 \pm 0.581) \times 10^{-3}\,\mathrm{eV}.$$

Cosmology plays a fundamental role in this respect. Gravity, which rules the evolution of matter perturbations on cosmic scales, is sensitive to the sum of neutrino masses, $\sum m_\nu$, rather than to their splitting. Consequently, the study of the effects that neutrinos imprint on the Large-Scale Structure (LSS) of the Universe, in combination with other cosmological probes, can lead to ruling out the inverted ordering, provided that the total neutrino mass is constrained below the minimal value allowed in that scenario. Furthermore, reaching an accurate description of the effects of massive neutrinos at all scales is important to avoid potential systematics in the inference of other cosmological parameters in present and upcoming Stage IV spectroscopic surveys, like Euclid (Euclid

Collaboration: Mellier et al. 2025), the Dark Energy Spectroscopic Instrument (DESI; DESI Collaboration: Aghamousa et al. 2016), and the Nancy Grace Roman Space Telescope (Dore et al. 2019).

The Big Bang paradigm predicts the existence of a Cosmic Neutrino Background, contributing to the total radiation energy density at very early times, and at later times, when neutrinos became non-relativistic, to the total matter density of both background and perturbations. Due to their large thermal velocities, perturbations of non-relativistic neutrinos are washed out below their free-streaming scale (Bond et al. 1980), while above it they behave like cold dark matter (CDM). As a consequence, in the presence of massive neutrinos, the growth rate of cosmic structures becomes scale-dependent, showing a few percent suppression below the free-streaming scale; this implies a lower number density of cosmic structures formed at a given time with respect to massless neutrino cosmologies sharing the same amplitude of the primordial field of perturbations (see Lesgourgues and Pastor 2006 for a review).

Cosmological perturbations are the seeds for the formation of observable structures, such as galaxies and galaxy clusters; consequently, *galaxy clustering* represents an ideal probe for investigating the imprint of massive neutrinos on structure formation. This field of Cosmology exploits the statistical properties of the fluctuations in the galaxy distribution on large scales, by measuring and analyzing its correlation functions, starting from the *2-point correlation function* (2PCF). The impact of massive neutrinos on 2-point statistics has been extensively studied both in configuration and Fourier space. Neutrinos suppress by a few percent the total matter and CDM power spectra below the free-streaming scale; moreover, results obtained from N-body simulations show that they induce a scale-dependent bias even on large scales, and affect redshift-space distortions (Marulli et al. 2011; Castorina et al. 2014; Villaescusa-Navarro et al. 2014; Castorina et al. 2015; Verdiani et al. 2025). 2PCF data are frequently exploited to constrain the sum of neutrino masses in combination with CMB, yielding the most stringent 95% upper bound on the sum of neutrino masses to date, $\sum m_\nu < 0.064\,\mathrm{eV}$, from DESI (DESI Collaboration: Abdul-Karim et al. 2025).

However, the 2PCF or the power spectrum can completely characterize only Gaussian fields. The non-Gaussian properties of the galaxy distribution are quantified by higher-order statistics, starting from the *3-point correlation function* (3PCF) in configuration space and its Fourier-space counterpart, called *bispectrum*. These statistics can prove useful for capturing the effects of neutrino free-streaming, since small scales are most strongly affected by nonlinearity, which is the main source of non-Gaussianity in the LSS. The first measurement of bispectrum from N-body simulations including massive neutrinos, and the derivation of constraints on bias parameters from it, has been published in Ruggeri et al. (2018), where the bispectrum shows the expected suppression as in the case of the power spectrum, with terms containing neutrino perturbations being highly suppressed below the neutrino free streaming scale.

Early studies about 3PCF focused only on a subset of all the potentially measurable configurations (e.g., Marín et al. 2013; Moresco et al. 2014 to cite, respectively, a study on real data and simulations), due to the high computational cost required to measure it through estimators based on direct triplet counting. The introduction of an estimator based on spherical harmonic expansion (Slepian and Eisenstein 2015a; Slepian and Eisenstein 2018) has significantly reduced the computational burden, allowing for the extension of measurements to larger scales and datasets.

Aim of the thesis

A measurement of the 3PCF in simulations that implement massive neutrino cosmologies has never been performed to date. Bridging this gap is important, as configuration space provides a natural framework for handling survey geometry effects. Conversely, such effects are more challenging to account for in Fourier space, where convolution with complex window functions is employed. More generally, addressing this challenge is fundamental in preparation for the full scientific exploitation of the wealth of data that will be delivered by Stage IV surveys.

The primary objective of this work is to explore the additional information gained by adding 3-point statistics in configuration space in the study of the effects imprinted by massive neutrinos on the LSS, providing the first measurement, to our knowledge, of 3PCF on simulations implementing massive neutrino cosmologies. This thesis will focus on the analysis of the connected 3PCF and reduced 3PCF, searching for signatures that can provide improved constraints on neutrino masses with respect to 2-point statistics.

In addition, we will explore the possibility of breaking the degeneracy between the sum of neutrino masses and other cosmological parameters that could mimick neutrino effects at lower orders, among all σ_8 (the RMS amplitude of present-day linear fluctuations on a scale of 8 $h^{-1}\,\mathrm{Mpc}$), whose degeneracy with $\sum m_\nu$ affects 2-point statistics (Viel et al. 2010; Villaescusa-Navarro et al. 2014; Villaescusa-Navarro et al. 2018).

Methods

We measure and analyze the monopole, quadrupole, and hexadecapole of the 2PCF and the isotropic multipoles of the 3PCF up to $\ell_{\max} = 10$, in a set of 4 000 halo catalogues from simulations that include cosmologies with massless and massive neutrinos, from fully nonlinear scales to large, linear scales. We extract the catalogues from the Quijote simulations suite (Villaescusa-Navarro et al. 2020). We define and apply several parameters, building a framework to systematically identify the triangle shapes and scales for which the neutrino signal is higher, and quantify its strength.

To study the $\sum m_\nu - \sigma_8$ degeneracy, we complement our dataset with measurements of the 2PCF and 3PCF from 1500 halo catalogs extracted from ΛCDM simulations (also taken from the Quijote suite) in which the value of σ_8 varies by ± 0.015 around a fiducial value, and apply part of the framework previously developed.

Outline

This thesis is structured as follows:

- In Chapter 1, we briefly review the fundamental equations of Cosmology in a homogeneous and isotropic Universe, and introduce the ΛCDM model and its open problems. Then, we discuss the clustering of matter and galaxies, starting from a brief overview of bias, the 2PCF, and the power spectrum, followed by an in-depth discussion about the isotropic 3PCF and its multipole moment decomposition. We then address redshift-space distortions and their impact on clustering statistics. Finally, we summarize the main results obtained so far in the context of massive neutrino cosmologies.

- In Chapter 2, we describe our dataset and the estimators we have used for measuring the 2PCF and 3PCF multipoles. Then, we introduce the functions that we have reconstructed from the measured multipole moments and describe the estimation of covariant errors. Finally, we define the parameters and indicators we have developed and applied in our data analysis.

- In Chapter 3, we present the results of the analysis of massive neutrino simulations. We start from the 2PCF, discussing the role of the scalar amplitude A_s and bias in producing the obtained results. We then proceed with the analysis of the multipoles of the 3PCF, followed by the study of the connected and reduced 3PCF for different neutrino masses as a function of the angle between two sides of the triangles, and by the identification of the triangle shapes that yield the strongest signal in the presence of massive neutrinos.

- In Chapter 4, we address the degeneracy between the neutrino mass and σ_8. First, we concisely present the literature results on the topic involving 2- and 3-point statistics. We proceed to describe the additional set of simulations that we have considered in our analysis. Then, we discuss the role of halo bias in the space of degenerate parameters. We conclude by comparing the effects of the considered neutrino masses and variations in σ_8 on the connected and reduced 3PCF, in terms of both triangle shape and scale.

- In Chapter 5, we draw our conclusions and discuss future prospects.

Chapter 1

Theoretical background

In this chapter, we review the main theoretical concepts on which this work is based.

In Sec. 1.1, we begin by recalling the fundamental equations that are valid in a homogeneous and isotropic Universe (i.e., initially neglecting perturbations), and provide an overview of the ΛCDM model, highlighting the reasons for its success as well as the current open problems. We proceed by discussing the framework needed for the study of clustering in Sec. 1.2, focusing on the concept of bias and introducing the necessary statistical tools in configuration and Fourier space, such as the 2-point correlation function/power spectrum and the 3-point correlation function/bispectrum, with the analysis of the latter restricted to its isotropic component. We briefly recall the fundamentals of Lagrangian Perturbation Theory in Sec. 1.3 since it enters the initial conditions of the simulations considered in this work. We then describe the effect of redshift-space distortions on clustering in redshift space in Sec. 1.4. We summarize the effects of massive neutrinos on clustering, providing a state-of-the-art overview of current research in this field in Sec. 1.5. Finally, we discuss the motivations behind this work in Sec. 1.6.

1.1 A homogeneous and isotropic Universe

1.1.1 The fundamental equations of Cosmology

We provide here a summary of the most fundamental equations of Cosmology for a homogeneous and isotropic Universe, primarily following Coles and Lucchin (2002), which will serve as a reference throughout the section.

Cosmology studies the properties and evolution of the Universe on large scales, where the dominant force is gravity. The most accurate theory of gravitation developed so far is Einstein's General Relativity (GR), whose field equations are the following:

$$R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}. \tag{1.1.1}$$

Here, $g_{\mu\nu}$ is the metric tensor, $R_{\mu\nu}$ and R are respectively the Ricci tensor and its trace, the Ricci scalar (which are functions of $g_{\mu\nu}$ and its derivatives), Λ is called Cosmological Constant and $T_{\mu\nu}$ is the energy momentum tensor. The left-hand side of the equations is made up of geometric entities only, whereas the right-hand side encodes information about the distribution of mass and energy of spacetime.

By adopting Einstein's summation convention, the conservation of energy and momentum in GR takes the form

$$\nabla_\mu T^{\mu\nu} = 0, \tag{1.1.2}$$

where ∇_μ denotes the operation of covariant derivative in differential geometry.

In Eqs. 1.1.1, the unknown is the metric tensor $g_{\mu\nu}$. The solution of these equations is, in general, very difficult due to their nonlinearity, and often analytical solutions are obtained after imposing some symmetry properties on $g_{\mu\nu}$.

In a cosmological context, the almost perfect isotropy of the Cosmic Microwave Background (CMB) and the observation of the Large Scale Structure (LSS) on scales larger than $\sim 250\,\mathrm{Mpc}$ have led to the formulation of the Cosmological Principle, which states that, *on sufficiently large scales*, the Universe is *homogeneous* and *isotropic*. The assumptions of homogeneity and isotropy of the Universe translate mathematically into symmetry constraints on the metric tensor $g_{\mu\nu}$. The most general space-time metric describing a Universe in which the Cosmological Principle is obeyed is called the Friedmann-Lemaître-Robertson-Walker (FLRW) metric. Its invariant line element $ds^2 = g_{\mu\nu}dx^\mu dx^\nu$ can be expressed as

$$ds^2 = c^2dt^2 - a(t)^2\left[\frac{dr^2}{1-kr^2} + r^2(d\theta^2 + \sin^2\theta\, d\phi^2)\right], \tag{1.1.3}$$

where r, θ and ϕ are called *comoving coordinates*, t is called *cosmic time* and $a(t)$ is a function of cosmic time called *cosmic scale factor*. The *curvature parameter* k takes the values -1 for a hyperbolic spatial geometry, +1 for spherical geometry, and 0 for flat geometry.

A direct consequence of the Cosmological Principle is the possibility for space to expand or contract. To understand this, let us introduce the concept of *proper distance*, i.e., the distance between two events measured at fixed cosmic time. If we fix the origin at the spatial position of one of the events, while the radial comoving coordinate of the other event is r, then we can write the comoving distance at cosmic time t as

$$d_p \equiv a(t)\int_0^r \frac{dr'}{\sqrt{1-kr'^2}} = a(t)f(r;k), \tag{1.1.4}$$

where f is the function of r (depending on the curvature k), resulting from the evaluation of the integral. Since the proper distance is time-dependent, any comoving observer (i.e., an observer whose comoving coordinates do not change with cosmic time) has a nonzero velocity with respect to an arbitrary origin, which is radially directed and is given by

$$v_{rec} \equiv \dot{d}_p = \frac{\dot{a}}{a}d_p = H(t)d_p. \tag{1.1.5}$$

The quantity v_{rec} is called *recessional velocity*[1], $H(t) \equiv \dot{a}/a$ is called *Hubble parameter*

---

[1] Strictly speaking, an observer sitting in the origin sees the other comoving observers receding from him at a given cosmic time t only if $H > 0$.

and Eq. 1.1.5 is known as *Hubble-Lemaître Law*. The present-day value of the Hubble parameter is called *Hubble constant*, it is denoted by H_0 and is usually parameterized by

$$H_0 = 100\, h \,\mathrm{km\,s^{-1}\,Mpc^{-1}}, \tag{1.1.6}$$

where h is observationally bounded between ~ 0.67 and 0.73.

A fundamental quantity in Cosmology, which makes it possible to translate cosmic time into an observable quantity, is *cosmological redshift*. Suppose that two comoving observers in the Universe are separated by some distance, and one of the two observers sends to the other one a monochromatic electromagnetic wave generated with a wavelength λ_e (where the subscript "e" stands for "emitted"). Let us suppose that this electromagnetic wave is received by the second observer at the present time t_0. In the relativistic framework, electromagnetic signals travel along *null geodesics*, characterized by $ds^2 = 0$. By imposing this condition in the FLRW metric (Eq. 1.1.3), we get that the time interval needed for the arrival of two consecutive wave crests at the second observer, δt_o, is related to the time interval for the generation of the two crests at the position of the emitter, δt_e, by

$$\frac{\delta t_o}{\delta t_e} = \frac{a_0}{a_e}, \tag{1.1.7}$$

where a_0 and a_e are, respectively, the values of the scale factor at time t_0 and at the time of emission t_e. In an expanding Universe, $a_0 > a_e$ and so $\delta t_o > \delta t_e$. Since $\lambda \propto \delta t$, the receiver observes a signal which is shifted towards the redder regions of the spectrum. In particular, we define the cosmological redshift as

$$z \equiv \frac{\lambda_o - \lambda_e}{\lambda_e} = \frac{a_0}{a} - 1. \tag{1.1.8}$$

In the last step, we have used Eq. 1.1.7.

Another important point is how the scale factor a evolves with cosmic time in a homogeneous and isotropic Universe. Its evolution is determined, once we choose a $T_{\mu\nu}$ to describe the large-scale mass-energy distribution in the Universe, by solving Eqs. 1.1.1 with the metric in the FLRW form 1.1.3. Typically, the Universe is treated as a perfect fluid specified by mass-energy density ρ and pressure p, in which thermal conduction and viscous processes are neglected. The energy-momentum tensor then reads

$$T^{\mu}{}_{\nu} = g_{\nu\sigma} T^{\mu\sigma} = \begin{pmatrix} -\rho c^2 & 0 & 0 & 0 \\ 0 & p & 0 & 0 \\ 0 & 0 & p & 0 \\ 0 & 0 & 0 & p \end{pmatrix}, \tag{1.1.9}$$

where the summation over the repeated index σ is understood. By combining Eqs. 1.1.1, 1.1.3 and 1.1.9 one gets two relations called *Friedmann equations*:

$$\dot{a}^2 + kc^2 = \frac{8\pi G}{3}\tilde{\rho} a^2 \tag{1.1.10}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}\left(\tilde{\rho} + \frac{3\tilde{p}}{c^2}\right), \tag{1.1.11}$$

where $\tilde{\rho}$ and $\tilde{p}$ are, respectively, the effective density and pressure containing the contribution of the cosmological constant Λ, in the form

$$\tilde{\rho} \equiv \rho + \rho_\Lambda = \rho + \frac{\Lambda c^2}{8\pi G}, \tag{1.1.12}$$

$$\tilde{p} \equiv p + p_\Lambda = p - \frac{\Lambda c^4}{8\pi G}. \tag{1.1.13}$$

From Eq. 1.1.10, the value of $\tilde{\rho}$ for which $k = 0$ is given by

$$\rho_c = \frac{3H^2}{8\pi G}. \tag{1.1.14}$$

This quantity is called *critical density* of the Universe and makes it possible to rewrite Eq. 1.1.10 in the form

$$\Omega - 1 = \frac{kc^2}{a^2H^2}, \tag{1.1.15}$$

where $\Omega \equiv \tilde{\rho}/\rho_c$ is called *density parameter*. Note that the sign of k (and consequently the curvature of the Universe) *at any time* is determined by the difference between the *present-day* value of the density parameter Ω_0 and unity: $\Omega_0 > 1, \Omega_0 < 1$ and $\Omega_0 = 1$ correspond, respectively, to $k = +1, k = -1$ and $k = 0$.

By combining Eqs. 1.1.2 and 1.1.9 one gets the continuity equation

$$\dot{\rho} + 3\frac{\dot{a}}{a}\left(\rho + \frac{p}{c^2}\right) = 0. \tag{1.1.16}$$

The evolution of ρ with the scale factor a can be obtained from Eq. 1.1.16 after specifying how the variables ρ and p are intrinsically related, i.e., after writing down the *equation of state* (EOS) of the fluid. Typically, the EOS for a given component is expressed in the Zel'dovich form

$$p = w\rho c^2, \tag{1.1.17}$$

where w is called *equation-of-state parameter* (EOS *parameter*). For a single-component Universe and constant w, Eq. 1.1.16 yields:

$$\rho = \rho_0 \left(\frac{a}{a_0}\right)^{-3(1+w)}, \tag{1.1.18}$$

where ρ_0 and a_0 are, respectively, the present-day values of the density and of the scale factor[2].

For a Universe dominated by non-relativistic matter (dust), $w = 0$ and $\rho_m \propto a^{-3}$. For a radiation-dominated Universe, $w = 1/3$ and $\rho_r \propto a^{-4}$. For a fluid dominated by the Cosmological Constant, $w = -1$ and $\rho_\Lambda = \frac{\Lambda c^2}{8\pi G} = \text{const}$. In the case of a multi-component Universe, where a generic species is labeled with i, then $\rho = \sum_i \rho_i$ and we can define

---

[2] The scale factor can be normalized such that $a_0 = 1$.

a density parameter for each component as $\Omega_i \equiv \rho_i/\rho_c$. By combining Eq. 1.1.18 for each component i with the Friedmann Equations, the evolution of the scale factor in a multi-component homogeneous and isotropic Universe is given by

$$\left(\frac{\dot{a}}{a_0}\right)^2 = H_0^2 \left[\sum_i \Omega_{0,i} \left(\frac{a_0}{a}\right)^{1+3w_i} + \left(1 - \sum_i \Omega_{0,i}\right) \right], \tag{1.1.19}$$

where w_i and $\Omega_{0,i}$ are, respectively, the EOS parameter and the present-day density parameter of the species i.

1.1.2 The ΛCDM model

The ΛCDM model or Standard Cosmological Model is the best model to date reproducing cosmological observations. This model assumes GR and the Cosmological Principle, and describes a flat Universe ($\Omega_0 = 1$) in an expansion state starting from an initial singularity, the Big Bang, characterized by divergent temperature and density. During the expansion history of the Universe, the temperature progressively decreased, allowing the different components of the mass-energy fluid, initially coupled, to decouple at different cosmic times.

The name refers to the components which constitute the matter-energy budget of the Universe, specifically:

Dark Energy (DE). This component is responsible for the accelerated expansion of the Universe (Riess et al. 1998; Perlmutter et al. 1999). In the model, it is described in terms of the Cosmological Constant Λ, characterized by a negative fluid pressure $p_\Lambda = -\rho_\Lambda c^2$. DE dominates the present-day total mass-energy budget, with a percent contribution of about 68.3% and a density of $\rho_\Lambda \approx 7 \times 10^{-30}\,\mathrm{g\,cm^{-3}}$.

Cold Dark Matter (CDM). According to current estimates, it accounts for the $\approx 27\%$ of the mass-energy content of the Universe. It is modeled as a non-relativistic, collisionless component that interacts with ordinary baryonic matter and radiation only through gravity. The adjective “cold” comes from the fact that this component is already non-relativistic at decoupling. Although dark matter has never been directly observed due to its non-interacting nature, it is currently required to reproduce several observables such as the rotation curves of galaxies, the dynamics of galaxies in galaxy clusters, gravitational lensing, and LSS. In particular, CDM is believed to build up the scaffolding for the growth of observable cosmic structures (galaxies and galaxy clusters).

Baryonic Matter. the ordinary matter component, building up the remaining $\approx 5\%$ of the mass-energy content. It is mainly composed of light elements, such as Hydrogen and Helium.

The Standard Model also predicts a radiation and neutrino background, i.e., the already mentioned CMB and the so-called Cosmic Neutrino Background (CνB), which are relics of the Big Bang. In particular, the effects of the CνB on the LSS, the central topic of this work, are discussed in Sec. 1.5.

The success of the ΛCDM model stems from the fact that it is able to account for the accelerated expansion of the Universe, the generation and the characteristics of the CMB, the properties of the LSS, baryon acoustic oscillations (BAO; see Sect. 1.2.5), and light-element abundances (Hydrogen, Helium, Lithium), in a quite simple framework based on six cosmological parameters: the baryon and total matter density parameters $\Omega_b h^2$ and $\Omega_m h^2$, the dimensionless Hubble parameter h, the scalar amplitude of the primordial field of density fluctuations A_s, the spectral index of the primordial power spectrum n_s and the ratio of the sound horizon to the angular diameter distance to the surface of last scattering θ_*.

Despite the effectiveness of the ΛCDM model in explaining several results from cosmological observations, several open problems or tensions remain that need to be addressed, some of which may suggest the existence of new physics. We give here a short summary of some of the main open questions, without entering the discussion about their proposed solutions, which is beyond the scope of this work:

H_0 tension. It arises from a discrepancy between the values of the current expansion rate of the Universe, the Hubble constant H_0, estimated either indirectly from CMB observables and BAO data, or directly from standard candles. The most recent estimate, which assumes ΛCDM, coming from the analysis of CMB data released by the space observatory *Planck* is $H_0 = (67.36 \pm 0.54)\,\mathrm{km\,s^{-1}Mpc^{-1}}$ at the 68% C.L. (Planck Collaboration: Aghanim et al. 2020). The second method relies on the calibration of Type Ia Supernovae (SNIa) through the observation of Cepheids, especially employing the *Hubble Space Telescope* (HST): in this respect, the main contender is the SH0ES team, which recently published a measurement of $H_0 = (73.29 \pm 0.90)\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ (Murakami et al. 2023), in 5.6σ discrepancy with Planck result.

S_8 tension. The S_8 parameter is defined as $S_8 \equiv \sqrt{\Omega_m/0.3}\,\sigma_8$ (see Sec. 1.2.4 for the definition of σ_8), and quantifies the level of matter fluctuations on scales of $\sim 8\,h^{-1}\,\mathrm{Mpc}$. The tension arises between more direct measurements of S_8 and the analysis of CMB data assuming ΛCDM. For example, the value obtained by Planck in Planck Collaboration: Aghanim et al. (2020), $S_8 = 0.834 \pm 0.016$, is in 2.9σ discrepancy with the directly measured values from the *Kilo Degree Survey* (KiDS) and the *Dark Energy Survey* (DES), which give $S_8 = 0.769 \pm 0.016$ after combining KiDS-1000 (Heymans et al. 2021) and DES Y3 (DES Collaboration: Abbott et al. 2022) results.

Nature of Dark Matter. The ΛCDM model assumes the existence of CDM, but its particle nature has not been discovered yet. Several experiments aiming at directly detecting CDM particles are ongoing, like *XENONnT*[3] at *Laboratori Nazionali del Gran Sasso* (LNGS) in Italy.

---

[3] https://xenonexperiment.org/

Nature of Dark Energy. The vacuum energy predicted by Quantum Field Theory (QFT) has often been invoked as the main contribution to DE. However, QFT predicts a theoretical contribution of zero-point energy to the Cosmological Constant Λ from 50 to 120 orders of magnitude larger than has actually been observed, a problem known as *Cosmological Constant problem* or *Vacuum catastrophe* (see Adler 2021 for a review). Moreover, recent results from the *Dark Energy Spectroscopic Instrument* (DESI) Collaboration confirm a preference for dynamical (i.e., time evolving) DE over a cosmological constant: when using the parameterization $w = w_0 + w_a(1 - a)$ for the DE EOS parameter, best-fit values are obtained in the quadrant $w_0 > -1$ and $w_a < 0$ (DESI Collaboration: Abdul-Karim et al. 2025).

Small-scale problems. At small scales, CDM N-body simulations produce some results which are incompatible with observations, for example cusped density profiles versus observed cored profiles in dwarfs (*cusp-core problem*, CC), a large number of subhalos than observationally inferred (*missing satellite problem*, MSP) and an overabundance of massive, dense subhaloes compared to the observed number of relatively luminous satellite galaxies of the Milky Way or the Local Group (*Too Big To Fail* problem, TBTF). See Bullock and Boylan-Kolchin (2017) for a review.

Large-scale anomalies. The existence of structures made of galaxy clusters, galaxies, gas, and dust with Gpc sizes, like the Giant Arc (Lopez et al. 2022), has been claimed to challenge the Cosmological Principle, together with possible discrepancies between measurements of the dipole in the angular distribution of quasars and the CMB dipole (Secrest et al. 2021).

1.2 Galaxy clustering

Present-day structures, such as galaxies and galaxy clusters, are the final products of a process of structure formation involving dark matter and baryonic matter, which begins with a primordial field of density fluctuations. In this section, we will describe the necessary statistical tools for analyzing the phenomenon of clustering, up to its present-day evolution, characterized by a nonlinear growth regime.

1.2.1 The need for a statistical approach

The current model of cosmological structure formation states that the matter distribution which we observe today originates from the growth of primordial fluctuations (or perturbations), amplified by gravitational instability, in an otherwise homogeneous and isotropic Universe. The most widely accepted cosmological models attribute the origin of the primordial perturbation field to quantum fluctuations stretched to cosmological scales by *inflation*, a phase of rapid expansion that occurred in the very early stages of the history of the Universe.

Although the equations governing the evolution of this field are deterministic in nature, its study requires a statistical framework. In this context, the Universe is treated as a *stochastic realization of a statistical ensemble*. The motivation for adopting such a perspective is based on the following considerations:

- The primordial distribution of density fluctuations (immediately after the inflationary epoch), providing definite conditions to the evolution equations, is not accessible to direct observation.
- It is impossible to follow the growth of a given cosmic structure because the timescales involved in such a growth are too long. Rather, we can observe different objects at different evolutionary phases.

A field whose statistical properties we want to describe is the *density contrast field* $\delta(\boldsymbol{x}, t)$, defined as

$$\delta(\boldsymbol{x}, t) \equiv \frac{\rho(\boldsymbol{x}, t) - \bar{\rho}(t)}{\bar{\rho}(t)}, \tag{1.2.1}$$

where $\rho(\boldsymbol{x}, t)$ is the matter density at the position $\boldsymbol{x}$ and at cosmic time t and $\bar{\rho}(t)$ is the average matter density (also called *background density*) of the Universe at cosmic time t.

We assume in what follows that cosmic fields are statistically homogeneous and isotropic. A random field δ is *statistically homogeneous* if all the joint multi-point probability distribution functions $p(\delta_1, \delta_2, ...)$[4] and their moments are invariant under translation of the coordinates $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots$. A stochastic field is *statistically isotropic* if $p(\delta_1, \delta_2, ...)$ is invariant under rotations (Bernardeau et al. 2002). In a statistically homogeneous and isotropic field, p and its moments are functions of the norm of the differences between the different points $|\boldsymbol{x}_i - \boldsymbol{x}_j|$. We anticipate that redshift-space distortions, discussed in Sec. 1.4.1, break the isotropy of the redshift-space contrast field, introducing a privileged direction corresponding to the observer's line of sight.

Inflationary models suggest that the primordial contrast field has a nearly Gaussian distribution. This means that given N points $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N$, the probability of obtaining the values of primordial contrasts $\delta_1, \ldots, \delta_N$ at $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N$ is well described by a multivariate Gaussian distribution. If we denote $\boldsymbol{V} \equiv (\delta_1, \ldots, \delta_N)$, this distribution writes

$$P_N(\boldsymbol{V}) = \frac{\det(M)^{\frac{1}{2}}}{(2\pi)^{\frac{N}{2}}} \, e^{-\frac{1}{2}\boldsymbol{V}^T M \boldsymbol{V}}, \tag{1.2.2}$$

where M is the inverse of the covariance matrix C, whose element C_{ij} is given by the expectation value of the product of δ at the points $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$, and the superscript T denotes the transpose of a matrix.

1.2.2 Biased tracers

When dealing with real observations, statistical tools must be applied to galaxies or galaxy clusters, as the matter density field is not directly observable. In cosmological simulations, these tools can also be applied to dark matter halos. All these objects are not faithful

---

[4] here $\delta_i \equiv \delta(\boldsymbol{x}_i)$.

representatives of the underlying matter density field, but only tracers of it; this leads to the problem of how statistics measured using tracers must be interpreted, a problem known as the *bias problem*.

The evolution of the bias between tracers and the underlying field depends on cosmology, and in the case of luminous tracers, it involves the complex and not yet well-understood phenomenon of galaxy formation. Galaxies are at the endpoint of the process of structure formation and are in a nonlinear growth regime. Moreover, their formation and evolution involve baryon-physics processes, non-local phenomena such as tidal field interactions, and non-deterministic processes in δ, with a dependence on other physical quantities, including temperature and metallicity.

From a mathematical point of view, the emergence of bias induces a map between the matter contrast field and the contrast field of the tracers δ_t:

$$\delta \quad \xrightarrow{\text{bias}} \quad \delta_t(\boldsymbol{x}) \equiv \frac{N_t(\boldsymbol{x}) - \bar{N}_t}{\bar{N}_t} \tag{1.2.3}$$

where $N_t(\boldsymbol{x})$ is the number of tracers in a given volume around $\boldsymbol{x}$, and $\bar{N}_t$ is the average number of tracers inside that volume. Unfortunately, a theory that can model such a relation from first principles does not exist. Usually, phenomenological approaches are adopted, in which we enclose our ignorance about the physical phenomena leading to the emergence of bias into some parameters to be fitted to the data.

A simple and often adopted bias relation is the *nonlinear and local bias relation*, which is expressed in terms of a power series (Bernardeau et al. 2002, Eq. 525):

$$\delta_t = \sum_{m=1}^{\infty} \frac{b_m}{m!} \delta^m = b \left(\delta + \frac{\gamma}{2} \delta^2 + \dots \right), \tag{1.2.4}$$

where $\gamma \equiv b_2/b$ and the coefficient $b \equiv b_1$ denotes the *linear bias* parameter. An example of *nonlinear* and *non-local* bias term is associated with tidal forces: it can be written as a function of the tidal tensor s_{ij}, which is defined in terms of the gravitational potential ϕ and δ as

$$s_{ij}(\boldsymbol{x}, t) = \partial_i \partial_j \phi(\boldsymbol{x}, t) + \delta^K_{ij}\, \delta(\boldsymbol{x}, t), \tag{1.2.5}$$

where δ^K denotes the Kronecker delta. In particular, the bias relation in McDonald and Roy (2009) takes the form

$$\delta_t(\boldsymbol{x}, t) = b\, \delta(\boldsymbol{x}, t) + \frac{b_2}{2} \delta^2(\boldsymbol{x}, t) + \frac{b_{s^2}}{2} s^2(\boldsymbol{x}, t) \tag{1.2.6}$$

where $s^2(\boldsymbol{x}, t) = s^{ij}(\boldsymbol{x}, t) s_{ij}(\boldsymbol{x}, t)$ and b_{s^2} is the so-called *tidal bias parameter*.

Additionally, galaxy formation depends on other variables besides the density field, and consequently, the bias relation is *stochastic*. In mathematical terms, if we denote a given deterministic bias relation with $\mathcal{F}$ (like the one in Eq. 1.2.3), we can add the stochastic component by writing the final relation in the form (Bernardeau et al. 2002, Eq. 529)

$$\delta_t(\boldsymbol{x}) = \mathcal{F}[\delta(\boldsymbol{x})] + \varepsilon_\delta(\boldsymbol{x}) \tag{1.2.7}$$

in which a scatter ε_δ in the bias relation at a given δ is introduced.

We now begin to discuss the lowest-order clustering statistics, i.e., 2-point statistics.

1.2.3 2-point statistics

2-point statistics takes into account pairs of objects in space, and it is the simplest one that can be constructed.

Suppose we want to answer the following question: what is the *excess probability*, with respect to a completely random distribution of matter particles (i.e., in the absence of clustering), of finding two of these particles separated by a vector $\boldsymbol{r}$? The joint probability $dP_{1,2}$ of finding two particles in the comoving volume elements dV_1 and dV_2, centered at the positions $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$, and separated by a vector $\boldsymbol{r} = \boldsymbol{x}_2 - \boldsymbol{x}_1$, can be expressed, under the assumptions of statistical homogeneity and isotropy of the Universe, as:

$$dP_{1,2} \propto \bar{n}^2 \, dV_1 dV_2 \, [1 + \xi(r)] \, , \tag{1.2.8}$$

where $r = |\boldsymbol{r}|$, $\bar{n}$ is the average number density of particles and $\xi(r)$ is the so-called *2-point correlation function* (2PCF), defined as the joint ensemble average of the density contrast at two different locations (Bernardeau et al. 2002, Eq. 115):

$$\xi(r) \equiv \langle \delta(\boldsymbol{x}) \delta(\boldsymbol{x} + \boldsymbol{r}) \rangle \, , \tag{1.2.9}$$

where we have omitted, as we will do in the following if not otherwise specified, the time dependence in δ.

It is also convenient from the theoretical point of view to consider the Fourier space analogue of the 2PCF, called *power spectrum*. If we denote with $\tilde{\delta}(\boldsymbol{k})$ a Fourier mode of the density contrast field with wavevector $\boldsymbol{k}$, namely

$$\tilde{\delta}(\boldsymbol{k}) = \int d^3x \, \delta(\boldsymbol{x}) \, e^{-i\boldsymbol{k}\cdot\boldsymbol{x}} \, , \tag{1.2.10}$$

the power spectrum $P(k)$ is defined in terms of correlators in Fourier space through the relation

$$\left\langle \tilde{\delta}(\boldsymbol{k}) \tilde{\delta}(\boldsymbol{k'}) \right\rangle \equiv (2\pi)^3 \delta^{(D)}(\boldsymbol{k} + \boldsymbol{k'}) P(k), \tag{1.2.11}$$

where the symbol $\delta^{(D)}$ denotes the Dirac delta. It can be shown, by expanding the density contrast in Eq. 1.2.9 into Fourier modes and using Eq. 1.2.11, that the 2PCF and the power spectrum are a Fourier pair, i.e., they satisfy the relation

$$\xi(r) = \int \frac{d^3k}{(2\pi)^3} P(k) \, e^{i\boldsymbol{k}\cdot\boldsymbol{r}} \, . \tag{1.2.12}$$

The power spectrum allows us to easily express the variance σ^2 of the fluctuations of the density field δ, defined as $\sigma^2 \equiv < \delta^2 >$ (we recall that the field δ has zero mean by definition). Indeed, starting from the definition, we get

$$\begin{aligned}\sigma^2 \equiv \left\langle \delta(\boldsymbol{x})^2 \right\rangle &= \frac{1}{(2\pi)^6} \int d^3k\, d^3k' e^{i(\boldsymbol{k}+\boldsymbol{k}')\cdot\boldsymbol{x}} \left\langle \tilde{\delta}(\boldsymbol{k})\, \tilde{\delta}(\boldsymbol{k}') \right\rangle \\ &= \frac{1}{(2\pi)^3} \int d^3k\, d^3k' e^{i(\boldsymbol{k}+\boldsymbol{k}')\cdot\boldsymbol{x}}\, \delta^{(D)}(\boldsymbol{k}+\boldsymbol{k}') P(k) \\ &= \frac{1}{2\pi^2} \int_0^\infty dk\, k^2 P(k),\end{aligned} \tag{1.2.13}$$

where we used Eqs. 1.2.10 and 1.2.11, and in the last step, we switched to polar coordinates.

If the field δ were *perfectly* Gaussian, then as a consequence of Wick's theorem (see App. A) any ensemble average of the product of variables $\left\langle \tilde{\delta}(\boldsymbol{k}_1) \cdot \ldots \cdot \tilde{\delta}(\boldsymbol{k}_n) \right\rangle$ could be written as a function of the ensemble average of the product of pairs $\left\langle \tilde{\delta}(\boldsymbol{k}_i)\ \tilde{\delta}(\boldsymbol{k}_j) \right\rangle$, with $i, j \in \{1, \ldots, n\}$. This means, as clearly shown by Eq. 1.2.11 and 1.2.12, that in the case of a Gaussian field with zero mean, all the statistical information is enclosed in the 2PCF or, equivalently, in the power spectrum. However, the following reasons explain why 2-point statistics is not sufficient to characterize the density field of the Universe:

- Primordial fluctuations might not be perfectly Gaussian. We refer to this phenomenon as *primordial non-Gaussianity* (PNG; see Celoria and Matarrese 2018 for a review).
- The dynamics of gravitational instability is nonlinear, so even if the primordial density field were perfectly Gaussian, nonlinear evolution would lead to the development of non-Gaussianity with cosmic time. Another source of nonlinearity in the field of matter tracers is introduced by bias, as is evident from the fact that Eqs. 1.2.4 and 1.2.6 are nonlinear in δ.

This requires going beyond the study of 2-point statistics, starting from 3-point statistics, which we describe in Sec. 1.2.6. Before entering this discussion, we briefly describe the filtering of the density field and the motivation behind its introduction, and the shape and features of the power spectrum and the 2PCF.

1.2.4 Density field filtering

From a practical point of view, we cannot measure the density contrast field with "infinite" resolution. To infer the properties of the underlying continuous matter density field, one must consider discrete objects, like dark matter halos or galaxies. Since these objects are characterized by a given spatial scale, it is useful to consider the density contrast field δ smoothed below a given scale R. This filtering action produces another field, which we denote with δ_R, whose value at a given position is obtained by suitably averaging the original density field over a region of size R around that position, whose shape is determined by the choice of a filter. Mathematically, we can write

$$\delta_R(\boldsymbol{x}) = \int d^3x'\, \delta(\boldsymbol{x}')\, W(\boldsymbol{x}'; \boldsymbol{x}, R), \tag{1.2.14}$$

where $W(\boldsymbol{x}'; \boldsymbol{x}, R)$ denotes a generic filtering function centered on $\boldsymbol{x}$ and acting on the scale R. The operation performed in Eq. 1.2.14 is a convolution. Common choices for W are the Gaussian filter or the top-hat filter, respectively defined (in 3D) by

$$W_G(\boldsymbol{x}; \boldsymbol{x}', R) \equiv \frac{1}{(2\pi R^2)^{3/2}} \exp\left\{ -\frac{||\boldsymbol{x} - \boldsymbol{x}'||^2}{2R^2} \right\}, \tag{1.2.15}$$

$$W_{TH}(\boldsymbol{x}; \boldsymbol{x}', R) \equiv \begin{cases} \frac{3}{4\pi R^3} & \text{if } ||\boldsymbol{x} - \boldsymbol{x}'|| < R, \\ 0 & \text{otherwise.} \end{cases} \tag{1.2.16}$$

The variance of the filtered density field σ_R^2 can be computed with a procedure analogous to Eq. 1.2.13, by replacing δ with δ_R and recalling that the Fourier transform of the convolution of two functions is equal to the product of their Fourier transforms. This yields

$$\sigma_R^2 = \frac{1}{2\pi^2} \int_0^\infty dk\, k^2\, P(k) \left| \tilde{W}(k; R) \right|^2, \tag{1.2.17}$$

where $\tilde{W}$ denotes the Fourier transform of the filter. Notice that Eq. 1.2.17 shows that σ_R gives information on the amplitude of perturbations on the scale R. For historical reasons, a common choice to specify the density perturbation amplitude is through the root mean square (RMS) of the present linear-theory density field filtered on a scale $R = 8\, h^{-1}\mathrm{Mpc}$. This quantity is one of the fundamental cosmological parameters, known as σ_8.

1.2.5 2-point statistics shape: Fourier modes and physical separations

According to current cosmological models, the formation of the primordial distribution of potential (i.e, metric) fluctuations, whose subsequent evolution gave rise to the present-day cosmic structures, is due to a rapid phase of accelerated expansion that occurred in the early Universe, known as inflation.

The absence of preferred scales predicted by standard models of inflation translates into a power-law primordial power spectrum P_{in},

$$P_{\mathrm{in}}(k) = A_s \left(\frac{k}{k_p} \right)^{n_s}, \tag{1.2.18}$$

where A_s and n_s are referred to as the *scalar amplitude* and *spectral index*, respectively, and k_p is the so-called *scalar pivot*. Actually, several inflationary models predict a scale-independent amplitude of primordial potential fluctuations; this is realized for $n_s = 1$, and the corresponding power spectrum $P_{\mathrm{in}} \propto k$ is referred to as *Harrison-Zel'dovich power spectrum*[5] (Harrison 1970; Zeldovich 1972).

---

[5] Notice that, strictly speaking, this is true for a pure exponential inflationary epoch (see the discussion in Coles and Lucchin 2002, Sec. 13.6). Current constraints from CMB are consistent with little scale dependence with $n_s = 0.9652 \pm 0.0042$ (Planck Collaboration: Aghanim et al. 2020).

As perturbations evolve from inflation to the present epoch, their amplitudes are modified in a scale-dependent way. Perturbations on scale R are in *linear regime* if $\sigma_R^2 \ll 1$. The linear matter power spectrum can be expressed at a generic epoch as

$$P_{m,\mathrm{lin}}(k,a) = D_1^2(a)\, T^2(k)\, P_{\mathrm{in}}(k), \tag{1.2.19}$$

where the time and scale dependence can be separated and enclosed, respectively, in the functions $D_1(a)$ and $T(k)$, called, respectively, *linear growth factor* (see Sec. 1.3.2) and *transfer function*.

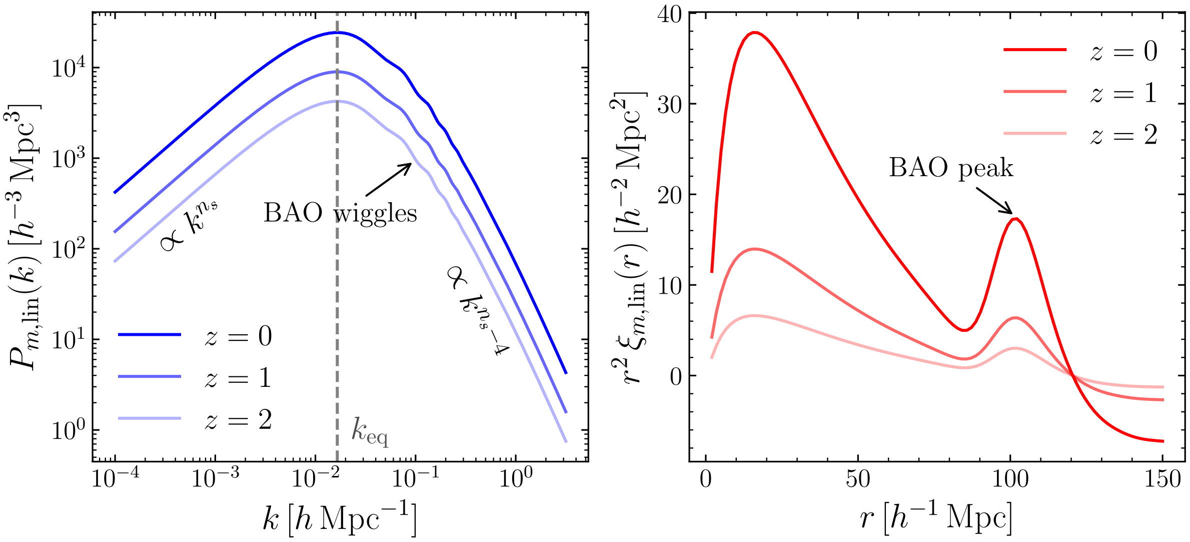

Figure 1.1: The linear matter power spectra (*left*) at $z = 0, 1$, and 2 for a *Planck18* cosmology, generated with the code CAMB (Lewis et al. 2000), and their corresponding matter 2PCFs (*right*). The wavenumber k_{eq} corresponding to the horizon scale at equivalence is marked with a vertical dashed line, and the k dependence of the power spectrum for $k < k_{\mathrm{eq}}$ and $k > k_{\mathrm{eq}}$ is reported. The positions of the BAO wiggles in the power spectra and the BAO peak in the 2PCFs are indicated with arrows.

The linear matter power spectra at $z = 0, 1$, and 2 are plotted in Fig. 1.1, together with the corresponding linear 2PCFs. On large scales (small k), the linear power spectrum preserves the primordial logarithmic slope n_s, while on small scales (large k), it decreases with slope $n_s - 4 \approx -3$ (this value for a Harrison-Zel'dovich primordial power spectrum). These two regimes are separated by a maximum at $k_{\mathrm{eq}} = 2\pi/\lambda_{\mathrm{eq}} \approx 2 \times 10^{-2}\, h\,\mathrm{Mpc}^{-1}$, where λ_{eq} is the size of the comoving particle horizon[6] at the time of matter-radiation equivalence (i.e., when $\rho_r = \rho_m$), occurred at redshift $z_{\mathrm{eq}} \sim 3400$. Modes with $k > k_{\mathrm{eq}}$ entered the horizon at some time during radiation domination ($\rho_r \gg \rho_m$); in this

---

[6] The *comoving particle horizon* is defined as the maximum comoving distance that light can travel from the Big Bang ($t = 0$) up to a given cosmic time t. It thus delineates the boundary of the region of the Universe that can be causally connected to a given observer at time t (i.e., its *observable universe*). By setting $c = 1$, it is given by $\eta(t) = \int_0^t dt'/a(t')$. Notice that this quantity is a non-decreasing function of t.

period, the sub-horizon growth of the amplitude of matter perturbations was strongly suppressed, a phenomenon called *Mészáros effect* or *stagnation* (Meszaros 1974). Modes with larger k entered the horizon earlier, suffering for a longer time from stagnation, so that for increasing $k > k_{\mathrm{eq}}$ the power spectrum decreases. Modes with $k < k_{\mathrm{eq}}$ did not enter the horizon before equivalence. Being their growth factor scale-independent, the power spectrum preserved the primordial slope. For $z < z_{\mathrm{eq}}$, during matter domination, both super-horizon and sub-horizon fluctuations are characterized by the same scale-independent growth factor $D_1(a > a_{\mathrm{eq}}) \propto a$, and the linear power spectrum evolution simply corresponds to a k-independent rescaling.

The wiggles visible for k just above k_{eq}, which are Fourier-transformed into the peak of the 2PCF at $\approx 105\, h^{-1}\,\mathrm{Mpc}$, are the signature of *baryon acoustic oscillations* (BAO), generated by the propagation of sound waves in the photon-baryon plasma prior to recombination ($z \gtrsim 1100$). After recombination and structure formation, BAO remained imprinted in the LSS as overdensities of structures on the comoving scale of sound horizon at recombination r_s. BAO are a major cosmological probe: since their comoving size is known, they are used as *standard rulers* to study the expansion history of the Universe, especially to shed light on the nature of dark energy, and constrain cosmological parameters. Entire surveys have been and are currently dedicated to the study of BAO features, as BOSS (*Baryon Oscillation Spectroscopic Survey*; Schlegel et al. 2007), or the already mentioned DESI.

Several phenomena change the shape of the power spectrum (and, consequently, the 2PCF) as perturbations evolve. Modes with large $k = 2\pi/R$ are the first to become nonlinear ($\sigma_R^2 \sim 1$), and this manifests with an increase of power at large k. In mixed dark matter scenarios, where a component is made up of light, nonrelativistic particles, the phenomenon of *free streaming* produces a relative suppression on small scales, as in the case of massive neutrinos, as detailed in Sec. 1.5.

1.2.6 3-point statistics

The lowest-order clustering statistics that can encode information about non-Gaussian properties of the LSS is 3-point statistics.

By following the approach of Sec. 1.2.3, we can try this time to answer this question: what is the excess probability, with respect to a perfectly Gaussian distribution of matter particles, of finding three particles at the vertices of a triangle of given shape and size? This time, the joint probability $dP_{1,2,3}$ of finding three particles in the comoving volume elements dV_1, dV_2 and dV_3 at the positions $\boldsymbol{x}_1$, $\boldsymbol{x}_2$ and $\boldsymbol{x}_3$, separated by the vectors $\boldsymbol{r} = \boldsymbol{x}_2 - \boldsymbol{x}_1$, $\boldsymbol{s} = \boldsymbol{x}_3 - \boldsymbol{x}_1$, $\boldsymbol{t} = \boldsymbol{x}_3 - \boldsymbol{x}_2$, can be expressed, under the assumptions of statistical homogeneity and isotropy of the Universe, as:

$$dP_{1,2,3} \propto \bar{n}^3\, dV_1 dV_2 dV_3 \left[1 + \xi(r) + \xi(s) + \xi(t) + \zeta(r,s,t)\right], \tag{1.2.20}$$

where $r = |\boldsymbol{r}|$, $s = |\boldsymbol{s}|$, $t = |\boldsymbol{t}|$ and $\zeta(r,s,t)$ is the so-called *connected 3-point correlation function* (3PCF), defined as the joint ensemble average of the density contrast at three different locations:

$$\zeta(r,s,t) \equiv \langle \delta(\boldsymbol{x})\delta(\boldsymbol{x}+\boldsymbol{r})\delta(\boldsymbol{x}+\boldsymbol{s})\rangle\,, \tag{1.2.21}$$

with the constraint $|\boldsymbol{s}-\boldsymbol{r}| = t$.

Following a completely analogous line of reasoning as in Sec. 1.2.3, one can introduce the Fourier space analogue of the 3PCF, the *bispectrum* $B(\boldsymbol{k}_1,\boldsymbol{k}_2,\boldsymbol{k}_3)$, in terms of correlators, so that

$$\left\langle \tilde{\delta}(\boldsymbol{k}_1)\tilde{\delta}(\boldsymbol{k}_2)\tilde{\delta}(\boldsymbol{k}_3)\right\rangle \equiv (2\pi)^3\delta^{(D)}(\boldsymbol{k}_1+\boldsymbol{k}_2+\boldsymbol{k}_3)B(\boldsymbol{k}_1,\boldsymbol{k}_2,\boldsymbol{k}_3). \tag{1.2.22}$$

With this definition, the 3PCF and the bispectrum are a Fourier pair since we have

$$\zeta(r,s,t) = \int \frac{d^3k_2}{(2\pi)^3}\frac{d^3k_3}{(2\pi)^3} B(\boldsymbol{k}_1,\boldsymbol{k}_2,\boldsymbol{k}_3)\, e^{i(\boldsymbol{k}_2\cdot\boldsymbol{r}+\boldsymbol{k}_3\cdot\boldsymbol{s})}\,. \tag{1.2.23}$$

The connected 3PCF or, equivalently, the bispectrum encodes the non-Gaussian information present in the density field. Indeed, notice that, from Wick's theorem, under the hypothesis of a perfectly Gaussian field, the left-hand side of Eq. 1.2.22 vanishes, consequently $\zeta(r,s,t)=0$ identically.

1.2.7 Configuration space vs Fourier space

The information encoded in a given clustering statistics is the same in configuration space and Fourier space. Indeed, as we have seen for the particular cases of 2-point statistics with the 2PCF and the power spectrum, and 3-point statistics with the 3PCF and the bispectrum, these functions are Fourier pairs.

In practice, performing measurements in a real survey is easier in configuration space than in Fourier space. Indeed, a survey covers only a limited volume and exhibits an irregular geometry, with incomplete regions and masks applied to foreground-contaminated areas. Its shape is often characterized by jagged boundaries, resulting from the finite field of view of the instrument, which requires splitting the surveyed area into smaller subregions, or "tiles".

In full generality, the observed galaxy field $\delta_{g,\mathrm{obs}}$ can be written in configuration space in terms of the full galaxy field δ_g as

$$\delta_{g,\mathrm{obs}}(\boldsymbol{r}) = \delta_g(\boldsymbol{r})\,\mathcal{W}(\boldsymbol{r}), \tag{1.2.24}$$

where $\mathcal{W}(\boldsymbol{r})$ is the so-called *window function*, which takes into account all the effects mentioned above. In Fourier space, the product of Eq. 1.2.24 becomes a convolution between the Fourier transforms of δ_g and $\mathcal{W}$, namely

$$\tilde{\delta}_{g,\mathrm{obs}}(\boldsymbol{k}) = \int d^3k'\,\tilde{\delta}_g(\boldsymbol{k'})\,\tilde{\mathcal{W}}(\boldsymbol{k}-\boldsymbol{k'}). \tag{1.2.25}$$

Eq. 1.2.25 shows that survey geometry introduces spurious mode coupling (i.e., not related to physical processes) in Fourier space, through the Fourier transform of the window function. Removing this effect from the observed power spectrum or bispectrum through a deconvolution requires the difficult task of computing $\tilde{\mathcal{W}}$.

This problem is absent in configuration space. However, in general, estimating the 2PCF and 3PCF with a brute-force approach is computationally more expensive than estimating their Fourier counterparts, as it relies on counting pairs and triplets of objects. The computational cost of this counting scales with the number of objects N as $\mathcal{O}(N^2)$ and $\mathcal{O}(N^3)$, respectively, making the estimation of the 3PCF particularly prohibitive for large datasets. The solution, as we shall discuss in Sec. 1.2.8 and 1.2.9, has been provided by the development of algorithms based on the spherical harmonic decomposition of the density field around each object, enabling a full exploitation of the information potentially encoded in the 3PCF. These methods reduce the scaling to $\mathcal{O}(N^2)$, based on the identification of pairs formed by each object with its neighbors, with the correlation then computed as a final, fast step enabled by the decomposition.

The generation of 3PCF models also suffers from computational cost issues. In general, models are developed in Fourier space, exploiting the greater level of decoupling of the k-modes with respect to configuration scales, and then Fourier-transformed to obtain configuration-space models. This operation becomes particularly prohibitive in 3-point statistics, where obtaining a 3PCF model from the corresponding bispectrum requires the numerical evaluation of a 2D Fourier transform. The generation of a single model can demand up to $\mathcal{O}(10)$ CPU hours, effectively preventing a proper exploration of the cosmological parameter space. Recently, this gap has begun to be bridged by the development of emulators (Euclid Collaboration: Guidi et al. 2025), i.e., neural networks trained on libraries of 3PCF models obtained using the standard approach, which can generate accurate 3PCF predictions in under a second per model, given the cosmological parameters as input.

All these advancements contribute to an increasingly efficient exploitation of the information encoded in higher-order statistics in configuration space, which remains the most natural framework from an observational perspective.

1.2.8 Measuring the isotropic 3-point correlation function

In general, a triangle in space is specified by nine variables, e.g., the coordinates of its vertices. However, under the assumption of statistical homogeneity and isotropy, the 3PCF *in real space* is a function of only three variables. Indeed, from the hypothesis of homogeneity, we can conclude that two triangles that differ from each other only by a translation contribute in the same way to 3-point statistics, while from the hypothesis of isotropy, we conclude the same for triangles that differ only by rotations. In other words, under our assumptions, the 3PCF in real space depends only on the shape and scale of a given triangle, and averaging over translations and rotations does not result in a loss of information.

A problem in dealing with 3-point statistics is how to visualize the information. A possibility is to consider a fixed shape, for example, by focusing only on equilateral triangles, so that one can study the dependence of the 3PCF on the scale. Another approach to addressing the problem is to fix the length of two given sides of a triangle, compute the 3PCF by varying the angle between them, and repeat the procedure for different side lengths to study both the dependence on scale and shape.

A more convenient way to mathematically describe the dependence of the 3PCF on its variables is represented by multipole decomposition in different bases. This approach proves to be highly effective in rapidly translating results between configuration space and Fourier space, and it is also well-suited for graphical representation. The reason for this will be clear in the subsequent discussion.

However, the most important reason for the success of multipole-based approaches is the reduction in computational time required to obtain a 3PCF measurement for a given sample of objects. As anticipated in the previous section, the most straightforward way to evaluate the spatial correlation of triplets of objects is to perform a direct count of triplets. Assuming that we have a sample of N objects with a mean number density $\bar{n}$, and that we wish to measure the 3PCF up to a certain scale $R_{\max}$, then there are $N(\bar{n}V_{R_{\max}})$ relevant triangles, where $V_{R_{\max}} = (4/3)\pi R_{\max}^3$: this means that an algorithm based on direct triplet counting is characterized by an $\mathcal{O}(N^3)$ scaling. A fundamental contribution came from Slepian and Eisenstein (2015a), in which the authors introduced the first algorithm for the computation of the isotropic 3PCF characterized by an $\mathcal{O}(N^2)$ scaling, based on a multipole-decomposition approach, which is the subject of the following section.

1.2.9 The Slepian-Eisenstein 2015 algorithm

In Slepian and Eisenstein (2015a), the authors present the first algorithm for the computation of the isotropic 3PCF with $\mathcal{O}(N^2)$ scaling. Previous algorithms scaled as the number of triangles in a survey, i.e., with $\mathcal{O}(N^3)$ scaling. This algorithm employs the spherical harmonic decomposition (SHD) of the density field surrounding the objects in a given survey.

The algorithm adopts the decomposition of the 3PCF first advanced in Szapudi (2004), which we briefly describe. If we label the triangle vertices with $1, 2, 3$, then triangle configurations can be parametrized by the length of the two sides with extremes 1-2 and 1-3, r_{12} and r_{13}, and the angle θ between them. It is natural to decompose the 3PCF into Legendre polynomials in this angular dependence:

$$\zeta(r_{12}, r_{13}, \cos\theta) = \sum_{\ell=0}^{\infty} \zeta_\ell(r_{12}, r_{13}) P_\ell(\cos\theta) \tag{1.2.26}$$

where P_ℓ denotes the Legendre polynomial of degree ℓ and the scale-dependent coefficients $\zeta_\ell(r_{12}, r_{13})$ are given by

$$\zeta_\ell(r_{12}, r_{13}) = \frac{2\ell+1}{2} \int_{-1}^{+1} d\cos\theta \; \zeta(r_{12}, r_{13}, \cos\theta) P_\ell(\cos\theta). \tag{1.2.27}$$

A completely analogous decomposition holds for the bispectrum (simply replace ζ with B in the previous equations, and the configuration space variables with Fourier space variables). It is possible to show that the 3PCF and the bispectrum multipoles are related through a Hankel transform (Szapudi 2004, Eq. 3):

$$\zeta_\ell(r_{12}, r_{13}) = (-1)^\ell \int \frac{dk_1}{2\pi^2}\, k_1^2\, j_\ell(k_1 r_{12}) \int \frac{dk_2}{2\pi^2}\, k_2^2\, j_\ell(k_2 r_{13})\, B_\ell(k_1, k_2), \tag{1.2.28}$$

where j_ℓ are the spherical Bessel functions of order ℓ. However, we expect that only a few multipoles are needed to capture the angle dependence, since the shape of the 3PCF for fixed side lengths as a function of the angle is smooth and slowly varying.

The algorithm works as follows: let

$$\bar{\delta}(r_i;\hat{r}_i;\boldsymbol{s}) \equiv \int dr\; r^2 \Phi(r;r_i)\delta(\boldsymbol{r}_i+\boldsymbol{s}) \tag{1.2.29}$$

be the radially binned density field about a given origin $\boldsymbol{s}$ (a given object of the survey), where Φ is a binning function whose second argument specifies the bin of interest. Let us define the following binned estimator:

$$\bar{\hat{\zeta}}_\ell(r_{12},r_{13};\boldsymbol{s}) \equiv \frac{2\ell+1}{(4\pi)^2}\int d\hat{r}_{12}\, d\hat{r}_{13}\; \delta(\boldsymbol{s})\, \bar{\delta}(r_{12};\hat{r}_{12};\boldsymbol{s})\, \bar{\delta}(r_{13};\hat{r}_{13};\boldsymbol{s})\, P_\ell(\hat{r}_{12}\cdot\hat{r}_{13}), \tag{1.2.30}$$

Notice that once the $\bar{\hat{\zeta}}_\ell(r_{12},r_{13};\boldsymbol{s})$ have been obtained for each element of the survey, the radially binned estimated multipole coefficients can be computed with a translation average:

$$\bar{\zeta}_\ell(r_{12},r_{13}) = \frac{1}{V}\int d^3s\; \bar{\hat{\zeta}}_\ell(r_{12},r_{13};\boldsymbol{s}), \tag{1.2.31}$$

where V is the survey volume. The point is to obtain the estimators $\bar{\hat{\zeta}}_\ell(r_1,r_2;\boldsymbol{s})$ with a method which requires an $\mathcal{O}(N^2)$ time.

We can use the spherical harmonic addition theorem to rewrite the Legendre polynomials in Eq. 1.2.30, i.e., the relation (the superscript * denotes the complex conjugation)

$$P_\ell(\hat{r}_{12}\cdot\hat{r}_{13}) = \frac{4\pi}{2\ell+1}\sum_{m=-l}^{+l} Y_{\ell m}(\hat{r}_{12})Y^*_{\ell m}(\hat{r}_{13}) \tag{1.2.32}$$

where $Y_{\ell m}(\hat{r})$ denotes the spherical harmonic function of degree ℓ ($\ell = 0,1,2,\ldots$) and order m ($m \in \{-\ell,-\ell+1,\ldots,\ell-1,\ell\}$), which can be represented in terms of a complex exponential and the associated Legendre polynomial $P_{\ell m}$ as

$$Y_{\ell m}(\hat{r}) = (-1)^m\sqrt{\frac{2\ell+1}{4\pi}\frac{(\ell-m)!}{(\ell+m)!}}\,P_{\ell m}(\cos\theta)\,e^{im\varphi}, \quad \hat{r}\equiv(\theta,\varphi). \tag{1.2.33}$$

With the normalization above, $\{Y_{\ell,m}\}$ are orthonormal.

Eq. 1.2.30 then becomes:

$$\bar{\hat{\zeta}}_\ell(r_{12},r_{13};\boldsymbol{s}) = \frac{1}{4\pi}\delta(\boldsymbol{s})\sum_{m=-\ell}^{+\ell}\int d\hat{r}_{12}\; \bar{\delta}(r_{12};\hat{r}_{12};\boldsymbol{s})\; Y_{\ell m}(\hat{r}_{12}) \times \int d\hat{r}_{13}\; \bar{\delta}(r_{13};\hat{r}_{13};\boldsymbol{s})\; Y^*_{\ell m}(\hat{r}_{13}). \tag{1.2.34}$$

Let us introduce the coefficients

$$
\begin{aligned}
a_{\ell m}(r;\boldsymbol{s}) &\equiv \int d\hat{r}\ \bar{\delta}(r;\hat{r};\boldsymbol{s})\ Y^*_{\ell m}(\hat{r}) \\
&= \int d\hat{r}\ Y^*_{\ell m}(\hat{r}) \int dr'\ r'^2 \Phi(r';r)\delta(\boldsymbol{r'}+\boldsymbol{s}). \qquad (1.2.35)
\end{aligned}
$$

Then, Eq. 1.2.30 can be written in the form

$$
\hat{\bar{\zeta}}_\ell(r_{12}, r_{13}; \boldsymbol{s}) = \frac{1}{4\pi}\delta(\boldsymbol{s}) \sum_{m=-\ell}^{+\ell} a_{\ell m}(r_{12};\boldsymbol{s})\, a^*_{\ell m}(r_{13};\boldsymbol{s}). \qquad (1.2.36)
$$

In practice, one chooses a maximum scale $r_{\max}$ and a radial binning, sits on each object of the survey (whose position is indicated by $\boldsymbol{s}$) and computes, for each radial bin and up to $r_{\max}$, the coefficients $a_{\ell m}(r;\boldsymbol{s})$ with Eq. 1.2.35. These coefficients allow building the $\hat{\bar{\zeta}}_\ell$ using Eq. 1.2.36. Notice that if the number of radial bins is N_{bins}, we need only to construct $N_{\text{bins}}(N_{\text{bins}}+1)/2$ combinations of these coefficients. Eventually, the translation-averaged multipoles of the 3PCF $\bar{\zeta}_\ell$ are derived through Eq. 1.2.31, which reduces to a sum with N terms since the objects are discrete. Given that estimating the coefficients $a_{\ell m}(r;\boldsymbol{s})$ takes only as long as computing the integral in Eq. 1.2.35, which scales as the mean number of objects contained inside a sphere of radius $r_{\max}$, this process plus the translation average scale as $\mathcal{O}(N^2)$.

A completely analogous approach is followed for the projected 3PCF, defined as the 3PCF integrated along the line of sight (l.o.s.), described in detail in Slepian and Eisenstein (2015a), Sec. 3.

1.3 Basics of Lagrangian Perturbation Theory

Here, we briefly recall the fundamental concepts of Lagrangian Perturbation Theory (LPT) for completeness, as it is involved in the methods for generating the initial conditions for the simulations used in this work, as detailed in Ch. 2. We will refer to Bernardeau et al. (2002) throughout this section, and in place of cosmic time we will use the *conformal time* τ, defined such that $d\tau \equiv dt/a$.

1.3.1 The Lagrangian framework

In the Lagrangian framework, the evolution of the cosmological fluid is studied by following the trajectories of particles or fluid elements. In this respect, the object of interest is the so-called *displacement field* $\boldsymbol{\Psi}(\boldsymbol{q},\tau)$: this field relates the initial positions of the fluid elements $\boldsymbol{q}$ to the final position $\boldsymbol{x}$, i.e.,

$$
\boldsymbol{x}(\tau) = \boldsymbol{q} + \boldsymbol{\Psi}(\boldsymbol{q},\tau). \qquad (1.3.1)
$$

The density field obeys the condition $\bar{\rho}\,[1+\delta(\boldsymbol{x})]\,d^3x = \bar{\rho}\,d^3q$, so the Jacobian of this transformation is

$$J(\boldsymbol{q},\tau) \equiv \mathrm{Det}(\delta_{ij} + \Psi_{i,j}) = [1+\delta(\boldsymbol{x},\tau)]^{-1}, \tag{1.3.2}$$

where $\Psi_{i,j} \equiv \partial \Psi_i / \partial \boldsymbol{q}_j$.

The equation of motion for the trajectories $\boldsymbol{x}(\tau)$ is

$$\frac{\mathrm{d}^2\boldsymbol{x}}{\mathrm{d}\tau^2} + \mathcal{H}(\tau)\frac{\mathrm{d}\boldsymbol{x}}{\mathrm{d}\tau} = -\nabla_{\boldsymbol{x}}\Phi, \tag{1.3.3}$$

where $\mathcal{H} \equiv aH$ denotes the *conformal expansion rate*, Φ is the cosmological gravitational potential and $\nabla_{\boldsymbol{x}} \equiv \partial/\partial\boldsymbol{x}$. The potential Φ satisfies the Poisson equation,

$$\nabla^2_{\boldsymbol{x}}\Phi(\boldsymbol{x},\tau) = \frac{3}{2}\Omega_m(\tau)\,\mathcal{H}^2(\tau)\,\delta(\boldsymbol{x},\tau). \tag{1.3.4}$$

By taking the divergence of Eq. 1.3.3 and using Eqs. 1.3.4 and 1.3.2, we get

$$J\left(\delta_{ij} + \Psi_{i,j}\right)^{-1}\frac{\partial}{\partial q_j}\left[\frac{\mathrm{d}^2\Psi_i}{\mathrm{d}\tau^2} + \mathcal{H}(\tau)\frac{\mathrm{d}\Psi_i}{\mathrm{d}\tau}\right] = \frac{3}{2}\Omega_m\,\mathcal{H}^2(J-1) \tag{1.3.5}$$

This nonlinear equation must be solved perturbatively around its linear solution.

1.3.2 Linear solutions and Zel'dovich approximation

The linear solution of Eq. 1.3.5 is

$$\nabla_{\boldsymbol{q}} \cdot \boldsymbol{\Psi}^{(1)} = -D_1(\tau)\,\delta(\boldsymbol{q}), \tag{1.3.6}$$

where $\delta(\boldsymbol{q})$ is the initial density field and $D_1(\tau)$ is the *linear growth factor*, obeying

$$\frac{\mathrm{d}^2 D_1(\tau)}{\mathrm{d}\tau^2} + \mathcal{H}(\tau)\frac{\mathrm{d}D_1(\tau)}{\mathrm{d}\tau} = \frac{3}{2}\Omega_m(\tau)\mathcal{H}^2(\tau)D_1(\tau). \tag{1.3.7}$$

This solution is *local*, i.e., it does not depend on the behavior of the rest of the fluid elements. If the curl of $\boldsymbol{\Psi}^{(1)}$ vanishes, then the evolution of the displacement field to linear order is completely determined by this equation. Moreover, linear solutions have the property of being exact for one-dimensional motions.

It is convenient for the subsequent discussion to introduce the logarithmic derivative of the linear growth factor with respect to the scale factor, called *linear growth rate* and usually denoted with f:

$$f \equiv \frac{\mathrm{d}\ln D_1}{\mathrm{d}\ln a} \approx \Omega_m^{5/9}, \tag{1.3.8}$$

where the approximation holds for cosmological models with $k=0$, $\Omega_\Lambda \neq 0$ and $0.01 \leq \Omega_m \leq 1$.

In the *Zel'dovich approximation* (ZA; Zel'dovich 1970), the linear displacement field is used as an approximate solution for the dynamical equations. From Eqs. 1.3.6, 1.3.2, the density field reads

$$1+\delta(\boldsymbol{x},\tau) = \frac{1}{[1-\lambda_1 D_1(\tau)][1-\lambda_2 D_1(\tau)][1-\lambda_3 D_1(\tau)]}, \tag{1.3.9}$$

where λ_1, λ_2 and λ_3 are the (local) eigenvalues of $\Psi_{i,j}$ at $\tau = 0$. The ZA leads to *planar collapse* (one positive eigenvalue larger than the rest), *filamentary collapse* (two positive eigenvalues larger than the third), or *spherical collapse* (all eigenvalues positive and equal). If all eigenvalues are negative, then the evolution corresponds to an underdense region.

The ZA breaks down completely when *shell-crossing* occurs, i.e., fluid elements with different initial positions $\boldsymbol{q}$ end up at the same position $\boldsymbol{x}$. A region where shell-crossing occurs is called a *caustic*. In this case, particles continue to move through the caustic in the same direction as they did before. In reality, matter entering a caustic would feel the strong gravitational pull and would be attracted back before it could escape.

1.3.3 Second-order Lagrangian Perturbation Theory

Second-order Lagrangian Perturbation Theory (2LPT) offers a significant improvement over ZA, as it accounts for the non-local nature of gravitational instability, which also depends on tidal effects. The divergence of the second-order correction to the displacement field reads:

$$\nabla_{\boldsymbol{q}} \cdot \boldsymbol{\Psi}^{(2)} = \frac{1}{2} D_2(\tau) \sum_{i \neq j} \left(\Psi^{(1)}_{i,i} \Psi^{(1)}_{j,j} - \Psi^{(1)}_{i,j} \Psi^{(1)}_{j,i} \right) , \tag{1.3.10}$$

where $D_2(\tau)$ denotes the *second-order growth factor*. For cosmological models with $k = 0$ and $\Omega_\Lambda \neq 0$, a good approximation for $D_2(\tau)$ is given by (Bouchet et al. 1994):

$$D_2(\tau) \approx -\frac{3}{7} D_1^2(\tau)\, \Omega_m^{-1/143}, \qquad 0.01 \leq \Omega_m \leq 1. \tag{1.3.11}$$

The perturbative approach can be continued by moving to third-order LPT (and higher orders), but it has been shown that the improvement over 2LPT given by third-order theory is marginal in most practical cases (Buchert et al. 1994). When density fluctuations become strongly nonlinear, the perturbative approach breaks down, and numerical simulations are necessary to study the evolution of the density field.

1.4 From real space to redshift space

We have seen that in real space, where positions are measured in terms of comoving coordinates $\boldsymbol{x}$, the usual assumptions of homogeneity and isotropy leave only the dependence on the pair separation in the 2PCF and on the scale and shape of triangles in the 3PCF. For instance, in Slepian and Eisenstein (2015a) the final functional dependence of the 3PCF on three variables (two scales and an angle) is obtained after a rotation and translation averaging procedure: the presumed losslessness of this average is based precisely on the two assumptions of homogeneity and isotropy.

Comoving coordinates $\boldsymbol{x}$ of galaxies are not directly measurable: what we can instead directly measure are the angular coordinates of a galaxy on the plane of the sky (e.g., right ascension and declination) and its redshift z, up to a limiting magnitude set by the characteristics of the instrument and the duration of the observations. Neglecting to exclude the contribution of peculiar velocities from measured redshifts, when converting these latter into distances, gives rise to the so-called *redshift-space distortions.* The impact of these distortions on clustering functions makes them a major probe for extracting information from the LSS: indeed, the peculiar velocity of a given cosmic structure arises from its interaction with the surrounding cosmic perturbations, whose evolution depends on cosmological parameters. The following section is dedicated to the discussion of this topic.

1.4.1 Redshift-space distortions

The redshift of an object, such as a galaxy, is generically measured by comparing the observed wavelength $\lambda_{\rm obs}$ of distinctive spectral features with their rest-frame wavelength $\lambda_{\rm rest}$ (known from laboratory experiments):

$$z = \frac{\lambda_{\rm obs} - \lambda_{\rm rest}}{\lambda_{\rm rest}}. \tag{1.4.1}$$

The wavelength shift is determined by two different contributions:

- The Hubble flow, which makes the rest-frame wavelength of an electromagnetic wave emitted from a source at cosmological redshift z_c arrive at the observer at the present time increased by a factor $1 + z_c$. The Hubble flow is responsible for the recessional velocity of the source $\boldsymbol{v_{rec}}$, described by the Hubble-Lemaître law (Sec. 1.1).

- The Doppler shift associated with the l.o.s. component $\delta v_{||}$ of the peculiar velocity $\delta \boldsymbol{v}$ of the galaxy, namely $\delta v_{||} \equiv \hat{n} \cdot \delta \boldsymbol{v}$, where $\hat{n}$ is the l.o.s. unit vector.

Since in general $\delta v_{||}$ is different than zero, the redshift of an object measured by applying Eq. 1.4.1 does not correspond to z_c. The relation between the two is

$$z = z_c + \frac{\delta v_{||}}{c}(1 + z_c). \tag{1.4.2}$$

Peculiar velocities originate from gravitational interactions with the surrounding fluctuations of the density field. In general, peculiar velocities of galaxies in a survey are unknown, but for increasing distances, they are less dominant than the Hubble flow. Therefore, the usual approximation in this context is to relate the projection along the l.o.s. of the total velocity of the galaxy entirely to the Hubble flow. Although peculiar velocities are small compared to the recessional velocities for sufficiently far objects, they are coherent on cosmological scales, so this approximation introduces systematic distortions on large scales, known as *redshift-space distortions* (RSD).

RSD manifest in two particular ways: galaxies bound in virialized (and so nonlinear) structures, such as galaxy clusters, show an elongated distribution in redshift space, with the elongation axis pointing towards the observer. This distortion is the so-called *fingers of God* (FoG) effect. Another effect known as *Kaiser effect* originates from coherent motions of galaxies towards the center of a gravitational potential well: these motions

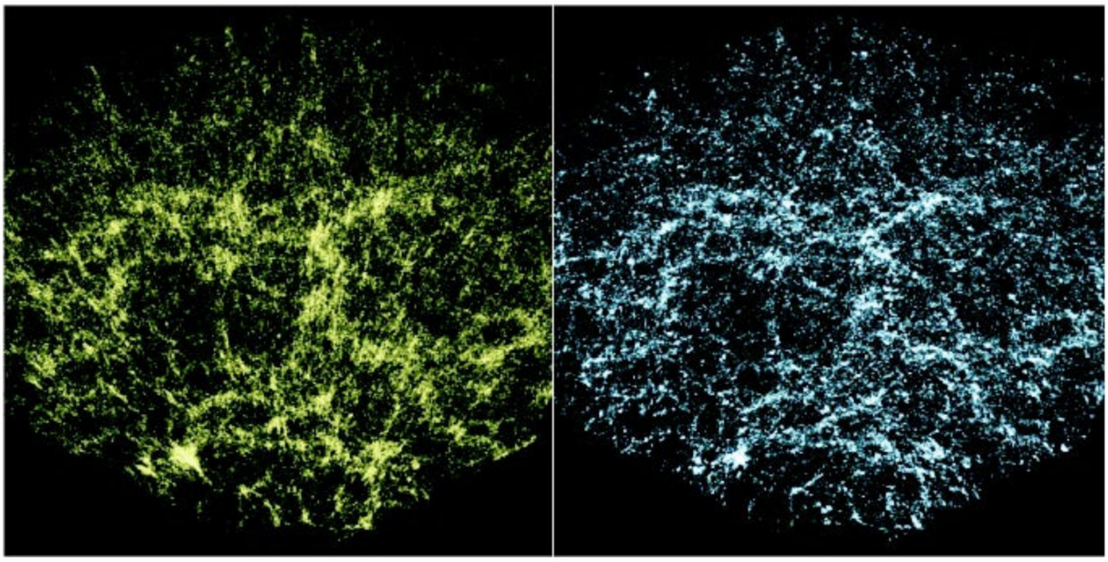

Figure 1.2: Comparison between the same survey light cone in redshift space (*left*) and real space (*right*), for the galaxy distribution in the *Sloan Digital Sky Survey* (SDSS). The left image is clearly a radially distorted version of the right one. Credits: SubbaRao et al. (2008).

lead to a flattening along the l.o.s. of the distribution in redshift space (*pancakes of God effect*), and occur on larger scales (almost linear) as compared to the FoG effect. In Fig. 1.2, we show an illustration taken from literature of the same galaxy survey light cone in real and redshift space. The figure greatly helps to visualize these effects, showing how the redshift-space galaxy distribution is a radially distorted version (with respect to the observer located in the lower corner) of the real-space one.

We now want to obtain the relation between the Fourier modes of the overdensity field $\tilde{\delta}_s(\boldsymbol{k})$ in redshift space and the analogous ones $\tilde{\delta}(\boldsymbol{k})$ in real space. The transformation between the comoving coordinates in real space $\boldsymbol{x}$ and in redshift space $\boldsymbol{s}$ of an object with velocity $\boldsymbol{v}(\boldsymbol{x})$ is

$$\boldsymbol{s} = \boldsymbol{x} - f u_n(\boldsymbol{x})\,\hat{n}, \tag{1.4.3}$$

where f is the linear growth rate and $u_n(\boldsymbol{x}) \equiv -\frac{\boldsymbol{v}(\boldsymbol{x})}{faH}$.

Let us work in the plane-parallel approximation, in which we assume that the l.o.s. $\hat{n}$ is global for the entire survey. The number of objects inside an infinitesimal volume must be conserved by the map 1.4.3, so

$$(1+\delta_s)d^3\boldsymbol{s} = (1+\delta)d^3\boldsymbol{x}. \tag{1.4.4}$$

Since the Jacobian of the map is

$$J(\boldsymbol{x}) \equiv \frac{d\boldsymbol{s}}{d\boldsymbol{x}} = 1 - f\frac{\partial u_n}{\partial \hat{n}}(\boldsymbol{x}), \tag{1.4.5}$$

Eq. 1.4.4 becomes

$$\delta_s = \frac{\delta(\boldsymbol{x}) + f\frac{\partial u_n}{\partial \hat{n}}(\boldsymbol{x})}{1 - f\frac{\partial u_n}{\partial \hat{n}}(\boldsymbol{x})}. \tag{1.4.6}$$

Eventually, the Fourier transform of this expression is

$$\tilde{\delta}_s(\boldsymbol{k}) = \int d^3\boldsymbol{x}\, e^{-i\boldsymbol{k}\cdot\boldsymbol{x}} e^{ifk_n u_n(\boldsymbol{x})} \left[\delta(\boldsymbol{x}) + f\frac{\partial u_n}{\partial \hat{n}}(\boldsymbol{x})\right]. \tag{1.4.7}$$

The term within square brackets is associated with the Kaiser effect, whereas the exponential factor is related to the FoG effect.

1.4.2 Clustering statistics in redshift space

We briefly describe, for completeness, how Eq. 1.4.7 can be treated analytically, through a perturbative approach introduced in Scoccimarro et al. (1999). This is useful because it allows us to express the already described statistics (power spectrum/2PCF, bispectrum/3PCF) in redshift space in terms of their analogues in real space.

First, the FoG exponential is expressed in terms of its Taylor expansion, and the Fourier transform of $\tilde{\delta}_s(\boldsymbol{k})$ is explicitly computed, yielding:

$$\begin{aligned}\tilde{\delta}_s(\boldsymbol{k}) = \sum_{n=1}^{\infty} \int d^3k_1 \ldots \int \frac{d^3k_n}{(2\pi)^{3n-3}}\, \delta^{(D)}(\boldsymbol{k} - \boldsymbol{k_1} - \boldsymbol{\ldots} - \boldsymbol{k_n})\times \\ \times \left[\tilde{\delta}(k_1) + f\mu_1^2\tilde{\theta}(k_1)\right] \frac{(f\mu k)^{n-1}}{(n-1)!} \frac{\mu_2}{k_2}\tilde{\theta}(k_2)\ldots\frac{\mu_n}{k_n}\tilde{\theta}(k_n),\end{aligned} \tag{1.4.8}$$

where $\mu_i \equiv \boldsymbol{k_i}\cdot\hat{\boldsymbol{n}}/k_i$ and $\tilde{\theta}$ is the Fourier transform of the divergence of the peculiar velocity field $\theta \equiv \nabla\cdot\boldsymbol{u}$. We then write the matter density and velocity divergence fields $\tilde{\delta}_m$ and $\tilde{\theta}_m$ in terms of the following perturbative expansions:

$$\tilde{\delta}_m(\boldsymbol{k}) = \sum_{i=1}^{\infty} D_1^i(\tau)\delta_i(\boldsymbol{k}), \tag{1.4.9}$$

$$\tilde{\theta}_m(\boldsymbol{k}) = -\mathcal{H}(\tau)f\sum_{i=1}^{\infty} D_1^i(\tau)\theta_i(\boldsymbol{k}), \tag{1.4.10}$$

which are the equations at the basis of the Standard Perturbation Theory (SPT; Bernardeau et al. 2002). The coefficients δ_i and θ_i can be obtained by solving perturbatively the continuity and Euler equations for the cosmological fluid, and expressed in terms of the linear matter overdensity field δ_1 as:

$$\begin{aligned}\delta_i(\boldsymbol{k}) &= \int d^3q_1\ldots\int\frac{d^3q_i}{(2\pi)^{3n-3}}\,\delta^{(D)}(\boldsymbol{k}-\boldsymbol{q_1}-\ldots-\boldsymbol{q_i})\,F_i(\boldsymbol{q_1},\ldots,\boldsymbol{q_i})\,\delta_1(\boldsymbol{q_1})\ldots\delta_1(\boldsymbol{q_i})\\ \theta_i(\boldsymbol{k}) &= \int d^3q_1\ldots\int\frac{d^3q_i}{(2\pi)^{3n-3}}\,\delta^{(D)}(\boldsymbol{k}-\boldsymbol{q_1}-\ldots-\boldsymbol{q_i})\,G_i(\boldsymbol{q_1},\ldots,\boldsymbol{q_i})\,\delta_1(\boldsymbol{q_1})\ldots\delta_1(\boldsymbol{q_i})\end{aligned} \tag{1.4.11}$$

where F_i and G_i are convolutive kernels (for the full derivation we refer the reader to Goroff et al. 1986). By plugging the expansions 1.4.9 and 1.4.10 into Eq. 1.4.8, after adopting the bias relation in Eq. 1.2.4 and no bias for the velocity field of the tracers ($\theta = \theta_m$), and using the expressions for the coefficients in Eq. 1.4.11, $\tilde{\delta}_s(\boldsymbol{k})$ can be written in the form

$$\begin{aligned}\tilde{\delta}_s(\boldsymbol{k}) = \sum_{n=1}^{\infty} \int d^3k_1 \dots \int \frac{d^3k_n}{(2\pi)^{3n-3}}\, \delta^{(D)}(\boldsymbol{k} - \boldsymbol{k_1} - \dots - \boldsymbol{k_n}) \times \\ \times Z_n(\boldsymbol{k_1}, \dots, \boldsymbol{k_n})\, \delta_1(\boldsymbol{k_1}) \dots \delta_1(\boldsymbol{k_n}),\end{aligned} \tag{1.4.12}$$

where Z_n are the corresponding convolutive kernels in redshift space, which depend on the real-space kernels, on f, and on the bias coefficients in Eq. 1.2.4.

This formalism can be applied to obtain the expression for the redshift-space power spectrum/2PCF or bispectrum/3PCF. Let us first focus on 2-point statistics, in particular on the *linear* power spectrum, computed by applying Eq. 1.2.11 at the lowest possible order, i.e., by keeping in the δ terms of the correlator only the $n = 1$ term $\tilde{\delta}_{s,1}(\boldsymbol{k})$ of the sum in Eq. 1.4.12. We get

$$\begin{aligned}\left\langle \tilde{\delta}_{s,1}(\boldsymbol{k})\tilde{\delta}_{s,1}(\boldsymbol{k'}) \right\rangle &= Z_1(k)Z_1(k') \left\langle \delta_1(\boldsymbol{k})\delta_1(\boldsymbol{k'}) \right\rangle \\ &= Z_1(k)Z_1(k')(2\pi)^3\delta^{(D)}(\boldsymbol{k} + \boldsymbol{k'})P_{m,\mathrm{lin}}(k)\end{aligned} \tag{1.4.13}$$

where $P_{m,\mathrm{lin}}(k)$ is the linear matter power spectrum. Since it can be shown that, under the previous hypotheses, $Z_1(k) = (b + f\mu^2)$, it is clear from Eq. 1.4.13 that the linear power spectrum $P_{s,\mathrm{lin}}$ of biased tracers in redshift space is given by

$$P_{s,\mathrm{lin}}(k, \mu) = (b + f\mu^2)^2 P_{m,\mathrm{lin}}(k) = (1 + \beta\mu^2)^2 P_{\mathrm{lin}}(k), \tag{1.4.14}$$

where $P_{\mathrm{lin}}(k)$ denotes the linear power spectrum of the tracers in real space and $\beta \equiv f/b$. This is exactly the famous result obtained in Kaiser (1987).

Fourier transforming Eq. 1.4.14 gives the redshift-space 2PCF $\xi(s, \mu)$, which can be written as

$$\xi(s, \mu) = \xi_0(s)P_0(\mu) + \xi_2(s)P_2(\mu) + \xi_4(s)P_4(\mu), \tag{1.4.15}$$

where the multipoles $\xi_\ell(s)$, with $\ell = 0, 2, 4$ (respectively, the monopole, quadrupole, and hexadecapole of the 2PCF) can be written in terms of the real-space 2PCF $\xi(r)$ and β as follows (Hamilton 1992):

$$\xi_0(s) = \left(1 + \frac{2\beta}{3} + \frac{\beta^2}{5}\right) \xi(r), \tag{1.4.16}$$

$$\xi_2(s) = \left(\frac{4\beta}{3} + \frac{4\beta^2}{7}\right) [\xi(r) - \bar{\xi}(r)], \tag{1.4.17}$$

$$\xi_4(s) = \frac{8\beta^2}{35} \left(\xi(r) + \frac{5}{2}\bar{\xi}(r) - \frac{7}{2}\bar{\bar{\xi}}(r)\right), \tag{1.4.18}$$

where the *barred* correlation functions are defined as:

$$\bar{\xi}(r) \equiv \frac{3}{r^3} \int_0^r dr' \xi(r') r'^2, \tag{1.4.19}$$

$$\bar{\bar{\xi}}(r) \equiv \frac{5}{r^5} \int_0^r dr' \xi(r') r'^4. \tag{1.4.20}$$

Eq. 1.4.16 clearly shows for the particular case of the 2PCF that RSD have an impact also on the isotropic component (i.e., the monopole) of clustering functions in redshift space.

Similarly, an analogous computation can be done for the bispectrum. To obtain a result to the lowest order (the so-called *tree level*), it is necessary to expand the density contrast field up to second order in SPT. We do not show the computations, which include the application of Wick's theorem, but we report the result, which is

$$B_{s,\mathrm{TL}}(\boldsymbol{k}_1, \boldsymbol{k}_2, \boldsymbol{k}_3) = 2\, Z_2(\boldsymbol{k}_1, \boldsymbol{k}_2) Z_1(\boldsymbol{k}_1) Z_1(\boldsymbol{k}_2)\, P_{m,\mathrm{lin}}(k_1) P_{m,\mathrm{lin}}(k_2) + \mathrm{cycl.\,perm.} \tag{1.4.21}$$

Notice that if one adopts a non-local bias relation as the one in Eq. 1.2.6, the tidal bias term enters the computation of the bispectrum already at the tree-level, through the convolutive kernel Z_2 which depends on b_{s^2}: so, in the context of 3-point statistics, tidal effects are important on all scales. We refer the reader to Baldauf et al. (2012) for the full computation and the expression of Z_2 including the tidal bias term. The corresponding expression for the 3PCF can be obtained by taking the Fourier transform of Eq. 1.4.21. In general, higher-order corrections can be incorporated by progressively including more terms of the SPT expansion of the density contrast field in the correlators defining a given statistics.

A fundamental effect of RSD, which is evident by looking at the expressions of clustering functions in redshift space, is that they break the statistical isotropy of the density field by introducing a privileged direction, i.e., the l.o.s. direction. In particular, clustering statistics become dependent on the orientation of pairs (in the case of 2-point statistics) or triplets (for 3-point statistics) with respect to the l.o.s. vector. In redshift space, this dependence makes the 2PCF/power spectrum depend on two variables: the norm of the vector $\boldsymbol{s}$ separating a given pair and the cosine of the angle between $\boldsymbol{s}$ and the l.o.s. $\hat{\boldsymbol{n}}$, $\mu \equiv \boldsymbol{s} \cdot \hat{\boldsymbol{n}} / |\boldsymbol{s}|$. Two additional variables (e.g., two angles) are instead needed to describe the relative orientation of a triangle with respect to a fixed direction. These two variables add to the three of the isotropic case, which constrain the size and shape of a triangle, so the redshift space bispectrum/3PCF depends on five variables. Different bases and multipole decompositions have been proposed to deal with anisotropic 3-point statistics (Scoccimarro et al. 1999; Slepian and Eisenstein 2018; Sugiyama et al. 2019): in particular, a very convenient formalism is the one presented in Sugiyama et al. (2019), based on the tensor product of three spherical harmonics and called TripoSH (Tripolar Spherical Harmonics; Varshalovich et al. 1988). This formalism makes it possible to naturally translate the symmetry properties of the Universe into constraints on the allowed multipoles, to easily isolate the anisotropic contribution of RSD, and to recover the isotropic multipoles described in Sec. 1.2.8. We do not describe these formalisms since this work focuses on the study of the isotropic 3PCF.

RSD can be exploited in large deep redshift surveys to constrain the growth rate f. This parameter contains valuable information, since it is a direct probe of DE and allows cosmologists to test GR against modified gravity theories (Dodelson and Schmidt 2020).

1.5 Clustering in massive neutrino cosmologies

We now discuss the differences between massive neutrino cosmologies and cosmological models in which neutrinos are assumed to be massless, as in the case of the standard ΛCDM model. We mainly refer to the well-known review by Lesgourgues and Pastor (2006).

The existence of a relic background of neutrinos is a prediction of the standard Hot Big Bang model. Neutrinos were produced in the very early phases of the history of the Universe by frequent weak interactions. They decoupled from the cosmological fluid before the onset of Big Bang Nucleosynthesis (BBN), when they were still relativistic particles: being relativistic fermions at decoupling, their equilibrium phase-space distribution in a homogeneous Universe, independently of their subsequent evolution, is described by a Fermi-Dirac distribution:

$$f_0(q) = \frac{1}{1+\exp(q/aT_\nu)}, \tag{1.5.1}$$

where q is the comoving momentum (if p is the physical momentum, then $q = ap$) and T_ν is the neutrino temperature, evolving with respect to the scale factor as $T_\nu \propto a^{-1}$. Before becoming non-relativistic, neutrinos contributed to the total energy density of radiation along with photons, and the growth of their density fluctuations was suppressed.

A given species of neutrino with mass m_ν became non-relativistic when the mean thermal energy $\sim k_\mathrm{B}T_\nu$ decreased below its rest mass energy $m_\nu c^2$, an event occurred at a cosmological redshift z_{nr} approximately given by

$$1 + z_{nr} \approx 1890 \left(\frac{m_\nu}{1\,\mathrm{eV}}\right). \tag{1.5.2}$$

For $m_\nu = 50\,\mathrm{meV}$, $z_{nr} \sim 100$. The transition from relativistic to non-relativistic neutrinos, thus, happens in the time span between recombination and LSS formation.

From this redshift onward, neutrinos contribute to the total matter density of the Universe together with CDM and baryons, in the form of *hot dark matter*[7] (the adjective *hot* comes from their being relativistic at decoupling). So, the total matter density parameter must be computed in terms of the density parameters of CDM, baryons, and neutrinos Ω_c, Ω_b, and Ω_ν as $\Omega_m = \Omega_c + \Omega_b + \Omega_\nu$. In particular, the present-day value of Ω_ν is given by (Particle Data Group: Navas et al. 2024)

$$\Omega_{\nu,0} = \frac{\rho_{\nu,0}}{\rho_{c,0}} \equiv \frac{\sum_\nu m_\nu}{93.12\,h^2\,\mathrm{eV}}, \tag{1.5.3}$$

---

[7] Notice that neutrinos are so far the only dark matter candidates that we actually know to exist.

where the total present-day neutrino density $\rho_{\nu,0}$ is computed from the neutrino-to-photon number density ratios per species n_i/n_γ (obtained from decoupling studies) and knowing n_γ from CMB temperature measurements.

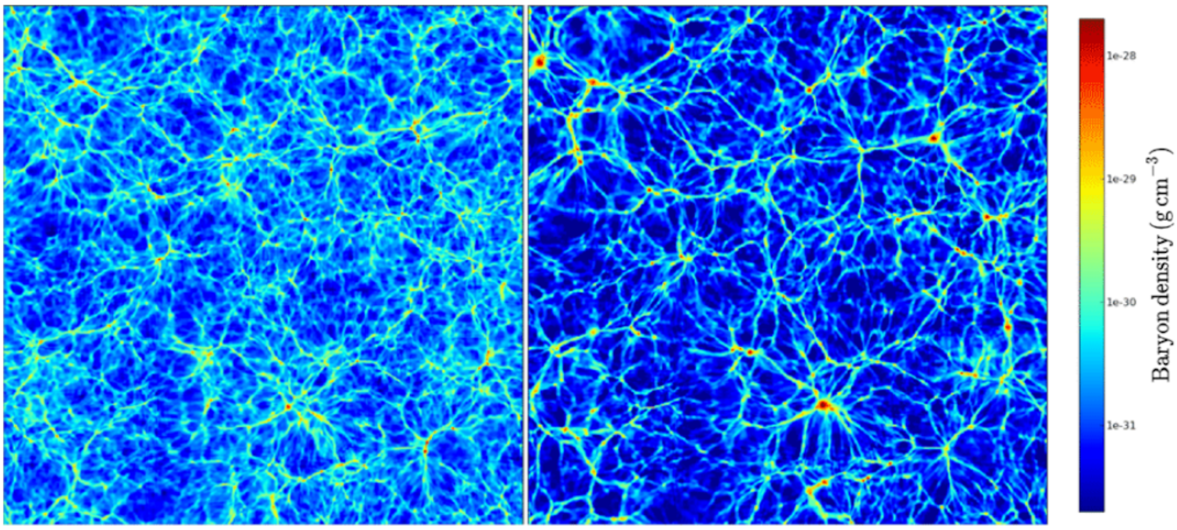

Figure 1.3: Density distribution at $z = 0$ of the baryonic matter in two N-body simulated Universes, started from the same initial conditions, without massive neutrinos (*right*) and with massive neutrinos with $\sum m_\nu = 1.9\,\mathrm{eV}$ (*left*). The slices are $200\,h^{-1}$ Mpc wide and show the baryonic mass averaged over a depth of $\sim 391\,h^{-1}$ kpc. The unrealistic value of $\sum m_\nu = 1.9\,\mathrm{eV}$ was chosen to enhance the differences between the two snapshots. Credits: Agarwal and Feldman (2011).

Long after the non-relativistic transition, a collisionless fluid approximation can be adopted for neutrinos (Shoji and Komatsu 2010). Real fluids react to perturbations, potentially leading to gravitational instability and collapse, through internal pressure, with a speed which can be assumed to be the sound speed. Roughly speaking, a region of fluid is stable if it extends on a scale such that the time needed for a sound wave to cross it is lower than the time this region would take to collapse freely. In a pressureless fluid, as in the case of the effective fluid description for neutrinos, the role of the sound speed is assumed by the (large) thermal velocity of the particles. This implies that below a given threshold scale, perturbations in the neutrino fluid are damped. This scale is known as *free-streaming* scale and its expression as a function of neutrino mass and redshift is given by

$$\lambda_{\mathrm{FS}}(m_\nu, z) = a\left(\frac{2\pi}{k_{\mathrm{fs}}}\right) \approx 7.7\,\frac{H_0(1+z)}{H(z)}\left(\frac{1\,\mathrm{eV}}{m_\nu}\right) h^{-1}\mathrm{Mpc}. \tag{1.5.4}$$

where k_{fs} is the associated free-streaming wavenumber. On scales larger than the free-streaming scale, neutrinos behave like CDM. The minimum value of the free-streaming wavenumber, corresponding to the largest scale that can be affected by neutrinos, is

$$k_{\mathrm{nr}} \approx 0.018\,\Omega_m^{1/2}\left(\frac{m_\nu}{1\,\mathrm{eV}}\right) h\,\mathrm{Mpc}^{-1}. \tag{1.5.5}$$

The existence of this scale, across which neutrinos exhibit different clustering properties, impacts several observables and quantities, including clustering statistics, RSD, bias, and the halo mass function. Qualitatively, the impact of massive neutrinos on the LSS can be appreciated in Fig. 1.3, which shows the baryonic matter distribution in two snapshots from N-body simulations: one with $\sum m_\nu = 0\,\mathrm{eV}$ and another with an unrealistically large $\sum m_\nu = 1.9\,\mathrm{eV}$, chosen to make the comparison clear. Even if the two simulations share the same initial conditions, differences can be noticed in the distribution of structures, especially for filaments and voids.

In what follows, we summarize the findings of previous studies in configuration and Fourier space about the effects of massive neutrinos on the quantities listed above. These studies include results obtained both analytically and through cosmological simulations, the latter allowing the exploration of the fully nonlinear regime.

1.5.1 Linear regime

We start by describing the effects of neutrinos in the linear regime. It is convenient to write a generic matter perturbation δ_m as a weighted sum of neutrino and CDM (+baryon) fluctuations, respectively denoted as δ_ν and δ_c[8]:

$$\delta_m = f_\nu \delta_\nu + (1 - f_\nu)\delta_c\,, \qquad f_\nu \equiv \Omega_\nu / \Omega_m. \tag{1.5.6}$$

By indicating with P_{mm}, P_{cc} and $P_{\nu\nu}$, respectively, the power spectra of the total matter, CDM, and neutrino fields, and with $P_{c\nu}$ the cross-power spectrum of the CDM and neutrino fields, Eq. 1.5.6 together with the definition of the power spectrum in Eq. 1.2.11 yields:

$$P_{mm} = f_\nu^2 P_{\nu\nu} + (1 - f_\nu)^2 P_{cc} + 2 f_\nu (1 - f_\nu) P_{c\nu}. \tag{1.5.7}$$

Due to free-streaming, for $k \gg k_{\rm nr}$ neutrino perturbations are suppressed, and consequently the power spectra $P_{\nu\nu}, P_{c\nu} \to 0$. For $k < k_{\rm nr}$, for what we have already discussed $P_{\nu\nu}, P_{c\nu} = P_{cc}$, so we have

$$P_{mm}(k) = \begin{cases} P_{cc}(k) & \text{if } k < k_{\rm nr} \\ (1 - f_\nu)^2 P_{cc}(k) & \text{if } k \gg k_{\rm nr}, \end{cases} \tag{1.5.8}$$

implying $P_{cc} \geq P_{mm}$ on all scales.

It is also possible to evaluate the effect of massive neutrinos on the matter and CDM power spectra determined for wavenumbers $k \gg k_{\rm nr}$, with respect to the massless case, under the hypothesis that the total matter density in Eq. 1.5.6 is fixed. The result is a suppression of power, which for small f_ν is

$$\frac{P_{mm}(k; f_\nu)}{P_{mm}(k; f_\nu = 0)} \approx 1 - 8 f_\nu\,, \qquad \frac{P_{cc}(k; f_\nu)}{P_{cc}(k; f_\nu = 0)} \approx 1 - 6 f_\nu. \tag{1.5.9}$$

---

[8] After decoupling from radiation, perturbations δ_b in the baryon fluid increase, finally reaching the amplitude of CDM fluctuations, so we can assume $\delta_b \approx \delta_c$.

Another important consequence determined by the different behavior of neutrino perturbations as a function of scale is that they induce a different growth of CDM perturbations below the free-streaming scale. In particular, it is possible to show that during the matter domination era, for $k \gg k_{\rm nr}$ and small f_ν,

$$\delta_c \propto a^{1-\frac{3}{5}f_\nu} \qquad \longrightarrow \qquad \frac{f_N}{f_{\Lambda\rm CDM}} \approx 1 - \frac{3}{5}f_\nu . \tag{1.5.10}$$

where f_N and $f_{\Lambda\rm CDM}$ are, respectively, the CDM linear growth rates in the massive and massless neutrino cases. Since above the free-streaming scale neutrinos behave as CDM, in general (i.e, including also the non-asymptotic behavior), the linear growth rate f_N is scale-dependent (as well as time dependent) to match the two asymptotic behaviors, i.e., $f_N = f_N(k,a)$. During Λ domination, the combined effect of Λ and neutrinos on the growth of δ_c can be well approximated by $\delta_c \propto [a\, g(a)]^{1-3/5 f_\nu}$, where $g(a)$ is a scale-independent damping factor regulating the decay of the gravitational potential (see Lesgourgues and Pastor 2006, Eq. 134).

1.5.2 Nonlinear regime

As discussed in the previous section, the shape of the linear power spectrum is sensitive to the value of neutrino masses. This dependence becomes stronger in the mildly and fully nonlinear regimes. The accuracy of analytical predictions in the nonlinear regime is usually tested on simulations. In particular, by measuring the different components of the matter power spectrum for a set of N-body simulations, Castorina et al. (2015) show that both in the linear ($k \lesssim 0.1\, h^{-1}$Mpc) and nonlinear regime most of the contribution to P_{mm} comes from P_{cc}, which is the only component significantly deviating from the linear prediction. Their conclusion shows that in massive neutrino cosmologies, the nonlinear P_{mm} can be described at the 1% level by accounting for the nonlinear evolution of CDM perturbations alone, while adopting the linear prediction for the other components.

Nonlinearity, as already pointed out, leads to the emergence of non-Gaussianity, quantified in the first place by a non-vanishing bispectrum. By using Eq. 1.5.6 applied to the definition of bispectrum Eq. 1.2.22, the analogue of Eq. 1.5.7 for the matter bispectrum is (Ruggeri et al. 2018)

$$\begin{aligned} B_{mmm}(k_1,k_2,k_3) = (1-f_\nu)^3 B_{ccc}(k_1,k_2,k_3) + (1-f_\nu)^2 f_\nu\, B^{(s)}_{cc\nu}(k_1,k_2,k_3) \\ + (1-f_\nu) f_\nu^2\, B^{(s)}_{\nu\nu c}(k_1,k_2,k_3) + f_\nu^3\, B_{\nu\nu\nu}(k_1,k_2,k_3), \end{aligned} \tag{1.5.11}$$

where the symmetrized versions of the cross CDM-CDM-neutrino bispectrum $B^{(s)}_{cc\nu}$ and neutrino-neutrino-CDM bispectrum $B^{(s)}_{\nu\nu c}$ are defined by the following relations:

$$\begin{aligned} \delta^{(D)}(\boldsymbol{k_1}+\boldsymbol{k_2}+\boldsymbol{k_3})\, B^{(s)}_{cc\nu}(k_1,k_2,k_3) \equiv \left\langle \tilde{\delta}_c(\boldsymbol{k_1})\tilde{\delta}_c(\boldsymbol{k_2})\tilde{\delta}_\nu(\boldsymbol{k_3}) \right\rangle \\ + \left\langle \tilde{\delta}_c(\boldsymbol{k_3})\tilde{\delta}_c(\boldsymbol{k_1})\tilde{\delta}_\nu(\boldsymbol{k_2}) \right\rangle + \left\langle \tilde{\delta}_c(\boldsymbol{k_2})\tilde{\delta}_c(\boldsymbol{k_3})\tilde{\delta}_\nu(\boldsymbol{k_1}) \right\rangle , \end{aligned} \tag{1.5.12}$$

$$\delta^{(D)}(\boldsymbol{k_1}+\boldsymbol{k_2}+\boldsymbol{k_3})\, B^{(s)}_{\nu\nu c}(k_1,k_2,k_3) \equiv \left\langle \tilde{\delta}_\nu(\boldsymbol{k_1})\tilde{\delta}_\nu(\boldsymbol{k_2})\tilde{\delta}_c(\boldsymbol{k_3})\right\rangle + \left\langle \tilde{\delta}_\nu(\boldsymbol{k_3})\tilde{\delta}_\nu(\boldsymbol{k_1})\tilde{\delta}_c(\boldsymbol{k_2})\right\rangle + \left\langle \tilde{\delta}_\nu(\boldsymbol{k_2})\tilde{\delta}_\nu(\boldsymbol{k_3})\tilde{\delta}_c(\boldsymbol{k_1})\right\rangle . \tag{1.5.13}$$

Also in the case of the bispectrum, neutrinos induce a suppression below their free-streaming scale. For the equilateral configurations ($k_1 = k_2 = k_3 \equiv k$), the suppression induced on the leading-order matter bispectrum on scales $k \gg k_{\rm nr}$ is (Levi and Vlah 2016)

$$\frac{B_{mmm}(k,k,k;f_\nu)}{B_{mmm}(k,k,k;f_\nu=0)} \approx 1 - 13.5\, f_\nu . \tag{1.5.14}$$

We discuss more about the neutrino-induced suppression on the bispectrum measured in numerical simulations in Sec. 1.6.

1.5.3 Halo mass function, bias and redshift-space distortions

The free-streaming of non-relativistic neutrinos opposes the gravitational collapse, which constitutes the basis of cosmic structure formation. As an immediate consequence, a significant suppression in the average number density of massive structures is expected in massive neutrino cosmologies, with respect to a ΛCDM cosmology sharing the same amplitude of the primordial matter power spectrum. A way to quantify this effect is through the halo mass function (HMF), defined as the function giving the comoving number density of dark matter halos per mass interval (Dodelson and Schmidt 2020). A fundamental result obtained in Press and Schechter (1974) is that the HMF can be written as

$$\frac{M dM}{\bar{\rho}}\frac{{\rm d}n(M,z)}{{\rm d}M} = \zeta f(\zeta)\frac{d\zeta}{\zeta}, \tag{1.5.15}$$

with $\zeta \equiv [\delta_{\rm sc}(z)/\sigma(M,z)]^2$ (not to be confused with the 3PCF), and where $\delta_{\rm sc}$ is the linearly extrapolated overdensity at the time of collapse (assuming spherical symmetry). This parameter shows a very weak dependence on cosmology (Kitayama and Suto 1996), so one can safely assume its value to be, even in the presence of massive neutrinos, the one for the $f_\nu = 0$ case, that is $\delta_{\rm sc} \approx 1.686$. The quantity $\bar{\rho}$ is the background matter density. Here $dn(M,z)$ is the number density of halos in the mass interval $[M, M+dM]$, and $\sigma^2(M)$ is the variance of the linear matter density field smoothed below scale M (Eq. 1.2.17, with $M = 4\pi\bar{\rho}R^3/3$). Several fits have been proposed for the function $f(\zeta)$ (Sheth and Tormen 2002; Tinker et al. 2008; Crocce et al. 2010) whose expressions we do not report here.

In Marulli et al. (2011), the effect of suppression, progressively larger for larger neutrino masses, is visible in the high-mass tail of the HMF measured from simulations (mainly in the mass range $10^{14}\, h^{-1}M_\odot < M < 10^{15}\, h^{-1}M_\odot$). A crucial point is to accurately reproduce this trend with models. Notice that in massless neutrino cosmologies, the linear power spectrum of CDM and total matter coincide, while in the massive neutrino scenario, it is not clear in principle which power spectrum (CDM+ν or only CDM) is more relevant to describe halo clustering. A very important result assessed in different studies (Costanzi et al. 2013; Castorina et al. 2014; Castorina et al. 2015) is that the modeled HMFs better reproduce the results from numerical simulations if the *cold dark matter prescription* is adopted, i.e., if models are computed by using the CDM linear power spectrum to compute $\sigma(M, z)$, and $\bar{\rho}$ is set to be equal to the background density of the CDM component.

The definition of bias in the presence of massive neutrinos is not unambiguous: one can measure different kinds of bias depending on whether the total matter, rather than the CDM field, is chosen as the underlying field. For example, for dark matter halos, one can define halo auto- and cross-bias coefficients with respect to CDM as

$$b_c^{(hh)} \equiv \sqrt{\frac{P_{hh}(k)}{P_{cc}(k)}}\,, \qquad b_c^{(hc)} \equiv \frac{P_{hc}(k)}{P_{cc}(k)} \tag{1.5.16}$$

and with respect to the total matter field as

$$b_m^{(hh)} \equiv \sqrt{\frac{P_{hh}(k)}{P_{mm}(k)}}\,, \qquad b_m^{(hm)} \equiv \frac{P_{hm}(k)}{P_{mm}(k)}, \tag{1.5.17}$$

where P_{mm}, P_{cc}, P_{hh} are, respectively, the auto matter-matter, CDM-CDM and halo-halo power spectra, while P_{hm}, P_{hc} represent, respectively, the cross halo-matter and halo-CDM power spectra. Analogous quantities can be defined in configuration space by replacing power spectra with 2PCFs.

Measurements performed on simulations show that neutrinos produce a significant scale dependence in the bias coefficients (Castorina et al. 2014). On large scales, the b_c coefficients converge to the same values assumed by the b_m ones, reflecting the fact that sufficiently above the free-streaming scale, neutrinos behave as CDM. Conversely, on scales smaller than the free-streaming scale, the bias coefficients $b_m^{(hh)}, b_c^{(hh)}$ and $b_m^{(hm)}, b_c^{(hc)}$ differ roughly by a rescaling, as predicted by combining Eq. 1.5.7 with the definitions in Eqs. 1.5.16 and 1.5.17 since

$$b_m^{(hh)} \equiv \sqrt{\frac{P_{hh}(k)}{P_{mm}(k)}} = b_c^{(hh)}\sqrt{\frac{P_{cc}(k)}{P_{mm}(k)}} \xrightarrow{k \gg k_{\mathrm{nr}}} b_c^{(hh)}(1 - f_\nu) \tag{1.5.18}$$

$$b_m^{(hm)} \equiv \frac{P_{hm}(k)}{P_{mm}(k)} = b_c^{(hc)}\frac{P_{cm}(k)}{P_{cc}(k)} \xrightarrow{k \gg k_{\mathrm{nr}}} b_c^{(hc)}(1 - f_\nu). \tag{1.5.19}$$

Although the scale dependence in the bias coefficients is always present, it has been found that it is reduced when they are measured with respect to the CDM field. Moreover, in Villaescusa-Navarro et al. (2014) it has been shown that the coefficients $b_m^{(hh)}$ and $b_m^{(hm)}$ measured from simulations show a better agreement with models if, as in the case of the HMF, modeling is carried out by following the CDM prescription.

A significant result is obtained for massive neutrino cosmologies in the scenario of the so-called *universality* of the halo bias. In the context of bias modeling, universality means that the bias coefficients, as a function of mass, cosmology, and redshift, can be written in terms of the *peak height* $\nu \equiv \delta_{\rm sc}/\sigma(M,z)$. An example is represented by the *Tinker fitting formula* (Tinker et al. 2010) for the linear bias which can be written as

$$b_{\rm Tinker}(\nu) = 1 - A\frac{\nu^a}{\nu^a + \delta_{\rm sc}^a} + B\nu^b + C\nu^c, \tag{1.5.20}$$

where B, b and c are constants and A, a and C are functions of Δ, the halo mean overdensity within its virial radius in units of the mean matter density of the Universe. The calibrated values of these coefficients can be found in Tab. 2 of the cited work. In the presence of massive neutrinos, universality is valid if the peak height is computed from the variance of the CDM field (Castorina et al. 2014).

Massive neutrinos also impact RSD through their induced scale dependence in the growth rate f. However, translating the results of Sec. 1.4.1 into massive neutrino cosmologies gives rise to an ambiguity similar to that of bias, even in the simplest, linear regime, where the Kaiser relation (Eq. 1.4.14) holds. Indeed, Eq. 1.4.14 in a massless neutrino cosmology assumes that the peculiar velocity field of halos precisely follows the underlying matter velocity field (i.e., the halo velocities are unbiased). In the massive neutrino case, one can define the bias of the halo peculiar velocity field with respect to the total matter velocity field or the CDM velocity field alone. It is not clear a priori whether halo velocities can be similarly assumed to be unbiased with respect to the total matter field (including neutrinos) or only to the CDM component, or to neither of these two limiting cases. In the first two (extreme) cases, one can write the Kaiser formula for the redshift-space halo power spectrum $P_{hh,s}(k,\mu)$ as a function of the real-space one $P_{hh,r}(k)$ as, respectively,

$$P_{hh,s}(k,\mu) = \left[1+\beta_c(k)\mu^2\right]^2 P_{hh,r}(k), \quad P_{hh,s}(k,\mu) = \left[1+\beta_m(k)\mu^2\right]^2 P_{hh,r}(k), \tag{1.5.21}$$

where $\beta_i = f_i/b_i$, with $i = c$ or m, and $f_i(k,a) \equiv {\rm d}\ln D_{1,i}(k,a)/{\rm d}\ln a$ is the scale-dependent linear growth rate of the field i.

Verdiani et al. (2025) showed, by estimating f_c and f_m in simulations, that on large scales, a proper match between numerical and theoretical results is obtained by assuming no bias with respect to the CDM velocity field.

Another effect on RSD in massive neutrino cosmologies manifests through a modification of the RMS of galaxy peculiar velocities. This quantity is smaller than in a pure ΛCDM cosmology sharing the same power spectrum amplitude, because the power spectrum and linear growth rate suppressions decrease the bulk flow predicted by linear theory, given by (Elgarøy and Lahav 2005)

$$\left\langle v^2(R)\right\rangle = \frac{1}{2\pi^2}H_0^2\int dk\, f^2 P_{m,{\rm lin}}(k) W_G^2(kR), \tag{1.5.22}$$

where $W_G(kR)$ is the window function, e.g., for a Gaussian sphere of radius R, $W(kR) \equiv \exp(-k^2R^2/2)$. Conversely, the suppression of the power spectrum below the free-streaming scale induces an increase in halo bias, as this quantity is larger for higher values of the peak height. These two competing effects result in a redshift-space halo 2PCF only slightly suppressed in a ΛCDM+ν cosmology on small scales, as measured in Marulli et al. (2011) in N-body simulations. Moreover, the authors demonstrate that in the case of massive neutrinos, spatial halo clustering is less enhanced in redshift space than in real space, primarily on large scales due to the lower value of bulk flow velocity, and on small scales due to a decrease in the FoG effect.

1.6 Why go beyond 2-point statistics? Tightening the constraints on $\sum m_\nu$ with 3-point statistics

3-point statistics represent a promising tool for investigating massive neutrino cosmologies for several reasons. First, they are naturally sensitive to non-linear signals characterizing small scales, since, as already discussed, non-linearity induces deviations from the Gaussianity of cosmological fields. These are precisely the scales where massive neutrinos, due to the free-streaming effect, produce the largest alterations in clustering. Secondly, Fig. 1.3 clearly shows that the effects imprinted by massive neutrinos on the growth of perturbations have an impact on the shape of the cosmic web. In this regard, 3-point statistics are the lowest-order statistics that encode not only scale, but also shape information (i.e., in the shape of the triangles).

The study of massive neutrino cosmologies through the bispectrum has received increasing attention in recent years. The first measurement of bispectrum in N-body simulations of massive neutrino cosmologies can be found in Ruggeri et al. (2018). The authors quantify the impact of neutrino perturbations in the different terms of Eq. 1.5.11, finding that B_{ccc} and $B^{(s)}_{cc\nu}$ account for $\sim 90\%$ and $\sim 10\%$, respectively, of the total matter bispectrum for equilateral configurations, with $B_{\nu\nu c}$ and $B_{\nu\nu\nu}$ contributing less than 1%. Moreover, they find a shape-dependent response, with squeezed triangles in Fourier space (with sides (q, k, k), where $q \ll k$) less affected by neutrinos. They show that the behavior of the measured bispectrum is better reproduced if B_{ccc} is modeled up to one-loop corrections (the next-to-leading order with respect to tree-level), keeping the other terms, containing neutrino perturbations, at tree-level.[9] This work also extends the result discussed in Sec. 1.5.3, about the universality of bias in the CDM component in massive neutrino cosmologies, also beyond linear bias, including the local quadratic bias b_2 and the tidal bias b_{s^2} (see Eq. 1.2.6).

A further, more general advantage of including 3-point statistics in clustering analysis is their ability to break degeneracies, due to their distinct dependence on cosmological parameters compared to 2-point statistics. In addition, 3-point statistics are particularly powerful at tightening parameter constraints, as the number of triangular configurations increases much faster with increasing k than the number of modes in the power spectrum.

---

[9] Notice that all their predictions are computed by using the approximation consisting of considering the perturbative kernels in cosmologies with massive neutrinos to have the same form as in ΛCDM cosmology and assuming that all the effects induced by neutrinos are encoded in the linear power spectrum (Saito et al. 2009).

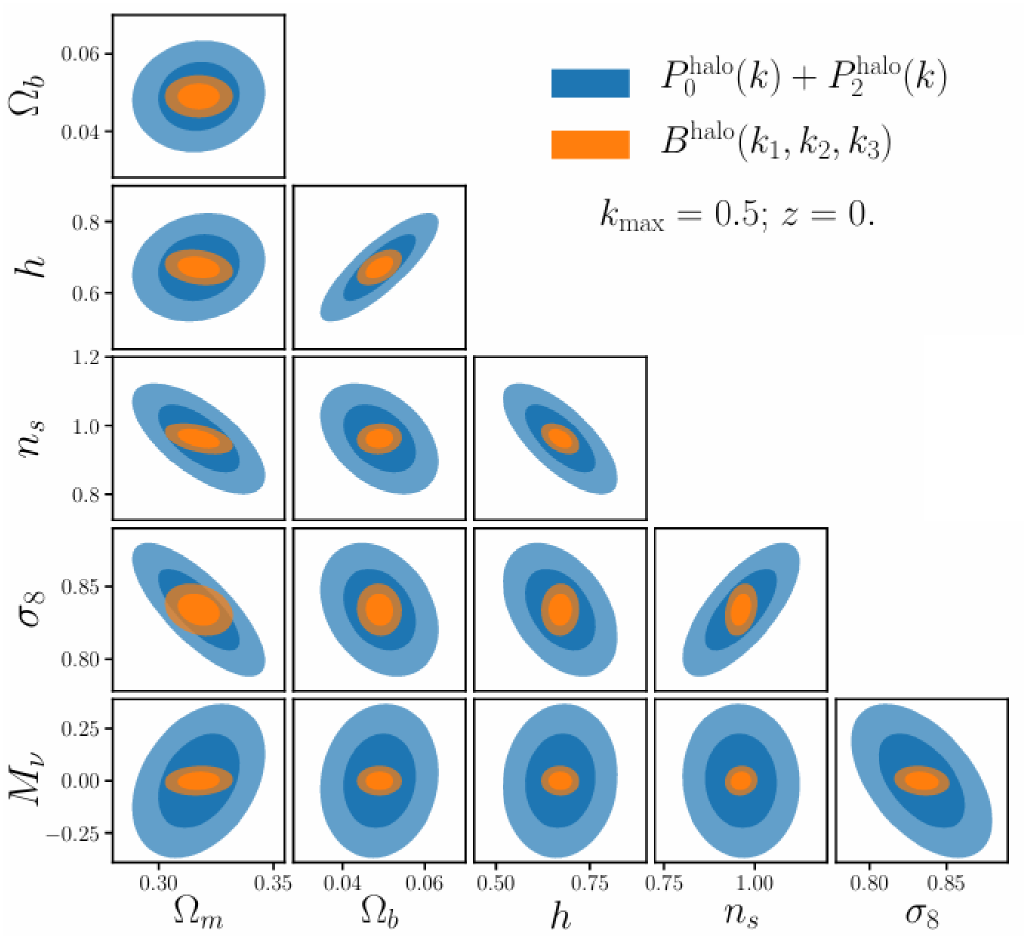

Figure 1.4: Forecasted constraints on $\sum m_\nu$ (here reported as M_ν) and other cosmological parameters, from Hahn et al. (2020). Blue and orange contours represent, respectively, the constraints obtained from the redshift-space halo power spectrum monopole $P_0^{\text{halo}}(k)$ and quadrupole $P_2^{\text{halo}}(k)$, and from the redshift-space halo bispectrum monopole $B_0^{\text{halo}}(k_1, k_2, k_3)$. Dark and light areas mark, respectively, the 68% and 95% confidence intervals. Forecasts are carried out at $z = 0$ for a volume of $1\,h^{-3}\,\text{Gpc}^3$, and exploit modes up to $k_{\text{max}} = 0.5\,h\,\text{Mpc}^{-1}$. Both derivatives of clustering statistics with respect to parameters and covariance have been numerically estimated using the QUIJOTE mock halo catalogues (Villaescusa-Navarro et al. 2020).

This allows information to be extracted from scales where the informative gain from 2-point statistics has already saturated.

As can be seen from 1.4, taken from Hahn et al. (2020) (which shows forecasts for several cosmological parameters, including $\sum m_\nu$), the contours significantly tighten in the case of the bispectrum monopole with respect to the joint analysis of the power spectrum monopole and quadrupole, and also have different orientations for some parameters. Specifically, the 1σ constraint on $\sum m_\nu$ improves by a factor of ≈ 5 when scales up to $k_{\text{max}} = 0.5\,h\,\text{Mpc}^{-1}$ are included.

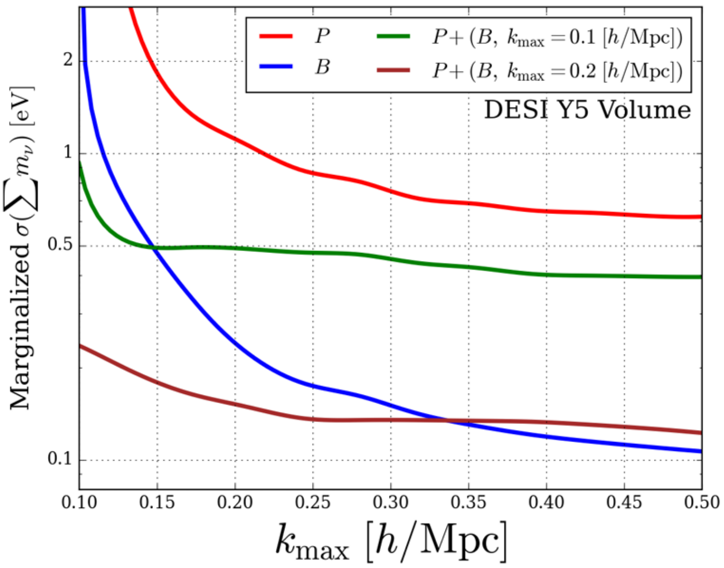

Figure 1.5: The figure, taken from Kamalinejad and Slepian (2025a), shows the marginalized 1σ error on the sum of neutrino masses $\sum m_\nu$, as a function of the maximum considered wavenumber $k_{\rm max}$, forecasted at $z = 0$ for a volume of $25\,h^{-3}\,{\rm Gpc}^3$. Forecasts are obtained using the real-space power spectrum P (red), bispectrum B (blue), and the joint power spectrum + bispectrum, with a scale cut for the bispectrum at $k_{\rm max} = 0.1\,h\,{\rm Mpc}^{-1}$ (green) or $k_{\rm max} = 0.2\,h\,{\rm Mpc}^{-1}$ (dark red). SPT models of one-loop power spectrum and tree-level bispectrum are employed in the forecasts, while covariance is numerically estimated from the Quijote simulations (Villaescusa-Navarro et al. 2020).

For the particular case of $\sum m_\nu$, the strong improvement as a function of $k_{\rm max}$ in the constraining power of the bispectrum with respect to the power spectrum can be clearly seen in Fig. 1.5, taken from the forecasts in Kamalinejad and Slepian (2025a). In particular, the marginalized 1σ uncertainty on $\sum m_\nu$ from the bispectrum (blue curve) in the shown range $0.10 \leq k_{\rm max}[h\,{\rm Mpc}^{-1}] \leq 0.5$ is always lower than the same quantity forecasted from the power spectrum (red curve). Moreover, their difference increases with $k_{\rm max}$, since above $\sim 0.3\,h\,{\rm Mpc}^{-1}$, the power spectrum curve essentially saturates.

Despite the promising results obtained from bispectrum analyses, there is still no published analysis of the 3PCF in massive neutrino cosmologies, despite the observational advantages offered by configuration space over Fourier space, and the substantial reduction in the computational cost introduced by the new 3PCF estimation algorithms, as discussed in Sec. 1.2.7. Not least, the 3PCF provides the additional benefit of a more intuitive interpretation, as visualizing triplets in real space is considerably simpler than triangular configurations of Fourier-space modes.

These motivations, therefore, lead us to perform, as far as we know, the first measurement of the 3PCF in simulations with massive neutrino cosmologies. Using metrics we specifically developed for this work, we quantify how the neutrino signal varies with triplet configurations. We also investigate the capability of the 3PCF to break parameter degeneracies, applying our developed framework, in particular, on the $\sum m_\nu/\sigma_8$ degeneracy. This work thus represents the first step toward bridging the existing gap between configuration and Fourier spaces in the study of massive neutrino cosmologies through higher-order clustering statistics.

The remainder of this work is devoted to describing the analysis we have carried out, discussing the results obtained, and outlining the future prospects they open.

Chapter 2

Extraction of massive neutrino signal: measurements and methodology

In this chapter, we describe how we performed the measurements of the 2PCFs and 3PCFs, estimated their covariance, and developed methods to extract the neutrino signal.

Our starting dataset consists of mock halo catalogues from the QUIJOTE N-body simulations (Villaescusa-Navarro et al. 2020). We summarize the main features of the QUIJOTE suite in Sec. 2.1, and specifically of our dataset in 2.2. In Sec. 2.3 and 2.4, we respectively describe the code and the estimators we used to perform the measurements, and how we derived from those the quantities needed for our analysis, including covariance. Finally, in Sec. 2.5 we focus on the metrics we developed and applied to quantify the detectability of the signal from massive neutrinos.

Hereafter, to refer to the sum of neutrino masses, we will use the notation M_ν rather than $\sum m_\nu$, corresponding to the one adopted in the QUIJOTE suite.

2.1 The Quijote simulations

The QUIJOTE simulations[1] (Villaescusa-Navarro et al. 2020) are a suite of more than 82 000 full N-body simulations, designed to allow the scientific community to quantify the information content on different statistics up to the fully nonlinear regime. These simulations have been run using the code GADGET-III, an improved version of GADGET-II (Springel 2005). Each set of simulations is characterized by specific values of the cosmological parameters listed in Tab. 2.1. The suite includes simulations with massive neutrinos, assuming three degenerate neutrino masses.

The initial conditions (ICs) of the simulations are generated at $z_i = 127$, from an input matter power spectrum and transfer functions obtained by rescaling the $z = 0$ matter power spectrum and transfer functions from CAMB (Lewis et al. 2000). Due to the scale dependence introduced by massive neutrinos in the growth factor, for models in

---

[1] `https://quijote-simulations.readthedocs.io/en/latest/`

Parameter(s)	Definition
Ω_m, Ω_b	Total matter (m) and baryon (b) density parameters
h	Reduced Hubble constant
n_s	Spectral index of the primordial power spectrum
σ_8	Present-day RMS amplitude of linear fluctuations on scale $8\,h^{-1}$Mpc
M_ν	Sum of neutrino masses
w	DE EOS parameter
δ_b	background overdensity in the simulated box
$f_{\rm NL}^{\rm loc}, f_{\rm NL}^{\rm equ}, f_{\rm NL}^{\rm ort}$	amplitude of the local, equilateral and orthogonal shapes of the primordial bispectrum in presence of PNG (Coulton et al. 2023)
$p_{\rm NL}$	parity-violation amplitude
f_{R_0}	present-day amplitude of the f_R field in the modified-gravity model by Hu and Sawicki (2007)

Table 2.1: The parameters characterizing the simulations in the Quijote suite and their corresponding description.

which $M_\nu > 0$, the rescaled method described in Zennaro et al. (2017) is used, while for all cases in which $M_\nu = 0$, a scale-independent rescaling is adopted, namely

$$P_{m,\rm lin}(k, z_i) = \left(\frac{D_1(z_i)}{D_1(z=0)}\right)^2 P_{m,\rm lin}(k, z=0). \tag{2.1.1}$$

Matter particles are initially distributed on a regular grid. In models with massive neutrinos, two grids are used, one for CDM and one for neutrinos, offset by half a grid step. From the input power spectrum, initial displacements and peculiar velocities are computed through ZA (Sec. 1.3.2) for massive neutrino models, and with 2LPT (Sec. 1.3.3) for cosmologies with massless neutrinos, and assigned to particles. In simulations with $M_\nu > 0$, the cosmic neutrino background is modelled by adopting the particle-based method (Brandbyge et al. 2008; Viel et al. 2010), in which neutrinos are described as a collisionless and pressureless fluid, discretized into a set of particles. On top of peculiar velocities, neutrino particles are also assigned thermal velocities by randomly sampling their Fermi-Dirac distribution at $z = z_i$.

All simulations have been run in a periodic box with cosmological volume of $1\,(h^{-1}\,\rm Gpc)^3$, and for the vast majority follow the evolution of 512^3 CDM particles (with additional 512^3 particles for simulations including massive neutrinos): this is referred to as *fiducial resolution*. The Quijote suite also contains a lower number of simulations run with 256^3 and 1024^3 CDM particles (respectively *low* and *high resolution*). The gravitational softening length is set to $1/40$ of the mean interparticle distance for both CDM and neutrino particles (if present): for the fiducial resolution, this corresponds to $50\,h^{-1}$kpc.

ICs are generated in different ways depending on whether the simulation is *standard*, *fixed*, or *paired fixed*. Let $\tilde{\delta}(\boldsymbol{k})$ be a generic Fourier mode of the initial density contrast field: being a complex number, it can be written in terms of an amplitude and a phase as $\tilde{\delta}(\boldsymbol{k}) = Ae^{i\theta}$, where A and θ depend on $\boldsymbol{k}$. For a Gaussian density field, A follows a *Rayleigh distribution*,

$$p(A;\sigma) \equiv \frac{A}{\sigma^2} e^{-A^2/(2\sigma^2)}, \tag{2.1.2}$$

while θ is drawn from a uniform distribution between 0 and 2π. This corresponds to standard ICs. In fixed simulations, the difference with respect to the standard case is that the value of A is fixed to the square root of the variance of the Rayleigh distribution, which from Eq. 2.1.2 can be computed to be $(4-\pi)\sigma^2/2$. Finally, paired fixed simulations are two fixed simulations where the phases of the two pairs differ by π. It has been shown (e.g., Angulo and Pontzen 2016) that fixed and paired fixed simulations significantly reduce the amplitude of cosmic variance in different statistics without biasing the results, making them particularly useful in the computation of numerical derivatives. However, for the same reason, they cannot be used for covariance estimations.

Halos are identified by running the Friends-of-Friends (FoF) algorithm (Davis et al. 1985) with linking length equal to $b = 0.2$ times the mean interparticle distance, and saved into catalogues only if the number of linked particles is above 20. For fiducial resolution, this corresponds to saving halos with masses above $\approx 1.3 \times 10^{13}\,h^{-1}M_\odot$. Only CDM particles are linked, even in the presence of neutrino particles, since it has been shown that this affects halo masses negligibly (Villaescusa-Navarro et al. 2011) and avoids the risk of linking unclustered neutrinos below their free-streaming scale.

Snapshots and halo catalogues are saved for redshifts $z = 3, 2, 1, 0.5$ and 0: at these values of z, for simulations run at fiducial resolution, the number of halos contained in the catalogues is of the order of $\sim 5 \times 10^3, 4.4 \times 10^4, 2 \times 10^5, 3.1 \times 10^5$ and 4×10^5 respectively.

2.2 Simulated dataset

The simulations considered in this work, together with their corresponding parameters, are listed in Tab. 2.2.

All the simulations shown in the table share the same fiducial resolution and differ in the values of the parameters M_ν and σ_8, as well as in the ICs. The values of the parameters $\Omega_m, \Omega_b, h, n_s$ and w do not vary across the different simulations and are, respectively, 0.3175, 0.049, 0.6711, 0.9624 (compatibly with *Planck18* values given in Planck Collaboration: Aghanim et al. 2020), while the value of Ω_Λ is given by $1 - \Omega_m = 0.6825$ since they simulate a flat Universe. The other parameters listed in Tab. 2.1 are all 0 for the considered set of simulations.

We use 2000 realizations from the set of `fiducial` simulations, run with standard 2LPT ICs, for the estimation of covariance matrices. For the analysis of the effect of massive neutrinos on clustering, we consider the simulations named `Mnu_ppp`, `Mnu_pp`,

Name	M_ν (eV)	σ_8	Realizations	Type	ICs	$N_c^{1/3}$	$N_\nu^{1/3}$
`fiducial`	0	0.834	2 000	standard	2LPT	512	0
`fiducial`	0	0.834	500	paired fixed	2LPT	512	0
`fiducial_ZA`	0	0.834	500	standard	ZA	512	0
`Mnu_ppp`	0.4	0.834	500	standard	ZA	512	512
`Mnu_pp`	0.2	0.834	500	standard	ZA	512	512
`Mnu_p`	0.1	0.834	500	standard	ZA	512	512
`s8_p`	0	0.849	500	paired fixed	2LPT	512	0
`s8_m`	0	0.819	500	paired fixed	2LPT	512	0

Table 2.2: The simulations for which the measurements have been performed and their corresponding parameters. The fourth column refers to the number of realizations measured in this work: this number is equal to the total number of realizations present in the QUIJOTE suite per simulation type, with the only exception of `fiducial` standard simulations, for which the total number of realizations available in the suite is 15 000. N_c and N_ν are, respectively, the number of CDM and neutrino particles in the grid. All these simulations share the same values of the parameters Ω_m, Ω_b, h, n_s, w respectively equal to 0.3175, 0.049, 0.6711, 0.9624, -1. The parameters δ_b, $(f_{\rm NL}^{\rm loc}, f_{\rm NL}^{\rm equ}, f_{\rm NL}^{\rm ort})$, $p_{\rm NL}$ and f_{R_0} are 0. All the listed realizations simulate a flat Universe, therefore $\Omega_\Lambda = 1 - \Omega_m$.

and `Mnu_p`, in which three degenerate neutrino masses are simulated, with M_ν respectively equal to 0.4, 0.2, and 0.1 eV. To compare the results with the massless case, we select the set of simulations `fiducial_ZA`, characterized by the same initial conditions (standard ZA) as the massive neutrino simulations. The neutrino masses in these simulations, being more physically plausible, are significantly smaller than those chosen to produce the snapshot in Fig. 1.3. In our case, therefore, the effects imprinted on the LSS are much less evident by a simple visual inspection of the snapshots. We illustrate this in Fig. 2.1, where we compare two $z = 0$ snapshots from the `fiducial_ZA` and `Mnu_ppp` simulations showing the CDM distribution. As can be seen, the differences are minimal and are noticeable mainly in the most overdense regions along the filaments of the cosmic web.

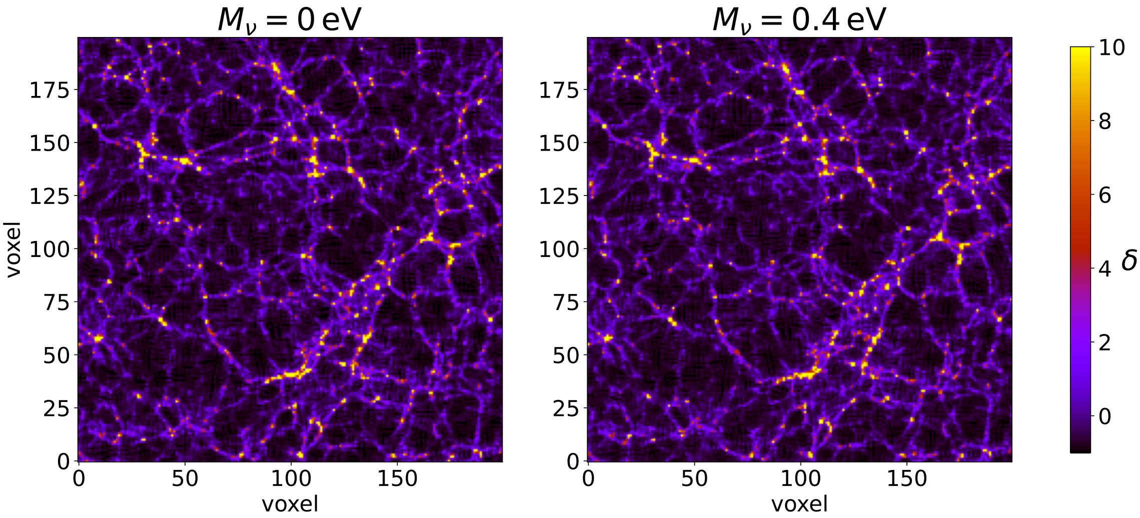

Figure 2.1: Detail of two snapshots from QUIJOTE `fiducial_ZA` simulations (with $M_\nu = 0\,\text{eV}$, *left*) and `Mnu_ppp` simulations (with $M_\nu = 0.4\,\text{eV}$, *right*) showing the dark matter density contrast field averaged over a slice with depth $10\,h^{-1}\,\text{Mpc}$, at $z = 0$. The two simulations share the same initial conditions and the same value of $\sigma_8 = 0.834$. The panels show a close-up of a region of width $0.4\,h^{-1}$Gpc in both horizontal and vertical directions.

Finally, to carry out the analysis of degeneracy between M_ν and σ_8 reported in Ch. 4, we consider the simulations named `s8_p` and `s8_m` run with paired fixed 2LPT ICs, along with the `fiducial` ones with the same ICs, for comparison with the fiducial model. For `fiducial`, `fiducial_ZA`, `Mnu_ppp`, `Mnu_pp` and `Mnu_p` simulations, the value of the parameter σ_8 is equal to 0.834, whereas it differs by $\Delta\sigma_8 = +0.015$ for `s8_p` and $\Delta\sigma_8 = -0.015$ for `s8_m`, resulting in $\sigma_8(\texttt{s8\_p}) = 0.849$ and $\sigma_8(\texttt{s8\_m}) = 0.819$. Except for the simulations used to compute covariance matrices, we employ the total number of available realizations for each type, i.e., 500 mocks.

2.3 Measurements

We measure the monopole, quadrupole, and hexadecapole of the halo 2PCF and the isotropic multipoles of the halo 3PCF ζ_ℓ for ℓ in the range $0 \leq \ell \leq 10$, in real and redshift space, and at redshifts $z = 0, 1$, and 2.

We move halos to redshift space along the z-axis. To read the positions and peculiar velocities of the halos from the corresponding catalogues, we use the set of Python libraries `Pylians3`[2], and employ the routine `pos_redshift_space`: this routine simply transforms the component r_i of real-space halo positions (in units of length/h) along an axis i parallel to one of the box sides to the corresponding redshift-space component s_i by applying the usual relation

$$s_i = r_i + \frac{v_i}{H(z)}(1+z), \tag{2.3.1}$$

where v_i is the component of the physical peculiar velocity of a given halo along the i axis in units of $\mathrm{km\,s^{-1}}$, and $H(z)$ is the value of the Hubble parameter in units of $\mathrm{h\,km\,s^{-1}\,Mpc^{-1}}$, computed with the relation

$$H(z) = 100\,\sqrt{\Omega_m(1+z)^3 + \Omega_\Lambda}\;\; h\,\mathrm{km\,s^{-1}\,Mpc^{-1}}. \tag{2.3.2}$$

2.3.1 Estimators

To measure correlation functions from observations (e.g., a survey) or simulations, one must use estimators. For a generic N-point correlation function ξ^N, an unbiased and minimum-variance estimator has been proposed in Szapudi and Szalay (1998), and reads:

$$\hat{\xi}^N(\boldsymbol{x_1}, \ldots, \boldsymbol{x_N}) \equiv \frac{(D-R)^N}{R^N}, \tag{2.3.3}$$

where D and R indicate, respectively, the counts in the data and the random sample (properly normalized, see below): this latter sample indicates a set of unclustered objects distributed in the same volume and sharing the same selection effects of the data. For example, for the 2PCF ξ and the 3PCF ζ, this estimator takes respectively the forms

$$\hat{\xi} = \frac{DD - 2DR + RR}{RR}, \qquad \hat{\zeta} = \frac{DDD - 3DDR + 3DRR - RRR}{RRR}. \tag{2.3.4}$$

The DD, DR and RR counts are normalized by $N_D^2/2, N_D N_R$ and $N_R^2/2$ respectively, while the DDD, DDR, DRR and RRR counts by $N_D^3/6, N_D^2 N_R/2, N_D N_R^3/2$ and $N_R^3/6$ respectively, where N_D and N_R denote the number of objects in the data and random catalogue. Notice that the expression for $\hat{\xi}$ is merely the well-known Landy and Szalay estimator (Landy and Szalay 1993). The estimator in Eq. 2.3.3 automatically subtracts the disconnected part of the 3PCF and corrects for selection and edge effects due to survey

---

[2] https://github.com/franciscovillaescusa/Pylians3

boundaries. However, these effects are absent in our catalogues since they are obtained from simulations run in periodic boxes.

When adopting the decomposition of ζ into Legendre multipoles, as in Eq. 1.2.26, the analogue of the estimator for ζ in Eq. 2.3.4 for the 3PCF isotropic multipoles ζ_ℓ can be expressed as

$$\hat{\zeta}_\ell = \frac{DDD_\ell - 3DDR_\ell + 3DRR_\ell - RRR_\ell}{RRR_0}, \tag{2.3.5}$$

where $DDD_\ell, RRR_\ell, DDR_\ell, DRR_\ell$ are the coefficients of the Legendre expansion of the triplet counts from the data, from the random, and the two mixed terms. The implementation of the estimator of Eq. 2.3.5 in the approach of Slepian and Eisenstein (2015a), used in our 3PCF measurements, is detailed in the following section.

2.3.2 Measuring the correlation functions

To perform the measurements on the QUIJOTE halo catalogues, we use the C++/Python-wrapped package `MeasCorr`[3] (Farina et al. 2024). The code implements the estimators in Eq. 2.3.4. For a periodic box, the Landy-Szalay estimator for the 2PCF reduces to the "natural" one first introduced in Peebles and Groth (1975), i.e.,

$$\hat{\xi}(r,\mu) = \frac{DD(r,\mu)}{RR(r,\mu)} - 1, \tag{2.3.6}$$

where r and μ are, respectively, the binned pair separation and the cosine of the angle between the pair and the l.o.s. direction. The 2PCF is measured for a given range of separations, $r_{\min} \leq r \leq r_{\max}$, and by choosing a binning for r and μ, i.e., the respective bin sizes Δr and $\Delta\mu$. To estimate the DD vector, the code distributes the objects in the catalogue over a grid and then iterates over the particles contained in each non-empty cell, by finding the neighbouring particles whose separations are within the limit set by $r_{\max}$. Only the neighbouring non-empty cells within $r_{\max}$ are considered in the search, to speed up computations. The DD vector is increased by one count at the index of the corresponding (r,μ) bin whenever a pair in that bin is found, and normalized at the end of the cycles.

The normalized RR term is instead estimated analytically due to periodicity: in the bin $[r_i; r_{i+1}] \times [\mu_j; \mu_{j+1}]$, for a box whose side length is L, it is simply given by

$$RR_{ij} = \frac{1}{L^3}\frac{1}{N_{\mu\,\mathrm{bins}}} \times \frac{4}{3}\pi(r_{i+1}^3 - r_i^3), \tag{2.3.7}$$

where $N_{\mu\,\mathrm{bins}}$ is the number of μ bins.

Once the estimator is computed, the estimated monopole, quadrupole, and hexadecapole of the 2PCF are obtained from it by projecting $\hat{\xi}$ onto the corresponding Legendre polynomials, i.e., by evaluating numerically the integrals

---

[3] https://gitlab.com/veropalumbo.alfonso/meascorr

$$\hat{\xi}_\ell(r) = \frac{2\ell+1}{2}\int_{-1}^{1} d\mu\, P_\ell(\mu)\,\hat{\xi}(r,\mu), \tag{2.3.8}$$

with $\ell = 0, 2, 4$.

The isotropic 3PCF is estimated following the approach in Slepian and Eisenstein (2015a), which we have already described in Sec. 1.2.9. The code computes separately the contributions coming from triplets in which data points and randoms occupy the primary vertex. These contributions are respectively given by $DDD - DDR - DRD + DRR$ and $RRD + RDR - RDD - RRR$. The basic parameters to be chosen are the maximum and minimum triangle side $r_{\rm min}$ and $r_{\rm max}$, the radial binning Δr, and the maximum order of the expansion, $\ell_{\rm max}$.

We briefly outline the computation of the first contribution. The implementation of the computation of the spherical harmonic coefficients $a_{\ell m}$ in Eq. 1.2.35 relies on a pair-searching scheme analogous to the one described above for the 2PCF estimator. Before looping over the cells, the normalizations of the spherical harmonics $Y_{\ell m}$ in Eq. 1.2.33 are pre-computed up to $\ell_{\rm max}$ for all the radial bins, and stored in a vector with size $N_{r\,\rm bins} \times (\ell_{\rm max}+1)(\ell_{\rm max}+2)/2$ [4]. When the search finds a pair formed by a given object and a neighbour separated by a vector $\boldsymbol{r}$, with $r_{\rm min} \leq r \leq r_{\rm max}$, the code computes the terms $P_{\ell m}(\cos\theta)\, e^{im\varphi}$, with $\hat{r} \equiv \boldsymbol{r}/r = (\theta, \varphi)$, up to $\ell = \ell_{\rm max}$ and updates the vector elements corresponding to the radial bin in which r falls. For a periodic box, when the loop over the neighbours of a given particle ends, an isotropic signal representing the randoms is added to the $\ell = 0, m = 0$ components of all the radial bins. The number of random neighbours R_i present in the radial bin $[r_i; r_{i+1}]$ around a given particle is estimated analytically, and in order to match the number density of data points, it is computed as

$$R_i = -\frac{N_D}{L^3} \times \frac{4}{3}\pi\,(r_{i+1}^3 - r_i^3), \tag{2.3.9}$$

The negative sign accounts for the fact that the R term is preceded by a minus sign in the numerator of Eq. 2.3.3. Finally, the different bins are correlated according to Eq. 1.2.36.

For the second contribution, i.e., the one in which randoms occupy the first vertex, an analogous procedure is implemented, with the difference that neighbours in the data catalogue are searched around particles randomly generated in each cell with a uniform spatial distribution. The number of random particles per cell $N_{\rm rand}$ to be generated can be specified by the user. These particles are assigned a weight $w = -N_D/(N_{\rm rand}\, N_{\rm cells})$, where $N_{\rm cells}$ is the number of cells of the grid, to match the density of the data; the minus sign is included for the same reason mentioned above. To minimize fluctuations in the result, the procedure can be repeated $N_{\rm repeat}$ times, and the results of the correlation from each repetition are eventually averaged.

The RRR term is fully analytic due to the periodicity of the box. In a generic (r_{12}, r_{13}) bin $[r_{12,i}; r_{12,i+1}] \times [r_{13,j}; r_{13,j+1}]$, the number of RRR triplets normalized to the number density of the data catalogue is given by

---

[4] The quantity $(\ell_{\rm max}+1)(\ell_{\rm max}+2)/2$ corresponds to the total number of spherical harmonics $Y_{\ell m}$ with $\ell \in \{0, 1, \ldots, \ell_{\rm max}\}$ and $m \in \{0, 1, \ldots, \ell\}$. Spherical harmonics with $m \in \{-1, -2, \ldots, -\ell\}$ can be easily retrieved with the relation $Y_{\ell\,-m}(\theta,\varphi) = (-1)^m Y^*_{\ell m}(\theta,\varphi)$.

2PCF

Parameter	$r_{\min}$ (h^{-1}Mpc) $\Delta\mu$	$r_{\max}$ (h^{-1}Mpc) ε	Δr (h^{-1}Mpc)
Value	1 0.01	150 0.15	1

3PCF

Parameter	$r_{\min}$ (h^{-1}Mpc) $\ell_{\max}$ N_{repeat}	$r_{\max}$ (h^{-1}Mpc) ε	Δr (h^{-1}Mpc) N_{rand}
Value	2.5 10 5	147.5 0.15	5 10

Table 2.3: Parameter settings for the measurements of the 2PCF and 3PCF.

$$RRR_{ij} = N_D \times \left(\frac{N_D}{L^3}\right)^2 \times \frac{4}{3}\pi\,(r^3_{12,i+1} - r^3_{12,i}) \times \frac{4}{3}\pi\,(r^3_{13,i+1} - r^3_{13,i}). \tag{2.3.10}$$

The computational workload is efficiently parallelized using `OpenMP`[5]. Each thread independently handles a subset of grid cells and computes partial correlations, which are combined into the final result with minimal synchronization overhead.

Tab. 2.3 lists the parameter settings used for running the measurements. All the listed parameters have been defined above, with the only exception of ε, which represents the side length of the cells of the grid in units of $r_{\max}$.

We choose a finer radial binning and a smaller minimum separation $r_{\min}$ for the 2PCF compared to the 3PCF, i.e., $\Delta r^{(2\text{pt})}, r_{\min}^{(2\text{pt})} = 1\,h^{-1}\,\text{Mpc}$ versus $\Delta r^{(3\text{pt})}, r_{\min}^{(3\text{pt})} = 5\,h^{-1}\,\text{Mpc}$, to ensure an accurate linear interpolation of $\hat{\xi}_0$ in the computation of the reduced 3PCF (see Sec. 2.4), and set the parameter $\Delta\mu$ to 0.01 to guarantee an accurate numerical evaluation of the integrals in Eq. 2.3.8. We choose the values of $r_{\min}$ and $r_{\max}$ (for this latter, $r_{\max}^{(2\text{pt})} = 150\,h^{-1}\,\text{Mpc}$, $r_{\max}^{(3\text{pt})} = 147.5\,h^{-1}\,\text{Mpc}$) to include both small scales, in fully nonlinear regime, and BAO scales. We adopt the default values for $N_{\text{rand}} = 10$, $N_{\text{repeat}} = 5$ and $\varepsilon = 0.15$. We choose to measure the isotropic multipoles of the 3PCF up to $\ell_{\max} = 10$ because for the vast majority of scales (r_{12}, r_{13}), this is sufficient to reconstruct the angular dependence of the 3PCF accurately (see the discussion in Sec. 2.4), without requiring excessive computational time.

---

[5] `https://www.openmp.org/`

We run all the measurements on the *Matrix* partition of the *Open Physics Hub* (OPH) Cluster[6] at the Physics and Astronomy Department (DIFA) of the University of Bologna, employing approximately 330 000 CPU hours.

2.4 Extraction of clustering statistics from mock catalogues

2.4.1 Lower- and higher-order clustering statistics

From our measurements, we estimate the following quantities:

2PCF and 3PCF multipoles. We simply take the averages over each set of mocks of the multipoles of the 2PCF (monopole, quadrupole, and hexadecapole) and 3PCF ($0 \leq \ell \leq 10$) measured for each mock. In particular, we also use the average monopole of the 2PCF to estimate the denominator of the reduced 3PCF, as detailed below.

Single-scale connected 3PCF. This function allows for the inspection of the angular behavior of the 3PCF by fixing two sides of a triangle, e.g., r_{12} and r_{13}, and varying the angle θ between them. It can be obtained from the isotropic multipoles $\zeta_\ell(r_{12}, r_{13})$ through the already discussed expansion in Eq. 1.2.26. Since only a finite number of measured multipoles can be available, $\zeta(r_{12}, r_{13}, \theta)$ is estimated by truncating the sum at $\ell = \ell_{\rm max}$, yielding

$$\hat{\zeta}(r_{12}, r_{13}, \theta) = \sum_{\ell=0}^{\ell_{\rm max}} \bar{\hat{\zeta}}_\ell(r_{12}, r_{13}) P_\ell(\cos\theta), \tag{2.4.1}$$

where $\bar{\hat{\zeta}}_\ell(r_{12}, r_{13})$ denotes the average over the realizations of a given simulation of the estimated multipoles $\hat{\zeta}_\ell$ in the (r_{12}, r_{13}) bin.

All-scales connected 3PCF. This function is useful to show the information contained in the 3PCF for all the considered side-binned triangles. Suppose the centers of the side bins are $C = \{r_{\rm min} + n \times \Delta r/2,\ n = 1, 2, \ldots, N_{\rm r\,bins}\}$. The *all-scales* 3PCF is estimated by constructing the following function:

$$\hat{\zeta}_{\rm all}(r_{12}, r_{13}, r_{23}) = \sum_{\ell=0}^{\ell_{\rm max}} \bar{\hat{\zeta}}_\ell(r_{12}, r_{13})\, P_\ell(r_{12}, r_{13}, r_{23}), \tag{2.4.2}$$

where the triplets (r_{12}, r_{13}, r_{23}) satisfy the following conditions:

---

[6] https://apps.difa.unibo.it/wiki/oph:cluster

$$\begin{cases} r_{12}, r_{13}, r_{23} \in C \\ r_{12} \leq r_{13} \leq r_{23} \\ r_{23} \leq r_{12} + r_{13} \end{cases} \tag{2.4.3}$$

and

$$P_\ell(r_{12}, r_{13}, r_{23}) \equiv P_\ell(\cos\theta), \quad \cos\theta = \frac{r_{12}^2 + r_{13}^2 - r_{23}^2}{2\, r_{12} r_{13}}. \tag{2.4.4}$$

The second condition in Eq. 2.4.3 avoids redundancy, while the third one ensures that the triplet (r_{12}, r_{13}, r_{13}) represents the sides of a triangle. Notice that one can associate an index to order each triplet (r_{12}, r_{13}, r_{23}) satisfying the above conditions by deciding how to loop over the sides. In this work, we always loop through r_{23} in the innermost loop and r_{12} in the outermost loop.

Single-scale reduced 3PCF. The reduced 3PCF Q is defined as (Groth and Peebles 1977)

$$Q(r_{12}, r_{13}, r_{23}) \equiv \frac{\zeta(r_{12}, r_{13}, r_{23})}{\xi_0(r_{12})\xi_0(r_{13}) + \xi_0(r_{13})\xi_0(r_{23}) + \xi_0(r_{23})\xi_0(r_{12})}. \tag{2.4.5}$$

This quantity guarantees an optimal combination of 3PCF and 2PCF: indeed Q exhibits less variation with scale, since it can be shown that in hierarchical scenarios $\zeta \propto \xi^2$ (Peebles and Groth 1975), and it is explicitly independent of σ_8 since $\xi \propto \sigma_8^2$ and $\zeta \propto \sigma_8^4$ (see for example the discussion in Moresco et al. 2021).

We estimate the single-scale reduced 3PCF in the following way:

$$\hat{Q}(r_{12}, r_{13}, \theta) = \frac{\hat{\zeta}(r_{12}, r_{13}, \theta)}{\bar{\hat{\xi}}_0^{\,\text{interp.}}(r_{12})\, \bar{\hat{\xi}}_0^{\,\text{interp.}}(r_{13}) \; + \text{cyc.perm.}}, \tag{2.4.6}$$

where $\bar{\hat{\xi}}_0^{\,\text{interp.}}(r)$ denotes the average estimated monopole of the 2PCF linearly iterpolated at r.

All-scales reduced 3PCF. This quantity is estimated as

$$\hat{Q}_{\text{all}}(r_{12}, r_{13}, r_{23}) = \frac{\hat{\zeta}_{\text{all}}(r_{12}, r_{13}, r_{23})}{\bar{\hat{\xi}}_0^{\,\text{interp.}}(r_{12})\, \bar{\hat{\xi}}_0^{\,\text{interp.}}(r_{13}) \; + \text{cyc.perm.}}, \tag{2.4.7}$$

with triplets (r_{12}, r_{13}, r_{23}) satisfying the same conditions in Eq. 2.4.3.

Notice that we could have also decided to estimate the average reduced 3PCF by computing $\hat{Q}$ for each realization, and then by averaging over the set of realizations. However, in this case it is impossible to avoid the noise associated with the zero crossing of the denominator of Q, due to the change of sign of $\xi_0(r)$ for $r \gtrsim 120\, h^{-1}$Mpc, which occurs at different triangular configurations depending on the particular realization. Taking the ratio between pre-averaged quantities (over a large number of realizations, as in our case) significantly reduces this noise and allows for the avoidance of the problem of zero crossing in the denominator by simply excluding triangle sides above a given threshold. For this reason, we estimate Q only for triangles whose side lengths do not exceed the value $r = 110\,h^{-1}\mathrm{Mpc}$, for which the estimated function (both single-scale and all-scales) remains stable.

A comment is needed about Eq. 2.4.1: the reconstruction is accurate only if the multipoles with $\ell > \ell_{\max}$ are negligible. This in general depends on r_{12} and r_{13}. A comparison can be made with measurements obtained from direct triplet counts instead of the SHD estimator: in particular, Veropalumbo et al. (2021) show that when $r_{12} \approx r_{13}$, the 3PCF estimated from direct triplet counts rapidly steepens for $\theta \to 0$ (meaning that the third side $r_{23} \to 0$), but this steepening cannot be accurately reconstructed using the SHD estimator even for $\ell_{\max} \sim 30$. This lack of accuracy decreases for increasing difference between r_{13} and r_{12}, i.e., for increasing values of the parameter (Veropalumbo et al. 2022)

$$\eta \equiv \frac{|r_{13} - r_{12}|}{\Delta r}. \tag{2.4.8}$$

In this respect, setting $\eta \geq \eta_{\min}$ serves to exclude pathological configurations depending on the value of $\eta_{\min}$.

2.4.2 Covariance matrices

We estimate the covariance matrices numerically, from the 2PCF and 3PCF multipoles measured in the 2000 QUIJOTE `fiducial` mock catalogues. The covariance matrix of the multipoles is estimated as follows:

$$\hat{C}_{ij} = \frac{1}{N_m - 1} \sum_{k=1}^{N_m} \left(d_i^k - \bar{d}_i\right) \left(d_j^k - \bar{d}_j\right), \tag{2.4.9}$$

where N_m is the number of mocks, while d_i^k represents the data vector (i.e., the measured multipoles of the 2PCF or 3PCF), with indexes i and k identifying the bin and the mock sample, respectively. Here, $\bar{d}_i = (1/N_m) \sum_{k=1}^{N_m} d_i^k$. From $\hat{C}_{ij}$, the corresponding correlation matrix $\hat{c}_{ij}$ is estimated as

$$\hat{c}_{ij} \equiv \frac{\hat{C}_{ij}}{\sqrt{\hat{C}_{ii}\hat{C}_{jj}}}. \tag{2.4.10}$$

The correlation matrices of the 2PCF and 3PCF multipoles, in redshift space and at $z = 0$, are shown in Fig. 2.2. These matrices are block-structured. Each block shows the correlation between two multipoles, ℓ and ℓ' (the pairs of multipoles corresponding to a given block are obtained by crossing the values indicated along the top and right edges of each matrix). The total number of blocks is therefore N_{mult}^2, where N_{mult} denotes the number of measured multipoles: 3 for the 2PCF and 11 for the 3PCF. Each block has

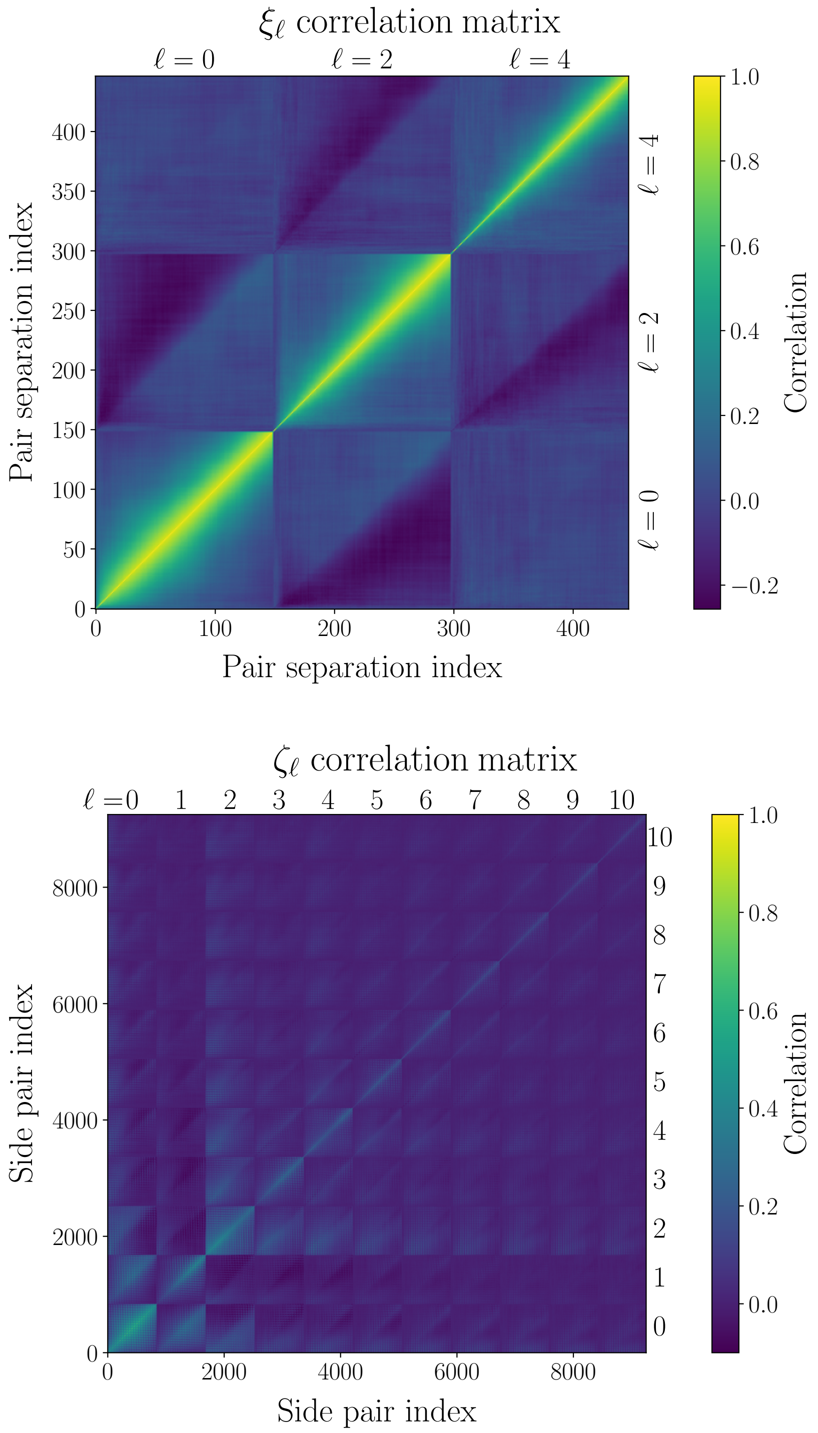

Figure 2.2: The correlation matrices of the multipoles of the 2PCF (top) and connected 3PCF (bottom), in redshift space and at $z = 0$. The multipole indices are displayed along the top and left sides in correspondence with each block.

size $N_{\mathrm{r\,bins}}$ for the 2PCF and $N^2_{\mathrm{r\,bins}}$ for the 3PCF (since in the latter case the multipoles are defined over pairs of sides). As can be derived from Tab. 2.3, the number of radial bins $N_{\mathrm{r\,bins}}$ is 149 for the 2PCF and 29 for the 3PCF. For the 3PCF, within each block, the side pairs (r_{12}, r_{13}) are ordered for increasing r, by looping internally on one of the sides and externally on the remaining one.

From the covariance of the multipoles of the 3PCF, we get the covariance matrix for the single-scale and all-scales ζ, through the following expansions into Legendre polynomials:

$$\hat{C}_{ij}^{(\mathrm{single-sc.})}(r_{12}, r_{13}) = \sum_{\ell,\ell'=0}^{\ell_{\max}} \hat{C}_{\ell,\ell'}(r_{12}, r_{13}; r_{12}, r_{13})\, P_\ell(\cos\theta_i)\, P_{\ell'}(\cos\theta_j), \tag{2.4.11}$$

$$\hat{C}_{ij}^{(\mathrm{all-sc.})} = \sum_{\ell,\ell'=0}^{\ell_{\max}} \hat{C}_{\ell,\ell'}(r_{12}, r_{13}; r'_{12}, r'_{13})\, P_\ell(r_{12}, r_{13}, r_{23})\, P_{\ell'}(r'_{12}, r'_{13}, r'_{23}). \tag{2.4.12}$$

In Eq. 2.4.11 the quantity $\hat{C}_{\ell,\ell'}(r_{12}, r_{13}; r_{12}, r_{13})$ is the estimated covariance of the multipoles $\{\zeta_\ell(r_{12}, r_{13})\}_{\ell=0}^{\ell_{\max}}$, while in Eq. 2.4.12 the quantity $\hat{C}_{\ell,\ell'}(r_{12}, r_{13}; r'_{12}, r'_{13})$ denotes the cross-covariance between the multipoles $\{\zeta_\ell(r_{12}, r_{13})\}_{\ell=0}^{\ell_{\max}}$ and $\{\zeta_{\ell'}(r'_{12}, r'_{13})\}_{\ell'=0}^{\ell_{\max}}$, and i, j represent the triangle indices associated to the triplets (r_{12}, r_{13}, r_{23}) and $(r'_{12}, r'_{13}, r'_{23})$.

The covariance of the single-scale and all-scales Q must be estimated by using error propagation (see, e.g., Cowan 1998), because of the way in which we estimate Q itself. It is known that if $\boldsymbol{x} = (x_1, \ldots, x_n)$ is a set of n random variables whose estimated means and covariance matrix are respectively $\hat{\boldsymbol{\mu}} = (\hat{\mu}_1, \ldots, \hat{\mu}_n)$ and $\hat{C}$, then the covariance matrix $\hat{U}$ of a set of m functions $y_1(\boldsymbol{x}), \ldots, y_m(\boldsymbol{x})$ of $\boldsymbol{x}$ can be estimated to first order as

$$\hat{U} = J\hat{C}\, J^T, \tag{2.4.13}$$

where

$$J_{ij} \equiv \left(\frac{\partial y_i}{\partial x_j}\right)_{\boldsymbol{x}=\boldsymbol{\mu}}. \tag{2.4.14}$$

This procedure has the advantage of involving only average quantities, which is ideal for avoiding the previously discussed noise in Q estimated from single mocks.

In the single-scale approach, if Q is evaluated at the angles $\theta_1, \ldots, \theta_n$, we take $\boldsymbol{x}$ to be

$$\boldsymbol{x} \equiv (\zeta(\theta_1), \ldots, \zeta(\theta_n), \xi_0(\theta_1), \ldots, \xi_0(\theta_n)), \tag{2.4.15}$$

where we have omitted the dependence on r_{12} and r_{13} in ζ since the sides are fixed and we have indicated with $\xi_0(\theta)$ the 2PCF monopole evaluated at a separation equal to the length of the side opposite to θ. We apply the procedure outlined above for the functions $Q(\theta_1), \ldots, Q(\theta_n)$.

In the all-scales approach, if the triangles are ordered with indices $1, \ldots, n$ and the centers of the radial bins are $r_1, \ldots, r_s$, we define $\boldsymbol{x}$ to be

$$\boldsymbol{x} \equiv (\zeta_{\mathrm{all},1}, \ldots, \zeta_{\mathrm{all},n}, \xi_0(r_1), \ldots, \xi_0(r_s)), \tag{2.4.16}$$

and we apply the procedure above for the functions $Q_{\text{all},1}, \ldots, Q_{\text{all},n}$. In both cases, the relatively simple expression of Q as a function of ζ and ξ_0 makes it possible to compute analytically the derivatives needed to construct the Jacobian matrix J in equation Eq. 2.4.14.

The correlation matrices of single-scale and all-scales connected and reduced 3PCF are shown in App. B (the former for some selected scales), in redshift space and at $z = 0$.

2.5 Methods for signal extraction

We employ several metrics, applied to both ζ and Q, to quantify the signal produced by massive neutrinos relative to the massless case, also accounting for errors. We emphasize that, due to the dimensionality of the data, in order to perform a systematic analysis without limiting it to particular cases, metrics introduced for this purpose should efficiently compress information, both in the single-scale and all-scales approaches. The outcomes produced by these compressed metrics can then be employed to guide the extraction of particular configurations for interpreting results.

In the following, we will use the letter f to refer equivalently to the functions ζ or Q. The easiest metric that can be defined is an error-weighted difference, which we refer to as *detectability*, in the form

$$\mathrm{DET}_i \equiv \frac{\hat{f}_i^{(\nu)} - \hat{f}_i^{(\text{fiducial})}}{\sqrt{2}\,\sigma_{f,i}}, \tag{2.5.1}$$

where $\hat{f}_i^{(\nu)}$ and $\hat{f}_i^{(\text{fiducial})}$ are, respectively, the functions estimated for a massive neutrino and the fiducial cosmologies in a given bin i, while $\sigma_{f,i}$ denotes the standard deviation of the function f in the same bin, extracted from the square root of the diagonal of the covariance matrix of f. The factor $\sqrt{2}$ in the denominator takes into account both the uncertainties on the average values of the functions from fiducial and massive neutrino mocks. We also use this parameter in the analysis of the 2PCF multipoles.

The detectability exploits only the diagonal elements of the covariance matrix of a given quantity, since $\sigma_{f,i} = \sqrt{\hat{C}_{f,ii}}$, i.e., it does not take into account the correlation between different bins, which is in general nonzero. For these reasons, in both the single-scale and all-scales approaches, we define additional parameters explicitly built to take into account the contribution of correlation, generalizing the detectability in the case of multiple bins whose covariance is nonzero.

For the single-scale approach, we introduce a metric that compresses the neutrino signal contained in the different triangles obtained by varying θ, resulting in a function of the fixed sides r_{12} and r_{13} only. We define this parameter as

$$\begin{aligned}\chi^2(r_{12}, r_{13}) \equiv \sum_{i,j=1}^{N_\theta} \left(\hat{f}^{(\nu)} - \hat{f}^{(\text{fiducial})}\right)(r_{12}, r_{13}, \theta_i) \times \left[\hat{C}^{(\text{single-sc.})}(r_{12}, r_{13})\right]^{-1}_{ij} \\ \times \left(\hat{f}^{(\nu)} - \hat{f}^{(\text{fiducial})}\right)(r_{12}, r_{13}, \theta_j),\end{aligned} \tag{2.5.2}$$

where N_θ is the number of angles at which we evaluate the single-scale function $\hat{f}$. The choice of the parameter's name reflects its common mathematical structure with the usual chi-square used in statistics. The correlation between different angles is now introduced by contracting the differences with the inverse covariance matrix.

We also extract from this metric a parameter which quantifies the significance of detection of the neutrino signal, in terms of the level of "disagreement" between data (the differences $\Delta\hat{f}$) and the hypothesis of no signal ($\Delta f^{(\mathrm{hyp})} = 0$). To do this, we compute the p-value associated with the quantity $\chi^2(r_{12}, r_{13})$, defined as

$$p(r_{12}, r_{13}) \equiv \int_{\chi^2(r_{12}, r_{13})}^{\infty} d\chi^2 \; G(\chi^2; N_\theta), \tag{2.5.3}$$

where $G(\chi^2; N_\theta)$ is the probability density function of a chi-square random variable with N_θ degrees of freedom. We translate this quantity into an equivalent number Z of Gaussian standard deviations, i.e., a *significance*, by following a common approach (Cowan et al. 2011): we define Z such that the probability of finding a Gaussian-distributed variable at a distance of Z standard deviations from its mean is equal to p. That is, if $\mathcal{N}$ denotes the standard Gaussian distribution (with zero mean and unitary standard deviation), Z is implicitly defined by the expression[7]

$$2 \int_{Z}^{\infty} dz \; \mathcal{N}(z) = p. \tag{2.5.4}$$

We compare the values of Z from measurements with standard thresholds $\bar{Z}$: common choices for $\bar{Z}$ are 1, 3, 5. For example, if $Z > \bar{Z} = 3$, then one says that a given null hypothesis (like the one of no signal) can be rejected with significance larger than 3σ. Tab. 2.4 shows the value of p in correspondence with these choices for $\bar{Z}$.

$\bar{Z}$	$\bar{p}$
1	0.3174
3	0.0027
5	5.73×10^{-7}

Table 2.4: The values of p (*right column*), indicated with $\bar{p}$, which solve Eq. 2.5.4 for the listed values of $\bar{Z}$ (*left column*).

We also define, within the all-scale framework, a cumulative indicator which provides information, starting from the largest scales, on how the inclusion of progressively smaller scales contributes to the increase of the signal. The choice to start from the largest scales down to smaller ones and not vice versa is motivated by the fact that the effects of nonlinearity are progressively larger for smaller separations: since models are accurate up to mildly nonlinear scales, let us say, for $r \gtrsim r_{\min}$ (because they are computed from

---

[7] Notice that another convention exists, in which the integral is performed only on the positive side of the Gaussian. In this case, Eq. 2.5.4 must be modified by removing the factor 2 in front of the integral. We adopt the choice already described as it gives $Z = 0$ for $p = 1$.

perturbation theories), such a cumulative indicator can be used to quantify the "loss" consequent to the exclusion of scales smaller than $r_{\rm min}$.

Our cumulative indicator is defined as

$$\chi^2_{\rm cum}(r_{\rm min}) \equiv \sum_{\substack{(i,j)\,\in\,\mathcal{I}\times\mathcal{I}\\ r_{12}^{(i)},r_{12}^{(j)}\geq r_{\rm min}}} \left(\hat{f}^{(\nu)}_{\rm all} - \hat{f}^{(\rm fiducial)}_{\rm all}\right)_i \left[\hat{C}^{(\rm all-sc.)}\right]^{-1}_{ij} \left(\hat{f}^{(\nu)}_{\rm all} - \hat{f}^{(\rm fiducial)}_{\rm all}\right)_j , \tag{2.5.5}$$

where $\mathcal{I}$ is the set of triangle indices associated to side triplets in the way described in Sec. 2.4, while $r_{12}^{(i)}$ and $r_{12}^{(j)}$ are the values of the smallest side in the triplets indexed by i and j, respectively.

The indicators 2.5.2 and 2.5.5 require the inversion of covariance matrices for their computation. For numerically estimated covariance matrices, this requires some care: indeed, it is known that, although the covariance matrix of a sample extracted from a given statistical population is an unbiased estimator of the covariance matrix of the population itself[8], its inverse is a biased estimator of the population precision matrix (Anderson 2003). Moreover, in the case in which also the means are estimated from data, if the number of data points per sample element n_d is greater than $N_m - 1$, the sample covariance matrix is singular. An unbiased estimator for the population precision matrix $\widehat{C^{-1}}^{(\rm unb)}$ can be obtained from the inverse of the sample covariance matrix $\hat{C}^{-1}$ as follows (Hartlap et al. 2007):

$$\widehat{C^{-1}}^{(\rm unb)} = \frac{N_m - n_d - 2}{N_m - 1}\,\hat{C}^{-1}, \tag{2.5.6}$$

where the factor in front of $\hat{C}^{-1}$ is called *Hartlap correction*. In our case, neglecting the bias in the inversion of $\hat{C}$ would lead to an overestimation of the neutrino signal.

As a final note, we emphasize that, although the various indicators defined may appear quite different at face value (Eqs. 2.5.1, 2.5.2, and 2.5.3), they are all based on the same underlying principle: they quantify the deviation of a given measurement from the fiducial one, weighted by the associated uncertainty[9]. It is, however, useful to define and use all of them, as they provide complementary information. The detectability (2.5.1) quantifies the signed deviation for a specific configuration, hence allowing us to understand not only the amplitude of the difference, but also where the clustering signal is higher or lower. The χ^2 value (2.5.2), instead, measures the average deviation over multiple configurations (e.g., for a given triangle shape as a function of the angle between the two main sides), also accounting for the covariance between different scales, which, as we will show, can be significant. Finally, the p-value (2.5.3) provides a single homogeneous metric that enables a straightforward comparison between different setups.

---

[8] If the means are also estimated from data, than this is strictly true if the normalization at the denominator of covariance is $N_m - 1$, where N_m is the number of elements of the sample.

[9] We notice that if we sum the square of Eq. 2.5.1 over multiple configurations, we obtain Eq. 2.5.2 (under the assumption of diagonal errors).

Chapter 3

Disentangling massive neutrino signal in lower- and higher-order clustering statistics

In this chapter, we discuss the results of applying our indicators to the measurements of 3PCF obtained from the simulations of massive neutrino cosmologies.

In Sec. 3.1, we describe the measurements of 2PCF (required for estimating the reduced 3PCF Q), discussing the role of variations in the scalar amplitude A_s and bias. In the following sections, we present the 3PCF analysis of the massive neutrino simulations. In particular, in Sec. 3.2, we focus on the detectability analysis of the multipoles of the isotropic 3PCF; then, we present the results obtained by applying the indicators defined in Sec. 2.5 for both ζ and Q in the single-scale and all-scales frameworks, respectively in Sec. 3.3 and 3.4, discussing the impact of massive neutrinos on triangles with different scales and shapes. For the reasons explained in the previous chapter, for Q we focus only on scales below $110\,h^{-1}\,\mathrm{Mpc}$. Finally, in Sec. 3.5, we summarize the main results found in the analysis.

In what follows, unless otherwise specified, we analyze averaged quantities over the realizations of each simulation, and compute the covariance for a volume $V = 10\,h^{-3}\,\mathrm{Gpc}^3$, representative of a generic Stage IV survey like Euclid. For instance, for the *Euclid Wide Survey*, that will cover an area of approximately $15\,000\ \mathrm{deg}^2$ (Euclid Collaboration: Scaramella et al. 2022), the comoving volume in the redshift slices $z \in [0.5, 0.7]$, $[0.9, 1.1]$, and $[1.5, 1.7]$ is, respectively, $\approx 4.5, 7.8$, and $10.6\,h^{-3}\,\mathrm{Gpc}^3$, assuming the QUIJOTE ΛCDM cosmology ($h = 0.6711$, $\Omega_m = 0.3175$, $\Omega_\Lambda = 0.6825$). With respect to our choice, errors would be larger by a factor of $\sqrt{10/4.5} \approx 1.5$ and $\sqrt{10/7.8} \approx 1.1$ in the first two cases, and smaller by a factor $\sqrt{10.6/10} \approx 1.03$ in the third one, without qualitatively impacting our results.

Since, from now on, we will only deal with estimated quantities, for notational simplicity, we omit the “hat” symbol used throughout the previous chapter to identify estimators.

3.1 Neutrino signal in the 2-point correlation function

We show in this section the measurements of the multipoles of the 2PCF for the ΛCDM+ν simulations and their comparison with the massless case represented by the set of `fiducial_ZA` mocks described in Sec. 2.2.

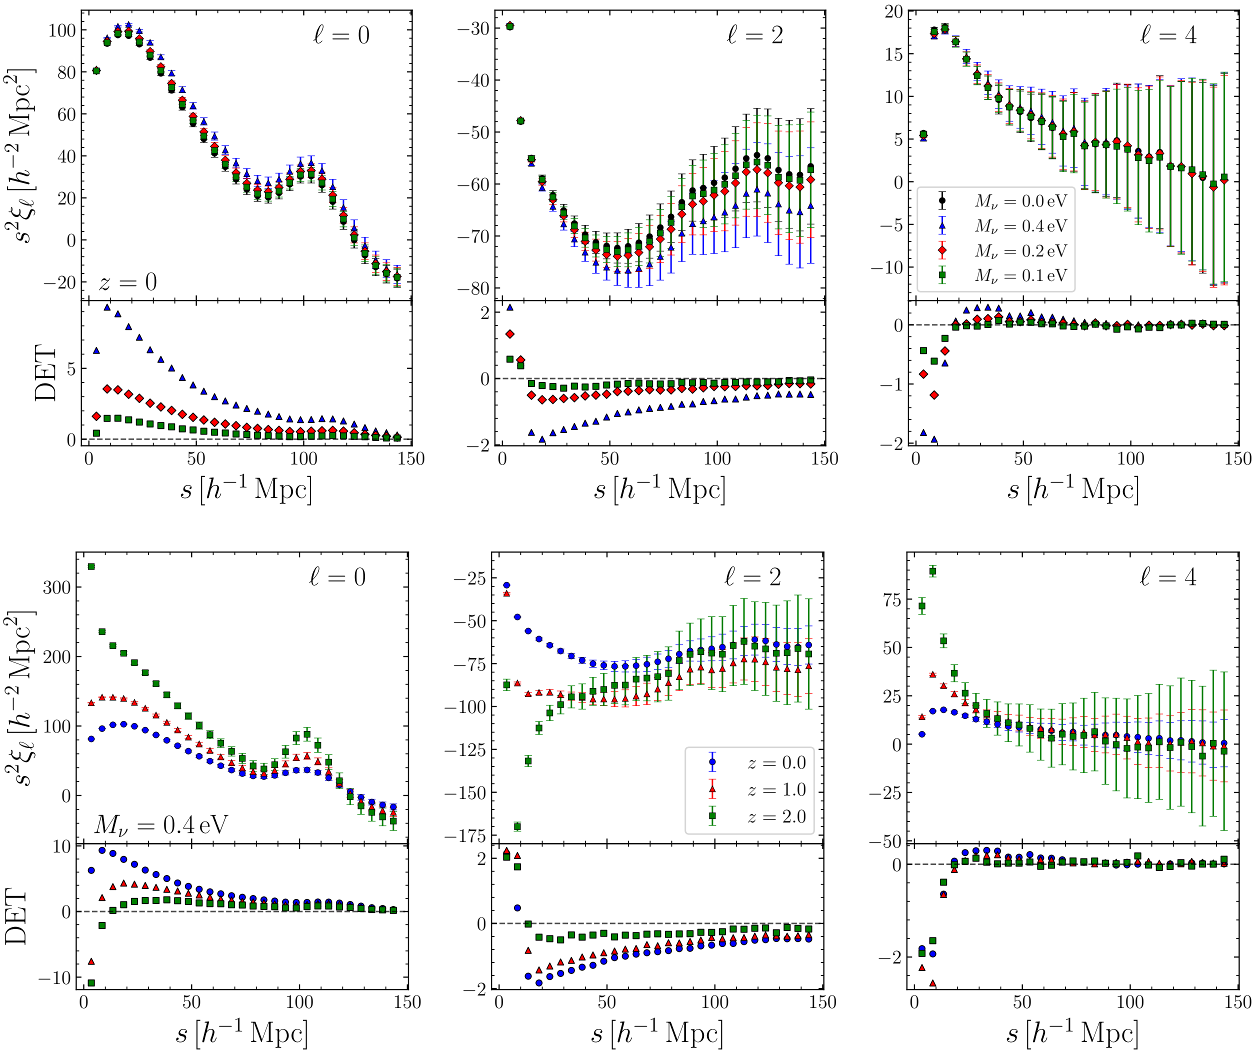

Figure 3.1: *Top row*: the measured monopoles (*left*), quadrupoles (*center*) and hexadecapoles (*right*) of the 2PCF, together with their detectabilities, at $z = 0$ for the simulations with total neutrino mass $M_\nu = 0.4, 0.2,$ and $0.1\,\mathrm{eV}$ (blue triangles, red diamonds, and green squares, respectively) and `fiducial_ZA` ones ($M_\nu = 0\,\mathrm{eV}$, black circles) averaged over the realizations of each simulation. Data are rebinned from $1\,h^{-1}$Mpc to $5\,h^{-1}$ Mpc for better readability. *Bottom row*: same as in the top row, but for fixed neutrino mass $M_\nu = 0.4\,\mathrm{eV}$ and for all the considered redshifts $z = 0, 1$ and 2 (blue circles, red triangles, and green squares, respectively).

In Fig. 3.1, we plot the average redshift-space monopoles, quadrupoles, and hexadecapoles of the 2PCF, and the corresponding detectabilities (Eq. 2.5.1). For illustration purposes, we show the multipoles at fixed redshift $z = 0$ for the different neutrino masses $M_\nu = 0.4, 0.2, 0.1\,\mathrm{eV}$ and the massless case, and for fixed mass $M_\nu = 0.4\,\mathrm{eV}$, at the different redshifts $z = 0, 1, 2$.

At fixed M_ν, the monopoles show an increase of amplitude with redshift, a behavior caused by the increase of bias associated with the larger rarity of halos with fixed minimum mass at higher z (see, e.g., Sridhar et al. 2017); at fixed redshift, the monopoles increase for progressively larger neutrino masses. Strictly speaking, this latter trend is not a simple amplitude rescaling: this can be noticed by inspecting the average values of the monopoles on scales larger than their zero crossing, occurring roughly at $\sim 120\,h^{-1}\,\mathrm{Mpc}$: although the monopoles are all negative, they still show an increase for increasing M_ν. This is somewhat expected since neutrinos affect the shape of the power spectrum/2PCF as already discussed in Sec. 1.5. The detectabilities decrease with increasing pair separation, and at BAO scales, they are, for the chosen volume, below 2 for all neutrino masses.

Let us extend the discussion on the trend exhibited by the monopoles for different M_ν. Since in our set of simulations the value of σ_8 is fixed, this behavior must result from the interplay between the scalar amplitude of the primordial power spectrum A_s, the halo bias, and, in redshift space, the effect of RSD. We have seen in Sec. 1.5 that the relative suppression induced below the free-streaming scale by massive neutrinos on the total matter and CDM power spectra, compared to the massless case, increases with the neutrino density fraction f_ν (see Eq. 1.5.9). As a consequence, it is evident from Eq. 1.2.17 that, to keep σ_8 fixed, simulations run with a higher value of M_ν (and f_ν since in our case Ω_m is fixed) are characterized by higher values of A_s.

In Villaescusa-Navarro et al. (2014), to estimate the bias, the authors measure both the matter and halo 2PCF in N-body simulations, by selecting halos above a given minimum mass. Their dataset includes simulation runs with different neutrino masses at fixed σ_8. In this case, they show that the measured bias has only a very weak dependence on the neutrino mass, even for the highest value in their dataset $M_\nu = 0.6\,\mathrm{eV}$, which is above our highest total mass $M_\nu = 0.4\,\mathrm{eV}$. This fact, when applied to the context of our measurements, implies that an increase moving to larger M_ν should already be present in the matter 2PCF, driven by the increase in A_s.

To obtain an estimate of the bias in the different simulations, along with the measured halo 2PCFs, we have included the 2PCFs of the total matter field in our dataset, which are stored in the data products of the Quijote suite. These matter correlation functions are not computed through direct pair counting, but rather by determining the Fourier transform of the density contrast field after assigning the particle masses to a regular grid with N^3 cells (where N is the cubic root of the number of particles in the simulation) with the Cloud-in-Cell (CIC) mass assignment scheme (Hockney and Eastwood 1988). The square modulus of each Fourier mode $|\tilde{\delta}(\boldsymbol{k})|^2$ is computed, and the field is Fourier-transformed back in configuration space. Eventually, the correlation functions are computed by averaging modes falling within a given radial interval.

The real-space averages of the total matter 2PCFs are shown in Fig. 3.2, for all the values of M_ν and redshifts considered in this work. As claimed above, the matter correlation functions are larger for increasing M_ν at fixed z; due to the growth of matter

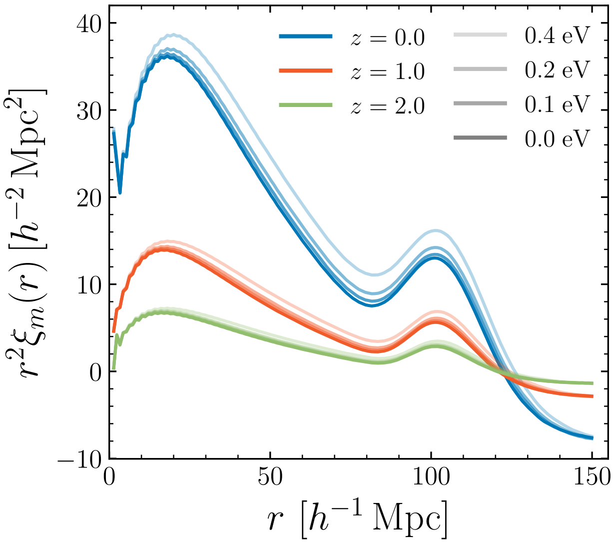

Figure 3.2: The average real-space 2PCFs of the total matter field for the `Mnu_ppp`, `Mnu_pp`, `Mnu_p` and `fiducial_ZA` simulations, with $M_\nu = 0.4, 0.2, 0.1$, and $0.0\,\mathrm{eV}$, respectively, at $z = 0, 1$, and 2 (blue, orange, and green curves, respectively). At fixed z, the 2PCFs of the different simulations are displayed with increasing opacity for decreasing M_ν. We have taken these measurements from the data products of the Quijote suite.

perturbations, their amplitude increases with decreasing redshift, going as $D_1^2(z)$ on large scales (where linear theory is still valid).

We estimate the bias at a given scale with the relation

$$b(r) = \sqrt{\frac{\xi_h(r)}{\xi_m(r)}}, \tag{3.1.1}$$

where $\xi_h(r)$ and $\xi_m(r)$ are, respectively, the real-space halo and matter 2PCFs.

We show the values of $b(r)$ in Fig. 3.3 for all the values of M_ν and z, up to a scale of $110\,h^{-1}\,\mathrm{Mpc}$ to avoid the zero crossing of ξ_m. The associated error bars are computed from the errors in the halo 2PCFs, so they correspond to the same volume of $10\,h^{-3}\,\mathrm{Gpc}^3$. At fixed redshift, the bias decreases with increasing M_ν, as opposed to A_s: indeed, for larger values of A_s, more halos are expected to form, with the consequence that the underlying matter field is better traced, resulting in lower values of bias. Conversely, as noticed before, the increase of bias with redshift reflects the fact that halos are rarer at higher z. We find a non-trivial scale dependence even on large scales, with a decreasing bias for increasing r, steeper for larger values of M_ν, which reaches a minimum in correspondence with the left dip of the BAO peak ($r \approx 85\,h^{-1}\,\mathrm{Mpc}$), and then increases again. The decrease with r of the bias even on large scales (smaller than BAO) in the presence of massive neutrinos is supported by the results obtained in Villaescusa-Navarro et al. (2014) up to a scale of $60\,h^{-1}\,\mathrm{Mpc}$ i.e., the largest scale for which the halo bias has been measured through the 2PCF in the cited work. Due to this scale limit, further comparison with this

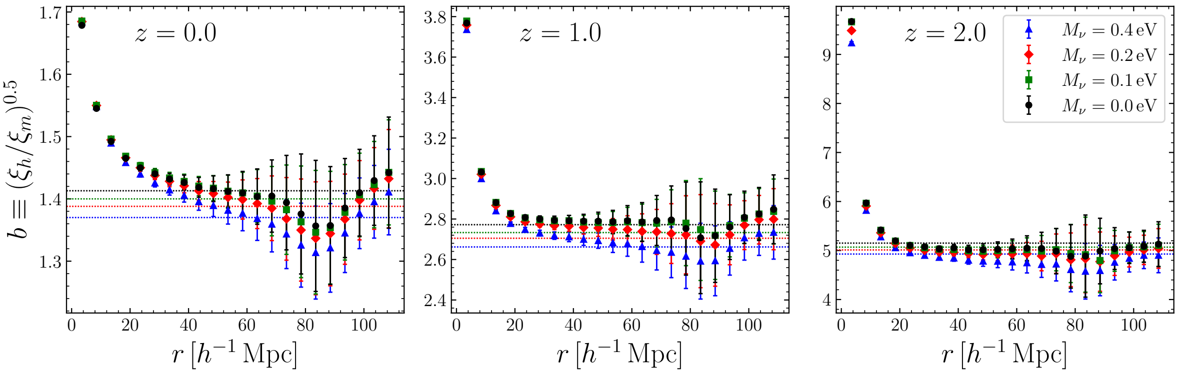

Figure 3.3: The halo bias computed as $b(r) = \sqrt{\xi_h(r)/\xi_m(r)}$, where ξ_h and ξ_m are the real-space halo and matter 2PCFs, for the `Mnu_ppp`, `Mnu_pp`, `Mnu_p`, and `fiducial_ZA` simulations (blue triangles, red diamonds, green squares, and black circles, respectively). Each plot corresponds to a different redshift, from left to right: $z = 0, 1$, and 2. The associated error bars are computed from the errors in the halo 2PCFs, so they correspond to a volume of $10\,h^{-3}\,\mathrm{Gpc}^3$. The dotted lines correspond to the effective bias obtained by weighting the Tinker bias with the average HMFs of the simulations (Eq. 3.1.2), and adopting the cold dark matter prescription.

paper on even larger scales is not possible.

When computing errors on the full effective volume of 500 $h^{-3}\,\mathrm{Gpc}^3$, the presence of the minimum is confirmed (notice, by contrast, that the chosen volume of $10\,h^{-3}\,\mathrm{Gpc}^3$ would not be enough to rule out the hypothesis of constant bias at those scales). We do not display the corresponding panel with the errors computed for our full effective volume, as it would be analogous to Fig. 3.3 with all the error bars rescaled by a factor of $1/\sqrt{50}$.

A possible origin of this feature can be found in the work of Wang and Zhan (2013), who have estimated the bias in ΛCDM N-body simulations for halos with fixed minimum mass up to very large scales ($\sim 180\,h^{-1}\,\mathrm{Mpc}$) using Eq. 3.1.1. They find a decrease of bias on scales below $90\,h^{-1}\,\mathrm{Mpc}$, as in our case near the trough of the BAO signal, compared to the average bias between 50 and 70 $h^{-1}\,\mathrm{Mpc}$, and a positive offset with respect to this average on larger scales. Interestingly, they show that the amplitude of the drop is positively correlated with the minimum halo mass chosen for bias estimation, and that the drop and its mass dependence persist even if the bias is estimated as $b(r) = \xi_h(r)/\xi_{hm}(r)$, where ξ_{hm} is the halo-matter cross-correlation function: this suggests a physical origin for this feature, ruling out the possibility of being an effect related to a particular choice of the bias estimator. The authors ultimately associate this particular behavior of the BAO-scale bias with the damping of the BAO signal, a phenomenon related to non-linear evolution and due to the motions of matter or tracers relative to their initial separation (Eisenstein et al. 2007), which has a larger extent in the matter field or in the distribution of tracers with a smaller minimum mass. However, the forecasted errors for the chosen volume make the large-scale halo biases for the different simulations indistinguishable from each other and from the hypothesis of constant bias.

We have also computed the effective bias for the halos in our simulations, by weighting the Tinker bias at a given mass scale $b_{\mathrm{Tinker}}(M, z)$ (Eq. 1.5.20) with the average HMF

$\bar{n}(M,z)$ we have measured from our set of simulations, i.e., by computing the quantity

$$b_{\rm eff}(z) = \frac{\int_{M_{\rm min}}^{M_{\rm max}} \bar{n}(M,z)\, b_{\rm Tinker}(\nu(M,z))\, dM}{\int_{M_{\rm min}}^{M_{\rm max}} \bar{n}(M,z)\, dM}, \tag{3.1.2}$$

In measuring the HMFs, we have applied the relation given in Warren et al. (2006) to correct for systematics affecting the masses of halos identified via the FoF method, which depend on the number N_p of particles within a given halo. This correction consists of multiplying the mass $M_h = N_p m_p$ of each halo (with m_p the mass of a CDM particle) by the factor $1 - N_p^{-0.6}$. The HMFs for the different neutrino masses at fixed $z = 0$, and at the different redshifts for $M_\nu = 0.4\,\mathrm{eV}$, are shown in Fig. 3.4. While the HMFs show an expected increase in the abundance of high-mass halos for decreasing redshift, at a fixed redshift and varying M_ν, their differences are limited.

We generate the function $b_{\rm Tinker}(M,z)$ for the different cosmologies of our simulations through the `CosmoBolognaLib`[1] (CBL; Marulli et al. 2016), a large set of free-software `C++` libraries for cosmological calculations. In particular, from the CBL we find that the values of A_s for cosmologies with the same parameters of the `Mnu_ppp`, `Mnu_pp`, `Mnu_p` and `fiducial_ZA` simulations are, respectively, $10^9 A_s = 2.72$, 2.39, 2.24 and 2.12 (evaluated at a scalar pivot $k_p = 0.05\,h\,\mathrm{Mpc}^{-1}$). Larger values of A_s are expected to produce smaller biases, given that the variance $\sigma^2(M,z)$ increases for increasing A_s, and $b_{\rm Tinker}$ is an increasing function of the peak height $\nu(M,z) \propto \sigma(M,z)^{-1}$. We adopt the cold dark matter prescription for the computation of $\sigma(M,z)$ (see Sec. 1.5), and set $\Delta = 200$.

We show the obtained values of effective bias with horizontal dotted lines in Fig. 3.3. The values of $b_{\rm eff}(z)$ agree well with the measured bias on large scales. The relative differences between their maximum and minimum values with respect to the fiducial cosmology are within $\sim 3\%, 4\%, 4.5\%$, at $z = 0, 1, 2$, respectively.

One could argue that we have used a bias relation calibrated for halos identified with spherical overdensity (SO) group-finding algorithms (Lacey and Cole 1994), whereas in the Quijote simulations, halos are identified using FoF. However, as commented in Tinker et al. (2010), it is well known that the overdensity of FoF halos with a linking length of 0.2 (as the Quijote halos) is close to 200. Furthermore, when applying other bias models available in the CBL, i.e., the ones proposed in Sheth and Tormen (1999) and Sheth et al. (2001), we obtain worse results regardless of whether we adopt the matter prescription or the cold dark matter prescription.

We also show in Fig. 3.5 the ratio between the redshift-space and real-space monopoles as a function of the pair separation r, for varying neutrino mass and at all the considered redshifts. We plot scales below $110\,h^{-1}\,\mathrm{Mpc}$ to avoid the noisy behavior associated with the zero crossing of the real-space monopole. The ratios at fixed z are compatible with each other, following very similar trends from small to large scales. This can be understood from Eq. 1.4.16, at least on sufficiently large scales where this relation holds, based on what we have already shown about the weak dependence of bias on the neutrino mass at fixed z and σ_8, and considering that the suppression induced by massive neutrinos in the linear growth rate f with respect to the massless case (Eq. 1.5.10) is $\lesssim 2\%$ even for our

---

[1] `https://github.com/federicomarulli/CosmoBolognaLib`

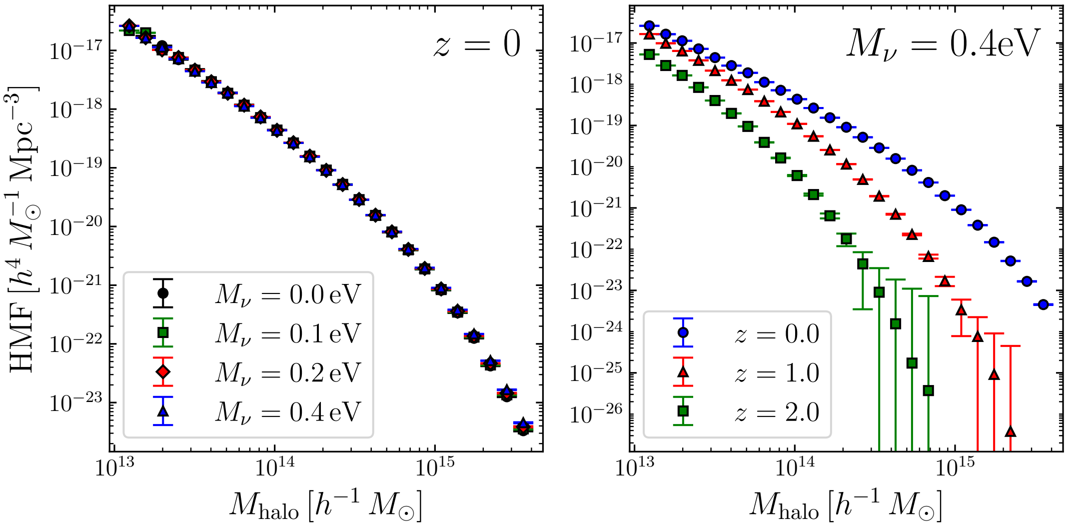

Figure 3.4: The averaged halo mass functions for the simulations with total neutrino mass $M_\nu = 0.4, 0.2, 0.1\,\mathrm{eV}$ (blue triangles, red diamonds, and green squares, respectively) and `fiducial_ZA` ones ($M_\nu = 0\,\mathrm{eV}$, black circles), at $z = 0$ (*left*), and at $z = 0, 1, 2$ (blue circles, red triangles, and green squares, respectively) for fixed $M_\nu = 0.4\,\mathrm{eV}$ (*right*). The error bars show the standard deviation of the mean in each mass bin.

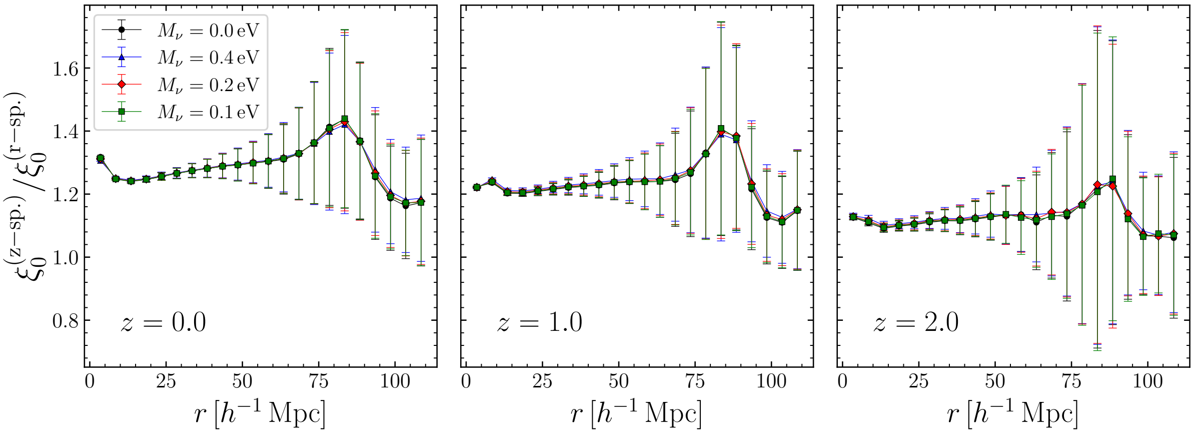

Figure 3.5: The ratio between the average redshift-space and real-space monopoles of the 2PCF as a function of the pair separation r, rebinned to $\Delta r = 5\,h^{-1}\,\mathrm{Mpc}$, for varying neutrino mass, at $z = 0, 1$, and 2 (from left to right). Data are shown for $r < 110\,h^{-1}\,\mathrm{Mpc}$ to avoid the scales characterized by the zero crossing of the denominator.

largest mass $M_\nu = 0.4\,\mathrm{eV}$. The ratios also show a weak redshift evolution, consistent with the fact that both f and b increase for increasing redshift, making their ratio β weakly dependent on z.

At fixed z, the quadrupoles show an opposite trend as a function of neutrino masses with respect to the monopoles, decreasing for increasing M_ν, which can be explained considering that in Eq. 1.4.17 the term $\xi(r) - \bar{\xi}(r)$ must be negative and the amplitude of $\xi(r)$ increases for increasing M_ν. Additionally, the quadrupole shows a transition from positive to negative slope on scales $\lesssim 40\,h^{-1}\,\mathrm{Mpc}$ for fixed M_ν and decreasing redshift. The large-scale amplitude of the halo quadrupoles shows a nonmonotonic relation with redshift, which can be accounted for by the different time evolution of the growth factor D_1, the growth rate f, and halo bias b. In the linear regime, the amplitude goes as $D_1^2 b^2(4\beta/3 + 4\beta^2/7)$: while both b and f increase with z, D_1 decreases with it.

The measured hexadecapoles appear insensitive to the neutrino mass at fixed redshift. On scales $\lesssim 30\,\mathrm{Mpc}$, they show a dependence on redshift, with a progressive flattening going to smaller z. For larger scales at fixed M_ν, the hexadecapoles rapidly flatten to 0.

For the considered volume, all the neutrino masses are poorly detectable in the 2PCF quadrupoles and hexadecapoles at all redshifts. Even considering scales of $\lesssim 10\,h^{-1}\,\mathrm{Mpc}$, the detectability does not exceed 2, and rapidly vanishes for increasing scale.

3.2 Neutrino signal in the isotropic 3-point correlation function multipoles

We performed a detectability analysis (Eq. 2.5.1) to compare the 3PCF isotropic multipoles measured from the massive neutrino simulations with the measurements from the `fiducial_ZA` mocks. To avoid excessive redundancy, we do not show the outcomes of this analysis for all the possible combinations of spaces, redshifts, and neutrino masses, but we choose as a representative case the results obtained in redshift space, at $z = 0$, and for $M_\nu = 0.4\,\mathrm{eV}$.

We show in Fig. 3.6, for illustrative purposes, the average 3PCF isotropic multipoles measured from the set of `fiducial_ZA` mock catalogues, while presenting in Fig. 3.7 the detectabilities for the selected case. In both figures, we display all the information encoded in all the scales probed in our analysis through a 2-D colormap for each ℓ, where a given pixel corresponds to a binned side pair (s_{12}, s_{13}), and the color scale represents, respectively in Fig. 3.6 and 3.7, the values of the corresponding average multipole and detectability.

Fig. 3.6 shows what we have already noticed in Sec. 2.4 regarding the drop in efficiency of the SHD estimator in the reconstruction of the shape of the 3PCF for isosceles and quasi-isosceles configurations. It is clearly noticeable that for increasing ℓ the values of the multipoles in the off-diagonal regions rapidly approach zero for $\ell > 6$ (see, for example, the lower-right and, by symmetry, the top-left corner of the colormaps), while on the diagonal they remain positive on scales below $s_{\mathrm{cross}} \approx 60\,h^{-1}\,\mathrm{Mpc}$ and negative above (with s_{cross} slightly increasing with ℓ), without significant differences approaching $\ell_{\max} = 10$. Notice that along the diagonal and above s_{cross}, the multipoles show a local maximum around

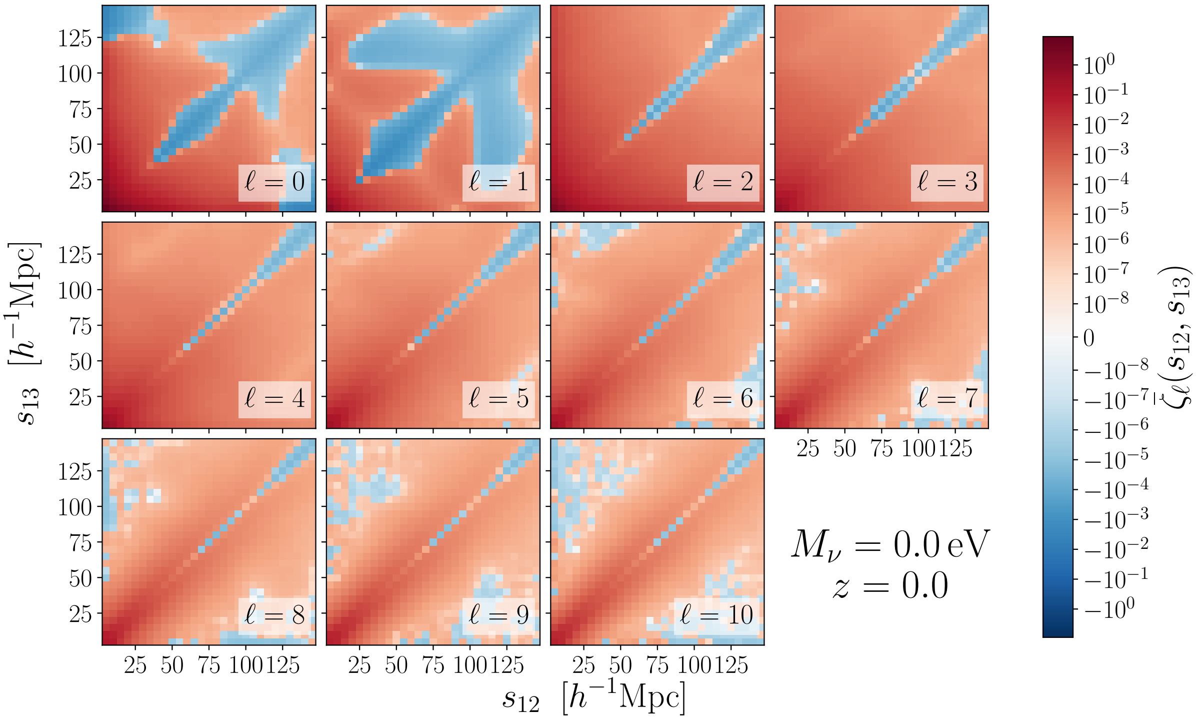

Figure 3.6: The measured redshift-space multipoles ζ_ℓ of the isotropic 3PCF as a function of the two triangle sides s_{12} and s_{13}, at $z = 0$, for the `fiducial_ZA` mocks, averaged over the 500 realizations. Each pixel of the colormaps corresponds to a (s_{12}, s_{13}) bin with size $5 \times 5\,h^{-1}\,\mathrm{Mpc}$.

$s \approx 100\,h^{-1}\,\mathrm{Mpc}$ associated with the BAO peak.

The fact that neutrinos mostly affect small, nonlinear scales as described in Sec. 1.5 is evident by inspecting the colormaps in Fig. 3.7, where the highest, positive values of detectability are mostly concentrated on small scales for all the considered ℓ. For $\ell = 0$, whose corresponding monopole ζ_0 controls the offset of ζ, a positive detectability is present also for squeezed triangles (corresponding to the regions adjacent to the bottom and, by symmetry with respect to the diagonal, left sides of the colormap), decreasing for increasing scale.

The monopole $\ell = 0$ and the dipole $\ell = 1$ exhibit a mildly negative detectability (for which, in both cases $-1 < \mathrm{DET} < 0$) on a large region of the probed (s_{12}, s_{13}) area. For $\ell = 0$, this feature is mostly located along the diagonal, divided into two features with $50 \lesssim s_{12}, s_{13}\,[h^{-1}\,\mathrm{Mpc}] \lesssim 95$ and $s_{12}, s_{13}\,[h^{-1}\,\mathrm{Mpc}] \gtrsim 105$, separated on BAO scales. For $\ell = 1$, it shows a fork-shaped pattern, with three features located on the diagonal on scales $s_{12}, s_{13} \gtrsim 40\,h^{-1}\,\mathrm{Mpc}$, and on scales $s_{12}\,(s_{13}) \gtrsim 30\,h^{-1}\,\mathrm{Mpc}$ when $s_{13}\,(s_{12}) \sim 125\,h^{-1}\,\mathrm{Mpc}$.

The multipole that, in the presence of massive neutrinos, shows the most significant departures from the fiducial cosmology is the quadrupole ζ_2, which imprints the characteristic parabolic shape to ζ as a function of the angle between two fixed sides, as shown in Sec. 3.3. For ζ_2, on scales $s_{12}, s_{13} \lesssim 25\,h^{-1}\,\mathrm{Mpc}$ the detectability for $M_\nu = 0.4\,\mathrm{eV}$ is around 8 and for squeezed configurations it remains $\gtrsim 2$ even for the largest scales probed in our measurement, vanishing only for triangles with large s_{12} and s_{13}. Starting

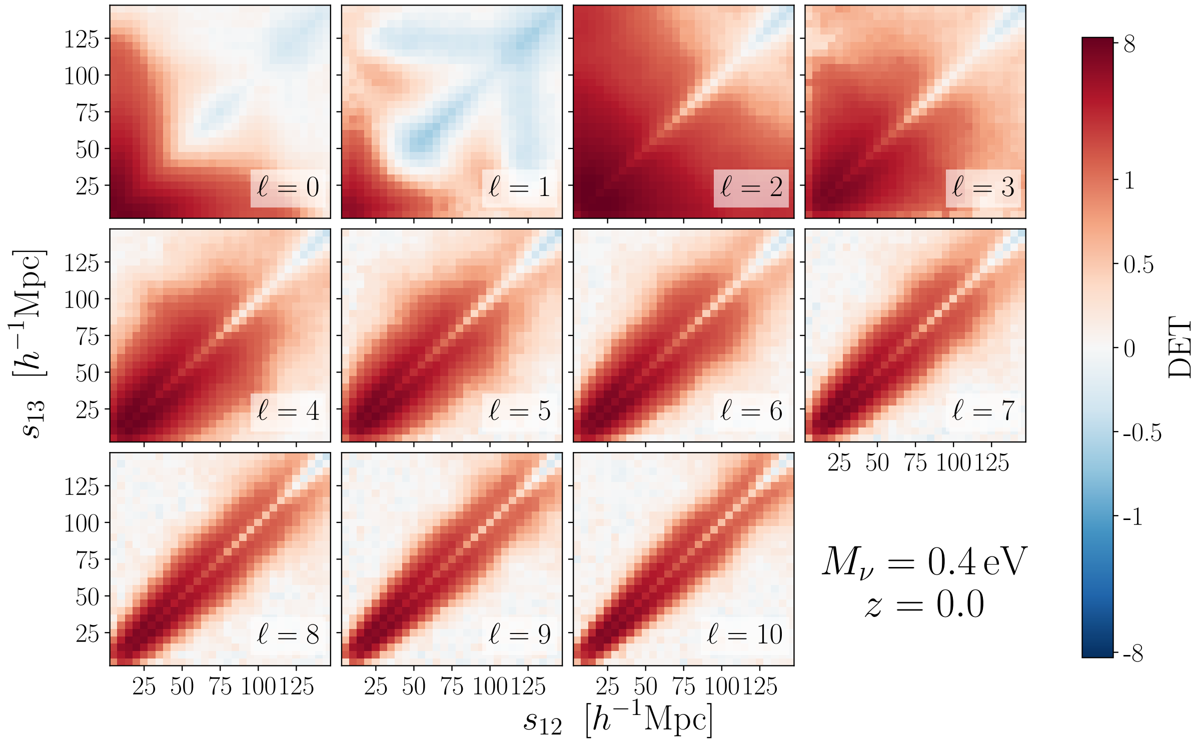

Figure 3.7: The detectabilities, computed with Eq. 2.5.1, in the case of total neutrino mass $M_\nu = 0.4\,\mathrm{eV}$, for the redshift-space multipoles ζ_ℓ of the 3PCF at $z = 0$. The scale on the colorbar is logarithmic.

from $\ell = 3$, the highest-detectability pattern progressively concentrates on quasi-isosceles configurations, decreasing with increasing s_{12} and s_{13}.

Interestingly, for $\eta = 0$ and $\ell \geq 2$, the detectability patterns show a local minimum. We find that this drop is driven by the presence of a local maximum in the profile of the standard deviations of all multipoles. Although this maximum is also present for $\ell = 0$ and $\ell = 1$, the complex detectability pattern associated with these two multipoles prevents it from being apparent in the plots of Fig. 3.7. We suggest that this feature originates from autocorrelation in the same radial bin: in this case, the expected variance is larger than in the case of cross-correlation between two different radial bins.

3.2.1 Neutrino signal on baryon acoustic oscillations

We also investigate the BAO signal encoded in the 3PCF multipoles, which is not confined to a single, well-defined feature, at odds with the distinctive peak in the 2PCF, but rather spread across several regions of the (s_{12}, s_{13}) plane. This is due to the fact that the multipoles depend on two scales, namely the two sides s_{12} and s_{13}, and, at the same time, encode the dependence of ζ on the third side s_{23}. As a consequence, each of the three sides can, in principle, contribute to the BAO signal.

The positions of several BAO features in the monopole, dipole, and quadrupole of the redshift-space 3PCF in ΛCDM cosmologies have been predicted in the perturbative

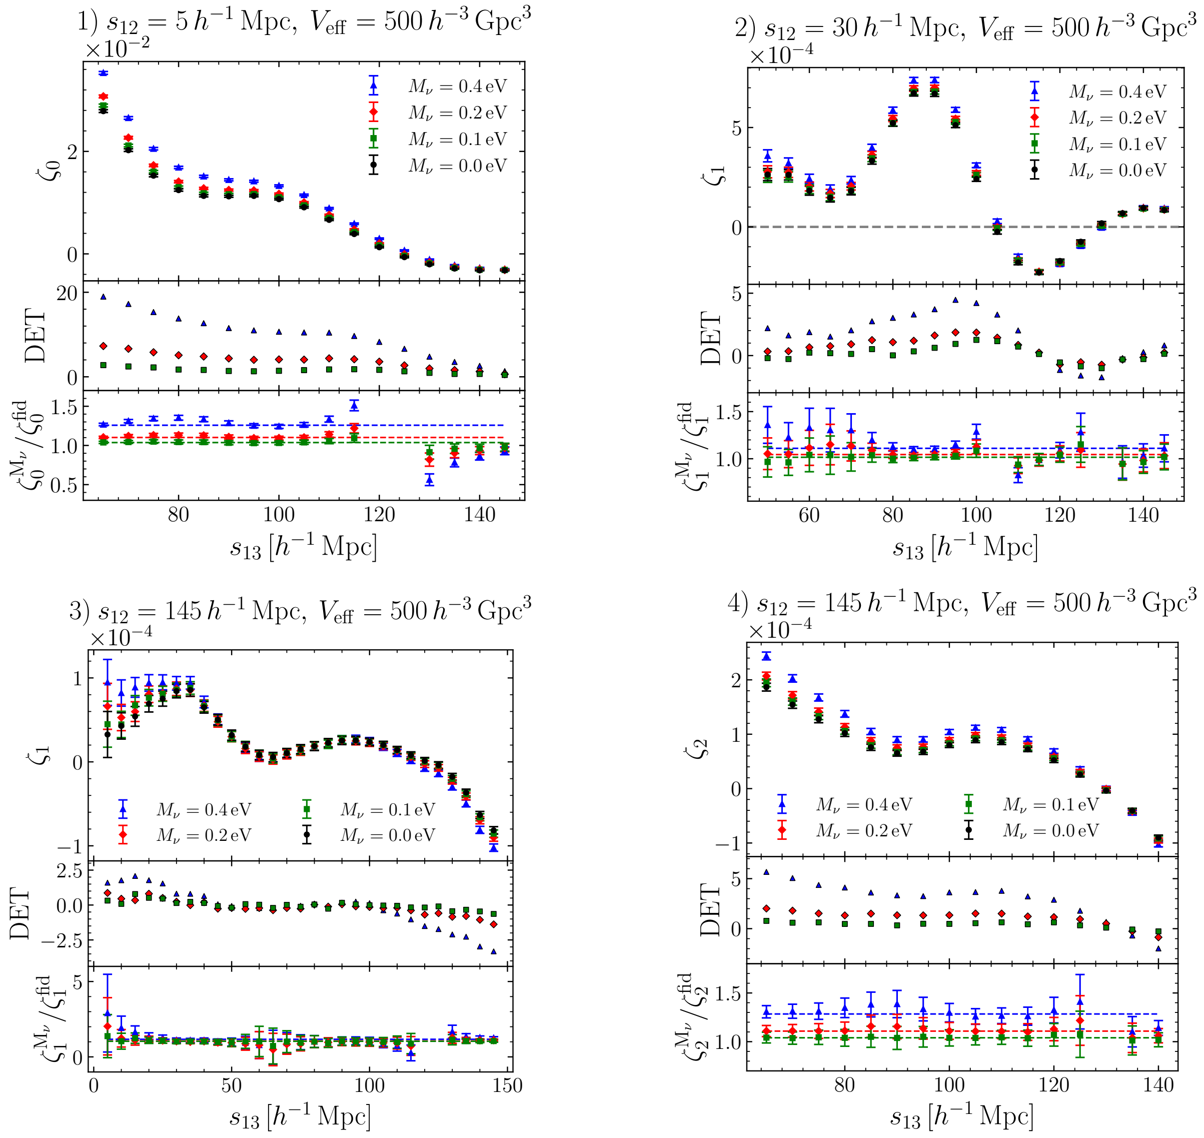

Figure 3.8: The average redshift-space, $z = 0$ 3PCF monopole ζ_0 (*upper left*), dipole ζ_1 (*upper right and lower left*), and quadrupole ζ_2 (*lower right*) for each of the considered neutrino masses, at the scales of some of the BAO features predicted in Slepian and Eisenstein (2017). The values of the multipoles are shown with markers in the upper plot of each panel as a function of s_{13}, with s_{12} fixed to the value indicated above it. Errors are estimated on the full effective volume $V_{\mathrm{eff}} = 500\,h^{-3}\,\mathrm{Gpc}^3$. A dashed gray line marks the dipole zero crossing in the upper-right panel, occurring at the BAO position. The middle and bottom plots of each panel show, respectively, the corresponding detectabilities and the ratios between the massive neutrino multipoles and the fiducial one. The ratios are fitted with a constant, and the best fits are reported with dashed lines. Bins associated with zero crossings of the fiducial multipoles, where ratios are unstable, have been excluded from the fit (and are not plotted). All the shown quantities follow the same color code as a function of M_ν.

model of Slepian and Eisenstein (2017) up to very large scales ($\sim 200\,h^{-1}\,\mathrm{Mpc}$). It is particularly interesting to detect them in our data and explore their possible dependence on M_ν.

We focus on the following predicted features, which we report as they are described in the cited work (we refer the reader to Sec. 4.3 and Fig. 8 of the cited paper for further details):

1. For the monopole ζ_0, an increase driven by the density field, when one of the two sides crosses BAO scales ($\sim 100\,h^{-1}\,\mathrm{Mpc}$) while the other remains on smaller scales.

2. For the dipole ζ_1, an increase below the BAO scale and a decrease slightly above it (corresponding, in the blue fork-shaped area in Fig. 3.6 and Fig. 3.7 at $\ell = 1$, to the inner edge of the two most external regions), with a zero crossing at BAO, sourced by the gradient of the halo density field (which therefore vanishes at the BAO position).

3. A series of local maxima in ζ_1 placed slightly below the line $s_{13} \approx s_{12} - 100\,h^{-1}\,\mathrm{Mpc}$ connecting the point $(s_{12}, s_{13}) \approx (100, 0)\,h^{-1}\,\mathrm{Mpc}$ to $(s_{12}, s_{13}) \sim (200, 100)\,h^{-1}\,\mathrm{Mpc}$ (and similarly with s_{12} and s_{13} switched). This is a feature arising from the fact that, when the angle between s_{12} and s_{13} goes to 0, the remaining side $s_{23} = |s_{13} - s_{12}|$ is on BAO scales.

4. For the quadrupole ζ_2, a bump when one of the two sides is above BAO scales and the other crosses them.

For their identification, we fix the side s_{12} to some values that allow s_{13}, free to vary, to cross them. The results are shown in Fig. 3.8, where each panel corresponds to one of the features described above, numbered in the same way as they are listed. The value at which we fix s_{12} is shown above each panel. Notice that the large scales we consider yield mutually consistent results across all neutrino masses for the volume of $V = 10\,h^{-3}\,\mathrm{Gpc}^3$, making any possible difference in the multipoles with M_ν undetectable. Therefore, in Fig. 3.8, we assign to the average multipoles the errors estimated on our total effective volume $V_{\mathrm{eff}} = 500\,h^{-3}\,\mathrm{Gpc}^3$.

All features discussed above are evident in the figure, with the exception of $M_\nu = 0.4\,\mathrm{eV}$ in case 3). In this case, whereas for the other masses a maximum at $s_{13} \approx 40\,h^{-1}\,\mathrm{Mpc}$ can be clearly identified, as expected from the choice $s_{12} = 145\,h^{-1}\,\mathrm{Mpc}$, the uncertainties on ζ_1 for this mass do not allow us to rule out a constant behavior. In case 2), the data suggest a shift towards larger scales of the dipole zero-crossing (and hence of the BAO location), at least for the highest mass. Indeed, for $s_{13} = 105\,h^{-1}\,\mathrm{Mpc}$, while the fiducial dipole lies below zero by 2.5σ, $\zeta_1(M_\nu = 0.4\,\mathrm{eV})$ is above zero at approximately 3σ. This conclusion is also consistent with the trend observed in the detectability as a function of s_{23}, which qualitatively reproduces the dipole shape with a forward shift.

Along with the average multipoles and the detectabilities, for each M_ν we also show the ratios between the massive-neutrino and fiducial multipoles, in order to investigate whether any shape variation of the multipoles with M_ν can be distinguished from a simple linear bias rescaling. To this end, we fit these ratios with a constant by adopting the least-squares algorithm. The best-fit horizontal lines are shown in the Fig. 3.8, while the corresponding reduced χ^2 values are reported in Table 3.1. We do not fit the points close to any zero crossings of the fiducial multipoles, where the ratios become unstable, and do not show them in the figure.

Panel	$M_\nu = 0.1\,\text{eV}$		$M_\nu = 0.2\,\text{eV}$		$M_\nu = 0.4\,\text{eV}$	
	Best-fit ratio	χ^2/ndof	Best-fit ratio	χ^2/ndof	Best-fit ratio	χ^2/ndof
1	1.037(7)	0.66	1.102(7)	3.60	1.258(8)	22.91
2	1.016(14)	0.22	1.044(14)	0.37	1.110(15)	1.86
3	1.036(38)	0.05	1.054(38)	0.17	1.168(41)	0.66
4	1.040(22)	0.02	1.109(22)	0.09	1.284(24)	0.54

Table 3.1: The best-fit values (and corresponding 1σ uncertainties) for the constant with which the ratios in Fig. 3.8 have been fitted, together with the minimum reduced chi-square at the fit. The values are reported for a different value of M_ν in each distinct column. Each row corresponds to a panel of Fig. 3.8, numbered as in the figure.

The values of reduced chi-square for $M_\nu = 0.4$ and $0.2\,\text{eV}$ clearly indicate a change in the shape of the 3PCF monopole imprinted by massive neutrinos, similarly to the 2PCF case discussed in Sec. 3.1. Unfortunately, the uncertainties prevent a distinction between an analogue effect and a simple rescaling in the other cases, even though, at least for case 2) with $M_\nu = 0.4\,\text{eV}$, data systematically lying above or below the best-fit values across most bins point in the direction of a nontrivial change in shape compared to the fiducial case.

3.3 Neutrino signal in the isotropic connected and reduced 3-point correlation function: single scale

To study how massive neutrinos impact clustering for triangular configurations with two fixed sides r_{12} and r_{13}, depending on the angle θ between them, we reconstructed the single-scale 3PCFs for each possible side pair (r_{12}, r_{13}) by applying Eq. 2.4.1 to the averages of the estimated multipoles. We also obtained the corresponding reduced 3PCF Q by dividing by the cyclic permutation of the 2PCF monopole evaluated at the sides of the triangle, as detailed in Eq. 2.4.6.

The high number of independent combinations of side pairs (r_{12}, r_{13}) required the development of an efficient procedure to determine the values of the triangle sides at which massive neutrinos mostly impact the angular behavior of the 3PCF. Indeed, the number of single-scale functions which can be computed at fixed space, redshift and M_ν is 435 for ζ and 121 for Q[2]. This number is lower for Q because, to ensure that all triangle

---

[2] Since the total number of radial bins $N_\text{bins} = 29$ and $f(r_{12}, r_{13}, \theta)$, with $f = \zeta$ or Q, is symmetric under

sides remain below $110\,h^{-1}$ Mpc for the reasons discussed in Sec. 2.4, we selected only side pairs (r_{12}, r_{13}) for which $r_{12}+r_{13} \leq 110\,h^{-1}$ Mpc ($r_{12}+r_{13}$ is the maximum length of the third side r_{23} of a triangle whose other two sides are r_{12} and r_{13}).

We adopted the strategy to compute, for each space, redshift and M_ν, the chi-square indicator defined in Eq. 2.5.2, normalizing it by the number of chosen angular samples N_θ to form a reduced chi-square, and then inspect the single-scale $\zeta(r_{12}, r_{13}, \theta)$ and $Q(r_{12}, r_{13}, \theta)$ for some selected pairs (r_{12}, r_{13}).

The number of angular samples N_θ in Eq. 2.5.2 cannot be arbitrary, since the numerically estimated covariance matrix $\hat{C}^{(\mathrm{single-sc.})}$ is invertible if N_θ does not exceed the number of multipoles, equal to 11 in our analysis. So we applied the reconstruction in Eqs. 2.4.1 and 2.4.6 by choosing 10 evenly spaced samples in the angular range $[0, \pi]$.

For small values of η and $\theta = 0$, the length of the third side $r_{23} = |r_{12} - r_{13}|$ can assume values below the minimum radial bin in which we measured the 2PCF, $r_{\mathrm{min}} = 1\,h^{-1}$ Mpc: to avoid extrapolating the behavior of ξ_0 below r_{min}, in reconstructing Q we discarded the angle $\theta = 0$ whenever r_{23} dropped below r_{min}.

The reduced chi-square indicator $\tilde{\chi}^2$ for ζ and Q, in redshift space and for $M_\nu = 0.4$ eV, is shown in Fig. 3.9, through a colormap for each redshift $z = 0, 1, 2$.

For ζ, at $z = 2$ the maximum values are localized along the diagonal on scales $\lesssim 30\,h^{-1}$ Mpc. At lower redshifts, quasi-isosceles and squeezed configurations acquire relevance, with the highest values of $\tilde{\chi}^2$ below $\sim 30\,h^{-1}$ Mpc at $z = 0$, progressively decreasing on larger scales. Such a behavior is expected for the detectability of the neutrino signal in the 3PCF; the phenomenon of free-streaming affects small scales, which are also those most strongly impacted by the onset of nonlinearity. In addition, the impact of nonlinearities at a given scale grows as the redshift decreases. Since nonlinearity breaks the Gaussianity of perturbations, it naturally generates a signal in the 3PCF. This picture is consistent with the overall increase of $\tilde{\chi}^2$ values as z decreases, and with the configurations that maximize it at fixed redshift, since isosceles, quasi-isosceles, and squeezed triangles are precisely those for which at least one of the three sides can probe nonlinear or mildly nonlinear scales. The effect of the increase in the errors on the measured multipoles commented at the end of Sec. 3.2 is visible along the diagonal, where a local minimum is evident above $\sim 30\,h^{-1}$ Mpc.

Compared with ζ, the values of $\tilde{\chi}^2$ for Q are lower at fixed (s_{12}, s_{13}), due to the larger errors from the propagation described in Sec. 2.4. At all redshifts, the largest values are localized along the diagonal. At $z = 0$, they decrease moving to larger scales, while for $z = 1, 2$ they show a maximum for $s_{12} = s_{13} \approx 20\,h^{-1}$ Mpc.

The described trends are also present in real space and for the remaining neutrino masses $M_\nu = 0.1$ and 0.2 eV.

We cut the maps for ζ in Fig. 3.9 along some lines representative of the configurations for which we observe an enhanced signal. These lines are overplotted on the leftmost colormaps, and are described by the equations $\eta = 0$ (isosceles configurations), $\eta = 4$

---

exchange of r_{12} and r_{13}, the total number of independent single-scale 3PCFs which can be reconstructed at fixed space, z and M_ν is given by $N_{\mathrm{bins}} \times (N_{\mathrm{bins}} + 1)/2 = 435$. For the single-scale reduced 3PCF, the adoption of the constraint $r_{12} + r_{13} \leq 110\,h^{-1}$ Mpc reduces this number to $[(N_{\mathrm{bins,\,red}} + 1)/2]^2 = 121$, where $N_{\mathrm{bins,\,red}} = 21$ is the number of radial bins with centers from 5 to 105 h^{-1} Mpc.

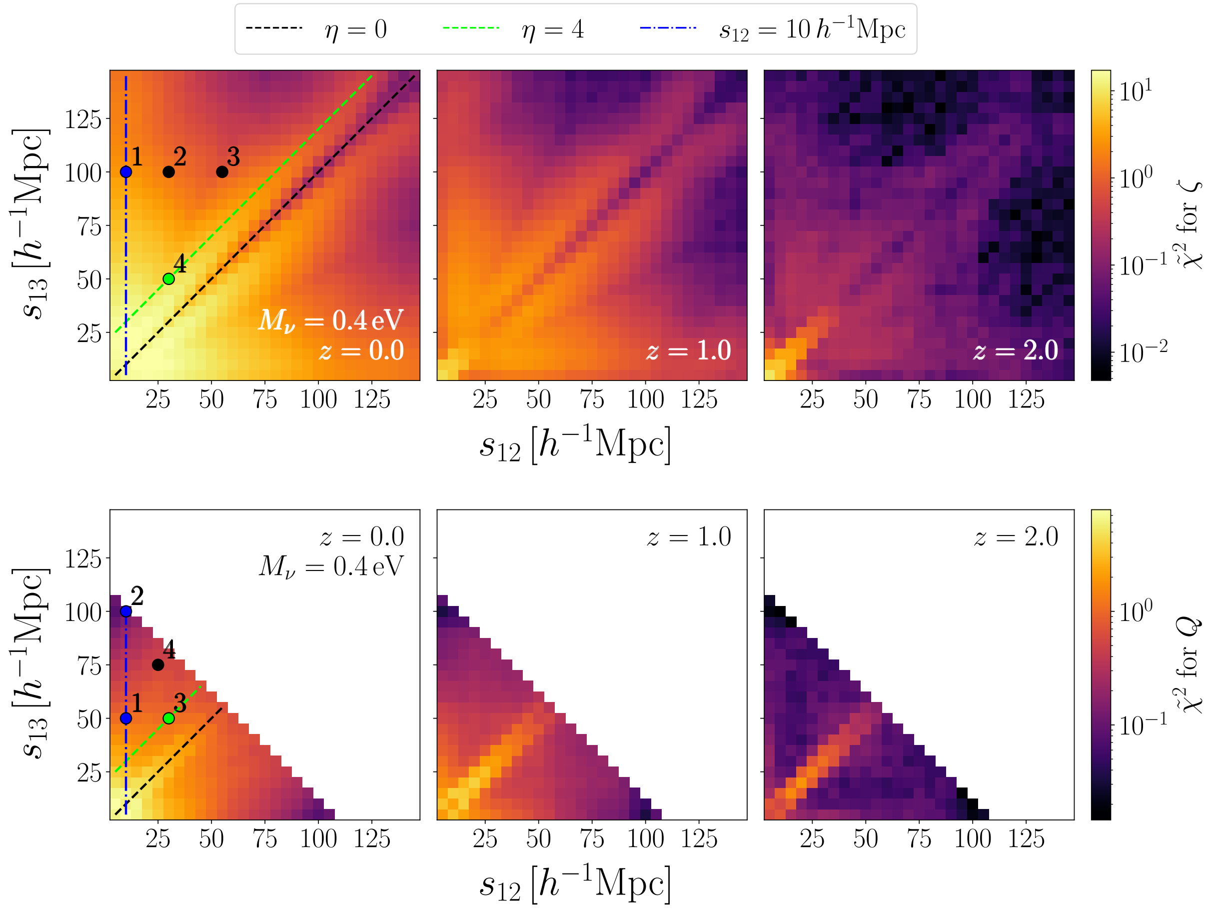

Figure 3.9: *Top panel*: the colormaps show the reduced chi-square indicator $\tilde{\chi}^2(s_{12}, s_{13})$ for the redshift-space connected 3PCF as a function of the triangle sides s_{12} and s_{13}, computed by dividing the chi-square indicator $\chi^2(s_{12}, s_{13})$ in Eq. 2.5.2 by the number of angular samples. The results are shown for $M_\nu = 0.4\,$eV, and each colormap refers to a different redshift: from left to right, $z = 0, 1, 2$. The lines overplotted on the leftmost colormap identify isosceles and quasi-isosceles configurations satisfying $\eta = 0, 4$ and squeezed configurations with $s_{12} = 10\,h^{-1}\,$Mpc (black and green dashed lines, and blue dash-dotted line, respectively). The numbered points on the leftmost panels correspond to the scales selected for inspection in Figs. 3.10 and 3.11. *Bottom panel*: same as in the top panel, but for the reduced 3PCF Q. White regions correspond to scales excluded by the constraint $s_{12} + s_{13} \leq 110\,h^{-1}\,$Mpc.

(representative of quasi-isosceles configurations) and $s_{12} = 10\,h^{-1}\,$Mpc (representative of squeezed triangles); for consistency, we adopted the same cuts also for Q. We selected some side pairs on these lines and off them (for comparison), corresponding to the numbered points in the figure, namely $(10, 100), (30, 100), (55, 100), (30, 50)\,h^{-1}\,$Mpc for ζ and $(10, 50), (10, 100), (30, 50), (25, 75)\,h^{-1}\,$Mpc for Q. These points are selected to directly inspect which configurations, among all the possible ones given by varying θ, drive the enhancement of the signal. Furthermore, the first three pairs chosen for ζ and the first one for Q are such that the third side crosses BAO scales, allowing us to possibly detect its signal.

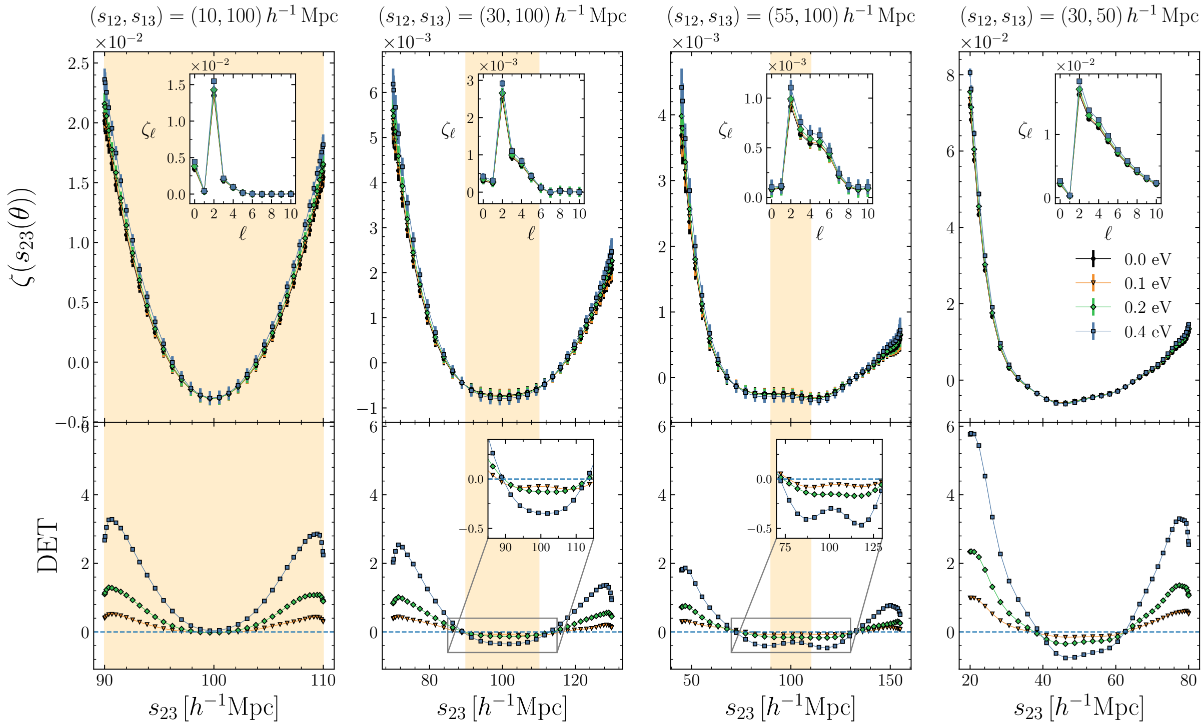

Figure 3.10: The redshift-space single-scale 3PCF for the (s_{12}, s_{13}) pairs selected in Fig. 3.9 (upper plots of each panel), and indicated at the top, as a function of the length of the third side s_{23}. Here, ζ is plotted at $z = 0$, with a different color for each of the considered neutrino masses, as shown in the legend. The inset plots in the upper row show the multipoles $\zeta_\ell(s_{12}, s_{13})$ from which the functions have been computed. The lower plots show the corresponding detectabilities as a function of s_{23} (Eq. 2.5.1). A blue dashed line marks the zero detectability level. Both the multipoles and detectabilities follow the same color coding of ζ. The orange shaded areas show the region $90 \leq s_{23}\ [h^{-1}\ \mathrm{Mpc}] \leq 110$, corresponding to the BAO location. Where the BAO scales do not cover the full s_{23} range, we include a zoom on the detectabilities to improve visualization in the BAO region.

The single-scale ζ and Q associated with the numbered points of Fig. 3.9 are shown in Fig. 3.10 and 3.11 for varying M_ν, alongside their detectabilities. To make the range of scales spanned by the third side s_{23} immediately clear, especially for visualizing the BAO signal, we show the data directly as a function of this side, converting the angle θ through the cosine rule. We also increase the number of points with respect to those used to compute the $\tilde{\chi}^2$ parameter in Fig. 3.9 to improve visualization. For a given single-scale $\zeta(s_{12}, s_{13}; s_{23}(\theta))$, also the corresponding multipoles $\zeta_\ell(s_{12}, s_{13})$ are displayed. The corresponding correlation matrices are shown in App. B (Fig. B.1). We did not select isosceles configurations to be shown in the figure because the SHD estimator, as expected, provides an inaccurate angular reconstruction, due to the fact that multipoles higher than $\ell = 10$ still retain a large fraction of the signal. For the selected (s_{12}, s_{13}) pairs, instead, ζ (and, consequently, Q) is optimally reconstructed since the value of the multipoles with $\ell > \ell_{\max}$ is negligible.

The connected 3PCFs shown in the figure span two orders of magnitude, decreasing

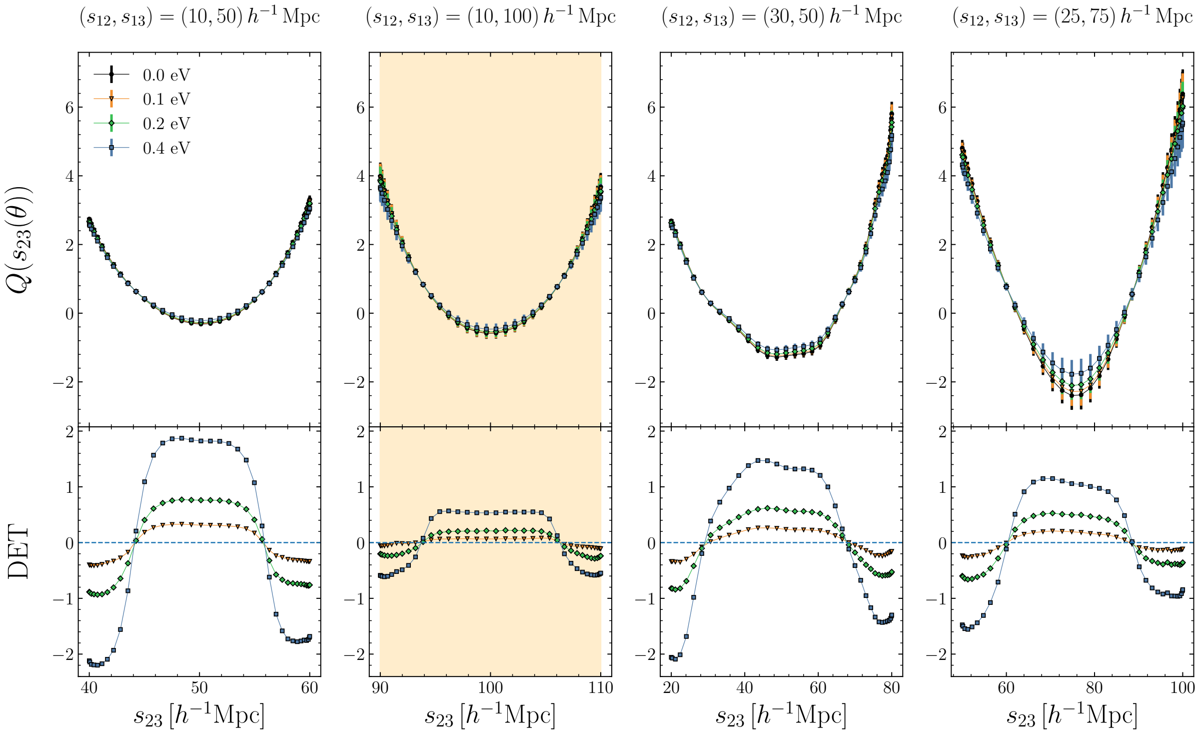

Figure 3.11: The redshift-space single-scale reduced 3PCF (upper plot of each panel), for the (s_{12}, s_{13}) pairs selected in Fig. 3.9 and indicated at the top, as a function of the third side s_{23}. Q is plotted at $z = 0$, with a different color for each of the considered neutrino masses, as shown in the legend. The lower plots show the corresponding detectabilities as a function of s_{23} (Eq. 2.5.1), with the same color coding as Q. A blue dashed line marks the zero detectability level. The orange shaded area shows the region $90 \leq s_{23}\,[h^{-1}\,\mathrm{Mpc}] \leq 110$, corresponding to the location of the BAO peak.

as the scale increases, while, as expected, Q is always of order unity. Moreover, as already noticed in Moresco et al. (2014), the shape of Q, higher for elongated triangles than for right-angled ones, is given to a well-understood effect (Gaztañaga and Scoccimarro 2005; Marín et al. 2008) due to the fact that cosmic structures preferentially move along gradients of the density field within nonlinear gravitational instabilities. Furthermore, going from smaller to larger scales, it exhibits a transition from a *U-shape* to a *V-shape* (consistently with the already cited works), as a consequence of the fact that, for progressively larger scales, halos reside preferentially in more filamentary structures.

A notable effect for varying M_ν is observed in the concavity of the functions, with opposite trends in ζ and Q as M_ν increases. Specifically, concavity increases for the former for all the considered configurations, and decreases for the latter. This produces a U-shaped trend in the detectability of ζ, which is rather symmetric for squeezed triangles, where the range of scales spanned by s_{23} is quite limited (for instance, in the case of $s_{12} = 10\,h^{-1}$ Mpc, it varies in the range $s_{13} \pm 10\,h^{-1}$ Mpc), and more asymmetric for $\eta = 4$. In the latter case, the detectability at $s_{23} = |s_{13} - s_{12}|$ (i.e., $\theta = 0$) is higher than in the opposite case, $s_{23} = s_{12} + s_{13}$ ($\theta = \pi$). The small decrease in detectability at the extremes

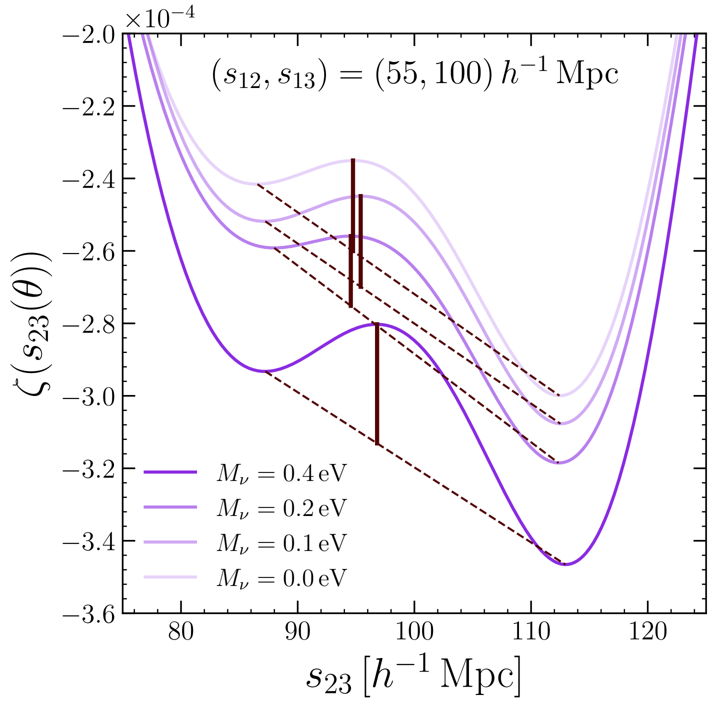

Figure 3.12: Close-up of the BAO region in the redshift-space single-scale 3PCF for $(s_{12}, s_{13}) = (55, 100)\, h^{-1}\,\mathrm{Mpc}$, at $z = 0$, plotted against the third, variable side s_{23}. The neutrino mass increases going from the lightest to the darkest curve. The figure also shows the vertical segment joining the local maximum corresponding to the BAO peak to the line interpolating between its two adjacent minima. The length of this segment corresponds to the *BAO contrast* parameter defined in Moresco et al. (2021).

of the s_{23} intervals in Fig. 3.10 is driven by an increase in the errors. In Q, on the other hand, the flattening induced by increasing M_ν makes the detectabilities most negative at the extremes, and maximizes them for approximately perpendicular configurations ($\theta \sim \pi/2$). Notice that for these intermediate configurations, the detectabilities are quite flat: this enables the extraction of the signal also from them, which are instead more weakly, or not affected at all by the neutrino signal in ζ (e.g., compare the leftmost panels of Fig. 3.11 and 3.10).

In any case, it is evident that the elongated configurations ($s_{23} = |s_{12} - s_{13}|$ and $s_{23} = s_{12}+s_{13}$) drive the signal in terms of $\tilde{\chi}^2$ in ζ, with equal weight for squeezed configurations, and more strongly when $s_{23} = |s_{12} - s_{13}|$ for nearly isosceles ones. This suggests that neutrinos predominantly affect clustering along the filaments of the cosmic web. This conclusion highlights the usefulness of the 3PCF in this analysis, as it is the lowest-order statistic capable of encoding information about shape, unlike the 2PCF, which only captures the amplitude of clustering at a given distance scale between two objects. We focus in detail on the extraction of shape information in Sec. 3.4. Perpendicular configurations also produce some of the signal shown in Fig. 3.9 in the case of Q.

We now comment on the BAO signal in the single-scale approach. The single-scale 3PCF exhibits a small peak at $\sim 100\, h^{-1}\,\mathrm{Mpc}$ when, for increasing θ, the two fixed sides

s_{12}, s_{13} allow the third one s_{23} to cross BAO scales (Gaztañaga et al. 2009). This peak can possibly cancel out with the 3PCF dip, giving ζ a flat shape for $\theta \sim \pi/2$, or resulting in a small peak embedded in the dip. Notice that this interplay can also influence the position of the peak compared to that of 2PCF (see, e.g., Fig. 6 in Guidi et al. 2023). The behavior depends on the choice of the (s_{12}, s_{13}) pairs, since the BAO feature is more prominent if BAO scales in s_{23} are concentrated in a progressively smaller range in the full s_{23} interval $[|s_{12} - s_{13}|, s_{12} + s_{13}]$ (Moresco et al. 2021). This is clearly noticeable by comparing the first three plots from the left in Fig. 3.10, in which the expected location of the BAO feature ($90 \lesssim s_{23}\,[h^{-1}\,\mathrm{Mpc}] \lesssim 110$) is marked with an orange shaded area. For $(s_{12}, s_{13}) = (10, 100)\,h^{-1}\,\mathrm{Mpc}$, the area spans the entire s_{23} interval, so no distinctive feature is observable. The same consideration holds for Q at the same scale in Fig. 3.11. For $(s_{12}, s_{13}) = (30, 100)\,h^{-1}\,\mathrm{Mpc}$, ζ is flatter at the minimum with respect to the previous case, without showing any local maximum, and an extremely shallow peak is visible in the $M_\nu = 0.1\,\mathrm{eV}$ detectabilities at $\sim 100\,h^{-1}\,\mathrm{Mpc}$. Finally, for $(s_{12}, s_{13}) = (55, 100)\,h^{-1}\,\mathrm{Mpc}$, a small peak is visible in both ζ for all the values of M_ν and in the detectabilities of $M_\nu = 0.4$ and $0.1\,\mathrm{eV}$.

A zoom on ζ in the BAO region for $(s_{12}, s_{13}) = (55, 100)\,h^{-1}\,\mathrm{Mpc}$ is shown in Fig. 3.12, where we also illustrate the estimation of the *BAO contrast*, defined in Moresco et al. (2021). It corresponds to the vertical distance between the local maximum at the BAO position and the line interpolating between the two adjacent local minima, and has the advantage of being based only on the measured curves. We follow the same procedure adopted in the cited work and detect the local minima and the maximum with the modules `argrelmin` and `argrelmax` of the `scipy.signal` Python library[3]. Overall, the values of BAO contrast we obtain range from ≈ 1.9 to 3.3×10^{-5}, and are not monotonic with M_ν (for instance, $M_\nu = 0.2\,\mathrm{eV}$ gives a BAO contrast lower than any other mass). However, this non-monotonicity can be simply an effect of the uncertainties on ζ, which, even when considering the full effective volume $V_{\mathrm{eff}} = 500\,h^{-3}\,\mathrm{Gpc}^3$, are $\mathcal{O}(10^{-5})$, making impossible to establish a statistically significant ordering with M_ν. Notice also that a non-monotonic trend with M_ν of the peak amplitude could be intrinsically due to the nature of neutrinos themselves, as shown for the 2PCF through perturbative models in Peloso et al. (2015). The authors show that, while in linear theory the 2PCF BAO peak decreases for increasing neutrino masses, due to the suppression of P_{lin} (for fixed A_s), the effect of nonlinear structure formation goes in the opposite direction, since the peak broadening by large-scale flows is less effective. As a result of this effect, they predict a decrease by $\sim 0.6\%$ and an increase by $\sim 1.2\%$ of the 2PCF BAO peak, respectively for $M_\nu = 0.15\,\mathrm{eV}$ and $M_\nu = 0.3\,\mathrm{eV}$, when fixing the other cosmological parameters.

We quantitatively assessed the significance of the deviation from the massless case of the single-scale ζ and Q in massive neutrino cosmologies, a piece of information that can be extracted from the χ^2 indicator through the computation of the p-values in Eq. 2.5.3. We computed the p-values for every side pair falling onto the lines in Fig. 3.9, as a function of one of the sides, being the other one specified by the line equations. The p-values evaluated at $z = 0$ for varying M_ν, and by keeping M_ν fixed at the value $0.4\,\mathrm{eV}$ for varying z, are shown in redshift space, respectively in the upper and lower panels of

---

[3] `https://docs.scipy.org/doc/scipy/reference/signal.html`

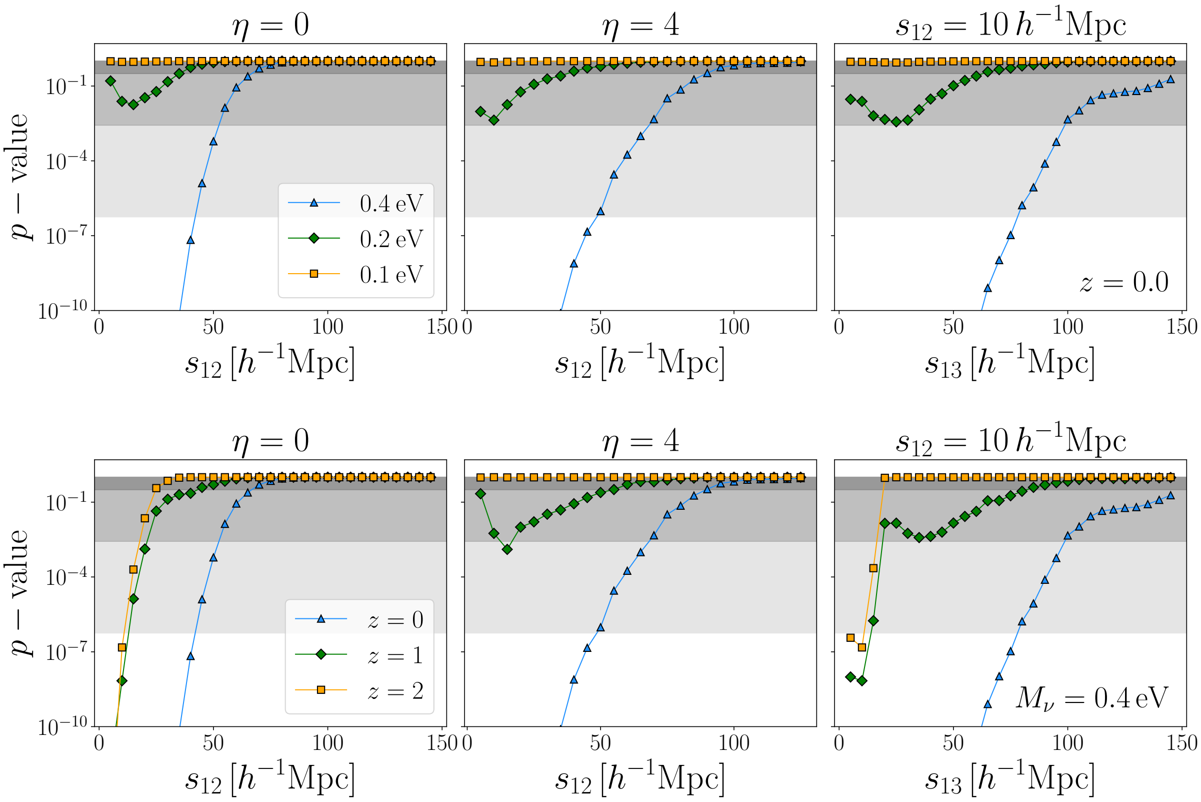

Figure 3.13: *Top panel*: the p-values computed for ζ through Eq. 2.5.3, for the choices of η and s_{12} reported above each plot, as a function of s_{12} (in the case of fixed η) and s_{13} (for fixed s_{12}). We fix the redshift to $z = 0$ and plot the p-values for varying M_ν. The gray bands, from the darkest to the lightest, individuate the ranges of p-values for which the significance computed through Eq. 2.5.4 is between 0 and 1σ, 1σ and 3σ, and 3σ and 5σ. *Bottom panel*: same as in the top panel, but for fixed mass $M_\nu = 0.4\,$eV for all the considered redshifts $z = 0, 1, 2$. The vertical axis shows only data with p-values above 10^{-10}.

Fig. 3.13 (for ζ) and 3.14 (for Q). We fix three thresholds, at 1σ, 3σ, and 5σ (shown in the figure with gray bands) by converting the p-values through Eq. 2.5.4 (see Tab. 2.4 for the corresponding p-values).

Let us first focus on the case $z = 0$. For both ζ and Q, the points associated with $M_\nu = 0.1\,$eV always have a statistical significance below 1σ, a result present at all redshifts.

For $M_\nu = 0.2\,$eV the points cross the 1σ threshold, remaining always below 3σ for ζ. For Q, they reach $\sim 1\sigma$ on small scales but only for isosceles and squeezed configurations, while remaining below 1σ for $\eta = 4$. The only mass detectable above 3σ and 5σ is the largest one $M_\nu = 0.4\,$eV, with the only exception of $\eta = 4$ in the case of Q, where it crosses 3σ but remains below 5σ.

For ζ, the scales at which the p-values cross a given significance threshold increase moving from $\eta = 0$ to $\eta = 4$ and then to the squeezed triangles described by $s_{12} = 10\,h^{-1}\,$Mpc. For example, for $M_\nu = 0.2\,$eV the 1σ crossing occurs at $s_{12} \approx 35\,h^{-1}\,$Mpc, $s_{12} \approx 40\,\mathrm{h}^{-1}\,$Mpc and $s_{13} \approx 60\,h^{-1}\,$Mpc respectively for $\eta = 0$, $\eta = 4$ and $s_{12} = 10\,h^{-1}\,$Mpc. Remarkably, for squeezed configurations, the statistical significance for $M_\nu = 0.4\,$eV is above 1σ at all

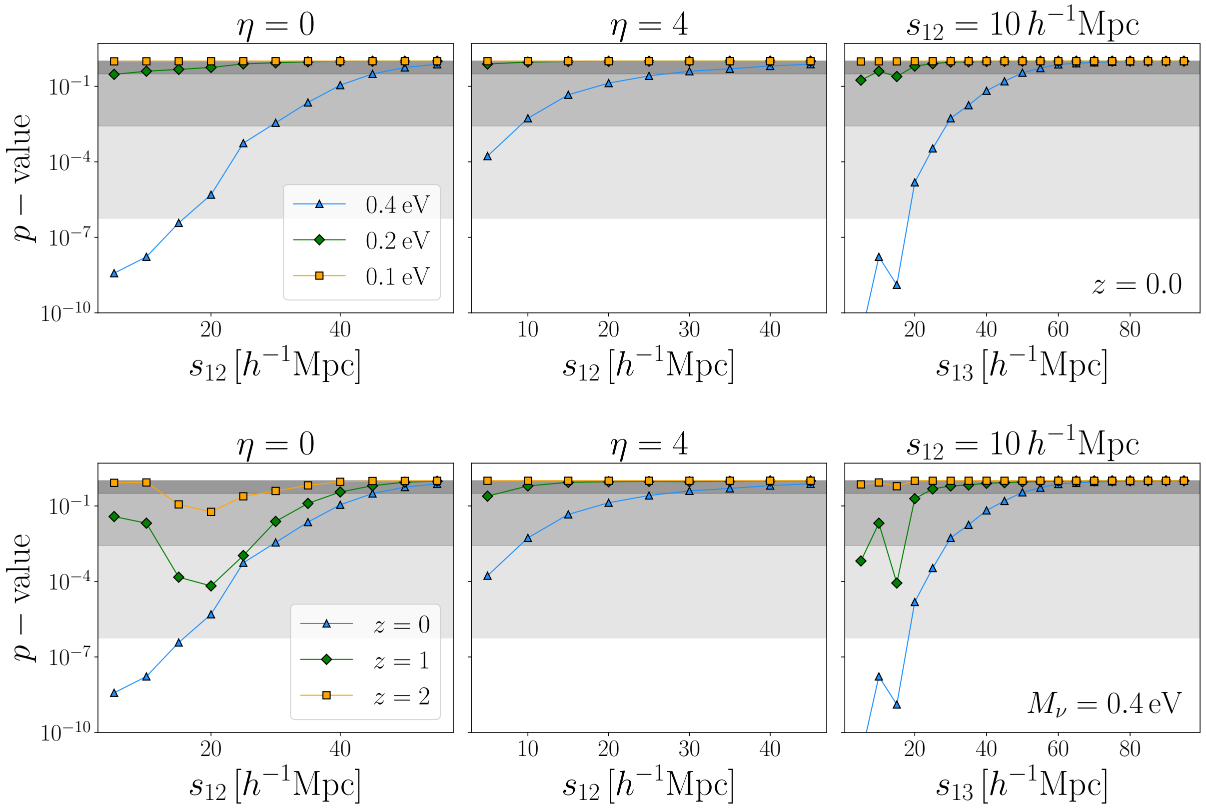

Figure 3.14: Same as Fig. 3.13, but for the reduced 3PCF Q.

the considered scales.

For Q, in the case $M_\nu = 0.4\,$eV, the 1σ, 3σ and 5σ crossings occur at very similar scales for $s_{12} = 10\,h^{-1}\,$Mpc and $\eta = 0$, while they move to smaller scales for $\eta = 4$. However, the statistical significance of a given scale for fixed M_ν is lower for Q as compared to ζ.

Limited to the case $z = 0$ shown in the figures, these results indicate that for ζ, squeezed configurations have higher statistical significance in the detection of massive neutrinos, followed by quasi-isosceles configurations and finally by isosceles ones. On the other hand, lower statistical significances overall characterize Q compared to ζ, with the largest values for $\eta = 0$ and squeezed triangles (comparably), followed by quasi-isosceles configurations. Obviously, these significances increase with M_ν, or otherwise stated, it is necessary to move to smaller scales to reach a fixed threshold of significance for progressively lower M_ν.

Let us now move to the case in which the neutrino mass is fixed to $M_\nu = 0.4\,$eV and the variable is redshift. A result valid for all the selected configurations is that the p-value corresponding to a given scale decreases as the redshift decreases. In the case of ζ, all the redshifts cross the 5σ threshold for $\eta = 0$ and $s_{12} = 10\,h^{-1}\,$Mpc, and only in the case $z = 0$ for $\eta = 4$. For $\eta = 4$, $z = 2$ remains always below 1σ, while $z = 1$ crosses the 3σ threshold at a scale $s_{12} \approx 15\,h^{-1}$Mpc, even if in the remaining two lower s_{12} bins the p-value rises again to significances between 1σ and 3σ.

For Q, the points at $z = 0$ always cross the 5σ threshold with the only exception of

$\eta = 4$, where they reach significances between 3σ and 5σ. For $z = 1$, the p-values show a minimum at $s_{12} \approx 20\,h^{-1}$ Mpc for $\eta = 0$, reaching significances between 3σ and 5σ. At $\eta = 4$, $z = 1$ remains always below 1σ except in the very first bin $s_{12} = 5\,h^{-1}$Mpc, where it reaches 1σ significance, while for squeezed triangles it fluctuates around 3σ on scales $s_{13} \lesssim 20\,h^{-1}$ Mpc. The points at $z = 2$ never cross the 1σ threshold with the only exception of $\eta = 0$, where the crossing occurs (remaining below the 3σ significance) in correspondence of a minimum around $s_{12} = 20\,h^{-1}$Mpc.

For Q, as in the previous case, the highest significances are reached comparably for $\eta = 0$ and squeezed configurations, followed by $\eta = 4$. In the case of ζ, the behavior appears to vary with redshift; moving from $z = 0$ to $z = 2$, quasi-isosceles configurations lose relevance, whereas isosceles and squeezed configurations become comparable.

We provide a summary of these trends at the end of this chapter, schematically represented in Fig. 3.19.

3.4 Neutrino signal in the isotropic connected and reduced 3-point correlation function: all scales

The three variables characterizing the single-scale approach considered in the previous section, i.e., two triangle sides r_{12}, r_{13} and the angle θ between them, although effective in examining the amplitude of clustering on a given scale for varying triangle shape, do not allow an efficient isolation of the contribution coming from triangles of fixed shape, from small to large scales. For example, if r_{12} and r_{13} are chosen such that $r_{12} \ll r_{13}$ or $r_{12} \gg r_{13}$, then squeezed triangles are obtained for both small and large θ. The all-scales approach, by ordering the sides of triangles, such that $r_{12} \leq r_{13} \leq r_{23}$ (as detailed in Sec. 2.4), allows for clear shape classification based on the values assumed by two independent side ratios, as we are about to discuss.

A useful tool for shape analysis is the triangular plot used in Hahn et al. (2020) (for the bispectrum), which we schematize in Fig. 3.15.

The horizontal and vertical coordinates of this plot are, respectively, the ratios r_{12}/r_{23} and r_{13}/r_{23} between the sides r_{12}, r_{13}, r_{23} of a triangle, following the order relation $r_{12} \leq r_{13} \leq r_{23}$. Fixing the values of these two ratios identifies a given triangle up to a similarity transformation (since similarity rescales all the sides by the same quantity). Otherwise stated, it identifies all triangles that share the same shape, determined by the values of the two ratios.

By construction, $r_{12}/r_{23}, r_{13}/r_{23} \in [0, 1]$, although not all the possible pairs in the rectangle $[0, 1] \times [0, 1]$ are allowed: indeed, the conditions $r_{23} \leq r_{12} + r_{13}$ and $r_{12} \leq r_{13}$ restrict the allowed pairs to the closed triangle of vertices $(0, 1), (1/2, 1/2)$ and $(1, 1)$ (the colored region in Fig. 3.15).

Let us now describe the triangle shapes close to the border of this triangle, starting from the lower corner and moving clockwise. The lower corner contains triangles with $r_{12}/r_{23} \approx r_{13}/r_{23} \approx 1/2$. These triangles are isosceles (at the exact location of the lower vertex) or quasi-isosceles, with the angle between the two smaller sides $r_{12} \approx r_{13}$ close to π (and exactly equal to π at the vertex), so that they exhibit an elongated shape. We

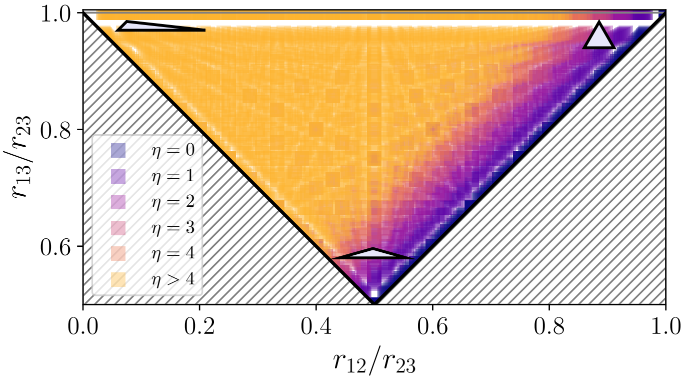

Figure 3.15: Generic scheme of a shape plot described in this section. A given point inside the triangle delimited by the shaded regions represents a set of similar triangles, sharing the same values of the side ratios r_{12}/r_{23} and r_{13}/r_{23}, where $r_{12} \leq r_{13} \leq r_{23}$. Shaded regions mark the values of the ratios for which no triangles can exist based on the closure condition $r_{23} \leq r_{12} + r_{13}$ and the ordering of the sides. The small white triangles drawn close to the vertices of the larger triangle represent the locations in the plot of squeezed (upper-left corner), equilateral (upper-right corner), and folded (lower corner) triangles. Different regions are colored depending on the value of the parameter η.

refer to these triangles as *folded* triangles.

Moving towards the upper-left corner, along the region adjacent to the left oblique side of the triangle plot, the side r_{12} decreases and r_{13} increases, but both satisfy the relation $r_{13} + r_{12} \approx r_{23}$: as a consequence of this latter constraint, triangles maintain an elongated shape (degenerating to a segment for $r_{13} + r_{12} = r_{23}$, i.e., on the exact side of the colored triangle). Points located close to the upper-left corner are characterized by $r_{12} \ll r_{13}$ and $r_{13} \approx r_{23}$. These triangles, in which the angles between r_{12} and r_{13}, and between r_{13} and r_{23} are, respectively, close to π and 0, are the already-mentioned squeezed triangles.

The upper-left vertex corresponds to the degenerate isosceles triangle in which $r_{13} = r_{23}$ and $r_{12} = 0$. Moving towards the upper-right corner across the region adjacent to the horizontal line described by $r_{13}/r_{23} = 1$, triangles are again isosceles or quasi-isosceles, this time in the two largest sides since $r_{13} \approx r_{23}$. Along this region, since r_{12} becomes progressively larger, the angle between r_{13} and r_{23} progressively increases from 0 to $\pi/3$ at the location of the upper-right vertex. Here, $r_{12} = r_{13} = r_{23}$, so this vertex describes *equilateral* triangles. Descending back to the lower vertex along the line $r_{12} = r_{13}$, the angle between two arbitrarily chosen sides of the equilateral triangle[4] progressively increases, describing isosceles triangles in r_{12} and r_{13}, with these two sides progressively folding over the largest one r_{23}. Right-angled triangles satisfy the identity $(r_{12}/r_{23})^2 + (r_{13}/r_{23})^2 = 1$, so they are located along the intersection between the colored triangle and the unitary circle centered at $(0, 0)$. The points far from the boundary describe shapes with intermediate features between the "extreme" cases we have just described.

Fig. 3.15 also shows, with different colors, the regions of the plot corresponding to

---

[4] Since the sides are all equal, everyone of them can represent r_{12}, r_{13} or r_{23}.

possible choices of $\eta \equiv (r_{13} - r_{12})/\Delta r$. The sides of the triangles characterized by a given η satisfy the equation $r_{13}/r_{23} = r_{12}/r_{23} + \eta \Delta r / r_{23}$. If $\eta = 0$, then this equation simply identifies the right oblique side of the triangular plot. If $\eta \neq 0$, since r_{23} is binned, this equation identifies $N_{\mathrm{r\,bins}}$ lines, one for each value of r_{23} (which we assume to represent the center of the bin). Naturally, triangles characterized by a given value of η lie on the $N_{\mathrm{r\,bins}}$ segments resulting from the intersection between these lines and the triangular region. This explicit dependence on r_{23} binning implies that a given choice of η does not translate into a partition of the triangle shapes, since triangles with different values of η can have the same shape: indeed, for a given side bin Δr, a pair (η, r_{23}), and an integer k, also $r'_{23} \equiv k \times r_{23}$ can be the center of another bin (if $r'_{23} < r_{\max}$), and the line associated with r_{23} and η coincides with the one related to r'_{23} and $\eta' \equiv k\eta$.

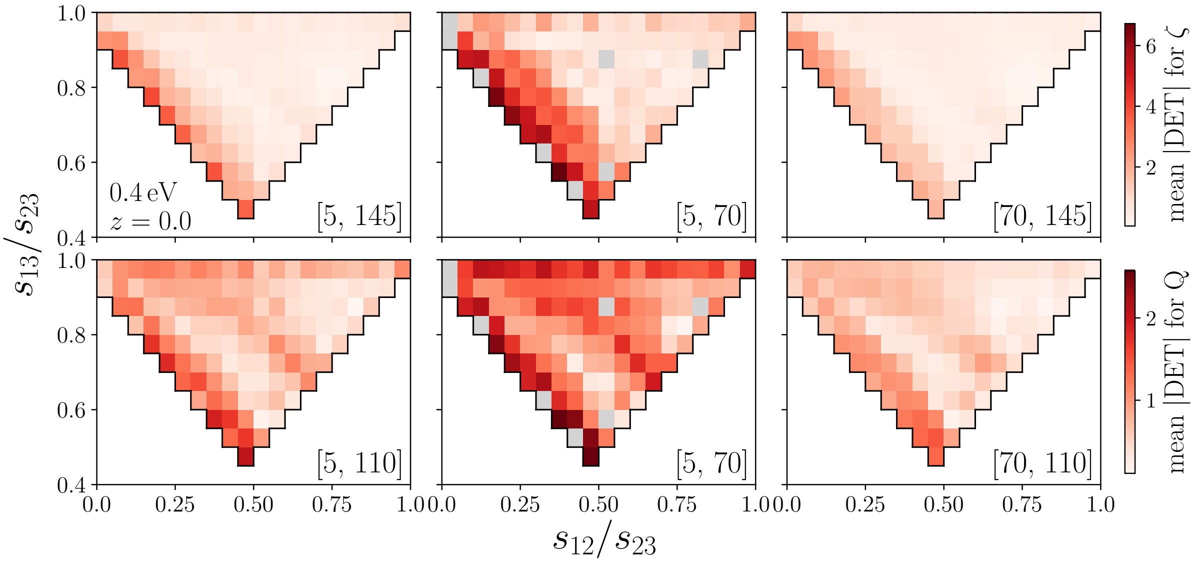

Figure 3.16: The detectability of the redshift-space halo connected 3PCF (*upper panels*) and reduced 3PCF (*lower panels*) at $z = 0$, as a function of triangle shape, for $M_\nu = 0.4\,\mathrm{eV}$. The triangle shapes are identified by the side ratios s_{12}/s_{23} and s_{13}/s_{23}, with $s_{12} \leq s_{13} \leq s_{23}$. In each panel, a given shape bin displays the absolute value of detectability averaged over all triangles available in the all-scales approach with that shape and different sizes. The white region corresponds to pairs $(s_{12}/s_{23}, s_{13}/s_{23})$ with which it is not possible to construct a triangle. The panels belonging to the same row differ by the range of s_{23} considered, specified in the intervals $[a, b]$ shown in the bottom-right corner (with a, b in units of $h^{-1}\,\mathrm{Mpc}$). The bins that admit triangles but are unpopulated due to the choice of the s_{23} range are colored in grey.

We employed this type of plot for a shape analysis based on detectability values, as follows.

Since the triangle sides r_{12}, r_{13} and r_{23} are binned, the ratios r_{12}/r_{23} and r_{13}/r_{23} computed for our triangles do not uniformly cover the triangular region shown in Fig. 3.15. Hence, we divided the range $[0, 1]$ into equal-width bins of size 0.05, a value which guarantees a sufficiently high shape "resolution" and at least one triangle falling in each bin (if no constraints on triangles in the all-scales framework are adopted). Then, after fixing a given space, redshift, and M_ν, for each shape bin we averaged the absolute values of the detectabilities over the triangles falling in that bin, for both $\zeta_{\rm all}$ and $Q_{\rm all}$. We performed this procedure for all the triangles considered in the all-scales framework, and separately for triangles with $5 \leq r_{23} \leq 70\,h^{-1}\,{\rm Mpc}$ and $70 \leq r_{23} \leq 145\,h^{-1}\,{\rm Mpc}$ for $\zeta_{\rm all}$ and $5 \leq r_{23} \leq 70\,h^{-1}\,{\rm Mpc}$ and $70 \leq r_{23} \leq 110\,h^{-1}\,{\rm Mpc}$ for $Q_{\rm all}$, to keep track of the evolution with scale of possible shape contributions, which risk, in considering all the triangles together, to be diluted by averaging over triangles covering a large range of scales.

The plots in Fig. 3.16 report the results of this procedure in redshift space, at $z = 0$ and for $M_\nu = 0.4\,{\rm eV}$, for both $\zeta_{\rm all}$ (in the top panel) and $Q_{\rm all}$ (in the bottom panel).

The average absolute values of detectability are overall lower for $s_{23} \geq 70\,h^{-1}\,{\rm Mpc}$, compatibly with the results of all the previous analyses, which show that the largest effects of massive neutrinos are mostly concentrated on small scales.

For $\zeta_{\rm all}$, as can be seen from the upper-left plot in the figure, which includes all the triangles, the most impacted regions are those corresponding to the left oblique side of the plot, containing, as extensively discussed above, elongated triangle shapes gradually transitioning, for increasing s_{13}/s_{23}, from folded to squeezed. The two adjacent plots show that this feature arises from two distinct behaviors characterizing the triangles falling within the two s_{23} ranges. For smaller s_{23}, the contribution from folded and squeezed shapes is comparable, with averages reaching values $\gtrsim 6$ in some shape bins for the selected configuration, and extends also to less elongated triangles (i.e., also towards more central regions of the plot). The upper region, including shapes that transition from squeezed to equilateral, also exhibits a lower but non-vanishing detectability. For larger s_{23}, this latter feature disappears, and the detectability pattern on the left oblique side of the plot migrates towards squeezed configurations.

These changes are naturally explained by considering that the larger values of detectability trace the location of triangles whose shapes, for a given range of s_{23}, allow at least one side to be small: for $5 \leq s_{23} \leq 70\,h^{-1}\,{\rm Mpc}$, this is possible for both squeezed and folded shapes (some examples are, for these latter, triangles like $(s_{12}, s_{13}, s_{23}) = (5, 5, 10), (10, 10, 20)\,h^{-1}\,{\rm Mpc}$), and for triangles with $s_{13}/s_{23} \approx 1$ simply because it is sufficient that s_{23} decreases for increasing s_{12}/s_{23} to guarantee that at least the value of s_{12} is small. For $70 \leq s_{23} \leq 145\,h^{-1}\,{\rm Mpc}$, the smallest folded triangle is $(s_{12}, s_{13}, s_{23}) = (45, 45, 70)\,h^{-1}\,{\rm Mpc}$, significantly larger as compared to the previous examples. In contrast, squeezed triangles can still have small s_{12} values even for the largest s_{13} and s_{23} values, as in the case of $(s_{12}, s_{13}, s_{23}) = (5, 140, 145)\,h^{-1}\,{\rm Mpc}$. Moreover, for $s_{13}/s_{23} \approx 1$ (and consequently $s_{13}, s_{23} \geq 70\,h^{-1}\,{\rm Mpc}$), only triangles with small s_{12}/s_{23} (i.e., with squeezed shape) still preserve small values of s_{12}.

For $Q_{\rm all}$ and including all the triangles, we can instead identify three regions showing the highest detectabilities: the one associated with squeezed and folded shapes (as in the case of $\zeta_{\rm all}$), the isosceles triangles with shape transitioning from squeezed to equilateral (having $s_{13}/s_{23} \approx 1$), and a region crossing obliquely the colormaps, roughly from $(s_{12}/s_{23}, s_{13}/s_{23}) \approx (0, 1)$ to $(s_{12}/s_{23}, s_{13}/s_{23}) \approx (0.65, 0.70)$, compatibly with the location of right-angled triangles. The fact that this latter region is detectable in $Q_{\rm all}$ but not in $\zeta_{\rm all}$ is attributable to the flattening noticed on the single-scale Q caused by an increase in M_ν, an opposite trend as compared with single-scale ζ. Indeed, one of the consequences of this flattening is the increase in detectability around $\theta \approx \pi/2$, i.e., in correspondence

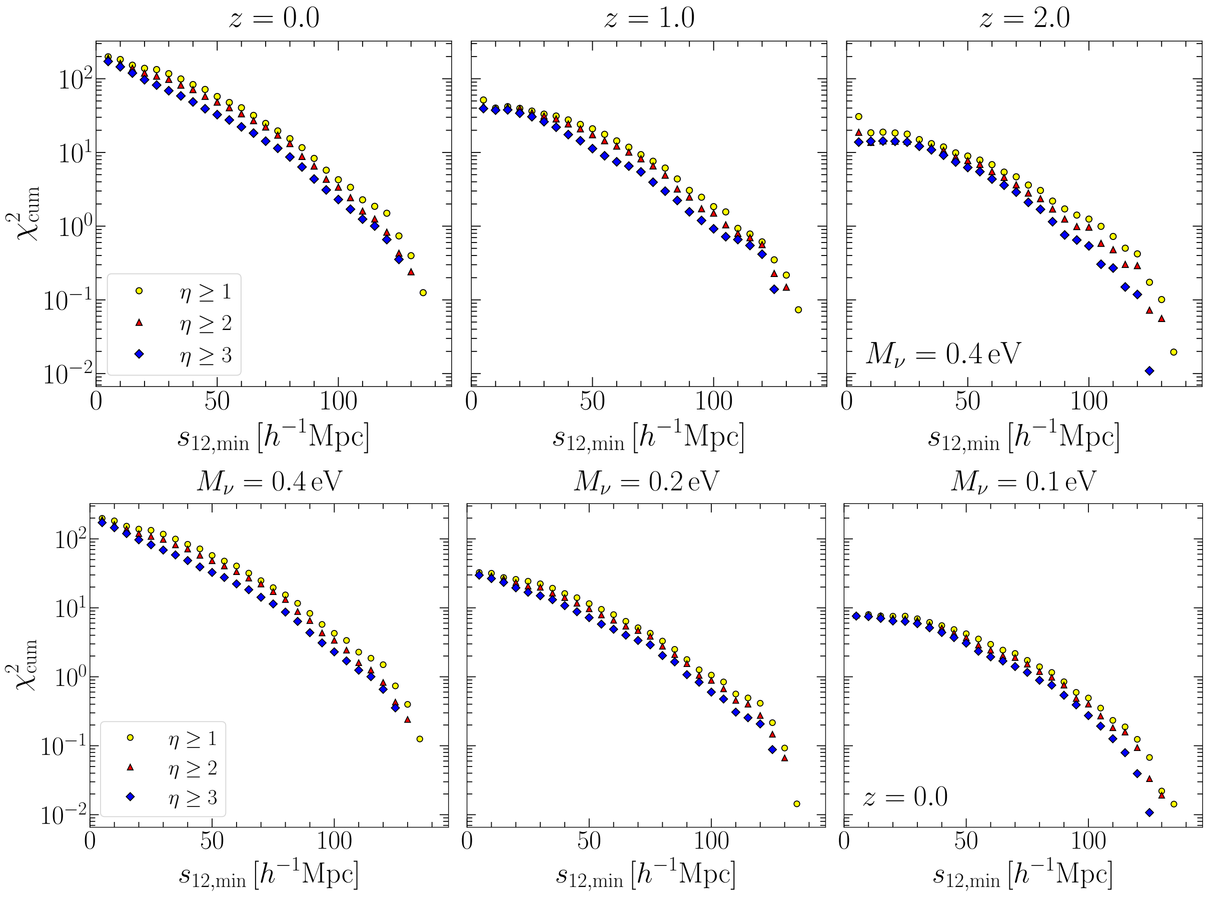

Figure 3.17: *Top panel*: the cumulative chi-square indicator $\chi^2_{\rm cum}$ for the all-scales 3PCF $\zeta_{\rm all}$ as a function of $s_{12,\rm min}$, computed through Eq. 2.5.5, for the redshifts z reported above each plot. We fix $M_\nu = 0.4\,\mathrm{eV}$ and plot the results for $\eta \geq \eta_{\rm min}$, with $\eta_{\rm min} = 1, 2, 3$. The maximum value of $s_{12,\rm min}$ depends on $\eta_{\rm min}$, since triangles must satisfy the constraint $\eta = (s_{13} - s_{12})/\Delta s \geq \eta_{\rm min}$, and corresponds to $145 - \eta_{\rm min}\Delta s$ (in units of $h^{-1}\,\mathrm{Mpc}$; $145\,h^{-1}\,\mathrm{Mpc}$ is our largest measured scale, and $\Delta s = 5\,h^{-1}\,\mathrm{Mpc}$ our binning). *Bottom panel*: same as in the top panel, but at fixed redshift $z = 0$, for varying M_ν.

with right-angled triangles.

The three features are all present for $5 \leq s_{23} \leq 70\,h^{-1}\,\mathrm{Mpc}$, where, in the first two ones, the highest detectabilities are reached, respectively, for folded and squeezed shapes. Moving to $70 \leq s_{23} \leq 110\,h^{-1}\,\mathrm{Mpc}$, the feature present for $s_{13}/s_{23} \approx 1$ migrates towards smaller values of s_{12}/s_{23} (more squeezed triangles) with respect to the previous case where it extends up to equilateral triangles. The other two features are still present, with increasing detectability going from squeezed to folded triangles.

To conclude this analysis, we present the results of the computation of the cumulative chi-square indicator $\chi^2_{\rm cum}$ as a function of $r_{12,\rm min}$ as defined in Eq. 2.5.5.

Due to the large number of triangles determined by the chosen range of scales and binning (2571 triangles for ζ and 1199 for Q), we performed a pseudo-inversion (Strang 1980) of the covariance matrix in Eq. 2.5.5 rather than a full inversion, including all the

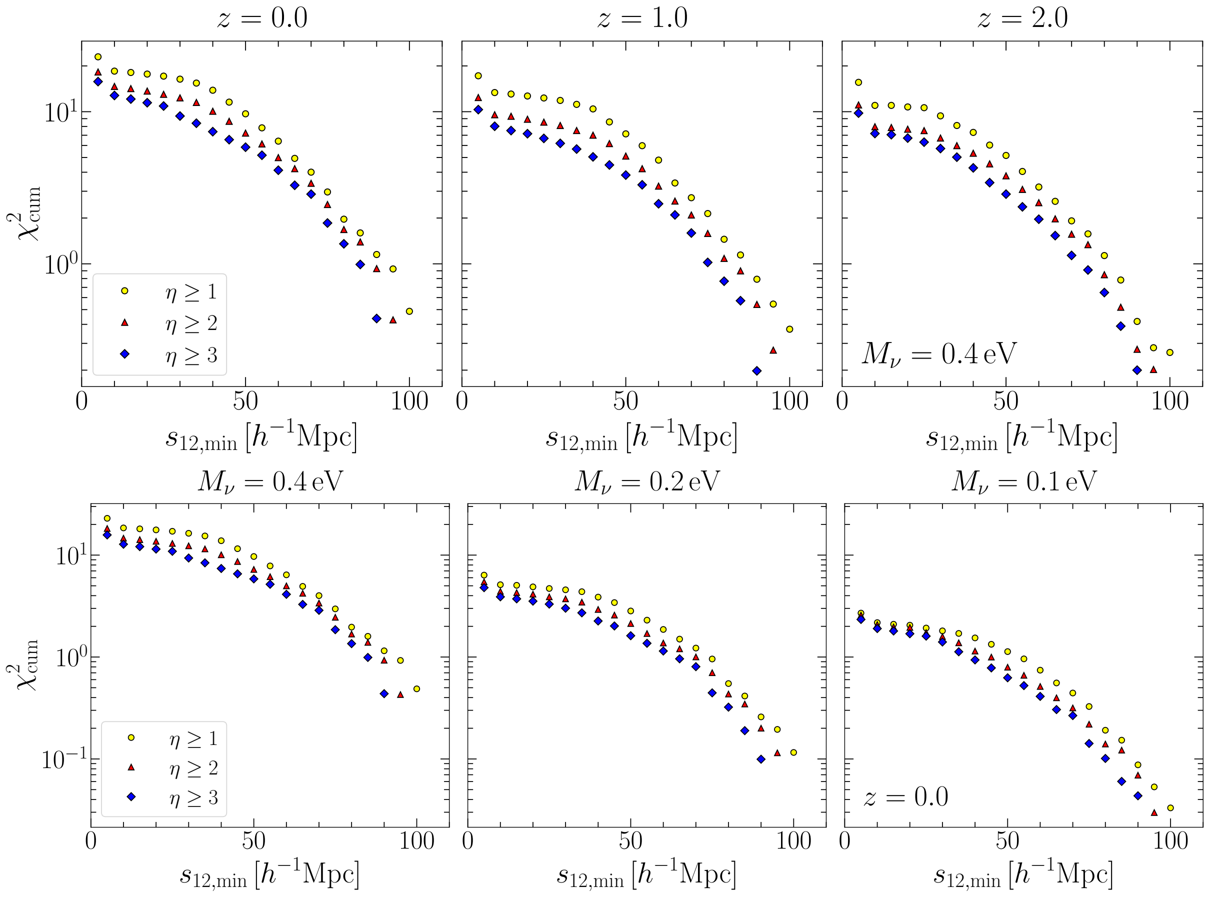

Figure 3.18: Same as Fig. 3.17, but for the reduced 3PCF Q. Here, for the same reason as in Fig. 3.18, the maximum value of $s_{12,\rm min}$ depends on $\eta_{\rm min}$ as $110 - \eta_{\rm min}\Delta s$ (n units of h^{-1} Mpc; $110\,h^{-1}$ Mpc is the largest scale we consider for Q, and $\Delta s = 5\,h^{-1}$ Mpc our binning).

eigenvalues of $\hat{C}$ which, when accumulated from largest to smallest, account for 99.5% of the total sum of the eigenvalues. This cut excludes the smallest eigenvalues, which are also the most affected by noise. We consequently computed the Hartlap corrections in Eq. 2.5.6 by replacing the total number of data points n_d with the number of selected eigenvalues in each pseudo-inversion.

The results are displayed in redshift space in Fig. 3.17 and 3.18 for $\zeta_{\rm all}$ and $Q_{\rm all}$, respectively. In both cases, we show in the upper panels the values of $\chi^2_{\rm cum}$ computed at fixed $M_\nu = 0.4\,\mathrm{eV}$ at the three redshifts $z = 0, 1, 2$ (one subplot per redshift), while in the lower panels we fix $z = 0$ and show the results for varying $M_\nu = 0.4, 0.2$ and $0.1\,\mathrm{eV}$ (one subplot per mass). To study the dependence of the indicator on the parameter η, we evaluated it on triangles satisfying $\eta \geq \eta_{\rm min}$, with $\eta_{\rm min} = 1, 2, 3$. These three cases are shown in each subplot.

For both $\zeta_{\rm all}$ and $Q_{\rm all}$, the values of $\chi^2_{\rm cum}$ at fixed $s_{12,\rm min}$ increase for increasing M_ν at fixed z and $\eta_{\rm min}$, and for decreasing z at fixed M_ν and $\eta_{\rm min}$. An increase in $\chi^2_{\rm cum}$ is also present for decreasing $\eta_{\rm min}$ at fixed M_ν and z, compatibly with the inclusion of a larger number of triangles in the analysis.

The values of $\chi^2_{\rm cum}$ for ζ steepen in the range $s_{12,\rm min} \lesssim 30$ for decreasing z at fixed M_ν, or for increasing M_ν at fixed z. Such a behavior is not evident in Q. This quantitatively shows that discarding these scales leads to a significant loss of the informative content of massive neutrinos in ζ, especially at low z. Therefore, it will be crucial to include those scales in future modeling and survey analysis.

3.5 Summary of the main results

Here, we summarize the main results of the analysis presented in this chapter.

Neutrino signal in the 3PCF isotropic multipoles. We find that the detectability of massive neutrinos (Eq. 2.5.1), shown in Fig. 3.7, increases at small separations for all the multipoles. The quadrupole $\ell = 2$ is the most affected by neutrinos, with vanishing detectabilities only for large values of both sides s_{12} and s_{13}. This multipole represents the quadratic term in the expansion of Eq. 1.2.26, which encodes most of the information of the 3PCF, giving ζ a shape roughly similar to that of a parabola as a function of the angle θ between two sides.

Impact of small scales and filamentary structures. We find that squeezed, quasi-isosceles, and isosceles triangles up to $s_{12} = s_{13} \approx 30\,h^{-1}$ Mpc are most impacted by neutrinos for ζ, while for Q, isosceles triangles are dominant (see Fig. 3.9 and Eq. 2.5.2). The signal increases for decreasing redshift. The main contribution to the signal from neutrinos comes from squeezed or elongated triangles in ζ (i.e., respectively, $\theta \approx 0$ and $\theta \approx \pi$), i.e., from the filamentary structure of the cosmic web, translating into an increase of concavity for increasing M_ν (Fig. 3.10). The signal is negligible for perpendicular configurations ($\theta \sim \pi/2$). In Q, the opposite trend is found, resulting in a flattening for increasing M_ν (Fig. 3.11). In this latter case, the result is that, in addition to elongated triangles, the signal is also driven by perpendicular configurations. Furthermore, we confirm also for the 3PCF that smaller scales are particularly sensitive to neutrinos (Eqs. 2.5.3 and 2.5.4, and Figs. 3.13 and 3.14)

These scales are schematically represented in Fig. 3.19 for ζ and Q, for the chosen 1σ, 3σ, and 5σ significance thresholds, as a function of M_ν, redshift, and the configurations maximizing our $\tilde{\chi}^2$ parameter (squeezed and quasi-isosceles configurations are respectively represented by fixing $s_{12} = 10\,h^{-1}$ Mpc and $\eta \equiv |s_{13} - s_{12}|/\Delta r = 4$). We find that $M_\nu = 0.1\,\rm eV$ is never detectable above 1σ, neither with ζ nor Q. Given that the overall signal is lower for Q than for ζ, going to smaller scales is needed to reach a given statistical significance of detection for the former with respect to the latter.

Shape dependence of massive neutrino signal. The analysis of the impact of neutrinos on ζ and Q (Fig. 3.16) shows that while the signal in ζ is essentially concentrated on elongated shapes ($s_{12} + s_{13} \approx s_{23}$), the one in Q presents additional neutrino signatures from right-angled triangles ($s_{23}^2 \approx s_{12}^2 + s_{13}^2$), and isosceles triangles with $s_{13} \approx s_{23} > s_{12}$ (we assume the ordering $s_{12} \leq s_{13} \leq s_{23}$). As s_{23} increases, in ζ the signal on elongated triangles migrates towards squeezed shapes ($s_{12} \ll s_{13} \approx s_{23}$), while in Q towards folded isosceles triangles ($s_{12} \approx s_{13} \approx s_{23}/2$).

Impact of massive neutrinos on BAO through the 3PCF. BAO features are identified in the 3PCF monopole $\ell = 0$, dipole $\ell = 1$, and quadrupole $\ell = 2$ (Fig. 3.8), compatibly with their predicted positions in Slepian and Eisenstein (2017). For $\ell = 0$ with $M_\nu = 0.2 - 0.4\,\mathrm{eV}$, and $\ell = 1$ with $M_\nu = 0.4\,\mathrm{eV}$, the uncertainties computed for our full effective volume of $V_{\mathrm{eff}} = 500\,h^{-3}\,\mathrm{Gpc}^3$ are sufficiently small to detect a change of shape incompatible with a simple linear bias rescaling. In particular, the zero crossing of the dipole on scales $\approx 105\,h^{-1}\,\mathrm{Mpc}$, marking the BAO position, shifts toward higher scales for increasing M_ν (top-right panel in Fig. 3.8).

For the single-scale 3PCF, we find a BAO feature in the form of a small peak embedded in the 3PCF minimum for $(s_{12}, s_{13}) = (55, 110)\,h^{-1}\,\mathrm{Mpc}$, when the third side s_{23} crosses BAO scales (Fig. 3.12). Unfortunately, due to the too large errors on ζ, no particular trend in BAO amplitude variation with M_ν can be significantly detected in this case.

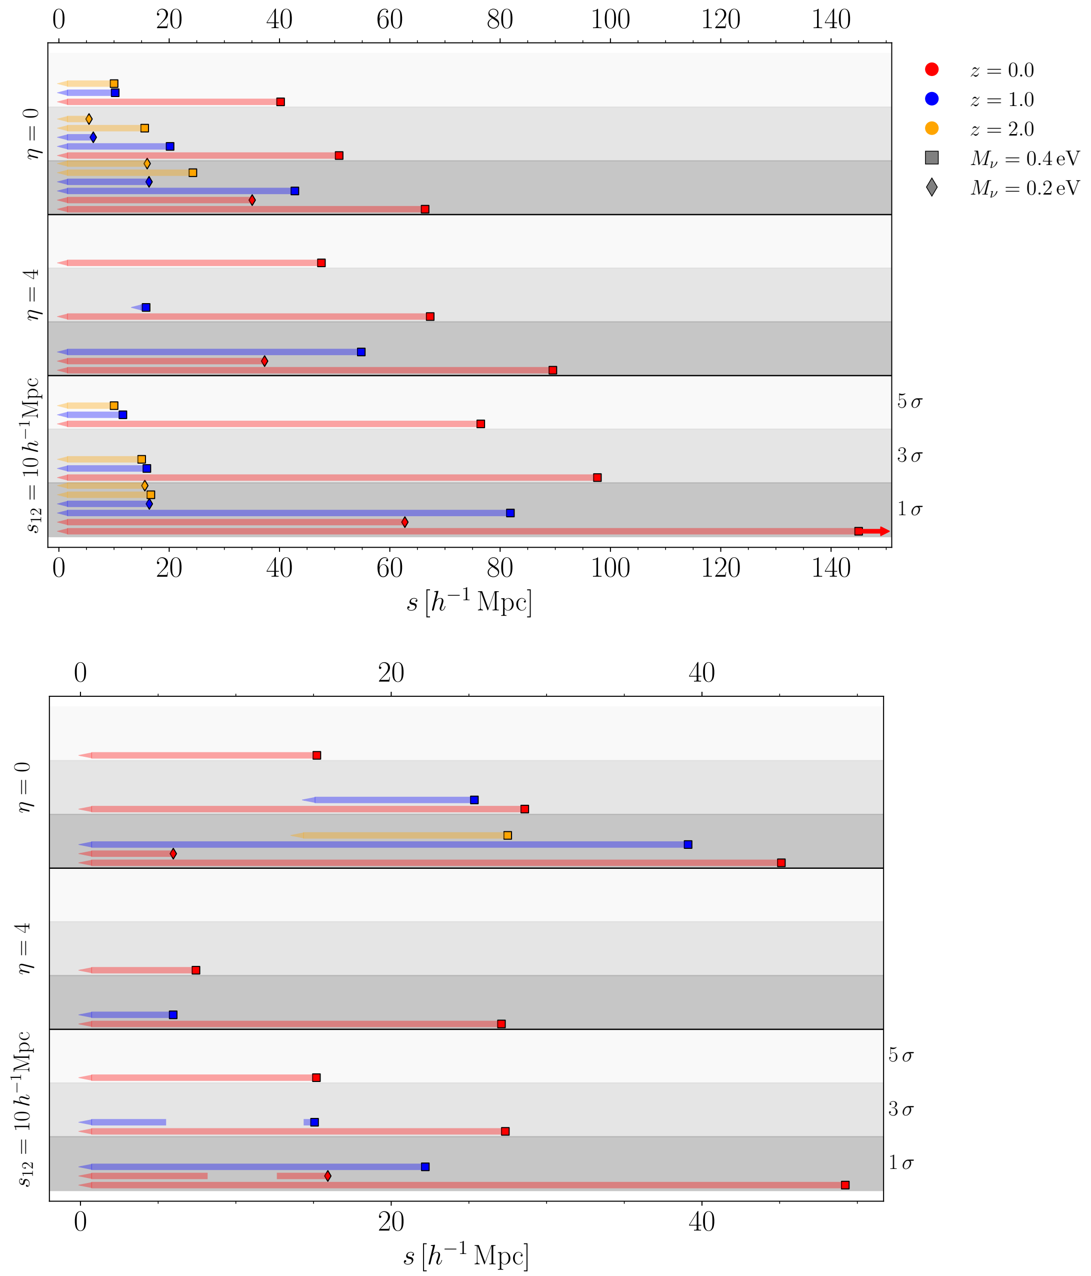

Figure 3.19: *Top panel*: the scales s (horizontal colored bands), at which neutrinos are detectable above a given significance threshold in the redshift-space connected 3PCF, depending on the gray band on which they are plotted: from the darkest to the lightest: 1σ, 3σ, and 5σ. Each color and marker corresponds, respectively, to a different redshift and M_ν, as shown in the legend. $M_\nu = 0.1\,\mathrm{eV}$ is absent because it never gives signal above 1σ. The panel is divided into three stacked sub-panels, one for each set of considered triangles: isosceles ($\eta = 0$), nearly-isosceles with $\eta = 4$, and squeezed with $s_{13} = 10\,h^{-1}\,\mathrm{Mpc}$. The variable s corresponds to s_{12} if η is fixed, and to s_{13} for squeezed triangles. The case $M_\nu = 0.4\,\mathrm{eV}$, $z = 0$, $s_{12} = 10\,h^{-1}\,\mathrm{Mpc}$ is always above 1σ significance, so the largest scale we probe $s = 145\,h^{-1}\,\mathrm{Mpc}$ is only a lower limit (represented with a right-pointing arrow). *Bottom panel*: same as in the top panel, but for the reduced 3PCF Q.

Chapter 4

Mitigating the M_ν-σ_8 degeneracy with higher-order clustering statistics

In this chapter, we explore the possibility of breaking the well-known M_ν-σ_8 degeneracy affecting 2-point statistics through the connected and reduced 3PCF. For this purpose, we enlarge our dataset by incorporating measurements of halo ζ and Q that we perform on ΛCDM simulations from the Quijote suite, with σ_8 varied by ± 0.015 relative to the fiducial value.

In Sec. 4.1, we frame the topic in the broader context of the degeneracies involving neutrino mass in massive-neutrino cosmologies. Subsequently, in Sec. 4.2, we provide a concise overview of the literature results on this degeneracy concerning 2-point statistics, and we discuss a promising result obtained through the bispectrum. In Sec. 4.3, we describe our new dataset and show some comparisons of average single-scale ζ and Q we estimated in our massive neutrino and varying σ_8 simulations. We then show how halo bias enters the space of degenerate parameters in both the 2PCF and 3PCF in Sec. 4.4. We investigate whether neutrinos impact the triangles as a function of shape and scale differently from σ_8, and present the results of this analysis in Sec. 4.5. Finally, we summarize the main results reported in this chapter in Sec. 4.6.

4.1 Degeneracies in massive neutrino cosmologies

The peculiar thermal history of the Cosmic Neutrino Background and the fact that at late times, despite being non-relativistic, the particles that make it up still free-stream on small scales, imprint characteristic signatures in several cosmological observables, as described in Sec. 1.5 and studied in the previous chapter. However, although the possibility of constraining neutrino masses through cosmological studies relies on detecting these effects, several degeneracies are present with other parameters.

Without aiming to be exhaustive, we provide here some examples that we will not address in the remainder of this work. In the rest of the chapter, we will focus specifically on the degeneracy between the sum of neutrino masses and σ_8.

Geometrical effect. Neutrinos affect the background evolution of the Universe during their transition to the non-relativistic regime, in particular its expansion rate, described in terms of the Hubble parameter $H(z)$ (a phenomenon referred to as the *geometrical effect*, see the recent discussion in Chebat et al. 2025). At the time of the non-relativistic transition of a given neutrino species, the matter density parameter $\Omega_m(z)$ increases, and its variation affects the Hubble parameter. The dependence on redshift of $H(z)$, assuming a flat background cosmology with a cosmological constant Λ, can be expressed by rearranging Eq. 1.1.19 as

$$H(z) = H_0\sqrt{\Omega_m(1+z)^3 + \Omega_r(1+z)^4 + \Omega_\Lambda}, \tag{4.1.1}$$

where Ω_m, Ω_r and Ω_Λ are the present-day values of, respectively, the matter, radiation, and dark energy density parameters. A change in the redshift evolution of the Hubble parameter (i.e., a variation in the expansion history of the Universe) caused by massive neutrinos affects the angular diameter distance to the surface of last scattering. This quantity is measurable from CMB anisotropies, but it is degenerate with the effects of other cosmological parameters such as Ω_m and H_0, as can be realized by looking at Eq. 4.1.1 [1]. This implies that constraints on the neutrino masses exploiting the geometrical effect can be obtained by combining CMB measurements with datasets able to independently constrain Ω_m, like the BAO measurements from DESI (which in turn constrain Ω_m but not H_0).

$\Omega_\nu - \Omega_m$ **degeneracy.** This degeneracy is related to the already-introduced parameter $f_\nu \equiv \Omega_\nu/\Omega_m$, which rules the effects imprinted by neutrinos on the matter power spectrum (rather than Ω_ν alone) even in the nonlinear regime, as shown by Viel et al. (2010): this implies that Ω_ν, thus M_ν, and Ω_m are degenerate in reproducing a given small-scale suppression in the power spectrum.

$\Omega_\nu - w$ **degeneracy.** If the EOS parameter of DE is allowed to take values different from that of a cosmological constant ($w = -1$), it becomes degenerate with M_ν. This degeneracy, in the form of a $w - \Omega_\nu$ anticorrelation, has already been pointed out in Hannestad (2005), where it was shown that it is responsible for relaxing the constraints on the total neutrino mass from the combination of CMB, SNIa, and galaxy clustering data (specifically, for these latter, the galaxy power spectrum). Indeed, when Ω_ν increases, Ω_m must increase to produce the same power spectrum (i.e., to produce the same f_ν). However, a larger Ω_m would slow down the expansion of the Universe probed by SNIa, an effect which can be remedied by simultaneously decreasing w.

Modified gravity scenarios. The fact that the growth rate of cosmic structures in the presence of massive neutrinos differs with respect to the case of a standard ΛCDM cosmology (as already discussed in Sec. 1.5) can induce a bias in the inferred growth rate from data analysis (e.g., from measurements of the clustering

---

[1] The other parameters in the expression, Ω_r and Ω_Λ, are constrained respectively by the CMB mean temperature and, in a flat cosmology, by the requirement that $\Omega_m + \Omega_r + \Omega_\Lambda = 1$.

anisotropies induced by RSDs), leading to a potentially false signature of modified gravity (Simpson et al. 2011).

We now proceed to discuss the main focus of this chapter, i.e., the degeneracy between M_ν and σ_8, providing a concise overview of the results obtained so far from clustering studies, at the level of 2-point statistics and, more recently, 3-point statistics.

4.2 The M_ν - σ_8 degeneracy: an overview

4.2.1 2-point statistics

We have pointed out, in Sec. 1.5, that massive neutrinos induce a suppression on the total matter and CDM power spectra on scales below their free-streaming scale, with respect to the power spectrum of a massless neutrino cosmology *normalized to the same amplitude*. The neutrino-to-total matter density ratio f_ν regulates this suppression, which, in the nonlinear regime, can be larger than the linear theory prediction by about 25%, while remaining, with good approximation, fully dependent on f_ν only (Viel et al. 2010).

For fixed amplitude, an increase in the neutrino mass naturally produces a decrease in the value of σ_8 (see Eq. 1.2.17). This effect can be mimicked, in a massless neutrino cosmological model, by an overall rescaling of the matter power spectrum, corresponding to a change in the amplitude of the primordial curvature perturbations A_s.

We point out from the beginning that this degeneracy is not perfect, in the sense that the power spectrum or, equivalently, the 2PCF in the presence of massive neutrinos cannot be perfectly mimicked at all scales simply with a σ_8 change of the background cosmology, as noticed, e.g., in Marulli et al. (2011). Indeed, such a renormalization would not affect the shape of the power spectrum/2PCF, at odds with the scale-dependent effect of neutrino free-streaming. This implies that it is not impossible to obtain an upper bound on the sum of neutrino masses and, at the same time, constraining σ_8 without CMB priors: for example, in the already cited work by Viel et al. (2010), this has been done by using the *Sloan Digital Sky Survey* 3 (SDSS-3)[2] Lyman-α flux power spectrum alone. Hopefully, upcoming galaxy surveys like Euclid will provide a significant step forward in breaking this degeneracy, measuring the overall shape of the power spectrum (see e.g., the forecasts contained in Euclid Collaboration: Archidiacono et al. 2025).

In massive neutrino cosmologies, similarly to bias, σ_8 can be defined either with respect to the total matter field or the CDM field only: we refer to these quantities with $\sigma_8^{(m)}$ and $\sigma_8^{(c)}$ respectively. In Villaescusa-Navarro et al. (2018), the authors have measured the redshift-space total matter, CDM+baryon, and halo power spectra at different redshifts using a subset of the HADES simulations[3]. Their suite includes massive neutrino simulations run with fixed value of A_s with $M_\nu = 0.06, 0.10$, and $0.15\,\mathrm{eV}$, and ΛCDM simulations with value of σ_8 matching either the value of $\sigma_8^{(m)}$ or $\sigma_8^{(c)}$ of the massive neutrino simulations. They have compared the ratios between the monopoles and quadrupoles of

---

[2] https://www.sdss3.org/index.php

[3] https://franciscovillaescusa.github.io/hades.html

the power spectrum in massive and massless neutrino cosmologies and the same quantities measured for a set of fiducial simulations, finding for the matter power spectrum a strong degeneracy for $k \gtrsim 0.1\,h\,\mathrm{Mpc}^{-1}$, particularly when the value of σ_8 of a massless neutrino simulation matches the value of $\sigma_8^{(c)}$ of a massive neutrino simulation. On scales $k \lesssim 0.1\,h\,\mathrm{Mpc}^{-1}$, the effects are different, (but within $\sim 6\%$ up to $10^{-2}\,h\,\mathrm{Mpc}^{-1}$ for their largest mass $M_\nu = 0.15\,\mathrm{eV}$), with a larger decrease of power for the ΛCDM simulations with decreased σ_8. At the level of total matter, the quadrupoles of massive neutrino simulations are distinguishable from the ΛCDM quadrupoles at all scales, while degeneracy is still present at small scales for the CDM+baryon field. A degeneracy is found on small scales ($k > 0.1\,h\,\mathrm{Mpc}^{-1}$) also for the linear growth rate f (Eq. 1.3.8). In all the cases, the different time evolution of the models with massive and massless neutrinos partially breaks the degeneracy. Similar results have been found in the same work for the halo power spectra, with a stronger level of degeneracy for $\sigma_8^{(c)}$ rather than $\sigma_8^{(m)}$, although weaker than in the case of the matter/CDM+baryon power spectra.

4.2.2 The importance of adding 3-point statistics

3-point statistics can play an important role in breaking the $M_\nu - \sigma_8$ degeneracy, as suggested by Hahn et al. (2020), where the authors have measured the redshift-space halo bispectrum monopole $\hat{B}_0$ for the same suite of simulations, finding that in simulations with matching $\sigma_8^{(c)}$ it has larger amplitude variations than the power spectrum. In particular, they have used the shape plots described in Sec. 3.4 in Fourier space, to plot the relative differences of the halo bispectrum monopole with respect to their fiducial simulations. They have proved that different shapes react differently to changes of M_ν in massive neutrino cosmologies compared to changes of σ_8 in the ΛCDM cosmologies. More specifically, if σ_8 is varied to match the values of the relative differences for a given subset of triangle shapes between a massive and massless neutrino cosmology, then the match cannot be reproduced for another subset of shapes.

This is evident in Fig. 4.1, which has been taken from the cited work: by comparing the colormaps for the simulation with $M_\nu = 0.1\,\mathrm{eV}$ and the ΛCDM simulation with $\sigma_8 = 0.818$ (top center and bottom right panels, respectively), it can be noticed that near-equilateral triangles (top-right corner of the colormaps) in the two panels have similar $\Delta\hat{B}_0/\hat{B}_0^{\mathrm{fid}}$, while these values are significantly different for squeezed-folded triangles (left side of the colormaps).

This result motivates us to address the problem in configuration space through the study of the triangle shape dependence of the connected and reduced 3PCF, similarly to what we have already discussed for the neutrino detectability patterns as a function of triangle shape in Sec. 3.4.

For all the following discussion, it is important to emphasize that, although the reduced 3PCF does not explicitly depend on σ_8, it becomes indirectly sensitive to variations of this parameter when estimated for the halo distribution, through the effects such variations induce on halo bias, depending on how halos are selected. Consequently, we defer the details of our shape analysis to Sec. 4.5. Before that, after presenting the new dataset that complements our previous measurements, we discuss the role of bias in investigating

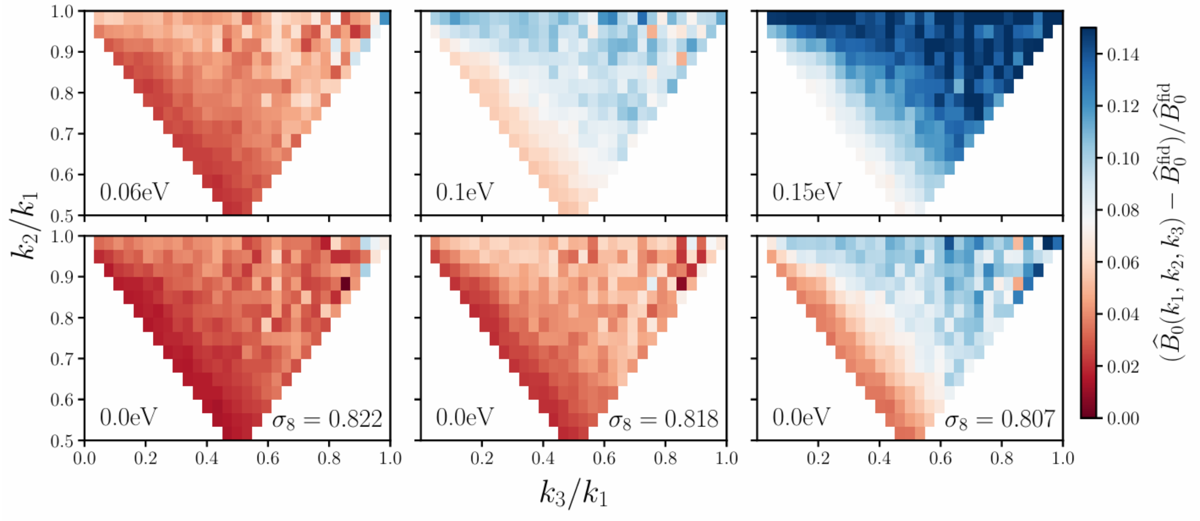

Figure 4.1: The figure, taken from Hahn et al. (2020), shows the differences of the redshift-space halo bispectrum monopole measured in the HADES simulations as a function of the triangle shapes, with respect to a fiducial simulation ($M_\nu = 0.0\,$eV and $\sigma_8 = 0.833$). The simulations with $M_\nu = 0.06, 0.10,$ and $0.15\,$eV in the upper panels are aligned with the ones with $M_\nu = 0.0\,$eV and $\sigma_8 = 0.822, 0.818,$ and 0.807 in the bottom panel such that the simulations in each column have matching $\sigma_8^{(c)}$, which produce mostly degenerate imprints on the reshift-space power spectrum. The ordering of the triangle sides is $k_1 \geq k_2 \geq k_3$.

the degeneracy between M_ν and σ_8 through halo clustering.

4.3 The simulated dataset

Together with the massive neutrino simulations analyzed in the previous chapter (run with a fixed value of $\sigma_8 = 0.834$), we have included two sets of 500 paired-fixed ΛCDM simulations each, taken also in this case from thw QUIJOTE suite, characterized by the same resolution and run with $M_\nu = 0.0\,$eV, and two different values of σ_8, namely $\sigma_8 = 0.819$ and 0.849. Further details about the simulations (e.g., the ICs) can be found in Tab. 2.2 contained in Ch. 2, where they have already been mentioned with the name `s8_p` (for the simulation with $\sigma_8 = 0.849$) and `s8_m` (for $\sigma_8 = 0.819$). The corresponding control sample is made up of 500 additional paired-fixed ΛCDM simulations with $M_\nu = 0.0\,$eV and $\sigma_8 = 0.834$, labelled with the name `fiducial` in the same table (second row). Also for these simulations, we measured the halo 2PCF and 3PCF using the procedure detailed in Ch. 2 and computed the corresponding errors for a volume of $10\,h^{-3}\,\mathrm{Gpc}^3$.

The small variation in σ_8 considered (±0.015) is of the same order of magnitude or larger than the uncertainties with which this cosmological parameter is currently constrained by several cosmological probes. For instance, exploiting CMB data, including polarization and lensing, *Planck18* (Planck Collaboration: Aghanim et al. 2020) obtained a 68% estimate of $\sigma_8 = 0.8111 \pm 0.0060$ for a ΛCDM cosmology[4] . The full-shape mod-

---

[4] Since CMB photons are scattered by the free electrons present after reionization, the CMB tem-

elling of the power spectrum in a flat ΛCDM cosmology applied to DESI DR1 clustering data yielded a 68% estimate of $\sigma_8 = 0.842 \pm 0.034$ (DESI Collaboration: Adame et al. 2025). The uncertainties tighten by one order of magnitude if external *Planck18* CMB data (temperature, polarization, and lensing) are included, giving $\sigma_8 = 0.8121 \pm 0.0053$. Interestingly, when a flat $w_0 w_a$CDM model is considered, results are compatible with the flat ΛCDM scenario: in this case, data are complemented with external SNIa datasets, obtaining 68% uncertainties of $\approx 9 \times 10^{-3}$, depending on the particular SNe dataset included in the fit.

We show some examples of connected and reduced 3PCFs for the new set of simulations, respectively in Fig. 4.2 and 4.3 (in redshift space and at $z = 0$), and compare them to the massive neutrino simulations at the same scales. We choose the same scales (s_{12}, s_{13}) of Fig. 3.10 and 3.11 and the same angular sampling for consistency. Also in this case, we convert the value of the angle θ between s_{12} and s_{13} into the corresponding length of the third side s_{23}. We subdivide the plots for each of the selected scales into two columns, showing ζ or Q of the `Mnu_ppp`, `Mnu_pp`, `Mnu_p`, and `fiducial_ZA` simulations in the left column and the `s8_p`, `s8_m`, and `fiducial` simulations in the right one. We also plot the corresponding detectabilities as a function of s_{23}.

Due to the small variation in σ_8, the effects on both ζ and Q are very small, and the data points in Fig. 4.2 and 4.3 for the `s8_p` and `s8_m` simulations are almost visually indistinguishable from the fiducial ones. Moreover, the effect that, at fixed halo minimum mass, varying σ_8 has on bias is responsible for reducing the differences in ζ between the `s8_p` and `s8_m` simulations. Indeed, the amplitude of the halo 3PCF is positively correlated both to the amplitude of the matter 3PCF and bias: while increasing σ_8 with respect to a fiducial value leads to a growth of the amplitude of the matter 3PCF, it conversely produces a decrease in the halo bias (for halos with fixed minimum mass).

For the chosen volume and scales, the detectabilities are $\lesssim 1$ across the entire s_{23} range. Their values show that for the selected scales, the considered variation of σ_8 is less or comparably detectable to a neutrino mass $M_\nu \sim 0.1\,\mathrm{eV}$.

By comparing the detectability values for the `s8_p` and `s8_m` simulations, it can be noticed that the data points for the `s8_p` and `s8_m` simulations are symmetric with respect to the fiducial ones. A mild growth of ζ for $s_{23} \approx |s_{13} - s_{12}|$ $(\theta \approx 0)$ and $s_{23} \approx s_{12} + s_{13}$ $(\theta \approx \pi)$ and a decrease at intermediate s_{23} for $\sigma_8 = 0.849$ (and vice versa for $\sigma_8 = 0.819$) is noticeable.

In Q, for all the selected scales, detectabilities indicate a flattening of Q for the lowest value of σ_8 (and vice versa, a steepening for the largest one) with respect to the fiducial case, suggesting that a massive neutrino signature in Q can be mimicked, in a ΛCDM cosmology, by a decrease of σ_8. Indeed, we anticipate here, and show in more detail in Sec. 4.4.2, that Q progressively flattens for more biased fields, as it occurs when σ_8 is smaller.

We stress that Fig. 4.2 and 4.3 show a low number of scales, so the above considerations are limited to them; a more detailed analysis exploiting triangle shapes, and based on the

---

perature power spectrum allows one to constrain the quantity $A_s e^{-2\tau}$, where τ is the optical depth of reionization; the optical depth τ can be well constrained by polarization data. CMB lensing instead is sensitive to the combination $\sigma_8 \Omega_m^{0.25}$.

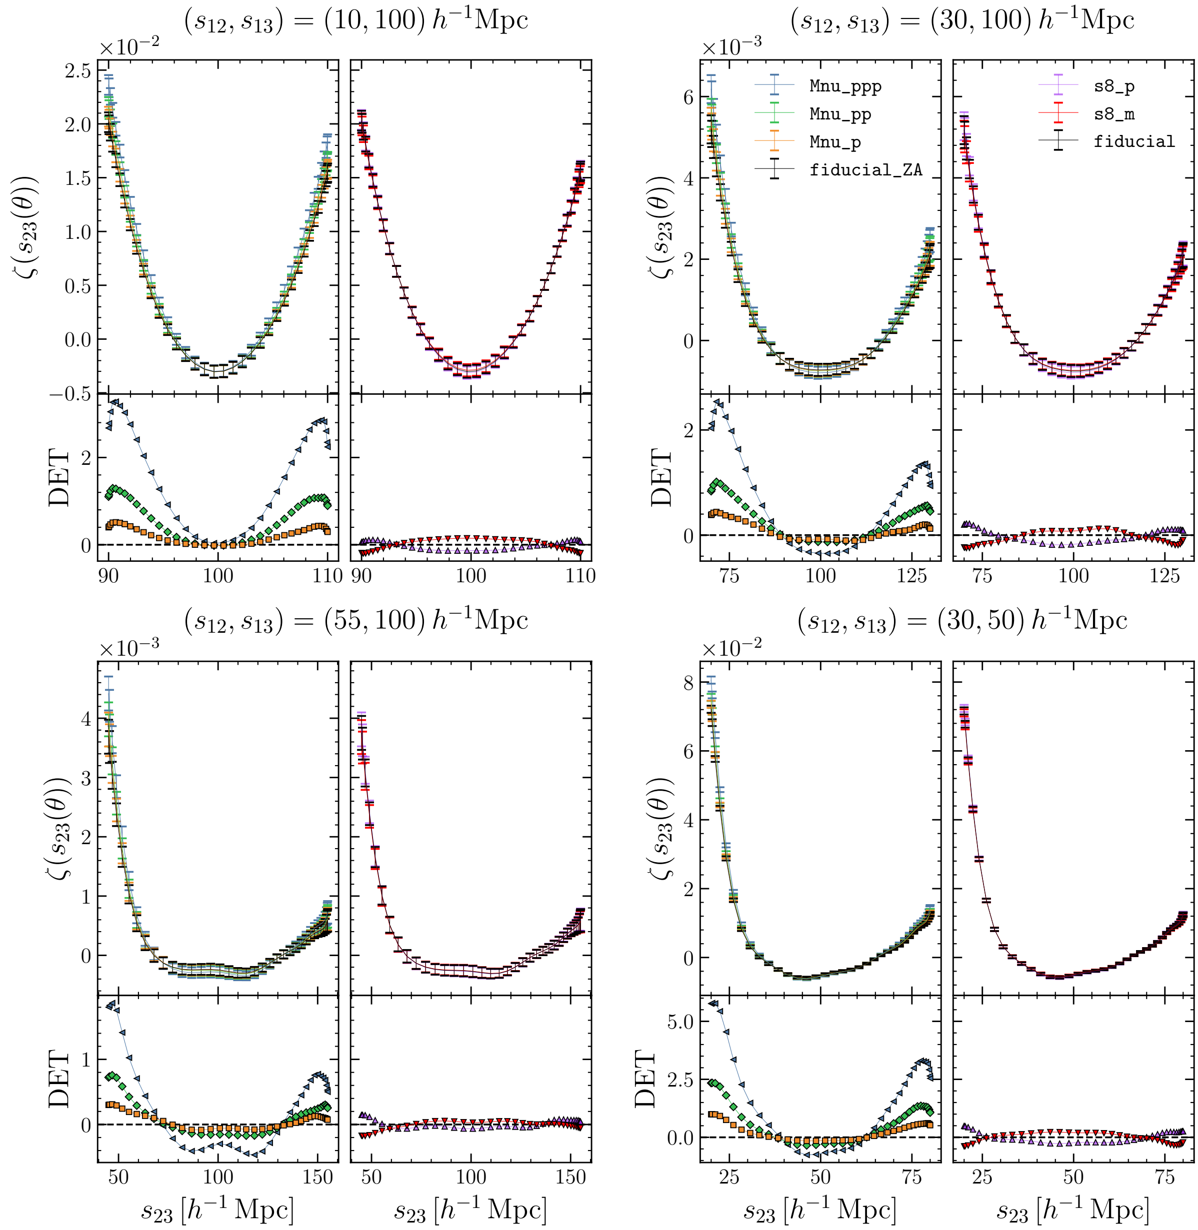

Figure 4.2: Comparison between the 3PCF of simulations with massive neutrinos and varying σ_8 simulations. The four panels show the redshift-space, single-scale 3PCFs ζ, computed at the same scales (s_{12}, s_{13}) of Fig. 3.10 (displayed at the top of each panel). The 3PCFs of the massive neutrino simulations `Mnu_ppp`, `Mnu_pp`, `Mnu_p`, together with `fiducial_ZA`, are plotted in the left column of each panel, while `s8_p`, `s8_m`, and `fiducial` simulations are in the right one. The corresponding detectabilities are shown in the second row of each panel, with the same color code as the above 3PCFs.

all-scales ζ and Q will be conducted in Sec. 4.5.

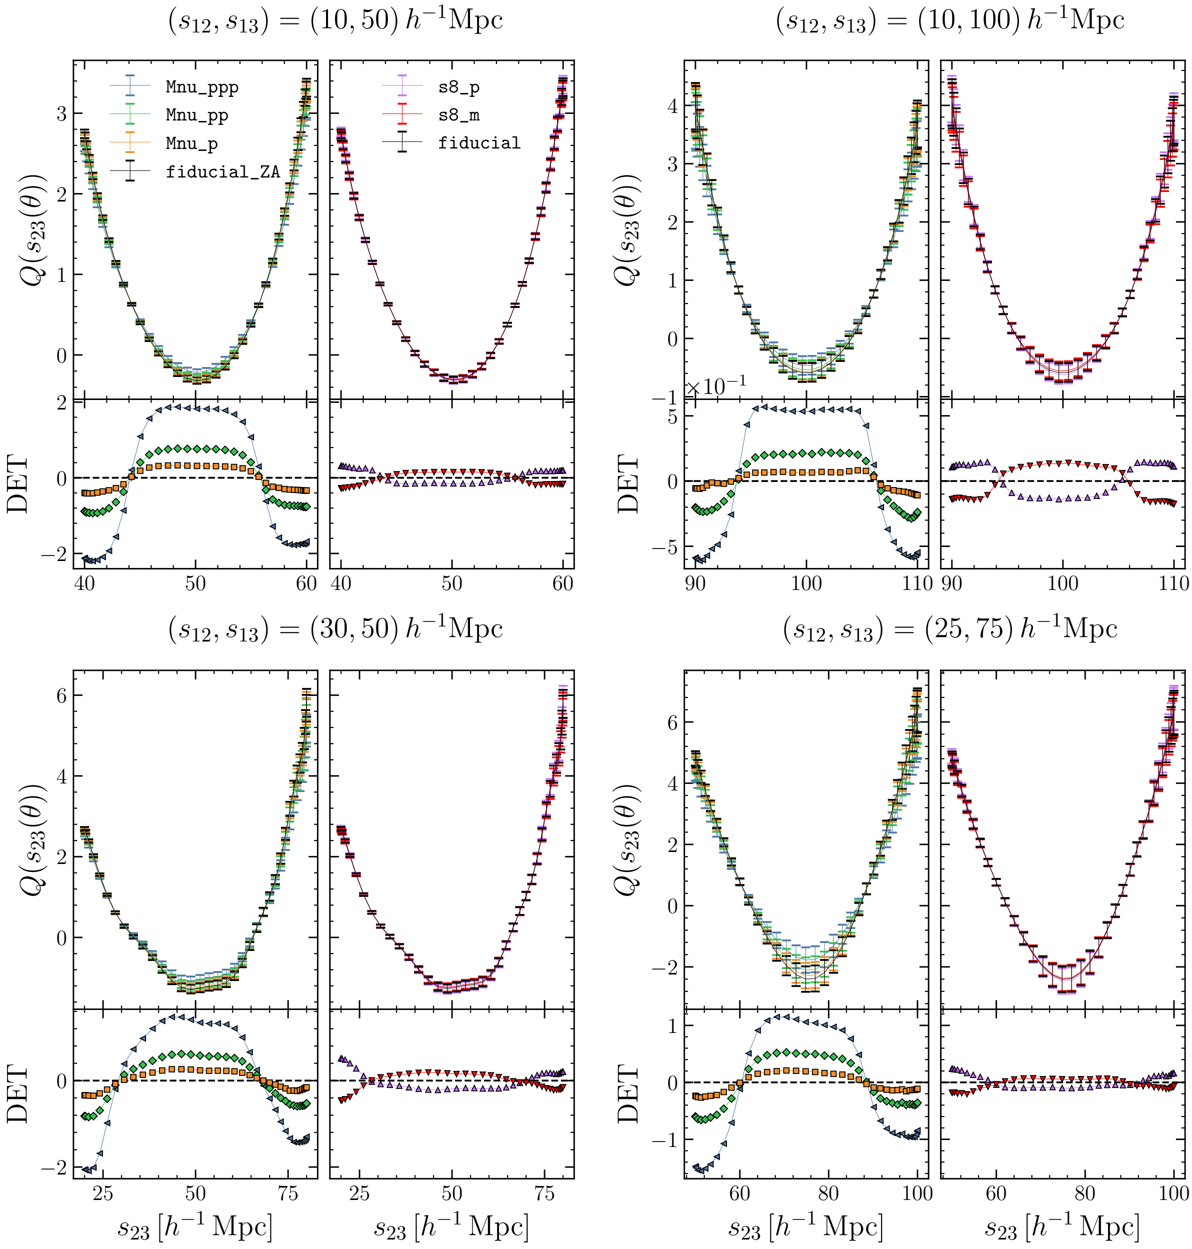

Figure 4.3: Same as in Fig. 4.2, but for the reduced 3PCF Q. The scales displayed in the figure are the same as those shown in Fig. 3.11.

4.4 The role of halo bias

The new dataset is characterized by the same halo minimum mass as the one analyzed in Ch. 3. As a consequence, the different values of σ_8 of the `s8_p` and `s8_m` simulations are responsible for different values of halo bias: in particular, for larger σ_8 with respect to the fiducial value, more halos are formed at a given redshift and the bias is lower (and viceversa for smaller values of σ_8). This complicates the analysis, because it forces us to include the bias in the space of degenerate parameters in addition to M_ν and σ_8, as we show in the following two sections for both the 2PCF and the 3PCF.

4.4.1 2-point statistics

We have quantified the bias as a function of scale for the `s8_p`, `s8_m`, and `fiducial` simulations by estimating it through Eq. 3.1.1, where for ξ_m we have taken the matter 2PCFs already available in the data products of the Quijote suite, analogously to what done for the massive neutrino simulations (Fig. 3.3). The real-space matter 2PCFs (averaged over all the realizations of each simulation) at $z = 0, 1, 2$ for the three simulations are shown in Fig. 4.4. The amplitude increase for decreasing redshift is a consequence of the fact that matter clustering becomes stronger over time, while at fixed redshift, the three different curves differ by a $\propto \sigma_8^2$ rescaling as expected.

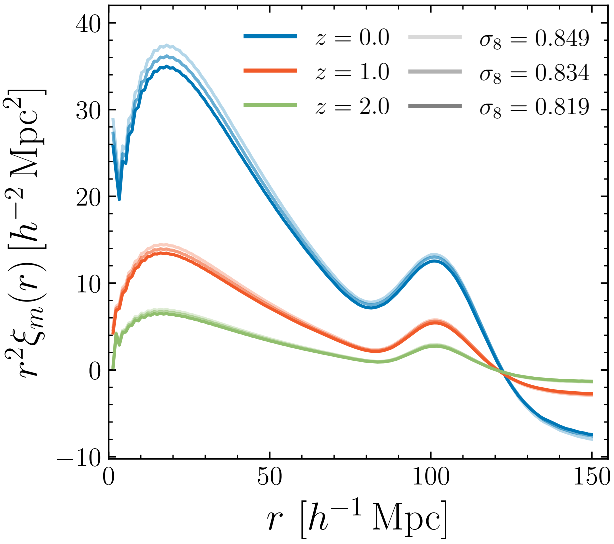

Figure 4.4: The average real-space matter 2PCFs for the `s8_p`, `s8_m`, and `fiducial` simulations, with $\sigma_8 = 0.849, 0.819$ and 0.834, respectively, at $z = 0, 1$, and 2 (blue, orange, and green curves, respectively). At fixed z, the 2PCFs of the different simulations are displayed with increasing opacity for decreasing σ_8. We have taken these measurements from the data products of the Quijote suite.

The values of bias for each simulation are shown in Fig. 4.5, separated into three different panels for each redshift. As expected and explained above, lower values of σ_8 lead to large values of bias. Compared to the analogous plots for the massive neutrino simulations (Fig. 3.3), b shows a lower scale dependence on pre-BAO scales, and is essentially constant for $25 \lesssim r[h^{-1}\,\mathrm{Mpc}] \lesssim 75$ at $z = 1$ and 2. The other features, like the increase of bias with redshift and the presence of a minimum at scales $\sim 85\,h^{-1}\,\mathrm{Mpc}$, have already been discussed in Sec. 3.1; we refer the reader to that section for details. The horizontal dotted lines in Fig. 4.5 correspond to the values of effective bias computed with the procedure outlined in Sec. 3.1, consisting of weighting the Tinker bias at a given mass

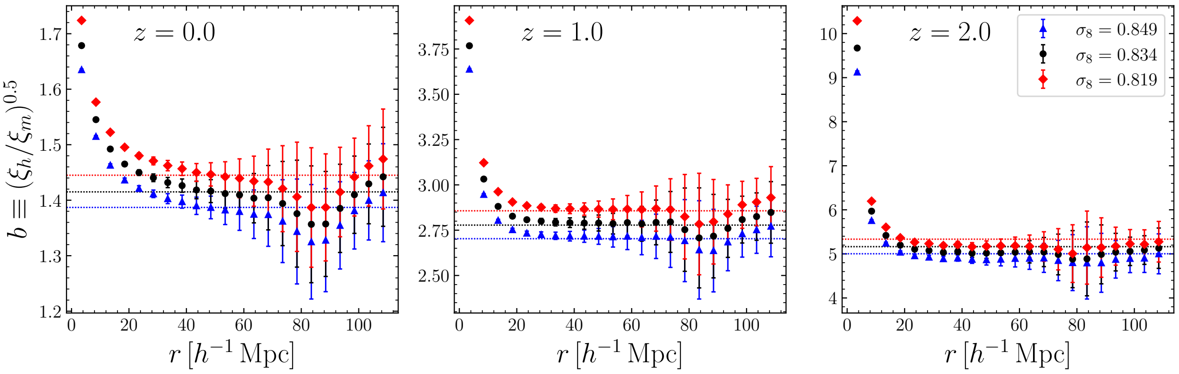

Figure 4.5: The halo bias computed as $b(r) = \sqrt{\xi_h(r)/\xi_m(r)}$, where ξ_h and ξ_m are the real-space halo and matter 2PCFs, for the `s8_p`, `fiducial`, and `s8_m` simulations (blue triangles, black circles and red squares, respectively). Each plot corresponds to a different redshift, from left to right: $z = 0, 1$, and 2. The associated error bars are computed from the errors in the halo 2PCFs, so they correspond to a volume of $10\,h^{-3}\,\mathrm{Gpc}^3$. The dotted lines correspond to the large-scale bias prediction obtained by weighting the Tinker bias with the average HMFs of the simulations (Eq. 3.1.2).

with the average HMF extracted from the simulations (Eq. 3.1.2). The values of effective bias show the best agreement at $z = 1$ above $30h^{-1}\,\mathrm{Mpc}$ and at $z = 0$ above $40h^{-1}\,\mathrm{Mpc}$. At $z = 2$, they slightly overpredict the measured bias between 20 and $60\,h^{-1}\,\mathrm{Mpc}$.

On sufficiently large scales, the matter and halo 2PCFs are, with good approximation, proportional, since bias is almost scale-independent and linear. This implies that the effects of a variation of σ_8 (i.e., of the normalization of the matter 2PCF) can be compensated by a variation of bias (increasing b for decreasing σ_8) to reproduce the same large-scale halo 2PCF. This gives rise to the "triple" degeneracy involving bias, M_ν, and σ_8, as claimed at the beginning of this section.

We illustrate this point in a simple but effective way by considering the matter 2PCFs of the `s8_p`, `s8_m`, and `fiducial` simulations, and the halo 2PCFs of the massive neutrino simulations. Given one matter and one halo 2PCF, we search for a value of linear bias b to rescale the former, such that the rescaled curve $b^2\xi_m$ overlaps (within errors) with the latter on large scales. This procedure is shown in Fig. 4.6 in real space and at $z = 0$, applied by considering the halo 2PCFs in the two extreme cases of the largest (0.4 eV) and smallest (0.1 eV) total neutrino masses (panels in the first and second row, respectively) and the matter 2PCFs for all the available values of σ_8 (changing across the panels in different columns). We have obtained the values of b by performing a least-square fitting in the range $r > 50\,h^{-1}\,\mathrm{Mpc}$; the residuals are shown in the bottom part of each panel in units of σ (only the errors on halo 2PCFs are considered): the dark and light gray bands represent the range where residuals are within, respectively, 1σ and 2σ.

Obviously, larger values of b for decreasing σ_8 are needed to fit the halo 2PCFs at fixed M_ν, since the amplitude of the matter 2PCFs decreases for decreasing σ_8: in particular, since the amplitude scales as σ_8^2, the best-fit values of b scale approximately as σ_8^{-1}. We find an excellent agreement between the rescaled matter 2PCF of the `s8_p`, `fiducial`,

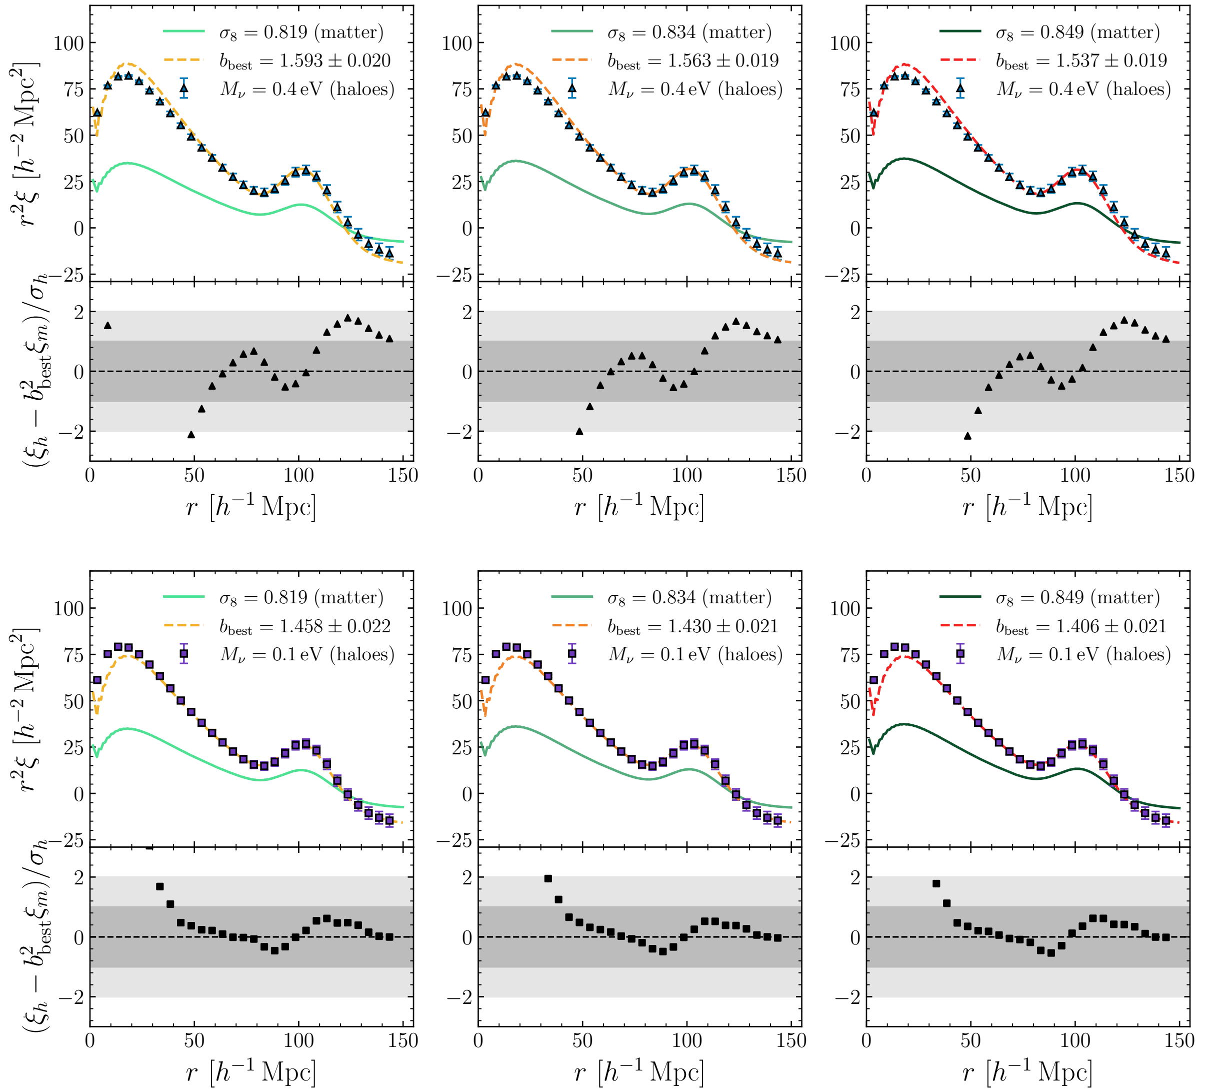

Figure 4.6: The upper plots in each panel show with a solid curve the real-space average matter 2PCF ξ_m of the ΛCDM simulations with the value of σ_8 indicated in the legend: $\sigma_8 = 0.819, 0.834$, and 0.849 for the panels located in the leftmost, central, and rightmost column, respectively. The markers correspond to the real-space average halo 2PCFs ξ_h for the massive neutrino simulations, shown for the smallest and largest available total neutrino mass $M_\nu = 0.1\,$eV and $M_\nu = 0.4\,$eV (panels in the lower and upper row, respectively). Here, the redshift is fixed to $z = 0$. The dashed line shows the least-squares best fit to the halo 2PCFs of the function $b^2\xi_m(r)$, in the large-scale range $r > 50\,h^{-1}\,$Mpc, and the best-fit values $b_{\rm best}$ are reported in the legend. The bottom plots in each panel show the difference between ξ_h and $b^2_{\rm best}\xi_m$, in units of the errors σ_h on the halo 2PCF. Dark and light gray bands correspond to 1σ and 2σ, respectively.

and `s8_m` simulations and the halo 2PCF for $M_\nu = 0.1\,$eV, with residuals always below 1σ in the fitting range. Increasing the neutrino mass to $M_\nu = 0.4\,$eV (significantly larger compared with the current estimates for the upper limit of the total neutrino mass) naturally leads to a worsening of the fit, with residuals systematically above 1σ on scales larger than the BAO peak; in the fitting range, residuals are within 2σ.

The key message that we want to convey with this discussion and Fig. 4.6 is that *on sufficiently large scales, the halo 2PCF in a massive neutrino cosmology with a fixed* σ_8 *(and a realistic value of* M_ν*) can be mimicked, inside the typical volume of a redshift bin of a Stage IV survey, by a* Λ*CDM cosmology with a different value of* σ_8 *through an appropriate linear bias rescaling.* In this sense, Fig. 4.6 gives an intuitive, but effective idea of the triple degeneracy projected onto the $b-M_\nu$ and $b-\sigma_8$ planes, when considering, respectively, the panels in the same column, where σ_8 is fixed, and in the same row, where instead M_ν is fixed.

4.4.2 3-point statistics: connected and reduced 3-point correlation function

We illustrate and compare the effects of variations in the halo bias and σ_8 on the connected and reduced 3PCF by exploiting the modeling implemented in the CBL. The library offers a complete implementation of the real-space, tree-level 3PCF model of Slepian and Eisenstein (2015b), based on perturbation theory, which is accurate on large scales and far from isosceles configurations. We describe the model in App. C, while here we report simply the assumed bias relation and the expressions for the halo connected and reduced 3PCF. The model assumes the following nonlinear bias relation between the halo and matter overdensity fields δ_h and δ_m:

$$\delta_h = b_1\delta_m + \frac{b_2}{2}\left(\delta_m^2 - \langle\delta_m^2\rangle\right), \tag{4.4.1}$$

where b_1 and b_2 are, respectively, the linear and nonlinear bias coefficients. Such a bias relation neglects possible tidal bias contributions. Using this approximation, it is possible to express the halo 3PCF as a function of the dark matter one as follows:

$$\zeta_h(r_{12}, r_{13}, \theta) = \left(\frac{\sigma_8}{\sigma_{8,\mathrm{fid}}}\right)^4 \left\{b_1^3\,\zeta_{m,\mathrm{fid}}(r_{12}, r_{13}, \theta) \; + b_1^2 b_2\left[\,\xi_{m,\mathrm{fid}}(r_{12})\xi_{m,\mathrm{fid}}(r_{13}) + \mathrm{cyc.\,perm.}\,\right]\right\}, \tag{4.4.2}$$

where $\zeta_{m,\mathrm{fid}}$ and $\xi_{m,\mathrm{fid}}$ are, respectively, the matter 3PCF and 2PCF for the fiducial model, while the factor multiplying the term in curly brackets is the ratio between the value of σ_8 in the considered cosmological model and in the fiducial model.

By adopting the same bias relation as in Eq. 4.4.1, the real-space halo reduced 3PCF Q_h is given by

$$Q_h(r_{12}, r_{13}, \theta) = \frac{1}{b_1}Q_m(r_{12}, r_{13}, \theta) + \frac{b_2}{b_1^2}, \tag{4.4.3}$$

where Q_m is the reduced matter 3PCF, i.e., the ratio between ζ_m and the term within square brackets of Eq. 4.4.2.

To illustrate the behavior of the models for different values of the parameters, we choose some physically plausible reference values for b_1, b_2 and σ_8, and show in Fig. 4.7

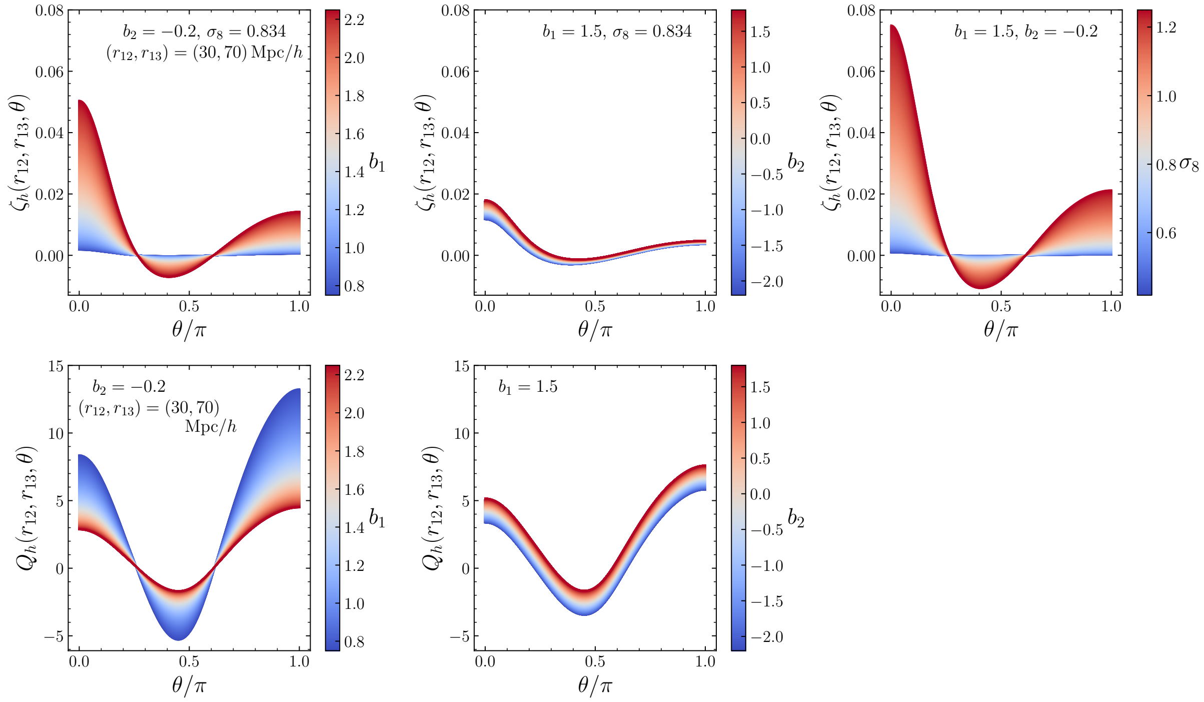

Figure 4.7: The effect of varying the linear and nonlinear bias parameters b_1 and b_2, and σ_8, on single-scale ζ (top row) and Q (bottom row), computed from the real-space isotropic multipoles of the model in Slepian and Eisenstein (2015b), at $z = 0$. Only the $0 \leq \ell \leq 10$ multipoles are included, consistently with our dataset. In each plot, a single parameter is varied at a time, with respect to the reference model $b_1 = 1.5$, $b_2 = -0.2$, $\sigma_8 = 0.834$ (in white in the colormaps). The varying parameter is indicated in the colorbar label, alongside its corresponding values, while the values of the frozen parameters are reported at the top of each plot. In the bottom row, the last column is left empty because Q does not depend explicitly on σ_8. The functions are plotted for an illustrative scale $(r_{12}, r_{13}) = (30, 70)\, h^{-1}\,\mathrm{Mpc}$.

the single-scale ζ and Q obtained from the CBL when these parameters are varied in an interval around the reference values. The model allows one to obtain the multipoles $\zeta_{h,\ell}$ of the halo 3PCF of Eq. 4.4.2. For consistency with our dataset, we sum the multipoles up to $\ell = 10$. In the plots we have fixed (r_{12}, r_{13}) to $(30, 70)\, h^{-1}\,\mathrm{Mpc}$ and chosen $b_1 = 1.5, b_2 = -0.2$, and $\sigma_8 = 0.834$ as reference values. Each plot shows the effect of varying the value of one parameter at a time, while keeping the others fixed at their reference values. We vary the parameters b_1 and σ_8 by 50% of their reference values, whereas we let b_2 take values between -2 and 2 to produce a visually appreciable effect.

We see that, also in the 3PCF, a variation of b_1 (mostly a rescaling) produces very similar effects to those of varying σ_8: for values of $b_1 > 1$, for which the amplitude of ζ_h scales roughly as $\sim \sigma_8^4 b_1^3$, a relative variation $\mathrm{d}b_1/b_1$ can be roughly compensated by a smaller (by a factor $\sim 3/4$) relative variation $\mathrm{d}\sigma_8/\sigma_8$ with opposite sign. This shows that a strong degeneracy exists between these two parameters.

Contrary to ζ_h, whose amplitude grows for increasing b_1, Q_h flattens because of the presence of the factor b_1^{-1} multiplying Q_m in Eq. 4.4.3. Incidentally, this indicates that differences between the "U" and "V" shapes of the single-scale Q also depend on bias, with "U" shapes characterizing more biased fields, and not only on what has been noticed in Moresco et al. (2014), i.e., the shape of the structures within which halos reside at a given scale. Varying b_1 also produces an offset effect, whose extent becomes less relevant as b_1 increases.

The same equation shows that the effect of changing the value of b_2 is a simple, angle-independent offset on Q_h. Conversely, ζ_h increases at all angles when b_2 increases (and vice versa when b_2 decreases), but its variation is more pronounced for $\theta \to 0$. This also has a simple explanation: the term involving the cyclic sum of ξ_m, which multiplies b_2 in Eq. 4.4.2, depends on θ only through r_{23}. When θ decreases, r_{23} decreases, leading to an increase of $\xi_m(r_{23})$: this, together with the fact that the chosen r_{12} and r_{13} guarantee that all the terms in the cyclic sum are positive (all the sides are well below the zero crossing of ξ_m), and consequently the entire cyclic sum increases for $\theta \to 0$, makes variations of b_2 more evident at smaller opening angles. However, the contribution of b_2 is much smaller compared to b_1 and/or σ_8; this can be more generally seen at the level of the bias-weighted post-cyclic multipoles shown in Fig. 9 in Slepian and Eisenstein (2015b).

4.5 Sensitivity to triangle shapes

We apply the same kind of analysis, presented in Hahn et al. (2020) for the bispectrum (see Sec. 4.2.2), in configuration space, to study whether it is possible to disentangle the effects produced on the halo distribution by massive neutrinos and σ_8 through the triangle shape dependence of both ζ and Q.

For our new dataset, we have computed, as in Sec. 3.4, the absolute values of the detectability (Eq. 2.5.1) for the same set of triangle shapes, averaging over the contributions of triangles with the same shape but different sizes. The results are shown in Fig. 4.8, following the same layout of Fig. 4.1: we show two panels, the upper one for ζ and the lower one for Q, each composed of five shape plots (for an accurate description of this kind of plots, see again Sec. 3.4). The plots in the top row of each panel correspond to the massive neutrino simulations `Mnu_p`, `Mnu_pp`, and `Mnu_ppp`, while the ones in the bottom row to the `s8_m` and `s8_p` simulations. In the figure, detectabilities are computed for the entire set of available triangles, without constraints on scale or configuration, with the usual upper limit $s_{23} \leq 110\, h^{-1}$ Mpc for Q adopted throughout this work.

Examining this figure, we can make a few preliminary considerations, leaving more quantitative analysis for the subsequent discussion. We notice that, for ζ, also in the case of the `s8_p` and `s8_m` simulations, the largest detectabilities correspond to elongated triangles (squeezed/folded). However, the effects of massive neutrinos on these triangles appear stronger and more extended away from the left oblique side of the plot than those resulting from a ±0.015 variation in σ_8, even for the lowest mass $M_\nu = 0.1\,\mathrm{eV}$. Interestingly, in Q, the feature that obliquely crosses the colormaps from the upper-left vertex to the midpoint of the opposite side for $M_\nu \neq 0$ (which we have associated with right-angled triangles in the previous chapter), seems to be absent for `s8_p` and `s8_m`. Moreover, for these two simulations, the imprint on elongated shapes appears more

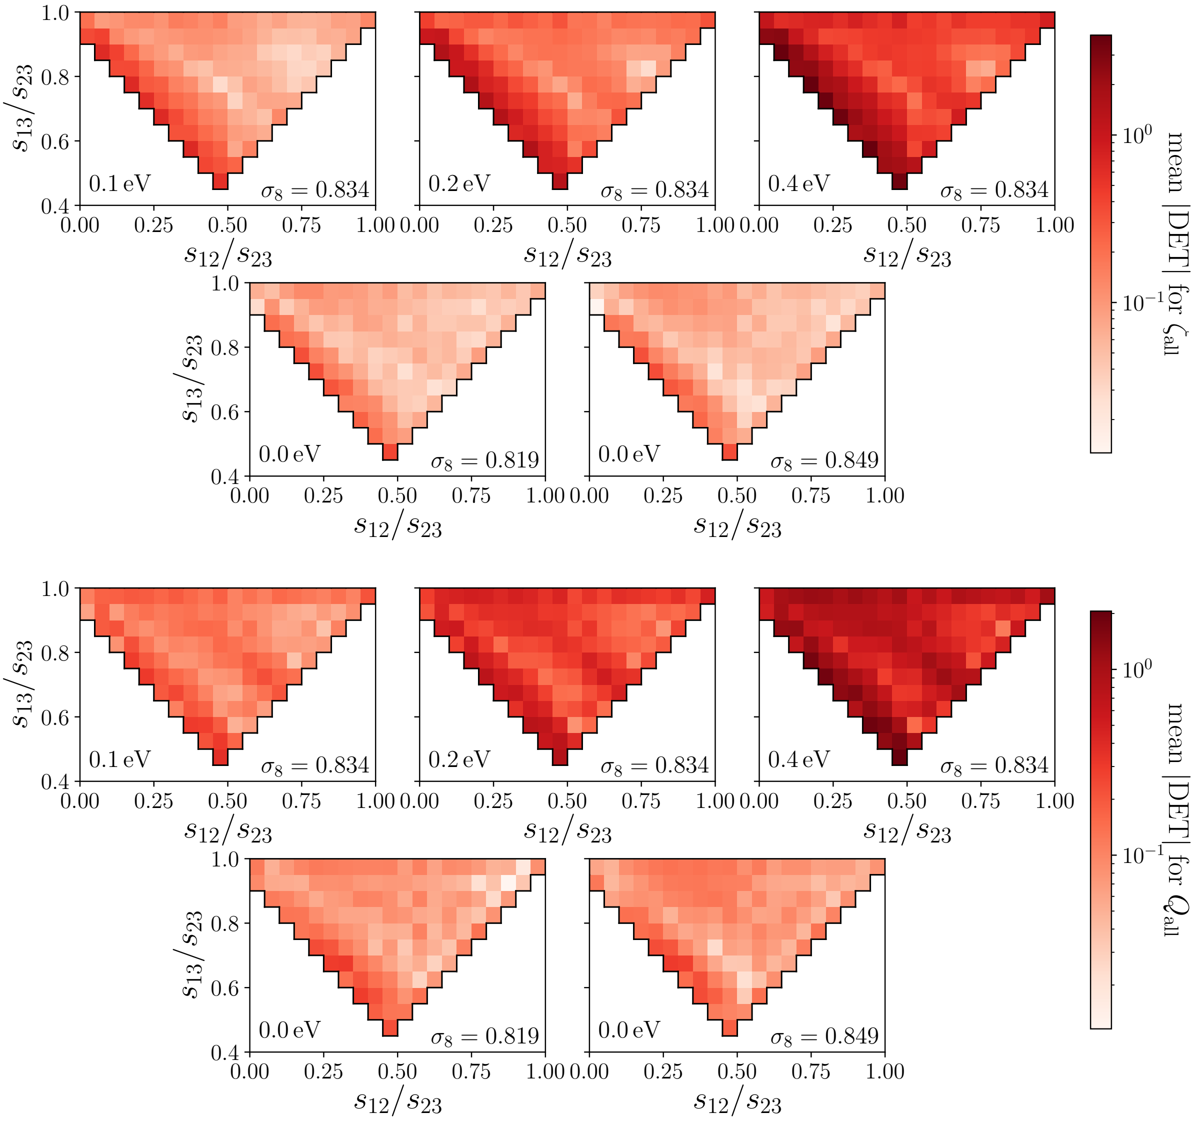

Figure 4.8: The detectability of the redshift-space, $z = 0$ halo connected 3PCF (*top panel*) and reduced 3PCF (*bottom panel*) as a function of triangle shape, for the simulations with massive neutrinos, and variable σ_8. The triangle shapes are identified by the side ratios s_{12}/s_{23} and s_{13}/s_{23}, with $s_{12} \leq s_{13} \leq s_{23}$. In each panel, a given shape bin displays the absolute value of detectability averaged over all triangles available in the all-scales approach with that shape and different sizes. The white regions correspond to pairs $(s_{12}/s_{23}, s_{13}/s_{23})$ with which it is not possible to construct a triangle. In both groups of five plots, from left to right, the top row corresponds to the simulations `Mnu_p`, `Mnu_pp`, and `Mnu_ppp`, while the bottom row corresponds to the simulations `s8_m` and `s8_p`.

concentrated at smaller values of s_{13}/s_{23}.

To quantitatively determine how the results change by excluding isosceles or quasi-isosceles configurations or particular scales, we analyze these plots by applying several cuts in $\eta_{\rm min}$ and select different ranges of s_{23}, extracting the mean and standard deviations of

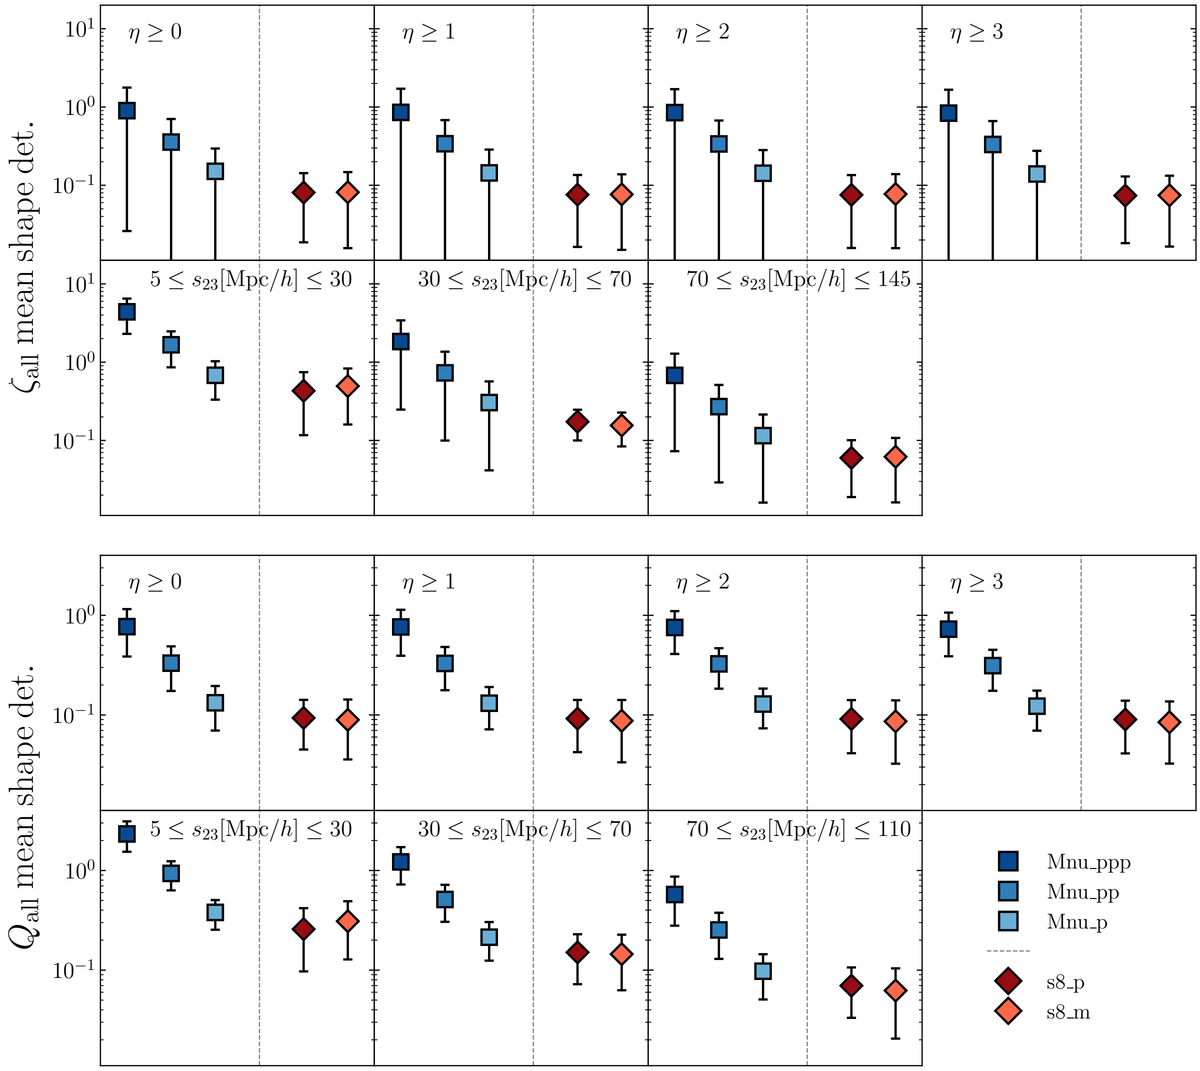

Figure 4.9: Mean values (markers) and standard deviations (errorbars) of the values of shape detectability measured from Fig. 4.8, for the connected 3PCF (upper panel), and the reduced 3PCF (lower panel). Simulations are distinguished by color and marker, as shown in the legend in the lower-right corner. In each plot, the vertical dashed line separates the simulations `Mnu_ppp`, `Mnu_pp`, and `Mnu_p` from `s8_p` and `s8_m`. In each panel, the top and bottom rows show, respectively, the results for different lower bounds on η (from 0 to 3), and for different ranges of s_{23} ([5,30], [30,70], and [70, $s_{23,\rm max}$] h^{-1} Mpc, with $s_{23,\rm max} = 145$ or $110\,h^{-1}$ Mpc for ζ and Q, respectively).

the values displayed in the corresponding shape plot. While the means provide an overall amplitude of the detectability level from triangle shapes (i.e., the possibility to distinguish an effect due to neutrinos), the standard deviations quantify how much the signal changes between different triangle shapes.

We explored values of $\eta_{\rm min}$ in the range $0 \leq \eta_{\rm min} \leq 3$ and s_{23} in the ranges $[5, 30]$, $[30, 70]$, and $[70, s_{23,\rm max}]\ h^{-1}$ Mpc, with $s_{23,\rm max} = 145$ or $110\,h^{-1}$ Mpc respectively for ζ and Q. This allowed us to test whether the neutrino signal provides peculiar imprints at given scales that are not reproducible by a σ_8 variation.

In Fig. 4.9 we plot, for each choice of $\eta_{\rm min}$ or s_{23} range, the means (square markers) of the redshift-space, $z = 0$ shape detectabilities for all the simulations. The corresponding standard deviations are displayed using a bar on each marker. In the figure, the two groups of simulations are identified with different colors.

As expected, for increasing M_ν, the means follow the growth of neutrino detectabilities, becoming larger regardless of the particular choice of $\eta_{\rm min}$ or the s_{23} range. On the other hand, the variation in σ_8 of ± 0.015 provides similar (since `s8_p` and `s8_m` present almost the same means in each subplot) and significantly smaller detectabilities compared to the M_ν case. For any choice of $\eta_{\rm min}$ or s_{23} range in the figure, the means for the `Mnu_ppp`, `Mnu_pp`, and `Mnu_p` simulations are larger than the ones for the `s8_p` and `s8_m` simulations, even if only slightly for $M_\nu = 0.1\,\mathrm{eV}$ (compare the light blue markers with the dark red and orange ones). For a given simulation, regardless of the group, the overall detectability level appears quite insensitive to variations of $\eta_{\rm mean}$ within the chosen range, while reducing with growing s_{23}, as indicated by the nearly absent dependence on $\eta_{\rm min}$ of the means and by their decrease with increasing scale.

For varying $\eta_{\rm min}$, we have $\sigma(\texttt{s8\_p}) < \sigma(\texttt{s8\_m})$[5] for both ζ and Q, with the difference decreasing for Q as $\eta_{\rm min}$ increases. However, the differences are always within 5% and 11% of the mean standard deviation for ζ and Q, respectively. For ζ we have $\sigma(\texttt{s8\_p}) < \sigma(\texttt{s8\_m})$ for small and large values of s_{23}, (with differences of $\approx 7\%$ and 11%, respectively), while at intermediate scales $\sigma(\texttt{s8\_p}) > \sigma(\texttt{s8\_m})$, even if only slightly ($\approx 2.5\%$). For Q, $\sigma(\texttt{s8\_p}) < \sigma(\texttt{s8\_m})$ always, but with differences at small and large scales larger than at intermediate ones ($\approx 13\%$ versus 4%). For the massive neutrino simulations, we always have $\sigma(\texttt{Mnu\_p}) < \sigma(\texttt{Mnu\_pp}) < \sigma(\texttt{Mnu\_ppp})$, for both ζ and Q, regardless of the value of $\eta_{\rm min}$ or the s_{23} range.

Comparing the variation in the detectability of the two groups, we derive interesting conclusions. To better quantify this, we complement Fig. 4.9 with Tab. 4.1, in which we report, as a function of $\eta_{\rm min}$ and s_{23} range, the ratios between the standard deviations of the `Mnu_ppp`, `Mnu_pp`, and `Mnu_p` and the average standard deviation between the `s8_p` and `s8_m` simulations shown in Fig. 4.9. Furthermore, in order to identify more clearly the triangle shapes responsible for the different behaviors when comparing the two groups of simulations for varying constraints on η and s_{23}, we have also generated the shape plots of ζ and Q for all the values of $\eta_{\rm min}$ and s_{23} bin. Although these plots are undoubtedly informative, we deemed that including them in the main text could overburden the discussion. Therefore, the interested reader is invited to consult them in App. D, where we have chosen to compare the `s8_m` simulation with the `Mnu_p` simulation, being characterized, as discussed above, by similar values of mean detectability. The plots are grouped by s_{23} bin and values of $\eta_{\rm min}$ in Fig. D.1 and D.2, respectively.

We outline here our main findings.

- Qualitatively, the fact that the pattern in detectability induced by neutrinos and σ_8 differs means that a combination of different triangle shapes will be impacted dif-

---

[5] Here $\sigma(\texttt{sim})$ denotes the standard deviation for a given simulation `sim`.

$\sigma(\texttt{Mnu})/[0.5\,(\sigma(\texttt{s8\_p}) + \sigma(\texttt{s8\_m}))]$

	$\eta_{\rm min}$				s_{23} range (h^{-1} Mpc)		
	0	1	2	3	[5,30]	[30,70]	[70,145]
`Mnu_ppp`	13.64	14.16	13.90	14.55	6.45	21.85	13.97
`Mnu_pp`	5.41	5.63	5.52	5.79	2.49	8.68	5.57
`Mnu_p`	2.25	2.34	2.30	2.39	1.07	3.63	2.29

(a)

	$\eta_{\rm min}$				s_{23} range (h^{-1} Mpc)		
	0	1	2	3	[5,30]	[30,70]	[70,110]
`Mnu_ppp`	7.53	7.19	6.67	6.70	4.59	6.18	7.52
`Mnu_pp`	3.08	2.94	2.74	2.73	1.77	2.57	3.15
`Mnu_p`	1.23	1.15	1.07	1.05	0.73	1.12	1.19

(b)

Table 4.1: Each table reports, for ζ (a) and Q (b), the standard deviations of the massive neutrino simulations normalized by the average of the standard deviations of the simulations `s8_p` and `s8_m` derived from Fig. 4.9. The values are reported as a function of the choices of $\eta_{\rm min}$ and s_{23} adopted in Fig. 4.9, displayed at the top of each table.

ferently by these two parameters, and hence can be used to discriminate a neutrino effect.

- The ratios reported in Tab. 4.1 are overall larger for ζ than for Q, while for Q the standard deviations of `s8_p` and `s8_m` are comparable to the `Mnu_p` ones. This points toward a larger signal variation expected in ζ than in Q. For ζ, the ratios are up to a factor 20 (for the more massive neutrinos at intermediate scales), suggesting a well distinguishable imprint of the neutrinos on the 3PCF.

- The ratios are almost stable for ζ at increasing $\eta_{\rm min}$, while they decrease for Q. This is due to the fact that in Q a particular range of shapes contributes to the detectability (the diagonal regions in Fig. 4.8, in particular at intermediate scales, corresponding to folded and right-angled triangles), which gets progressively discarded with the increasing cut in $\eta_{\rm min}$.

- The ratios are significantly larger at intermediate and large scales rather than at small separations s_{23}. While the overall amplitude is more significant for ζ rather than Q, this behavior can be observed in both. This suggests that at those scales in particular, the effect of neutrinos cannot be mimicked by a σ_8 variation.

The results discussed in this section, although certainly not definitive, are encouraging, as they suggest, at least for the variations in σ_8 considered, that the effect of such variations on ζ and Q impacts triangles in configuration space, depending on their shape and scale, differently from massive neutrinos. This outcome therefore aligns with what has been found by Hahn et al. (2020), when examining larger variations in σ_8 in Fourier space, and it opens the way to more robust analyses in configuration space in the future, possibly considering larger variations in σ_8 and controlling bias effects, for instance through a selection of the number of halos that keeps the bias fixed.

4.6 Summary of the main results

In this section, we summarize the main results obtained in the analysis presented in this chapter.

Degeneracy between linear bias, M_ν, and σ_8. In the case of dark matter halos, the study of the $M_\nu - \sigma_8$ degeneracy is complicated by the presence of the bias as an additional degenerate parameter.

For the halo 2PCF, on large scales (we consider $s > 50\,h^{-1}\,\mathrm{Mpc}$) the effect of massive neutrinos is degenerate with a simple linear bias rescaling of the matter 2PCF of a massless neutrino cosmology, even for Stage IV surveys (see Fig. 4.6). This degeneracy reduces for increasing M_ν due to additional shape effects on the 2PCF induced by massive neutrinos, but remains within 2σ for $M_\nu = 0.4\,\mathrm{eV}$ (our largest mass).

Since in our analysis the minimum halo mass is fixed ($M_{\mathrm{min}} \approx 1.3\times10^{13}\,h^{-1}\,M_\odot$), our simulations are characterized by a different halo bias (see Fig. 4.5). For this reason, our analyses of both ζ and Q are dependent on variations of σ_8 (only indirectly for the latter), even if Q should formally be independent of σ_8 (but it depends on bias parameters, see Eq. 4.4.3). It is particularly interesting to note that the single-scale ζ and Q shown in Fig. 4.2 and 4.3 prove that a $+0.015$ increase in σ_8 for ζ or a -0.015 decrease for Q can potentially mimic the signature of neutrinos of mass $\sim 0.1\,\mathrm{eV}$, but for larger masses the two effects can be clearly discriminated.

Disentangling the effect of M_ν and σ_8 with the 3PCF. We explored the possibility of distinguishing an effect of the variation of M_ν and σ_8 by quantifying its impact on different triangle shapes in ζ and Q (see Sec. 4.5).

For the reduced 3PCF Q, we find that in massive neutrino cosmologies, a significant contribution to the detectability variance comes from elongated and right-angled triangles; this feature appears to be absent in the ΛCDM simulations with varying σ_8 we considered, pointing toward the possibility of using these triangles to discriminate between the two effects. Moreover, we find a different response as a function of scale of the detectability in both ζ and Q, when comparing massive neutrinos and varying σ_8 simulations.

Our results show that, also in configuration space, 3-point statistics respond differently (as a function of triangle shape) to variations in M_ν and σ_8, making them a promising tool for breaking the degeneracy between these two parameters.

Chapter 5

Conclusions and future prospects

> *Everything comes to him who knows how to wait.*
>
> Text of the telegram sent by Wolfgang E. Pauli to Frederick Reines and Clyde Cowan after the first detection of neutrinos, 1956

Stage IV spectroscopic surveys (Euclid Collaboration: Mellier et al. 2025; DESI Collaboration: Aghamousa et al. 2016; Dore et al. 2019) are expected to map galaxy positions on an unprecedented volume, thereby enabling significant improvements in constraining cosmological parameters through galaxy clustering studies.

One of these parameters is the sum of neutrino masses from the Cosmic Neutrino Background. Constraining it is important not only to shed light on the expansion history of the Universe and the growth of cosmic structures, but also because it provides crucial insights into beyond-the-Standard-Model particle physics. Since the confirmation that neutrinos are massive (Fukuda et al. 1998), cosmology has been able to place the most stringent upper limits on the sum of their masses by combining several probes (see Particle Data Group: Navas et al. 2024 for a summary), including clustering, affected by the phenomenon of non-relativistic neutrino free-streaming (Bond et al. 1980).

Constraints from clustering come from the extensive study and modeling of massive neutrino effects in 2-point statistics, both in configuration (2PCF) and in Fourier space (power spectrum), including the study of small-scale power spectrum suppression, halo mass function, bias, and redshift-space distortions through simulations (e.g., Viel et al. 2010; Marulli et al. 2011; Costanzi et al. 2013; Castorina et al. 2014; Villaescusa-Navarro et al. 2014; Castorina et al. 2015; Villaescusa-Navarro et al. 2018; Verdiani et al. 2025).

Tackling the problem within the framework of higher-order statistics is timely, as they are sensitive to nonlinearity characterizing small scales, at which neutrinos leave their imprint through free-streaming. Advances in algorithms for estimating the 3PCF (Slepian and Eisenstein 2015a; Slepian and Eisenstein 2018) have significantly reduced computational costs, enabling its measurement across a wide range of scales. The information content of the 3PCF is equivalent to that of its Fourier-space counterpart, the bispectrum. However, being defined in configuration space, the 3PCF presents several advantages, offering a particularly suitable tool in view of current and future surveys, accounting more naturally for geometric effects.

Despite this potential, a very limited number of studies have been carried out so far, and these have focused only on the bispectrum (Ruggeri et al. 2018; Hahn et al. 2020).

In this thesis, we have studied for the first time the halo connected 3PCF ζ and reduced 3PCF Q in simulations of massive neutrino cosmologies. Q complements the information gained from ζ because it optimally combines it with the 2PCF, normalizing the former by the cyclic permutation over the triangle sides of the products of the latter at consecutive triangle side pairs.

In particular, we have measured the isotropic 3PCF multipoles ζ_ℓ of the Szapudi (2004) decomposition, in which the 3PCF at a given triangle, parameterized by two sides (s_{12}, s_{13}) and the angle θ between them, is expanded into Legendre polynomials in $\cos\theta$. We have estimated the multipoles up to $\ell_{\rm max} = 10$, together with the monopole, quadrupole, and hexadecapole of the 2PCF, for a set of 2 000 halo catalogues extracted from the QUIJOTE N-body simulations (Villaescusa-Navarro et al. 2020). The simulations we considered are characterized by values of the sum of neutrino masses $M_\nu = 0.0, 0.1, 0.2$, and $0.4\,\mathrm{eV}$, and halo minimum mass to $1.3 \times 10^{13}\,h^{-1}\,M_\odot$ (set by resolution). We have estimated the covariance numerically, on a Stage IV-like volume of $10\,h^{-3}\,\mathrm{Gpc}^3$ from an additional set of 2 000 fiducial simulations.

We have run our measurements with the code `MeasCorr` (Farina et al. 2024), probing a wide range of scales, from 1 to $150\,h^{-1}\,\mathrm{Mpc}$ (with $\Delta r = 1\,h^{-1}\,\mathrm{Mpc}$) for the 2PCF and from 5 to $145\,h^{-1}\,\mathrm{Mpc}$ (with $\Delta r = 5\,h^{-1}\,\mathrm{Mpc}$) for the 3PCF, in the latter case thanks to the implementation of the Slepian and Eisenstein (2015a) spherical harmonic decomposition estimator.

We have developed a framework to quantify the neutrino signal over the massless case, including a detectability defined as the ratio between the mean differences of ζ and Q in the massive and massless neutrino cases and the standard deviation on the chosen volume, and a χ^2-like detectability obtained by contracting the precision matrix with these differences.

In the last part of this work, we have applied part of this framework to study whether ζ and Q can break the degeneracy between M_ν and σ_8, affecting 2-point statistics. We have thus complemented our previous dataset (in which σ_8 is fixed at the fiducial value 0.834) with the same kind of measurements performed on three sets of 500 QUIJOTE massless neutrino simulations each, with $\sigma_8 = 0.819, 0.834$, and 0.849.

Results

We summarize below the key results obtained in our analysis. For ζ we considered scales up to $145\,h^{-1}\,\mathrm{Mpc}$, while for Q we considered scales up to $110\,h^{-1}\,\mathrm{Mpc}$, to avoid zero crossings of its denominator.

Imprint of neutrinos on the 3PCF multipoles. We find that the neutrino detectability increases moving from large to small scales for all the multipoles ζ_ℓ, compatibly with the fact that neutrinos affect the clustering of dark matter halos at small scales due to the phenomenon of free-streaming. The multipole most affected by the presence of massive neutrinos, even at the largest scales probed by our measurements, is $\ell = 2$, containing most of the shape information encoded in the 3PCF. For higher multipoles, neutrinos become progressively detectable only on quasi-isosceles triangles. We observe an increase in variance for isosceles triangles, potentially related to autocorrelation effects.

The effect of neutrinos from small to intermediate scales. We have studied the behavior of ζ and Q for all the possible combinations of two fixed side pairs (s_{12}, s_{13}), as a function of the angle between them.

For ζ, we find that the effect of neutrinos (assessed with our χ^2 indicator) has a larger impact for isosceles triangles below $\sim 30\,h^{-1}\,\mathrm{Mpc}$, and for quasi-isosceles (up to $\eta \equiv |s_{12} - s_{13}|/\Delta s \approx 4$) and squeezed triangles ($s_{12} \ll s_{13}$). For Q, we observe a maximum for isosceles triangles, with detectabilities overall smaller with respect to ζ. In correspondence with these configurations, we find that the concavity of ζ increases, while Q flattens for growing M_ν.

Given the general trend of decreasing detectability going to large scales, we have also predicted the largest scale at which, for the selected configurations, a mass M_ν at a given redshift is detectable above $1, 3$, and 5σ in the chosen volume. As a common trend, a given significance threshold is reached at smaller scales for decreasing mass or at lower redshifts for fixed mass, showing that the signal follows the evolution of nonlinearity. However, we find that the only masses that are detectable above 1σ are $M_\nu = 0.2$ and $0.4\,\mathrm{eV}$. The only mass detectable also above 5σ is the largest one $M_\nu = 0.4\,\mathrm{eV}$ (which, however, is significantly larger than the current estimates of the upper limits on the sum of neutrino masses).

The effect of neutrinos on filaments and BAO. The changes in the shape of ζ in the presence of massive neutrinos described above indicate that neutrinos mostly impact the filaments of the cosmic web. In particular, for fixed σ_8, as in the case of our simulations, an increase in M_ν produces an overdensity of triplets along them.

Although, due to free-streaming, most of the signal from massive neutrinos comes from small scales, our analysis proves that BAO scales are also affected by their presence, but these effects are largely undetectable within the volume probed by Stage-IV surveys. Considering a full effective volume of $500\,h^{-3}\,\mathrm{Gpc}^3$ (the largest available in our simulations), we detect a shift towards larger scales of the zero crossing of the 3PCF dipole ($\ell = 1$): since this multipole is sourced by the gradient of the density field, this shift corresponds to a variation of the BAO position.

Triangle shapes and configurations maximizing neutrino detectability. We have studied the dependence of ζ and Q also on all the possible side-binned triangles, adopting the ordering $s_{12} \leq s_{13} \leq s_{23}$. In particular, we have examined the dependence of detectability on shape by averaging its absolute values over triangles with the same shape but different sizes.

We find in ζ the highest signal for elongated triangles, confirming that filamentary shapes contain the largest signal. We explored how the signal varies at different scales (small-intermediate $5 \leq s_{23}\,[h^{-1}\,\mathrm{Mpc}] \leq 70$ and large $70 \leq s_{23}\,[h^{-1}\,\mathrm{Mpc}] \leq 145$), finding a preference for squeezed triangles when including larger scales. For Q, we find that a complementary source of signal, with respect to ζ, is represented by right-angled triangles ($s_{23}^2 \approx s_{12}^2 + s_{13}^2$).

The study of the cumulative detectability on the triangles with sides larger than a given threshold s_{min}, for decreasing values of s_{min}, quantitatively confirms that it is crucial to include progressively smaller scales at lower redshift, especially $s \lesssim 30\,h^{-1}\,\mathrm{Mpc}$, in future modeling and survey analysis.

Assessing the role of bias in the $M_\nu - \sigma_8$ degeneracy. Since in our dataset M_{min} is fixed, an increase in σ_8 produces a decrease in the halo bias. This makes Q indirectly dependent on σ_8 through its effects on halo bias.

If bias is variable, it enters the space of degenerate parameters. Considering scales above $50\,h^{-1}\,\mathrm{Mpc}$, we find that the halo 2PCF for $M_\nu = 0.1\,\mathrm{eV}$ is very well reproduced, within the errors for the considered volume of $10\,h^{-3}\,\mathrm{Gpc}^3$, by a linear bias rescaling of the of total matter 2PCF of the varying σ_8 simulations (taken from the Quijote data products). By studying the dependence of the single-scale ζ and Q on σ_8 and on the linear and nonlinear bias parameters b_1 and b_2 in the Slepian and Eisenstein (2015b) model, we find a strong degeneracy in ζ between σ_8 and b_1. The flattening of Q for increasing b_1 shows that on a sample of halos with fixed M_{min}, the imprint of neutrinos can be mimicked by lower values of σ_8 (producing larger biases).

Breaking the $M_\nu - \sigma_8$ degeneracy with the 3PCF. We have computed the neutrino detectability of ζ and Q considering different shapes and scales as in the analysis presented above, and compared it with the effect induced on the 3PCF due to a σ_8 variation. We considered different configurations, namely $0 \leq \eta_{\mathrm{min}} \leq 3$ or scales, with $s_{23}\,[h^{-1}\,\mathrm{Mpc}]$ in the ranges $[0, 30]$, $[30, 70]$, $[70, s_{23,\mathrm{max}}]$, with $s_{23,\mathrm{max}} = 145$ and $110\,h^{-1}\,\mathrm{Mpc}$ respectively for ζ and Q.

We find that, for Q, right-angled triangles exhibit the highest potential for breaking the $M_\nu - \sigma_8$ degeneracy, as they produce a pronounced response to variations in the neutrino mass, while being comparatively less sensitive to changes in σ_8. We also find differences in the configurations affected by variations in M_ν and σ_8 in both ζ and Q as a function of scale.

We notice that a major limitation arises from the small variation in σ_8 considered, due to the available simulations. Should this result be confirmed also with a larger difference in σ_8, the 3PCF could provide a promising tool to break the M_ν–σ_8 degeneracy through distinct dependences on triangle shape, in line with what has been similarly suggested for the bispectrum in Hahn et al. (2020). However, the result appears promising since the range of σ_8 explored is compatible with current constraints on this parameter available from the analysis of real data.

This work represents a crucial step in the framework of higher-order statistics, as it shows, for the first time in the literature for the 3PCF, the possibility of employing it as a cosmological tool beyond the standard ΛCDM model. In particular, it highlights its potential for constraining the sum of neutrino masses. This is fundamental not only in preparation of the data that will be released in the near future by the already mentioned Stage-IV surveys (first and foremost Euclid, Euclid Collaboration: Mellier et al. 2025), but also, in the longer term, for the future Stage-V surveys. An example is represented by the Widefield Spectroscopic Telescope (WST; Mainieri et al. 2024), which aims to provide an unprecedented map of the Universe up to $z \sim 5.5$.

Future prospects

Building on the analysis presented in this Thesis, several future directions can be pursued, including maximizing the extraction of the anisotropic information provided by RSDs and developing a proper perturbative modeling of neutrino effects. We outline below the potential extensions of this study.

Neutrino signatures in the anisotropic 3PCF. In this thesis, we have measured only the isotropic multipoles of the 3PCF. In redshift space, however, the 3PCF (and all the other statistics) also has an anisotropic component, because the action of RSDs depends on the orientation of triplets with respect to the line of sight. An extension of this work aiming to maximize the extraction of information in redshift space would therefore naturally include the measurement of the anisotropic multipoles of the 3PCF. The software `MeasCorr` already employed in this thesis provides an implementation of the algorithm for estimating the anisotropic 3PCF multipoles presented in Slepian and Eisenstein (2018), which generalizes the isotropic case while retaining the same $\mathcal{O}(N^2)$ scaling. The software also converts these multipoles into the TripoSH decomposition multipoles (Sugiyama et al. 2019), which provide a more natural encoding of symmetry properties, such as rotational invariance around any possible line of sight and parity symmetry.

Quantifying the 3PCF information content through Fisher forecast. The Quijote suite, with its large number of simulations in which only one parameter at a time is varied with respect to the fiducial cosmology, offers the possibility to numerically compute partial derivatives of a given statistic with respect to cosmological parameters, needed in a Fisher forecast (Dodelson and Schmidt 2020). This is particularly useful, as it allows the extraction of information content from clustering statistics up to nonlinear scales (Hahn et al. 2020). This approach can be exploited to compare the forecasted constraints on M_ν and σ_8 from the 2PCF and the 3PCF, thereby firmly assessing the potential of the latter to break the degeneracy between these two parameters.

In particular, linear bias can be kept fixed (a condition not met in our analysis) when taking derivatives with respect to other parameters by tuning the number of most massive halos in each simulation across the different cosmologies[1].

Nonlinear neutrino modeling. In view of the forthcoming surveys mentioned above, a broader future direction is represented by the full exploitation of low- and higher-order statistics, combining them in a joint likelihood analysis. Achieving this type of analysis requires addressing both theoretical and computational challenges. A theoretical challenge that still needs to be addressed is the 3PCF nonlinear modelling in the presence of massive neutrinos. The current state of the art is represented by already-developed perturbative models in Fourier space, particularly relevant for the bispectrum, which take into account the effects of neutrinos on mode coupling by introducing terms dependent on M_ν in the perturbative kernels (Aviles et al. 2021; Kamalinejad and Slepian 2025a; Kamalinejad and Slepian 2025b). A straightforward strategy to get their counterpart in configuration space, i.e., the 3PCF including massive neutrinos, would rely on evaluating the Fourier inverse transform.

Then, a first and computationally accessible possibility is represented by a joint 2+3PCF *template-fitting* likelihood analysis, performed by compressing the cosmological information to the distortions around a template model: *dynamical* ones (i.e., RSD) governed by the linear growth rate f, and *geometrical* ones, quantifying departures from the assumed fiducial cosmology, encoded in the so-called Alcock–Paczyński parameters (Alcock and Paczynski 1979). Usually, the extraction of cosmological information from dynamical distortions allows one to derive constraints on the growth rate of cosmic structures, also treating σ_8 as a free parameter, through the parameter combination $f\sigma_8$ (see, e.g., the analysis presented in Veropalumbo et al. 2021).

Conversely, a second method for extracting cosmological information is the so-called *full-shape analysis*, which directly exploits the full dependence of the 3PCF shape on cosmological parameters. However, when implementing this type of analysis, severe computational limitations, primarily due to the computational time required to evaluate 3PCF modelling predictions at each sampling of the cosmological parameter space, make the approach highly demanding and unfeasible. To overcome this issue, one can rely on emulators that, once trained on extensive libraries of theoretical models, are capable of rapidly producing accurate and fast predictions even on standard computing hardware. Such an approach has been recently adopted in real space by Euclid Collaboration: Guidi et al. (2025) for the validation of full-shape ΛCDM models. Their work presents the first joint full-shape analysis of the 2PCF and 3PCF, limited to the parameters A_s, h, and $\Omega_{\rm cdm}h^2$, and demonstrates the advantages of this combined approach, which yields significantly tighter constraints than those obtained from the 2PCF alone. A similar approach, performed by opening the parameter space to M_ν, would allow inferring the sum of neutrino masses for the first time in an LSS analysis involving the 3PCF, and directly ex-

---

[1] Tabulated numbers maximizing the selected halos, and keeping bias fixed, are available for a subset of the Quijote simulations and redshifts at `https://quijote-simulations.readthedocs.io/en/latest/halo_statistics.html`

ploring degeneracies between M_ν and other cosmological parameters, with the goal of alleviating them when adding higher-order statistics. This serves as an important source of information to disentangle massive and massless neutrino cosmologies when combining different cosmological datasets.

Extending the analysis to other phenomena. The methodology developed to study the information content of the 3PCF in the presence of massive neutrinos can be extended to other possible phenomena that produce small but relevant signals in the LSS. Among these, the small departures from Gaussianity imprinted by inflation in the primordial field of fluctuations, usually referred to as Primordial non-Gaussianities (PNG; reviewed in Celoria and Matarrese 2018), are, by definition, particularly suited to be studied with higher-order statistics, an aspect that still is largely unexplored in configuration space.

As a final remark, we point out that our study has focused on analyzing halos in a simulation. In general, translating this kind of analysis to real data from a survey is complicated by the need to deal with multiple tracers of the underlying matter field, each characterized by different selection functions, and distributed within a light cone rather than in a simple box at fixed redshift.

Appendix A

Wick's Theorem

Wick's theorem (Isserlis 1918; Wick 1950) makes it possible to compute the higher-order moments of a multivariate Gaussian distribution in terms of its covariance matrix. We state it in Fourier space.

Theorem *Let δ be a multivariate Gaussian field with zero mean. Then the following relation holds:*

$$\left\langle \tilde{\delta}(\mathbf{k}_1)\dots\tilde{\delta}(\mathbf{k}_n)\right\rangle = \sum_{p\in P_{n,2}} \prod_{\{i,j\}\in p} \left\langle \tilde{\delta}(\mathbf{k}_i)\tilde{\delta}(\mathbf{k}_j)\right\rangle , \tag{A.0.1}$$

where $P_{n,2}$ represents the family of all the possible partitions of $\{\tilde{\delta}(\mathbf{k}_1), \dots, \tilde{\delta}(\mathbf{k}_n)\}$ into pairs. □

Notice that, according to Wick's theorem, odd statistics of a Gaussian field, namely $\left\langle \tilde{\delta}(\mathbf{k}_1)\dots\tilde{\delta}(\mathbf{k}_n)\right\rangle$ with n odd, are zero, since for a set with an odd number of elements it is not possible to find a partition into pairs. Conversely, even statistics, i.e. $\left\langle \tilde{\delta}(\mathbf{k}_1)\dots\tilde{\delta}(\mathbf{k}_n)\right\rangle$ with n even, are all functions of 2-point statistics $\left\langle \tilde{\delta}(\mathbf{k}_i)\tilde{\delta}(\mathbf{k}_j)\right\rangle$.

Appendix B

Correlation matrices

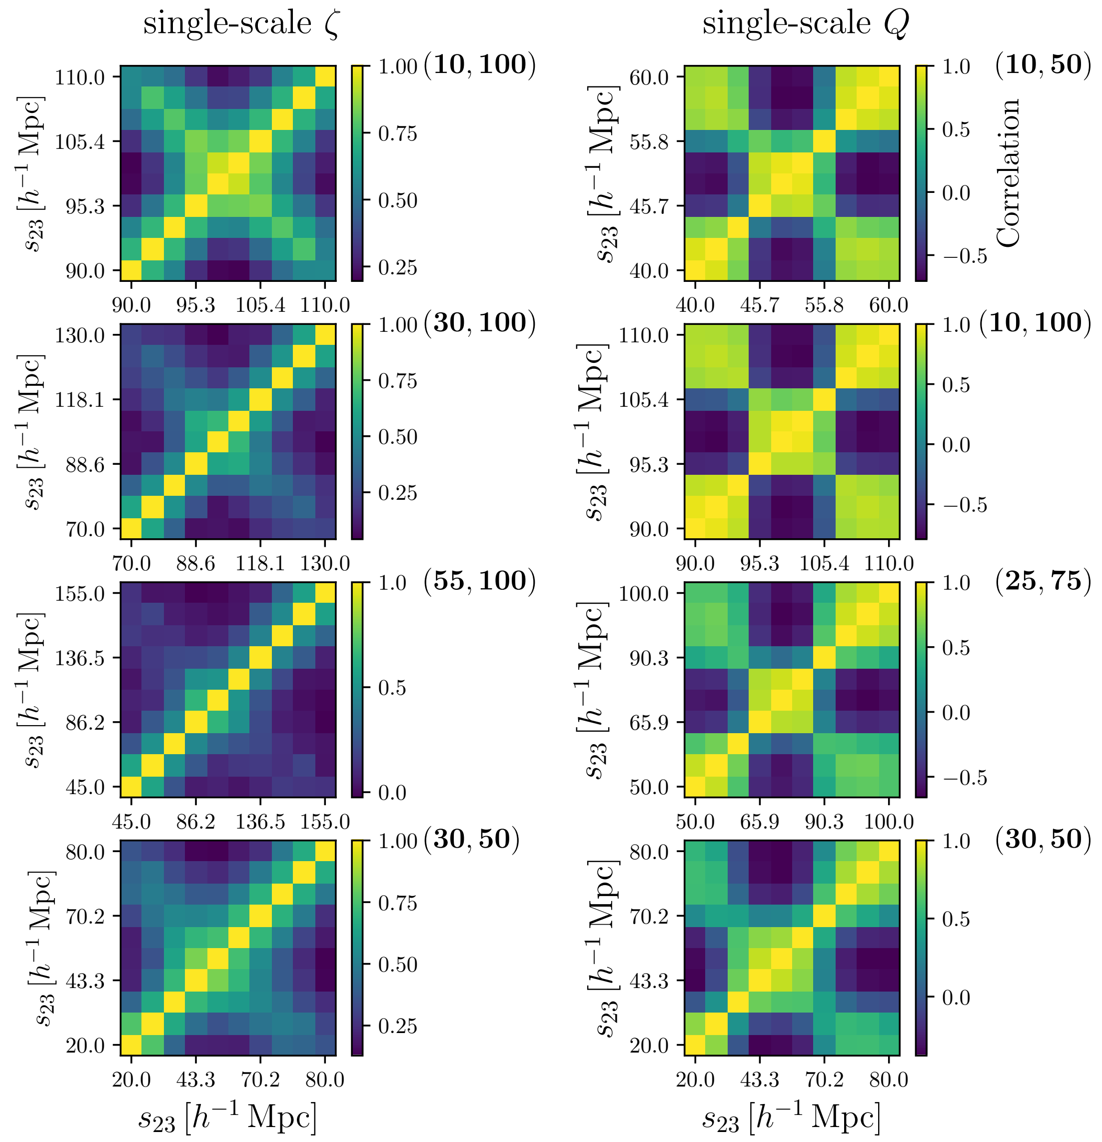

Figure B.1: The correlation matrices of the single-scale connected 3PCF (left column) and reduced 3PCF (right column), in redshift space and at $z = 0$, for the scales (s_{12}, s_{13}) indicated to the right of each plot (in units of h^{-1} Mpc). These scales correspond to those selected in Fig. 3.10 and 3.11. The angle θ between s_{12} and s_{13} has been converted into the value of the third side s_{23} through the cosine rule.

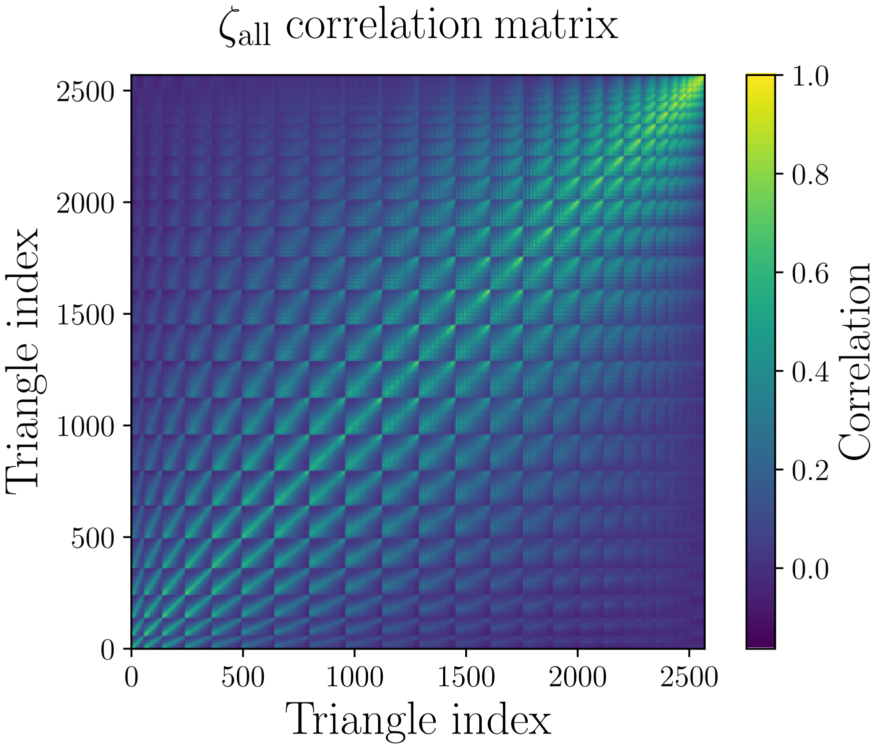

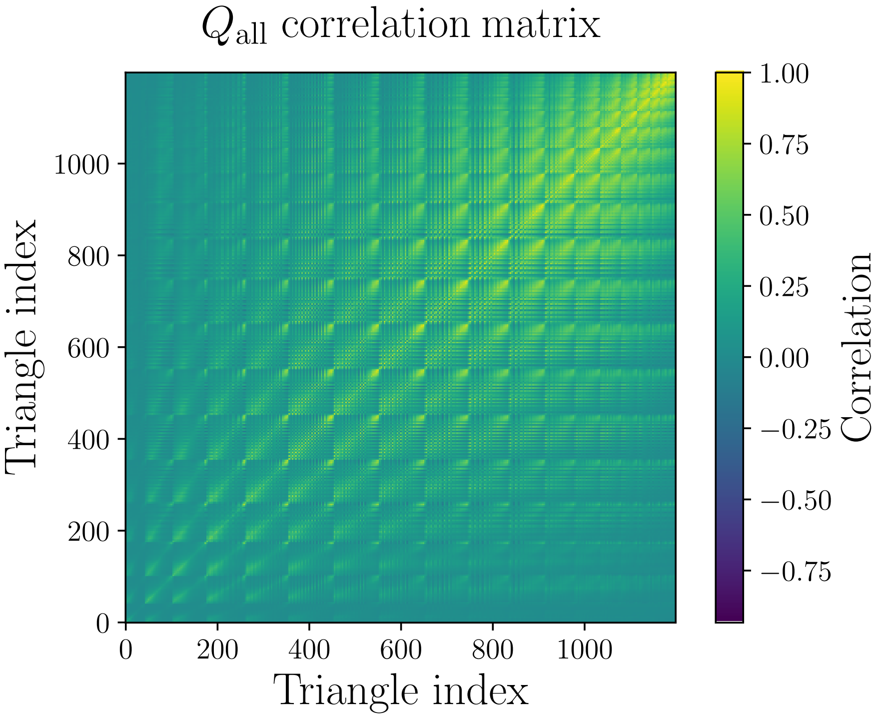

Figure B.2: The correlation matrices of the all-scales connected 3PCF $\zeta_{\rm all}$ (top) and reduced 3PCF $Q_{\rm all}$ (bottom). Triangle indices are assigned in the way detailed in Ch. 2. The matrices are shown in redshift space and at $z = 0$.

Fig. B.1 shows the correlation matrices of the single-scale connected 3PCF ζ and reduced 3PCF Q, in redshift-space and at $z = 0$, for the (s_{12}, s_{13}) pairs selected in Figs. 3.10 and 3.11. The matrix elements are identified by pairs of possible values of the third side s_{23} for each (s_{12}, s_{13}) pair.

Fig. B.2 shows the correlation matrices of the all-scales connected 3PCF $\zeta_{\rm all}$ and reduced 3PCF $Q_{\rm all}$. The matrix elements are identified by pairs of triangle indices, obtained by assigning a progressive number to each triplet (s_{12}, s_{13}, s_{23}), with $s_{12} \leq s_{13} \leq s_{23}$ and $s_{23} \leq s_{12} + s_{13}$, by looping in the order $s_{23} \rightarrow s_{13} \rightarrow s_{12}$ (see Ch. 2). For $\zeta_{\rm all}$, the side-bin centers range from 5 to $145\,h^{-1}\,$Mpc, while for $Q_{\rm all}$ they range from 5 to $110\,h^{-1}\,$Mpc, with a bin width of $\Delta s = 5\,h^{-1}\,$Mpc.

Appendix C

The Slepian and Eisenstein 3PCF model in real space

We detail in this appendix the real-space 3PCF perturbative model of Slepian and Eisenstein (2015b), accurate on large scales ($r \gtrsim 20\,h^{-1}\,\mathrm{Mpc}$) and far from isosceles configurations. The model assumes the nonlinear bias relation of Eq. 4.4.1 between the halo and matter overdensity fields δ_h and δ_m, which we report here for clarity:

$$\delta_h = b_1 \delta_m + \frac{b_2}{2}\left(\delta_m^2 - \langle \delta_m^2 \rangle\right), \tag{C.0.1}$$

with b_1 and b_2 respectively, the linear and nonlinear bias coefficients.

In this model, a given isotropic multipole $\zeta_{h,\ell}(r_{12}, r_{13})$ of the halo 3PCF can be obtained as the projection of the modeled connected 3PCF $\zeta_{h,c}(r_{12}, r_{13}, \theta)$ onto the corresponding Legendre polynomial, namely

$$\zeta_{h,\ell}(r_{12}, r_{13}) = \frac{2\ell+1}{2} \int_{-1}^{1} d\cos\theta \; \zeta_{h,c}(r_{12}, r_{13}, \theta)\; P_\ell(\cos\theta). \tag{C.0.2}$$

To compute the function $\zeta_{h,c}$, first a particular tracer is chosen to contribute each term in the expectation value $\langle \delta_h(\boldsymbol{x_i})\delta_h(\boldsymbol{x_j})\delta_h(\boldsymbol{x_k})\rangle$ needed to form the 3PCF: this greatly simplifies the computations, especially in dealing with quadratic terms in the linear matter density field or second-order terms. Then, to account for the indistinguishability of triangle vertices, the result is cyclically summed around the triangle. Following this procedure, $\zeta_{h,c}$ can be written as

$$\zeta_{h,c} = \sum_\ell \zeta_{\mathrm{pc},\ell}(r_{12}, r_{13})\; P_\ell(\cos\theta) + \text{cyc.perm.}, \tag{C.0.3}$$

where $\zeta_{\mathrm{pc},\ell}$ are the pre-cyclic (i.e., before cyclic summing) multipoles of the 3PCF. In this formula, when cycling over the sides of the triangle, the angular dependence inside the Legendre polynomial must also be cycled accordingly.

The authors show that only for $\ell = 0, 1, 2$, the pre-cyclic multipoles are nonzero: for two given sides, say r_{12} and r_{13}, their values are:

$$\begin{aligned}
\ell = 0 :&\; \left[2b_1^2 b_2 + \frac{34}{21} b_1^3\right] \xi_{12}^{[0]}\, \xi_{13}^{[0]}, \\
\ell = 1 :&\; -b_1^3 \left[\xi_{12}^{[1+]}\, \xi_{13}^{[1-]} + \xi_{13}^{[1+]}\, \xi_{12}^{[1-]}\right], \\
\ell = 2 :&\; \frac{8}{21} b_1^3\, \xi_{12}^{[2]}\, \xi_{13}^{[2]},
\end{aligned} \tag{C.0.4}$$

where the subscripts 12 and 13 denote evaluations of the functions at r_{12} and r_{13}, respectively, and the functions $\xi_{ij}^{[n]}$, $\xi_{ij}^{[n\pm]}$ are integrals of the linear matter power spectrum given by

$$\begin{aligned} \xi_{ij}^{[n]} &\equiv \int \frac{k^2 dk}{2\pi^2} P_{m,\mathrm{lin}}(k) j_n(k r_{ij}), \\ \xi_{ij}^{[n\pm]} &\equiv \int \frac{k^2 dk}{2\pi^2} k^{\pm 1} P_{m,\mathrm{lin}}(k) j_n(k r_{ij}), \end{aligned} \tag{C.0.5}$$

where j_n denotes the n-th order spherical Bessel function of the first kind; the term $\xi^{[0]}$ simply corresponds to the linear matter 2PCF.

It is worth pointing out that, although the only contributions to the pre-cyclic multipoles come from $\ell = 0, 1$ and 2, the cyclic sum of Eq. C.0.3 introduces higher orders, generically onto an arbitrary number of Legendre polynomials. This means that the $\zeta_{h,\ell}$ of Eq. C.0.2 with $\ell > 2$ are not zero, contrary to what one might erroneously infer from the expression of the pre-cyclic multipoles.

Eventually, the modeled single-scale halo 3PCF ζ_h can be obtained by expanding the modeled multipoles up to a given $\ell_{\max}$ into a Legendre series. By performing the operations above, with the expressions in Eq. C.0.4 for the pre-cyclic multipoles, one can see that ζ_h can be written as

$$\zeta_h(r_{12}, r_{13}, \theta) = b_1^3 \, \zeta_m(r_{12}, r_{13}, \theta) \; + b_1^2 b_2 \, [\, \xi_m(r_{12}) \xi_m(r_{13}) + \mathrm{cyc.\,perm.}\,], \tag{C.0.6}$$

where ζ_m and ξ_m are, respectively, the matter 3PCF and 2PCF. The function ζ_m is obtained by applying the above equations with $b_1 = 1$ and $b_2 = 0$).

Since both ζ_m and the term within square brackets scale as σ_8^4, to model their σ_8 dependence it is sufficient to compute ζ_h from the linear matter power spectrum generated with a fiducial value of this parameter, say $\sigma_{8,\mathrm{fid}}$, and then multiply Eq. C.0.6 by the factor $(\sigma_8/\sigma_{8,\mathrm{fid}})^4$ to obtain the model for a new value σ_8. This precisely yields Eq. 4.4.2.

By recalling that, for the bias relation given by Eq. C.0.1, the halo 2PCF ξ_h can be written as $\xi_h = (\sigma_8/\sigma_{8,\mathrm{fid}})^2 \, b_1^2 \xi_{m,\mathrm{fid}}$, with $\xi_{m,\mathrm{fid}}$ the fiducial matter 2PCF, then for the model halo reduced 3PCF one has

$$Q_h(r_{12}, r_{13}, \theta) = \frac{1}{b_1} Q_m(r_{12}, r_{13}, \theta) + \frac{b_2}{b_1^2}, \tag{C.0.7}$$

which corresponds to Eq. 4.4.3.

Appendix D

Shape plots for different choices of $\eta_{\min}$ and s_{23} range

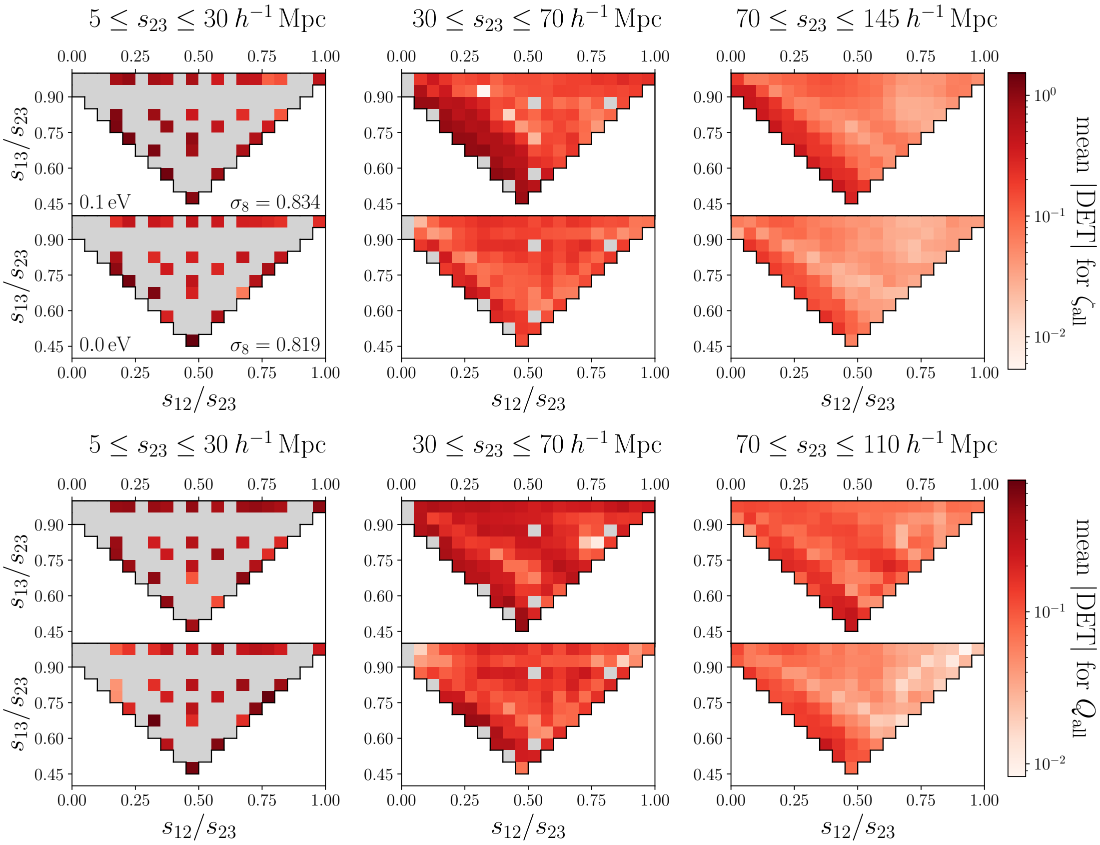

Figure D.1: The absolute values of detectability as a function of the triangle side ratios s_{13}/s_{23} versus s_{12}/s_{23} (with $s_{12} \leq s_{13} \leq s_{23}$), averaged, at each $(s_{12}/s_{23}, s_{13}, s_{23})$ bin, over all the triangles falling in the bin. The upper panel (first and second rows) and the lower panel (third and fourth rows) display, respectively, the results for ζ and Q. Each panel compares the `Mnu_p` simulation (upper row) with the `s8_m` (lower row) simulation. In each column, only the triangles with s_{23} falling in a specific range are considered: $[5, 30]$, $[30, 70]$, and $[70, s_{23,\max}]\,h^{-1}\,\mathrm{Mpc}$, with $s_{23,\max} = 145\,h^{-1}\,\mathrm{Mpc}$ for ζ and $110\,h^{-1}\,\mathrm{Mpc}$ for Q. Gray regions indicate unpopulated bins. Detectabilities are computed in redshift space at $z = 0$.

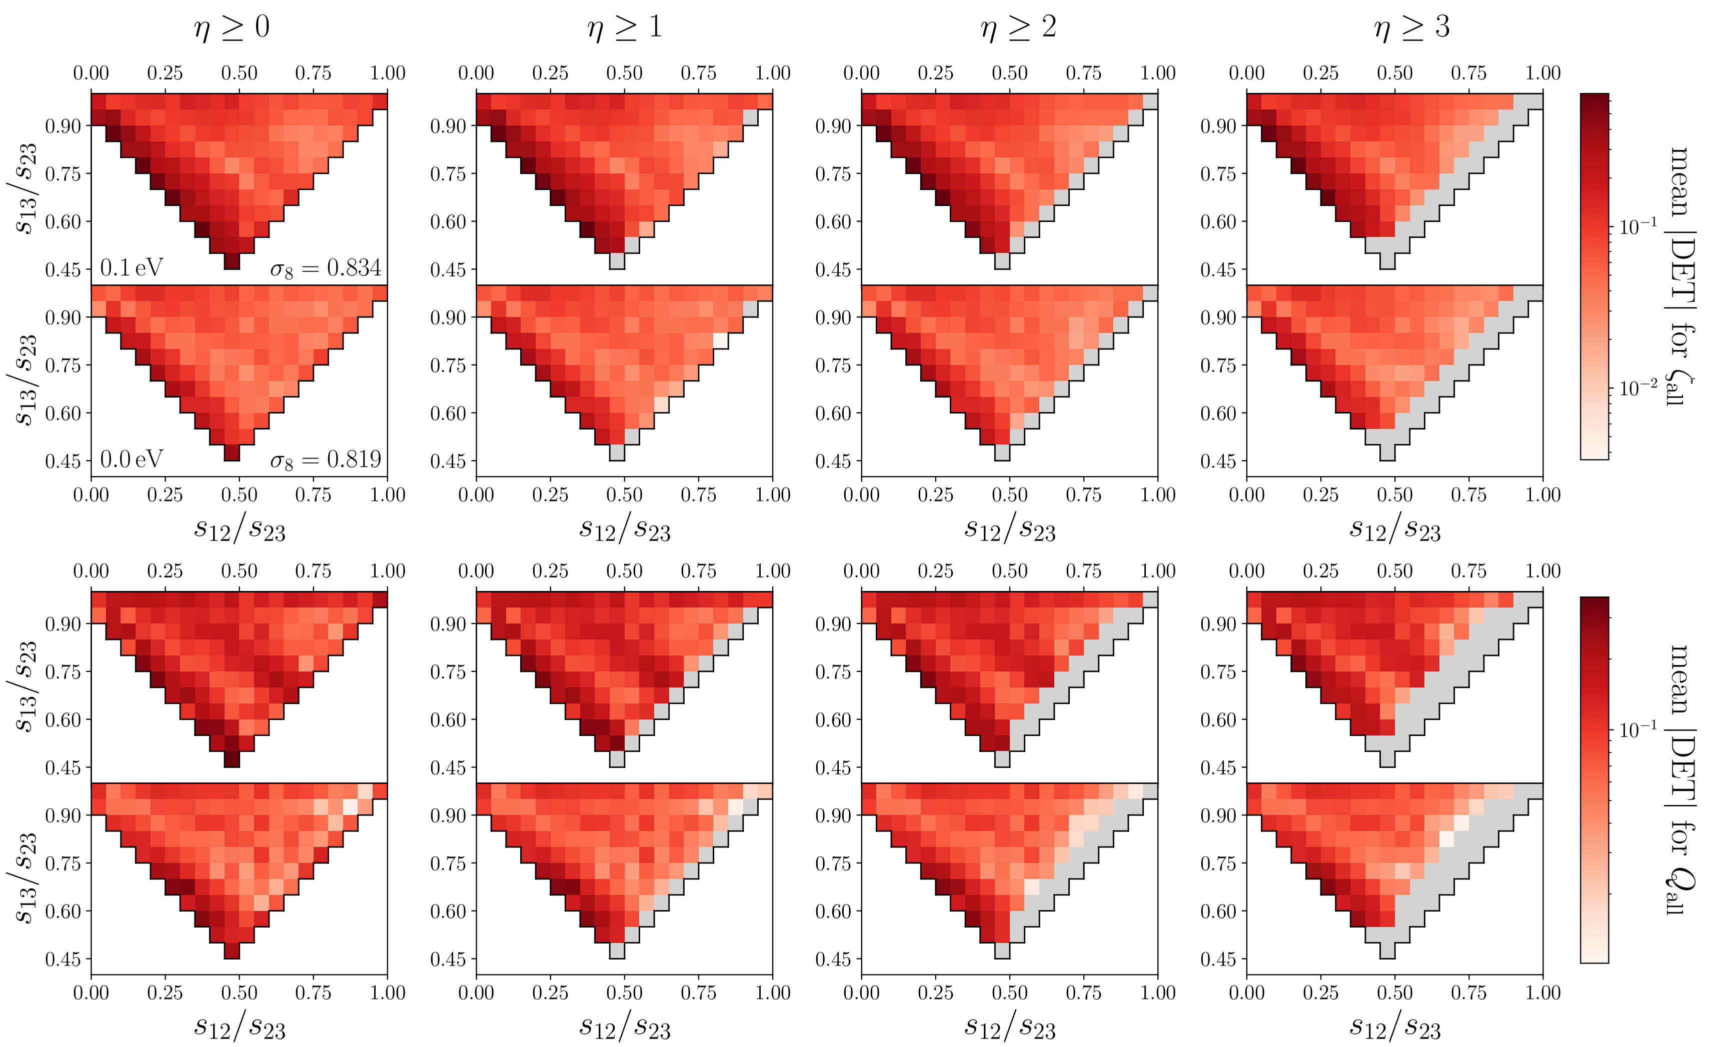

Figure D.2: Same as in Fig. D.1, but columns are now organized by considering only triangles for which the parameter η (Eq. 2.4.8) is above or equal to a threshold $\eta_{\rm min}$, regardless of their sizes. The values of $\eta_{\rm min}$ are chosen in the range $0 \leq \eta_{\rm min} \leq 3$.

Fig. D.1 compares the absolute values of the detectabilities of both the connected and reduced 3PCF, as a function of triangle shape, for the simulation `Mnu_p` ($M_\nu = 0.1\,\mathrm{eV}$, $\sigma_8 = 0.834$) and `s8_m` ($M_\nu = 0\,\mathrm{eV}$, $\sigma_8 = 0.819$). Different intervals for the largest side s_{23} are considered, namely $5 \leq s_{23}[h^{-1}\,\mathrm{Mpc}] \leq 30$, $30 \leq s_{23}[h^{-1}\,\mathrm{Mpc}] \leq 70$, and $70 \leq s_{23}[h^{-1}\,\mathrm{Mpc}] \leq s_{23,\mathrm{max}}$, with $s_{23,\mathrm{max}} = 145\,h^{-1}\,\mathrm{Mpc}$ for the connected 3PCF and $s_{23,\mathrm{max}} = 110\,h^{-1}\,\mathrm{Mpc}$ for thre reduced 3PCF.

Fig. D.2 shows the same quantities of Fig. D.1, for the same simulations, but for different choices of $\eta_{\rm min} = 0, 1, 2$, and 3. The detectabilities shown in each plot refer to the full set of triangle configurations, with $5 \leq s_{23}[h^{-1}\,\mathrm{Mpc}] \leq 145$ and $5 \leq s_{23}[h^{-1}\,\mathrm{Mpc}] \leq 110$ for connected and reduced 3PCF, respectively.

References

Chebat, D. et al. (2025). "Cosmological neutrino mass: a frequentist overview in light of DESI." In: *arXiv e-prints*, arXiv:2507.12401, arXiv:2507.12401. DOI: 10.48550/arXiv.2507.12401. arXiv: 2507.12401 [astro-ph.CO].

DESI Collaboration: Abdul-Karim, M. et al. (2025). "DESI DR2 Results II: Measurements of Baryon Acoustic Oscillations and Cosmological Constraints." In: *arXiv e-prints*, arXiv:2503.14738, arXiv:2503.14738. DOI: 10.48550/arXiv.2503.14738. arXiv: 2503.14738 [astro-ph.CO].

DESI Collaboration: Adame, A.G. et al. (2025). "DESI 2024 VII: cosmological constraints from the full-shape modeling of clustering measurements." In: *JCAP* 2025.07, p. 028. DOI: 10.1088/1475-7516/2025/07/028. URL: https://dx.doi.org/10.1088/1475-7516/2025/07/028.

Euclid Collaboration: Archidiacono, M. et al. (2025). "Euclid preparation - LIV. Sensitivity to neutrino parameters." In: *A&A* 693, A58. DOI: 10.1051/0004-6361/202450859. URL: https://doi.org/10.1051/0004-6361/202450859.

Euclid Collaboration: Guidi, M. et al. (2025). "Euclid preparation. Full-shape modelling of 2-point and 3-point correlation functions in real space." In: *arXiv e-prints*, arXiv:2506.22257, arXiv:2506.22257. DOI: 10.48550/arXiv.2506.22257. arXiv: 2506.22257 [astro-ph.CO].

Euclid Collaboration: Mellier, Y. et al. (2025). "Euclid: I. Overview of the Euclid mission." In: *A&A* 697, A1, A1. DOI: 10.1051/0004-6361/202450810. arXiv: 2405.13491 [astro-ph.CO].

Kamalinejad, F. and Slepian, Z. (2025a). "From Theory to Forecast: Neutrino Mass Effects on Mode-Coupling Kernels and Their Observational Implications." In: *arXiv e-prints*, arXiv:2508.06759, arXiv:2508.06759. DOI: 10.48550/arXiv.2508.06759. arXiv: 2508.06759 [astro-ph.CO].

— (2025b). "Neutrino mass signatures in the galaxy bispectrum." In: *Phys. Rev. D* 112.8, 083501, p. 083501. DOI: 10.1103/ckfx-r4q9.

Verdiani, F. et al. (2025). "Redshift-space distortions in massive neutrinos cosmologies." In: *Phys. Rev. D* 112.4, 043545, p. 043545. DOI: 10.1103/mjsv-885q. arXiv: 2503.06655 [astro-ph.CO].

Farina, A. et al. (2024). “Modeling and measuring the anisotropic halo 3-point correlation function: a coordinated study.” In: *arXiv e-prints*, arXiv:2408.03036, arXiv:2408.03036. DOI: `10.48550/arXiv.2408.03036`. arXiv: `2408.03036 [astro-ph.CO]`.

Mainieri, V. et al. (2024). “The Wide-field Spectroscopic Telescope (WST) Science White Paper.” In: *arXiv e-prints*, arXiv:2403.05398, arXiv:2403.05398. DOI: `10.48550/arXiv.2403.05398`. arXiv: `2403.05398 [astro-ph.IM]`.

Particle Data Group: Navas, S. et al. (2024). “Review of particle physics.” In: *Phys. Rev. D* 110.3, p. 030001. DOI: `10.1103/PhysRevD.110.030001`.

Coulton, W. R. et al. (2023). “Quijote-PNG: Simulations of Primordial Non-Gaussianity and the Information Content of the Matter Field Power Spectrum and Bispectrum.” In: *Ap. J.* 943.1, p. 64. DOI: `10.3847/1538-4357/aca8a7`. URL: `https://dx.doi.org/10.3847/1538-4357/aca8a7`.

Guidi, M. et al. (2023). “Modelling the next-to-leading order matter three-point correlation function using FFTLog.” In: *JCAP* 2023.8, 066, p. 066. DOI: `10.1088/1475-7516/2023/08/066`. arXiv: `2212.07382 [astro-ph.CO]`.

Murakami, Y. S. et al. (2023). “Leveraging SN Ia spectroscopic similarity to improve the measurement of H0.” In: *JCAP* 2023.11, p. 046. DOI: `10.1088/1475-7516/2023/11/046`. URL: `https://dx.doi.org/10.1088/1475-7516/2023/11/046`.

DES Collaboration: Abbott, T. M. C. et al. (2022). “Dark Energy Survey Year 3 results: Cosmological constraints from galaxy clustering and weak lensing.” In: *Phys. Rev. D* 105, p. 023520. DOI: `10.1103/PhysRevD.105.023520`. URL: `https://link.aps.org/doi/10.1103/PhysRevD.105.023520`.

Euclid Collaboration: Scaramella, R. et al. (2022). “Euclid preparation. I. The Euclid Wide Survey.” In: *A&A* 662, A112, A112. DOI: `10.1051/0004-6361/202141938`. arXiv: `2108.01201 [astro-ph.CO]`.

Lopez, A. M., Clowes, R. G., and Williger, G. M. (2022). “A Giant Arc on the Sky.” In: *MNRAS* 516.2, pp. 1557–1572. DOI: `10.1093/mnras/stac2204`. URL: `https://doi.org/10.1093/mnras/stac2204`.

Veropalumbo, A. et al. (2022). “The halo 3-point correlation function: a methodological analysis.” In: *JCAP* 2022.9, 033, p. 033. DOI: `10.1088/1475-7516/2022/09/033`. arXiv: `2206.00672 [astro-ph.CO]`.

Adler, S. L. (2021). “Is «dark energy» a quantum vacuum energy?” In: *MPLA* 36.40, 2130027-87, pp. 2130027–87. DOI: `10.1142/S0217732321300275`. arXiv: `2111.12576 [gr-qc]`.

Aviles, A. et al. (2021). “Clustering in massive neutrino cosmologies via Eulerian Perturbation Theory.” In: *JCAP* 2021.11, p. 028. DOI: `10.1088/1475-7516/2021/11/028`. URL: `https://dx.doi.org/10.1088/1475-7516/2021/11/028`.

Heymans, C. et al. (2021). "KiDS-1000 Cosmology: Multi-probe weak gravitational lensing and spectroscopic galaxy clustering constraints." In: *A&A* 646, A140. DOI: `10.1051/0004-6361/202039063`. URL: `https://doi.org/10.1051/0004-6361/202039063`.

Moresco, M. et al. (2021). "C3: Cluster Clustering Cosmology. II. First Detection of the Baryon Acoustic Oscillations Peak in the Three-point Correlation Function of Galaxy Clusters." In: *Ap. J.* 919.2, p. 144. DOI: `10.3847/1538-4357/ac10c9`. URL: `https://dx.doi.org/10.3847/1538-4357/ac10c9`.

Secrest, N. J. et al. (2021). "A Test of the Cosmological Principle with Quasars." In: *Ap. J. L.* 908.2, p. L51. DOI: `10.3847/2041-8213/abdd40`. URL: `https://dx.doi.org/10.3847/2041-8213/abdd40`.

Veropalumbo, A. et al. (2021). "A joint 2- and 3-point clustering analysis of the VIPERS PDR2 catalogue at z $\sim$ 1: breaking the degeneracy of cosmological parameters." In: *MNRAS* 507.1, pp. 1184–1201. DOI: `10.1093/mnras/stab2205`. arXiv: `2106.12581 [astro-ph.CO]`.

Dodelson, S. and Schmidt, F. (2020). *Modern Cosmology.* 2nd ed. Academic Press. DOI: `https://doi.org/10.1016/C2017-0-01943-2`.

Hahn, C. et al. (2020). "Constraining M_ν with the bispectrum. Part I. Breaking parameter degeneracies." In: *JCAP* 2020.3, 040, p. 040. DOI: `10.1088/1475-7516/2020/03/040`. arXiv: `1909.11107 [astro-ph.CO]`.

Planck Collaboration: Aghanim, N. et al. (2020). "Planck2018 results: VI. Cosmological parameters." In: *A&A* 641, A6. ISSN: 1432-0746. DOI: `10.1051/0004-6361/201833910`. URL: `http://dx.doi.org/10.1051/0004-6361/201833910`.

Villaescusa-Navarro, F. et al. (2020). "The Quijote Simulations." In: *Ap. J. S.* 250.1, p. 2. DOI: `10.3847/1538-4365/ab9d82`. URL: `https://dx.doi.org/10.3847/1538-4365/ab9d82`.

Dore, O. et al. (2019). "WFIRST: The Essential Cosmology Space Observatory for the Coming Decade." In: *BAAS* 51.3, 341, p. 341. DOI: `10.48550/arXiv.1904.01174`. arXiv: `1904.01174 [astro-ph.CO]`.

Sugiyama, N. S. et al. (2019). "A complete FFT-based decomposition formalism for the redshift-space bispectrum." In: *MNRAS* 484.1, pp. 364–384. ISSN: 1365-2966. DOI: `10.1093/mnras/sty3249`. URL: `http://dx.doi.org/10.1093/mnras/sty3249`.

Celoria, M. and Matarrese, S. (2018). *Primordial Non-Gaussianity.* arXiv: `1812.08197 [astro-ph.CO]`. URL: `https://arxiv.org/abs/1812.08197`.

Ruggeri, R. et al. (2018). "DEMNUni: massive neutrinos and the bispectrum of large scale structures." In: *JCAP* 2018.03, p. 003. DOI: `10.1088/1475-7516/2018/03/003`. URL: `https://dx.doi.org/10.1088/1475-7516/2018/03/003`.

Slepian, Z. and Eisenstein, D. J. (2018). "A practical computational method for the anisotropic redshift-space three-point correlation function." In: *MNRAS* 478.2, pp. 1468–

1483. ISSN: 1365-2966. DOI: 10.1093/mnras/sty1063. URL: http://dx.doi.org/10.1093/mnras/sty1063.

Villaescusa-Navarro, F. et al. (2018). "The Imprint of Neutrinos on Clustering in Redshift Space." In: *Ap. J.* 861.1, p. 53. DOI: 10.3847/1538-4357/aac6bf. URL: https://dx.doi.org/10.3847/1538-4357/aac6bf.

Bullock, J. S. and Boylan-Kolchin, M. (2017). "Small-Scale Challenges to the ΛCDM Paradigm." In: *ARA&A* 55, pp. 343–387. ISSN: 1545-4282. DOI: https://doi.org/10.1146/annurev-astro-091916-055313. URL: https://www.annualreviews.org/content/journals/10.1146/annurev-astro-091916-055313.

Slepian, Z. and Eisenstein, D. J. (2017). "Modelling the large-scale redshift-space 3-point correlation function of galaxies." In: *MNRAS* 469.2, pp. 2059–2076. DOI: 10.1093/mnras/stx490. arXiv: 1607.03109 [astro-ph.CO].

Sridhar, S. et al. (2017). "Evolution of the real-space correlation function from next generation cluster surveys. Recovering the real-space correlation function from photometric redshifts." In: *A&A* 600, A32, A32. DOI: 10.1051/0004-6361/201629369. arXiv: 1612.02821 [astro-ph.CO].

Zennaro, M. et al. (2017). "Initial conditions for accurate N-body simulations of massive neutrino cosmologies." In: *MNRAS* 466.3, pp. 3244–3258. DOI: 10.1093/mnras/stw3340. arXiv: 1605.05283 [astro-ph.CO].

Angulo, R. E. and Pontzen, A. (2016). "Cosmological N-body simulations with suppressed variance." In: *MNRAS* 462.1, pp. L1–L5. DOI: 10.1093/mnrasl/slw098. arXiv: 1603.05253 [astro-ph.CO].

DESI Collaboration: Aghamousa, A. et al. (2016). "The DESI Experiment Part I: Science,Targeting, and Survey Design." In: *arXiv e-prints*, arXiv:1611.00036, arXiv:1611.00036. DOI: 10.48550/arXiv.1611.00036. arXiv: 1611.00036 [astro-ph.IM].

Levi, M. and Vlah, Z. (2016). "Massive neutrinos in nonlinear large scale structure: A consistent perturbation theory." In: *arXiv: Cosmology and Nongalactic Astrophysics*. URL: https://api.semanticscholar.org/CorpusID:118647078.

Marulli, F., Veropalumbo, A., and Moresco, M. (2016). "CosmoBolognaLib: C++ libraries for cosmological calculations." In: *Astron. Comput.* 14, pp. 35–42. ISSN: 2213-1337. DOI: 10.1016/j.ascom.2016.01.005. URL: http://dx.doi.org/10.1016/j.ascom.2016.01.005.

Castorina, E. et al. (2015). "DEMNUni: the clustering of large-scale structures in the presence of massive neutrinos." In: *JCAP* 2015.07, p. 043. DOI: 10.1088/1475-7516/2015/07/043. URL: https://dx.doi.org/10.1088/1475-7516/2015/07/043.

Peloso, M. et al. (2015). "The effect of massive neutrinos on the BAO peak." In: *JCAP* 2015.7, pp. 001–001. DOI: 10.1088/1475-7516/2015/07/001. arXiv: 1505.07477 [astro-ph.CO].

Slepian, Z. and Eisenstein, D. J. (2015a). “Computing the three-point correlation function of galaxies in $\mathcal{O}(N^2)$ time.” In: *MNRAS* 454.4, pp. 4142–4158. ISSN: 1365-2966. DOI: `10.1093/mnras/stv2119`. URL: `http://dx.doi.org/10.1093/mnras/stv2119`.

— (2015b). “On the signature of the baryon-dark matter relative velocity in the two- and three-point galaxy correlation functions.” In: *MNRAS* 448.1, pp. 9–26. DOI: `10.1093/mnras/stu2627`. arXiv: `1411.4052 [astro-ph.CO]`.

Castorina, E. et al. (2014). “Cosmology with massive neutrinos II: on the universality of the halo mass function and bias.” In: *JCAP* 2014.02, p. 049. DOI: `10.1088/1475-7516/2014/02/049`. URL: `https://dx.doi.org/10.1088/1475-7516/2014/02/049`.

Moresco, M. et al. (2014). “Disentangling interacting dark energy cosmologies with the three-point correlation function.” In: *MNRAS* 443.4, pp. 2874–2886. DOI: `10.1093/mnras/stu1359`. arXiv: `1312.4530 [astro-ph.CO]`.

Villaescusa-Navarro, F. et al. (2014). “Cosmology with massive neutrinos I: towards a realistic modeling of the relation between matter, haloes and galaxies.” In: *JCAP* 2014.3, 011, p. 011. DOI: `10.1088/1475-7516/2014/03/011`. arXiv: `1311.0866 [astro-ph.CO]`.

Costanzi, M. et al. (2013). “Cosmology with massive neutrinos III: the halo mass function and an application to galaxy clusters.” In: *JCAP* 2013.12, p. 012. DOI: `10.1088/1475-7516/2013/12/012`. URL: `https://dx.doi.org/10.1088/1475-7516/2013/12/012`.

Marín, F. A. et al. (2013). “The WiggleZ Dark Energy Survey: constraining galaxy bias and cosmic growth with three-point correlation functions.” In: *MNRAS* 432.4, pp. 2654–2668. DOI: `10.1093/mnras/stt520`. arXiv: `1303.6644 [astro-ph.CO]`.

Wang, Q. and Zhan, H. (2013). “Mass-dependent Baryon Acoustic Oscillation Signal and Halo Bias.” In: *Ap.J.L.* 768.2, L27, p. L27. DOI: `10.1088/2041-8205/768/2/L27`. arXiv: `1302.1640 [astro-ph.CO]`.

Baldauf, T et al. (2012). “Evidence for quadratic tidal tensor bias from the halo bispectrum.” In: *Phys. Rev. D* 86 (8), p. 083540. DOI: `10.1103/PhysRevD.86.083540`. URL: `https://link.aps.org/doi/10.1103/PhysRevD.86.083540`.

Agarwal, S. and Feldman, H. A. (2011). “The effect of massive neutrinos on the matter power spectrum.” In: *MNRAS* 410.3, pp. 1647–1654. DOI: `10.1111/j.1365-2966.2010.17546.x`. arXiv: `1006.0689 [astro-ph.CO]`.

Cowan, G. et al. (2011). “Asymptotic formulae for likelihood-based tests of new physics.” In: *Eur. Phys. J. C* 71.2. ISSN: 1434-6052. DOI: `10.1140/epjc/s10052-011-1554-0`. URL: `http://dx.doi.org/10.1140/epjc/s10052-011-1554-0`.

Marulli, F. et al. (2011). “Effects of massive neutrinos on the large-scale structure of the Universe.” In: *MNRAS* 418.1, pp. 346–356. DOI: `10.1111/j.1365-2966.2011.19488.x`. arXiv: `1103.0278 [astro-ph.CO]`.

Simpson, F., Jackson, B., and Peacock, J. A. (2011). “Unmodified gravity.” In: *MNRAS* 411.2, pp. 1053–1058. DOI: 10.1111/j.1365-2966.2010.17734.x. arXiv: 1004.1920 [astro-ph.CO].

Villaescusa-Navarro, F. et al. (2011). “Neutrino halos in clusters of galaxies and their weak lensing signature.” In: *JCAP* 2011.6, 027, p. 027. DOI: 10.1088/1475-7516/2011/06/027. arXiv: 1104.4770 [astro-ph.CO].

Crocce, M. et al. (2010). “Simulating the Universe with MICE: the abundance of massive clusters.” In: *MNRAS* 403.3, pp. 1353–1367. DOI: 10.1111/j.1365-2966.2009.16194.x. arXiv: 0907.0019 [astro-ph.CO].

Shoji, M. and Komatsu, E. (2010). “Massive neutrinos in cosmology: Analytic solutions and fluid approximation.” In: *Phys. Rev. D* 81 (12), p. 123516. DOI: 10.1103/PhysRevD.81.123516. URL: https://link.aps.org/doi/10.1103/PhysRevD.81.123516.

Tinker, J. L. et al. (2010). “The Large-scale Bias of Dark Matter Halos: Numerical Calibration and Model Tests.” In: *Ap. J.* 724.2, pp. 878–886. DOI: 10.1088/0004-637X/724/2/878. arXiv: 1001.3162 [astro-ph.CO].

Viel, M., Haehnelt, M. G., and Springel, V. (2010). “The effect of neutrinos on the matter distribution as probed by the intergalactic medium.” In: *JCAP* 2010.6, 015, p. 015. DOI: 10.1088/1475-7516/2010/06/015. eprint: 1003.2422 (astro-ph.CO).

Gaztañaga, E. et al. (2009). “Clustering of luminous red galaxies - III. Baryon acoustic peak in the three-point correlation.” In: *MNRAS* 399.2, pp. 801–811. DOI: 10.1111/j.1365-2966.2009.15313.x. arXiv: 0807.2448 [astro-ph].

McDonald, P. and Roy, A. (2009). “Clustering of dark matter tracers: generalizing bias for the coming era of precision LSS.” In: *JCAP* 2009.8, 020, p. 020. DOI: 10.1088/1475-7516/2009/08/020. arXiv: 0902.0991 [astro-ph.CO].

Saito, S., Takada, M., and Taruya, A. (2009). “Nonlinear power spectrum in the presence of massive neutrinos: Perturbation theory approach, galaxy bias, and parameter forecasts.” In: *Phys. Rev. D* 80 (8), p. 083528. DOI: 10.1103/PhysRevD.80.083528. URL: https://link.aps.org/doi/10.1103/PhysRevD.80.083528.

Brandbyge, J. et al. (2008). “The effect of thermal neutrino motion on the non-linear cosmological matter power spectrum.” In: *JCAP* 2008.8, 020, p. 020. DOI: 10.1088/1475-7516/2008/08/020. arXiv: 0802.3700 [astro-ph].

Marín, F. A. et al. (2008). “Modeling the Galaxy Three-Point Correlation Function.” In: *Ap. J.* 672.2, pp. 849–860. DOI: 10.1086/523628. arXiv: 0704.0255 [astro-ph].

SubbaRao, M. U. et al. (2008). “Visualization of large scale structure from the Sloan Digital Sky Survey.” In: *New J. Phys.* 10.12, p. 125015. DOI: 10.1088/1367-2630/10/12/125015. URL: https://doi.org/10.1088/1367-2630/10/12/125015.

Tinker, J. et al. (2008). “Toward a Halo Mass Function for Precision Cosmology: The Limits of Universality.” In: *Ap. J.* 688.2, pp. 709–728. DOI: `10.1086/591439`. arXiv: `0803.2706 [astro-ph]`.

Eisenstein, D. J., Seo, H.-J., and White, M. (2007). “On the Robustness of the Acoustic Scale in the Low-Redshift Clustering of Matter.” In: *Ap.J.* 664.2, pp. 660–674. DOI: `10.1086/518755`. arXiv: `astro-ph/0604361 [astro-ph]`.

Hartlap, J., Simon, P., and Schneider, P. (2007). “Why your model parameter confidences might be too optimistic – unbiased estimation of the inverse covariance matrix.” In: *A&A* 464.1, pp. 399–404. DOI: `10.1051/0004-6361:20066170`.

Hu, W. and Sawicki, I. (2007). “Models of $f(R)$ cosmic acceleration that evade solar system tests.” In: *Phys. Rev. D* 76 (6), p. 064004. DOI: `10.1103/PhysRevD.76.064004`. URL: `https://link.aps.org/doi/10.1103/PhysRevD.76.064004`.

Schlegel, David J. et al. (2007). “SDSS-III: The Baryon Oscillation Spectroscopic Survey (BOSS).” In: *American Astronomical Society Meeting Abstracts.* Vol. 211. American Astronomical Society Meeting Abstracts, 132.29, p. 132.29.

Lesgourgues, J. and Pastor, S. (2006). “Massive neutrinos and cosmology.” In: *Phys. Rep.* 429.6, pp. 307–379. ISSN: 0370-1573. DOI: `https://doi.org/10.1016/j.physrep.2006.04.001`. URL: `https://www.sciencedirect.com/science/article/pii/S0370157306001359`.

Warren, M. S. et al. (2006). “Precision Determination of the Mass Function of Dark Matter Halos.” In: *Ap. J.* 646.2, pp. 881–885. DOI: `10.1086/504962`. arXiv: `astro-ph/0506395 [astro-ph]`.

Elgarøy, Ø. and Lahav, O. (2005). “Neutrino masses from cosmological probes.” In: *New J. Phys.* 7, pp. 61–61. ISSN: 1367-2630. DOI: `10.1088/1367-2630/7/1/061`. URL: `http://dx.doi.org/10.1088/1367-2630/7/1/061`.

Gaztañaga, E. and Scoccimarro, R. (2005). “The three-point function in large-scale structure: redshift distortions and galaxy bias.” In: *MNRAS* 361.3, pp. 824–836. DOI: `10.1111/j.1365-2966.2005.09234.x`. arXiv: `astro-ph/0501637 [astro-ph]`.

Hannestad, S. (2005). “Neutrino Masses and the Dark Energy Equation of State:Relaxing the Cosmological Neutrino Mass Bound.” In: *Phys. Rev. Lett.* 95.22, 221301, p. 221301. DOI: `10.1103/PhysRevLett.95.221301`. arXiv: `astro-ph/0505551 [astro-ph]`.

Springel, V. (2005). “The cosmological simulation code GADGET-2.” In: *MNRAS* 364.4, pp. 1105–1134. DOI: `10.1111/j.1365-2966.2005.09655.x`. arXiv: `astro-ph/0505010 [astro-ph]`.

Szapudi, I. (2004). “Three-Point Statistics from a New Perspective.” In: *Ap. J. L.* 605.2, pp. L89–L92. DOI: `10.1086/420894`. arXiv: `astro-ph/0404476 [astro-ph]`.

Anderson, T. W. (2003). *An introduction to multivariate statistical analysis.* 3rd ed. Wiley-Interscience.

Bernardeau, F. et al. (2002). "Large-scale structure of the Universe and cosmological perturbation theory." In: *Phys. Rep.* 367.1-3, pp. 1–248. DOI: `10.1016/S0370-1573(02)00135-7`. arXiv: `astro-ph/0112551 [astro-ph]`.

Coles, P. and Lucchin, F. (2002). *Cosmology. The Origin and Evolution of Cosmic Structures.* 2nd ed. Wiley.

Sheth, R. K. and Tormen, G. (2002). "An excursion set model of hierarchical clustering: ellipsoidal collapse and the moving barrier." In: *MNRAS* 329.1, pp. 61–75. DOI: `10.1046/j.1365-8711.2002.04950.x`. arXiv: `astro-ph/0105113 [astro-ph]`.

Sheth, R. K., Mo, H. J., and Tormen, G. (2001). "Ellipsoidal collapse and an improved model for the number and spatial distribution of dark matter haloes." In: *MNRAS* 323.1, pp. 1–12. DOI: `10.1046/j.1365-8711.2001.04006.x`. arXiv: `astro-ph/9907024 [astro-ph]`.

Lewis, A., Challinor, A., and Lasenby, A. (2000). "Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models." In: *Ap.J.* 538.2, pp. 473–476. DOI: `10.1086/309179`. arXiv: `astro-ph/9911177 [astro-ph]`.

Perlmutter, S. et al. (1999). "Measurements of Ω and Λ from 42 High-Redshift Supernovae." In: *Ap. J.* 517.2, pp. 565–586. DOI: `10.1086/307221`. arXiv: `astro-ph/9812133 [astro-ph]`.

Scoccimarro, R., Couchman, H. M. P., and Frieman, J. A. (1999). "The Bispectrum as a Signature of Gravitational Instability in Redshift Space." In: *Ap.J.* 517.2, pp. 531–540. DOI: `10.1086/307220`. arXiv: `astro-ph/9808305 [astro-ph]`.

Sheth, R. K. and Tormen, G. (1999). "Large-scale bias and the peak background split." In: *MNRAS* 308.1, pp. 119–126. DOI: `10.1046/j.1365-8711.1999.02692.x`. arXiv: `astro-ph/9901122 [astro-ph]`.

Cowan, G. (1998). *Statistical Data Analysis.* Oxford University Press. DOI: `10.1093/oso/9780198501565.001.0001`. URL: `https://doi.org/10.1093/oso/9780198501565.001.0001`.

Fukuda, Y. et al. (1998). "Evidence for Oscillation of Atmospheric Neutrinos." In: *Phys. Rev. Lett.* 81.8, pp. 1562–1567. DOI: `10.1103/PhysRevLett.81.1562`. arXiv: `hep-ex/9807003 [hep-ex]`.

Riess, A. G. et al. (1998). "Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant." In: *The Astron. J.* 116.3, p. 1009. DOI: `10.1086/300499`. URL: `https://dx.doi.org/10.1086/300499`.

Szapudi, I. and Szalay, A. S. (1998). "A New Class of Estimators for the N-Point Correlations." In: *Ap. J.* 494.1, p. L41. DOI: `10.1086/311146`. URL: `https://dx.doi.org/10.1086/311146`.

Kitayama, T. and Suto, Y. (1996). “Semianalytic Predictions for Statistical Properties of X-Ray Clusters of Galaxies in Cold Dark Matter Universes.” In: *Ap.J.* 469, p. 480. DOI: 10.1086/177797. arXiv: astro-ph/9604141 [astro-ph].

Bouchet, F. R. et al. (1994). “Perturbative Lagrangian approach to gravitational instability.” In: *A&A* 296. URL: https://www.osti.gov/biblio/1847343.

Buchert, T., Melott, A.L., and Weiss, A.G. (1994). “Testing higher-order Lagrangian perturbation theory against numerical simulations I. Pancake models.” In: *A&A* 288, pp. 349–364. DOI: 10.48550/arXiv.astro-ph/9309056. arXiv: astro-ph/9309056 [astro-ph].

Lacey, C. and Cole, S. (1994). “Merger Rates in Hierarchical Models of Galaxy Formation - Part Two - Comparison with N-Body Simulations.” In: *MNRAS* 271, p. 676. DOI: 10.1093/mnras/271.3.676. arXiv: astro-ph/9402069 [astro-ph].

Landy, S. D. and Szalay, A. S. (1993). “Bias and Variance of Angular Correlation Functions.” In: *Ap. J.* 412, p. 64. DOI: 10.1086/172900.

Hamilton, A. J. S. (1992). “Measuring Omega and the Real Correlation Function from the Redshift Correlation Function.” In: *Ap. J. L.* 385, p. L5. DOI: 10.1086/186264.

Hockney, R.W. and Eastwood, J.W. (1988). *Computer Simulation Using Particles*. Taylor & Francis.

Varshalovich, D. A., Moskalev, A. N., and Khersonskii, V. K. (1988). *Quantum Theory of Angular Momentum*. World Scientific. DOI: 10.1142/0270. eprint: https://www.worldscientific.com/doi/pdf/10.1142/0270. URL: https://www.worldscientific.com/doi/abs/10.1142/0270.

Kaiser, N. (1987). “Clustering in real space and in redshift space.” In: *MNRAS* 227, pp. 1–21. DOI: 10.1093/mnras/227.1.1.

Goroff, M. H. et al. (1986). “Coupling of modes of cosmological mass density fluctuations.” In: *Ap. J.* 311, pp. 6–14. DOI: 10.1086/164749.

Davis, M. et al. (1985). “The evolution of large-scale structure in a universe dominated by cold dark matter.” In: *Ap.J.* 292, pp. 371–394. DOI: 10.1086/163168.

Bond, J. R., Efstathiou, G., and Silk, J. (1980). “Massive Neutrinos and the Large-Scale Structure of the Universe.” In: *Phys. Rev. Lett.* 45 (24), pp. 1980–1984. DOI: 10.1103/PhysRevLett.45.1980. URL: https://link.aps.org/doi/10.1103/PhysRevLett.45.1980.

Strang, G. (1980). *Linear Algebra and its Applications*. Second Edition. Academic Press. ISBN: 978-0-12-673660-1.

Alcock, C. and Paczynski, B. (1979). “An evolution free test for non-zero cosmological constant.” In: *Nat.* 281, p. 358. DOI: 10.1038/281358a0.

Groth, E. J. and Peebles, P.J.E. (1977). “Statistical analysis of catalogs of extragalactic objects. VII. Two- and three-point correlation functions for the high-resolution Shane-Wirtanen catalog of galaxies.” In: *Ap.J.* 217, pp. 385–405. DOI: 10.1086/155588.

Peebles, P.J.E. and Groth, E. J. (1975). “Statistical analysis of catalogs of extragalactic objects. V. Three- point correlation function for the galaxy distribution in the Zwicky catalog.” In: *Ap. J.* 196 (1), pp. 1–11. DOI: 10.1086/153390. URL: https://www.osti.gov/biblio/4222134.

Meszaros, P. (1974). “The behaviour of point masses in an expanding cosmological substratum.” In: *A&A* 37.2, pp. 225–228.

Press, W. H. and Schechter, P. (1974). “Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation.” In: *Ap. J.* 187, pp. 425–438. DOI: 10.1086/152650.

Zeldovich, Y. B. (1972). “A hypothesis, unifying the structure and the entropy of the Universe.” In: *MNRAS* 160, 1P. DOI: 10.1093/mnras/160.1.1P.

Harrison, E.R. (1970). “Fluctuations at the Threshold of Classical Cosmology.” In: *Phys. Rev. D* 1.10, pp. 2726–2730. DOI: 10.1103/PhysRevD.1.2726.

Zel’dovich, Y. B. (1970). “Gravitational instability: An approximate theory for large density perturbations.” In: *A&A* 5, pp. 84–89.

Wick, G. C. (1950). “The Evaluation of the Collision Matrix.” In: *Phys. Rev.* 80 (2), pp. 268–272. DOI: 10.1103/PhysRev.80.268. URL: https://link.aps.org/doi/10.1103/PhysRev.80.268.

Isserlis, L. (1918). “On a Formula for the Product-Moment Coefficient of any Order of a Normal Frequency Distribution in any Number of Variables.” In: *Biometrika* 12 (1-2), pp. 134–139. URL: https://doi.org/10.1093/biomet/12.1-2.134.