

Physics Area - PhD course in
Astroparticle and Gravitational Physics

Constraining the Dark Sector with Galaxy Clustering

Candidate:
Francesco Verdiani

Advisors:
Emiliano Sefusatti
Matteo Viel

Academic Year 2025-26



Abstract

The nature of dark matter and the value of the neutrino masses remain among the central open questions in physics, and the clustering of galaxies has become one of the most promising probes for addressing them. The statistical power of current surveys, however, is outrunning our ability to fully exploit it. On the theory side, the galaxy power spectra analysis methods have been developed assuming a purely cold and non-interacting dark sector, so that the control of theoretical systematics is already a limiting factor. On the observational side, analyses restricted to the galaxy power spectrum alone are suboptimal, and ultimately limited by degeneracies that single statistics cannot break. This thesis addresses these limitations along three complementary directions. First, the redshift-space clustering of biased tracers in massive neutrino cosmologies is examined on N-body simulations, clarifying an ambiguity in the connection between halos and underlying matter. Second, a first step is taken towards extending the Effective Field Theory of Large-Scale Structure to cosmologies containing a non-cold dark matter subcomponent, including a self-consistent treatment of the coupled nonlinear evolution at one-loop level. Third, the projected bispectrum of the galaxy overdensity in combination with CMB lensing is validated and employed, for the first time, as a cosmological probe. The methods developed along these lines are then brought together in an original analysis of DESI DR1 full-shape measurements combined with Planck and ACT DR6, including lensing. This places the most stringent constraints to date on ultra-light axions as a subcomponent of dark matter, excluding, over a range of masses, subpercent fractions of the total matter density. Taken together, these studies pave the way for a more systematic and controlled exploration of the dark sector with the data of the coming decade.

Acknowledgments

My first thanks go to Emiliano and Matteo, my supervisors, for their guidance, for all I have learned from them, and for the trust and wide breadth they allowed me throughout. I extend them as well to the wider Large Scale Structure group in Trieste; being constantly involved in an engaging environment gave me much of the drive that made this work possible.

Immense thanks go to Emanuele, Ennio and Diego, for involving me in their research, excitement and flood of ideas. It really shaped the way I see cosmology.

Thanks to David, Léa and Matteo, for their hospitality and involvement, for sharing with me their expertise, and for giving me the chance to approach research from a different angle. It has been an incredibly fun and formative experience.

Outside the research environment, I thank my old friends, for always making me feel I never left home. I owe much to the new friends I met at SISSA as well, for our time together, the infinite chats, the drinking nights, and much more. They are what made my time in Trieste.

A huge thank you to my family, for supporting and encouraging me at every step.

Finally, thank you Elisa for always being by my side along this journey, for believing in me, and at times bringing me back down to earth, it wouldn't have been the same without you.

Foreword

This thesis presents part of the research carried out during doctoral studies at SISSA, which has led to the following works:

- [1] **F. Verdiani**, E. Bellini, C. Moretti, E. Sefusatti, C. Carbone, M. Viel, “Redshift-space distortions in massive neutrinos cosmologies”, *Phys. Rev. D* 112.4 (2025) [1].
- [2] **F. Verdiani**, E. Castorina, E. Salvioni, E. Sefusatti, “The Effective Field Theory of Large Scale Structure for mixed dark matter scenarios”, *JCAP* 01 (2026) [2].
- [3] **F. Verdiani**, L. Harscouet, M. Zennaro, D. Alonso, B. Hadzhiyska, “Cosmological constraints from the angular power spectrum and bispectrum of luminous red galaxies and CMB lensing”, *JCAP* 06 (2026) [3].
- [4] **F. Verdiani**, E. Castorina, E. Salvioni, E. Sefusatti, M. Viel “Constraints on Ultra-Light Axions from the DESI DR1 Full Shape, Planck and ACT”, accepted for publication in *Phys. Rev. D* [4].

Chapters 3, 4, 5 and 6 are based, respectively, on [1], [2], [3] and [4]. The corresponding material has been adapted to provide a unified presentation and to place each study within the broader context of the thesis.

Contents

Acknowledgments	iv
Foreword	v
1. Introduction	1
2. Galaxy Clustering as a Cosmological Probe	5
2.1. Theory: predicting correlation functions from first-principles in Λ CDM	6
2.1.1. Linear theory: from initial conditions to growth of perturbations	6
2.1.2. Nonlinear clustering: perturbative approach and EFTofLSS	7
2.2. Data: estimating correlation functions from simulated and real data catalogs	13
3. Massive neutrinos and Galaxy Clustering	17
3.1. Review of Massive Neutrinos effect on the large-scale structure	17
3.2. Redshift-space distortions in massive neutrino cosmologies	20
3.2.1. Linear theory	21
3.2.2. Nonlinear perturbation theory in redshift space	23
3.3. N-body simulations and measurements	25
3.4. Qualitative tests of linear RSD theory	26
3.4.1. Halo power spectrum and halo-matter crosspower spectrum	26
3.4.2. Momentum bias	29
3.5. Nonlinear theory fits	29
3.5.1. Data and analysis setup	29
3.5.2. Analytical Cross-power spectra Gaussian covariance matrices	31
3.5.3. Results: model selection and velocity bias	33
3.6. Concluding remarks	34
4. EFTofLSS and Mixed Dark Matter	37
4.1. Motivation	37
4.2. Mixed DM Linear Cosmology	39
4.2.1. General linear solutions	42
4.2.2. Matching to microscopic models	44
4.2.3. Ultra-light axions	44
4.2.4. Light but massive thermal relics	46
4.3. Galaxy Power Spectrum in Redshift Space	49
4.3.1. Matter nonlinearities	49
4.3.2. Bias expansion with two fluids	54
4.3.3. Redshift space distortions	56

4.3.4.	Prescription for nonlinear kernels	57
4.4.	A first application: Planck+BOSS constraints on Ultra-Light Axions	64
4.4.1.	Data specification and modeling choices	64
4.4.2.	Parameters and priors	66
4.4.3.	Results	67
4.5.	Concluding remarks	71
Appendix 4.A.	Further details on perturbation theory in Mixed DM cosmologies	71
4.A.1.	Universality of Perturbation Dynamics in Scale-free Beyond- Λ CDM Scenarios	71
4.A.2.	More Analytical Expressions for Nonlinear Kernels	72
4.A.3.	Impact of Intrinsic Nonlinearities	73
5.	Cross-Correlation with CMB Lensing and Higher Orders	79
5.1.	Motivation	79
5.2.	Data	80
5.2.1.	Public galaxy and CMB lensing data	80
5.2.2.	Power spectrum and bispectrum estimation	81
5.3.	Methodology	83
5.3.1.	Simulations	83
5.3.2.	Theory model	84
5.3.3.	Accelerating theory predictions	87
5.3.4.	Likelihood	88
5.4.	Results	90
5.4.1.	Validation on simulations	90
5.4.2.	Cosmological constraints from 2- and 3-point correlators	94
5.4.3.	Galaxy bias relations from projected clustering	99
5.5.	Concluding remarks	101
Appendix 5.A.	Additional tests	103
5.A.1.	Treatment of Redshift-space distortions	103
5.A.2.	Simpler bias parametrisations	104
5.A.3.	Emulators accuracy	105
6.	DESI DR1 Full-Shape Constraints on Ultra-Light Axions	107
6.1.	Motivation	107
6.2.	Data and Methods	108
6.2.1.	Datasets	108
6.2.2.	CMB theoretical predictions	108
6.2.3.	Galaxy clustering model	111
6.2.4.	Parameters and priors	111
6.3.	Results	112
6.4.	Concluding remarks	116
7.	Conclusion and outlooks	119
	Bibliography	123

1. Introduction

Cosmology has turned into a precision science over the past decades. From limited order-of-magnitude estimates, a succession of increasingly ambitious observational campaigns transformed it into a field where the parameters of a cosmological model are routinely quoted at the percent level. The cosmic microwave background (CMB) has been the flagship of this transformation: measurements of the temperature and polarization anisotropies, culminating in the *Planck* mission [5] and further sharpened by ground-based experiments [6, 7], have provided an exquisitely detailed snapshot of the Universe at the epoch of recombination and pinned down the six parameters of the standard Λ CDM model with remarkable accuracy. Yet the CMB is, in an important sense, a finite resource: it is a two-dimensional image of a single surface, and with the primary anisotropies now measured close to the cosmic-variance limit over a wide range of scales, the information that can still be extracted from it is beginning to saturate.

The distribution of matter traced by galaxies offers a complementary and, in the long run, possibly larger source of information. As spectroscopic galaxy surveys probe a three-dimensional volume rather than a surface, and they sample the Universe across a range of cosmic epochs, the number of accessible modes grows rapidly with survey depth and area. This makes galaxy clustering one of the most promising probes of cosmology for the coming decades. Building on the legacy of SDSS-era spectroscopy [8, 9], a new generation of surveys, including DESI [10], Euclid [11], and the Vera C. Rubin Observatory [12], is now mapping the large-scale structure with unprecedented volume and precision, and the resulting datasets have the potential to be transformative.

Crucially, this observational leap is matched by comparable advances on the theoretical side. The clustering of galaxies on large scales is now understood from first principles within cosmological perturbation theory [13] and, in particular, within the Effective Field Theory of Large-Scale Structure (EFTofLSS) [14–17], which provides controlled, systematically improvable predictions for correlation functions in the mildly non-linear regime; predictions that have by now been successfully confronted with data [18–21]. In this case, within Λ CDM, the theory is arguably ahead of the data: we possess robust predictions well before the surveys that will fully exploit them have completed their observations.

It is worth being precise, however, about what it is that we understand so well. The theoretical control achieved so far is control *within* Λ CDM: the framework tells us how to predict clustering statistics given a cosmological model, and it allows us to measure the parameters of that model with ever-greater precision. The deeper questions in cosmology, however, are not questions of parameter estimation but concern the nature of the model itself: the fundamental nature of dark matter (DM), the origin of the accelerated expansion, and whether the neutrino and dark sectors are really as minimal as Λ CDM assumes. Pinning down the parameters of a model we take for granted is, in the end, less interesting than discovering that the model is incomplete. It is precisely in the search for physics beyond Λ CDM that galaxy clustering has the most to offer, and it is here that this thesis seeks to contribute.

Despite focusing primarily on galaxy clustering, it should be mentioned that much of the future

1. Introduction

constraining power will come not from any single probe in isolation but from the combination of complementary datasets. The joint analysis of galaxy clustering with CMB lensing [22–24], for instance, breaks parameter degeneracies that neither probe can resolve on its own. This cross-correlation programme is still in its early stages, and there remains substantial room for improvement, both in the statistical exploitation of overlapping surveys and in the theoretical modelling required to interpret them consistently. The work presented in this thesis is situated within this broader effort to extract new physics by combining probes of the large-scale structure of the Universe.

Among the many extensions of the minimal model, the mass of neutrinos occupies a special place. Neutrino masses are not a speculative addition to the Standard Model: neutrino oscillation experiments have established beyond doubt that neutrinos are massive [25], making $\sum m_\nu$ a parameter that we already know to be non-zero. In this sense it is not exotic new physics at all, but a guaranteed ingredient whose value simply remains to be measured. This makes it exceptionally well motivated as a target: the question is not whether the effect is there, but at what level it will be detected. Remarkably, cosmology is poised to measure this fundamental particle-physics parameter [26] before terrestrial laboratory experiments do [27, 28], inferring the absolute neutrino mass scale from its imprint on the growth of cosmic structure rather than from kinematics in the laboratory.

A further peculiarity makes the neutrino-mass problem uniquely sharp. Oscillation experiments measure the squared-mass splittings between the eigenstates, and these translate into an absolute *lower* bound on the sum, $\sum m_\nu \gtrsim 60$ meV in the normal ordering (and $\gtrsim 100$ meV in the inverted ordering). We therefore know, before looking, that the signal must be present at or above a definite floor. This has two important consequences. First, it tells us precisely where to search: there is a well-defined target value that upcoming surveys are designed to reach [29, 30]. Second, and more subtly, it provides a rare and powerful consistency test of the entire cosmological inference pipeline. Any robust cosmological measurement must return $\sum m_\nu \gtrsim 60$ meV; a confident determination that fell below this floor, or that preferred the unphysical region $\sum m_\nu < 0$, would signal not a discovery about neutrinos but a failure of our modelling or our assumptions. Few parameters in cosmology come with such an external, independent benchmark against which the analysis can be validated.

We are, right now, on the verge of this measurement. The tightest cosmological constraints have pushed down to the point where they are beginning to probe the very floor set by oscillations [31], and a genuine detection appears to be within reach of the current and next generation of surveys. This proximity, however, is precisely what makes the situation delicate. The most aggressive bounds already sit in mild tension with the oscillation lower limit, and some data combinations even drift toward unphysical negative values of $\sum m_\nu$ [32–35]. Whether this reflects a statistical fluctuation, unaccounted-for systematics, or the propagation into the neutrino sector of the broader cosmological tensions concerning H_0 [36, 37] and the dark energy equation of state [31, 38], remains unresolved. These tensions must be understood and resolved before any near-threshold signal can be interpreted as a robust detection rather than an artifact. It is in this context of a guaranteed but yet unmeasured signal, sitting just at the edge of detectability, that the careful modelling of neutrino effects in galaxy clustering, to which this thesis contributes, becomes very relevant.

If neutrinos represent a case in which we know that the signal exists, and even where to look for it, with dark matter we are in precisely the opposite situation. Although it is a pillar of the standard cosmological model, the fundamental nature of dark matter has so far remained elusive. All current

experimental evidence stems solely from its gravitational interactions, and the dark sector may be entirely decoupled from the Standard Model, leaving astrophysics and cosmology as possibly our only windows into its properties. Without a guaranteed target value or a floor to anchor the search, there is no single decisive measurement to aim for; instead, every available probe must be exploited to look for any deviation from the simple cold, collisionless picture assumed by Λ CDM.

Different probes offer deeply complementary perspectives. Astrophysics targets localized, often highly non-linear systems, yielding tests of the small-scale behaviour of the dark sector [39], while cosmology provides a highly robust, statistically grounded probe on the largest scales. Small-scale cosmological probes such as the Lyman- α forest are sensitive to the *bulk* nature of the dark matter: a pronounced suppression of clustering, for instance, would signal a warm, or otherwise non-cold, component making up the whole of the DM [40–42]. Galaxy clustering, conversely, on large and mildly non-linear scales offers something different in character. Its strength lies not in reaching the deeply non-linear regime but in *precision*: the combination of the transformative statistical power of upcoming datasets with the robustness and systematic control of modern perturbative theory makes it possible to perform genuine precision tests of the dark sector. Rather than probing order-unity modifications, this allows one to search for and constrain subdominant, percent-level subcomponents of the dark matter; for instance, a small admixture of light relics or of a non-cold species hiding within an otherwise standard cold dark matter background [43–45].

It is this regime, that is that of subtle departures from Λ CDM, accessible only through the joint precision of data and theory, that motivates much of the work in this thesis. Light relics of the early Universe provide a particularly compelling and well-defined target in this respect [30, 46]: whether or not they constitute the dominant dark matter, their presence would leave a characteristic imprint on the growth of structure, and searching for it is a natural application of the precision galaxy-clustering programme developed in the following chapters.

So far this discussion has implicitly centered on the two-point statistics of the galaxy field, the power spectrum, which is certainly the prime observable of large-scale structure analyses. But the information contained in the cosmic web is far richer than its two-point function alone, and moving forward the programme of precision dark-sector tests will require going beyond it. The gravitational evolution that shapes the large-scale structure is non-linear, and it therefore imprints non-Gaussian features that are invisible to the power spectrum but accessible through higher-order statistics such as the bispectrum [47–50]. Equally, the same underlying density field is probed by more than one observable: cross-correlations between galaxy clustering and, for instance, CMB lensing tie the galaxy distribution directly to the total matter field, breaking degeneracies, possibly being less prone to certain systematics, and adding statistical weight that neither probe carries alone. Projected clustering statistics, which trade some radial information for robustness and for access to wider sky areas, offer in this sense a complementary handle [51]. Each of these directions enlarges the space of observables in which a signal of new physics might reveal itself, and each demands that the beyond- Λ CDM modelling be extended accordingly. Part of this thesis takes a concrete step in this direction of developing the analysis of new observables, through a study of the galaxy bispectrum in combination with CMB lensing, which stands out as one of the central directions to pursue in the future.

More broadly, the search for the nature of dark matter through cosmological structure can be seen as a guiding thread of the programme to which this thesis belongs. The dark sector may never be brought

1. Introduction

into a laboratory; what remains possible is to exploit, to their fullest, the probes the Universe itself provides, and to extract from them every constraint they are capable of yielding. The coming decades will deliver an extraordinary wealth of data deeper, wider, and more precise than anything available today. Buried within it there may well be some signal of physics beyond the cold, collisionless paradigm. Whether or not such a signal is present, it will only be recognised if the tools are ready for it. That preparation is, above all, a theoretical task: it requires robust, systematically controlled, beyond- Λ CDM modelling capabilities, across the full range of observables the data will populate, so that any departure from the standard model can be identified for what it is rather than, for instance, absorbed into a mis-modelled nuisance parameter. Building and sharpening a part of that toolkit, concerning neutrinos and for non-cold species and across power spectrum, bispectrum, and cross-correlations, is the aim of the present thesis.

This thesis is structured as follows. In Chapter 2 we offer a brief review of the canonical way in which galaxy clustering observations are exploited to infer fundamental cosmology, connecting theory to observations.

Chapter 3 extends this to the peculiarities of cosmologies with massive neutrinos [26], and presents an original study of the redshift-space distortions of biased tracers in such cosmologies, carried out through the analysis of N-body simulations implementing this physics [52, 53].

Chapter 4 presents the main theoretical development of this thesis, laying the groundwork for extending the EFTofLSS to cosmologies containing a small fraction of non-cold dark matter, in generic scenarios rather than specialising to massive neutrinos. The resulting framework is then applied to real data, as a proof-of-principle, to the already exploited BOSS dataset [8].

Turning to new observables, Chapter 5 presents a Λ CDM analysis of projected clustering in combination with CMB lensing, extended for the first time to include a, compressed, bispectrum correlator.

Chapter 6 brings these directions together, deriving novel constraints on ultra-light axions [54–57] from the combination of DESI DR1 spectroscopic galaxy clustering [38] with CMB data including lensing [5, 23]. The analysis draws on essentially every element developed in the preceding chapters, the treatment of a non-cold component within the EFTofLSS, the problem of biased tracers in the presence of a scale-dependent growth [58], and the modelling of CMB lensing, and is in this sense the point at which the tools developed along the way are finally put to work together.

Finally, Chapter 7 presents the conclusions of this work and outlines directions for future research.

2. Galaxy Clustering as a Cosmological Probe

This chapter briefly reviews the method by which galaxies, stochastic tracers of a random field, can be used to constrain fundamental cosmology.

The distribution of galaxies across the sky contains a vast amount of cosmological information, which is statistical in nature. The relevant quantity is therefore not the position of a single galaxy, that reflects both the underlying structure and the stochastic process of galaxy formation, but the fluctuation of the galaxy number density around its mean value,

$$\delta_g(\mathbf{x}) = \frac{n_g(\mathbf{x}) - \bar{n}_g}{\bar{n}_g}, \quad (2.1)$$

which defines a continuous random field encoding the spatial clustering of galaxies. Since δ_g is a realization of a stochastic process, the natural objects to study are not the field values themselves but their statistical moments, characterized by the hierarchy of correlation functions $\langle \delta_g(\mathbf{x}_1) \cdots \delta_g(\mathbf{x}_n) \rangle$, where the angle brackets denote an ensemble average over all possible realizations of the Universe. This immediately raises a conceptual problem: what theory can predict ensemble averages, yet our Universe is a single realization of this ensemble, making a direct comparison seemingly impossible. The way out relies on the ergodic hypothesis: we assume that for a statistically homogeneous random field, ensemble averages coincide with spatial averages over a sufficiently large volume V , i.e.

$$\langle \delta_g(\mathbf{x}_1) \cdots \delta_g(\mathbf{x}_n) \rangle = \lim_{V \rightarrow \infty} \frac{1}{V} \int_V d^3y \delta_g(\mathbf{x}_1 + \mathbf{y}) \cdots \delta_g(\mathbf{x}_n + \mathbf{y}). \quad (2.2)$$

In practice, this means that a survey covering a large enough volume allows us to treat different spatial regions as approximately independent realizations of the same statistical process, and thereby to perform meaningful inference with statistical methods.

As it will become clear in what follows, it is however more natural to work in Fourier space for both theoretical predictions and data analysis, since translational invariance makes different Fourier modes statistically independent and, at linear order, dynamically decoupled, so that each wavenumber evolves independently and carries its own cosmological information. Hence the field we will deal with is actually

$$\delta_g(\mathbf{k}) = \int d^3x \delta_g(\mathbf{x}) e^{-i\mathbf{k} \cdot \mathbf{x}} \quad (2.3)$$

Consequently, as the first moment is zero by definition Eq. (2.1), the most relevant correlator is the power spectrum $P_{gg}(k)$, defined by

$$\langle \delta_g(\mathbf{k}) \delta_g^*(\mathbf{k}') \rangle = (2\pi)^3 \delta^D(\mathbf{k} - \mathbf{k}') P_{gg}(k), \quad (2.4)$$

where statistical homogeneity, i.e. invariance under translations, enforces the Dirac delta setting $\mathbf{k} = \mathbf{k}'$, while isotropy ensures that P_{gg} depends only on the magnitude $k = |\mathbf{k}|$.

Eq. (2.4) will be the most important quantity for this work, as the connecting point of theory and galaxy surveys datasets for the purpose of cosmological inference. In this chapter, we first lay out the

theoretical framework used to predict galaxy correlation functions from first principles, from the linear evolution of cosmological perturbations to the nonlinear, EFT-based, description of biased tracers; then we outline basic details on how such correlation functions are estimated in practice from simulated and observational data.

2.1. Theory: predicting correlation functions from first-principles in Λ CDM

In this section we review how the statistics of the large-scale structure can be predicted from first principles within the standard Λ CDM cosmological model. We begin in Sect. 2.1.1 with linear theory, following the evolution of density perturbations from their primordial initial conditions to the growth of structure at later times. We then move beyond the linear regime in Sect. 2.1.2, introducing the perturbative treatment of non-linear clustering and its state-of-the-art formulation within the EFTofLSS. Building on this, we discuss the steps needed to connect the theory to observations: the bias expansion relating galaxies to the underlying matter field, the mapping to redshift space, and the stochastic contributions. Together, these elements provide the theoretical backbone for the correlation functions and power spectra used throughout this work.

2.1.1. Linear theory: from initial conditions to growth of perturbations

The seeds of cosmic structure are, supposedly, laid during inflation, which generates a nearly scale-invariant spectrum of primordial curvature perturbations [59]. In the simplest scenario, this is characterised by a primordial power spectrum of the form

$$\mathcal{P}_0(k) = A_s \left(\frac{k}{k_*} \right)^{n_s-1}, \quad (2.5)$$

where A_s is the amplitude, n_s the spectral index, and k_* a reference pivot scale. These perturbations then evolve through the subsequent cosmological epochs, steered by gravity and the microphysics of the various species that make up the Universe. In the early Universe, density perturbations are small, $|\delta_i| \ll 1$, and their evolution can be studied by linearising the equations of motion around the homogeneous background, a regime known as linear perturbation theory. The resulting Einstein-Boltzmann system [59] couples the density contrasts and velocity divergences θ_i of each species i (i.e. photons, baryons, cold dark matter, and neutrinos in the standard Λ CDM model) to the metric perturbations via the perturbed Einstein equations. This system encodes a rich variety of physical processes: the oscillation and diffusion damping of the baryon-photon plasma before recombination, the decay of gravitational potentials on scales that entered the horizon during radiation domination, and the gravitational coupling between the different species. The full solution is in general too complex to be obtained analytically, and in practice it is computed numerically using Einstein-Boltzmann codes such as CLASS [60] or CAMB [61]. The net effect of this early-universe evolution on the matter distribution is customarily encoded in the transfer function $T(k)$, which describes the scale-dependent evolution of perturbations relative to their primordial values.

Once matter domination is established and the relevant scales are well within the Hubble horizon, $k \gg \mathcal{H}$, the evolution simplifies considerably: photons and baryons have decoupled after recombination,

and the growth of structure is driven by the matter component. In this regime, the linear continuity and Euler equations for a cold matter fluid reduce to

$$\begin{cases} \delta'_c + \theta_c = 0 \\ \theta'_c + \mathcal{H}\theta_c + \frac{3}{2}\mathcal{H}^2\Omega_m\delta_m = 0 \end{cases} \quad (2.6)$$

where primes denote conformal time derivatives, $\mathcal{H} = aH$ is the conformal Hubble rate, and δ_m is the total matter density contrast sourcing the gravitational potential and $\theta_c = \nabla \cdot \mathbf{v}_c$, being v_c the velocity field. In the rest of this section we will specialize Eqs. (2.6) to a massless neutrinos scenario, i.e. $\delta_m = \delta_c$ as there is only a cold matter component, leaving the interesting phenomenology that $\delta_m \neq \delta_c$ entails as the subject of Chapter 3 and 4.

The system (2.6) admits a growing-mode solution $\delta_c \propto D(a)$, where $D(a)$ is the linear growth factor, normalised to unity today, with $\theta_c = -\mathcal{H}f\delta_c$ and the growth rate

$$f \equiv d \log D / d \log a \simeq \Omega_m(a)^\gamma \quad (2.7)$$

for some power $\gamma \approx 0.55$ in Λ CDM [13, 62]. Combining the primordial spectrum, the transfer function, and the late-time growth, one can give a *linear* prediction for the matter power spectrum as

$$P^L(k, z) \propto D^2(z)T^2(k)\mathcal{P}_0(k), \quad (2.8)$$

with its characteristic shape: a broad peak near the scale of matter-radiation equality $k_{\text{eq}} \sim 0.01 h \text{ Mpc}^{-1}$, a suppression on smaller scales due to the decay of potentials during radiation domination, and the baryon acoustic oscillation (BAO) feature imprinted by sound waves in the baryon-photon plasma, visible as a series of oscillations around $k \sim 0.1 h \text{ Mpc}^{-1}$.

2.1.2. Nonlinear clustering: perturbative approach and EFTofLSS

While Eq. (2.8) is a remarkable first-principles prediction, it is not what can be directly compared to galaxy clustering observations for the purpose of cosmological inference. There are two main reasons for this, which we address in turn. First, at the low redshifts where galaxies are surveyed, the density perturbations are no longer small, invalidating the assumption of linear evolution. Second, there is no straightforward connection between the observed galaxy overdensity field and the underlying matter overdensity that the theory predicts in Eq. (2.8).

In principle, both complications can be tackled fully numerically via N -body simulations. This is indeed a very active field of research, that of Simulation-Based Inference (e.g., for galaxy clustering, see [63, 64]), whose main challenge lies in the fact that simulations are computationally expensive and thus difficult to embed directly within a parameter inference pipeline. An analytical route is, at the same time, that of cosmological non-linear perturbation theory, which treats the density and velocity fields as a controlled expansion around the linear solution. Valid on scales where δ is not yet of order unity but where corrections to linear theory are nonetheless important, it captures the leading effects of non-linear gravitational evolution and provides a systematic framework for connecting the galaxy field to the underlying matter distribution through a bias expansion. This makes it a powerful and controlled tool, well suited to searches for beyond- Λ CDM physics, and it is the approach we adopt in the following.

Standard Perturbation Theory

Standard Perturbation Theory (SPT) [13] provides a systematic expansion of the matter density contrast and velocity divergence, conveniently through the quantity $\Theta \equiv -\theta/(f\mathcal{H})$, in powers of the linear density field $\delta^{[1]}$,

$$\delta(\mathbf{k}) = \delta^{[1]}(\mathbf{k}) + \delta^{[2]}(\mathbf{k}) + \delta^{[3]}(\mathbf{k}) + \dots \quad (2.9)$$

where the n -th order term $\delta^{[n]}$ is an n -fold convolution of the linear field with mode-coupling kernels F_n and G_n ,

$$\delta^{[n]}(\mathbf{k}) = \int_{\mathbf{q}} dq_{1\dots n} F_n(\mathbf{q}_1, \dots, \mathbf{q}_n) \delta^{[1]}(\mathbf{q}_1) \dots \delta^{[1]}(\mathbf{q}_n) \quad (2.10)$$

$$\Theta^{[n]}(\mathbf{k}) = \int_{\mathbf{q}} dq_{1\dots n} G_n(\mathbf{q}_1, \dots, \mathbf{q}_n) \delta^{[1]}(\mathbf{q}_1) \dots \delta^{[1]}(\mathbf{q}_n) \quad (2.11)$$

where we have defined the shorthand notation

$$\int_{\mathbf{q}} dq_{1\dots j} \equiv \int \frac{d^3 q_1 \dots d^3 q_j}{(2\pi)^{3(j-1)}} \delta^D\left(\mathbf{k} - \sum_{i=1}^j \mathbf{q}_i\right), \quad (2.12)$$

The kernels F_n and G_n are fully determined by the nonlinear equations of motion of a pressureless fluid in an expanding universe, i.e. the Eqs. (2.6) but without neglecting the quadratic source terms as in linear theory. This is conveniently rewritten in terms of Θ and switching time variable, via the linear growth factor D appearing in Eq. (2.7), to

$$\eta \equiv \log D(\tau) \quad (2.13)$$

as [13]

$$\begin{cases} \partial_\eta \delta - \Theta = \int_{\mathbf{q}} dq_{12} \alpha(\mathbf{q}_1, \mathbf{q}_2) \Theta(\mathbf{q}_1) \delta(\mathbf{q}_2) \\ \partial_\eta \Theta + \frac{1}{2} \Theta - \frac{3}{2} \delta = \int_{\mathbf{q}} dq_{12} \beta(\mathbf{q}_1, \mathbf{q}_2) \Theta(\mathbf{q}_1) \Theta(\mathbf{q}_2) \end{cases} \quad (2.14)$$

having defined

$$\alpha(\mathbf{k}_1, \mathbf{k}_2) \equiv 1 + \frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{k_1^2}, \quad \beta(\mathbf{k}_1, \mathbf{k}_2) \equiv \frac{(\mathbf{k}_1 + \mathbf{k}_2)^2 \mathbf{k}_1 \cdot \mathbf{k}_2}{2k_1^2 k_2^2}, \quad (2.15)$$

In Eqs. (2.14) we also used the excellent approximation $\Omega_m(a)/f^2(a) \simeq 1$, valid at the sub-percent level in Λ CDM, which renders the kernels time-independent. The kernels can be computed recursively giving, for instance, at second order

$$F_2(\mathbf{k}_1, \mathbf{k}_2) = \frac{5}{7} \alpha_s(\mathbf{k}_1, \mathbf{k}_2) + \frac{2}{7} \beta(\mathbf{k}_1, \mathbf{k}_2) \quad (2.16)$$

$$G_2(\mathbf{k}_1, \mathbf{k}_2) = \frac{3}{7} \alpha_s(\mathbf{k}_1, \mathbf{k}_2) + \frac{4}{7} \beta(\mathbf{k}_1, \mathbf{k}_2) \quad (2.17)$$

with the symmetric component as

$$\alpha_s(\mathbf{k}_1, \mathbf{k}_2) \equiv \frac{1}{2} [\alpha(\mathbf{k}_1, \mathbf{k}_2) + \alpha(\mathbf{k}_2, \mathbf{k}_1)]. \quad (2.18)$$

Expanding the power spectrum to one-loop order, i.e. retaining contributions up to fourth order in the linear density field, one obtains

$$P_{mm}(k) = P_{11}(k) + P^{1\text{-loop}}(k), \quad (2.19)$$

where the one-loop correction is made up of two contributions stemming from two combinations of Eq. (2.9)

$$P^{1\text{-loop}}(k) = P_{22}(k) + P_{13}(k), \quad (2.20)$$

with

$$P_{22}(k) = 2 \int \frac{d^3q}{(2\pi)^3} [F_2(\mathbf{q}, \mathbf{k} - \mathbf{q})]^2 P_{11}(q) P_{11}(|\mathbf{k} - \mathbf{q}|), \quad (2.21)$$

$$P_{13}(k) = 6 P_{11}(k) \int \frac{d^3q}{(2\pi)^3} F_3(\mathbf{k}, \mathbf{q}, -\mathbf{q}) P_{11}(q). \quad (2.22)$$

The P_{22} term arises from the product of two second-order fields, while P_{13} comes from the cross-correlation between the linear and third-order fields. Together, these corrections are the first step to extend the range of validity of the perturbative description to mildly nonlinear scales. Formally, the expansion is controlled by the ratio k/k_{NL} , where $k_{\text{NL}} \sim 0.3\text{--}0.5 h \text{ Mpc}^{-1}$ at $z = 0$ is defined by $k_{\text{NL}}^3 P^L(k_{\text{NL}})/(2\pi^2) \sim 1$. In practice, one-loop analyses of galaxy clustering in redshift-space are typically restricted to $k_{\text{max}} \simeq 0.2 h \text{ Mpc}^{-1}$, a choice validated on simulations.

A well-known limitation of SPT, however, is the behavior of Eqs. (2.21)–(2.22) at the integration boundaries, which exposes the framework to both infrared (IR) and ultraviolet (UV) sensitivities. In the IR limit ($q \rightarrow 0$), the integrands of P_{22} and P_{13} individually exhibit strong divergences. Physically, these long-wavelength modes correspond to large-scale bulk flows. By virtue of the Galilean invariance, a uniform large-scale flow simply translates small-scale structures and cannot alter the internal dynamics. Consequently, when P_{22} and P_{13} are summed, these leading IR divergences exactly cancel one another, ensuring that the total one-loop power spectrum remains well-behaved and safely decoupled from arbitrarily large scales [65, 66].

Conversely, the UV limit ($q \gg k$) presents a fundamental inconsistency of the theory, if constructed as above. The loop integrals receive contributions from arbitrarily large wavenumbers, probing scales where the density field is highly nonlinear and the perturbative expansion itself is invalid. Specifically, by taking the small-scale limit of the F_3 kernel, the P_{13} integral exhibits the leading UV divergence [47, 66]

$$P_{13}(k) \Big|_{q \gg k} = -\frac{61}{630\pi^2} k^2 P_{11}(k) \int_{k \ll q} \frac{q^2 dq}{q^2} P_{11}(q). \quad (2.23)$$

This is an integral that scales as $\int dq P_{11}(q)$, leaving the final result heavily dependent on the uncontrolled, highly nonlinear dynamics of small-scale structures. Operatively, in our standard cosmology the integrals actually happen to mathematically converge: this stems from the fact that the linear power spectrum integrand is strongly suppressed at small scales due to sub-horizon evolution during the radiation-dominated epoch, specifically known as the Mészáros effect. However, it is clear that this convergence is entirely fortuitous: had there been no radiation domination, the integral would formally diverge, therefore signaling in any event a breakdown of the theory. Also, even though it evaluates to a finite value in practice, integral Eq. (2.23) amounts to spurious contribution that heavily weights uncontrolled small-scale dynamics and would affect inference. It is indeed now well established that UV divergence is not a physical phenomenon but an artifact of treating the dark matter as a perfect, pressureless fluid on all scales, improperly neglecting the finite extent of virialized halos, shell-crossing, and the resulting small-scale velocity dispersion.

The Effective Field Theory of Large-Scale Structure

The Effective Field Theory of Large-Scale Structure (EFTofLSS) [14–16, 67, 68] provides a systematic framework to make SPT self-consistent, including tackling the UV sensitivity presented at 1-loop order by Eq (2.23). The key idea is that, on large scales, the details of the small-scale dynamics (such as, the formation of halos, the non-linear collapse of dark matter, the complex baryonic physics) cannot be resolved, but their effect on large-scale perturbations can be captured by a set of effective fluid equations. This is achieved by coarse-graining the full equations of motion over a smoothing scale Λ^{-1} , larger than the nonlinear scale k_{NL}^{-1} . The resulting equations describe a fluid that, even in the case of starting from an ideal fluid, acquires an effective stress tensor τ_{ij} , whose long-wavelength contributions encode the backreaction of small-scale dynamics on the large-scale modes.

In practice, the effect of this stress tensor at leading order in derivatives and at one-loop level introduces corrections to the equations of motion that take the form of a pressure-like term proportional to $k^2 \delta^{[1]}$. These are the EFT counterterms, which appear in the power spectrum as

$$P_{mm}(k) = P^L(k) + P^{1\text{-loop}}(k) - 2c_m^2 k^2 P^L(k) + \dots \quad (2.24)$$

where c_m^2 is a free parameter whose value is not predicted, and must be fitted to data or from numerical simulations. From a complementary point of view, one can notice that the counterterm $-2c_m^2 k^2 P^L(k)$ has precisely the effect to absorb the UV sensitivity Eq. (2.23): it has the same $k^2 P^L(k)$ scaling, and its free coefficient can be tuned to cancel the dependence on the cutoff Λ , making the final result UV-finite and independent of the arbitrary smoothing scale. Additional counterterms arise at higher orders in derivatives, suppressed by further powers of k^2/k_{NL}^2 , and are typically neglected at the order of accuracy we work at. The EFTofLSS therefore provides a systematic, self-consistent perturbative expansion for the matter power spectrum, whose validity can be systematically extended to higher accuracy by including higher-order loop corrections.

Galaxy Bias

So far, we have presented the galaxy power spectrum Eq. (2.4) as the primary statistics for doing cosmology with galaxy surveys, but have derived the theoretical predictions for the matter fields, culminating in the EFTofLSS prediction Eq. (2.24). A key step is still missing, namely the galaxy–matter connection; as we show below, the EFT framework remarkably accommodates this step too, again in a controlled, first-principles fashion.

In the same regime where perturbation theory for matter holds, the relation between galaxies and matter can, in full generality, be treated as the most general local functional of the long-wavelength gravitational and matter fields, evaluated along the fluid trajectory and consistent with the symmetries of the problem [69]. Locality here is understood in the sense dictated by the equivalence principle: the galaxy density at a given point depends on the second derivatives of the gravitational and velocity potentials along the past trajectory *of the fluid element*, which in Eulerian coordinates generates a set of otherwise non-local operators.

This so-called *bias expansion* is, at any given perturbative order, organised as a sum over all operators $\mathcal{O}$ compatible with the symmetries, with unknown bias coefficients,

$$\delta_g(\mathbf{x}) = \sum_{\mathcal{O}} b_{\mathcal{O}} \mathcal{O}(\mathbf{x}). \quad (2.25)$$

At first order, the only relevant operator is the linear density contrast $\delta^{[1]}$, giving $\delta_g = b_1 \delta_m$ with the linear bias parameter b_1 . At second order, two new operators appear: the quadratic density $(\delta^{[1]})^2$ and the squared tidal field $\mathcal{G}_2 \equiv (\partial_i \partial_j \phi)^2 - (\nabla^2 \phi)^2$, where ϕ is the gravitational potential. At third order the operators δ^3 , $\delta \mathcal{G}_2$, $\mathcal{G}_3$ and Γ_3 appear. Only those correlating with $\delta^{[1]}$ contribute to the one-loop power spectrum, and their contributions are all proportional to $P^L(k)$ times a loop integral: they are therefore mutually degenerate, and can be captured by a single effective parameter, conventionally chosen as b_{Γ_3} . The full bias expansion at one-loop order therefore involves the parameters $\{b_1, b_2, b_{\mathcal{G}_2}, b_{\Gamma_3}\}$, where we follow the notation of [19].

In addition, the EFT framework comprises higher-derivative bias operators. These arise because galaxy formation is not strictly local, but responds to the environment over a finite spatial scale R_* of order the Lagrangian radius of the host halo; each additional derivative is therefore suppressed by a power of $(kR_*)^2$, defining a controlled expansion

$$\delta_g \supset b_{\nabla^2} \nabla^2 \delta^{[1]} + b_{\nabla^4} \nabla^4 \delta^{[1]} + \dots, \quad (2.26)$$

whose leading term is $b_{\nabla^2} \nabla^2 \delta^{[1]}$. At the level of the power spectrum this leading operator contributes a term proportional to $k^2 P^L(k)$, with exactly the same scale dependence as the EFT counterterm for matter; the two are therefore degenerate and can be reabsorbed into a single effective parameter.

The one-loop galaxy power spectrum, assembling all contributions from the bias expansion and the EFT counterterms, can then be written schematically as

$$P_{gg}(k) = b_1^2 P_{mm}^{1\text{-loop}}(k) + P_{\text{bias}}^{1\text{-loop}}(k) - 2c_{gg}^2 k^2 P^L(k). \quad (2.27)$$

Eq. (2.27) is thus a fully theoretical prediction for P_{gg} that can be directly tested on simulations. For real data, and hence for cosmological inference, one complication (which will actually turn out to be an opportunity for additional information) still remains, i.e. that the position of galaxies along the line of sight can only be inferred through the redshift. We treat this in detail in what follows.

Redshift-space distortions

Observational galaxy catalogs do not natively provide three-dimensional spatial maps; their fundamental data products actually consist of a pair of angular coordinates and a measured redshift for each target. To map the large-scale structure in 3D and construct the density field of Eq. (2.1), one must assume a fiducial cosmological model to convert these observed redshifts into radial (comoving) distances. However, this transformation is complicated by the fact that the observed redshift does not stem exclusively from the cosmological Hubble expansion, but also has a Doppler contribution from the galaxy's peculiar velocity $\mathbf{v}$ along the line of sight $\hat{\mathbf{z}}$. The mapping between the true real-space position $\mathbf{x}$ and the inferred redshift-space position $\mathbf{s}$ is therefore

$$\mathbf{s} = \mathbf{x} + \frac{v_z(\mathbf{x})}{\mathcal{H}} \hat{\mathbf{z}}, \quad (2.28)$$

where $v_z = \mathbf{v} \cdot \hat{\mathbf{z}}$ is the line-of-sight peculiar velocity. This displacement distorts the observed galaxy distribution with respect to the true one, an effect known as redshift-space distortions (RSD). As galaxies are treated as a set of discrete tracers, the density computed from observations $\delta_g^s(\mathbf{s})$ is necessarily related to the real- (as opposed to redshift-) space one by number conservation

$$[1 + \delta_g^s(\mathbf{s})] d^3 s = [1 + \delta_g(\mathbf{x})] d^3 x \quad (2.29)$$

2. Galaxy Clustering as a Cosmological Probe

At the linear level, one can further make use of Eq. (2.6) to obtain, in the plane-parallel approximation where $\hat{\mathbf{z}}$ is taken as fixed,

$$\delta_g^s(\mathbf{k}) = \delta_g(\mathbf{k}) - \frac{ik_z}{\mathcal{H}} v_z(\mathbf{k}) = (b_1 + f\mu^2) \delta^{[1]}(\mathbf{k}), \quad (2.30)$$

where $\mu \equiv k_z/k = \hat{\mathbf{k}} \cdot \hat{\mathbf{z}}$ is the cosine of the angle between the wavevector and the line of sight, and f is the linear growth rate defined in Eq. (2.7). This is the Kaiser formula [70], which shows that RSD introduce an anisotropy, i.e. a μ dependence that was absent in e.g. Eq. (2.27), in the galaxy power spectrum: at linear order,

$$P_{gg}^s(k, \mu) = (b_1 + f\mu^2)^2 P_{11}(k). \quad (2.31)$$

The angular dependence of Eq. (2.31) can be decomposed into Legendre multipoles

$$P_\ell(k) \equiv \frac{2\ell+1}{2} \int_{-1}^{+1} d\mu P_{gg}^s(k, \mu) \mathcal{L}_\ell(\mu), \quad (2.32)$$

where $\mathcal{L}_\ell$ are the Legendre polynomials, of which the monopole ($\ell = 0$), quadrupole ($\ell = 2$), and hexadecapole ($\ell = 4$) are non-zero at linear order. This anisotropy is a powerful cosmological observable, as it allows the growth rate f , or better the combination $f\sigma_8$, to be measured independently of the bias parameter b_1 through the ratio of multipoles.

Beyond linear order, the RSD mapping generates additional mode-coupling contributions. In the SPT framework, these are systematically incorporated by replacing the real-space bias kernels $F_{g,r}^{[n]}$ with their redshift-space counterparts $F_{g,s}^{[n]}$ [13, 69], which depend on both the momenta and the angle μ . In the literature, such as in [13], $F_{g,r}^{[n]}$ and $F_{g,s}^{[n]}$ are usually labeled K_n and Z_n , respectively. In this thesis, however, we prefer a more explicit, although heavier, notation, especially in view of the two-fluid expansion derived in Chapt. 4. For instance, at linear order $F_{g,s}^{[1]}(\mathbf{k}) = b_1 + f\mu^2$, consistently with the Kaiser formula. The one-loop redshift-space galaxy power spectrum then takes the same schematic form as Eq. (2.27), but with all kernels replaced by their redshift-space equivalents and with additional counterterms accounting for the UV sensitivity introduced by the velocity field, the leading ones being proportional to $k^2\mu^2 P_{11}(k)$ and $k^2\mu^4 P_{11}(k)$.

Stochastic contributions

Beyond the deterministic bias expansion, the relationship between galaxies and matter also has an irreducible stochastic component. This arises because galaxy formation is a discrete, non-linear process sensitive to the small-scale environment, which is not captured by any finite number of long-wavelength operators. In the EFT framework, this stochasticity is modelled as an additional noise field $\epsilon(\mathbf{x})$, uncorrelated with the long-wavelength density field by construction,

$$\delta_g(\mathbf{x}) = \sum_{\mathcal{O}} b_{\mathcal{O}} \mathcal{O}(\mathbf{x}) + \epsilon(\mathbf{x}). \quad (2.33)$$

The power spectrum of ϵ contributes as an additive, scale-independent (white noise) term to the galaxy power spectrum,

$$P_{gg}(k) \supset P_{\text{stoch}} = \frac{1 + \alpha_P}{\bar{n}}, \quad (2.34)$$

where $\bar{n}$ is the mean number density of tracers and α_P parametrizes deviations from the Poisson expectation of galaxies being discrete, independent points. In general, small-separation effects such as

halo exclusion or the finite size of the host haloes correlate the tracers over a characteristic scale, so the noise cannot remain exactly white beyond the very largest scales; expanding this scale dependence in powers of k , the leading correction to the constant Poisson term is quadratic. Small-separation effects such as exclusion between haloes can thus lead to either super- or sub-Poissonian noise [71], and at higher order in derivatives additional scale-dependent stochastic contributions appear, the leading one being proportional to $k^2/\bar{n}$, parametrised by a free coefficient ϵ_{k^2} and leaking also to higher multipoles [62]. These stochastic terms must be included in any accurate model of the galaxy power spectrum, and their free parameters are marginalised over in the cosmological analyses.

The full one-loop model in redshift space,

$$P_{gg}(k) = b_1^2 P_{mm}^{1\text{-loop}}(k) + P_{\text{bias}}^{1\text{-loop}}(k) + P_{\text{count}}^{1\text{-loop}}(k) + P_{\text{stoch}}(k), \quad (2.35)$$

with its complete set of bias, EFT, and stochastic parameters constitutes the theoretical backbone of the analyses presented in this thesis.

2.2. Data: estimating correlation functions from simulated and real data catalogs

In this section we, very briefly, focus on the left hand side of Eq. (2.4), that is how the ensemble average is performed in practice to measure the galaxy power spectrum. An *estimator* is a practical way of computing the ensemble average $\langle \delta_g(\mathbf{k}) \delta_g^*(\mathbf{k}') \rangle$ with a finite operation on a single realisation of the observed density field, such that the result converges to the true power spectrum in the limit of a large survey volume.

In a cubic box A cubic box, such as the product of N -body simulations, is conceptually the cleanest environment to give basic details on how power spectrum estimators work. The first step is to obtain the Fourier space overdensity $\delta(\mathbf{k})$ from the point, being this galaxies or particles of simulations, catalog. In principle, this can be done via a straightforward application of the definition Eq. (2.3), summing over pairs

$$\delta_g(\mathbf{k}) = \frac{V}{N} \sum_{i=1}^N e^{-i\mathbf{k} \cdot \mathbf{x}_i}, \quad (2.36)$$

which is indeed dubbed as “direct summation” method for a grid of N_k values in wavenumbers. This, however, falls short operationally, as it scales as $\mathcal{O}(NN_k)$ and becomes computationally prohibitive for modern surveys with $N \sim 10^6$ – 10^8 objects. This issue is overcome by first interpolating the points onto a Cartesian grid of N_{grid}^3 cells using a mass-assignment scheme (MAS) such as Nearest Grid Point (NGP), Cloud-In-Cell (CIC), or the higher-order Piecewise Cubic Spline (PCS). With that, one can then exploit Fast Fourier Transform algorithms to get the Fourier space quantity evaluated on the same grid, reducing the computational cost to $\mathcal{O}(N_{\text{grid}}^3 \log N_{\text{grid}})$. The introduction of a grid, however, entails two complications. The first is the appearance of a Nyquist frequency $k_{\text{Nyq}} = \pi N_{\text{grid}}/L$, with L the box size, setting the maximum frequency that can be accurately sampled on the grid, while the fundamental frequency $k_f = 2\pi/L$ sets the minimum. The mass-assignment scheme also introduces a smoothing kernel $W_{\text{MAS}}(\mathbf{k})$ in Fourier space that suppresses power near the Nyquist frequency; since this kernel is known analytically for any given MAS, the measured $\tilde{\delta}_g(\mathbf{k})$ is corrected by dividing out $W_{\text{MAS}}(\mathbf{k})$ before forming the power spectrum estimate.

2. Galaxy Clustering as a Cosmological Probe

The second complication is aliasing. The discretisation of the density field onto a grid causes power from modes with wavenumbers $|\mathbf{k} + 2nk_{\text{Nyq}}\hat{\mathbf{n}}|$, for integers n , to fold back and contribute to the modes $\mathbf{k}$, that is

$$\delta_g^{(\text{grid})}(\mathbf{k}) = \sum_{\mathbf{n}} \delta_g^{(\text{dir. summ.})}(\mathbf{k} + 2nk_{\text{Nyq}}\hat{\mathbf{n}}), \quad (2.37)$$

A particularly effective and computationally feasible remedy to this is the interlacing technique [72]: two grids are assigned with a half-cell offset $L/2N_{\text{grid}}$, their Fourier transforms are combined as

$$\delta_g^{(\text{interlac.})}(\mathbf{k}) = \frac{1}{2} \left[\delta_g^{(1)}(\mathbf{k}) + \delta_g^{(2)}(\mathbf{k}) \right] = \sum_{\mathbf{n} \text{ even}} \delta_g^{(\text{dir. summ.})}(\mathbf{k} + 2nk_{\text{Nyq}}\hat{\mathbf{n}}) \quad (2.38)$$

This combination thus effectively eliminates the odd $\mathbf{n}$ aliasing contributions, including the leading $n = 1$ one, at the cost of a single additional FFT, and is now standard practice in state-of-the-art power spectrum pipelines [72].

Once one has the corrected Fourier-space density evaluated on a 3D grid, the redshift-space power spectrum multipoles are estimated by weighting each mode with the appropriate Legendre polynomial $\mathcal{L}_\ell(\mu)$ before averaging over shells in k ,

$$\hat{P}_\ell(k) = \frac{(2\ell + 1)V}{N_k} \sum_{\mathbf{k}_n \in k \pm \Delta k/2} |\delta_g(\mathbf{k}_n)|^2 \mathcal{L}_\ell(\hat{\mathbf{k}}_n \cdot \hat{\mathbf{z}}), \quad (2.39)$$

where N_k is the number of modes in the shell, $\hat{\mathbf{z}}$ is the (fixed) line-of-sight direction, and $\mu = \hat{\mathbf{k}}_n \cdot \hat{\mathbf{z}}$. The $\ell = 0$ monopole reduces to the standard angle-averaged power spectrum, while the $\ell = 2$ quadrupole and $\ell = 4$ hexadecapole capture the anisotropic signal induced by redshift-space distortions, as discussed in the previous section. In this idealised case the estimator is unbiased, and its variance is dominated by the finite number of independent modes, a fundamental limit known as *cosmic variance*. Note that, with the normalisation adopted above, the estimator receives the self-pair contribution $1/\bar{n}$ in the monopole. This Poisson shot noise is either subtracted explicitly or, as is standard in full-shape analyses, absorbed into the free stochastic amplitude α_P introduced in Sect. 2.1.2 and marginalised over.

Building an estimator for higher-order correlations, most importantly the bispectrum, carries further challenges that are not covered here, and we refer the reader to [73] for a full treatment. In the context of the doctoral work PBI4 [74], a code that implements the described steps, including interlacing building upon powerI4¹, for power spectrum and bispectrum multipoles has been tested, made public², and employed for measuring the simulation data in the following chapters.

Real observations Real galaxy surveys depart from a periodic box in several important ways, each of which requires dedicated treatment. Here below we just outline the main further effects to take into account without entering into technical details.

First, in a real survey the mean galaxy density $\bar{n}(\mathbf{x})$ varies across the volume, so modes are sampled with unequal noise. [75] showed that assigning each galaxy a weight $w_{\text{FKP}}(\mathbf{x}) = [1 + \bar{n}(\mathbf{x}) P_0]^{-1}$, for a fiducial amplitude P_0 , minimises the variance of $\hat{P}_g(k)$ in the limit of a Gaussian field.

¹<https://github.com/sefusatti/PowerI4>

²<https://github.com/matteobiagetti/pbi4>

Redshift-space distortions (RSD) break isotropy along the line of sight $\hat{\mathbf{n}}$, making the power spectrum anisotropic: $P_g(\mathbf{k}) = \sum_{\ell} P_{\ell}(k) \mathcal{L}_{\ell}(\hat{\mathbf{k}} \cdot \hat{\mathbf{n}})$. For a wide survey the line of sight varies across the field, so the assumption of a fixed $\hat{\mathbf{n}}$ (plane-parallel approximation) breaks down. The [76] estimator accounts for this by expanding the pair-wise line of sight, yielding multipole estimates that remain accurate even at large opening angles.

A real survey covers only a fraction of the sky and applies a spatially varying selection function, encoded in the window $W(\mathbf{x})$. The observed field is $\delta_g^{\text{obs}}(\mathbf{x}) = W(\mathbf{x}) \delta_g(\mathbf{x})$, so its power spectrum is the convolution $\hat{P}_g^{\text{obs}} = P_g * |\tilde{W}|^2$. This couples different k -modes and introduces a characteristic suppression and ringing at scales comparable to the survey size. The window must be modelled and the theoretical predictions convolved, i.e. a forward-modelling of the window effects, before comparing with $\hat{P}_g$ [77–79].

3. Massive neutrinos and Galaxy Clustering

In this chapter, we focus on the effects of massive neutrinos on linear cosmological perturbations and galaxy clustering, whose treatment in Chapter 2 has been presented in a massless scenario. After reviewing the main aspects of linear theory in Sect. 3.1, which will also be at the basis of the theoretical study of Chapter 4, we present in Sect. 3.2-3.6 an original study on redshift-space distortions for biased tracers, clarifying some subtleties arising when massive neutrinos are taken into account and that have so far been overlooked.

3.1. Review of Massive Neutrinos effect on the large-scale structure

Massive neutrinos are the only particles in the Standard Model whose mass remains unmeasured. Neutrino oscillation experiments have established that at least two of the three mass eigenstates are non-zero, measuring the two squared-mass splittings $\Delta m_{21}^2 \approx 7.5 \times 10^{-5} \text{ eV}^2$ and $|\Delta m_{3\ell}^2| \approx 2.5 \times 10^{-3} \text{ eV}^2$ with $\ell = 1, 2$ depending on the ordering [80]. These splittings set an absolute lower bound on the sum of neutrino masses of $\sum m_\nu \gtrsim 0.059 \text{ eV}$ in the normal ordering and $\sum m_\nu \gtrsim 0.10 \text{ eV}$ in the inverted ordering, while leaving the overall mass scale and the ordering itself undetermined. Current laboratory bounds from tritium beta decay place the mass of a particular flavour combination below 0.45 eV [28].

Despite their small fraction with respect to total matter $f_\nu \equiv \Omega_\nu/\Omega_m \lesssim 0.01$, massive neutrinos leave a distinctive and potentially detectable imprint on the large-scale structure of the Universe, making cosmology one of the most sensitive probes of the neutrino mass scale. This sensitivity arises mainly through two, intertwined, effects. The first is a background effect. At early times neutrinos are relativistic and contribute to the radiation density, but as the Universe expands they become non-relativistic once the temperature drops below their mass, transitioning to behave as matter. This shifts the redshift of matter–radiation equality and modifies the expansion history $H(z)$, so that for a fixed set of cosmological parameters a larger $\sum m_\nu$ alters the background evolution and the associated geometric distances probed by, e.g., baryon acoustic oscillations. The second is the suppression of perturbations. Because of their large thermal velocities, neutrinos free-stream out of gravitational potential wells on scales below the free-streaming length, and therefore do not cluster there. This has two consequences: neutrino perturbations do not contribute to the growth of structure on small scales, and their free-streaming also slows the growth of the cold dark matter and baryon perturbations through the reduced gravitational potential. The net result is a characteristic, scale-dependent suppression of the matter power spectrum on scales below the free-streaming scale, with a relative amplitude $\Delta P/P \approx -8f_\nu$ in the linear regime. It is precisely this scale-dependent imprint that makes galaxy clustering one of the most powerful probes of the neutrino mass scale.

Cosmology indeed turns out to be incredibly competitive with laboratory experiments concerning upper-limits on $\sum m_\nu$, where the state-of-the-art constraints push this down even to $\lesssim 70 \text{ meV}$ with the Planck+ACT+DESI (BAO) combination [5, 6, 24, 31]. This figure actually sits in mild tension

3. Massive neutrinos and Galaxy Clustering

with the oscillation lower bounds: the tightest cosmological constraints now approach, and in some data combinations formally exclude at the $\sim 2\sigma$ level, the minimal value allowed for the inverted ordering, and are beginning to press against the normal-ordering floor as well. It should be noted that this observables combination mainly constrains M_ν through the background effect, i.e. $\Omega_m - \Omega_c$, which is not a signature as distinctive as the scale-dependent suppression. Some analyses even prefer negative values of $\sum m_\nu$ when the parameter is allowed to float unphysically [32]. This hints either at a statistical fluctuation, unaccounted-for systematics, or a possible breakdown of the Λ CDM assumptions underlying these bounds. Whether this reflects new physics or the residual tensions degenerate with neutrino mass effects remains an open question.

In what follows, we briefly review the theoretical treatment of massive neutrinos effect on linear cosmological perturbations and show quantitatively the suppression it entails on the power spectrum. Then, we turn to the effects on tracers, with an in-depth study on the redshift-space power spectrum in presence of massive neutrinos, highlighting some subtleties that have so far been overlooked and could potentially bias inference in case of a future non-null detection of $\sum m_\nu$.

Boltzmann hierarchy and fluid approximation

Unlike cold dark matter, neutrinos decouple from the thermal bath while still relativistic, at temperatures $T \sim \text{MeV}$, and thereafter free-stream with a thermal velocity dispersion that remains cosmologically relevant long after they become non-relativistic. That is, their background phase space distribution remains a relativistic Fermi-Dirac distribution at all times, i.e.

$$f_0(q, \tau) = \frac{1}{1 + e^{\frac{q}{aT_\nu(a)}}} \quad (3.1)$$

where q is the comoving momentum and T_ν is the neutrino temperature parameter. Actually, it is worth noting that $T_{\nu,0}$ is not strictly a thermodynamic temperature, as neutrinos are completely out of thermal equilibrium following their decoupling; rather, it functions as a formal parameter characterizing the shape of the frozen phase-space distribution, that decreases as $T_\nu(a) = (a_0/a)T_0$ even after neutrinos are fully non-relativistic.

The evolution of the full neutrino phase-space distribution $f(\mathbf{x}, \mathbf{q}, \tau)$ is governed by the collisionless Boltzmann, or Vlasov, equation [26]. For what concerns linear cosmology, it is convenient to consider the Legendre multipoles of the Fourier space, hence introducing the wavenumber $\mathbf{k}$, phase-space distribution perturbations

$$f(\mathbf{k}, \mathbf{q}, \tau) = f_0(q, \tau) \left[1 + \sum_{l=0}^{\infty} (-i)^l (2l+1) \Psi_l(k, q, \tau) P_l\left(\frac{\mathbf{k} \cdot \mathbf{q}}{kq}\right) \right] \quad (3.2)$$

This turns the Boltzmann equation in a set of coupled equations, the so-called Boltzmann hierarchy, as [81]

$$\dot{\Psi}_0 = -\frac{qk}{\epsilon} \Psi_1 - \dot{\phi} \frac{d \log f_0}{d \log q}, \quad (3.3)$$

$$\dot{\Psi}_1 = \frac{qk}{3\epsilon} (\Psi_0 - 2\Psi_2) - \frac{\epsilon k}{3q} \psi \frac{d \log f_0}{d \log q}, \quad (3.4)$$

$$\dot{\Psi}_\ell = \frac{qk}{(2\ell+1)\epsilon} [\ell \Psi_{\ell-1} - (\ell+1) \Psi_{\ell+1}], \quad \ell \geq 2, \quad (3.5)$$

where dots denote derivatives with respect to conformal time, ϕ and ψ are the conformal Newtonian potentials [81], and $\epsilon(q) = \sqrt{q^2 + m_\nu^2 a^2}$. In principle this hierarchy must be truncated at some maximum multipole ℓ_{max} to be solved numerically, and this is the approach implemented in Boltzmann codes such as CLASS [60] and CAMB [61].

The first two moments of the phase-space distributions, the density contrast δ_ν and the velocity divergence θ_ν , do thus not form a closed system like in Eq. (2.6) but are coupled to higher multipoles via the shear stress $\sigma(\mathbf{k}, \tau)$. However, it has been shown that at late times the effect of the shear can be approximated as an additional contribution to the pressure term, effectively rescaling it by a factor 9/5 [82]. This is a particularly useful approximation, as in this way the $\ell = 1, 2$ equations of the Boltzmann hierarchy (3.3) effectively decouple from the higher multipoles. We thus get a system very similar to Eqs. 2.6, where all the free-streaming physics is captured in the Euler equation by an effective sound speed

$$\theta'_\nu + \mathcal{H}\theta_\nu - k^2 c_s^2 \delta_\nu + \frac{3}{2} \mathcal{H}^2 \Omega_m \delta_m = 0 \quad (3.6)$$

with $c_s^2 = \frac{9}{5} c_g^2$, where $c_g^2 \equiv \bar{\mathcal{P}}'/\bar{\rho}' \simeq \frac{5}{9} \sigma_\nu^2$ is the squared adiabatic sound speed, so that $c_s^2 \simeq \sigma_\nu^2 = \frac{15\zeta(5)}{\zeta(3)} T_{\nu,0}^2 / (m_\nu^2 a^2)$.¹ Eq. (3.6) is describing the evolution of a perfect fluid, which will further be at the basis of the theoretical study presented in Chapt. 4 for generic behaviours of the sound speed. Notably, Eq. (3.6) entails a characteristic wavenumber, which is interpreted here as the neutrino free streaming scale, that reads

$$k_s(a) \approx a^{1/2} 3.7 \times 10^{-2} \frac{h}{\text{Mpc}} \left(\frac{m_\nu}{0.1 \text{ eV}} \right) \left(\frac{\Omega_m^0}{0.3} \right)^{1/2}. \quad (3.7)$$

The physical interpretation is transparent: neutrinos with thermal velocities $v_\nu \sim \sigma_\nu$ free-stream out of potential wells on scales smaller than the free-streaming length $\lambda_{\text{fs}} \sim v_\nu/H$, preventing the formation of structure on those scales.

Effect on the matter power spectrum and growth rate

The cosmological impact of massive neutrinos on large-scale structure is best understood by separating the total matter field into a cold component δ_c (CDM and baryons) and the neutrino component δ_ν , with the total matter overdensity given by

$$\delta_m = (1 - f_\nu) \delta_c + f_\nu \delta_\nu, \quad (3.8)$$

where $f_\nu = \Omega_\nu/\Omega_m$ is the neutrino fraction. On scales much larger than the free-streaming scale, $k \ll k_s$, neutrinos cluster like cold matter, and the growth of perturbations proceeds as in the massless case. On scales $k \gg k_s$, however, neutrino perturbations are strongly suppressed, $\delta_\nu \ll \delta_c$, and the gravitational potential sourcing the growth of δ_c is reduced by a factor $(1 - f_\nu)$. This leads to a suppression of the growth rate of cold perturbations, which on small scales evolves as [26]

$$\delta_c \propto D(a)^{1-3f_\nu/5}, \quad (3.9)$$

so that the growth factor becomes scale-dependent, $D \rightarrow D(k, a)$, and similarly the growth rate

$$f(k) \equiv \frac{d \log D(k, a)}{d \log a} \simeq \begin{cases} f_{\Lambda\text{CDM}} & k \ll k_s, \\ f_{\Lambda\text{CDM}} (1 - \frac{3}{5} f_\nu) & k \gg k_s, \end{cases} \quad (3.10)$$

¹Alternative definitions of the velocity dispersion, differing by $\mathcal{O}(1)$ factors, are common in the literature. In particular, $\sigma_\nu^2 = \frac{3\zeta(3)}{2 \log 2} T_{\nu,0}^2 / (m_\nu^2 a^2)$ in Refs. [45, 83, 84], which would lead to $k_s(a)$ being larger by a factor ≈ 2.2 compared to our expression in Eq. (3.7).

3. Massive neutrinos and Galaxy Clustering

transitions from its standard Λ CDM value on large scales to a suppressed value on small scales. This scale dependence of the growth rate is a distinctive signature of massive neutrinos in redshift-space clustering, as it imprints a scale-dependent anisotropy in the observed power spectrum multipoles, and its detection and characterisation is one of the main subjects of this thesis.

The cumulative effect on the linear matter power spectrum is a step-like suppression on scales $k \gtrsim k_s$. Relative to a massless neutrino cosmology with the same values of the other parameters, the suppression of the total matter power spectrum is approximately [26, 85]

$$\frac{\Delta P_{mm}}{P_{mm}} \approx -8f_\nu, \quad (3.11)$$

valid on small scales and at low redshift. This result has a simple physical interpretation: a fraction f_ν of the matter does not contribute to the gravitational clustering on small scales, reducing the amplitude of the power spectrum by a factor $(1 - f_\nu)^2 \simeq 1 - 2f_\nu$ at the level of the density field, and the additional suppression factor of $\sim 6f_\nu$ comes from the reduced growth of cold perturbations due to the lower gravitational source. The location of the transition encodes information about the neutrino mass through k_s , while the depth of the step is proportional to $f_\nu \propto \sum m_\nu$. This makes the full-shape power spectrum, measured over a wide range of scales spanning both $k \ll k_s$ and $k \gg k_s$, a powerful tool to constrain $\sum m_\nu$ from galaxy clustering data, as exploited in the analyses presented in the following chapters.

3.2. Redshift-space distortions in massive neutrino cosmologies

As described, the peculiar suppression of small-scale growth of matter perturbations due to massive neutrinos leads to a potentially distinctive feature in the shape of the galaxy power spectrum. To connect theory to observations, with e.g. the goal of inferring $\sum m_\nu$, it is thus essential to extend the theory of bias and redshift space distortions (RSDs) presented for Λ CDM in Sect. 2.1.2 to underlying matter power spectra exhibiting a suppression given by Eq. (3.11), which is not trivial already at linear level.

In massive neutrino cosmologies, a natural ambiguity arises indeed already w.r.t. the definition of the matter density field. One option is to consider the combination of dark matter and baryons in standard Λ CDM models, hereafter denoted as “cold matter”, since baryonic effects can be neglected at large scales. In addition, since massive neutrinos behave as dark matter at low redshift, an alternative definition of matter could include dark matter, baryons, and neutrinos. In what follows, we will denote this “total matter”.

As far as the galaxy density is concerned, i.e. the first term of Eq. (2.30), numerical simulations have shown that a proper definition of galaxy bias expansion in terms of *constant* bias parameters must be defined w.r.t. the cold matter density [52, 86]. The alternative option of a galaxy bias defined w.r.t. the total matter perturbations would lead to bias parameters affected by a spurious scale dependence.

The second term in the Kaiser formula is also affected by the same ambiguity. The galaxy velocity field is usually presumed to be unbiased w.r.t. the matter field velocity, but for massive neutrinos cosmologies the cold matter and total matter velocities are not the same. Although it is often assumed, sometimes implicitly following the results of [52, 86], that the galaxy velocity coincides with the cold-matter velocity, this has not found a proper, quantitative confirmation from numerical simulations.

Some early work looked at the issue, but with a relatively limited cumulative simulation volume [53, 87]. In particular, in [53] a few quantities such as the growth rate f , scale-dependent when neutrinos are massive, have been directly estimated from simulations and compared to linear theory predictions under different assumptions. The lack of statistical significance could not lead to any clear result.

Perhaps the most directly relevant work in the past literature is [88], taking advantage of a much larger set of simulations, then denoted as HADES, now part of the even larger QUIJOTE suite [88]. In [88] the bias between the halo and cold-matter *momenta* is estimated and shown to approach unity at large scales. However this result, while not directly dealing with the velocity field, is also not compared to the same bias defined w.r.t. the total matter momentum. As we will show, such comparison would have led, taken alone, to the same qualitative conclusion.

In this chapter we revisit some of these prior results and extend them in order to obtain a firm, quantitative assessment on the galaxy velocity field in massive neutrino cosmologies. In the first place, we reproduce some of the results of both [53] and [88], in some cases with better statistics or with more complete comparisons. In addition, we perform full fits of a theoretical model based on Perturbation Theory (PT), for both the cold-matter or total matter velocity options, to a set of measurements of halo auto-power spectra and halo-cold matter cross-power spectra both in real and redshift space. This allows to make a quantitative assessment on which assumption provides the best fit to the simulated data.

This study is presented as follows. In Sect. 3.2 we introduce the theoretical background both at linear and nonlinear level in PT leading to predictions for both auto- and cross-power spectra. In Sect. 3.3 we briefly present the simulations sets employed and the power spectra measurements. In Sect. 3.4 we perform a direct, qualitative comparison of measurements from the halo catalogs and linear theory predictions, while in Sect. 3.5 we present the results for the fit of the nonlinear model. We conclude in Sect. 3.6.

3.2.1. Linear theory

As shown by Eq. (2.30), in the large-scale limit the density contrast of tracers (e.g. galaxies or dark-matter halos) in redshift space is distorted by the effect of peculiar velocities [70].

In the standard Λ CDM picture, with massless neutrinos and baryons approximated as cold dark matter, the velocity bias of tracers with respect to matter perturbations has been shown to be a higher-derivative effect, with the large-scale limit of unity bias dictated by the equivalence principle [69, 89]. In linear theory we can thus set $\theta_h = \theta_m$ and, along with the linear bias relation $\delta_h = b\delta_m$, rewrite Eq. (2.30) as [70]

$$\delta_{h,s}(\mathbf{k}) = (b + f_{\Lambda\text{CDM}}\mu^2) \delta_m(\mathbf{k}). \quad (3.12)$$

In the context of massive neutrino cosmologies, however, matter perturbations are made up of two components with different properties, even at large scales. These are the *cold* component, composed of dark matter and baryons, here assimilated to DM, and the neutrino component. In what follows, we will denote the cold component with the subscript “c” while we will denote total matter perturbations with the subscript “m”.

A first problem is given by the definition of the bias expansion in real space, since already at the linear level this can be defined w.r.t. the cold matter density,

$$\delta_h(\mathbf{k}) \simeq b_c \delta_c(\mathbf{k}), \quad (3.13)$$

3. Massive neutrinos and Galaxy Clustering

or the total matter one

$$\delta_h(\mathbf{k}) \simeq b_m \delta_m(\mathbf{k}) , \quad (3.14)$$

where $\delta_m = (1 - f_\nu)\delta_c + f_\nu\delta_\nu$, with $f_\nu \equiv \Omega_\nu/\Omega_m$ representing the relative neutrino fraction. References [52, 53, 86] have shown that, in numerical simulations where the neutrino component is described in terms of a distinct particle species, the first relation leads to a scale-independent linear bias parameter, as expected in the standard perturbative description of galaxy bias. This is likely due to the relatively small effect of neutrino perturbations on structure formation and to the low value of the neutrino fraction. Still, one can clearly expect, in principle, that a general description of halo bias would include, to some extent, also neutrino perturbations. We will therefore consider the assumption of Eq. (3.13) as a good approximation to simulation results in the linear and mildly non-linear regime. We note that massive neutrinos have been shown [58, 90] to introduce a non-locality in halo formation inducing a further, subtle, scale dependence, nevertheless, this correction is sub-leading at the scales considered here.

In the interpretation of the Kaiser formula in massive neutrino cosmologies, a similar problem arises at the level of the peculiar velocity field θ_h . Again, while a general model could be more complex, one can start by considering the simple case of the halo velocity coinciding with the cold matter one θ_c , so that

$$\delta_{h,s} = b_c \delta_c - \mu^2 \frac{\theta_c}{\mathcal{H}} . \quad (3.15)$$

Introducing the growth rate of the cold component perturbations, scale-dependent in massive neutrino cosmologies,

$$f_c(k) \equiv \frac{d \log D_c(a, k)}{d \log a} , \quad (3.16)$$

from the continuity equation for δ_c one gets

$$\delta_{h,s}(\mathbf{k}) = \left[b_c + f_c(k) \mu^2 \right] \delta_c(\mathbf{k}) . \quad (3.17)$$

The opposite, limiting case of $\theta_h = \theta_m$ leads to

$$\delta_{h,s}(\mathbf{k}) = b_c \delta_c - \mu^2 \frac{\theta_m}{\mathcal{H}} \quad (3.18)$$

$$= b_c \delta_c(\mathbf{k}) + f_m(k) \mu^2 \delta_m(\mathbf{k}) , \quad (3.19)$$

where now

$$f_m(k) \equiv \frac{d \log D_m(a, k)}{d \log a} , \quad (3.20)$$

is the growth rate of the total matter perturbations. We keep the standard notation for the growth rates f_c and f_m , not to be confused with the *neutrino fraction* f_ν , also in the standard use.

We can summarize the two prescriptions for the redshift-space halo overdensity $\delta_{h,s}$ in terms of

$$\delta_{h,s}(\mathbf{k}) = \left[1 + \beta_c(k) \mu^2 \right] \delta_h(\mathbf{k}) , \quad (3.21)$$

with

$$\beta_c(k) \equiv \frac{f_c(k)}{b_c} , \quad (3.22)$$

or, alternatively, in terms of

$$\delta_{h,s}(\mathbf{k}) = \left[1 + \beta_m(k) \mu^2 \right] \delta_h(\mathbf{k}) , \quad (3.23)$$

with

$$\beta_m(k) \equiv \frac{f_m(k)}{b_m(k)} = \frac{f_m(k)}{b_c} \frac{D_m(k)}{D_c(k)}, \quad (3.24)$$

where we assume a constant linear bias defined w.r.t. cold matter perturbations b_c , while the one defined w.r.t. total matter includes the spurious scale dependence described as $b_m(k) = b_c [D_c(k)/D_m(k)]$.

The relations above allow to write the redshift-space power spectrum, $P_{hh,s}(k, \mu)$, in terms of the real-space one, $P_{hh,r}(k)$ generically as

$$P_{hh,s}(k, \mu) = [1 + \beta(k)\mu^2]^2 P_{hh,r}(k), \quad (3.25)$$

where the parameter $\beta(k)$ is given in the two limit cases by Eq. (3.22) or (3.24). As usual, the monopole, quadrupole, and hexadecapole will read

$$P_{hh,0}(k) = \left[1 + \frac{2}{3}\beta(k) + \frac{1}{5}\beta^2(k)\right] P_{hh,r}(k), \quad (3.26)$$

$$P_{hh,2}(k) = \left[\frac{4}{3}\beta(k) + \frac{4}{7}\beta^2(k)\right] P_{hh,r}(k), \quad (3.27)$$

$$P_{hh,4}(k) = \frac{8}{35}\beta^2(k)P_{hh,r}(k). \quad (3.28)$$

In the case of the halo-cold dark matter cross-power spectrum the relation between redshift space and real space is

$$P_{hc,s}(k, \mu) = [1 + \beta_c(k)\mu^2] P_{hc,r}(k), \quad (3.29)$$

so that monopole and quadrupole read

$$P_{hc,0}(k) = \left[1 + \frac{1}{3}\beta_c(k)\right] P_{hc,r}(k), \quad (3.30)$$

$$P_{hc,2}(k) = \frac{2}{3}\beta_c(k)P_{hc,r}(k) \quad (3.31)$$

with a vanishing hexadecapole. Identical expressions can be written in terms of $P_{hm,s}(k)$, $P_{hm}(k)$ and $\beta_m(k)$, for both the anisotropic power spectrum and for the redshift-space multipoles.

3.2.2. Nonlinear perturbation theory in redshift space

In the perturbative approach presented in Sect. 2.1.2, the halo auto-power spectrum can be written at 1-loop as

$$\begin{aligned} P_{hh,s}^{\text{SPT}}(\mathbf{k}) &= F_{h,s}^{[1]2}(\mathbf{k})P_c^L(k) + \\ &+ 2 \int d^3\mathbf{q} \left[F_{h,s}^{[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}) \right]^2 P_c^L(q)P_c^L(k - q) + \\ &+ 6F_{h,s}^{[1]}(\mathbf{k})P_c^L(k) \int d^3\mathbf{q} F_{h,s}^{[3]}(\mathbf{k}, \mathbf{k}, -\mathbf{q})P_c^L(q), \end{aligned} \quad (3.32)$$

where the $F_{h,s}^{[n]}$ are the “biased” redshift-space kernels, whose expressions can be found, labeled as Z_n , in [13]. In our analysis, we will compare our theoretical predictions with measurements of the cross power spectrum between halos and cold-matter densities as well. In this case, the one-loop expression has two main contributions

$$P_{hc,s}(\mathbf{k}) = P_{hc,s}^{\text{SPT}}(\mathbf{k}) + P_{hc,s}^{\text{ctr}}(\mathbf{k}), \quad (3.33)$$

3. Massive neutrinos and Galaxy Clustering

with the SPT term given by (see, for instance [91])

$$\begin{aligned}
P_{hc,s}^{\text{SPT}}(\mathbf{k}) &= F_{h,s}^{[1]}(\mathbf{k}) P_c^L(k) + 2 \int d^3\mathbf{q} F_{h,s}^{[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}) \\
&\times F_2(\mathbf{q}, \mathbf{k} - \mathbf{q}) P_c^L(q) P_c^L(k - q) + 3 P_c^L(k) \\
&\times \int d^3\mathbf{q} \left[F_{h,s}^{[3]}(\mathbf{k}, \mathbf{k}, -\mathbf{q}) + F_3(\mathbf{k}, \mathbf{k}, -\mathbf{q}) \right] P_c^L(q), \tag{3.34}
\end{aligned}$$

where the F_n are the *matter* kernels, given at second order by Eq. (2.16), and with the counterterms contribution

$$P_{hc,s}^{\text{ctr}}(k, \mu) = -c_0^\times k^2 P_c^L(k) - c_2^\times k^2 \mu^2 P_c^L(k), \tag{3.35}$$

where we assume the parameters $c_0^\times$ and $c_2^\times$ in general distinct from those describing the halo auto-power spectrum.

From the expressions above for both the auto- and cross-power spectra, one can readily obtain the real-space predictions replacing each $F_{h,s}^{[n]}$ occurrence with the real-space equivalent $F_{h,r}^{[n]}$ kernels, with no contributions from terms depending on the angle with the line-of-sight (see e.g. [92, 93]) and removing all anisotropic counterterms.

In addition to the linear effects discussed above, the presence of massive neutrinos modifies the kernels already at the matter density level [94, 95]. However, due to the very low abundance and the relatively large value of the free-streaming scale, the approximation [94] that only accounts for standard one-loop corrections to the cold-matter components, computed from its linear power spectrum and in terms of the usual Einstein-deSitter kernels, performs very well compared to numerical simulations [53, 96]. For biased tracers we further assume nonlinear bias to be defined as well w.r.t. cold-matter perturbations, as for linear theory. It should however be stressed again that these approximations work for massive neutrinos with very low m_ν , as constrained by the data, and thus entail a very large free streaming scale as in Eq. (3.7); we will go beyond this, considering generic hierarchy of scales, in Chapter 4.

Concerning instead the redshift-space contributions to the nonlinear kernels $F_{h,s}^{[n]}$, we replace every occurrence of the growth rate f as

$$f \rightarrow f_c(k),$$

or

$$f \rightarrow f_m(k) \frac{D_m(k)}{D_c(k)},$$

according to the modeling assumption on the halo velocity field θ_m , with the first option corresponding to the scheme of [97, 98]. However, for every instance of f appearing in a loop integrand, we consider its constant, small-scale, asymptotic value, making it feasible to compute the integrals with FFTLog methods [99].

Finally, the counterterms in the auto- and cross-power spectrum models may also explicitly involve the scale-dependent growth rate $f(k)$ [97]. The choice on the halo velocity θ_h thus reflects, in principle, here as well. However, by directly testing alternative options we found that this has negligible impact on our results, opting then for the agnostic choice presented in, e.g., Eqs. (2.27) and (3.35).

We will make use of the PBJ code [62, 92] that we extend to cover cross-power spectra predictions. We therefore adopt its specific implementation, in a standard Λ CDM cosmology, of the nonlinear galaxy bias expansion (see, e.g., [69]) and Infrared Resummation [100], and we refer the reader to [62] for any further detail.

Type	N_{realiz}	Ω_c	A_s	$\sigma_{8,m}$	$\sigma_{8,c}$	M_ν [eV]
QUIJOTE set						
ΛCDM	500	0.2685	2.13×10^{-9}	0.834	0.834	0
M_ν	500	0.2590	2.74×10^{-9}	0.834	0.855	0.4
DEMNUiCoV set						
ΛCDM	50	0.270	2.1265×10^{-9}	0.8301	0.8301	0
M_ν	50	0.266	2.1265×10^{-9}	0.7926	0.8013	0.16

Table 3.1.: Summary of the cosmological parameters of the simulations employed. Values for the QUIJOTE are detailed in [101], while those for the DEMNUiCoV in [103, 104]. The QUIJOTE share the values $h = 0.6711$, $\Omega_b = 0.049$ and $n_s = 0.9624$, while the DEMNUiCoV have $h = 0.67$, $\Omega_b = 0.05$ and $n_s = 0.96$.

3.3. N-body simulations and measurements

For the baseline analysis, we will employ the QUIJOTE simulations [101], specifically the subsets with total neutrino mass $M_\nu = 0.4$ eV and the corresponding ΛCDM runs with same initial seeds. These sets of 500 realizations each have been run with 512^3 cold dark matter particles plus an additional 512^3 for the neutrino component with initial conditions set up according to the linear rescaling proposed by [102], in a box with a side of $L = 1 h^{-1}$ Gpc, corresponding to a fundamental frequency $k_f = 2\pi/L \simeq 8 \times 10^{-3} h \text{ Mpc}^{-1}$. The parameter describing the amplitude of the linear scalar perturbations, A_s , is chosen here in order to obtain the same value of $\sigma_{8,m}$ – the r.m.s. of matter perturbation in spheres of $8 h^{-1}$ Mpc of radius – for both massive and massless cosmologies (the cold matter equivalent, $\sigma_{8,c}$ is therefore different in the two cases). Halos have been identified using the Friends-of-Friends (FoF) algorithm with linking length $b = 0.2$ resolving halos of masses $M_h \geq 10^{13} h^{-1} M_\odot$. This choice coincides with the mass cut considered in [53]: when we mention halo bias, in what follows, it is understood that we refer to the effective bias value corresponding to an integral of the proper halo bias $b(M)$, weighted by the mass function, over the relevant mass range.

In addition, we will make use of the DEMNUiCov [103, 104], also with a particle-based neutrino component and sharing the same initial conditions set-up of [102]. They are a suite of 50 realizations characterized by the same box side size, $L = 1 h^{-1}$ Gpc, as the QUIJOTE suite, but eight times larger mass resolution, having a particle number of $N = 2 \times 1024^3$ (where the factor of 2 stands for same number of CDM and ν particles). The DEMNUiCov subset is available for a lower value of the total neutrino mass $M_\nu = 0.16$ eV in addition to the massless neutrino case. For the DEMNUiCov simulations the A_s parameter is kept to the same value for both cosmologies, leading to different values of $\sigma_{8,c}$ and $\sigma_{8,m}$. The halo catalogs used in this study are produced with the same FoF algorithm linking only the CDM particles, where for a consistent comparison with the QUIJOTE we apply the same mass cut of $10^{13} h^{-1} M_\odot$.

The values of the relevant cosmological parameters for the different sets of simulations are reported in Table 3.1.

We will focus on $z = 0.5$ and $z = 1$, as relevant for recent and ongoing surveys. For each snapshot we construct the density fields of the cold matter δ_c , of the total matter δ_m and of the halo number density in real space δ_h and redshift-space $\delta_{h,s}$. This is performed using a fourth-order density interpolation and

interlacing scheme as described in [72] and implemented in the PBI4 code [74], also briefly discussed in Sect. 2.2. We then estimate, extending the estimator of Eq. (2.39) also to cross-power spectra, the multipoles as

$$\hat{P}_{AB,\ell}(k) = (2\ell + 1) \frac{k_f^3}{N_k} \sum_{\mathbf{q} \in k} \delta_A(\mathbf{q}) \delta_B^*(\mathbf{q}) \mathcal{L}_\ell(\mathbf{q} \cdot \hat{z}), \quad (3.36)$$

where $\hat{z}$ represent the line-of-sight, here in the plane-parallel approximation, the sum is intended over all wavenumbers $k - \Delta k/2 < |\mathbf{q}| < k + \Delta k/2$ with $\Delta k = k_f$ is the bin size, and $N_k \equiv \sum_{\mathbf{q} \in k}$ is the number of modes per bin.

3.4. Qualitative tests of linear RSD theory

In this Section, we revisit some qualitative tests already explored in the literature comparing linear theory predictions to simulations measurements. Such tests are usually limited by the short range of scales where nonlinear effects can safely be ignored in simulation boxes of few h^{-1} Gpc of side, particularly in the case of a large set of realizations and low statistical errors. We will consider in the next Section proper fits of nonlinear models including shot-noise corrections, also relevant at linear scales for the power spectrum of halos.

3.4.1. Halo power spectrum and halo-matter crosspower spectrum

Given measurements of both the real-space power spectrum and the redshift-space monopole, one can test the relation (3.26) comparing the estimated l.h.s. of the expression

$$\frac{\hat{P}_{hh,0}}{\hat{P}_{hh,r}} - 1 \simeq \frac{2}{3}\beta(k) + \frac{1}{5}\beta^2(k), \quad (3.37)$$

with a prediction of r.h.s., which is, however, still taking advantage of an estimate of the linear bias from data, and see whether at least at large scales there is agreement between the two. Such comparison is shown on the left panel of Fig. 3.1, qualitatively reproducing the results of Fig. 13 in [53] and Fig. 8 in [88]. We subtract a Poisson shot-noise contribution from both power spectra. This results in a systematic error possibly contributing to the mismatch between measurements and prediction at large scales where, on the other hand, statistical errors are also quite significant. With respect to the specific goal of distinguishing between the θ_c vs. θ_m scenarios in the Kaiser formula, these results are clearly inconclusive.

A more direct comparison can be obtained from measurements of cross power spectra, not affected by shot-noise, with (3.30) leading to

$$\frac{\hat{P}_{hc,0}}{\hat{P}_{hc,r}} - 1 \simeq \frac{1}{3}\beta(k), \quad (3.38)$$

where the measured l.h.s. is compared to the r.h.s in the right panel of Fig. 3.1. In this case the predictions match the measurements at the largest scales, but it is still not possible to distinguish the two prescriptions for the velocity contribution.

One can take further advantage of cross-power spectrum measurements, in particular considering ratios of a given quantity estimated from a massive neutrino simulation to the Λ CDM simulation. In




Figure 3.1.: Tests of the linear relation between redshift-space monopole and real-space halo power spectra. *Left panel:* comparison of the l.h.s. and r.h.s. of (3.37) based on auto-power spectrum measurements. This reproduces the results of Fig. 8 in [88]. *Right panel:* comparison of the l.h.s. and r.h.s. of (3.31) based on halo-matter cross-power spectrum measurements. In both cases we employ the mean of 200 QUIJOTE realizations at $z = 1$. The tests clearly cannot distinguish the two halo velocity scenarios.

this way we partly remove cosmic variance as the two sets with different cosmology share the same initial seeds. More specifically, we can build an estimator of the linear growth rate based on linear theory. This would be given, under the assumption of cold dark matter velocity (3.17), by

$$\hat{f}_c(k) = 3 \frac{\hat{P}_{hc,r}(k)}{\hat{P}_{cc}(k)} \left[\frac{\hat{P}_{hc,0}(k)}{\hat{P}_{hc,r}(k)} - 1 \right], \quad (3.39)$$

which follows from (3.30), with $\hat{P}_{hc,r}/\hat{P}_{cc}$ providing an estimate of the linear bias b_c , in this case as a function of scale. Conversely, if the halo velocity field coincides with the *total* matter velocity (3.19) the correct combination to estimate f reads

$$\hat{f}_m(k) = 3 \frac{\hat{P}_{hm,r}(k)}{\hat{P}_{mm}(k)} \left[\frac{\hat{P}_{hm,0}(k)}{\hat{P}_{hm,r}(k)} - 1 \right], \quad (3.40)$$

where now the bias estimates $b_m(k) = \hat{P}_{hm,r}(k)/\hat{P}_{mm}(k)$ includes any scale-dependence from the different linear growth factors.

The results shown in Fig. 3.2 provide a more significant test with respect to an early attempt in Fig. 14 from [53], which was based on a single realization. Measurements of the ratio $\hat{f}(k)/\hat{f}_{\Lambda\text{CDM}}$ (*connected data points*) are compared to the linear theory prediction $f(k)/f_{\Lambda\text{CDM}}$ (*dashed curves*) for the cold-matter velocity case (*blue*) and total matter case (*green*). Top panels show the results from 40 realizations of the DEMNUniCoV simulations for $M_\nu = 0.16 \text{ eV}$ at redshift $z = 1$ (*left*) and $z = 0.5$ (*right*) while the bottom panel is obtained from 200 realizations of the QUIJOTE simulations for $M_\nu = 0.4 \text{ eV}$ at $z = 1$. We remind the reader that in addition to the different value for M_ν , the two sets of simulation assume different cosmologies and different assumptions on the amplitude of perturbations in the comparison with the ΛCDM case (see Sect. 3.3 and table 3.1). This could in part explain the different departure from linearity, which, we should stress, is shown here as the ratio of nonlinear quantities.

3. Massive neutrinos and Galaxy Clustering




Figure 3.2.: Growth rate estimated from simulation data using the halo-matter cross-power spectrum, assuming (3.39) in the cold-matter velocity assumption (*in blue*) and assuming (3.40) (*in green*). The measurements (*connected data points*) are compared to the linear theory prediction (*dashed curves*). All the quantities are reported as the ratio with respect to their correspondent Λ CDM values. Top panels show the results from 40 realizations of the DEMNUniCoV simulations for $M_\nu = 0.16$ eV at redshift $z = 1$ (*left*) and $z = 0.5$ (*right*) while the bottom panel is obtained from 200 realizations of the QUIJOTE simulations for $M_\nu = 0.4$ eV at $z = 1$. The assumption that the velocity of halos traces the one of the *cold* matter component clearly provides a good match at large scales, while the alternative assumption is clearly failing.

In all cases, the agreement of the data with the $\theta_h = \theta_c$ ansatz of (3.17) is evident. At the same time, the discrepancy between theory and simulations seems to exclude the description that assumes $\theta_h = \theta_m$ given by (3.19).

3.4.2. Momentum bias

Given the well-known difficulties in measuring the velocity field from discrete tracers (see, e.g., [105] for a recent work and references therein), in [88] the momentum field

$$\mathbf{J} = (1 + \delta)\mathbf{u} \quad (3.41)$$

is considered as a proxy for the properties of the halo velocity field, since it can be easily estimated from a simulation snapshot. The same reference estimates then the momentum bias of halos with respect to the cold matter

$$b_{J_c}(k) \equiv \frac{P_{J_h, J_c}(k)}{P_{J_c, J_c}(k)} \quad (3.42)$$

defined in terms of the halo momentum-cold matter momentum cross power spectrum and the auto power spectrum of the cold matter momentum. This is expected to approximate, in the limit $\delta_c, \delta_h \ll 1$, the velocity bias. The cold-matter momentum bias $b_{J_c}(k)$ is there shown to approach unity in the large-scale limit, qualitatively pointing to no velocity bias between halo and cold matter velocity fields.

In the same way, one can compare with the momentum of the total matter

$$J_m = (1 - f_\nu)(1 + \delta_c)u_c + f_\nu(1 + \delta_\nu)u_\nu \quad (3.43)$$

and estimate a momentum bias w.r.t. total matter defined as

$$b_{J_m}(k) \equiv \frac{P_{J_h, J_m}(k)}{P_{J_m, J_m}(k)}. \quad (3.44)$$

In Fig. 3.3 we show estimates from both definitions, with the results for b_{J_c} reproducing those in Fig. 7 of [88]. We notice that both quantities qualitatively approach unity at large scale, with b_{J_c} closer to the limiting value by only a small difference.

3.5. Nonlinear theory fits

The results of the previous Section, although giving some clear indications in favor of the assumption $\theta_h = \theta_c$, rely on simple, qualitative comparisons of numerical estimates with linear theory predictions. All measured quantities present, on the other hand, important nonlinear evolution. In this Section we take full advantage of these measurements, providing fits in terms of models at one-loop in perturbation theory. This will result in a more quantitative discrimination between the two options based on a Bayesian analysis of the simulation data.

3.5.1. Data and analysis setup

In light of the results of the previous Section, the quantities most informative on the problem at hand are cross-power spectra between the halo distribution and the underlying matter density. We therefore perform a joint fit of the halo power spectrum and cross-power spectrum between halos and cold matter

3. Massive neutrinos and Galaxy Clustering




Figure 3.3.: Estimates of the momentum bias of halos, when defined with respect to the cold-matter momentum (*blue*) or the total matter one (*green*), obtained as the mean over 50 realizations at $z = 1$. The lower panel shows the ratio with respect to the halo momentum bias in a massless neutrino cosmology, in order to highlight the scale dependence. Data from the QUIJOTE set, to compare with and generalize [88]; error bars are present in both panels, yet too small to be visible.

in real and redshift space, although limited to the redshift-space monopole. The data-vector we consider is thus

$$\left\{ \hat{P}_{hr,c,0}(k), \hat{P}_{hs,c,0}(k), \hat{P}_{hr,hr,0}(k), \hat{P}_{hs,hs,0}(k) \right\},$$

with each measurement corresponding to the mean over 200 realizations of the QUIJOTE simulations, for a cumulative volume of $200 h^{-3} \text{ Gpc}^3$. The corresponding covariance is given by $C_{\text{mean}} = C/R$, with $R = 200$ and C being the covariance of the measurements in a single realization. The latter, which includes all cross-covariances among different spectra, is computed analytically in the Gaussian approximation according to the expressions described in the next Section.

We perform a Bayesian analysis and explore the posterior distribution by means of [106]. We adopt a Gaussian likelihood and sample the following parameters describing bias, counterterms and shot noise:

$$\left\{ b_1, b_2, b_{\mathcal{G}_2}, c_0^\times, c_2^\times, c_0, c_2, \alpha_P^{(r)}, \epsilon_{k^2}^{(r)}, \alpha_P^{(s)}, \epsilon_{k^2}^{(s)} \right\}$$

while keeping fixed cosmological parameters to their fiducial values. Notice that, as already mentioned, we assume the cross-power spectra counterterms to be independent from the auto-power spectra ones. The nonlocal bias term dependent on b_{Γ_3} is fixed to the value predicted by coevolution expectation $b_{\Gamma_3} = 23(b_1 - 1)/42$ [107–109], having checked that this does not affect our result. We use wide, uninformative priors for the parameters we vary.

Before applying it to the simulation with massive neutrinos, we first assess the validity range of the model on the ΛCDM simulations, where the interpretation of the Kaiser formula does not present ambiguity. To this end, we perform two tests: on one hand, we look at the running of the inferred value of the parameters with the maximum wavenumber of the fit while, on the other hand, we consider the goodness of the fit in terms of the χ^2 value averaged over the chain (as in, e.g., [92]). We find that up to

$k_{\max} = 0.1 \, h \, \text{Mpc}^{-1}$ both tests indicate that the model adequately describes the simulation data, and we will therefore adopt this value as a maximum in our analysis.

3.5.2. Analytical Cross-power spectra Gaussian covariance matrices

We provide here explicit expressions for the analytical covariance and cross-covariance of all power spectra and cross-power spectra in real and redshift space considered in the study. These results can be found, in part, in [110–112] while the extension to include redshift-space multipoles of cross-correlations has not been considered, to the best of our knowledge, in the literature so far.

A generic definition of covariance encompassing all cases considered in this study can be written as

$$C_{AB,CD}^{(\ell_1,\ell_2)}(k_1, k_2) \equiv \text{Cov} \left[\hat{P}_{AB,\ell_1}(k_1), \hat{P}_{CD,\ell_2}(k_2) \right], \quad (3.45)$$

where capital letters denote (cold) matter or halo densities in real and redshift space, e.g. $A \in \{c, h_r, h_s\}$, while the index ℓ refers to the power spectrum multipole estimator considered, with $\ell \in \{0, 2, 4\}$. Notice that, for simplicity, we assume the notation for the real-space estimator to coincide with the redshift-space, monopole estimator. In the Gaussian approximation we have

$$\begin{aligned} C_{AB,CD}^{(\ell_1,\ell_2)} &\equiv \left\langle \hat{P}_{AB,\ell_1}(k_1) \hat{P}_{CD,\ell_2}(k_2) \right\rangle - P_{AB,\ell_1}(k_1) P_{CD,\ell_2}(k_2) \\ &\simeq \frac{(2\ell_1 + 1)(2\ell_2 + 1)}{N_k^2} \sum_{|\mathbf{q}| \in k_1, |\mathbf{p}| \in k_2} \left[\delta_{\mathbf{q}+\mathbf{p}}^{(K)} P_{AC}(\mathbf{q}) P_{BD}(\mathbf{q}) + \delta_{\mathbf{q}-\mathbf{p}}^{(K)} P_{AD}(\mathbf{q}) P_{CB}(\mathbf{q}) \right] \times \\ &\quad \times \mathcal{L}_{\ell_1}(\mu_q) \mathcal{L}_{\ell_2}(\mu_p), \end{aligned} \quad (3.46)$$

which, in the thin-shell approximation, leads to

$$\begin{aligned} C_{AB,CD}^{(\ell_1,\ell_2)} &\simeq \frac{(2\ell_1 + 1)(2\ell_2 + 1)}{N_k^2} \int_{-1}^1 d\mu \left[P_{AC}(k, \mu) P_{BD}(k, \mu) + P_{AD}(k, \mu) P_{CB}(k, \mu) \right] \times \\ &\quad \times \mathcal{L}_{\ell_1}(\mu) \mathcal{L}_{\ell_2}(\mu), \end{aligned} \quad (3.47)$$

where we implicitly assume that all auto-power spectra $P_{AA}(k)$ include a shot noise contribution.

Specializing in (3.47) the tracer indexes to the relevant cases, one can therefore obtain all the combinations relevant for this study. The full covariance matrix for our data-vector

$$\left\{ \hat{P}_{h_r c, 0}, \hat{P}_{h_s c, 0}, \hat{P}_{h_r h_r}, \hat{P}_{h_s h_s, 0} \right\}$$

will be given by

$$C = \begin{pmatrix} C_{h_r c, h_r c}^{(0,0)} & C_{h_r c, h_s c}^{(0,0)} & C_{h_r c, h_r h_r}^{(0,0)} & C_{h_r c, h_s h_s}^{(0,0)} \\ C_{h_s c, h_r c}^{(0,0)} & C_{h_s c, h_s c}^{(0,0)} & C_{h_s c, h_r h_r}^{(0,0)} & C_{h_s c, h_s h_s}^{(0,0)} \\ C_{h_r h_r, h_r c}^{(0,0)} & C_{h_r h_r, h_s c}^{(0,0)} & C_{h_r h_r, h_r h_r}^{(0,0)} & C_{h_r h_r, h_s h_s}^{(0,0)} \\ C_{h_s h_s, h_r c}^{(0,0)} & C_{h_s h_s, h_s c}^{(0,0)} & C_{h_s h_s, h_r h_r}^{(0,0)} & C_{h_s h_s, h_s h_s}^{(0,0)} \end{pmatrix}, \quad (3.48)$$

where, clearly, $C_{AB,CD} = C_{CD,AB}$ and each block is diagonal and given by

$$\begin{aligned} C_{h_r c, h_r c}^{(0,0)} &= \frac{1}{N_k} (P_{h_r h_r} P_{cc} + P_{h_r c}^2) \\ C_{h_r c, h_s c}^{(0,0)} &= \frac{1}{N_k} (P_{h_r h_s, 0} P_{cc} + P_{h_r c} P_{h_s c, 0}) \\ C_{h_s c, h_s c}^{(0,0)} &= \frac{1}{N_k} \left(P_{h_s h_s, 0} P_{cc} + P_{h_s c, 0}^2 + \frac{1}{5} P_{h_s c, 2}^2 + \frac{1}{9} P_{h_s c, 4}^2 \right) \end{aligned}$$

3. Massive neutrinos and Galaxy Clustering




Figure 3.4.: Comparison between analytical (red curves), Eq. (3.47), and numerical covariance (black curves). In each panel, the upper half shows the diagonals (i.e. the variance) of some relevant blocks of the full covariance matrix, Eq. (3.48). The lower half shows the relative differences between the analytical prediction and the numerical expectation, with the gray band indicating for reference a 10% deviation.

$$\begin{aligned}
C_{h_r c, h_r h_r}^{(0,0)} &= \frac{2}{N_k} P_{h_r h_r} P_{h_r c} \\
C_{h_s c, h_r h_r}^{(0,0)} &= \frac{2}{N_k} P_{h_s h_r, 0} P_{h_r c} \\
C_{h_r c, h_s h_s}^{(0,0)} &= \frac{2}{N_k} \left(P_{h_s h_r, 0} P_{h_s c, 0} + \frac{1}{5} P_{h_s h_r, 2} P_{h_s c, 2} + \frac{1}{9} P_{h_s h_r, 4} P_{h_s c, 4} \right) \\
C_{h_s c, h_s h_s}^{(0,0)} &= \frac{2}{N_k} \left(P_{h_s h_s, 0} P_{h_s c, 0} + \frac{1}{5} P_{h_s h_s, 2} P_{h_s c, 2} + \frac{1}{9} P_{h_s h_s, 4} P_{h_s c, 4} \right) \\
C_{h_r h_r, h_r h_r}^{(0,0)} &= \frac{2}{N_k} (P_{h_r h_r}^2) \\
C_{h_r h_r, h_s h_s}^{(0,0)} &= \frac{2}{N_k} \left(P_{h_s h_r, 0}^2 + \frac{1}{5} P_{h_s h_r, 2}^2 + \frac{1}{9} P_{h_s h_r, 4}^2 \right) \\
C_{h_s h_s, h_s h_s}^{(0,0)} &= \frac{2}{N_k} \left(P_{h_s h_s, 0}^2 + \frac{1}{5} P_{h_s h_s, 2}^2 + \frac{1}{9} P_{h_s h_s, 4}^2 \right)
\end{aligned}$$

where N_k is the number of modes per bin. A comparison of the analytical prediction to their numerical estimate obtained from the 200 Qui jote simulations, for some relevant blocks of the full covariance, is reported in Fig. 3.4, showing a good agreement at least up to $0.2 h \text{ Mpc}^{-1}$, well beyond the scales relevant for our analysis.




Figure 3.5.: Results of the nonlinear fit alternatively assuming Eq. (3.17) ($\theta_h = \theta_c$, blue) and Eq. (3.19) ($\theta_h = \theta_m$, green). The datavector, common to both, consists of the average of 200 simulation boxes with $M_\nu = 0.4 \text{ eV}$ at $z = 1$. *Left panel:* Posterior-averaged reduced χ^2 for the two models as a function of the maximum wavenumber k_{max} with the shaded areas denoting values beyond the 95% confidence interval for the given number of degrees of freedom. *Right panel:* 1- and 2- σ posterior constraints on the c_2^x counterterm as a function of k_{max} .

3.5.3. Results: model selection and velocity bias

We first compare the two options introduced in Sect. 3.2, that is $\mathbf{u}_h = \mathbf{u}_c$ vs. $\mathbf{u}_h = \mathbf{u}_m$. On the left panel of Fig. 3.5 we plot the reduced posterior-averaged χ^2 as a function of the maximum wavenumber k_{max} adopted in the analysis for the two models, with the shaded areas denoting values beyond the 95% confidence interval for the given number of degrees of freedom. We notice how the total matter assumption, $\theta_h = \theta_m$, fails already at $0.1 h \text{ Mpc}^{-1}$ while the alternative provides an acceptable fit for all the intervals considered. Clearly, given that the datavector and the covariance are exactly the same, any difference in the χ^2 is due to the modeling. The right panel shows, as a prominent example, the 1- and 2- σ posterior constraints on the c_2^x counterterm as a function of k_{max} . In addition to the better goodness of fit, the cold-matter velocity model provides as well a less pronounced running of some parameters.

Clearly there are no reasons to consider the halo velocity field to coincide *either* with the cold-matter one *or* the total matter one: it is easy to see that this cannot be the case for an arbitrarily large neutrino fraction. To consider a more general, phenomenological model we consider the relation

$$\mathbf{u}_h = \mathbf{u}_m + b_\theta (\mathbf{u}_c - \mathbf{u}_m) \quad (3.49)$$

where the halo velocity is a linear combination of the two, with the parameter b_θ interpolating between the extremes of total-matter velocity ($b_\theta = 0$) and cold-matter velocity ($b_\theta = 1$). This gives the data freedom to potentially point at some intermediate combination of the two.

The left panel of Fig. 3.6 reports the inferred value of the b_θ parameter as a function of the k_{max} of the fit, for two values of redshift snapshot from the QUIJOTE simulations. We find that for $k_{\text{max}} > 0.08 h \text{ Mpc}^{-1}$ the data exclude $b_\theta = 0$ quite clearly, while being broadly consistent with $b_\theta = 1$. The




Figure 3.6.: *Left*: Marginalized posteriors of the parameter b_θ . This shows that, even in this general scenario of mixed velocity, there is a clear preference for the pure $\mathbf{u}_h = \mathbf{u}_c$ case ($b_\theta = 1$). *Right*: 2D posterior for the fit using (3.49) with $k_{\max} = 0.08 h/\text{Mpc}$, showing the strong degeneracy between the two parameters.

right panel of Fig. 3.6 reports the marginalized 2D posterior for $k_{\max} = 0.08 h \text{ Mpc}^{-1}$ to highlight the most relevant degeneracy between the velocity bias and the redshift-space counterterm.

3.6. Concluding remarks

In this chapter we presented a test of the linear Kaiser prediction [70] for the galaxy density in redshift space, Eq. (2.30), against numerical simulations of massive neutrino cosmologies. The first contribution to Eq. (2.30) has been shown to be accurately described in terms of the product of a constant bias parameter and the cold-matter (cold dark matter and baryons) density [52, 86]. The second contribution depends on the galaxy velocity field, usually assumed to coincide with the cold-matter velocity, justified by the results of [52, 86]. Such assumption has not been thoroughly explored with N-body simulations, despite some early attempts characterized by low statistics [53] and a more recent test in terms of halo momentum without a quantitative comparison against the opposite assumption of a galaxy velocity following the total matter velocity [88].

Taking advantage of two large sets of 200 (QUIJOTE, [101]) and 50 (DEMNuNCov, [103, 104]) simulations we first provide, in Section 3.4, comparisons of several measured quantities to their linear theory prediction. We find, in particular, that the direct estimate of the linear growth rate $f(k)$ provides the most clear indication that a proper match between numerical and theoretical results is obtained only under the assumption $\theta_h = \theta_c$, see Fig. 3.2.

In Section 3.5 we explored instead a full fit of the nonlinear model in perturbation theory to measurements of the auto- and cross-power spectra of the halo density with cold-matter density, both in real and redshift space. This corresponds to a more quantitative assessment, based on a Bayesian analysis, showing the cold-matter velocity assumption providing a better fit to the numerical data and avoiding running of nuisance parameters, a typical sign of systematic errors in the theoretical model (see Fig. 3.5).

The full-shape analysis of current galaxy surveys as BOSS [62, 113, 114] or DESI [115] already assumed a cold-matter velocity for the galaxy distribution and likely no appreciable difference would result from the opposite choice given the current error bars. Yet, these findings contribute to the understanding of galaxy clustering in massive neutrino cosmologies *at the linear level* and might be relevant for a correct analysis of future galaxy redshift surveys (see, e.g., [116]). Naturally, they rely on the validity of the N-body simulations, that in the case of massive neutrino simulations are particularly challenging (see, e.g., [96, 102, 117–131]).

As a final remark, while our conclusions in the specific case of massive neutrinos cosmologies might come as expected, a more general description of both matter and velocity bias in more exotic scenarios of mixed cold and hot dark matter is at the moment an open and interesting problem, with different venues to be explored both numerically and analytically.

4. EFTofLSS and Mixed Dark Matter

In this chapter we present a theoretical work on the EFTofLSS, in particular a first attempt to extend its formulation to cosmologies containing a small fraction of non-cold dark matter, without specializing to the case of massive neutrinos. We highlight the main complication introduced by such a component, namely the scale dependence of the linear growth, which invalidates the usual separation of the time and scale dependence of the perturbation theory kernels, and we provide analytical insights which are then compared with the fully numerical results.

4.1. Motivation

In the Λ CDM scenario Dark matter (DM) is described on large scales as a collisionless, cold fluid whose energy density redshifts like the volume while the Universe expands. In other words, the clustering of DM particles is a process with no physical scale, fully controlled by gravity, and structure formation can proceed indefinitely down to the smallest scales.

This sharp prediction of the Λ CDM scenario can be tested against observations. If DM growth is strongly suppressed below some characteristic scale, we expect $\mathcal{O}(1)$ changes in the abundance and clustering properties of astrophysical objects whose size is right at or smaller than that scale. Counts of the Milky Way satellites [132–135], measurements of the Lyman- α transmitted flux through the Intergalactic Medium [42, 136–139], or strong lensing of galaxies [140–142] already place important bounds on scenarios where all of DM has vanishing fluctuations at wavenumbers larger than some characteristic value. Prototypical models realizing such suppressed growth are thermal warm DM, where the DM particles have a sufficiently high temperature that their random velocities prevent structure from forming below the free-streaming scale, and ultra-light scalar DM, where power is suppressed below a scale determined by quantum pressure support. The above-mentioned astrophysical probes are sensitive to characteristic wavenumbers k_s in the ballpark of $10 h \text{ Mpc}^{-1}$.

If the scale of the DM power suppression is cosmologically relevant, i.e. much larger than individual astrophysical objects, then it can be probed using the anisotropy of the Cosmic Microwave Background (CMB) and measurements of the LSS obtained from the clustering of galaxies. In this regime, the possibility that all of DM has suppressed fluctuations above k_s is already ruled out by data. This is reminiscent of the constraints on massive neutrinos [26], whose characteristic wavenumber is roughly $0.01 h \text{ Mpc}^{-1}$: given current uncertainties, neutrinos can form at most a $\sim 0.5\%$ fraction of the matter density [5, 31, 143–145].

Nevertheless, scenarios where a small fraction of the DM energy density has suppressed growth of fluctuations above a certain cosmologically relevant k_s , while the remainder is cold, are still allowed and highly plausible from a particle physics perspective. Using a common terminology, we will refer to

such models as mixed DM scenarios, where the fraction of non-cold component χ

$$f_\chi \equiv \frac{\Omega_\chi}{\Omega_m} \quad (4.1)$$

would set the magnitude of the expected deviations from pure CDM. The χ component can consist of, to name but a few possibilities, of ultra-light axions produced via vacuum misalignment in the early Universe [146–148], perhaps originating from string theory [55], or by light but massive thermal relics, including fermions such as gravitinos [149, 150] or bosons such as hot axions [46]. Furthermore, a suppression of clustering can be realized even for a cold DM component, in the presence of interactions with additional dark sector species [151, 152].

As the size of f_χ probed by cosmological observations decreases with increasing sensitivity, it becomes imperative to refine the theoretical predictions in order to avoid a biased inference. This is especially true in the current era of galaxy surveys, with DESI [10], Euclid [11], and the Vera Rubin Observatory [153] poised to revolutionize our understanding of the dark sector over the next few years [154–157]. Of particular interest are scenarios where the new characteristic scale k_s is within the perturbative range, and the theory is therefore amenable to ab-initio theoretical predictions for the LSS observables.

In this chapter, focusing on such perturbative scales, we present a first attempt to extend the EFTofLSS, whose Λ CDM formulation has been introduced in Sect. 2.1.2, to mixed DM scenarios where a fraction f_χ of the DM energy density is not cold. Our working assumption is that the non-cold species can be modeled, within linear perturbation theory, as a perfect fluid with non-negligible pressure perturbation, parametrized by a finite sound speed c_s^2 . As a consequence, its fluctuations are suppressed below a characteristic scale $\sim c_s/\mathcal{H}$, where $\mathcal{H}$ is the Hubble scale. We consider general forms for the time and scale dependences of c_s^2 . These encompass, for example, both ultra-light axions and light thermal relics, which we adopt as benchmarks to illustrate our results, as well as many other microscopic dark sector models whose detailed study we leave for future work.

The novel two-fluid extension of the EFTofLSS we develop here is most impactful when the characteristic wavenumber is within the observational window of galaxy surveys, very roughly $0.01 h \text{ Mpc}^{-1} \lesssim k_s \lesssim 1 h \text{ Mpc}^{-1}$. It is constructed via a perturbative expansion for small f_χ , which is justified a posteriori by the constraints we derive. Similar questions have been studied in previous literature, our work indeed heavily draws inspiration from, and is indebted to, the nearly two-decades-long effort of the LSS community to model the clustering of massive neutrinos in perturbation theory beyond the linear regime [1, 88, 95–97, 158–171]. Related efforts have also been carried out to consistently extend the EFTofLSS to the CDM-baryon system [172, 173].

As mentioned, at linear level we model the non-cold species as a perfect fluid, i.e., the presence of a finite sound speed is the only difference with respect to the CDM description. It should, however, be kept in mind that the phenomenology of some BSM scenarios goes beyond the perfect fluid regime. For instance, light relics that decoupled from thermal equilibrium while relativistic have a non-zero anisotropic shear. Although it decays with time, the shear can still contribute to the evolution of the density perturbations of the warm species, see Ref. [26] for a review. Ultra-light scalar fields are an exception in this sense as the linear theory dynamics can be exactly mapped onto that of a perfect fluid with non-zero sound speed [174]. In the following we assume that such additional effects, unless they can be approximately reabsorbed into an effective value for c_s^2 (as it is indeed the case for the shear of thermal relics), are negligible.

Beyond linear level, in general each BSM model can introduce intrinsic nonlinearities associated with its microphysical description. These need to be added to the nonlinearities of pure gravitational origin, which are universal and are the only ones included in our framework. Nevertheless, in Sect. 4.A.3 we study the intrinsic nonlinearities in the two scenarios we adopt as benchmarks in this paper, namely ultra-light scalar fields and thermal relics. In the former case the intrinsic nonlinearities are associated with the wave-like nature of the field [175], whereas in the latter they originate from the Vlasov equation. We show that the effects of intrinsic nonlinearities are minor in both scenarios, thus supporting our approach.

The study presented in this chapter is organized as follows. First, Sec. 4.2 reviews and presents new analytical linear solutions that accurately describe the evolution of the perturbations of the two coupled fluids, for any power-law dependence of the sound speed on both time and scale. These linear solutions, which are obtained by means of an expansion for small f_χ , are useful to elucidate the dependence of the observables on the underlying model parameters, but also represent the input for the subsequent nonlinear calculations.

Moving beyond linear order, it is well known that analytical solutions in the presence on a finite sound speed cannot be attained for all kinematic configurations. Nevertheless, in Sec. 4.3.1 we present nonlinear analytical expressions valid in the kinematic regimes where the involved momenta are well separated from k_s , which could be relevant for future measurements of the matter and galaxy bispectra. We do so for a specific choice of the time dependence of c_s^2 , which applies to both ultra-light axions and light thermal relics. Furthermore, in Sec. 4.3.1 we identify for the first time the exact solutions in the infrared limit. In Sec. 4.3.1 we clarify the ultraviolet structure of the two-fluid model and the associated new EFTofLSS counterterms required to absorb any possible divergences.

As an important goal of this study is to apply the new two-fluid EFTofLSS to perform cosmological tests of mixed DM scenarios, in Sec. 4.3.4 we devise a prescription, which allows us to evaluate the perturbative solutions for both CDM and the non-cold species with sufficient speed to be implemented in a Markov Chain Monte Carlo parameter exploration. When compared to the exact, but computationally very expensive, numerical solutions, our prescription makes an error significantly smaller than the uncertainties in the measurements by current and upcoming surveys, as we show in Sec. 4.3.4.

As a first, proof-of-principle application of our framework, in Sec. 4.4 we present updated constraints on the energy density in ultra-light axions from a combination of Planck CMB and BOSS LSS data, highlighting that a consistent description of BSM effects is essential to ensure the robustness of potential discoveries by galaxy surveys.

4.2. Mixed DM Linear Cosmology

We assume the presence of a non-cold DM species χ and a cold component c , where we will always refer to χ as the “warm” component, even though this may technically be a misnomer for some microscopic realizations. In this two-fluid context, the pure CDM Eqs. (2.6), keeping also the non-linear sources as in Eqs. 2.14, extend to [81]

$$\delta'_\chi + \theta_\chi = -\nabla_i(\delta_\chi v_\chi^i), \quad (4.2a)$$

$$\theta'_\chi + \mathcal{H}\theta_\chi + \frac{1}{\bar{\rho}_\chi}\nabla^2\delta\mathcal{P}_\chi + \frac{3}{2}\mathcal{H}^2\Omega_m[f_\chi\delta_\chi + (1-f_\chi)\delta_c] = -\nabla_i(v_\chi^j\nabla_j v_\chi^i), \quad (4.2b)$$

$$\delta'_c + \theta_c = -\nabla_i(\delta_c v_c^i), \quad (4.2c)$$

$$\theta'_c + \mathcal{H}\theta_c + \frac{3}{2}\mathcal{H}^2\Omega_m[f_\chi\delta_\chi + (1-f_\chi)\delta_c] = -\nabla_i(v_c^j\nabla_j v_c^i), \quad (4.2d)$$

where, importantly, now also the pressure field is appearing. As in Chapt. 2, the subset c refers to the combined CDM+baryon fluid: $\delta_c = (f_{\text{cdm}}\delta_{\text{cdm}} + f_b\delta_b)/(f_{\text{cdm}} + f_b)$ and similarly for v_c^i . The background pressure is assumed to be negligible for both species, hence the background energy densities redshift as $\bar{\rho}_x \propto a^{-3}$ at all times $a > a_{\text{ini}}$ where the description in Eqs. (4.2) applies. It is also useful to define the total and relative density perturbations as

$$\delta_m \equiv f_\chi\delta_\chi + (1-f_\chi)\delta_c, \quad \delta_r \equiv \delta_\chi - \delta_c, \quad (4.3)$$

with analogous expressions for the velocity fields.

In the perturbed Euler equation for the warm component (4.2b), along the pressure term $\nabla^2\delta\mathcal{P}_\chi/\bar{\rho}_\chi$, we do not include a possible anisotropic shear stress term, which would manifest as $-\nabla^2\sigma_\chi$ on the left-hand side of Eq. (4.2b). This is fully accurate for some microscopic scenarios described by our framework (such as ultra-light axions [56]), whereas in other scenarios it is possible to absorb at least partially the effect of shear stress into a suitable modification of the pressure term, as it was discussed for the massive neutrino case in Sect. 3.1.

To solve Eqs. (4.2) we change the time variable to $\eta \equiv \log D_{\Lambda\text{CDM}}$, where $D_{\Lambda\text{CDM}}$ is the linear growth factor of a ΛCDM cosmology with the same total matter energy density as the mixed warm/cold scenario we are considering [169]. We obtain

$$\partial_\eta\delta_\chi - \Theta_\chi = \nabla_i(\delta_\chi V_\chi^i), \quad (4.4a)$$

$$\begin{aligned} \partial_\eta\Theta_\chi + \left(\frac{3\Omega_m}{2f_{\Lambda\text{CDM}}^2} - 1\right)\Theta_\chi - \frac{\nabla^2\delta\mathcal{P}_\chi}{f_{\Lambda\text{CDM}}^2\mathcal{H}^2\bar{\rho}_\chi} - \frac{3\Omega_m}{2f_{\Lambda\text{CDM}}^2}[f_\chi\delta_\chi + (1-f_\chi)\delta_c] \\ = \nabla_i(V_\chi^j\nabla_j V_\chi^i), \end{aligned} \quad (4.4b)$$

$$\partial_\eta\delta_c - \Theta_c = \nabla_i(\delta_c V_c^i), \quad (4.4c)$$

$$\partial_\eta\Theta_c + \left(\frac{3\Omega_m}{2f_{\Lambda\text{CDM}}^2} - 1\right)\Theta_c - \frac{3\Omega_m}{2f_{\Lambda\text{CDM}}^2}[f_\chi\delta_\chi + (1-f_\chi)\delta_c] = \nabla_i(V_c^j\nabla_j V_c^i), \quad (4.4d)$$

where the Λ CDM growth factor is defined in Eq. (2.7), and we introduce generically for two fluids

$$V_x^i = -\frac{v_x^i}{f_{\Lambda\text{CDM}}\mathcal{H}}, \quad \Theta_x = \nabla_i V_x^i. \quad (4.5)$$

Henceforth, we work in Fourier space. Setting $\Omega_m/f_{\Lambda\text{CDM}}^2 \rightarrow 1$, which is known to be a highly accurate approximation in ΛCDM cosmologies even during vacuum energy domination (see e.g. Ref. [169] for an appraisal), we rewrite the equations as

$$\partial_\eta\delta_\chi(\mathbf{k}, \eta) - \Theta_\chi(\mathbf{k}, \eta) = \int_{\mathbf{k}} dk_{12}\alpha(\mathbf{k}_2, \mathbf{k}_1)\delta_\chi(\mathbf{k}_1, \eta)\Theta_\chi(\mathbf{k}_2, \eta), \quad (4.6a)$$

$$\begin{aligned} \partial_\eta\Theta_\chi(\mathbf{k}, \eta) + \frac{1}{2}\Theta_\chi(\mathbf{k}, \eta) + \frac{3}{2}c_{s,\text{eff}}^2k^2\delta_\chi(\mathbf{k}, \eta) - \frac{3}{2}[f_\chi\delta_\chi(\mathbf{k}, \eta) + (1-f_\chi)\delta_c(\mathbf{k}, \eta)] \\ = \int_{\mathbf{k}} dk_{12}\beta(\mathbf{k}_1, \mathbf{k}_2)\Theta_\chi(\mathbf{k}_1, \eta)\Theta_\chi(\mathbf{k}_2, \eta), \end{aligned} \quad (4.6b)$$

$$\partial_\eta\delta_c(\mathbf{k}, \eta) - \Theta_c(\mathbf{k}, \eta) = \int_{\mathbf{k}} dk_{12}\alpha(\mathbf{k}_2, \mathbf{k}_1)\delta_c(\mathbf{k}_1, \eta)\Theta_c(\mathbf{k}_2, \eta), \quad (4.6c)$$

$$\begin{aligned} \partial_\eta \Theta_c(\mathbf{k}, \eta) + \frac{1}{2} \Theta_c(\mathbf{k}, \eta) - \frac{3}{2} [f_\chi \delta_\chi(\mathbf{k}, \eta) + (1 - f_\chi) \delta_c(\mathbf{k}, \eta)] \\ = \int_{\mathbf{k}} dk_{12} \beta(\mathbf{k}_1, \mathbf{k}_2) \Theta_c(\mathbf{k}_1, \eta) \Theta_c(\mathbf{k}_2, \eta), \end{aligned} \quad (4.6d)$$

where $c_{s,\text{eff}}^2 = 2c_s^2/(3\mathcal{H}^2\Omega_m)$ and

$$c_s^2 \equiv \frac{\delta P_\chi(\mathbf{k}, \eta)}{\delta \rho_\chi(\mathbf{k}, \eta)} \quad (4.7)$$

is the sound speed of the warm fluid. The standard mode-coupling functions appearing in Eqs. (4.6) have been defined in Eq. (2.15), and the shorthand notation of Eq. (2.12) has been used.

The temporal and spatial variations of the sound speed in Eq. (4.7) encode the model dependence. We assume power-law scalings with wavenumber and scale factor,

$$c_s^2 \propto k^p a^{-\gamma}, \quad (4.8)$$

where the coefficients p and γ are a priori arbitrary. The matching to several concrete microphysical scenarios will be discussed later in Sec. 4.2.2. Since $\mathcal{H}^2\Omega_m \propto a^{-1}$, we parametrize

$$c_{s,\text{eff}}^2 \equiv \frac{k^p}{k_{s*}^{2+p}} \left(\frac{a_*}{a} \right)^{\gamma-1} \simeq \frac{k^p}{k_{s*}^{2+p}} e^{(\eta_* - \eta)(\gamma-1)} \quad (4.9)$$

where we found it convenient to define a reference time a_* . In the last equality we have made the Einstein-de Sitter (EdS) approximation $D_{\Lambda\text{CDM}} = a$, which implies $\eta = \log a$. This is the only instance where the EdS approximation (which also implies $\Omega_m/f_{\Lambda\text{CDM}}^2 \rightarrow 1$, but is more restrictive) is used and it will be justified a posteriori, by showing in Sec. 4.2.4 that our analytical linear solutions accurately reproduce exact numerical results down to the redshift $z \sim 0.5$ of observational relevance for galaxy surveys.

After substituting Eq. (4.9) into Eqs. (4.6), we observe that the explicit dependence on k introduced by the χ sound speed term can be removed (for $\gamma \neq 1$) by switching to the new momentum-dependent time variable

$$\tilde{\eta} \equiv \eta - \eta_* - \frac{2+p}{\gamma-1} \log \frac{k}{k_{s*}}. \quad (4.10)$$

Since $d\tilde{\eta} = d\eta$, the equations are then cast in their final form

$$\partial_{\tilde{\eta}} \delta_\chi(\mathbf{k}, \tilde{\eta}) - \Theta_\chi(\mathbf{k}, \tilde{\eta}) = \int_{\mathbf{k}} dk_{12} \alpha(\mathbf{k}_2, \mathbf{k}_1) \delta_\chi(\mathbf{k}_1, \tilde{\eta}) \Theta_\chi(\mathbf{k}_2, \tilde{\eta}), \quad (4.11a)$$

$$\begin{aligned} \partial_{\tilde{\eta}} \Theta_\chi(\mathbf{k}, \tilde{\eta}) + \frac{1}{2} \Theta_\chi(\mathbf{k}, \tilde{\eta}) + \frac{3}{2} e^{-(\gamma-1)\tilde{\eta}} \delta_\chi(\mathbf{k}, \tilde{\eta}) - \frac{3}{2} [f_\chi \delta_\chi(\mathbf{k}, \tilde{\eta}) + (1 - f_\chi) \delta_c(\mathbf{k}, \tilde{\eta})] \\ = \int_{\mathbf{k}} dk_{12} \beta(\mathbf{k}_1, \mathbf{k}_2) \Theta_\chi(\mathbf{k}_1, \tilde{\eta}) \Theta_\chi(\mathbf{k}_2, \tilde{\eta}), \end{aligned} \quad (4.11b)$$

$$\partial_{\tilde{\eta}} \delta_c(\mathbf{k}, \tilde{\eta}) - \Theta_c(\mathbf{k}, \tilde{\eta}) = \int_{\mathbf{k}} dk_{12} \alpha(\mathbf{k}_2, \mathbf{k}_1) \delta_c(\mathbf{k}_1, \tilde{\eta}) \Theta_c(\mathbf{k}_2, \tilde{\eta}), \quad (4.11c)$$

$$\begin{aligned} \partial_{\tilde{\eta}} \Theta_c(\mathbf{k}, \tilde{\eta}) + \frac{1}{2} \Theta_c(\mathbf{k}, \tilde{\eta}) - \frac{3}{2} [f_\chi \delta_\chi(\mathbf{k}, \tilde{\eta}) + (1 - f_\chi) \delta_c(\mathbf{k}, \tilde{\eta})] \\ = \int_{\mathbf{k}} dk_{12} \beta(\mathbf{k}_1, \mathbf{k}_2) \Theta_c(\mathbf{k}_1, \tilde{\eta}) \Theta_c(\mathbf{k}_2, \tilde{\eta}), \end{aligned} \quad (4.11d)$$

In what follows, we solve Eqs. (4.11) perturbatively in the density and velocity fluctuations. At each perturbative order, we further perform an expansion for $f_\chi \ll 1$. Hence, we write for the density fields

$$\delta_x(\mathbf{k}, \tilde{\eta}) = \sum_{n=1}^{\infty} \delta_x^{[n]}(\mathbf{k}, \tilde{\eta}) = \sum_{n=1}^{\infty} \left[\delta_x^{(0)[n]}(\mathbf{k}, \tilde{\eta}) + f_\chi \delta_x^{(1)[n]}(\mathbf{k}, \tilde{\eta}) + \mathcal{O}(f_\chi^2) \right], \quad (4.12)$$

and analogously for the velocity fields Θ_x , with $x = \chi, c$. As in Eq. (2.9) the $[n]$ index denotes the perturbative order, while here we introduce a new index (0) or (1) to denote the order in the small f_χ expansion. In practice, we will find that the key physical effects are captured by truncating the small f_χ expansion to the first nontrivial order for each species: $\mathcal{O}(f_\chi^0)$ for χ and $\mathcal{O}(f_\chi)$ for c , which behaves as CDM at zeroth order.

As common practice, we also define a time-dependent characteristic wavenumber $k_s(a)$ via the relation $c_{s,\text{eff}}^2 k^2 \equiv [k/k_s(a)]^{2+p}$, so that the sound speed term in the Euler equation for χ takes the form $\frac{3}{2}[k/k_s(a)]^{2+p}\delta_\chi$. From Eq. (4.9) we thus obtain

$$k_s(a) = k_{s*} \left(\frac{a}{a_*} \right)^{\frac{\gamma-1}{2+p}} \quad (4.13)$$

and $\tilde{\eta}$ can be written as

$$\tilde{\eta} \simeq -\frac{2+p}{\gamma-1} \log \frac{k}{k_s(a)} \quad (4.14)$$

by using the EdS approximation. From a microscopic perspective, $k_s(a)$ can have a variety of origins. Some examples will be provided in Sec. 4.2.2.

4.2.1. General linear solutions

First, we focus on solving Eqs. (4.11) at linear level, keeping the parameters γ and p arbitrary. At zeroth order in f_χ , the equations for c are combined into

$$\left(\frac{\partial^2}{\partial \tilde{\eta}^2} + \frac{1}{2} \frac{\partial}{\partial \tilde{\eta}} - \frac{3}{2} \right) \delta_c^{(0)[1]} = 0, \quad (4.15)$$

where we recall that [1] indicates the first (linear) perturbative order, while (0) indicates the zeroth order in f_χ . The above equation is solved by the standard growing mode for CDM, which in this notation reads $\delta_c^{(0)[1]}(\mathbf{k}, \tilde{\eta}) = \Theta_c^{(0)[1]}(\mathbf{k}, \tilde{\eta}) = e^{\tilde{\eta} - \tilde{\eta}_{\text{ini}}} \delta_0(\mathbf{k})$. Plugging this result into the equations for χ gives, still at zeroth order in f_χ ,

$$\left(\frac{\partial^2}{\partial \tilde{\eta}^2} + \frac{1}{2} \frac{\partial}{\partial \tilde{\eta}} + \frac{3}{2} e^{-(\gamma-1)\tilde{\eta}} \right) \delta_\chi^{(0)[1]} = \frac{3}{2} e^{\tilde{\eta} - \tilde{\eta}_{\text{ini}}} \delta_0(\mathbf{k}). \quad (4.16)$$

This is solved by

$$\delta_\chi^{(0)[1]}(\mathbf{k}, \tilde{\eta}) = \int_{-\infty}^{+\infty} d\tilde{\eta}' G_\chi(\tilde{\eta}, \tilde{\eta}') \frac{3}{2} e^{\tilde{\eta}' - \tilde{\eta}_{\text{ini}}} \delta_0(\mathbf{k}), \quad \Theta_\chi^{(0)[1]}(\mathbf{k}, \tilde{\eta}) = \partial_{\tilde{\eta}} \delta_\chi^{(0)[1]}(\mathbf{k}, \tilde{\eta}), \quad (4.17)$$

where the Green's function is expressed in terms of the homogeneous solutions as

$$G_\chi(\tilde{\eta}, \tilde{\eta}') = \frac{A_2(\tilde{\eta}) - A_1(\tilde{\eta}) \frac{A_2(\tilde{\eta}')}{A_1(\tilde{\eta}')}}{\left(\frac{\partial A_2}{\partial \tilde{\eta}} - \frac{\partial A_1}{\partial \tilde{\eta}} \frac{A_2}{A_1} \right)(\tilde{\eta}')} \theta_H(\tilde{\eta} - \tilde{\eta}'), \quad A_{1,2}(\tilde{\eta}) = e^{-\tilde{\eta}/4} J_{\mp \frac{1}{2(\gamma-1)}} \left(\frac{\sqrt{6} e^{-(\gamma-1)\tilde{\eta}/2}}{\gamma-1} \right), \quad (4.18)$$

with J_α denoting the Bessel function of order α and θ_H the Heaviside step function. Introducing a notation that will be used throughout this chapter, we define the ratios of linear perturbations

$$g(\tilde{\eta}) \equiv \frac{\delta_\chi^{[1]}(\mathbf{k}, \tilde{\eta})}{\delta_c^{[1]}(\mathbf{k}, \tilde{\eta})}, \quad h(\tilde{\eta}) \equiv \frac{\Theta_\chi^{[1]}(\mathbf{k}, \tilde{\eta})}{\delta_c^{[1]}(\mathbf{k}, \tilde{\eta})}, \quad s(\tilde{\eta}) \equiv \frac{\Theta_c^{[1]}(\mathbf{k}, \tilde{\eta})}{\delta_c^{[1]}(\mathbf{k}, \tilde{\eta})}. \quad (4.19)$$




Figure 4.1.: *Left*: The ratio of the linear density perturbations of χ and c , $g = \delta_\chi^{[1]}/\delta_c^{[1]}$, evaluated at zeroth order in the small f_χ expansion as a function of the k -dependent time variable $\tilde{\eta}$, for some illustrative choices of the parameter γ that sets the time dependence of the characteristic scale $k_s(a)$. *Right*: For $\gamma = 2$, the ratios of linear transfer functions g and $h = \Theta_\chi^{[1]}/\delta_c^{[1]}$ at zeroth order in f_χ , while $s = \Theta_c^{[1]}/\delta_c^{[1]}$ is evaluated up to first order with $f_\chi = 0.05$ or 0.15 . One has $s^{(0)} = 1$ identically, whereas $s^{(1)} = -3/5$ for $k \gg k_s(a)$ (namely, $\tilde{\eta} \ll 0$).

Assuming adiabatic initial conditions, as we do in this study, the quantities above are simply ratios of linear transfer functions. Notice that

$$s(\tilde{\eta}) = \partial_{\tilde{\eta}} \log \delta_c^{[1]} \equiv \frac{f(k)}{f_{\Lambda\text{CDM}}}, \quad (4.20)$$

where the last equality defines the scale-dependent linear growth rate $f(k)$.

In the left panel of Fig. 4.1 we show $g^{(0)}$, evaluated at the leading order in the small f_χ expansion, for some (integer) choices of γ . We see that our framework encompasses very different possible time evolutions of the characteristic scale $k_s(a)$ in Eq. (4.13), which can either increase with time (as observed for $\gamma = 2, 3$ in the figure) or decrease ($\gamma = 0$).¹ We also have $h^{(0)} = g^{(0)} + \partial_{\tilde{\eta}} g^{(0)}$, whereas for the cold species we simply have $s^{(0)} = 1$ at this order.

Next, we solve the equations for the cold fluid at $\mathcal{O}(f_\chi)$. The correction to the growth rate, $s(\tilde{\eta}) = 1 + f_\chi s^{(1)}(\tilde{\eta})$, satisfies the first order differential equation

$$\left(\partial_{\tilde{\eta}} + \frac{5}{2}\right)s^{(1)} = \frac{3}{2}(g^{(0)} - 1), \quad (4.21)$$

whose solution reads

$$s^{(1)}(\tilde{\eta}) = \int_{-\infty}^{+\infty} d\tilde{\eta}' G_s(\tilde{\eta}, \tilde{\eta}') \frac{3}{2}(g^{(0)}(\tilde{\eta}') - 1), \quad G_s(\tilde{\eta}, \tilde{\eta}') = e^{-\frac{5}{2}(\tilde{\eta} - \tilde{\eta}')} \theta_H(\tilde{\eta} - \tilde{\eta}'). \quad (4.22)$$

We then invert the relation $s(\tilde{\eta}) = \partial_{\tilde{\eta}} \log \delta_c^{[1]}$ to arrive at

$$\delta_c^{(1)[1]}(\mathbf{k}, \tilde{\eta}) = \int_{\tilde{\eta}_{\text{ini}}}^{\tilde{\eta}} d\tilde{\eta}' s^{(1)}(\tilde{\eta}') \delta_c^{(0)[1]}(\mathbf{k}, \tilde{\eta}), \quad (4.23a)$$

¹For $\gamma = 1$ the characteristic scale is time-independent. In this case we solve the equations directly in η , obtaining at zeroth order in f_χ the solutions $\delta_c^{(0)[1]}(\mathbf{k}, \eta) = \Theta_c^{(0)[1]}(\mathbf{k}, \eta) = e^{\eta - \eta_{\text{ini}}} \delta_0(\mathbf{k})$ and $\delta_\chi^{(0)[1]}(\mathbf{k}, \eta) = \Theta_\chi^{(0)[1]}(\mathbf{k}, \eta) = \delta_c^{(0)[1]}(\mathbf{k}, \eta) / [1 + (k/k_{s*})^{2+p}]$.

$$\Theta_c^{(1)[1]}(\mathbf{k}, \tilde{\eta}) = \left(\int_{\tilde{\eta}_{\text{ini}}}^{\tilde{\eta}} d\tilde{\eta}' s^{(1)}(\tilde{\eta}') + s^{(1)}(\tilde{\eta}) \right) \delta_c^{(0)[1]}(\mathbf{k}, \tilde{\eta}). \quad (4.23b)$$

4.2.2. Matching to microscopic models

In the remainder of this section we focus the discussion on $\gamma = 2$, corresponding to a time dependence of the sound speed $c_s^2 \propto k^p a^{-2}$. This case encompasses two widely-studied scenarios, which we adopt as benchmarks to illustrate our results: (i) ultra-light axions (ULAs, $p = 2$), where k_s is identified with the Jeans scale induced by quantum pressure support; and (ii) thermal relics ($p = 0$), including the standard neutrinos, where k_s is interpreted as the free streaming scale remnant of a thermal distribution. Anyway, it should be emphasized that the framework developed in this paper is expected to be applicable to a much broader range of dark sector models.

For $\gamma = 2$, Eq. (4.17) gives the closed-form expressions

$$\delta_\chi^{(0)[1]}(\mathbf{k}, \tilde{\eta}) = g^{(0)}(\tilde{\eta}) \delta_c^{(0)[1]}(\mathbf{k}, \tilde{\eta}), \quad \Theta_\chi^{(0)[1]}(\mathbf{k}, \tilde{\eta}) = h^{(0)}(\tilde{\eta}) \delta_c^{(0)[1]}(\mathbf{k}, \tilde{\eta}), \quad (4.24)$$

where $h^{(0)} = g^{(0)} + \partial_{\tilde{\eta}} g^{(0)}$ and

$$g^{(0)}(\tilde{\eta}) = G^{(0)}(\sqrt{6}e^{-\tilde{\eta}/2}), \quad G^{(0)}(x) \equiv 1 + x^2 \left[\left(\text{Si}(x) - \frac{\pi}{2} \right) \sin x + \text{Ci}(x) \cos x \right]. \quad (4.25)$$

Si and Ci denote the sine integral and cosine integral functions.² For $k \gg k_s(a)$ (namely $\tilde{\eta} \ll 0$) the χ perturbations are extremely suppressed and $g^{(0)}, h^{(0)} \rightarrow 0$ with $h^{(0)}/g^{(0)} \rightarrow 2$. Conversely, for $k \ll k_s(a)$ (corresponding to $\tilde{\eta} \gg 0$) the χ species becomes indistinguishable from the cold component and $g^{(0)}, h^{(0)} \rightarrow 1$.

Including $\mathcal{O}(f_\chi)$ corrections, the growth rate of the cold species takes the constant value $s = 1 - 3f_\chi/5$ for $k \gg k_s(a)$. This reproduces, as it should, the result found in the case of massive neutrinos presented in Eq. (3.10). By contrast, again just rephrasing in a different fashion the other limit of Eq. (3.10), on large scales $k \ll k_s(a)$ the growth rate takes the Λ CDM value, $s = 1$. In the right panel of Fig. 4.1 we show the behavior of the ratios of transfer functions g, h and s for $\gamma = 2$. Each of them is evaluated to the first non-trivial order in the small f_χ expansion, namely $\mathcal{O}(f_\chi^0)$ for g and h and $\mathcal{O}(f_\chi)$ for s .

We now turn to a brief description of our two benchmarks of choice, ULAs and light thermal relics, as well as their matching to our framework.

4.2.3. Ultra-light axions

In this section, we briefly introduce the physics of ultra-light axions, a class of models that will be central to this work as the benchmark scenario constrained with real data at the end of this chapter and again in Chapter 6. This fundamental physics realization is indeed ideal as a proxy for searches of new physics with galaxy clustering, due the rich phenomenology arising from its broad parameter space, e.g. compared to thermal relics. While we give the basic concepts related to the cosmological evolution of ULAs, for a more complete introduction and the particle physics motivation we refer the reader to the relevant literature [54, 55, 146–148, 177–183].

ULAs are described by a scalar field ϕ , minimally coupled to gravity but otherwise decoupled from all the other cosmic species, with a potential that is well approximated by a pure mass term,

²Solutions similar to Eqs. (4.24) and (4.25) were given in Ref. [176] where, however, Heaviside function-like initial conditions were assumed. Our results are consistent with Eq. (4.9) in Ref. [176], but we automatically implement the natural initial conditions, which yield the far simpler expression for $g^{(0)}$ in Eq. (4.25) above.




Figure 4.2.: Schematic representation of the cosmological features of ULAs, in the range of masses for which the suppression scale k_{drop} , as defined in Eqs. (4.28)–(4.29), is on cosmological scales (middle row). Which observables carry the strongest observational imprints (bottom row) mainly depends on the background evolution after the redshift of matter-radiation equality, z_{eq} (top row).

$V(\phi) = m_a^2 \phi^2 / 2$.³ The evolution of the background value of the field, $\bar{\phi}$, is governed by the Klein-Gordon equation

$$\bar{\phi}'' + 2\mathcal{H}\bar{\phi}' + a^2 m_a^2 \bar{\phi} = 0, \quad (4.26)$$

where primes denote, again, derivatives with respect to conformal time. The initial conditions are $\phi(a_i)$, which is traded in practice for the present-day fraction of energy density in ULAs, $\Omega_a h^2$, and $\bar{\phi}'(a_i) = 0$. At first, the field remains stuck due to Hubble friction and acts as a dark energy component, $\bar{\rho}_a \propto a^0$. Once the expansion of the Universe has slowed down to a sufficient degree, namely around $a = a_{m_a}$ with $3H(a_{m_a}) = m_a$, the field begins to oscillate rapidly around the minimum of its quadratic potential. Upon averaging over the coherent oscillations, the energy density in ULAs redshifts as nonrelativistic matter, $\bar{\rho}_a \propto a^{-3}$. For this reason, the phenomenology is strictly tied to the value of the mass of the field, as briefly sketched in Fig. 4.2.

At the level of perturbations, in this same regime one can accurately describe the evolution of the energy density and pressure of ϕ using an effective fluid approximation [54, 183, 186, 187]. The description in our Eqs. (4.6), with $\gamma = p = 2$, starts applying at $a_{ini} = a_{m_a}$ if the ULA begins to oscillate in matter domination (MD), whereas it starts at $a_{ini} = a_{eq}$ if the oscillations begin during radiation domination (RD). Taking $c_s^2 = k^2 / (4m_a^2 a^2)$ [54, 183, 188], we obtain for the characteristic (Jeans) scale

$$k_s(a) \approx a^{1/4} 3.9 \times 10^{-4} \frac{h}{\text{Mpc}} \left(\frac{m_a}{H_0} \right)^{1/2} \left(\frac{\Omega_m^0}{0.3} \right)^{1/4}, \quad (4.27)$$

where $H_0 = 2.1 \times 10^{-33} h \text{ eV}$. Notice that the assumption $\sigma(\mathbf{k}, \tau) = 0$ made in our framework is fully justified for ULAs, because a scalar field minimally coupled to gravity has zero anisotropic stress at linear order in perturbation theory [174]. However, a complete fluid-dynamical description for ULAs would also include intrinsic wave nonlinearities, which stem from the higher-order terms in a perturbative expansion of the quantum pressure (see for instance Ref. [175]). In Appendix 4.A.3 we show that the impact on our results of these intrinsic nonlinearities is mild, which justifies our choice to omit them.

The matter power spectrum for ULA cosmologies has a drop-like feature, as a consequence of the Jeans scale. The location of the feature is determined by m_a , whereas its depth is fixed by f_χ . To obtain

³Following the implementation in the Boltzmann solver axiCLASS [184, 185], the ULA potential is modeled as $V(\phi) = m_a^2 f_a^2 [1 - \cos(\phi/f_a)]$, where f_a is set to $0.4 M_{\text{Pl}}$. We have verified that this produces negligible differences with respect to a pure mass term in the parameter space we consider.

4. EFTofLSS and Mixed Dark Matter

an analytical estimate of the location, we need to distinguish between ULAs that became nonrelativistic during MD or RD. For ULAs that became nonrelativistic during MD, the condition $3H(a_{m_a}) = m_a$ is satisfied at $a_{m_a} = [3(\Omega_m^0)^{1/2}H_0/m_a]^{2/3}$. Imposing that $a_{m_a} > a_{\text{eq}}$ then sets the upper edge of this region as $m_a/H_0 < 3 \times 10^5$ or equivalently $m_a < 4 \times 10^{-28}$ eV. The drop in the power spectrum is located at the smallest k that is affected by the suppressed growth [181], namely $k_s(a_{m_a})$. From Eq. (4.27) we find

$$k_{\text{drop}} = k_s(a_{m_a}) \approx 4.2 \times 10^{-4} \frac{h}{\text{Mpc}} \left(\frac{m_a}{H_0} \right)^{1/3} \left(\frac{\Omega_m^0}{0.3} \right)^{1/3}, \quad (m_a/H_0 < 3 \times 10^5) \quad (4.28)$$

with a $m_a^{1/3}$ scaling.⁴ For ULAs that became nonrelativistic during RD, k_{drop} saturates at the value of k_s at matter-radiation equality, which yields

$$k_{\text{drop}} = k_s(a_{\text{eq}}) \approx 5.2 \times 10^{-2} \frac{h}{\text{Mpc}} \left(\frac{m_a}{10^6 H_0} \right)^{1/2}, \quad (m_a/H_0 > 3 \times 10^5) \quad (4.29)$$

scaling as $m_a^{1/2}$.

4.2.4. Light but massive thermal relics

Light but massive thermal relics [44, 45, 83] correspond to $\gamma = 2$ and $p = 0$. These species decouple from the thermal bath while relativistic, preserving their unperturbed phase space distribution, which takes a Fermi-Dirac or Bose-Einstein form. Eventually, they become nonrelativistic at $a = a_{\text{nr}}$, but they retain a cosmologically relevant free-streaming scale. In general, light relics are characterized by their mass, number of degrees of freedom and temperature. The temperature, in turn, depends on when the species decoupled from the thermal bath.

In this study, to illustrate the main features of the application to light relics we focus exclusively on the most familiar among them, namely the standard massive neutrinos. The phenomenology has already been introduced in Sect. 3.1, and importantly for the present purpose also the expression that links the fundamental physics parameter m_ν to the free-streaming scale $k_s(a)$, namely Eq. (3.7). Connecting that to the language of this chapter, for massive neutrinos the location k_{drop} of the drop-like feature in the power spectrum is approximately given by $k_s(a_{\text{nr}})$ if the transition to the nonrelativistic regime happens during MD and by $k_s(a_{\text{eq}})$ if it happens during RD.

For the purpose of going beyond linear theory, a general fluid description of thermal relics would also include intrinsic nonlinearities, which arise upon taking moments of the Vlasov equation to derive the continuity and Euler equations. In Sect. 4.A.3 we show that these intrinsic nonlinearities have a mild impact on our results, supporting our approach where only nonlinearities of pure gravitational origin are accounted for. In this chapter we view the standard neutrinos as a convenient proxy for more general light thermal relics, hence for the sake of illustration we show results for a variety of neutrino masses, including values up to $m_\nu \sim 1$ eV that have already been excluded by data.




Figure 4.3.: *Left*: Time evolution of the characteristic wavenumber $k_s(a)$ for two example models described by our framework: an ULA field with $m_a = 10^{-27}$ eV and $f_\chi = 0.1$ (purple line) and a thermal relic massive neutrino with $m_\nu = 1$ eV and $f_\chi \approx 0.08$ (dashed green line). The black curves show the evolution of the Hubble scale $\mathcal{H}$ for the ULA (solid) and massive neutrino (dashed) cosmologies. The dotted horizontal line indicates a representative k mode that evolves differently in the two models. *Right*: The ratio of linear perturbations $\delta_\chi^{[1]}/\delta_c^{[1]}$ as a function of k/k_s , evaluated at $z = 0.5$, for the same cosmologies shown in the left panel. Dashed curves correspond to our universal closed-form approximation $g^{(0)}\left(-\frac{2+p}{\gamma-1} \log \frac{k}{k_s}\right)$, with $\gamma = 2$ and $p = 2$ ($p = 0$) for ULAs (massive neutrinos). Solid curves show the same ratios as obtained from CLASS.

Accuracy of the analytical linear solutions and discussion

The time evolution of k_s is illustrated in the left panel of Fig. 4.3 for two example cosmologies, one containing ULAs and the other containing thermal relics in the form of massive neutrinos. For the former we assume $m_a = 10^{-27}$ eV with $f_\chi = 0.1$, whereas for the latter we take one neutrino of mass $m_\nu = 1$ eV, corresponding to a warm fraction $f_\chi \approx 0.08$.

For the same two cosmologies, in the right panel of Fig. 4.3 we compare our closed-form approximation for the ratio of transfer functions $\delta_\chi^{[1]}/\delta_c^{[1]}$, which recalling Eq. (4.14) we write as $g^{(0)}\left(-\frac{2+p}{\gamma-1} \log \frac{k}{k_s}\right)$ with the $g^{(0)}$ function defined in Eq. (4.25), to the numerical value obtained from the numerical Einstein-Boltzmann solver CLASS [60, 191]. Specifically, for massive neutrinos we use the standard CLASS release [192] with `ncdm_fluid_approximation = 1` [174], while for ULAs we employ AxICLASS [184, 185]. Excellent agreement is observed in both cases: our universal description of linear fluctuations accurately predicts the ratios of linear transfer functions. The ULA model has larger p , therefore a sharper transition around $k \sim k_s$ is observed.

As can be seen in the left panel of Fig. 4.3, the parameters of the two benchmark cosmologies were chosen to yield similar values of $k_s(z_{\text{obs}})$ at the low redshift $z_{\text{obs}} \sim 0.5$ observed by galaxy surveys. The warm fraction f_χ also takes similar values in the two scenarios. Nevertheless, we expect different imprints on the matter power spectrum due to the different time dependence of k_s : for instance, the $k = 0.08 h \text{ Mpc}^{-1}$ mode (dotted horizontal line in the left panel of Fig. 4.3) starts growing later in the

⁴An alternative definition of the location of the drop is the horizon crossing scale $k_m = a_{m_a} H(a_{m_a})$, which was used for instance in Ref. [189]. As it was pointed out in Ref. [190], the two definitions agree aside from $\mathcal{O}(1)$ numbers: we find $k_m = (2/9)^{2/3} k_{\text{drop}} \approx 0.37 k_{\text{drop}}$.




Figure 4.4.: *Left*: Ratio of the linear matter power spectrum to Λ CDM (solid purple) compared to the ratio of linear perturbations $\delta_\chi^{[1]}/\delta_c^{[1]}$ (solid blue) for the same ULA cosmology as in Fig. 4.3: $m_a = 10^{-27}$ eV and $f_\chi = 0.1$. Both ratios are evaluated at $z = 0.5$ and obtained from CLASS. The dot-dashed vertical lines show analytical approximations for the two relevant dynamical scales: $k_{\text{drop}} = k_s(z_{\text{eq}})$, which determines the beginning of the drop in the linear power spectrum, and $k_s(z)$, which marks the end of the step. *Right*: Same as in the left panel, but for the massive neutrino cosmology considered in Fig. 4.3: $m_\nu = 1$ eV and $f_\chi \approx 0.08$. Here the dot-dashed vertical lines show $k_{\text{drop}} \simeq k_s(z_{\text{nr}})$ and $k_s(z)$.

massive neutrino model, hence the power spectrum will be more suppressed at this wavenumber than in the ULA scenario. These expectations are confirmed by Fig. 4.4, where the CLASS output for the linear power spectrum of total matter, $P_{\delta_m \delta_m}^L$, is shown for the ULA (left panel) and massive neutrino (right panel) cosmologies, both normalized to the Λ CDM prediction. For comparison, we also show the ratio of transfer functions $\delta_\chi^{[1]}/\delta_c^{[1]}$ that appeared in the right panel of Fig. 4.3.

Figure 4.4 highlights the clear separation between the two scales at play: k_{drop} , which corresponds to the smallest k where $P_{\delta_m \delta_m}^L$ is suppressed relative to Λ CDM, and $k_s(z_{\text{obs}})$, which is relevant for dynamics at the low redshift z_{obs} and marks the end of the step-like feature. The ratio between the two scales, which determines the width of the step, is bounded from above by

$$\frac{k_s(z_{\text{obs}})}{k_{\text{drop}}} \leq \left(\frac{1 + z_{\text{eq}}}{1 + z_{\text{obs}}} \right)^{\frac{1}{2+p}}, \quad (4.30)$$

where the inequality is saturated when $k_{\text{drop}} = k_s(z_{\text{eq}})$. Taking a representative $z_{\text{obs}} = 0.5$, we find an upper bound of ~ 7 for ULAs ($p = 2$), which is indeed saturated in the left panel of Fig. 4.4 because for the considered mass the axion field starts oscillating before matter-radiation equality. The hierarchy between the two scales can be much larger for light thermal relics ($p = 0$), with an upper bound of ~ 50 . In the cosmology assumed in the right panel of Fig. 4.4 the neutrino becomes nonrelativistic during MD, hence $k_{\text{drop}} \simeq k_s(z_{\text{nr}})$ and the ratio is about 30.

On the other hand, the depth of the step in the matter power spectrum is determined by the fraction f_χ . At the linear level, using our analytical solutions we readily estimate the depth as

$$\left. \frac{P_{\delta_m \delta_m}^L}{P_{\delta_m \delta_m}^{L, \Lambda\text{CDM}}} \right|_{k \gg k_s(z_{\text{obs}})} = (1 - f_\chi)^2 \exp \left(-\frac{6}{5} f_\chi \log \frac{1 + z_{\text{ini}}}{1 + z_{\text{obs}}} \right). \quad (4.31)$$

For the ULA (neutrino) cosmology in Fig. 4.4 this gives ≈ 0.32 (0.42), in $\sim 20\%$ agreement with the numerical result ≈ 0.28 (0.34). This is reasonable, if one considers that in the estimate for the ULA

we assumed instantaneous transition from radiation to matter domination ($z_{\text{ini}} = z_{\text{eq}}$), whereas for the neutrino we assumed instantaneous transition from the relativistic to the nonrelativistic regime ($z_{\text{ini}} = z_{\text{nr}}$).

Importantly, we note that expanding the linear power spectrum for small f_χ would not be appropriate, due to the large logarithm appearing in the exponent of Eq. (4.31) (and the relatively large size of f_χ considered). It is well known that such large logarithm, which corresponds to a time-integrated effect, can invalidate the expansion of the linear perturbations of the cold species [163]. For this reason, in our approach the expansion in $f_\chi \ll 1$ is never performed on $\delta_c^{[1]}$, but only on quantities that are by construction free from large time logarithms, namely *ratios* of linear perturbations ($\delta_\chi^{[1]}/\delta_c^{[1]}$ and $\Theta_x^{[1]}/\delta_c^{[1]}$ for $x = \chi, c$) and nonlinear kernels. In our framework and results the linear power spectrum of the cold species, $P_{\delta_c \delta_c}^L$, is always obtained from a full numerical solution computed with CLASS.

4.3. Galaxy Power Spectrum in Redshift Space

In this section we present all the theoretical ingredients needed to predict the 1-loop power spectrum of galaxies in redshift space, which can be compared to data.

4.3.1. Matter nonlinearities

We begin by discussing solutions beyond the linear order to Eqs. (4.11), which describe the dynamics of the two coupled fluids χ and c . Using the notation introduced in Eq. (4.12), we trivially extend the perturbative expansion Eq. (2.10) to the two-fluid system, which now in full generality entails time-dependent couplings, as

$$\delta_\chi^{[n]}(\mathbf{k}, \tilde{\eta}) = \int_{\mathbf{k}} dk_{1\dots n} F_\chi^{[n]}(\mathbf{k}_1, \dots, \mathbf{k}_n; \tilde{\eta}) \delta_\chi^{[1]}(\mathbf{k}_1, \tilde{\eta}) \dots \delta_\chi^{[1]}(\mathbf{k}_n, \tilde{\eta}), \quad (4.32a)$$

$$\Theta_\chi^{[n]}(\mathbf{k}, \tilde{\eta}) = \int_{\mathbf{k}} dk_{1\dots n} G_\chi^{[n]}(\mathbf{k}_1, \dots, \mathbf{k}_n; \tilde{\eta}) \delta_\chi^{[1]}(\mathbf{k}_1, \tilde{\eta}) \dots \delta_\chi^{[1]}(\mathbf{k}_n, \tilde{\eta}), \quad (4.32b)$$

$$\delta_c^{[n]}(\mathbf{k}, \tilde{\eta}) = \int_{\mathbf{k}} dk_{1\dots n} F_c^{[n]}(\mathbf{k}_1, \dots, \mathbf{k}_n; \tilde{\eta}) \delta_c^{[1]}(\mathbf{k}_1, \tilde{\eta}) \dots \delta_c^{[1]}(\mathbf{k}_n, \tilde{\eta}), \quad (4.32c)$$

$$\Theta_c^{[n]}(\mathbf{k}, \tilde{\eta}) = \int_{\mathbf{k}} dk_{1\dots n} G_c^{[n]}(\mathbf{k}_1, \dots, \mathbf{k}_n; \tilde{\eta}) \delta_c^{[1]}(\mathbf{k}_1, \tilde{\eta}) \dots \delta_c^{[1]}(\mathbf{k}_n, \tilde{\eta}). \quad (4.32d)$$

Note that we chose to expand the nonlinear corrections for the χ species in terms of its own linear field $\delta_\chi^{[1]}$, rather than $\delta_c^{[1]}$. Importantly, our definition of the momentum-dependent time variable $\tilde{\eta}$ always refers to the external momentum $\mathbf{k} = \mathbf{k}_1 + \dots + \mathbf{k}_n$. The linear solutions found in Sec. 4.2.1 set

$$F_\chi^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) = F_c^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) = 1, \quad G_\chi^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) = \frac{h(\tilde{\eta}_{k_1})}{g(\tilde{\eta}_{k_1})}, \quad G_c^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) = s(\tilde{\eta}_{k_1}), \quad (4.33)$$

where we defined

$$\tilde{\eta}_q \equiv \tilde{\eta} - \frac{2+p}{\gamma-1} \log \frac{q}{k} \quad (4.34)$$

for an arbitrary vector $\mathbf{q}$.

At nonlinear order, differential equations for the kernels are found by substituting the expansions (4.32) into Eqs. (4.11). For example, at second order ($n = 2$) we obtain

$$\left(\partial_{\tilde{\eta}} + \frac{h}{g}(\tilde{\eta}_{k_1}) + \frac{h}{g}(\tilde{\eta}_{k_2}) \right) F_\chi^{[2]} - G_\chi^{[2]} = \frac{1}{2} \left(\alpha(\mathbf{k}_1, \mathbf{k}_2) \frac{h}{g}(\tilde{\eta}_{k_1}) + \alpha(\mathbf{k}_2, \mathbf{k}_1) \frac{h}{g}(\tilde{\eta}_{k_2}) \right), \quad (4.35a)$$

4. EFTofLSS and Mixed Dark Matter

$$\begin{aligned} \left(\partial_{\tilde{\eta}} + \frac{1}{2} + \frac{h}{g}(\tilde{\eta}_{k_1}) + \frac{h}{g}(\tilde{\eta}_{k_2}) \right) G_{\chi}^{[2]} + \frac{3}{2} \left(e^{-(\gamma-1)\tilde{\eta}} - f_{\chi} \right) F_{\chi}^{[2]} - \frac{3}{2} (1 - f_{\chi}) \frac{F_c^{[2]}}{g(\tilde{\eta}_{k_1})g(\tilde{\eta}_{k_2})} \\ = \beta(\mathbf{k}_1, \mathbf{k}_2) \frac{h}{g}(\tilde{\eta}_{k_1}) \frac{h}{g}(\tilde{\eta}_{k_2}), \end{aligned} \quad (4.35b)$$

$$\left(\partial_{\tilde{\eta}} + s(\tilde{\eta}_{k_1}) + s(\tilde{\eta}_{k_2}) \right) F_c^{[2]} - G_c^{[2]} = \frac{1}{2} \left(\alpha(\mathbf{k}_1, \mathbf{k}_2) s(\tilde{\eta}_{k_1}) + \alpha(\mathbf{k}_2, \mathbf{k}_1) s(\tilde{\eta}_{k_2}) \right), \quad (4.35c)$$

$$\begin{aligned} \left(\partial_{\tilde{\eta}} + \frac{1}{2} + s(\tilde{\eta}_{k_1}) + s(\tilde{\eta}_{k_2}) \right) G_c^{[2]} - \frac{3}{2} f_{\chi} F_{\chi}^{[2]} g(\tilde{\eta}_{k_1}) g(\tilde{\eta}_{k_2}) - \frac{3}{2} (1 - f_{\chi}) F_c^{[2]} \\ = \beta(\mathbf{k}_1, \mathbf{k}_2) s(\tilde{\eta}_{k_1}) s(\tilde{\eta}_{k_2}), \end{aligned} \quad (4.35d)$$

where the argument of the kernels is understood to be $(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta})$.

At $\mathcal{O}(f_{\chi}^0)$ we have $s^{(0)}(\tilde{\eta}) = 1$ and the equations for c are solved by the standard, $\tilde{\eta}$ -independent EdS kernels for all n ,

$$F_c^{(0)[n]}(\mathbf{k}_1, \dots, \mathbf{k}_n; \tilde{\eta}) = F_n(\mathbf{k}_1, \dots, \mathbf{k}_n), \quad G_c^{(0)[n]}(\mathbf{k}_1, \dots, \mathbf{k}_n; \tilde{\eta}) = G_n(\mathbf{k}_1, \dots, \mathbf{k}_n). \quad (4.36)$$

All the remaining pieces of the nonlinear solutions depend on the value of γ . In the rest of this section we focus on $\gamma = 2$, building upon the linear results presented in Sec. 4.2.2.

Analytical solutions for $\gamma = 2$

Although a general solution for the nonlinear kernels defined in Eq. (4.32) requires numerical integration, analytical expressions can be derived for kinematic configurations where all the momenta are well separated from the characteristic scale k_s . In that respect, three distinct physical regimes can be identified depending on the size of k_s relative to the observational window of galaxy surveys, which is roughly given by $k_{\text{eq}} \lesssim k \lesssim k_{\text{NL}}$ (with $k_{\text{eq}} \equiv \mathcal{H}_{\text{eq}}$):

- When $k_{\text{drop}} \gtrsim k_{\text{NL}}$ (and therefore $k_s(z_{\text{obs}}) \gg k_{\text{NL}}$), the χ and c species are indistinguishable on all scales probed by galaxy surveys, hence a standard description in terms of a single CDM fluid is adequate.
- When $k_s(z_{\text{obs}}) \lesssim k_{\text{eq}}$, the χ perturbations are strongly suppressed in the whole region of wavenumbers relevant to LSS data. Massive neutrinos with $m_{\nu} \lesssim 0.1$ eV are a well-known realization of this case, as Eq. (3.7) shows. A description in terms of a single fluid c is appropriate, however the presence of the non-clustering species χ affects the evolution of the c perturbations in a nontrivial way.
- When $k_{\text{eq}} < k_s(z_{\text{obs}})$ and $k_{\text{drop}} < k_{\text{NL}}$ are both satisfied,⁵ the characteristic scale leaves direct imprints in the fluctuations of both χ and c . A computationally efficient yet accurate modeling of the 1-loop galaxy power spectrum in this regime is the main goal of this paper.

We now present analytical results for the nonlinear matter kernels corresponding to each of the three regimes.

If all momenta are much smaller than $k_s(a)$, then the χ and c species are indistinguishable. This is a trivial limit where we just have a single fluid behaving as CDM, with kernels given by the standard F_n and G_n .

⁵For instance, in the case of ULAs this regime corresponds, recalling Eq. (4.30), to $k_{\text{eq}} < k_s(z_{\text{obs}}) < 7 k_{\text{NL}}$.

If all momenta are much larger than $k_s(a)$, the χ perturbations are entirely negligible and we only need to solve for the dynamics of the cold species c . In Appendix 4.A.1 we show that, perhaps surprisingly, in this limit the equations obeyed by the c perturbations are tightly related to the equations found in Refs. [193, 194] for a different type of new dynamics affecting the CDM+baryon fluid: DM self-interactions mediated by a massless (scalar) field. As a consequence, and as detailed in Appendix 4.A.1, solutions for this regime where all momenta are large are immediately obtained from the results of Ref. [194]. For instance, up to second perturbative order for $k_1, k_2, k \gg k_s(a)$, with $\mathbf{k} = \mathbf{k}_1 + \mathbf{k}_2$, we find

$$F_c^{[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) = F_2(\mathbf{k}_1, \mathbf{k}_2) + \frac{6f_\chi}{245} [\alpha_s(\mathbf{k}_1, \mathbf{k}_2) - \beta(\mathbf{k}_1, \mathbf{k}_2)] + \mathcal{O}(f_\chi^2), \quad (4.37a)$$

$$G_c^{[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) = \left(1 - \frac{3f_\chi}{5}\right) G_2(\mathbf{k}_1, \mathbf{k}_2) + \frac{12f_\chi}{245} [\alpha_s(\mathbf{k}_1, \mathbf{k}_2) - \beta(\mathbf{k}_1, \mathbf{k}_2)] + \mathcal{O}(f_\chi^2). \quad (4.37b)$$

These expressions can, of course, also be obtained by solving directly Eqs. (4.35) in the appropriate limit. The dominant beyond- Λ CDM imprint is the suppression of the linear growth rate $f(k)/f_{\Lambda\text{CDM}} \simeq 1 - 3f_\chi/5$, which results in the factor multiplying G_2 in Eq. (4.37b). Additional effects are strongly suppressed by accidentally small coefficients, such as $6/245 \approx 0.024$, which make them practically unobservable. This picture is consistent with the results obtained in Refs. [97, 98] for massive neutrino cosmologies.⁶

Finally, we turn to the regime where some momenta are large relative to $k_s(a)$, while others are small. This is the most relevant situation for our purposes, as well as the most complex one. At second perturbative order and $\mathcal{O}(f_\chi^0)$, Eqs. (4.35a) and (4.35b) yield in the $k_1 \ll k_s(a) \ll k_2$ configuration

$$F_\chi^{(0)[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) = F_2(\mathbf{k}_1, \mathbf{k}_2), \quad (4.38a)$$

$$G_\chi^{(0)[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) = 2G_2(\mathbf{k}_1, \mathbf{k}_2) - \frac{3}{14}\alpha_s(\mathbf{k}_1, \mathbf{k}_2) + \frac{1}{2}\alpha_a(\mathbf{k}_1, \mathbf{k}_2) - \frac{2}{7}\beta(\mathbf{k}_1, \mathbf{k}_2), \quad (4.38b)$$

where the symmetric component of α is defined in Eq. (2.18) and, interestingly, now the antisymmetric component

$$\alpha_a(\mathbf{k}_1, \mathbf{k}_2) \equiv \frac{1}{2} [\alpha(\mathbf{k}_1, \mathbf{k}_2) - \alpha(\mathbf{k}_2, \mathbf{k}_1)] \quad (4.39)$$

is not anymore forbidden in this scale-dependent scenario. For the cold species at $\mathcal{O}(f_\chi)$ we obtain

$$F_c^{(1)[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) = -\frac{69}{490}\alpha_s(\mathbf{k}_1, \mathbf{k}_2) + \frac{3}{14}\alpha_a(\mathbf{k}_1, \mathbf{k}_2) - \frac{18}{245}\beta(\mathbf{k}_1, \mathbf{k}_2), \quad (4.40a)$$

$$G_c^{(1)[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) = -\frac{3}{5}G_2(\mathbf{k}_1, \mathbf{k}_2) - \frac{15}{98}\alpha_s(\mathbf{k}_1, \mathbf{k}_2) + \frac{9}{70}\alpha_a(\mathbf{k}_1, \mathbf{k}_2) + \frac{6}{245}\beta(\mathbf{k}_1, \mathbf{k}_2). \quad (4.40b)$$

We have chosen to write the right-hand sides of Eqs. (4.38) and (4.40) so that the infrared (IR) divergence for $\mathbf{k}_1 \rightarrow 0$, where present, is fully captured by the term proportional to F_2 or G_2 . As mentioned, the appearance of the antisymmetric mode coupling function α_a , which is absent in Λ CDM, is enabled by the presence of the new scale $k_s(a)$. The source terms of the continuity equations (4.35a) and (4.35c), which must be symmetric under $\mathbf{k}_1 \leftrightarrow \mathbf{k}_2$ exchange, contain indeed products of α_a times antisymmetric combinations of the new functions of momentum h/g and s . For example, for the χ species one finds

$$\alpha_a(\mathbf{k}_1, \mathbf{k}_2) \left[\frac{h}{g}(\tilde{\eta}_{k_1}) - \frac{h}{g}(\tilde{\eta}_{k_2}) \right] \subset \alpha(\mathbf{k}_1, \mathbf{k}_2) \frac{h}{g}(\tilde{\eta}_{k_1}) + \alpha(\mathbf{k}_2, \mathbf{k}_1) \frac{h}{g}(\tilde{\eta}_{k_2}), \quad (4.41)$$

⁶Equations (4.37) also agree with Eq. (15) in Ref. [195], after correcting obvious typos in that reference.

which leads to the appearance of α_a in Eq. (4.38) after using the limits $h^{(0)}/g^{(0)} \rightarrow 1$ for $k_1 \ll k_s(a)$ and $h^{(0)}/g^{(0)} \rightarrow 2$ for $k_2 \gg k_s(a)$. It should be noted that this momentum coupling is not present in bootstrap approaches to LSS [196, 197], as the latter are constructed for scale-free theories where such a structure cannot appear.

The simplest observables where the new dynamics we derived above for the $k_1 \ll k_s(a) \ll k_2$ configuration will appear, are tree-level bispectra. However, since the LSS data analysis of Sec. 4.4 considers only the galaxy power spectrum, we will in the following primarily focus on this observable only.

Infrared structure

Here we discuss the IR behavior of the matter solutions. We begin with χ at $\mathcal{O}(f_\chi^0)$. Taking $\mathbf{k}_1 = \mathbf{q}$ and $\mathbf{k}_2 = \mathbf{k} - \mathbf{q}$ with $\mathbf{q} \rightarrow 0$, the second-order equations (4.35a) and (4.35b) are solved by

$$F_\chi^{(0)[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta}) \stackrel{\mathbf{q} \rightarrow 0}{=} F_2(\mathbf{q}, \mathbf{k} - \mathbf{q}), \quad G_\chi^{(0)[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta}) \stackrel{\mathbf{q} \rightarrow 0}{=} \frac{h^{(0)}}{g^{(0)}}(\tilde{\eta}) G_2(\mathbf{q}, \mathbf{k} - \mathbf{q}), \quad (4.42)$$

which should be understood as equalities of the $1/q$ poles.⁷ They were derived by setting $g^{(0)}(\tilde{\eta}_q), h^{(0)}(\tilde{\eta}_q) \stackrel{\mathbf{q} \rightarrow 0}{\rightarrow} 1$ and employing the linear equation (4.16) for $g^{(0)}(\tilde{\eta})$. Hence, at $\mathcal{O}(f_\chi^0)$ the IR divergence in $F_\chi^{[2]}$ is the same as in a standard EdS cosmology. By contrast, the residual of the IR pole in $G_\chi^{[2]}$ is time dependent and proportional to the χ linear growth rate h/g . A similar calculation at the third perturbative order results in the equalities

$$F_\chi^{(0)[3]}(\mathbf{k}, \mathbf{q}, -\mathbf{q}; \tilde{\eta}) \stackrel{\mathbf{q} \rightarrow 0}{=} F_3(\mathbf{k}, \mathbf{q}, -\mathbf{q}), \quad G_\chi^{(0)[3]}(\mathbf{k}, \mathbf{q}, -\mathbf{q}; \tilde{\eta}) \stackrel{\mathbf{q} \rightarrow 0}{=} \frac{h^{(0)}}{g^{(0)}}(\tilde{\eta}) G_3(\mathbf{k}, \mathbf{q}, -\mathbf{q}), \quad (4.43)$$

which hold for the $1/q^2$ poles.

An analogous IR structure is observed in the solutions for the cold species. At $\mathcal{O}(f_\chi)$, the second-order equations (4.35c) and (4.35d) are solved by

$$F_c^{(1)[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta}) \stackrel{\mathbf{q} \rightarrow 0}{=} 0, \quad G_c^{(1)[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta}) \stackrel{\mathbf{q} \rightarrow 0}{=} s^{(1)}(\tilde{\eta}) G_2(\mathbf{q}, \mathbf{k} - \mathbf{q}), \quad (4.44)$$

obtained by setting $s^{(1)}(\tilde{\eta}_q) \stackrel{\mathbf{q} \rightarrow 0}{\rightarrow} 0$ and $g^{(0)}(\tilde{\eta}_q) \stackrel{\mathbf{q} \rightarrow 0}{\rightarrow} 1$ and making use of the linear equation (4.21) for $s^{(1)}(\tilde{\eta})$. Therefore, up to $\mathcal{O}(f_\chi)$ the IR divergence of $F_c^{[2]}$ is not modified relative to standard cosmology, whereas the residual of the IR pole in $G_c^{[2]}$ is proportional to the linear growth rate s . Similarly, at third order we find

$$F_c^{(1)[3]}(\mathbf{k}, \mathbf{q}, -\mathbf{q}; \tilde{\eta}) \stackrel{\mathbf{q} \rightarrow 0}{=} 0, \quad G_c^{(1)[3]}(\mathbf{k}, \mathbf{q}, -\mathbf{q}; \tilde{\eta}) \stackrel{\mathbf{q} \rightarrow 0}{=} s^{(1)}(\tilde{\eta}) G_3(\mathbf{k}, \mathbf{q}, -\mathbf{q}). \quad (4.45)$$

It is easy to guess, then, that Eqs. (4.42)–(4.45) generalize to all orders in f_χ , namely

$$F_\chi^{[n]} \stackrel{\mathbf{q} \rightarrow 0}{=} F_n, \quad G_\chi^{[n]} \stackrel{\mathbf{q} \rightarrow 0}{=} \frac{h}{g}(\tilde{\eta}) G_n, \quad F_c^{[n]} \stackrel{\mathbf{q} \rightarrow 0}{=} F_n, \quad G_c^{[n]} \stackrel{\mathbf{q} \rightarrow 0}{=} s(\tilde{\eta}) G_n, \quad (4.46)$$

where the arguments of the kernels are understood to be $(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta})$ for $n = 2$ and $(\mathbf{k}, \mathbf{q}, -\mathbf{q}; \tilde{\eta})$ for $n = 3$.

⁷Equation (4.42) is found to be consistent with Eq. (4.38) after recalling that in the former $h^{(0)}/g^{(0)} \rightarrow 2$ for $k \gg k_s(a)$, while in the latter only the terms proportional to F_2 or G_2 are IR divergent.

Equation (4.46) guarantees, at all orders in f_χ , the absence of IR divergences in all 1-loop power spectra, as a direct consequence of their well-known cancellation in Λ CDM. For instance, the 1-loop correction to the δ_c - Θ_χ cross power spectrum reads

$$P_{\delta_c \Theta_\chi}^{1\text{-loop}}(k) = 4 \int \frac{d^3 q}{(2\pi)^3} F_c^{[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}) G_\chi^{[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}) P_{\delta_c \delta_\chi}^L(q) P_{\delta_c \delta_\chi}^L(|\mathbf{k} - \mathbf{q}|) \Theta_H(|\mathbf{k} - \mathbf{q}| - q) \\ + 3 P_{\delta_c \Theta_\chi}^L(k) \int \frac{d^3 q}{(2\pi)^3} F_c^{[3]}(\mathbf{k}, \mathbf{q}, -\mathbf{q}) P_{\delta_c \delta_c}^L(q) + 3 P_{\delta_c \delta_\chi}^L(k) \int \frac{d^3 q}{(2\pi)^3} G_\chi^{[3]}(\mathbf{k}, \mathbf{q}, -\mathbf{q}) P_{\delta_\chi \delta_\chi}^L(q), \quad (4.47)$$

where the dependence on the time variable $\tilde{\eta}$ is understood. Exploiting Eq. (4.46), it is immediate to verify that in the $\mathbf{q} \rightarrow 0$ limit the integrand in $P_{\delta_c \Theta_\chi}^{1\text{-loop}}$ is proportional to

$$4F_2(\mathbf{q}, \mathbf{k} - \mathbf{q}) G_2(\mathbf{q}, \mathbf{k} - \mathbf{q}) + 3F_3(\mathbf{k}, \mathbf{q}, -\mathbf{q}) + 3G_3(\mathbf{k}, \mathbf{q}, -\mathbf{q}), \quad (4.48)$$

where divergences cancel exactly.

Furthermore, Eq. (4.46) ensures that all tree-level bispectra are free from IR poles. For example, consider

$$B_{\delta_c \delta_c \Theta_\chi}(\mathbf{q}, \mathbf{p}_1, \mathbf{p}_2) = 2 \left[F_c^{[2]}(\mathbf{p}_1, \mathbf{p}_2; \tilde{\eta}_q) P_{\delta_c \delta_c}^L(p_1) P_{\delta_c \Theta_\chi}^L(p_2) + F_c^{[2]}(\mathbf{q}, \mathbf{p}_2; \tilde{\eta}_{p_1}) P_{\delta_c \delta_c}^L(q) P_{\delta_c \Theta_\chi}^L(p_2) \right. \\ \left. + G_\chi^{[2]}(\mathbf{q}, \mathbf{p}_1; \tilde{\eta}_{p_2}) P_{\delta_c \delta_\chi}^L(q) P_{\delta_c \delta_\chi}^L(p_1) \right]. \quad (4.49)$$

Making use of Eq. (4.46) one easily verifies that $B_{\delta_c \delta_c \Theta_\chi}$ is proportional to $F_2(\mathbf{q}, -\mathbf{p}_1) + G_2(\mathbf{q}, \mathbf{p}_1)$ in the $\mathbf{q} \rightarrow 0$ limit, hence the IR divergence cancels.

Ultraviolet structure

The sensitivity of the 1-loop matter power spectra to UV modes is dominated by the P^{13} contributions, which involve the third-order kernels. For the χ species we find the following behavior at $\mathcal{O}(f_\chi^0)$,

$$F_\chi^{(0)[3]}(\mathbf{k}, \mathbf{q}, -\mathbf{q}; \tilde{\eta}) \stackrel{k_s(a) \ll q}{\simeq} \frac{u^{(0)}(\tilde{\eta}) F_3(\mathbf{k}, \mathbf{q}, -\mathbf{q})}{g^{(0)}(\tilde{\eta}) g^{(0)}(\tilde{\eta}_q)^2}, \quad (4.50a)$$

$$G_\chi^{(0)[3]}(\mathbf{k}, \mathbf{q}, -\mathbf{q}; \tilde{\eta}) \stackrel{k_s(a) \ll q}{\simeq} \frac{(\partial_{\tilde{\eta}} u^{(0)} + 3u^{(0)})(\tilde{\eta}) F_3(\mathbf{k}, \mathbf{q}, -\mathbf{q})}{g^{(0)}(\tilde{\eta}) g^{(0)}(\tilde{\eta}_q)^2}, \quad (4.50b)$$

where

$$u^{(0)}(\tilde{\eta}) = U^{(0)}(\sqrt{6}e^{-\tilde{\eta}/2}), \quad \text{with} \quad U^{(0)}(x) = \frac{1}{7} - \frac{x^2}{140} + \frac{x^4}{840} G^{(0)}(x). \quad (4.51)$$

Accordingly, the 1-loop density power spectra involving the χ species scale as

$$P_{\delta_c \delta_\chi}^{(0)13}(k) \xrightarrow{\text{UV}} -\frac{61}{630\pi^2} \frac{1}{2} \left(1 + \frac{u^{(0)}}{g^{(0)}}(\tilde{\eta}) \right) k^2 P_{\delta_c \delta_\chi}^{(0)L}(k) \int^{\Lambda_{\text{UV}}} dq P_{\delta_c \delta_c}^{(0)L}(q), \quad (4.52)$$

$$P_{\delta_\chi \delta_\chi}^{(0)13}(k) \xrightarrow{\text{UV}} -\frac{61}{630\pi^2} \frac{u^{(0)}}{g^{(0)}}(\tilde{\eta}) k^2 P_{\delta_\chi \delta_\chi}^{(0)L}(k) \int^{\Lambda_{\text{UV}}} dq P_{\delta_c \delta_c}^{(0)L}(q). \quad (4.53)$$

Thus, the standard UV dependence $\sim k^2 P^L(k)$ of Eq. (2.23) is further modulated, already at $\mathcal{O}(f_\chi^0)$, by the ratio $u^{(0)}/g^{(0)}$, which smoothly transitions from $1/7$ at $k \ll k_s(a)$ to 1 at $k \gg k_s(a)$.

For the cold species, we can evaluate analytically the UV behavior at $\mathcal{O}(f_\chi)$ when the external wavenumber is well separated from $k_s(a)$,

$$F_c^{(1)[3]}(\mathbf{k}, \mathbf{q}, -\mathbf{q}) \xrightarrow{k \ll k_s(a) \ll q} \frac{1}{1890} (-116\bar{\mu}^4 + 397\bar{\mu}^2 - 71) \left(\frac{k}{q} \right)^2 + \mathcal{O}((k/q)^4), \quad (4.54a)$$




Figure 4.5.: Coefficient of the UV divergence for the $P_{\delta_c \delta_c}^{13}(k)$ loop integral, as a function of $k/k_s(a)$, normalized to the standard Λ CDM value. The solid blue line shows the numerical result for $p = 0$, obtained by setting $f_\chi = 0.1$. The two dotted lines show the analytical limits for $k \ll k_s(a)$ and $k \gg k_s(a)$, calculated using Eqs. (4.54). The dashed purple line corresponds to the analytical interpolation given within curly brackets in Eq. (4.55). The residual discrepancies observed between numerical and analytical results are due to terms of $\mathcal{O}(f_\chi^2) \sim 1\%$, which were neglected in the analytical expansion for small f_χ .

$$F_c^{(1)[3]}(\mathbf{k}, \mathbf{q}, -\mathbf{q}) \xrightarrow{k_s(a) \ll k \ll q} \frac{4}{6615} (49\bar{\mu}^4 - 71\bar{\mu}^2 + 22) \left(\frac{k}{q}\right)^2 + \mathcal{O}((k/q)^4), \quad (4.54b)$$

where $\bar{\mu} \equiv \mathbf{k} \cdot \mathbf{q}/(kq)$. As a consequence, at $\mathcal{O}(f_\chi)$ the UV limit of the 1-loop power spectrum of δ_c deviates from the standard $k^2 P_{\delta_c \delta_c}^L(k)$ dependence. Interpolating between the two limits in Eq. (4.54) gives the approximate expression

$$P_{\delta_c \delta_c}^{13}(k) \xrightarrow{\text{UV}} -\frac{61}{630\pi^2} \left\{ 1 - f_\chi \left[\frac{16}{105} + \frac{3028}{6405} g^{(0)} \left(- (2+p) \log \frac{k}{k_s(a)} \right) \right] \right\} \times k^2 P_{\delta_c \delta_c}^L(k) \int^{\Lambda_{\text{UV}}} dq P_{\delta_c \delta_c}^L(q), \quad (4.55)$$

which indeed agrees with the numerical solution up to tiny $\mathcal{O}(f_\chi^2)$ corrections, as shown in Fig. 4.5. In principle, then, the standard $k^2 P_{\delta_c \delta_c}^L(k)$ counterterm should be modified to account for this new UV sensitivity at first order in f_χ .

Finally, analogous UV structures are found for the power spectra involving the velocity divergences: for instance,

$$P_{\delta_c \Theta_\chi}^{(0)13}(k) \xrightarrow{\text{UV}} -\frac{61}{630\pi^2} \frac{1}{2} \left(1 + \frac{\partial_{\tilde{\eta}} u^{(0)} + 3u^{(0)}}{h^{(0)}}(\tilde{\eta}) \right) k^2 P_{\delta_c \Theta_\chi}^{(0)L}(k) \int^{\Lambda_{\text{UV}}} dq P_{\delta_c \delta_c}^{(0)L}(q), \quad (4.56)$$

where $(\partial_{\tilde{\eta}} u^{(0)} + 3u^{(0)})/h^{(0)}$ smoothly transitions from $3/7$ at $k \ll k_s(a)$ to 2 at $k \gg k_s(a)$.

4.3.2. Bias expansion with two fluids

As dealing with a two underlying fluid scenarios, with generic suppression scale, the Λ CDM treatment of galaxy bias of Sect. 2.1.2, as well as the special case of neutrinos with very low, allowed, mass of

Sect. 3.1 given by Eq. (3.13), has to be generalized. The response of galaxy (and halo) formation to the presence of multiple fluid species has been studied extensively in the literature [58, 83, 86, 91, 166, 194, 198, 199]. The bias expansion for two fluids is naturally expressed in terms of the total matter (m) and relative (r) perturbations, defined in Eq. (4.3). However, in this study we limit a full description of bias to linear operators, for which the map between the (m, r) and (χ, c) field bases is straightforward. Therefore we prefer to work with the latter basis, which we have used thus far to express our solutions.

The galaxy number density perturbation $\delta_g \equiv n_g/\bar{n}_g - 1$ admits the bias expansion [166, 194, 198]

$$\delta_g = b_c \delta_c + b_\chi \delta_\chi + b_\Theta (\Theta_\chi - \Theta_c) + \dots, \quad (4.57)$$

where the dots stand for nonlinear operators. What is the expected size of b_χ ? A simple argument indicates that b_χ must vanish as $f_\chi \rightarrow 0$. If this were not the case, the galaxy perturbation would experience a finite jump from $\delta_g = (b_c + b_\chi)\delta_c$ to $\delta_g = b_c\delta_c$ around $k_s(a)$ despite the χ energy density being negligible, which is unphysical. Furthermore, studies of the scale dependence of bias in the presence of massive neutrinos [58, 83, 86] indicate that b_ν is of order f_ν and has negative sign, $b_\nu = -f_\nu \hat{b}_\nu$ with $0 < \hat{b}_\nu < b_c$.

However, in general it is not guaranteed that an expansion of b_χ for $f_\chi \ll 1$ is meaningful when f_χ is as large as $\mathcal{O}(0.1)$, much bigger than the realistic values for the neutrino fraction f_ν , as it will be considered in this work. For this reason, as our baseline choice we conservatively treat b_χ as an $\mathcal{O}(1)$ parameter and similarly for b_Θ . Nonetheless, in Sec. 4.4 we will also discuss the impact of tightening the prior on b_χ to $\mathcal{O}(f_\chi)$ size. Under our baseline assumptions, a formally consistent analytical calculation of galaxy correlators would require one to solve for the χ perturbations up to $\mathcal{O}(f_\chi)$, which is beyond the scope of this study. In practice, however, we expect the key physical effects to be already captured by the leading non-trivial order, which for χ is $\mathcal{O}(f_\chi^0)$ and which we have retained throughout. Expanding the galaxy density perturbation as

$$\delta_g^{[n]}(\mathbf{k}, \tilde{\eta}) = \int_{\mathbf{k}} dk_{1\dots n} F_g^{[n]}(\mathbf{k}_1, \dots, \mathbf{k}_n; \tilde{\eta}) \delta_c^{[1]}(\mathbf{k}_1, \tilde{\eta}) \dots \delta_c^{[1]}(\mathbf{k}_n, \tilde{\eta}), \quad (4.58)$$

we find at linear level (recall that $s(\tilde{\eta}) = f(k)/f_{\Lambda\text{CDM}}$)

$$F_g^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) = b_c + b_\chi g(\tilde{\eta}_{k_1}) + b_\Theta \left[h(\tilde{\eta}_{k_1}) - \frac{f(k_1)}{f_{\Lambda\text{CDM}}} \right]. \quad (4.59)$$

The $\mathcal{O}(f_\chi^0)$ results $s = 1$ and $h \sim g$, the latter meaning that the two functions have similar qualitative behaviors (see the right panel of Fig. 4.1), suggest that the effect of b_Θ can be largely absorbed into redefinitions of b_c and b_χ . This is confirmed by Fig. 4.6, which shows that the impact of a non-zero b_Θ on the linear galaxy power spectrum can be approximately captured by shifting $b_c \rightarrow b_c - b_\Theta$ and $b_\chi \rightarrow b_\chi + b_\Theta$.

Nonlinear bias operators can be systematically included in Eq. (4.57) by applying the general results of Ref. [194] (which built upon Refs. [166, 198]), and are constructed using the perturbations of both species. However, given that the impact of the nonlinear bias coefficients on the extraction of cosmological parameters is, in general, numerically smaller than for linear bias, here we adopt a simplified approach where only nonlinear bias operators containing the cold species alone are included. The prescription for kernels presented in Sec. 4.3.4 below also implies that nonlinear bias operators involving δ_χ and v_χ^i can be viewed as scale-dependent corrections to the already largely uncertain nonlinear bias parameters with respect to the CDM plus baryon fluid.




Figure 4.6.: The ratio $P_{\delta_g \delta_g}^L / P_{\delta_c \delta_c}^L$ calculated at linear order from the bias expansion in Eq. (4.57), as a function of $\tilde{\eta}$. Solid lines correspond to a set of bias coefficients (b_c, b_χ, b_Θ) . Dotted lines correspond to a different set, where $b_\Theta = 0$ but b_c and b_χ have been appropriately shifted, which reproduces well the shape of the galaxy power spectrum. The small residual difference depends on the size of b_Θ .

4.3.3. Redshift space distortions

Our next step is to perform the mapping between real and redshift space. To derive the redshift-space expression of the galaxy number density perturbation, which we denote $\delta_{g,s}$, we also require the bias expansion of the galaxy velocity v_g^i ,

$$v_g^i = (1 - b_v)v_c^i + b_v v_\chi^i + \dots, \quad (4.60)$$

where the dots indicate again nonlinear operators. The presence of a single linear velocity bias coefficient b_v is fixed by invariance under generalized Galilean transformations, which boost all velocities by the same (spatially constant) amount.

In Chapter 3 we presented indeed a study, involving a comparison with N-body simulations, that considered the extreme cases $b_v = 0$ and $b_v = f_\nu$ in the context of massive neutrinos, which, as mentioned, represent a limiting case of the present study. Here, instead, we aim at a more general treatment, leaving b_v in principle free, while limiting ourselves to purely theoretical considerations.

Expanding the divergence of the galaxy velocity in analogy with Eq. (4.58), we write

$$\Theta_g^{[n]}(\mathbf{k}, \tilde{\eta}) = \int_{\mathbf{k}} dk_{1\dots n} G_g^{[n]}(\mathbf{k}_1, \dots, \mathbf{k}_n; \tilde{\eta}) \delta_c^{[1]}(\mathbf{k}_1, \tilde{\eta}) \dots \delta_c^{[1]}(\mathbf{k}_n, \tilde{\eta}). \quad (4.61)$$

At linear level we obtain

$$G_g^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) = (1 - b_v) \frac{f(k_1)}{f_{\Lambda\text{CDM}}} + b_v h(\tilde{\eta}_{k_1}), \quad (4.62)$$

and by comparison with Eq. (4.59) we expect the effect of b_v to be largely degenerate with the one of b_Θ . We do not consider any nonlinear velocity bias operators in this work, since their effects are expected to be very suppressed. However, if desired they could be included in Eq. (4.60) by exploiting previous results [194, 198].

Defining the expansion of $\delta_{g,s}$ as in Eq. (4.58), but with redshift space galaxy kernels as introduced in Sect. 2.1.2, we obtain up to third perturbative order

$$F_{g,s}^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) = F_g^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) + f_{\Lambda\text{CDM}} \mu_1^2 G_g^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}), \quad (4.63a)$$

$$\begin{aligned} F_{g,s}^{[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) &= F_g^{[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) + f_{\Lambda\text{CDM}} \mu^2 G_g^{[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) \\ &\quad + f_{\Lambda\text{CDM}} k \mu \left[\frac{\mu_1}{2k_1} F_g^{[1]}(\mathbf{k}_2; \tilde{\eta}_{k_2}) G_g^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) + \mathbf{k}_1 \leftrightarrow \mathbf{k}_2 \right] \\ &\quad + f_{\Lambda\text{CDM}}^2 k^2 \mu^2 \frac{\mu_1 \mu_2}{2k_1 k_2} G_g^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) G_g^{[1]}(\mathbf{k}_2; \tilde{\eta}_{k_2}), \end{aligned} \quad (4.63b)$$

$$\begin{aligned} F_{g,s}^{[3]}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3; \tilde{\eta}) &= F_g^{[3]}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3; \tilde{\eta}) + f_{\Lambda\text{CDM}} \mu^2 G_g^{[3]}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3; \tilde{\eta}) \\ &\quad + \frac{1}{3} \left\{ f_{\Lambda\text{CDM}} k \mu \left[\frac{\mu_1}{k_1} G_g^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) F_g^{[2]}(\mathbf{k}_2, \mathbf{k}_3; \tilde{\eta}_{k_{23}}) + \frac{\mu_{23}}{k_{23}} F_g^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) G_g^{[2]}(\mathbf{k}_2, \mathbf{k}_3; \tilde{\eta}_{k_{23}}) \right] \right. \\ &\quad + f_{\Lambda\text{CDM}}^2 k^2 \mu^2 \left[\frac{\mu_1}{k_1} \frac{\mu_{23}}{k_{23}} G_g^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) G_g^{[2]}(\mathbf{k}_2, \mathbf{k}_3; \tilde{\eta}_{k_{23}}) \right. \\ &\quad \left. \left. + \frac{\mu_2 \mu_3}{2k_2 k_3} F_g^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) G_g^{[1]}(\mathbf{k}_2; \tilde{\eta}_{k_2}) G_g^{[1]}(\mathbf{k}_3; \tilde{\eta}_{k_3}) \right] \right\} + \mathbf{k}_1 \leftrightarrow \mathbf{k}_2 + \mathbf{k}_1 \leftrightarrow \mathbf{k}_3 \\ &\quad + f_{\Lambda\text{CDM}}^3 k^3 \mu^3 \frac{\mu_1 \mu_2 \mu_3}{6k_1 k_2 k_3} G_g^{[1]}(\mathbf{k}_1; \tilde{\eta}_{k_1}) G_g^{[1]}(\mathbf{k}_2; \tilde{\eta}_{k_2}) G_g^{[1]}(\mathbf{k}_3; \tilde{\eta}_{k_3}), \end{aligned} \quad (4.63c)$$

where $\mu \equiv \mathbf{k} \cdot \hat{\mathbf{z}}/k$ with $\mathbf{k} = \mathbf{k}_1 + \dots + \mathbf{k}_n$, $\mu_{i_1 \dots i_m} \equiv \mathbf{k}_{i_1 \dots i_m} \cdot \hat{\mathbf{z}}/k_{i_1 \dots i_m}$ with $\mathbf{k}_{i_1 \dots i_m} = \mathbf{k}_{i_1} + \dots + \mathbf{k}_{i_m}$ and $\hat{\mathbf{z}}$ is the unit vector along the line of sight, which we take as our z axis. In the ΛCDM limit, Eqs. (4.63) reproduce standard results [18, 19]. The 1-loop contribution to the redshift-space galaxy power spectrum is therefore

$$\begin{aligned} P_{\delta_g \delta_{g,s}}^{1\text{-loop}}(k, \mu) &= 2 \int \frac{d^3 q}{(2\pi)^3} [F_{g,s}^{[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta})]^2 P_{\delta_c \delta_c}^L(q) P_{\delta_c \delta_c}^L(|\mathbf{k} - \mathbf{q}|) \\ &\quad + 6 F_{g,s}^{[1]}(\mathbf{k}; \tilde{\eta}) P_{\delta_c \delta_c}^L(k) \int \frac{d^3 q}{(2\pi)^3} F_{g,s}^{[3]}(\mathbf{k}, \mathbf{q}, -\mathbf{q}; \tilde{\eta}) P_{\delta_c \delta_c}^L(q). \end{aligned} \quad (4.64)$$

4.3.4. Prescription for nonlinear kernels

The analytical understanding of the matter kernels we obtained in Sec. 4.3.1 by means of the small f_χ expansion, allows us to design a prescription that strongly simplifies the evaluation of $P_{\delta_g \delta_{g,s}}(k, \mu)$ up to 1-loop order, while retaining sufficient accuracy. The prescription is obtained by imposing that all IR limits in Eq. (4.46) are reproduced exactly, leading us to define

$$F_\chi^{[n]}(\mathbf{k}_1, \dots, \mathbf{k}_n; \tilde{\eta})_{\text{p}} = \frac{g(\tilde{\eta}) F_n(\mathbf{k}_1, \dots, \mathbf{k}_n)}{g(\tilde{\eta}_{k_1}) \dots g(\tilde{\eta}_{k_n})}, \quad (4.65a)$$

$$G_\chi^{[n]}(\mathbf{k}_1, \dots, \mathbf{k}_n; \tilde{\eta})_{\text{p}} = \frac{h(\tilde{\eta}) G_n(\mathbf{k}_1, \dots, \mathbf{k}_n)}{g(\tilde{\eta}_{k_1}) \dots g(\tilde{\eta}_{k_n})}, \quad (4.65b)$$

$$F_c^{[n]}(\mathbf{k}_1, \dots, \mathbf{k}_n; \tilde{\eta})_{\text{p}} = F_n(\mathbf{k}_1, \dots, \mathbf{k}_n), \quad (4.65c)$$

$$G_c^{[n]}(\mathbf{k}_1, \dots, \mathbf{k}_n; \tilde{\eta})_{\text{p}} = s(\tilde{\eta}) G_n(\mathbf{k}_1, \dots, \mathbf{k}_n). \quad (4.65d)$$

We will denote any quantities evaluated according to this prescription with a “p” subscript.

We see that the standard ΛCDM kernels F_n and G_n are the only nonlinear quantities appearing on the right-hand sides of Eqs. (4.65). They multiply the ratios of linear transfer functions g , h and s evaluated on the *external* momentum $\mathbf{k} = \mathbf{k}_1 + \dots + \mathbf{k}_n$, whereas the products of g functions in the denominators of Eqs. (4.65a) and (4.65b) simply rephrase the perturbative expansions of Eqs. (4.32) in

terms of only one linear field, namely $\delta_c^{[1]}$. As a consequence of these properties, the application of the prescription casts the 1-loop corrections to the matter power spectra into a simple factorized form. For instance, Eq. (4.47) becomes

$$P_{\delta_c \Theta_\chi}^{1\text{-loop}}(k)_p = h(\tilde{\eta}) \left\{ 4 \int \frac{d^3 q}{(2\pi)^3} F_2(\mathbf{q}, \mathbf{k} - \mathbf{q}) G_2(\mathbf{q}, \mathbf{k} - \mathbf{q}) P_{\delta_c \delta_c}^L(q) P_{\delta_c \delta_c}^L(|\mathbf{k} - \mathbf{q}|) \right. \\ \left. \times \Theta_H(|\mathbf{k} - \mathbf{q}| - q) + 3 P_{\delta_c \delta_c}^L(k) \int \frac{d^3 q}{(2\pi)^3} [F_3(\mathbf{k}, \mathbf{q}, -\mathbf{q}) + G_3(\mathbf{k}, \mathbf{q}, -\mathbf{q})] P_{\delta_c \delta_c}^L(q) \right\}, \quad (4.66)$$

where the quantity in curly brackets is simply the standard form of the 1-loop correction calculated with Λ CDM kernels, taking as input the linear power spectrum for the cold species, $P_{\delta_c \delta_c}^L(k)$, as obtained from CLASS. Standard methods developed for Λ CDM then enable a fast evaluation of Eq. (4.66) suitable for MCMC analysis.

One immediate consequence of the prescription is that, at all perturbative orders, the χ fields are obtained from the c fields through multiplication by ratios of linear transfer functions, namely

$$\delta_\chi(\mathbf{k}, \tilde{\eta})_p = g(\tilde{\eta}) \delta_c(\mathbf{k}, \tilde{\eta})_p, \quad \Theta_\chi(\mathbf{k}, \tilde{\eta})_p = \frac{h}{s}(\tilde{\eta}) \Theta_c(\mathbf{k}, \tilde{\eta})_p. \quad (4.67)$$

The first relation in Eq. (4.67) was indeed assumed in some previous studies of massive neutrino cosmologies [97, 158, 162].

In the rest of this section we thoroughly analyze the accuracy and implications of our prescription in Eqs. (4.65). Before doing so, however, we stress again that this prescription is motivated by the need to efficiently compute the galaxy power spectrum at 1-loop order. Indeed, since analytical expressions for the kernels covering the entire kinematic regime where the loops have support are not available, we are naturally prompted to adopt the above ansatz in order to manageably perform the required hundreds of thousands of evaluations. Nevertheless, the analytical expressions we presented in Sec. 4.3.1 (which were based on the small f_χ expansion, but did not rely on any ansätze) may be used to compute other observables exactly, for instance the tree-level galaxy bispectrum in some kinematic configurations, as argued in Sec. 4.3.1.

Accuracy for kernels

While the prescription reproduces exactly, by construction, the behavior of the complete solutions in the IR limit, it does introduce inaccuracies for general configurations of the momenta. Nevertheless, it preserves the correct structure of the kernels even far away from the IR, modulo momentum-dependent corrections that are a priori of $\mathcal{O}(f_\chi^0)$ for the χ fields and $\mathcal{O}(f_\chi)$ for the c fields.

In Fig. 4.7 we compare numerical and analytical solutions for the second order kernels of the cold fluid in the $(\mathbf{q}, \mathbf{k} - \mathbf{q})$ configuration of momenta relevant for the P^{22} contribution to the 1-loop power spectrum. Both numerical and analytical solutions are normalized to the value prescribed by Eqs. (4.65c) or (4.65d). We set $f_\chi = 0.1$ and evaluate the numerical results for $p = 0$ (as well as $\gamma = 2$, as we always assumed from Sec. 4.3.1 onwards). We choose $\mathbf{k}$ and $\mathbf{q}$ to be orthogonal in order to remove the IR pole at small q (by construction, the IR pole is exactly reproduced by our prescription). The figure shows two conceptually distinct aspects. (i) First, we observe that the numerical solutions (solid curves) are in good agreement with the analytical approximations (dashed lines) when all momenta are separated from $k_s(a)$. The small residual differences are due to corrections of $\mathcal{O}(f_\chi^2)$ that we have not calculated. This validates the analytical results presented in Sec. 4.3.1. (ii) Second, the deviations of the solid curves




Figure 4.7.: Comparison of numerical solutions and analytical limits for the second order kernels of the cold species, $F_c^{[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta})$ and $G_c^{[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta})$, normalized to the respective values $F_{c,p}^{[2]}$ and $G_{c,p}^{[2]}$ defined by the prescription in Eqs. (4.65). We set $f_\chi = 0.1$ and choose $\bar{\mu} = \mathbf{k} \cdot \mathbf{q}/(kq) = 0$ to remove the $1/q$ IR pole, which is exactly reproduced by the prescription. Solid curves correspond to the numerical solution for $p = 0$, while dashed lines show analytical limits. *Left:* $k \ll k_s(a)$. The $q \ll k_s(a)$ limits are trivial, recovering Λ CDM. The $q \gg k_s(a)$ analytical limits are derived from Eqs. (4.86) in Appendix 4.A.2. *Right:* $k \gg k_s(a)$. The $q \ll k_s(a)$ [$q \gg k_s(a)$] analytical limits are obtained from Eqs. (4.40) [Eqs. (4.37)].

from unity show that our prescription does indeed introduce errors of $\mathcal{O}(f_\chi)$ with respect to the exact numerical result, as expected. This is the price we accept to pay, in exchange for the fast evaluation of the 1-loop corrections illustrated by Eq. (4.66).

In Fig. 4.8 we perform an analogous comparison, but for the second order kernels of the warm species χ . Here we observe that: (i) numerical solutions (solid curves) and analytical limits (dashed lines) agree, modulo $\mathcal{O}(f_\chi)$ corrections which we have not calculated, thus confirming the results of Sec. 4.3.1; and (ii) the prescription introduces $\mathcal{O}(f_\chi^0)$ errors with respect to the exact numerical result in some regions of momenta, as the deviations of the solid curves from unity show. The deviations are especially large for $G_\chi^{[2]}$ (solid purple), whereas the prescription provides a better approximation to $F_\chi^{[2]}$ (solid blue).

Ultraviolet behavior and counterterms

Next, we focus on the behavior of the prescription in the UV limit. For the 1-loop power spectra of matter density perturbations we find

$$\{P_{\delta_c \delta_\chi}^{13}(k), P_{\delta_\chi \delta_\chi}^{13}(k), P_{\delta_c \delta_c}^{13}(k)\}_p \xrightarrow{\text{UV}} -\frac{61k^2}{630\pi^2} \{P_{\delta_c \delta_\chi}^L(k), P_{\delta_\chi \delta_\chi}^L(k), P_{\delta_c \delta_c}^L(k)\} \int^{\Lambda_{\text{UV}}} dq P_{\delta_c \delta_c}^L(q). \quad (4.68)$$

By comparing these to the exact perturbative results for the UV limits in Eqs. (4.52), (4.53) and (4.55), respectively, we see that the prescription does not capture the additional scale-dependent effects seen




Figure 4.8.: Comparison of numerical solutions and analytical limits for the second order kernels of the warm species, $F_\chi^{[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta})$ and $G_\chi^{[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta})$, normalized to the respective values $F_{\chi, \text{p}}^{[2]}$ and $G_{\chi, \text{p}}^{[2]}$ defined by the prescription in Eqs. (4.65). We set $f_\chi = 0.1$ and choose $\bar{\mu} = \mathbf{k} \cdot \mathbf{q}/(kq) = 0$ to remove the $1/q$ IR pole, which is exactly reproduced by the prescription. Solid curves correspond to the numerical solution for $p = 0$, while dashed lines show analytical limits. *Left:* $k \ll k_s(a)$. The $q \ll k_s(a)$ limits are trivial, recovering Λ CDM. The $q \gg k_s(a)$ analytical limits are derived from Eqs. (4.84) in Appendix 4.A.2. *Right:* $k \gg k_s(a)$. The $q \ll k_s(a)$ [$q \gg k_s(a)$] analytical limits are obtained from Eqs. (4.38) [Eqs. (4.84) in Appendix 4.A.2].

in those equations. Nevertheless, we expect these additional scale-dependent effects to be small for two separate reasons. For power spectra involving the χ fields, where the additional scale dependence formally arises at $\mathcal{O}(f_\chi^0)$, because the χ transfer functions are suppressed for $k \gtrsim k_s(a)$. For power spectra involving only the CDM+baryon perturbations, because the additional scale dependence appears only at $\mathcal{O}(f_\chi)$.

Owing to the presence of both δ_c and δ_χ in the two-fluid bias expansion of Eq. (4.57) and in accordance with the UV behavior obtained with the prescription, Eq. (4.68), our prediction for the galaxy power spectrum multipoles will need to include the following counterterms

$$P_\ell^{\text{ctr}}(k) = -2c_{cc, \ell} k^2 P_{\delta_c \delta_c}^L(k) - 2c_{c\chi, \ell} k^2 P_{\delta_c \delta_\chi}^L(k) - 2c_{\chi\chi, \ell} k^2 P_{\delta_\chi \delta_\chi}^L(k). \quad (4.69)$$

When $k_s \gg k_{\text{NL}}$ these reduce to the usual counterterms of Λ CDM analyses [19] expressed in terms of the total matter power spectrum, namely $-2c_{mm, \ell} k^2 P_{\delta_m \delta_m}^L(k)$. Conversely, when $k_s \ll k_{\text{NL}}$ the power spectra involving χ are strongly suppressed and one is left with $-2c_{cc, \ell} k^2 P_{\delta_c \delta_c}^L(k)$. This is the relevant situation for the standard massive neutrinos [97].

Qualitatively similar results are obtained for power spectra of the velocity fields, for instance

$$P_{\delta_c \Theta_\chi}^{13}(k)_\text{p} \xrightarrow{\text{UV}} -\frac{25k^2}{126\pi^2} P_{\delta_c \Theta_\chi}^L(k) \int^{\Lambda_{\text{UV}}} dq P_{\delta_c \delta_c}^L(q), \quad (4.70)$$

which should be compared with the exact perturbative expression in Eq. (4.56). In principle, the presence of the velocity divergences in the bias expansion of Eq. (4.57) would then require the introduction of

new counterterms. In practice, since $f(k) = f_{\Lambda\text{CDM}}$ and $h(\tilde{\eta}) \sim g(\tilde{\eta})$ at zeroth order in f_χ , we expect the set of counterterms in Eq. (4.69) to adequately capture all the UV sensitivity.

Accuracy for power spectra

In Fig. 4.9 we compare the 1-loop matter power spectra computed with our prescription to those obtained from exact numerical solutions of the fluid equations. We do so for the two benchmark cosmologies we have already considered in Sec. 4.2, namely an ULA with $m_a = 10^{-27}$ eV (top row) and a neutrino with $m_\nu = 1$ eV (bottom row). The low-redshift characteristic scale is $k_s(z_{\text{obs}}) \sim 0.3 h \text{ Mpc}^{-1} \sim k_{\text{NL}}$ for both models, corresponding to the physical regime where accurately modeling the nonlinear corrections is most challenging.

We discuss first the ULA cosmology. The quantitative comparison confirms the expectations based on the previous discussion: the relative deviations of the 1-loop terms $(P^{1\text{-loop}} - P_{\text{p}}^{1\text{-loop}})/P^{1\text{-loop}}$ observed in the top row of Fig. 4.9 are somewhat smaller than 1 for the $\chi\chi$ and $c\chi$ power spectra and somewhat smaller than f_χ for the cc and mm power spectra. This happens because, although the prescription introduces inaccuracies of $\mathcal{O}(f_\chi^0)$ in the χ kernels and $\mathcal{O}(f_\chi)$ in the c kernels, by construction these inaccuracies do not affect the IR limits of the kernels, which give the largest contributions to the loop integrals.

As a result, our prescription ensures a relative accuracy *better* than $f_\chi P^{1\text{-loop}}/P^L$ in the predictions of the cc and mm power spectra. For the benchmark ULA cosmology with $f_\chi = 0.1$ shown in the top row of Fig. 4.9, we obtain a $\lesssim 2\%$ accuracy up to $k = 0.4 h \text{ Mpc}^{-1}$. This implies, in particular, that our prescription can be used to accurately evaluate lensing observables, which are only sensitive to $P_{\delta_m\delta_m}$. On the other hand, for the $\chi\chi$ and $c\chi$ power spectra the prescription gives a relative accuracy that is better than $P^{1\text{-loop}}/P^L$, but is *not* proportional to f_χ . In the top row of Fig. 4.9 we find $\lesssim 20\%$ for $c\chi$ and $\lesssim 30\%$ for $\chi\chi$ (recall that $f_\chi = 0.1$ here). These results are confirmed by Fig. 4.10, where the same ULA mass as in the top row of Fig. 4.9 is assumed but the fraction is smaller, $f_\chi = 0.02$. The real-space galaxy power spectrum is evaluated up to 1-loop order, starting from the simplified two-fluid linear bias expansion $\delta_g = b_c\delta_c + b_\chi\delta_\chi$. In the left panel of Fig. 4.10 we observe that the prescription reproduces the cc component of $P_{\delta_g\delta_g}$ with negligible error. However, since based on the discussion in Sec. 4.3.2 we conservatively assume $b_\chi \sim \mathcal{O}(1)$, the prescription introduces relative errors in the $c\chi$ and $\chi\chi$ components that are not proportional to f_χ , and would remain constant even for vanishingly small f_χ .

Nonetheless, this limitation of the prescription does not hinder our ability to extract reliable cosmological constraints. The reason is that the spurious terms introduced by the prescription in the $c\chi$ and $\chi\chi$ components of the galaxy power spectrum are largely reabsorbed into the free coefficients of the $k^2 P_{\delta_c\delta_\chi}^L(k)$ and $k^2 P_{\delta_\chi\delta_\chi}^L(k)$ counterterms in Eq. (4.69). This is demonstrated by the right panel of Fig. 4.10: after summing $c\chi$ and $\chi\chi$ counterterms with fitted coefficients to $P_{\delta_g\delta_g, \text{p}}(k)$, the residual difference with the full numerical calculation of $P_{\delta_g\delta_g}(k)$ is entirely negligible relative to the BOSS error bars. This is then expected to hold also for DESI DR1 [201], as this comprises 6 bins spanning a wider range of redshift, but the error bars of power spectra of each bin are comparable to BOSS. The $c\chi$ counterterm gives the most important effect. We view this as strong evidence that our prescription does not introduce unwanted spurious effects in the derivation of constraints on f_χ from the galaxy power

4. EFTofLSS and Mixed Dark Matter




Figure 4.9.: *Top row:* Comparison of the computations of the 1-loop terms of the matter power spectra using our prescription for the nonlinear kernels, Eqs. (4.65) (dashed lines) to the full numerical results (solid lines). From left to right: $\chi\chi$, $c\chi$, cc and mm density power spectra at $z_{\text{obs}} = 0.5$, representative of the redshift observed by galaxy surveys. Red, blue and black correspond to the 22, (minus) 13 and total loop contributions. The solid gray line indicates the linear power spectrum. The bottom inset shows $(P^{1\text{-loop}} - P_p^{1\text{-loop}})/P^{1\text{-loop}}$ in black and $(P^{1\text{-loop}} - P_p^{1\text{-loop}})/(P^L + P^{1\text{-loop}})$ in gray. We have used as inputs the linear power spectra for an ULA cosmology with $m_a = 10^{-27}$ eV and $f_\chi = 0.1$, corresponding to $p = 2$ and $k_s(z_{\text{obs}}) \approx 0.29 h \text{Mpc}^{-1}$. *Bottom row:* Same as in the top row, but using as inputs the linear power spectra for a massive neutrino cosmology. We assume a single neutrino species with $m_\nu = 1$ eV, corresponding to $f_\chi \approx 0.08$, $p = 0$ and $k_s(z_{\text{obs}}) \approx 0.30 h \text{Mpc}^{-1}$.

spectrum, when b_χ is treated as an $\mathcal{O}(1)$ coefficient in the two-fluid bias expansion.⁸ The prescription remains valid even when intrinsic nonlinearities associated to the wave-like nature of ULAs are included, see Appendix 4.A.3 for a discussion.

In the bottom row of Fig. 4.9 we check the accuracy of the prescription for our benchmark massive

⁸If $b_\chi \sim \mathcal{O}(f_\chi)$ is assumed, the prescription automatically yields an accuracy on $P_{\delta_g\delta_g}$ comparable to the one for the total matter power spectrum.




Figure 4.10.: *Left*: Fractional contributions of the cc , $c\chi$ and $\chi\chi$ components to the real space power spectrum of $\delta_g = b_c\delta_c + b_\chi\delta_\chi$ up to 1-loop order. Solid lines show the full numerical results, e.g. $b_c^2 P_{\delta_c\delta_c}/P_{\delta_g\delta_g}$, whereas dashed lines correspond to evaluating the 1-loop terms in the numerator with the prescription, e.g. $b_c^2 P_{\delta_c\delta_c,p}/P_{\delta_g\delta_g}$. The bottom inset shows the difference between solid and dashed curves, e.g. $b_c^2 (P_{\delta_c\delta_c}^{1\text{-loop}} - P_{\delta_c\delta_c,p}^{1\text{-loop}})/P_{\delta_g\delta_g}$. In black we show the difference for the full galaxy power spectrum, $(P_{\delta_g\delta_g}^{1\text{-loop}} - P_{\delta_g\delta_g,p}^{1\text{-loop}})/P_{\delta_g\delta_g}$, which is dominated by the $c\chi$ component and would remain constant even for $f_\chi \rightarrow 0$. *Right*: Comparison between the full numerical result for $P_{\delta_g\delta_g}$, its value obtained with the prescription, and the additional effect of summing $c\chi$ and $\chi\chi$ counterterms with fitted coefficients. The bottom inset shows the differences between the predictions, compared to the BOSS Q_0 errorbars [200] for reference. In both panels we have set $b_c = 1.5$, $b_\chi = 0.5$ and used as inputs the linear power spectra for an ULA cosmology with $m_a = 10^{-27}$ eV and $f_\chi = 0.02$.

neutrino cosmology. Relative to the ULA scenario, the massive neutrino yields a smaller $P^{1\text{-loop}}/P^L$ ratio for the $c\chi$ and $\chi\chi$ power spectra, hence our prescription performs better. Nevertheless, the genuine 1-loop inaccuracy $(P^{1\text{-loop}} - P_p^{1\text{-loop}})/P^{1\text{-loop}}$ has comparable size to the ULA scenario and all the lessons we learned there still apply. Furthermore, the prescription is robust against the inclusion of intrinsic nonlinearities originating from the Vlasov equation, as shown in Sect. 4.A.3.

Redshift space treatment

The 1-loop contribution to the redshift-space galaxy power spectrum, which was given in Eq. (4.64), simplifies considerably after the prescription (4.65) is applied to the redshift-space galaxy kernels (4.63). The resulting form can be worked out explicitly by means of the formulas for the real-space galaxy kernels, which we report in Eqs. (4.87) of Sect. 4.A.2 (retaining only linear bias operators for brevity). Even after the application of the prescription, the evaluation of momentum integrals over the scale-dependent linear functions that appear in $F_g^{[1]}$ and $G_g^{[1]}$ is still required. For instance, the third line of Eq. (4.63c) yields the following contribution to the 13 piece,

$$P_{\delta_g\delta_g,s}^{13}(k, \mu)_p \supset 4 F_{g,s}^{[1]}(\mathbf{k}; \tilde{\eta}) P_{\delta_c\delta_c}^L(k) f_{\Lambda\text{CDM}}^2 k^2 \mu^2$$

$$\times \int \frac{d^3q}{(2\pi)^3} \frac{\mu_{\mathbf{q}}}{q} \frac{\mu_{\mathbf{k}-\mathbf{q}}}{|\mathbf{k}-\mathbf{q}|} G_g^{[1]}(\mathbf{q}; \tilde{\eta}_q) G_g^{[1]}(\mathbf{k}-\mathbf{q}; \tilde{\eta}_{|\mathbf{k}-\mathbf{q}|}) G_2(\mathbf{k}, -\mathbf{q}) P_{\delta_c \delta_c}^L(q), \quad (4.71)$$

where the expression of $G_g^{[1]}$ was given in Eq. (4.62). Analogous terms appear in the 22 piece. Additionally, there are terms built out of linear fields only, which are insensitive to the application of the prescription. In the 13 piece, these arise from the fourth and fifth lines of Eq. (4.63c) and yield

$$\begin{aligned} P_{\delta_g \delta_g, s}^{13}(k, \mu) \supset & -F_{g, s}^{[1]}(\mathbf{k}; \tilde{\eta}) P_{\delta_c \delta_c}^L(k) f_{\Lambda\text{CDM}}^2 k^2 \mu^2 \left[F_g^{[1]}(\mathbf{k}; \tilde{\eta}) + f_{\Lambda\text{CDM}} \mu^2 G_g^{[1]}(\mathbf{k}; \tilde{\eta}) \right] \\ & \times \int \frac{d^3q}{(2\pi)^3} \frac{\mu_{\mathbf{q}}^2}{q^2} [G_g^{[1]}(\mathbf{q}; \tilde{\eta}_q)]^2 P_{\delta_c \delta_c}^L(q). \end{aligned} \quad (4.72)$$

In principle, the fast evaluation of momentum convolution integrals such as those in Eqs. (4.71) and (4.72) could be achieved through a new application of FFTLog methods [99, 202, 203]. A similar strategy was followed in Ref. [97] for massive neutrino cosmologies. In the analysis of Sec. 4.4 as well as the following with DESI data in Chapt. 6, however, we adopt a simplified approach where the scale dependence of the linear functions is fully retained outside the q integrals, but the scale-independent expressions for standard CDM are used inside the integrals. Namely, in the integrands we set $g, h \rightarrow 1$ and $f \rightarrow f_{\Lambda\text{CDM}}$. This implies in particular $G_g^{[1]} \rightarrow 1$, hence the integrals in Eqs. (4.71) and (4.72) reduce to standard expressions for which the FFTLog procedure is well known.⁹

4.4. A first application: Planck+BOSS constraints on Ultra-Light Axions

The in-depth discussion of the previous sections on the extension of the EFTofLSS in the presence of multiple fluids with finite sound speed provides the theoretical framework to investigate the parameter spaces of such models with cosmological data. The two benchmark scenarios considered so far in this chapter are ULAs and light thermal relics. In the latter case the relic abundance and the suppression of the matter and galaxy power spectra depend on the mass of the new particle, the number of its bosonic or fermionic degrees of freedom and its temperature. For illustrative purposes, in Sec. 4.2.1 the latter two were identified with those of the standard neutrinos. A complete analysis would require varying the freeze-out temperature and matching to relevant particle physics models. As already mentioned, for this reason, in this first proof-of-principle analysis we only consider ULAs, our goal being to constrain their energy density $\Omega_a h^2$ as a function of their mass m_a .

4.4.1. Data specification and modeling choices

The data considered in this analysis are the CMB temperature, polarization and lensing power spectra as measured by Planck [5, 23], combined with the measurement of the full shape of the galaxy power spectrum performed by spectroscopic surveys. To enable a direct comparison with the previous cosmological constraints on ULAs presented in Ref. [204], we here restrict ourselves to SDSS/BOSS data [49, 205], and postpone an original study for DESI to Chapter 6. The BOSS sample consists of approximately one million galaxies divided in two redshift slices, centered at $z_{\text{eff}} = 0.38$ and $z_{\text{eff}} = 0.61$,

⁹In fact, Eq. (4.72) is fully degenerate with a linear combination of counterterms and it is therefore dropped from the theory prediction.

and in two regions above and below the galactic plane, for a total of four sets of measured galaxy power spectra.

Similarly to previous full-shape analyses of the Λ CDM model [18, 19, 21, 49], when inferring cosmological parameters we marginalize over four bias parameters, one EFT counterterm for each multipole of the power spectrum, and three free parameters to account for the amplitude and the scale dependence of the galaxy shot noise. These nuisance parameters completely characterize the dependence of the galaxy fluctuations on the CDM+baryon density and velocity fields (labeled by c in this thesis).

In addition, and drawing from the discussion in Sec. 4.3, we include the following new ingredients. The galaxy bias model contains one extra free parameter quantifying the response of the galaxy number density to the ULA density field, from now on specializing the generic χ index to a and the fraction f_χ to f_{ax} ,

$$\delta_g \supset b_c \delta_c + b_a \delta_a, \quad (4.73)$$

Importantly, in Eq. (4.73) the axion density field δ_a is computed up to n -th order in perturbation theory using the prescription in Eqs. (4.65), i.e.

$$\delta_a^{[n]}(\mathbf{k}) = \left(\frac{\delta_a^{[1]}(\mathbf{k})}{\delta_c^{[1]}(\mathbf{k})} \right)_{\text{CLASS}} \delta_c^{[n]}(\mathbf{k}) \quad (4.74)$$

in terms of the linear theory transfer functions.¹⁰ As noted in Sec. 4.3.1, for a proper parameter exploration that involves δ_a it is imperative to include nonlinearities of this field, at least in an approximate way. Neglecting those, e.g. keeping δ_a only at linear level, would lead to spurious effects in the regime of small deviations from pure standard CDM, which is where precision cosmological data are the most powerful. As discussed in Sec. 4.3, we expect Eq. (4.73) evaluated according to Eq. (4.74) to capture the main novelty of the two-fluid description.

The new dependence of the galaxy field on δ_a also requires the inclusion of two new classes of counterterms proportional to the linear power spectra $P_{\delta_c \delta_a}^L(k)$ and $P_{\delta_a \delta_a}^L(k)$, see Eq. (4.69). Then all other effects arising at the nonlinear level are very degenerate with b_a , or too small given BOSS/DESI DR1 errorbars. A similar argument applies to any difference between our prescription in Eq. (4.74) and the numerical solution, see Sec. 4.3.4. In total we vary, at fixed m_a , the six Λ CDM cosmological parameters plus the energy density in the axion, $\Omega_a h^2$, and 16 nuisance parameters for each of the four sets of BOSS galaxy power spectra. Finally, the new scale dependences introduced by the axion perturbations and by the linear growth rate of the cold species are neglected inside redshift space loops, as discussed in Sec. 4.3.4.

Concerning galaxy clustering data we follow Refs. [62, 204] in using the windowless measurements of the BOSS DR12 samples [205] at $z_{\text{eff}} = 0.38$ and $z_{\text{eff}} = 0.61$ provided in Ref. [49]. As in Ref. [204] we fit the galaxy power spectrum multipoles P_ℓ up to $k_{\text{max}} = 0.2 h \text{ Mpc}^{-1}$ and the RSD-free combination Q_0 [200] for $k \in [0.2, 0.4] h \text{ Mpc}^{-1}$. Differently from Ref. [204], we do not use post-reconstruction BAO information from BOSS or BAO data from other surveys. Theoretical predictions for the galaxy power spectrum are obtained with the PBJ code [62, 92], which has been integrated in the Cobaya framework [206] for the purpose.

¹⁰The linear power spectra of the CDM+baryon component, δ_c , and of the axion field, δ_a , are obtained directly from the Einstein-Boltzmann solver AxICLASS [184, 185].

4.4.2. Parameters and priors

Summing up, we fix m_a and vary the parameters

$$\begin{aligned} & (\Omega_a h^2, \Omega_c h^2, \Omega_b h^2, h, \log(10^{10} A_s), n_s, \tau_{\text{reio}}) \\ & \times \prod_{i=1}^4 (b_c, b_a, b_2, b_{\mathcal{G}_2}, b_{\Gamma_3}, c_{cc,\ell}, c_{ca,\ell}, c_{aa,\ell'}, P_{\text{shot}}, a_0, a_2)_i \end{aligned} \quad (4.75)$$

where i indexes the redshift/sky cut bins, while $\ell = 0, 2, 4$ and $\ell' = 0, 2$ run over the power spectrum multipoles. We assume one massive neutrino with fixed $m_\nu = 0.06$ eV. The first line of Eq. (4.75) contains the cosmological parameters that are varied in the fit, whereas the second line contains the EFTofLSS parameters. The linear bias coefficients b_c and b_a are defined in Eq. (4.73). We set to zero b_Θ and b_v (see Eqs. (4.57) and (4.60), respectively) because, based on the discussion in Sec. 4.3, we expect that their effects can be largely absorbed into redefinitions of b_c and b_a . The nonlinear bias coefficients $b_2, b_{\mathcal{G}_2}$ and b_{Γ_3} are given in the basis of Ref. [207]; we recall that in our analysis only the perturbations of the cold species c (but not a) are used to build the associated nonlinear operators. The counterterms are introduced in Eq. (4.69), those proportional to the $c_{ca,\ell}$ and $c_{aa,\ell'}$ coefficients being new additions associated to our two-fluid treatment. The $c_{aa,4}$ counterterm is not included in our analysis, because the associated operator $P_{\delta_g \delta_g, s}(k, \mu) \supset \mu^4 k^2 P_{\delta_a \delta_a}^L(k)$ is not generated under our assumption that $b_\Theta = b_v = 0$ (see the discussion in Sec. 4.3.4). We have checked that adding also $c_{aa,4}$ or, vice versa, removing $c_{aa,2}$, affects the bounds on $\Omega_a h^2$ by less than 5%.

While adopt wide uninformative priors for the cosmological parameters, as they are all well constrained by our combination of observables, the priors we choose for the EFTofLSS parameters require a detailed discussion. For the linear bias coefficients, by default we assume the wide uniform priors $b_c \sim \mathcal{U}(-1.5, 3)$ and $b_a \sim \mathcal{U}(-3, 3)$. This allows us to capture, in particular, the perfect degeneracy between the two parameters for large values of m_a : since all observationally relevant scales satisfy $k \ll k_s$, both underlying fluids are effectively cold and only the sum $b_c + b_a$ is constrained. This is shown in Fig. 4.11 for $m_a = 10^{-25}$ eV, corresponding to $k_s(z_{\text{obs}}) \approx 3 \text{ h Mpc}^{-1} \gg k_{\text{NL}}$. For comparison we also show the degeneracy for a much smaller mass, $m_a = 10^{-30}$ eV: in this case $k \gg k_s(z_{\text{obs}}) \approx 0.01 \text{ h Mpc}^{-1}$ for all relevant scales, hence the ULA clustering is nearly negligible and only b_c is significantly constrained. Finally, $m_a = 10^{-28}$ eV shows an intermediate situation, with $k_s(z_{\text{obs}}) \approx 0.1 \text{ h Mpc}^{-1}$ falling in the middle of the observational window: in this case both linear bias parameters are significantly constrained. The impact of modifying the prior on b_a to a narrow informative one, $b_a \sim \mathcal{U}(-0.2, 0.2)$, is discussed in Sec. 4.4.

For the nonlinear bias coefficients we adopt the same Gaussian priors as in Λ CDM analyses [49], as we also do for the counterterms of the cold species $c_{cc,\ell}$. For the new counterterms, on the other hand, we perform tests to select priors that are wide enough to be uninformative. In summary, for the counterterms we assume

$$\frac{c_{cc,0}, c_{cc,4}}{[\text{Mpc}/h]^2} \sim \mathcal{N}(0, 30^2), \quad \frac{c_{cc,2}}{[\text{Mpc}/h]^2} \sim \mathcal{N}(30, 30^2), \quad \frac{c_{ca,\ell}, c_{aa,\ell'}}{[\text{Mpc}/h]^2} \sim \mathcal{N}(0, 60^2). \quad (4.76)$$

Notice that we perform an analytical marginalization over the parameters that enter linearly in the model. This is the case of b_{Γ_3} , all the counterterm coefficients $c_{cc,\ell}$, $c_{ca,\ell}$ and $c_{aa,\ell'}$ and the stochastic parameters P_{shot} , a_0 and a_2 .




Figure 4.11.: Posteriors for the linear bias coefficients b_c and b_a of one redshift/sky cut bin ($z_{\text{eff}} = 0.38$, NGC), for different ULA masses, in our Planck+BOSS analysis. The different degeneracy directions are manifest, depending on whether $k_s \gg k_{\text{NL}}$ (blue), $k_s \sim k_{\text{NL}}$ (purple) or $k_s \ll k_{\text{NL}}$ (red).

As a validation of our setup, we reproduced with high accuracy the Planck constraints presented in Ref. [204]. In addition, we reproduced with high accuracy their Planck+BOSS constraints by restricting the theoretical model presented above to a single-fluid description and setting the priors to the standard Λ CDM values [49, 204].

4.4.3. Results

The constraints from Planck CMB data in the $(m_a, \Omega_a h^2)$ plane are shown by the red curve in Fig. 4.12 and reported in Table 4.1. The shape of these bounds is well known [183, 204, 208, 209] and represents the baseline which the addition of LSS data can further improve upon.

For $m_a \gg H(a_{\text{eq}}) \sim 10^{-28}$ eV the ULAs transition to a DM behavior before matter-radiation equality. At the high-mass end the transition happens early during radiation domination, thus leaving the pre- and post-recombination physics unchanged compared to a Λ CDM Universe. Consequently, the constraining power of the primary CMB evaporates quickly, as we observe in Fig. 4.12 for $m_a \gtrsim 10^{-26}$ eV. For lighter ULAs, such that the transition from $w_a \simeq -1$ to $w_a \simeq 0$ happens close to equality and hence recombination, the diffusion damping scale is modified and the evolution of the potential wells in which the baryon-photon plasma propagates is affected, leading to an early Integrated Sachs-Wolfe (ISW) contribution. These effects impact the relative heights of the acoustic peaks. The CMB lensing power spectrum is not measured with sufficient precision by Planck to alter significantly the above conclusions, though this is expected to change [208] with the recent ACT results [210] as will be presented in Chapt. 6, and also future data from the Simons Observatory [211] and CMB-S4 [30].

For $m_a < H(a_{\text{eq}})$ the ULAs retain $w_a \simeq -1$ until some time during matter domination (or even Λ domination, for the lightest masses). If the redshift of matter-radiation equality is kept fixed, the presence of the axion modifies the expansion history post-recombination and in particular the distance to the last scattering surface, thereby shifting the location of the CMB acoustic peaks. This shift can be




Figure 4.12.: Our final constraints on ULAs from Planck+BOSS (blue), compared to those from Planck only (red). For each value of the ULA mass, m_a , we report the 95% CL upper bound on its energy density today, $\Omega_a h^2$.

compensated by a different value of the Hubble constant H_0 , and consequently (assuming flatness) of the energy density in Λ , leaving a residual modification of the late ISW amplitude. These effects lead to the strongest Planck constraints on ULAs, $f_{\text{ax}} \lesssim 0.01$ for $10^{-30} < m_a/\text{eV} < 10^{-28}$. For even smaller masses (not shown in our plots), the modification of the distance to last scattering is progressively reduced and the constraints arise primarily from the ISW effect, resulting in weaker sensitivity [189, 204]. Finally, as m_a approaches $H_0 \sim 10^{-33}$ eV the constraint on $\Omega_a h^2$ disappears as the ULAs cannot be distinguished from a Cosmological Constant.

The LSS data are sensitive to a similar range of masses, bracketed by the ULA becoming indistinguishable from a Cosmological Constant at the lower end and from standard CDM at the upper end. The latter happens when the scale k_{drop} (and therefore also the late-time Jeans scale $k_s(z_{\text{obs}})$) becomes larger than the nonlinear scale k_{NL} , hence the ULA behaves as CDM in the whole perturbative range where the EFTofLSS applies. Our Planck+BOSS constraints are shown by the blue line in Fig. 4.12 and listed in Table 4.1. We find that the addition of BOSS data improves the bounds by roughly a factor of 2 for lower ULA masses, mostly driven by a better determination of the matter density, which helps reduce the geometric degeneracies of CMB data. At the higher end, by setting $k_{\text{drop}} \sim k_{\text{NL}} \sim 0.4 h \text{ Mpc}^{-1}$ we estimate that $m_a \sim 10^{-25}$ eV is the largest mass that can be meaningfully constrained, as confirmed by Fig. 4.12. A subset of the contour plots for $m_a = 10^{-25}$ eV is shown in the left panel of Fig. 4.13: the constraints are dominated by BOSS (cyan), while Planck data cannot distinguish between ULAs and standard CDM, as shown by the very elongated red contour in the direction of constant $\Omega_a h^2 + \Omega_c h^2$. The BOSS sensitivity arises from shape information, as it probes the suppression of the galaxy power spectrum at the largest perturbative k included in our analysis. The combination of both datasets is slightly less constraining than BOSS alone, which we attribute to projection effects. Similarly to Ref. [204], we find a preference for non-zero $\Omega_a h^2$ at $m_a = 10^{-26}$ eV, whose posterior is shown in the

m_a [eV]	Planck	Planck + BOSS (FS)
10^{-24}	0.11	0.11
10^{-25}	0.099	0.034
10^{-26}	0.0063	0.0066
10^{-27}	0.0033	0.0031
10^{-28}	0.0020	0.0013
10^{-29}	0.00158	0.00087
10^{-30}	0.00179	0.00089

Table 4.1.: 95% CL upper bounds on the ULA energy density obtained in this chapter, as a function of the ULA mass, for Planck and BOSS. An updated version with DESI and ACT follows in Chapt. 6, precisely in Table 6.2.




Figure 4.13.: Constraints on the parameter space of ULAs, at 68% and 95% CL, for $m_a = 10^{-25}$ eV (left) and $m_a = 10^{-26}$ eV (right). Red contours correspond to Planck alone, cyan contours to BOSS only, and blue contours to our final Planck+BOSS combination.

right panel of Fig. 4.13. For BOSS-only, this is likely due again to projection effects after marginalizing over the bias and EFT nuisance parameters [204], since not only is the value of $\Omega_a h^2$ preferred by BOSS incompatible with Planck, but also the other Λ CDM parameters, especially A_s and n_s , present significant shifts compared to the best fit CMB values.

In Fig. 4.14 we quantify the effect on the bounds of our refined model of the galaxy power spectrum. In the displayed ULA mass range, the addition of the new bias parameter b_a and of the accompanying new counterterms (solid blue curve, matching Fig. 4.12) markedly reduces the improvement over Planck brought by a single-fluid analysis of LSS data (dashed blue curve) we performed following closely those already present in the literature [189, 204].¹¹ As expected, the effect of marginalizing over a larger set of EFT nuisance parameters starts being important once the Jeans scale is well inside the region of wavenumbers probed by galaxy surveys, $k_s(z_{\text{obs}}) \gtrsim 0.05 h \text{ Mpc}^{-1}$. The effect is most pronounced

¹¹We checked that our single fluid analysis reproduces to high accuracy the results presented in Ref. [204].




Figure 4.14.: Impact of our refined theoretical modeling of the galaxy power spectrum on the constraints on the ULA energy density. All curves show 95% CL upper bounds. Results from the Planck+BOSS combination are reported in blue, while the red curve corresponds to Planck alone. Relative to Fig. 4.12, we zoom in on the mass range where the effect of our two-fluid EFTofLSS extension is important. The solid blue curve shows our baseline results, whereas the dashed blue curve corresponds to a single-fluid analysis following Ref. [204]. The dotted blue curve reports the constraints when restricting the prior on b_a from $\mathcal{O}(1)$ to a tighter interval of $\mathcal{O}(f_{\text{ax}})$, see the main text for details. The Jeans scale k_s is evaluated at $z_{\text{obs}} = 0.5$.

for $k_s(z_{\text{obs}}) \sim k_{\text{NL}}$, corresponding to $m_a \sim 10^{-27}$ eV, and eventually becomes negligible again when the mass is so large that k_{drop} is beyond the smallest scales included in our analysis, hence BOSS cannot distinguish the ULA fluid from standard CDM ($k_{\text{drop}} > 0.4 h \text{ Mpc}^{-1}$, corresponding to $k_s(z_{\text{obs}}) \gtrsim 3 h \text{ Mpc}^{-1}$).

Finally, some of the information gain over Planck is recovered if one is able to impose a less conservative prior on the linear bias parameter b_a , which one could expect to be of $\mathcal{O}(f_{\text{ax}})$ based on the discussion in Sec. 4.3.2. This is shown by the dotted curve in Fig. 4.14, obtained by restricting the prior on b_a from $\mathcal{U}(-3, 3)$, as assumed in our baseline analysis, to $\mathcal{U}(-0.2, 0.2)$.

Since we focus on the region of ULA masses where our novel two-fluid modeling has an impact, we do not show here results for $m_a < 10^{-30}$ eV. In this range $k_s(z_{\text{obs}})$ is so small that a single-fluid analysis is fully appropriate. Moreover, the accidental suppression of the coefficients multiplying new momentum structures in, e.g., Eqs. (4.37) implies that using the standard EdS kernels supplemented by the $1 - 3f_{\text{ax}}/5$ suppression of the linear growth rate suffices for all practical purposes.

4.5. Concluding remarks

The two-fluid extension of the EFTofLSS presented in this chapter lays the foundation to test mixed DM scenarios with future LSS data. This includes the ultra-light axions as we discussed here, but also massive thermal relics as well as models where a fraction of DM develops a high temperature due to interactions with other species.

More generally, the framework we have introduced opens up the possibility to derive “model-independent” constraints in the space of mixed DM models. For each choice of the parameters γ and p that determine the time- and scale-dependence of the sound speed (see Eq. (4.9)), one may express, with minimal assumptions, the cosmological constraints in terms of the fraction of the energy density of the non-cold species, f_χ , versus its characteristic scale, k_s . Such constraints could then be projected onto any microscopic models of interest that feature a sound speed with the assumed scaling dependences. This strategy is part of a broader program aimed at identifying a representative set of scenarios for BSM physics in the dark sector, which can be tested with LSS data. Carrying it out will likely require one to solve the nonlinear equations considered in this work for other time dependences of the sound speed term, going beyond the $\gamma = 2$ case we have focused on.

Furthermore, the exact analytical expressions we have derived for nonlinear kernels in certain kinematic configurations provide the starting point to investigate the signatures of mixed DM scenarios in higher-point correlation functions. Finally, this effort can serve as a stepping stone to investigate the effects of BSM dynamics beyond the perfect fluid approximation.

Appendix 4.A Further details on perturbation theory in Mixed DM cosmologies

4.A.1 Universality of Perturbation Dynamics in Scale-free Beyond- Λ CDM Scenarios

The “hot” limit of the two-fluid setup we consider, namely k modes for which the clustering of χ is entirely negligible, is a regime where the dynamics of the CDM+baryon fluid c is affected non-trivially by new physics, but without the appearance of a new scale. Here we show that, remarkably, the dynamics in this regime is tightly related to the one found in Refs. [193, 194] for another example of scale-free new dynamics affecting the CDM+baryon fluid: DM self-interactions mediated by a massless (scalar) field. Such tight relation is present even though the background dynamics is modified away from Λ CDM in the self-interacting DM scenario, while it is not modified in the mixed warm/cold DM scenario we study here. This result suggests a universal nature of the dynamics of cosmological perturbations in the scale-free regime, even beyond Λ CDM.

For k modes where χ does not cluster, $\delta_\chi \rightarrow 0$, the Eqs. (4.4c) and (4.4d) describing the dynamics of the cold species read, after setting $\Omega_m/f_{\Lambda\text{CDM}}^2 \rightarrow 1$,

$$\partial_\eta \delta_c - \Theta_c = \nabla_i (\delta_c V_c^i), \quad (4.77a)$$

$$\partial_\eta \Theta_c + \frac{1}{2} \Theta_c - \frac{3}{2} \delta_c (1 - f_\chi) = \nabla_i (V_c^j \nabla_j V_c^i). \quad (4.77b)$$

We now define $V_c^i \equiv (1 - 3f_\chi/5) \bar{V}_c^i$, leading to the following form for the equations up to $\mathcal{O}(f_\chi)$,

$$\frac{\partial \delta_c}{\partial [(1 - \frac{3}{5} f_\chi) \eta]} - \bar{\Theta}_c = \nabla_i (\delta_c \bar{V}_c^i), \quad (4.78a)$$

$$\frac{\partial \bar{\Theta}_c}{\partial[(1 - \frac{3}{5}f_\chi)\eta]} + \frac{1}{2}\bar{\Theta}_c\left(1 + \frac{3}{5}f_\chi\right) - \frac{3}{2}\delta_c\left(1 + \frac{1}{5}f_\chi\right) = \nabla_i(\bar{V}_c^j \nabla_j \bar{V}_c^i). \quad (4.78b)$$

Next, we consider the evolution of the matter perturbations in a scenario where 100% of DM χ self-interacts via an effectively massless scalar mediator s . We start from the first two of Eqs. (A5) in Ref. [194] for $\delta_m = F_\chi \delta_\chi + (1 - F_\chi)\delta_b$ and the analogously-defined v_m^i , where to avoid confusion we have introduced one modification to the notation used in that reference: here we denote $F_\chi = \bar{\rho}_\chi/(\bar{\rho}_\chi + \bar{\rho}_b)$. We define $V_m^i \equiv -v_m^i/(f_m \mathcal{H})$ and take $\eta_m \equiv \log D_{1m}$ as time variable, obtaining

$$\frac{\partial \delta_m}{\partial \eta_m} - \Theta_m = \nabla_i(\delta_m V_m^i), \quad (4.79a)$$

$$\frac{\partial \Theta_m}{\partial \eta_m} + \Theta_m \left[\frac{3}{2} \frac{\Omega_m}{f_m^2} (1 + F_\chi \varepsilon) - 1 \right] - \frac{3}{2} \frac{\Omega_m}{f_m^2} \delta_m (1 + F_\chi \varepsilon) = \nabla_i(V_m^j \nabla_j V_m^i), \quad (4.79b)$$

where $\varepsilon = \beta \tilde{m}_s^2 F_\chi$. Here β is the self-interaction strength normalized to gravity and $\tilde{m}_s = d \log m_\chi(s)/ds$, where $m_\chi(s)$ is the field-dependent DM mass. Finally, in matter domination we have $\Omega_m = (\bar{\rho}_\chi + \bar{\rho}_b)/\bar{\rho}_{\text{tot}} = 1 - F_\chi \varepsilon/3$ and $f_m = 1 + F_\chi \varepsilon$ [194],¹² leading us to the expressions

$$\frac{\partial \delta_m}{\partial \eta_m} - \Theta_m = \nabla_i(\delta_m V_m^i), \quad (4.80a)$$

$$\frac{\partial \Theta_m}{\partial \eta_m} + \frac{1}{2}\Theta_m(1 - 4F_\chi \varepsilon) - \frac{3}{2}\delta_m\left(1 - \frac{4}{3}F_\chi \varepsilon\right) = \nabla_i(V_m^j \nabla_j V_m^i). \quad (4.80b)$$

We observe that making the following replacements in Eqs. (4.80a) and (4.80b) gives exactly Eqs. (4.78a) and (4.78b):

$$\delta_m \rightarrow \delta_c, \quad \Theta_m \rightarrow \bar{\Theta}_c, \quad D_{1m} \rightarrow e^{(1 - \frac{3}{5}f_\chi)\eta}, \quad F_\chi \varepsilon \rightarrow -\frac{3}{20}f_\chi. \quad (4.81)$$

Thus, at all orders in perturbation theory, the solutions for the hot limit of the scenario we study in this paper are immediately obtained from the results for the m perturbations given in Ref. [194], by substituting

$$\delta_m \rightarrow \delta_c, \quad -\frac{\theta_m}{f_m \mathcal{H}} \rightarrow \left(1 + \frac{3}{5}f_\chi\right)\Theta_c, \quad D_{1m} \rightarrow e^{(1 - \frac{3}{5}f_\chi)\eta}, \quad F_\chi \varepsilon \rightarrow -\frac{3}{20}f_\chi, \quad (4.82)$$

where we recall that $F_\chi \varepsilon$ was called $f_\chi \varepsilon$ in Ref. [194].

4.A.2 More Analytical Expressions for Nonlinear Kernels

This appendix contains additional analytical results for the nonlinear kernels in the two-fluid EFTofLSS, which did not find space in the main text. We start by recalling the expressions of the standard EdS second-order kernels in terms of the mode-coupling functions α_s and β ,

$$F_2(\mathbf{k}_1, \mathbf{k}_2) = \frac{5}{7}\alpha_s(\mathbf{k}_1, \mathbf{k}_2) + \frac{2}{7}\beta(\mathbf{k}_1, \mathbf{k}_2), \quad G_2(\mathbf{k}_1, \mathbf{k}_2) = \frac{3}{7}\alpha_s(\mathbf{k}_1, \mathbf{k}_2) + \frac{4}{7}\beta(\mathbf{k}_1, \mathbf{k}_2). \quad (4.83)$$

At second perturbative order, the configurations of the χ kernels that enter the 1-loop power spectra have the following UV behavior at $\mathcal{O}(f_\chi^0)$,

$$F_\chi^{(0)[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta}) \stackrel{k_s(a) \ll q}{\simeq} \frac{t^{(0)}(\tilde{\eta}) F_2(\mathbf{q}, \mathbf{k} - \mathbf{q})}{g^{(0)}(\tilde{\eta}_q) g^{(0)}(\tilde{\eta}_{|\mathbf{k}-\mathbf{q}|})}, \quad (4.84a)$$

¹²We thank the authors of Ref. [212] for pointing out to us that in the analytical treatment of sub-horizon perturbations in Refs. [193, 194], Ω_m was incorrectly set to 1. Rectifying this leads to small changes in the coefficients of some formulas. These corrections have already been made in arXiv v4 of Ref. [193], and will also be implemented in an upcoming v3 of Ref. [194].

$$G_{\chi}^{(0)[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta}) \stackrel{k_s(a) \ll q}{=} \frac{(\partial_{\tilde{\eta}} t^{(0)} + 2t^{(0)})(\tilde{\eta}) F_2(\mathbf{q}, \mathbf{k} - \mathbf{q})}{g^{(0)}(\tilde{\eta}_q) g^{(0)}(\tilde{\eta}_{|\mathbf{k}-\mathbf{q}|})}, \quad (4.84b)$$

where

$$t^{(0)}(\tilde{\eta}) = T^{(0)}(\sqrt{6}e^{-\tilde{\eta}/2}), \quad \text{with} \quad T^{(0)}(x) = \frac{3}{10} - \frac{x^2}{20} G^{(0)}(x). \quad (4.85)$$

The ratio $t^{(0)}/g^{(0)} [(\partial_{\tilde{\eta}} t^{(0)} + 2t^{(0)})/h^{(0)}]$ smoothly transitions from $3/10$ $[3/5]$ at $k \ll k_s(a)$ to 1 $[3/2]$ at $k \gg k_s(a)$.

For the cold species at $\mathcal{O}(f_{\chi})$ we find, in the $k \ll k_s(a) \ll q$ configuration,

$$F_c^{(1)[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta}) = \frac{3}{490} [19\alpha_s(\mathbf{q}, \mathbf{k} - \mathbf{q}) + 2\beta(\mathbf{q}, \mathbf{k} - \mathbf{q})], \quad (4.86a)$$

$$G_c^{(1)[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta}) = -\frac{6}{245} [\alpha_s(\mathbf{q}, \mathbf{k} - \mathbf{q}) + 13\beta(\mathbf{q}, \mathbf{k} - \mathbf{q})], \quad (4.86b)$$

whereas in the $k_s(a) \ll k, q$ configuration one finds the results already reported in Eq. (4.37), with $\mathbf{k}_1 = \mathbf{q}$ and $\mathbf{k}_2 = \mathbf{k} - \mathbf{q}$. Equations (4.86) and (4.84) correspond to some of the analytical limits shown in Figs. 4.7 and 4.8 in the main text, respectively.

Considering for brevity only linear bias parameters, the real space galaxy kernels at second and third order are found to be

$$F_g^{[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) = b_c F_c^{[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) + b_{\chi} F_{\chi}^{[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) g(\tilde{\eta}_{k_1}) g(\tilde{\eta}_{k_2}) + b_{\Theta} \left[G_{\chi}^{[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) g(\tilde{\eta}_{k_1}) g(\tilde{\eta}_{k_2}) - G_c^{[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) \right], \quad (4.87a)$$

$$G_g^{[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) = (1 - b_v) G_c^{[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) + b_v G_{\chi}^{[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) g(\tilde{\eta}_{k_1}) g(\tilde{\eta}_{k_2}), \quad (4.87b)$$

$$F_g^{[3]}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3; \tilde{\eta}) = b_c F_c^{[3]}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3; \tilde{\eta}) + b_{\chi} F_{\chi}^{[3]}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3; \tilde{\eta}) g(\tilde{\eta}_{k_1}) g(\tilde{\eta}_{k_2}) g(\tilde{\eta}_{k_3}) + b_{\Theta} \left[G_{\chi}^{[3]}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3; \tilde{\eta}) g(\tilde{\eta}_{k_1}) g(\tilde{\eta}_{k_2}) g(\tilde{\eta}_{k_3}) - G_c^{[3]}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3; \tilde{\eta}) \right], \quad (4.87c)$$

$$G_g^{[3]}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3; \tilde{\eta}) = (1 - b_v) G_c^{[3]}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3; \tilde{\eta}) + b_v G_{\chi}^{[3]}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3; \tilde{\eta}) g(\tilde{\eta}_{k_1}) g(\tilde{\eta}_{k_2}) g(\tilde{\eta}_{k_3}). \quad (4.87d)$$

We recall that our definition of $\tilde{\eta}$ always refers to the external momentum $\mathbf{k} = \mathbf{k}_1 + \dots + \mathbf{k}_n$. Applying our prescription (4.65) to Eqs. (4.87) is the first step needed to derive the exemplary redshift-space expression shown in Eq. (4.71) in the main text.

4.A.3 Impact of Intrinsic Nonlinearities

In this chapter we described the warm species as a perfect fluid characterized by a sound speed c_s^2 at linear level, and with nonlinearities of pure gravitational origin. This is, however, an approximation, as in general each microscopic model can predict intrinsic nonlinearities, too. In this Appendix we discuss such intrinsic nonlinearities for the two benchmark models studied in this paper: ULAs, where they arise from the wave nature of the scalar field, and light but massive thermal relics, where they originate from the Vlasov (or collisionless Boltzmann) equation. In both cases, we demonstrate quantitatively that neglecting the intrinsic nonlinearities has a minor impact on our results.

Wave Nonlinearities for Ultra-Light Axions

In this chapter we approximated ULAs as a perfect fluid characterized by a sound speed $c_s^2 = k^2/(4m_a^2 a^2)$ at linear level, and with nonlinearities given only by the gravitational mode coupling. However, it is well

known that, after taking the nonrelativistic approximation for the axion field ϕ , the resulting Schrödinger-Poisson system can be recast into a fluid dynamics formulation that also includes nonlinearities intrinsic to the wave-like nature of the ULAs (see Refs. [175, 213] for previous discussions in the scenario where heavier scalars with $m_a \sim 10^{-22}$ eV form the entirety of DM). In our framework, this formulation is reproduced by making in the Euler equation for χ , Eq. (4.2b), the substitution

$$+\frac{1}{\bar{\rho}_\chi} \nabla^2 \delta \mathcal{P}_\chi \rightarrow -\frac{1}{2m_a^2 a^2} \nabla^2 \frac{\nabla^2 \sqrt{\rho_\chi}}{\sqrt{\rho_\chi}} = -\frac{1}{4m_a^2 a^2} \nabla^2 \left[\frac{\nabla^2 \delta_\chi}{1 + \delta_\chi} - \frac{1}{2} \frac{\nabla_i \delta_\chi \nabla^i \delta_\chi}{(1 + \delta_\chi)^2} \right]. \quad (4.88)$$

The continuity equation (4.2a) is unchanged. The quantity $-\nabla^2 \sqrt{\rho_\chi} / \sqrt{\rho_\chi}$ is known as quantum pressure and arises from a stress-energy tensor containing off-diagonal terms [175]. In this paper we retained only the linear piece of the quantum pressure in Eq. (4.88), which led us to the perfect fluid formulation. Here, we quantitatively show that neglecting the intrinsic wave nonlinearities has a minor effect on our results.

After changing the time variable to $\eta = \log D_{\Lambda\text{CDM}}$, setting $\Omega_m / f_{\Lambda\text{CDM}}^2 \rightarrow 1$ and going to Fourier space, we find that Eq. (4.6b) is extended to

$$\begin{aligned} & \partial_\eta \Theta_\chi(\mathbf{k}, \eta) + \frac{1}{2} \Theta_\chi(\mathbf{k}, \eta) - \frac{3}{2} [f_\chi \delta_\chi(\mathbf{k}, \eta) + (1 - f_\chi) \delta_c(\mathbf{k}, \eta)] \\ &= \int_{\mathbf{k}} dk_{12} \beta(\mathbf{k}_1, \mathbf{k}_2) \Theta_\chi(\mathbf{k}_1, \eta) \Theta_\chi(\mathbf{k}_2, \eta) \\ &+ \frac{c_s^2 k^2}{\mathcal{H}^2 \Omega_m} \sum_{r=1}^{\infty} (-1)^r \int_{\mathbf{k}} dk_{1\dots r} P_r(\mathbf{k}_1, \dots, \mathbf{k}_r) \delta_\chi(\mathbf{k}_1, \eta) \dots \delta_\chi(\mathbf{k}_r, \eta), \end{aligned} \quad (4.89)$$

where we defined

$$P_r(\mathbf{k}_1, \dots, \mathbf{k}_r) \equiv \frac{1}{2^r} \left[1 + \frac{\sum_{i=1}^r k_i^2}{(\sum_{i=1}^r \mathbf{k}_i)^2} \right]. \quad (4.90)$$

The $r = 1$ term in the third line of Eq. (4.89) reproduces the linear sound speed term in Eq. (4.6b), whereas the $r \geq 2$ terms represent the intrinsic ULA nonlinearities.

Finally, switching to the k -dependent time variable $\tilde{\eta}$ we arrive at the extended version of Eq. (4.11b),

$$\begin{aligned} & \partial_{\tilde{\eta}} \Theta_\chi(\mathbf{k}, \tilde{\eta}) + \frac{1}{2} \Theta_\chi(\mathbf{k}, \tilde{\eta}) + \frac{3}{2} e^{-\tilde{\eta}} \delta_\chi(\mathbf{k}, \tilde{\eta}) - \frac{3}{2} [f_\chi \delta_\chi(\mathbf{k}, \tilde{\eta}) + (1 - f_\chi) \delta_c(\mathbf{k}, \tilde{\eta})] \\ &= \int_{\mathbf{k}} dk_{12} \beta(\mathbf{k}_1, \mathbf{k}_2) \Theta_\chi(\mathbf{k}_1, \tilde{\eta}) \Theta_\chi(\mathbf{k}_2, \tilde{\eta}) + \frac{3}{2} e^{-\tilde{\eta}} \int_{\mathbf{k}} dk_{12} P_2(\mathbf{k}_1, \mathbf{k}_2) \delta_\chi(\mathbf{k}_1, \tilde{\eta}) \delta_\chi(\mathbf{k}_2, \tilde{\eta}) \\ &- \frac{3}{2} e^{-\tilde{\eta}} \int_{\mathbf{k}} dk_{123} P_3(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) \delta_\chi(\mathbf{k}_1, \tilde{\eta}) \delta_\chi(\mathbf{k}_2, \tilde{\eta}) \delta_\chi(\mathbf{k}_3, \tilde{\eta}) + \dots \end{aligned} \quad (4.91)$$

where we used the fact that $\gamma = 2$ for ULAs. The dots in Eq. (4.91) indicate terms with at least four fields, which we omitted since they are not needed for the calculation of the power spectra up to 1-loop order. In summary, to include intrinsic wave nonlinearities in a mixed DM scenario containing a fraction f_χ of ULAs, one needs to solve the four coupled equations (4.11a), (4.91), (4.11c) and (4.11d).

To gain some insight on the impact of the intrinsic ULA nonlinearities, we derive analytical solutions for the kernels at second perturbative order. For $n = 2$, Eq. (4.91) reads

$$\begin{aligned} & \left(\partial_{\tilde{\eta}} + \frac{1}{2} + \frac{h}{g}(\tilde{\eta}_{k_1}) + \frac{h}{g}(\tilde{\eta}_{k_2}) \right) G_\chi^{[2]} + \frac{3}{2} \left(e^{-(\gamma-1)\tilde{\eta}} - f_\chi \right) F_\chi^{[2]} - \frac{3}{2} (1 - f_\chi) \frac{F_c^{[2]}}{g(\tilde{\eta}_{k_1})g(\tilde{\eta}_{k_2})} \\ &= \beta(\mathbf{k}_1, \mathbf{k}_2) \frac{h}{g}(\tilde{\eta}_{k_1}) \frac{h}{g}(\tilde{\eta}_{k_2}) + \frac{3}{2} e^{-\tilde{\eta}} P_2(\mathbf{k}_1, \mathbf{k}_2), \end{aligned} \quad (4.92)$$

which extends Eq. (4.35b). At $\mathcal{O}(f_\chi^0)$ and considering the $k_1 \ll k_s \ll k_2$ kinematic configuration, the solutions for the χ kernels in Eqs. (4.38) are extended to

$$F_\chi^{(0)[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) = F_2(\mathbf{k}_1, \mathbf{k}_2) + P_2(\mathbf{k}_1, \mathbf{k}_2), \quad (4.93a)$$

$$G_\chi^{(0)[2]}(\mathbf{k}_1, \mathbf{k}_2; \tilde{\eta}) = 2G_2(\mathbf{k}_1, \mathbf{k}_2) - \frac{3}{14}\alpha_s(\mathbf{k}_1, \mathbf{k}_2) + \frac{1}{2}\alpha_a(\mathbf{k}_1, \mathbf{k}_2) - \frac{2}{7}\beta(\mathbf{k}_1, \mathbf{k}_2) + 3P_2(\mathbf{k}_1, \mathbf{k}_2), \quad (4.93b)$$

whereas the solutions for the cold species in Eqs. (4.40) are unaffected, because they are sourced only by gravity in this configuration. The solutions for c in the hot limit, Eqs. (4.37), are also unaffected. Importantly, since the $P_2(\mathbf{q}, \mathbf{k} - \mathbf{q})$ kernel does not have a $1/q$ pole, the IR structure discussed in Sec. 4.3.1 is not altered. The IR structure is unchanged even at higher perturbative orders, hence one still arrives at the prescription we presented in Sec. 4.3.4.

Finally, we revisit the UV structure of the second-order χ kernels. At $\mathcal{O}(f_\chi^0)$, Eqs. (4.84) are extended to

$$F_\chi^{(0)[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta}) \stackrel{k_s \ll q}{=} \frac{t^{(0)}(\tilde{\eta})F_2(\mathbf{q}, \mathbf{k} - \mathbf{q})}{g^{(0)}(\tilde{\eta}_q)g^{(0)}(\tilde{\eta}_{|\mathbf{k}-\mathbf{q}|})} + c^{(0)}(\tilde{\eta})P_2(\mathbf{q}, \mathbf{k} - \mathbf{q}), \quad (4.94a)$$

$$G_\chi^{(0)[2]}(\mathbf{q}, \mathbf{k} - \mathbf{q}; \tilde{\eta}) \stackrel{k_s \ll q}{=} \frac{(\partial_{\tilde{\eta}}t^{(0)} + 2t^{(0)})(\tilde{\eta})F_2(\mathbf{q}, \mathbf{k} - \mathbf{q})}{g^{(0)}(\tilde{\eta}_q)g^{(0)}(\tilde{\eta}_{|\mathbf{k}-\mathbf{q}|})} + (\partial_{\tilde{\eta}}c^{(0)} + 4c^{(0)})(\tilde{\eta})P_2(\mathbf{q}, \mathbf{k} - \mathbf{q}), \quad (4.94b)$$

where

$$c^{(0)}(\tilde{\eta}) = C^{(0)}(\sqrt{6}e^{-\tilde{\eta}/2}), \quad \text{with} \quad C^{(0)}(x) = \frac{x^2}{42} - \frac{x^4}{840} + \frac{x^6}{5040}G^{(0)}(x). \quad (4.95)$$

We find that $c^{(0)}[\partial_{\tilde{\eta}}c^{(0)} + 4c^{(0)}] \rightarrow e^{-\tilde{\eta}}/7 [3e^{-\tilde{\eta}}/7]$ for $k \ll k_s$ and $\rightarrow 1 [4]$ for $k \gg k_s$. We observe that, even in the $k \gg k_s$ regime where the wave nonlinearities are relevant, the new pieces on the right-hand sides of Eqs. (4.94) are strongly suppressed by the linear transfer functions, relative to the perfect-fluid solutions proportional to F_2 . We thus conclude that the wave nonlinearities give negligible contributions to the kernels in the UV limit.

In summary, by solving the second-order equations we have shown that the intrinsic wave nonlinearities do not modify the IR and UV limits of the kernels with respect to the perfect fluid description assumed in this paper. As illustrated by Eqs. (4.93), away from the IR and UV limits modifications of $\mathcal{O}(1)$ are expected for the χ kernels, as long as the external momentum is comparable to or larger than the characteristic scale. We expect these results to hold even at higher orders. This is confirmed by Fig. 4.15, where we compare the 1-loop matter power spectra obtained by solving numerically the equations with (purple) and without (black and gray) the wave nonlinearities associated with the quantum pressure. As expected, we observe mild differences in the $\chi\chi$ and $c\chi$ power spectra at $k \gtrsim k_s$, which will be anyway reabsorbed by counterterms (recall Fig. 4.10). There are also small differences at $k \lesssim 0.1 h \text{ Mpc}^{-1}$, but these are due to the finite numerical precision. The impact on the cc power spectrum, instead, is tiny. These results demonstrate that our prescription for the evaluation of the 1-loop galaxy power spectrum remains valid even when wave nonlinearities are considered. Nonetheless, it would be interesting to study the impact of the latter on other observables that can be calculated exactly, such as the tree-level bispectrum. We leave this for future work.

4. EFTofLSS and Mixed Dark Matter




Figure 4.15.: Impact of including intrinsic wave nonlinearities due to quantum pressure in the computation of the 1-loop terms of the matter power spectra, for a mixed DM scenario containing ULAs. Black and gray lines are identical to the top row of Fig. 4.9, while the purple lines correspond to the full numerical solutions including wave nonlinearities. From left to right: $\chi\chi$, $c\chi$, cc and mm density power spectra at $z_{\text{obs}} = 0.5$, representative of the redshift observed by galaxy surveys. We have used as inputs the linear power spectra for an ULA cosmology with $m_a = 10^{-27}$ eV and $f_\chi = 0.1$, corresponding to $p = 2$ and $k_s(z_{\text{obs}}) \approx 0.29 h \text{ Mpc}^{-1}$.

Nonlinearities from the Vlasov Equation

To examine the role of intrinsic nonlinearities for light thermal relics, we start from the Vlasov equation for the particle phase space density $f_\chi(\mathbf{x}, t, \mathbf{p})$,

$$\frac{df_\chi}{dt} = \frac{\partial f_\chi}{\partial t} + \frac{dx^i}{dt} \frac{\partial f_\chi}{\partial x^i} + \frac{dp^i}{dt} \frac{\partial f_\chi}{\partial p^i} = 0, \quad (4.96)$$

where $\mathbf{x}$ denotes comoving coordinates while $\mathbf{p}$ is the physical momentum. In the nonrelativistic limit and in the sub-horizon regime, taking the first two moments of the Vlasov equation yields

$$\rho'_\chi + 3\mathcal{H}\rho_\chi + \nabla_i(\rho_\chi v_\chi^i) = 0, \quad (4.97)$$

$$(\rho_\chi v_\chi^i)' + \nabla_j \Sigma_\chi^{ij} + 4\mathcal{H}\rho_\chi v_\chi^i + \rho_\chi \nabla^i \Psi = 0, \quad (4.98)$$

where primes denote conformal time derivatives, $\rho_\chi = m_\chi \int d^3p f_\chi / (2\pi)^3$ is the energy density, $v_\chi^i = \rho_\chi^{-1} \int d^3p p^i f_\chi / (2\pi)^3$ is the peculiar velocity, and $\Sigma_\chi^{ij} = m_\chi^{-1} \int d^3p p^i p^j f_\chi / (2\pi)^3$ is related to the stress tensor σ_χ^{ij} by $\Sigma_\chi^{ij} = \rho_\chi v_\chi^i v_\chi^j + \sigma_\chi^{ij}$. From these one obtains the standard continuity equation (4.2a) and the Euler equation in the form

$$\theta'_\chi + \mathcal{H}\theta_\chi + \nabla_i \left(\frac{\nabla_j \sigma_\chi^{ij}}{\rho_\chi} \right) + \nabla^2 \Psi = -\nabla_i (v_\chi^j \nabla_j v_\chi^i). \quad (4.99)$$

Extensions of perturbation theory for CDM, where the velocity dispersion tensor $\sigma_\chi^{ij} / \rho_\chi$ is included as a dynamical variable, were developed in Refs. [214–219].

We now assume that the stress tensor of the warm species contains a non-negligible pressure perturbation, while dropping viscous terms (and the negligible background pressure): $\sigma_\chi^{ij} = \delta \mathcal{P}_\chi \delta^{ij}$.

After application of the Poisson equation, our final form of the Euler equation is Eq. (4.2b) with the substitution

$$+\frac{1}{\bar{\rho}_\chi}\nabla^2\delta\mathcal{P}_\chi \quad \rightarrow \quad +\frac{1}{\bar{\rho}_\chi}\nabla_i\left(\frac{\nabla^i\delta\mathcal{P}_\chi}{1+\delta_\chi}\right) = +\frac{1}{\bar{\rho}_\chi}\left[\frac{\nabla^2\delta\mathcal{P}_\chi}{1+\delta_\chi} - \frac{\nabla_i\delta\mathcal{P}_\chi\nabla^i\delta_\chi}{(1+\delta_\chi)^2}\right]. \quad (4.100)$$

In this study, only the linear piece of the pressure term on the right-hand side of Eq. (4.100) was retained. We now demonstrate that including the additional intrinsic nonlinearities has a minor impact on our results.

After changing the time variable to η , setting $\Omega_m/f_{\Lambda\text{CDM}}^2 \rightarrow 1$ and going to Fourier space, we assume that the χ sound speed is k -independent, $\delta\mathcal{P}_\chi(\mathbf{k}, \eta) = c_s^2(\eta)\delta\rho_\chi(\mathbf{k}, \eta)$. This is appropriate, in particular, when the perturbations are adiabatic, as it is the case for thermal relics (recall that in Sec. 4.2.4 we took c_s^2 proportional to the adiabatic sound speed). We thus arrive at an equation that is formally identical to Eq. (4.89) as derived above for ULAs, but with a different (and simpler) form for the intrinsic nonlinear kernels, $P_r(\mathbf{k}_1, \dots, \mathbf{k}_r) = 1/r$. Notice that this result was already obtained in Ref. [220]. Since $\gamma = 2$ for thermal relics, Eq. (4.91) also holds. Then, the solutions to the equations including intrinsic nonlinearities for ULAs, which were given in Eqs. (4.93) - (4.95), also apply here provided the appropriate form of the P_r kernels is used. As we did for ULAs, we can thus conclude that the prescription we propose in this paper remains valid even when nonlinearities intrinsic to the Vlasov equation are considered.

5. Cross-Correlation with CMB Lensing and Higher Orders

In this chapter we present a study belonging to the other line of this thesis, the search for new observables beyond the galaxy power spectrum. Specifically, we consider a “compressed” projected bispectrum: we briefly review its definition and estimation, validate its modelling on N -body simulations, and for the first time include it in a cosmological analysis in combination with the cross-correlation with the CMB lensing convergence.

5.1. Motivation

The projected clustering of galaxies is one of the oldest probes of the large-scale structure (LSS) in existence. Early measurements of the angular correlation function of galaxies and clusters of galaxies from optical and radio surveys were vital in establishing the concept of galaxy bias, providing evidence for the large-scale homogeneity of the Universe, and confirming the basic predictions of gravitationally-led structure formation [221–225].

The amount of cosmological information that can be extracted from the angular clustering of galaxies in photometric surveys is relatively limited by comparison with their spectroscopic counterparts, where information from the full three-dimensional distribution of galaxies can be exploited [49, 226–228]. Nevertheless, photometric surveys have a number of advantages, most importantly the large number density of sources that can be photometrically identified. This allows us to study their clustering over smaller scales before reaching the shot noise-dominated regime, to increase our sensitivity to important effects, such as lensing magnification, and to obtain high-precision measurements of the cross-correlation between galaxies and other projected tracers of the LSS [229–234]. In particular, the combination of these cross-correlations with the auto-correlation of galaxies at different redshifts enables a *tomographic* reconstruction of the key physical quantities that these probes are sensitive to [235–239]. Within this general framework, *lensing tomography*, the combination of galaxy clustering and their cross-correlation with measurements of the gravitational lensing effect, has now become an established and high-precision cosmological probe. Lensing tomography analyses have allowed us to reconstruct the universe’s growth history at high precision, reaching redshifts $z \sim 3$ [240–242], as well as measuring the cosmological parameters governing the distribution evolution of cosmic structures at late times, such as the matter fraction Ω_m , or the expansion rate H_0 .

Lensing tomography, as well as other data combinations employing the angular clustering of galaxies, are set to become more important with the advent of Stage-IV [243] imaging surveys, such as Euclid [244], the Rubin Observatory [12, 245], and the future Roman Space telescope [246], as well as surveys at other wavelengths. It is therefore important to ensure that as much information as possible is recovered from these observations, including any potentially new effect that could become observable

given the increased sensitivity that these experiments will achieve. In particular, the study of higher-order correlators of the projected galaxy distribution, such as the bispectrum, has lagged behind the use of similar statistics in other probes, such as spectroscopic clustering [247–250], the Cosmic Microwave Background (CMB) [251, 252], and cosmic shear [253, 254]. In part this is due to the relative complexity inherent to bispectrum analysis. First, bispectrum estimators are typically significantly more computationally costly than power spectrum methods. Secondly, the resulting data vector is often much larger, involving a large number of triangle configurations. Thirdly, estimating the covariance matrix of the measurements using simulations can be prohibitively expensive. Finally, the regime over which different models for the galaxy bispectrum are sufficiently accurate has not been as thoroughly studied as the case of the power spectrum, and obtaining theoretical predictions is also often expensive. The first three problems in this list may be tackled through the use of novel bispectrum estimators [255–258], such as the “filter-square” approach first presented in [259]. Likewise, advances in machine learning-based emulators may be used to accelerate the theoretical calculations.

We hereby carry out the first comprehensive analysis of the angular power spectrum and bispectrum of galaxies in combination with CMB lensing. Our work will build on the analysis of [260] (H25 hereafter), where the measurements of the galaxy power- and bispectrum used here, as well as their first rudimentary analysis, were presented. Crucially, we will use simulations to quantify the regime over which a perturbative treatment of the galaxy bias may be safely used in a joint analysis of 2- and 3-point correlators, before applying this framework to the real data. We will show that including the bispectrum allows us to self-consistently measure the free parameters of the bias expansion at high precision, thereby improving the final cosmological constraints, while demonstrating the internal consistency of the model.

This Chapter is structured as follows. Section 5.2 presents the datasets used, as well as the methods used to generate the power spectrum and bispectrum measurements used in our analysis. The model used to describe these measurements, the simulations used to validate this model, and the method used to obtain constraints on the model parameters are described in Section 5.3. The results of our analysis, including both the validation on simulations and the constraints found on real data, are presented in Section 5.4. We summarise our main results and conclude in Section 5.5.

5.2. Data

5.2.1. Public galaxy and CMB lensing data

We use the publicly available *Planck* PR4 CMB lensing map, specifically the Generalized Minimum Variance (GMV) map, constructed from a joint Wiener filtering of the temperature and polarisation maps. The angular multipoles of the map are first rotated to Equatorial coordinates before being low-pass filtered with a top-hat filter that nulls all harmonic coefficients with $\ell > 3N_{\text{side}}$ and being transformed into a real-space map. Here $N_{\text{side}} = 2048$ is the HEALPix pixelisation resolution parameter used, corresponding to a pixel size of $\delta\theta \simeq 1.7'$. Similarly, the publicly available sky mask associated with this map is rotated to Equatorial coordinates in pixel space, and then apodised using a “C1” apodisation kernel with angular scale 0.2° , as implemented in NaMaster¹. We then make the resulting mask binary, setting all values above 0.5 to 1, and all other values to zero. The impact of contamination from thermal

¹<https://github.com/LSSTDESC/NaMaster>




Figure 5.1.: Redshift distribution and scale cuts. *Left*: redshift distribution for each tomographic bin used here (filled histograms). The distribution of the simulated ABACUSUMMIT sample is shown as dashed lines. *Right*: multipole-wavenumber relation, in the Limber approximation, for the filtered bispectrum data. The gray horizontal line marks the fiducial scale cut used here for the F²B measurements.

Sunyaev-Zel’dovich and Cosmic Infrared Background foregrounds on CMB lensing cross-correlations with the *Planck* data were found to be largely negligible in [51, 232].

We also use galaxy overdensity maps constructed from the photometric luminous red galaxy (LRG) sample from the 9th data release (DR9) of the Legacy Survey [261]. Specifically, we use the galaxy overdensity maps and associated masks of the “main” sample made publicly available by [262], and used in the analysis of [232] (S25 hereafter), which reports the angular number densities along other properties of the sample in Table 1. We do not perform any further post-processing steps on these maps, which are provided natively at our chosen resolution $N_{\text{side}} = 2048$. Maps are provided for four different redshift bins, spanning the redshift range $0.3 \lesssim z \lesssim 1.3$. The footprint of the LRGs is near-identical in all redshift bins, and its overlap with the CMB lensing footprint covers $\sim 42\%$ of the sky, which corresponds to about 17340 deg^2 . The redshift distributions of each bin are also publicly available, and shown for completeness in the left panel of Fig. 5.1. We use these redshifts to build our theoretical predictions for the angular power spectra and bispectra. The final ingredient needed for this analysis is an estimate of the slope in the cumulative number counts of sources as a function of magnitude s_μ , which we use to quantify the impact of magnification bias. We use the values of s_μ listed in Table 1 of S25. For further details regarding the sample selection and characterisation, the derivation of contaminant weights, redshift distributions, and number count slopes, we refer readers to [262] and S25.

5.2.2. Power spectrum and bispectrum estimation

We estimate the galaxy-galaxy angular power spectrum C_ℓ^{gg} in each redshift bin, as well as its cross-correlation with CMB lensing convergence, $C_\ell^{g\kappa}$, using the pseudo- C_ℓ algorithm PCL [263] as implemented in NaMaster [264]. All power spectra were estimated using bandpowers of width $\Delta\ell = 9$, in the range $2 \leq \ell \leq 6140$. The impact of residual mode coupling due to the presence of sky masks is forward-modelled in the theoretical predictions using the bandpower window functions predicted by

5. Cross-Correlation with CMB Lensing and Higher Orders

NaMaster. As per the procedure laid out in Appendix C of S25, an additional correction must be performed on the galaxy-convergence cross-spectrum in order to account for different masking procedures and inhomogeneous normalization introduced in the making of the CMB lensing convergence map κ . These effects are corrected using a simulation-based transfer function, computed over 480 simulated CMB lensing reconstructions [231]. Once computed, we divide the $C_\ell^{g\kappa}$ measurements by the transfer function (whose inverse is given by Eq. C.1 in S25).

We also estimate the angular galaxy bispectrum in each redshift bin, using the filtered-squared bispectrum (F²B) estimator presented in [259]. In this framework, the bispectrum of a given field δ_g is calculated as the cross-correlation between δ_g and the square of a filtered version of itself $(\delta_{g,L})^2 \equiv (W_L \otimes \delta_g)^2$, where the filter W_L suppresses all modes outside a range of multipoles around the angular scale L . The F²B is thus defined as

$$\Phi_{LL\ell}^{ggg} \equiv \langle (\delta_{g,L})_{\ell m}^2 \delta_{g,\ell m}^* \rangle, \quad (5.1)$$

and is related to the reduced bispectrum $b_{\ell_1\ell_2\ell_3}^{ggg}$ via

$$\Phi_{LL\ell}^{ggg} = \sum_{\ell_1\ell_2\ell_3} K_{LL\ell}^{\ell_1\ell_2\ell_3} b_{\ell_1\ell_2\ell_3}^{ggg}, \quad (5.2)$$

where the F²B kernel $K_{LL\ell}^{\ell_1\ell_2\ell_3}$ is

$$K_{LL\ell}^{\ell_1\ell_2\ell_3} \equiv h_{\ell_1\ell_2\ell_3}^2 \frac{W_{\ell_1}^L W_{\ell_2}^L W_{\ell_3}^{\ell_b}}{N_{\ell_b}}, \quad (5.3)$$

where we have explicitly noted that we use bandpowers, labelled by ℓ_b , with N_{ℓ_b} the total number of modes in the bandpower ($N_{\ell_b} \equiv \sum_{\ell \in \ell_b} (2\ell + 1) W_\ell^{\ell_b}$). Finally, the prefactor $h_{\ell_1\ell_2\ell_3}^2$ is

$$h_{\ell_1\ell_2\ell_3}^2 \equiv \frac{(2\ell_1 + 1)(2\ell_2 + 1)(2\ell_3 + 1)}{4\pi} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}^2. \quad (5.4)$$

Note that the cross-correlation between $(\delta_{g,L})^2$ and δ_g is estimated using the pseudo- C_ℓ method to account for mask-related mode coupling. As shown in [259], the residual mode coupling due to carrying out the filtering operation on an incomplete sky is negligibly small. Thus defined, the F²B therefore recovers the angular bispectrum for triangular configurations that are close to isosceles, although, as noted in [259], this is nevertheless able to capture a large fraction of all available triangles. The estimator can also be easily generalised to capture arbitrary triangle configurations, as well as correlations between different fields [260].

The filter W_L defined above can take a variety of forms; in this study, we define three equal-width top-hat filters which span the range $2 \leq \ell \leq 272$, each $\Delta L = 90$ wide. For each filter W_L , there is a corresponding F²B measurement over the range $2 \leq \ell \leq 6140$, binned in 682 bandpowers of width $\Delta\ell = 9$. The data vector therefore contains the gg and $g\kappa$ power spectra, as well as a few F²B measurements – the number of F²B included in the data vector will depend on the scale cuts we implement in the analysis. Indeed, a F²B whose filter selects scales beyond the regime where our theoretical predictions are reliable should be discarded; the right panel of Fig. 5.1 shows the redshift-dependent relation between Fourier modes (where we define theoretical scale cuts) and harmonic multipoles (where the filters W_L are defined) which we use to inform our filter selection in each redshift bin. For instance, for a scale cut at $k_{\max} \sim 0.08h/\text{Mpc}$ (shown as the horizontal dotted line in Fig. 5.1), it

is safe to include the first F²B measurement for the lowest redshift bin, and the first two measurements for the highest redshift bin.

The covariance of the data vector is estimated using the procedure first described in [259] – for a detailed account of the various terms entering the estimate of the covariance for cross-correlations, please refer to Appendix C of H25. It is a fully analytical approximation of the covariance, easily estimated from the data itself in a model-independent way. Reinterpreting the bispectrum as a power spectrum of fields δ_g and $\delta_{g,L}^2$ allows us to reuse already-existing covariance estimation methods for power spectra (see e.g. [265]); in doing so (which means, in practice, simply implementing the Gaussian power spectrum covariance from NaMaster), we recover all Gaussian and non-Gaussian diagonal contributions of the F²B and F²B-PCL covariance. Additional non-negligible off-diagonal contributions, called N_{222} and N_{32} for the F²B auto-covariance and F²B-PCL cross-covariance respectively, can be derived analytically, and are rescaled by the inverse of the effective sky fraction (to account for mode loss) before being added to the diagonal “Gaussian” covariance. In this analysis, we only include the N_{222} term and neglect the N_{32} term. This is because the latter, which is considerably smaller in amplitude than N_{222} in the $[\Phi_{LL\ell}^{ggg}, C_\ell^{gg}]$ covariance block, and simply negligible for the $[\Phi_{LL\ell}^{ggg}, C_\ell^{g\kappa}]$ block (see Appendix C of [260]), takes longer to compute as it contains a generalized F²B, $\Phi_{L\ell\ell}$, in which one of the legs entering the squared field is filtered on a range of scales corresponding to the bandpowers. As shown in [259], adding N_{32} to the N_{222} correction only marginally affects the F²B residual distribution.

5.3. Methodology

5.3.1. Simulations

In order to validate the theoretical model described below, and quantify the range of scales over which we may be able to use it safely, we make use of simulation-based mock observations mimicking the data we use here.

ABACUSSUMMIT is a suite of cosmological N -body simulations designed to meet and exceed the Cosmological Simulation Requirements of the DESI survey [266]. The simulations were run with ABACUS [267, 268], a high-accuracy, high-performance cosmological N -body code optimized for GPU architectures and large-volume runs, on the Summit supercomputer at the Oak Ridge Leadership Computing Facility. In this work, we use the halo light cone catalogs from the huge resolution box, which spans $7.5 h^{-1} \text{Gpc}$ on a side with 8640^3 particles, each with mass $5.7 \cdot 10^{10} h^{-1} M_\odot$ (corresponding to a particle mass 27 times larger than the base resolution). These catalogs include the past light cone geometry needed for mock survey generation, and allow us to cover the full DESI footprint with sufficient volume [269]. The simulations adopt the fiducial *Planck* 2018 cosmology ($\Omega_b h^2 = 0.02237$, $\Omega_c h^2 = 0.12$, $h = 0.6736$, $10^9 A_s = 2.0830$, $n_s = 0.9649$).

To mimic the DESI Extended photometric LRG sample from Data Release 10 (DR10) [262, 270], which is based on imaging from DECaLS on Blanco, MzLS on Mayall, and BASS on Bok, we generate galaxy mocks from the ABACUSSUMMIT halo light cones using the AbacusHOD package. This method populates halos with galaxies according to a halo occupation distribution (HOD), for which we adopt the best-fit parameters for the DESI 1% LRGs at $z = 0.5$ and $z = 0.8$ from [271], linearly interpolated between these redshifts to provide continuous evolution across $0.3 < z < 1.1$. The resulting catalog is then randomly downsampled to match the observed number densities in the four redshift bins of our target dataset.

The galaxies are binned into tomographic bins matching roughly the redshift range and distribution of the target sample. For this, each galaxy is assigned a Gaussian photometric redshift error with standard deviation $\sigma_z = 0.024(1+z)$, and then binned according to the resulting photometric redshift using the same bin edges that define our target LRG sample. The magnitude of the Gaussian photo- z error used was selected to approximately reproduce the width and tails of the redshift distributions in the real data. The resulting redshift distributions are shown in the left panel of Fig. 5.1.

We construct full-sky CMB lensing convergence maps using the density shells of the ABACUSUMMIT huge light cone. The matter distribution is projected into ~ 1000 shells of thickness $\sim 10 h^{-1}\text{Mpc}$ between $z = 0.1$ and $z = 2.3$, and mapped onto HEALPix pixels at $N_{\text{side}} = 16384$ (corresponding to $0.23'$ resolution). We compute the convergence field under the Born approximation, with sources placed at $z = 1089.3$ corresponding to the period of recombination.

5.3.2. Theory model

Angular power spectra and bispectra

In this work we make use of two projected tracers of the large-scale structure: the angular overdensity of LRGs δ_g in different redshift bins, and maps of the CMB lensing convergence κ .

The angular galaxy overdensity, accounting for the impact of lensing magnification, can be related to the three-dimensional matter and galaxy overdensities (δ_g and δ_m , respectively) via

$$\delta_g(\hat{\mathbf{n}}) = \int d\chi [q_g(\chi)\delta_g(\chi\hat{\mathbf{n}}, z(\chi)) + q_\mu(\chi)\delta_m(\chi\hat{\mathbf{n}}, z(\chi))], \quad (5.5)$$

where χ is the radial comoving distance, $\hat{\mathbf{n}}$ is a line-of-sight unit vector, and the radial kernels associated with intrinsic clustering (q_g) and magnification (q_μ) are, assuming a spatially flat universe,

$$q_g(\chi) = H(z(\chi))p(z(\chi)), \quad (5.6)$$

$$q_\mu(\chi) = \frac{3}{2}H_0^2\Omega_m(1+z)\chi \int_{z(\chi)}^\infty dz' p(z') (5s(z') - 2) \left(1 - \frac{\chi}{\chi(z')}\right). \quad (5.7)$$

Here $p(z)$ is the redshift distribution of the galaxy sample, Ω_m and H_0 are the matter fraction and the current expansion rate, respectively, and $s(z)$ is the slope of the cumulative number count of sources as a function of magnitude at the magnitude limit [262]. Following S25, we will assume that s does not vary significantly within the redshift bin, and adopt the values presented in [51]. It is worth noting that, since for this sample $s > 0.4$, magnification leads to a *positive* contribution to the observed galaxy overdensity (i.e. regions of positive matter overdensity will contribute towards a positive δ_g).

The CMB lensing convergence traces the matter overdensities directly via, again valid only for a spatially flat universe,

$$\kappa(\hat{\mathbf{n}}) = \int d\chi q_\kappa(\chi) \delta_m(\chi\hat{\mathbf{n}}, z(\chi)), \quad q_\kappa(\chi) \equiv \frac{3}{2}H_0^2\Omega_m(1+z)\chi \left(1 - \frac{\chi}{\chi_{\text{LSS}}}\right), \quad (5.8)$$

where χ_{CMB} is the distance to the last-scattering surface at $z_{\text{CMB}} \simeq 1089$. Note that this expression, as well as the description of magnification used above, neglects a scale-dependent prefactor in harmonic space, $f_\ell \equiv \ell(\ell+1)/(\ell+1/2)^2$, which is negligibly different from 1 on the scales used here.

Within the Limber approximation adopted here, the angular power spectra of δ_g and κ can be related to the 3D power spectra of δ_g and δ_m via

$$C_\ell^{gg} = \int \frac{d\chi}{\chi^2} [q_g^2(\chi)P_{gg}(k_\ell, z(\chi)) + 2q_g(\chi)q_\mu(\chi)P_{gm}(k_\ell, z(\chi)) + q_\mu^2(\chi)P_{mm}(k_\ell, z(\chi))], \quad (5.9)$$

$$C_\ell^{g\kappa} = \int \frac{d\chi}{\chi^2} q_\kappa(\chi) [q_g(\chi) P_{gm}(k_\ell, z(\chi)) + q_\mu(\chi) P_{mm}(k_\ell, z(\chi))], \quad (5.10)$$

where $P_{gm}(k, z)$ and $P_{gg}(k, z)$ are the 3D galaxy-matter and galaxy-galaxy power spectra, respectively, and $k_\ell \equiv (\ell + 1/2)/\chi$. At this point it is worth noting that the impact of lensing magnification is negligibly small in the galaxy auto-correlation, but becomes relevant at the $\sim 10\%$ level in the CMB lensing cross-correlation, leading to parameter shifts of order $\sim 1\sigma$ if neglected.

For simplicity, the expressions above have assumed the so-called Limber approximation [272], which is a rather good approximation given the relatively wide redshift bins used here. Nevertheless, we will calculate all power spectrum templates described in Section 5.3.3 without adopting this approximation on scales $\ell < 100$, where non-Limber effects are most relevant [232]. Another potentially important large-scale effect is redshift-space distortions (RSDs) [273]. In the context of projected clustering, RSDs cause a spatially-correlated migration of galaxies in and out of the selected redshift bin, leading to additional large-scale power [274]. We incorporated the impact of linear RSDs in the galaxy auto-correlation, where it is most relevant. As we show in Appendix 5.A.1, the impact of these large-scale effects is relatively small, and our results are consistent with those found after removing the scales that are most affected by them.

Finally, our analysis will incorporate measurements of the projected galaxy bispectrum $b_{\ell_1 \ell_2 \ell_3}^{ggg}$, related to the F²B estimator through Eq. (5.2). The projected and three-dimensional bispectra are then related via the Limber integral

$$b_{\ell_1 \ell_2 \ell_3}^{ggg} = \int \frac{d\chi}{\chi^4} q_g^3(\chi) B_{ggg}(k_{\ell_1}, k_{\ell_2}, k_{\ell_3}, z(\chi)). \quad (5.11)$$

As mentioned above, the impact of magnification is negligible in the galaxy auto-spectrum and, for the same reason, we will ignore its contribution to the galaxy bispectrum. This is further justified by the significantly larger statistical uncertainties of the bispectrum measurements.

Shot noise

We add a shot noise contribution to the model galaxy auto-spectrum of the form

$$C_\ell^{gg, \text{SN}} = \frac{1 + \alpha_P}{\bar{n}}, \quad (5.12)$$

where $\bar{n}$ is the angular number density of sources, and α_P is a free parameter characterising deviations from purely Poissonian shot noise. This parameter thus allows us to account for other stochastic contributions to the small-scale clustering of galaxies that are effectively uncorrelated (and hence leading to a white noise component) within the scales studied.

We will include a similar contribution to the projected bispectrum, of the form

$$b_{\ell_1 \ell_2 \ell_3}^{ggg, \text{SN}} = \left(\frac{1 + \alpha_P}{\bar{n}} \right)^2 + \frac{1 + \alpha_P}{\bar{n}} \left[C_{\ell_1}^{gg(1)} + C_{\ell_2}^{gg(1)} + C_{\ell_3}^{gg(1)} \right], \quad (5.13)$$

where $C_\ell^{gg(1)}$ is the tree-level galaxy power spectrum without shot noise. The two terms in Eq. (5.13) stem from different contribution of the finite number density to a three-point correlator, and depend on different powers of $\bar{n}$. For the explicit computation that leads to Eq. (5.13) for the projected bispectra, we refer the reader to App. D of [259]. Note that our model assumes implicitly that the impact of non-Poissonian shot noise can be taken into account in both the 2- and 3-point correlators by simply scaling

5. Cross-Correlation with CMB Lensing and Higher Orders

the effective number density of sources assuming Poisson statistics ($\bar{n}_{\text{eff}} = \bar{n}/(1 + \alpha_P)$). Although this is not the most general model describing stochastic bias components in galaxy clustering, we find this approximation to be sufficiently accurate in this analysis.

Three-dimensional clustering

The final ingredient needed in our model is a description of the three-dimensional clustering statistics of the galaxy and matter overdensities, namely $P_{gg}(k, z)$, $P_{gm}(k, z)$, and $B_{ggg}(k_1, k_2, k_3, z)$.

Here we will use a perturbative bias expansion in Eulerian space, within the EFTofLSS. However, here we use a different bias expansion at non-linear level, that is consistent with [232] and corresponds to the one implemented in pycc1 [275]. That is, we describe the galaxy and matter overdensities up to third order as

$$\delta_m = \delta^{[1]} + \delta^{[2]} + \delta^{[3]}, \quad \delta_g = b_1 \delta_m + \frac{b_2}{2} (\delta^{[1]})^2 + \frac{b_s}{2} s^2 + \frac{b_{3nl}}{2} \psi_{nl} + \frac{b_{k^2}}{2} \nabla^2 \delta^{[1]}, \quad (5.14)$$

where $\delta^{[1]}$ is the linear matter overdensity, and $\delta^{(n)}$ is the contribution to the matter overdensity at n -th order in standard perturbation theory. The squared traceless tidal field s^2 is defined as

$$s^2 \equiv s_{ij} s^{ij}, \quad s_{ij} = (\partial_i \partial_j \nabla^{-2} - \delta_{ij}^K \delta^{[1]}), \quad (5.15)$$

and the third-order operator ψ_{nl} is defined in e.g. [276].

We model the galaxy-galaxy and galaxy-matter power spectra at one-loop order in standard perturbation theory (SPT), including EFT counterterm corrections:

$$P_{gg}(k) = b_1^2 P_{1\text{-loop}}(k) + b_1 b_2 P_{\delta^{[2]}(\delta^{[1]})^2}(k) + b_1 b_s P_{\delta^{[2]}s^2}(k) + b_1 b_{3nl} P_{\delta^{[1]}\psi_{nl}}(k) + \frac{b_2^2}{4} P_{(\delta^{[1]})^2(\delta^{[1]})^2}(k) + \frac{b_s^2}{4} P_{s^2 s^2}(k) + \frac{b_2 b_s}{2} P_{s^2(\delta^{[1]})^2}(k) + b_1 b_{k^2} P_{\delta^{[1]}\nabla^2 \delta^{[1]}}(k) \quad (5.16a)$$

$$P_{gm}(k) = b_1 P_{1\text{-loop}}(k) + \frac{b_2}{2} P_{\delta^{[2]}(\delta^{[1]})^2}(k) + \frac{b_s}{2} P_{\delta^{[2]}s^2}(k) + \frac{b_{3nl}}{2} P_{\delta^{[1]}\psi_{nl}}(k) + \left(\frac{b_{k^2}}{2} + b_1 c_m \right) P_{\delta^{[1]}\nabla^2 \delta^{[1]}}(k) \quad (5.16b)$$

$$P_{mm}(k) = P_{1\text{-loop}}(k) + 2c_m P_{\delta^{[1]}\nabla^2 \delta^{[1]}}(k). \quad (5.16c)$$

This modelling is analogous to that adopted in Chapters 3 and 4, although with a different bias expansion, and neglecting here the complexities of redshift-space distortions. $P_{1\text{-loop}}(k)$ is the 1-loop SPT matter power spectrum, and all other power spectra of the form $P_{\mathcal{O}_1 \mathcal{O}_2}(k)$ are the power spectra between operators $\mathcal{O}_1$ and $\mathcal{O}_2$ at lowest non-zero order in perturbation theory. For combinations involving the higher derivatives term, we only retain the leading order contribution given by $\langle \delta^{[1]} \nabla^2 \delta^{[1]} \rangle$, ignoring the contributions from $\langle (\delta^{[1]})^2 \nabla^2 \delta^{[1]} \rangle$, $\langle s^2 \nabla^2 \delta^{[1]} \rangle$ and $\langle \nabla^2 \delta^{[1]} \nabla^2 \delta^{[1]} \rangle$ [69]. Note that we have added a counterterm for the matter field proportional to k^2 , with a free parameter c_m . A similar counterterm would also be present for the galaxy overdensity, although its amplitude is fully degenerate with the non-local bias parameter b_{k^2} .

In turn, we model the galaxy bispectrum at tree level in SPT, in which case:

$$B_{ggg}(k_1, k_2, k_3) = 2P(k_1)P(k_2) \left[b_1^3 F_2(\mathbf{k}_1, \mathbf{k}_2) + \frac{b_1^2 b_2}{2} Q_2(\mathbf{k}_1, \mathbf{k}_2) + \frac{b_1^2 b_s}{2} T_2(\mathbf{k}_1, \mathbf{k}_2) \right] + (\text{c.c.}), \quad (5.17)$$

where “c.c.” stands for the other two cyclic permutations of the three wavevectors, $P(k) \equiv P_{\delta^{[1]}\delta^{[1]}}(k)$ is the linear matter power spectrum, $F_2(\mathbf{k}_1, \mathbf{k}_2)$ is the second-order density SPT kernel of Eq. (2.16), and Q_2 and T_2 are the corresponding kernels for the quadratic and tidal fields

$$Q_2(\mathbf{k}_1, \mathbf{k}_2) \equiv 1, \quad T_2(\mathbf{k}_1, \mathbf{k}_2) \equiv \mu_{12}^2 - \frac{1}{3}, \quad (5.18)$$

with $\mu_{12} \equiv \mathbf{k}_1 \cdot \mathbf{k}_2 / k_1 k_2$. Note that we do not include non-local or higher-derivative contributions to the bispectrum, since these are mostly relevant in the presence of non-linear redshift-space distortions and at one-loop order [277].

As in S25, we neglect the redshift evolution of nuisance parameters in the Limber integrals, and evaluate them at the effective redshift defined by the galaxy clustering kernel

$$z_{\text{eff}} = \frac{\int (d\chi/\chi^2) q_g^2(\chi) z(\chi)}{\int (d\chi/\chi^2) q_g^2(\chi)} \quad (5.19)$$

In principle the other combinations of kernels, such as $q_g q_\kappa$ and q_g^3 , would entail a slightly different effective value. However, given the mild redshift evolution of the nuisance parameters observed in our results, we verified that this effect is negligible.

5.3.3. Accelerating theory predictions

Although the calculation of the theoretical predictions for the angular power spectra can be carried out with relative efficiency, particularly using fast schemes to calculate PT convolutions, such as FAST-PT [99, 278], the time to complete a fully converged Markov chain (see Section 5.3.4) can still be significant (e.g. hours or days). This problem is exacerbated in the case of the angular bispectrum, given the large number of Limber integrals needed to cover all required triangle configurations, and the additional complication of convolving the result with the F²B kernel (Eq. (5.3)). In order to increase the efficiency of our theory predictions as a function of cosmological parameters, we employ machine-learning based emulators. Concretely, we will emulate projected quantities, integrated over the radial kernels of the tracers under study. Thus, our emulator is strictly tied to the specific dataset considered in this work.

In both the 2-point and 3-point observables studied here, the bias coefficients and counterterms (b_1 , b_2 , b_s , b_{3nl} , b_{k^2} , c_m) are always multiplicative and can be brought outside the Limber integrals. For instance, consider the galaxy-convergence power spectrum in Eq. (5.10). Within the perturbative bias expansion used here, it can be written as

$$C_\ell^{g\kappa} = b_1 C_{1-\text{loop},\ell}^{g\kappa} + \frac{b_2}{2} C_{\delta^{[2]}(\delta^{[1]})^2,\ell}^{g\kappa} + \frac{b_s}{2} C_{\delta^{[2]}s^2,\ell}^{g\kappa} + \frac{b_{3nl}}{2} C_{\delta^{[1]}\psi_{nl},\ell}^{g\kappa} \quad (5.20)$$

$$+ \left(\frac{b_{k^2}}{2} + b_1 c_m \right) C_{\delta^{[1]}\nabla^2\delta^{[1]},\ell}^{g\kappa} + C_{1-\text{loop},\ell}^{\mu\kappa} + c_m C_{\delta^{[1]}\nabla^2\delta^{[1]},\ell}^{\mu\kappa}, \quad (5.21)$$

where we have defined the angular power spectrum templates corresponding to the different PT power spectra in Eq. (5.16):

$$C_{\mathcal{O}_1\mathcal{O}_2,\ell}^{ab} \equiv \int \frac{d\chi}{\chi^2} q_a(\chi) q_b(\chi) P_{\mathcal{O}_1\mathcal{O}_2}(k_\ell, z(\chi)). \quad (5.22)$$

A similar expression can be derived for C_ℓ^{gg} , with the addition of the shot noise component. Crucially, these templates depend only on a small number of cosmological parameters (two in our case, see below), and can therefore be emulated a priori.

5. Cross-Correlation with CMB Lensing and Higher Orders

In the same way for the galaxy bispectrum, where the acceleration enabled by emulators is even more essential, we may write

$$\Phi_{LL\ell}^{ggg} = b_1^3 F_{LL\ell}^{ggg} + \frac{b_1^2 b_2}{2} Q_{LL\ell}^{ggg} + \frac{b_1^2 b_s}{2} T_{LL\ell}^{ggg} + b_1^2 (1 + \alpha_P) \text{SN}_{2,LL\ell}^{ggg} + (1 + \alpha_P)^2 \text{SN}_{1,LL\ell}^{ggg}, \quad (5.23)$$

where we have defined the F²B templates:

$$\mathbf{x}_{LL\ell}^{ggg} \equiv \sum_{\ell_1 \ell_2 \ell_3} K_{LL\ell}^{\ell_1 \ell_2 \ell_3} \mathbf{b}_{\ell_1 \ell_2 \ell_3}^{ggg, X}, \quad (5.24)$$

with the F²B kernel defined in Eq. (5.3), and

$$\mathbf{b}_{\ell_1 \ell_2 \ell_3}^{ggg, X} \equiv \int \frac{d\chi}{\chi^4} q_g^3(\chi) [2P(k_{\ell_1})P(k_{\ell_2})X(\mathbf{k}_{\ell_1}, \mathbf{k}_{\ell_2}) + (\text{c.c.})] \quad \text{for } X \in \{F, Q, T\}, \quad (5.25)$$

$$\mathbf{b}_{\ell_1 \ell_2 \ell_3}^{ggg, \text{SN}_2} \equiv \frac{1}{\bar{n}} \left[\mathbf{C}_{\delta^{[1]}\delta^{[1]}, \ell_1}^{gg} + \mathbf{C}_{\delta^{[1]}\delta^{[1]}, \ell_2}^{gg} + \mathbf{C}_{\delta^{[1]}\delta^{[1]}, \ell_3}^{gg} \right], \quad \mathbf{b}^{ggg, \text{SN}_1} \equiv \frac{1}{\bar{n}^2} \quad (5.26)$$

As before, the X templates depend only on cosmological parameters.

We build one emulator per template and per redshift bin. The architecture of each of our emulators is a Feed-Forward Neural Network, with 2 hidden layers of 128 nodes each. The input layer of the network has size 2, corresponding to $\{\log(10^{10} A_s), \Omega_c h^2\}$. The output layer has the same size as our ℓ vector. For each emulator, we create a sample of 10,000 input-output pairs, which we split into a training set (holding 90% of the sample) and a test set (with the remaining 10%). We build and train our models with Keras3 [279] and TensorFlow [280], using a Rectified Linear Unit (ReLU) activation function, and optimizing our weights and biases with the Adam gradient descent algorithm [281], with exponentially decreasing learning rate. We use mean squared errors between the training and test sets as a loss function and, during training, also compute an accuracy metric that accounts for the percentage of output nodes differing less than 0.5% from the ground truth. We use this metric to establish an early stopping criterion. The emulators we obtain with this procedure are able to reproduce the corresponding test sets at better than 1% accuracy on all of the considered scales, see App. 5.A.3.

The numerical computation of the templates appearing in Eq. (5.25), with the method developed in H25, takes ≈ 5 s when already emulating the linear power spectrum [282], rather than calling the Einstein-Boltzmann solver, and interpolating the evaluation of Eq. (5.24). Still, this makes posterior sampling, which we find requires evaluating the likelihood $\mathcal{O}(10,000 - 100,000)$ times, practically unfeasible, especially when combining the redshift bins. Our approach of emulating the templates reduces the single evaluation time of Eq. (5.25) by three orders of magnitude, down to ≈ 5 ms, enabling a feasible, and extremely fast, exploration of the posterior with MCMC methods.

5.3.4. Likelihood

We derive parameter constraints from the measurements described in Section 5.2 assuming a Gaussian likelihood of the form

$$-2 \log p(\mathbf{d}|\theta) = (\mathbf{d} - \mathbf{t}(\theta))^T \mathbf{C}^{-1} (\mathbf{d} - \mathbf{t}(\theta)) + K, \quad (5.27)$$

where $\mathbf{d}$ is the data vector, $\mathbf{t}(\theta)$ is the theory prediction, described in the preceding sections and dependent upon a set of free parameters θ , and K is an arbitrary normalisation constant.

parameter	prior	parameter	prior
$\Omega_c h^2$	$\mathcal{U}(0.1, 0.14)$	b_{3nl}	$\mathcal{U}(-10, 10)$
$\log(10^{10} A_s)$	$\mathcal{U}(2.3, 3.5)$	b_{k^2}	$\mathcal{N}(0, 15 \text{ Mpc}^2)$
b_1	$\mathcal{U}(1.2, 3.2)$	c_m	$\mathcal{N}(0, 10 \text{ Mpc}^2)$
b_2	$\mathcal{U}(-5, 5)$	s_μ	$\mathcal{N}(s_\mu^{\text{fid}}, 0.1)$
b_s	$\mathcal{U}(-10, 10)$	α_P	$\mathcal{N}(0, 0.3)$

Table 5.1.: Model parameters and priors used.

In our analysis, $\mathbf{d}$ will comprise a set of power spectrum (C_ℓ^{gg} and/or $C_\ell^{g\kappa}$) and F²B measurements ($\Phi_{LL\ell}^{gg}$), with the covariance matrix estimated as described in Section 5.2.2. The range of scales over which different power spectra and bispectra are used was determined following the validation against simulations presented in Section 5.4.1. In our fiducial analysis we will use power spectrum data on angular scales $\ell < \ell_{\text{max}}^P \equiv k_{\text{max}}^P \chi(\bar{z})$, where $\chi(\bar{z})$ is the comoving distance to the mean redshift of the sample, and $k_{\text{max}}^P = 0.2 \text{ hMpc}^{-1}$. Likewise, we will include F²B measurements on scales $(L_{\text{mean}}, \ell) < \ell_{\text{max}}^B \equiv k_{\text{max}}^B \chi(\bar{z})$, with $k_{\text{max}}^B = 0.08 \text{ hMpc}^{-1}$. Here L_{mean} is the mean multipole of the top-hat filter defined in Sect. 5.2.2; our fiducial criterion thus corresponds to including the first filter for all bins, as well as the second filter for z_3 and z_4 . Note that we find that the constraints on Ω_m and σ_8 do not improve significantly by including smaller scales. In addition to this, we remove the largest angular scales in the power spectrum measurements, with $\ell < 29$ (i.e. we discard the first three bandpowers).

The free parameters of our model will be the full set of bias expansion parameters, including the counter-term and shot-noise amplitude, $\{b_1, b_2, b_s, b_{k^2}, c_m, \alpha_P\}$, as well as two cosmological parameters, the physical abundance of cold dark matter $\Omega_c h^2$, and the amplitude of primordial scalar fluctuations $\log(10^{10} A_s)$. We fix all other Λ CDM cosmological parameters as follows: following [232], we use $\Omega_m h^3 = 0.09633$ [5] as a proxy for the angular acoustic scale, measured by *Planck* at very high precision, as well as $\Omega_b h^2 = 0.02236$, $n_s = 0.9649$ (also measured by *Planck*). We also fix the sum of neutrino masses to $M_\nu = 0.06 \text{ eV}$. Although the parameters we will sample are $\Omega_c h^2$ and $\log(10^{10} A_s)$, we will most commonly present our constraints in terms of the matter fraction Ω_m and the linear matter power spectrum amplitude σ_8 , as derived parameters. When considering multiple redshift bins in the same dataset, we will assume that the clustering of galaxies and matter in each bin is described by a different set of bias parameters (including α_P and c_m).

We sample the posterior distribution using *nautilus*, a nested sampling algorithm that optimises the way its live points explore the parameter space thanks to the use of Neural Networks [106]. We use 1000 live points, with 4 Neural Networks, and with convergence criterion $f_{\text{live}} = 0.01$ and $n_{\text{eff}} = 5000$. To construct this posterior we impose the priors on our model parameters listed in Table 5.1. The flat priors on $\Omega_c h^2$ and $\log(10^{10} A_s)$ coincide with the parameter ranges over which the emulator described in Section 5.3.3 is trained, and were chosen to ensure that they enclose the region of parameter space that can be constrained by the data in each redshift bin. The priors for b_2 , b_s , and b_{3nl} were chosen to be sufficiently broad to encompass the ranges over which they may be constrained by two-point function data alone. The prior on b_{k^2} mimics the prior used in [232] for the counterterm amplitude and the prior

5. Cross-Correlation with CMB Lensing and Higher Orders

on c_m is based on the arguments described in [283], whereby

$$c_m \simeq \mathcal{O}(1) \times \frac{1}{k_{NL}^2} \approx \mathcal{O}(1) \times 5 h^{-2} \text{Mpc}^2 \quad (5.28)$$

for $k_{NL}(z = 0.5) \simeq 0.45 h \text{Mpc}^{-1}$.

5.4. Results

We now present the results of our analysis, beginning with a validation of the model using simulations before applying it to the real LRG data.

5.4.1. Validation on simulations

Simulation data provides a controlled setting, where the true underlying cosmology is known, to assess the performance of the theoretical model employed.

A key aspect of full-shape fits of galaxy clustering data is understanding the regime of validity of the model, i.e. defining respectively the largest and smallest scales that can be employed for the purpose of parameter inference. While the limiting factor on the largest scales is usually the presence of systematics, or the potential mismodelling of large-angle effects (e.g. the Limber approximation, redshift-space distortions, etc.), the modelling on smaller scales is limited by non-linearities in the clustering. The latter is particularly subtle to assess, and several approaches have been proposed in the literature, most of which rely on *a posteriori* validation by varying the small-scale cut k_{max} in the fit. In general, we expect the fully theoretical, first-principle, model based on 1-loop (tree level) EPT to be able to describe the mildly non-linear scales in the real-space halo and galaxy power spectrum up to a k_{max}^P of order $\approx 0.3 h \text{Mpc}^{-1}$ ($\approx 0.1 h \text{Mpc}^{-1}$), as extensively validated for 3D clustering on N-body simulations [284–286] for the redshifts of interest to this work.

With this aim, we thus start by analysing the mock LRG catalogue and CMB lensing map generated from the ABACUSUMMIT simulation as described in Section 5.3.1. We analyse the simulated data using the same pipeline as for the real data, measuring the galaxy auto-spectrum, the galaxy- κ cross-spectrum, and the ggg F²B. We will then extract parameter constraints from these measurements using the covariance matrix estimated for the real data, allowing us to quantify the amplitude of any potential parameter biases against their expected measurement uncertainties in the real data. Note, however, that the simulated data vector has intrinsically smaller statistical uncertainties than the data, since the catalogue and κ map cover the full celestial sphere, and the κ map does not include any noise contribution from lensing reconstruction. This means that we should expect the scatter in the best-fit parameters extracted from the simulation to be much smaller than the posterior uncertainties, so any significant deviation from the true cosmology of the simulation will be likely due to the limitations of the theory model used.

To quantify the adequacy of the theoretical model, we consider two key performance metrics commonly employed in the literature [287]. Specifically, we compute the Figure of Bias (FoB), which quantifies the unbiasedness of the model in the recovery of the fiducial parameters, and the Figure of Merit (FoM), which measures statistical power in constraining the parameters. These are computed for the two cosmological parameters $\theta = \{\Omega_m, \sigma_8\}$ as a function of the small-scale cut. The FoB and FoM




Figure 5.2.: Model performance in simulated data. Figure of Bias (FoB, top panel), and Figure of Merit (FoM, bottom panel) for the galaxy-galaxy and galaxy- κ angular power spectra in each redshift bin, for the two cosmological parameters $\{\Omega_m, \sigma_8\}$, as a function of the small-scale cut $k_{\max}^P$. Note that the FoM is reported relative to its value for $k_{\max}^P = 0.08 \, h\text{Mpc}^{-1}$, in order to show results from all bins on the same scale. The solid lines refer to the nonlinear bias model assumed in our fiducial analysis, while dashed lines refer to a simplistic linear bias model. While linear bias can only give unbiased cosmological results up to $k_{\max}^P = 0.1 \, h\text{Mpc}^{-1}$, the nonlinear model can extend to $k_{\max}^P = 0.3 \, h\text{Mpc}^{-1}$ (but without significant gains in terms of extracted information for $k_{\max}^P > 0.2 \, h\text{Mpc}^{-1}$).

are defined as

$$\text{FoB} \equiv \sqrt{(\langle\theta\rangle - \theta_{\text{true}})^T \text{Cov}_\theta^{-1} (\langle\theta\rangle - \theta_{\text{true}})}, \quad \text{FoM} \equiv \frac{1}{\sqrt{\det(\text{Cov}_\theta)}}, \quad (5.29)$$

where $\langle\theta\rangle$ is the mean of the posterior distribution for the model parameters θ , θ_{true} is the true value of the parameter θ , and Cov_θ is the covariance of the posterior. The FoB thus quantifies the bias induced on the input parameters as a fraction of their uncertainties, while FoM quantifies the constraining power of the model.

We begin by quantifying the scales over which the one-loop bias model used here is able to describe the projected power spectra C_ℓ^{gg} and $C_\ell^{g\kappa}$. For this, we analyse the simulated measurements excluding the F²B for now. Fig. 5.2 reports the result of the fits for each individual redshift bin. For reference, we also report the FoB for a “linear bias” model in which b_1 is the only nuisance parameter. This simpler model, while recovering unbiased results at very large scales ($k_{\max}^P \lesssim 0.1 \, h\text{Mpc}^{-1}$), fails shortly afterwards. The apparently counter-intuitive fact that the metric, at a given scale cut, is worse for higher-redshift bins is an artifact of plotting as a function of comoving scales (as opposed to angular scales): a given $k_{\max}^P$ corresponds to a higher multipole at higher redshifts, and the measurements thus achieve a stronger constraining power (i.e. smaller uncertainties). In contrast, the higher-order bias model used here (see Section 5.3) consistently returns unbiased constraints on the two cosmological parameters for the various scale cuts and redshift bins that we consider. Furthermore, while there is a notable increase in constraining power when including mildly non-linear scales, this seemingly saturates already at $k_{\max}^P \sim 0.3 \, h\text{Mpc}^{-1}$, or earlier ($k_{\max}^P \sim 0.2 \, h\text{Mpc}^{-1}$) for the higher-redshift bins. This saturation can be related to the impact of shot noise and CMB lensing reconstruction noise on




Figure 5.3.: Consistency with bispectrum scale cut on simulation data. Posterior constraints on the model parameters as a function of the small-scale cut used for the galaxy bispectrum $k_{\max}^B$. Results are shown for two redshift bins, z_3 and z_4 (left and right panels, respectively). Similar results are found for the first two redshift bins. The shaded coloured bands show the 68% constraints on each parameter using the data covariance, with the gray band showing the constraints found when using only power spectrum data with $k_{\max}^P = 0.08 \, h\text{Mpc}^{-1}$ (and $\ell_{\min} = 29$). Vertical lines represent, incrementally, the high- ℓ bound of the different filters used.




Figure 5.4.: Constraints from the ABACUSUMMIT simulated data. *Left*: posterior, for single and combined redshift bins, of the cosmological parameters in the combined $gg + g\kappa + ggg$ analysis. The dashed lines show the true cosmological parameters of the ABACUSUMMIT simulation, while the yellow cross marks the best-fit parameters found using the data covariance in the likelihood. The solid purple contours show the constraints obtained from the combination of all redshift bins, with the dotted contours showing the constraints found from individual redshift bins, all using the data covariance. The constraints found using the simulation covariance (significantly smaller due to the larger sky area and lack of lensing reconstruction noise) are shown in pink. The true cosmological parameters are recovered in all cases with no significant bias. *Right*: posterior, for the highest redshift bin, showing the relevant degeneracies of σ_8 with the bias parameters with and without the 3-point information, using the data covariance. This shows how the 1D posterior for σ_8 is affected by the marginalization over the bias parameters, especially along the degeneracy with b_2 , when new information is added, leading also to a small shift in the central value.

the small-scale statistical uncertainties, as well as the model flexibility allowed by the multiple bias parameters in the model, particularly the counter-term contributions proportional to $k^2 P(k)$. This is explored further in Appendix 5.A.2. Given these results, we make the conservative choice of setting the same small-scale cut of $k_{\max}^P = 0.2 \, h\text{Mpc}^{-1}$ for all the redshift bins, which we apply to both C_ℓ^{gg} and $C_\ell^{g\kappa}$.

To estimate a reliable scale cut for the three-point correlation function, we carry out a separate, dedicated analysis. The main difficulty in doing so is the fact that, on its own, $\Phi_{LL\ell}^{ggg}$ is not able to place meaningful constraints on cosmological and bias parameters simultaneously. For the same reason, the same metrics used above (FoB and FoM), if applied to the combination of C_ℓ^{gg} , $C_\ell^{g\kappa}$ and $\Phi_{LL\ell}^{ggg}$, would be dominated by the two-point functions, making them almost insensitive to the bispectrum. Thus, we adopt a different strategy, combining $\Phi_{LL\ell}^{ggg}$ with C_ℓ^{gg} and $C_\ell^{g\kappa}$, but including the latter two only on the largest scales ($k_{\max}^P = 0.08 \, h\text{Mpc}^{-1}$), thus minimising the information extracted by the two-point function. Furthermore, rather than simply examining the FoB and FoM directly, we will instead quantify the stability of the posterior constraints on all parameters as a function of $k_{\max}^B$. Fig. 5.3 shows the constraints on all cosmological and bias parameters for the two highest redshift bins as a function of $k_{\max}^B$ in this setup. Similar results were found for the other two redshift bins. We find that the preferred

values of both Ω_m and σ_8 are largely stable with respect to the choice of k_{max}^B , and their statistical uncertainties only vary slowly within the range of scales studied. In turn, the constraints on the bias parameters benefit strongly from the inclusion of the three-point function. Their preferred value is also stable with respect to k_{max}^B . Given these results, we choose a scale cut of $k_{\text{max}}^B = 0.08 \, h\text{Mpc}^{-1}$, compatible with the values assumed in three-dimensional clustering studies [73, 283], as a compromise between the final size of the data vector and the desire to improve the constraints on the higher-order bias parameters.

Having chosen our fiducial scale cuts, $(k_{\text{max}}^P, k_{\text{max}}^B) = (0.2, 0.08) \, h\text{Mpc}^{-1}$, we now derive constraints from the full data vector, combining all two- and three-point correlations from the four redshift bins. The resulting constraints on Ω_m and σ_8 are shown in the left panel of Fig. 5.4. The constraints found using the data covariance matrix are compatible with the true cosmology of the ABACUSUMMIT simulations within 1σ , and the best-fit parameter values (defined as the point in the MCMC chain with the lowest χ^2) are very close to the truth. Similar results are found for each redshift bin when analysed individually, and these are also shown in the figure for comparison. To determine whether the small shift observed between the posterior constraints and the true parameter values is a consequence of volume/projection effects or a genuine parameter bias, we repeat our analysis using the covariance matrix of the simulation (which, as discussed above, leads to significantly smaller data errors). The parameter constraints found in this case, shown in Fig. 5.4, are significantly tighter and perfectly compatible with the true parameter values. This suggests that the small parameter shift observed with the data covariance is indeed a result of projection effects. To confirm this further, the right panel of Fig. 5.4 shows the constraints on σ_8 and the b_1 and b_2 bias parameters of the fourth redshift bins. We can see that the multi-parameter posterior distribution is clearly non-Gaussian, and marginalisation over the bias parameters, particularly in the case of b_2 , can lead to a shift in the marginalised posterior distribution. We also verified that with the true simulation covariance this shift moves the posterior closer to the true cosmology, the $gg + g\kappa + ggg$ result being shown in pink in the left panel of Fig. 5.4.

In conclusion, we find that the bias model used here can describe our measurements with sufficient accuracy within our fiducial scale cuts. In Appendix 5.A.2 we explored the ability of simpler bias parametrisations to recover unbiased cosmological constraints, using the mock data and validation procedure described here. The results of these simplified models will be briefly discussed in the next section.

5.4.2. Cosmological constraints from 2- and 3-point correlators

Having validated our bias model and scale cuts, we now apply them to the analysis of the real data. Our constraints on Ω_m and σ_8 for different data combinations, as well as the best-fit χ^2 and associated probability-to-exceed (PTE, i.e. the probability of obtaining a χ^2 value above that observed in the data, given the number degrees of freedom of the model), are summarised in Table 5.2. We discuss these results in detail in the remainder of this section.

Figure 5.5 shows our power spectrum and F²B measurements for the fourth redshift bin, together with the best-fit theory prediction. The galaxy bispectrum is measured at high significance in the two filter scales used in the analysis (covering $2 \leq \ell < 92$ and $92 \leq \ell < 182$). Similar results are found for the other three redshift bins. As can be seen in the figure, and as quantified in Table 5.2, we find that the model is able to fit all components of the data vector well in all cases explored here, with PTE

z_{bin}	Dataset	Ω_m	σ_8	fit χ^2 (PTE)
z_1	$gg + g\kappa$	0.296 ± 0.034	0.686 ± 0.058	1.07 (0.35)
	$gg + g\kappa + ggg$	0.294 ± 0.033	0.686 ± 0.055	0.98 (0.50)
z_2	$gg + g\kappa$	0.329 ± 0.035	0.772 ± 0.050	0.96 (0.55)
	$gg + g\kappa + ggg$	0.337 ± 0.036	0.761 ± 0.046	0.91 (0.69)
z_3	$gg + g\kappa$	0.297 ± 0.025	0.773 ± 0.042	1.02 (0.44)
	$gg + g\kappa + ggg$	0.302 ± 0.024	0.758 ± 0.036	0.99 (0.52)
z_4	$gg + g\kappa$	0.290 ± 0.023	0.785 ± 0.040	0.93 (0.65)
	$gg + g\kappa + ggg$	0.296 ± 0.022	0.765 ± 0.035	0.89 (0.78)
$z_2 + z_3 + z_4$	$gg + g\kappa$	0.301 ± 0.016	0.774 ± 0.023	0.96 (0.64)
	$gg + g\kappa + ggg$	0.307 ± 0.015	0.761 ± 0.020	0.96 (0.65)
All bins	$gg + g\kappa$	0.299 ± 0.014	0.759 ± 0.020	0.99 (0.54)
	$gg + g\kappa + ggg$	0.303 ± 0.014	0.753 ± 0.018	0.97 (0.63)
All bins $\{b_1, b_2, b_s\}$	$gg + g\kappa$	0.299 ± 0.011	0.778 ± 0.020	0.95 (0.72)
	$gg + g\kappa + ggg$	0.300 ± 0.010	0.764 ± 0.018	0.95 (0.75)
All bins $\{b_1, b_{k^2}, c_m\}$	$gg + g\kappa$	0.295 ± 0.011	0.772 ± 0.021	0.94 (0.74)
	$gg + g\kappa + ggg$	0.293 ± 0.009	0.785 ± 0.016	0.98 (0.61)

Table 5.2.: Cosmological parameter constraints found from different data combinations. Results are shown for each redshift bin independently, as well as for the combination of the last three bins and all four bins. The last two sets of rows show the constraints found when analysing the data under the simplified bias models explored in Appendix 5.A.2. In each case, we show constraints from both two-point and two- and three-point correlators. The last column shows the reduced χ^2 and probability-to-exceed for the best-fit model.

values ranging from 0.35 to 0.78.

Figure 5.6, left panel, shows the two-dimensional marginalised constraints on Ω_m and σ_8 obtained from each individual redshift bin, compared with the constraints found by *Planck*. Results are shown for two-point functions only and including the F²B data. We find that the inclusion of three-point information leads to consistently smaller uncertainties by ~ 10 -20%, in addition to a preference for a smaller value of σ_8 and larger Ω_m . Interestingly, this preference is also found in full-shape analyses of the three-dimensional clustering of galaxies when including the bispectrum [49, 50, 288]. We also find that the data in the first redshift bin favours a value of σ_8 that is significantly lower than the *Planck* constraint, by approximately 2.3σ . This is a known result, found in past analyses of this sample by [232] (S25 hereafter) and [51]. As reported in H25 [260], we find that the galaxy bispectrum measurements are also compatible with this result.

We compare these tomographic constraints on σ_8 with previous results in the top panel of Fig. 5.7. Our results are in general good agreement with those found by S25 using the same sample. Our reported uncertainties are $\sim 20\%$ larger than those of S25 for this tomographic analysis, although the latter are supplemented by external constraints on Ω_m from baryon acoustic oscillation data. The constraints found by H25 for two-point data are systematically lower than those found here by about $\sim 0.5\sigma$, rising




Figure 5.5.: Power spectrum and F²B measurements, and model fit. Measurements of C_{ℓ}^{gg} , $C_{\ell}^{g\kappa}$, and $\Phi_{LL\ell}^{ggg}$ in the fourth redshift bin. Results are shown over the scales used in the analysis, with the F²B measurements covering the first two filters (the corresponding scale ranges are shown in the legend). The red line shows the best-fit model found from the combination of all data shown here, while the blue line shows the prediction found by fitting only the two-point functions. The lower panel reports the residuals with respect to the best-fit model normalized by the statistical uncertainties.

to $\sim 1\sigma$ when including bispectrum data. These differences are not unexpected, however. First, H25 analysed a previous version of the LRG catalogue compared to that used here and in S25, where a number of shortcomings were addressed (see [262] for further details). Secondly, the constraints obtained by H25 from two-point data modelled the galaxy power spectrum at the tree level, equivalent to a linear bias model. Although this approach may lead to unbiased constraints using sufficiently conservative scale cuts, we saw in Section 5.4.1 that the model is significantly less stable, and past analyses have found consistent downward shifts in amplitude parameters, such as σ_8 , when using a linear bias prescription (e.g. the “model independent” analysis of S25). Finally, when including bispectrum data, H25 made use of the coevolution relations of [289, 290] to relate the values of the higher-order bias parameters to b_1 . Any deviation from these relations would therefore result in additional bias in the inferred cosmological parameters.

Since, strictly speaking, the clustering measurements made in each redshift bin are only sensitive to the amplitude of matter fluctuations at the redshift of the galaxies that make up the sample, the quantity that is directly measured by our data is $\sigma_8(z) \equiv \sigma_8 D(z)$, where $D(z)$ is the large-scale linear growth factor². We present the tomographic constraints on this parameter obtained from our data in the bottom panel of Fig. 5.7. Our constraints are largely compatible with the growth history predicted by *Planck* (also shown in the figure), with the exception of the mild $\sim 2\sigma$ downwards deviation observed in the first redshift bins. This is also in agreement with other constraints on the growth history from CMB lensing tomography, including S25 (using the same galaxy sample), [230] using a variety of galaxy samples, [231, 291] using the unWISE catalog, [241, 242] using high-redshift quasars from *Quia*, [233]

²Note that, in the presence of massive neutrinos, structure growth is scale dependent at the linear level. The scale-dependent features induced by the value of M_ν assumed here ($M_\nu = 0.06$ eV) are negligible given our statistical uncertainties, however.




Figure 5.6.: Cosmological constraints. *Left*: two-dimensional posterior constraint from each redshift bin when analysing C_ℓ^{gg} and $C_\ell^{g\kappa}$ (dashed lines), and with the addition of the F²B measurements (filled contours). For reference, the *Planck* constraints are shown in yellow in all panels. *Right*: Constraints on σ_8 , Ω_m , and S_8 from the combination of the three highest redshift bins analysed here. Results are shown using only two-point information (C_ℓ^{gg} and $C_\ell^{g\kappa}$, turquoise), and with the addition of the bispectrum measurements ($\Phi_{LL\ell}^{ggg}$, dark blue).

combining photometric LRGs and the DESI Bright Galaxy Sample and, most recently, [234] using the 2MPZ and WISE $\times$ SuperCOSMOS galaxy samples at low redshifts. The addition of the galaxy bispectrum results in a ~ 10 -15% reduction in the statistical uncertainties of these measurements, particularly at the higher redshifts, accompanied by the slight downwards shift in the value of $\sigma_8(z)$ discussed earlier.

Finally, we combine the measurements in different redshift bins to place joint constraints on σ_8 and Ω_m . Given the mild tension between the constraints found in the first redshift bin and the rest, we choose our fiducial constraints to correspond to the combination of the last three redshift bins, although we will also present the results of combining all redshift bins. Fig. 5.6, right panel, shows these combined constraints, including those found using only two-point correlations, as well as those incorporating the galaxy bispectrum. Additionally, Fig. 5.8 shows the same constraints, as well as those arising from the combination of all redshift bins, in comparison with other results found in the literature. Our constraint on Ω_m is compatible with *Planck* within $\sim 1\sigma$. As shown in [234, 260, 292], this constraint is dominated by the large-scale broadband shape of the power spectrum, as constrained by C_ℓ^{gg} and $C_\ell^{g\kappa}$. Perhaps interestingly, our constraints on this parameter lie on the low side of the *Planck* measurements, as do most estimates derived from CMB lensing tomography at redshifts $z \lesssim 1$, as well as the latest BAO constraints on Λ CDM from DESI [31]. Nevertheless, the measurement uncertainties are too large to draw any definitive conclusion regarding a potential parameter tension.

In turn, our fiducial constraints on σ_8 are lower than *Planck* at the 2.4σ level when including power spectrum and bispectrum measurements. Compared to the 1.6σ deviation obtained using only two-point functions, this result is driven by the $\sim 15\%$ smaller error bars combined with the $\sim 0.5\sigma$ downward shift in σ_8 we discussed above. Combining the measurements from all redshift bins, this tension in the value of σ_8 rises to the 3σ level (2.5σ from two-point functions only), driven by the significantly lower amplitude of matter fluctuations favoured by the first redshift bin.

We can also derive constraints on the so-called ‘‘clumping parameter’’ $S_8 \equiv \sigma_8 \sqrt{\Omega_m/0.3}$, which can




Figure 5.7.: Tomographic constraints on σ_8 and growth. *Top*: constraints on σ_8 found in each redshift bin using only two-point information and adding the bispectrum measurements (thick points with error bars). These results are compared with those found from the analysis of the same data in S25 [232] and H25 [260], the latter assuming coevolution relations for the bias parameters. *Bottom*: tomographic reconstruction of the growth history in terms of $\sigma_8(z)$. The *Planck* best-fit prediction is shown in yellow for comparison.

be accurately constrained by cosmic shear analyses. Combining different redshift bins, we obtain:

$$S_8 = 0.769 \pm 0.020 \ (z_2 + z_3 + z_4), \quad S_8 = 0.756 \pm 0.018 \ (\text{all bins}). \quad (5.30)$$

Compared to the *Planck* constraints on this parameter ($S_8 = 0.832 \pm 0.013$), these measurements lie 2.6 and 3.4 σ lower, respectively. The tension in the latter case is driven by the inclusion of the first redshift bin, where the constraints on both Ω_m and σ_8 lie below the *Planck* values (see top left panel of Fig. 5.6). This level of tension is marginally higher than that found in [232] using lensing from *Planck* and ACT (2.3 σ), driven primarily by the shift induced by the addition of bispectrum data, as well as the different modelling choices. Fig. 5.8 shows our constraints on Ω_m and σ_8 in the context of other CMB lensing tomography analyses, including S25 [232], the analysis of [230] combining multiple galaxy samples, the constraints of [293] using an alternative sample from the Legacy Survey, the recent analysis of [234] using low-redshift data, and the high-redshift constraints of [292] using quasars from *Quaia*. The constraints found here are in broad agreement with these works, although it is interesting to note that the analyses showing the strongest tension in σ_8/S_8 with the CMB are those using LRG samples. Additionally, with the exception of [292], most lensing tomography analyses, in the absence of external BAO data, seem to favour a marginally low value of Ω_m . Given the current uncertainties, this could be simply a mild statistical anomaly, although misspecification in the non-linear bias model, or the presence of systematic contamination in the observed fluctuations in galaxy number counts, could also contribute to a low S_8 or Ω_m .

To further explore the limits of the bias model used here, and the ability of the bispectrum measurements to improve the final cosmological constraints, we have explored the possibility of using simpler bias models to describe the data. In particular, we considered two cases: one in which the counter-term parameters $\{c_m, b_{k^2}\}$ and the third-order bias parameter b_{3nl} are set to zero, and one in




Figure 5.8.: Comparison of our constraints on $\{\Omega_m, \sigma_8, S_8\}$ with past constraints derived from lensing tomography [230, 232, 234, 292, 293] as well as the *Planck* results.

which the higher-order bias parameters $\{b_2, b_s, b_{3nl}\}$ are set to zero. In the first case, the power spectra and bispectra are both described by the same set of bias parameters, thus maximising the information content in the bispectrum. The details of this study are discussed in Appendix 5.A.2, where we show that these simplified models are able to describe the data well over the same range of scales as our fiducial bias model. When applied to the real data, we (unsurprisingly) recover tighter constraints on both Ω_m and σ_8 , with a small shift towards the values preferred by *Planck*. Interestingly, as shown in Table 5.2, both models (indicated by the subset of parameters that we vary) are able to fit the data well, with a very small penalty in the best-fit χ^2 with respect to the fiducial model, commensurate with the reduced number of degrees of freedom. This suggests the additional complexity of the fiducial model is not required by the data. Nevertheless, a more detailed validation of these simpler models, and their ability to describe the clustering of generic galaxy samples, would be required before they can be used reliably in cosmological analyses.

5.4.3. Galaxy bias relations from projected clustering

Galaxy bias is a fundamental element of perturbation theory models for non-linear clustering. In this first-principles approach, the bias coefficients have to be treated as nuisance parameters in cosmological inference, introducing degeneracies that inevitably limit the precision of parameter estimation. Moreover, non-linear bias parameters notably contribute to the power spectrum in a largely degenerate way, making it challenging to independently constrain them using power spectrum alone in 3D analyses. This problem becomes more pronounced for projected clustering, where linear features such as the BAO wiggles are highly suppressed. In this context, the inclusion of higher-order correlation functions helps mitigate this limitation: the need to simultaneously describe two- and three-point functions helps break the degeneracies between different bias parameters, allowing us to measure them at higher precision. A better constraint on the higher-order bias parameters may then allow us to use the galaxy-galaxy and galaxy- κ power spectra on mildly non-linear scales to constrain cosmology.




Figure 5.9.: Impact of additional three-point information on the posterior of the linear and higher-order bias parameters $\{b_1, b_2, b_s\}$ that contribute to our model of the bispectrum (Eq. (5.17)). Results are shown for the two highest redshift bins, with similar results found for the first two bins. The inclusion of three-point information leads to a factor $\mathcal{O}(5 - 10)$ improvement in the FoM for this set of bias parameters. The coevolution relations obtained from [289, 290] for haloes are shown as dotted lines, while the relations found by [294] for galaxies are shown as dashed lines.

Fig. 5.9 shows the constraints on the bias parameters using only power spectra or power spectra and bispectrum measurements. Results are shown for the last two redshift bins, where we are able to use a larger number of angular modes at a fixed k_{max}^B and where, in consequence, this improvement is more evident. Quantitatively, the inclusion of bispectrum information, even on the conservatively large scales used here, leads to a remarkable $\mathcal{O}(10)$ improvement in the FoM for the three bias parameters $\{b_1, b_2, b_s\}$, with b_2 and b_s experiencing the strongest reduction in their statistical uncertainties. Importantly, the constraints obtained from the full data vector are compatible with those obtained from power spectra alone, with the bispectrum favouring the region of parameter space with larger values of b_1 and b_2 (and marginally lower b_s). This trend leads to the mild downwards shift in σ_8 observed in the previous section when including bispectrum data (as also observed in [288]). Also, and reassuringly, the best-fit bias parameters $\{b_1, b_2, b_s\}$ found using power spectrum data alone provide a reasonable prediction for the bispectrum measurements (blue lines in Fig. 5.5), demonstrating the internal self-consistency of our data and model. Having improved constraints allows us to examine the relationships between the different bias parameters. In particular, measurements of the clustering of galaxies and haloes in N -body simulations have found that the higher-order bias parameters exhibit a correlation with the linear bias b_1 . These often called “coevolution relations” are not completely universal, differing slightly between haloes and galaxies, and showing a non-negligible scatter. Nevertheless, it is interesting to confirm whether the clustering of galaxies observed in real data also obeys these relations. Fig. 5.9 shows the relations found by [289] and [290] for haloes (dotted lines), and by [294] for galaxies (dashed lines). The preference of the bispectrum data for a larger b_2 pushes our constraints marginally away from the coevolution relations. This is shown in more detail in Fig. 5.10, which collects the constraints found in all redshift bins when analysed jointly. With the inclusion of bispectrum data, our constraints




Figure 5.10.: Relation between the quadratic (b_2 , *left*) and tidal (b_s , *right*) bias parameters and the linear bias b_1 found in each redshift bin. Results are shown using only two-point information (*thinner, shaded markers*), and with the inclusion of bispectrum data (*thicker markers*). These constraints are compared against the coevolution relations for haloes ([289, 290], *dotted lines*), and galaxies ([294], *dashed lines*).

show a mild preference (at the $\sim 2\sigma$ level) for upwards (downwards) deviations in b_2 (b_s) for larger values of b_1 . These deviations are not entirely unexpected, given the relatively large scatter found in the case of galaxies [294–296]. Note also that our constraints use data from a single galaxy type (LRGs), spanning a relatively narrow range of b_1 values.

5.5. Concluding remarks

In this chapter, we have carried out the first comprehensive cosmological analysis of lensing tomography incorporating information from the angular galaxy bispectrum, complementing the first tentative analysis of [260]. A key aspect of such an analysis is the high redundancy of the full data vector, forcing our model of the galaxy-matter connection to simultaneously describe the two- and three-point clustering of galaxies as well as their cross-correlation with the matter overdensities. This redundancy improves both the precision and robustness of the resulting cosmological constraints.

We have used a perturbative bias expansion within the Effective Field Theory of LSS framework, modelling the galaxy-galaxy and galaxy-convergence power spectra at one-loop order, and the galaxy bispectrum at tree level. We have validated this model against N -body-based mock observations, and determined the scales over which it can safely recover cosmological constraints. Compared to studies of the three-dimensional clustering of galaxies, this task is made simpler in the case of projected statistics due to the negligible impact of non-linear redshift-space distortions. We have shown that the model is able to recover unbiased constraints down to scales where cosmological information is in fact saturated for the LRG sample studied here.

Applying this model to the analysis of the four different redshift bins considered here, we find constraints on σ_8 and Ω_m that are consistent with past works analysing this sample, and have shown that the inclusion of the galaxy bispectrum leads to a ~ 10 -15% improvement on the cosmological constraints. The constraints on Ω_m and σ_8 found from the three highest-redshift bins are compatible with each other and in reasonable agreement with *Planck*, as is the growth history reconstructed from these data (see Fig. 5.7). In turn, the first redshift bin shows a preference for lower σ_8 , and is in tension with *Planck* at the $\sim 2.5\sigma$ level.

We have then derived cosmological constraints from the joint analysis of the last three redshift bins. The measured value of σ_8 (S_8) is moderately lower than the *Planck* value at the 2.4σ (2.6σ) level. This is due to the mild but coherent preference for a lower σ_8 displayed by the two-point clustering of galaxies in all redshift bins, combined with the additional $\sim 0.5\sigma$ downwards shift observed when including the galaxy bispectrum. Interestingly, a similar shift has been observed in the three-dimensional clustering of galaxies when combining power spectrum and bispectrum data [288]. This tension is then exacerbated, above the 3σ level, when combined with the first redshift bin. In agreement with the results obtained in our per-bin analysis, the inclusion of the projected bispectrum leads to a $\sim 10\%$ improvement in the statistical uncertainties on cosmological parameters.

Including the bispectrum has a much stronger impact on the higher-order bias parameters of our model, b_2 and b_s , with a reduction in their statistical uncertainties of more than a factor 2 in some cases. This allows us to study the relations between the different bias parameters with higher precision, as well as their agreement with coevolution relations found in simulations. The constraints obtained after including bispectrum information are in good agreement with those from power spectra alone, but show a consistent preference for higher b_1 and b_2 , and lower b_s . This manifests itself as a mild departure from the coevolution relations of [294] for high values of b_1 . This is not entirely surprising, however, given the large scatter expected around these relations.

Adding bispectrum data improves the robustness and precision of the resulting cosmological constraints. The improvement in the latter case is relatively disappointing, however ($\sim 10\text{-}15\%$), and it is interesting to reflect on the potential reasons for this. First, although the bispectrum is able to significantly improve the constraints on b_2 and b_s when modelled at the tree level, the counter-term parameters $\{b_{k^2}, c_m\}$ and the third-order parameter b_{3nl} are only constrained by the power spectrum measurements. This leaves significant freedom to describe the featureless clustering of galaxies on mildly non-linear scales. In Appendix 5.A.2 we have explored the possibility of using a simpler bias model, setting $\{b_{k^2}, c_m, b_{3nl}\}$ to zero, and thus allowing us to maximally exploit the information contained in the bispectrum measurements. Interestingly, we have shown that this simplified model is still able to recover unbiased constraints in mock data and, when applied to the data, the aforementioned tensions are somewhat relieved. Nevertheless, the relative improvement in the final constraints from including the bispectrum is still limited to $\sim 10\%$, suggesting that the limiting factor is in fact not the additional freedom allowed by the counterterm parameters and b_{3nl} .

Another potential cause for the mild improvement from the bispectrum is the large uncertainties in $C_\ell^{g\kappa}$ and $\Phi_{LL\ell}^{ggg}$ relative to the galaxy power spectrum. Both of these measurements are crucial to break degeneracies in the bias model and constrain cosmology, and the final constraints may be limited by these uncertainties, rather than freedom in the model itself. Future work could improve both of these aspects of our analysis using existing data. First, replacing the CMB lensing convergence with cosmic shear data from existing imaging surveys [297, 298] could improve the sensitivity of the galaxy-lensing power spectrum. It could also allow us to exploit the galaxy-galaxy-lensing bispectrum, which we have discarded in this analysis due to the low signal-to-noise ratio on the large scales over which it can be modelled accurately, as shown in [260]. Secondly, our measurements of the angular bispectrum have been based on the F^2B estimator presented in [259], which only recovers triangle configurations that are close to isosceles. As discussed in Section 5.2.2 and [259], the F^2B estimator can be easily generalised to recover the remaining non-zero bispectrum configurations, thus increasing the total three-point statistical power. Finally, our use of the bispectrum is limited by the range of scales over




Figure 5.11.: *Left*: contribution of RSDs to the galaxy power spectrum on large scales. The vertical dotted line represents the large-scale cut that we adopt in our fiducial analysis. *Right*: systematic shift on the inferred Ω_m parameter for the LRGs z_3 bin only when ignoring the RSD contribution. Using our fiducial large-scale cut, the shift is relatively minor ($\sim 0.3\sigma$).

which our tree-level model is sufficiently accurate. In the future, this could be significantly improved by resorting to more sophisticated approaches to galaxy bias, such as hybrid effective field theory (HEFT, [299–301]), which could remain accurate on significantly smaller scales without a substantial increase in the number of bias parameters required.

This work demonstrates that a joint analysis of the projected power spectrum and bispectrum of galaxies and weak lensing can be efficiently and reliably carried out with existing data. The deployment of these methods could also enhance the science returns of future imaging surveys, such as the Rubin Observatory [12, 245], maximising the information extracted from the projected clustering of galaxies.

Appendix 5.A Additional tests

5.A.1 Treatment of Redshift-space distortions

Redshift-space distortions are the most relevant secondary effect modifying the galaxy angular power spectrum on large scales, particularly for relatively narrow redshift bins [274, 275]. As described in Section 5.3.2, we include RSDs at the linear level in the galaxy auto-spectrum, including only the correlations between the RSD term and the linear density term proportional to b_1 and the RSD auto-correlation. Given the results of our simulation-based validation, this is sufficient to obtain unbiased cosmological constraints from our data. Nevertheless, to gain a better intuitive understanding of the impact of RSDs on our analysis, the left panel of Fig. 5.11 shows the large-scale power spectrum with and without the RSD contributions. Within the scale cuts used here, RSDs lead to deviations in the power spectrum amplitude of up to $\sim 10\%$, peaking on the largest scales used and becoming negligible at high ℓ .

Repeating our analysis omitting these contributions, the right panel of Fig. 5.11 shows that the largest effect is a relatively mild ($\sim 0.3\sigma$) downwards shift in the preferred value of Ω_m . Thus, although relevant for this analysis, a precise modelling of RSDs (e.g. including the impact of non-linear “fingers of God”) is not required.




Figure 5.12.: Constraints on Ω_m and σ_8 obtained from the ABACUSSUMMIT simulation, combining all redshift bins, with simpler bias models. Our baseline constraints are shown as black lines. The brighter and darker contours show the constraints found using only two-point functions and combining them with the bispectrum, respectively. The *left panel* shows the constraints found using a simpler bias model in which the scale-dependent counterterms $\{b_{k^2}, c_m\}$, and the third-order parameter b_{3nl} are set to zero. The right panel shows the constraints found setting b_2, b_s , and b_{3nl} to zero instead, in a model inspired by the phenomenological parametrisation of [239]. Both restricted parametrizations are able to find unbiased constraints on the cosmological parameters.

5.A.2 Simpler bias parametrisations

The bias model used here (see Section 5.3.2), in which power spectra are modelled at the 1-loop level including EFT counterterms, and the bispectrum is modelled at the tree level, involves free parameters, such as c_m, b_{k^2} , and b_{3nl} , which can only be constrained by a single type of correlation (e.g. two-point functions, but not the bispectrum). Although this level of freedom is theoretically necessary and physically motivated, it is interesting to explore whether simpler versions of this model are still able to provide unbiased cosmological constraints. Furthermore, recent works [239] have shown that, in the case of real-space galaxy clustering, a simplified phenomenological parametrization of galaxy bias is still able to provide an adequate and unbiased description of the data. In this section we explore how incorporating bispectrum information impacts parameter constraints when working within two restricted bias parametrizations.

Specifically, we test two complementary simplified models. First, we explore a parametrization that only retains the parameters that enter in Eq. (5.17), i.e. varying $\{b_1, b_2, b_s, \alpha_P\}$, while fixing the counterterm values, and third order Eulerian bias, to zero. Then, we also tested a theoretical model that, in the spirit of [239], effectively only parametrizes the deviations from a scale-independent linear bias through counterterm-like coefficients. In our setup, we achieve this by only varying the subset of nuisance parameters $\{b_1, b_{k^2}, c_m, \alpha_P\}$, fixing to zero the remaining bias parameters³.

We begin by testing its performance on simulated data. As shown in Fig. 5.12, we find that both simplified parametrizations are able to recover the fiducial values of the cosmological parameters when

³Note that the model in [239] also allows for a stochastic noise component in the galaxy-matter power spectrum, making it more general than the model studied here.




Figure 5.13.: As Fig. 5.12, now showing the constraints found when analysing the real data. The *Planck* constraints are also shown for comparison.

applied to the combination of all redshift bins, both using only two-point correlators and including the bispectrum. This is likely due to the relatively featureless nature of the projected galaxy statistics, making it possible to describe them with a relatively small number of smooth “basis” functions. Notably, in this simulated data, the addition of *ggg* information is perfectly compatible with the 2-point posterior, tightening the constraints around the (Ω_m, σ_8) fiducial point, improving their statistical uncertainties by (26%, 8%) using the $\{b_1, b_2, b_s, \alpha_P\}$ model (left panel) and (17%, 24%) using the $\{b_1, b_k^2, c_m, \alpha_P\}$ model (right panel).

Then, having validated both on simulations, we apply them to the analysis of the real clustering data. Fig. 5.13 shows the posterior distribution of the cosmological parameters, as well as the constraints found using our baseline model for reference. We find that, using the $\{b_1, b_2, b_s, \alpha_P\}$, the uncertainty on Ω_m shrinks by $\sim 25\%$ with respect to our baseline model, while the σ_8 remains largely unchanged. The value also favours a slightly lower value of Ω_m and a larger value of σ_8 , reducing the tension with *Planck* found in our baseline analysis. A similar shift is found in the case of the $\{b_1, b_k^2, c_m, \alpha_P\}$ model, again relieving this tension. The inclusion of the bispectrum in this model also leads to a significant $\sim 20\%$ reduction in the σ_8 uncertainty.

It is worth stressing the fact that these two simpler bias models, defined as subsets of the full EFT bias model presented, are not, strictly speaking, well motivated or complete from a physical point of view. Thus, their ability to describe our data, as tested on simulations, might be purely accidental and only applicable to the specific galaxy sample we have simulated, which would lead to artificially reduced cosmological constraints. Therefore, a more comprehensive validation of the models, considering a large suite of non-linear galaxy clustering scenarios, would be necessary before they can be safely applied to cosmological analyses.

5.A.3 Emulators accuracy

The approach of emulating, rather than computing *ab initio* at each likelihood step, the theory templates enables fast and feasible MCMC sampling as detailed in Sect. 5.3.3. While essential, this additional step may introduce systematic errors; therefore, the accuracy of the emulators must be carefully validated.




Figure 5.14.: Emulating F^2B templates. Training history (*left*) of the emulator for the $F_{LL\ell}^{ggg}$ (first filter) template, and performance (*right*) on some elements of the test set. All the emulators for the other templates present an analogue training history and a similar performance.

To this purpose, we focus on the emulator of the $F_{LL\ell}^{ggg}$ template as a representative example, as it exhibits consistent behaviour with the others. Specifically, in Fig. 5.14 we show the evolution of the loss function (mean squared errors) and of our accuracy metric throughout the training epochs. We find that in a few hundred epochs the loss function improves by almost 5 orders of magnitude, and virtually 100% of the output layer neurons differ less than 0.5% from the ground truth (which prompts us to stop the training). We repeat this exercise for both the training set and the test set (which is never seen by the Neural Network while training), finding similar results and no indication of overfitting.

Finally, in the right panel of Fig. 5.14, we show the overall accuracy of this emulator, in terms of percentage differences between a subsample of the $F_{LL\ell}^{ggg}$ template values used in the test set and their corresponding emulator predictions. We find that the emulator reproduces the test set at better than 1% over the whole considered ℓ range, with no significant scale dependence. We find similar performance for the emulators of all of the templates of section 5.3.3, which is accurate enough given the statistical error of the dataset used in this work.

6. DESI DR1 Full-Shape Constraints on Ultra-Light Axions

Building upon the theoretical model developed in Chapter 4, in this chapter we present novel constraints on ultra-light axions from the combination of CMB, including lensing, and DESI DR1 full-shape data.

6.1. Motivation

In Chapter 4 it was shown that there exists a mass window of about 8 orders of magnitude, $10^{-32} \text{ eV} \leq m_a \leq 10^{-24} \text{ eV}$, where a DM abundance made of ULAs can be probed with a joint analysis of CMB and spectroscopic galaxy clustering. There, we provided refined bounds on already exploited BOSS data using the novel two-fluid EFTofLSS framework.

The latest Data Release 1 (DR1) from DESI [10, 201] provides a timely opportunity to update that analysis with new data. Comprising spectroscopic measurements of luminous red galaxies (LRG), emission line galaxies (ELG) and quasars (QSO), as well as a sample of bright galaxies (BGS), DR1 maps nearly 6 million LSS tracers spanning an effective cosmic volume of approximately 18 Gpc^3 . This is an improvement by a factor of ~ 3 compared to BOSS DR12 [9] data. The official DESI analysis [38, 115] and later reanalyses [288, 302, 303] of DR1 full-shape data within the EFTofLSS framework have demonstrated its remarkable constraining power, but have so far been restricted to the baseline ΛCDM model and a few of its standard extensions. Given the pristine quality of the data, there is strong motivation to push these boundaries and exploit this dataset as a laboratory to probe fundamental physics.

In this chapter, we therefore present the first application of DESI DR1 full-shape data to constrain non-cold components of the dark sector, adopting again ULAs as a benchmark scenario. Rather than relying on standard pipelines hard-coded for ΛCDM , we perform an independent reanalysis of the DESI dataset employing the self-consistent EFTofLSS framework explicitly tailored to capture the signatures of mixed (i.e., cold + non-cold) DM models developed in Chapte. 4. Furthermore, we also complement to the previous analysis with ACT DR6 [6, 24, 143, 210] data.

This Chapter organized as follows. Section 6.2 then presents the datasets employed in the present study, and the methods with which they were analyzed. Our main results, including updated upper limits on the abundance of an ULA DM component as a function of the ULA mass, are presented in Sect. 6.3. We focus first on CMB data alone, then on the DESI DR1 full shape, and finally on their combination. Section 6.4 highlights the outcomes of this work and outlines future directions.

6.2. Data and Methods

6.2.1. Datasets

CMB We employ the state-of-the-art combination of ACT Data Release 6 (DR6) [24, 143, 210] and Planck [5, 304] legacy CMB datasets, supplemented by reconstructed lensing measurements. The optimal primary CMB combination that we consider, dubbed P-ACT in [6], merges the high-resolution temperature and polarization measurements from ACT with the large-scale full-sky observations from Planck. To avoid overlap, the Planck high- ℓ likelihood is explicitly truncated at $\ell_{\max} = 1000$ (600) for TT (TE, EE) [6].

Analogously, we make use of the CMB lensing combination of Planck PR4 + ACT DR6 in the same setup of the official ACT analysis [24, 210, 305]. While Planck lensing covers a larger sky area, ACT probes smaller scales with better precision; we consider the baseline range of multipoles $40 < L < 763$, while fitting the extended spectrum (up to $L = 1300$) [306] is left to future work, for reasons outlined in what follows.

DESI DR1 full-shape Starting from the DESI DR1 catalogs [201] we measured the power spectra with `pypower`¹ [307], which implements the θ -cut, removing all pair counts with angular separation $\theta < 0.05^\circ$ [308] to mitigate fiber assignment incompleteness. The numerical covariances have been measured from the approximate EZmocks catalogs [227]. The mixing matrices relevant for the convolution of the theory model with the survey mask, are then obtained following standard methods [79, 227] with the `convo`² code.

For each tracer, we make use of the $P_{gg,\ell}$ multipoles with $\ell = 0$ and 2, discarding the hexadecapole due to the poor signal-to-noise ratio. We adopt a fixed scale cut of $k_{\max} = 0.2 \, h \, \text{Mpc}^{-1}$ as in previous analyses [115, 227, 288]. Compared to the official analysis, we do not apply a further rotation to the data vector, as it has been observed that the quite non-diagonal structure of the window matrix does not affect inference [227, 308], performing the convolution up to $k_{\text{conv}} = 0.35 \, h \, \text{Mpc}^{-1}$ as discussed in [308]. We anyway tested the stability of our main results for the joint analysis against further increases of k_{conv} , confirming the lack of appreciable differences in the cosmological parameter posteriors.

6.2.2. CMB theoretical predictions

As in Sect. 4.4, linear power spectra are computed using the Boltzmann solver `axiCLASS` [184, 185]. However, CMB lensing is sensitive to nonlinearities in P_{mm} . The previous ULA analyses based on Planck+BOSS data [204] (see also [189, 209]) computed the lensing potential power spectrum using *linear* predictions, which was shown not to bias the results for Planck errorbars. The ACT DR6 four-point function measurements from which the lensing potential is extracted, however, probe smaller (hence, more nonlinear) scales with better precision, therefore this aspect needs to be revisited [208, 312, 313].

The left panel of Fig. 6.1 compares two different predictions for the lensing convergence power spectrum, $C_L^{\kappa\kappa}$, both of which include nonlinear corrections through halo model-based fitting functions: the standard Λ CDM-based HMcode [309] (to which we give the total P_{mm}^L , comprising the ULA component, as input) and `axionHMcode`, which was developed in [310] specifically to include ULA

¹<https://github.com/cosmodesi/pypower>

²<https://jacoposalvalaggio.gitlab.io/convo>




Figure 6.1.: Fractional difference between C_L^{KK} computed with different prescriptions for the nonlinear corrections to P_{mm} , and the result obtained using the standard HMcode [309]. The ACT DR6 errorbars are also shown for reference. Two values of the ULA mass are considered, while f_{ax} is set close to the upper limit allowed by the data in each case (see Table 6.2). The extended range $763 < L < 1300$, not included in the present study, is shaded in gray. *Left:* $m_a = 10^{-27}$ eV. The solid blue line employs axionHMcode in the version of [310] for the nonlinear corrections. We do not show the version of [311], because it was calibrated to simulations at the much larger mass $m_a = 10^{-24.5}$ eV. The linear result is reported in orange. *Right:* $m_a = 10^{-25}$ eV. Solid blue and green lines employ the two different versions of axionHMcode [310, 311] for the nonlinear corrections. Dashed lines correspond to our EFTofLSS predictions, with the counterterm $c_m(0)$ adjusted to reproduce the axionHMcode results.

nonlinearities. The ACT DR6 lensing errorbars are also displayed, for reference. For sufficiently light ULAs, $m_a \lesssim 10^{-27}$ eV, the fraction f_{ax} is constrained to be subpercent and, as a consequence, genuine ULA nonlinearities are negligible for ACT scales and errorbars. On the other hand, we see that a linear prediction is not sufficient, as it entails a systematic error of the order of the statistical one.

Larger ULA masses, for which the allowed fraction is larger, pose a theoretical challenge in the description of genuine ULA nonlinearities. This is shown in the right panel of Fig. 6.1, taking $m_a = 10^{-25}$ eV, for which $f_{\text{ax}} = 15\%$ is allowed by the combination of CMB (excluding lensing) and DESI data. We see that, as already noted in [312], the two axionHMcode halo-model implementations available for this heavier ULA regime [310, 311] yield different predictions. Furthermore, [312] has shown that cosmological inference of $\Omega_a h^2$ is highly sensitive to the choice of the halo-model-based nonlinear prescription, already with the baseline scale window $L < 763$. This can even produce an artificial preference for ULAs with $m_a \approx 10^{-24}$ eV [312]. Hence, nonlinear modeling is a theoretical systematic uncertainty that, at present, challenges a robust exploration of the $m_a \gtrsim 10^{-25}$ eV regime.

In this work, we therefore adopt a perturbative modeling of CMB lensing, focusing on multipoles with $L < 763$, for which the bulk of the CMB lensing amplitude lies on quasi-linear scales. The 1-loop prediction for the matter power spectrum in the framework of the EFTofLSS is

$$P_{mm}(k, z) = P_{mm}^L(k, z) + P_{mm}^{1\text{-loop}}(k, z) - 2c_m(z)k^2 P_{mm}^L(k, z), \quad (6.1)$$

where $P_{mm}^{1\text{-loop}}(k, z)$ is the next-to-leading-order correction, while the last term is proportional to the EFT coefficient (counterterm) c_m that absorbs any small-scale dependence, beyond the perturbative regime, of the loops. Notice that standard CDM and baryons are treated as indistinguishable on large




Figure 6.2.: Additional test, performed in Λ CDM and given ACT DR6 errorbars, concerning the time-dependence of the counterterm for our perturbative modeling. We display two different values for the exponent of Eq. (6.2), $p = 3$ and $p = 1$, showing that the fit is not very sensitive to this assumption. The linear and the bare 1-loop ($c_m = 0$) predictions are also reported.

scales. The challenge is that the broad redshift span of the CMB lensing kernel requires knowledge of the time dependence of c_m . To address this, several approaches have been put forward [173, 314–317] that still allow one to use the bulk of information on large scales from perturbation theory. Here, as a proof of principle, we show that for ACT errorbars the 1-loop model above is already sufficient to capture these deviations from linear theory. Following [314], we assume

$$c_m(z) = c_m(0)D(z)^p, \quad (6.2)$$

where $D(z)$ is the linear growth factor, normalized to unity at $z = 0$. We take $p = 3$ as motivated by comparison with N-body simulations [314]. In the right panel of Fig. 6.1 we observe that by adjusting the value of $c_m(0)$ (red and dark green dashed lines) we are able to reproduce both halo-model prescriptions that include ULA nonlinearities, shown by the blue and light green solid lines. In this respect, our analysis is more conservative than the recent one in [313], which employed the halo model, as we marginalize over one extra free parameter, namely the value of $c_m(0)$.

Concerning instead the p parameter appearing in Eq. (6.2), we fixed $p = 3$ as motivated by numerical studies [314]. We tested that other values of p produce, for what concerns $C_L^{\kappa\kappa}$, an equally good fit to the HMcode prediction when letting $c_m(0)$ free to vary, as shown in the left panel of Fig. 6.2 assuming a Λ CDM cosmology. This is expected, as for the range $L < 763$ considered in this work the lensing power spectrum integral is dominated by a limited range in redshift, rendering the precise time dependence of c_m a minor effect. This statement may need to be revised if the lensing analysis is extended to smaller scales, or new data become available with significantly smaller errorbars. In summary, for $m_a \geq 10^{-26}$ eV we adopt the conservative EFTofLSS modeling, while for lighter ULAs ($m_a \leq 10^{-27}$ eV) the standard HMcode is sufficient, due to the smallness of the allowed f_{ax} .

Achieving a computationally efficient, yet accurate, evaluation of $P_{mm}^{1\text{-loop}}(k, z)$ in Eq. (6.1) is a separate challenge, as the presence of the ULA Jeans scale increases the complexity of cosmological perturbation theory beyond linear order. For this purpose we again apply our results of Chapter 4, where we presented the extension of the EFTofLSS to general mixed DM scenarios, including cosmologies with a DM subcomponent made of ULAs. In particular, we use the prescription of Eq. (4.74) to compute the ULA density component in the total matter, leaving no further ambiguities in Eq. (6.1).

6.2.3. Galaxy clustering model

To analyze the DESI full-shape spectra we will employ the refined modelling, based on a two-fluid extension of the EFTofLSS, developed in Chapter 4. That is, we compute non-linearities for the non-cold field using Eq. (4.74) and then add the further bias parameter as in Eq. (4.73) since galaxies in full generality may trace also this subcomponent, as was extensively discussed. This was applied in Chapter 4 to BOSS data as a first proof-of-principle study for this model, and to compare to previous one-fluid analysis.

Concerning, now, DESI data, we fit the monopole and quadrupole of the galaxy power spectrum up to $k_{\max} = 0.2 \, h \, \text{Mpc}^{-1}$ for all six redshift bins. This matches the baseline choice of previous analyses [115, 227] and [288]; the latter reference also includes the hexadecapole, which however has a very low signal-to-noise ratio. We do not perform a combination with post-reconstruction BAO information.

Our theoretical model, which is nearly identical to the one used for BOSS and we repeat here for convenience and to highlight the few differences, is given by

$$P_{gg,\ell} = P_{gg,\ell}^L + P_{gg,\ell}^{1\text{-loop}} + P_{gg,\ell}^{\text{ctr}} + P_{gg,\ell}^{\text{SN}}, \quad (6.3)$$

Concerning the stochastic contribution, before expanding in multipoles it is now given by [227]

$$P_{gg}^{\text{SN}}(k, \mu) = \frac{1}{\bar{n}} \left[P_{\text{shot}} + a_2 \mu^2 \left(\frac{k}{k_{\text{NL}}} \right)^2 \right], \quad (6.4)$$

as the Poissonian prediction is already subtracted from the measurements. Here $\mu \equiv \mathbf{k} \cdot \hat{\mathbf{z}}/k$, where $\hat{\mathbf{z}}$ is the unit vector along the line of sight. We set $k_{\text{NL}} = 0.45 \, h \, \text{Mpc}^{-1}$, although its exact value is not important, as a_2 is marginalized over with a wide prior.

The theoretical model is then convolved with the survey window function, accounting also for the θ -cut, and compared to data with a Gaussian likelihood. We employ the covariance numerically computed from the EZmocks mocks [318] and rescaled by the appropriate factor to match the analytic-based estimates [201, 227]. The likelihood is computed, again, within the PBJ³ [62] code, embedded in the Cobaya framework [206] for the purpose of combination with the other datasets, and then sampled with the Nautilus sampler [106]. We tested our pipeline by comparing Λ CDM results against both the official [227] and independent [288] analyses, finding remarkable consistency, both for single redshift bins as well as for the joint case.

6.2.4. Parameters and priors

In the analysis we consider a set of ULA mass values spaced logarithmically in the range $[10^{-32}, 10^{-24}] \, \text{eV}$, and for each m_a we vary the parameters reported in Table 6.1.

Concerning the cosmological parameters, when combining with the CMB we use wide uninformative priors on all of them, as they are well constrained by the data. In the DESI-only analysis, as in [227], we instead set a Big Bang Nucleosynthesis Gaussian prior on $\Omega_b h^2 \sim \mathcal{N}(0.02218, 0.00055^2)$, and an informative, but still wide, prior on $n_s \sim \mathcal{N}(0.9649, 0.0422^2)$, while fixing $\tau_{\text{reio}} = 0.0544$. The sum of neutrino masses is always fixed to its minimum allowed value, $\sum m_\nu = 0.06 \, \text{eV}$. While the possible interplay of light scalar field dynamics and neutrino mass inference is receiving attention in light of

³<https://chiaramoretti.gitlab.io/pbj>

Cosmological						
$\Omega_a h^2$	$\Omega_{\text{cdm}} h^2$	$\Omega_b h^2$	h	$\log(10^{10} A_s)$	n_s	τ_{reio}
CMB nuisance (sampled)			EFTofLSS nuisance (sampled)			
parameter	prior/value		parameter	prior		
A_{ACT}	$\mathcal{N}(1, 0.03^2)$		$b_c \sigma_{8,c}$	$[0, 3]$		
p_{ACT}	$[0.9, 1.1]$		$b_a \sigma_{8,c}$	$[-3, 3]$		
A_{Planck}	A_{ACT}		$b_2 \sigma_{8,c}^2$	$\mathcal{N}(0, 5^2)$		
$c_m(0)$	$\mathcal{N}(0, 10^2)$		$b_{\mathcal{G}_2} \sigma_{8,c}^2$	$\mathcal{N}(0, 5^2)$		
EFTofLSS nuisance (analytically marginalized)						
parameter	prior			parameter	prior	
$c_{cc,\ell}$	$\mathcal{N}(0, 30^2)_{\ell=0}, \mathcal{N}(30, 30^2)_{\ell=2}$			b_{Γ_3}	$\mathcal{N}(0, 5^2)$	
$c_{ca,\ell}$	$\mathcal{N}(0, 60^2)$			P_{shot}	$\mathcal{N}(0, 1^2)$	
$c_{aa,\ell}$	$\mathcal{N}(0, 60^2)$			a_2	$\mathcal{N}(0, 5^2)$	

Table 6.1.: Parameters and priors used in the analysis, for each assumed value of the ULA mass m_a . For the cosmological parameters we always use wide uninformative priors, with the exception of $\Omega_b h^2$, n_s and τ_{reio} when not combining with the CMB, as in [227]. The counterterms $c_m(0)$, $c_{cc,\ell}$, $c_{ca,\ell}$ and $c_{aa,\ell}$ have dimensions of $h^{-2} \text{Mpc}^2$. Since we fit the monopole and quadrupole of the galaxy power spectrum, the index ℓ in the table runs over 0 and 2.

current tensions [319], we note that allowing the neutrino masses to take larger values would yield tighter constraints on the ULA density.

In the baseline analysis, we assume wide Gaussian priors for the bias and counterterm parameters related to the axion component, analogously to the cold fluid. However, following the discussion of Sect. 4.4, we also consider $\mathcal{O}(f_{\text{ax}})$ theoretical priors on the ULA parameters,

$$\begin{aligned}
 b_a &\sim \mathcal{N}(0, (b_c^{\Lambda\text{CDM}} f_{\text{ax}})^2), \\
 c_{ca,\ell}, c_{aa,\ell} &\sim \mathcal{N}(0, (60 f_{\text{ax}})^2),
 \end{aligned} \tag{6.5}$$

where f_{ax} is fixed at the conservative 95% CL upper bound reported in Table 6.2, while $b_c^{\Lambda\text{CDM}}$ is the linear bias parameter inferred from a ΛCDM analysis.

6.3. Results

The results for the analysis setup discussed above are presented first for the CMB alone, followed by the DESI DR1 full shape as the main novelty of this work, and finally their combination. As we never find a




Figure 6.3.: Marginalized posteriors for Planck alone, Planck + ACT, and Planck + DESI full-shape data. *Left:* $m_a = 10^{-25}$ eV. The ULA field begins to redshift as nonrelativistic matter early in radiation domination, hence $\Omega_a h^2$ shares a strong degeneracy with $\Omega_{\text{cdm}} h^2$ and the ULA signatures are mainly in perturbations. To remain agnostic over theoretical systematics, nonlinearities for ACT CMB lensing have been modeled in a perturbative EFTofLSS approach. *Right:* $m_a = 10^{-29}$ eV. The ULA field behaves as a dark energy component until well after recombination, therefore it is mainly constrained by the modification of the distance to last scattering, as evidenced by the degeneracy between $\Omega_a h^2$ and h .

preference for deviations from pure Λ CDM in the joint analysis, the final results are summarized in Fig. 6.5 as an exclusion plot in the $\Omega_a h^2$ vs. m_a plane. The corresponding numerical values are reported in Table 6.2.

Planck and ACT DR6 Figure 6.3 shows the posterior distributions for the CMB-only dataset given by the combination of Planck and ACT DR6, including lensing. In the left panel we take $m_a = 10^{-25}$ eV, i.e., a case where the ULA background evolution transitions to nonrelativistic matter early in radiation domination. As such, lensing is expected to be the main source of constraining power for this data combination. As argued in Sect. 6.2.2, matter nonlinearities are relevant for CMB lensing predictions and when the allowed ULA fraction f_{ax} is rather large, the axion’s own nonlinear effects may also become important for properly evaluating the tails of the posterior. For this reason we adopted a treatment based on perturbation theory, which can effectively marginalize over the theoretical systematics due to different halo-model assumptions. Compared to employing an HMcode, we found that the perturbative treatment widens the constraints on $\Omega_a h^2$ (whereas the impact is negligible for the other cosmological parameters). This can be explained by the degeneracy between the ULA energy density and the counterterm: notice that, according to Eq. (6.1), a negative $c_m(0)$ counteracts the shape suppression caused by the ULAs.

We also comment on $m_a = 10^{-24}$ eV, the largest mass considered in our analysis. In this case the ULA subcomponent is even less constrained by our conservative modeling of reconstructed lensing for the baseline scale cut of ACT ($L_{\text{max}} = 763$). Nevertheless, we include this point in Fig. 6.5 as a dashed line, as an indication of the sensitivity of the dataset, albeit with a possibly large theoretical error.




Figure 6.4.: Results involving the full shape of DESI DR1 galaxy power spectra. *Left*: The full-shape constraint for a model with $m_a = 10^{-25}$ eV is compared against a combination of “geometry” probes given by CMB primaries + BAO. This shows that for the largest ULA masses the shape suppression of the galaxy power spectrum is the prime source of constraining power. *Middle*: Single-redshift-bin results for ULAs with $m_a = 10^{-26}$ eV. Two LRG bins hint at a preference for a non-zero value of $\Omega_a h^2$. However, the statistical significance is mild, gets further reduced in the joint-bin analysis and disappears when combining with the CMB. *Right*: Comparison of the constraints obtained for $m_a = 10^{-27}$ eV by employing wide agnostic priors, vs. theoretically-motivated $\mathcal{O}(f_{ax})$ priors, on the ULA linear bias parameters and counterterms. See Sect. 6.2.4 for further details.

We further note that our Planck+ACT constraints on $\Omega_a h^2$ for $m_a = 10^{-26}$ (10^{-25}) eV are weaker by a factor of 1.6 (2.6) compared to the results of [313], which employed the version of axionHMcode from [311] with an extended scale cut ($L_{\max} = 1250$). While the analysis presented here is agnostic to different halo-model assumptions, we leave for future work a more direct comparison of the two approaches, for instance by considering an informative prior on the c_m nuisance parameter.

In the right panel of Fig. 6.3 we consider much lighter ULAs, with $m_a = 10^{-29}$ eV, which behave as a dark energy component until $z_{m_a} \approx 250$, well after recombination. In this case, the constraining power is dominated by primary CMB anisotropies, especially through the impact on the angular size of the sound horizon at recombination, θ_* , and on the integrated Sachs-Wolfe effect at very low ℓ [183]. The addition of ACT to Planck does not lead to an appreciable change in the axion density marginalized posterior, despite the small effects it entails for other parameters. The visible $\Omega_a h^2$ - h degeneracy direction can be readily interpreted as stemming from keeping θ_* , which is measured very accurately by the CMB primaries, fixed. It is, therefore, a geometric constraint.

More generally, for $m_a \lesssim 10^{-26}$ eV ACT has additional constraining power compared to Planck-only, as reported by the orange vs. red exclusions in Fig. 6.5. As noted in [143], this is partially driven also by the degeneracy with n_s , for which ACT favors a slightly higher value.

DESI DR1 full shape Figure 6.4 reports the posterior for $\Omega_a h^2$ obtained from the full shape of DESI DR1 galaxy power spectra, for three masses where DESI is expected to be informative independently of the CMB. Remarkably, as shown in the left panel for $m_a = 10^{-25}$ eV, for the heaviest ULAs considered in this work the galaxy power spectrum alone is more constraining than the combination of primary CMB anisotropies and BAO geometric information. This genuinely stems from the fact that the small-

scale $P_{gg,\ell}$ datapoints probe a suppression of power compared to the amplitude on large scales, and would strongly benefit from reaching even larger k , such as could be achieved with the real-space combination Q_0 [200, 320]. However, that is beyond the recommended DESI DR1 data cuts and therefore not included here.

For $m_a = 10^{-26}$ eV, the middle panel of Fig. 6.4 shows the posterior obtained from single-redshift-bin analyses, as well as their combination. As can be seen, two LRG bins (LRG2 and, albeit less strongly, LRG3) show a preference for a non-zero value of $\Omega_a h^2$. Interestingly, for the same m_a a similar preference was found in the full-shape analysis of BOSS LRG of [204] and in Sect. 4.4, see for instance Fig. 4.13. This should not be totally unexpected, as DESI DR1 shares 27% of targets with the BOSS LRG catalog [227]. Nevertheless, all other tracers do not exhibit a preference for any deviation from pure CDM and, accordingly, the joint-bin analysis only exhibits a very mild preference, fully compatible with a statistical fluctuation.

Combined The combination of CMB and galaxy power spectrum full shape is, of course, expected to be the most informative one. Qualitatively, the heavier ULAs considered in this work are mainly probed via the shape suppression, while CMB primary anisotropies calibrate the standard cosmological parameters. For lighter ULAs, which are already well-constrained by primary CMB anisotropies, galaxy clustering is also able to add further information about the expansion at late times.

We first derive Planck+DESI constraints without including ACT, to highlight the genuine constraining power of the full-shape analysis. These are reported in purple in Fig. 6.3, showing that both at the high and low ends of the mass range we consider, the improvement from adding DESI is considerable. For $m_a = 10^{-29}$ eV, shown in the right panel of Fig. 6.3, beyond the non-degenerate information that is added the improvement is also driven by the DESI preference for a higher value of the Hubble parameter. Indeed, along the degeneracy direction at fixed θ_* imposed by CMB data alone, a higher h is only compatible with a lower $\Omega_a h^2$.

The right panel of Fig. 6.4 shows, for $m_a = 10^{-27}$ eV, the further gain in constraining power one can obtain by using the theoretically-motivated priors given in Eq. (6.5). This improvement, which from Fig. 4.14 is indeed expected to be most relevant in a mass window around $m_a \sim 10^{-27}$ eV, would be important especially if data were to show a preference for an ULA subcomponent. The validity of such narrower priors may be assessed, for instance, using N-body simulations, as done for massive neutrinos [53, 88]. We leave this for future investigation.

Finally, Fig. 6.5 and Table 6.2 summarize the main results of this work, in the form of upper limits on $\Omega_a h^2$ as functions of the ULA mass. We see that in large portions of the parameter space the addition of DESI data, and of ACT data to a lesser extent, reduce the allowed ULA energy density by more than a factor of two compared to previous analyses. These are among the leading constraints on ULAs in the considered mass range. We do not include here results from probes of smaller scales, whose modeling lies beyond the perturbative regime. These include for example the Lyman- α forest [42, 137, 138, 321, 322], which has anyway not been studied in the range of ULA masses relevant to this work, or counts of high-redshift galaxies in the ultraviolet band [323–325].




Figure 6.5.: 95% CL upper bounds on the energy density in ULAs today as functions of the ULA mass, obtained in this work from large-scale perturbative probes. The results of DESI, ACT and BOSS are individually combined only with Planck. The hatched areas correspond to the novel exclusions derived in this work. The combined constraints for $m_a > 10^{-25}$ eV are shown by dashed lines, as they are subject to theoretical systematics due to the large allowed value of the ULA fraction f_{ax} . The gray horizontal lines mark the different orders of magnitude for f_{ax} .

6.4. Concluding remarks

In this chapter we presented the first joint analysis of DESI DR1 full-shape and state-of-the-art CMB data aimed at probing a non-cold DM component, taking ULAs as our benchmark scenario.

We performed an independent reanalysis of the DESI DR1 galaxy power spectra, consistently modeling the observables from first principles. Our implementation builds on the two-fluid EFTofLSS framework developed in Chapt. 4, which allowed us to constrain the signatures of $\Omega_a h^2$ without resorting to phenomenological assumptions (such as on bias parameters) calibrated only on Λ CDM or massive neutrino cosmologies. Our analysis shows that the combination of galaxy clustering with CMB primary anisotropies and CMB lensing achieves exquisite sensitivity to deviations from pure CDM, restricting the fraction of matter energy density in ULAs, f_{ax} , to be below 1% for all masses $m_a \lesssim 10^{-27}$ eV, and as small as 0.3% for $m_a \sim 10^{-29}$ eV. For the lightest ULAs considered in this work, the addition of DESI DR1 full-shape information to Planck+ACT tightens the bounds on $\Omega_a h^2$ by a factor of ~ 2 , as the galaxy power spectrum breaks the geometric degeneracies that otherwise limit the CMB-only constraints. The DESI DR1 full shape provides new constraining power also at the highest end of the mass range, where it directly probes the suppression of power caused by the ULA Jeans scale: as shown in Fig. 6.5 and Table 6.2, for $m_a = 10^{-25}$ eV the combination of DESI with Planck gives the tightest bound. Our results represent the most stringent limits on the abundance of ULAs from the combination of galaxy clustering and CMB observables, which admit a perturbative modeling.

An intriguing feature of the DESI DR1 full-shape analysis is a mild preference for a non-zero ULA

m_a [eV]	Planck + DESI (FS)	Planck + ACT + DESI (FS)
10^{-25}	0.017 _(12%)	0.022 _(16%)
10^{-26}	0.0051 _(3.6%)	0.0048 _(3.4%)
10^{-27}	0.0015 _(1.1%)	0.0011 _(0.81%)
10^{-28}	0.00055 _(0.39%)	0.00064 _(0.46%)
10^{-29}	0.00038 _(0.27%)	0.00047 _(0.33%)
10^{-30}	0.00044 _(0.31%)	0.00045 _(0.32%)
10^{-31}	0.00061 _(0.43%)	0.00063 _(0.45%)
10^{-32}	0.0012 _(0.84%)	0.0012 _(0.83%)

Table 6.2.: 95% CL upper limits on the fraction of energy density in ULAs, $\Omega_a h^2$. The corresponding upper limit on f_{ax} is expressed as a percentage.

fraction at $m_a \approx 10^{-26}$ eV, see the middle panel of Fig. 6.4. This is driven primarily by the LRG2, and partially the LRG3, redshift bins and follows a previous hint from BOSS analyses. However, the preference is substantially diluted once all DESI tracers are combined, and disappears entirely upon combination with CMB data. We thus interpret it as a statistical fluctuation, rather than a hint of new physics.

Several developments would further sharpen our results and extend their reach. On the galaxy clustering side, at the upper end of the ULA mass window the constraints are driven by the shape suppression and would therefore benefit considerably from the inclusion of information from smaller scales, beyond the $k_{\text{max}} = 0.2 h \text{ Mpc}^{-1}$ recommended by the DESI collaboration but still in the perturbative regime. On the CMB side, extending the ACT DR6 lensing analysis to the extended multipole range with $L_{\text{max}} = 1300$ [306, 313] would add valuable constraining power, particularly for the largest masses, provided the nonlinear modeling of the ULA contribution to the matter power spectrum is brought under sufficient theoretical control. Indeed, pushing this framework to masses above 10^{-25} eV, where the allowed f_{ax} is large enough for genuine ULA nonlinearities to be important, will require a dedicated study. Our 1-loop perturbative treatment of CMB lensing is a first step in this direction, but a more complete analysis will be needed to robustly extend the exclusion limits into that regime. Simulations could also be leveraged, for instance to devise physically motivated priors on the axion bias and counterterm parameters b_a , $c_{ca,\ell}$, and $c_{aa,\ell}$, refining the conservative priors adopted here.

7. Conclusion and outlooks

The overarching aim of this thesis has been to sharpen the theoretical tools with which galaxy clustering can be turned into precision tests of the dark sector, and to apply them to the concrete cases of massive neutrinos and ultra-light axions across a widening range of observables.

The motivation could hardly be more fundamental: the nature of dark matter and the value of the neutrino masses remain among the most consequential open questions in physics, and with the CMB approaching its fundamental limits in constraining them, the three-dimensional clustering of galaxies is now the most promising probe. There, however, the statistical power of Stage-IV surveys is outrunning the accuracy of the models available to interpret them, mainly developed under the assumption of a purely cold and non-interacting dark sector. The chapters that precede this one are among the first attempts at tackling these limitations, along three different, but complementary, directions: the response of biased tracers to the massive neutrino component at the level of the velocity field (Chapter 3), a first extension of the Effective Field Theory of Large-Scale Structure to mixed dark matter cosmologies (Chapter 4), and the enlargement of the set of observables through cross-correlations and higher-order statistics (Chapter 5). These methods are then brought together in the original analysis of DESI DR1 and ACT DR6 data presented in Chapter 6. What follows summarises the main results, considers what they establish jointly, and indicates the main outlooks.

Chapter 3 cleared an ambiguity that arises in the modelling of redshift-space clustering in massive neutrino cosmologies, namely which velocity field biased tracers should be taken to follow, in analogy to what is done for the real-space density field. On the basis of large ensembles of N-body simulations, it was shown that the Kaiser formula reproduces the measured multipoles only when halo velocities are assumed to trace the baryon-plus-CDM component, rather than the total matter field. The distinction is not a refinement of the non-linear description: it operates already at linear order, and thus affects the very leading-order term on which any treatment of redshift-space distortions is built. Adopting the total-matter convention would therefore induce a scale dependence in the inferred growth rate that has no physical origin, and which is liable to misinterpretation in analyses aimed at detecting precisely such scale-dependent effects.

Chapter 4 constitutes the major theoretical development of this thesis. We started the construction of a framework in which the nonlinearities for two species, namely the cold and a non-cold with small abundance, are treated self-consistently. Analytic solutions were obtained in the relevant limiting regimes, leading to a numerical prescription for the one-loop galaxy power spectrum sufficiently efficient for use within a full-shape likelihood. Importantly, the resulting formalism applies not to a single fundamental physics model, but to the whole class of scenarios sharing this phenomenology.

Chapter 5 turned to the constraining power available from observables beyond the galaxy power spectrum alone, working within Λ CDM, as in this case the novelty lies in the probe rather than in the physics being searched for. Employing the angular power spectrum and bispectrum of luminous red galaxies and the cross-correlation with CMB lensing convergence, accelerating the theoretical

7. Conclusion and outlooks

predictions with a neural-network emulator in order to render the inference tractable, we demonstrated that the inclusion of the bispectrum provides an independent calibration of the galaxy bias parameters and a corresponding, mild, improvement in the cosmological constraints. This is the first validation on N-body simulations, and consistent application to data, of the projected galaxy bispectrum–CMB lensing combination as a cosmological probe.

Chapter 6 represents, in some sense, a culmination of the work presented in this thesis, in that it draws upon both the expertise gained with, and the theoretical developments of, the previous chapters. It presents an original analysis of state-of-the-art data, and the first constraints on light species obtained from the DESI full-shape measurements: the extended framework of Chapter 4 supplies the theoretical template, applied to the DESI DR1 galaxy power spectra in combination with CMB data from Planck and ACT, over the mass window $10^{-32} \text{ eV} \leq m_a \leq 10^{-24} \text{ eV}$. For the heavier masses, the CMB lensing likelihood required a novel treatment of its own, modelled here entirely within perturbation theory and thus also consistent with the galaxy clustering likelihood, drawing on the expertise developed in Chapter 5: this is a first attempt at a controlled description beyond CDM, the alternative being phenomenological extensions of halo-model codes. For the lightest masses, the joint analysis improves upon CMB-only constraints by more than a factor of two, establishing the most stringent limits to date: at $m_a \sim 10^{-29} \text{ eV}$, the ULA subcomponent is bounded to a fraction as small as 0.3% of the total matter energy density. The DESI Luminous Red Galaxy sample has been found to display a mild preference for an ULA subcomponent with $m_a \approx 10^{-26} \text{ eV}$, mirroring earlier hints from BOSS, though this preference does not survive the combination with CMB data.

It must be emphasised that the developments of this work apply to a specific class of searches for new physics in cosmology, with limitations that should be stated plainly. The ubiquitous one is the perturbative reach of the modelling: since the analysis is restricted to scales where the EFT expansion is under control, the sensitivity to a non-cold component degrades heavily once its characteristic suppression scale falls below that range. For ultra-light axions this sets the upper edge of the accessible mass window at $m_a \sim 10^{-25} \text{ eV}$, precisely where the axion becomes increasingly dark-matter-like, and where the constraints would be of greatest interest. Extending the reach to smaller scales, where possible, is therefore not a refinement but a prerequisite for progress in this framework.

A limitation of the mixed DM framework developed in this thesis concerns the description of the non-cold species itself. The study of Chapter 4 treats the subcomponent as a fluid, an approximation that need not capture the details of other BSM models, and that may not be sufficient once nonlinearities are included; its failures are expected to appear first in higher-order correlators. It further retains only the growing mode, whereas the relative modes between the two species have been shown, for some models, to contribute at a level that is not always negligible. Neither approximation is expected to compromise the constraints obtained here, but both would require scrutiny before the framework is applied to observables, or to scales, more sensitive to them.

A third limitation concerns the degree of model dependence of the bounds obtained in this thesis. The constraints of Chapter 6 are quoted in the parameters of a specific scenario, and while the underlying formalism applies to the wider class of models sharing the same phenomenology, that mapping has not been carried out here, and it is not clear whether that can be done in full generality. These bounds cannot therefore, at present, be straightforwardly reinterpreted in terms of different fundamental physics scenarios.

A further caveat is of a different character, in that it concerns not the modelling but the data

themselves. Constraints on light species are extracted under the assumption that the datasets entering the combination are mutually consistent, and this assumption has come under strain. The case of the neutrino mass is the most apparent: the combination of CMB and recent BAO measurements now favours values of $\sum m_\nu$ compatible with zero, and in some dataset combinations formally below the lower bound established by oscillation experiments, when this is not imposed as an external prior for the cosmological analysis. Whether this reflects a statistical fluctuation, an unaccounted systematic in one of the datasets, or a genuine departure from Λ CDM in the late-time expansion history remains unsettled, and the answer bears directly on the interpretation of any bound derived from the same combinations. Constraints on a non-cold subcomponent are not immune to this, resting as they do on the same measurements and on the same assumed background evolution.

Several of the directions opened by this work will progress with the new data that will be incoming. The methods implemented for the studies of this thesis are, in this respect, an investment whose return is also quite immediate: for instance, the constraints of Chapter 6 can be updated with each successive DESI data release at almost no theoretical cost, and the same applies to the spectroscopic samples that Euclid will deliver. The gains will not be merely statistical, as larger volume and a wider redshift baseline will yield non-degenerate information on a scale-dependent suppression. Comparable progress can be expected on the CMB side, where forthcoming lensing measurements of substantially higher signal-to-noise will sharpen the $\kappa\kappa$ constraints and, notably, also the cross-correlation with galaxy clustering.

On the theoretical side, a natural next step is to widen the class of scenarios to which the framework of Chapter 4 is applied. Some require little more than a change of variables: light thermal relics, for instance, are described by a similar background evolution, and a mapping of their particle physics parameters might be attempted. Others demand genuinely new work, for instance decaying subcomponents, in that their nontrivial background evolution at late times cannot be accommodated within the present construction. More ambitious, and in the longer term more valuable, would be to try to move away from a model-by-model treatment. The phenomenology exploited throughout this thesis is governed by two quantities, a characteristic suppression scale and the abundance of the non-cold species, and constraints expressed directly in that plane would be reusable by anyone whose model can be located within it, rather than requiring a dedicated analysis for each scenario in turn. A necessary ingredient of such a programme is a cleaner separation between the suppression induced at the level of the background and that arising from the perturbations, which the present formalisms do not make fully manifest.

A main outlook of this work, moreover, lies in the fact that the observational reach of these searches for beyond-CDM signals remains far from exhausted. The combination of galaxy clustering with CMB lensing exploited in Chapter 5 has not yet been brought into the beyond-CDM context, where the independent calibration of the bias parameters it allows would translate directly into tighter constraints on the non-cold fraction. Furthermore, the case of the bispectrum is stronger than degeneracy-breaking alone would suggest. As noted above, the fluid description of the non-cold species may be expected to fail first in higher-order correlators, so that the bispectrum contains information with no counterpart at the two-point level. Whether such signatures are large enough to be detectable is unknown; the fact that the question has scarcely been posed is itself a strong motivation for a novel exploration. Beyond spectroscopic clustering, weak lensing statistics and the kinematic Sunyaev–Zel’dovich effect have received comparatively little attention in this context, and their systematic development, individually and in cross-correlation, would open a genuinely new front in the search for physics beyond the

7. Conclusion and outlooks

Standard Model.

It is worth closing on a more general note. The questions that motivated this work have been with us for decades, especially concerning the fundamental nature of dark matter, and it is clear that they are unlikely to be settled by any single decisive measurement. What has changed is that the data now arriving give a reason to keep asking them: the volumes being mapped spectroscopically, and the sensitivity now reached in secondary CMB anisotropies, give access to a regime of statistical accuracy that simply has not been probed before. The signals to be found there might not be the striking ones of earlier decades; they are small, and extracting them demands a control of the modelling, and a careful combination of probes. That is a very technical, and somewhat hard, effort, but that is precisely what the frontier of precision cosmology consists of. Furthermore, alongside the known unknowns which the community is actively searching for, there remains the genuinely exciting possibility of some unknown unknown in the forthcoming data. Whether or not such a surprise will materialize, we hope that this thesis has contributed its small piece to the wider effort of the whole cosmology community.

Bibliography

- [1] Francesco Verdiani et al. “Redshift-space distortions in massive neutrinos cosmologies”. In: *Phys. Rev. D* 112.4 (2025), p. 043545. DOI: [10.1103/mjsv-885q](#). arXiv: [2503.06655 \[astro-ph.CO\]](#).
- [2] Francesco Verdiani et al. “The Effective Field Theory of Large Scale Structure for mixed dark matter scenarios”. In: *JCAP* 2601 (2026), p. 047. DOI: [10.1088/1475-7516/2026/01/047](#). arXiv: [2507.08792 \[astro-ph.CO\]](#).
- [3] Francesco Verdiani et al. “Cosmological constraints from the angular power spectrum and bispectrum of luminous red galaxies and CMB lensing”. In: *JCAP* 06 (2026), p. 027. DOI: [10.1088/1475-7516/2026/06/027](#). arXiv: [2510.17796 \[astro-ph.CO\]](#).
- [4] Francesco Verdiani et al. “Constraints on Ultra-Light Axions from the DESI DR1 Full Shape, Planck and ACT”. In: (July 2026). arXiv: [2607.11876 \[astro-ph.CO\]](#).
- [5] N. Aghanim et al. “Planck 2018 results. VI. Cosmological parameters”. In: *Astron. Astrophys.* 641 (2020). [Erratum: *Astron.Astrophys.* 652, C4 (2021)], A6. DOI: [10.1051/0004-6361/201833910](#). arXiv: [1807.06209 \[astro-ph.CO\]](#).
- [6] Thibaut Louis et al. “The Atacama Cosmology Telescope: DR6 power spectra, likelihoods and Λ CDM parameters”. In: *JCAP* 2511 (2025), p. 062. DOI: [10.1088/1475-7516/2025/11/062](#). arXiv: [2503.14452 \[astro-ph.CO\]](#).
- [7] D. Dutcher et al. “Measurements of the E-mode polarization and temperature-E-mode correlation of the CMB from SPT-3G 2018 data”. In: *Phys. Rev. D* 104.2 (2021), p. 022003. DOI: [10.1103/PhysRevD.104.022003](#). arXiv: [2101.01684 \[astro-ph.CO\]](#).
- [8] Shadab Alam et al. “The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample”. In: *Mon. Not. Roy. Astron. Soc.* 470.3 (2017), pp. 2617–2652. DOI: [10.1093/mnras/stx721](#). arXiv: [1607.03155 \[astro-ph.CO\]](#).
- [9] Shadab Alam et al. “Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological implications from two decades of spectroscopic surveys at the Apache Point Observatory”. In: *Phys. Rev. D* 103.8 (2021), p. 083533. DOI: [10.1103/PhysRevD.103.083533](#). arXiv: [2007.08991 \[astro-ph.CO\]](#).
- [10] M. Abdul Karim et al. “Data Release 1 of the Dark Energy Spectroscopic Instrument”. In: (Mar. 2025). arXiv: [2503.14745 \[astro-ph.CO\]](#).
- [11] Y. Mellier et al. “Euclid. I. Overview of the Euclid mission”. In: *Astron. Astrophys.* 697 (2025), A1. DOI: [10.1051/0004-6361/202450810](#). arXiv: [2405.13491 \[astro-ph.CO\]](#).
- [12] Željko Ivezić et al. “LSST: from Science Drivers to Reference Design and Anticipated Data Products”. In: *Astrophys. J.* 873.2 (2019), p. 111. DOI: [10.3847/1538-4357/ab042c](#). arXiv: [0805.2366 \[astro-ph\]](#).

- [13] F. Bernardeau et al. “Large scale structure of the universe and cosmological perturbation theory”. In: *Phys. Rept.* 367 (2002), pp. 1–248. DOI: [10.1016/S0370-1573\(02\)00135-7](https://doi.org/10.1016/S0370-1573(02)00135-7). arXiv: [astro-ph/0112551](https://arxiv.org/abs/astro-ph/0112551).
- [14] Daniel Baumann et al. “Cosmological Non-Linearities as an Effective Fluid”. In: *JCAP* 07 (2012), p. 051. DOI: [10.1088/1475-7516/2012/07/051](https://doi.org/10.1088/1475-7516/2012/07/051). arXiv: [1004.2488](https://arxiv.org/abs/1004.2488) [[astro-ph.CO](#)].
- [15] John Joseph M. Carrasco, Mark P. Hertzberg, and Leonardo Senatore. “The Effective Field Theory of Cosmological Large Scale Structures”. In: *JHEP* 1209 (2012), p. 082. DOI: [10.1007/JHEP09\(2012\)082](https://doi.org/10.1007/JHEP09(2012)082). arXiv: [1206.2926](https://arxiv.org/abs/1206.2926) [[astro-ph.CO](#)].
- [16] Rafael A. Porto, Leonardo Senatore, and Matias Zaldarriaga. “The Lagrangian-space Effective Field Theory of Large Scale Structures”. In: *JCAP* 05 (2014), p. 022. DOI: [10.1088/1475-7516/2014/05/022](https://doi.org/10.1088/1475-7516/2014/05/022). arXiv: [1311.2168](https://arxiv.org/abs/1311.2168) [[astro-ph.CO](#)].
- [17] Ashley Perko et al. “Biased Tracers in Redshift Space in the EFT of Large-Scale Structure”. In: (Oct. 2016). arXiv: [1610.09321](https://arxiv.org/abs/1610.09321) [[astro-ph.CO](#)].
- [18] Guido D’Amico et al. “The Cosmological Analysis of the SDSS/BOSS data from the Effective Field Theory of Large-Scale Structure”. In: *JCAP* 05 (2020), p. 005. DOI: [10.1088/1475-7516/2020/05/005](https://doi.org/10.1088/1475-7516/2020/05/005). arXiv: [1909.05271](https://arxiv.org/abs/1909.05271) [[astro-ph.CO](#)].
- [19] Mikhail M. Ivanov, Marko Simonović, and Matias Zaldarriaga. “Cosmological Parameters from the BOSS Galaxy Power Spectrum”. In: *JCAP* 05 (2020), p. 042. DOI: [10.1088/1475-7516/2020/05/042](https://doi.org/10.1088/1475-7516/2020/05/042). arXiv: [1909.05277](https://arxiv.org/abs/1909.05277) [[astro-ph.CO](#)].
- [20] Thomas Colas et al. “Efficient Cosmological Analysis of the SDSS/BOSS data from the Effective Field Theory of Large-Scale Structure”. In: *JCAP* 06 (2020), p. 001. DOI: [10.1088/1475-7516/2020/06/001](https://doi.org/10.1088/1475-7516/2020/06/001). arXiv: [1909.07951](https://arxiv.org/abs/1909.07951) [[astro-ph.CO](#)].
- [21] Shi-Fan Chen, Zvonimir Vlah, and Martin White. “A new analysis of galaxy 2-point functions in the BOSS survey, including full-shape information and post-reconstruction BAO”. In: *JCAP* 02.02 (2022), p. 008. DOI: [10.1088/1475-7516/2022/02/008](https://doi.org/10.1088/1475-7516/2022/02/008). arXiv: [2110.05530](https://arxiv.org/abs/2110.05530) [[astro-ph.CO](#)].
- [22] Antony Lewis and Anthony Challinor. “Weak gravitational lensing of the CMB”. In: *Phys. Rept.* 429 (2006), pp. 1–65. DOI: [10.1016/j.physrep.2006.03.002](https://doi.org/10.1016/j.physrep.2006.03.002). arXiv: [astro-ph/0601594](https://arxiv.org/abs/astro-ph/0601594).
- [23] N. Aghanim et al. “Planck 2018 results. VIII. Gravitational lensing”. In: *Astron. Astrophys.* 641 (2020), A8. DOI: [10.1051/0004-6361/201833886](https://doi.org/10.1051/0004-6361/201833886). arXiv: [1807.06210](https://arxiv.org/abs/1807.06210) [[astro-ph.CO](#)].
- [24] Mathew S. Madhavacheril et al. “The Atacama Cosmology Telescope: DR6 Gravitational Lensing Map and Cosmological Parameters”. In: *Astrophys. J.* 962.2 (2024), p. 113. DOI: [10.3847/1538-4357/acff5f](https://doi.org/10.3847/1538-4357/acff5f). arXiv: [2304.05203](https://arxiv.org/abs/2304.05203) [[astro-ph.CO](#)].
- [25] Y. Fukuda et al. “Evidence for oscillation of atmospheric neutrinos”. In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. DOI: [10.1103/PhysRevLett.81.1562](https://doi.org/10.1103/PhysRevLett.81.1562). arXiv: [hep-ex/9807003](https://arxiv.org/abs/hep-ex/9807003).
- [26] Julien Lesgourgues and Sergio Pastor. “Massive neutrinos and cosmology”. In: *Phys. Rept.* 429 (2006), pp. 307–379. DOI: [10.1016/j.physrep.2006.04.001](https://doi.org/10.1016/j.physrep.2006.04.001). arXiv: [astro-ph/0603494](https://arxiv.org/abs/astro-ph/0603494).
- [27] M. Aker et al. “Direct neutrino-mass measurement with sub-electronvolt sensitivity”. In: *Nature Phys.* 18.2 (2022), pp. 160–166. DOI: [10.1038/s41567-021-01463-1](https://doi.org/10.1038/s41567-021-01463-1). arXiv: [2105.08533](https://arxiv.org/abs/2105.08533) [[hep-ex](#)].

- [28] Max Aker et al. “Direct neutrino-mass measurement based on 259 days of KATRIN data”. In: *Science* 388.6743 (2025), adq9592. DOI: [10.1126/science.adq9592](https://doi.org/10.1126/science.adq9592). arXiv: [2406.13516](https://arxiv.org/abs/2406.13516) [nucl-ex].
- [29] Amir Aghamousa et al. “The DESI Experiment Part I: Science, Targeting, and Survey Design”. In: (Oct. 2016). arXiv: [1611.00036](https://arxiv.org/abs/1611.00036) [astro-ph.IM].
- [30] Kevork N. Abazajian et al. “CMB-S4 Science Book, First Edition”. In: (Oct. 2016). arXiv: [1610.02743](https://arxiv.org/abs/1610.02743) [astro-ph.CO].
- [31] M. Abdul Karim et al. “DESI DR2 results. II. Measurements of baryon acoustic oscillations and cosmological constraints”. In: *Phys. Rev. D* 112.8 (2025), p. 083515. DOI: [10.1103/tr6y-kpc6](https://doi.org/10.1103/tr6y-kpc6). arXiv: [2503.14738](https://arxiv.org/abs/2503.14738) [astro-ph.CO].
- [32] Nathaniel Craig et al. “No ν_s is Good News”. In: *JHEP* 09 (2024), p. 097. DOI: [10.1007/JHEP09\(2024\)097](https://doi.org/10.1007/JHEP09(2024)097). arXiv: [2405.00836](https://arxiv.org/abs/2405.00836) [astro-ph.CO].
- [33] Daniel Green and Joel Meyers. “Cosmological preference for a negative neutrino mass”. In: *Phys. Rev. D* 111.8 (2025), p. 083507. DOI: [10.1103/PhysRevD.111.083507](https://doi.org/10.1103/PhysRevD.111.083507). arXiv: [2407.07878](https://arxiv.org/abs/2407.07878) [astro-ph.CO].
- [34] Willem Elbers et al. “Negative neutrino masses as a mirage of dark energy”. In: *Phys. Rev. D* 111.6 (2025), p. 063534. DOI: [10.1103/PhysRevD.111.063534](https://doi.org/10.1103/PhysRevD.111.063534). arXiv: [2407.10965](https://arxiv.org/abs/2407.10965) [astro-ph.CO].
- [35] Daniel Naredo-Tuero et al. “Critical look at the cosmological neutrino mass bound”. In: *Phys. Rev. D* 110.12 (2024), p. 123537. DOI: [10.1103/PhysRevD.110.123537](https://doi.org/10.1103/PhysRevD.110.123537). arXiv: [2407.13831](https://arxiv.org/abs/2407.13831) [astro-ph.CO].
- [36] L. Verde, T. Treu, and A. G. Riess. “Tensions between the Early and the Late Universe”. In: *Nature Astron.* 3 (2019), p. 891. DOI: [10.1038/s41550-019-0902-0](https://doi.org/10.1038/s41550-019-0902-0). arXiv: [1907.10625](https://arxiv.org/abs/1907.10625) [astro-ph.CO].
- [37] Eleonora Di Valentino et al. “In the realm of the Hubble tension—a review of solutions”. In: *Class. Quant. Grav.* 38.15 (2021), p. 153001. DOI: [10.1088/1361-6382/ac086d](https://doi.org/10.1088/1361-6382/ac086d). arXiv: [2103.01183](https://arxiv.org/abs/2103.01183) [astro-ph.CO].
- [38] A. G. Adame et al. “DESI 2024 VI: cosmological constraints from the measurements of baryon acoustic oscillations”. In: *JCAP* 2502 (2025), p. 021. DOI: [10.1088/1475-7516/2025/02/021](https://doi.org/10.1088/1475-7516/2025/02/021). arXiv: [2404.03002](https://arxiv.org/abs/2404.03002) [astro-ph.CO].
- [39] James S. Bullock and Michael Boylan-Kolchin. “Small-Scale Challenges to the Λ CDM Paradigm”. In: *Ann. Rev. Astron. Astrophys.* 55 (2017), pp. 343–387. DOI: [10.1146/annurev-astro-091916-055313](https://doi.org/10.1146/annurev-astro-091916-055313). arXiv: [1707.04256](https://arxiv.org/abs/1707.04256) [astro-ph.CO].
- [40] Matteo Viel et al. “Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman- α forest data”. In: *Phys. Rev. D* 88 (2013), p. 043502. DOI: [10.1103/PhysRevD.88.043502](https://doi.org/10.1103/PhysRevD.88.043502). arXiv: [1306.2314](https://arxiv.org/abs/1306.2314) [astro-ph.CO].
- [41] Vid Iršič et al. “New Constraints on the free-streaming of warm dark matter from intermediate and small scale Lyman- α forest data”. In: *Phys. Rev. D* 96.2 (2017), p. 023522. DOI: [10.1103/PhysRevD.96.023522](https://doi.org/10.1103/PhysRevD.96.023522). arXiv: [1702.01764](https://arxiv.org/abs/1702.01764) [astro-ph.CO].

- [42] Keir K. Rogers and Hiranya V. Peiris. “Strong Bound on Canonical Ultralight Axion Dark Matter from the Lyman-Alpha Forest”. In: *Phys. Rev. Lett.* 126.7 (2021), p. 071302. DOI: [10 . 1103 / PhysRevLett.126.071302](https://doi.org/10.1103/PhysRevLett.126.071302). arXiv: [2007.12705](https://arxiv.org/abs/2007.12705) [astro-ph.CO].
- [43] Alexey Boyarsky et al. “Lyman-alpha constraints on warm and on warm-plus-cold dark matter models”. In: *JCAP* 05 (2009), p. 012. DOI: [10 . 1088/1475-7516/2009/05/012](https://doi.org/10.1088/1475-7516/2009/05/012). arXiv: [0812.0010](https://arxiv.org/abs/0812.0010) [astro-ph].
- [44] Nicholas DePorzio et al. “Finding eV-scale light relics with cosmological observables”. In: *Phys. Rev. D* 103.2 (2021), p. 023504. DOI: [10 . 1103 / PhysRevD . 103 . 023504](https://doi.org/10.1103/PhysRevD.103.023504). arXiv: [2006 . 09380](https://arxiv.org/abs/2006.09380) [astro-ph.CO].
- [45] Weishuang Linda Xu, Julian B. Muñoz, and Cora Dvorkin. “Cosmological constraints on light but massive relics”. In: *Phys. Rev. D* 105.9 (2022), p. 095029. DOI: [10 . 1103/PhysRevD . 105 . 095029](https://doi.org/10.1103/PhysRevD.105.095029). arXiv: [2107.09664](https://arxiv.org/abs/2107.09664) [astro-ph.CO].
- [46] Daniel Baumann, Daniel Green, and Benjamin Wallisch. “New Target for Cosmic Axion Searches”. In: *Phys. Rev. Lett.* 117.17 (2016), p. 171301. DOI: [10 . 1103/PhysRevLett . 117 . 171301](https://doi.org/10.1103/PhysRevLett.117.171301). arXiv: [1604.08614](https://arxiv.org/abs/1604.08614) [astro-ph.CO].
- [47] Tobias Baldauf et al. “The Bispectrum in the Effective Field Theory of Large Scale Structure”. In: *JCAP* 05 (2015), p. 007. DOI: [10 . 1088 / 1475 - 7516 / 2015 / 05 / 007](https://doi.org/10.1088/1475-7516/2015/05/007). arXiv: [1406 . 4135](https://arxiv.org/abs/1406.4135) [astro-ph.CO].
- [48] Raul E. Angulo et al. “The One-Loop Matter Bispectrum in the Effective Field Theory of Large Scale Structures”. In: *JCAP* 10 (2015), p. 039. DOI: [10 . 1088/1475-7516/2015/10/039](https://doi.org/10.1088/1475-7516/2015/10/039). arXiv: [1406.4143](https://arxiv.org/abs/1406.4143) [astro-ph.CO].
- [49] Oliver H. E. Philcox and Mikhail M. Ivanov. “BOSS DR12 full-shape cosmology: Λ CDM constraints from the large-scale galaxy power spectrum and bispectrum monopole”. In: *Phys. Rev. D* 105.4 (2022), p. 043517. DOI: [10 . 1103 / PhysRevD . 105 . 043517](https://doi.org/10.1103/PhysRevD.105.043517). arXiv: [2112 . 04515](https://arxiv.org/abs/2112.04515) [astro-ph.CO].
- [50] Guido D’Amico et al. “The BOSS bispectrum analysis at one loop from the Effective Field Theory of Large-Scale Structure”. In: *JCAP* 05 (2024), p. 059. DOI: [10 . 1088/1475-7516/2024/05/059](https://doi.org/10.1088/1475-7516/2024/05/059). arXiv: [2206.08327](https://arxiv.org/abs/2206.08327) [astro-ph.CO].
- [51] Martin White et al. “Cosmological constraints from the tomographic cross-correlation of DESI Luminous Red Galaxies and Planck CMB lensing”. In: *JCAP* 2022.2, 007 (Feb. 2022), p. 007. DOI: [10 . 1088/1475-7516/2022/02/007](https://doi.org/10.1088/1475-7516/2022/02/007). arXiv: [2111 . 09898](https://arxiv.org/abs/2111.09898) [astro-ph.CO].
- [52] Francisco Villaescusa-Navarro et al. “Cosmology with massive neutrinos I: towards a realistic modeling of the relation between matter, haloes and galaxies”. In: *JCAP* 03 (2014), p. 011. DOI: [10 . 1088/1475-7516/2014/03/011](https://doi.org/10.1088/1475-7516/2014/03/011). arXiv: [1311 . 0866](https://arxiv.org/abs/1311.0866) [astro-ph.CO].
- [53] Emanuele Castorina et al. “DEMNUi: The clustering of large-scale structures in the presence of massive neutrinos”. In: *JCAP* 1507 (2015), p. 043. DOI: [10 . 1088/1475-7516/2015/07/043](https://doi.org/10.1088/1475-7516/2015/07/043). arXiv: [1505.07148](https://arxiv.org/abs/1505.07148) [astro-ph.CO].
- [54] Wayne Hu, Rennan Barkana, and Andrei Gruzinov. “Cold and fuzzy dark matter”. In: *Phys. Rev. Lett.* 85 (2000), pp. 1158–1161. DOI: [10 . 1103 / PhysRevLett . 85 . 1158](https://doi.org/10.1103/PhysRevLett.85.1158). arXiv: [astro-ph/0003365](https://arxiv.org/abs/astro-ph/0003365).

- [55] Asimina Arvanitaki et al. “String Axiverse”. In: *Phys. Rev. D* 81 (2010), p. 123530. DOI: [10.1103/PhysRevD.81.123530](https://doi.org/10.1103/PhysRevD.81.123530). arXiv: [0905.4720](https://arxiv.org/abs/0905.4720) [hep-th].
- [56] David J. E. Marsh. “Axion Cosmology”. In: *Phys. Rept.* 643 (2016), pp. 1–79. DOI: [10.1016/j.physrep.2016.06.005](https://doi.org/10.1016/j.physrep.2016.06.005). arXiv: [1510.07633](https://arxiv.org/abs/1510.07633) [astro-ph.CO].
- [57] Lam Hui et al. “Ultralight scalars as cosmological dark matter”. In: *Phys. Rev. D* 95.4 (2017), p. 043541. DOI: [10.1103/PhysRevD.95.043541](https://doi.org/10.1103/PhysRevD.95.043541). arXiv: [1610.08297](https://arxiv.org/abs/1610.08297) [astro-ph.CO].
- [58] Marilena LoVerde. “Halo bias in mixed dark matter cosmologies”. In: *Phys. Rev. D* 90.8 (2014), p. 083530. DOI: [10.1103/PhysRevD.90.083530](https://doi.org/10.1103/PhysRevD.90.083530). arXiv: [1405.4855](https://arxiv.org/abs/1405.4855) [astro-ph.CO].
- [59] Scott Dodelson and Fabian Schmidt. *Modern Cosmology*. Academic Press, 2020. DOI: [10.1016/C2017-0-01943-2](https://doi.org/10.1016/C2017-0-01943-2).
- [60] Julien Lesgourgues. “The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview”. In: (Apr. 2011). arXiv: [1104.2932](https://arxiv.org/abs/1104.2932) [astro-ph.IM].
- [61] Antony Lewis, Anthony Challinor, and Anthony Lasenby. “Efficient computation of CMB anisotropies in closed FRW models”. In: *Astrophys. J.* 538 (2000), pp. 473–476. DOI: [10.1086/309179](https://doi.org/10.1086/309179). arXiv: [astro-ph/9911177](https://arxiv.org/abs/astro-ph/9911177).
- [62] Chiara Moretti et al. “Modified gravity and massive neutrinos: constraints from the full shape analysis of BOSS galaxies and forecasts for Stage IV surveys”. In: *JCAP* 12 (2023), p. 025. DOI: [10.1088/1475-7516/2023/12/025](https://doi.org/10.1088/1475-7516/2023/12/025). arXiv: [2306.09275](https://arxiv.org/abs/2306.09275) [astro-ph.CO].
- [63] ChangHoon Hahn et al. “A forward modeling approach to analyzing galaxy clustering with SimBIG”. In: *Proc. Nat. Acad. Sci.* 120.42 (2023), e2218810120. DOI: [10.1073/pnas.2218810120](https://doi.org/10.1073/pnas.2218810120). arXiv: [2211.00723](https://arxiv.org/abs/2211.00723) [astro-ph.CO].
- [64] ChangHoon Hahn et al. “Cosmological constraints from non-Gaussian and nonlinear galaxy clustering using the SimBIG inference framework”. In: *Nature Astron.* 8.11 (2024), pp. 1457–1467. DOI: [10.1038/s41550-024-02344-2](https://doi.org/10.1038/s41550-024-02344-2). arXiv: [2310.15246](https://arxiv.org/abs/2310.15246) [astro-ph.CO].
- [65] Diego Blas, Mathias Garny, and Thomas Konstandin. “On the non-linear scale of cosmological perturbation theory”. In: *JCAP* 09 (2013), p. 024. DOI: [10.1088/1475-7516/2013/09/024](https://doi.org/10.1088/1475-7516/2013/09/024). arXiv: [1304.1546](https://arxiv.org/abs/1304.1546) [astro-ph.CO].
- [66] Mikhail M. Ivanov. “Effective Field Theory for Large-Scale Structure”. In: 2023. DOI: [10.1007/978-981-19-3079-9_5-1](https://doi.org/10.1007/978-981-19-3079-9_5-1). arXiv: [2212.08488](https://arxiv.org/abs/2212.08488) [astro-ph.CO].
- [67] Matthew Lewandowski et al. “EFT of large scale structures in redshift space”. In: *Phys. Rev. D* 97.6 (2018), p. 063526. DOI: [10.1103/PhysRevD.97.063526](https://doi.org/10.1103/PhysRevD.97.063526). arXiv: [1512.06831](https://arxiv.org/abs/1512.06831) [astro-ph.CO].
- [68] Zvonimir Vlah, Martin White, and Alejandro Aviles. “A Lagrangian effective field theory”. In: *JCAP* 09 (2015), p. 014. DOI: [10.1088/1475-7516/2015/09/014](https://doi.org/10.1088/1475-7516/2015/09/014). arXiv: [1506.05264](https://arxiv.org/abs/1506.05264) [astro-ph.CO].
- [69] Vincent Desjacques, Donghui Jeong, and Fabian Schmidt. “Large-Scale Galaxy Bias”. In: *Phys. Rept.* 733 (2018), pp. 1–193. DOI: [10.1016/j.physrep.2017.12.002](https://doi.org/10.1016/j.physrep.2017.12.002). arXiv: [1611.09787](https://arxiv.org/abs/1611.09787) [astro-ph.CO].
- [70] N. Kaiser. “Clustering in real space and in redshift space”. In: *Mon. Not. Roy. Astron. Soc.* 227 (1987), pp. 1–27. DOI: [10.1093/mnras/227.1.1](https://doi.org/10.1093/mnras/227.1.1).

Bibliography

- [71] Tobias Baldauf et al. “Halo stochasticity from exclusion and nonlinear clustering”. In: *Phys. Rev. D* 88.8 (2013), p. 083507. DOI: [10.1103/PhysRevD.88.083507](https://doi.org/10.1103/PhysRevD.88.083507). arXiv: 1305.2917 [astro-ph.CO].
- [72] Emiliano Sefusatti et al. “Accurate Estimators of Correlation Functions in Fourier Space”. In: *Mon. Not. Roy. Astron. Soc.* 460.4 (2016), pp. 3624–3636. DOI: [10.1093/mnras/stw1229](https://doi.org/10.1093/mnras/stw1229). arXiv: 1512.07295 [astro-ph.CO].
- [73] Andrea Oddo et al. “Toward a robust inference method for the galaxy bispectrum: likelihood function and model selection”. In: *JCAP* 03 (2020), p. 056. DOI: [10.1088/1475-7516/2020/03/056](https://doi.org/10.1088/1475-7516/2020/03/056). arXiv: 1908.01774 [astro-ph.CO].
- [74] Matteo Biagetti, Francesco Verdiani, and Emiliano Sefusatti. *PBI4*. Version v1.0.0. Oct. 2023. URL: <https://zenodo.org/records/10008045>.
- [75] Hume A. Feldman, Nick Kaiser, and John A. Peacock. “Power spectrum analysis of three-dimensional redshift surveys”. In: *Astrophys. J.* 426 (1994), pp. 23–37. DOI: [10.1086/174036](https://doi.org/10.1086/174036). arXiv: astro-ph/9304022.
- [76] Kazuhiro Yamamoto et al. “A Measurement of the quadrupole power spectrum in the clustering of the 2dF QSO Survey”. In: *Publ. Astron. Soc. Jap.* 58 (2006), pp. 93–102. DOI: [10.1093/pasj/58.1.93](https://doi.org/10.1093/pasj/58.1.93). arXiv: astro-ph/0505115.
- [77] M. J. Wilson et al. “Rapid modelling of the redshift-space power spectrum multipoles for a masked density field”. In: *Mon. Not. Roy. Astron. Soc.* 464.3 (2017), pp. 3121–3130. DOI: [10.1093/mnras/stw2576](https://doi.org/10.1093/mnras/stw2576). arXiv: 1511.07799 [astro-ph.CO].
- [78] Florian Beutler and Patrick McDonald. “Unified galaxy power spectrum measurements from 6dFGS, BOSS, and eBOSS”. In: *JCAP* 11 (2021), p. 031. DOI: [10.1088/1475-7516/2021/11/031](https://doi.org/10.1088/1475-7516/2021/11/031). arXiv: 2106.06324 [astro-ph.CO].
- [79] Florian Beutler, Emanuele Castorina, and Pierre Zhang. “Interpreting measurements of the anisotropic galaxy power spectrum”. In: *JCAP* 1903 (2019), p. 040. DOI: [10.1088/1475-7516/2019/03/040](https://doi.org/10.1088/1475-7516/2019/03/040). arXiv: 1810.05051 [astro-ph.CO].
- [80] Ivan Esteban et al. “NuFit-6.0: updated global analysis of three-flavor neutrino oscillations”. In: *JHEP* 12 (2024), p. 216. DOI: [10.1007/JHEP12\(2024\)216](https://doi.org/10.1007/JHEP12(2024)216). arXiv: 2410.05380 [hep-ph].
- [81] Chung-Pei Ma and Edmund Bertschinger. “Cosmological perturbation theory in the synchronous and conformal Newtonian gauges”. In: *Astrophys. J.* 455 (1995), pp. 7–25. DOI: [10.1086/176550](https://doi.org/10.1086/176550). arXiv: astro-ph/9506072.
- [82] Masatoshi Shoji and Eiichiro Komatsu. “Massive Neutrinos in Cosmology: Analytic Solutions and Fluid Approximation”. In: *Phys. Rev. D* 81 (2010). [Erratum: *Phys.Rev.D* 82, 089901 (2010)], p. 123516. DOI: [10.1103/PhysRevD.81.123516](https://doi.org/10.1103/PhysRevD.81.123516). arXiv: 1003.0942 [astro-ph.CO].
- [83] Julian B. Muñoz and Cora Dvorkin. “Efficient Computation of Galaxy Bias with Neutrinos and Other Relics”. In: *Phys. Rev. D* 98.4 (2018), p. 043503. DOI: [10.1103/PhysRevD.98.043503](https://doi.org/10.1103/PhysRevD.98.043503). arXiv: 1805.11623 [astro-ph.CO].
- [84] Caio Nascimento. “Accurate fluid approximation for massive neutrinos in cosmology”. In: *Phys. Rev. D* 108.2 (2023), p. 023505. DOI: [10.1103/PhysRevD.108.023505](https://doi.org/10.1103/PhysRevD.108.023505). arXiv: 2303.09580 [astro-ph.CO].

- [85] Wayne Hu, Daniel J. Eisenstein, and Max Tegmark. “Weighing neutrinos with galaxy surveys”. In: *Phys. Rev. Lett.* 80 (1998), pp. 5255–5258. doi: [10.1103/PhysRevLett.80.5255](#). arXiv: [astro-ph/9712057](#).
- [86] Emanuele Castorina et al. “Cosmology with massive neutrinos II: on the universality of the halo mass function and bias”. In: *JCAP* 02 (2014), p. 049. doi: [10.1088/1475-7516/2014/02/049](#). arXiv: [1311.1212 \[astro-ph.CO\]](#).
- [87] Federico Marulli et al. “Effects of Massive Neutrinos on the Large-Scale Structure of the Universe”. In: *Mon. Not. Roy. Astron. Soc.* 418 (2011), p. 346. doi: [10.1111/j.1365-2966.2011.19488.x](#). arXiv: [1103.0278 \[astro-ph.CO\]](#).
- [88] Francisco Villaescusa-Navarro et al. “The imprint of neutrinos on clustering in redshift-space”. In: *Astrophys. J.* 861.1 (2018), p. 53. doi: [10.3847/1538-4357/aac6bf](#). arXiv: [1708.01154 \[astro-ph.CO\]](#).
- [89] Mehrdad Mirbabayi, Fabian Schmidt, and Matias Zaldarriaga. “Biased Tracers and Time Evolution”. In: *JCAP* 07 (2015), p. 030. doi: [10.1088/1475-7516/2015/07/030](#). arXiv: [1412.5169 \[astro-ph.CO\]](#).
- [90] Chi-Ting Chiang et al. “Scale-dependent bias and bispectrum in neutrino separate universe simulations”. In: *Phys. Rev. D* 97.12 (2018), p. 123526. doi: [10.1103/PhysRevD.97.123526](#). arXiv: [1710.01310 \[astro-ph.CO\]](#).
- [91] Raul Angulo et al. “On the Statistics of Biased Tracers in the Effective Field Theory of Large Scale Structures”. In: *JCAP* 09 (2015), p. 029. doi: [10.1088/1475-7516/2015/9/029](#). arXiv: [1503.08826 \[astro-ph.CO\]](#).
- [92] Andrea Oddo et al. “Cosmological parameters from the likelihood analysis of the galaxy power spectrum and bispectrum in real space”. In: *JCAP* 11 (2021), p. 038. doi: [10.1088/1475-7516/2021/11/038](#). arXiv: [2108.03204 \[astro-ph.CO\]](#).
- [93] Lucía Fonseca de la Bella et al. “The matter power spectrum in redshift space using effective field theory”. In: *JCAP* 11 (2017), p. 039. doi: [10.1088/1475-7516/2017/11/039](#). arXiv: [1704.05309 \[astro-ph.CO\]](#).
- [94] Shun Saito, Masahiro Takada, and Atsushi Taruya. “Impact of massive neutrinos on nonlinear matter power spectrum”. In: *Phys. Rev. Lett.* 100 (2008), p. 191301. doi: [10.1103/PhysRevLett.100.191301](#). arXiv: [0801.0607 \[astro-ph\]](#).
- [95] Yvonne Y. Y. Wong. “Higher order corrections to the large scale matter power spectrum in the presence of massive neutrinos”. In: *JCAP* 10 (2008), p. 035. doi: [10.1088/1475-7516/2008/10/035](#). arXiv: [0809.0693 \[astro-ph\]](#).
- [96] Amol Upadhye et al. “Large-Scale Structure Formation with Massive Neutrinos and Dynamical Dark Energy”. In: *Phys. Rev. D* 89.10 (2014), p. 103515. doi: [10.1103/PhysRevD.89.103515](#). arXiv: [1309.5872 \[astro-ph.CO\]](#).
- [97] Alejandro Aviles et al. “Clustering in massive neutrino cosmologies via Eulerian Perturbation Theory”. In: *JCAP* 11 (2021), p. 028. doi: [10.1088/1475-7516/2021/11/028](#). arXiv: [2106.13771 \[astro-ph.CO\]](#).

- [98] Hernán E. Noriega et al. “Fast computation of non-linear power spectrum in cosmologies with massive neutrinos”. In: *JCAP* 11 (2022), p. 038. DOI: [10.1088/1475-7516/2022/11/038](https://doi.org/10.1088/1475-7516/2022/11/038). arXiv: [2208.02791](https://arxiv.org/abs/2208.02791) [astro-ph.CO].
- [99] Joseph E. McEwen et al. “FAST-PT: a novel algorithm to calculate convolution integrals in cosmological perturbation theory”. In: *JCAP* 09 (2016), p. 015. DOI: [10.1088/1475-7516/2016/09/015](https://doi.org/10.1088/1475-7516/2016/09/015). arXiv: [1603.04826](https://arxiv.org/abs/1603.04826) [astro-ph.CO].
- [100] Mikhail M. Ivanov and Sergey Sibiryakov. “Infrared Resummation for Biased Tracers in Redshift Space”. In: *JCAP* 07 (2018), p. 053. DOI: [10.1088/1475-7516/2018/07/053](https://doi.org/10.1088/1475-7516/2018/07/053). arXiv: [1804.05080](https://arxiv.org/abs/1804.05080) [astro-ph.CO].
- [101] Francisco Villaescusa-Navarro et al. “The Quijote simulations”. In: *Astrophys. J. Suppl.* 250.1 (2020), p. 2. DOI: [10.3847/1538-4365/ab9d82](https://doi.org/10.3847/1538-4365/ab9d82). arXiv: [1909.05273](https://arxiv.org/abs/1909.05273) [astro-ph.CO].
- [102] Matteo Zennaro et al. “Initial Conditions for Accurate N-Body Simulations of Massive Neutrino Cosmologies”. In: *Mon. Not. Roy. Astron. Soc.* 466.3 (2017), pp. 3244–3258. DOI: [10.1093/mnras/stw3340](https://doi.org/10.1093/mnras/stw3340). arXiv: [1605.05283](https://arxiv.org/abs/1605.05283) [astro-ph.CO].
- [103] G. Parimbelli et al. “The effects of massive neutrinos on the linear point of the correlation function”. In: *JCAP* 01 (2021), p. 009. DOI: [10.1088/1475-7516/2021/01/009](https://doi.org/10.1088/1475-7516/2021/01/009). arXiv: [2007.10345](https://arxiv.org/abs/2007.10345) [astro-ph.CO].
- [104] Philippe Baratta et al. “COVMOS: a new Monte Carlo approach for galaxy clustering analysis”. In: *Astron. Astrophys.* 673 (2023), A1. DOI: [10.1051/0004-6361/202245683](https://doi.org/10.1051/0004-6361/202245683). arXiv: [2211.13590](https://arxiv.org/abs/2211.13590) [astro-ph.CO].
- [105] Matteo Esposito et al. “Evolution mapping – II. Describing statistics of the non-linear cosmic velocity field”. In: *Mon. Not. Roy. Astron. Soc.* 534.4 (2024), pp. 3906–3915. DOI: [10.1093/mnras/stae2351](https://doi.org/10.1093/mnras/stae2351). arXiv: [2406.08539](https://arxiv.org/abs/2406.08539) [astro-ph.CO].
- [106] Johannes U. Lange. “Nautilus: boosting Bayesian importance nested sampling with deep learning”. In: *Mon. Not. Roy. Astron. Soc.* 525.2 (2023), pp. 3181–3194. DOI: [10.1093/mnras/stad2441](https://doi.org/10.1093/mnras/stad2441). arXiv: [2306.16923](https://arxiv.org/abs/2306.16923) [astro-ph.IM].
- [107] Kwan Chuen Chan, Roman Scoccimarro, and Ravi K. Sheth. “Gravity and Large-Scale Non-local Bias”. In: *Phys. Rev. D* 85 (2012), p. 083509. DOI: [10.1103/PhysRevD.85.083509](https://doi.org/10.1103/PhysRevD.85.083509). arXiv: [1201.3614](https://arxiv.org/abs/1201.3614) [astro-ph.CO].
- [108] Tobias Baldauf et al. “Evidence for Quadratic Tidal Tensor Bias from the Halo Bispectrum”. In: *Phys. Rev. D* 86 (2012), p. 083540. DOI: [10.1103/PhysRevD.86.083540](https://doi.org/10.1103/PhysRevD.86.083540). arXiv: [1201.4827](https://arxiv.org/abs/1201.4827) [astro-ph.CO].
- [109] Alexander Eggemeier, Roman Scoccimarro, and Robert E. Smith. “Bias Loop Corrections to the Galaxy Bispectrum”. In: *Phys. Rev. D* 99.12 (2019), p. 123514. DOI: [10.1103/PhysRevD.99.123514](https://doi.org/10.1103/PhysRevD.99.123514). arXiv: [1812.03208](https://arxiv.org/abs/1812.03208) [astro-ph.CO].
- [110] Robert E. Smith. “Covariance of cross-correlations: towards efficient measures for large-scale structure”. In: *Mon. Not. Roy. Astron. Soc.* 400 (2009), p. 851. DOI: [10.1111/j.1365-2966.2009.15490.x](https://doi.org/10.1111/j.1365-2966.2009.15490.x). arXiv: [0810.1960](https://arxiv.org/abs/0810.1960) [astro-ph].

- [111] Jan Niklas Grieb et al. “Gaussian covariance matrices for anisotropic galaxy clustering measurements”. In: *Mon. Not. Roy. Astron. Soc.* 457.2 (2016), pp. 1577–1592. DOI: [10.1093/mnras/stw065](https://doi.org/10.1093/mnras/stw065). arXiv: 1509.04293 [astro-ph.CO].
- [112] A. Fumagalli et al. “Euclid preparation - XXXV. Covariance model validation for the two-point correlation function of galaxy clusters”. In: *Astron. Astrophys.* 683 (2024), A253. DOI: [10.1051/0004-6361/202245540](https://doi.org/10.1051/0004-6361/202245540). arXiv: 2211.12965 [astro-ph.CO].
- [113] Mikhail M. Ivanov, Marko Simonović, and Matias Zaldarriaga. “Cosmological Parameters and Neutrino Masses from the Final Planck and Full-Shape BOSS Data”. In: *Phys. Rev. D* 101.8 (2020), p. 083504. DOI: [10.1103/PhysRevD.101.083504](https://doi.org/10.1103/PhysRevD.101.083504). arXiv: 1912.08208 [astro-ph.CO].
- [114] Agne Semenaite et al. “Beyond Λ CDM constraints from the full shape clustering measurements from BOSS and eBOSS”. In: *Mon. Not. Roy. Astron. Soc.* 521.4 (2023), pp. 5013–5025. DOI: [10.1093/mnras/stad849](https://doi.org/10.1093/mnras/stad849). arXiv: 2210.07304 [astro-ph.CO].
- [115] A. G. Adame et al. “DESI 2024 VII: cosmological constraints from the full-shape modeling of clustering measurements”. In: *JCAP* 2507 (2025), p. 028. DOI: [10.1088/1475-7516/2025/07/028](https://doi.org/10.1088/1475-7516/2025/07/028). arXiv: 2411.12022 [astro-ph.CO].
- [116] Vincenzo Mainieri et al. “The Wide-field Spectroscopic Telescope (WST) Science White Paper”. In: (Mar. 2024). arXiv: 2403.05398 [astro-ph.IM].
- [117] Jacob Brandbyge and Steen Hannestad. “Grid Based Linear Neutrino Perturbations in Cosmological N-body Simulations”. In: *JCAP* 05 (2009), p. 002. DOI: [10.1088/1475-7516/2009/05/002](https://doi.org/10.1088/1475-7516/2009/05/002). arXiv: 0812.3149 [astro-ph].
- [118] Jacob Brandbyge and Steen Hannestad. “Resolving Cosmic Neutrino Structure: A Hybrid Neutrino N-body Scheme”. In: *JCAP* 01 (2010), p. 021. DOI: [10.1088/1475-7516/2010/01/021](https://doi.org/10.1088/1475-7516/2010/01/021). arXiv: 0908.1969 [astro-ph.CO].
- [119] Matteo Viel, Martin G. Haehnelt, and Volker Springel. “The effect of neutrinos on the matter distribution as probed by the Intergalactic Medium”. In: *JCAP* 06 (2010), p. 015. DOI: [10.1088/1475-7516/2010/06/015](https://doi.org/10.1088/1475-7516/2010/06/015). arXiv: 1003.2422 [astro-ph.CO].
- [120] Yacine Ali-Haïmoud and Simeon Bird. “An efficient implementation of massive neutrinos in non-linear structure formation simulations”. In: *Mon. Not. Roy. Astron. Soc.* 428 (2012), pp. 3375–3389. DOI: [10.1093/mnras/sts286](https://doi.org/10.1093/mnras/sts286). arXiv: 1209.0461 [astro-ph.CO].
- [121] Jacob Brandbyge et al. “Cosmological N -body simulations including radiation perturbations”. In: *Mon. Not. Roy. Astron. Soc.* 466 (2017), pp. L68–L72. DOI: [10.1093/mnrasl/slw235](https://doi.org/10.1093/mnrasl/slw235). arXiv: 1610.04236 [astro-ph.CO].
- [122] Arka Banerjee and Neal Dalal. “Simulating nonlinear cosmological structure formation with massive neutrinos”. In: *JCAP* 11 (2016), p. 015. DOI: [10.1088/1475-7516/2016/11/015](https://doi.org/10.1088/1475-7516/2016/11/015). arXiv: 1606.06167 [astro-ph.CO].
- [123] Julian Adamek, Ruth Durrer, and Martin Kunz. “Relativistic N-body simulations with massive neutrinos”. In: *JCAP* 11 (2017), p. 004. DOI: [10.1088/1475-7516/2017/11/004](https://doi.org/10.1088/1475-7516/2017/11/004). arXiv: 1707.06938 [astro-ph.CO].
- [124] Jia Liu et al. “MassiveNuS: Cosmological Massive Neutrino Simulations”. In: *JCAP* 03 (2018), p. 049. DOI: [10.1088/1475-7516/2018/03/049](https://doi.org/10.1088/1475-7516/2018/03/049). arXiv: 1711.10524 [astro-ph.CO].

- [125] Jacob Brandbyge, Steen Hannestad, and Thomas Tram. “Momentum space sampling of neutrinos in N -body simulations”. In: *JCAP* 03 (2019), p. 047. DOI: [10.1088/1475-7516/2019/03/047](https://doi.org/10.1088/1475-7516/2019/03/047). arXiv: [1806.05874](https://arxiv.org/abs/1806.05874) [astro-ph.CO].
- [126] Arka Banerjee et al. “Weighing neutrinos with the halo environment”. In: *JCAP* 06 (2020), p. 032. DOI: [10.1088/1475-7516/2020/06/032](https://doi.org/10.1088/1475-7516/2020/06/032). arXiv: [1907.06598](https://arxiv.org/abs/1907.06598) [astro-ph.CO].
- [127] Willem Elbers et al. “An optimal non-linear method for simulating relic neutrinos”. In: *Mon. Not. Roy. Astron. Soc.* 507.2 (2021), pp. 2614–2631. DOI: [10.1093/mnras/stab2260](https://doi.org/10.1093/mnras/stab2260). arXiv: [2010.07321](https://arxiv.org/abs/2010.07321) [astro-ph.CO].
- [128] Willem Elbers et al. “Higher order initial conditions with massive neutrinos”. In: *Mon. Not. Roy. Astron. Soc.* 516.3 (2022), pp. 3821–3836. DOI: [10.1093/mnras/stac2365](https://doi.org/10.1093/mnras/stac2365). arXiv: [2202.00670](https://arxiv.org/abs/2202.00670) [astro-ph.CO].
- [129] Willem Elbers. “Geodesic motion and phase-space evolution of massive neutrinos”. In: *JCAP* 11 (2022), p. 058. DOI: [10.1088/1475-7516/2022/11/058](https://doi.org/10.1088/1475-7516/2022/11/058). arXiv: [2207.14256](https://arxiv.org/abs/2207.14256) [astro-ph.CO].
- [130] Beatriz Hernández-Molinero et al. “Neutrino halo profiles: HR-DEMNUi simulation analysis”. In: *JCAP* 09 (2024), p. 033. DOI: [10.1088/1475-7516/2024/09/033](https://doi.org/10.1088/1475-7516/2024/09/033). arXiv: [2407.12694](https://arxiv.org/abs/2407.12694) [astro-ph.CO].
- [131] César Hernández-Aguayo et al. “The MillenniumTNG project: impact of massive neutrinos on the cosmic large-scale structure and the distribution of galaxies”. In: *Mon. Not. Roy. Astron. Soc.* 540.4 (2025), pp. 3642–3660. DOI: [10.1093/mnras/staf864](https://doi.org/10.1093/mnras/staf864). arXiv: [2407.21103](https://arxiv.org/abs/2407.21103) [astro-ph.CO].
- [132] Ethan O. Nadler et al. “Dark Matter Constraints from a Unified Analysis of Strong Gravitational Lenses and Milky Way Satellite Galaxies”. In: *Astrophys. J.* 917.1 (2021), p. 7. DOI: [10.3847/1538-4357/abf9a3](https://doi.org/10.3847/1538-4357/abf9a3). arXiv: [2101.07810](https://arxiv.org/abs/2101.07810) [astro-ph.CO].
- [133] Oliver Newton et al. “Constraints on the properties of warm dark matter using the satellite galaxies of the Milky Way”. In: *JCAP* 08 (2021), p. 062. DOI: [10.1088/1475-7516/2021/08/062](https://doi.org/10.1088/1475-7516/2021/08/062). arXiv: [2011.08865](https://arxiv.org/abs/2011.08865) [astro-ph.CO].
- [134] Ariane Dekker et al. “Warm dark matter constraints using Milky Way satellite observations and subhalo evolution modeling”. In: *Phys. Rev. D* 106.12 (2022), p. 123026. DOI: [10.1103/PhysRevD.106.123026](https://doi.org/10.1103/PhysRevD.106.123026). arXiv: [2111.13137](https://arxiv.org/abs/2111.13137) [astro-ph.CO].
- [135] Chin Yi Tan, Ariane Dekker, and Alex Drlica-Wagner. “Mixed warm dark matter constraints using Milky Way satellite galaxy counts”. In: *Phys. Rev. D* 111.6 (2025), p. 063079. DOI: [10.1103/PhysRevD.111.063079](https://doi.org/10.1103/PhysRevD.111.063079). arXiv: [2409.18917](https://arxiv.org/abs/2409.18917) [astro-ph.CO].
- [136] Vid Iršič et al. “Unveiling dark matter free streaming at the smallest scales with the high redshift Lyman-alpha forest”. In: *Phys. Rev. D* 109.4 (2024), p. 043511. DOI: [10.1103/PhysRevD.109.043511](https://doi.org/10.1103/PhysRevD.109.043511). arXiv: [2309.04533](https://arxiv.org/abs/2309.04533) [astro-ph.CO].
- [137] Takeshi Kobayashi et al. “Lyman- α constraints on ultralight scalar dark matter: Implications for the early and late universe”. In: *Phys. Rev. D* 96.12 (2017), p. 123514. DOI: [10.1103/PhysRevD.96.123514](https://doi.org/10.1103/PhysRevD.96.123514). arXiv: [1708.00015](https://arxiv.org/abs/1708.00015) [astro-ph.CO].
- [138] Olga Garcia-Gallego et al. “Constraining mixed dark matter models with high-redshift Lyman-alpha forest data”. In: *Phys. Rev. D* 112.4 (2025), p. 043502. DOI: [10.1103/4k29-h991](https://doi.org/10.1103/4k29-h991). arXiv: [2504.06367](https://arxiv.org/abs/2504.06367) [astro-ph.CO].

- [139] Eric Armengaud et al. “Constraining the mass of light bosonic dark matter using SDSS Lyman- α forest”. In: *Mon. Not. Roy. Astron. Soc.* 471.4 (2017), pp. 4606–4614. DOI: [10.1093/mnras/stx1870](#). arXiv: [1703.09126 \[astro-ph.CO\]](#).
- [140] Ioana A. Zelko et al. “Constraints on Sterile Neutrino Models from Strong Gravitational Lensing, Milky Way Satellites, and the Lyman- α Forest”. In: *Phys. Rev. Lett.* 129.19 (2022), p. 191301. DOI: [10.1103/PhysRevLett.129.191301](#). arXiv: [2205.09777 \[hep-ph\]](#).
- [141] S. Vegetti et al. “Strong Gravitational Lensing as a Probe of Dark Matter”. In: *Space Sci. Rev.* 220.5 (2024), p. 58. DOI: [10.1007/s11214-024-01087-w](#). arXiv: [2306.11781 \[astro-ph.CO\]](#).
- [142] Ryan E. Keeley et al. “Pushing the limits of detectability: mixed dark matter from strong gravitational lenses”. In: *Mon. Not. Roy. Astron. Soc.* 524.4 (2023), pp. 6159–6166. DOI: [10.1093/mnras/stad2251](#). arXiv: [2301.07265 \[astro-ph.CO\]](#).
- [143] Erminia Calabrese et al. “The Atacama Cosmology Telescope: DR6 constraints on extended cosmological models”. In: *JCAP* 11 (2025), p. 063. DOI: [10.1088/1475-7516/2025/11/063](#). arXiv: [2503.14454 \[astro-ph.CO\]](#).
- [144] C. Garcia-Quintero et al. “Cosmological implications of DESI DR2 BAO measurements in light of the latest ACT DR6 CMB data”. In: *Phys. Rev. D* 112.8 (2025), p. 083529. DOI: [10.1103/d6yc-xpqb](#). arXiv: [2504.18464 \[astro-ph.CO\]](#).
- [145] E. Camphuis et al. “SPT-3G D1: CMB temperature and polarization power spectra and cosmology from 2019 and 2020 observations of the SPT-3G Main field”. In: (June 2025). arXiv: [2506.20707 \[astro-ph.CO\]](#).
- [146] John Preskill, Mark B. Wise, and Frank Wilczek. “Cosmology of the Invisible Axion”. In: *Phys. Lett. B* 120 (1983). Ed. by M. A. Srednicki, pp. 127–132. DOI: [10.1016/0370-2693\(83\)90637-8](#).
- [147] L. F. Abbott and P. Sikivie. “A Cosmological Bound on the Invisible Axion”. In: *Phys. Lett. B* 120 (1983). Ed. by M. A. Srednicki, pp. 133–136. DOI: [10.1016/0370-2693\(83\)90638-X](#).
- [148] Michael Dine and Willy Fischler. “The Not So Harmless Axion”. In: *Phys. Lett. B* 120 (1983). Ed. by M. A. Srednicki, pp. 137–141. DOI: [10.1016/0370-2693\(83\)90639-1](#).
- [149] Steven Weinberg. “Cosmological Constraints on the Scale of Supersymmetry Breaking”. In: *Phys. Rev. Lett.* 48 (1982), p. 1303. DOI: [10.1103/PhysRevLett.48.1303](#).
- [150] Jonathan L. Feng, Marc Kamionkowski, and Samuel K. Lee. “Light Gravitinos at Colliders and Implications for Cosmology”. In: *Phys. Rev. D* 82 (2010), p. 015012. DOI: [10.1103/PhysRevD.82.015012](#). arXiv: [1004.4213 \[hep-ph\]](#).
- [151] Francis-Yan Cyr-Racine et al. “ETHOS—an effective theory of structure formation: From dark particle physics to the matter distribution of the Universe”. In: *Phys. Rev. D* 93.12 (2016), p. 123527. DOI: [10.1103/PhysRevD.93.123527](#). arXiv: [1512.05344 \[astro-ph.CO\]](#).
- [152] Niklas Becker et al. “Cosmological constraints on multi-interacting dark matter”. In: *JCAP* 02 (2021), p. 019. DOI: [10.1088/1475-7516/2021/02/019](#). arXiv: [2010.04074 \[astro-ph.CO\]](#).
- [153] LSST Science Collaboration, P. A. Abell, J. Allison, et al. “LSST Science Book, Version 2.0”. In: (2009). arXiv: [0912.0201 \[astro-ph.IM\]](#).

- [154] Kimberly K. Boddy et al. “Snowmass 2021 theory frontier white paper: Astrophysical and cosmological probes of dark matter”. In: *JHEAp* 35 (2022), pp. 112–138. doi: [10.1016/j.jheap.2022.06.005](https://doi.org/10.1016/j.jheap.2022.06.005). arXiv: [2203.06380](https://arxiv.org/abs/2203.06380) [hep-ph].
- [155] Yao-Yuan Mao et al. “Snowmass 2021: Vera C. Rubin Observatory as a Flagship Dark Matter Experiment”. In: (Mar. 2022). arXiv: [2203.07252](https://arxiv.org/abs/2203.07252) [hep-ex].
- [156] J. Lesgourgues et al. “Euclid preparation - LVI. Sensitivity to non-standard particle dark matter models”. In: *Astron. Astrophys.* 693 (2025), A249. doi: [10.1051/0004-6361/202451611](https://doi.org/10.1051/0004-6361/202451611). arXiv: [2406.18274](https://arxiv.org/abs/2406.18274) [astro-ph.CO].
- [157] Cora Dvorkin et al. “Snowmass 2021: The Physics of Light Relics”. In: (). arXiv: [2203.07943](https://arxiv.org/abs/2203.07943) [hep-ph].
- [158] Julien Lesgourgues et al. “Non-linear Power Spectrum including Massive Neutrinos: the Time-RG Flow Approach”. In: *JCAP* 06 (2009), p. 017. doi: [10.1088/1475-7516/2009/06/017](https://doi.org/10.1088/1475-7516/2009/06/017). arXiv: [0901.4550](https://arxiv.org/abs/0901.4550) [astro-ph.CO].
- [159] Diego Blas et al. “Structure formation with massive neutrinos: going beyond linear theory”. In: *JCAP* 11 (2014), p. 039. doi: [10.1088/1475-7516/2014/11/039](https://doi.org/10.1088/1475-7516/2014/11/039). arXiv: [1408.2995](https://arxiv.org/abs/1408.2995) [astro-ph.CO].
- [160] Florian F hrer and Yvonne Y. Y. Wong. “Higher-order massive neutrino perturbations in large-scale structure”. In: *JCAP* 03 (2015), p. 046. doi: [10.1088/1475-7516/2015/03/046](https://doi.org/10.1088/1475-7516/2015/03/046). arXiv: [1412.2764](https://arxiv.org/abs/1412.2764) [astro-ph.CO].
- [161] Maria Archidiacono and Steen Hannestad. “Efficient calculation of cosmological neutrino clustering in the non-linear regime”. In: *JCAP* 06 (2016), p. 018. doi: [10.1088/1475-7516/2016/06/018](https://doi.org/10.1088/1475-7516/2016/06/018). arXiv: [1510.02907](https://arxiv.org/abs/1510.02907) [astro-ph.CO].
- [162] Michele Levi and Zvonimir Vlah. “Massive neutrinos in nonlinear large scale structure: A consistent perturbation theory”. In: (May 2016). arXiv: [1605.09417](https://arxiv.org/abs/1605.09417) [astro-ph.CO].
- [163] Leonardo Senatore and Matias Zaldarriaga. “The Effective Field Theory of Large-Scale Structure in the presence of Massive Neutrinos”. In: (July 2017). arXiv: [1707.04698](https://arxiv.org/abs/1707.04698) [astro-ph.CO].
- [164] Rossana Ruggeri et al. “DEMNUi: Massive neutrinos and the bispectrum of large scale structures”. In: *JCAP* 03 (2018), p. 003. doi: [10.1088/1475-7516/2018/03/003](https://doi.org/10.1088/1475-7516/2018/03/003). arXiv: [1712.02334](https://arxiv.org/abs/1712.02334) [astro-ph.CO].
- [165] Roger de Belsunce and Leonardo Senatore. “Tree-Level Bispectrum in the Effective Field Theory of Large-Scale Structure extended to Massive Neutrinos”. In: *JCAP* 02 (2019), p. 038. doi: [10.1088/1475-7516/2019/02/038](https://doi.org/10.1088/1475-7516/2019/02/038). arXiv: [1804.06849](https://arxiv.org/abs/1804.06849) [astro-ph.CO].
- [166] Shi-Fan Chen, Emanuele Castorina, and Martin White. “Biased Tracers of Two Fluids in the Lagrangian Picture”. In: *JCAP* 06 (2019), p. 006. doi: [10.1088/1475-7516/2019/06/006](https://doi.org/10.1088/1475-7516/2019/06/006). arXiv: [1903.00437](https://arxiv.org/abs/1903.00437) [astro-ph.CO].
- [167] Hong-Ming Zhu and Emanuele Castorina. “Measuring dark matter-neutrino relative velocity on cosmological scales”. In: *Phys. Rev. D* 101.2 (2020), p. 023525. doi: [10.1103/PhysRevD.101.023525](https://doi.org/10.1103/PhysRevD.101.023525). arXiv: [1905.00361](https://arxiv.org/abs/1905.00361) [astro-ph.CO].

- [168] Alejandro Aviles and Arka Banerjee. “A Lagrangian Perturbation Theory in the presence of massive neutrinos”. In: *JCAP* 10 (2020), p. 034. DOI: [10.1088/1475-7516/2020/10/034](https://doi.org/10.1088/1475-7516/2020/10/034). arXiv: [2007.06508](https://arxiv.org/abs/2007.06508) [astro-ph.CO].
- [169] Mathias Garny and Petter Taule. “Loop corrections to the power spectrum for massive neutrino cosmologies with full time- and scale-dependence”. In: *JCAP* 01 (2021), p. 020. DOI: [10.1088/1475-7516/2021/01/020](https://doi.org/10.1088/1475-7516/2021/01/020). arXiv: [2008.00013](https://arxiv.org/abs/2008.00013) [astro-ph.CO].
- [170] Mathias Garny and Petter Taule. “Two-loop power spectrum with full time- and scale-dependence and EFT corrections: impact of massive neutrinos and going beyond EdS”. In: *JCAP* 09 (2022), p. 054. DOI: [10.1088/1475-7516/2022/09/054](https://doi.org/10.1088/1475-7516/2022/09/054). arXiv: [2205.11533](https://arxiv.org/abs/2205.11533) [astro-ph.CO].
- [171] Caio Nascimento and Marilena Loverde. “Neutrino winds on the sky”. In: *JCAP* 11 (2023), p. 036. DOI: [10.1088/1475-7516/2023/11/036](https://doi.org/10.1088/1475-7516/2023/11/036). arXiv: [2307.00049](https://arxiv.org/abs/2307.00049) [astro-ph.CO].
- [172] Matthew Lewandowski, Ashley Perko, and Leonardo Senatore. “Analytic Prediction of Baryonic Effects from the EFT of Large Scale Structures”. In: *JCAP* 05 (2015), p. 019. DOI: [10.1088/1475-7516/2015/05/019](https://doi.org/10.1088/1475-7516/2015/05/019). arXiv: [1412.5049](https://arxiv.org/abs/1412.5049) [astro-ph.CO].
- [173] Diogo P. L. Bragança et al. “Baryonic effects in the Effective Field Theory of Large-Scale Structure and an analytic recipe for lensing in CMB-S4”. In: *JCAP* 10 (2021), p. 074. DOI: [10.1088/1475-7516/2021/10/074](https://doi.org/10.1088/1475-7516/2021/10/074). arXiv: [2010.02929](https://arxiv.org/abs/2010.02929) [astro-ph.CO].
- [174] Wayne Hu. “Structure formation with generalized dark matter”. In: *Astrophys. J.* 506 (1998), pp. 485–494. DOI: [10.1086/306274](https://doi.org/10.1086/306274). arXiv: [astro-ph/9801234](https://arxiv.org/abs/astro-ph/9801234).
- [175] Xinyu Li, Lam Hui, and Greg L. Bryan. “Numerical and Perturbative Computations of the Fuzzy Dark Matter Model”. In: *Phys. Rev. D* 99.6 (2019), p. 063509. DOI: [10.1103/PhysRevD.99.063509](https://doi.org/10.1103/PhysRevD.99.063509). arXiv: [1810.01915](https://arxiv.org/abs/1810.01915) [astro-ph.CO].
- [176] Farshad Kamalinejad and Zachary Slepian. “A Simple Analytic Treatment of Neutrino Mass Impact on the Full Power Spectrum Shape via a Two-Fluid Approximation”. In: (Mar. 2022). arXiv: [2203.13103](https://arxiv.org/abs/2203.13103) [astro-ph.CO].
- [177] R. D. Peccei and Helen R. Quinn. “CP Conservation in the Presence of Instantons”. In: *Phys. Rev. Lett.* 38 (1977), pp. 1440–1443. DOI: [10.1103/PhysRevLett.38.1440](https://doi.org/10.1103/PhysRevLett.38.1440).
- [178] Steven Weinberg. “A New Light Boson?” In: *Phys. Rev. Lett.* 40 (1978), pp. 223–226. DOI: [10.1103/PhysRevLett.40.223](https://doi.org/10.1103/PhysRevLett.40.223).
- [179] Frank Wilczek. “Problem of Strong P and T Invariance in the Presence of Instantons”. In: *Phys. Rev. Lett.* 40 (1978), pp. 279–282. DOI: [10.1103/PhysRevLett.40.279](https://doi.org/10.1103/PhysRevLett.40.279).
- [180] Peter Svrcek and Edward Witten. “Axions In String Theory”. In: *JHEP* 0606 (2006), p. 051. DOI: [10.1088/1126-6708/2006/06/051](https://doi.org/10.1088/1126-6708/2006/06/051). arXiv: [hep-th/0605206](https://arxiv.org/abs/hep-th/0605206).
- [181] Luca Amendola and Riccardo Barbieri. “Dark matter from an ultra-light pseudo-Goldstone-boson”. In: *Phys. Lett. B* 642 (2006), pp. 192–196. DOI: [10.1016/j.physletb.2006.08.069](https://doi.org/10.1016/j.physletb.2006.08.069). arXiv: [hep-ph/0509257](https://arxiv.org/abs/hep-ph/0509257).
- [182] David J. E. Marsh and Pedro G. Ferreira. “Ultra-Light Scalar Fields and the Growth of Structure in the Universe”. In: *Phys. Rev. D* 82 (2010), p. 103528. DOI: [10.1103/PhysRevD.82.103528](https://doi.org/10.1103/PhysRevD.82.103528). arXiv: [1009.3501](https://arxiv.org/abs/1009.3501) [hep-ph].

- [183] Renée Hlozek et al. “A search for ultralight axions using precision cosmological data”. In: *Phys. Rev. D* 91.10 (2015), p. 103512. DOI: [10.1103/PhysRevD.91.103512](https://doi.org/10.1103/PhysRevD.91.103512). arXiv: [1410.2896](https://arxiv.org/abs/1410.2896) [astro-ph.CO].
- [184] Vivian Poulin et al. “Cosmological implications of ultralight axionlike fields”. In: *Phys. Rev. D* 98.8 (2018), p. 083525. DOI: [10.1103/PhysRevD.98.083525](https://doi.org/10.1103/PhysRevD.98.083525). arXiv: [1806.10608](https://arxiv.org/abs/1806.10608) [astro-ph.CO].
- [185] Tristan L. Smith, Vivian Poulin, and Mustafa A. Amin. “Oscillating scalar fields and the Hubble tension: a resolution with novel signatures”. In: *Phys. Rev. D* 101.6 (2020), p. 063523. DOI: [10.1103/PhysRevD.101.063523](https://doi.org/10.1103/PhysRevD.101.063523). arXiv: [1908.06995](https://arxiv.org/abs/1908.06995) [astro-ph.CO].
- [186] Tessa Cookmeyer et al. “How sound are our ultralight axion approximations?” In: *Phys. Rev. D* 101.2 (2020), p. 023501. DOI: [10.1103/PhysRevD.101.023501](https://doi.org/10.1103/PhysRevD.101.023501). arXiv: [1909.11094](https://arxiv.org/abs/1909.11094) [astro-ph.CO].
- [187] Samuel Passaglia and Wayne Hu. “Accurate effective fluid approximation for ultralight axions”. In: *Phys. Rev. D* 105.12 (2022), p. 123529. DOI: [10.1103/PhysRevD.105.123529](https://doi.org/10.1103/PhysRevD.105.123529). arXiv: [2201.10238](https://arxiv.org/abs/2201.10238) [astro-ph.CO].
- [188] L. Arturo Ureña-López and Francisco X. Linares Cedeño. “Cosmological evolution of scalar field dark matter in the CLASS code: Accuracy and precision of numerical solutions”. In: *Phys. Rev. D* 109.2 (2024), p. 023512. DOI: [10.1103/PhysRevD.109.023512](https://doi.org/10.1103/PhysRevD.109.023512). arXiv: [2307.05600](https://arxiv.org/abs/2307.05600) [astro-ph.CO].
- [189] Alex Laguë et al. “Constraining ultralight axions with galaxy surveys”. In: *JCAP* 01.01 (2022), p. 049. DOI: [10.1088/1475-7516/2022/01/049](https://doi.org/10.1088/1475-7516/2022/01/049). arXiv: [2104.07802](https://arxiv.org/abs/2104.07802) [astro-ph.CO].
- [190] David J. E. Marsh et al. “Ultra-light Axions: Degeneracies with Massive Neutrinos and Forecasts for Future Cosmological Observations”. In: *Phys. Rev. D* 85 (2012), p. 103514. DOI: [10.1103/PhysRevD.85.103514](https://doi.org/10.1103/PhysRevD.85.103514). arXiv: [1110.0502](https://arxiv.org/abs/1110.0502) [astro-ph.CO].
- [191] Diego Blas, Julien Lesgourgues, and Thomas Tram. “The Cosmic Linear Anisotropy Solving System (CLASS) II: Approximation schemes”. In: *JCAP* 07 (2011), p. 034. DOI: [10.1088/1475-7516/2011/07/034](https://doi.org/10.1088/1475-7516/2011/07/034). arXiv: [1104.2933](https://arxiv.org/abs/1104.2933) [astro-ph.CO].
- [192] Julien Lesgourgues and Thomas Tram. “The Cosmic Linear Anisotropy Solving System (CLASS) IV: efficient implementation of non-cold relics”. In: *JCAP* 09 (2011), p. 032. DOI: [10.1088/1475-7516/2011/09/032](https://doi.org/10.1088/1475-7516/2011/09/032). arXiv: [1104.2935](https://arxiv.org/abs/1104.2935) [astro-ph.CO].
- [193] Maria Archidiacono et al. “Unveiling dark fifth forces with linear cosmology”. In: *JCAP* 10 (2022), p. 074. DOI: [10.1088/1475-7516/2022/10/074](https://doi.org/10.1088/1475-7516/2022/10/074). arXiv: [2204.08484](https://arxiv.org/abs/2204.08484) [astro-ph.CO].
- [194] Salvatore Bottaro et al. “Unveiling Dark Forces with Measurements of the Large Scale Structure of the Universe”. In: *Phys. Rev. Lett.* 132.20 (2024), p. 201002. DOI: [10.1103/PhysRevLett.132.201002](https://doi.org/10.1103/PhysRevLett.132.201002). arXiv: [2309.11496](https://arxiv.org/abs/2309.11496) [astro-ph.CO].
- [195] Farshad Kamalinejad and Zachary Slepian. “A Non-Degenerate Neutrino Mass Signature in the Galaxy Bispectrum”. In: (Nov. 2020). arXiv: [2011.00899](https://arxiv.org/abs/2011.00899) [astro-ph.CO].
- [196] Guido D’Amico et al. “The Large Scale Structure bootstrap: perturbation theory and bias expansion from symmetries”. In: *JCAP* 10 (2021), p. 069. DOI: [10.1088/1475-7516/2021/10/069](https://doi.org/10.1088/1475-7516/2021/10/069). arXiv: [2109.09573](https://arxiv.org/abs/2109.09573) [astro-ph.CO].

- [197] Marco Marinucci, Kevin Pardede, and Massimo Pietroni. “Bootstrapping Lagrangian perturbation theory for the Large Scale Structure”. In: *JCAP* 10 (2024), p. 051. doi: [10.1088/1475-7516/2024/10/051](https://doi.org/10.1088/1475-7516/2024/10/051). arXiv: [2405.08413](https://arxiv.org/abs/2405.08413) [astro-ph.CO].
- [198] Fabian Schmidt. “Effect of relative velocity and density perturbations between baryons and dark matter on the clustering of galaxies”. In: *Phys. Rev. D* 94.6 (2016), p. 063508. doi: [10.1103/PhysRevD.94.063508](https://doi.org/10.1103/PhysRevD.94.063508). arXiv: [1602.09059](https://arxiv.org/abs/1602.09059) [astro-ph.CO].
- [199] Chi-Ting Chiang, Marilena LoVerde, and Francisco Villaescusa-Navarro. “First detection of scale-dependent linear halo bias in N -body simulations with massive neutrinos”. In: *Phys. Rev. Lett.* 122.4 (2019), p. 041302. doi: [10.1103/PhysRevLett.122.041302](https://doi.org/10.1103/PhysRevLett.122.041302). arXiv: [1811.12412](https://arxiv.org/abs/1811.12412) [astro-ph.CO].
- [200] Mikhail M. Ivanov et al. “Cosmological constraints without nonlinear redshift-space distortions”. In: *Phys. Rev. D* 105.4 (2022), p. 043531. doi: [10.1103/PhysRevD.105.043531](https://doi.org/10.1103/PhysRevD.105.043531). arXiv: [2110.00006](https://arxiv.org/abs/2110.00006) [astro-ph.CO].
- [201] A. G. Adame et al. “DESI 2024 II: sample definitions, characteristics, and two-point clustering statistics”. In: *JCAP* 2507 (2025), p. 017. doi: [10.1088/1475-7516/2025/07/017](https://doi.org/10.1088/1475-7516/2025/07/017). arXiv: [2411.12020](https://arxiv.org/abs/2411.12020) [astro-ph.CO].
- [202] Marcel Schmittfull, Zvonimir Vlah, and Patrick McDonald. “Fast large scale structure perturbation theory using one-dimensional fast Fourier transforms”. In: *Phys. Rev. D* 93.10 (2016), p. 103528. doi: [10.1103/PhysRevD.93.103528](https://doi.org/10.1103/PhysRevD.93.103528). arXiv: [1603.04405](https://arxiv.org/abs/1603.04405) [astro-ph.CO].
- [203] Marko Simonović et al. “Cosmological perturbation theory using the FFTLog: formalism and connection to QFT loop integrals”. In: *JCAP* 04 (2018), p. 030. doi: [10.1088/1475-7516/2018/04/030](https://doi.org/10.1088/1475-7516/2018/04/030). arXiv: [1708.08130](https://arxiv.org/abs/1708.08130) [astro-ph.CO].
- [204] Keir K. Rogers et al. “Ultra-light axions and the S_8 tension: joint constraints from the cosmic microwave background and galaxy clustering”. In: *JCAP* 06 (2023), p. 023. doi: [10.1088/1475-7516/2023/06/023](https://doi.org/10.1088/1475-7516/2023/06/023). arXiv: [2301.08361](https://arxiv.org/abs/2301.08361) [astro-ph.CO].
- [205] Shadab Alam et al. “The Eleventh and Twelfth Data Releases of the Sloan Digital Sky Survey: Final Data from SDSS-III”. In: *Astrophys. J. Suppl.* 219.1 (2015), p. 12. doi: [10.1088/0067-0049/219/1/12](https://doi.org/10.1088/0067-0049/219/1/12). arXiv: [1501.00963](https://arxiv.org/abs/1501.00963) [astro-ph.IM].
- [206] Jesus Torrado and Antony Lewis. “Cobaya: Code for Bayesian Analysis of hierarchical physical models”. In: *JCAP* 05 (2021), p. 057. doi: [10.1088/1475-7516/2021/05/057](https://doi.org/10.1088/1475-7516/2021/05/057). arXiv: [2005.05290](https://arxiv.org/abs/2005.05290) [astro-ph.IM].
- [207] Anton Chudaykin et al. “Nonlinear perturbation theory extension of the Boltzmann code CLASS”. In: *Phys. Rev. D* 102.6 (2020), p. 063533. doi: [10.1103/PhysRevD.102.063533](https://doi.org/10.1103/PhysRevD.102.063533). arXiv: [2004.10607](https://arxiv.org/abs/2004.10607) [astro-ph.CO].
- [208] Renée Hložek et al. “Future CMB tests of dark matter: Ultralight axions and massive neutrinos”. In: *Phys. Rev. D* 95.12 (2017), p. 123511. doi: [10.1103/PhysRevD.95.123511](https://doi.org/10.1103/PhysRevD.95.123511). arXiv: [1607.08208](https://arxiv.org/abs/1607.08208) [astro-ph.CO].
- [209] Renée Hložek, David J. E. Marsh, and Daniel Grin. “Using the Full Power of the Cosmic Microwave Background to Probe Axion Dark Matter”. In: *Mon. Not. Roy. Astron. Soc.* 476.3 (2018), pp. 3063–3085. doi: [10.1093/mnras/sty271](https://doi.org/10.1093/mnras/sty271). arXiv: [1708.05681](https://arxiv.org/abs/1708.05681) [astro-ph.CO].

- [210] Frank J. Qu et al. “The Atacama Cosmology Telescope: A Measurement of the DR6 CMB Lensing Power Spectrum and Its Implications for Structure Growth”. In: *Astrophys. J.* 962.2 (2024), p. 112. DOI: [10.3847/1538-4357/acfe06](https://doi.org/10.3847/1538-4357/acfe06). arXiv: [2304.05202](https://arxiv.org/abs/2304.05202) [astro-ph.CO].
- [211] M. Abitbol et al. “The Simons Observatory: science goals and forecasts for the enhanced Large Aperture Telescope”. In: *JCAP* 08 (2025), p. 034. DOI: [10.1088/1475-7516/2025/08/034](https://doi.org/10.1088/1475-7516/2025/08/034). arXiv: [2503.00636](https://arxiv.org/abs/2503.00636) [astro-ph.IM].
- [212] Marco Costa et al. “Dark forces suppress structure growth”. In: (Sept. 2025). arXiv: [2510.00098](https://arxiv.org/abs/2510.00098) [astro-ph.CO].
- [213] Hamed Manouchehri Kousha, Sina Hooshangi, and Aliakbar Abolhasani. “The Effective Field Theory of Large Scale Structures of a Fuzzy Dark Matter Universe”. In: *Astrophys. J.* 961.1 (2024), p. 131. DOI: [10.3847/1538-4357/ad148f](https://doi.org/10.3847/1538-4357/ad148f). arXiv: [2305.20075](https://arxiv.org/abs/2305.20075) [astro-ph.CO].
- [214] Patrick McDonald. “How to generate a significant effective temperature for cold dark matter, from first principles”. In: *JCAP* 04 (2011), p. 032. DOI: [10.1088/1475-7516/2011/04/032](https://doi.org/10.1088/1475-7516/2011/04/032). arXiv: [0910.1002](https://arxiv.org/abs/0910.1002) [astro-ph.CO].
- [215] Alejandro Aviles. “Dark matter dispersion tensor in perturbation theory”. In: *Phys. Rev. D* 93 (2016), p. 063517. DOI: [10.1103/PhysRevD.93.063517](https://doi.org/10.1103/PhysRevD.93.063517). arXiv: [1512.07198](https://arxiv.org/abs/1512.07198) [astro-ph.CO].
- [216] Giulia Cusin, Vittorio Tansella, and Ruth Durrer. “Vorticity generation in the Universe: A perturbative approach”. In: *Phys. Rev. D* 95.6 (2017), p. 063527. DOI: [10.1103/PhysRevD.95.063527](https://doi.org/10.1103/PhysRevD.95.063527). arXiv: [1612.00783](https://arxiv.org/abs/1612.00783) [astro-ph.CO].
- [217] Alaric Erschfeld and Stefan Floerchinger. “Evolution of dark matter velocity dispersion”. In: *JCAP* 06 (2019), p. 039. DOI: [10.1088/1475-7516/2019/06/039](https://doi.org/10.1088/1475-7516/2019/06/039). arXiv: [1812.06891](https://arxiv.org/abs/1812.06891) [astro-ph.CO].
- [218] Mathias Garny, Dominik Laxhuber, and Roman Scoccimarro. “Perturbation theory with dispersion and higher cumulants: Framework and linear theory”. In: *Phys. Rev. D* 107.6 (2023), p. 063539. DOI: [10.1103/PhysRevD.107.063539](https://doi.org/10.1103/PhysRevD.107.063539). arXiv: [2210.08088](https://arxiv.org/abs/2210.08088) [astro-ph.CO].
- [219] Mathias Garny, Dominik Laxhuber, and Roman Scoccimarro. “Perturbation theory with dispersion and higher cumulants: Nonlinear regime”. In: *Phys. Rev. D* 107.6 (2023), p. 063540. DOI: [10.1103/PhysRevD.107.063540](https://doi.org/10.1103/PhysRevD.107.063540). arXiv: [2210.08089](https://arxiv.org/abs/2210.08089) [astro-ph.CO].
- [220] Masatoshi Shoji and Eiichiro Komatsu. “Third-order Perturbation Theory With Non-linear Pressure”. In: *Astrophys. J.* 700 (2009), pp. 705–719. DOI: [10.1088/0004-637X/700/1/705](https://doi.org/10.1088/0004-637X/700/1/705). arXiv: [0903.2669](https://arxiv.org/abs/0903.2669) [astro-ph.CO].
- [221] N. Kaiser. “On the spatial correlations of Abell clusters.” In: *Astrophys. J. Lett.* 284 (Sept. 1984), pp. L9–L12. DOI: [10.1086/184341](https://doi.org/10.1086/184341).
- [222] G. Efstathiou, W. J. Sutherland, and S. J. Maddox. “The cosmological constant and cold dark matter”. In: *Nature* 348.6303 (Dec. 1990), pp. 705–707. DOI: [10.1038/348705a0](https://doi.org/10.1038/348705a0).
- [223] J. A. Peacock. “The power spectrum of galaxy clustering”. In: *Mon. Not. Roy. Astron. Soc.* 253 (Nov. 1991), 1P–5P. DOI: [10.1093/mnras/253.1.1P](https://doi.org/10.1093/mnras/253.1.1P).
- [224] J. A. Peacock and D. Nicholson. “The large-scale clustering of radio galaxies.” In: *Mon. Not. Roy. Astron. Soc.* 253 (Nov. 1991), pp. 307–319. DOI: [10.1093/mnras/253.2.307](https://doi.org/10.1093/mnras/253.2.307).

- [225] C. M. Baugh and G. Efstathiou. “The three-dimensional power spectrum measured from the APM galaxy survey - I. Use of the angular correlation function.” In: *Mon. Not. Roy. Astron. Soc.* 265 (Nov. 1993), pp. 145–156. DOI: [10.1093/mnras/265.1.145](https://doi.org/10.1093/mnras/265.1.145).
- [226] Mikhail M. Ivanov et al. “Full-shape analysis with simulation-based priors: Constraints on single field inflation from BOSS”. In: *Phys. Rev. D* 110.6, 063538 (Sept. 2024), p. 063538. DOI: [10.1103/PhysRevD.110.063538](https://doi.org/10.1103/PhysRevD.110.063538). arXiv: [2402.13310](https://arxiv.org/abs/2402.13310) [astro-ph.CO].
- [227] A. G. Adame et al. “DESI 2024 V: Full-Shape galaxy clustering from galaxies and quasars”. In: *JCAP* 2509 (2025). [Erratum: *JCAP* 02, E02 (2026)], p. 008. DOI: [10.1088/1475-7516/2025/09/008](https://doi.org/10.1088/1475-7516/2025/09/008). arXiv: [2411.12021](https://arxiv.org/abs/2411.12021) [astro-ph.CO].
- [228] S. Novell-Masot et al. “Full-Shape analysis of the power spectrum and bispectrum of DESI DR1 LRG and QSO samples”. In: *JCAP* 2025.6, 005 (June 2025), p. 005. DOI: [10.1088/1475-7516/2025/06/005](https://doi.org/10.1088/1475-7516/2025/06/005). arXiv: [2503.09714](https://arxiv.org/abs/2503.09714) [astro-ph.CO].
- [229] Andrina Nicola et al. “Tomographic galaxy clustering with the Subaru Hyper Suprime-Cam first year public data release”. In: *JCAP* 2020.3, 044 (Mar. 2020), p. 044. DOI: [10.1088/1475-7516/2020/03/044](https://doi.org/10.1088/1475-7516/2020/03/044). arXiv: [1912.08209](https://arxiv.org/abs/1912.08209) [astro-ph.CO].
- [230] Carlos García-García et al. “The growth of density perturbations in the last 10 billion years from tomographic large-scale structure data”. In: *JCAP* 2021.10, 030 (Oct. 2021), p. 030. DOI: [10.1088/1475-7516/2021/10/030](https://doi.org/10.1088/1475-7516/2021/10/030). arXiv: [2105.12108](https://arxiv.org/abs/2105.12108) [astro-ph.CO].
- [231] Gerrit S. Farren et al. “The Atacama Cosmology Telescope: Cosmology from Cross-correlations of unWISE Galaxies and ACT DR6 CMB Lensing”. In: *Astrophys. J.* 966.2, 157 (May 2024), p. 157. DOI: [10.3847/1538-4357/ad31a5](https://doi.org/10.3847/1538-4357/ad31a5). arXiv: [2309.05659](https://arxiv.org/abs/2309.05659) [astro-ph.CO].
- [232] Noah Sailer et al. “Cosmological constraints from the cross-correlation of DESI Luminous Red Galaxies with CMB lensing from Planck PR4 and ACT DR6”. In: *JCAP* 2025.6, 008 (June 2025), p. 008. DOI: [10.1088/1475-7516/2025/06/008](https://doi.org/10.1088/1475-7516/2025/06/008). arXiv: [2407.04607](https://arxiv.org/abs/2407.04607) [astro-ph.CO].
- [233] Noah Sailer et al. “Evolution of structure growth during dark energy domination: Insights from the cross-correlation of DESI galaxies with CMB lensing and galaxy magnification”. In: *Phys. Rev. D* 111.10, 103540 (May 2025), p. 103540. DOI: [10.1103/27rg-tq8z](https://doi.org/10.1103/27rg-tq8z). arXiv: [2503.24385](https://arxiv.org/abs/2503.24385) [astro-ph.CO].
- [234] Andrea Rubiola et al. “Low-redshift constraints on structure growth from CMB lensing tomography”. In: *arXiv e-prints*, arXiv:2510.09563 (Oct. 2025), arXiv:2510.09563. DOI: [10.48550/arXiv.2510.09563](https://doi.org/10.48550/arXiv.2510.09563). arXiv: [2510.09563](https://arxiv.org/abs/2510.09563) [astro-ph.CO].
- [235] Yi-Kuan Chiang, Brice Ménard, and David Schiminovich. “Broadband Intensity Tomography: Spectral Tagging of the Cosmic UV Background”. In: *Astrophys. J.* 877.2, 150 (June 2019), p. 150. DOI: [10.3847/1538-4357/ab1b35](https://doi.org/10.3847/1538-4357/ab1b35). arXiv: [1810.00885](https://arxiv.org/abs/1810.00885) [astro-ph.CO].
- [236] Nick Koukoufilippas et al. “Tomographic measurement of the intergalactic gas pressure through galaxy-tSZ cross-correlations”. In: *Mon. Not. Roy. Astron. Soc.* 491.4 (Feb. 2020), pp. 5464–5480. DOI: [10.1093/mnras/stz3351](https://doi.org/10.1093/mnras/stz3351). arXiv: [1909.09102](https://arxiv.org/abs/1909.09102) [astro-ph.CO].
- [237] Yi-Kuan Chiang et al. “The Cosmic Thermal History Probed by Sunyaev-Zeldovich Effect Tomography”. In: *Astrophys. J.* 902.1, 56 (Oct. 2020), p. 56. DOI: [10.3847/1538-4357/abb403](https://doi.org/10.3847/1538-4357/abb403). arXiv: [2006.14650](https://arxiv.org/abs/2006.14650) [astro-ph.CO].

- [238] Yi-Kuan Chiang, Ryu Makiya, and Brice Ménard. “Cosmic Infrared Background Tomography and a Census of Cosmic Dust and Star Formation”. In: *Astrophys. J.* 992.1, 65 (Oct. 2025), p. 65. DOI: [10.3847/1538-4357/adfb6a](https://doi.org/10.3847/1538-4357/adfb6a). arXiv: 2504.05384 [astro-ph.CO].
- [239] Sara Maleubre et al. “The impact of galaxy bias on cross-correlation tomography”. In: *arXiv e-prints*, arXiv:2508.05319 (Aug. 2025), arXiv:2508.05319. DOI: [10.48550/arXiv.2508.05319](https://doi.org/10.48550/arXiv.2508.05319). arXiv: 2508.05319 [astro-ph.CO].
- [240] Hironao Miyatake et al. “First Identification of a CMB Lensing Signal Produced by 1.5 Million Galaxies at $z \sim 4$: Constraints on Matter Density Fluctuations at High Redshift”. In: *Phys. Rev. Lett.* 129.6, 061301 (Aug. 2022), p. 061301. DOI: [10.1103/PhysRevLett.129.061301](https://doi.org/10.1103/PhysRevLett.129.061301). arXiv: 2103.15862 [astro-ph.CO].
- [241] Giulia Piccirilli et al. “Growth history and quasar bias evolution at $z < 3$ from Quiaia”. In: *JCAP* 2024.6, 012 (June 2024), p. 012. DOI: [10.1088/1475-7516/2024/06/012](https://doi.org/10.1088/1475-7516/2024/06/012). arXiv: 2402.05761 [astro-ph.CO].
- [242] Carmen Embil Villagra et al. “The Atacama Cosmology Telescope: High-redshift measurement of structure growth from the cross-correlation of Quiaia quasars and CMB lensing from ACT DR6 and *Planck* PR4”. In: *arXiv e-prints*, arXiv:2507.08798 (July 2025), arXiv:2507.08798. DOI: [10.48550/arXiv.2507.08798](https://doi.org/10.48550/arXiv.2507.08798). arXiv: 2507.08798 [astro-ph.CO].
- [243] Andreas Albrecht et al. “Report of the Dark Energy Task Force”. In: *arXiv e-prints*, astro-ph/0609591 (Sept. 2006), astro-ph/0609591. DOI: [10.48550/arXiv.astro-ph/0609591](https://doi.org/10.48550/arXiv.astro-ph/0609591). arXiv: astro-ph/0609591 [astro-ph].
- [244] R. Laureijs et al. “Euclid Definition Study Report”. In: *arXiv e-prints*, arXiv:1110.3193 (Oct. 2011), arXiv:1110.3193. DOI: [10.48550/arXiv.1110.3193](https://doi.org/10.48550/arXiv.1110.3193). arXiv: 1110.3193 [astro-ph.CO].
- [245] The LSST Dark Energy Science Collaboration et al. “The LSST Dark Energy Science Collaboration (DESC) Science Requirements Document”. In: *arXiv e-prints*, arXiv:1809.01669 (Sept. 2018), arXiv:1809.01669. DOI: [10.48550/arXiv.1809.01669](https://doi.org/10.48550/arXiv.1809.01669). arXiv: 1809.01669 [astro-ph.CO].
- [246] D. Spergel et al. “Wide-Field Infrared Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report”. In: *arXiv e-prints*, arXiv:1503.03757 (Mar. 2015), arXiv:1503.03757. DOI: [10.48550/arXiv.1503.03757](https://doi.org/10.48550/arXiv.1503.03757). arXiv: 1503.03757 [astro-ph.IM].
- [247] Héctor Gil-Marín et al. “The power spectrum and bispectrum of SDSS DR11 BOSS galaxies - I. Bias and gravity”. In: *Mon. Not. Roy. Astron. Soc.* 451.1 (July 2015), pp. 539–580. DOI: [10.1093/mnras/stv961](https://doi.org/10.1093/mnras/stv961). arXiv: 1407.5668 [astro-ph.CO].
- [248] Héctor Gil-Marín et al. “The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: RSD measurement from the power spectrum and bispectrum of the DR12 BOSS galaxies”. In: *Mon. Not. Roy. Astron. Soc.* 465.2 (Feb. 2017), pp. 1757–1788. DOI: [10.1093/mnras/stw2679](https://doi.org/10.1093/mnras/stw2679). arXiv: 1606.00439 [astro-ph.CO].
- [249] Naonori S. Sugiyama et al. “Towards a self-consistent analysis of the anisotropic galaxy two- and three-point correlation functions on large scales: application to mock galaxy catalogues”. In: *Mon. Not. Roy. Astron. Soc.* 501.2 (Feb. 2021), pp. 2862–2896. DOI: [10.1093/mnras/staa3725](https://doi.org/10.1093/mnras/staa3725). arXiv: 2010.06179 [astro-ph.CO].

- [250] Oliver H. E. Philcox et al. “Cosmology with the redshift-space galaxy bispectrum monopole at one-loop order”. In: *Phys. Rev. D* 106.4, 043530 (Aug. 2022), p. 043530. DOI: [10.1103/PhysRevD.106.043530](https://doi.org/10.1103/PhysRevD.106.043530). arXiv: [2206.02800](https://arxiv.org/abs/2206.02800) [astro-ph.CO].
- [251] Martin Bucher, Benjamin Racine, and Bartjan van Tent. “The binned bispectrum estimator: template-based and non-parametric CMB non-Gaussianity searches”. In: *JCAP* 2016.5, 055 (May 2016), p. 055. DOI: [10.1088/1475-7516/2016/05/055](https://doi.org/10.1088/1475-7516/2016/05/055). arXiv: [1509.08107](https://arxiv.org/abs/1509.08107) [astro-ph.CO].
- [252] Planck Collaboration et al. “Planck 2018 results. IX. Constraints on primordial non-Gaussianity”. In: *Astron. Astrophys.* 641, A9 (Sept. 2020), A9. DOI: [10.1051/0004-6361/201935891](https://doi.org/10.1051/0004-6361/201935891). arXiv: [1905.05697](https://arxiv.org/abs/1905.05697) [astro-ph.CO].
- [253] M. Gatti et al. “Dark Energy Survey Year 3 results: Cosmology with moments of weak lensing mass maps”. In: *Phys. Rev. D* 106.8, 083509 (Oct. 2022), p. 083509. DOI: [10.1103/PhysRevD.106.083509](https://doi.org/10.1103/PhysRevD.106.083509). arXiv: [2110.10141](https://arxiv.org/abs/2110.10141) [astro-ph.CO].
- [254] Sven Heydenreich et al. “A roadmap to cosmological parameter analysis with third-order shear statistics. I. Modelling and validation”. In: *Astron. Astrophys.* 672, A44 (Apr. 2023), A44. DOI: [10.1051/0004-6361/202244820](https://doi.org/10.1051/0004-6361/202244820). arXiv: [2208.11686](https://arxiv.org/abs/2208.11686) [astro-ph.CO].
- [255] Gabriel Jung et al. “The integrated angular bispectrum”. In: *JCAP* 2020.6, 035 (June 2020), p. 035. DOI: [10.1088/1475-7516/2020/06/035](https://doi.org/10.1088/1475-7516/2020/06/035). arXiv: [2004.03574](https://arxiv.org/abs/2004.03574) [astro-ph.CO].
- [256] Priyesh Chakraborty, Shu-Fan Chen, and Cora Dvorkin. “Skewing the CMB×LSS: a fast method for bispectrum analysis”. In: *JCAP* 2022.7, 038 (July 2022), p. 038. DOI: [10.1088/1475-7516/2022/07/038](https://doi.org/10.1088/1475-7516/2022/07/038). arXiv: [2202.11724](https://arxiv.org/abs/2202.11724) [astro-ph.CO].
- [257] Oliver H. E. Philcox. “Optimal estimation of the binned mask-free power spectrum, bispectrum, and trispectrum on the full sky: Scalar edition”. In: *Phys. Rev. D* 107.12, 123516 (June 2023), p. 123516. DOI: [10.1103/PhysRevD.107.123516](https://doi.org/10.1103/PhysRevD.107.123516). arXiv: [2303.08828](https://arxiv.org/abs/2303.08828) [astro-ph.CO].
- [258] Gerrit S. Farren et al. “Detection of the CMB lensing – galaxy bispectrum”. In: *arXiv e-prints*, arXiv:2311.04213 (Nov. 2023), arXiv:2311.04213. DOI: [10.48550/arXiv.2311.04213](https://doi.org/10.48550/arXiv.2311.04213). arXiv: [2311.04213](https://arxiv.org/abs/2311.04213) [astro-ph.CO].
- [259] Lea Harscouet et al. “Fast Projected Bispectra: the filter-square approach”. In: *The Open Journal of Astrophysics* 8, 6 (Jan. 2025), p. 6. DOI: [10.33232/001c.128309](https://doi.org/10.33232/001c.128309). arXiv: [2409.07980](https://arxiv.org/abs/2409.07980) [astro-ph.CO].
- [260] Lea Harscouet et al. “Constraints from CMB lensing tomography with projected bispectra”. In: *arXiv e-prints*, arXiv:2507.07968 (July 2025), arXiv:2507.07968. DOI: [10.48550/arXiv.2507.07968](https://doi.org/10.48550/arXiv.2507.07968). arXiv: [2507.07968](https://arxiv.org/abs/2507.07968) [astro-ph.CO].
- [261] Haojie Xu et al. “DESI Legacy Imaging Surveys Data Release 9: Cosmological constraints from galaxy clustering and weak lensing using the minimal bias model”. In: *Science China Physics, Mechanics, and Astronomy* 66.12, 129811 (Dec. 2023), p. 129811. DOI: [10.1007/s11433-023-2242-8](https://doi.org/10.1007/s11433-023-2242-8). arXiv: [2310.03066](https://arxiv.org/abs/2310.03066) [astro-ph.CO].
- [262] Rongpu Zhou et al. “DESI luminous red galaxy samples for cross-correlations”. In: *JCAP* 2023.11, 097 (Nov. 2023), p. 097. DOI: [10.1088/1475-7516/2023/11/097](https://doi.org/10.1088/1475-7516/2023/11/097). arXiv: [2309.06443](https://arxiv.org/abs/2309.06443) [astro-ph.CO].

- [263] Benjamin D. Wandelt, Eric Hivon, and Krzysztof M. Górski. “Cosmic microwave background anisotropy power spectrum statistics for high precision cosmology”. In: *Phys. Rev. D* 64.8, 083003 (Oct. 2001), p. 083003. DOI: [10.1103/PhysRevD.64.083003](https://doi.org/10.1103/PhysRevD.64.083003). arXiv: [astro-ph/0008111](https://arxiv.org/abs/astro-ph/0008111) [astro-ph].
- [264] David Alonso et al. “A unified pseudo- C_ℓ framework”. In: *Mon. Not. Roy. Astron. Soc.* 484.3 (Apr. 2019), pp. 4127–4151. DOI: [10.1093/mnras/stz093](https://doi.org/10.1093/mnras/stz093). arXiv: [1809.09603](https://arxiv.org/abs/1809.09603) [astro-ph.CO].
- [265] Carlos García-García, David Alonso, and Emilio Bellini. “Disconnected pseudo- C_ℓ covariances for projected large-scale structure data”. In: *JCAP* 2019.11, 043 (Nov. 2019), p. 043. DOI: [10.1088/1475-7516/2019/11/043](https://doi.org/10.1088/1475-7516/2019/11/043). arXiv: [1906.11765](https://arxiv.org/abs/1906.11765) [astro-ph.CO].
- [266] Nina A. Maksimova et al. “ABACUSSUMMIT: a massive set of high-accuracy, high-resolution N-body simulations”. In: *Mon. Not. Roy. Astron. Soc.* 508.3 (Dec. 2021), pp. 4017–4037. DOI: [10.1093/mnras/stab2484](https://doi.org/10.1093/mnras/stab2484). arXiv: [2110.11398](https://arxiv.org/abs/2110.11398) [astro-ph.CO].
- [267] Lehman H. Garrison, Daniel J. Eisenstein, and Philip A. Pinto. “A high-fidelity realization of the Euclid code comparison N-body simulation with ABACUS”. In: *Mon. Not. Roy. Astron. Soc.* 485.3 (May 2019), pp. 3370–3377. DOI: [10.1093/mnras/stz634](https://doi.org/10.1093/mnras/stz634). arXiv: [1810.02916](https://arxiv.org/abs/1810.02916) [astro-ph.CO].
- [268] Lehman H. Garrison et al. “The ABACUS cosmological N-body code”. In: *Mon. Not. Roy. Astron. Soc.* 508.1 (Nov. 2021), pp. 575–596. DOI: [10.1093/mnras/stab2482](https://doi.org/10.1093/mnras/stab2482). arXiv: [2110.11392](https://arxiv.org/abs/2110.11392) [astro-ph.CO].
- [269] Boryana Hadzhiyska et al. “The halo light-cone catalogues of ABACUSSUMMIT”. In: *Mon. Not. Roy. Astron. Soc.* 509.2 (Jan. 2022), pp. 2194–2208. DOI: [10.1093/mnras/stab3066](https://doi.org/10.1093/mnras/stab3066). arXiv: [2110.11413](https://arxiv.org/abs/2110.11413) [astro-ph.CO].
- [270] Rongpu Zhou et al. “Target Selection and Validation of DESI Luminous Red Galaxies”. In: *Astron. J.* 165.2, 58 (Feb. 2023), p. 58. DOI: [10.3847/1538-3881/aca5fb](https://doi.org/10.3847/1538-3881/aca5fb). arXiv: [2208.08515](https://arxiv.org/abs/2208.08515) [astro-ph.CO].
- [271] Sihan Yuan et al. “The DESI one-per cent survey: exploring the halo occupation distribution of luminous red galaxies and quasi-stellar objects with ABACUSSUMMIT”. In: *Mon. Not. Roy. Astron. Soc.* 530.1 (May 2024), pp. 947–965. DOI: [10.1093/mnras/stae359](https://doi.org/10.1093/mnras/stae359). arXiv: [2306.06314](https://arxiv.org/abs/2306.06314) [astro-ph.CO].
- [272] D. Nelson Limber. “The Analysis of Counts of the Extragalactic Nebulae in Terms of a Fluctuating Density Field.” In: *Astrophys. J.* 117 (Jan. 1953), p. 134. DOI: [10.1086/145672](https://doi.org/10.1086/145672).
- [273] A. J. S. Hamilton. “Linear Redshift Distortions: a Review”. In: *The Evolving Universe*. Ed. by Donald Hamilton. Vol. 231. Astrophysics and Space Science Library. Jan. 1998, p. 185. DOI: [10.1007/978-94-011-4960-0_17](https://doi.org/10.1007/978-94-011-4960-0_17). arXiv: [astro-ph/9708102](https://arxiv.org/abs/astro-ph/9708102) [astro-ph].
- [274] Konstantinos Tanidis and Stefano Camera. “Developing a unified pipeline for large-scale structure data analysis with angular power spectra - I. The importance of redshift-space distortions for galaxy number counts”. In: *Mon. Not. Roy. Astron. Soc.* 489.3 (Nov. 2019), pp. 3385–3402. DOI: [10.1093/mnras/stz2366](https://doi.org/10.1093/mnras/stz2366). arXiv: [1902.07226](https://arxiv.org/abs/1902.07226) [astro-ph.CO].
- [275] Nora Elisa Chisari et al. “Core Cosmology Library: Precision Cosmological Predictions for LSST”. In: *Astrophys. J.* 242.1, 2 (May 2019), p. 2. DOI: [10.3847/1538-4365/ab1658](https://doi.org/10.3847/1538-4365/ab1658). arXiv: [1812.05995](https://arxiv.org/abs/1812.05995) [astro-ph.CO].

- [276] Patrick McDonald and Arabindo Roy. “Clustering of dark matter tracers: generalizing bias for the coming era of precision LSS”. In: *JCAP* 2009.8, 020 (Aug. 2009), p. 020. DOI: [10.1088/1475-7516/2009/08/020](https://doi.org/10.1088/1475-7516/2009/08/020). arXiv: [0902.0991](https://arxiv.org/abs/0902.0991) [astro-ph.CO].
- [277] Mikhail M. Ivanov et al. “Cosmology with the galaxy bispectrum multipoles: Optimal estimation and application to BOSS data”. In: *Phys. Rev. D* 107.8, 083515 (Apr. 2023), p. 083515. DOI: [10.1103/PhysRevD.107.083515](https://doi.org/10.1103/PhysRevD.107.083515). arXiv: [2302.04414](https://arxiv.org/abs/2302.04414) [astro-ph.CO].
- [278] Xiao Fang et al. “FAST-PT II: an algorithm to calculate convolution integrals of general tensor quantities in cosmological perturbation theory”. In: *JCAP* 2017.2, 030 (Feb. 2017), p. 030. DOI: [10.1088/1475-7516/2017/02/030](https://doi.org/10.1088/1475-7516/2017/02/030). arXiv: [1609.05978](https://arxiv.org/abs/1609.05978) [astro-ph.CO].
- [279] François Chollet et al. *Keras*. <https://keras.io>. 2015.
- [280] Martín Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015. URL: <https://www.tensorflow.org/>.
- [281] Diederik P Kingma. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014).
- [282] Giovanni Aricò, Raul E. Angulo, and Matteo Zennaro. “Accelerating Large-Scale-Structure data analyses by emulating Boltzmann solvers and Lagrangian Perturbation Theory”. In: *arXiv e-prints*, arXiv:2104.14568 (Apr. 2021), arXiv:2104.14568. DOI: [10.48550/arXiv.2104.14568](https://doi.org/10.48550/arXiv.2104.14568). arXiv: [2104.14568](https://arxiv.org/abs/2104.14568) [astro-ph.CO].
- [283] Mikhail M. Ivanov et al. “Precision analysis of the redshift-space galaxy bispectrum”. In: *Phys. Rev. D* 105.6, 063512 (Mar. 2022), p. 063512. DOI: [10.1103/PhysRevD.105.063512](https://doi.org/10.1103/PhysRevD.105.063512). arXiv: [2110.10161](https://arxiv.org/abs/2110.10161) [astro-ph.CO].
- [284] Lucía Fonseca de la Bella et al. “Impact of bias and redshift-space modelling for the halo power spectrum: testing the effective field theory of large-scale structure”. In: *JCAP* 2020.7, 011 (July 2020), p. 011. DOI: [10.1088/1475-7516/2020/07/011](https://doi.org/10.1088/1475-7516/2020/07/011). arXiv: [1805.12394](https://arxiv.org/abs/1805.12394) [astro-ph.CO].
- [285] Samuel Goldstein et al. “Perturbation theory models for LSST-era galaxy clustering: Tests with subpercent mock catalog measurements in Fourier and configuration space”. In: *Phys. Rev. D* 105.12, 123518 (June 2022), p. 123518. DOI: [10.1103/PhysRevD.105.123518](https://doi.org/10.1103/PhysRevD.105.123518). arXiv: [2111.00501](https://arxiv.org/abs/2111.00501) [astro-ph.CO].
- [286] Andrina Nicola et al. “Galaxy bias in the era of LSST: perturbative bias expansions”. In: *JCAP* 2024.2, 015 (Feb. 2024), p. 015. DOI: [10.1088/1475-7516/2024/02/015](https://doi.org/10.1088/1475-7516/2024/02/015). arXiv: [2307.03226](https://arxiv.org/abs/2307.03226) [astro-ph.CO].
- [287] Euclid Collaboration et al. “Euclid preparation. XLI. Galaxy power spectrum modelling in real space”. In: *Astron. Astrophys.* 687, A216 (July 2024), A216. DOI: [10.1051/0004-6361/202348939](https://doi.org/10.1051/0004-6361/202348939). arXiv: [2312.00679](https://arxiv.org/abs/2312.00679) [astro-ph.CO].
- [288] Anton Chudaykin, Mikhail M. Ivanov, and Oliver H. E. Philcox. “Reanalyzing DESI DR1. 1. Λ CDM constraints from the power spectrum and bispectrum”. In: *Phys. Rev. D* 113.6 (2026), p. 063502. DOI: [10.1103/qsnt-dppc](https://doi.org/10.1103/qsnt-dppc). arXiv: [2507.13433](https://arxiv.org/abs/2507.13433) [astro-ph.CO].
- [289] Titouan Lazeyras et al. “Precision measurement of the local bias of dark matter halos”. In: *JCAP* 2016.2 (Feb. 2016), pp. 018–018. DOI: [10.1088/1475-7516/2016/02/018](https://doi.org/10.1088/1475-7516/2016/02/018). arXiv: [1511.01096](https://arxiv.org/abs/1511.01096) [astro-ph.CO].

- [290] Titouan Lazeyras and Fabian Schmidt. “Beyond LIMD bias: a measurement of the complete set of third-order halo bias parameters”. In: *JCAP* 2018.9, 008 (Sept. 2018), p. 008. DOI: [10.1088/1475-7516/2018/09/008](https://doi.org/10.1088/1475-7516/2018/09/008). arXiv: 1712.07531 [astro-ph.CO].
- [291] Gerrit S. Farren et al. “Atacama Cosmology Telescope: Multiprobe cosmology with unWISE galaxies and ACT DR6 CMB lensing”. In: *Phys. Rev. D* 111.8, 083516 (Apr. 2025), p. 083516. DOI: [10.1103/PhysRevD.111.083516](https://doi.org/10.1103/PhysRevD.111.083516). arXiv: 2409.02109 [astro-ph.CO].
- [292] David Alonso et al. “Constraining cosmology with the Gaia-unWISE Quasar Catalog and CMB lensing: structure growth”. In: *JCAP* 2023.11, 043 (Nov. 2023), p. 043. DOI: [10.1088/1475-7516/2023/11/043](https://doi.org/10.1088/1475-7516/2023/11/043). arXiv: 2306.17748 [astro-ph.CO].
- [293] Frank J. Qu et al. “Atacama Cosmology Telescope DR6 and DESI: Structure growth measurements from the cross-correlation of DESI legacy imaging galaxies and CMB lensing from ACT DR6 and Planck PR4”. In: *Phys. Rev. D* 111.10, 103503 (May 2025), p. 103503. DOI: [10.1103/PhysRevD.111.103503](https://doi.org/10.1103/PhysRevD.111.103503). arXiv: 2410.10808 [astro-ph.CO].
- [294] Alexandre Barreira, Titouan Lazeyras, and Fabian Schmidt. “Galaxy bias from forward models: linear and second-order bias of IllustrisTNG galaxies”. In: *JCAP* 2021.8, 029 (Aug. 2021), p. 029. DOI: [10.1088/1475-7516/2021/08/029](https://doi.org/10.1088/1475-7516/2021/08/029). arXiv: 2105.02876 [astro-ph.CO].
- [295] Matteo Zennaro et al. “Priors on Lagrangian bias parameters from galaxy formation modelling”. In: *Mon. Not. Roy. Astron. Soc.* 514.4 (Aug. 2022), pp. 5443–5456. DOI: [10.1093/mnras/stac1673](https://doi.org/10.1093/mnras/stac1673). arXiv: 2110.05408 [astro-ph.CO].
- [296] Mikhail M. Ivanov et al. “The Millennium and Astrid galaxies in effective field theory: comparison with galaxy-halo connection models at the field level”. In: *arXiv e-prints*, arXiv:2412.01888 (Dec. 2024), arXiv:2412.01888. DOI: [10.48550/arXiv.2412.01888](https://doi.org/10.48550/arXiv.2412.01888). arXiv: 2412.01888 [astro-ph.CO].
- [297] M. Gatti et al. “Dark energy survey year 3 results: weak lensing shape catalogue”. In: *Mon. Not. Roy. Astron. Soc.* 504.3 (July 2021), pp. 4312–4336. DOI: [10.1093/mnras/stab918](https://doi.org/10.1093/mnras/stab918). arXiv: 2011.03408 [astro-ph.CO].
- [298] Benjamin Stölzner et al. “KiDS-Legacy: Consistency of cosmic shear measurements and joint cosmological constraints with external probes”. In: *arXiv e-prints*, arXiv:2503.19442 (Mar. 2025), arXiv:2503.19442. DOI: [10.48550/arXiv.2503.19442](https://doi.org/10.48550/arXiv.2503.19442). arXiv: 2503.19442 [astro-ph.CO].
- [299] Takahiko Matsubara. “Nonlinear perturbation theory with halo bias and redshift-space distortions via the Lagrangian picture”. In: *Phys. Rev. D* 78.8, 083519 (Oct. 2008), p. 083519. DOI: [10.1103/PhysRevD.78.083519](https://doi.org/10.1103/PhysRevD.78.083519). arXiv: 0807.1733 [astro-ph].
- [300] Chirag Modi, Shi-Fan Chen, and Martin White. “Simulations and symmetries”. In: *Mon. Not. Roy. Astron. Soc.* 492.4 (Mar. 2020), pp. 5754–5763. DOI: [10.1093/mnras/staa251](https://doi.org/10.1093/mnras/staa251). arXiv: 1910.07097 [astro-ph.CO].
- [301] Matteo Zennaro et al. “The BACCO simulation project: biased tracers in real space”. In: *Mon. Not. Roy. Astron. Soc.* 524.2 (Sept. 2023), pp. 2407–2419. DOI: [10.1093/mnras/stad2008](https://doi.org/10.1093/mnras/stad2008). arXiv: 2101.12187 [astro-ph.CO].

- [302] Anton Chudaykin, Mikhail M. Ivanov, and Oliver H. E. Philcox. “Reanalyzing DESI DR1. 2. Constraints on dark energy, spatial curvature, and neutrino masses”. In: *Phys. Rev. D* 113.12 (2026), p. 123506. DOI: [10.1103/42f4-grkj](https://doi.org/10.1103/42f4-grkj). arXiv: 2511.20757 [astro-ph.CO].
- [303] Anton Chudaykin, Mikhail M. Ivanov, and Oliver H. E. Philcox. “Reanalyzing DESI DR1. 3. Constraints on inflation from galaxy power spectra and bispectra”. In: *Phys. Rev. D* 113.6 (2026), p. 063552. DOI: [10.1103/fhj3-6q4x](https://doi.org/10.1103/fhj3-6q4x). arXiv: 2512.04266 [astro-ph.CO].
- [304] N. Aghanim et al. “Planck 2018 results. V. CMB power spectra and likelihoods”. In: *Astron. Astrophys.* 641 (2020), A5. DOI: [10.1051/0004-6361/201936386](https://doi.org/10.1051/0004-6361/201936386). arXiv: 1907.12875 [astro-ph.CO].
- [305] Julien Carron, Mark Mirmelstein, and Antony Lewis. “CMB lensing from Planck PR4 maps”. In: *JCAP* 2209 (2022), p. 039. DOI: [10.1088/1475-7516/2022/09/039](https://doi.org/10.1088/1475-7516/2022/09/039). arXiv: 2206.07773 [astro-ph.CO].
- [306] Niall MacCrann et al. “The Atacama Cosmology Telescope: Mitigating the Impact of Extragalactic Foregrounds for the DR6 Cosmic Microwave Background Lensing Analysis”. In: *Astrophys. J.* 966.1 (2024), p. 138. DOI: [10.3847/1538-4357/ad2610](https://doi.org/10.3847/1538-4357/ad2610). arXiv: 2304.05196 [astro-ph.CO].
- [307] Nick Hand et al. “An optimal FFT-based anisotropic power spectrum estimator”. In: *JCAP* 1707 (2017), p. 002. DOI: [10.1088/1475-7516/2017/07/002](https://doi.org/10.1088/1475-7516/2017/07/002). arXiv: 1704.02357 [astro-ph.CO].
- [308] M. Pinon et al. “Mitigation of DESI fiber assignment incompleteness effect on two-point clustering with small angular scale truncated estimators”. In: *JCAP* 2501 (2025), p. 131. DOI: [10.1088/1475-7516/2025/01/131](https://doi.org/10.1088/1475-7516/2025/01/131). arXiv: 2406.04804 [astro-ph.CO].
- [309] Alexander Mead et al. “HMcode-2020: improved modelling of non-linear cosmological power spectra with baryonic feedback”. In: *Mon. Not. Roy. Astron. Soc.* 502.1 (2021), pp. 1401–1422. DOI: [10.1093/mnras/stab082](https://doi.org/10.1093/mnras/stab082). arXiv: 2009.01858 [astro-ph.CO].
- [310] Sophie M. L. Vogt, David J. E. Marsh, and Alex Laguë. “Improved mixed dark matter halo model for ultralight axions”. In: *Phys. Rev. D* 107.6 (2023), p. 063526. DOI: [10.1103/PhysRevD.107.063526](https://doi.org/10.1103/PhysRevD.107.063526). arXiv: 2209.13445 [astro-ph.CO].
- [311] Tibor Dome et al. “Improved halo model calibrations for mixed dark matter models of ultralight axions”. In: *Mon. Not. Roy. Astron. Soc.* 537.1 (2025), pp. 252–271. DOI: [10.1093/mnras/staf005](https://doi.org/10.1093/mnras/staf005). arXiv: 2409.11469 [astro-ph.CO].
- [312] Lauren Gaughan, Anne M. Green, and Adam Moss. “Ultra-light axion constraints from Planck and ACT: the role of nonlinear modelling”. In: (). arXiv: 2605.12054 [astro-ph.CO].
- [313] Alex Laguë et al. “The Atacama Cosmology Telescope: Probing new signatures of ultralight axions with gravitational lensing”. In: (). arXiv: 2606.06410 [astro-ph.CO].
- [314] Simon Foreman and Leonardo Senatore. “The EFT of Large Scale Structures at All Redshifts: Analytical Predictions for Lensing”. In: *JCAP* 1604 (2016), p. 033. DOI: [10.1088/1475-7516/2016/04/033](https://doi.org/10.1088/1475-7516/2016/04/033). arXiv: 1503.01775 [astro-ph.CO].
- [315] Oliver H. E. Philcox, David N. Spergel, and Francisco Villaescusa-Navarro. “Effective halo model: Creating a physical and accurate model of the matter power spectrum and cluster counts”. In: *Phys. Rev. D* 101.12 (2020), p. 123520. DOI: [10.1103/PhysRevD.101.123520](https://doi.org/10.1103/PhysRevD.101.123520). arXiv: 2004.09515 [astro-ph.CO].

- [316] Shi-Fan Chen et al. “Suppression without Thawing: Constraining Structure Formation and Dark Energy with Galaxy Clustering”. In: *Phys. Rev. Lett.* 133.23 (2024), p. 231001. DOI: [10.1103/PhysRevLett.133.231001](#). arXiv: [2406.13388 \[astro-ph.CO\]](#).
- [317] Mikhail M. Ivanov et al. “Reanalyzing DESI DR1: 4. Percent-Level Cosmological Constraints from Combined Probes and Robust Evidence for the Normal Neutrino Mass Hierarchy”. In: (). arXiv: [2601.16165 \[astro-ph.CO\]](#).
- [318] Chia-Hsun Chuang et al. “EZmocks: extending the Zel’dovich approximation to generate mock galaxy catalogues with accurate clustering statistics”. In: *Mon. Not. Roy. Astron. Soc.* 446 (2015), pp. 2621–2628. DOI: [10.1093/mnras/stu2301](#). arXiv: [1409.1124 \[astro-ph.CO\]](#).
- [319] Peter W. Graham, Daniel Green, and Joel Meyers. “New interpretations of the cosmological preference for a negative neutrino mass”. In: *Phys. Rev. D* 113.4 (2026), p. 043514. DOI: [10.1103/1bqb-qlrj](#). arXiv: [2508.20999 \[astro-ph.CO\]](#).
- [320] Guido D’Amico et al. “Taming redshift-space distortion effects in the EFTofLSS and its application to data”. In: *JCAP* 2401 (2024), p. 037. DOI: [10.1088/1475-7516/2024/01/037](#). arXiv: [2110.00016 \[astro-ph.CO\]](#).
- [321] Vid Iršič et al. “First Constraints on Fuzzy Dark Matter from Lyman- α Forest Data and Hydrodynamical Simulations”. In: *Phys. Rev. Lett.* 119.3, 031302 (July 2017), p. 031302. DOI: [10.1103/PhysRevLett.119.031302](#). arXiv: [1703.04683 \[astro-ph.CO\]](#).
- [322] Matteo Nori et al. “Lyman α forest and non-linear structure characterization in Fuzzy Dark Matter cosmologies”. In: *Mon. Not. Roy. Astron. Soc.* 482.3 (2019), pp. 3227–3243. DOI: [10.1093/mnras/sty2888](#). arXiv: [1809.09619 \[astro-ph.CO\]](#).
- [323] Vid Iršič, Huangyu Xiao, and Matthew McQuinn. “Early structure formation constraints on the ultralight axion in the postinflation scenario”. In: *Phys. Rev. D* 101.12, 123518 (June 2020), p. 123518. DOI: [10.1103/PhysRevD.101.123518](#). arXiv: [1911.11150 \[astro-ph.CO\]](#).
- [324] Harrison Winch et al. “High-redshift, Small-scale Tests of Ultralight Axion Dark Matter Using Hubble and Webb Galaxy UV Luminosities”. In: *Astrophys. J.* 976.1 (2024), p. 40. DOI: [10.3847/1538-4357/ad7a73](#). arXiv: [2404.11071 \[astro-ph.CO\]](#).
- [325] Jackson Sipple et al. “Fuzzy dark matter constraints from the Hubble Frontier Fields”. In: *Mon. Not. Roy. Astron. Soc.* 538.3 (2025), pp. 1830–1842. DOI: [10.1093/mnras/staf340](#). arXiv: [2407.17059 \[astro-ph.CO\]](#).