INSIGHT ON THE DARK UNIVERSE FROM COSMOLOGICAL PERTURBATIONS

Von der Fakultät für Mathematik, Informatik und Naturwissenschaften der RWTH AachenUniversity zur Erlangung des akademischen Grades eines Doktors der Naturwissenschaftengenehmigte Dissertation

vorgelegt von

Thejs Brinckmann, M.Sc.

aus

Kopenhagen, Dänemark

Berichter*:* *Prof. Dr. Julien Lesgourgues*

Prof. Dr. Philipp Mertsch

Tag der mündlichen Prüfung*:* *20.09.2018*

SUPERVISOR:
Prof. Julien Lesgourgues

REFEREE:
Prof. Philipp Mertsch

CHAIR:
Prof. Achim Stahl

COMMITTEE:
Prof. Christopher Wiebusch

DEFENSE DATE:
20.09.2018

LOCATION:
Aachen

THESIS TIME FRAME:
September 2015 – October 2018

ABSTRACT

Cosmology presents the best hope of measuring the sum of neutrino masses in the future. The Cosmic Microwave Background (CMB) has already been a treasure trove of information; from Planck, BICEP/KECK, and their predecessors; and will continue to provide ever more precise information with upcoming or proposed CMB experiments, such as LiteBird, CMB-S4, CORE and PICO. These missions will have great synergy with other branches of cosmology. In particular, massive neutrinos leave a distinct imprint on the matter distribution of the universe, which upcoming large-scale structure experiments such as Euclid and the Square Kilometre Array will observe with unprecedented levels of precision. The uncertainty in modelling of non-linear structure formation is often neglected in other forecasts, or scales corresponding to this regime are entirely removed. In this work, we take into account that our understanding of non-linear modelling is imperfect. We show that a neutrino mass sum measurement is all but guaranteed from cosmology in the next decade and that this statement is robust to choice of cosmological model or modelling of non-linear effects.

PUBLICATIONS

1. CORE cosmological parameters: Di Valentino, Brinckmann, Gerbino, Poulin et al. 2016 [50]
   In the first paper, we studied a number of different possible instrumental settings for the CORE-M5 CMB satellite proposal to ESA for a myriad of common extended cosmological models, e.g. varying neutrino mass sum, alone and in combination with varying effective relativistic degrees of freedom, curvature, primordial helium fraction, and equation of state of dark energy. Additionally, we included future large-scale structure surveys in the analysis: BAO with DESI and galaxy clustering with Euclid, and compared to a mock Planck likelihood intended to mimic the sensitivity of the full Planck mission.

2. Archidiacono, Brinckmann, Lesgourgues, & Poulin 2016 [157]
   For the second paper, we conducted a detailed study of the physical effects involved in breaking degeneracies between the sum of neutrino masses and other cosmological parameters, explaining the great improvement in sensitivity for the neutrino mass sum when considering complementary probes. For this study, we considered the CORE-M5 CMB satellite proposal, BAO with DESI, as well as galaxy clustering and cosmic shear with Euclid.

3. Sprenger, Archidiacono, Brinckmann, Clesse, & Lesgourgues 2018 [102]
   In the third paper, we improve on the Euclid likelihoods employed in the first two papers and develop galaxy clustering, cosmic shear, and 21cm intensity mapping likelihoods for the Square Kilometre Array (SKA). We consider a few extended models relevant to the sensitivity of future surveys to the sum of neutrino masses, varying neutrino mass sum alone and in combination with varying number of relativistic degrees of freedom, equation of state of dark energy, and dynamical dark energy, for many combinations of the Euclid and SKA likelihoods with the full-mission Planck mock likelihood.

4. Vagnozzi, Brinckmann, Archidiacono, Freese, Gerbino, Lesgourgues, & Sprenger 2018 [1]
   With the fourth paper, we clarify a subtlety related to the galaxy or halo bias when considering a cosmology with massive neutrinos and definitively demonstrate that the bias must be correctly accounted for when considering future surveys such as Euclid.

5. Brinckmann, Hooper, Archidiacono, Lesgourgues, & Sprenger (in prep.)
   In the fifth paper, we compile a reference paper for the sensitivity of future large-scale structure surveys (Euclid, SKA, DESI) and CMB missions (full-mission mock Planck, Litebird, CMB ground based stage 4, CORE, and PICO), and relevant combinations thereof, to the neutrino mass sum, considering varying sum of neutrino masses alone and in combination with varying equation of state of dark energy, dynamical dark energy, and number of relativstic degrees of freedom.

6. Brinckmann & Lesgourgues 2018 [175]
   Finally, in the sixth paper, we review relevant concepts on the topic of statistical inference and outline recent developments of the cosmological sampling package `MontePython`. The results of papers 1, 3 and 5 would not have been feasible without the improvements in computational efficiency outlined in this paper.

CONTENTS

1

INTRODUCTION

Chapter 1 is based on Vagnozzi, Brinckmann, Archidiacono, Freese, Gerbino & Lesgourgues [1] and Brinckmann, Hooper, Archidiacono, Lesgourgues & Sprenger (in prep.)

During the last three decades, the increasingly accurate observations of the Cosmic Microwave Background (CMB) anisotropies have promoted Cosmology to a precision science. Combined with other probes such as Type-1a supernovae, the statistical distribution of large-scale structure (LSS), the weak gravitational lensing, and the Lyman-α forest, the Planck satellite has measured the standard cosmological parameters with an accuracy down to the percent level [2]. Until now, however, the nature of Dark Energy and Dark Matter, accounting for about 95% of the density of the Universe, still remains a deep mystery.

In the next decade, galaxy surveys like Euclid [3, 4] and the Square Kilometre Array (SKA) [5, 6] will probe the growth of large-scale structure with an unprecedented precision, up to redshifts $z \sim 3$. Compared to the CMB, which is a snapshot of the early Universe, galaxy surveys will achieve a tomography of the Universe over its last twelve billion years. In addition, the SKA will also achieve a precise map of neutral hydrogen through 21-cm intensity mapping; a tracer of the large-scale structure distribution; back to the reionization era and the cosmic dawn [7], up to redshift $z \sim 20$. Euclid and the SKA will detect billions of galaxies, over a large fraction of the sky, and will set unprecedented constraints on the various cosmological scenarios, through galaxy clustering and weak gravitational lensing. Three major challenges of Euclid and SKA are to reveal the properties of Dark Energy and Dark Matter, and to measure the cosmological neutrino mass.

Neutrinos are one of the most abundant particle species in the Universe, yet remain among the least understood. While the Standard Model of Particle Physics treats neutrinos as fundamental massless particles, neutrino oscillation experiments have shown that at least two out of the three neutrino mass eigenstates are massive. Therefore, massive neutrinos represent the only direct evidence for physics beyond the Standard Model. In fact, neutrino oscillation experiments have measured two mass-squared splittings: the solar splitting $\Delta m^2_{21} \equiv m^2_2 - m^2_1 \simeq 7.6 \times 10^{-5}\,\text{eV}^2$ and the atmospheric splitting $|\Delta m^2_{31}| \equiv |m^2_3 - m^2_1| \simeq 2.5 \times 10^{-3}\,\text{eV}^2$ [8–12], where m_1, m_2, m_3 denote the masses of the three mass eigenstates. The uncertainty in the sign of the atmospheric splitting leaves two possibilities open for the neutrino mass hierarchy: a *normal hierarchy* (NH) where $\Delta m^2_{31} > 0$ and $m_1 < m_2 < m_3$, and an *inverted hierarchy* (IH) where $\Delta m^2_{31} < 0$ and $m_3 < m_1 < m_2$.

Cosmological data currently provide the most stringent upper limits on the sum of the neutrino masses, $M_\nu \equiv m_1 + m_2 + m_3$, suggesting $M_\nu \lesssim 0.15\,\text{eV}$ at 95% confidence level (C.L.) [2, 11, 13–32], with potentially interesting implications concerning the determination of the mass hierarchy [20, 33–43]. These bounds are somewhat model-dependent and rely on assuming an underlying background standard ΛCDM cosmological model. The upper limits

on M_ν typically degrade when considering extensions to this baseline model (e.g. when the dark energy equation of state is allowed to vary), see e.g. [20, 24, 25, 31, 32, 44–49] for recent investigations.

Moving from these upper limits to a robust detection of M_ν is a key goal of upcoming cosmological surveys. Cosmological observables are known to be very sensitive to neutrino mass, with the potential to detect the neutrino mass scales well before laboratory experiments. As mentioned, this conclusion is often tempered by the fact that cosmological constraints depend on underlying assumptions on the cosmological model, since all constraints are derived from global fits to the observed data set. Fortunately, different ingredients in the cosmological model usually have distinct effects, and the cosmological data sets result in thousands of independent data points, such that in many cases parameter degeneracies can be broken. Still, constraints on parameters like the total neutrino mass can weaken significantly when more complicated cosmological models with extra free parameters are considered. Over the next decades, we expect increasingly precise data on the power spectrum of CMB anisotropies and large-scale structure. This will allow us to get not just stronger bounds on the total neutrino mass, but also more robust and model-independent detection.

Therefore, a very interesting question to address is: which combination of datasets do we need in order to detect the total neutrino mass at a given confidence level, not just assuming a minimal underlying cosmological model, but also extended models? And at which point will we be able to resolve parameter degeneracies?

This question can be addressed by performing sensitivity forecasts for parameter inference from future cosmological data. The literature already presents many such forecasts and one may wonder why we are doing new forecasts. The reason is that, these forecasts are usually based on many different methods (e.g. Fisher matrix approaches with different prescriptions or different Markov Chain Monte Carlo methods), assumptions related to future observations (list of observables included for a given experiment, assumed instrumental sensitivities, list of systematic errors taken into account and marginalized over) and assumptions related to cosmology (list of parameters being varied). To give examples, when doing forecasts for a given CMB experiment, one may choose whether to include information from CMB lensing extraction; for a given redshift survey, one may use different schemes to remove information from non-linear scales, model bias and redshift space distortions, etc.

We would like to compare the ability of different datasets to resolve a given parameter degeneracy between the total neutrino mass and another cosmological parameter, and this can be difficult since we usually have to refer to different published forecasts, where many of the assumptions differ. What is important in this thesis is not that our assumptions are the best (although we made an effort to implement as realistic assumptions as possible in our pipeline), but that they are the same across the variety of datasets and cosmological models that we consider. As such, we hope to provide a robust frame for comparing the merits of different combinations of future experiments as far as the neutrino mass measurement is concerned.

The goals of this thesis are to improve on current analysis methods, both in terms of modelling and methodology, to better understand the physical effects involved in neutrino cosmology, and finally the culmination of this work is to run a three-dimensional grid of forecasts, with homogeneous methodology and assumptions. The three axes of the grid are: underlying cosmological models, with more or fewer free parameters; CMB experiments (or combination of them when relevant); Large Scale Structure surveys or combination of them. We perform our forecasts with an MCMC exploration of the parameter space with a mock likelihood

describing future data and instrumental sensitivities. Using modern tools this method is reasonably fast and considerably more robust than Fisher matrix forecasts. Still, running a three-dimensional grid of MCMC parameter inference run is a significant computational effort, and we must carefully choose the number of cases that we want to consider for each axis of the grid. Along the way, we outline fundamental aspects of the underlying theory (chapters 2 and 3), document improvements in modelling (chapters 3 and 4) and methodology (chapter 6), and make a study of the physical effects in play in neutrino cosmology (chapter 5), before finally presenting the results of our forecasts (chapters 7 to 9).

2

COSMIC MICROWAVE BACKGROUND

Sections 2.1, 2.4 and 2.5 are based on Di Valentino, Brinckmann, Gerbino, Poulin et al. [50] and sections 2.4 and 2.5 also on Brinckmann, Hooper, Archidiacono, Lesgourgues & Sprenger (in prep.)

2.1 INTRODUCTION

In the quarter century since their first firm detection by the COBE satellite [51], Cosmic Microwave Background (CMB) anisotropies have revolutionized the field of cosmology with an enormous impact on several branches of astrophysics and particle physics. From observations made by ground-based experiments such as TOCO [52], DASI [53] and ACBAR [54], balloon-borne experiments like BOOMERanG [55, 56], MAXIMA [57] and Archeops [58], and satellite experiments such as COBE, WMAP [59, 60] and, more recently, Planck [2, 61], a cosmological "concordance" model has emerged, in which the need for new physics beyond the standard model of particles is blatantly evident. The impressive experimental progress in detector sensitivity and observational techniques, combined with the accuracy of linear perturbation theory, have clearly identified the CMB as the "sweet spot" from which to accurately constrain cosmological parameters and fundamental physics. Such a fact calls for new and significantly improved measurements of CMB anisotropies, to continue mining their scientific content.

DARK MATTER. In particular, observations of the CMB angular power spectrum are not only in impressive agreement with the expectations of the so-called ΛCDM model, based on cold dark matter (CDM hereafter), inflation and a cosmological constant, but they now also constrain several parameters with exquisite precision. For example, the cold dark matter density is now constrained to 1.25% accuracy using recent Planck measurements, naively yielding an evidence for CDM at about ~ 80 standard deviations (see [2]). Cosmology is indeed extremely powerful in identifying CDM, since on cosmological scales the gravitational effect of CDM are cleaner and can be precisely discriminated from those of standard baryonic matter. In this respect, no other cosmological observable aside from the CMB could show, if considered alone, the need for CDM to such a level of significance. Moreover, the cosmological signatures of CDM rely mainly on gravity, while astrophysical searches of DM annihilating or decaying into standard model particles depend on the strength of the interaction. Similarly, a possible signal in underground laboratory experiments depends on the coupling between CDM particles and ordinary matter (nuclei and electrons). It is possible to construct CDM models that could interact essentially just through gravity, and the current lack of detection of CDM in underground and astrophysics experiments is leaving this possibility open. If this is the case, structure formation on cosmological scales could result in the best observatory we have for

studying CDM properties, and a further improvement from future CMB measurements will clearly play a crucial and complementary role.

BARYONS. CMB measurements also provide an extremely stringent constraint on standard baryonic matter. The recent results from Planck constrain the baryonic content with a 0.7% accuracy, nearly a factor 2 better than the present constraints derived from primordial deuterium measurements [62], obtained assuming standard Big Bang Nucleosynthesis. In this respect, the experimental uncertainties on nuclear rates like $d(p,\gamma)^3$He that enter in BBN computations are starting to be relevant for accurate estimates of the baryon content from measurements of primordial nuclides. A combination of CMB and primordial deuterium measurements is starting to produce independent bounds on these quantities (see e.g. [63, 64]). As a matter of fact, a further improvement in the determination of the baryon density is mainly expected from future CMB anisotropy measurements and could help not only in testing the BBN scenario but also in providing independent constraints on nuclear physics.

LIGHT RELICS. The CMB is also a powerful probe of the density and properties of "light" particles, i.e. particles with masses below ~ 1 eV that become non-relativistic between recombination (at redshift $z \sim 1100$, when the primary CMB anisotropies are visible) and today. Such particles may affect primary and secondary CMB anisotropies, as well as structure formation. In particular, this can change the amplitude of gravitational lensing produced by the intervening matter fluctuations [65] and leave clear signatures in the CMB power spectra. Neutrinos are the most natural candidate to leave such an imprint (see e.g. [66, 67]). From neutrino oscillation experiments we indeed know that neutrinos are massive and that their total mass summed over the three eigenstates should be *larger* than $M_\nu > 60$ meV in the case of a normal hierarchy and of $M_\nu > 100$ meV in the case of an inverted hierarchy (see e.g. [68–70] for recent reviews of the current data). The most recent constraints from Planck measurements (temperature, polarization and CMB lensing) bound the total mass to $M_\nu < 140$ meV [71] at 95% c.l. Clearly, an improvement of the constraint towards a sensitivity of $\sigma(M_\nu) \sim 30$ meV will provide a *guaranteed discovery* for the neutrino absolute mass scale and for the neutrino mass hierarchy (see e.g. [33, 72–74]). Neutrinos are firmly established in the standard model of particle physics and a non-detection of the neutrino mass would cast serious doubts on the ΛCDM model, opening the window to new physics in the dark sector, such as, for instance, interactions between neutrinos and new light particles [75].

Additionally, several extensions of the standard model of particle physics feature light relic particles that could produce effects similar to massive neutrinos, and might be detected or strongly constrained by future CMB measurements. Thermal light axions (see e.g. [14, 76, 77]), for example, can produce very similar effects. Axions change the growth of structure formation after decoupling and increase the energy density in relativistic particles at early times[1], parametrized by the quantity $N_{\rm eff}$. Models of thermal axions will be difficult to accommodate with a value of $N_{\rm eff} < 3.25$, and a CMB experiment with a sensitivity of $\Delta N_{\rm eff} = 0.04$ could significantly rule out or confirm their existence. Other possible candidates are light sterile neutrinos and asymmetric dark matter (see e.g. [78–80] and [81]). More generally, a sensitivity to $\Delta N_{\rm eff} = 0.04$ could rule out the presence of any thermally-decoupled Goldstone boson that decoupled after the QCD phase transition (see e.g. [82]). The same sensitivity would also

1 The effective neutrino number $N_{\rm eff}$ is normally defined at times such that all "light" particles (neutrinos, axions, etc.) are still ultra relativistic.

probe non-standard neutrino decoupling (see e.g. [83]) and the possibility of a low reheating temperature of the order of $\mathcal{O}$(MeV) [84].

In combination with galaxy clustering and type Ia luminosity distances, CMB measurements from Planck have also provided the tightest constraints on the dark energy equation of state w [2]. In particular, the current tension between the Planck value and the HST value of the Hubble constant from Riess et al. 2016 [85] could be resolved by invoking an equation of state $w < -1$ [86]. Planck alone is currently unable to constrain the equation of state w and the Hubble constant H_0 independently, due to a "geometrical degeneracy" between the two parameters. An improved measurement of the CMB anisotropies could break this degeneracy, produce two independent constraints on w and H_0, and possibly resolve the current tension on the value of the Hubble constant.

WHAT'S NEXT? In order to further improve current measurements and provide deeper insight on neutrinos and the nature of dark matter and dark energy, a high precision CMB mission is clearly an important goal. This does, however, raise two fundamental questions. The first one is whether we really need to go to space and launch a new satellite, given that several other ground-based and balloon-borne experiments are under discussion or already under construction (see e.g. [87]). In fifteen years, it is certainly reasonable to assume that these experiments will collect excellent data that could, in principle, constrain cosmological parameters to similar precision. However, there is a fundamental aspect to consider: ground-based experiments have very limited frequency coverage and sample just a portion of the CMB sky. Contaminations from unknown foregrounds can be extremely dangerous for ground-based experiments, and can easily fool us. The claimed detection of a primordial Gravitational Waves (GW) background from the BICEP2 experiment [88] was latter ruled out by Planck observations at high frequencies, showing that contaminations from thermal dust in our Galaxy are far more severe than anticipated. This shows that unprecedented control of systematics and a wide frequency coverage are required, both of which call for a space-based mission. In fact, future ground-based and satellite experiments must be seen as complementary: while ground-based experiments could provide a first hint for primordial GWs or neutrino masses, a satellite experiment could monitor the frequency dependence of the corresponding signal with the highest possible accuracy, and unambiguously confirm its primordial nature.

Moreover, most of the future galaxy and cosmic shear surveys will sample several extended regions of the sky. Cross correlations with CMB data in the same sky area will offer a unique opportunity to test for systematics and new physics. It is, therefore, clear that a full sky survey from a satellite will offer much more complete, consistent and homogeneous information than several ground based observations of sky patches. Moreover, an accurate full-sky map of CMB polarisation on large angular scales can provide extremely strong constraints on the reionization optical depth, breaking degeneracies with other parameters such as neutrino masses.

The second fundamental question related to a new CMB satellite proposal arises from the fact that after increasing sensitivity and frequency coverage, one has to deal with the intrinsic limit of cosmic variance. At a certain point, no matter how much we increase the instrumental sensitivity, we reach the cosmic variance limit and stop improving the precision of parameter estimates. This opens the following issue: how close are we from cosmic variance with current CMB data? The Planck satellite measured the temperature angular spectrum up to the limit of cosmic variance in a wide range of angular scales; however, we are far from this limit when

we consider polarization spectra. But how much can current constraints improve with a future CMB satellite or ground based experiment?

This is one of the questions we want to address in this thesis, with a particular focus on neutrino-related parameters (mass sum and effective number of relativistic species) and parameters correlated with these.

This chapter is structured as follows: in sections 2.2 and 2.3 we review cosmological perturbation theory, with a particular focus on the CMB, and in sections 2.4 and 2.5 we describe the experimental setup and analysis method.

2.2 RELATIVISTIC PERTURBATION THEORY

In order to understand the CMB we make use of cosmological perturbation theory. This thesis is on the topic of neutrinos in cosmology and so, since neutrinos behave relativistically at early times, we need to consider cosmological perturbation theory in full general relativity.

2.2.1 *Metric perturbations*

We start from a small perturbation $\delta g_{\mu\nu}$ of the Friedmann-Robertson-Walker (FRW) metric,

$$g_{\mu\nu} = \bar{g}_{\mu\nu} + \delta g_{\mu\nu} \, . \tag{1}$$

The metric perturbations are then coupled to perturbations in the matter distribution via the Einstein equations.

In the following, we will disregard curvature and only consider a flat FRW background spacetime. As usual, latin indices indicate spatial vectors and tensors, which can be raised or lowered with δ_{ij} . In the presence of perturbations, the homogenous FRW metric, expressed in terms of conformal time τ ,

$$ds^2 = a^2(\tau) \left[d\tau^2 - \delta_{ij} dx^i dx^j \right] \, , \tag{2}$$

becomes

$$ds^2 = a^2(\tau) \left[(1+2A) d\tau^2 - 2B_i dx^i d\tau - (\delta_{ij} + h_{ij}) dx^i dx^j \right] \, , \tag{3}$$

where the perturbations h_{ij} , $A = h_{00}$ and $B_i = h_{0i} = h_{i0}$ are functions of both space and time.

2.2.2 *Metric degrees of freedom*

We would like to obtain quantities that are irreducable representations of the rotation group and therefore perform a scalar-vector-tensor (SVT) decomposition of the perturbations h_{ij} and B_i . We can decompose 3-vectors like B_i (the *shift*) into a longitudinal part and a transverse part,

$$B_i = \partial_i B + \hat{B}_i \, , \tag{4}$$

where the irrotational (curl-free) longitudinal part is the gradiant of a scalar with shorthand notation $\partial_i B = \frac{\partial B}{\partial x^i}$ and transverse part is the divergenceless (denoted by a hat) vector $\hat{B}_i$ (i.e. $\partial^i \hat{B}_i = 0$).

Analogously, the rank-2 symmetric tensor h_{ij} can be decomposed into the trace of h_{ij} and a traceless part s_{ij} (also called the *shear* or the *strain*),

$$h_{ij} = 2C\delta_{ij} + 2s_{ij} \ , \tag{5}$$

where the scalar C encodes the trace. The traceless part s_{ij} can be further decomposed into a longitudinal, a transverse and a solenoidal part.

$$h_{ij} = 2C\delta_{ij} + 2\partial_{<i}\partial_{j>}E + 2\partial_{(i}\hat{E}_{j)} + 2\hat{E}_{ij} \ , \tag{6}$$

where we have an anti-symmetrized part given by

$$\partial_{<i}\partial_{j>}E \equiv \left(\partial_i\partial_j - \frac{1}{3}\delta_{ij}\nabla^2\right)E \tag{7}$$

and a symmetrized part defined as

$$\partial_{(i}\hat{E}_{j)} \equiv \frac{1}{2}\left(\partial_i\hat{E}_j + \partial_j\hat{E}_i\right) \ . \tag{8}$$

Note that this time we have both a divergenceless vector $\partial^i\hat{E}_i = 0$ and a divergenceless tensor $\hat{E}^i_i = 0$.

This leaves us with a total of ten degrees of freedom for the metric, four scalar, four vector, and two tensor degrees of freedom:

Scalars

- A, i.e. the δg_{00} component of the perturbed metric (Eq. 3), corresponding to the *generalized gravitational potential* ψ , also called the *lapse*,
- the *potential* B from the irrotational part of δg_{0i} (first term of Eq. 4),
- C, so the trace of δg_{ij} (first term of Eq. 6), which is the *local distortion* ϕ of the average scale factor $a(t)$, also called the *spatial curvature perturbation*, acting as a kind of local perturbation of the scale factor, giving $(1-\phi)a$,
- E from the traceless longitudinal part of δg_{ij} (second term of Eq. 6, i.e. Eq. 8), which is the *potential* μ of the metric shear tensor.

The scalar modes can be seen as a generalization of the Newtonian potential and govern growth of perturbations, CMB temperature anisotropies, non-linear evolution, etc..

Vectors

- Two from $\hat{B}_i$,
- Two from $\hat{E}_i$.

Vector modes account for Gravito-Magnetic effects generated by matter distributions with vorticity. They correspond to perturbations that decay quickly with the expansion of the universe, so vector modes are not important for cosmology, but only for some astrophysical GR topics, such as black hole mergers.

Tensors

- Two from $\hat{E}_{ij}$.

Tensor modes are predicted by inflation and account for primordial gravitational waves. However, they are tiny perturbations that are not important for understanding the CMB or e.g. galaxy clustering. Although tensor modes can bring small corrections, they can, usually, be safely neglected, unless we want to specifically study primordial gravitational waves, models that predict them, or quantities that are degenerate with them in parameter space, which is not the case in this thesis. For that reason, tensor modes will only be discussed briefly.

2.2.3 *Gauge transformations*

The perturbations in the previous section are defined with respect to a particular set of spatial coordinates and time slicing. Performing a coordinate transformation can change the value of the perturbation variables and introduce fictitious perturbations. For this reason, we are interested in quantities that do not change under a coordinate transformation. We will consider a change of coordinates

$$\begin{aligned} X^\mu \to \tilde{X}^\mu &\equiv X^\mu + \xi^\mu(\tau, x) \text{ , where} \\ \xi^0 &\equiv T \text{ ,} \\ \xi^i &\equiv L^i = \partial^i L + \hat{L}^i \text{ ,} \end{aligned} \tag{9}$$

where the spatial translation L^i is separated into a divergenceless vector, $\hat{L}^i$, and an irrotational part with the scalar L, $\partial^i L$.

The spacetime line element ds^2 is invariant under coordinate transformation,

$$ds^2 = g_{\mu\nu}(X) dX^\mu dX^\nu = \tilde{g}_{\alpha\beta}(\tilde{X}) d\tilde{X}^\alpha d\tilde{X}^\beta \text{ ,} \tag{10}$$

which we can use to determine how the metric transforms under this coordinate change. Rewriting this expression, we get an expression relating the new metric $\tilde{g}_{\alpha\beta}(\tilde{X})$ to the old one $g_{\mu\nu}(X)$,

$$g_{\mu\nu}(X) = \frac{\partial \tilde{X}^\alpha}{\partial X^\mu} \frac{\partial \tilde{X}^\beta}{\partial X^\nu} \tilde{g}_{\alpha\beta}(\tilde{X}) \text{ .} \tag{11}$$

We want to see what this metric transformation means for the perturbation quantities in the perturbed metric (Eq. 3). To do this, we first consider the metric transformation in terms of temporal and spatial indices and neglect higher order terms.

Starting with $\mu = \nu = 0$ we get

$$g_{00}(X) = \frac{\partial \tilde{X}^\alpha}{\partial \tau} \frac{\partial \tilde{X}^\beta}{\partial \tau} \tilde{g}_{\alpha\beta}(\tilde{X}) \text{ .} \tag{12}$$

We will neglect terms that are higher than first order in perturbation, leaving only $\alpha = \beta = 0$:

$$\alpha = i, \beta = j \Rightarrow g_{00}(X) = \frac{\partial \tilde{X}^i}{\partial \tau} \frac{\partial \tilde{X}^j}{\partial \tau} \tilde{g}_{ij}(\tilde{X}) \text{ ,} \tag{13}$$

$$\alpha = i, \beta = 0 \Rightarrow g_{00}(X) = \frac{\partial \tilde{X}^i}{\partial \tau} \frac{\partial \tilde{\tau}}{\partial \tau} \tilde{g}_{i0}(\tilde{X}) \text{ ,} \tag{14}$$

$$\alpha = 0, \beta = 0 \Rightarrow g_{00}(X) = \left(\frac{\partial \tilde{\tau}}{\partial \tau} \right)^2 \tilde{g}_{00}(\tilde{X}) \text{ .} \tag{15}$$

The partial derivatives $\frac{\partial \tilde{X}^{i/j}}{\partial \tau}$ are each proportional to ξ^i and $\tilde{g}_{i0}(\tilde{X})$ is proportional to $\tilde{B}_i$, both of which are first order in perturbation, making Eqs. 13 and 14 higher order in perturbation and we are left with Eq. 15.

Next, let us consider $\mu = 0$ and $\nu = i$, so we have

$$g_{0i}(X) = \frac{\partial \tilde{X}^\alpha}{\partial \tau} \frac{\partial \tilde{X}^\beta}{\partial X^i} \tilde{g}_{\alpha\beta}(\tilde{X}) \, . \tag{16}$$

Neglecting the higher order term ($\alpha = j$, $\beta = 0$), where we have $\frac{\partial \tilde{\tau}}{\partial X^i} = \partial_i T$ times $\tilde{g}_{j0} \propto \tilde{B}_j$, which is second order in perturbation, we get contributions from the terms corresponding to ($\alpha = 0$, $\beta = 0$), ($\alpha = 0$, $\beta = j$) and ($\alpha = l$, $\beta = j$)

$$g_{0i}(X) = \frac{\partial \tilde{\tau}}{\partial \tau} \frac{\partial \tilde{\tau}}{\partial X^i} \tilde{g}_{00}(\tilde{X}) + \frac{\partial \tilde{\tau}}{\partial \tau} \frac{\partial \tilde{X}^j}{\partial X^i} \tilde{g}_{0j}(\tilde{X}) + \frac{\partial \tilde{X}^l}{\partial \tau} \frac{\partial \tilde{X}^j}{\partial X^i} \tilde{g}_{lj}(\tilde{X}) \, . \tag{17}$$

For $\mu = i$ and $\nu = j$ we have

$$g_{ij}(X) = \frac{\partial \tilde{X}^\alpha}{\partial X^i} \frac{\partial \tilde{X}^\beta}{\partial X^j} \tilde{g}_{\alpha\beta}(\tilde{X}) \, . \tag{18}$$

Neglecting higher order terms we are left with only ($\alpha = k$, $\beta = l$), because for ($\alpha = 0$, $\beta = 0$) we have $\frac{\partial \tilde{\tau}}{\partial X^i} \frac{\partial \tilde{\tau}}{\partial X^j} = \partial_i T \partial_j T$ and for ($\alpha = 0$, $\beta = l$) we get $\tilde{g}_{0l} \propto \tilde{B}_l$ times $\partial_i T$, both of which are second order in perturbation. This means we have

$$g_{ij}(X) = \frac{\partial \tilde{X}^k}{\partial X^i} \frac{\partial \tilde{X}^l}{\partial X^j} \tilde{g}_{kl}(\tilde{X}) \, . \tag{19}$$

To summarize, the metric transforms as

$$g_{00}(X) = \left(\frac{\partial \tilde{\tau}}{\partial \tau} \right)^2 \tilde{g}_{00}(\tilde{X}) \, , \tag{20}$$

$$g_{0i}(X) = \frac{\partial \tilde{\tau}}{\partial \tau} \frac{\partial \tilde{\tau}}{\partial X^i} \tilde{g}_{00}(\tilde{X}) + \frac{\partial \tilde{\tau}}{\partial \tau} \frac{\partial \tilde{X}^j}{\partial X^i} \tilde{g}_{0j}(\tilde{X}) + \frac{\partial \tilde{X}^l}{\partial \tau} \frac{\partial \tilde{X}^j}{\partial X^i} \tilde{g}_{lj}(\tilde{X}) \, , \tag{21}$$

$$g_{ij}(X) = \frac{\partial \tilde{X}^k}{\partial X^i} \frac{\partial \tilde{X}^l}{\partial X^j} \tilde{g}_{kl}(\tilde{X}) \, . \tag{22}$$

Let us see what this means for the transformation of our perturbation quantities.

We start with A. From Eq. 9 we have

$$\begin{aligned} \tilde{X}^0 &= X^0 + \xi^0 \, , \text{ i.e.} \\ \tilde{\tau} &= \tau + T \, , \end{aligned} \tag{23}$$

giving

$$\left(\frac{\partial \tilde{\tau}}{\partial \tau} \right)^2 = (1 + T')^2 \, , \tag{24}$$

where the prime indicates a derivative with respect to conformal time. Remembering that $g_{00} = a^2(\tau)(1 + 2A)$ (from Eq. 3) and using Eqs. 20, 23 and 24 we get

$$g_{00} = a^2(\tau)(1 + 2A) = (1 + T')^2 a^2(\tau + T)(1 + 2\tilde{A}) \, . \tag{25}$$

Taylor expanding the first and second factor on the RHS

$$a^2(\tau + T) = (a(\tau) + a(\tau)'T + \ldots)^2 \,, \tag{26}$$
$$(1 + T')^2 = (1 + 2T' + \ldots) \,, \tag{27}$$

and neglecting terms that are higher order in perturbation we are left with

$$a^2(\tau)(1 + 2A) = a^2(\tau) + 2a(\tau)a(\tau)'T + 2a^2(\tau)T' + 2a^2(\tau)\tilde{A} \,. \tag{28}$$

Defining the Hubble parameter in conformal time $\mathcal{H} = \frac{a'}{a}$ we can rewrite this to

$$A = \mathcal{H}T + T' + \tilde{A} \,. \tag{29}$$

We now have the transformation for A

$$A \to \tilde{A} = A - T' - \mathcal{H}T \,. \tag{30}$$

For the transformation for B_i, we remember from Eq. 9 that

$$\tilde{X}^i = X^i + \xi^i = X^i + L^i \,. \tag{31}$$

Rewriting Eq. 21 using Eqs. 23, 24 and 31 we get

$$g_{0i}(X) = \frac{\partial(\tau + T)}{\partial\tau}\frac{\partial(\tau + T)}{\partial X^i}\tilde{g}_{00}(\tilde{X}) + \frac{\partial(\tau + T)}{\partial\tau}\frac{\partial(X^j + L^j)}{\partial X^i}\tilde{g}_{0j}(\tilde{X}) + \frac{\partial(X^l + L^l)}{\partial\tau}\frac{\partial(X^j + L^j)}{\partial X^i}\tilde{g}_{lj}(\tilde{X}) \,, \tag{32}$$
$$= (1 + T')(\partial_i\tau + \partial_i T)\tilde{g}_{00} + (1 + T')(\delta_i{}^j + \partial_i L^j)\tilde{g}_{0j} + (\partial_\tau X^l + \partial_\tau L^l)(\delta_i{}^j + \partial_i L^j)\tilde{g}_{lj} \,, \tag{33}$$

where we introduced the shorthand notation $\partial_\tau = \frac{\partial}{\partial\tau}$. We have $\partial_i\tau = \partial_\tau X^l = 0$, so the expression reduces to

$$g_{0i}(X) = (1 + T')(\partial_i T)\tilde{g}_{00} + (1 + T')(\delta_i{}^j + \partial_i L^j)\tilde{g}_{0j} + (\partial_\tau L^l)(\delta_i{}^j + \partial_i L^j)\tilde{g}_{lj} \,. \tag{34}$$

Neglecting terms that are higher order in perturbation we are left with

$$g_{0i}(X) = \partial_i T\tilde{g}_{00} + \delta_i{}^j\tilde{g}_{0j} + \partial_\tau L^l\delta_i{}^j\tilde{g}_{lj} \,. \tag{35}$$

We have from Eq. 3

$$\tilde{g}_{00} = a^2(\tilde{\tau})(1 + 2\tilde{A}) \,, \tag{36}$$
$$\tilde{g}_{0j} = -a^2(\tilde{\tau})\tilde{B}_j \,, \tag{37}$$
$$\tilde{g}_{lj} = -a^2(\tilde{\tau})(\delta_{lj} + \tilde{h}_{lj}) \,. \tag{38}$$

Inserting this into Eq. 35 we get

$$g_{0i}(X) = \partial_i T a^2(\tilde{\tau})(1 + 2\tilde{A}) - \delta_i{}^j a^2(\tilde{\tau})\tilde{B}_j - \partial_\tau L^l\delta_i{}^j a^2(\tilde{\tau})(\delta_{lj} + \tilde{h}_{lj}) \,. \tag{39}$$

Discarding higher order terms we have

$$g_{0i}(X) = \partial_i T a^2(\tilde{\tau}) - \delta_i{}^j a^2(\tilde{\tau})\tilde{B}_j - \partial_\tau L^l\delta_i{}^j a^2(\tilde{\tau})\delta_{lj} \,. \tag{40}$$

We can rewrite this to

$$g_{0i}(X) = -a^2(\tau)B_i = a^2(\tilde{\tau})\partial_i T - a^2(\tilde{\tau})\tilde{B}_i - a^2(\tilde{\tau})\partial_\tau L_i \,. \tag{41}$$

Again Taylor expanding $a^2(\tilde{\tau}) = a^2(\tau + T) = (a(\tau) + a(\tau)'T + \ldots)^2$ we see the terms with $a(\tau)'T$ would be second order in perturbation, so we are left with

$$-a^2(\tau)B_i = a^2(\tau)\partial_i T - a^2(\tau)\tilde{B}_i - a^2(\tau)\partial_\tau L_i \,. \tag{42}$$

So we find the transformation for B_i is

$$B_i \to \tilde{B}_i = B_i + \partial_i T - \partial_\tau L_i \,. \tag{43}$$

However, we are more interested in the transformation of the perturbation quantities B and $\hat{B}_i$. Recalling from Eqs. 4 and 9 that $B_i = \partial_i B + \hat{B}_i$ and $L_i = \partial_i L + \hat{L}_i$, we can rewrite Eq. 43 to

$$\begin{aligned} \tilde{B}_i &= B_i + \partial_i T - \partial_\tau(\partial_i L + \hat{L}_i) \\ &= \partial_i B + \hat{B}_i + \partial_i T - \partial_i L' - \hat{L}_i' \,, \end{aligned} \tag{44}$$

which we can separate into scalars and vectors

$$\tilde{B}_i = \partial_i \tilde{B} + \hat{\tilde{B}}_i = \underbrace{\partial_i(B + T - L')}_{\text{scalars}} + \underbrace{(\hat{B}_i - \hat{L}_i')}_{\text{vectors}} \,, \tag{45}$$

allowing us to determine the transformations for B and $\hat{B}_i$,

$$B \to \tilde{B} = B + T - L' \,, \tag{46}$$

$$\hat{B}_i \to \hat{\tilde{B}}_i = \hat{B}_i - \hat{L}_i' \,. \tag{47}$$

For h_{ij} we start from Eq. 22

$$\begin{aligned} g_{ij}(X) &= \frac{\partial \tilde{X}^k}{\partial X^i}\frac{\partial \tilde{X}^l}{\partial X^j}\tilde{g}_{kl}(\tilde{X}) \\ &= \frac{\partial(X^k + \xi^k)}{\partial X^i}\frac{\partial(X^l + \xi^l)}{\partial X^j}\tilde{g}_{kl}(\tilde{X}) \\ &= (\delta_i{}^k + \partial_i L^k)(\delta_j{}^l + \partial_j L^l)\tilde{g}_{kl}(\tilde{X}) \,. \end{aligned} \tag{48}$$

Inserting Eq. 38 we get

$$g_{ij}(X) = -(\delta_i{}^k + \partial_i L^k)(\delta_j{}^l + \partial_j L^l)a^2(\tilde{\tau})(\delta_{kl} + \tilde{h}_{kl}) \,. \tag{49}$$

Neglecting higher order perturbation terms we have

$$\begin{aligned} g_{ij}(X) &= -a^2(\tilde{\tau})(\delta_i{}^k\delta_j{}^l(\delta_{kl} + \tilde{h}_{kl}) + \delta_i{}^k\delta_{kl}\partial_j L^l + \delta_j{}^l\delta_{kl}\partial_i L^k) \\ &= -a^2(\tilde{\tau})(\delta_{ij} + \tilde{h}_{ij} + \partial_j L_i + \partial_i L_j) \,. \end{aligned} \tag{50}$$

Expanding $a^2(\tilde{\tau})$ again and disregarding higher order terms we are left with

$$g_{ij}(X) = -a^2(\tau)(\delta_{ij} + h_{ij}) = -a^2(\tau)(\delta_{ij} + \tilde{h}_{ij} + \partial_j L_i + \partial_i L_j) - 2a(\tau)a(\tau)'T\delta_{ij} \,. \tag{51}$$

From which we find

$$h_{ij} = \tilde{h}_{ij} + \partial_j L_i + \partial_i L_j + 2\mathcal{H}T\delta_{ij} \,, \tag{52}$$

giving the transformation for $\tilde{h}_{ij}$

$$h_{ij} \to \tilde{h}_{ij} = h_{ij} - \partial_j L_i - \partial_i L_j - 2\mathcal{H}T\delta_{ij} \,. \tag{53}$$

Again, we are more interested in the transformations for the perturbation quantities C, E, $\hat{E}_i$ and $\hat{E}_{ij}$. From Eq. 6, we have $h_{ij} = 2C\delta_{ij} + 2\partial_{<i}\partial_{j>}E + 2\partial_{(i}\hat{E}_{j)} + 2\hat{E}_{ij}$, which we can insert in the transformation from Eq. 53 to obtain

$$\tilde{h}_{ij} = 2C\delta_{ij} + 2\partial_{<i}\partial_{j>}E + 2\partial_{(i}\hat{E}_{j)} + 2\hat{E}_{ij} - \partial_j L_i - \partial_i L_j - 2\mathcal{H}T\delta_{ij} \,. \tag{54}$$

Recalling from Eq. 9 that $L_i = \partial_i L + \hat{L}_i$ and reordering terms we have

$$\tilde{h}_{ij} = 2\delta_{ij}(C - \mathcal{H}T) + 2\partial_{<i}\partial_{j>}E + 2\partial_{(i}\hat{E}_{j)} - \partial_j(\partial_i L + \hat{L}_i) - \partial_i(\partial_j L + \hat{L}_j) + 2\hat{E}_{ij} \,, \tag{55}$$

which we can rewrite to

$$\tilde{h}_{ij} = 2\delta_{ij}(C - \mathcal{H}T) + 2\partial_{<i}\partial_{j>}E + 2\partial_{(i}\hat{E}_{j)} - 2\partial_i\partial_j L - 2\partial_{(i}\hat{L}_{j)} + 2\hat{E}_{ij} \,. \tag{56}$$

Making use of the anti-symmetrized expression $\partial_{<i}\partial_{j>}L \equiv \left(\partial_i\partial_j - \frac{1}{3}\delta_{ij}\nabla^2\right)L$ we obtain

$$\tilde{h}_{ij} = 2\delta_{ij}(C - \mathcal{H}T) + 2\partial_{<i}\partial_{j>}E + 2\partial_{(i}\hat{E}_{j)} - 2(\partial_{<i}\partial_{j>} + \frac{1}{3}\delta_{ij}\nabla^2)L - 2\partial_{(i}\hat{L}_{j)} + 2\hat{E}_{ij} \tag{57}$$

$$= 2\delta_{ij}(C - \mathcal{H}T - \frac{1}{3}\nabla^2 L) + 2\partial_{<i}\partial_{j>}(E - L) + 2\partial_{(i}(\hat{E} - \hat{L})_{j)} + 2\hat{E}_{ij} \,. \tag{58}$$

Comparison to Eq. 6 yields the transformations of the SVT perturbation quantities C, E, $\hat{E}_i$ and $\hat{E}_{ij}$. *Summarizing all the transformations*, we have

$$\begin{aligned}
&A \to \tilde{A} = A - T' - \mathcal{H}T \,, \\
&B \to \tilde{B} = Bf + T - L' \,, \\
&B_i \to \tilde{B}_i = B_i + \partial_i T - \partial_\tau L_i \,, \\
&\hat{B}_i \to \hat{\tilde{B}}_i = \hat{B}_i - \hat{L}_i \,, \\
&C \to \tilde{C} = C - \mathcal{H}T - \frac{1}{3}\nabla^2 L \,, \\
&E \to \tilde{E} = E - L \,, \\
&\hat{E}_i \to \hat{\tilde{E}}_i = \hat{E}_i - \hat{L}_i \,, \\
&\hat{E}_{ij} \to \hat{\tilde{E}}_{ij} = \hat{E}_{ij} \,, \\
&h_{ij} \to \tilde{h}_{ij} = h_{ij} - \partial_j L_i - \partial_i L_j - 2\mathcal{H}T\delta_{ij} \,.
\end{aligned} \tag{59}$$

This illustrates the gauge problem: if we perform a coordinate change, we may change the perturbation quantities. There are two solutions to this problem: define gauge-invariant perturbation variables or fix the gauge.

2.2.4 *Gauge-invariant metric perturbations*

An example of gauge-invariant perturbation quantities are the so-called Bardeen variables [89], which are combinations of metric perturbations that are invariant under coordinate transformation,

$$\Psi \equiv A + \mathcal{H}(B - E') + (B - E')' \,, \tag{60}$$

$$\Phi \equiv -C - \mathcal{H}(B - E') + \frac{1}{3}\nabla^2 E \,, \tag{61}$$

$$\hat{\Psi}_i \equiv \hat{E}_i' - \hat{B}_i \,, \tag{62}$$

$$\hat{E}_{ij} \,. \tag{63}$$

From Eq. 59 we can show that these variables indeed do not transform under a coordinate change. Beginning with Ψ, we have

$$\begin{aligned}\tilde{\Psi} &\equiv \tilde{A} + \mathcal{H}(\tilde{B} - \tilde{E}') + (\tilde{B} - \tilde{E}')' \\ &= \left[A - T' - \mathcal{H}T\right] + \mathcal{H}(\left[B + T - L'\right] - \left[E - L\right]') + (\left[B + T - L'\right] - \left[E - L\right]')' \\ &= A + \mathcal{H}(B - E') + (B - E')' \,,\end{aligned} \tag{64}$$

which indeed is invariant under coordinate transformation. Next we will consider Φ. From Eqs. 59 and 60 we find Φ does not change under a gauge transformation,

$$\begin{aligned}\tilde{\Phi} &\equiv -\tilde{C} - \mathcal{H}(\tilde{B} - \tilde{E}') + \frac{1}{3}\nabla^2\tilde{E} \\ &= -\left[C - \mathcal{H}T - \frac{1}{3}\nabla^2 L\right] - \mathcal{H}\left(\left[B + T - L'\right] - \left[E - L\right]'\right) + \frac{1}{3}\nabla^2\left[E - L\right] \\ &= -C - \mathcal{H}\left(B - E'\right) + \frac{1}{3}\nabla^2 E \,.\end{aligned} \tag{65}$$

Finally, from Eqs. 47, 59 and 60 we can show that $\hat{\Psi}_i$ is also invariant under gauge transformation,

$$\begin{aligned}\hat{\tilde{\Psi}}_i &\equiv \hat{\tilde{E}}_i' - \hat{\tilde{B}}_i \\ &= \left[\hat{E}_i - \hat{L}_i\right]' - \left[\hat{B}_i - \hat{L}_i'\right] \\ &= \hat{E}_i' - \hat{B}_i \,.\end{aligned} \tag{66}$$

This shows that we have a set of variables that are invariant under gauge transformation and could therefore be considered the true perturbation quantitites.

2.2.5 *Stress-energy perturbations*

We describe the stress-energy tensor with the density ρ, pressure p, 4-velocity u^μ, and anisotropic stress $\sigma^{\mu\nu}$. The perturbed quantities are defined through

$$\rho(\tau, x_i) \equiv \bar{\rho}(\tau) + \delta\rho(\tau, x^i) \,, \tag{67}$$

$$p(\tau, x^i) \equiv \bar{p}(\tau) + \delta p(\tau, x^i) \,, \tag{68}$$

$$u_\mu \equiv (-1 - \psi, a v_i) \,, \tag{69}$$

$$u^\mu \equiv (1 - \psi, a^{-1}(v^i - B^i)) \,, \tag{70}$$

where we recall the metric perturbation $A = \psi$, that v_i and v^i are the 3-velocities, and that the anisotropic stress is already first order in perturbation. The trace of the anisotropic stress can be absorbed into the isotropic pressure, so we can choose the anisotropic stress to be traceless, i.e. $\sigma^i_i = 0$. We also know that $\sigma^{\mu\nu}u_\nu = \sigma^\mu_\mu = 0$, which in turn implies that $\sigma^{00} = \sigma^{0j} = 0$, so only σ^{ij} is non-zero, making the anisotropic stress a traceless, symmetric 3-tensor.

The perturbed stress-energy tensor is given by

$$T^0_0 = -(\bar{\rho} + \delta\rho) \,, \tag{71}$$
$$T^0_i = (\bar{\rho} + \bar{p})av_i \,, \tag{72}$$
$$T^i_0 = -\frac{1}{a}(\bar{\rho} + \bar{p})(v^i - B^i) \,, \tag{73}$$
$$T^i_j = \delta^i_j(\bar{p} + \delta p) + \sigma^i_j \,. \tag{74}$$

The universe consists of many different species of particles, but the contribution from each species I to the perturbed stress-energy tensor, and therefore also the density and pressure perturbations, as well as the velocity anisotropy, is simply given by the sum over all species,

$$T_{\mu\nu} = \sum_I T^I_{\mu\nu} \,, \tag{75}$$
$$\delta\rho = \sum_I \delta\rho_I \,, \tag{76}$$
$$\delta p = \sum_I \delta p_I \,, \tag{77}$$
$$\sigma^{ij} = \sum_I \sigma^{ij}_I \,. \tag{78}$$

The 3-velocity is the exception, so we define the 3-momentum density $\delta q^i \equiv (\bar{\rho} + \bar{p})av^i$, which will be useful later,

$$(\bar{\rho} + \bar{p})v^i = \sum_I (\bar{\rho}_I + \bar{p}_I)v^i_I \,, \tag{79}$$
$$\delta q^i = \sum_I \delta q^i_I \,. \tag{80}$$

Just like the metric perturbations, the matter perturbations are also gauge-dependent, e.g. consider the temporal gauge transformations of the density and pressure perturbations,

$$\delta\rho \to \delta\rho - \dot{\bar{\rho}}\alpha \,, \tag{81}$$
$$\delta p \to \delta p - \dot{\bar{p}}\alpha \,. \tag{82}$$

We want to define gauge-invariant quantities. First, we define the adiabatic pressure perturbation

$$\delta p_{\rm ad} \equiv \frac{\dot{\bar{p}}}{\dot{\bar{\rho}}}\delta\rho \tag{83}$$

and split the pressure perturbation into an adiabatic and a non-adiabatic pressure perturbation $\delta p = \delta p_{\rm non-ad} + \delta p_{\rm ad}$, where the non-adiabatic pressure perturbation is gauge-invariant

$$\delta p_{\rm non-ad} = \delta p - \delta p_{\rm ad} \,. \tag{84}$$

The longitudinal part of the 3-momentum density perturbation $\partial_i \delta q$ transforms as

$$\delta q \to \delta q + (\bar{\rho} + \bar{p})\alpha \ , \tag{85}$$

but we want a gauge-invariant quantity, so we define the comoving density perturbation

$$\delta\rho_m = \delta_\rho - 3H\delta q \ , \tag{86}$$

where H is the Hubble parameter. From the metric and matter perturbations we define the comoving curvature perturbation

$$\mathcal{R} \equiv \phi - \frac{H}{\bar{\rho} + \bar{p}}\delta q \ , \tag{87}$$

recalling that $C = \phi$, and the curvature perturbation on uniform density hypersurfaces

$$\zeta \equiv -\phi - \frac{H}{\dot{\bar{\rho}}}\delta\rho \ . \tag{88}$$

2.2.6 *Perturbed Einstein equations*

We want to relate the metric perturbations to the stress-energy via the *perturbed Einstein equations*

$$\delta G_{\mu\nu} = 8\pi G \delta T_{\mu\nu} \ . \tag{89}$$

As usual we are only interested in linear order in perturbation. From the Einstein equations we get two *evolution equations*

$$\ddot{\phi} + 3H\dot{\phi} + H\dot{\psi} + \left(3H^2 + 2\dot{H}\right)\psi = 4\pi G\left(\delta p - \frac{2}{3}k^2\delta\sigma\right) \ , \tag{90}$$

$$(\partial_t + 3H)\left(\dot{E} - \frac{B}{a}\right) + \frac{\phi - \psi}{a^2} = 8\pi G\delta\sigma \ , \tag{91}$$

the second of which we can write in terms of the Bardeen variables Φ and Ψ

$$\Phi - \Psi = 8\pi G a^2 \delta\sigma \ . \tag{92}$$

We see that the Bardeen variables are equal if there is no anisotropic stress (and, by extension, in the Newtonian gauge the metric perturbations ϕ and ψ are equal). From the Einstein equations we also get the *energy and momentum constraint equations*

$$3H\left(\dot{\phi} + H\psi\right) + \frac{k^2}{a^2}\left[\phi + H\left(a^2\dot{E} - aB\right)\right] = -4\pi G\delta\rho \ , \tag{93}$$

$$\dot{\phi} + H\psi = -4\pi G\delta q \ , \tag{94}$$

from which we construct the *Poisson equation*,

$$\frac{k^2}{a^2}\Phi = -4\pi G\delta\rho_m \ , \tag{95}$$

also expressed in terms of one of the Bardeen variables, and, in fact, the Poisson equation is gauge-invariant. The last two equations we wish to obtain from the Einstein equations is the *continuity equation* and the *Euler equation*, which we get from energy-momentum conservation, i.e. $\nabla_\mu T^{\mu\nu} = 0$,

$$\dot{\delta\rho} + 3H(\delta\rho + \delta p) = \frac{k^2}{a^2}\delta q + (\bar{\rho} + \bar{p})\left[3\dot{\phi} + k^2\left(\dot{E} + \frac{B}{a}\right)\right] \ , \tag{96}$$

$$\dot{\delta q} + 3H\delta q = -\delta p + \frac{2}{3}k^2\delta\sigma - (\bar{\rho} + \bar{p})\psi \ . \tag{97}$$

2.2.7 *Gauge fixing*

Instead of defining gauge-invariant quantities, we can choose to fix the gauge. We have ten metric and ten matter degrees of freedom and can fix two of them and keep track of the remaining metric and matter perturbations.

There are many possibilities for how to fix the gauge. Popular choices include the synchronous gauge $\phi = B = 0$ (which is, in fact, not a gauge at all, as it is not uniquely defined. However, synchronous coordinates are efficient for numerical computations), the uniform density gauge $\delta\rho = E = 0$ with $\phi = -\zeta$ (useful for super-horizon scales), the comoving gauge $\delta q = E = 0$ with $\phi = -\mathcal{R}$ and the spatially-flat gauge $\phi = E = 0$ (either of which can be useful for studying inflation). However, in the following, we will only discuss the Newtonian gauge, which is what we will later use.

NEWTONIAN GAUGE. This gauge gets its name from the fact that it reduces to Newtonian gravity in the small-scale limit. A convenient advantage of this gauge is that the Bardeen variables Φ and Ψ reduce to the metric perturbations ϕ and ψ, making it an efficient choice for obtaining gauge-invariant perturbations. In the Newtonian gauge we set two metric degrees of freedom to zero $B = E = 0$, so we have the metric

$$ds^2 = a^2(\tau)\left[(1+2\psi)d\tau^2 - (1-2\phi)\delta_{ij}dx^i dx^j\right] . \tag{98}$$

The Einstein equations become

$$3H(\dot{\phi} + H\psi) + \frac{k^2}{a^2}\phi = -4\pi G\delta\rho , \tag{99}$$

$$\dot{\phi} + H\psi = -4\pi G\delta q , \tag{100}$$

$$\ddot{\phi} + 3H\dot{\phi} + H\dot{\psi}(3H^2 + 2\dot{H})\psi = 4\pi G\left(\delta\rho - \frac{2}{3}k^2\delta\sigma\right) , \tag{101}$$

$$\frac{\phi - \psi}{a^2} = 8\pi G\delta\sigma , \tag{102}$$

and the continuity and Euler equations are

$$\dot{\delta\rho} + 3H(\delta\rho + \delta p) = \frac{k^2}{a^2}\delta q + 3(\bar{\rho} + \bar{p})\dot{\phi} , \tag{103}$$

$$\dot{\delta q} + 3H\delta q = -\delta p + \frac{2}{3}k^2\delta\sigma - (\bar{\rho} + \bar{p})\psi . \tag{104}$$

2.2.8 *Power spectra*

Cosmological perturbation theory is a stochastic theory, which means we can describe perturbations in terms of probability distributions. The probability distribution for some perturbation $A(t, \vec{x})$ is given by the dimensionless power spectrum $\mathcal{P}_A(A, t, \vec{x})$. We can study the evolution of the fluctuation A through the usual statistical quantities, such as the mean or variance, which can then be interpreted as e.g. the mean over different realizations of the stochastic theory.

As a concrete example, let us consider curvature perturbations. If we define the normalization of the Fourier transform as

$$\mathcal{R}_{\vec{k}} = \int d^3x \; \mathcal{R}(\vec{x}) \; \exp\left[-i\vec{k}\cdot\vec{x}\right] \; , \tag{105}$$

$$\mathcal{R}(\vec{x}) = \frac{1}{2\pi^3}\int d^3k \; \mathcal{R}_{\vec{k}} \; \exp\left[i\vec{k}\cdot\vec{x}\right] \; , \tag{106}$$

and the two-point correlation function assuming isotropy as

$$\xi_{\mathcal{R}}(r) \equiv \langle \mathcal{R}(\vec{x})\mathcal{R}(\vec{x}+\vec{r})\rangle \; , \tag{107}$$

where $r = |\vec{r}|$, then we can define the Fourier power spectrum

$$P_{\mathcal{R}}(k) \equiv \int d^3r \; \xi_{\mathcal{R}}(r) \; \exp\left[-i\vec{k}\cdot\vec{r}\right] \tag{108}$$

as the Fourier transform of the two-point correlation function. This quantity is related to the ensemble average through

$$\left\langle \mathcal{R}_{\vec{k}}\mathcal{R}_{\vec{k}'}\right\rangle = (2\pi)^3\delta(\vec{k}+\vec{k}')\int d^3r \; \xi_{\mathcal{R}}(r) \; \exp\left[-i\vec{k}\cdot\vec{r}\right] \tag{109}$$

$$= (2\pi)^3\delta(\vec{k}+\vec{k}')P_{\mathcal{R}}(k) \; , \tag{110}$$

where the δ is the Dirac δ, and to the variance through

$$\sigma^2_{\mathcal{R}} \equiv \left\langle \mathcal{R}^2(x)\right\rangle = \frac{1}{(2\pi)^3}\int d^3k \; P_{\mathcal{R}}(k) \tag{111}$$

$$= \int d\ln k \; \mathcal{P}_{\mathcal{R}}(k) \; , \tag{112}$$

where in the second equality we defined the dimensionless power spectrum

$$\mathcal{P}_{\mathcal{R}}(k) \equiv \frac{k^3}{2\pi^2}P_{\mathcal{R}}(k) \; . \tag{113}$$

2.3 TEMPERATURE ANISOTROPIES

2.3.1 *The Boltzmann equation*

The photons in the early universe are described by the phase-space distribution $f(t,\vec{x},\vec{p})$, with proper time t, spatial coordinate $\vec{x}$ and three-momentum $\vec{p}$, and the equation of motion is the Boltzmann equation

$$\frac{df(t,\vec{x},\vec{p})}{dt} = c[f_\gamma, f_e] \; , \tag{114}$$

where $c[f_\gamma, f_e]$ is the interaction term, which is a function of the photon and electron phase-space distributions, f_γ and f_e. For CMB photons, the only relevant interaction is Thomson scattering with free electrons. From the phase-space distribution, we have seven degrees of freedom, leading to a computationally challenging problem. However, we can simplify the computation by considering two regimes, *I) tightly coupled regime* and *II) after tight coupling*.

I) TIGHTLY COUPLED REGIME, $T > T_{\rm dec}$, $z > z_{dec} \sim 1100$.

At early times the photons, electrons and baryons behave as one effective fluid. Because the photons are in thermal equilibrium, the phase-space distribution is given by a Bose-Einstein distribution,

$$f(t, \vec{x}, \vec{p}) = \frac{1}{\exp(p/T(t, \vec{x})) - 1} \, . \tag{115}$$

We want to expand the temperature and phase-space distribution into a homogeneous and a perturbed part,

$$T(t, \vec{x}) = \bar{T}(t, \vec{x}) + \delta T(t, \vec{x}) \, , \tag{116}$$

$$f(t, \vec{x}, \vec{p}) = \bar{f}(t, \vec{x}, \vec{p}) + \delta f(t, \vec{x}, \vec{p}) \, , \tag{117}$$

where

$$\bar{f}(t, \vec{p}) = \frac{1}{\exp(p/\bar{T}(t, \vec{x})) - 1} \, , \tag{118}$$

$$\delta f(t, \vec{x}, \vec{p}) = \frac{d\bar{f}}{d \ln p} \frac{\delta T}{\bar{T}} \, . \tag{119}$$

We define the temperature anisotropy

$$\Theta(t, \vec{x}) \equiv \frac{\delta T}{\bar{T}} \, , \tag{120}$$

so we have

$$\delta f(t, \vec{x}, \vec{p}) = \frac{d\bar{f}}{d \ln p} \Theta(t, \vec{x}) \, . \tag{121}$$

This means the equation of motion for $\Theta(t, \vec{x})$, i.e. the Boltzmann equation, now is a lower dimensionality problem and therefore easier to solve.

II) AFTER TIGHT COUPLING, $T \lesssim T_{\rm dec}$, $z \lesssim z_{\rm dec}$.

Later on, the photons will decouple. Each decoupled photon follows a geodesic path described by the geodesic equation

$$\frac{dp^\mu}{dt} = -\Gamma^\mu_{\alpha\beta} p^\alpha p^\beta \, , \tag{122}$$

where $\Gamma^\mu_{\alpha\beta}$ is the Christoffel symbol. From this we get, in Newtonian gauge and for conformal time η (with primes indicating derivatives with respect to conformal time),

$$\frac{d(ap)}{d\eta} = ap\phi' - a\epsilon \hat{n} \cdot \vec{\nabla} \psi \, . \tag{123}$$

Here the $\hat{n} = \vec{p}/p$ is the direction of propagation of the photon. For massless photons, the energy $\epsilon \equiv \sqrt{p^2 + m^2}$ is equal to the momentum p, leaving us with

$$\frac{d \ln(ap)}{d\eta} = \phi' - \hat{n} \cdot \vec{\nabla} \psi \, . \tag{124}$$

The first term accounts for the distortion of the "local scale factor" $a(1 + \phi)$ and the second for gravitational blue/redshifting due to the generalized gravitational potential ψ. From this

expression we can see that for photons there can be no deviation from a Bose-Einstein distribution, but we need to account for a directional dependence of the temperature $T(t,\vec{x}) \to T(t,\vec{x},\hat{n})$, i.e.

$$f(t,\vec{x},\vec{p}) = \frac{1}{\exp(p/T(t,\vec{x},\hat{n})) - 1} \,. \tag{125}$$

This dependence will, of course, propagate to the mean and perturbed quantities as well, so that $\Theta(t,\vec{x}) \to \Theta(t,\vec{x},\hat{n})$, picking up two new degrees of freedom (for six in total), i.e. the two angles from $\hat{n}$. From the Boltzmann equation we obtain the new equation of motion for $\Theta(t,\vec{x},\hat{n})$,

$$f(t,\vec{x},\vec{p}) = \bar{f}(t,\vec{x},\vec{p}) \left(1 + \frac{d\ln \bar{f}(t,\vec{x},\vec{p})}{d\ln p}\Theta(t,\vec{x},\hat{n})\right) \,. \tag{126}$$

We can reduce the number of degrees of freedom by transforming to Fourier space $\Theta(t,\vec{x},\hat{n}) \to \Theta(t,\vec{k},\hat{n})$ and by expressing $\hat{n}$ in spherical coordinates with the z axis aligned with $\vec{k}$, so that $\vec{k}\cdot\hat{n} = k\cos\theta$. This leaves us with five degrees of freedom for $\Theta(t,\vec{k},\theta)$.

Next we want to expand the temperature anisotropy in θ with a Legendre multipole expansion,

$$\Theta(t,\vec{k},\theta) = \sum_\ell (-i)^\ell (2\ell+1)\Theta_\ell(t,\vec{k})P_\ell(\cos\theta) \,, \tag{127}$$

where ℓ is the multipole number, P_ℓ are the Legendre polynomials and Θ_ℓ are the temperature anisotropy multipoles, i.e.

$$\Theta_0(t,\vec{k}) = \frac{1}{4}\delta_\gamma(t,\vec{k}) \quad \text{is the monopole, related to the photon density perturbations,} \tag{128}$$

$$\Theta_1(t,\vec{k}) = \frac{1}{3k}\theta_\gamma(t,\vec{k}) \quad \text{is the dipole, which has to do with the velocity divergence,} \tag{129}$$

$$\Theta_2(t,\vec{k}) = \frac{1}{2}\sigma_\gamma(t,\vec{k}) \quad \text{is the quadrupole, linked to the anisotropic stress.} \tag{130}$$

In real space the equation of motion for Θ is

$$\Theta' + \hat{n}\cdot\vec{\nabla}\Theta - \phi' + \hat{n}\cdot\vec{\nabla}\phi = -\Gamma(\Theta - \Theta_0 - \hat{n}\cdot\hat{v}_e) \,, \tag{131}$$

where Γ is the conformal Thomson scattering rate, ignoring small corrections that are only relevant for CMB polarization, and the third and fourth terms on the LHS come from Eq. 124. We now have a general expression for the Boltmann equation for photon anisotropies.

2.3.2 *Directional CMB temperature anisotropy*

We observe the CMB today at t_0 as a temperature map in direction $-\hat{n}$, which we write as a spatial average and a deviation from the average, i.e. a temperature perturbation or anisotropy,

$$T^{\text{obs}}(t_0,\vec{x}_0,-\hat{n}) = \bar{T}(t_0)\left[1 + \Theta(t_0,\vec{x}_0,\hat{n})\right] \tag{132}$$

$$= \bar{T}(t_0)\left[1 + \Theta_0(t_0,\vec{x}_0)\right] + \bar{T}(t_0)\left[\Theta(t_0,\vec{x}_0,\hat{n}) - \Theta_0(t_0,\vec{x}_0)\right] \,, \tag{133}$$

where the first term is the mean temperature of the CMB, $\langle T^{\text{obs}}\rangle = 2.726\ K$, and the second term gives the deviation from the mean δT^{obs}, which is of order 10 μK.

We cannot measure the average temperature across all of the universe, $\bar{T}(t_0) = \langle T \rangle_{\text{all space}}$, but rather we measure the average temperature over all possible directions for our universe,

$$\left\langle T^{\text{obs}} \right\rangle_{\text{all directions}} = \bar{T}(t_0)\left[1 + \Theta_0(t_0, \vec{x}_0)\right] \ , \tag{134}$$

and the temperature anisotropy

$$\delta T^{\text{obs}} = T^{\text{obs}} - \left\langle T^{\text{obs}} \right\rangle_{\text{all directions}} \tag{135}$$

$$= \bar{T}(t_0)\left[\Theta(t_0, \vec{x}_0, \hat{n}) - \Theta_0(t_0, \vec{x}_0)\right] \ . \tag{136}$$

For brevity, in the following we will drop the "all directions" subscript and refer to $\langle T \rangle$ as the directional average and $\bar{T}$ as the spatial average. If we take the ratio of the two, the spatial mean will cancel and we are left with

$$\frac{\delta T^{\text{obs}}}{\langle T^{\text{obs}} \rangle}(-\hat{n}) = \frac{\Theta(t_0, \vec{x}_0, \hat{n}) - \Theta_0(t_0, \vec{x}_0)}{1 + \Theta_0(t_0, \vec{x}_0)} \tag{137}$$

$$\approx \Theta(t_0, \vec{x}_0, \hat{n}) - \Theta_0(t_0, \vec{x}_0) \ , \tag{138}$$

where we used that $\Theta_0(t_0, \vec{x}_0) \sim 10^{-5} \ll 1$.

2.3.3 *The Sachs-Wolfe effect*

The Sachs-Wolfe effect accounts for gravitational red/blue-shifting of photons propagating from the last scattering surface. Photons are gravitationally blue/red-shifted when entering/leaving an overdense part of the universe (or inversely, for an underdense part). This effect is called the Sachs-Wolfe effect and is the reason for the hot and cold spots we see in the CMB (i.e. hot/cold spots consist of photons originating from an under/over-dense region at decoupling). We also need to remember that, the photons we are observing have traversed a universe with an inhomogeneous mass distribution that is evolving in time. As time passes, the over- and under-densities grow more pronounced and the photons take a long time to traverse these regions (e.g. at late times the characteristic scale of a galaxy cluster is of order 1 Mpc or $3.26 \times 10^6 Ly$), meaning the over- and under-densities will have had time to grow. The effect of this is that photons will be more red/blue-shifted when leaving an over/under-density than they were blue/red-shifted going in, which is a cumulative effect called the integrated Sachs-Wolfe (ISW) effect. We start with the line of sight integral, i.e. the Sachs-Wolfe formula [90],

$$\Theta^{\text{obs}}(\eta_0, \vec{x}_0, \hat{n}) = \int_{\eta_{dec}}^{\eta_0} d\eta \ \{g(\eta)(\hat{\Theta} + \psi + \hat{n} \cdot \vec{\nabla} v_b) + \exp\left[-\tau(\eta)\right](\psi' + \phi')\} - \psi(\eta_0, \vec{x}_0) \ , \tag{139}$$

where $\hat{\Theta}$ is intrinsic temperature anisotropy (i.e. before corrections), $g(\eta) = -\tau'(\eta) \ \exp\left[-\tau(\eta)\right]$ is the visibility function, and τ is the optical depth, i.e. the depth of the "diffusion" or "fog" effect caused by photon interactions, which is given by the integral over the scattering rate, $\tau(\eta) = \int_{\eta}^{\eta_0} d\tilde{\eta} \ \Gamma(\tilde{\eta})$.

The last term only contributes to the monopole, so we can safely disregard it. In order to better understand this effect, we can make some simplifying assumptions that will allow us to solve the integral analytically: we assume *instant decoupling* and *neglect reionization*, i.e. the optical depth is replaced by the Heaviside function $e^{-\tau(\eta)} \approx \text{Hea}(\eta - \eta_{\text{dec}})$ and the visibility

function by a delta function $g(\eta) \approx \delta(\eta - \eta_{\rm dec})$. We can now re-write the Sachs-Wolfe formula to

$$\Theta^{\rm obs}(\eta_0, \vec{x}_0, \hat{n}) \approx \hat{\Theta}\left(\eta_{\rm dec}, \vec{x}_{\rm lss}\right) + \psi\left(\eta_{\rm dec}, \vec{x}_{\rm lss}\right) + \hat{n} \cdot \vec{v}_b(\eta_{\rm dec}, \vec{x}_{\rm lss}) \tag{140}$$

$$+ \int_{\eta_{dec}}^{\eta_0} d\eta \, \left\{\phi'(\eta, \vec{x}_0 - \hat{n}[\eta_0 - \eta]) + \psi'(\eta, \vec{x}_0 - \hat{n}[\eta_0 - \eta])\right\} , \tag{141}$$

where the subscript lss denotes the last scattering surface. We can now consider the different terms of this equation. We see the first term is the intrinsic temperature anisotropy from the previous sections, the second term is the Sachs-Wolfe correction, the third is the Doppler term, and the integral is the ISW term.

The line-of-sight integral in real space is convenient for understanding the equations, but in reality we need to go to Fourier space for the actual computations, as they will otherwise be too time consuming. The line-of-sight integral in Fourier space is

$$\Theta^{\rm obs}(\eta_0, \vec{x}_0, \hat{n}) = \int_{\eta_{dec}}^{\eta_0} d\eta \, \left\{g(\eta)(\hat{\Theta} + \psi) + g(\eta)\frac{\theta_b}{k^2} + \exp\left[-\tau(\eta)\right](\psi' + \phi')\right\} j_\ell(k[\eta_0 - \eta]) , \tag{142}$$

where θ_b is the velocity divergence, $j_\ell(x) = \frac{\pi}{2x} J_{\ell+\frac{1}{2}}(x)$ is the spherical Bessel function, and, as before, we disregarded the term that only contributes to the monopole.

The ISW effect is usually separated into two different regimes, even though it is technically the same effect: I) *the early integrated Sachs-Wolfe effect* impacting sub-sound horizon scales before the universe transitions deep into matter domination, where the metric fluctuations are static and contribution to the ISW term is negligible, and II) *the late integrated Sachs-Wolfe effect*, when the universe is dominated by a cosmological constant or by dark energy and the equation of state changes again, leading to varying metric perturbations on all scales. Since the early ISW effect is important in the time shortly after decoupling and on scales smaller than the sound horizon, the effect is one of increasing power of multipoles of around $\ell \sim 200$, corresponding to around the first acoustic peak of the CMB. In contrast, the late ISW effect impacts the largest scales ($\ell \lesssim 50$) and tilts the Sachs-Wolfe plateau.

2.3.4 *Expansion in spherical harmonics*

We expand the CMB temperature anisotropies in spherical harmonics

$$\frac{\delta T^{\rm obs}}{\langle T^{\rm obs}\rangle}(-\hat{n}) = \Theta(\eta_0, \vec{x}_0, -\hat{n}) - \Theta_0(\eta_0, \vec{x}_0) \tag{143}$$

$$= \sum_{\substack{\ell \geq 1, \\ -\ell \leq m \leq \ell}} a_{\ell m} Y_{\ell m}(\hat{n}) , \tag{144}$$

where the ℓ's encode the scale and the m's give the direction. We can invert this expression to obtain

$$a_{\ell m} = \sqrt{\frac{2}{\pi}}\, i^\ell \int d^3\vec{k}\; Y_{\ell m}(\hat{k}) \Theta_\ell(\eta_0, \vec{k}) , \tag{145}$$

Each $\Theta_\ell(\eta_0, \vec{k})$ is a Gaussian random number and we can consider the two-point correlation function

$$\left\langle \Theta^*_\ell(\eta_0, \vec{k}) \Theta_{\ell'}(\eta_0, \vec{k}')\right\rangle = \Theta^*_\ell(\eta_0, k)\Theta_{\ell'}(\eta_0, k') P_{\mathcal{R}}(k) \delta^{(3)}(\vec{k}' - \vec{k}) , \tag{146}$$

where we have defined the transfer function

$$\Theta_\ell(\eta_0, k) \equiv \frac{\Theta_\ell(\eta_0, \vec{k})}{\mathcal{R}(\eta_{\rm ini}, \vec{k})} \,, \tag{147}$$

which does not have a directional dependence. The sum of many Gaussian random numbers is also a Gaussian (central limit theorem), so the $a_{\ell m}$ in Eq. 145 is also a Gaussian random number. We can compute the two-point correlation function

$$\langle a^*_{\ell m} a_{\ell' m'} \rangle = \frac{2}{\pi}\, i^{\ell-\ell'} \int d^3\vec{k} \int d^3\vec{k}'\; Y^*_{\ell m}(\hat{k}) Y_{\ell' m'}(\hat{k}') \left\langle \Theta^*_\ell(\eta_0, \vec{k}) \Theta_{\ell'}(\eta_0, \vec{k}') \right\rangle \,. \tag{148}$$

Inserting Eq. 146, the Kronecker δ takes care of one integral and we are left with

$$\langle a^*_{\ell m} a_{\ell' m'} \rangle = \frac{2}{\pi}\, i^{\ell-\ell'} \int d^3\vec{k}\; Y^*_{\ell m}(\hat{k}) Y_{\ell' m'}(\hat{k}) \Theta^*_\ell(\eta_0, k) \Theta_{\ell'}(\eta_0, k) P_{\mathcal{R}}(k) \,. \tag{149}$$

We can rewrite the integration to

$$\langle a^*_{\ell m} a_{\ell' m'} \rangle = \frac{2}{\pi}\, i^{\ell-\ell'} \int k^2 dk\, d\hat{k}\; Y^*_{\ell m}(\hat{k}) Y_{\ell' m'}(\hat{k}) \Theta^*_\ell(\eta_0, k) \Theta_{\ell'}(\eta_0, k) P_{\mathcal{R}}(k) \,. \tag{150}$$

We have $d\hat{k}\; Y^*_{\ell m}(\hat{k}) Y_{\ell' m'}(\hat{k}) = \delta_{\ell\ell'}\delta_{mm'}$, where the δ's are Kronecker δ's, so we are left with

$$\begin{aligned}
\langle a^*_{\ell m} a_{\ell' m'} \rangle &= \frac{2}{\pi}\, i^{\ell-\ell'} \int k^2 dk\; \delta_{\ell\ell'}\delta_{mm'} \Theta^*_\ell(\eta_0, k) \Theta_{\ell'}(\eta_0, k) P_{\mathcal{R}}(k) && (151)\\
&= \frac{2}{\pi}\, i^{\ell-\ell'} \delta_{\ell\ell'}\delta_{mm'} \int k^2 dk\; \Theta^2_\ell(\eta_0, k) P_{\mathcal{R}}(k) && (152)\\
&= \frac{2}{\pi} \delta_{\ell\ell'}\delta_{mm'} \int k^2 dk\; \Theta^2_\ell(\eta_0, k) P_{\mathcal{R}}(k) \,, && (153)
\end{aligned}$$

where we were able to rewrite the expression because of the Kronecker δ's. Now, replacing the scale dependent primordial power spectrum of curvature perturbations with the scale-independent one $P_{\mathcal{R}}(k) = \frac{2\pi^2}{k^3}\mathcal{P}_{\mathcal{R}}(k)$, we arrive at the final expression

$$\langle a^*_{\ell m} a_{\ell' m'} \rangle = \delta_{\ell\ell'}\delta_{mm'} \left[4\pi \int \frac{dk}{k}\; \Theta^2_\ell(\eta_0, k) \mathcal{P}_{\mathcal{R}}(k) \right] \,. \tag{154}$$

The part in the brackets is called the power spectrum of temperature anisotropies in harmonic space, i.e.

$$C_\ell \equiv 4\pi \int \frac{dk}{k}\; \Theta^2_\ell(\eta_0, k) \mathcal{P}_{\mathcal{R}}(k) \,, \tag{155}$$

where $\Theta_\ell(\eta_0, k)$ plays the role of the transfer function, like $T_m(\eta_0, k)$ in the expression for the matter power spectrum, connecting the primordial power spectrum from inflation with the observed CMB anisotropies.

2.3.5 *Cosmic variance*

The C_ℓ's contain all the information we can extract from the CMB temperature anisotropies about our cosmological model, if the perturbations are linear and Gaussian. Primordial non-Gaussianity has been widely studied (see e.g. [91–94]) and requires measuring the three-point correlation function or bispectrum, but will not be addressed in this thesis. At CMB times, the

growth of perturbations is linear, but the effect of non-linear formation on large-scale structure observables will be an important topic throughout this work. However, even with Gaussian perturbations that only experience linear growth, the information we can extract from the CMB about our model is limited. This would be the case even with the single most sensitive experiment imaginable. The reason is cosmic variance, i.e. we are limited to only one universe and on large scales (small multipoles) we have a small number of Gaussian random numbers on which to base our conclusions. We consider $C_\ell^{\rm obs}$ as an estimator of $C_\ell^{\rm theory}$,

$$C_\ell^{\rm obs} \equiv \frac{1}{2\ell+1} \sum_{-\ell \leq m \leq \ell} \left| a_{\ell m}^{\rm obs} \right|^2 , \tag{156}$$

from which we can estimate the error due to cosmic variance

$$\left\langle \left(C_\ell^{\rm obs} - C_\ell^{\rm theory} \right)^2 \right\rangle = \left\langle \left(\frac{1}{2\ell+1} \sum_{-\ell \leq m \leq \ell} a_{\ell m}^* a_{\ell m} - C_\ell^{\rm theory} \right)^2 \right\rangle \tag{157}$$

$$= \frac{2}{2\ell+1} \left(C_\ell^{\rm theory} \right)^2 , \tag{158}$$

where $\sqrt{\frac{2}{2\ell+1}}$ is then the maximum precision with which we can estimate the power spectrum of the underlying cosmological model from the observed power spectrum, even with infinitely precise observations.

2.4 EXPERIMENTAL CONFIGURATIONS

We are primarily interested in the comparison of four future CMB experiments, based on current expectations for instrumental sensitivities. These are:

- the Litebird[2] satellite project of JAXA [95], currently in phase A, optimized for primordial B-modes, with very good sensitivity but modest resolution.
- the CORE-M5 satellite project [50] recently submitted to the M5 call of ESA and not approved within this call, but still being considered for future applications: CORE-M5 would have a slightly better sensitivity and significantly better resolution than LiteBird. We will also compare the baseline CORE-M5 configuration to four other possible versions: LiteCORE-80, LiteCORE-120, LiteCORE-150, and COrE+ (i.e. the specifications of the CORE proposal to ESA for the M4 call).
- the CMB Stage Four[3] (CMB-S4) project [87], an ambitious project gathering many ground-based detectors to be deployed over the next decade, with outstanding resolution and sensitivity, but smaller sky coverage than satellites.
- and the PICO satellite project[4] that may be submitted to NASA in the future, which would improve over the sensitivity of LiteBird by a factor of 3 to 4.

2 http://litebird.jp/eng/

3 https://cmb-s4.org

4 See https://zzz.physics.umn.edu/ipsig/start and https://zzz.physics.umn.edu/ipsig/_media/pico_science_aas_v11.pdf; channel resolution and sensitivity taken from https://zzz.physics.umn.edu/ipsig/baseline.

For each of these experiments we assume resolutions, sensitivity parameters and sky fractions summarised in Table 1 for LiteBird, CMB-S4, and PICO, in Table 2 for CORE-M5 and in Table 3 for the other CORE-related configurations. The two experiments likely to provide results on the shortest time scale, Litebird and CMB-S4, will be complementary, since they are optimised respectively for large and small angular scales. As such, it is natural to combine them. We then make the same assumption as in [87]: we consider that the optimal combination will consist in LiteBird data for $\ell \leq 50$, CMB-S4 data for $\ell > 50$ in the region covered by the experiment (40% of the sky) and additional high-ℓ data from LiteBird in the region covered by the satellite but not by CMB-S4 (30% of the sky).

Channel [GHz]	FWMH [arcmin]	ΔT [μK arcmin]	ΔP [μK arcmin]
1. LiteBird, $\ell_{\rm max} = 1350, f_{\rm sky} = 0.7$			
140	31	4.1	5.8
2. CMB-S4, $\ell_{\rm min} = 30, \ell_{\rm max} = 3000, f_{\rm sky} = 0.4$			
150	3.0	1.0	1.41
3. LiteBird + CMB-S4 in combination			
low-ℓ from LiteBird, $\ell_{\rm max} = 50, f_{\rm sky} = 0.7$			
140	31	4.1	5.8
high-ℓ from CMB-S4, $\ell_{\rm min} = 51, \ell_{\rm max} = 3000, f_{\rm sky} = 0.4$			
150	3.0	1.0	1.41
additional high-ℓ from LiteBird, $\ell_{\rm min} = 51, \ell_{\rm max} = 1350, f_{\rm sky} = 0.3$			
140	31	4.1	5.8
4. PICO, $\ell_{\rm max} = 3000, f_{\rm sky} = 0.7$			
62.2	12.8	2.76	3.9
74.6	10.7	2.26	3.2
89.6	9.5	1.41	2.0
107.5	7.9	1.20	1.7
129.0	7.4	1.13	1.6
154.8	6.2	0.99	1.4
185.8	4.3	1.84	2.6
222.9	3.6	2.19	3.1

Table 1: Experimental specifications for the CMB experiments used in this work. From left to right, frequency channel(s) dedicated to cosmology, beam width, temperature and polarization sensitivities for this/these channel(s). See the text for references to each experiment or combination of experiments.

It is useful to add Planck to this list of experiments, in order to quantify the progress that can be made compared to the present situation. We choose not to use the actual Planck likelihood, in order to enjoy one feature of MontePython: when running only mock likelihoods, the code automatically creates mock spectra at the beginning of the first run (for fiducial parameter values specified by the user), with, of course, a single fiducial model being used across all likelihoods. Therefore, in our grid of forecast, it is technically easier to use a mock Planck likelihood, with resolution and sensitivity assumptions close to those of the full Planck mission [2]. It is impossible to exactly mimick the real sensitivity of the Planck results, due to our assumption of a Gaussian CMB likelihood with uncorrelated $a_{\ell m}$'s (breaking mainly at low

ℓ's), of uncorrelated temperature and polarisation noise, and of perfect foreground cleaning up to ℓ_{max}. Nevertheless, we made an educated guess for the noise level in our mock Planck likelihood, leading to sensitivities very close to the real ones, with the exception of the error on τ_{reio}, which is smaller in our forecast than in reality by about 50%.

channel GHz	beam arcmin	N_{det}	ΔT μK.arcmin	ΔP μK.arcmin	ΔI μK_{RJ}.arcmin	ΔI kJy/sr.arcmin	$\Delta y \times 10^6$ y_{SZ}.arcmin	PS (5σ) mJy
60	17.87	48	7.5	10.6	6.81	0.75	-1.5	5.0
70	15.39	48	7.1	10	6.23	0.94	-1.5	5.4
80	13.52	48	6.8	9.6	5.76	1.13	-1.5	5.7
90	12.08	78	5.1	7.3	4.19	1.04	-1.2	4.7
100	10.92	78	5.0	7.1	3.90	1.2	-1.2	4.9
115	9.56	76	5.0	7.0	3.58	1.45	-1.3	5.2
130	8.51	124	3.9	5.5	2.55	1.32	-1.2	4.2
145	7.68	144	3.6	5.1	2.16	1.39	-1.3	4.0
160	7.01	144	3.7	5.2	1.98	1.55	-1.6	4.1
175	6.45	160	3.6	5.1	1.72	1.62	-2.1	3.9
195	5.84	192	3.5	4.9	1.41	1.65	-3.8	3.6
220	5.23	192	3.8	5.4	1.24	1.85	-	3.6
255	4.57	128	5.6	7.9	1.30	2.59	3.5	4.4
295	3.99	128	7.4	10.5	1.12	3.01	2.2	4.5
340	3.49	128	11.1	15.7	1.01	3.57	2.0	4.7
390	3.06	96	22.0	31.1	1.08	5.05	2.8	5.8
450	2.65	96	45.9	64.9	1.04	6.48	4.3	6.5
520	2.29	96	116.6	164.8	1.03	8.56	8.3	7.4
600	1.98	96	358.3	506.7	1.03	11.4	20.0	8.5
Array		2100	1.2	1.7			0.41	

Table 2: Proposed CORE-M5 frequency channels. The sensitivity is calculated assuming $\Delta\nu/\nu = 30\%$ bandwidth, 60% optical efficiency, total noise of twice the expected photon noise from the sky and the optics of the instrument at 40K temperature. This configuration has 2100 detectors, about 45% of which are located in CMB channels between 130 and 220 GHz. Those six CMB channels yield an aggregated CMB sensitivity of 2 μK.arcmin (1.7 μK.arcmin for the full array).

The specifications for the CORE-M5 proposal, a complete survey of polarised sky emission in 19 frequency bands, with sensitivity and angular resolution requirements is summarized in Table 2. Obviously, data from low (60-115 GHz) and high frequencies (255-600 GHz) channels will be mainly used for monitoring foreground contaminations (and deliver rich related science). In our forecasts, we therefore use only the six channels in the frequency range of $130 - 220$ GHz and refer to this experimental configuration as **CORE-M5**.

Channel [GHz]	FWMH [arcmin]	ΔT [μK arcmin]	ΔP [μK arcmin]
LiteCORE-80, $\ell_{\rm max} = 2400, f_{\rm sky} = 0.7$			
80	20.2	8.8	12.5
90	17.8	7.1	10.0
100	15.8	8.5	12.0
120	13.2	6.7	9.5
140	11.2	5.3	7.5
166	8.5	5.0	7.0
195	8.1	3.6	5.0
LiteCORE-120, $\ell_{\rm max} = 3000, f_{\rm sky} = 0.7$			
80	13.5	8.8	12.5
90	11.9	7.1	10.0
100	10.5	8.5	12.0
120	8.8	6.7	9.5
140	7.4	5.3	7.5
166	6.3	5.0	7.0
195	5.4	3.6	5.0
LiteCORE-150, $\ell_{\rm max} = 3000, f_{\rm sky} = 0.7$			
80	10.8	8.8	12.5
90	9.5	7.1	10.0
100	8.4	8.5	12.0
120	7.0	6.7	9.5
140	5.9	5.3	7.5
166	5.0	5.0	7.0
195	4.3	3.6	5.0
COrE+, $\ell_{\rm max} = 3000, f_{\rm sky} = 0.7$			
100	8.4	6.0	8.5
115	7.3	5.0	7.0
130	6.5	4.2	5.9
145	5.8	3.6	5.0
160	5.3	3.8	5.4
175	4.8	3.8	5.3
195	4.3	3.8	5.3
220	3.8	5.8	8.1

Table 3: Experimental specifications for LiteCORE-80, LiteCORE-120, LiteCORE-150 and COrE+: frequency channels dedicated to cosmology, beam width, temperature and polarization sensitivities for each channel.

2.5 LIKELIHOOD AND METHODOLOGY

For CMB experiments, we assume a Gaussian likelihood for the multipole coefficients of temperature, polarization and CMB lensing potential maps, described by equations (3.1) to (3.7) of [96]. The noise spectrum of temperature and polarization are inferred from the resolution and sensitivity parameters expected to reflect the instrumental characteristics according to standard approximations (see e.g. equation (2.2) in [96]). The temperature and polarization

noises are assumed to be statistically independent, which means that the noise spectrum N_ℓ^{TE} is approximated as zero. Given the fiducial model and noise spectra, one can estimate the error that would be performed on the measurement of the lensing potential spectra by running a quadratic estimator [97] (using products of four multipoles, each of the T, E or B type, but avoiding any auto-correlation of the B-mode maps due to the non-Gaussianity of the $a_{\ell m}^B$ multipoles. All quadratic estimators can then be combined in order to minimize the combined noise: this defined the minimum variance estimator [97]. We use the FuturCMB[5] code [96] to compute the noise spectrum of the CMB lensing potential expected from this minimum variance estimator technique.

We use the temperature, cross temperature-polarization, and E-mode polarization spectra, C_ℓ^{TT}, C_ℓ^{TE}, and C_ℓ^{EE}, but do not consider the B-mode channel. The missions considered are also sensitive to the BB lensing polarization signal, but we take the conservative approach to not include it in the forecasts. This leaves open the possibility of using the channel for further checks for foregrounds contamination and systematics. Note that in this work, we consider fiducial models with negligible primordial gravitational waves from inflation. Otherwise, the BB channel would contain primary signal on large angular scales and could not be neglected.

In the mock likelihoods, the variance of the "observed" multipoles a_{lm}'s is given by the sum of the fiducial C_ℓ's and of an instrumental noise spectrum given by

$$N_\ell = w^{-1} \exp(\ell(\ell+1)\theta^2/8\ln 2) \,, \tag{159}$$

where θ is the FWHM of the beam, assuming a Gaussian profile, and where w^{-1} is the experimental power noise related to the detectors sensitivity σ by $w^{-1} = (\theta\sigma)^2$.

We assume that beam uncertainties are small and that uncertainties due to foreground removal are smaller than statistical errors. In Figure 1 we show, for each CORE-related configuration, the variance $C_l + N_l$ compared to the fiducial model C_l for the temperature (left) and polarisation (middle) auto-correlation spectra. The data are cosmic-variance-limited up to the multipole at which this variance departs from the fiducial model. The figure also shows that the lensing reconstruction noise is different on all scales for the various configurations.

Together with the primary anisotropy signal, we also take into account information from CMB weak lensing, considering the power spectrum of the CMB lensing potential C_ℓ^{PP}. In what follows, we use the quadratic estimator method of Hu & Okamoto [97], that provides an algorithm for estimating the corresponding noise spectrum N_ℓ^{PP} from the observed CMB primary anisotropy and noise power spectra. Like in [98], we use, for chapter 5 and sections 7.1, 8.1 and 9.1 to 9.3, the noise spectrum N_ℓ^{PP} associated to the EB estimator of lensing, which is the most sensitive one for all CORE configurations (out of all pairs of maps). We occasionally repeated the analysis with the actual minimum variance estimator, and found very similar results. For section 9.5 we always use the minimum variance estimator based on all pairs of maps [97, 99].

$$-2\ln\mathcal{L} = \sum_l (2l+1) f_{\rm sky} \left(\frac{D}{|\bar{C}|} + \ln\frac{|\bar{C}|}{|\hat{C}|} - 3 \right) \,, \tag{160}$$

5 http://lpsc.in2p3.fr/perotto/

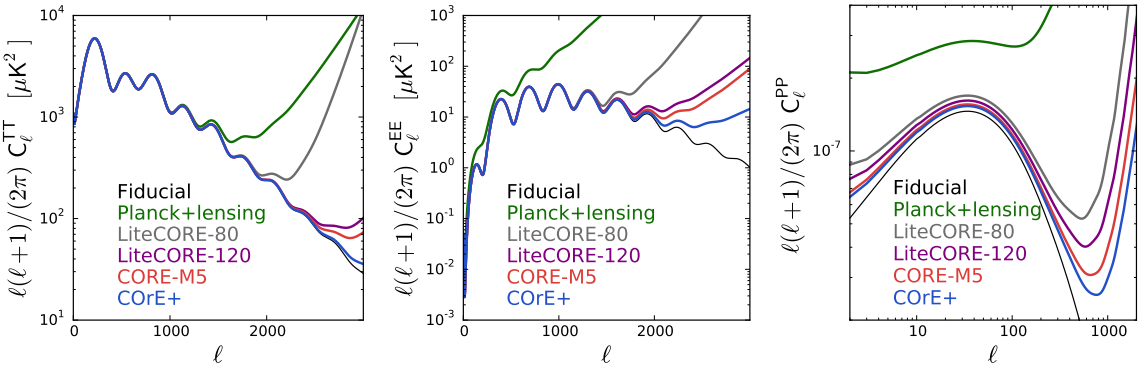

Figure 1: Fiducial model and variance $C_l + N_l$ of each data point a_{lm}, given the sensitivity of each CORE configuration (Planck is also shown for comparison). As long as the variance traces the fiducial model, the data is cosmic variance limited. This happens down to different angular scales for the temperature (left) and E-mode polarisation (middle). For CMB lensing extraction (right), on all scales, there is a substantial difference between the noise level of the different configurations.

where $\bar{C}_l$ and $\hat{C}_l$ are the fiducial and theoretical spectra plus noise respectively, $|\bar{C}|$, $|\hat{C}|$ denote the determinants of the theoretical and observed data covariance matrices respectively,

$$|\bar{C}| = \bar{C}_\ell^{TT}\bar{C}_\ell^{EE}\bar{C}_\ell^{PP} - \left(\bar{C}_\ell^{TE}\right)^2\bar{C}_\ell^{PP} - \left(\bar{C}_\ell^{TP}\right)^2\bar{C}_\ell^{EE}\,, \tag{161}$$

$$|\hat{C}| = \hat{C}_\ell^{TT}\hat{C}_\ell^{EE}\hat{C}_\ell^{PP} - \left(\hat{C}_\ell^{TE}\right)^2\hat{C}_\ell^{PP} - \left(\hat{C}_\ell^{TP}\right)^2\hat{C}_\ell^{EE}\,, \tag{162}$$

D is defined as

$$\begin{aligned} D = {} & \hat{C}_\ell^{TT}\bar{C}_\ell^{EE}\bar{C}_V^{PP} + \bar{C}_\ell^{TT}\hat{C}_\ell^{EE}\bar{C}_\ell^{PP} + \bar{C}_\ell^{TT}\bar{C}_\ell^{EE}\hat{C}_\ell^{PP} \\ & -\bar{C}_\ell^{TE}\left(\bar{C}_\ell^{TE}\hat{C}_\ell^{PP} + 2\hat{C}_\ell^{TE}\bar{C}_\ell^{PP}\right) \\ & -\bar{C}_\ell^{TP}\left(\bar{C}_\ell^{TP}\hat{C}_\ell^{EE} + 2\hat{C}_\ell^{TP}\bar{C}_\ell^{EE}\right)\,, \end{aligned} \tag{163}$$

and finally f_{sky} is the sky fraction sampled by the experiment after foregrounds removal.

Note that for temperature and polarization, $\bar{C}_l$ and $\hat{C}_l$ could be defined to include the lensed or unlensed fiducial and theoretical spectra, and in both cases the above likelihood is slightly incorrect. If we use the unlensed spectra, we optimistically assume that we will be able to do a perfect de-lensing of the T and E map, based on the measurement of the lensing map with quadratic estimators. If we use the lensed spectra, we take the risk of double-counting the same information in two observables which are not statistically independent: the lensing spectrum, and the lensing corrections to the TT, EE and TE spectra. To deal with this issue, one could adopt a more advanced formalism including non-Gaussian corrections, like in [100, 101]. However, we performed dedicated forecasts to compare the two approximate Gaussian likelihoods, and even with the best sensitivity settings of *COrE+* we found nearly indistinguishable results (at least for the ΛCDM+M_ν model). The reconstructed parameter errors change by negligible amounts between the two cases. The biggest impact is on the error on the sound horizon angular scale $\sigma(\theta_s)$, which is 5% smaller when using unlensed spectra, because perfect de-lensing would allow to better identify the primary peak scales. When using the lensed spectra, we do not observe any statistically significant reduction of the error bars, and we

conclude that over-counting the lensing information is not important for an experiment with the sensitivity of *COrE+*.

Hence, in the rest of this work we choose to always use the version of the Gaussian likelihood that includes lensed TT, EE and TE spectra. We will occasionally refer to our full CMB likelihoods with the acronym "TEP", standing for "Temperature, E-polarisation and lensing Potential data".

3

LARGE-SCALE STRUCTURE

Chapter 3 is based on Sprenger, Archidiacono, Brinckmann, Clesse & Lesgourgues [102]

Future surveys like Euclid and SKA will reach the required sensitivity and angular resolution to probe the non-linear growth of structures on small scales. Their figure of merit will therefore strongly depend on our understanding of the various, often complex physical processes at play on these scales. This includes general relativistic corrections to nonlinear structure formation (see e.g. [103] and references therein), the accuracy of Newtonian N-body simulations [104] and fitting methods [105, 106] or emulators [107], the galaxy non-linear bias [108], the baryonic feedback [109–111], the intrinsic alignment of galaxies [112], etc. Usually, when doing forecasts, these considerations lead to the introduction of a maximal wavenumber $k_{\rm max}$, below which one trusts the theoretical prediction of the matter power spectrum, and above which the physical uncertainties are expected to exceed the experimental noise. For instance, $k_{\rm max} = 0.2\ h\ {\rm Mpc}^{-1}$ (h defining the Hubble expansion rate today $H_0 = h \times 100$ km/s/Mpc) was often used in Euclid forecasts [3, 4]. Introducing such a cut-off scale means that, all the information obtained by the experiment on smaller scales is simply unexploited. However, this information is crucial: indeed, the non-linear growth of structure can be significantly altered in theories of modified gravity [113–115], interacting/decaying dark matter, and for massive neutrinos [116]. Moreover, a larger lever arm allows for better constraints on the initial shape of the power spectrum of density fluctuations from inflation [117].

Therefore, new methods need to be developed and tested in order to take into account the non-linear theoretical uncertainties, while optimizing the amount of information relevant for cosmology. For instance, it has been proposed to introduce either an error on the power spectrum, totally uncorrelated between wavelength modes, or, on the contrary, a correlated error increasing at small scales [118]. Significant differences in forecasts were found between these two extremes. A totally uncorrelated error would underestimate the theoretical error, given that non-linear physics should affect the matter power spectrum in a relatively smooth way, as discussed in Ref. [119], whereas a correlated error may overestimate it, given that different physical processes will have a different impact on different scales. As an example, baryonic feedback is expected to dominate on scales $0.5 \lesssim k_{\rm max} \lesssim 2\ h\ {\rm Mpc}^{-1}$, whereas the numerical uncertainty of N-body simulations increases monotonically when going to smaller scales.

In this chapter, we introduce a new numerical method to take into account the theoretical uncertainties on the non-linear spectra in a more realistical way. Two cases, based on current knowledge and expected improvements, will be presented: first, a *conservative* case, based on the present theoretical non-linear uncertainties combined with a conservative redshift dependent cut-off scale, and second, a *realistic* case in which we consider some expected and realistic refinement in the modelisation of nonlinear effects, e.g. simply due to the increasing

numerical resources by the time the real data will be available in the 2020's. These two cases allow us to present what we think are the most realistic range for the future constraints on the cosmological parameters of the standard cosmological model, including the sum of neutrino masses, as well as on the parameters of common extended cosmological scenarios. For the first time, we use such a method to derive realistic forecasts for both Euclid and the SKA, and for the combination of them, using three probes: galaxy clustering power spectrum, cosmic shear angular power spectrum, and 21-cm intensity mapping subsequent to reionization. For this purpose, we have used the Bayesian Markov Chain Monte-Carlo (MCMC) technique, rather than the Fisher matrix formalism that might be subject to numerical instabilities, particularly in non-standard cosmological scenarios.

The chapter is organized as follows: in sections 3.1 to 3.3, we introduce respectively the calculations of the galaxy clustering power spectrum, the weak lensing angular power spectrum, and the 21cm intensity mapping power spectrum, as well as the related experimental uncertainties due to Euclid and SKA specifications, and a description of the likelihood computations. Finally, in section 3.4, the method used to model the nonlinear theoretical uncertainties is described.

3.1 GALAXY CLUSTERING

3.1.1 *Galaxy power spectrum*

The spatial distribution of galaxies represents a biased tracer of the underlying dark matter distribution. Therefore, various effects have to be taken into account when converting the matter power spectrum[1] P_m into the observed galaxy power spectrum P_g, i.e.

$$P_g(k,\mu,z) = f_{\rm AP}(z) \times f_{\rm res}(k,\mu,z) \times f_{\rm RSD}(\hat{k},\hat{\mu},z) \times b^2(z) \times P_m(\hat{k},z) \ . \tag{164}$$

Before proceeding with explaining the different effects and the associated functions f_i contributing to this formula, let us notice that we have employed a flat-sky approximation [122, 123] that allows for an unambiguous definition of the angle between the Fourier modes $\boldsymbol{k}$ and the line of sight distance vector $\boldsymbol{r}$. The observer's fixed point of view breaks the isotropy of the matter power spectrum, but symmetry in perpendicular directions to the line of sight is preserved. Hence, the following coordinates are sufficient to describe all the effects,

$$k = |\boldsymbol{k}| \ , \ \mu = \frac{\boldsymbol{k}\cdot\boldsymbol{r}}{kr} \ . \tag{165}$$

The parallel part of a mode is given by $k_{\parallel} = \mu k$ and the perpendicular one by $k_{\perp} = k\sqrt{1-\mu^2}$.

Since we can observe only the redshift and the position in the sky, in order to get a distribution in three-dimensional space, we need to make assumptions on the underlying cosmology[2]. However, physical quantities calculated within this fiducial cosmology may differ from the corresponding values in the true/real cosmology (hereafter denoted by $\hat{}$, e.g. $\hat{H}$). The Fourier modes of real space can be related to those of the fiducial space via

$$\hat{k}^2 = \left[\left(\frac{\hat{H}}{H}\right)^2 \mu^2 + \left(\frac{D_A}{\hat{D}_A}\right)^2 (1-\mu^2)\right] k^2 \tag{166}$$

1 As explained later in section 8.4, when the model features massive neutrinos, we don't plug in here the total matter power spectrum $P_m(k,z)$, but only the power spectrum of baryons and CDM $P_{cb}(k,z)$, since the galaxy power spectrum is more a tracer of the latter quantity, see e.g. [1, 120, 121].

2 This is actually one of the reasons for which alternative methods to express the 2-point statistics of galaxy distributions are being discussed in the literature, see e.g. [103] and references therein.

and

$$\hat{\mu}^2 = \left(\frac{\hat{H}}{H}\right)^2 \mu^2 \cdot \left[\left(\frac{\hat{H}}{H}\right)^2 \mu^2 + \left(\frac{D_A}{\hat{D}_A}\right)^2 (1-\mu^2)\right]^{-1} , \tag{167}$$

where H and D_A are, respectively, the Hubble parameter and the angular diameter distance as functions of redshift z. The change in the power spectrum when extracted from the same data but assuming different cosmologies, the so-called Alcock-Paczinsky effect, gives rise to the first term in Eq. 164,

$$f_{\mathrm{AP}}(z) = \frac{D_A^2 \hat{H}}{\hat{D}_A^2 H} . \tag{168}$$

The second term in Eq. 164 is due to the limited resolution of instruments suppressing the apparent perturbations on small scales. Assuming Gaussian errors $\sigma_{\parallel}(z)$ and $\sigma_{\perp}(z)$ on coordinates parallel and perpendicular to the line of sight at redshift z, the suppression factor turns out to be exponential,

$$f_{\mathrm{res}}(k,\mu,z) = \exp\left(-k^2\left[\mu^2 \cdot \left(\sigma_{\parallel}^2(z) - \sigma_{\perp}^2(z)\right) + \sigma_{\perp}^2(z)\right]\right) . \tag{169}$$

Since Fourier modes scale inversely w.r.t. spatial distances under a change of cosmology, the above factor is independent of the assumed cosmology.

The cosmological redshift, that is used to obtain the spatial coordinates, is not the only source of redshift. The classical Doppler effect induces an apparent anisotropy in the redshift-space power spectrum. On large scales within the linear regime, this effect is described by the Kaiser formula [124]. On top of this large scale infall, additional random peculiar velocities of the galaxies further distort the redshift information on smaller scales, leading to features in redshift-space called fingers of God [125]. Following Ref. [126], we describe this additional suppression with an exponential factor. To sum up, the redshift effects encoded in the third term of Eq. 164 are given by

$$f_{\mathrm{RSD}}(\hat{k},\hat{\mu},z) = \left(1+\beta(\hat{k},z)\,\hat{\mu}^2\right)^2 e^{-\hat{k}^2\hat{\mu}^2\sigma_{\mathrm{NL}}^2} , \tag{170}$$

where the first term in parentheses corresponds to the Kaiser formula and the exponential accounts for the fingers of God. In particular, σ_{NL} has a fiducial value of 7 Mpc and we allow it to vary between 4-10 Mpc in our forecasts, while β is the (possibly scale-dependent) growth rate $f(\hat{k},z)$ corrected by the galaxy bias $b(z)$,

$$\beta(\hat{k},z) \equiv \frac{f(\hat{k},z)}{b(z)} \equiv \frac{1}{b(z)} \cdot \frac{d\ln\left(\sqrt{P_m(\hat{k},z)}\right)}{d\ln a} = -\frac{1+z}{2b(z)} \cdot \frac{d\ln P_m(\hat{k},z)}{dz} . \tag{171}$$

The bias is a function of redshift which relates density perturbations in the galaxy field to dark matter density perturbations. We will assume the linear approximation $\delta_g = b(z) \times \delta_m$ where the bias is scale independent[3]. Approximate formulas for the bias are obtained by populating cosmological simulations with galaxies which will then be measured. In this case, δ_g and δ_m can be identified separately [120, 127, 128].

We divide the surveys into bins of width $\Delta z = 0.1$ with mean redshift $\bar{z}$. Correlation functions are defined inside the bin's data and are approximated to probe the power spectrum at a fixed redshift $\bar{z}$. The volume of one redshift bin can be computed via

$$V_r(\bar{z}) = 4\pi f_{\mathrm{sky}} \cdot \int_{\Delta r(\bar{z})} r^2 dr = \frac{4\pi}{3} f_{\mathrm{sky}} \cdot \left[r^3\left(\bar{z}+\frac{\Delta z}{2}\right) - r^3\left(\bar{z}-\frac{\Delta z}{2}\right)\right] , \tag{172}$$

3 We will take into account the consequences of non-linear bias later, either in our choice of a cut-off k_{max} or through our ansatz for the theoretical error function: this will be discussed in section 3.4.2.

Table 4: Euclid specifications. [4, 118]

parameter	$z_{\min}$	$z_{\max}$	$f_{\rm sky}$	σ_z	σ_θ [']
Euclid	0.45	2.05	0.3636	$0.001(1+z)$	0

where $f_{\rm sky}$ is the fraction of the sky covered by the survey. The distribution of galaxies is discrete, rather than continuous like the density field δ_g. Therefore, we have to take into account the experimental shot noise in each redshift bin,

$$P_N(\bar{z}) = \frac{1}{\bar{n}(\bar{z})} = \frac{V_r(\bar{z})}{N(\bar{z})} \,, \tag{173}$$

where $N(\bar{z})$ is the number of galaxies in the bin, $V_r(\bar{z})$ the volume of the bin and $\bar{n}(\bar{z})$ the galaxy number density. Taking this shot noise into account, the quantity actually measured by the experiment in each bin is

$$P_{\rm obs}(k,\mu,\bar{z}) = P_g(k,\mu,\bar{z}) + P_N(\bar{z}) \,. \tag{174}$$

3.1.2 *Euclid specifications*

The redshift range accessible to Euclid is roughly $0.45 - 2.05$. Hence, mean redshifts of $\bar{z} = 0.5, 0.6, ..., 2.0$ are used. The error on spectroscopic redshift measurements is assumed to be $\sigma_z = 0.001(1+z)$, as in Refs. [4, 118]. The effect of angular resolution is neglected. Thus, $\sigma_\perp$ is set to 0. The specifications for Euclid are summarized in Table 4. The redshift error can be propagated to the error on radial distance,

$$\sigma_{\shortparallel} = \frac{c}{H}\sigma_z \,. \tag{175}$$

The galaxy number count distribution $\frac{dN(z)/dz}{1\text{deg}^2}$ has been taken from Table 2 of Ref. [129] assuming a limiting flux of $3\times10^{-16}\,\text{erg}\,\text{s}^{-1}\,\text{cm}^{-2}$. The values have been divided by 1.37 as recommended for conservative forecasts. We use a sky fraction of $f_{\rm sky} = 0.3636$. The total number of detected galaxies in a given redshift bin can be inferred from the given values,

$$N(\bar{z}) = 41253\, f_{\rm sky}\, \text{deg}^2 \cdot \int_{\bar{z}-\frac{\Delta z}{2}}^{\bar{z}+\frac{\Delta z}{2}} \frac{dN(z)/dz}{1\text{deg}^2} dz \,. \tag{176}$$

As done in Ref. [118], the bias factor corresponding to galaxies detected by Euclid is assumed to be approximately given by the simple relation (see, however, Ref. [130] for discussion about a more realistic scale-dependent galaxy bias relation),

$$b(z) = \sqrt{1+z} \,. \tag{177}$$

In order to account for inaccuracies in this relation, we have introduced two nuisance parameters with mean value 1 in the form of the relation

$$b(z) = \beta_0^{\rm Euclid}(1+z)^{0.5\beta_1^{\rm Euclid}} \,, \tag{178}$$

where a 5%-precision (2σ) is taken as a prior on the β-factors.

3.1.3 *SKA specifications*

Currently, SKA1-MID Band 2 is the most promising option for a galaxy survey with SKA1. We use specifications according to the baseline design defined in Ref. [131] (see Table 6). We assume a survey area $S_{\rm area} = f_{\rm sky} \times 41253\,{\rm deg}^2$ in agreement with the optimization procedure described in Ref. [132], while the frequency range of SKA2 is the same of Ref. [133].

The target signal of SKA is the HI line or 21cm line of cold neutral hydrogen with a rest frequency of $\nu_0 = 1420\,{\rm MHz}$. The frequency range translates through the redshifting of ν_0 into a redshift range which has been rounded to fit redshift bins of width $\Delta z = 0.1$. The frequency ν and its error translate into redshifts

$$z = \frac{\nu_0}{\nu} - 1\,, \tag{179}$$

$$\sigma_z = (1+z)^2 \frac{\sigma_\nu}{\nu_0}\,. \tag{180}$$

These relations are independent of cosmology. Therefore it is valid to treat z as a direct observable as was done in the case of Euclid.

The number counts of detected galaxies and their bias w.r.t. the underlying dark matter distribution have to be extracted from simulations. This was done in Ref. [132] using the following fitting formula

$$\frac{dN(z)/dz}{1{\rm deg}^2} = 10^{c_1} z^{c_2} \exp(-c_3 z)\,, \tag{181}$$

$$b_{\rm HI}(z) = c_4 \exp(c_5 z)\,. \tag{182}$$

Ref. [133] adapted the results to the current baseline design, obtaining the parameters listed in Table 5. Ref. [132] used a frequency resolution of 10 kHz. Band 2 divided into 64,000 channels (see [131]) yields a bandwidth of $\delta\nu = 12.7\,{\rm kHz}$ per channel, which verifies this number. The same approach yields $\delta\nu = 12.8\,{\rm kHz}$ for SKA2. By equating $\delta\nu$ to the full width at half maximum (FWHM), the approximation of a Gaussian error $\sigma_\nu = \delta\nu/\sqrt{8\ln 2}$ can be made. This determines the error on the redshift measurement as described in Eq. 180. Note that we are equating here the frequency sensitivity to the sensitivity of the mean frequency of a galaxy's signal. However, the fingers of God described by σ_{NL}, whose effect is indistinguishable from the redshift resolution, dominate the suppression of the power spectrum. Hence, this approximation is good enough.

The inaccuracy of the theoretical bias formula can be accommodated for with similar nuisance parameters as in the case of Euclid,

$$b(z) = c_4 \beta_0^{\rm SKA1/2} \exp(c_5 \beta_1^{\rm SKA1/2} z)\,, \tag{183}$$

where $\beta_0^{\rm SKA1/2}$, $\beta_1^{\rm SKA1/2}$ are assigned Gaussian priors with mean value 1 and standard deviation 0.025. We also include the effect of angular resolution as a Gaussian error,

$$\sigma_\perp = (1+z) D_A \sigma_\theta\,, \tag{184}$$

$$\sigma_\theta = \frac{1}{\sqrt{8\ln 2}} \frac{\lambda_0}{B} (1+z)\,. \tag{185}$$

The FWHM of an interferometer is approximately given by the wavelength divided by the maximum baseline B. In the case of SKA, the wavelength is the redshifted rest wavelength

Table 5: Fitting parameters. [133]

parameter	c_1	c_2	c_3	c_4	c_5
SKA1 band 2 (5σ)	5.450	1.310	14.394	0.616	1.017
SKA2 (10σ)	6.319	1.736	5.424	0.554	0.783

Table 6: SKA specifications. [131, 133]

parameter	$\nu_{\min}$ [MHz]	$\nu_{\max}$ [MHz]	$z_{\min}$	$z_{\max}$	S_{area} [deg^2]	$\delta\nu$ [kHz]	B [km]
SKA1 band 2	950	1760	0.00	0.50	5,000	12.7	150 (5)
SKA2	470	1290	0.10	2.00	30,000	12.8	3000 (5)

$\lambda_0 = 21.11\,\text{cm}$. The maximum baseline is $B \approx 150\,\text{km}$ for SKA1 and $B \approx 3000\,\text{km}$ for SKA2. However, given the large sky fraction, the survey is not expected to exploit the maximum resolution the array is capable of. The simulated number counts are valid for a 10,000 hour survey. Taking a conservative approach, we use a maximum baseline of 5 km for both SKA1 and SKA2, corresponding to the diameter of the inner core of the array with a high density of dishes. Even with this approach, the effect of angular resolution remains insignificant for a galaxy survey.

3.1.4 *Galaxy clustering likelihood*

The galaxy power spectrum is defined as a function of a continuous density field, which represents the probability density of finding a galaxy at some position $\boldsymbol{r}$. The galaxy density perturbation δ_g is then a perturbation of this probability density p_g,

$$p_g(\boldsymbol{r}) = \bar{n}(\boldsymbol{r})(1 + \delta_g(\boldsymbol{r}))\,, \tag{186}$$

where $\bar{n}(\boldsymbol{r})$ is the expected number density of galaxies on a homogeneous background; it is calculated as the mean density over a sufficiently large volume. In our case this will be the volume corresponding to one redshift bin.

Starting from this idea, Ref. [134] derived a method to estimate the galaxy power spectrum P_g, with an error that will be used in the following to build up the likelihood. However, it is worth noting that this is a simplified approach ignoring the possible effect of galaxy properties such as luminosity (see e.g. [135]). The result is a Gaussian error

$$\sigma_P^2 = \frac{(2\pi)^3}{V_k V_r} P_{\text{obs}}^2(\boldsymbol{k}) = \frac{(2\pi)^3}{V_k V_r} \left(P_g(\boldsymbol{k}) + P_N\right)^2\,. \tag{187}$$

The quantity in parentheses on the RHS is the observable power spectrum of Eq. (174) split into the part proportional to the matter power spectrum and the shot noise $P_N = 1/\bar{n} = V_r/N$, where N is the total number of detected galaxies inside the observed volume V_r. The volume V_k of the shell in k-space over which the estimator is averaged has to be big enough to cancel the effects of performing a Fourier transform on a finite volume. It can be chosen to be as small as $V_k = \frac{(2\pi)^3}{V_r}$, where V_r is the volume of the single redshift bin. Inside this volume Fourier

modes are not independent. This means the likelihood is the product of as many Gaussians as there are independent Fourier modes (one per V_k),

$$\mathcal{L} = \mathcal{N} \exp\left[-\frac{1}{2} \sum_{\text{independent } \boldsymbol{k}} \frac{V_k V_r}{(2\pi)^3} \frac{\left(\hat{P}_{\text{obs}}(\boldsymbol{k}) - P_{\text{obs}}(\boldsymbol{k})\right)^2}{P^2_{\text{obs}}(\boldsymbol{k})}\right] . \tag{188}$$

In the case of forecasts, $\hat{P}_{\text{obs}}$ does not represent actual measurements, but it is mock data corresponding to some fiducial cosmology, computed in the same way as the theoretical P_{obs}. Hence, $\hat{P}_{\text{obs}}$ and P_{obs} do not suffer from finite volume effects and are thus smooth, i.e. approximately constant, inside V_k. To replace the sum by an integral over the whole k-space an additional factor of $1/2$ has to be introduced to account for the fact that the power spectrum is the Fourier transform of a real quantity, $P_{\text{obs}}(\boldsymbol{k}) = P_{\text{obs}}(-\boldsymbol{k})$. Reformulated in terms of $\chi^2 = -2 \ln \mathcal{L}$, the result is

$$\chi^2 = \sum_{\bar{z}} \int d^3k \frac{V_r(\bar{z})}{2(2\pi)^3} \frac{\left(\hat{P}_{\text{obs}}(\boldsymbol{k}, \bar{z}) - P_{\text{obs}}(\boldsymbol{k}, \bar{z})\right)^2}{P^2_{\text{obs}}(\boldsymbol{k}, \bar{z})} . \tag{189}$$

The dependencies of all quantities are shown here for clarity. Inside a redshift bin, all quantities are evaluated at the mean redshift $\bar{z}$ of that bin. In other words, anything is evaluated at the same time. This approximation has to be made to get information on the equal-time three-dimensional power spectrum $P_{\text{obs}}(\boldsymbol{k}, \bar{z})$.

The change of coordinates and Fourier modes depending on the choice of the cosmological model has already been discussed in Sec. 3.1. Here, in order to estimate the χ^2 we have to deal with three different cosmologies: the fiducial one used to compute $\hat{P}_{\text{obs}}$, the one used to compute P_{obs}, which we want to compare to the mock data, and, finally, the one used to interpret the observations, in which all quantities of Eq. 189 are defined. The choice of the last one is arbitrary since all conversion factors in Eq. 189 cancel. Therefore, we can assume it to be equal to the fiducial one. To clarify what was done in the computation of forecasts, eq. (189) can be rewritten such that every quantity depending on the underlying cosmology is labelled with either f (fiducial cosmology) or s (sample cosmology):

$$\chi^2 = \sum_{\bar{z}} \int (k^f)^2 dk^f \int_{-1}^{1} d\mu^f \frac{V_r^f}{2(2\pi)^2} \left[\frac{\frac{H^f}{(D_A^f)^2} P^f_{\text{obs}}(k^f, \mu^f) - \frac{H^s}{(D_A^s)^2} P^s_{\text{obs}}(k^s, \mu^s)}{\frac{H^s}{(D_A^s)^2} P^s_{\text{obs}}(k^s, \mu^s)} \right]^2 . \tag{190}$$

We can now replace the observed power spectrum as a function of the galaxy power spectrum and of the shot noise, and make use of the relation between the volumes in the two different spaces:

$$\frac{H^s}{(D_A^s)^2} \frac{V_r^s}{N} = \frac{H^f}{(D_A^f)^2} \frac{V_r^f}{N} . \tag{191}$$

Then the shot noise exactly cancels from the numerator and we are left with

$$\chi^2 = \sum_{\bar{z}} \int (k^f)^2 dk^f \int_{-1}^{1} d\mu^f \frac{V_r^f}{2(2\pi)^2} \left[\frac{\frac{H^f}{(D_A^f)^2} P^f_g(k^f, \mu^f) - \frac{H^s}{(D_A^s)^2} P^s_g(k^s, \mu^s)}{\frac{H^s}{(D_A^s)^2} P^s_g(k^s, \mu^s) + \frac{H^f}{(D_A^f)^2} \frac{V_r^f}{N}} \right]^2 . \tag{192}$$

The prefactor $f_{\rm AP}$ of Eq. 164 has been written explicitly so every power spectrum can be evaluated in the same cosmology that was used to produce it. To compute the integral, k^s and μ^s must be expressed in terms of k^f and μ^f as described in Eqs. (166, 167).

Note that in this formula we correct some small inaccuracies present in the previous work by Ref. [118]. In that reference, the pre-factor V_r^f was incorrectly replaced by $V_r^s(1+\bar{z})^{-3}$, the second argument of P_g^s was approximated as μ^f instead of μ^s, and the volume in the last term of the denominator was V_r^s instead of V_r^f. We checked explicitly that these inaccuracies led to slightly over-conservative error forecasts in Ref. [118].

3.2 COSMIC SHEAR

3.2.1 *Angular power spectrum*

A cosmic shear survey maps the alignments of galaxies induced by weak gravitational lensing caused by large scale structures along the line of sight. The cosmological information is extracted from auto- and cross-correlations of alignment maps at different redshifts.

The projected shear power spectrum of the redshift bins i and j at multipoles ℓ can be inferred from the three-dimensional matter power spectrum via

$$C_\ell^{ij} = \frac{9}{16}\Omega_m^2 H_0^4 \int_0^\infty \frac{dr}{r^2} g_i(r) g_j(r) P\left(k=\frac{\ell}{r}, z(r)\right) \ . \tag{193}$$

The functions $g_i(r)$ depend on the radial distribution of galaxies in the redshift bin i, i.e. on the convolution of the distribution of detected galaxies with the corresponding redshift errors,

$$g_i(r) = 2r(1+z(r)) \int_r^\infty dr' \frac{\eta_i(r')(r'-r)}{r'} \ , \tag{194}$$

$$\eta_i(r) = H(r) n_i(z(r)) \ , \tag{195}$$

$$n_i(z) = \frac{D_i(z)}{\int_0^\infty D_i(z')dz'} \ , \tag{196}$$

$$D_i(z) = \int_{z_i^{\rm min}}^{z_i^{\rm max}} \mathcal{P}(z,z') \frac{dn_{\rm gal}}{dz}(z')\, dz' \ . \tag{197}$$

Due to the intrinsic alignment of galaxies, there is also a noise contribution N_ℓ. The noise spectrum added to the theoretical C_ℓ^{ij} is

$$N_\ell^{ij} = \delta_{ij}\sigma_{\rm shear}^2 n_i^{-1} \ , \tag{198}$$

where $\sigma_{\rm shear}$ is the root mean square of the galaxy intrinsic ellipticity and is set to 0.3, and n_i is the number of galaxies per steradian in the i'th redshift bin. We divide the redshift range into ten redshift bins with equal number of galaxies. Therefore, for every redshift bin, we have

$$n_i = \frac{n_{\rm gal}}{10} \times 3600 \left(\frac{180}{\pi}\right)^2 \ . \tag{199}$$

Table 7: Sky coverage and cosmic shear specifications for Euclid (see [118]) and for SKA (see [136]). Here $n_{\rm gal}$ is in units of arcmin^{-2}.

Experiment	$f_{\rm sky}$	$n_{\rm gal}$	z_m	α	β	γ	$f_{\text{spec-z}}$	$z_{\text{spec-max}}$	$\sigma_{\text{photo-z}}$	$z_{\text{photo-max}}$	$\sigma_{\text{no-z}}$
SKA1	0.1212	2.7	1.1	$\sqrt{2}$	2	1.25	0.15	0.6	0.05	2.0	0.3
SKA2	0.7272	10	1.3	$\sqrt{2}$	2	1.25	0.5	2.0	0.03	2.0	0.3
Euclid	0.3636	30	0.9	$\sqrt{2}$	2	1.5	0.0	0.0	0.05	4.0	0.3

3.2.2 *Euclid and SKA specifications*

The number density of sources and the corresponding redshift errors for Euclid and SKA are taken, respectively, from Ref. [118] and from Ref. [136][4]. The unnormalized redshift number density distribution is given by

$$\frac{dn_{\rm gal}}{dz} = z^\beta \exp\left[-\left(\frac{z}{\alpha z_m}\right)^\gamma\right] . \tag{200}$$

The redshift uncertainty is parameterized as follows,

$$\mathcal{P}(z,z') = \begin{cases} \dfrac{1-f_{\text{spec-z}}}{\sqrt{2\pi}\sigma_{\text{photo-z}}} \exp\left[-\dfrac{(z-z')^2}{2\sigma^2_{\text{photo-z}}(1+z)^2}\right] + f_{\text{spec-z}}\delta(z-z') & , z \le z_{\text{spec-max}} \\ \dfrac{1}{\sqrt{2\pi}\sigma_{\text{photo-z}}} \exp\left[-\dfrac{(z-z')^2}{2\sigma^2_{\text{photo-z}}(1+z)^2}\right] & , z \le z_{\text{photo-max}} \\ \dfrac{1}{\sqrt{2\pi}\sigma_{\text{no-z}}} \exp\left[-\dfrac{(z-z')^2}{2\sigma^2_{\text{no-z}}(1+z)^2}\right] & , z \ge z_{\text{photo-max}} \, , \end{cases} \tag{201}$$

where the measured redshift is denoted as z', while the true one is denoted as z. Some of these Gaussians span a big redshift range, but they are anyway multiplied by the distribution $\frac{dn_{\rm gal}}{dz}$ that is almost zero outside the range of interest.

The number counts of Euclid are negligible above $z = 3.5$, so here the error function becomes

$$\mathcal{P}(z,z') = \frac{1}{\sqrt{2\pi}\sigma_{\text{photo-z}}} \exp\left[-\frac{(z-z')^2}{2\sigma^2_{\text{photo-z}}(1+z)^2}\right] . \tag{202}$$

The sky coverage $f_{\rm sky}$ is the same as in the case of galaxy clustering (table 7).

3.2.3 *Cosmic shear likelihood*

The likelihood for lensing surveys is taken from Ref. [118],

$$-2\ln\mathcal{L} \equiv \sum_l (2l+1) f_{\rm sky}\left(\frac{d_l^{\rm mix}}{d_l^{\rm th}} + \ln\frac{d_l^{\rm th}}{d_l^{\rm obs}} - N\right) , \tag{203}$$

4 Ref. [136] provides values for Euclid which differ only in the value of $\sigma_{\text{photo-z}}$

where N is the number of redshift bins, which is equal to the dimension of the C_l-matrices whose determinants are denoted with d. The determinant of these $N \times N$ symmetric matrices can be defined as

$$d_l^{\text{th}} = \det\left(C_l^{\text{th}\,ij} + N_l^{ij}\right) , \tag{204}$$

$$d_l^{\text{obs}} = \det\left(C_l^{\text{fiducial}\,ij} + N_l^{ij}\right) , \tag{205}$$

$$d_l^{\text{mix}} = \sum_k \det\left(N_l^{ij} + \begin{cases} C_l^{\text{th}\,ij} & , j \neq k \\ C_l^{\text{fiducial}\,ij} & , j = k \end{cases}\right) . \tag{206}$$

Writing the theoretical angular power spectrum as C_l and the observational one as $\hat{C}_l$, both including noise, eq. (203) can be expressed as a multivariate Gaussian,

$$\mathcal{L} = \mathcal{N} \prod_{l,m} \left\{ \frac{1}{\sqrt{\det C_l}} \exp\left[-\frac{1}{2} \sum_{i,j} a_{lm}^{i*} (C_l^{-1})^{ij} a_{lm}^{j} \right] \right\} . \tag{207}$$

The observed angular power spectrum is defined as

$$\hat{C}_l^{ij} = \frac{1}{2l+1} \sum_{m=-l}^{l} a_{lm}^{i*} a_{lm}^{j} . \tag{208}$$

The inverse of a matrix A can be replaced with $A^{-1} = \text{adj}(A)/\det(A)$, where $\text{adj}(A)$ is the adjugate of A, i.e. the transpose of the cofactor matrix of A. The likelihood can then be rewritten by executing the sum over m,

$$\mathcal{L} = \mathcal{N} \prod_{l} \left\{ \left(\frac{1}{\sqrt{\det C_l}} \right)^{2l+1} \exp\left[-\frac{1}{2} \sum_{i,j} (2l+1) \hat{C}_l^{ij} \frac{(\text{adj}\, C_l)^{ij}}{\det C_l} \right] \right\} . \tag{209}$$

The remaining sum over i and j yields exactly d_l^{mix}. This can best be explained graphically,

$$\sum_{i,j} \hat{C}_l^{ij} (\text{adj}\, C_l)^{ij} = \sum_{i,j} \hat{C}_l^{ij} \begin{vmatrix} \dots & 0 & \dots \\ 0 & \hat{C}_l^{ij} & 0 \\ \dots & 0 & \dots \end{vmatrix} = \sum_j \begin{vmatrix} \dots & \hat{C}_l^{i_1 j} & \dots \\ \dots & \hat{C}_l^{i_2 j} & \dots \\ \dots & \dots & \dots \end{vmatrix} = d_l^{\text{mix}} . \tag{210}$$

We can now compute χ^2,

$$\chi^2 = -2 \ln \mathcal{L} = -2 \ln \mathcal{N} + \sum_l (2l+1) \left(\ln d_l^{\text{th}} + \frac{d_l^{\text{mix}}}{d_l^{\text{th}}} \right) . \tag{211}$$

Our best fit model, i.e. where the theoretical model gives the same as what is observed, should yield an effective χ^2 of zero. This means we should subtract the zero point from Eq. 211, where the zero point is given by Eq. 211 with $d_l^{\text{th}} = d_l^{\text{obs}}$,

$$\chi_0^2 = -2 \ln \mathcal{L} = -2 \ln \mathcal{N} + \sum_l (2l+1) \left(\ln d_l^{\text{obs}} + \frac{N \times d_l^{\text{obs}}}{d_l^{\text{obs}}} \right) . \tag{212}$$

Finally, after introducing an approximative correction for incomplete sky coverage [118], we obtain Eq. 203.

3.3 21CM INTENSITY MAPPING

3.3.1 *21cm power spectrum*

The goal of 21cm intensity mapping experiments is to measure the differential brightness temperature ΔT_b, defined as the difference between the observed brightness temperature T_b and the one expected for CMB photons only, T_γ.

The signal coming from 21cm hyperfine transitions of neutral hydrogen (HI) atoms is emitted at the frequency $\nu_0 = 1420.4057\,\mathrm{MHz}$, so the frequency measured today can be directly related to the redshift,

$$\Delta T_\mathrm{b} \equiv \frac{T_\mathrm{b}(z) - T_\gamma(z)}{1+z} \,. \tag{213}$$

Here, we focus only on the low redshift signal coming from the neutral hydrogen inside galaxies. A detailed description of cosmology with the 21cm-signal at high redshifts, both from reionization and from the cosmic dawn and dark ages, and how it can be used to probe modified gravity, can be found in Refs. [137–141]. At low frequencies, the mean differential brightness temperature is given by

$$\overline{\Delta T_b} \simeq 189 \left[\frac{H_0(1+z)^2}{H(z)}\right] \Omega_{\mathrm{HI}(z)}\, h\,\mathrm{mK} \,, \tag{214}$$

where H_0 is the Hubble constant $H_0 = h \times 100\,\mathrm{km/(s\,Mpc)}$ and $\Omega_\mathrm{HI}(z) = \rho_\mathrm{HI}(z)/\rho_c$ is the mass density of neutral hydrogen divided by the critical density of the present-day universe.

Let us first understand where this relation comes from. The total brightness temperature at redshift z is given by the background radiation field's temperature, with some fraction of it that is absorbed and re-emitted due to 21cm hyperfine transitions in neutral hydrogen atoms. The properties of HI in absorption and emission are described by the spin temperature T_S and the optical depth τ through

$$T_\mathrm{b} = T_S(1 - e^{-\tau}) + T_\gamma e^{-\tau} \,. \tag{215}$$

Due to the low probability of a 21cm transition, the optical depth is typically small. The differential brightness temperature can therefore be written linearly in τ,

$$\Delta T_\mathrm{b} = \frac{T_S - T_\gamma}{1+z}\left(1 - e^{-\tau}\right) \approx \frac{T_S - T_\gamma}{1+z}\tau \,. \tag{216}$$

In order to compxute τ, the absorption coefficient α has to be determined through the equation of radiative transfer

$$\frac{dI}{ds} = -\alpha I + j \,, \tag{217}$$

where s is the radial distance (in physical units) and I is the specific intensity, which is the energy flux per frequency and solid angle. Its radial derivative is given by

$$\frac{dI}{ds} = E_{10}\frac{\phi(\nu)}{4\pi}\frac{dn_0}{dt} \,. \tag{218}$$

Each atom falling from the exited state 1 into the ground state 0 emits a photon of energy E_{10}. The radial derivative of the energy flux is hence proportional to the time derivative of the number of atoms in the ground state per unit of physical volume, i.e. the number density n_0. Under the assumption of isotropy, the derivative with respect to solid angle becomes a factor

of $1/4\pi$. Due to line broadening, a single measured frequency corresponds to a small band of emitted frequencies described by the line profile $\phi(\nu)$, which is normalized to $\int \phi(\nu)d\nu = 1$.

In terms of Einstein coefficients, the time derivative of the number density can be written as $dn_0/dt = -n_0 B_{01} I + n_1 B_{10} I + n_1 A_{10}$. In the next steps, we use natural units where $\hbar = c = k_B = 1$, as well as the general relations $A_{10} = 4\pi\nu_0^3 B_{10}$ and $g_0 B_{01} = g_1 B_{10}$. For 21cm hyperfine transitions, the statistical weights are $g_0 = 1$ and $g_1 = 3$ and one gets $T_S \gg E_{10}$. Therefore, one gets the simplifications

$$\frac{n_1}{n_0} = \frac{g_1}{g_0}\exp(-\frac{E_{10}}{T_S}) \approx 3(1 - \frac{E_{10}}{T_S}) \,, \tag{219}$$

and $n_{\rm HI} \equiv n_0 + n_1 \simeq 4n_0 \simeq \frac{4}{3}n_1$. Put together, one gets

$$\frac{dn_0}{dt} = -\frac{A_{10}}{4\pi\nu_0^3}\frac{3}{4}n_{\rm HI}\frac{E_{10}}{T_S}I + \frac{3}{4}n_{\rm HI}A_{10} \,, \tag{220}$$

which gives an expression for α,

$$\alpha = \frac{3A_{10}}{16T_S}\frac{\phi(\nu)}{\nu_0}n_{\rm HI} \,. \tag{221}$$

The line profile will be described by the simple model of a constant distribution over some range $\delta\nu$, corresponding to a small Doppler shift caused by constant velocity dispersion $\frac{dv}{ds}$ over a region of HI of radial extend δs,

$$\phi(\nu) = \frac{1}{\delta\nu} = \frac{1}{\frac{dv}{ds}\delta s \cdot \nu_0} \,. \tag{222}$$

Averaged over big volumes, the approximation of a constant Hubble flow $\frac{dv}{ds} = H(z)$ can be used. The optical depth is then given by

$$\tau \equiv \int_{\delta s} \alpha ds = \frac{3A_{10}}{16\nu_0^2 T_S}\frac{1}{H(z)}n_{\rm HI} \,. \tag{223}$$

The number density of neutral hydrogen can be written as its background value plus a perturbation in the HI density field,

$$n_{\rm HI} = \frac{(1+z)^3}{m_{\rm H}}\frac{3H_0^2}{8\pi G}\Omega_{\rm HI}(z)(1+\delta_{\rm HI}) \,. \tag{224}$$

With all constants written explicitly, the differential brightness temperature is given by

$$\Delta T_b = \frac{3A_{10}}{16\nu_0^2}\frac{3H_0}{8\pi G h m_{\rm H}}\frac{\hbar c^3}{k_B}\left(\frac{H_0(1+z)^2}{H(z)}\right)\Omega_{\rm HI}(z)(1+\delta_{\rm HI})h\left(1 - \frac{T_\gamma}{T_S}\right) \,. \tag{225}$$

The last term can be neglected because $T_S \gg T_\gamma$ inside galaxies, whereby we arrive at Eq. 214 when we also replace the constants in front with a numerical prefactor. Modelling the differential brightness temperature in this way was also done by e.g. Refs. [126, 142, 143].

Deviations from this value are proportional to the density perturbations in neutral hydrogen, which can be related to the dark matter density perturbations via

$$\Delta T_b - \overline{\Delta T_b} = \overline{\Delta T_b}\delta_{\rm HI} = \overline{\Delta T_b} b_{\rm HI}\delta_m \,, \tag{226}$$

where, when computing the power spectrum of fluctuations in the differential brightness temperature, it is convenient to neglect the local fluctuations of $H(z)$. As a consequence, the power spectrum P_{21} is proportional to the matter power spectrum as

$$P_{21} = b_{21}^2 P_m \,, \tag{227}$$

with $b_{21} \equiv \overline{\Delta T_b} b_{\rm HI}$. The redshift dependence of $\Omega_{\rm HI}(z)$ and $b_{\rm HI}(z)$ are modelled following Ref. [144] as

$$\Omega_{\rm HI}(z) = \Omega_{\rm HI,o}(1+z)^{\alpha_{\rm HI}} \tag{228}$$

$$b_{\rm HI}(z) = 0.904 + 0.135(1+z)^{1.696} \,, \tag{229}$$

with $\Omega_{\rm HI,o}$ and $\alpha_{\rm HI}$ set to the fiducial values 4×10^{-4} and 0.6, respectively, and allowed to vary in the forecast. As in the case of galaxy clustering, we use nuisance parameters to describe the future accuracy of bias modeling,

$$b_{\rm HI}(z) = \beta_0^{\rm IM} \left[0.904 + 0.135(1+z)^{1.696\beta_1^{\rm IM}} \right] \,, \tag{230}$$

with mean value zero and a prior corresponding to a rms of 0.025.

Additionally, we have to consider observational effects analogous to those of the galaxy power spectrum of Eq. 164, i.e.

$$P_{21}(k,\mu,z) = f_{\rm AP}(z) \times f_{\rm res}(k,\mu,z) \times f_{\rm RSD}(\hat{k},\hat{\mu},z) \times b_{21}^2(z) \times P_m(\hat{k},z) \,. \tag{231}$$

The prefactors are the same as in the case of galaxy surveys, because the signal dominantly consists of radiation originating from galaxies. As such, it is affected by red/blue-shifting according to the movement of the galaxies. A power spectrum reconstructed from this map of intensities therefore suffers from redshift-space distortions, limited resolution, and the Alcock-Paczinsky effect in the same way as one reconstructed from a map of galaxy positions.

By considering $\delta(\boldsymbol{k})$ as a set of independent Gaussian random realizations, the variance is simply given by the power spectrum squared. In terms of independent modes of a finite volume survey, it is corrected by the volume of a single independent mode $(2\pi)^3/V_r$ and by the averaging volume V_k which determines the grid of sampled modes,

$$\sigma_P^2(\boldsymbol{k}) = \frac{(2\pi)^3}{V_k V_r} P^2(\boldsymbol{k}) \,. \tag{232}$$

So, the same formalism as described in section 3.1.4 can be used. The observed power spectrum also includes a noise spectrum,

$$P_{21}^{\rm obs}(k,\mu,z) = P_{21}(k,\mu,z) + P_N(z) \,, \tag{233}$$

where $P_N(z)$ will be described in the next subsection. Note that, although the noise power may mathematically take the role of the shot noise of galaxy power spectra, it is part of the power spectrum itself, whereas shot noise is only an artefact caused by the discrete nature of the signal, in contrast to the theoretical distribution used to describe it. In fact, there is also a shot noise in the case of intensity mapping, since the signal still originates from discretely spaced galaxies. Nevertheless, it is negligible (see e.g. [145]), because of the huge number of observed galaxies when no selection process is reducing their number.

By identifying galaxies, all unwanted contributions to the power spectrum except for the shot noise could be removed. Instead, the correlations of unprocessed intensity include correlations in the foregrounds and in random noise in the sky or in the experimental setup itself. Hence, the biggest disadvantage of 21cm intensity mapping surveys, in comparison to 21cm galaxy surveys, is the high contamination of the signal with telescope noise and foreground signals. If the latter are sufficiently smooth in frequency, they are nevertheless removable.

3.3.2 *SKA specifications*

We start by describing the noise power spectrum P_N in Eq. 231. We are here considering a survey executed in *single dish* mode; this enhances the speed of the survey, but dismisses the advantages of radio interferometry. In this case, the noise power is given by

$$P_N(z) = T_{\rm sys}^2 \frac{4\pi f_{\rm sky} r^2(z)\,(1+z)^2}{2H(z) t_{\rm tot} \nu_0 N_{\rm dish}}\,, \tag{234}$$

where $T_{\rm sys}$ is the system temperature, $t_{\rm tot}$ is the total observation time and $N_{\rm dish}$ is the number of dishes. The noise power originates from random uncorrelated fluctuations in the intensity of single pixels. Their amplitude is given by $T_{\rm sys}$. Since this noise is independent of the signal it can be described as an additional perturbation field whose power spectrum P_N is added to the power spectrum of the signal. We adopt $t_{\rm tot} = 10000\,{\rm h}$ and $N_{\rm dish} = 200$, i.e. the same noise power as in Ref. [144], where one can also find a derivation of Eq. 234.

Since there is no need to resolve a single galaxy, SKA1 has access to signals from higher redshift. Therefore, band 1 can also be used for intensity mapping. Following Ref. [145], the system temperature is defined as the sum of the instrument's temperature $T_{\rm inst}$ and the sky temperature,

$$T_{\rm sky} = 20\,{\rm K}\left(\frac{408\,{\rm MHz}}{\nu}\right)^{2.75}\,. \tag{235}$$

The Gaussian suppressions of the power spectrum are quantified using the relation between the full width at half maximum to the rms, given by ${\rm FWHM} = \sqrt{8\ln 2}\,\sigma$. In the case of frequency, the channel width due to band separation into 64,000 channels is used as FWHM,

$$\sigma_\theta = \frac{1}{\sqrt{8\ln 2}}\frac{\lambda_0}{D}(1+z)\,, \tag{236}$$

$$\sigma_\nu = \frac{\delta\nu}{\sqrt{8\ln 2}}\,. \tag{237}$$

In single dish mode, the angular resolution is determined by the diameter $D = 15$ m of a single dish. As a result, σ_θ is as big as $0.34^\circ(1+z)$.

As we already mentioned, foregrounds are expected to be much larger than the 21cm signal itself. Nevertheless, foregrounds are expected to be sufficiently spectrally smooth to be removed. Here, correlations between remnants of the foregrounds or artefacts of their removal are expected to be negligible. Yet, the effect of foreground removal is taken into account in two ways. First, the observed part of the sky is decreased to the regions with the lowest foreground intensity. Following Ref. [144], the probed fraction of the sky is reduced to $f_{\rm sky} = 0.58$. Second, as discussed in Ref. [146], foreground removal does not work close to the edges of the frequency band (i.e. $\lesssim 50$ MHz). Therefore, we reduce the redshift range to exclude information from the edges.

The specifications used for intensity mapping forecasts are listed in Table 8.

Table 8: IM specifications. [131, 145]

parameter	$\nu_{\min}$ [MHz]	$\nu_{\max}$ [MHz]	$z_{\min}$	$z_{\max}$	$\delta\nu$ [kHz]	T_{inst} [K]
SKA1 band 1	∼400 (350)	∼1000 (1050)	0.45	2.65	10.9	23
SKA1 band 2	∼1000 (950)	1421 (1760)	0.05	0.45	12.7	15.5

3.4 NON-LINEAR THEORETICAL UNCERTAINTY

Euclid and the SKA will survey a large sky volume and detect a huge number of galaxies. This will dramatically decrease the size of sampling variance and shot noise compared to current surveys. Therefore, on small scales, theoretical errors will be the leading source of uncertainty and the limiting factor for parameter extraction, at least for analyses based on three-dimensional power spectra, such as galaxy surveys and intensity mapping. Here, we describe our strategy for modelling the theoretical error. Since it is easier to deal with the theoretical error of the bi-dimensional angular power spectrum, we start by discussing the weak lensing case.

3.4.1 *Cosmic shear error modelling*

The simplest way to model the theoretical uncertainty is to introduce a cutoff. This means neglecting all theoretical uncertainties up to a wavenumber k_{NL}, while dismissing all information above that wavenumber. This scheme is a good approximation when the result does not depend strongly on the region where the uncertainty increases from almost zero to infinity. Since non-linear effects increase with time, the cutoff scale should then decrease with redshift. Following Ref. [147], the redshift dependence of non-linear effects can be parametrised as

$$k_{\text{NL}}(z) = k_{\text{NL}}(z=0) \cdot (1+z)^{2/(2+n_s)} \,. \tag{238}$$

The quantity of interest for weak lensing surveys is the shear power spectrum C_ℓ, which is given by Eq. 193 as a weighted integral of $P(k)$ convoluted with a window function spanning a large range in k. Therefore, there is no simple equivalent ℓ_{NL} of k_{NL}. Our approach consists in identifying values of ℓ above which most information comes from wavenumbers $k > k_{\text{NL}}$. First, we find the value r corresponding to the maximum of the product of the window functions for a pair (i, j) of redshift bins,

$$r^{ij}_{\text{peak}} = \frac{\int_0^\infty \frac{dr \cdot r}{r^2} g_i(r) g_j(r)}{\int_0^\infty \frac{dr}{r^2} g_i(r) g_j(r)} \,. \tag{239}$$

This value mostly depends on the lower redshift bin of i and j, so an average over higher bins can be performed to get $\bar{r}^i_{\text{peak}} \equiv (\sum_{j>i} r^{ij}_{\text{peak}})/(N-i)$ where N is the number of bins. This can be related to a maximum ℓ through

$$\ell^i_{\max} = k_{\text{NL}}(z) \cdot \bar{r}^i_{\text{peak}} \,. \tag{240}$$

All C^{ij}_ℓ at ℓ larger than $\ell^i_{\max}$ or $\ell^j_{\max}$ are discarded. The resulting C_ℓ's are still quadratic.

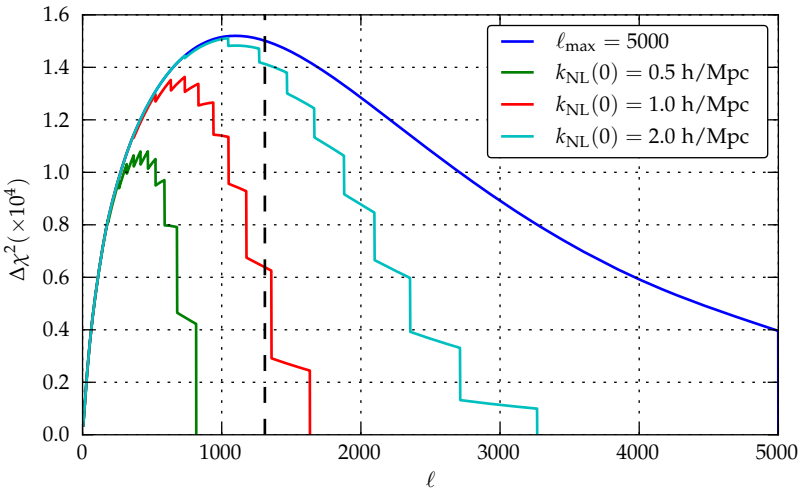

Figure 2: Euclid cosmic shear combined with Planck (see section 3.5 for details): sensitivity to a 0.1% variation of $P(k)$ for different cut-off wavenumbers (always scaled with redshift). The flat $\ell_{max} = 5000$ cut-off (blue) shows the amount of information available in absence of a cut-off. The second (green) and third (red) cases are more conservative than a sharp cut-off at $\ell = 1310$ would be. For comparison, the dashed line marks $\ell = 1310$, corresponding to the $\ell_{\rm max}$ used by the KiDS collaboration in Ref. [148] as a reasonable cut-off producing stable results. The last case (cyan) is a little more constraining than this sharp cut-off, intended to reflect improvements in non-linear modeling in the analysis of future data. For our analysis we will use $k_{\rm NL}(0) = 0.5\ h/\text{Mpc}$ (conservative) and $k_{\rm NL}(0) = 2.0\ h/\text{Mpc}$ (realistic) as our non-linear cut-off wavenumbers. The corresponding 1-σ sensitivity of our MCMC forecasts can be seen in table 9.

One way to better understand the likelihood consists in splitting it in contributions from each ℓ. The likelihood for cosmic shear can be expanded as

$$-2\ln\mathcal{L} = \sum_{\ell} \Delta\chi^2_\ell(\Delta P)\ , \tag{241}$$

where δP is the difference between the fiducial and sampled power spectrum. To understand the weight of each multipole, we can plot $\chi^2_\ell(\delta P)$ versus ℓ while assuming that the fiducial and sampled power spectra differ by the same relative factor $\Delta P = \Delta P(\boldsymbol{k}, \bar{z}) = P - \hat{P} = 0.001P$ for every redshift and wavenumber. The resulting contributions $\Delta\chi^2_\ell$ solely depend on the characteristics of the likelihood.

In Figure 2, we see the $\Delta\chi^2_\ell$ contributions to the Euclid cosmic shear likelihood for different choices of $k_{NL}(0)$. A comparison of forecasts for Planck + Euclid cosmic shear for these values of $k_{\rm NL}(0)$ is shown in Table 9. We see that the sensitivity does not differ by a large amount despite great changes in the non-linear cut-off, with only n_s and M_ν showing non-negligible improvement in sensitivity with increasing cut-off values. Since the results do not depend strongly on the choice of $k_{\rm NL}$, the cut-off approximation is accurate enough.

For our analysis, we will adopt two values: a "conservative" cut-off $k_{\rm NL}(0) = 0.5\ h/\text{Mpc}$, and a "realistic" cut-off $k_{\rm NL}(0) = 2.0\ h/\text{Mpc}$. The realistic case is supposed to reflect improvements in the modelling of non-linear scales in the analysis of future data. Previous analyses, like that of Ref. [148] with a sharp bin-independent cut-off at $l = 1310$, used an amount of information somewhere "in between" our conservative and realistic assumptions.

In Figure 3, we show the sensitivity distribution for Euclid, SKA1 and SKA2 for the realistic (left) and conservative (right) non-linear cut-offs. We see that SKA1 is not competitive and

Table 9: Planck (see section 3.5) plus Euclid cosmic shear 1-σ sensitivity (normalized by corresponding Planck-only values) of MCMC forecasts for the non-linear cut-off values used in Figure 2. We see that most sensitivities do not depend strongly on the choice of a given $k_{\rm NL}(0)$. Only n_s and M_ν show a non-negligible improvement in sensitivity, despite the large changes in the cut-off. Therefore, we find that the non-linear cut-off scheme is appropriate for our analysis.

$k_{\rm max}$	$100\omega_b$	$\omega_{\rm cdm}$	θ_s	$\ln(10^{10}A_s)$	n_s	$\tau_{\rm reio}$	M_ν [eV]
0.5 h/Mpc	0.77	0.27	0.97	0.94	0.72	0.96	0.50
1.0 h/Mpc	0.76	0.27	0.94	0.95	0.70	0.98	0.41
2.0 h/Mpc	0.76	0.25	0.97	0.94	0.65	0.97	0.36
$l_{\rm max} = 5000$	0.74	0.24	0.94	0.94	0.58	0.96	0.30
Planck only	1.00	1.00	1.00	1.00	1.00	1.00	1.00

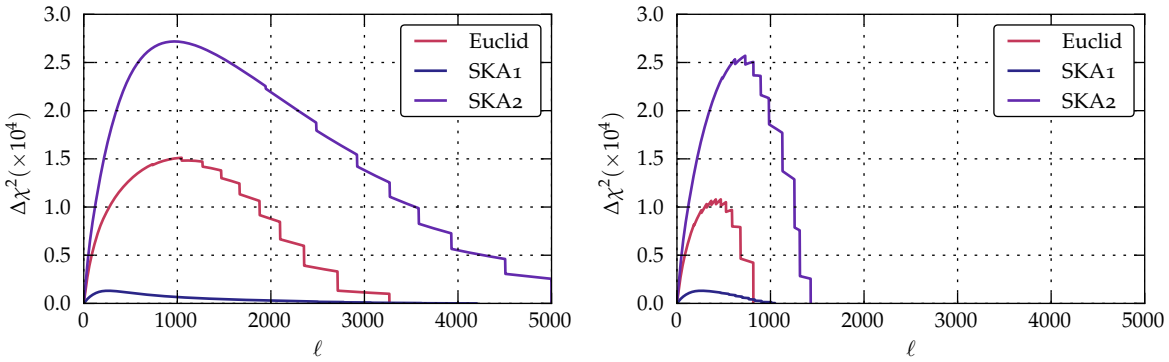

Figure 3: Sensitivity distribution for all cosmic shear likelihoods. The left panel shows the realistic approach and the right panel the conservative one. The $\Delta\chi^2$-values are contributions for each multipole l obtained by setting $\Delta P = 0.001P$ for all k. We find that SKA1 is not competitive, but that SKA2 will out-perform Euclid.

that SKA2 will be more constraining than Euclid, because of the better accuracy of redshift measurements and of the greater sky coverage.

3.4.2 *Power spectrum error modelling (galaxy clustering + intensity mapping)*

For the case of the three-dimensional galaxy power spectrum $P_g(\boldsymbol{k}, z) = P_g(k, \mu, z)$, things are a bit more complicated. After binning in z-space, instead of dealing with a discrete expansion parameter ℓ in each bin of mean redshift $\bar{z}$, we have two continuous variables (k, μ). The traditional way to build a likelihood is recalled in section 3.1.4. The contribution of one redshift bin and of the interval $(k \pm \frac{dk}{2}, \mu \pm \frac{d\mu}{2})$ to $-2\ln\mathcal{L}$ (i.e. to the χ^2) can be written in a differential form,

$$\frac{d\chi^2}{dkd\mu} = k^2 \frac{V_r(\bar{z})}{2(2\pi)^2} \times \left[\frac{\Delta P_g(k, \mu, \bar{z})}{\sigma_{\rm obs}(k, \mu, \bar{z})}\right]^2 , \tag{242}$$

where $\Delta P_g(k, \mu, \bar{z})$ is the difference between the predicted and observed galaxy power spectrum, the prefactor proportional to k^2 accounts for the density of independent Fourier modes, and $V_r(\bar{z})$ is the volume of one redshift bin given by Eq. 172. The observational error is given by $\sigma_{\rm obs}(k, \mu, \bar{z}) = P_g(k, \mu, \bar{z}) + P_N$, where P_N is some constant noise. If we want to understand

how the experimental sensitivity depends on different scales, it is worth looking at the effective error,

$$\frac{d\chi^2}{dkd\mu} \equiv \left[\frac{\Delta P_g(k,\mu,\bar{z})}{\sigma_{\rm eff}(k,\mu,\bar{z})}\right]^2 \Rightarrow \qquad \sigma_{\rm eff}(k,\mu,\bar{z}) = \sigma_{\rm obs}(k,\mu,\bar{z}) \left[k^2 \frac{V_r(\bar{z})}{2(2\pi)^2}\right]^{-1/2} . \tag{243}$$

The power spectrum decreases when k increases, such that the effective error decreases approximately like $\propto k^{-2}$. This in turn means that the amount of accessible information grows to infinity. If a cut-off $k_{\rm NL}$ is used to prevent this, the region directly below this cut-off will be the one with the biggest weight in the likelihood, making the results very sensitive to the choice of $k_{\rm NL}$. Thus, a more realistic way to account for the theoretical error is needed.

We first review the approach of Ref. [118] to this problem. It starts from the assumption that for each (k,μ,z), we can reasonably estimate the 1σ uncertainty $\delta P_g(k,\mu,z)$ on the theoretical prediction for the galaxy power spectrum $P_g(k,\mu,z)$ (e.g. by comparing the output of various simulations) and define a 1σ envelope function α for the relative error,

$$\alpha(k,\mu,z) = \delta P_g(k,\mu,z)/P_g(k,\mu,z) . \tag{244}$$

The concrete implementation of this error in the likelihood is not trivial, because the errors made on $P_g(k,\mu,z)$ at different values of (k,μ,z) should in principle be correlated. If we assume that the whole (k,μ,z) volume probed by the experiment can be split in bins in which the errors are uncorrelated, we can introduce one independent nuisance parameter per bin, and marginalise over it. This approach is actually numerically expensive, but in good approximation, the marginalisation can be replaced by an analytic minimisation. This leads to a simple expression for the contribution of each uncorrelated bin $(k \pm \frac{dk}{2}, \mu \pm \frac{d\mu}{2}, \bar{z} \pm \frac{\Delta\bar{z}}{2})$:

$$\frac{d\chi^2}{dkd\mu} = k^2 \frac{V_r(\bar{z})}{2(2\pi)^2} \times \frac{\left[\Delta P_g(k,\mu,\bar{z})\right]^2}{\sigma^2_{\rm obs}(k,\mu,\bar{z}) + \sigma^2_{\rm th}(k,\mu,\bar{z})} \tag{245}$$

where $\sigma_{\rm th}(k,\mu,z) = \alpha(k,\mu,z)P_g(k,\mu,z)$. Note that the bin width in (k,μ)-space appears explicitly in the differential expression on the left-hand side, while the bin width in redshift space appears implicitly in the expression of $V_r(\bar{z})$ given in Equation (172). Equivalently, one can take an arbitrary binning in (k,μ,z)-space, provided that the theoretical error is rescaled self-consistently: this is the approach followed in Ref. [118].

The difficulty is then to evaluate the correlation length in (k,μ,z)-space. The authors of Ref. [118] chose a method that compares the effect of the theoretical error to a reference $\Delta\chi^2 = 1$, obtained by varying a single Gaussian nuisance parameter. However, this method makes the amplitude of the error dependent on the range of the integrals $[k_{\rm min}, k_{\rm max}]$, and on the number of redshift bins. As a result, the error depends on the survey specifications and cannot be used in our combined forecast of the future sensitivity of various experiments. It is therefore necessary to take a closer look at the correlation of the theoretical error.

The authors of Ref. [119] address this problem by introducing the full correlation matrix of the theoretical error. They write the contribution of each redshift bin to their log-likelihood as

$$\chi^2 = \sum_{i,j} \left[(\Delta P(k_i) - Q(k_i)) \frac{\delta_{ij}}{\sigma_{\rm obs}(k_i)\sigma_{\rm obs}(k_j)} (\Delta P(k_j) - Q(k_i)) + Q(k_i) C^{-1}_{ij} Q(k_j)\right] , \tag{246}$$

where the wavenumber range has been discretised arbitrarily, redshift-space distortions are neglected (thus quantitites do not depend on μ), $Q(k_i)$ is a single realization of the theoretical

error, and $C_{ij} = \langle Q(k_i)Q(k_j)\rangle$ is the error correlation matrix, for which one needs to make some assumption. Ref. [119] parametrises C_{ij} in terms of an error amplitude for each given k_i (equivalent to $\delta P(k_i) = \alpha(k_i)P(k_i)$ in our notation), and a correlation length Δk such that the correlation is exponentially suppressed for $|k_i - k_j| > \Delta k$,

$$C_{ij} = \alpha(k_i)P(k_i)\exp\left[\frac{-(k_i-k_j)^2}{2\Delta k^2}\right]\alpha(k_j)P(k_j)\ . \tag{247}$$

The forecasts can then be performed with a marginalisation over each nuisance parameter $Q(k_i)$. Compared to the previous method, this approach relies on one more assumption: one needs to postulate not only an error amplitude function $\alpha(k)$, but also a correlation length Δk, accounting for the minimum typical scale over which we allow the theoretical error to fluctuate randomly. This enlarged parametrisation compared to Ref. [118] is a good thing, because it makes more clear and explicit assumptions on the theoretical error. The authors of Ref. [119] argue that Δk can be matched with the Baryon Acoustic Oscillation (BAO) scale.

Unfortunately, there are several reasons why we cannot directly use this approach in our forecasts. First, the marginalisation over one nuisance parameter for each k_i is still tractable for the authors of Ref. [119] because their forecast is Fisher matrix based. In a full MCMC forecast like ours, it would introduce too many varying parameters and the convergence of the MCMC chains would be prohibitively slow. Additionally, Ref. [119] neglected RSD corrections and assumed no correlations between the error in different redshift bins. In the present work, we wish to incorporate RSD effects and to address the issue of error correlations in redshift space. In principle, this would imply a generalisation of Eq. 246 to one nuisance parameter for each argument $Q(k_i, \mu_j, z_k)$, with a six-dimensional correlation matrix. This problem is too heavy to be solved with an MCMC approach. Note that we cannot simply approximate the marginalisation over each Q by an analytic minimization, because there is no simple analytic solution in presence of an exponential correlation function.

Therefore, we chose to stick with the idea of Ref. [119] in formulating the problem in terms of correlation lengths, but, in order to make it computationally tractable, we assume that the (k, μ, z) space can be split into approximately uncorrelated bins. In other words, we replace the non-diagonal covariance matrix by a diagonal one with a bigger spacing $(\Delta k, \Delta\mu, \Delta z)$ between adjacent bin centers $(k_{i\pm1}, \mu_{j\pm1}, z_{k\pm1})$, such that the $Q(k_i, \mu_j, z_k)$ are statistically independent:

$$\chi^2 = \sum_{m,n}\left[\frac{(\Delta P_g(\mathbf{k}_m, \bar{z}_n) - Q(\mathbf{k}_m, \bar{z}_n))^2}{\sigma^2_{\rm obs}(\mathbf{k}_m, \bar{z}_n)}\right] + \sum_{i,j,k}\frac{Q^2(k_i, \mu_j, z_k)}{(\alpha(k_i, \mu_j, z_k)P_g(k_i, \mu_j, z_k))^2}\ . \tag{248}$$

As usual in the expression of a galaxy survey likelihood, the first sum runs over all independent Fourier modes $\mathbf{k}_m$ (see section 3.1.4 or seminal papers like Ref. [134]) and over redshift bins of mean $\bar{z}_n$ separated by $\Delta\bar{z}$. Instead, the second sum runs over the centers of the larger bins with uncorrelated theoretical errors, spaced by steps $(\Delta k, \Delta\mu, \Delta z)$ that play the role of correlation lengths. The function Q is assumed to be continuous, with a few nodes $Q(k_i, \mu_j, z_k)$ that are treated as nuisance parameters. In principle, the value of Q in an arbitrary point (k, μ, z), or equivalently $(\mathbf{k}, z)$, which appears in the first sum could be obtained by performing a smooth interpolation of the Q function between the node values. However, in practice, we do not need to perform any such interpolation in our numerical implementation: we shall see below that Q can be eliminated analytically and does not appear in our final expression (253).

It is conventional to assume that, in the first sum the volume of each independent mode is so small that the discrete sum can be represented by an integral (see section 3.1.4),

$$\sum_{m,n} \to \sum_{n} \int_{k_{\rm min}}^{k_{\rm max}} dk \cdot k^2 \int_{-1}^{1} d\mu \; \frac{V_r(\bar{z}_n)}{2(2\pi)^2} \; . \tag{249}$$

Actually, the second sum can also be replaced by an integral, provided that we rescale the steps of integration by the correlation lengths, in order to avoid counting the contribution of each independent nuisance parameter multiple times,

$$\sum_{i,j,k} \to \sum_{\bar{z}} \frac{\Delta\bar{z}}{\Delta z} \int_{k_{\rm min}}^{k_{\rm max}} \frac{dk}{\Delta k} \int_{-1}^{1} \frac{d\mu}{\Delta\mu} \; . \tag{250}$$

In other words, this is equivalent to increasing the number of nuisance parameters by some factor, while dividing their weight by the same factor in order to keep a fully equivalent expression. Then the χ^2 can be written with a single integral

$$\chi^2 = \sum_{n} \int_{k_{\rm min}}^{k_{\rm max}} dk \cdot k^2 \int_{-1}^{1} d\mu \; \frac{V_r(\bar{z}_n)}{2(2\pi)^2} \left[\frac{(\Delta P_g(k,\mu,\bar{z}_n) - Q(k,\mu,\bar{z}_n))^2}{\sigma^2_{\rm obs}(k,\mu,\bar{z}_n)} + \frac{Q^2(k,\mu,\bar{z}_n)}{\sigma^2_{\rm th}(k,\mu,\bar{z}_n)} \right] \; , \tag{251}$$

where we defined

$$\sigma_{\rm th}(k,\mu,z) = \left[\frac{V_r(z)}{2(2\pi)^2} k^2 \Delta k \Delta\mu \frac{\Delta z}{\Delta\bar{z}} \right]^{1/2} \alpha(k,\mu,z) P_g(k,\mu,z) \; . \tag{252}$$

Finally, we can approximate the marginalisation over nuisance parameters by an analytic minimisation[5] and obtain a computationally tractable expression,

$$\chi^2 = \sum_{n} \int_{k_{\rm min}}^{k_{\rm max}} dk \cdot k^2 \int_{-1}^{1} d\mu \; \frac{V_r(\bar{z}_n)}{2(2\pi)^2} \left[\frac{(\Delta P_g(k,\mu,\bar{z}_n))^2}{\sigma^2_{\rm obs}(k,\mu,\bar{z}_n) + \sigma^2_{\rm th}(k,\mu,\bar{z}_n)} \right] \; . \tag{253}$$

This expression differs from the usual likelihood derived in absence of a theoretical error only through the presence of the term $\sigma^2_{\rm th}$ in the denominator. The χ^2 defined in Eq. 253 has the same form as the one used in Ref. [118], but the new ingredient is the more rigorous definition of the quantity $\sigma_{\rm th}$ in Eq. 252, which comes from a precise discussion of the role of correlation lengths $(\Delta k, \Delta\mu, \Delta z)$. We must now specify this term and motivate some choices for the correlation lengths and for the error envelope function $\alpha(k,\mu,z)$. These choices should be guided by the types of errors which are expected to be made on theoretical predictions for non-linear corrections, and by the shape and amplitude of these errors.

The choice of $\Delta\mu$ relates to the question: for a given bin (k_i, z_k), how many independent nuisance parameters should describe the error for different μ_j values? A most reasonable answer is one. Indeed, the error made on the prediction of the non-linear power spectrum $P_m(\hat{k},z)$, and on the (possibly non-linear) bias $b(\hat{k},z)$, is isotropic in 3D Fourier space. When it propagates to $\alpha(k,z,\mu)$ (the relative error on $P_g(k,\mu,z)$), it only gives a small μ-dependence through the projection from $\hat{k}$ to k, with no further error introduced in this projection. On the other hand, the departure from our possibly too simplistic ansatz for RSD and fingers-of-God corrections, as well as instrumental resolution effects, could potentially be strongly μ-dependent and motivate the introduction of more than one independent nuisance parameter per (k_i, z_k) bin.

5 Like in Ref.[118], we perform the analytic minimization with a small approximation. Since $\sigma_{\rm obs}$ contains P_g, it should also contain a term Q added to it. This small dependence of the standard deviation on the theoretical error has an extremely small impact. We neglect it and stick to the definition $\sigma_{\rm obs} = P_g + P_N$.

However, there are some arguments in favor of neglecting the theoretical error from RSD modeling, compared to other theoretical error sources. They will be discussed later in the context of quantifying the total theoretical error. We conclude that, for a given bin (k_i, z_k), the errors made on $P_g(k, \mu, z)$ for different μ values can be considered as fully correlated with each other to very good approximation. Taking a single independent nuisance parameter per (k_i, z_k) bin is mathematically equivalent to setting $\Delta\mu = \mu_{\max} - \mu_{\min} = 2$ in Eq. 252. Therefore the theoretical error expression is reduced to

$$\sigma_{\rm th}(k, \mu, z) = \left[\frac{V_r(z)}{(2\pi)^2} k^2 \Delta k \frac{\Delta z}{\Delta \bar{z}}\right]^{1/2} \alpha(k, \mu, z) P_g(k, \mu, z) \; . \tag{254}$$

For the correlation length in wavenumber space we will use $\Delta k = 0.05\, h/\mathrm{Mpc}$, like in Ref. [119]. This is chosen due to being similar to the BAO scale, which is the smallest inherent scale in the matter power spectrum, and as a conservative guess for the correlation length in k-space.

The correlation length in redshift space is harder to guess. Theoretical errors are not necessarily correlated throughout the whole redshift range probed by the experiment. A value close to, but slightly smaller than, the total redshift range probed by the experiment (1.5 for Euclid, 1.9 for SKA2) should be a conservative guess. We assume $\Delta z = 1$, which is equivalent to assuming between one and two z-bins with independent theoretical errors for each k_i bin. Note that Ref. [119] did not discuss the issue of correlations in redshift space and used the χ^2 formula (248) in each redshift bin, each time with a new bunch of independent nuisance parameters $Q(k_i)$. This is equivalent to setting a correlation length Δz implicitly equal to the size of individual redshift bins $\Delta\bar{z}$. We do not adopt this approach, since the errors made by N-body simulations in the prediction of $P_m(\hat{k}, z)$ at a given scale $\hat{k}$, and for two nearby redshifts $\bar{z}_n$ and $\bar{z}_{n+1}$, should not be statistically independent. Moreover, the impact of the theoretical error should not directly depend on the number of bins in which one chooses to split the data.

The relative error envelope function $\alpha(k, \mu, z)$ should model uncertainties on three types of non-linear corrections: the prediction of the matter power spectrum itself, of the bias, and of redshift-space distortions. We will neglect the third contribution (i.e. uncertainties on the RSD correction term) for several reasons. First, the fingers-of-God correction term already leads to a strong suppression of the power spectrum, which results in a big relative observational error when the power spectrum becomes smaller than the noise power. At this point, any additional theoretical error becomes irrelevant. Furthermore, RSD corrections can be modelled up to higher order than used here [149] and their modelling is continuously improving (see e.g. Refs. [150, 151] and references therein). Therefore, we can focus on the theoretical error on the non-linear matter power spectrum and bias predictions. The bias is usually assumed to be linear up to scales $k < 0.2\,\mathrm{h/Mpc}$. Beyond, a non-linear treatment would be more realistic (see e.g. [108]), but non-linear bias can be predicted by future simulations up to some residual uncertainty. Hence, the theoretical uncertainty should account mainly for inaccuracies in matter power spectrum and bias predictions from simulations.

The `HALOFIT` semi-analytic formula [105, 152], which we use for the present forecasts, only reaches accuracies of 5% at $k < 1\,\mathrm{h/Mpc}$ and 10% at $k < 10\,\mathrm{h/Mpc}$, according to Ref. [105]. This error was estimated from a comparison with N-body simulations. The more recent HMcode [106] achieves better precision than `HALOFIT` for k values larger than the BAO scale ($k > 1\,\mathrm{h/Mpc}$), while the precision on BAO scales is a little worse. Overall, the effect is an error of 5% on all scales [106]. Finally, Ref. [104] found that present-day N-body codes (Ramses [153], Pkdgrav3 [154], Gadget3 [155]) agree to within 1% at $k = 1\,\mathrm{h/Mpc}$ and 3% at

$k = 10\,\mathrm{h/Mpc}$. Following the plots in Ref. [104], we assume an exponential growth of the uncertainty in the decimal logarithm of k crossing 0.33% at $k = 0.01\,\mathrm{h/Mpc}$. However, these are pure dark matter simulations and the effect of baryonic feedback, as well as the k-dependence of galaxy-to-mass bias on nonlinear scales, will increase the error. According to Ref. [109], the effect of baryonic feedback reaches the one percent level at $k = 0.3\,\mathrm{h/Mpc}$ and grows to a 30% suppression of the power spectrum at $k = 10\,\mathrm{h/Mpc}$. To account for uncertainties in the future modeling of baryonic feedback, which will of course be smaller than the effect itself, and to allow for remaining additional uncertainty from bias and RSD modeling at small scales, we increase the theoretical error to 1% at $k = 0.3\,\mathrm{h/Mpc}$ and 10% at $k = 10\,\mathrm{h/Mpc}$. Consequently, we define a relative error function passing through these three fixed points:

1. 0.33% error at $k = 0.01\,\mathrm{h/Mpc}$,
2. 1% error at $k = 0.3\,\mathrm{h/Mpc}$,
3. 10% error at $k = 10\,\mathrm{h/Mpc}$.

This can be achieved with the following ansatz,

$$\alpha(\boldsymbol{k}, z) = \begin{cases} a_1 \exp\left(c_1 \cdot \log_{10} \frac{k}{k_1(z)}\right) , & \frac{k}{k_1(z)} < 0.3 \\ a_2 \exp\left(c_2 \cdot \log_{10} \frac{k}{k_1(z)}\right) , & \frac{k}{k_1(z)} > 0.3 \,, \end{cases} \tag{255}$$

where the wavenumber k_1 coincides with 1 h/Mpc at redshift zero and scales with redshift like

$$k_1(z) = \frac{1h}{\mathrm{Mpc}} \cdot (1+z)^{2/(2+n_s)} \,. \tag{256}$$

The four free factors are fixed by the three fixed points defined above and the condition of continuity:

$$\begin{aligned} a_1 &= 1.4806\,\% \,, \; c_1 = 0.75056 \,, \\ a_2 &= 2.2047\,\% \,, \; c_2 = 1.5120 \,. \end{aligned}$$

What we will call later the "realistic case" amounts to trusting this error function up to large wavenumbers. Then, the information coming from small scales is suppressed gradually by the increasing relative error function, and the actual value of the cut-off k_{max} becomes effectively irrelevant. What we will instead call the "conservative case" is an analysis using this error function, while introducing a sharp cut-off at $k = 0.2\,\mathrm{h/Mpc}$ (i.e., the error is effectively infinite above this value), following the scaling in redshift as defined Eq. 238.

Examples of the effective error[6] for each likelihood, and for a few selected redshift bins, are shown in Figure 4 for galaxy clustering and Figure 5 for intensity mapping. On the same plot, we show which scales the experiments are most sensitive to, by plotting $\frac{d\chi^2}{dk d\mu}$ arbitrarily normalised to a constant relative difference between the theoretical and observed spectra ($\Delta P_g = \epsilon P_g$). The vertical line marks $k_{\mathrm{NL}}(\bar{z})$, which is used as a sharp cut-off for the conservative setting. Note the different $k_{\mathrm{NL}}(\bar{z})$ values, corresponding to the mean redshift of each bin. Both the realistic and conservative setting make use of the theoretical error (red). For galaxy clustering, the observational errors (blue) dominate the error in the radial direction (right panel), due to redshift space distortions, whereas the closer we get to $\mu = 0$ the observational

6 The effective theoretical error is defined analogously to the observational one (Eq. 243).

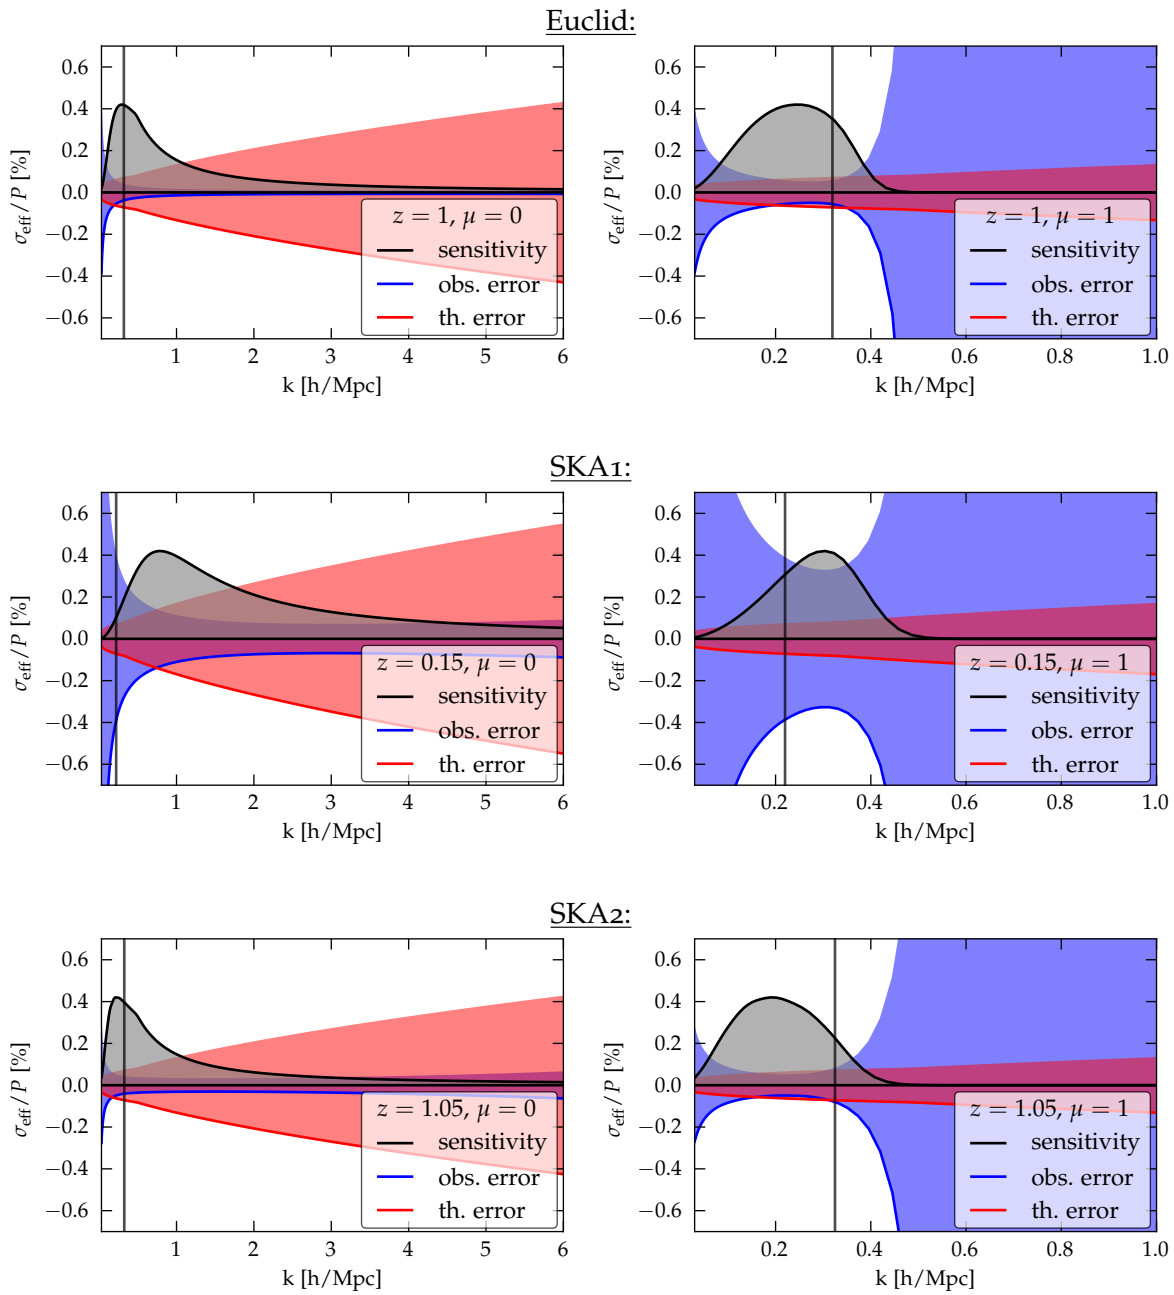

Figure 4: Galaxy clustering: examples of the relative effective error $\sigma_{\rm eff}/P_g$ in select redshift bins, decomposed into contributions from the observational error (blue) and theoretical error (red). To show which scales the experiment is most sensitive to (taking these errors into account), we show in grey the function $\sim d\chi^2/(dkd\mu)$, which is arbitrarily normalised to a constant relative difference between the theoretical and observed spectra ($\Delta P_g = \epsilon P_g$). The vertical line marks the sharp cut-off used for the conservative setting, $k_{\rm NL}(\bar{z})$.

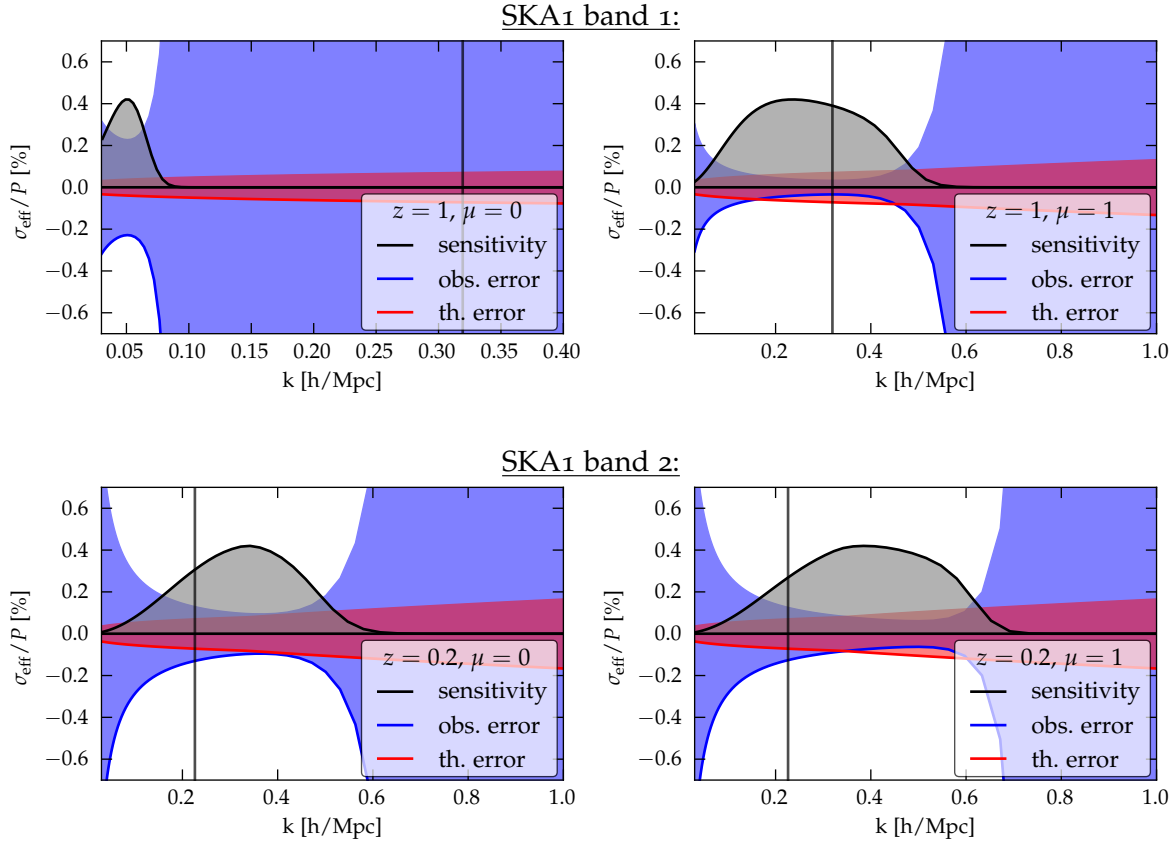

Figure 5: Intensity mapping: examples of relative effective errors $\sigma_{\rm eff}/P_{21}$ and sensitivity contributions $\sim d\chi^2/(dkd\mu)$, arbitrarily normalised to $\Delta P_{21} = \epsilon P_{21}$. The vertical line marks $k_{\rm NL}(\bar{z})$, which is used as a sharp cut-off for the conservative setting. Both the realistic and conservative setting make use of the theoretical error (red). Note the different $k_{\rm NL}(z)$ values, corresponding to the mean redshift of each bin.

error diminishes and the theoretical error takes over. This illustrates the necessity for introducing a measure of the error on non-linear modelling. The effect is also seen as a function of μ in the bottom panel of Figure 6. For intensity mapping, we notice that the trend is different from galaxy clustering (Figure 4). Due to the poor angular resolution, the observational error (blue) is larger in the transverse direction compared to the line-of-sight, especially for band 1, and dominates on most scales.

Just like we did for cosmic shear, we can illustrate the sensitivity of the likelihood to different parts of the data set by computing the χ^2 contribution projected on the parameters $\bar{z}$, k, or μ. The likelihood for galaxy surveys and intensity mapping can be written as

$$\chi^2 = \sum_{\bar{z}} \int_{k_{\min}}^{k_{\max}(\bar{z})} dk \int_{-1}^{+1} d\mu \; k^2 \frac{V_r(\bar{z})}{2(2\pi)^2} \times \frac{\Delta P^2}{\sigma^2_{\rm obs} + \sigma^2_{\rm th}} \,, \tag{257}$$

where ΔP is the difference between the fiducial and sampled power spectrum. Again, by omitting one of the sums or integrals, we obtain the desired projection. Similar to Figure 3 for cosmic shear, we see in Figure 6 the effect of varying the power spectrum by $\Delta P = \Delta P(\boldsymbol{k}, \bar{z}) = P - \hat{P} = 0.001P$, but now for galaxy clustering and intensity mapping and projected onto $\bar{z}$ (top row), k (middle row), or μ (bottom row) The left panel shows the realistic case and the right panel our conservative case.

From Figure 6, we see again that the sensitivity of galaxy survey forecasts for SKA1 are not competitive with the ones for Euclid and SKA2. This is because of the small sky coverage and low number of galaxies, which result in a large observational error. SKA1 will cover a much smaller redshift range than Euclid and SKA2. At low redshifts the constraining power of galaxy clustering data strongly depends on the amount of non-linear information available, i.e. on the non-linear uncertainty and cut-off. Since an SKA1 galaxy survey is limited to low redshifts, the difference between the sensitivity of the conservative scheme compared to the realistic one is expected to be quite large. We can see this effect in Figure 4 (middle row), where, for the conservative case, the sharp cut-off at $k_{\rm NL}(z)$ removes the k-range in which galaxy clustering with SKA1 is sensitive.

Intensity mapping with SKA1 is more promising than SKA1 galaxy clustering. In contrast to galaxy surveys, the observational error of intensity mapping dominates for high k for both $\mu = 0$ and $\mu = 1$ (Figure 5). As a result, there is less non-linear information available. The poor angular resolution limits the information gain to radial directions $\mu \simeq 1$ for large redshifts, i.e. for band 1 (Figure 6, bottom row). For the realistic case, band 1 provides a greater amount of information on linear and mildly non-linear scales, whereas band 2 is better at probing highly non-linear scales (Figure 5, left panel, with the cumulative effect summarized in Figure 6, middle row, left panel). In the conservative case, band 1 accesses more information than band 2, since the non-linear scales are largely removed by the sharp cut-off (Figure 5, with the cumulative effect in Figure 6, middle row, right panel).

Galaxy surveys with Euclid or SKA2 are much more sensitive than for SKA1. The excellent angular resolution, high sky coverage, and low shot noise levels strongly suppress observational errors. For $\mu = 0$ the information gain is almost entirely limited by theoretical accuracy, especially on non- and quasi-linear scales, so there is a lot of potential for increasing the sensitivity of predictions with improved modelling of non-linear effects on structure formation (Figure 4, left panel). At $\mu = 1$ redshift-space distortions suppress the power spectrum at high k to a level where shot noise again dominates (Figure 4, right panel). In comparison, Euclid is more sensitive at high redshifts than SKA2, because a higher density of detected galaxies is achieved, as the galaxies are brighter in the infrared than in the 21cm regime. SKA2 is

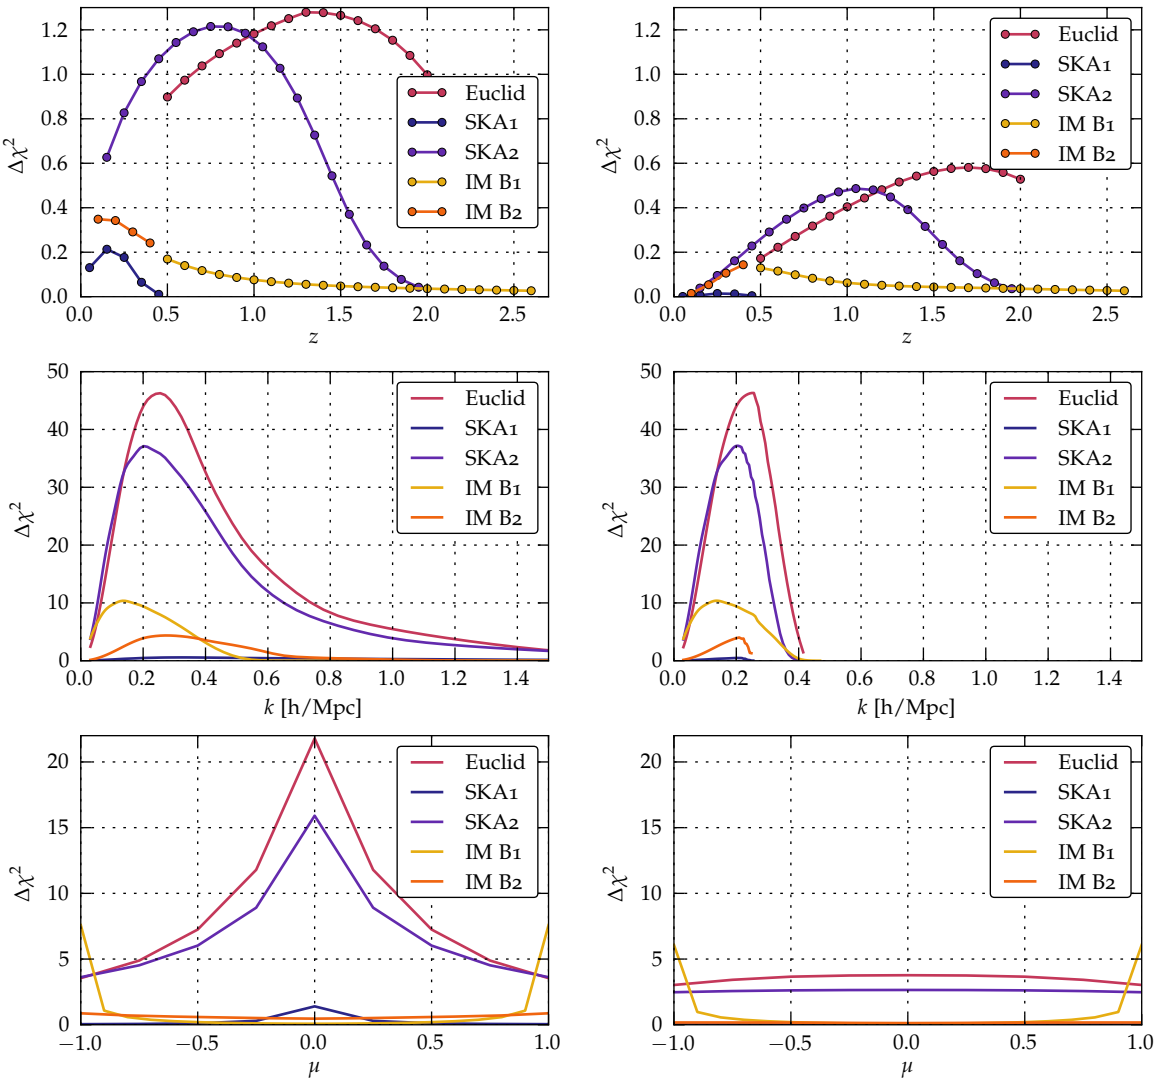

Figure 6: Sensitivity distribution for all three-dimensional power spectrum likelihoods. The left panel shows the realistic approach and the right panel the conservative one. The $\Delta\chi^2$-values are contributions when two of the quantities k, μ and z are integrated or summed over, and $\Delta P = 0.001P$ everywhere. For intensity mapping (IM), band 1 and 2 of SKA1 are considered.

more sensitive at lower redshifts, because of the greater sky coverage achieved due to the transparency of the Milky-Way to 21cm radiation (Figure 6, top row).

In chapter 8, section 8.4 and chapter 9, sections 9.4 and 9.5, we will use the results of this chapter to derive realistic (and conservative) sensitivities for Euclid and the SKA, with a particular focus on neutrino-related quantities and dark energy. However, first we will address a non-linear effect that is relevant when considering neutrinos: the scale-dependent galaxy bias induced by massive neutrinos.

3.5 EXPERIMENTAL SETUP

We run Markov Chain Monte Carlo (MCMC) forecasts for several possible experimental configurations, following the commonly used approach described for example in [96] and [156]. The method consists in generating mock data according to some fiducial model. One then postulates a Gaussian likelihood with some instrumental noise level, and fits theoretical predictions for various cosmological models to the mock data, using standard Bayesian extraction techniques. For the purpose of studying the sensitivity of the experiment to each cosmological parameter, as well as parameter degeneracies and possible parameter extraction biases, it is sufficient to set the mock data spectrum equal to the fiducial spectrum, instead of generating random realisations of the fiducial model.

In this section we summarize the mock data sets used in our forecasts in sections 8.4, 9.4 and 9.5, their shorthand names and relevant assumptions for non-linear modeling. These data sets can be combined in single forecasts. For this purpose we followed some rules to avoid double-counting of information:

- Only combine GC+GC, GC+IM or IM+IM if the redshift ranges do not overlap, since both take their information from the position of galaxies.
- Do not combine CS+CS because they all use information down to redshift zero.

Galaxy clustering (sections 3.1 and 3.4.2)

$k_{\rm min} = 0.02\,1/\text{Mpc}$ cuts off scales that are bigger than the bin width or violate the small angle approximation
Non-linear cut-off at $k_{\rm NL}(z) = k_{\rm NL}(0) \cdot (1+z)^{2/(2+n_s)}$ (see Eq. 238)
Theoretical uncertainty growing with k after $k = 0.01\,h/\text{Mpc}$ (see Eq. 255):
0.33% ($k \leq 0.01\,h/\text{Mpc}$), 1% ($k = 0.3\,h/\text{Mpc}$), 10% ($k = 10\,h/\text{Mpc}$)

- **Euclid GC cons.** Euclid galaxy clustering conservative
  Theoretical uncertainty, $k_{\rm NL}(0) = 0.2\,h/\text{Mpc}$
- **Euclid GC real.** Euclid galaxy clustering realistic
  Theoretical uncertainty, $k_{\rm max} = 10\,h/\text{Mpc}$
- **SKA1 GC cons.** SKA1-MID band 2 galaxy clustering conservative
  Theoretical uncertainty, $k_{\rm NL}(0) = 0.2\,h/\text{Mpc}$
- **SKA1 GC real.** SKA1-MID band 2 galaxy clustering realistic
  Theoretical uncertainty, $k_{\rm max} = 10\,h/\text{Mpc}$
- **SKA2 GC cons.** SKA2-MID galaxy clustering conservative
  Theoretical uncertainty, $k_{\rm NL}(0) = 0.2\,h/\text{Mpc}$

- **SKA2 GC real.** SKA2-MID galaxy clustering realistic
  Theoretical uncertainty, $k_{\rm max} = 10\,h/{\rm Mpc}$

Intensity mapping (sections 3.3 and 3.4.2)

Same $k_{\rm min}$, non-linear cut-off and theoretical uncertainty as for galaxy clustering

- **SKA1 IM1 cons.** SKA1-MID band 1 intensity mapping conservative
  Theoretical uncertainty, $k_{\rm NL}(0) = 0.2\,h/{\rm Mpc}$
- **SKA1 IM1 real.** SKA1-MID band 1 intensity mapping realistic
  Theoretical uncertainty, $k_{\rm max} = 10\,h/{\rm Mpc}$
- **SKA1 IM2 cons.** SKA1-MID band 2 intensity mapping conservative
  Theoretical uncertainty, $k_{\rm NL}(0) = 0.2\,h/{\rm Mpc}$
- **SKA1 IM2 real.** SKA1-MID band 2 intensity mapping realistic
  Theoretical uncertainty, $k_{\rm max} = 10\,h/{\rm Mpc}$

Cosmic shear (sections 3.2 and 3.4.1)

$\ell_{\rm min} = 5$ **Non-linear cut-off** at $\ell^i_{\rm max} = k_{\rm NL}(z) \cdot \bar{r}^i_{\rm peak}$ (see eqs. (238) and (240))

- **Euclid CS cons.** Euclid cosmic shear conservative
  $k_{\rm NL}(0) = 0.5\,h/{\rm Mpc}$
- **Euclid CS real.** Euclid cosmic shear realistic
  $k_{\rm NL}(0) = 2.0\,h/{\rm Mpc}$
- **SKA1 CS cons.** SKA1-MID cosmic shear conservative
  $k_{\rm NL}(0) = 0.5\,h/{\rm Mpc}$
- **SKA1 CS real.** SKA1-MID cosmic shear realistic
  $k_{\rm NL}(0) = 2.0\,h/{\rm Mpc}$
- **SKA2 CS cons.** SKA2-MID cosmic shear conservative
  $k_{\rm NL}(0) = 0.5\,h/{\rm Mpc}$
- **SKA2 CS real.** SKA2-MID cosmic shear realistic
  $k_{\rm NL}(0) = 2.0\,h/{\rm Mpc}$

Cosmic microwave background

- **Planck** Instead of the real Planck data, it is more convenient to run our forecasts with some mock temperature, polarization and CMB lensing data generated for the parameter values of our fiducial model. We go up to $\ell_{\rm max} = 3000$ and we use noise spectra matching the expected sensitivity of the final Planck data release, in particular improving constraints from polarization.

In section 9.5 we used a subset of the Euclid and SKA configurations, namely the "realistic" setup for galaxy clustering and the "conservative" setup for cosmic shear, and the DESI BAO setup from section 5.2, but we recomputed noise spectra for all CMB configurations from section 2.4 using the minimum variance estimator, in order to have a fair comparison between CMB experiments.

4

NEUTRINO INDUCED GALAXY BIAS

Chapter 4 is based on Vagnozzi, Brinckmann, Archidiacono, Freese, Gerbino & Lesgourgues [1]

Considerable effort is being devoted towards future large-scale structure surveys, such as galaxy redshift surveys, able to access smaller and increasingly non-linear scales. At the same time, achieving a reliable detection of M_ν requires an exquisite control of systematics, including both instrumental and modelling systematics. A potential particularly delicate modelling systematic stems from the fact that we only have access to *biased* tracers of the matter power spectrum, such as the clustering of galaxies. The underlying matter power spectrum is related to the measured power spectrum of a biased tracer $P_t(k,z)$ through

$$P_t(k,z) = b_m^2(k,z)P_m(k,z)\,, \tag{258}$$

where the bias b_m quantifies the statistical relation between the clustering of matter and of its luminous tracers (see e.g. [158] for a recent review). In this chapter, we will be concerned with galaxies as biased tracers of the underlying matter field.

The bias b_m, as appearing in Eq. (258), is defined with respect to the *total* matter field, comprising CDM, baryons, and massive neutrinos. When considering galaxies as tracers, we can already at this point appreciate an important subtlety: Eq. (258) implicitly assumes that galaxies trace the *total* matter field. However, the typical scales probed by galaxy clustering are below the neutrino free-streaming scale, where neutrinos do not cluster. Hence, one should actually expect galaxies to trace the CDM+baryons field instead of the total matter field. This expectation has been proved to be valid by several dedicated simulations [120, 121, 159, 160], as well as theoretical studies [161], and its consequences for parameter inference are the topic of investigation of this chapter. Our work is not the first one to investigate the impact of this aspect on parameter inference. As far as we are aware, such a study was first carried out in [127], which found that taking this effect into account when analysing future data will be crucial. In our work we confirm this finding, although there are a number of differences between our analysis and that of [127], which we shall comment more on later.

In cosmologies *without* massive neutrinos, the behaviour of the bias as a function of scale (wavenumber k) is well understood, albeit increasingly harder to model at smaller scales. On large, linear scales ($k \lesssim 0.15\,h\,\mathrm{Mpc}^{-1}$ at $z = 0$), the bias can be modelled to very good approximation as being constant [158]. On smaller, non-linear scales, complexities inherent to the process of galaxy formation and evolution make the bias intrinsically scale-dependent [158, 162]. On mildly non-linear scales, accurate phenomenological parametrizations of the scale-dependent bias exist (see e.g. [163] for constraints on scale-dependent galaxy bias models from current data, and [130] for a study on the scale-dependent galaxy bias from CMB lensing-galaxy cross-correlations, with applications to constraints on M_ν).

In cosmologies *with* massive neutrinos, the situation is more complex. In this case, the bias defined with respect to the *total* matter field is scale-dependent even on *large scales*, with the

form of the scale-dependence dictated by the value of M_ν: $b_m(k,z) = b_m(k,z,M_\nu)$. On the other hand, one can define the bias with respect to the CDM+baryons power spectrum $P_{cb}(k)$,

$$P_t(k,z) = b_{cb}^2(k,z)P_{cb}(k,z)\,. \tag{259}$$

In this case, the bias b_{cb} becomes scale-independent on large scales, and *universal* [120, 121, 159, 160]. That is, its scale-dependence is no longer determined by M_ν, but only by properties inherent to galaxy formation and evolution, and appears only on small scales[1]. Hereafter, all quantities (power spectra, biases, and so on) are implicitly considered to be redshift-dependent. Hence, for simplicity, we will drop all z-dependences in our equations from now on.

The common approach when deriving constraints on M_ν from clustering measurements is to define the bias with respect to the total matter power spectrum, as in Eq. (258). The bias is still approximated as scale-independent, at least on large scales. While formally incorrect, this choice does not affect analyses given the current sensitivity of cosmological data. However, this could be no longer true with upcoming high-precision cosmological data aiming to measure M_ν. Therefore, it is timely to explore the impact of neglecting the neutrino-induced scale-dependent bias (NISDB) in light of future clustering measurements.

In this chapter, we revisit the issue of the neutrino-induced scale-dependent bias in light of future data from the *Euclid* satellite [4, 168]. As we noted earlier, a previous similar study was performed in [127], where a Fisher matrix approach was adopted. Our work differs from [127] on at least three main aspects: I) we use a Markov Chain Monte Carlo (MCMC) approach, which is better suited to capture important parameter correlations, such as those involved in the measurement of the sum of the neutrino masses. The MCMC approach enables us to definitively show that *correcting for the NISDB will be important, in order not to bias the determination of cosmological parameters, in particular the sum of the neutrino masses*, II) we clarify a number of subtle issues concerning improper ways of correcting for the NISDB effect, and III) our MCMC forecasts includes an extensive modelling of systematic effects for the upcoming *Euclid* survey.

This chapter is organized as follows: in section 4.1, we show how to correct for the NISDB in practice, and clarify a number of subtle issues concerning this correction. In section 4.2, we describe the simulated datasets we consider and how we model them. In section 4.3, we present our findings. They indicate that correcting for the NISDB will be crucial for future cosmological surveys and are summarised in Tab. 11 for the reader's convenience. Finally, in section 4.4 we provide concluding remarks.

4.1 CORRECTING FOR NEUTRINO INDUCED BIAS

In this section, we describe in detail our prescription for correcting for the NISDB effect. We begin in Sec. 4.1.1 by considering a simple example to get a feeling for how the correction *could* be implemented. In section 4.1.2 we discuss a more general and correct approach in the presence of redshift-space distortions (RSD). We then comment on the impact of non-linear evolution and how it modifies our correction in section 4.1.3. Finally, we give a schematic

1 In reality, even defining the bias with respect to the CDM+baryons field still leaves a tiny residual scale-dependence in b_{cb} on large scales [128, 164–167]. Essentially, the reason is that the scale-dependent growth induced by massive neutrinos affects the critical collapse overdensity for halos at different scales. Heuristically, we can think of this effect as being due to the impact of massive neutrinos on halo formation, whereas the effect we will be concerned with in this paper is related to whether or not halos can form in first place. For simplicity, since the residual scale-dependence due to the effect of neutrinos on halo collapse is extremely small, we choose to ignore it here. This effect is nonetheless worth exploring in much more detail, and we plan to return to this issue in a future project.

summary of our final prescription for correcting for the NISDB effect, taking into account RSD and non-linear evolution, in section 4.1.4.

4.1.1 *Bias correction in the absence of redshift-space distortions*

In the absence of RSD, we can compare the two expressions for the power spectrum of a given tracer $P_t(k)$, Eqs. (258, 259), and obtain the following relation between the "meaningful" bias defined with respect to the CDM+baryons field b_{cb} and the "effective" one defined with respect to the total matter field b_m,

$$b_m(k, M_\nu) \quad = \quad b_{cb}(k)\sqrt{\frac{P_{cb}(k, M_\nu)}{P_m(k, M_\nu)}}\,. \tag{260}$$

At the level of linear perturbation theory, this can be expressed as

$$b_m(k, M_\nu) \quad = \quad b_{cb}(k)\frac{T_{cb}(k, M_\nu)}{T_m(k, M_\nu)}\,, \tag{261}$$

where $T_{cb}(k, M_\nu)$ and $T_m(k, M_\nu)$ are the linear transfer functions of the CDM+baryons and total matter components, respectively. In particular, the CDM+baryons transfer function is given by

$$T_{cb}(k, M_\nu) = \frac{\Omega_{cdm}T_{cdm}(k, M_\nu) + \Omega_b T_b(k, M_\nu)}{\Omega_{cdm} + \Omega_b}\,, \tag{262}$$

where Ω_{cdm} [Ω_b] and $T_{cdm}(k, M_\nu)$ [$T_b(k, M_\nu)$] denote the CDM [baryon] density parameter and transfer function, respectively. Notice that the definition of transfer function we adopt automatically includes the growth function $D(k, z)$. In Eq. (261), we have explicitly written all k and M_ν dependences to emphasize the important fact that b_{cb} is *universal* (its scale-dependence is not dependent on the value of M_ν), while b_m is not. The dependence of b_m on M_ν arises from the dependence of the transfer functions on the neutrino mass[2].

Equation (261) allows us to express the galaxy power spectrum P_t in terms of the total matter power spectrum P_m through a rescaling by a factor $(T_{cb}/T_m)^2$, since one can insert Eq. (261) into Eq. (258) to get

$$P_t(k, M_\nu) = b_{cb}^2(k)\left(\frac{T_{cb}(k, M_\nu)}{T_m(k, M_\nu)}\right)^2 P_m(k, M_\nu)\,. \tag{263}$$

Notice that all the terms appearing on the right-hand side of Eq. (263) are known: the linear total matter power spectrum P_m and transfer functions T_{cb} and T_m can be computed by using Einstein-Boltzmann codes, and we can model the universal bias b_{cb} accurately up to mildly non-linear scales.

4.1.2 *The impact of redshift-space distortions*

Redshift-space distortions (RSD) arise from the fact that galaxies are observed not in real space, but in redshift space. Thus, the redshift-space power spectrum we observe needs to be

2 We remind the reader that the scale-dependence of b_{cb} we explicitly introduced is relevant only on small non-linear scales [158], or equivalently at large k, which we will not be concerned with in this work.

corrected for the effect of peculiar velocities. Kaiser showed in a seminal paper [124] that, on large scales (small k) and when not including massive neutrinos, Eq. (258) is modified due to the effect of RSD to

$$P_t(k) = \left(b_m(k) + f_m(k)\mu^2\right)^2 P_m(k)\,, \tag{264}$$

where f is the growth rate, not to be confused with the growth factor that is included in the transfer functions. The growth rate f is defined as

$$f_m(k) = \frac{d\ln\left(\sqrt{P_m(k,z)}\right)}{d\ln a}\,, \tag{265}$$

with a being the scale factor (of course, the growth rate is implicitly redshift-dependent). Finally, the μ term in Eq. (264) is defined as the cosine of the angle between the Fourier mode $\mathbf{k}$ and the line-of-sight vector $\mathbf{r}$,

$$\mu \equiv \frac{\mathbf{k}\cdot\mathbf{r}}{kr}\,, \quad k \equiv |\mathbf{k}|\,, \quad r \equiv |\mathbf{r}|\,. \tag{266}$$

Note that the expression in Eq. (264) is not exact, as it lacks an exponential suppression due to the fingers-of-God (FoG) effect [126, 169]. However, the FoG correction is bias-independent and therefore does not impact our discussion. Hence, for simplicity, we drop it for the moment. It will be included later in the analysis.

Kaiser's result was derived for models in which the growth rate is scale-independent, which is true for the minimal ΛCDM model, but not in presence of massive neutrinos. In the latter case, one could expect again that RSD effects are driven solely by the baryon and cold dark matter fluctuations. This has been checked explicitly on the basis of simulations in [170], where it was proven that in the presence of RSD and massive neutrinos (but neglecting FoG effects), one has

$$P_t(k, M_\nu) = \left(b_{cb}(k) + f_{cb}(k, M_\nu)\mu^2\right)^2 P_{cb}(k, M_\nu)\,, \tag{267}$$

with

$$f_{cb}(k, M_\nu) \equiv \frac{d\ln\left(\sqrt{P_{cb}(k,z,M_\nu)}\right)}{d\ln a}\,. \tag{268}$$

As in the previous subsection, we can always use this result to define an effective bias $b_m(k, M_\nu)$ and an effective growth rate $f_m^{\rm eff}(k, z, M_\nu)$ such that the following holds,

$$P_t(k, M_\nu) = \left(b_m(k, M_\nu) + f_m^{\rm eff}(k, z, M_\nu)\mu^2\right)^2 P_m(k, M_\nu)\,, \tag{269}$$

provided that the "effective" bias $b_m(k, M_\nu)$ is related to the "meaningful" one by Eq. (261), while the effective growth factor should be defined as

$$f_m^{\rm eff}(k, z, M_\nu) \equiv \left(\frac{T_{cb}(k, z, M_\nu)}{T_m(k, z, M_\nu)}\right) f_{cb}(k, z, M_\nu)\,. \tag{270}$$

As a note of warning, the above should not be confused with the actual growth rate of the total matter spectrum

$$f_m(k, z, M_\nu) = \frac{d\ln\left(\sqrt{P_m(k,z,M_\nu)}\right)}{d\ln a}\,. \tag{271}$$

Both of these effective quantities, b_m and $f_m^{\rm eff}$, are easy to implement at the level of likelihood modules of an MCMC sampling package, with a proper computation of $f_{cb}(k,z,M_\nu)$ and $f_m^{\rm eff}(k,z,M_\nu)$ based on the total matter spectrum and on transfer functions. However, as we shall discuss very shortly, it is better to perform an implementation directly at the level of the Einstein-Boltzmann solver, which can be modified to output directly P_{cb} and f_{cb} on top of P_m and f_m: this allows to better take non-linear growth effects into account.

We emphasize that in the absence of RSD, the NISDB effect could be corrected equivalently by either rescaling the bias [Eq. (261)] or the power spectrum [Eq. (263)] by related quantities, (T_{cb}/T_m) or $(T_{cb}/T_m)^2$ respectively. The symmetry between bias and power spectrum is broken by the growth rate term in the RSD correction.

4.1.3 *The impact of non-linear effects*

Before discussing the details and results of our analysis, the impact of non-linear evolution on cosmological perturbations needs to be discussed. Usually, in Boltzmann solvers, non-linear effects are introduced via prescriptions such as `Halofit` [105] or `HMcode` [106]: these are accurate simulations-based fitting formulas for the non-linear matter power spectrum. In [152] `Halofit` has been revised and extended to describe non-linear evolution in cosmologies with massive neutrinos.

Let us consider the case in which the NISDB correction, envisaged in Sec. 4.1.2, is implemented at the level of likelihood module in the MCMC sampling package. In such a way, the correction would be implemented *after* applying the `Halofit` prescription to the linear evolution. However, the transfer functions entering Eqs. (261, 270) are by definition linear quantities. Hence, the NISDB correction, as described in Sec. 4.1.2, would rescale the non-linear power spectrum by a linear quantity, and would be somewhat inconsistent.

One way to address this issue, is to derive the right quantities not at the level of likelihood module, but rather at the level of the Einstein-Boltzmann solver. We modified the code `CLASS` [171] in such a way as to output at the same time $P_{cb}(k,z,M_\nu)$ and $P_m(k,z,M_\nu)$ for each given model. The former is used for computing the galaxy spectrum $P_t(k,z,M_\nu)$ [derived from Eqs. (267, 268)], while the latter can be used for computing other observables, for instance the Limber-approximated cosmic shear spectrum. For consistency, when computing non-linear spectra, the code processes the linear $P_{cb}(k,z,M_\nu)$ with a version of `Halofit` without massive neutrino corrections, and the linear $P_m(k,z,M_\nu)$ with a version of `Halofit` including the massive neutrino corrections[3] of Ref. [152]. These new features will be included in a forthcoming release of the public `CLASS` code.

4.1.4 *Prescription summary*

Let us briefly summarize our prescription for correcting for the NISDB effect, taking RSD and non-linear effects correctly into account. The prescription works according to the following three steps:

1. Compute the CDM+baryons power spectrum $P_{cb}(k)$ by tracking the evolution of the CDM and baryon overdensities in the Einstein-Boltzmann solver.

3 Note that another consistent way of getting the non-linear total matter power spectrum $P_m(k,z,M_\nu)$, previously investigated in Ref. [121], is to compute first the non-linear $P_{cb}(k)$ using `Halofit` without neutrino mass effects, and then to add the contribution of the linear neutrino spectrum $P_\nu(k)$ and of the cross-correlation term.

2. Use the `Halofit` prescription to calculate the non-linear corrections to $P_{cb}(k)$, and obtain the non-linear CDM+baryons power spectrum $P_{cb}^{\rm nl}(k)$. Pay attention to which version of `Halofit` is used. The version used in this case should *not* have the non-linear corrections due to massive neutrinos.

3. Finally, in the likelihood module, multiply the obtained non-linear CDM+baryons power spectrum $P_{cb}^{\rm nl}(k)$ by the *universal* b_{cb}-dependent RSD correction, to obtain the theoretical galaxy power spectrum $P_t^{\rm th}(k)$, i.e., to first approximation,

$$P_t^{\rm th}(k, M_\nu) = \left(b_{cb}(k) + f_{cb}(k, M_\nu)\mu^2\right)^2 P_{cb}^{\rm nl}(k, M_\nu)\,. \tag{272}$$

The obtained theoretical galaxy power spectrum $P_t^{\rm th}(k)$ can then be compared to the measured galaxy power spectrum through the likelihood function, in order to obtain constraints on the cosmological parameters through the usual MCMC analysis. In principle, a non-linear RSD correction should be used instead of the linear Kaiser formula in Eq. (272): for instance, see [172, 173] for a resummed perturbation theory approach to modelling non-linear RSD. However, the correct way of implementing non-linear RSD corrections is still a matter of debate in the community. In this work, we limit our analysis to linear scales and hence only model linear RSD (see section 4.2).

4.2 DATASETS AND ANALYSIS METHODOLOGY

In this section, we discuss the analysis we conduct in order to determine the impact of the NISDB correction on constraints from future galaxy clustering data. For each experiment, we construct a mock likelihood wherein the role of the mock data spectrum is played by the theoretical spectrum of a fiducial cosmological model. These likelihoods take into account the expected noise level and systematic uncertainties associated with the experiments.

We perform an MCMC parameter inference from these combined likelihoods in order to forecast the sensitivity of the future experiments to the model parameters (this approach was already followed e.g. in [96, 102, 118]). Our analysis pipeline is implemented in `MontePython` [174], an MCMC sampling and likelihood package that has recently been updated to `v3.0` [175]. The features of `v3.0`, most notably an expanded suite of likelihoods and a more efficient Metropolis-Hastings sampler, were employed for this work. `MontePython` is interfaced with the Boltzmann solver `CLASS` [171]. We choose a fiducial model where M_ν is pessimistically set to 0.06 eV (since this value of M_ν should in principle be the hardest to detect), the minimal value allowed by neutrino oscillation data in the normal ordering scenario [8–12].

On the CMB side, we use the likelihood `fake_planck_realistic` [50] included in `MontePython v3.0`, taking into account temperature, polarisation and CMB lensing extraction. We adopt noise spectra roughly matching those expected from the full *Planck* results[4]. For the purpose of forecasting sensitivities, it is easier to use a mock Planck likelihood rather than a real one, because we can then use the exact same fiducial model across all likelihoods.

On the side of Large Scale Structure (LSS), we employ the Fourier-space galaxy clustering mock likelihood `euclid_pk` presented in [102], with specifications for the *Euclid* mission taken from [4, 118, 129]. To reiterate, the likelihood employs the following corrections and approximations (for full details see section 3.1):

- Redshift-space distortions using the Kaiser formula [124].

4 Courtesy of Anthony Challinor, for the CORE-M5 proposal.

- Exponential suppression due to Fingers of God [126].
- Correction due to limited instrumental resolution.
- Correction for Alcock-Paczyński effect.
- Two nuisance parameters, β_0 and β_1, to account for inaccuracies in the bias evolution with redshift $b(z) = \beta_0(1+z)^{0.5\beta_1}$.
- Flat sky approximation [122, 123].
- NISDB correction discussed in this work.

We remark that the procedure to correct for the NISDB outlined in Sec. 4.1.4 is fully general once RSD and non-linear effects are properly accounted for. However, we will only use the linear power spectrum in this work. Our choice is motivated by considerations of simplicity, in addition to the fact that the correct way of accounting for non-linear RSD is still under debate. To restrict ourselves to linear scales, we impose a cut-off $k_{\rm max}$ in wavenumber space, scaling with redshift as in section 3.4,

$$k_{\rm max}(z) = k_{\rm max}(z=0) \times (1+z)^{\frac{2}{2+n_s}}\,, \tag{273}$$

where $k_{\rm max}(z=0) = 0.2\,h\,{\rm Mpc}^{-1}$ and n_s is the tilt of the primordial scalar power spectrum. We note that it will be interesting to check the impact of including non-linear scales, modulo a correct modelling thereof, on our results, and defer this analysis to future work.

We consider a seven-parameter cosmological model described by the six-parameter concordance ΛCDM model with the addition of the sum of the neutrino masses M_ν. The seven parameters are the physical baryon and CDM energy densities $\omega_b \equiv \Omega_b h^2$ and $\omega_{cdm} \equiv \Omega_{cdm} h^2$, the angular size of the sound horizon at decoupling θ_s, the redshift of reionization $z_{\rm reio}$ in place of the optical depth to reionization $\tau_{\rm reio}$, and the amplitude and tilt of the primordial scalar power spectrum A_s and n_s, in addition to the sum of the neutrino masses M_ν. The neutrino mass spectrum is approximated as three mass eigenstates degenerate in mass. It has been shown that this approximation is sufficiently accurate for the sensitivity of current and future data, which are strongly sensitive to the sum of the neutrino masses M_ν and negligibly sensitive to the masses of the individual eigenstates [18, 20, 50, 74, 99, 157, 176–180]. The values of the cosmological parameters defining our fiducial model are reported in Tab. 10.

Parameter	Fiducial value
ω_b	0.02218
ω_{cdm}	0.1205
$100\,\theta_s$	1.04146
$z_{\rm reio}$	8.24
$\ln[10^{10} A_s]$	3.056
n_s	0.9619
M_ν (eV)	0.06

Table 10: Fiducial values for the cosmological parameters adopted when simulating future cosmological data.

To explore the impact of not correctly accounting for the NISDB correction, we consider the following two cases, denoted by *correct* and *wrong* respectively:

- *Correct*: the NISDB correction is properly applied, at the level of Boltzmann solver, as discussed in Sec. 4.1.4.
- *Wrong*: the NISDB correction is not applied at all. We use the same fiducial mock data as in the *correct* case.

We run MCMC chains for the two cases considered above, monitoring the convergence of the generated chains through the Gelman-Rubin $R-1$ parameter [181], and requiring $R-1<0.01$ for the chains to be considered converged.

4.3 RESULTS

Here, we discuss the results of our MCMC analysis. The posterior distributions of M_ν are shown in Fig. 7, with the *correct* case corresponding to the blue solid curve and the *wrong* case to the red dashed curve. The vertical dot-dashed line at $M_\nu = 0.06\,\text{eV}$ corresponds to the input fiducial value of M_ν. The 68% C.L. bounds on M_ν from the two cases are reported in Tab. 11.

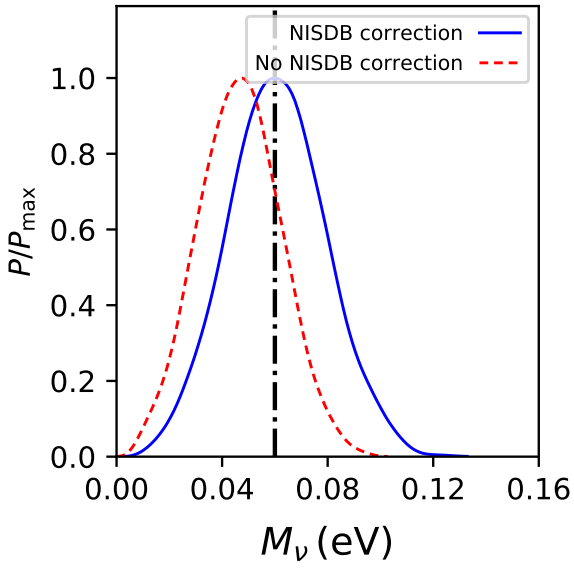

Figure 7: One-dimensional marginalized posterior probabilities, normalized to their maximum values, of the sum of the active neutrino masses M_ν (in eV) for the two cases considered in this analysis: the *correct* case (blue solid), where the neutrino-induced scale-dependent bias (NISDB) correction is properly applied, and the *wrong* case (red dashed), where the NISDB correction is not applied. The dot-dashed vertical line at $M_\nu = 0.06\,\text{eV}$ shows the input fiducial value of M_ν used in our analysis. It is clearly visible that the input fiducial value is perfectly recovered for the blue curve (*correct* case). When the correction is not applied, we fail in recovering the input fiducial value, as shown by the red curve.

We first consider the *correct* case, where the NISDB correction is properly applied. In this case, we can see that we correctly recover the fiducial value of M_ν, represented by the vertical dot-dashed line in Fig. 7. In fact, we find $M_\nu = (0.061 \pm 0.019)\,\text{eV}$, in perfect agreement with the fiducial value of $M_\nu = 0.06\,\text{eV}$.

Next, we consider the *wrong* case where the NISDB correction is not applied. In this case, we see that the choice of not applying the NISDB correction has biased our determination of M_ν. We find $M_\nu = (0.046 \pm 0.015)\,\text{eV}$, about 1σ away from the fiducial value of $M_\nu = 0.06\,\text{eV}$.

Parameter	NISDB correction	No NISDB correction
M_ν (eV)	0.061 ± 0.019	0.046 ± 0.015
$\omega_{cdm} \equiv \Omega_{cdm} h^2$	0.1205 ± 0.0003	0.1207 ± 0.0003
n_s	0.9621 ± 0.0014	0.9612 ± 0.0014

Table 11: Contraints at 68% C.L. on the sum of the neutrino masses M_ν and on the two cosmological parameters most correlated with M_ν, the cold dark matter physical density $\omega_{cdm} \equiv \Omega_{cdm} h^2$ and the scalar spectral index n_s. The constraints are obtained from the combination of mock CMB and galaxy clustering data described in Sec. 4.2. The datasets are simulated to match the expected sensitivity of the final data release from the *Planck* satellite and the upcoming *Euclid* satellite. The constraints are reported for the two cases considered in this analysis: the correction for the neutrino-induced scale-dependent bias is applied; the correction is not applied. By comparing the limits with the input fiducial values in Tab. 10, it clear that failure to apply the correction leads to biased determinations of cosmological parameters.

It is interesting to note that when the NISDB correction is not applied, the result is not only a biased determination of M_ν, but it also features a spurious $\sim 25\%$ decrease in the error bar σ_{M_ν}, consistent with the previous findings of [127].

These results are consistent with analytical expectations. It is well-known that on linear scales the effect of non-zero neutrino masses is to suppress the total matter power spectrum by an amount approximately given by [66, 67, 182]:

$$\frac{P_m(k, f_\nu)}{P_m(k, f_\nu = 0)} \simeq 1 - 8 f_\nu \,, \tag{274}$$

where f_ν is the neutrino fraction:

$$f_\nu \equiv \frac{\rho_\nu}{\rho_\nu + \rho_c + \rho_b} = \frac{\Omega_\nu}{\Omega_m} \,. \tag{275}$$

However, the CDM+baryons power spectrum is actually reduced by a smaller amount:

$$\frac{P_{cb}(k, f_\nu)}{P_{cb}(k, f_\nu = 0)} \simeq 1 - 6 f_\nu \,. \tag{276}$$

Comparing Eq. (274) and Eq. (276), we see that the impact of not applying the NISDB correction is to first approximation expected to lead to a decrease in both the inferred mean value M_ν and error σ_{M_ν} by a factor of 8/6. This ratio approximately matches our results for $M_\nu^{\rm correct} / M_\nu^{\rm wrong} = 0.061\,{\rm eV} / 0.046\,{\rm eV}$ and $\sigma_{M_\nu}^{\rm correct} / \sigma_{M_\nu}^{\rm wrong} = 0.019\,{\rm eV} / 0.015\,{\rm eV}$.

Finally, shifts in the inferred value of M_ν are expected to impact the inferred values of other cosmological parameters which are degenerate with M_ν. In particular, we have checked that the two most affected parameters are the CDM physical energy density $\omega_{cdm} \equiv \Omega_{cdm} h^2$ and the scalar spectral index n_s. The CDM physical energy density is negatively correlated with M_ν, while n_s is positively correlated. Both degeneracies are well-understood and documented in the literature [157, 180]. In Fig. 8 we show a triangular plot featuring the joint and one-dimensional posterior distributions of M_ν, ω_{cdm}, and n_s, for both the case where the NISDB correction is applied (blue solid curves/blue contours) and the case where it is not applied (red dashed curves/red contours).

The 68% C.L. bounds on ω_{cdm} and n_s are reported in Tab. 11. For the *correct* case, where the NISDB correction is applied, we find $\omega_{cdm} = 0.1205 \pm 0.0003$, perfectly recovering the input

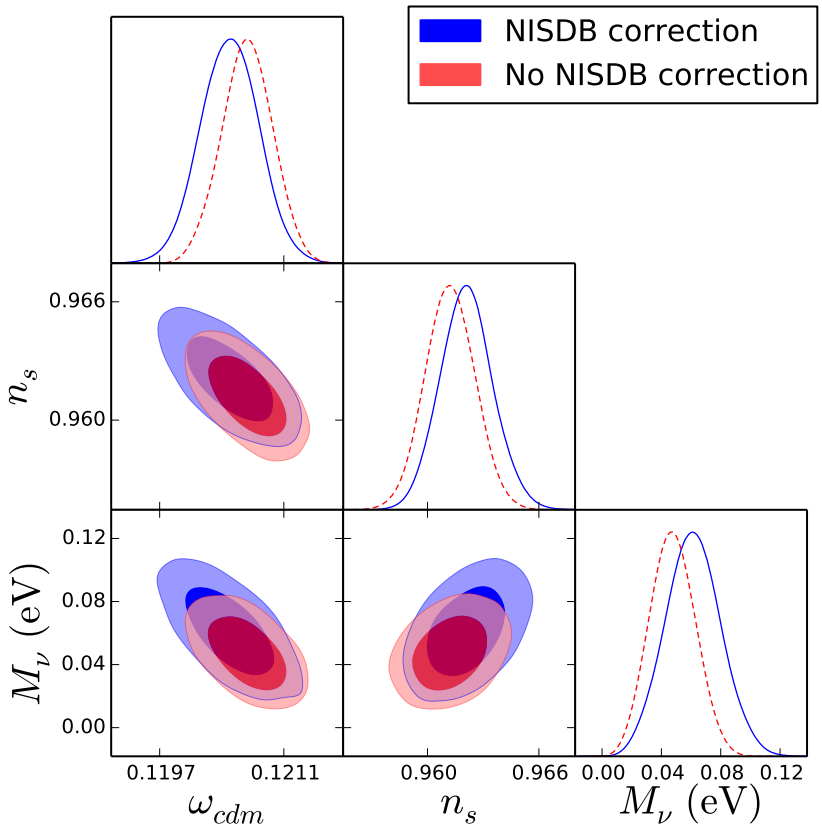

Figure 8: Triangular plot showing joint and one-dimensional marginalized posterior distributions for the parameters which are most correlated with the sum of the three active neutrino masses M_ν (in eV). These parameters are the dark matter physical density $\omega_{cdm} \equiv \Omega_{cdm}h^2$ and the scalar spectral index n_s. The panels along the diagonal show the one-dimensional probability distributions of the individual parameters. The remaining blocks show the 2D joint distributions. The distributions are shown for two cases: the case where the neutrino-induced scale-dependent bias (NISDB) correction is properly applied (blue, 1D posteriors are given by solid curves), and the case where the NISDB correction is not applied (red, 1D posteriors are given by dashed curves). The one-dimensional distributions along the diagonal represent normalized probability distributions and are hence in arbitrary units.

fiducial value. Similarly, we find $n_s = 0.9621 \pm 0.0014$, also perfectly recovering the input fiducial value. For the *wrong* case where the NISDB correction is not applied, the two previous values shift to $\omega_{cdm} = 0.1207 \pm 0.0003$ and $n_s = 0.9612 \pm 0.0014$ respectively. Although these correspond to $< 1\sigma$ shifts, they provide further indications that implementing the NISDB correction is important, not only for future determinations of the total neutrino mass, but also of other cosmological parameters, as found in [127].

Similar considerations concerning shifts in other parameters would hold in extended cosmologies as well, especially when considering additional parameters which are to some extent degenerate with M_ν (such as the dark energy equation of state parameter w and the curvature density parameter Ω_k), and could be explored in future work. As a final remark, we remind the reader that the details concerning the shifts and direction of degeneracy between the discussed parameters depend to some extent on the type of data used. For instance, when baryon acoustic oscillation distance measurements are considered, the degeneracy between M_ν and n_s

is expected to be reverted [180], i.e. the two parameters become negatively correlated instead of positively correlated.

4.4 CONCLUDING REMARKS

Cosmological data is exquisitely sensitive to the sum of the three active neutrino masses M_ν, and a combination of measurements from next-generation surveys is expected to provide the first measurement ever of M_ν, and thus of the absolute neutrino mass scale. Galaxy clustering is particularly sensitive to the effects of non-zero M_ν. However, galaxy clustering analyses also present significant challenges, such as the correct modeling of galaxy bias. Failure to do so could introduce significant model systematics which propagate to the determination of cosmological parameters, including M_ν. It is known that massive neutrinos introduce a scale-dependence in the galaxy bias *even on large scales*, if the bias is defined with respect to the total matter field. On the other hand, the bias defined with respect to the cold dark matter plus baryons field is *universal*, hence independent of the effects of M_ν. Most cosmological analyses in the presence of massive neutrinos so far have ignored this effect, defining the bias with respect to the total matter field and, at the same time, treating the bias as universal. Not accounting for this *neutrino-induced scale-dependent bias* (NISDB) could introduce severe model systematics in future analyses of galaxy clustering data.

In this work, we have quantified the importance of properly correcting for the NISDB effect when analysing galaxy clustering data. This issue was previously addressed in [127] using a Fisher matrix forecast. We revisit it through an MCMC sensitivity forecast and with an extended modeling of systematic effects. We have presented a simple prescription for correcting for the NISDB effect, summarized in Sec. 4.1.4. In doing so, we have also clarified some subtle issues concerning the correct way to implement the NISDB correction in the presence of redshift-space distortions and non-linearities.

We then presented a forecast based on mock cosmic microwave background and large-scale structure likelihoods, intended to mimic the legacy data release from the *Planck* satellite and measurements of the galaxy power spectrum from the *Euclid* satellite. We have shown that failure to implement the NISDB correction can introduce systematics in the inferred value of M_ν. In particular, we find that the value of M_ν inferred is a factor of $\sim 8/6$ lower than the fiducial value. At the same time, the 1σ uncertainty on the inferred value decreases by the same factor of $\sim 8/6$ with respect to the case where the correction is properly implemented. The latter effect represents a spurious increase in sensitivity. These results agree with the findings of [127] and match theoretical expectations. Finally, we have examined how the shift in the inferred value of M_ν correspondingly propagates to shifts in other cosmological parameters, such as $\Omega_{cdm}h^2$ and n_s.

We encourage the community to correctly account for the NISDB effect in future analyses of galaxy clustering data in the presence of massive neutrinos, in order to increase the robustness of the analyses and minimize the impact of modeling systematics. The tools necessary to easily correct for the neutrino-induced scale-dependent bias effect will be made publicly available in an upcoming release of the `CLASS` code.

5

NEUTRINO PHYSICS

Chapter 5 is based on Archidiacono, Brinckmann, Lesgourgues & Poulin [157]. Additionally, the introduction and section 5.0.1 is based on Di Valentino, Brinckmann, Gerbino, Poulin et al. [50]

Neutrino oscillation data show that neutrinos must be massive, but the data are insensitive to the absolute neutrino mass scale. For a normal hierarchy of masses ($m_1, m_2 \ll m_3$), the mass summed over all eigenstates is approximately at least 60 meV, while for an inverted hierarchy ($m_3 \ll m_1, m_2$) the minimal summed mass is approximately 100 meV [68–70].

The individual neutrino masses in these hierarchical limits are below the detection limit of current and future laboratory β-decay experiments, but they can remarkably be probed by cosmology [33, 66, 67, 72–74, 182–184]. The detection of the neutrino mass scale is even considered as one of the safest and most rewarding targets of future cosmological surveys, since we know that these masses are non-zero, that they have a significant impact on structure formation, and that their measurement will bring an essential clue for particle physicists to decipher the neutrino sector puzzle (origin of masses, leptogenesis and baryogenesis, etc.). Even the unlikely case of a non-detection would be interesting, since it would force us to revise fundamental assumptions in particle physics and/or cosmology, see e.g. Ref. [75].

Besides experimental sensitivity, parameter constraints are limited by degeneracies: a degeneracy indicates the ability of one parameter to mimic the effect of another parameter on a particular observable, making it impossible to disentangle them and to corner the value of each parameter separately. The key approach to tackle this problem consists in a joint analysis of complementary probes with different degeneracy directions in parameter space. For that reason, the next step in the era of precision cosmology will be based on the synergy of high- and low- redshift probes.

One of the parameters that will benefit from such an approach is the neutrino mass sum (hereafter M_ν). Indeed, the impact of massive neutrinos on cosmological observables comes from a very special effect: light massive neutrinos behave as radiation before their non-relativistic transition, while afterwards they gradually become a matter component; therefore their impact on cosmological probes at different redshifts is closely related to their mass.

For individual neutrino masses below 600 meV, the non-relativistic transition of neutrinos takes place after photon decoupling. After that time the neutrino density scales like matter instead of radiation, with an impact on the late expansion history of the universe. This is important for calculating the angular diameter distance to recombination, which determines the position of all CMB spectrum patterns in multipole space. At the time of the non-relativistic transition, metric fluctuations experience a non-trivial evolution which can potentially impact the observed CMB spectrum in the range $50 < \ell < 200$ due to the early ISW effect [66, 185, 186].

However, for individual neutrino masses below 100 meV, the non-relativistic transition happens at $z < 190$, hence too late to significantly affect the early ISW contribution.

Finally, massive neutrinos slow down gravitational clustering on scales below the horizon size at the non-relativistic transition, leaving a clear signature on the matter power spectrum [66, 182, 187]. The magnitude of this effect is controlled mainly by the summed neutrino mass M_ν. Roughly speaking, the suppression occurs on wavenumbers $k \geq 0.01\,h\,\mathrm{Mpc}$ (which means that even relatively large wavelengths are affected), and saturates for $k \geq 1\,h\,\mathrm{Mpc}$. Above this wavenumber and at redshift zero, the suppression factor is given in first approximation by $(M_\nu/10\,\mathrm{meV})\%$, i.e. at least 6% even for minimal normal hierarchy [66, 152, 182]. CMB lensing is expected to be a particularly clean probe of this effect [65, 99, 179, 188].

These neutrino mass effects have been widely studied in the literature [66, 185, 189, 190] and their impact on CMB and large scale structures on linear scales is well known. Even on non-linear scales, the neutrino mass effect is better understood thanks to recent progress in N-body simulations [152, 191–193].

However, neutrino cosmology is about to face a revolution for two reasons. First of all, current upper bounds on the neutrino mass sum are getting closer and closer to the minimum value allowed by the inverted hierarchy $M_\nu \sim 0.11$ eV [13, 15, 71]. Thus, future experiments will look at ultra-light neutrinos that became non-relativistic in a relatively recent cosmological epoch.

Secondly, future galaxy surveys will reach a very high sensitivity on very small scales. As for now, the use of small scale data is limited by the uncertainty on non linear structure formation, which is difficult to model, especially in presence of massive neutrinos [121, 194–197]. A major theoretical goal in the next few years will be to provide a better understanding of the processes governing clustering on small scales. Having the non linear effects under control, we will be able to exploit small scale data in order to break degeneracies. The neutrino mass effects are already important on linear scales, but by including smaller and smaller scales one would have a better lever arm and improve the constraints on M_ν.

The aim of this chapter is to investigate the physical effects induced by massive neutrinos as they will be unveiled by future cosmological data. We will pay specific attention to the correlation between M_ν and other cosmological parameters, and show that directions of degeneracy are very sensitive to probes of the cosmic history at different epochs. For some combinations of CMB and Large Scale Structure data sets, a correlation between M_ν and τ_{reio} has already been observed in Refs. [198, 199], but its interpretation is far from obvious and requires a detailed investigation. This correlation is very important, for the reason that independent measurements of the optical depth by 21cm surveys will lead to a remarkable improvement on the sensitivity to the neutrino mass [198]. We will confirm this expectation with a dedicated forecast showing that even the minimum allowed value of the summed neutrino mass could be detected at the 5σ level in a time scale of about ten years. Before proceeding, we will consider the question of neutrino mass splitting.

5.0.1 *Neutrino mass splitting*

Cosmology is mainly sensitive to the summed neutrino mass M_ν, but the mass splitting does play a small role, since the free-streaming length of each neutrino mass eigenstate is determined by the individual masses [33, 66, 67, 74, 176, 184]. Hence, before doing forecasts for

future high-precision experiments, it is worth checking the impact of making different assumptions of the mass splitting (for fixed total mass) on the results of a parameter extraction. If this impact is found to be small, we can perform generic forecasts sticking to one mass splitting scheme. Otherwise, several different cases should be considered separately.

We know from particle physics that there are two realistic neutrino mass schemes, NH and IH, both tending to a nearly-degenerate situation in the limit of large M_ν, but that limit is already contradicting current bounds ($M_\nu < 210$ meV from Planck 2015 TT+lowP+BAO [2], $M_\nu < 140$ meV when including the latest Planck polarisation data [71], $M_\nu < 130$ meV with recent BAO+galaxy survey data [15] and $M_\nu < 120$ meV with BOSS Lyman-α data [13], all at 95%CL). On top of NH and IH, the cosmological literature often discusses three unrealistic models (for the purpose of speeding up Boltzmann codes and integrating only one set of massive neutrino equations): the degenerate case with masses $(M_\nu/3, M_\nu/3, M_\nu/3)$, that we will call DEG; the case $(M_\nu/2, M_\nu/2, 0)$ that we will call 2M and the case $(M_\nu, 0, 0)$ that we will call 1M. These three unrealistic cases are potentially interesting to use as a fitting model in a forecast, because the total mass can be varied down to zero: thus, on top of estimating the value of M_ν, one can assess the significance of the neutrino mass detection by comparing the probability of $M_\nu = 0$ to that of the mean or best-fit value. Any of the DEG, 1M, or 2M models can achieve this purpose, however, we can already discard 1M and 2M, as a detailed inspection of the small difference between the matter power spectrum of these three models for fixed M_ν shows that the spectrum of the DEG model is much closer to that of the two realistic models (NH, IH) than the spectrum of 1M or 2M (see e.g. Figure 16 in [67]). Even current data starts to be slightly sensitive to the difference between 1M and (NH, IH) [18]. Hence we only need to address the question: can we fit future data with the DEG model, even if the true underlying model is probably either NH or IH, or does this lead to an incorrect parameter reconstruction?

Run	Fiducial model	Fitted model	posterior curve in Figure 9
1.	NH with $M_\nu = 0.06$ eV	DEG	top panels, green
2.	NH with $M_\nu = 0.06$ eV	NH	top panels, grey
3.	NH with $M_\nu = 0.10$ eV	DEG	bottom panels, solid green
4.	NH with $M_\nu = 0.10$ eV	NH	bottom panels, solid grey
5.	NH with $M_\nu = 0.10$ eV	IH	bottom panels, solid red
6.	IH with $M_\nu = 0.10$ eV	DEG	bottom panels, dashed green
7.	IH with $M_\nu = 0.10$ eV	NH	bottom panels, dashed grey
8.	IH with $M_\nu = 0.10$ eV	IH	bottom panels, dashed red

Table 12: List of fiducial and fitted model used to check for possible parameter reconstruction bias when using the wrong assumptions on neutrino mass splitting.

We first consider a fiducial model with a total mass $M_\nu = 60$ meV, thus necessarily given by NH. We generate mock data using the precise mass splitting of NH for such a value, with $\Delta m^2_{atm} = 2.45 \times 10^{-3}$ (eV)^2 and $\Delta m^2_{sol} = 7.50 \times 10^{-5}$ (eV)^2. We then compare the results of forecasts that assume either DEG or NH as a fitting model (still with fixed square mass differences). In both cases, the free parameters are the usual 6-parameter ΛCDM (with fiducial values given in footnote 1) and M_ν. These two forecasts correspond to the first two lines in Table 12. The results for the CORE-M5 satellite, alone or in combination with DESI BAOs and Euclid cosmic shear, are shown in the top three panels of Figure 9. This is the most pessimistic case for measuring the neutrino mass, since it corresponds to the minimal total mass allowed

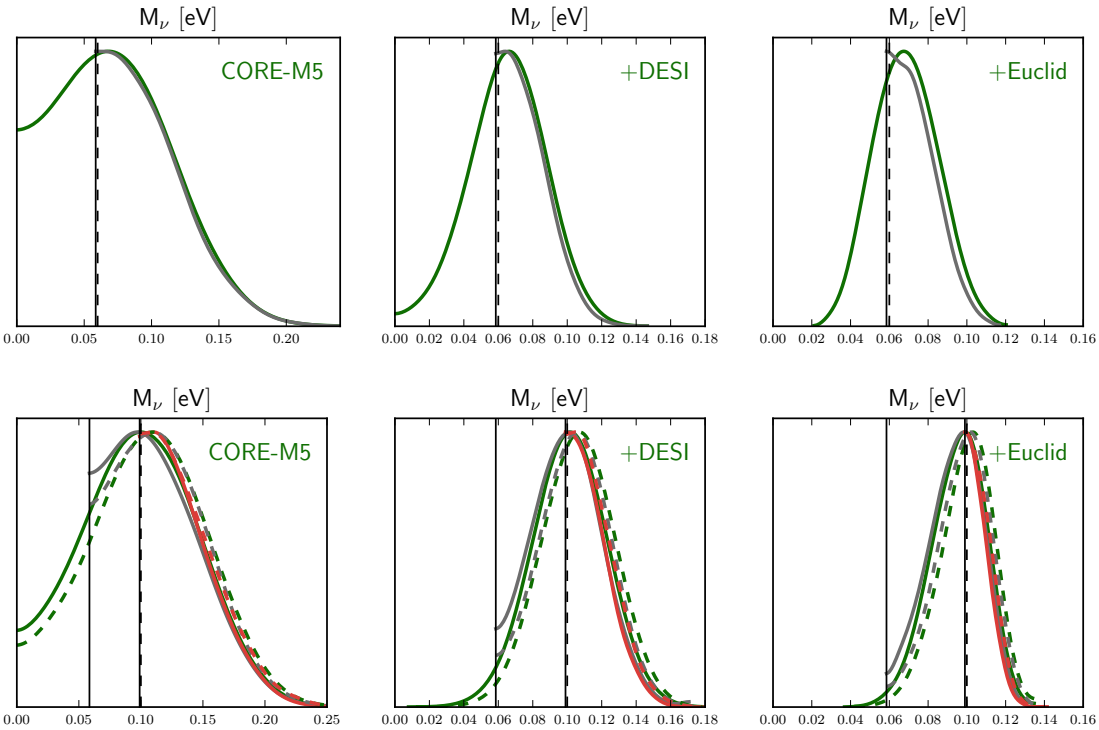

Figure 9: Reconstruction of the total neutrino mass using various schemes for the mass splitting (NH, IH, degenerate), not always matching the assumed fiducial model. Table 12 gives the explicit correspondence between the different curves and the assumptions made on the fiducial and fitting models. Vertical solid lines show lower prior edges in the NH and IH cases, while dashed lines show the fiducial values. The MCMC runs extracts the fiducial mass up to some reconstruction bias never exceeding 0.5σ.

by oscillation data. When looking at the results, one should keep in mind that we are fitting directly the fiducial spectrum, hence the posterior would peak at the fiducial value in absence of reconstruction bias; while with real scattered data the best fit would be shifted randomly, typically by one sigma. By looking at the results of the DEG fit (green curves in Figure 9 and numbers in Table 19), we see that CORE-M5 alone would not detect $M_\nu = 60$ meV with high significance, but it would typically achieve a 3σ detection in combination with DESI BAOs, or a 4σ detection when adding also Euclid cosmic shear data. There is a small offset between the mean value of M_ν found in the DEG fit and the fiducial value, corresponding respectively to 0.2σ, 0.2σ, 0.5σ in the CORE, CORE+DESI, and CORE+DESI+Euclid cases. This can be attributed to bias reconstruction from assuming the wrong fitting model. However, in this situation, the conclusion of fitting real data with DEG would be that the preferred scenario is NH, since $M_\nu = 100$ meV would be disfavoured typically at the 2σ level by CORE+DESI+Euclid, and one would then perform a second fit assuming NH in order to eliminate this reconstruction bias. More detailed discussions on the discrimination power of future data between NH and IH can be found e.g. in Refs. [33, 74].

Next, we considered a fiducial total mass $M_\nu = 100$ meV, which could be achieved either within the NH or IH model. We are not interested in the possibility of directly discriminating between these two models, because the sensitivity of CORE+DESI+Euclid is clearly too low for such an ambitious purpose. Instead we only want to check whether using the DEG model for the fits introduces significant parameter bias. For that purpose, we perform six forecasts for each data set, corresponding to the two possible fiducial models (NH or IH) fitted by each of the three models DEG, IH or NH. We see on the lower panels of Figure 9 that the fiducial mass is again correctly extracted by the DEG fits, up to a bias ranging from 0.1σ to 0.3σ: this is smaller than with a fiducial mass of 60 meV because masses are now larger and relative differences between NH, IH, and DEG are reduced. The error bars are always the same up to less than 0.1σ differences.

We have checked that regardless of the real mass splitting realised in nature, and with the experimental data sets discussed in this analysis, we can correctly reconstruct the mass simply by fitting the DEG model to the data. For the purpose of our forecasts, the most important things to check are that the error is stable under different assumptions, and that the reconstruction bias induced by fitting DEG to NH or DEG to IH is under control: this is found to be the case. So the next forecasts can be done using either NH or IH as a fiducial, and sticking to DEG as the fitted model. We can even do something simpler and use DEG as both fiducial and fitted model in the forecasts, since we know that if the fiducial model was NH or IH we would not have a large bias. This is exactly what we will do in the rest of this chapter and in chapters 8 and 9. However, we also see that in future analyses, we ought to be a little bit more careful, and compare the results of different fits using either NH or IH as a fitted model, to assess the impact of different assumptions on the posterior probability for M_ν.

5.1 EFFECT OF A SMALL NEUTRINO MASS ON THE CMB

5.1.1 *General parameter degeneracies for CMB data*

In the minimal, flat, 6-parameter ΛCDM model, it is well-known that the CMB temperature and polarisation unlensed spectra are determined by a number of effects[1], which remain

1 For a review of these effects, see e.g. Ref. [200], section 5.1 of Ref. [66], and Ref. [201].

identical as long as one fixes quantities usually depending on distance and density ratios, such as:

- the sound horizon angular scale $\theta_s(z_{\rm dec}) = \frac{d_s(z_{\rm dec})}{d_A(z_{\rm dec})}$ at decoupling,
- the diffusion angular scale $\theta_d(z_{\rm dec}) = \frac{d_d(z_{\rm dec})}{d_A(z_{\rm dec})}$ at decoupling,
- the baryon-to-photon ratio $R_{\rm dec} = \frac{3\rho_b}{4\rho_\gamma}\Big|_{\rm dec}$ at decoupling,
- the redshift of radiation-to-matter equality $z_{\rm eq} = \frac{\rho_m^0}{\rho_r^0} - 1$,
- the redshift of matter-to-cosmological-constant equality $z_\Lambda = \left(\frac{\rho_\Lambda^0}{\rho_r^0}\right)^{1/3} - 1$.

The CMB spectra also depend on a few extra parameters, like the scalar amplitude and tilt (A_s, n_s) and the optical depth at reionization $\tau_{\rm reio}$. However, bearing in mind that the small-ℓ (large angular) part of the spectra is loosely constrained due to cosmic variance, the parameters z_Λ, A_s and $\tau_{\rm reio}$ are always less constrained by CMB data than $(\theta_s, \theta_d, R, z_{\rm eq}, n_s)$, and also than the combination $A_s e^{-2\tau_{\rm reio}}$ giving the overall spectrum normalisation on small angular scales. The fact that we actually measure lensed CMB spectra gives extra information on the amplitude and slope of the matter power spectrum $P(k, z)$ at low redshift: in practice, this increases the sensitivity to the parameters (A_s, z_Λ), which enter into the normalisation of $P(k, z)$.

Adding neutrino masses into the model leads to several new effects studied extensively in the literature [66, 186, 201]:

(A) Neutrino masses affects the background expansion history. If we rely on standard assumptions for the photon and background densities ($T_{\rm cmb} = 2.726$ K, $N_{\rm eff} = 3.046$) and further fix ω_b and $\omega_{\rm cdm}$, the changes in the background evolution caused by neutrino masses are confined to late times. Then, the values of $d_s(z_{\rm dec})$, $d_d(z_{\rm dec})$, $R_{\rm dec}$ and $z_{\rm eq}$ are preserved, and the neutrino masses only impact the angular diameter distance (and, therefore, θ_s and θ_d in an equal way) and z_Λ (and hence, the loosely constrained late ISW effect). It is even possible to choose an appropriate value of the cosmological constant for each set of neutrino masses, in order to keep a fixed $d_A(z_{\rm dec})$: in that case, the impact of neutrino masses on the background is confined to variations of z_Λ and of the late ISW effect, and cannot be probed accurately due to cosmic variance, unless external non-CMB datasets come into play.

(B) At the perturbation level, massive neutrinos interact gravitationally with other species and produce small distortions in the CMB peaks. For individual neutrino masses m_ν smaller than ~ 600 meV, the neutrinos become non-relativistic after recombination: in that case the distortions can only be caused by the early ISW effect, and affect the CMB temperature spectrum in the multipole range $50 < \ell < 200$ [66, 185, 186]. Note that this neutrino-mass-induced early ISW effect takes place even if the redshift of equality is kept fixed: it is different from the redshift-of-equality-induced early ISW effect, which affects the height of the first CMB peak in the range $100 < \ell < 300$.

(C) Finally, at the lensing level, massive neutrinos slow down the growth of small-scale structure (leading to the well-known suppression factor $1 - 8\omega_\nu/\omega_m$ in the small-scale matter power spectrum at redshift zero) and globally decrease the impact of CMB lensing: the peaks are less smoothed and the damping tail less suppressed [65].

All these effects have played a role in previous constraints on neutrino masses from CMB data alone, or combined with other probes. Interestingly, while the sensitivity of CMB instruments increases with time, different effects come to dominate the neutrino mass constraints: early ISW effects **(b)** with WMAP alone [60], lensing effects **(c)** with Planck alone [2], and background effects **(a)** when combining CMB data with direct measurements of H_0 [85]. There are now several combinations of cosmological probes giving a 95%CL upper bound on the summed mass $M_\nu \equiv \sum m_\nu$ of the order of 120 meV to 150 meV [13, 15, 18, 71], while neutrino oscillation data enforces $M_\nu \geq 60$ meV at 95% CL [8]. The remaining conservatively allowed window is so narrow, $\Delta M_\nu \sim 90$ meV, that the impact of a realistic variation of the neutrino masses on the CMB is getting really small. Our purpose in section 5.1.3 is to study precisely this impact, and to understand the degeneracy between M_ν and other parameters when using future CMB data only, specifically for the very low mass range 60 meV$< M_\nu <$ 150 meV. This requires some preliminary remarks in section 5.1.2.

5.1.2 *CMB data definition*

The discussion of degeneracies is meaningless unless we specify which data set, and which experimental sensitivities, we are referring to. In this chapter, we take as a typical example of future CMB data a next-generation CMB satellite similar to the project COrE+, submitted to ESA for the call M4. A new version of CORE was recently submitted again for the call M5 [202, 203], with a small reduction of the instrumental performances, mainly in angular resolution. However, COrE+ and CORE-M5 are very similar, and the conclusions of this paper would not change significantly by adopting the CORE-M5 settings.

When displaying binned errors in C_l plots, and when doing MCMC forecasts with mock data and synthetic likelihoods, we assume that this CORE-like experiment is based on 9 frequency channels with sensitivity and beam angles given in section 2.4 We mimic the effect of sky masking by adopting a Gaussian likelihood with an overall rescaling by a sky fraction $f_{\rm sky} = 0.70$.

Our dataset consists primarily of temperature and E-mode polarisation auto-correlation and cross-correlation spectra $C_\ell^{TT}, C_\ell^{EE}, C_\ell^{TE}$. To get more information on CMB lensing, one can either analyse B-mode maps (in absence of significant primordial gravitational waves, the B-mode only comes from CMB lensing and foregrounds) and add the C_ℓ^{BB} spectrum to the list of observables; or perform lensing extraction with a quadratic or optimal estimator [204], and add the CMB lensing potential spectrum $C_\ell^{\phi\phi}$ to the list of observables (equivalently one could use the deflection spectrum $C_\ell^{dd} = \ell(\ell+1)C_\ell^{\phi\phi}$). We cannot use both C_ℓ^{BB} and $C_\ell^{\phi\phi}$ in the likelihood: the same information would be counted twice. Here we choose to use the lensing potential spectrum, which better separates the contribution of different scales to lensing, that would be mixed in the C_ℓ^{BB} spectrum by some integration kernel. To give an example, we will see in Figure 10 (bottom plots) that the neutrino mass effect on $C_\ell^{\phi\phi}$ is more pronounced on small angular scales, while in the lensed C_ℓ^{BB} this effect would be nearly independent of ℓ[2].

So the CMB data set that we have in mind consists in measurements for $C_\ell^{TT}, C_\ell^{EE}, C_\ell^{TE}, C_\ell^{\phi\phi}$, with a synthetic Gaussian likelihood similar to that in Ref. [96], and a lensing extraction error spectrum $N_\ell^{\phi\phi}$ based on the quadratic estimator method [97] for the EB estimator. In

2 However, we will also see that *within the range in which error bars are small*, the neutrino mass effect on $C_\ell^{\phi\phi}$ is also *nearly* ℓ-independent, so we may expect that trading $C_\ell^{\phi\phi}$ against C_ℓ^{BB} in the likelihood would have a minor impact on our conclusions.

the likelihood, we keep the lensed $C_\ell^{TT}, C_\ell^{EE}, C_\ell^{TE}$. Indeed, unlike C_ℓ^{BB}, these spectra are only weakly affected by lensing, and the lensing information redundency between lensed temperature/polarisation spectra and the $C_\ell^{\phi\phi}$ spectrum is small enough for being negligible at the instrumental sensitivity level of a CORE-like experiment [50].

5.1.3 *Degeneracies between very small M_ν's and other parameters with CMB data only*

We will discuss the impact of increasing the neutrino mass, while keeping various parameters or combination of parameters fixed. We illustrate this discussion with the plots of Figure 10, showing the spectrum ratio between different models sharing a summed mass $M_\nu = 150$ meV and a baseline model[3] with $M_\nu = 60$ meV. The plots shows the residuals of the lensed TT (top), lensed EE (middle) and lensing potential (bottom) power spectrum, as a function of multipoles ℓ with a linear (left) or logarithmic (right) scale. The light/pink and darker/green shaded rectangles refer respectively to the binned noise spectrum of a cosmic-variance-limited or CORE-like experiment, with linear bins of width $\Delta\ell = 25$. All spectra are computed with the Boltzmann solver CLASS [171, 205, 206], version 2.5.0, with the high precision settings `cl_permille.pre`.

Our discussion will also be illustrated by the results of Monte Carlo Markov Chains (MCMC) forecasts for our CORE-like experiment: Figure 11 gives the 2D probability contours for the pairs of parameters most relevant to our discussion. The MCMC forecasts are done with version 2.2 of the MontePython package [174].

The main conclusions can be reached in four steps:

1. We first assume that we increase neutrino masses with respect to the baseline model, while keeping the parameters $\{\omega_b, \omega_{\rm cdm}, h, n_s, A_s, \tau_{\rm reio}\}$ fixed (green solid curve in Figure 10). Given the discussion in point **(a)**, we expect that this is not a very clever choice, because the angular diameter distance is not preserved. So if the baseline model is a good fit to the data, the new model will be discrepant. Indeed, by looking especially at the top left and middle left plots in Figure 10, we see even-spaced oscillations signaling a change in the angular diameter distance, and the residuals are far above the instrumental noise.

2. We then perform the same increase in M_ν, but now with a fixed angular diameter distance to recombination, which means that $\{\omega_b, \omega_{\rm cdm}\}$ are still fixed, but h varies. With CLASS, this is easily achieved by keeping the input parameter $100\theta_s$ constant. Since the early cosmology and the sound horizon at decoupling are fixed, fixing θ_s means adjusting H_0 and the angular diameter distance for each M_ν. Then, the angular diffusion scale θ_d is also automatically fixed. In Figure 10, this transformation corresponds to the dashed red residuals. As expected, the previous oscillations disappear in the residuals. The only visible effects are much smaller oscillations, some tilt at large ℓ due to a different level of CMB lensing, and a tilt at small ℓ due to a different late ISW effect. However, both effects are below cosmic variance. We conclude that, the measurement of the temperature and E-mode spectra alone does not allow us to distinguish between $M_\nu = 60$ meV and

3 Our discussion is general and the value of cosmological parameters for the baseline model is unimportant. We choose Planck-inspired values, $\{\omega_b, \omega_{\rm cdm}, h, n_s, A_s, \tau_{\rm reio}, M_\nu\} = \{0.02214, 0.12070, 0.6663, 0.9624, 2.12\times10^{-9}, 0.0581, 0.06\,{\rm eV}\}$, giving an angular sound horizon at recombination θ_s (which defines the angular scale of the CMB acoustic peaks) roughly equal to $100\,\theta_s = 1.04075$. In this chapter, our total neutrino mass M_ν is assumed to be shared equally among the three species, like in the degenerate (DEG) model.

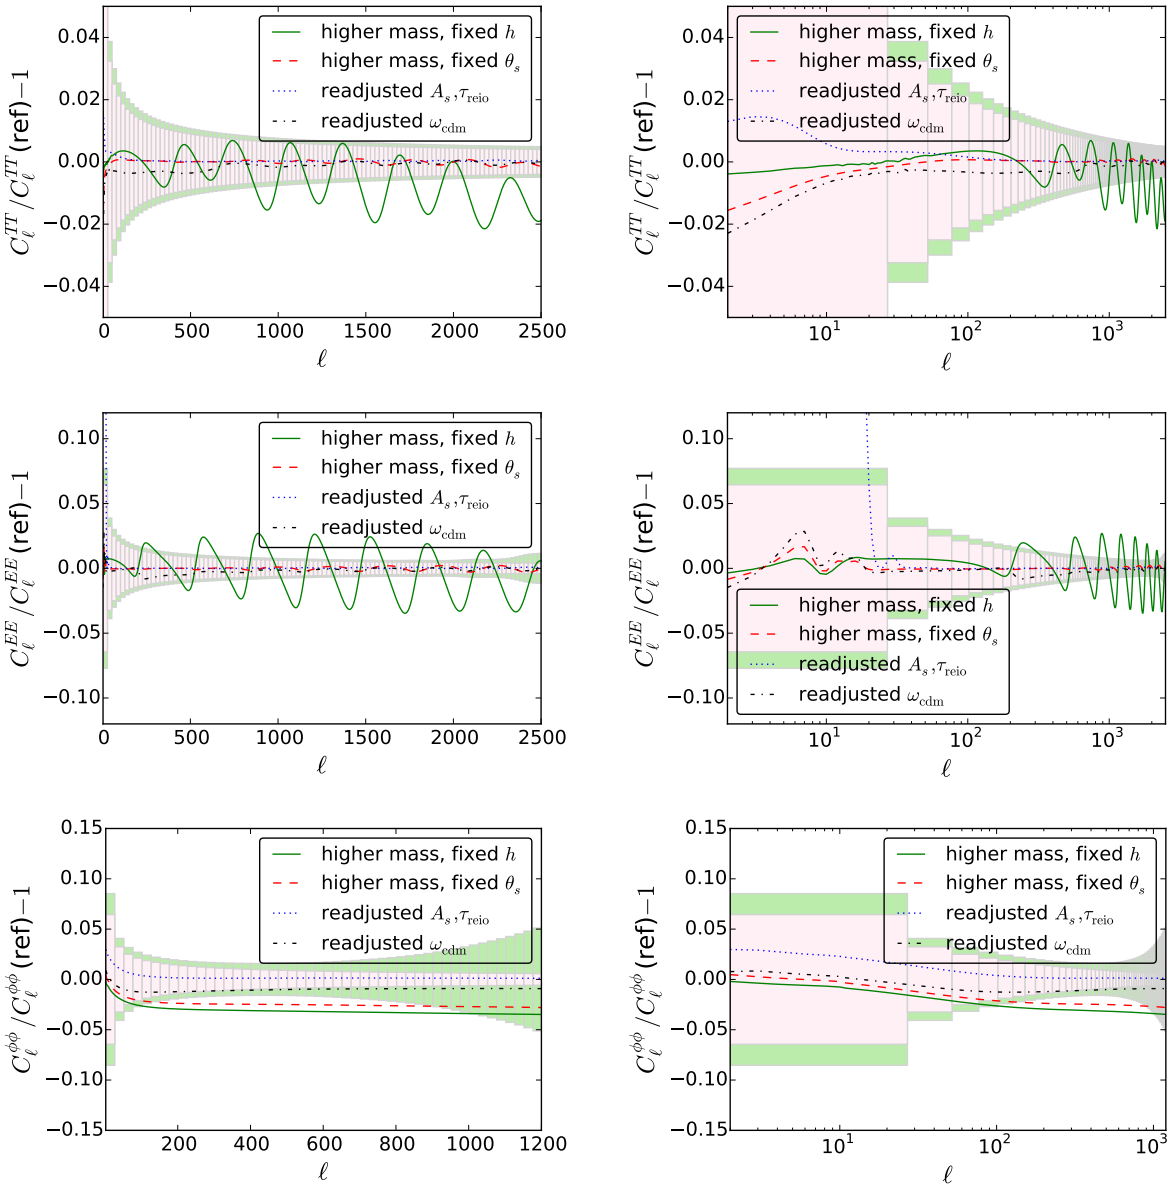

Figure 10: Relative change in the CMB spectra induced by increasing the summed neutrino mass from $M_\nu = 60$ meV to $M_\nu = 150$ meV. The plots show the residuals of the lensed TT (top), lensed EE (middle) and lensing potential (bottom) power spectrum, as a function of multipoles ℓ with a linear (left) or logarithmic (right) scale. The light/pink and darker/green shaded rectangles refer, respectively, to the binned noise spectrum of a cosmic-variance-limited or CORE-like experiment, with linear bins of width $\Delta\ell = 25$. The physical baryon density ω_b and the scalar spectral index n_s are kept fixed. In the first case (green solid line) the value of the Hubble constant is fixed at the reference value, while in all the other cases (labeled as fixed θ_s) h decreases in order to keep θ_s consistent with the reference model. Moreover, in the third case (dotted blue line), we tried to compensate for the changes in the lensing spectrum by increasing A_s, and in the fourth case (dotted-dashed black) we aim at the same result by increasing $\omega_{\rm cdm}$.

150 meV, and that in a CMB analysis the parameters (M_ν, H_0) are inevitably correlated, as it is well known, and illustrated by the upper left plot in Figure 11.

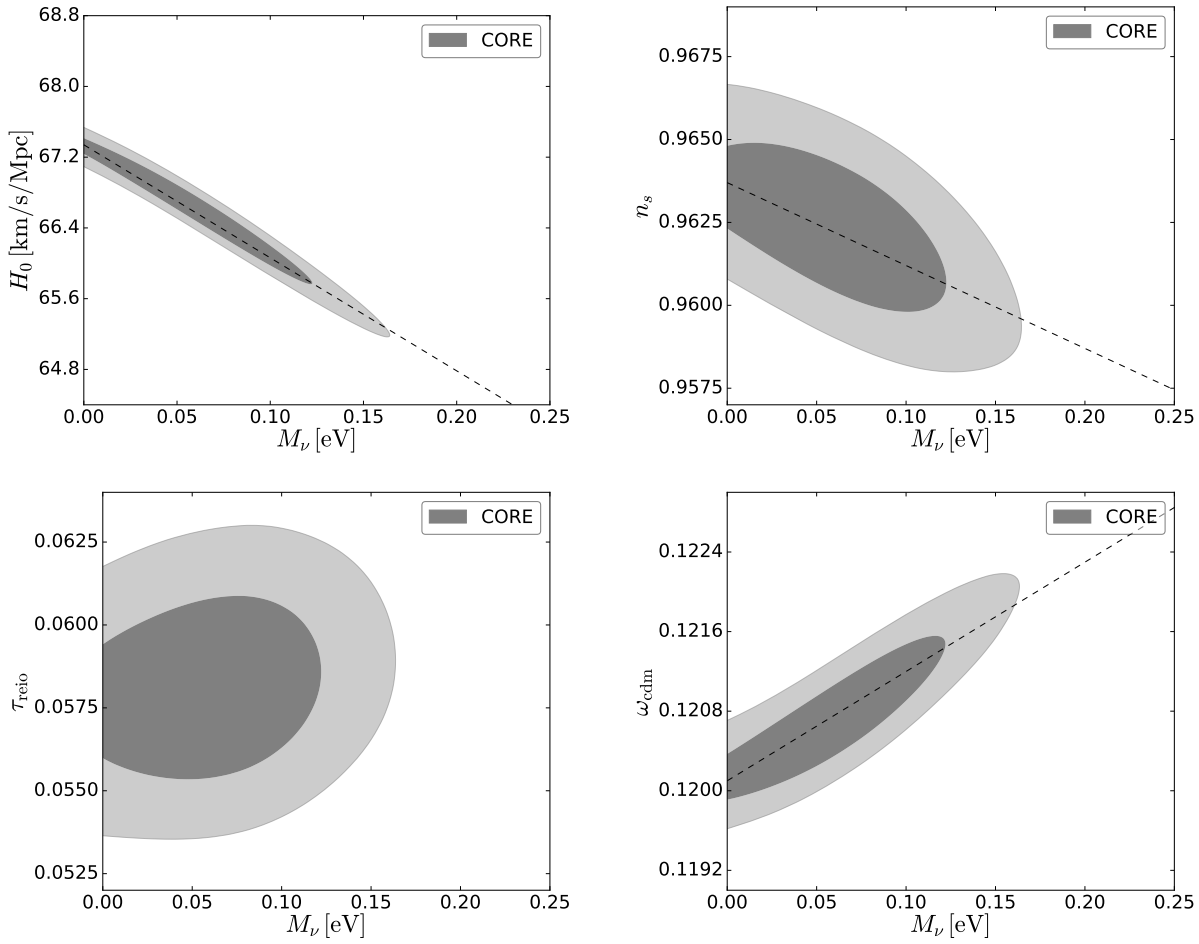

Figure 11: 68% and 95% CL posterior probability contour levels for different pairs of parameters, for an MCMC forecast of the sensitivity of a CORE-like experiment to the parameters of a 7-parameter model (ΛCDM plus total neutrino mass M_ν). The CMB data is assumed to consist of measurements of the TT, EE, TE and lensing potential spectra.

We can try to quantify this correlation. A simple numerical exercise shows that in order to keep the same value of θ_s while fixing $\{\omega_b, \omega_{\rm cdm}\}$ and varying M_ν, one finds a correlation

$$\Delta h \simeq -0.09 \left(\frac{\Delta M_\nu}{1\ \text{eV}}\right) . \tag{277}$$

We will come back to this relation later, and show that the correlation angle changes slightly when other effects are taken into account.

We now look at the bottom plots in Figure 10, showing variations in the lensing potential spectrum. The dashed red line is consistent with the fact that a higher neutrino mass implies more suppression in the small-scale matter power spectrum $P(k,z)$, and hence in the large-ℓ lensing potential spectrum $C_l^{\phi\phi}$. A comparison with the instrumental errors show that this effect is potentially relevant: the dashed red residual is outside the 68% error bars in about 30 consecutive bins, leading to a χ^2 increase by many units. It is also visible that the neutrino mass effect would be detectable only in a range given roughly by $50 \leq \ell \leq 800$, in which the effect is nearly equivalent to a suppression by some ℓ-independent factor (by about 3% in our example).

Hence, we see that a CORE-like CMB experiment could in principle discriminate between $M_\nu = 60$ meV and 150 meV, and that the effect of the neutrino mass with fixed $\{\omega_b, \omega_{\rm cdm}, \theta_s, n_s, A_s, \tau_{\rm reio}\}$ can be simply summarised as an apparent mismatch between the normalisation of the TT,TE,EE spectra and that of the CMB lensing spectrum. To check whether the distinction can be made in reality, and not just in principle, we must think whether the variation of other cosmological parameters could cancel this effect, and lead to new parameter correlations with M_ν.

As explained in references [207, 208], in a pure ΛCDM model with no massive neutrinos, the dependence of the global amplitude of $C_\ell^{\phi\phi}$ on the cosmological parameters is given approximately by:

$$\ell^4 C_\ell^{\phi\phi} \propto A_s \left(\Omega_m^{0.6}\, h\right)^{2.5} \quad (\ell > 200)\,,$$

and in terms of ω_m

$$\ell^4 C_\ell^{\phi\phi} \propto A_s \omega_m^{3/2} h^{-1/2} \quad (\ell > 200)\,,$$

plus an additional minor dependence on Ω_m. If we include massive neutrinos, the linear growth of structure becomes scale dependent, thus the exact impact of M_ν on $C_\ell^{\phi\phi}$ is ℓ-dependent, but only by a small amount in the range constrained by observations. Anyway, given that the neutrino mass slows down the growth of cold dark matter perturbations, we can generally assume:

$$\ell^4 C_\ell^{\phi\phi} \propto A_s \omega_m^{3/2} h^{-1/2} M_\nu^{-\alpha}\,, \tag{278}$$

with $\alpha > 0$. This qualitative result shows that in order to compensate an increase of M_ν, we have a priori two possibilities: increasing A_s, or increasing ω_m. We will explore them one after each other in the next points, and arrive at interesting conclusions.

3. We have the possibility to increase A_s in order to compensate for the neutrino mass effect in $C_\ell^{\phi\phi}$, while keeping $A_s e^{-2\tau_{\rm reio}}$ fixed, in order to have the same overall normalisation of the large-ℓ temperature and polarisation spectra. Hence, this transformation implies a higher reionisation optical depth $\tau_{\rm reio}$. We could expect that, this change in the optical depth is unobservable due to cosmic variance, which would mean that there is a parameter degeneracy at the level of CMB data, and that the three parameters $(M_\nu, A_s, \tau_{\rm reio})$ are correlated.

   This turns out not to be the case. In our example, the higher neutrino mass shifts the lensing potential down by 3%. This could be compensated by increasing A_s by 3% as well, and shifting $\tau_{\rm reio}$ by $\Delta\tau_{\rm reio} = \frac{1}{2}\log 1.03 \simeq 0.015$. This is a very big shift compared to the expected sensitivity of a CORE-like experiment, $\sigma(\tau_{\rm reio}) \simeq 0.002$. Hence this degeneracy should not be present.

   This is illustrated by the third set of curves (dotted blue) in Figure 10. We estimated numerically the reduction factor for $C_{400}^{\phi\phi}$ in the second model (red dashed). We increased A_s by exactly this factor, keeping $A_s e^{-2\tau_{\rm reio}}$ fixed. The new model has a much larger reionisation bump in C_ℓ^{EE}, with a residual largely exceeding the error bars.

   The lower left plot in Figure 11 brings the final confirmation that in a global fit of CMB data, with lensing extraction included, there is no significant correlation between M_ν and $\tau_{\rm reio}$.

   At this point, we still expect that very small neutrino masses could be accurately measured by CMB data alone, unless the other way to compensate for the neutrino mass

effect in the lensing potential (by increasing ω_m) works better than increasing A_s, and does lead to some parameter degeneracy. This is what we will explore in the final step of this discussion.

4. Considering that ω_b is accurately determined by the first peak ratios, we can increase $\omega_{\rm m} = \omega_{\rm b} + \omega_{\rm cdm}$ by enhancing $\omega_{\rm cdm}$ only. It is difficult to infer analytically from equation (278) the amount by which $\omega_{\rm cdm}$ should be enhanced in order to cancel the effect of M_ν in the lensing potential, because during the transformation, we must keep θ_s fixed; since θ_s depends on both h and $\omega_{\rm m}$, the Hubble parameter will also change. In the example displayed in Figure 10, we found numerically the factor by which we should increase $\omega_{\rm cdm}$ (with fixed $\omega_{\rm b}$ and θ_s), in order to nearly cancel the neutrino mass effect in the lensing power spectrum. This leads to the dotted-dashed black curve. In the lensing potential plots (bottom), the new residual is back inside the cosmic variance band.

   The problem with the previous attempt was that changing $\tau_{\rm reio}$ had "side effects" (namely, on the reionisation bump) potentially excluded by the data. Increasing $\omega_{\rm cdm}$ also has "side effects": it affects the redshift of radiation/matter equality $z_{\rm eq}$, and hence the amplitude of the first two peaks (through gravity boost effects and through the early ISW effect); it also affects the redshift of matter/Λ equality z_Λ and the late ISW effect; and finally, it has a small impact on the angular diameter distance. All these effects can be identified by looking at the details of the dotted-dashed black residuals in Figure 10. The key point is that a tiny enhancement of $\omega_{\rm cdm}$ is enough to compensate for the neutrino mass effect in $C_\ell^{\phi\phi}$, in such way that the "side effects" all remain well below cosmic variance. Hence, we expect a parameter degeneracy between M_ν and $\omega_{\rm cdm}$ when using CMB data alone, that will compromise the accuracy with which the neutrino mass can be pinned down, and lead to a correlation between these parameters. We notice that this correlation between M_ν and $\omega_{\rm cdm}$ is completely driven by CMB lensing. Removing lensing extraction would diminish the correlation factor. The residual correlation would be due to the lensing of the C_ℓ^{TT} spectrum (related to the tiny deviation of the black dot dashed line from the red dashed line on small scales in the top left panel of Figure 10), and it would disappear with delensing.

   This is confirmed by the lower right plot in Figure 11: in a global fit of CMB data, we obtain a degeneracy direction approximately parametrised by the slope of the dashed curve in that plot,

$$\Delta\omega_{\rm cdm} = 0.01\,\Delta M_\nu \sim \Delta\omega_\nu\ . \tag{279}$$

   Which is exactly the relation we used in Figure 10, when transforming to the fourth model (dotted-dashed black curves).

We can reach the main conclusion of this section: for CMB data alone (including lensing extraction), there is no significant parameter degeneracy between $(M_\nu, A_s, \tau_{\rm reio})$, but there is one between M_ν and $\omega_{\rm cdm}$. This is the most pronounced parameter degeneracy involving the neutrino mass when the cosmological model is parametrised by $\{\omega_b, \omega_{\rm cdm}, \theta_s, n_s, A_s, \tau_{\rm reio}\}$, and the correlation is given approximately by equation (279).

If instead the model is parametrised by $\{\omega_b, \omega_{\rm cdm}, h, n_s, A_s, \tau_{\rm reio}\}$, for the obvious reasons discussed previously, there is an additional clear correlation between M_ν and h. We return to the correlation factor, that we estimated before to be given by equation (277). This equation is

actually not a very good fit of the contours in the upper left plot of Figure 11: the dashed line in that plot corresponds to

$$\Delta h \simeq -0.13 \left(\frac{\Delta M_\nu}{1\,\mathrm{eV}}\right) . \tag{280}$$

The explanation for this mismatch is simple. Eq. (277) assumed fixed θ_s *and* $\omega_{\rm cdm}$ values. If instead we try to keep θ_s fixed while varying $\omega_{\rm cdm}$ according to equation (279), we see increased correlation between M_ν and h, as shown by equation (280)[4].

Hence, with CMB data only, the clearest and most important degeneracies involving the summed neutrino mass are between M_ν and $\omega_{\rm cdm}$ (due to lensing) and M_ν and h (due to the angular diameter distance). There are other correlations, but they are much less pronounced. The third one would be between M_ν and n_s [180]. This can be understood by looking closely at the dotted-dashed in Figure 10 (lower right plot). The variation of $\omega_{\rm cdm}$ did not only rescale the amplitude of the CMB lensing potential, it also generated a small positive tilt. The reason is that we have decreased the ratio $\omega_{\rm b}/\omega_{\rm cdm}$, thus changing the shape parameter controlling the effective spectral index of the matter power spectrum $P(k)$ for $k > k_{\rm eq}$: a smaller baryon amount relative to CDM implies a bluer spectrum. Hence, the (M_ν, $\omega_{\rm cdm}$) degeneracy is more pronounced when it goes together with a tiny decrease of the tilt n_s, by such a small amount that it would not conflict with temperature and polarisation data. This negative correlation is visible in Figure 11, upper right plot.

5.2 EFFECT OF NEUTRINO MASS ON THE BAO SCALE

The acoustic oscillations of the baryon-photon fluid that we observe in the CMB power spectrum produce a characteristic feature in the two point correlation function. In Fourier space the feature is located at a peculiar scale, the BAO scale, $k_{\rm BAO} = 2\pi/r_s(z_{\rm drag})$, where $r_s(z_{\rm drag})$ is the comoving sound horizon at baryon drag

$$r_s(z_{\rm drag}) = \int_0^{\tau_{\rm drag}} c_s d\tau = \int_{z_{\rm drag}}^{\infty} \frac{c_s}{H(z)} dz .$$

The observed scale, assuming an isotropic fit of a galaxy sample[5], provides the ratio $\dfrac{r_s(z_{\rm drag})}{D_V(z_{\rm BAO})}$, where D_V is the volume distance, defined as

$$D_V(z) = \left[z/H(z)(1+z)^2 d_A(z)^2\right]^{1/3} ,$$

and $D_A = (1+z)d_A(z)$ is the comoving angular diameter distance. In the ΛCDM model with massive neutrinos, the ratio $r_s(z_{\rm drag})/D_V(z_{\rm BAO})$ can only depend on the four parameters $\{\omega_{\rm b}, \omega_{\rm cdm}, \omega_\nu, h\}$. More precisely, $r_s(z_{\rm drag})$ depends on the three parameters $\{\omega_{\rm b}, \omega_{\rm cdm}, h^2\}$, while for redshifts below the non-relativistic transition, $z \ll z_{\rm nr} \sim 2\times 10^3 (m_\nu/1\,\mathrm{eV})$, $D_A(z)$ depends only on $\omega_{\rm tot} = \omega_{\rm b} + \omega_{\rm cdm} + \omega_\nu$ and on h, because it can be approximated as

$$D_A(z) = \int_0^z \frac{c dz'}{H(z')} \simeq 3000 \int_0^z \frac{dz'}{\sqrt{\omega_{\rm tot}(1+z')^3 + (h^2 - \omega_{\rm tot})}}\,\mathrm{Mpc} . \tag{281}$$

4 Note that we estimated the correlation factor in equation (279) with one significant digit, and in equation (280) with two significant digits: this is consistent with the fact that the correlation is much more clear and pronounced in the second case (the ratio of the minor over major axis is much smaller).

5 Anisotropic fit allow to disentangle the longitudinal information (i.e. the radial scale Hr_s) from the transverse one (i.e. the tangential scale D_A/r_s).

Note that the term inside the square root is a polynomial in z' in which the constant term is precisely h^2 (so as expected, for small redshifts $z \ll 1$, $D_A(z)$ depends *only* on the h parameter, like in a Hubble diagram).

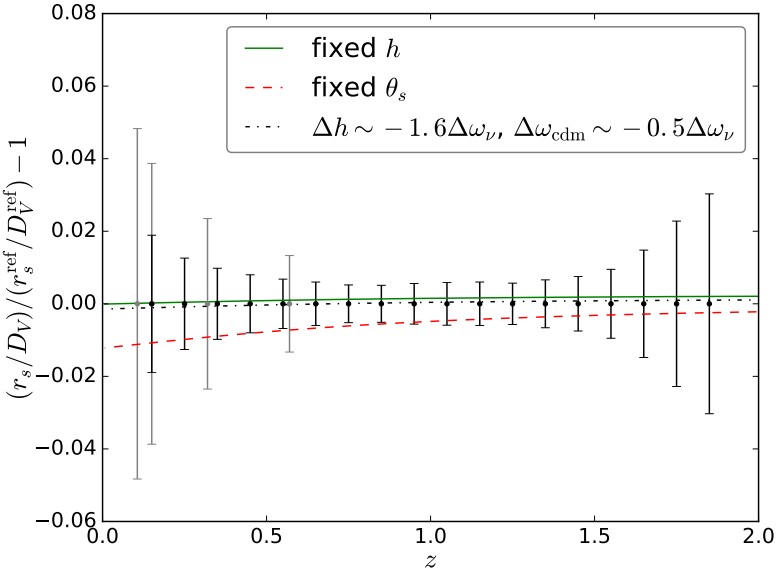

Figure 12: Relative error on r_s/D_V. Gray error bars refer to the current BAO measurements: from left to right 6dFGRS [209], SDSS MGS [210], LOW-Z, C-MASS [211]. Black error bars mark the expected sensitivity of the future DESI experiment [199, 212]. Green solid line and red dashed lines are the same as in figure 10, i.e. higher M_ν with fixed h (green solid line) and higher M_ν with fixed θ_s and varying h (red dashed line). However, here the black dot dashed line is obtained by increasing M_ν and varying h and $\omega_{\rm cdm}$ as in Eqs. (282).

In Figure 12 we show the residuals of current and future BAO measurements, taking as a reference the same model as before with $M_\nu = 60$ meV, as well as the relative difference on $r_s(z_{\rm drag})/D_V(z_{\rm BAO})$ between several models with a higher mass $M_\nu = 150$ meV (already introduced in section 5.1.3) and the reference model. For future measurement we take the example of DESI, assuming the same sensitivity as in Refs. [199, 212].

We first vary only M_ν with fixed $\{\omega_b, \omega_{\rm cdm}, h, n_s, A_s, \tau_{\rm reio}\}$ (green solid line). This means that the early cosmological evolution is identical, while the matter density is slightly enhanced at late times (after the neutrino non-relativistic transition), by about one percent. Thus $d_s(z_{\rm dec})$ and $r_s(z_{\rm drag})$ are fixed, but $d_A(z_{\rm dec})$, $d_A(z_{\rm BAO})$ and $D_V(z_{\rm BAO})$ are subject to change. We have seen that this transformation shifted the CMB peaks by a detectable amount. However, the accuracy with which CORE will measure θ_s ($< 0.01\%$) is much greater than that with which DESI will measure the BAO angular scales ($\sim 1\%$). From Eq. 281 we can see analytically that the typical variation of $D_A(z)$ between the two models is negligible for $z \ll 1$ and of the order of $\frac{1}{2}\frac{\Delta\omega_\nu}{\omega_{\rm tot}} \simeq 0.25\%$ for $1 < z < z_{\rm nr}$. This explains why the green curve in Figure 12 remains within the BAO error bars.

This preliminary discussion brings us to the key points of this section:

- *the BAO data alone can bound the neutrino mass, but not with great accuracy.* We showed previously that increasing M_ν with fixed $\{\omega_b, \omega_{\rm cdm}, h\}$ had no detectable effects, but this was because the mass variation was too small. If one keeps increasing ω_ν with the other parameters fixed, the function inside the square root in Eq. 281 keeps increasing

for the same value z', and $D_A(z)$ decreases. To avoid a BAO bound on M_ν, one could try to exactly compensate the variation $\Delta\omega_\nu$ by an opposite variation in either $\omega_{\rm b}$ or $\omega_{\rm cdm}$, to keep $D_A(z)$ exactly constant. But in that case, the early cosmological evolution would change (sound speed, redshift of equality, redshift of baryon drag) and the ratio $r_s(z_{\rm drag})/D_V(z_{\rm BAO})$ would be shifted anyway. Hence there is no parameter degeneracy cancelling exactly the effect of M_ν in BAO observables, at least in the ΛCDM+M_ν model. This explains why in Figures 13 and 14, the contours involving M_ν are closed for DESI data alone, setting an upper bound on the summed mass of a few hundreds of meV.

- *the strong degeneracy between M_ν and h observed in the CMB case cannot exist with BAO data.* This denegeracy came from the possibility to keep constant angular scales ($\theta_s(z_{\rm dec})$, $\theta_d(z_{\rm dec})$) by varying h with fixed $\{\omega_b, \omega_{\rm cdm}\}$. Indeed, when fitting CMB data with different neutrino masses, one can keep the same value of $d_A(z_{\rm dec})$ by altering the late time cosmological evolution: while M_ν tends to enhance the density at late times, one can decrease h and the cosmological constant in order to compensate for this effect. This cannot be done with BAO data, because they probe $d_A(z)$ at several small values of z, comparable to the redshift of the transition z_Λ. The proof is particularly obvious if we look at Eq. 281 again. Whatever change in h modifies the constant term inside the square root, and thus the value of $D_A(z)$ for $z \leq 1$. Thus the (M_ν, h) degeneracy discussed in the CMB section must be broken by BAO data. We get a first confirmation of this by looking at the red dashed curve in Figure 12, obtained by increasing M_ν with a constant $\theta_s(z_{\rm dec})$: the new model departs from the other one by a detectable amount, at least given BAO-DESI errors (especially at $z \ll 1$, as expected from this discussion). The second confirmation comes from the right plot in Figure 13, showing very different correlations between M_ν and h for CMB-CORE alone and BAO-DESI alone.

- *there exists, however, a correlation between M_ν, h and $\omega_{\rm cdm}$ with BAO data, but along different angles than with CMB data.* This comes from the possibility to modify parameters in such a way that both $r_s(z_{\rm drag})$ and $D_V(z_{\rm BAO})$ get shifted, but almost by the same relative amount. To compensate for the effect of an increasing ω_ν, one has three parameters to play with: $\{\omega_b, \omega_{\rm cdm}, h\}$. However, ω_b is precisely fixed by CMB data alone, and for that reason we keep it to its Planck best-fit value. We then find that variations of the other two parameters by approximately

$$\Delta\omega_{\rm cdm} \sim -0.5\Delta\omega_\nu\ , \qquad \Delta h \simeq -0.017\left(\frac{\Delta M_\nu}{1\ {\rm eV}}\right) \simeq -1.6\,\Delta\omega_\nu \tag{282}$$

  achieve a nearly constant ratio $r_s(z_{\rm drag})/D_V(z_{\rm BAO})$ in the redshift range best probed by the BAO-DESI experiment. As argued before, this ratio is more sensitive to h than $\omega_{\rm cdm}$ in that range, so the correlation between $\omega_{\rm cdm}$ and ω_ν is weak, while that between h and ω_ν is strong (see Figure 13).

The parameter correlations found in eq. (282) for BAO data are very different from those found in the previous section for CMB data:

$$\Delta\omega_{\rm cdm} \sim \Delta\omega_\nu\ , \qquad \Delta h \simeq -0.13\left(\frac{\Delta M_\nu}{1\ {\rm eV}}\right) \simeq -12\,\Delta\omega_\nu\ . \tag{283}$$

The combination of CMB and BAO data can thus break these degeneracies, as it is often the case when combining high and low redshift probes of the expansion history. The breaking

does not arise from the joint measurement of $\omega_{\rm cdm}$ and ω_ν (because BAO data are much less sensitive to $\omega_{\rm cdm}$ alone than CMB data), but from that of h and ω_ν, for which the different directions of degeneracy appear very clearly on Figure 13. Thus, the future BAO-DESI data will contribute to tighter constraints on M_ν.

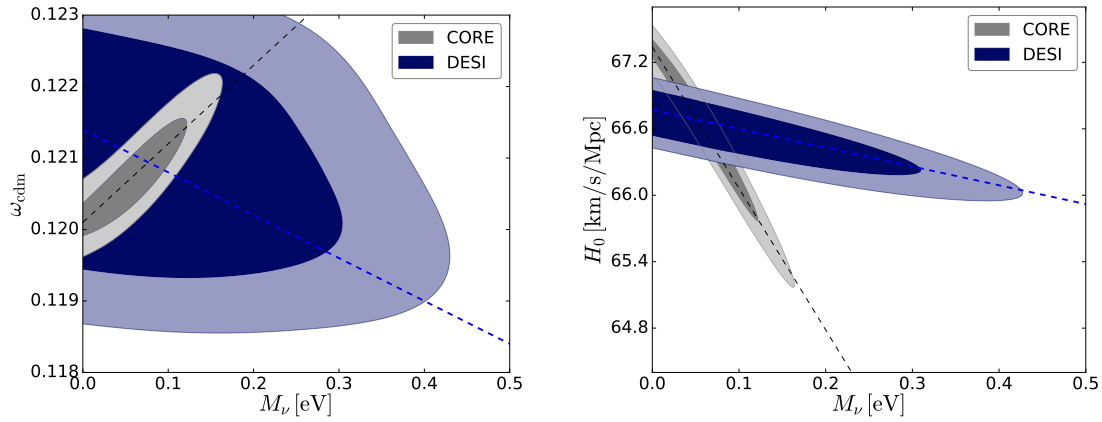

Figure 13: Marginalized one- and two- σ contours in the plane $(\omega_{\rm cdm}, M_\nu)$ (left panel) and (H_0, M_ν) (right panel), for CMB-CORE or BAO-DESI mock data. The black dashed lines show the directions of degeneracy given in Eqs. (283), and the blue ones in Eqs. (282).

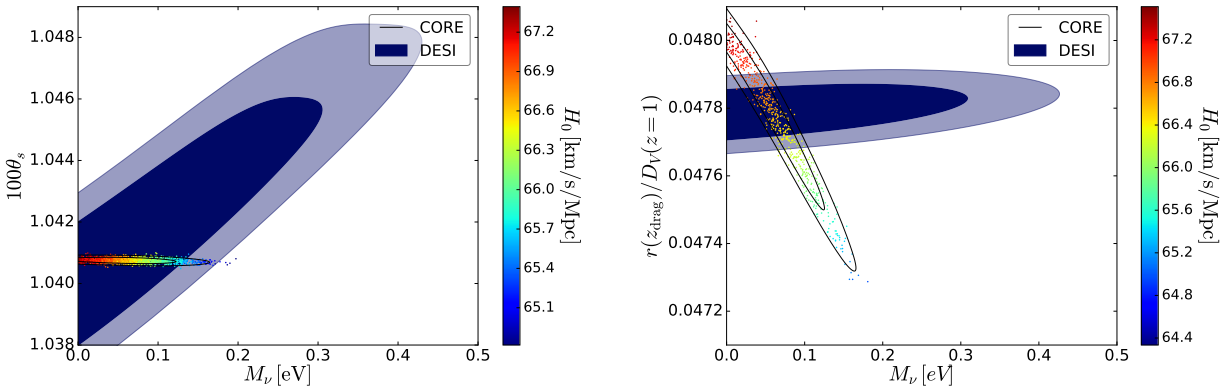

Figure 14: Marginalized one- and two- σ contours in the plane $(\theta_s(z_{\rm dec}), M_\nu)$ (left) and $\left(r(z_{\rm drag})/D_V(z=1), M_\nu\right)$ (right), for CMB-CORE or BAO-DESI mock data. In the CORE contours, samples are coloured according to the value of H_0.

Another way to illustrate the degeneracies discussed here is to fit CMB data or BAO data alone with a ΛCDM+M_ν model, and to plot the results in the space of parameters $(M_\nu, \theta_s(z_{\rm dec}))$ and $(M_\nu, r_s(z_{\rm drag})/D_V(z_{\rm BAO}))$ for a median redshift $z_{\rm BAO} = 1$. This is shown in Figure 14. When fitting CMB alone, thanks to the degeneracy of Eqs. (283), we can increase M_ν while keeping $d_A(z_{\rm dec})$ and $\theta_s(z_{\rm dec})$ fixed (left plot), but this is at the expense of decreasing the BAO angular scale by more than allowed by observational errors (right plot). Conversely, when fitting BAO data alone, we can play with the degeneracy of Eqs. (282) to keep the BAO angular scale fixed, but this requires $\theta_s(z_{\rm dec})$ to vary. The right plot in Figure 14 illustrates, in an alternative way to the right plot of Figure 13, how the combination of the two data sets can improve neutrino mass bounds.

Finally, we expect, as a secondary indirect effect, that the correlation between M_ν and $(A_s, \tau_{\rm reio})$ will be more noticeable in a combined analysis of CMB and BAO than for CMB alone. In section 5.1.3, we mentioned in points 3 and 4 that the impact of M_ν on CMB lensing could be compensated in two ways: by increasing either $(A_s, \tau_{\rm reio})$ (point 3) or $\omega_{\rm cdm}$ (point 4). We explained why the former option is favoured with CMB data alone. Since we just argued that BAO data can reduce the degeneracy between neutrino masses and $\omega_{\rm cdm}$, the latter option is more relevant when the data are combined with each other. Indeed, we will see a small correlation between $(M_\nu, \tau_{\rm reio})$ in the combined results presented in sections 8.2 and 8.3, one that was hardly noticeable with CMB alone. Of course, this degeneracy is not perfect, and extends only up to the point at which $\tau_{\rm reio}$ becomes too large to be compatible with CMB polarisation data.

5.3 EFFECT OF NEUTRINO MASS ON LARGE SCALE STRUCTURE OBSERVABLES

5.3.1 *Cosmic shear and galaxy clustering spectrum*

The Euclid satellite, whose launch is scheduled for 2020, will provide the most accurate ever galaxy redshift survey, measuring cosmological observables, such as cosmic shear and galaxy clustering, with 1% accuracy. Euclid data will certainly lead to a major breakthrough in precision cosmology thanks to very precise low redshift measurement which will break the CMB degeneracies among cosmological parameters (see Refs. [117, 118, 178, 183, 213–217]). Here we use the information extracted from the cosmic shear power spectrum projected in angular harmonics (2D) and the galaxy clustering power spectrum (3D). Both observable are related to the non-linear matter power spectrum depending on wavenumber and redshift, $P_{\rm m}(k, z)$. In our forecasts, we estimate this quantity using the HALOFIT algorithm, updated by Ref. [105] and also by Ref. [152] for the effect of neutrino masses, as implemented in CLASS v2.5.0.

Cosmic shear. The cosmic shear auto and cross correlation angular power spectrum in the i and j redshift bins is given in the Limber approximation by:

$$C_\ell^{ij} = H_0^4 \int_0^\infty \frac{dz}{H(z)} W_i(z) W_j(z) P_{\rm m}\left(k = \frac{l}{r(z)}, z\right), \tag{284}$$

where the window functions are given by

$$W_i(z) = \frac{3}{2}\Omega_m(1+z) \int_0^\infty dz_s \frac{n_i(z_s)(r(z_s) - r(z))}{r(z_s)}, \tag{285}$$

and the number of galaxies per steradian in the i bin is given by

$$n_i(z) = \frac{\int_{z_i^{\rm min}}^{z_i^{\rm MAX}} dn/dz \mathcal{P}(z, z_{\rm ph}) dz_{\rm ph}}{\int_0^\infty dn/dz \mathcal{P}(z, z_{\rm ph}) dz_{\rm ph}}$$

with $\mathcal{P}(z, z_{\rm ph})$ being the error function

$$\mathcal{P}(z, z_{\rm ph}) = \frac{1}{\sqrt{2\pi\sigma_{\rm ph}^2}} \exp\left[-\frac{1}{2}\left(\frac{z - z_{\rm ph}}{\sigma_{\rm ph}}\right)\right].$$

We use the Euclid prescription for the galaxy surface density

$$dn/dz = z^2 \exp\left[-(z/z_0)^{1.5}\right]$$

with $z_{\rm mean} = 1.412 z_0$. Finally, we consider a photometric redshift error $\sigma_{\rm ph} = 0.05(1+z)$, sky fraction $f_{\rm sky} = 0.3636$, mean internal ellipticity 0.22 and total number of observed galaxies 30 per arcmin^2.

Galaxy clustering. For galaxy clustering the observed power spectrum reads:

$$P(k_{\rm ref}, \mu, z) = \frac{D_A(z)^2_{\rm ref} H(z)}{D_A(z)^2 H(z)_{\rm ref}} b(z)^2 \left[1 + \beta(z, k(k_{\rm ref}, \mu, z))\mu^2\right]^2 \times P_{\rm m}(k(k_{\rm ref}, \mu, z)) e^{-k(k_{\rm ref}, \mu, z)^2, z)\mu^2 \sigma_r^2},$$

where μ is the cosine of the angle between the line of sight and the wavenumber in the reference cosmology (ref) $k_{\rm ref}$, k is the wavenumber in the true cosmology and it is defined as a function of $k_{\rm ref}$

$$k^2 = \left(\frac{(1-\mu^2) D_A(z)^2_{\rm ref}}{D_A(z)^2} + \frac{\mu^2 H(z)^2}{H(z)^2_{\rm ref}} \right) k^2_{\rm ref}.$$

The factor $\left[D_A(z)^2_{\rm ref} H(z)\right] / \left[D_A(z)^2 H(z)_{\rm ref}\right]$ encodes the geometrical distortions related to the Alcock-Paczynski effect. The bias can be written as $b = \sqrt{(1+z)}$, β encodes the redshift space distortions

$$\beta(k, z) = \frac{1}{2b(z)} \frac{d \ln P_{\rm m}(k, z)}{d \ln a},$$

and finally the spectroscopic redshift error is $\sigma_r = dr(z)/dz \sigma_z$.

Both the C_ℓ^{ij} and the $P(k)$ provide information on a broad range of scales; therefore, given the same survey sensitivity, they are more efficient than BAO in constraining cosmological parameters; however, for the very same reason, they are more prone to systematic effects such as residual errors in the estimate of non-linear corrections, non-linear light-to-mass bias or redshift space distortions (see e.g. Ref. [218, 219]). For that reason, we include in our forecast a theoretical error on the observable power spectrum, increasing above a given redshift-dependent scale of non-linearity (see Ref. [118], or Ref. [119] for a more refined treatment).

The assumed theoretical error amplitude has a direct impact on the galaxy clustering sensitivity to cosmological parameters. Here we stick to the approach of Ref. [118], and we refer to this work for details and equations. As emphasised in Ref. [119], this approach is extremely (and maybe overly) conservative, because the error is assumed to be uncorrelated between different k-bins. The error grows as a function of the ratio $k/k_{\rm nl}(z)$, where $k_{\rm nl}(z)$ is the redshift-dependent scale of non linearity, with a shape and amplitude inspired from the typical residuals between different N-body codes[6]. Choosing a value for the error amplitude parameter ϵ amounts to estimating the accuracy of grids of N-body simulations and of models for various non-linear and systematic effects in a few years from now. The baseline choice in Ref. [118] was $\epsilon = 0.05$. In this paper, given the progress in the field observed since 2012, we choose to reduce it to $\epsilon = 0.025$, meaning that the uncorrelated theoretical error saturates at the 2.5% level in the deep non-linear regime. This error is explicitly shown in Figure 16 for $z = 0.5$ and $z = 2$, and its impact on the lensing harmonic spectrum appears in Figure 15 for the lowest and highest redshift bins of the Euclid lensing survey. In presence of a theoretical error, the issue of where to cut the integrals in the galaxy clustering likelihood becomes hardly relevant, provided that the cut-off is chosen in the region where the theoretical error dominates. In what follows, we will cut the observable $P(k_{\rm ref}, \mu, z)$ at $k_{\rm max} = 0.6\, h/{\rm Mpc}$ for all redshifts. For

6 The error function is explicitly given by $\alpha(k,z) \equiv \frac{\Delta P_{\rm m}(k,z)}{P_{\rm m}(k,z)} = \frac{\ln[1+k/k_{\rm nl}(z)]}{1+\ln[1+k/k_{\rm nl}(z)]}\epsilon$, where $k_{\rm nl}(z)$ is identical to the quantity $k_\sigma(z)$ computed at each redshift by HALOFIT, and the error amplitude parameter ϵ is the unique free parameter in this model. The asymptotic error in the deep non-linear regime is then given by $100\epsilon\,\%$.

cosmic shear, the inclusion of the theoretical error is also important, although the observational error bar does not decreases indefinitely with ℓ due to the finite angular resolution of the shear maps. In our forecasts, we perform a cut at $\ell_{\rm max} = 2000$.

5.3.2 *Degeneracies between M_ν and other parameters*

In Figure 15 and 16 we show the relative shift in the shear power spectrum and in the galaxy power spectrum that is obtained when increasing the summed neutrino mass while keeping various quantities fixed. We also show for comparison the observational and theoretical errors computed in the same way as in Ref. [118], using the survey specifications listed above. We will study the impact of increasing the summed neutrino mass on these observables: (1) when keeping the usual cosmological parameters fixed, (2) when tuning h at the same time in order to keep the same angular peak scale in the CMB, and (3) when playing with other parameters in order to minimize the impact of neutrino mass on LSS observables. The discussion in (2) (respectively, (3)) is relevant for understanding the degeneracy between M_ν and other parameters when fitting CMB+LSS data (respectively, LSS data alone).

As in the previous sections, we will then check our theoretical conclusions through an MCMC forecast of the sensitivity of future experiments that will measure the spectra discussed above. In Figure 18 we plot the marginalized one- and two- σ contours showing the degeneracies at study: $(\omega_{\rm cdm}, M_\nu)$ (upper left panel), (H_0, M_ν) (upper right panel), (n_s, M_ν) (bottom left panel), (A_s, M_ν) (bottom right panel). The CORE only contours (in gray) are the same as in Figure 13. The Euclid related contours have been obtained through an MCMC forecast including either galaxy clustering (in green) or cosmic shear (in red), following the specifications listed in section 5.3.1. Fitting Euclid mock data alone would return wide contours in parameter space. Given that the two quantities best measured by CMB experiments are the angular scale of the acoustic horizon and the baryon density, the question in which we are most interested is: assuming that information on ω_b and θ_s is provided by a CORE-like CMB experiment, what is the pull on other parameters coming from Euclid alone? To address this, when fitting Euclid data, we impose two uncorrelated Gaussian priors on respectively ω_b and θ_s, with standard deviations taken from our previous CORE-CMB forecast, while keeping $\tau_{\rm reio}$ fixed, since the latter does not affect galaxy clustering and shear observables in any way.

1. *Neutrino mass effects with all standard cosmological parameters fixed: the usual neutrino–induced step–like suppression.*

   Like in the previous sections, we start by increasing the summed neutrino mass from $M_\nu = 0.06$ eV to $M_\nu = 0.15$ eV, keeping all the other cosmological parameters fixed $\{\omega_b, \omega_{\rm cdm}, h, n_s, A_s, \}$ (green solid line). Note that in most of the literature, the effect of neutrino masses on the matter power spectrum is discussed precisely in that way. One reason is that fixing $\{\omega_b, \omega_{\rm cdm}, n_s, A_s, \}$ amounts in keeping the same "early cosmological evolution" until the time of the neutrino non-relativistic transition. The choice to fix also h is mainly a matter of simplicity.

   As expected, the larger M_ν induces a relative suppression of power on small scales compared to large scales, visible both in the shear and in the galaxy power spectrum. To be precise, in the redshift range surveyed by Euclid, $0 < z < 2.5$, neutrinos with a mass $M_\nu > 0.05$ eV are already well inside the non-relativistic regime, thus, the spectrum is suppressed on scales smaller than the free-streaming scale $k > k_{\rm fs}(z)$. In the redshift range of interest, $0 < z < 2.5$, the free streaming wavenumber spans the range $[0.0077 -$

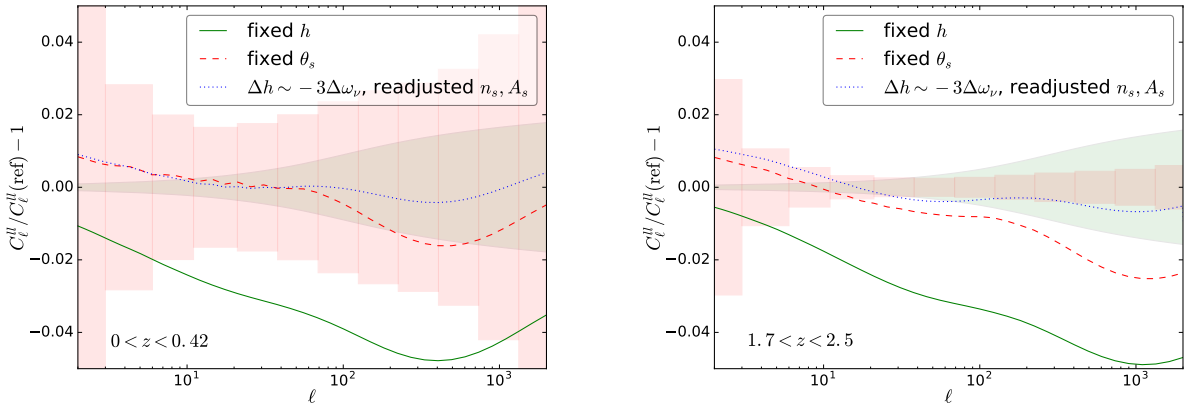

Figure 15: Relative error on the galaxy lensing C_ℓ^{ll} in the first redshift bin ($0 < z < 0.42$, left panel) and in the tenth redshift bin ($1.7 < z < 2.5$, right panel). Here the redshift range is $0 < z < 2.5$ and is divided in ten equi-populated redshift bins. The light pink rectangles refers to the observational error. The light green shaded area shows the relative error associated to our model for the theoretical uncertainty on $P_m(k, z)$. Green solid and red dashed lines are the same as in Figure 10, i.e. higher M_ν with fixed h (green solid line) and higher M_ν with fixed θ_s and varying h (red dashed line). The blue dotted line, besides the higher M_ν, implies a smaller value of h ($\Delta h \sim -3\Delta\omega_\nu$), an increase of n_s by 0.4% and of A_s by 2%.

$0.0041]\, h\,\mathrm{Mpc}^{-1}$ (respectively, $[0.0192 - 0.0103]\, h\,\mathrm{Mpc}^{-1}$) for $M_\nu = 0.06$ eV (respectively, $M_\nu = 0.15$ eV)[7]. The suppression in power makes both the C_ℓ^{ij} and the $P(k)$ directly sensitive to the neutrino mass sum, while this was not the case for the purely geometrical information encoded in BAO measurements.

This sensitivity is reinforced by non-linear effects which are well visible on Figures 15 and 16. In the shear spectrum of Figure 15, in absence of non-linear corrections, the green curve would be almost constant for $\ell > 100$. Non-linear gravitational clustering produces a characteristic "spoon shape" or dip [152]. The minimum of the dip is seen at $\ell \sim 40$ in the first redshift bin and $\ell \sim 1000$ in the last one. In Figure 16, non-linear effects are responsible for the further decrease of the green curve for $k \geq 0.1\, h/\mathrm{Mpc}$.

2. *Neutrino mass effects with h varied to keep the CMB angular scales fixed: why does LSS data lifts the (M_ν, h) degeneracy?*

   The second part of the discussion consists in increasing M_ν by the same amount, while varying h like in section 5.1.3, in such way as to keep a constant angular diameter distance to recombination, constant sound horizon angular scale, and constant damping angular scale (red dashed line). As we have seen in Section 5.1.3 this procedure leads to the well known (M_ν, h) CMB degeneracy.

   We showed that this degeneracy is broken by BAO data, because the lower value of h increases the angular diameter distance at low redshift (see Section 5.2). This conclusion is valid also for galaxy $P(k)$ and shear C_ℓ^{ij}, since the red dashed residuals in Figures 15, 16 are well outside the observational and theoretical error bars. For clarity, we should explain the shape of these red dashed lines, which is slightly counter-intuitive.

7 The free streaming length depends on the mass of each neutrino rather than on the sum. Here we have assumed three massive degenerate neutrinos.

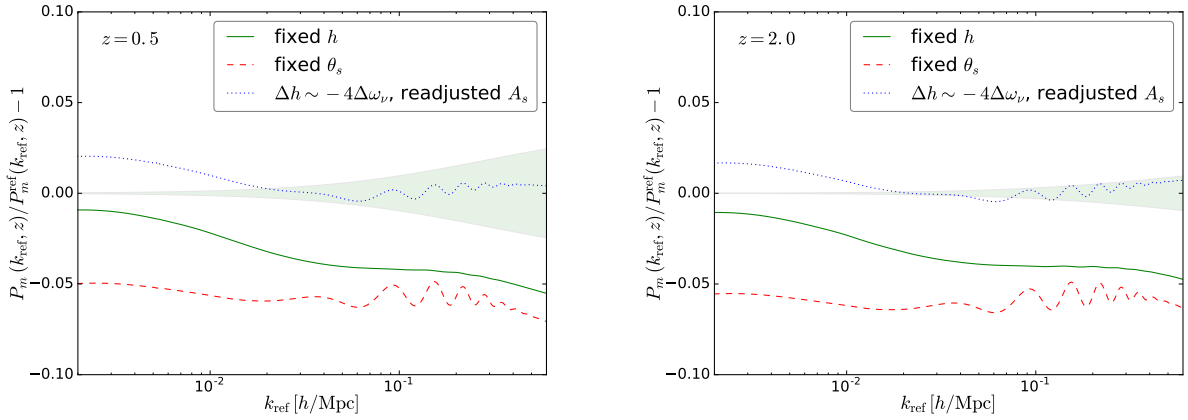

Figure 16: Relative error on the non linear matter power spectrum $P_{\rm m}(k_{\rm ref}, z)$ perpendicular to the line of sight ($\mu = 0$) at redshift $z = 0.5$ (left panel) and $z = 2$ (right panel). The light pink shaded area refers to the observational error, including cosmic variance. The light green shaded area shows our model for the theoretical uncertainty. Here the redshift range is $0.5 < z < 2$ and is divided in 16 redshift bins. Green solid and red dashed lines are the same as in Figure 10, i.e. higher M_ν with fixed h (green solid line) and higher M_ν with fixed θ_s and varying h (red dashed line). The blue dotted line, besides the higher M_ν, implies a smaller value of h ($\Delta h \sim -4\Delta\omega_\nu$) and an increase of A_s by 5%.

In the case of galaxy clustering, the higher value of M_ν and lower value of h lead to an almost constant suppression of power on every scale, plus some wiggles on small scales (see Figure 16). This may sound surprising since we are used to seeing more suppression on small scales when increasing the neutrino mass. This is true for fixed h, but here we are decreasing the Hubble rate at the same time. Since $\omega_{\rm m} = \Omega_{\rm m} h^2$ is kept fixed, this means that we are increasing $\Omega_{\rm m}$. For subtle reasons which can be understood analytically, the large-scale branch of the matter power spectrum is suppressed by the increase of $\Omega_{\rm m}$[8], while the small-scale branch is suppressed by massive neutrino free-streaming, coincidentally by roughly the same amount. This explains the almost constant suppression of power in the galaxy clustering spectrum (red dashed line, Figure 16). The wiggles located around $k \sim 0.1\,h\,{\rm Mpc}^{-1}$ are related to the shift of the BAO scale due to the different angular diameter distance at low redshift, as we have explained in section 5.2 (see also Ref. [220]).

The situation is a bit different for the galaxy lensing spectrum C_ℓ^{ij} (red dashed line, Figure 15) which probes metric fluctuations instead of matter fluctuations. As a result[9],

8 In order to understand the observed behaviour, we have to elaborate on the matter power spectrum $P_{\rm m}$ entering Eq. 5.3.1. An analytic study of the linear power spectrum expressed in units of $({\rm Mpc}/h)^3$ as a function of k in units of $h/{\rm Mpc}$ shows that at any given redshift, the large-scale branch ($k \ll k_{\rm eq}$) depends only on a factor $\left(g(\Omega_{\rm m}, z)/\Omega_{\rm m}\right)^2$ coming from the Poisson equation and from the behaviour of matter perturbations during Λ domination (see e.g. Ref. [66], equation (6.39)). The function $g(\Omega_{\rm m}, z) \leq 1$ is related to the decrease of matter perturbations during Λ domination. When increasing $\Omega_{\rm m}$, we decrease this factor $\left(g(\Omega_{\rm m}, z)/\Omega_{\rm m}\right)^2$ and we suppress the large-scale power spectrum, but not the small-scale one. Indeed, looking again at equation (6.39) in Ref. [66], the small-scale branch receives an extra factor $\tilde{k}_{\rm eq}^4$ (i.e. $k_{\rm eq}^4$ with $k_{\rm eq}$ in $h/{\rm Mpc}$). This new factor is actually proportional to $z_{\rm eq}^2 \Omega_{\rm m}^2$ (eq. (6.32) in the same reference), and the latter cancels the former $\Omega_{\rm m}^{-2}$ factor.

9 Since the lensing spectrum directly depends on metric fluctuations, it does not share with the matter power spectrum the factor $\Omega_{\rm m}^{-2}$ coming from the Poisson equation. Indeed, the factor $\Omega_{\rm m}^{-2}$ discussed in the previous footnote is exactly cancelled by a factor $\Omega_{\rm m}^2$ that appears in Eq. 284 when replacing the window functions with

the large-scale branch of the C_l^{ij}'s slightly increases when we decrease h and increase $\Omega_{\rm m}$. Instead, the small-scale branch remains nearly constant due to the antagonist effects of neutrino free-streaming and of the increase in $\Omega_{\rm m}$, but the neutrino effect wins on non-linear scales. As can be seen on the right panel of Figure 15, for the highest redshift bins, the lensing data is able to discriminate this effect and to lift the (M_ν, h) degeneracy, although with less significance than the galaxy clustering data.

These conclusions are confirmed by the (M_ν, h) joint probability contours presented in the upper right panel of Figure 18, for CORE, Euclid-lensing and Euclid-pk. Indeed, the slope of the (M_ν, h) degeneracy is different from the one observed in CMB data, and it is mainly driven by the CMB prior on θ_s.

3. *Degeneracy between M_ν and other parameters from Large Scale Structure data alone.*

Finally we increase M_ν, decrease h by a smaller amount than the one required for fixing θ_s, and, at the same time, we vary the primordial power spectrum parameters, the amplitude A_s and also the index n_s in the case of cosmic shear (blue dotted lines). It is clear from Figures 15 and 16 that this procedure can almost cancel the effect induced by a larger M_ν both in the shear C_ℓ^{ij} and in the galaxy $P(k)$, leading to a new degeneracy. We shall now explain the reasons for this degeneracy.

Considering that the primordial power spectrum of scalar perturbations is given by

$$\frac{k^3 \mathcal{P}_{\mathcal{R}}(k)}{2\pi^2} = A_s \left(\frac{k}{k_0}\right)^{n_s - 1}, \tag{286}$$

the matter power spectrum $P_{\rm m}$ can be written as

$$P_{\rm m}(k,z) \propto A_s \left(\frac{k}{k_0}\right)^{n_s} T(k,z)^2 , \tag{287}$$

where $T(k,z)$ is the time and scale dependent linear transfer function of matter density fluctuations (not separable in the case of massive neutrinos). As we have already explained, neutrinos induce a relative suppression of power on scales $k > k_{\rm fs}$; this suppression is encoded in the transfer function $T(k,z)$ of Eq. 287. In Figure 17 we show how $T(k,z)$ is suppressed by a larger neutrino mass sum on $k > k_{\rm fs}$ at redshift $z = 0$ and $z = 2$. Changing (n_s, A_s) affects only the primordial power spectrum, while leaving $T(k,z)$ unchanged, therefore, since we keep ω_b and $\omega_{\rm cdm}$ fixed, any deviation from the green solid line is due only to the variation of h and Ω_m. If, besides increasing M_ν, we decrease h to keep θ_s fixed (red dashed line), then the suppression of $T(k,z)$ extends to $k < k_{\rm fs}$ (because of the $(g(\Omega_{\rm m},z)/\Omega_{\rm m})^2$ factor) and the wiggles, due to the shift of the BAO scale, appear at smaller scales. This graphically explains what we have already discussed in point 2. Reducing the tweaking on h (blue dotted line) implies less reduction of power on the large scale branch and a smoothing of the wiggles; anyhow, the massive neutrino suppression of the transfer function is not fully compensated. However, if we look at Eq. 287 it is clear that a red tilt of the primordial power spectrum, combined with a smaller normalization, can mimic the same effect of a larger neutrino mass, reducing power on small scales respect to large scales.

Eq. 285. As a result, the large-scale branch of the C_l^{ij}'s depend on $g(\Omega_{\rm m},z)^2$ only, while the small-scale branch is proportional to $\Omega_{\rm m}^2$.

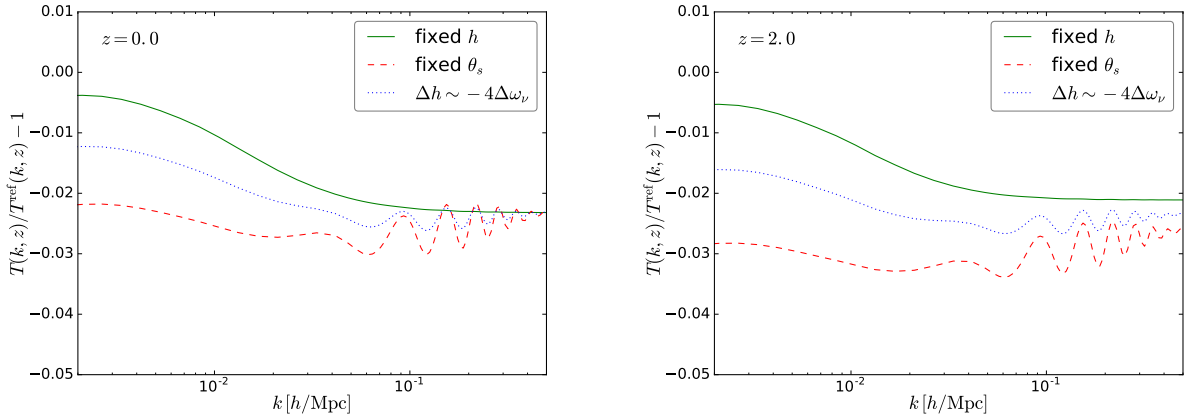

Figure 17: Relative error on the linear transfer function $T(k)$ at redshift $z = 0$ (left panel) and $z = 2$ (right panel). The line color/style - model correspondence is the same as in Figure 16.

The left and right bottom panels of Figure 18 show the degeneracies between M_ν and (n_s, A_s). We can see that the degeneracy between M_ν and n_s is mildly positive in galaxy lensing, as expected from the discussion above, while it is negative in CMB, as explained at the end of section 5.1.3, and mildly negative in galaxy clustering. The reason why this positive (M_ν, n_s) correlation emerges with cosmic shear, but not with galaxy correlation data, is related to the window function. Indeed, since the window function (Eq. 285) for each redshift bin is given by the integral over the line of sight, the C_ℓ^{ij}'s of Eq. 284 receive contributions from a larger range of scales. Therefore, being sensitive to a wider lever arm in k space, cosmic shear will be particularly sensitive to scale dependent variations of the power spectrum.

Notice that here the tweaking of A_s is larger than the one we performed at point 3 of section 5.1.3. Thus, the corresponding $\Delta\tau_{\rm reio} \sim 0.5\ln(1.05) \sim 0.027$ would lead to an enhancement of the reionization bump even bigger than the one we observed in the blue dotted line of the C_ℓ^{EE} plot (Figure 10, second row, right panel). This already shows that the degeneracy discussed here can be lifted by combining LSS data with CMB data. Nevertheless this discussion was important to understand the pulls in parameter space appearing when all data sets are combined with each other.

Figure 18 confirms the points discussed previously, and provides a comprehensive graphical summary of the complementarity between future CMB and LSS data in the context of neutrino mass measurement.

First, we see that even when adopting CMB-derived priors on ω_b and θ_s, the LSS data cannot efficiently constrain the neutrino mass, due to the degeneracy discussed in the previous paragraphs (point 3), involving mainly (M_ν, A_s, H_0), and to a lesser extent, n_s. We have seen that this degeneracy requires a milder correlation between M_ν and H_0 than the CMB data: $\Delta h \sim -3\Delta\omega_\nu$ for LSS alone, instead of $\Delta h \sim -12\Delta\omega_\nu$ for CMB alone. Since in Figure 18 the Euclid mock data was fitted together with a prior on θ_s, the final correlation angles represent compromises between these values. The lensing data also exhibits a negative correlation between M_ν and $\omega_{\rm cdm}$.

The CMB and LSS contours of Figure 18 clearly intersect each other for several pairs of parameters:

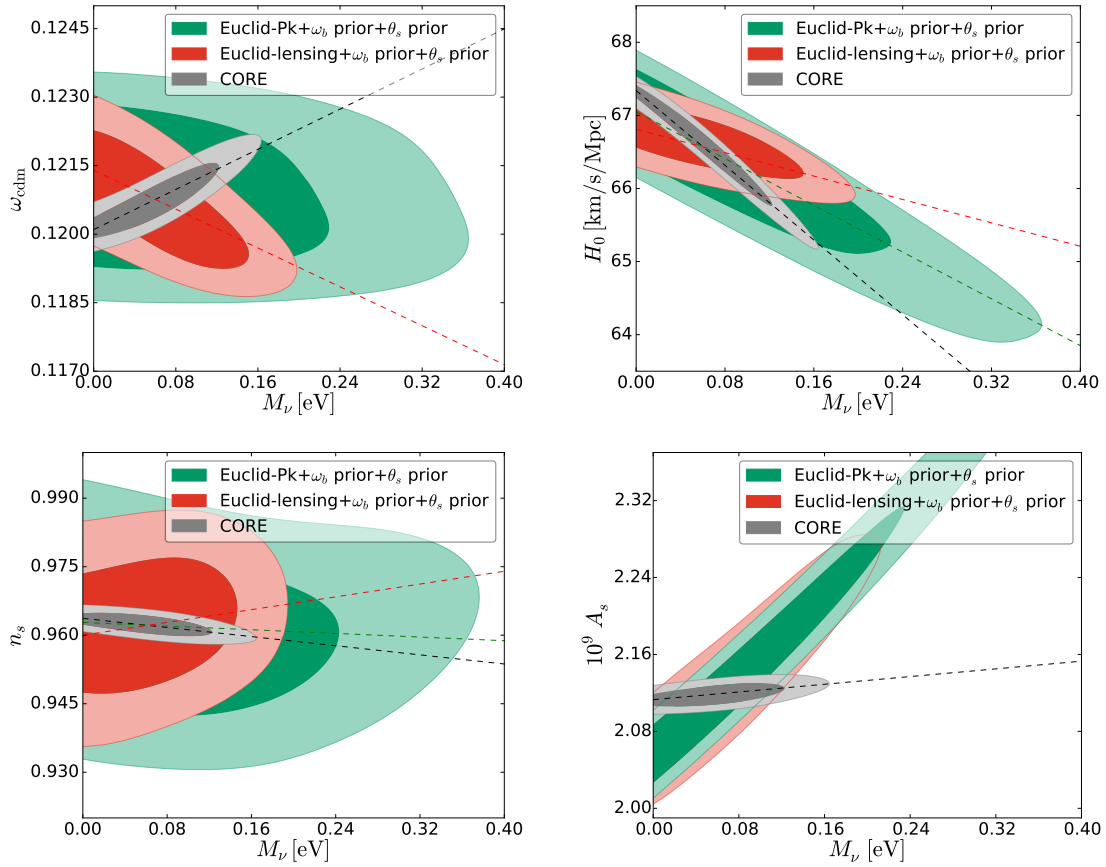

Figure 18: Marginalized one- and two- σ contours in the plane $(\omega_{\rm cdm}, M_\nu)$ (upper left panel), (H_0, M_ν) (upper right panel), (n_s, M_ν) (bottom left panel), (A_s, M_ν) (bottom right panel). The black dashed lines show the degeneracies encoded in CMB data, the red and green dashed lines account for some of the most prominent correlations arising from cosmic shear and galaxy clustering, respectively.

- The CMB and LSS data prefer different directions of degeneracy in (M_ν, H_0) space, hence the combination between them can strongly reduce the uncertainty on both M_ν and H_0.
- The CMB data lifts the (M_ν, A_s) degeneracy present in the LSS data, for the reason mentioned above: the shift in A_s would need to be compensated by a shift in $\tau_{\rm reio}$ producing a reionisation bump incompatible with the data. However, in the combined data set, the LSS data would keep pulling towards more positive correlation between M_ν and A_s.
- the very different correlations in $(M_\nu, \omega_{\rm cdm})$ space reduces uncertainties on $\omega_{\rm cdm}$, with a side effect on the CMB side. We have seen that the effect of neutrino masses on the CMB lensing spectrum can be compensated either by playing with $\omega_{\rm cdm}$, or with $(A_s, \tau_{\rm reio})$. The CMB alone would favour the first option. Like BAO data, weak lensing data breaks the $(M_\nu, \omega_{\rm cdm})$ degeneracy and leaves only the second option. This goes in the same direction as the previous point: pulling towards more positive correlation between M_ν and A_s.

Hence, we can already anticipate that the combination of CMB plus LSS data leads to an enhanced degeneracy between (M_ν, A_s) compared to CMB data alone. As a consequence, in

order to maintain a fixed combination $A_s e^{-2\tau_{\rm reio}}$, the combined data may generate a significant correlation in $(M_\nu, \tau_{\rm reio})$ space.

6

METHODOLOGY AND TOOLS

Chapter 6 is based on Brinckmann & Lesgourgues [175]

In this chapter, we review relevant concepts on the topic of statistical inference and outline recent developments of the cosmological sampling package `MontePython`, which is the framework within which the developments of chapters 2 and 3 were implemented. The results of chapters 7 to 9 were obtained with `MontePython` and would not have been feasible without the improvements in computational efficiency outlined in this chapter. Finally, in section 6.5 we provide a comprehensive overview of the current features and options of `MontePython`.

6.1 MONTEPYTHON INTRODUCTION

`MontePython` [174] is an MCMC sampling package in `Python` used for parameter inference in cosmology, similar to `CosmoMC` [221, 222] and `CosmoSIS` [223]. The modular nature of `MontePython` means modification of the code is particularly easy, and encourages implementation of specific modules to other Python sampling packages, e.g. the extensive library of cosmological likelihoods[1]. `MontePython` has two different modes: when running with

```
> python montepython/MontePython run <options>
```

it is a sampler (similar to `CosmoMC`), and when running with

```
> python montepython/MontePython info <options>
```

it is a tool for analyzing MCMC chains and plotting results (similar to `GetDist`[2]).

The code is currently interfaced with the Boltzmann code `CLASS` [171, 205, 206, 224] and extensions thereof, e.g. `HiCLASS` [225] and `SONG` [226]. There also exist some publicly available branches of CLASS achieving different purposes, e.g. `ExoCLASS` for advanced energy injection, recombination and reionisation features [227], `CLASS_SZ` for Sunyaev-Zel'dovich observables [228], `CLASSgal` for computing the number count C_ℓ's [214][3], or a branch incorporating non-local contributions to General Relativity [229][4].

In principle, it could easily be extended for use with e.g. `CAMB` [201, 230], via the new Python wrapper[5], or PyCosmo [231], a Boltzmann code in Python.

In this chapter, we present the latest development of `MontePython` over the past couple of years. In particular, we introduce two new ingredients that both contribute towards improving the performance of Metropolis-Hastings sampling. In section 6.2, after recalling the way in which the Metropolis-Hastings algorithm is implemented in `MontePython`, we present a

1 A practice the authors fully support and encourage, with proper citations and credits.
2 http://getdist.readthedocs.io/en/latest/
3 this feature has also been implemented in the main `CLASS`, but with small differences in the two implementations. The original `CLASSgal` code is still available at https://cosmology.unige.ch/content/classgal
4 see the pull request #86 in https://github.com/lesgourg/class_public
5 http://camb.readthedocs.io/en/latest/

new adaptation algorithm for the jumping factor. In section 6.3, we detail our strategy for calculating the Fisher matrix and its inverse, which can be used as a proposal density for a Metropolis-Hastings run. In section 6.4, we provide several examples of runs showing that these features speed up convergence and can save many hundreds of CPU-hours in the case of difficult runs, with a poor prior knowledge of the covariance matrix. IFinally, we summarise all the functionalities of `MontePython` in the current release, including extended cosmological parameter definitions with respect to `CLASS` in section 6.5.1, sampling options in section 6.5.2, analysis and plotting options in section 6.5.3 and likelihoods in section 6.5.4. Indeed, the new release of the code incorporates several new plotting options and even more new likelihoods based either on current or mock data.

The version of the code described in this paper has version number 3.0 and is available at `https://github.com/brinckmann/montepython_public`

6.2 METROPOLIS-HASTINGS SAMPLING STRATEGY

`MontePython` can switch between different ways to explore parameter space, which include Metropolis-Hastings, Nested Sampling, Cosmo Hammer, and a new Fisher sampling method described in section 6.3. These different algorithms are called `methods` in the code, and the same list of methods also includes post-processing algorithms like Importance Sampling or Adding Derived Parameter(s)[6].

The default method is the Metropolis-Hastings algorithm, working since `v2.0.0` (2013) with a *fast sampling* method quickly summarised in section 6.2.1, and since `v2.2.0` (October 2015) with a covariance matrix update method summarised in section 6.2.4. In this release `v3.0.0` we extend the latter to update also the jumping factor, as described in section 6.2.4, and we call the new approach `superupdate`.

6.2.1 *Fast sampling*

In `MontePython` the Metropolis-Hastings draws random jumps in parameter space from a Gaussian proposal density. The latter is encoded in a matrix $\mathbf{C}$, describing the parameter correlations and the standard deviations relative to each other, and an overall jumping parameter c, such that the parameter jumps $\Delta\mathbf{p}$ are generated randomly from the probability distribution $\mathcal{P} = \mathcal{N}\exp(-\frac{1}{2c}\Delta\mathbf{p}^T\mathbf{C}^{-1}\Delta\mathbf{p})$. Thus the actual covariance matrix of the proposal density is $c\mathbf{C}$. The standard way to generate random vectors from a multivariate Gaussian probability is to go to a basis of independent parameters, like the basis of the eigenvectors of $\mathbf{C}$; to generate independent random displacements along the eigenvectors; and to project back to the original space.

It is well-known that optimal proposal densities generate an acceptance rate of the order of 0.25, and that for Gaussian posterior distributions, this can be achieved when $\mathbf{C}$ is a good approximation to the covariance matrix of the *posterior distribution*, while the jumping parameter is fixed to $(2.4)^2$ [232]. Note that this jumping parameter applies when generating one single random number and moving in one single direction. Alternatively, for each jump, one can generate N random numbers and move in N directions simultaneously, but then each of these N random numbers should be drawn from a Gaussian distribution with variance $c = j^2/N$, where $j = 2.4$ is called the *jumping factor*.

6 The full list of methods can be viewed with `MontePython.py run --help`, and is of the form `-m {MH, NS, CH, IS, Der, Fisher}`.

Thus, in absence of fast sampling, the jumps could just be generated each time independently and randomly from $\mathcal{P}$ (global method), or in cycles of N draws along each of the N eigenvectors of $\mathbf{C}$ (sequential method). These methods can still be activated[7] in `MontePython`, but they are sub-optimal in presence of likelihoods with nuisance parameters.

Fast sampling was proposed by [222] for MCMC parameter estimation. For this sampling method, we separate the sampling of fast nuisance and slow cosmological parameters to optimize performance when dealing with a large number $N_{\rm fast}$ of nuisance parameters.

Generating displacements along the eigenvectors mixes slow and fast parameters and does not allow for high-speed explorations of the fast parameter space only. But introducing eigenvectors is not the only way to go to a parameter basis in which the proposal density is orthogonal. In particular, one can perform a Cholesky decomposition of the covariance matrix into $\mathbf{C} = \mathbf{L}\mathbf{L}^T$ where $\mathbf{L}$ is a lower triangular matrix. In the space of the vectors $\Delta\mathbf{p}'$ related to the physical parameters through $\Delta\mathbf{p}' = \mathbf{L}^{-1}\Delta\mathbf{p}$, the proposal density is orthogonal, so the jumps can easily be generated by drawing random numbers for each component of $\Delta\mathbf{p}'$ with a single one-dimensional Gaussian probability distribution of variance c, and projecting back to $\Delta\mathbf{p} = \mathbf{L}\,\Delta\mathbf{p}'$. The great advantage over the previous eigenvector-based scheme is that when $\Delta\mathbf{p}'$ only has non-zero components above a given index, this is true also for $\Delta\mathbf{p}$. Thus one can generate some jumps that will leave the slow parameters unchanged.

We begin by ordering our input parameters in blocks according to computational time. In practise, this is simply achieved by writing them in the right order in the input parameter file. The first block is that of cosmological parameters requiring new calls to the Boltzmann code. The next blocks are nuisance parameters for a given likelihood, which can be changed without requiring a Boltzmann code evaluation if the cosmological parameters are held fixed. The nuisance parameter blocks should be ordered from the slowest to fastest likelihood. When a nuisance parameters is common to several likelihoods, it should just be declared within the slowest block. We call M the number of blocks and d_j the number of parameters in the j-th block, with $d_1 = N_{\rm slow}$ being the number of cosmological parameters. `MontePython` will automatically detect the number M of blocks and will expect the user to pass an M-dimensional over-sampling vector $\mathbf{F}$. F_1 is the over-sampling factor of the cosmological parameter and is normally fixed to one. The other entries are the required number of redundent sampling for each of the other blocks.

When running chains, for $j = 1, ..., M$, we generate sequences of $F_j d_j$ random jumps in the d_j components of $\Delta\mathbf{p}'$ corresponding to the j-th block. In other words, during $F_j d_j$ steps, we generate d_j random numbers for each of the relevant components of $\Delta\mathbf{p}'$, drawn from a Gaussian distribution with standard variance $c = j^2/d_j$. Thus each full cycle consists of $\sum_j F_j d_j$ random jumps, with only $F_1 d_1 = N_s$ of them requiring a call to the Boltzmann code. Later we will call this number the Fast Parameter Multiplier (FPM):

$$\mathrm{FPM} = \sum_j F_j d_j. \tag{288}$$

There is no precise rule to fix the over-sampling factors $F_{2,...,N}$. These factors should be increased for faster likelihoods and/or larger numbers of nuisance parameters in the block. With too low numbers, one would not enjoy the advantages of the slow-fast parameter decomposition. With too high numbers, the time spent in the $(\sum_j F_j d_j - N_s)$ iterations over fast parameters could be significant compared to the time spent in the N_s iterations over slow parameters, and the convergence of the results for the cosmological parameter would be de-

7 with the *jumping* flag `-j {global, sequential}` instead of the default `-j fast`.

layed. For instance, for the nuisance parameters of the Planck likelihood, we usually apply an oversampling factor of 4.

6.2.2 *Update and Superupdate*

While the Metropolis-Hastings algorithm would in principle require Markov Chains, i.e. chains with a proposal density that is constant in time, it is highly desirable to implement some automatic update algorithms in order to get converged results even when starting from bad guesses for the Gaussian proposal density. We recall that the propoosal density is parametrised as $\mathcal{P} = \mathcal{N} \exp(-\frac{1}{2c}\Delta\mathbf{p}^T\mathbf{C}^{-1}\Delta\mathbf{p})$ and thus depends on two quantitites, the covariance matrix $\mathbf{C}$ and the jumping parameter c related to the jumping factor j. `MontePython` has two complementary options for speeding up the convergence of Metropolis-Hastings runs:

- `--update` U: update of covariance matrix, $\mathbf{C}$, every U cycles [default: $U = 50$]
- `--superupdate` SU: additionally, update of jumping factor, j, starting SU cycles after each covariance matrix update [default: $SU = 0$, meaning "deactivated"; recommended: 20]

Once certain criteria are met, the covariance matrix will be updated periodically and the jumping factor will be adapted every step. This leads to dramatic improvements in runtime, especially for runs with little prior knowledge in the form of an appropriate starting stepsize or a good initial proposal distribution. In the next two sections we describe the strategy chosen for these two schemes.

6.2.3 *Update strategy*

The covariance matrix update mechanism was implemented early on in `CosmoMC` and is part of `MontePython` since `v2.2.0` (October 2015), through the flag `--update`. The `MontePython` and `CosmoMC` implementations of this feature are very similar. For instance, in both codes, the decision to start or stop the update mechanism depends on the value of the Gelman-Rubin statistic, R [181], for the most poorly converged parameter. The update starts when the number $\max(R-1)$ computed from the second half of each chain goes below 3, and stops when it goes below 0.4. The difference between the two implementations only resides in two aspects:

- *Non-MPI-user friendliness.* The most straightforward way to launch multiple chains is to run `MontePython` with `MPI`, e.g. for 8 chains:
  `mpirun -np 8 python montepython/MontePython.py run ...`
  The alternative would be to launch 8 chains manually, or within a small shell script with a `for` loop. With `CosmoMC`, this second option would be incompatible with the covariance matrix updating. In `MontePython`, because installing `MPI` can sometimes be cumbersome, we chose to code the `--update` mechanism in such a way that it will work equally well with or without `MPI`. In the latter case, the user should just run several times `python montepython/MontePython.py run ...`, and the update mechanism will still start and use the information from *all* the chains running in the same directory[8].

8 There is no need to know how this is implemented in the code, because it is fully transparent to users. In brief, the key point is that `MontePython` reads and writes the covariance matrix in a file, rather than using the Message-Passing-Interface. The only difference between MPI and non-MPI runs is that in the former case, the code can

- *Keeping only Markovian steps in the final results.* To be rigorous, the user would like to base his/her final results and plots on true Markovian chains. This is what happens by default with `MontePython`. Indeed, every time that the covariance matrix is updated, the code writes in all chains files a comment line starting with
  `# After <k> accepted steps: update proposal...`
  which also contains information on the current convergence (thus these comments can also be used to scrutinise what is happening with the run). When analysing the chains with the `info` mode, by default, the code will only consider the part of the chains after the last update, i.e. the Markovian part. If the user wants to de-activate this behaviour in order to get more points in the chains, she/he can use the flag `--keep-non-markovian`.

Finally, the update periodicity is controlled by the U input parameter (default: $U = 50$), which is in units of cycles. Given that a cycle consists in FPM steps (see eq. (288)), the update takes place every

$$N_{\text{update}} = \text{U} \times \text{FPM} \tag{289}$$

steps (here we are referring to proposed steps, not accepted steps).

6.2.4 *Superupdate strategy*

The covariance matrix updating does not fully achieve the task of reaching optimal convergence conditions automatically. The other part of the proposal density is the jumping parameter c, related to the jumping factor j. If j is too large, the acceptance rate (*a.r.*) is too small and the number of accepted models remains insufficient to extract statistical information. If j is too small, the *a.r.* may get close to one. In that situation the chains would grow rapidly, but adjacent points would be very correlated, and the chains would not necessarily sample the full posterior distribution. Thus one should target a compromise value of the acceptance rate. Since the work of [232], cosmologists usually aim at $a.r. \simeq 0.25$ (although larger values in the range 0.3-0.5 would in principle still be acceptable).

While $f = 2.4$ matches this goal for a multivariate Gaussian posterior, many runs target non-Gaussian posteriors, e.g. due to non-trivial priors on cosmological parameters (like the requirement of a positive neutrino mass) or to strongly non-Gaussian posteriors for the nuisance parameters (like those of the Planck high-ℓ likelihoods). The current practise consists in training a bit and trying manually different values of f until the acceptance rate is correct. For instance, one quickly comes to know that, e.g., a given version of the Planck likelihood usually needs a given value of f to achieve $a.r. \simeq 0.25$. Of course, it would be better to let the code find this value automatically for each combination of a model and a dataset. An automatic jumping factor adaptation would also make the code more powerful when starting from a very bad proposal density (e.g., when adding many new free parameters to a previous run, or when investigating a very constraining set of likelihoods when only the covariance matrix of a much less constraining set is available).

---

define a "master chain" and some "slave chains", and only the master chain occasionally pauses in order to update the covariance matrix; in the latter case, all chains occasionally pause for the same purpose, but this does not affect convergence, and it only increases the total running time by a very small amount. Note that exchanging information on the covariance matrix through a file could have a potential inconvenience: When the user analyses an on-going run with the `info` mode, she/he could generate a new covariance matrix that may interfere with the automatic updating mechanism. This is not the case because when the user runs in `info` mode, the covariance matrix calculation is de-activated by default; it is only activated with the `--want-covmat` flag.

For that purpose, we added to version `v3.0.0` the new option `--superupdate`, which is complementary to `--update`: they should normally be used in combination, but then it is only necessary to pass the `--superupdate` flag since `--update` is activated by default.

Note that other schemes to update the full proposal density (rather than just the covariance matrix) have been investigated in the past. For instance, an Adaptative Metropolis algorithm for single-chain runs was proposed in [233], and a version of this algorithm was implemented in a cosmology MCMC code by the CAMEL collaboration [234]. After trying this method for single chains, we adapted it freely to the multi-chain case, in a way which remains compatible with the traditional covariance matrix update scheme (`--update`).

Overall strategy. When running with `--superupdate` SU, the code starts from a jumping factor that can be set manually with a command flag (e.g. `-f 2.2`), but when nothing is passed, 2.4 is used by default. The code keeps a record of the acceptance rate, $a.r.$, and of the jumping parameter, c, of the last SU cycles, i.e. of the last $SU \times$ FPM steps. It also keeps track of the average $\overline{a.r.}$ and $\bar{c}$ over these last SU cycles. This information is used to compute the start- and stop-criteria. Since `--superupdate` requires `--update` to be active, the run can be divided in several "update sequences", which are the ensemble of steps between two consecutive updates. The first "update sequence" is just the time until the first update. The basic principle of superupdate is to adapt the jumping factor at each step according to the recurrence relation

$$c_k = c_{k-1} + \frac{1}{(k - k_{\text{update}})}(\overline{a.r.} - 0.26) \,, \tag{290}$$

where k is the current step number, while k_{update} is the first step number of each new "update sequence". For the first sequence, $k_{\text{update}} = 0$. This recurrence relation leads to faster updating at the beginning of each new sequence, and to slower updating and safe convergence properties after some time.

Starting the jumping factor update. The code starts applying the recursion relation (290) when two conditions are met:

- We do not want to update the proposal distribution too early, as it could be based on chains still in the burn-in phase. For this reason, we wait until the chains have reached a certain level of convergence: the numbers $(R-1)$ computed from the second half of each chain should be below 10 for all parameters.

- We wait until we have done SU cycles since the beginning of the new "update sequence", or since the very beginning if we are still in the first sequence: $(k - k_{\text{update}}) \geq SU \times FPM$. Since the mean acceptance rate $\overline{a.r.}$ is computed over the last SU cycles, the recurrence will only take into account some steps from the same "update sequence". Choosing $SU \gtrsim 20$ ensures that the mean acceptance rate is not computed over too small a sample, where shot noise may lead to an acceptance rate significantly different from the target one (i.e. due to random fluctuations in the acceptance rate leading to prematurely stopping adaptation of the jumping factor). We recommend using $SU = 20$, as we found this to be a good compromise between efficiency and precision, but higher values can be considered in order to decrease the impact of superupdate (the jumping factor would start evolving later and would perform smaller excursions) or to further decrease the impact of shot noise on the determination of the jumping factor. Note that with $SU > U$, superupdate would sometimes only be active after the very last update of the covariance

matrix, in the final stage of the run, when the convergence is already good (or possibly before the chains are well enough converged for updating the covariance matrix to begin). Thus one should normally consider the range $20 \leq SU < U$ only.

Rescaling when the covariance matrix gets updated. Since the true covariance matrix of the Gaussian proposal density is in fact given by the product $c\mathbf{C}$, it would be sub-optimal to leave c unchanged when the matrix $\mathbf{C}$ is updated at the beginning of each new "update cycle". Suppose for instance that in the n-th "update cycle", a good jumping factor c_n has been found in combination with a covariance matrix $\mathbf{C}_n$ (by this we mean that the acceptance rate has the correct order of magnitude). If at the beginning of the next cycle the matrix is adapted to a $\mathbf{C}_{n+1}$ which is much smaller, while c_{n+1} restarts from the same value c_n, then obviously the whole proposal density will shrink and the acceptance rate will increase too much. We can limit this effect by requiring analytically that at each covariance matrix update, the volume probed by the full proposal density remains constant, which is achieved simply by imposing:

$$c_{\text{after}}^N \det(C_{\text{after}}) = c_{\text{before}}^N \det(C_{\text{before}}) \,, \tag{291}$$

where N is the number of free (slow+fast) parameters. This rescaling might not be very efficient when the evolution of the covariance matrix comes from only one or few parameters, but in general, it is the best simple guess that one can do. In terms of the jumping factor, this gives:

$$j_{\text{after}} = j_{\text{before}} \left[\frac{\det(C_{\text{before}})}{\det(C_{\text{after}})} \right]^{\frac{1}{2N}} \,. \tag{292}$$

Note that, for the first update of the covariance matrix, the logic behind this re-scaling does not hold. Indeed, if we started from a poor input covariance matrix, the first re-scaling of the jumping factor may be completely unrealistic. For safety, at the first update time, we reset the jumping factor to the input value (provided via `-f` [default: 2.4]).

Stopping the jumping factor update. We adapt the jumping parameter until three conditions are met:

- The *a.r.* should converge to 26% with a tolerance of 1 percent point[9] (in many cases, the *a.r.* starts low and increases to the optimal value, then the adaptation will stop when the code reaches 25%):

$$|\overline{a.r.} - 0.26| < 0.01 \,. \tag{293}$$

- In addition to the *a.r.* criterium, in order to stop adaptation of the jumping parameter we also require that it is stable,

$$\left| \frac{\bar{c}}{c_{k-1}} - 1 \right| < 0.01 \,, \tag{294}$$

  where $\bar{c}$ is the mean of the jumping parameter over the last $SU \times$ FPM steps. Also, we do not wish to allow the jumping parameter to converge to arbitrarily low values, as the risk of chains getting stuck in local minima would increase. Therefore, we introduce a minimum for the jumping factor corresponding to 10% of the initial jumping factor. If a small jumping parameter is desired, it is instead recommended to input a low value with `-f`.

9 This behaviour is controllable by the options `--superupdate-ar` and `--superupdate-ar-tol`.

- Finally, we require that the number $\max(R-1)$ computed from the second half of the chains is below 0.4 for all parameters: this is the same condition as for stopping the covariance matrix update. Thus the superupdate mechanism will only stop its activity in the final "update sequence", during which a large number of truly Markovian steps (generated with a constant proposal density) can be accumulated.

Non-MPI user friendliness. We implemented the superupdate mechanism in `MontePython` with the same coding principles as for the update mechanism. Thus it can also be used with or without MPI, thanks to the fact that the communication between chains works through files rather than MPI commands. In the running directory, a file `jumping_factors.txt` stores the sequence of all jumping factors that have been used, while the file `jumping_factor.txt` only contains the final one, that can be used as an input value in the next run. When chains are restarted in the same directory using the `-r` command, this will be done automatically.

Keeping only Markovian steps in the final results. When `--superupdate` is activated, the code still writes some comment lines in the chains at the beginning of each new "update cycle", with information on the current value of $\max(R-1)$, j and *a.r.*. Additionally, it writes an extra line of comments when the jumping factor updating stops. When analyzing the chains in `info` mode, and unless the user passes the option `--keep-non-markovian`, all the lines before will be discarded and the final numbers and plots will be based on purely Markovian chains.

Alternative implementation for single chain runs. The superupdate mechanism in principle requires multiple chains, since it uses convergence tests based on the Gelman-Rubin statistic. For single chain runs the code will split the chain into three separate chains in order to compute the Gelman-Rubin statistic, a practice that may be less reliable than running multiple chains. However, `MontePython` also has an alternative to superupdate that was previously implemented for single chain runs by [235]. This other mechanism is activated by the flag `--adaptive` instead of `--superupdate`. It does not use the Gelman-Rubin statistic, and it is slightly closer to the original Adaptive Metropolis algorithm of [233], with an update of the covariance matrix at each single step.

6.3 FISHER MATRIX

The well-known Fisher matrix is built from the second derivatives of the effective χ^2 with respect to the model parameters computed at a minimum of the χ^2, i.e. at the maximum of the likelihood (see e.g. [236]):

$$\mathbf{F}_{ij} = \frac{1}{2}\frac{\partial^2\chi^2}{\partial p_i \partial p_j} = -\frac{\partial^2 \ln \mathcal{L}}{\partial p_i \partial p_j} \,. \tag{295}$$

By definition of the maximum likelihood point, the Fisher matrix must be positive definite and invertible. Its inverse is the covariance matrix of the Gaussian approximation to the likelihood near the best-fit point. If the matrix of second derivatives is not computed at that point, it may not be invertible.

6.3.1 *Motivations for Fisher matrix computation*

We implemented in `MontePython v3.0.0` a calculation of the Fisher matrix directly from the likelihood and from eq. (295), using a finite difference method that we will detail in the next

section. The motivation behind this calculation is twofold:

Boosting MCMC runs. The inverse Fisher matrix can be used as the input covariance matrix of an MCMC run (e.g. Metropolis-Hastings). In that case we don't need a high-accuracy calculation of this matrix, because any approximate result will likely be a good enough guess, that the Metropolis-Hastings "update" mechanism will quickly improve anyway. This method leads to a very significant speed up in most cases, since one rarely starts an MCMC run with already at disposal a very good covariance matrix including all pairs of parameters. Still, the method can work only if the code finds an invertible Fisher matrix in the first place, and this is only guaranteed at the exact maximum likelihood point. This condition can be easy or difficult to achieve depending on the type of run:

- For parameter forecasts with mock data, we usually use the fiducial spectra in the role of the observed spectra, without generating a random realisation (see e.g. [96] for comments on this methodology). Thus the maximum likelihood exactly coincides with the fiducial model, known in advance by the user. Then the new Fisher method works particularly well.

- For parameter extraction from real data, we know at most an approximation to the best fit point. Then, one may hope that if the distance between the true and approximate best fit points is small compared to the steps used in the finite difference method, the approximate Fisher matrix computed in the latter point will still be positive definite and invertible. We found however that this does not happen very often, so the possibility to use this new method for real data remains somewhat random. To increase chances, we incorporated in `MontePython` a few minimum-finding algorithms taken from the `optimize` python library[10], that may at least help to get closer to the true best-fit point before trying the Fisher matrix computation. However, our tests show that, for the moment, the implemented minimization algorithms are not very robust, especially in the presence of many nuisance parameters.

In summary, the new Fisher calculation definitely improves all MCMC forecasts (as shown in section 6.4), and it may also improve MCMC runs with real data (see a few examples in section 6.4.2) unless one ends up with likelihood shapes that happen to be too complicated for finding the minimum and/or running the Fisher algorithm.

Replacing MCMC runs. When one knows that the posterior of a given run should be nearly Gaussian, or when one is not interested in the details of the posterior (e.g. non-trivial parameter correlations with some skewness, kurtosis, banana-shape, etc.), it is tempting to replace whole MCMC parameter extraction runs by simple Fisher matrix computations. Then the inverse Fisher matrix will give some approximate one-dimensional confidence regions and two-dimensional elliptic contours. This is particularly straighforward for sensitivity forecasts

10 If the user runs `MontePython` with the command flag `--minimize`, before using any engine (Metropolis-Hastings, Fisher, etc.), the code will re-evaluate the central starting point using a χ^2 minimization algorithm. This call is done in the routine `get_minimum()` of the module `montepython/sampler.py`. After loading some approximation for the best-fit point and for the iteration step size, this routine calls the python function `numpy.optimize.minimize()`, which accepts several values of the input parameter `method`, corresponding to different minimization algorithms. By default we did set `method='SLSQP'`, which calls the Sequential Least SQuares Programming algorithm, but the user is free to edit the module and change one line to try different methods. The algorithm stops when the χ^2 seems to be converged up to the tolerance passed through the `MontePython` input flag `--minimize-tol` (default 10^{-5}), but there is no guarantee that the algorithm leads to the true minimum.

since, in that case, the maximum likelihood point is known in advance. It can also be envisaged for real data if the maximum likelihood point is known up to good approximation, for instance, after a run with the new `--minimize` option of `MontePython` (whose success is not guaranteed).

In the cosmology literature, a vast majority of parameter forecasts are based on Fisher matrix calculations. These are usually performed by specific codes, using the fact that after a few steps of analytic calculations, F_{ij} can be re-expressed as a function of the derivative of the observable quantities with respect to the parameters (e.g. $\partial C_\ell / \partial p_i$ or $\partial P(k) / \partial p_i$). Instead, the Fisher matrix computation performed by `MontePython` is a direct likelihood-based evaluation, since we compute $\partial \ln \mathcal{L} / \partial p_i$. The two approaches are mathematically equivalent, but the latter may offer some practical avantages. Indeed, the quantity which is primarily build to model a given experiment is the likelihood. Skipping the analytical steps leading to derivatives like $\partial C_\ell / \partial p_i$ or $\partial P(k) / \partial p_i$ sometimes avoids complicated expressions, the need to introduce approximations, and further risks to make an error.

In both approaches, one has to compute some numerical derivatives with a given step size. For a purely Gaussian likelihood, the step size should be irrelevant, provided that it is not so small that numerical errors (from the Boltzmann code or from the likelihood code) come to dominate. Very often, Boltzmann codes are optimized in order to give an accuracy on the $\chi^2_{\rm eff}$ of the most constraining experiments (typically, nowadays, the Planck experiment) of the order of $\delta\chi^2 \sim \mathcal{O}(10^{-1})$, simply because achieving better precision would not change the results on confidence intervals, and would thus be a waste of computing time. Therefore, it is dangerous to use steps Δp_i such that

$$\Delta\chi^2 \equiv \frac{1}{2}\left(\ln \mathcal{L}(p_i^{\rm bestfit} + \Delta p_i) - \mathcal{L}(p_i^{\rm bestfit})\right) \tag{296}$$

is significantly smaller than 0.1. This provides roughly a lower bound on the Δp_i's. The question of the upper bound is more delicate, and especially important when the likelihood is not Gaussian: different choices can then return significantly different Fisher matrices and confidence limits. The community is split between different approaches on this issue. One school suggests to take the smallest possible steps until numerical noise comes into play, in order to be as close as possible to the mathematical definition of the second derivatives. Another school prefers to choose steps such that $\Delta\chi^2 \simeq 1$ (resp. 4) if the final goal is to deliver predictions for 68% (resp. 95%) confidence limits on the parameters, since in that case the Fisher matrix gives a Gaussian approximation of the likelihood valid precisely in the region that is relevant for the parameter bounds (see e.g. [96]). Many Fisher codes do not even target any particular given order of magnitude for $\Delta\chi^2$, and choose the steps Δp_i arbitrarily.

In `MontePython v3.0.0`, we approach this problem by letting the user choose a target value for $\Delta\ln\mathcal{L} = 2\Delta\chi^2$. By default, the code will first try to get an invertible Fisher matrix with $\Delta\ln\mathcal{L} \sim 0.1$, and will iteratively increase this value in case the result is non-invertible, as explained in the next section. However, the user can choose the first value of $\Delta\ln\mathcal{L}$, and may for instance set it to 0.5 or 2 (with the flag `--fisher-delta`).

6.3.2 *Iterative strategy for the fisher matrix computation*

The Fisher matrix calculation is a "method", like e.g. Metropolis-Hastings, Nested sampling, Importance Sampling, minimization, etc. It is activated by launching `MontePython v3.0.0` in `run` mode with the command flag:

- `--method Fisher` : calculate Fisher matrix

The calculation takes place around parameter values specified by the first entries of each list in the input file:

```
data.parameters['P1'] = [p1, .., .., .., .., ..] ,
```

unless another best-fit model is passed with the command line `-b path/to/file.bestfit`. The user may control the step size for the finite difference derivatives with two parameters:

- `--fisher-delta` D: target $\Delta\ln\mathcal{L}$ value for finding the steps Δp_i [default: $D = 0.1$]
- `--fisher-tol` T: tolerance for $\Delta\ln\mathcal{L}$ (note: decreasing slows down computation) [default: $T = 0.05$]

Then the code finds the step size for each parameter matching the target $\Delta\chi^2 = D \pm T$ by bisection. Sometimes the bisection struggles to converge (e.g. for non-Gaussian likelihoods, or if the calculation is not centered on the maximum of the likelihood). In this case, after 10 attempts, it gradually increases the tolerance T at each step until convergence is obtained. However, in such an event, it may be preferable to adjust the input parameters instead (e.g. target D or best fit parameter values).

Once the step sizes have been obtained, the code computes all the elements of the Fisher matrix. If the result is a non-invertible matrix (due to the non-Gaussianity of the likelihood or to a bad guess for the maximum likelihood), the code enters into a stage of iterations over the target value of $\Delta\chi^2$, which is steadily increased until the matrix becomes invertible, following the sequence $D, 2D, 3D, ...ND$. The maximum number of iterations can be controlled with

- `--fisher-step-it` N: number of step iterations attempted [default: $N = 10$]

If the matrix inversion still fails after the maximum number of iterations, the code stops and returns an explicit error message.

Whenever the code finds an invertible Fisher matrix, it stores both the Fisher matrix and its inverse in distinct files with the extension `.mat`. The inverse Fisher matrix file matches the usual format of any covariance matrix that the Metropolis-Hastings algorithm would take in input for the density proposal. Thus it can immediately be used in an MCMC run with the input flag `-c path/to/file.covmat`.

6.3.3 *Dealing with prior boundaries*

In the last section we mentioned that the code finds the step sizes Δp_i used in numerical derivatives with a bisection algorithm. The bisection starts with a first tentative step size given by the input $\sigma_i^{\rm input}$ value for a each parameter, as given by the input file or by the input covariance matrix specified by the input flag `-c path/to/file.covmat` (the second always has priority). In cases where a $p_i^{\rm bestfit} \pm \sigma_i^{\rm input}$ value exceeds the prior boundary, we change the initial step $\sigma_i^{\rm initial}$ according to the following criteria:

Case 0: When there are no boundaries, or the difference between the boundary and the center is greater than $\sigma_i^{\rm input}$, the initial step is given by $\sigma_i^{\rm input}$:

$$\sigma_i^{\rm input} < B_{i-} \;\text{ and }\; \sigma_i^{\rm input} < B_{i+} \quad \Rightarrow \quad \sigma_i^{\rm initial} = \sigma_i^{\rm input} \,,$$

where the lower boundary distance is $B_{i-} = p_i^{\rm bestfit} - p_i^{\rm lower-boundary}$ and the upper one is $B_{i+} = p_i^{\rm upper-boundary} - p_i^{\rm bestfit}$.

Case 1: When one or both of the boundary distances is smaller than $\sigma_i^{\rm input}$, but both are still larger than a tenth of $\sigma_i^{\rm input}$, we set the initial step to the smaller of the two boundary distances:

$$0.1\sigma_i^{\rm input} < B_{i-/+} < \sigma_i^{\rm input} \quad \Rightarrow \quad \sigma_i^{\rm initial} = \min(B_{i-}, B_{i+}) \; .$$

Case 2: When one or both of the boundary distances is smaller than a tenth of σ_{input}, we instead assume the likelihood is symmetric around the best-fit point, and we only compute steps in one direction (the one in which the distance to the boundary is the greatest), while mirroring the likelihood values to the other direction:

$$B_{i-/+} \leq 0.1\sigma_i^{\rm input} \quad \Rightarrow \quad \sigma_i^{\rm initial} = \min(\max(B_{i-}, B_{i+}), \sigma_i^{\rm input}) \; .$$

Once the steps have been settled in that way, the diagonal elements of the Fisher matrix are given by the numerical derivatives

$$\mathbf{F}_{ii} = \frac{\partial^2 \ln \mathcal{L}}{\partial p_i^2} \approx \frac{\ln \mathcal{L}(p_i{+}\Delta p_i) - 2\ln \mathcal{L}(p_i) + \ln \mathcal{L}(p_i{-}\Delta p_i)}{\Delta p_i^2} \tag{297}$$

and the the off-diagonal ones by

$$\begin{aligned}\mathbf{F}_{ij} = \frac{\partial^2 \ln \mathcal{L}}{\partial p_i \partial p_j} \approx & \frac{\ln \mathcal{L}(p_i{+}\Delta p_i, p_j{+}\Delta p_j) - \ln \mathcal{L}(p_i{+}\Delta p_i, p_j{-}\Delta p_j)}{4\Delta p_i \Delta p_j} \\ & - \frac{\ln \mathcal{L}(p_i{-}\Delta p_i, p_j{+}\Delta p_j) - \ln \mathcal{L}(p_i{-}\Delta p_i, p_j{-}\Delta p_j)}{4\Delta p_i \Delta p_j} \; .\end{aligned} \tag{298}$$

Asymmetric steps. In cases 0 and 1, the code always uses symmetric steps, and in case 2 it postulates a symmetry of the likelihood. In some situations the user may find it beneficial to use instead some asymmetric steps to compute the Fisher matrix. This can be activated with the input flag:

- `--fisher-asymmetric` : allow for asymmetric steps (note: slows down computation) [default: False]

Then the "case 1" and "case 2" rules are replaced with some evaluations of the likelihood at $p_i + \Delta p_{+i}$ and $p_i - \Delta p_{-i}$, with $\Delta p_{+i} = \min(\sigma_i^{\rm input}, B_{i+})$ and $\Delta p_{-i} = \min(\sigma_i^{\rm input}, B_{i-})$. In that case the diagonal terms of the Fisher matrix are given by

$$\frac{\partial^2 \ln \mathcal{L}}{\partial p_i^2} \approx 2\frac{\left(\frac{\Delta p_{-i}}{\Delta p_{+i}}\right) \ln \mathcal{L}(p_i{+}\Delta p_{+i}) - \left(\frac{\Delta p_{-i}}{\Delta p_{+i}} + 1\right) \ln \mathcal{L}(p_i) + \ln \mathcal{L}(p_i{-}\Delta p_{-i})}{\Delta p_{-i}\Delta p_{+i} + \Delta p_{-i}^2} \; , \tag{299}$$

and the off-diagonal ones by

$$\begin{aligned}\frac{\partial^2 \ln \mathcal{L}}{\partial p_i \partial p_j} \approx & \left(\frac{\Delta p_{-j}^2}{\Delta p_{+j}} + \Delta p_{-j}\right)^{-1} \left(\frac{\Delta p_{-i}^2}{\Delta p_{+i}} + \Delta p_{-i}\right)^{-1} \\ & \times \left\{ \left(\frac{\Delta p_{-j}}{\Delta p_{+j}}\right)^2 \left[\left(\frac{\Delta p_{-i}}{\Delta p_{+i}}\right)^2 \ln \mathcal{L}(p_i+\Delta p_{+i}, p_j+\Delta p_{+j}) - \ln \mathcal{L}(p_i-\Delta p_{-i}, p_j+\Delta p_{+j}) \right] \right. \\ & - \left(\frac{\Delta p_{-i}}{\Delta p_{+i}}\right)^2 \ln \mathcal{L}(p_i+\Delta p_{+i}, p_j-\Delta p_{-j}) + \ln \mathcal{L}(p_i-\Delta p_{-i}, p_j-\Delta p_{-j}) \\ & + \left[\left(\frac{\Delta p_{-j}}{\Delta p_{+j}}\right)^2 - 1\right] \left(\ln \mathcal{L}(p_i-\Delta p_{-i}, p_j) - \left(\frac{\Delta p_{-i}}{\Delta p_{+i}}\right)^2 \ln \mathcal{L}(p_i+\Delta p_{+i}, p_j) \right) \\ & + \left[\left(\frac{\Delta p_{-j}}{\Delta p_{+j}}\right)^2 - 1\right] \left[\left(\frac{\Delta p_{-i}}{\Delta p_{+i}}\right)^2 - 1\right] \ln \mathcal{L}(p_i, p_j) \\ & \left. + \left[\left(\frac{\Delta p_{-i}}{\Delta p_{+i}}\right)^2 - 1\right] \left(\ln \mathcal{L}(p_i, p_j-\Delta p_{-j}) - \left(\frac{\Delta p_{-j}}{\Delta p_{+j}}\right)^2 \ln \mathcal{L}(p_i, p_j+\Delta p_{+j}) \right) \right\} . \end{aligned} \tag{300}$$

6.3.4 *Efficient treatment of nuisance parameters*

Not counting the few intermediate steps necessary for the automatic determination of step sizes (which is typically around $2-4$ evaluations per parameter), the calculation of one Fisher matrix requires a number of likelihood evaluations equal to

$$N_{\text{evaluations}} = 1 + 2N_{\text{params}} + 4 \sum_{n=1}^{N_{\text{params}}-1} n \, , \tag{301}$$

i..e. one in the best-fit point, two for each diagonal element and four for each off-diagonal element, where $N_{\text{params}} = N_{\text{cosmo}} + N_{\text{nuisance}}$ is the total number of parameters, N_{cosmo} is the number of cosmological parameters and N_{nuisance} is the number of nuisance parameters. For a typical Planck run we have 6 cosmological parameters and 26 nuisance parameters, resulting in 2049 likelihood evaluations when the target $\Delta\chi^2$ is not iterated on.

The fact that varying only nuisance parameters does not require a call to the Boltzmann solver allows us to considerably optimize the computation. In `MontePython`, the routine calling the likelihoods always keeps a memory of the previous step (model parameters and cosmological observables). Therefore, if the likelihood is evaluated at a new point such that only nuisance parameters have changed, `MontePython` knows that the Boltzmann code should not be called again. To optimize the Fisher matrix calculation, we just need to arrange the $N_{\text{evaluations}}$ calls of the likelihood in a particular order minimizing the number of calls to the Boltzmann solver.

We loop over the parameters starting from the cosmological ones and ending with the nuisance ones.

For each parameter p_i, we first perform all the calculations involving the value $(p_i-\Delta p_i$, i.e. $\mathcal{L}(p_i-\Delta p_i)$ for the diagonal element and $\mathcal{L}(p_i-\Delta p_i, p_j \pm \Delta p_j)$ (for each $j > i$) for the non-diagonal elements. Then we perform all the calculations involving the value $(p_i+\Delta p_i$, i.e. $\mathcal{L}(p_i+\Delta p_i)$ for the diagonal element and $\mathcal{L}(p_i+\Delta p_i, p_j \pm \Delta p_j)$ (for each $j > i$) for the non-

diagonal elements. In that way, the number of calls to the Boltzmann solver is drastically reduced to

$$N_{\rm calls} = 1 + 4N_{\rm cosmo} + 4\sum_{n=1}^{N_{\rm cosmo}-1} n\,, \tag{302}$$

corresponding to $N_{\rm calls} = 85$ for a typical Planck run (again without step iteration). This vastly reduces the computational time[11].

6.3.5 *Plotting likelihood contours from inverse Fisher matrix*

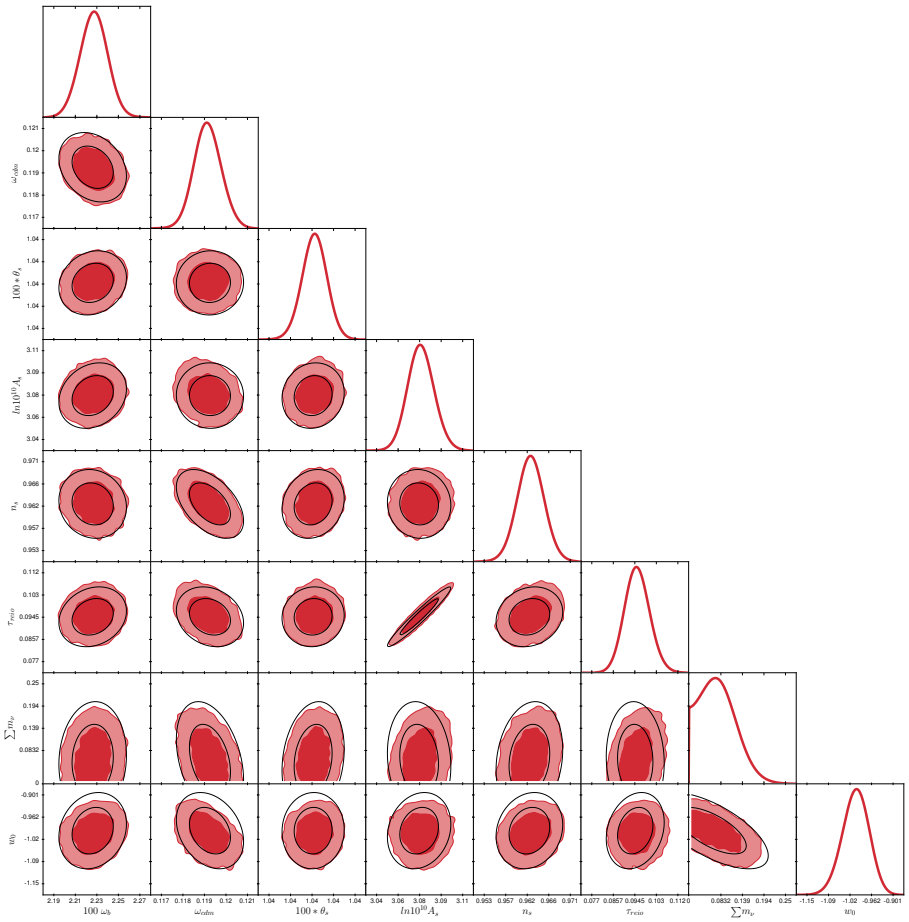

Figure 19: Confidence ellipses inferred from the inverse Fisher matrix (black lines) plotted on top of the 2d marginalized posterior distribution of a Metropolis-Hastings forecast. We fitted mock BAO data from the DESI survey combined with mock Planck data, assuming an 8 parameter cosmological model (νwCDM).

11 Note that if the number of operation was not reordered in such a special way, we would still get some gain, but the number of calls would still be as large as

$$N_{\rm calls} = 1 + 2N_{\rm cosmo} + 4\left(\sum_{n=1}^{N_{\rm params}-1} n - \sum_{m=1}^{N_{\rm nuisance}-1} m\right)\,.$$

This corresponds to $N_{\rm calls} = 697$ for a typical Planck run, without step iteration.

Having the Fisher matrix and its inverse, the user can easily write a small scripts to plot the ellipses corresponding to two-dimensional confidence levels in parameter space. This is not possible with `MontePython`, which can only do plots when some MCMC chains are present.

However, the `MontePython` plotting tools are useful for comparing the results of MCMC runs with the Gaussian posterior approximation given by the inverse Fisher matrix. For that purpose, one can analyse `MontePython` results with the usual `info` mode, adding just one input flag: `--plot-fisher`. Then the code will check whether a Fisher matrix has been computed and stored in the same directory as the chains that the user is trying to plot. If this is the case, the Fisher ellipses are drawn on top of the MCMC contours, like in figure 19.

This figure shows a sensitivity forecast based on mock BAO data from the DESI survey combined with Planck data, for a cosmological model with massive neutrinos and dynamical dark energy (νwCDM, 8 free parameters). The Fisher matrix was actually computed before launching the chains, and its inverse was used as an input covariance matrix. The final MCMC contours prove that in this case, the Fisher approximation is excellent. This inverse Fisher matrix does not only provide a good proposal density for MCMC runs, it also gives excellent estimates of parameter bounds, and it could be substituted to the whole MCMC results.

One could be in a situation in which a Fisher matrix is first computed around a guess for the best-fit point, and then used to launch MCMC chains that will be centered on the true best-fit point (in the case of a Gaussian posterior). In the comparison plot, one may find that the Fisher ellipses have the right shape, but are offset with respect to the true best-fit point. In order to get a nicer plot, the `MontePython` user can use the input flag `--center-fisher`. This will automatically center the Fisher ellipses on the maximum likelihood point extracted from the MCMC chains, instead of using the central values read in the `log.param` file, even if the Fisher matrix was actually computed in that point.

6.4 ILLUSTRATION OF PERFORMANCE

In order to illustrate the performance of `superupdate` and the impact of using an inverse Fisher matrix as input covariance matrix, we have chosen a few data sets and cosmological models, and performed some fits with or without these different options. The comparison is especially interesting in the most difficult situations: large number of free parameters, small prior knowledge (i.e. poor guess for the input covariance matrix), etc.

6.4.1 *Forecasts with a small prior knowledge*

We first run some MCMC forecasts for the combination of mock BAO data from the DESI[12] survey and Planck[13] data. We use the mock DESI likelihood called `fake_desi_vol` (documented in 6.5.4). For a forecast, we don't need to use real Planck data. We use instead a likelihood which simulates roughly the approximately of the Planck satellite, but uses some synthetic data corresponding to the Planck best-fit model. This likelihood is called `fake_planck_realistic` and is also documented in 6.5.4.

We fit these datasets with the minimal 6-parameter ΛCDM model and with several extended models featuring up to 12 free parameters. These extensions are listed in the first column of Table 13 and include massive neutrinos, dynamical dark energy with a constant equation of

12 http://desi.lbl.gov
13 http://sci.esa.int/planck/

Mock data: fake_planck_realistic, fake_desi_vol (see 6.5.4)

Running time: 12 hours

model	# param.	$R-1$: update	$R-1$: superupdate	$R-1$: Fisher + superupdate
ΛCDM	6	0.030	0.015	0.013
$+\sum m_\nu + w_0$	8	0.036	0.022	0.018
$+ N_{\rm eff}$ + running	10	not converged	not converged	0.040
$+ \Omega_k$	11	not converged	not converged	0.048
$+ w_a$	12	not converged	not converged	0.088

Running time: 48 hours

ΛCDM	6	0.0035	0.0029	0.0019
νw_0CDM + $N_{\rm eff}$ + running	10	0.014	0.0054	0.0038

Table 13: For mock data and several cosmological models, comparison of three sampling options, using the Gelman-Rubin convergence criterium. See text for details.

state, extra relativistic degrees of freedom, a running of the primordial spectrum index, spatial curvature, and finally dynamical dark energy with a CPL parametrisation [237, 238].

We run `MontePython` in these different cases with the Metropolis-Hastings algorithm and three different methods:

- `update` [`--update`]: periodical update of the covariance matrix,
- `superupdate` [`--superupdate`]: additional adaptation of the jumping factor,
- `superupdate` + Fisher [first `--method Fisher`; then `--superupdate`]: same but starting from the inverse Fisher matrix computed by `MontePython`.

For the `update` and `superupdate` runs, the proposal density is initialised as the "Planck 2015 covariance matrix", i.e. as the covariance matrix publicly distributed with the `MontePython` package, derived from the analysis of a well-converged run based on the Planck 2015 likelihoods and assuming the 6-parameter ΛCDM model. Also, in these two runs, the jumping factor is initially set to 2.4 (thus it remains equal to this value with the `update` method). For the Fisher matrix calculation, we pass to the code the exact best-fit model used to generate the mock data.

For each model and method, we launch the code with 8 chains, where each chain is running on 6 cores, using a total of 48 cores. After either 12 or 48 hours, we compute the worse [181] convergence criterium ($R-1$) over all parameters, removing the initial 10-20% of each chain (depending on the duration of the burn-in phase, but always the same for a given combination of models and experiments).

The difficulty of these runs reside in the poor guess for the input covariance matrix. In the six parameter runs, the input covariance matrix is derived from Planck data alone, while the DESI BAO data is very constraining. This means the proposal density is much too wide initially, and needs to shrink to the small region allowed by DESI data. When adding extra parameters, the situation is even worse. For the extra parameters, the code does not rely on

the input covariance matrix, but on the standard deviations written in the input file (for which we plug the Planck error bars). Therefore, the proposal density needs to learn both the correct order of magnitude for the jumps in these new directions, and the parameter correlations involving the extra parameters.

We find that, for the simplest models (6 and 8 parameters), all three methods successfully obtain at least a convergence of $R-1=0.03$, although `superupdate` and `superupdate` + Fisher perform better, obtaining an $R-1$ up to a factor 2 smaller. For the more complicated models (10, 11 and 12 parameters), starting from a Fisher matrix and using `superupdate` makes a big difference, as *only* the runs starting from a Fisher matrix managed to obtain any level of convergence, when limiting ourselves to only 12 hours of runtime. However, if we allow for longer runtime (48 hours) the `update` and `superupdate` methods also manage to converge, thanks to periodic updates of the covariance matrix, with the `superupdate` and `superupdate` + Fisher runs showing a factor of 2.6 to 3.7 better convergence than `update` alone. Figure 20 explicitely shows why the jumping factor adaptation and the Fisher matrix calculation result in a very significant speed up for the convergence of this run.

6.4.2 *Current data*

For a comparison of the efficiency of our new methods using current data, we consider only the 6-parameter ΛCDM model, that we fit to two data sets: a small set with just Planck and BAO likelihoods, and a larger one including Large Scale Structure (LSS) likelihoods (galaxy clustering from SDSS and weak lensing from CFHTLens). More details are given in Table 14. We perform again some fits in three different ways (`update`, `superupdate`, `superupdate`+Fisher), exactly like in the previous section (i.e. starting `update` and `superupdate` from the "Planck 2015 covariance matrix" distributed with the code, and from a jumping factor 2.4). We use the same number of chains and cores as in the forecasts, and allow the chains to run for 12 or 48 hours.

Like in the previous section, these runs illustrate the case of starting from a bad guess for the proposal density, because the input covariance matrix takes only Planck into account and needs to shrink to the smaller region compatible with BAO and LSS data. There are other significant differences with respect to the runs of the previous section. First, when we use the "small" dataset, we have all the nuisance parameters of the Planck high-ℓ TT likelihood, which have strongly non-Gaussian posteriors and are correlated with each other. This means that the optimal jumping factor is significantly different from 2.4 (it is actually closer to 1.9). It also means that the Fisher matrix calculation is difficult, due to the large number of parameters, the non-Gaussianity of the likelihood with respect to some parameters, and the fact that we only have a poor approximation of the best-fit point in parameter space (we compute the Fisher matrix in the approximate best-fit extracted from the chains of an earlier run with Planck data only). With the extended data set, the code actually fails to obtain an invertible Fisher matrix with the full Planck TTTEEE + BAO + LSS data, so we had to switch to the Planck-lite TTTEEE likelihood in order to get rid of nuisance parameters.

The run with the small dataset shows the impact of the automatic jumping factor update: with `superupdate`, the code rapidly adapts the jumping factor to about 1.9, while with `update` it remains stuck at 2.4, leading to a small acceptance rate. Table 14 shows a gain in $(R-1)$ by a factor of two when using `superupdate`. However, this run also shows that using the inverse Fisher matrix is not always a good idea with current data and many non-Gaussian parameters, because the Fisher Matrix can be such a poor approximation of the likelihood (especially in the

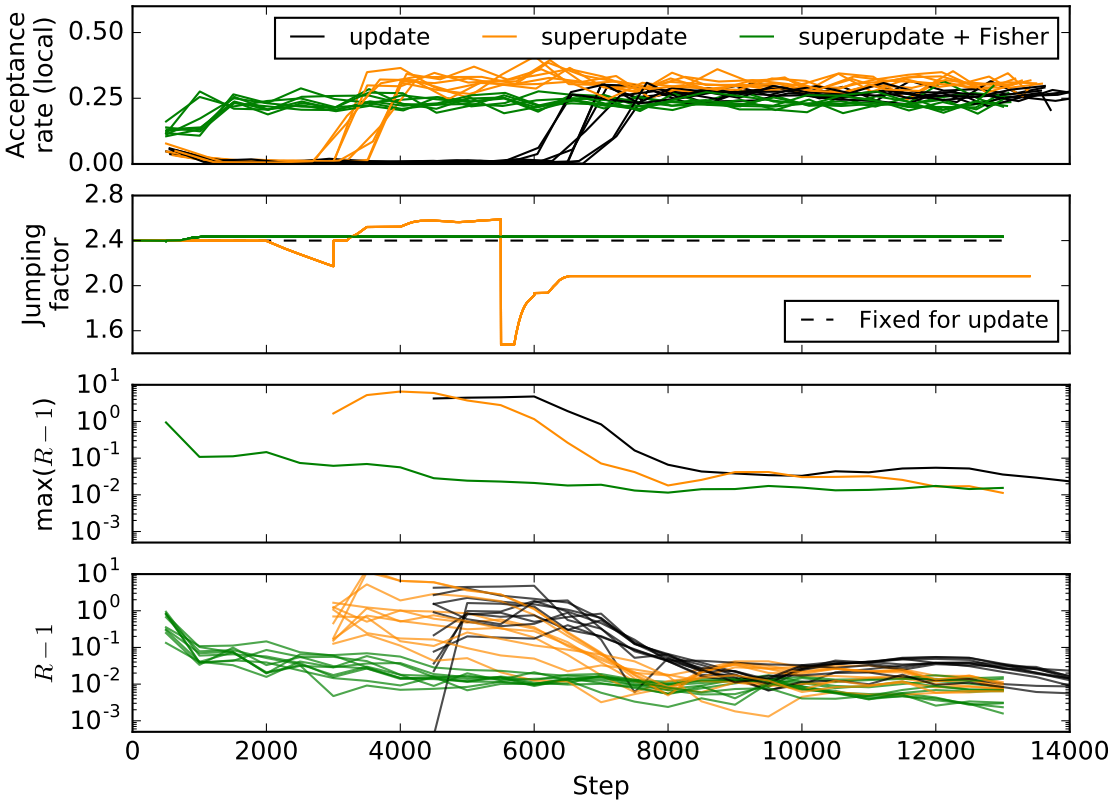

Figure 20: Evolution of the acceptance rate, jumping factor and convergence estimators for the last run from Table 1 (10 parameter model, mock Planck+DESI, 48h) using $U = 50$, $SU = 20$, FPM$= 10$, and thus $N_{\text{update}} = 500$. The jumping factor information is updated at each step, and the information on $(R-1)$ every N_{update} step (computed over the last 50% of the chain). This information can always be extracted from the code output. Instead, the local acceptance rate of the upper pannel was computed by post-processing the chains for the purpose of this plot, and was defined by averaging over about 500 steps (so this is slightly different from the quantity $\overline{a.r.}$ used by the superupdate algorithms, which is only averaged over $SU \times$ FPM=200 steps). The "Fisher" run essentially catches the right covariance matrix, jumping factor and acceptance rate from the beginning. The "superupdate" run reduces its jumping factor in order to quickly accumulate many points and get a good covariance matrix estimate; once this is done, it increases the jumping factor to avoid a too big acceptance rate. Finally, the "update" run needs about 3500 more steps before entering into an efficient sampling regime with a good covariance matrix and acceptance rate, which corresponds to about 12 hours on our 48 cores: thus we can say that in this particular example, "superupdate" saved about 600 core-hours for a single run.

direction of the non-Gaussian nuisance parameters) that it is actually a worse input covariance matrix than the one derived from MCMC chains for a previous Planck-only run. Therefore, the preliminary Fisher calculation degrades the performance by a factor three compared to `superupdate` + the "Planck 2015 covariance matrix".

The run with the large dataset (but with the Planck-lite TTTEEE likelihood) shows instead the same trend as the forecasts: both `superupdate` and Fisher bring significant improvement , by up to a factor two in $(R-1)$.

Planck 2015 (highl TT, lowℓ, lensing) + BAO (MGS, 6dFGS, LOWZ, CMASS)

Running time: 12 hours

model	# param.	$R-1$: update	$R-1$: superupdate	$R-1$: superupdate + Fisher
ΛCDM	6	0.019	0.0098	0.029

Planck 2015 (highl TTTEEE lite, lowℓ, lensing) + BAO (MGS, 6dFGS, LOWZ, CMASS) + galaxy clustering (SDSS DR7 LRG), weak lensing (CFHTLenS)

Running time: 12 hours

ΛCDM	6	0.042	0.032	0.018

Running time: 48 hours

ΛCDM	6	0.0062	0.0047	0.0038

Table 14: For current data and the ΛCDM model, comparison of three sampling options, using the Gelman-Rubin convergence criterium. See text for details.

These different situations bring us to the following conclusion, which match several other tests that we have performed and not included here:

- using `superupdate` is essentially always a good idea. The only situations in which one could consider sticking to `update` are the easiest ones, i.e. when a new run involves a dataset and a model so similar to a previous run that we already have an excellent knowledge of the covariance matrix and of the optimal jumping factor. In that case, `update` and `superupdate` are nearly equivalent, but in the most unlucky situations, `superupdate` could have a transitory phase during which the jumping factor would go away from the optimal value before going back to it asymptotically, and would lose a bit in efficiency. This is normally marginal and we can safely recommend to use `superupdate` in all cases: then, depending on the "difficulty" of the run, the improvement will range from negligible to large.
- when there are many non-Gaussian parameters, such as the Planck nuisance parameters, the Fisher matrix computation often fails, and even when it does not fail, the inverse Fisher matrix is often a bad approximation for the proposal density, compared to any input covariance matrix that was inferred from chains with the same nuisance parameters. For example, this means that when using the full Planck high-ℓ likelihoods one should use the distributed "Planck 2015 covariance matrices", or one's own covariance matrices from previous runs, instead of the Fisher option. In almost all other cases, we found that computing and starting from the inverse Fisher matrix is a very powerful way to speed up convergence.

6.5 CONCLUSION AND SUMMARY OF FEATURES

We find that using `superupdate` and an inverse Fisher matrix as input covariance matrix reduces convergence time in most cases, and makes the process of obtaining convergence significantly simpler, due to much fewer trial and error runs being necessary.

The Fisher matrix computation is very quick, so we recommend that for forecasts (where the minimum is known) with a Metropolis-Hastings algorithm are preceded by a Fisher matrix computation, when an accurate covariance matrix is not available beforehand. Likewise, `superupdate` generally performs equal to or better than `update` alone, as it optimizes the acceptance rate, and we recommend its use for *all* Metropolis-Hastings runs.

We acknowledge that the calculation the Fisher matrix is not entirely robust in `MontePython v3.0.0`, since in difficult cases with many non-Gaussian parameters (such as the Planck nuisance parameters), the Fisher matrix found by the code can be non-invertible, due to numerical errors and/or a poor estimate of the best-fit parameter values. We expect that further progress can be made in the future in the minimization and Fisher algorithms. However, we have tested our new features in hundreds of runs (including the few cases detailed in section 6.4) and found them extremely convenient for saving CPU time.

6.5.1 *Extra parameterisations*

By design, any function in `CLASS` that has been incorporated into the `CLASS` python wrapper[14] can be accessed directly from `MontePython`, without any additional coding required. However, sometimes it can be useful to define a specific parametrization within `MontePython`. This is easily done in the `python/data.py` module, where the `MontePython` input parameters can be intercepted and re-defined or re-named before being passed to `CLASS`. This happens in the `update_cosmo_arguments()` function, where all varying cosmological parameters are iterated through, and any additional parametrisations that are desired can be included by adding a simple `if` statement similar to existing ones.

`MontePython` includes several reparameterisation of this type. Some of them just deal with ordinary cosmological parameters, e.g.:

- if the parameter `Omega_Lambda` is used as a `MontePython` input parameter, instead of being passed to `CLASS`, it is used for defining `h` through $h = \sqrt{(\omega_b + \omega_{cdm})/(1-\Omega_\Lambda)}$.
- if the parameter `Omega_L` is used as a `MontePython` input parameter, instead of being passed to `CLASS`, it is used for defining `omega_cdm` via $\omega_{cdm} = (1-\Omega_\Lambda)h^2 - \omega_b$.
- if the parameter `ln10^{10}A_s` is used as a `MontePython` input parameter, instead of being passed to `CLASS`, it is used for defining `A_s`.
- if the parameter `exp_m_2_tau_As` ($\equiv e^{-2\tau_{reio}} A_s$) is used as a `MontePython` input parameter, instead of being passed to `CLASS`, it is used for defining `A_s` (assuming that `tau_reio` is also being used).

The code includes many other re-definitions related to isocurvature modes, neutrinos, dark energy, etc. Below, we expand the discussion concerning a few of the implemented neutrino and dark energy re-parameterisations. The user can easily extend the list of re-parameterisation for her/his specific cases.

14 Note that "incorporating" a new `CLASS` parameter in the wrapper just consists of adding one line in `python/cclassy.pxd` with just a declaration of this parameter (e.g. `double my_param`). The declaration must be done within the structure to which the parameter belongs. Incorporating a new `CLASS` function also boils down to declaring it in this file. New coding in the file `python/classy.pyx` is only required when one wants to create a new function specific to the wrapper itself, rather than just interfacing a `CLASS` function.

6.5.1.1 *Neutrino hierarchy*

`MontePython` can sample the total neutrino mass, with the individual neutrino masses arranged according to the Normal Hierarchy (NH, with two less massive and one more massive neutrino) or Inverted Hierarchy (IH, with one less massive and two more massive neutrinos).

The quantities passed to `CLASS` are the individual neutrino masses, but the quantity we are interested in sampling is the sum of neutrino masses. Formally, this is done using as a varying `MontePython` parameter `M_tot_NH` or `M_tot_IH`. Then, for each sampled value of the total neutrino mass (M_ν), the individual neutrino masses (m_i) are calculated by solving the system of equations (see e.g. [33])

$$\begin{aligned} M_\nu &= m_1 + m_2 + m_3 \,, \\ \Delta m^2_{\rm atm} &= m_3^2 - m_\ell^2 \,, \\ \Delta m^2_{\rm sol} &= m_2^2 - m_1^2 \,, \end{aligned}$$

where ℓ is 1 for NH and 2 for IH, and $\Delta m^2_{\rm atm}$ and $\Delta m^2_{\rm sol}$ are the current central values of the mass splittings obtained from neutrino oscillation experiments [10] (for a more recent study see e.g. [40])

$$\begin{aligned} \textbf{NH:} \quad & \Delta m^2_{\rm atm} = 2.524 \times 10^{-3}\ {\rm eV}^2 \,, \\ & \Delta m^2_{\rm sol} = 7.50 \times 10^{-5}\ {\rm eV}^2 \,, \\ \textbf{IH:} \quad & \Delta m^2_{\rm atm} = -2.514 \times 10^{-3}\ {\rm eV}^2 \,, \\ & \Delta m^2_{\rm sol} = 7.50 \times 10^{-5}\ {\rm eV}^2 \,. \end{aligned}$$

Additionally, the three parameters related to extra relativistic or non-cold dark matter (ncdm) species `N_ur=0.00641`, `N_ncdm=3` and `T_ncdm='0.71611,0.71611,0.71611'` should be fixed as `cosmo_arguments`, to reflect the fact that have three distinct standard active neutrino species and no extra relativistic degrees of freedom (unless one is studying a scenario with extra relativistic relics, in which case `N_ur` should be varied). Note that `T_ncdm` gives, for each species, the temperature of the neutrinos in units of photon temperature. In reality, the standard neutrinos distribution contains slightly non-thermal distribution, while by default `CLASS` will treat them as thermal species. The price to pay is to have a temperature ratio adjusted to 0.71611 in order to get the right neutrino density in the non-relativistic regime, and a small contribution to `N_ur` in order to get the right density in the relativistic one. This is why with three massive species we advise to take `N_ur=0.00641` instead of `N_ur=0`, but this correction is anyway well below the sensitivity of current experiments.

6.5.1.2 *Degenerate massive ν's and varying $N_{\rm eff}$*

In addition to arranging the mass of the neutrinos in a neutrino hierarchy, it is possible to sample the total neutrino mass for a case with three massive neutrinos with degenerate mass. Although not a realistic scenario, it is often sufficient to use three degenerate neutrinos (see e.g. [50]), speeding up computations in the Boltzmann solver.

This is done via the input parameter `M_tot`, remembering to specify the `cosmo_arguments` from before, but this time with only one type of neutrino species, `N_ur=0.00641`, `N_ncdm=1` and `T_ncdm=0.71611`, and instead specifying the degeneracy of the neutrino species, `deg_ncdm=3`. The total neutrino mass is then simply divided by the number of massive neutrino species and the resulting particle mass is passed to `CLASS`.

Additionally, this allows for varying the effective number of relativistic species, N_{eff}, by using the degeneracy of neutrino species, `deg_ncdm`, as a varying cosmological parameter instead of a fixed quantity.

For completeness, it is possible to use only one or two degenerate massive neutrinos and the rest massless, but this has been shown to be slightly inaccurate for the precision of current experiments [18].

6.5.1.3 *Dynamical dark energy*

Many phenomenological dark energy models can be treated using the fluid sector of `CLASS`, which has several free parameters labelled as `_fld`. By default, this sector uses the PPF parameterisation [239], although real fluid equation can be restored by setting `use_ppf` to `'no'` in the `cosmo_arguments`.

In principle, the dynamical dark energy equation of state parameters `w0_fld` and `wa_fld`, with a CPL parameterisation [237, 238] defined through $w(a) = p_{DE}/\rho_{DE} = w_0 + w_a(1 - a/a_0)$ (where, as usual, p is the pressure, ρ is the density, and a is the scale factor), can be passed directly to `CLASS`. However, it may be useful to sample the quantity $w_0 + w_a$ (implemented as `w0wa`) and w_0 (implemented `w0_fld`), in order to restrict the parameter space of $w_0 + w_a$ to only negative values (as in e.g. [26]).

6.5.1.4 *Sterile ν parametrization*

A final example of how a specific parameterisation can be introduced in MontePython is sterile neutrinos. In addition to degenerate massive neutrinos, one may wish to sample the sterile neutrino mass and the contribution of sterile neutrinos to N_{eff}, while avoiding the region of parameter space where the sterile neutrino mass becomes arbitrarily large and the contribution to N_{eff} becomes arbitrarily small (see Fig. 32. of [2]).

For this we defined the effective sterile neutrino mass $m_{s,\mathrm{eff}} \equiv m_s \Delta N_s$ as a possible varying cosmological parameter (`m_s_eff`), that should be used along with the parameter `deg_ncdm__2` standing for the contribution of sterile neutrinos to N_{eff}. This case is a bit more complicated than the others, as, in addition to setting `N_ncdm=2` and `T_ncdm='0.71611,0.71611'`, we also need to set the degeneracy of normal neutrinos `deg_ncdm__1` as a 'phantom' varying cosmological parameter, but with the parameter **fixed to 3**. The mass of the degenerate active neutrino species can also be varied as `m_ncdm__1`. This means the output of the chains will be the mass of a single active neutrino, rather than the sum, but of course we know that in this case $M_\nu =$ `m_ncdm__1` $\times$ `deg_ncdm__1` $= 3$ `m_ncdm__1`.

The effective sterile neutrino mass is then converted to physical sterile neutrino mass within the `data.py` module, in the function `update_cosmo_arguments()`, by dividing with ΔN_s (assuming that this is the ncdm species number 2 and that it is Dodelson-Widrow-like, i.e with the same temperature as active neutrinos). It is finally passed to `CLASS` along with the other neutrino masses.

6.5.2 *Sampling options*

`MontePython` has the following general sampling options

- `--method` : sampling method (MH, NS, CH, IS, Der, Fisher) [default: MH]
  which refer respectively to Metropolis-Hastings, Nested Sampling (= MultiNest), Cosmo

Hammer (= emcee), Importance Sampling, Derived (= reprocessing the chains to add columns with extra derived parameters requiring a new `CLASS` run for each model), and Fisher.

- `-T` : sample from the probability distribution $P^{1/T}$ instead of P [default: 1.0]

Options for Metropolis-Hastings and variants

- `--method MH` : Metropolis-Hastings sampling [default: MH]
- `--update` : proposal distribution update frequency in number of cycles [default: 50]
- `--superupdate` : also adapt jumping factor. Adaptation delay in number of cycles [default: 0] (i.e. deactivated by default. Recommended: 20)
- `--superupdate-ar` : target local acceptance rate [default: 0.26]
- `--superupdate-ar-tol` : tolerance for local acceptance rate [default: 0.01]
- `--adaptive` : running adaptation of covariance matrix and jumping factor (note: only suitable for single chain runs) [default: 0]
- `--adaptive-ts` : starting step for adapting the jumping factor [default: 1000]
- `-f` : jumping factor [default: 2.4]
- `--minimize` : attempt to re-evaluate starting point using a χ^2 minimization algorithm [by default uses `SLSQP` via `numpy.optiminize.minimize()`, can be changed in `sampler.py` function `get_minimum()`]
- `--minimize-tol` : tolerance for minimization [default: 10^{-5}]

Fisher matrix options

- `--method Fisher` : compute a Fisher matrix [default: MH]
- `--fisher-asymmetric` : allow for asymmetric steps (note: slows down computation) [default: False]
- `--fisher-step-it` : number of step iterations attempted [default: 10]
- `--fisher-delta` : target $\Delta \ln \mathcal{L}$ value for step iteration [default: 0.1]
- `--fisher-tol` : tolerance for $\Delta \ln \mathcal{L}$ (note: decreasing slows down computation) [default: 0.05]
- `--fisher-sym-lkl` : cut-off for switching to symmetric likelihood assumption in units of σ. Relevant when parameter space boundaries are close to the central value [default: 0.1]

`MontePython` also supports sampling with `MultiNest` (`--method NS`) [240–242] via a python wrapper [243] and `emcee` [244–246] via `CosmoHammer` (`--method CH`) [246]. For these sampling options we refer to the official documentation of those codes.

6.5.3 *Analyze and plotting options*

The range of plotting options and the general presentation of the plots has been significantly improved in `MontePython v3.0.0`. We should, however, point out that the user is free to use other plotting tools, if she/he prefers. In particular, the `MontePython` output is fully compatible with Antony Lewis's `GetDist`[15]. Note that `MontePython` writes in each output directory a file in the `.paramnames` format just for this purpose. `GetDist` has some very advanced plotting functionalities and a very nice graphical interface. However, the user will benefit from a few advantages when using the `MontePython` analyzing and plotting tools, such as: automatically evaluating the burn-in phase; automatically eliminating the non-Markovian part of the chains; and automatically reading information regarding the parameter names, ranges and scalings in the `log.param` file.

6.5.3.1 *Chain analysis*

When analyzing the chains, `MontePython` eliminates automatically the burn-in phase at the beginning of each chain, before applying additional cuts that can be customised with the options listed below. The burn-in phase of each chain is defined as: all the first points in the chains until an effective χ^2 value smaller than $\chi^2_{\rm min} + 6$ was reached for the first time. This number of 6 can be adjusted manually (it is equal to 2 `LOG_LKL_CUTOFF`, where `LOG_LKL_CUTOFF` is a parameter set in `montepython/analyze.py`, with a default value of 3). For runs in which a good estimate of the best-fit model was passed in input (with the option `-b <xxx>.bestfit`), the burn-in phase defined in this way may not exist at all.

Additionally, `MontePython` has the following options for analyzing chains (thus they should be written after the command line `python montepyhton/MontePython.py info`):

- `--keep-non-markovian` : keep the non-Markovian part of the chains [default: False].
- `--keep-fraction` : pass a decimal fraction, e.g. 0.8 to keep the last 80 % of the part of the chains that remain after the burn-in removal (note: redundant if non-Markovian points are discarded) [default: 1.0]
- `--want-covmat` : compute a covariance matrix based on the chains (note: this will overwrite the one produced by `--update`) [default: False]
- `--bins` : the number of bins for computing histograms [default: 20]
- `-T` : raise posteriors to the power T [default: 1.0]
- `--silent` : do not write any standard output (useful when running on clusters) [default: False]
- `--minimal` : use this flag to avoid computing posteriors, confidence limits and plots. The code just analyses the chains and outputs the files containing the convergence statistics, the best-fit parameters, and possibly the covariance matrix if `--want-covmat` is on [default: False]

15 http://getdist.readthedocs.io/en/latest/

Updating the proposal distribution or jumping parameter means that all prior steps in the chain are no longer Markovian, i.e. that each step should not depend on any prior steps. However, by using appropriate criteria for stopping adaptation of the jumping parameter and proposal distribution, and only including all steps after this point in our final analysis, we can ensure that our process was still Markovian. This is automatically done, but can be disabled with the command `--keep-non-markovian`, especially in slowly converging cases, when the user struggles to get a good covariance matrix that would stop the updating process, and wants to see some approximate results anyway. Although the burn-in phase is always removed, if non-Markovian steps are included the user may want to use the option `--keep-fraction <number>` in order to remove the first part of the chain.

6.5.3.2 *Basic plotting*

The most basic plotting features are implemented as command line options (but many of them can also be passed through an input customisation file, as we shall in 6.5.3.3):

- `--no-plot` : disable plotting [default: False]
- `--no-plot-2d` : only plot 1d posterior distributions [default: False]
- `--all` : output all individual 2D subplots and histogram files as separate files
- `--ext` : format and extension of the plot files (pdf, eps, png) [default: pdf]
- `--no-mean` : in 1D plot, do not plot the "mean likelihood" as dashed lines, only plot the posteriors as solid lines [default: False]
- `--contours-only` : line contours instead of filled contours [default: False]
- `--posterior-smoothing` : smoothing scheme for 1D posteriors: 0 means no smoothing, 1 means cubic interpolation, $n > 1$ means fitting $\ln(\mathcal{P})$ with a polynomial of order n [default: 5]
- `--interpolation-smoothing` : for 2D contours only, interpolation factor for getting a finer histogram before applying Gaussian smoothing and getting contours; 1 means no interpolation, increase for finer curves [default: 4]
- `--gaussian-smoothing` : for 2D contours only, width of Gaussian smoothing applied to histogram before getting contours, in units of bin size; increase for smoother contours, decrease for more exact results [default: 0.5]
- `--short-title-1d` : short 1D plot titles. Remove mean and confidence limits above each 1D plots. [default: False]
- `--num-columns-1d` : for 1D plots, number of plots per horizontal raw; if 'None' this is set automatically (trying to approach a square plot) [default: None]
- `--fontsize` desired fontsize [default: 16]
- `--ticksize` desired ticksize [default: 14]
- `--line-width` set line width [default: 4]

- `--decimal` number of decimal places on ticks [default to 3]
- `--ticknumber` number of ticks on each axis [default to 3]
- `--legend-style` specify the style of the legend, to choose from 'sides' or 'top' [default: sides]

When an Inverse Fisher matrix has been computed `--method Fisher`, the Fisher ellipses can be plotted on top of MCMC contours using the plotting options:

- `--plot-fisher` : plot inverse Fisher matrix contours [default: False]
- `--center-fisher` : centers Fisher ellipses on the parameters extracted from the best-fit model found in the chains, instead of the central starting values found in the input file [default: False]

6.5.3.3 *More advanced plot customisation*

Further options for customizing plots can be passed through a file with extension `.plot` called with the option `--extra`. All current functionalities are mentioned in the example file `plot_files/example.plot`, which the user would call with the plotting option `--extra plot_files/example.plot`. Although this file is self-explanatory, we list here the main functionalities provided by the use of `.plot` files. Several options have been present since the first release of `MontePython`:

- on-the-fly redefinition of the chain parameters with a simple syntax. For instance, if you know that there is a parameter called `A` and one called `B`, you can in principle replace the numbers in the column `A` by the result of any algebraic operation involving `A` alone, or `A` and `B`, or even more parameters, like e.g. `A + 3 A /B`. The file `plot_files/example.plot` provides the following example:

  ```
  info.redefine = {'omega_cdm': '(0.01*omega_b+omega_cdm)/(H0/100.)**2'}
  ```

  In this example, the code takes the numbers in the column `omega_b` and first multiplies them by 0.01, knowing that in the chains, ω_b was rescaled by 100 (this actually depends on what the user wrote in the input file). Thus, `0.01*omega_b` is the true ω_b, and `(0.01*omega_b+omega_cdm)/(H0/100.)**2` is in fact Ω_m. With the above command, each value of ω_{cdm} is replaced on-the-fly by Ω_m when the chains are read. The next necessary step is to change the name of the parameter for this column from `omega_cdm` to `Omega_m`, which can be done by the next functionality.

- redefinition of parameter name, for the purpose of redefinitions or making the parameter name better readable by the LaTeX routines of the plotting algorithm, e.g.

  ```
  info.to_change = {'omega_cdm': '$Omega_\mathrm{m}$' }
  ```

  will replace `omega_cdm` with `Omega_m`. Note that, for the purpose of getting a nice LaTeX format, `MontePython` already does several basic operations automatically, like identifying greek letters, subscripts and superscripts. Hence, at the time of producing a plot label, it would automatically convert `omega_cdm` into `$\omega_{cdm}$`. The functionality `info.to_change` is useful in order to further customise the LaTeX formatting.

- redefine the overall rescaling factor when the one from the input file is not optimal (scaling factors are useful e.g. to get rid of powers of ten in the plot captions, for very small or large parameters). This is done with the syntax `info.new_scales = {'A': 100}`.

- specify the list of parameters to be plotted (taking into account the new names, if there were name redefinitions). This is done with the `info.to_plot = [...]` syntax, which is very useful e.g. for getting rid of nuisance parameters in the 1D and 2D plots.

The new functionalities in `MontePython v3.0.0` are:

- parameters to control the legends: `info.plot_legend_1d`, `info.plot_legend_2d`, `info.legendnames` (see `plot_files/example.plot` for details).
- parameters to control the colors: `info.MP_color_cycle`, `info.MP_color`, `info.alphas` (see `plot_files/example.plot` for details).
- these lines simply overwrite the value of some parameters defined previously by the code within the python class `info`. Many other such lines can be added there, for instance `info.ticknumber = 5`, etc. Thus some of the options described previously as command line options can also be passed here, as lines of python.
- sometimes, the user would like to add some extra lines of python code in the plotting script, in order to further customise 1D or 2D plots, e.g. with vertical or horizontal lines, bands, arrows, labels, etc.. Usually, these lines are meant for only specific 1D or 2D plots. One can now achieve this by writing a few extra lines of python code in little files with a `.py` extension, which will be read and executed before finalizing the relevant plots. If they start with appropriate `if` statements, they will only be taken into account when plotting specific parameters. Some self-explanatory examples are provided together with the code in the files `plot_files/example.plot`, `add_h_contour.py`, and `add_sigma8_Omegam_contour.py`.

6.5.4 *Likelihoods*

Below is a comprehensive list of the likelihoods in the `MontePython` v3.0.0 package, as well as references to the paper(s) that should be cited when used (i.e. either where the likelihood was published and/or first used with `MontePython`). We recall that it is easy to modify these likelihoods or to create new ones. Some guidelines are given in the `MontePython` online documentation[16], in the section "Existing likelihoods, and how to create new ones". The column LU (Last Updated) shows the version number of the last modification.

In the column D (Dependencies), SC stands for self-contained; D means that some external data files must be downloaded; W means that we provide a wrapper to some external likelihood code that must be downloaded together with some data (as e.g. for Planck likelihoods); M means that this likelihood will automatically generate its own mock data, unless it has already been generated by a previous run. In the cases D, W, M, if you run the likelihood before downloading the required external files or before having created mock data, a self-explanatory message will tell you where to download from or what to do.

16 http://monte-python.readthedocs.io

Forecast likelihoods

name	description	type	LU	D	ref.(s)
core_m5	CORE M5 ESA proposal	CMB	3.0	M	[50] [247]
euclid_lensing	Euclid	Weak Lensing	3.0*	M	[118] [102]
euclid_pk	Euclid	Galaxy Clust.	3.0*	M	[118] [102]
fake_desi	DESI	BAO: d_A/r_s	3.0	M	[50]
fake_desi_euclid_bao	best from DESI + Euclid	BAO	3.0	M	[212]
fake_desi_vol	DESI	BAO: r_s/d_V	3.0	M	[157]
fake_planck_bluebook	Planck 2015 est.: TTTEEE	CMB	2.0	M	[248]
fake_planck_realistic	Planck 2018 est.: TTTEEE$\phi\phi$	CMB	3.0	M	[50]
litebird	LiteBIRD est.	CMB	3.0	M	[50] [95]
ska1_IM_band1	SKA1 band 1	21cm Int. Map.	3.0*	M	[102]
ska1_IM_band2	SKA1 band 2	21cm Int. Map.	3.0*	M	[102]
ska1_lensing	SKA1	Weak Lensing	3.0*	M	[102]
ska1_pk	SKA1	Galaxy Clust.	3.0*	M	[102]
ska2_lensing	SKA2	Weak Lensing	3.0*	M	[102]
ska2_pk	SKA2	Galaxy Clust.	3.0*	M	[102]

Table 15: Forecast likelihoods. * Euclid likelihoods will be updated and SKA likelihoods published when the relevant publication has been accepted for publication.

Current data likelihoods

name	description	type	LU	D	ref.(s)
acbar	ACBAR 2017	CMB	1.0	SC	[249]
bao	6dFGS BOSS DR9, SDSS DR7	BAO	1.1	SC	[209] [250] [210]
bao_known_rs	same as bao assuming known sound horizon value	BAO	1.1	SC	[251]
bao_angular	angular 2-point correlation function SDSS DR7: LRG BOSS DR10&11: CMASS BOSS DR12: QSO	BAO	3.0	SC	[252] [253] [254] [255] [256]
bao_boss	6dFGS, BOSS DR10&11: LOWZ, CMASS, SDSS DR7: MGS	BAO	2.0	SC	[209] [211] [210]
bao_boss_aniso	BOSS DR10&11: CMASS	BAO	2.0	SC	[211]
bao_boss_aniso_gauss_approx	BOSS DR10&11: CMASS	BAO	2.0	SC	[211]
bao_boss_dr12	BOSS DR12: LOWZ & CMASS	BAO	3.0	SC	[19] [257]
bao_fs_boss_dr12	BOSS DR12: LOWZ & CMASS	BAO+RSD	3.0	SC	[19] [257]
bao_smallz_2014	6dFGS, SDSS DR7: MGS	BAO	3.0	SC	[209] [210]
bicep	BICEP	CMB	1.0	SC	[258]
bicep2	BICEP2	CMB	2.0	SC	[88]
BK14	Bicep-Keck-Planck 2014	CMB	3.0	D	[259]
BK14priors	priors for the latter	CMB	3.0	D	[259]
boomerang	BOOMERanG	CMB	1.0	SC	[260]
cbi	CBIpol	CMB	1.0	SC	[261]
CFHTLens	CFHTLens as $\Omega_m^\alpha \sigma_8$ prior	Weak Lens.	2.1	SC	[262]
CFHTLens_correlation	full CFHTLens correlation	Weak Lens.	2.2	SC	[262]
clik_wmap_full	WMAP 7yr (through Planck wrapper)	CMB	1.2	W	[263]
clik_wmap_lowl	WMAP 7yr: low ℓ	CMB	1.2	W	[263]
cosmic_clocks_2016	cosmic clocks	$H(z)$	3.0	SC	[264]
cosmic_clocks_BC03	cosmic clocks	$H(z)$	2.1	SC	[265]
cosmic_clocks_MaStro	cosmic clocks	$H(z)$	2.1	SC	[265]

Table 16: Current data likelihoods (letters a-c)

name	description	type	LU	D	ref.(s)
cosmic_clocks_BC03_all	cosmic clocks	$H(z)$	2.1	SC	[265] [266] [267]
da_rec	prior on angular diameter distance	$d_A(z_{\rm rec})$	1.1	SC	[251]
gunn_peterson	constraints on reionization history	$x_e(z)$	1.0	SC	[268]
hst	Hubble Space Telescope	H_0 prior	3.0	SC	[85]
igm_temperature	constrains on baryon temperature	$T_b(z)$	1.0	SC	[269]
ISW	NVSS,2MPZ,WI×SC,SDSS/Planck	ISW	3.0	SC	[270]
JLA	full JLA likelihood	Supernovae	2.1	D	[271]
JLA_simple	simplified JLA likelihood	Supernovae	2.1	D	[271]
kids450_qe_likelihood_public	KiDS-450	Weak lensing	3.0	D	[148]
lowlike	Planck 2013 + WMAP 9: low-ℓ	CMB	1.2	W	[272]
Planck_actspt	ACT 2013, SPT 2011	CMB	2.0	W	[273] [274]
Planck_highl	Planck 2015: TT high ℓ	CMB	2.2	W	[275]
Planck_highl_lite	Planck 2015: TT high ℓ lite	CMB	2.2	W	[275]
Planck_highl_TTTEEE	Planck 2015: TTTEEE high ℓ	CMB	2.2	W	[275]
Planck_highl_TTTEEE_lite	Planck 2015: TTTEEE high ℓ lite	CMB	3.0	W	[275]
Planck_lensing	Planck 2015: lensing	CMB lensing	2.2	W	[208]
Planck_lowl	Planck 2015: TTTEEE low ℓ	CMB	2.2	W	[275]
Planck_SZ	Planck 2015: SZ cluster counts as $\Omega_m^\alpha\sigma_8$ prior	Cluster Count	2.2	SC	[276]
polarbear	Polarbear	CMB	2.1	SC	[277]
quad	QUAD DR3	CMB	1.0	SC	[278]
sdss_lrgDR4	SDSS DR4: LRG	Galaxy Clust.	3.0	SC	[279]
sdss_lrgDR7	SDSS DR7: LRG	Galaxy Clust.	3.0	SC	[280] [257]
simlow	from Planck 2016: TTTEEE low ℓ	$\tau_{\rm reio}$ prior	3.0	SC	[71] [281]
sn	Union2	Supernovae	1.0	SC	[282]
spt	SPT DR1	CMB	1.0	SC	[283]
spt_2500	SPT DR1, $\ell \leq 2500$	CMB	1.0	SC	[283]
timedelay	quasar time delays	Time Delay	1.1	SC	[284]
WiggleZ	WiggleZ power spectrum	Galaxy Clust.	2.0	SC	[285]
WiggleZ_bao	WiggleZ BAO	BAO	2.1	SC	[286]
wmap	WMAP 7yr (own wrapper)	CMB	1.0	D	[263]
wmap_9yr	WMAP 9yr (own wrapper)	CMB	1.2	D	[287]

Table 17: Current data likelihoods (letters c-z)

7

ΛCDM AND FUTURE SURVEYS

Chapter 7 is based on Di Valentino, Brinckmann, Gerbino, Poulin et al. [50]

7.1 ΛCDM AND DERIVED PARAMETERS

7.1.1 *Future constraints from CORE*

Adopting the method presented in section 2.4, here we forecast the achievable constraints on cosmological parameters from CORE in four configurations: LiteCORE-80, LiteCORE-120, CORE-M5 and COrE+. We work in the framework of the ΛCDM model, that assumes a flat universe with a cosmological constant, and is based on 6 parameters: the baryon $\Omega_b h^2$ and cold dark matter $\Omega_c h^2$ densities, the amplitude A_s and spectral index n_s of primordial inflationary perturbations, the optical depth to reionization τ, and the angular size of the sound horizon at recombination θ_s. Assuming ΛCDM, constraints can be subsequently obtained on "derived" parameters (i.e. that are not varied during the MCMC process) such as the Hubble constant H_0 and the r.m.s. amplitude of matter fluctuations on spheres of $8 Mpc^{-1}h$; σ_8. The ΛCDM model has been shown to be in good agreement with current measurements of CMB anisotropies (see e.g. [2]) and it is therefore necessary to first consider the improvements in sensitivity possible for a future CMB satellite experiment, such as CORE, on the parameters of the model.

Parameter	LiteCORE-80, TEP	LiteCORE-120, TEP	CORE-M5, TEP	COrE+, TEP
$\Omega_b h^2$	$0.022182 \pm 0.000052\,(2.9)$	$0.022180 \pm 0.000041\,(3.75)$	$0.022182 \pm 0.000037\,(4.0)$	$0.022180 \pm 0.000033\,(4.5)$
$\Omega_c h^2$	$0.12047 \pm 0.00033\,(4.1)$	$0.12049 \pm 0.00030\,(4.8)$	$0.12048 \pm 0.00026\,(5.4)$	$0.12048 \pm 0.00026\,(5.4)$
$100\theta_{MC}$	$1.040691 \pm 0.000097\,(3.2)$	$1.040691 \pm 0.000082\,(3.7)$	$1.040691 \pm 0.000078\,(4.0)$	$1.040693 \pm 0.000073\,(4.3)$
τ	$0.0598 \pm 0.0020\,(4.1)$	$0.0597 \pm 0.0020\,(4.5)$	$0.0597 \pm 0.0020\,(4.5)$	$0.0597 \pm 0.0020\,(4.5)$
n_s	$0.9619 \pm 0.0016\,(2.8)$	$0.9620 \pm 0.0015\,(3.0)$	$0.9619 \pm 0.0014\,(3.2)$	$0.9619 \pm 0.0014\,(3.2)$
$ln(10^{10} A_s)$	$3.0563 \pm 0.0037\,(3.9)$	$3.0562 \pm 0.0035\,(4.3)$	$3.0563 \pm 0.0035\,(5.1)$	$3.0562 \pm 0.0034\,(5.3)$
H_0[km/s/Mpc]	$66.96 \pm 0.14\,(4.4)$	$66.95 \pm 0.12\,(5.2)$	$66.96 \pm 0.11\,(5.6)$	$66.95 \pm 0.10\,(6.2)$
σ_8	$0.8173 \pm 0.0014\,(5.8)$	$0.8173 \pm 0.0012\,(7.4)$	$0.8172 \pm 0.0011\,(7.8)$	$0.8173 \pm 0.0010\,(8.6)$

Table 18: Forecasted sensitivity at 68% c.l. on cosmological parameters assuming standard ΛCDM for the CORE-M5 proposal and for three other possible CORE experimental configurations. The dataset used includes TT, EE, TE angular spectra and information from Planck CMB lensing. The numbers in parenthesis show the improvement $i = \sigma^{Planck}/\sigma^{CORE}$ with respect to the current constraints coming from the Planck satellite.

Our results are reported in Table 18, where we show the sensitivity at 68% c.l. on the cosmological parameters from CORE-M5 and we compare the results with three other possible experimental configurations: LiteCORE-80, LiteCORE-120 and COrE+. Besides the standard 6

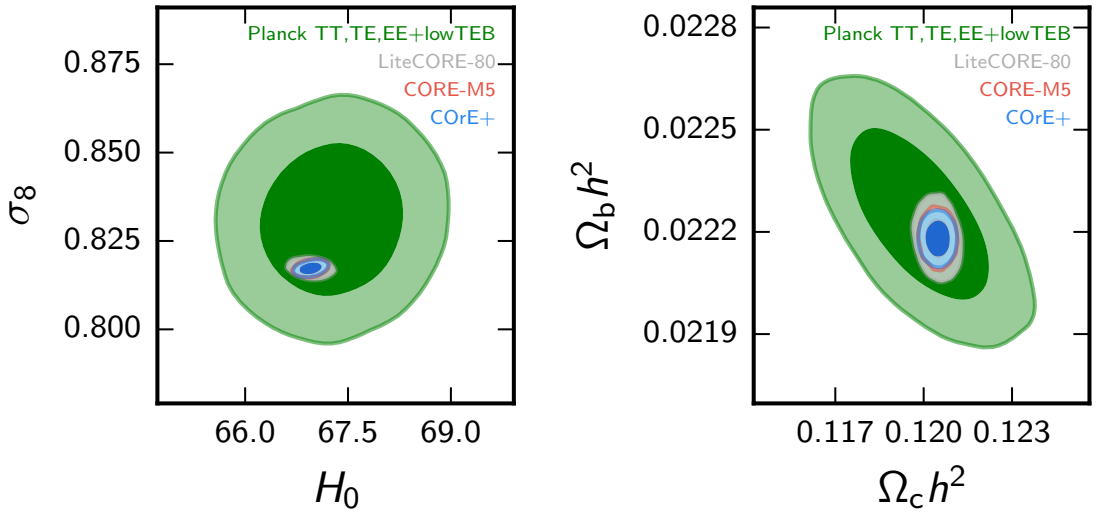

Figure 21: 2D posteriors in the σ_8 vs H_0 plane (left panel) and on the $\Omega_b h^2$ vs $\Omega_c h^2$ plane (right panel) from the recent Planck 2015 data release (temperature and anisotropy) and from the simulated LiteCORE-80, CORE-M5 and COrE+ experimental configurations. ΛCDM is assumed for the CORE simulations. The improvement of any CORE configuration in constraining parameters with respect to Planck is clearly visible.

parameters we also show the sensitivity obtained on derived parameters such as the Hubble constant H_0 and the amplitude of density fluctuations σ_8.

7.1.2 *Improvement with respect to the Planck 2015 release*

In Table 18, we show the improvement in the accuracy with respect to the most recent constraints coming from the TT, TE and EE angular spectra data from the Planck satellite [71], simply defined as $i = \sigma^{Planck}/\sigma^{CORE}$. As we can see, even the cheapest configuration of LiteCORE-80 could improve current constraints with respect to Planck by a factor that ranges between ~ 3, for the scalar spectral index n_s, and ~ 6, for the σ_8 density fluctuations amplitude. The most ambitious configuration, COrE+, could lead to even more significant improvements: up to a factor ~ 8 in σ_8 and up to a factor ~ 6 for H_0, for example. Similar constraints can be achieved by the proposed CORE-M5 configuration. The improvement with respect to current Planck measurements is clearly visible in Figure 21, where we show the 2D posteriors in the σ_8 vs H_0 plane (left panel) and on the $\Omega_b h^2$ vs $\Omega_c h^2$ plane (right panel) from the recent Planck 2015 data release (temperature and polarization) and from the LiteCORE-80, CORE-M5 and COrE+ experimental configurations. These numbers clearly indicate that, there is still a significant amount of information that can be extracted from the CMB angular spectra, even after the very precise Planck measurements.

It is important to note that, the most significant improvements are on two key observables: σ_8 and the Hubble constant H_0 that can be measured in several other independent ways. A precise measurement of these parameters, therefore, offers the opportunity for a powerful test of the standard cosmological model. It should indeed also be noted that the recent determination of the Hubble constant from observations of luminosity distances of Riess et al. (2016) [85] is in conflict at above 3 standard deviations with respect to the value obtained by Planck (see also [288, 289]). A significantly higher value of the Hubble constant has also recently been

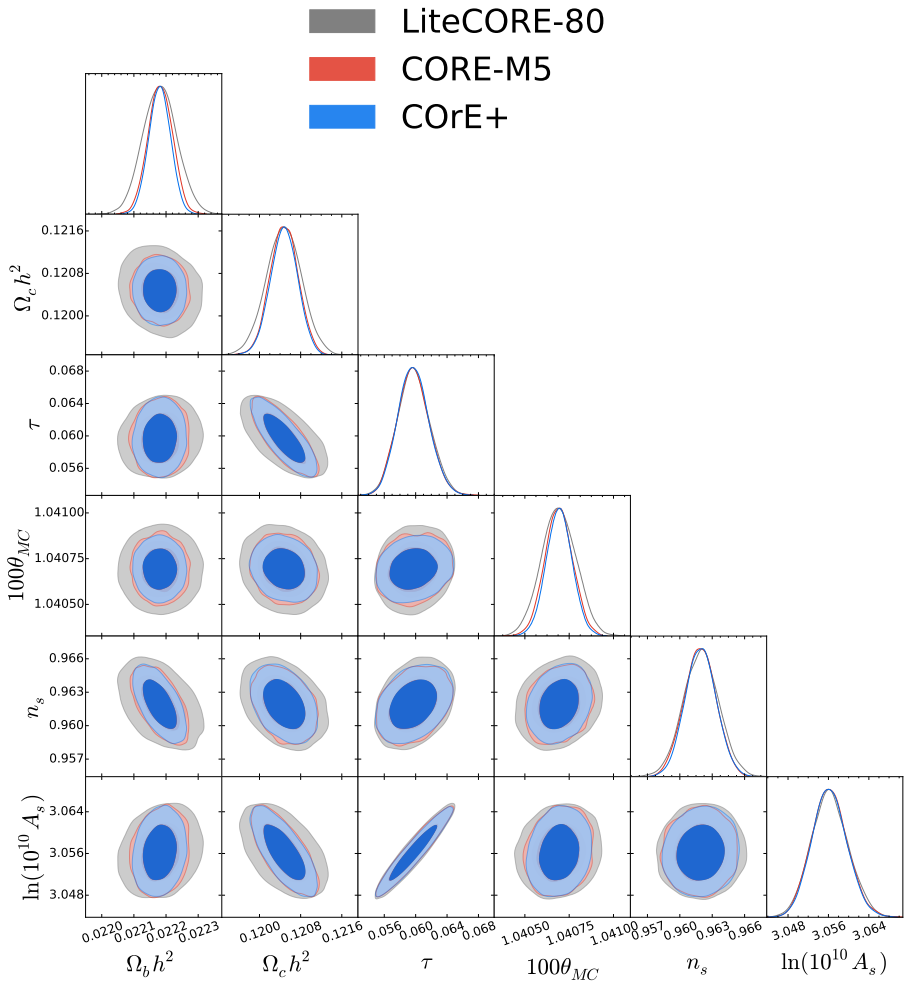

Figure 22: 2D posteriors for several combinations of parameters for the LiteCORE-80, CORE-M5 and COrE+ experimental configurations. ΛCDM is assumed as the underlying fiducial model.

reported by the HoLiCOW collaboration [290], from a joint analysis of three multiply-imaged quasar systems with measured gravitational time delays. Furthermore, values of σ_8 inferred from cosmic shear galaxy surveys such as CFHTLenS [291] and KiDS [292] are in tension above two standard deviations with Planck. While systematics can clearly play a role, new physics has been invoked to explain these tensions (see e.g. [86, 293–298]) and future and improved CMB determinations of H_0 and σ_8 are crucial in testing this possibility.

7.1.3 *Comparison between the different CORE configurations*

It is interesting to compare the results between the different experimental configurations as reported in Table 18 and as we can also visually see in Figure 22, where we show a triangular plot for the 2D posteriors from LiteCORE-80, CORE-M5 and COrE+.

We find four main conclusions from this comparison:

- When we move from LiteCORE-80 to COrE+ we notice an improvement of a factor ~ 1.6 on the determination of the baryon density $\Omega_b h^2$, and an improvement of a factor ~ 1.4 on the determination of the Hubble constant H_0 and the amplitude of matter fluctuations

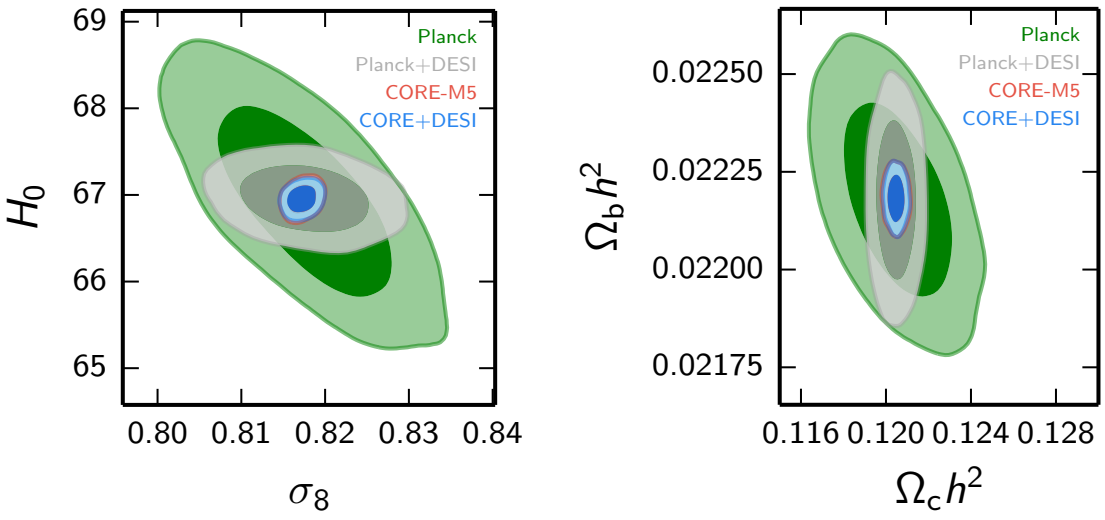

Figure 23: 2D posteriors in the H_0 vs σ_8 (left panel) and $\Omega_b h^2$ vs $\Omega_c h^2$ (right panel) planes from Planck (simulated), CORE-M5, and future BAO dataset from the DESI survey. ΛCDM is assumed as the underlying fiducial model.

σ_8. COrE+ is clearly the best experimental configuration in terms of constraints on these cosmological parameters. However, the CORE-M5 setup provides very similar bounds on these parameters as COrE+, with a degradation in the accuracy at the level of $\sim 10 - 12\%$.

- Moderate improvements are also present for the CDM density (of about ~ 1.3) and the spectral index (~ 1.14). The constraints from CORE-M5 and COrE+ are almost identical on these parameters.
- The constraints on the optical depth are identical for all four experimental configurations considered. This should not come as a surprise, since τ is mainly determined by the large angular scale polarization that is measured with almost the same accuracy with all the versions of CORE.
- Moving from COrE+ to CORE-M5 the maximum degradation on the constraints is about 12% (for the baryon density).

From these results, and considering also the contour plots in Figure 21 and Figure 22 that are almost identical between CORE-M5 and COrE+, we can conclude that CORE-M5, despite having a mirror of smaller size, will produce essentially the same constraints on the parameters with respect to COrE+ with, at worst, a degradation in the accuracy of just $\sim 12\%$.

7.1.4 *Constraints from CORE-M5 and future BAO datasets*

We have also considered the constraints achievable by a combination of the CORE-M5 data with information from Baryonic Acoustic Oscillation derived from a future galaxy survey as DESI. We found that the inclusion of this dataset will have minimal effect on the CORE-M5 constraints on ΛCDM parameters. This can clearly be seen in Figure 23, where we plot the 2D

posteriors in the H_0 vs σ_8 (left panel) and $\Omega_b h^2$ vs $\Omega_c h^2$ (right panel) planes. The CORE-M5 and the CORE+DESI contours are indeed almost identical.

It is also interesting to investigate whether the Planck dataset, when combined with future BAO datasets, could reach a precision on the ΛCDM parameters comparable with the one obtained by CORE-M5. To answer to this question we have simulated the Planck dataset with a noise consistent with the one reported in the 2015 release and combined it with our simulated DESI dataset. The 2D posteriors are reported in Figure 23: as we can see, while the inclusion of the DESI dataset with Planck will certainly help in constraining some of ΛCDM parameters, such as H_0 and the CDM density, the final accuracy will not be competitive with the one reachable by CORE-M5. In particular, there will be no significant improvement in the determination of σ_8 and the baryon density.

8

NEUTRINO MASS SENSITIVITY IN A MINIMAL MODEL

Section 8.1 is based on Di Valentino, Brinckmann, Gerbino, Poulin et al. [50], sections 8.2 and 8.3 are based on Archidiacono, Brinckmann, Lesgourgues & Poulin [157], and section 8.4 is based on Sprenger, Archidiacono, Brinckmann, Clesse & Lesgourgues [102]

8.1 NEUTRINO MASS SENSITIVITY IN A MINIMAL 7-PARAMETER MODEL

We fit the 7-parameter ΛCDM+M_ν model for different CORE settings, alone or in combination with mock DESI BAOs and Euclid cosmic shear data. Our forecasts consist in fitting these models to mock data, with a choice of fiducial parameters[1] slightly different from chapter 7, including in particular neutrino masses summing up to $M_\nu = 60$ meV.

Since we are looking at very small individual masses (mainly in the range $m_\nu < 100$ meV), we expect the sensitivity of the CMB to M_ν to be dominated by CMB lensing effects. The different CORE settings considered here lead to different sensitivities to the CMB lensing potential. However, we only observe marginal differences between the forecasted mass sensitivities shown in Table 19, with a symmetrized error ranging from 48 meV for LiteCORE-80 to 44 meV for CORE-M5 and COrE+. The reason is that the neutrino mass effect on the CMB lensing potential does not peak at the highest multipoles: rather it consists of a nearly constant suppression for a wide range of angular scales with $l > 100$. Hence, in order to achieve a good detection of M_ν, it is sufficient to have data in the region where the signal-to-noise ratio (S/N) is the largest, which is roughly from $\ell = 200$ to 700 for CMB lensing. Lensing extraction

1 The new choice of fiducial parameters is $\Omega_b h^2 = 0.022256$, $\Omega_c h^2 = 0.11976$, $100\theta_s = 1.0408$, $\tau = 0.06017$, $n_s = 0.96447$, $ln(10^{10} A_s) = 3.0943$, $M_\nu = 60$ meV, with neutrino masses ordered like in the Normal Hierarchy (NH) scenario.

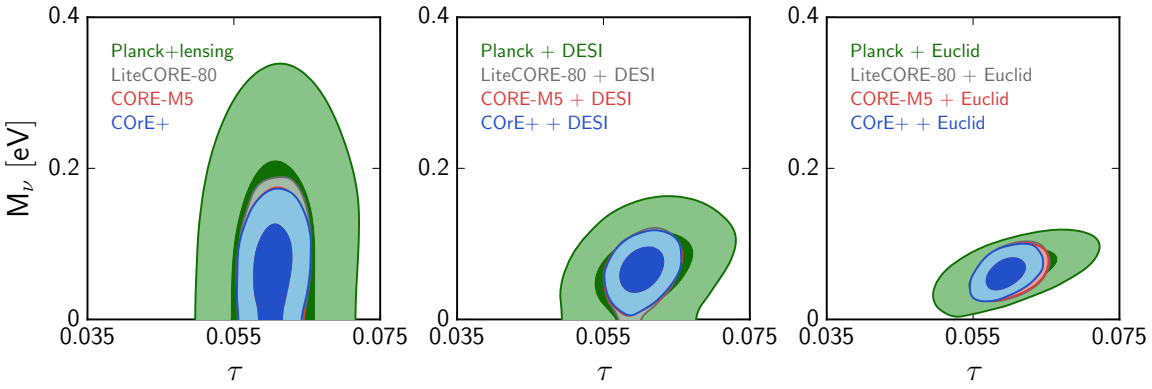

Figure 24: Results for the minimal model with massive neutrinos (discussed in section 8.1 and Table 19).

Parameter	Planck, TEP	LiteCORE-80, TEP	LiteCORE-120, TEP	CORE-M5, TEP	COrE+, TEP
M_ν (meV)	< 315 (68%CL)	78^{+36}_{-59}	74^{+38}_{-53}	72^{+37}_{-51}	72^{+38}_{-49}
$\Omega_b h^2$	0.02219 ± 0.00017	0.022250 ± 0.000059	0.022256 ± 0.000042	0.022256 ± 0.000039	0.022255 ± 0.000034
$\Omega_c h^2$	0.1198 ± 0.0015	$0.11987^{+0.00050}_{-0.00071}$	$0.11983^{+0.00048}_{-0.00067}$	$0.11980^{+0.00043}_{-0.00065}$	$0.11981^{+0.00044}_{-0.00066}$
$100\theta_s$	1.04069 ± 0.00036	1.04080 ± 0.00010	1.040800 ± 0.000087	1.040800 ± 0.000082	1.040800 ± 0.000079
τ	0.074 ± 0.017	0.0604 ± 0.0021	0.0604 ± 0.0021	0.0604 ± 0.0020	0.0604 ± 0.0021
n_s	0.9637 ± 0.0051	0.9643 ± 0.0020	0.9643 ± 0.0018	0.9644 ± 0.0018	0.9644 ± 0.0018
$ln(10^{10} A_s)$	3.081 ± 0.033	3.0951 ± 0.0042	3.0950 ± 0.0041	3.0949 ± 0.0040	3.0949 ± 0.0041
H_0 (km/s/Mpc)	$65.6^{+2.5}_{-1.4}$	$66.79^{+0.82}_{-0.44}$	$66.84^{+0.77}_{-0.40}$	$66.88^{+0.73}_{-0.39}$	$66.87^{+0.73}_{-0.39}$
σ_8	$0.783^{+0.040}_{-0.021}$	$0.828^{+0.011}_{-0.006}$	$0.8288^{+0.0098}_{-0.0059}$	$0.8293^{+0.0095}_{-0.0057}$	$0.8291^{+0.0093}_{-0.0058}$
Parameter	Planck, TEP + DESI	LiteCORE-80, TEP + DESI	LiteCORE-120, TEP + DESI	CORE-M5, TEP + DESI	COrE+, TEP + DESI
M_ν (meV)	72^{+36}_{-48}	65 ± 22	65^{+23}_{-20}	65^{+23}_{-19}	66^{+22}_{-20}
$\Omega_b h^2$	0.02226 ± 0.00012	0.022260 ± 0.000053	0.022259 ± 0.000041	0.022257 ± 0.000039	0.022258 ± 0.000033
$\Omega_c h^2$	0.11966 ± 0.00071	0.11970 ± 0.00031	0.11971 ± 0.00030	0.11971 ± 0.00028	0.11971 ± 0.00028
$100\theta_s$	1.04080 ± 0.00030	1.04080 ± 0.00010	1.040800 ± 0.000083	1.040800 ± 0.000078	1.040800 ± 0.000074
τ	0.0608 ± 0.0044	0.0603 ± 0.0021	0.0603 ± 0.0021	0.0602 ± 0.0021	0.0603 ± 0.0020
n_s	0.9647 ± 0.0027	0.9646 ± 0.0017	0.9646 ± 0.0015	0.9646 ± 0.0015	0.9645 ± 0.0015
$ln(10^{10} A_s)$	3.0954 ± 0.0086	3.0944 ± 0.0039	3.0944 ± 0.0039	3.0944 ± 0.0039	3.0944 ± 0.0039
H_0 (km/s/Mpc)	66.94 ± 0.28	66.98 ± 0.27	66.98 ± 0.26	66.98 ± 0.26	66.97 ± 0.26
σ_8	$0.829^{+0.012}_{-0.009}$	0.8309 ± 0.0043	0.8305 ± 0.0039	0.8305 ± 0.0038	0.8304 ± 0.0037
Parameter	Planck, TEP + DESI + Euclid	LiteCORE-80, TEP + DESI + Euclid	LiteCORE-120, TEP + DESI + Euclid	CORE-M5, TEP + DESI + Euclid	COrE+, TEP + DESI + Euclid
M_ν (meV)	65 ± 23	63^{+16}_{-18}	63^{+15}_{-17}	62 ± 16	62^{+15}_{-17}
$\Omega_b h^2$	0.02226 ± 0.00011	0.022256 ± 0.000052	0.022256 ± 0.000040	0.022256 ± 0.000037	0.022255 ± 0.000032
$\Omega_c h^2$	0.11977 ± 0.00032	0.11978 ± 0.00018	0.11977 ± 0.00018	0.11977 ± 0.00017	0.11979 ± 0.00018
$100\theta_s$	1.04080 ± 0.00031	1.040800 ± 0.000095	1.040800 ± 0.000084	1.040800 ± 0.000077	1.040800 ± 0.000072
τ	0.0606 ± 0.0046	0.0603 ± 0.0020	0.0603 ± 0.0021	0.0603 ± 0.0020	0.0595 ± 0.0020
n_s	0.9644 ± 0.0025	0.9645 ± 0.0016	0.9645 ± 0.0014	0.9645 ± 0.0014	0.9645 ± 0.0014
$ln(10^{10} A_s)$	3.0951 ± 0.0086	3.0945 ± 0.0039	3.0947 ± 0.0039	3.0947 ± 0.0038	3.0944 ± 0.0038
H_0 (km/s/Mpc)	$66.96^{+0.23}_{-0.17}$	$66.97^{+0.21}_{-0.14}$	$66.97^{+0.21}_{-0.13}$	$66.98^{+0.20}_{-0.13}$	$66.98^{+0.19}_{-0.13}$
σ_8	$0.8314^{+0.0039}_{-0.0030}$	$0.8316^{+0.0034}_{-0.0027}$	$0.8314^{+0.0033}_{-0.0026}$	0.8315 ± 0.0028	$0.8316^{+0.0031}_{-0.0024}$

Table 19: 68% CL constraints on cosmological parameters in the ΛCDM+M_ν model (accounting for the summed mass of standard neutrinos) from the different CORE experimental specifications and with or without external data sets (DESI BAOs, Euclid cosmic shear). For Planck alone, we quote the results from the 2015 data release, while for combinations of Planck with future surveys, we fit mock data with a fake Planck likelihood mimicking the sensitivity of the real experiment (although a bit more constraining).

on smaller angular scale will always have a smaller S/N and would bring little additional information. In the range $200 < \ell < 700$, LiteCORE-80 has a slightly worse sensitivity to the CMB lensing spectrum than other settings considered here, and hence a larger $\sigma(M_\nu)$; the other settings mainly differ for $\ell > 700$. We conclude that the determination of M_ν cannot drive the choice between different possible CORE settings, unlike the determination of other parameters (e.g. tensor-to-scalar ratio, $N_{\rm eff}$) that critically depend on the sensitivity and/or resolution of the instrument.

However, a next-generation CMB satellite is essential for getting such tight bounds on the summed neutrino mass, because of its potential to measure small-scale polarisation and to constrain the optical depth to reionization τ (this is true for all CORE configurations). Indeed, the suppression induced by neutrino masses in the CMB lensing potential could be nearly cancelled by an increase in the primordial spectrum amplitude A_s. Since the product $e^{-2\tau} A_s$

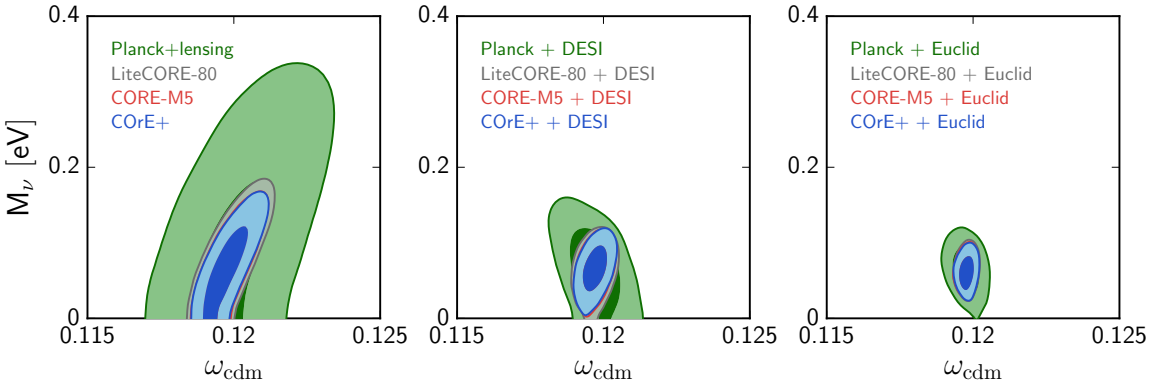

Figure 25: Results for the minimal model with massive neutrinos (discussed in section 8.1 and Table 19).

is fixed by the global amplitude of the CMB temperature/polarisation spectra, increasing A_s requires increasing τ. Future ground-based CMB experiments would only marginally improve on the τ determination from Planck, due to their limited sky coverage and large sampling variance for small multipoles. Hence, they would be affected by an (M_ν, τ) degeneracy for the reasons discussed above. To prove the importance of this effect we repeated the forecast for CORE-M5, but removing all polarisation information for $\ell < 30$ and replacing it with a Gaussian prior on τ with the sensitivity of Planck, $\sigma(\tau) \simeq 0.01$. We find a degeneracy between M_ν and τ and the error bar on the summed mass degraded by a factor 2. Instead, we can clearly see in the left panel of Figure 24 that there is no such degeneracy, neither in the Planck-alone contours, caused by too weak sensitivity to the CMB lensing spectrum, nor in CORE-alone contours because they break this degeneracy by measuring τ with good enough precision.

We can check how the combination of CMB data with other probes can achieve better constraints with CORE than with Planck. We find that CORE+DESI BAOs is about two times more constraining than Planck+DESI. This is related again to the better CMB lensing spectrum extraction *and* optical depth measurement by CORE. There are actually two ways to compensate the CMB lensing spectrum suppression induced by neutrino masses: by increasing A_s and τ, or by increasing $\omega_{\rm cdm}$ [157]. This leads to a strong $(M_\nu, \omega_{\rm cdm})$ degeneracy when using only CMB data (Figure 25, left plot). However, future BAO data will fix $\omega_{\rm cdm}$ with very good accuracy. In the Planck+BAO case, the (M_ν, τ) degeneracy would then still remain (Figure 24, middle plot). In the CORE+BAO case, with $\omega_{\rm cdm}$ fixed by BAOs and τ nearly fixed by polarisation measurements, very little degeneracies remain: in Figure 24, middle plot, we just see a small positive correlation controlled by the error bar on τ. Hence CORE will powerfully exploit the synergy between CMB and BAO measurements for measuring the neutrino mass. The combination with Euclid will further reduce degeneracies and errors by independently measuring the lensing spectrum at smaller redshifts than CORE. Even with very conservative assumptions on Euclid (i.e. including only cosmic shear data for $k \leq 0.5h/\text{Mpc}$) we find that CORE+DESI+Euclid would have a sensitivity of $\sigma(M_\nu) = 16$ meV, almost guaranteeing at least a 4σ detection.

8.2 COMBINATION OF CMB, BAO AND GALAXY SHEAR/CORRELATION DATA

In this section we will present the results of our Markov Chain Monte Carlo forecast of the combined sensitivity of future CMB, BAO and LSS experiments to the cosmological parame-

ters described in section 5.1, in particular to the neutrino mass sum.For these results we used the approach discussed in chapter 5. In particular, we already commented at the end of section 5.3.1 on our conservative choices for the precision parameters (following the approach of Ref. [118]): theoretical error parameter $\epsilon = 0.025$, cut-off at $k_{\rm max} = 0.6\,h/{\rm Mpc}$ for galaxy correlation, and at $\ell_{\rm max} = 2000$ for cosmic shear. Still this choice comes from a subjective estimate of the accuracy with which non linear corrections and systematic effects will be modelled in the future, and different assumptions would lead to different parameter sensitivities.

	$\sigma(M_\nu)/[\rm meV]$	$\sigma(\tau_{\rm reio})$	$\sigma(10^9 A_s)$	$\sigma(n_s)$	$\sigma(\omega_{\rm cdm})$	$\sigma(h)$
CORE	42	0.0020	0.0084	0.0018	0.00052	0.0052
CORE+DESI	19	0.0020	0.0080	0.0014	0.00026	0.0022
CORE+DESI+Euclid-lensing	16	0.0020	0.0078	0.0014	0.00023	0.0019
CORE+Euclid (lensing+pk)	14	0.0020	0.0079	0.0015	0.00025	0.0017
CORE+Euclid (lensing+pk)+21cm	12	−−	0.0042	0.0014	0.00021	0.0017

Table 20: Expected $1\,\sigma$ sensitivity of CORE, CORE + DESI, CORE + DESI + Euclid (lensing), CORE + Euclid (lensing+pk), CORE + Euclid (lensing+pk) + "21cm-motivated $\tau_{\rm reio}$ prior" to the parameters $\{M_\nu, \tau_{\rm reio}, 10^9 A_s, n_s, \omega_{\rm cdm}, h\}$. We did not combine DESI and Euclid-pk in order to avoid double counting the information coming from the wiggly part of the spectrum.

In the first four lines of table 20 we report the expected sensitivity of CORE, CORE+DESI, CORE+DESI+Euclid-lensing and CORE+Euclid (lensing+pk) to M_ν and other cosmological parameters playing a crucial role in our analysis of parameter degeneracies: $\tau_{\rm reio}$, $10^9 A_s$, n_s, $\omega_{\rm cdm}$ and h (the last independent parameter, $\omega_{\rm b}$, is always very well constrained by CMB data alone). In figure 26 we plot the one dimensional posteriors and the one- and two-σ marginalized contours for the same parameters.

First of all, we notice that the projected $1\,\sigma$ errors in table 20 and 1D distributions in figure 26 reflect the theoretical points we have discussed in the previous sections: both DESI and Euclid greatly improve the sensitivity to M_ν, $\omega_{\rm cdm}$ and h. The uncertainty on M_ν tightens by more than a factor two for CORE+DESI and a factor three for CORE+Euclid, compare to the CORE only sensitivity. The error on H_0 shrinks by a factor larger than two for CORE+DESI and a factor three for CORE+Euclid. However once more we want to stress that in the case of DESI the improved sensitivity arises from reducing the degeneracy between H_0 and M_ν, while in the case of Euclid the longer lever arm of the shear data is specifically sensitive to the suppression of power at small scales induced by M_ν.

The first column of figure 26 shows all the degeneracies with respect to M_ν. Let us describe the evolution of those correlations with the addition of the different datasets:

1. *CORE data only*. When only CMB data are considered, correlations follow the directions expected from our extensive discussion of section 5.1.3. Let us just note that contrarily to ΛCDM runs *without* neutrino mass as a free parameter, the mild correlation between A_s and n_s is negative, which is a result of the mild negative (resp. positive) correlation between M_ν and n_s (resp. A_s).

2. *Adding DESI data.* In general, the size of the 2D-distributions shrink by a factor ∼2. The extended regions defining the positive correlations between (M_ν, A_s) and $(M_\nu, \tau_{\rm reio})$ become steeper, since it is not possible anymore to play with H_0 or $\omega_{\rm cdm}$ to compensate the effect of the summed neutrino mass on the CMB lensing spectrum. Indeed, as described in section 5.2, moving along this degeneracy direction would lead to very different BAO

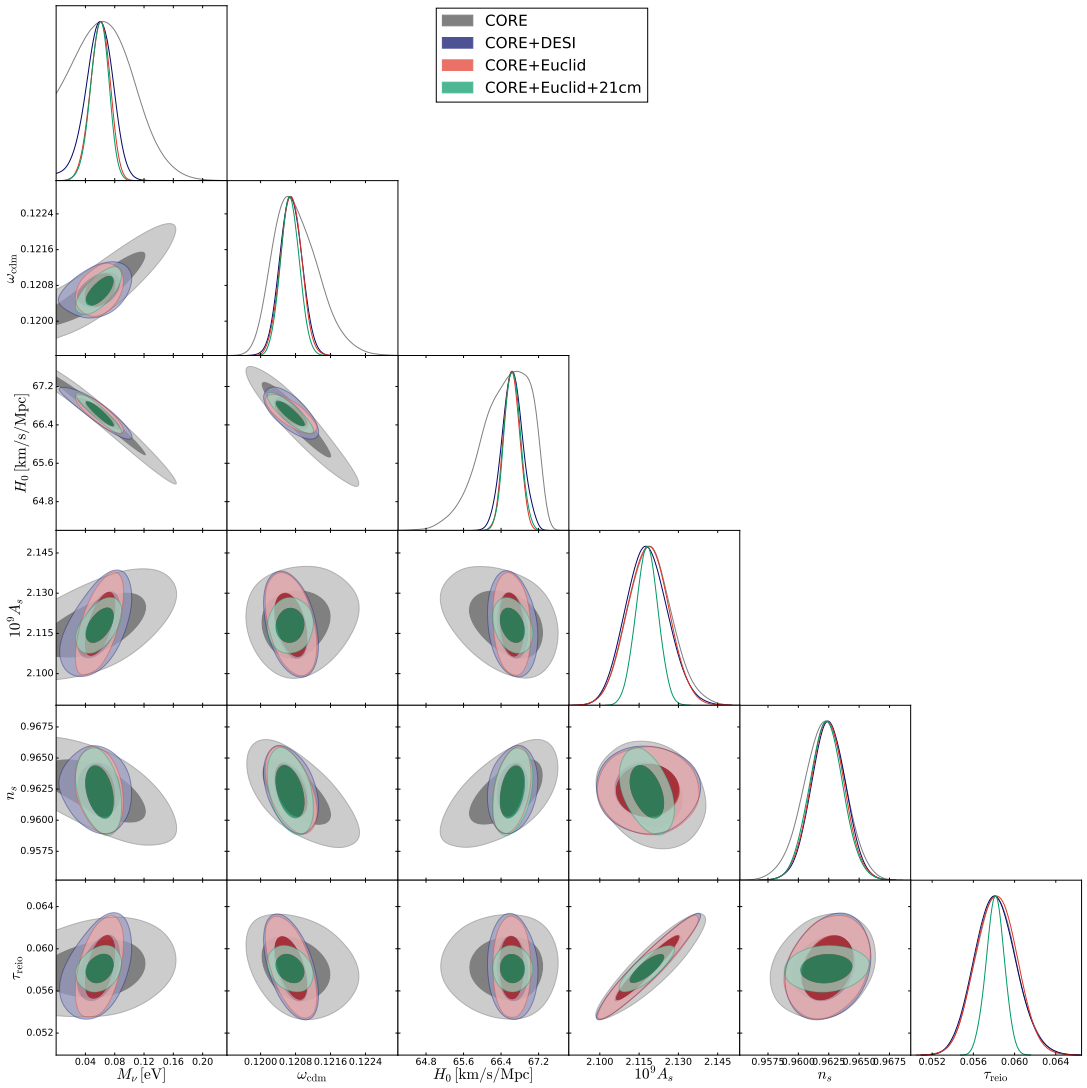

Figure 26: Marginalized one− or two−σ contours and one dimensional posteriors in the $(M_\nu, \omega_{\rm cdm}, H_0, A_s, n_s, \tau_{\rm reio})$ parameter space, showing the expected sensitivity of various future experiments: CORE only (gray contours), CORE+DESI (blue contours), CORE+Euclid (red contours) and CORE+Euclid+21cm (green contours). The last independent parameter, $\omega_{\rm b}$, is always very well constrained by CMB data alone.

angular scales. Thus, the effect of the summed neutrino mass on CMB lensing is rather compensated by playing with parameters to which BAO data are insensitive[2], namely A_s and $\tau_{\rm reio}$.

3. *Adding Euclid (lensing + $P(k)$) data*. Most of the discussion on the inclusion of DESI data still applies here, since Euclid data contains information on the BAO scale at different redshift. However the matter / shear power spectra contain extra information on cosmological perturbations, and lift or reinforce some parameter degeneracies, consistently with our previous discussion in section 4.2, point 3. The (M_ν, H_0) degeneracies get

2 As side remarks, note that such compensation cannot be done by playing with n_s: as a consequence, both the (M_ν, n_s) degeneracy and the (A_s, n_s) degeneracy are lifted when BAO data are added; finally, because of the different neutrino mass compensation driven by the inclusion of BAO data, the correlations of $\omega_{\rm cdm}$ and H_0 with respect to $A_s, n_s, \tau_{\rm reio}$ are lifted, as well.

considerably reduced because the LSS data would prefer a different correlation angle between these two parameters. As expected, the Euclid data considerably tightens the positive correlation between M_ν and A_s, and as a side effect the combined data leads to a clear positive correlation between M_ν and $\tau_{\rm reio}$. The degeneracy between M_ν and $\omega_{\rm cdm}$ is lifted by the weak lensing data. All these degeneracy reductions lead to an overall shrinking of all contours involving M_ν, H_0 and $\omega_{\rm cdm}$ by a factor of order 3 between CMB and CMB+LSS data. The neutrino mass value is accurately determined independently of the value of n_s, and the mild correlation between M_ν and n_s in CMB data disappears with additional LSS data.

Compared to figure 26, table 20 presents the results of one more MCMC run featuring CORE, DESI and Euclid weak lensing, but not the Euclid galaxy clustering information. The comparison of these results with those for CORE and Euclid weak lensing+galaxy clustering show the importance of geometrical information (BAO angular scales) versus shape information (full matter power spectrum), although both runs do contain some shape information coming from the weak lensing data. We clearly see that adding more shape information on the matter power spectrum benefits only to the determination of M_ν and H_0, and actually by a modest amount (10 to 15% per cent). At face value, this means that even if the analysis of future galaxy clustering data was plagued by unexpected systematics (besides the level that we conservatively took into account with our theoretical error bar), the prospects to accurately determine the summed neutrino mass with future surveys would not collapse.

In order to further improve the measurement of the neutrino mass with cosmological data, one should try to add independent constraints on the parameters that remain most strongly correlated with M_ν in the CMB+LSS contours: these are H_0, A_s and $\tau_{\rm reio}$. The role of a very precise determination of H_0, free of astrophysical systematics, for the measurement of the neutrino mass, has already been stressed e.g. in [17]. It was also previously noticed in Ref. [198] that 21cm surveys could improve the determination of the optical depth to reionization, and thus of the summed neutrino mass. Having understood the physical explanation for the $(M_\nu, \tau_{\rm reio})$ degeneracy, we wish to further investigate this possibility, while keeping our conservative assumption on the matter power spectrum theoretical error.

8.3 ADDING 21CM SURVEYS

In the near future, many experimental efforts will be devoted to measuring precisely the epoch of reionization (EoR), mostly through the 21 cm line created by the hyperfine transition of the Hydrogen atom[3], including the value of $\tau_{\rm reio}$. In general, details of the EoR are strongly connected to fundamental questions in cosmology and astrophysics. They could shed light on many properties of the first galaxies and quasars, measure the time at which they form, explain how the formation of very metal-poor stars proceeded, and reveal whether the first galaxies were indeed the only re-ionizing source.

However, these experiments can also have great implications for neutrino physics in cosmology. Indeed, the independent measurement of the epoch of reionization by 21cm surveys may break the degeneracy between A_s and $\tau_{\rm reio}$ [198, 299] which appears in combined analyses of future CMB+LSS data.

3 e.g. PAPER 64: http://eor.berkeley.edu, 21CMA: http://21cma.bao.ac.cn, MWA: http://www.mwatelescope.org, LOFAR: http://www.lofar.org, HERA: http://reionization.org or SKA: http://www.skatelescope.org.

To assess the impact of 21cm surveys on $\sigma(M_\nu)$, we performed a final MCMC run combining CORE+Euclid mock data with a gaussian prior on the value of $\tau_{\rm reio}$. In agreement with forecasts on the sensitivity of HERA or SKA, we fixed the prior variance to $\sigma(\tau_{\rm reio}) = 0.001$ [6, 198]. Note that by doing so, we are being conservative, since 21cm surveys will not only measure the evolution of the mean free electron fraction $x_e(z)$ (and thus the optical depth $\tau_{\rm reio}$), but also the power spectrum of the 21cm signal at different redshifts, $P_{21cm}(k, z)$, related to variations along the line of sight of the free electron fraction $x_e(\hat{n}, z)$ [300]. We are therefore using the minimal amount of information that one can extract from these experiments, and one could go beyond following e.g. the procedure of Refs. [198, 299, 301, 302].

The results of our MCMC forecast are summarized by the last line of table 20, and the green contours in figure 26.

The main impact of the $\tau_{\rm reio}-$prior is to reduce the possibility of varying of A_s, necessary to adjust the CMB parameter $A_s \exp(-2\tau_{\rm reio})$, by almost a factor two. Since M_ν was correlated directly with A_s and indirectly with $\tau_{\rm reio}$, the sensitivity to the summed neutrino mass also improves thanks to 21cm data, going from $\sigma(M_\nu) = 14$ meV for CORE+Euclid to 12 meV. As a side effect, the positive correlation between M_ν and $\omega_{\rm cdm}$ and the negative correlation between M_ν and h get steeper.

Thus, even if nature has chosen the summed neutrino mass to be close to the lower limit of the normal hierarchy, $M_\nu = 60$ meV, we expect that the joint analysis of CORE + Euclid + 21cm data will detect it at more than 5 σ.

8.4 IMPROVED TREATMENT OF LARGE-SCALE STRUCTURE

It is well known from neutrino oscillation experiments that at least two neutrinos are massive. However the absolute value of the mass has not been determined yet, neither by cosmology (for up-to-date upper limits see Refs. [15, 20, 40]), nor by β-decay experiments (in this regard see the sensitivity of the forthcoming KATRIN experiment [303]). One of the biggest achievements of Euclid and SKA will be to pin down the neutrino mass sum M_ν. Therefore, our baseline model will be ΛCDM+M_ν, parameterized as follows:

$$\{\omega_{\rm b},\ \omega_{\rm cdm},\ 100\times\theta_{\rm s},\ \tau_{\rm reio},\ \ln(10^{10}A_s),\ n_s,\ M_\nu\}. \tag{303}$$

Our fiducial model assumes a minimal value of the total neutrino mass and some Planck inspired values for other parameters:

$$\{0.02218,\ 0.1205,\ 1.04156,\ 0.0596,\ 3.056,\ 0.9619,\ 0.06\ \text{eV}\}. \tag{304}$$

We assume the total neutrino mass sum M_ν to be equally split among the three active neutrino species. This degenerate neutrino mass scheme is motivated by the fact that the deviation of its theoretical predictions both from the normal mass ordering and from the inverted mass ordering is negligible compared to the sensitivity of current and forthcoming cosmological data [18, 67, 74]. For a detailed discussion of the physical effects involved in the measurements of M_ν with galaxy clustering and cosmic shear, and for the impact of 21-cm surveys, see chapter 5 or Refs. [157, 304]. Here we just mention that low redshift measurements are sensitive to massive neutrinos because their free-streaming induces a relative suppression of the linear matter power spectrum on scales smaller than the free-streaming scale after the neutrino non-relativistic transition [66, 67, 189, 190]. On top of the linear suppression, an additional dip appears at non-linear scales, caused by the delay of the onset of the non-linear growth in

neutrino cosmologies [152]. Euclid and SKA span a broad range of redshifts and scales, where the aforementioned effects on the shape of the matter power spectrum can be detected, as long as an accurate theoretical prediction is provided [121, 193, 194, 196, 197].

This section makes use of the approach of chapter 3 and we take into consideration the results of chapter 4, i.e. remember that Refs. [120, 121] have shown that in massive neutrino cosmologies the clustering properties of halos are determined by cold dark matter and baryons only (hereafter, *cb*), rather than the total matter field (i.e., cold dark matter + baryons + massive neutrinos). Therefore, the galaxy power spectrum must be reconstructed by taking into account only the *cb* field, ignoring the contribution of light massive neutrinos with a free-streaming length far larger than the typical size of a galaxy. Neglecting this effect can lead to sizeable errors, as we showed in chapter 4 (see also Refs. [1, 127]). Following this approach, already used in Refs. [170, 305], we have modified Eq. 164 of the observed galaxy power spectrum as follows:

$$P_g(k,\mu,z) = f_{\rm AP}(z) \times f_{\rm res}(k,\mu,z) \times f_{\rm RSD}(\hat{k},\hat{\mu},z) \times b^2(z) \times P_{cb}(\hat{k},z)\ . \tag{305}$$

Moreover, the β factor of the Kaiser formula, i.e. the ratio between the growth rate and the bias, embedded in the third term of Eq. 305 ($f_{\rm RSD}(\hat{k},\hat{\mu},z)$) and originally defined in Eq. 171 has to be rewritten as:

$$\beta(\hat{k},z) = -\frac{1+z}{2b(z)} \cdot \frac{d\ln P_{cb}(\hat{k},z)}{dz}\ , \tag{306}$$

where the bias is now rightfully assumed to be scale independent, being defined as $\delta_g = b(z) \times \delta_{cb}$. The same considerations apply to the 21cm power spectrum, indeed the neutral hydrogen in low redshift galaxies is a biased tracer of the *cb* field only. Therefore, P_m has to be replaced with P_{cb} in Eq. 227 and Eq. 231, and, as in the case of galaxy clustering, the Kaiser formula has to take into account the *cb* growth rate, rather than the total matter one.

We will now evaluate the sensitivity of Euclid and SKA (combined with Planck) to cosmological parameters by performing a Markov Chain Monte Carlo (MCMC) forecast, i.e. using Monte-Carlo Markov Chains to fit the spectra of the fiducial model assuming the likelihood expressions discussed in the previous sections. Our MCMC forecasts are obtained with the MontePython package[4] [174, 175], implementing our new Euclid and SKA likelihoods to fit the theoretical spectra provided by the Boltzmann solver CLASS [171] to the mock data. In Table 21 we report the expected sensitivity of various probe combinations. In Figures 27 and 28 we depict the corresponding $1\,\sigma$ uncertainty, in order to visualize the impact of different experiments, probe combinations and theoretical error prescriptions.

We replace theoretical error approach of the last section with that of section 3.4, where we elaborated on our conservative and realistic approach to the implementation of the theoretical error in Section 3.4.1 and 3.4.2. The sensitivity of galaxy clustering measurements is affected by the choice of the theoretical error more than cosmic shear and intensity mapping. As explained in Section 3.4.2, SKA1 is more sensitive than SKA2 and Euclid to the theoretical error prescription, because of its narrow redshift range. Indeed, Figure 27 shows that only for SKA1 the uncertainty to every cosmological parameter shrinks in the realistic configuration with respect to the conservative one. In the case of SKA2 and Euclid, the improvement of the sensitivity due to the more optimistic theoretical error prescription mostly concerns the primordial power spectrum parameter n_s: In the absence of a sharp cut-off in the measured power spectrum, the extended lever arm in k provides more constraining power. Finally, notice that the sensitivity of SKA1 band 2 intensity mapping to the derived parameters Ω_m and σ_8

4 https://github.com/brinckmann/montepython_public

	CS	GC	$\sigma(100\times\omega_b)$	$\sigma(\omega_{cdm})$	$\sigma(\ln[10^{10}A_s])$	$\sigma(n_s)$	$\sigma(H_0)/[\frac{km}{s\,Mpc}]$	$\sigma(\tau_{reio})$	$\sigma(M_\nu)/[meV]$	$\sigma(\sigma_8)$
Planck	-	-	0.015	0.00133	0.0088	0.00351	1.220	0.0045	0.084	0.01810
Planck+SKA1	c	-	0.013	0.00071	0.0085	0.00286	0.892	0.0045	0.071	0.01480
	r		0.012	0.00061	0.0084	0.00275	0.734	0.0045	0.059	0.01160
	-	c	0.012	0.00072	0.0084	0.00279	0.460	0.0044	0.045	0.00986
	c		0.012	0.00046	0.0081	0.00257	0.439	0.0043	0.042	0.00802
	r		0.012	0.00044	0.0085	0.00260	0.439	0.0045	0.040	0.00733
	-	r	0.011	0.00041	0.0085	0.00184	0.215	0.0045	0.027	0.00474
	c		0.011	0.00036	0.0084	0.00183	0.214	0.0044	0.026	0.00432
	r		0.011	0.00034	0.0083	0.00180	0.217	0.0044	0.025	0.00390
Planck+SKA2	c	-	0.011	0.00036	0.0083	0.00240	0.339	0.0044	0.033	0.00544
	r		0.011	0.00033	0.0083	0.00199	0.259	0.0044	0.026	0.00385
	-	c	0.010	0.00030	0.0076	0.00152	0.083	0.0042	0.016	0.00199
	c		0.010	0.00029	0.0076	0.00151	0.080	0.0042	0.015	0.00112
	r		0.010	0.00029	0.0076	0.00142	0.081	0.0041	0.015	0.00092
	-	r	0.010	0.00026	0.0065	0.00058	0.071	0.0034	0.013	0.00110
	c		0.010	0.00025	0.0065	0.00055	0.070	0.0034	0.013	0.00085
	r		0.010	0.00025	0.0062	0.00055	0.072	0.0033	0.012	0.00064
Planck+Euclid	c	-	0.012	0.00038	0.0084	0.00253	0.446	0.0044	0.043	0.00810
	r		0.011	0.00034	0.0082	0.00233	0.305	0.0043	0.030	0.00507
	-	c	0.010	0.00032	0.0080	0.00168	0.115	0.0043	0.019	0.00221
	c		0.010	0.00030	0.0079	0.00168	0.109	0.0044	0.018	0.00214
	r		0.010	0.00030	0.0079	0.00165	0.105	0.0043	0.017	0.00157
	-	r	0.010	0.00027	0.0064	0.00060	0.077	0.0034	0.014	0.00133
	c		0.010	0.00026	0.0065	0.00056	0.078	0.0034	0.014	0.00121
	r		0.010	0.00024	0.0061	0.00051	0.077	0.0033	0.012	0.00096
Planck+Euclid +SKA1	c	c	0.010	0.00029	0.0077	0.00159	0.088	0.0042	0.017	0.00189
	IM2	r	0.010	0.00025	0.0062	0.00052	0.068	0.0033	0.013	0.00098
Planck+SKA1	IM1	c	0.011	0.00030	0.0081	0.00224	0.099	0.0044	0.018	0.00213
		r	0.011	0.00030	0.0079	0.00219	0.095	0.0043	0.017	0.00194
	IM2	c	0.011	0.00033	0.0084	0.00243	0.172	0.0045	0.027	0.00513
		r	0.011	0.00031	0.0086	0.00234	0.109	0.0046	0.019	0.00261
	IM1 +IM2	c	0.011	0.00030	0.0081	0.00216	0.090	0.0044	0.017	0.00194
		r	0.011	0.00029	0.0079	0.00205	0.075	0.0043	0.016	0.00116

Table 21: Expected 1σ sensitivity of Planck, Euclid and SKA to the cosmological parameters. For each probe combination in the first and in the second column we indicate whether the cosmological probe is present or not, and whether the theoretical error is described with a conservative (c) or a realistic (r) approach.

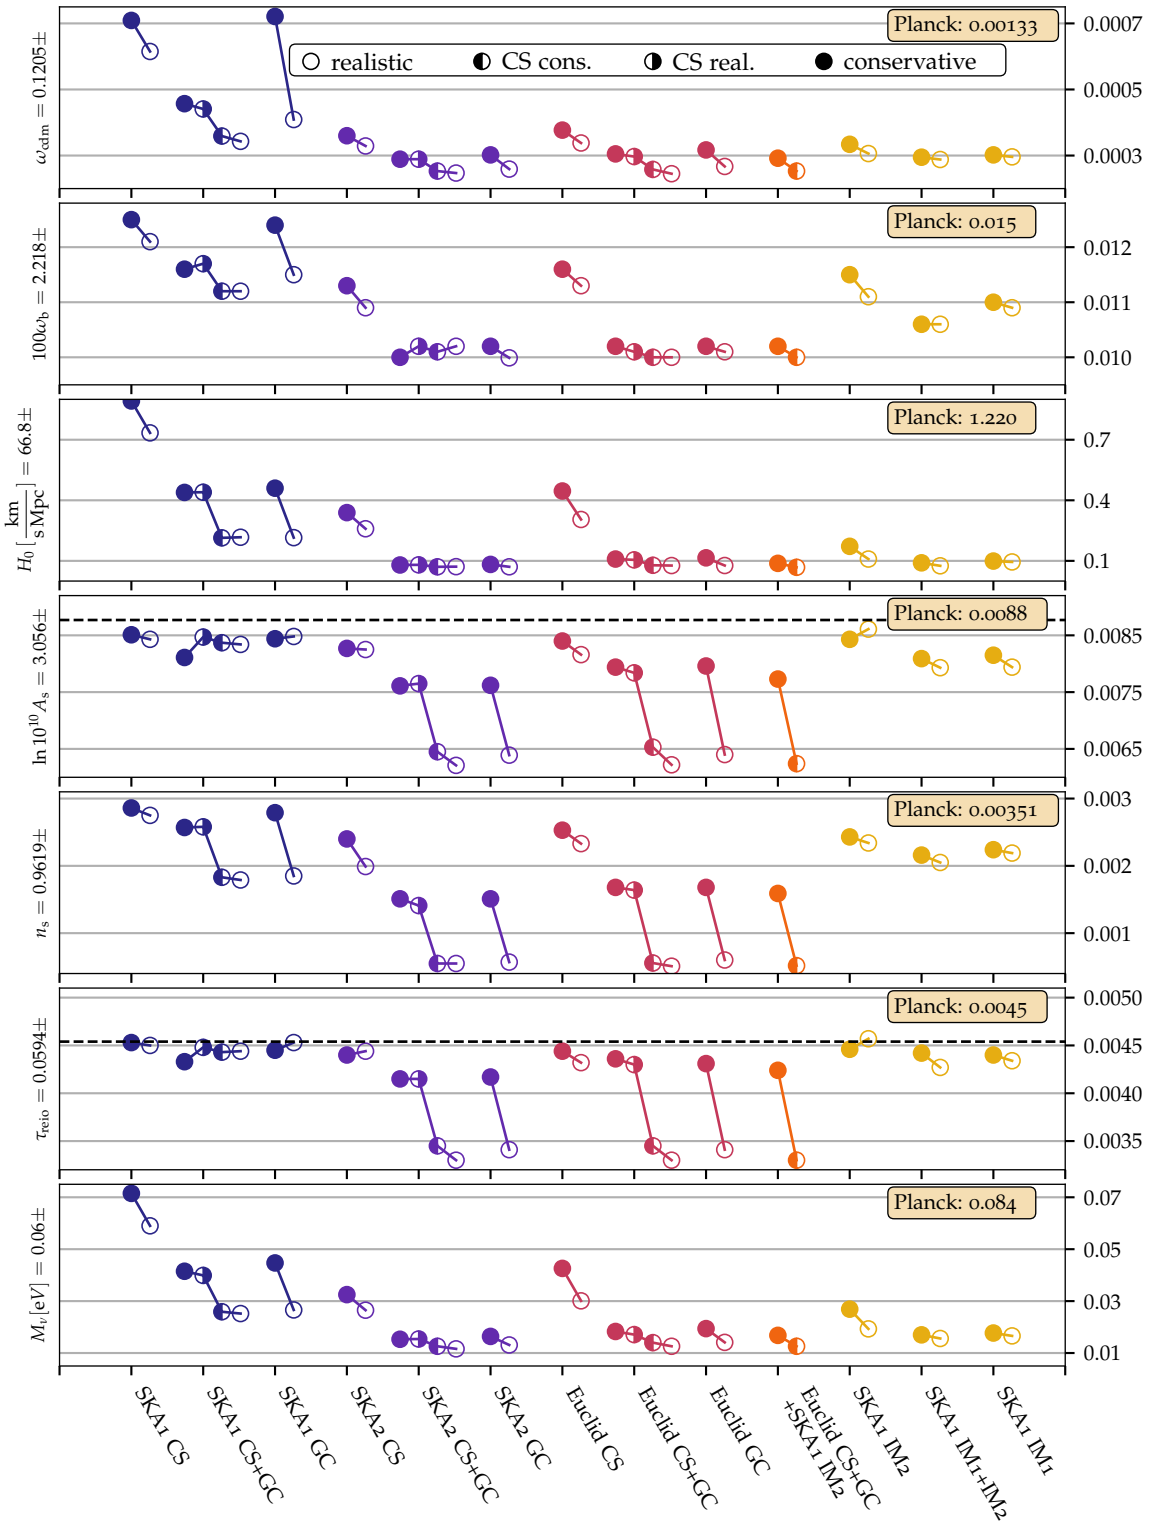

Figure 27: $1\,\sigma$ uncertainty for various combinations of experiments, all including Planck, to the baseline cosmological parameters. The combinations of experiments are specified on the x axis at the bottom of the plot. Empty circles denote realistic settings, while filled circles denote conservative settings. In the case of combinations of cosmic shear with other likelihoods, the left part of the circle describes the setting for the cosmic shear likelihood. The fiducial value is given on the left axis, while the scale of the $1\,\sigma$ uncertainty is written on the right axis. The Planck sensitivity is written inside the box in the upper right corner, and is also shown as a dashed line when within the range of the y axis on the right hand side.

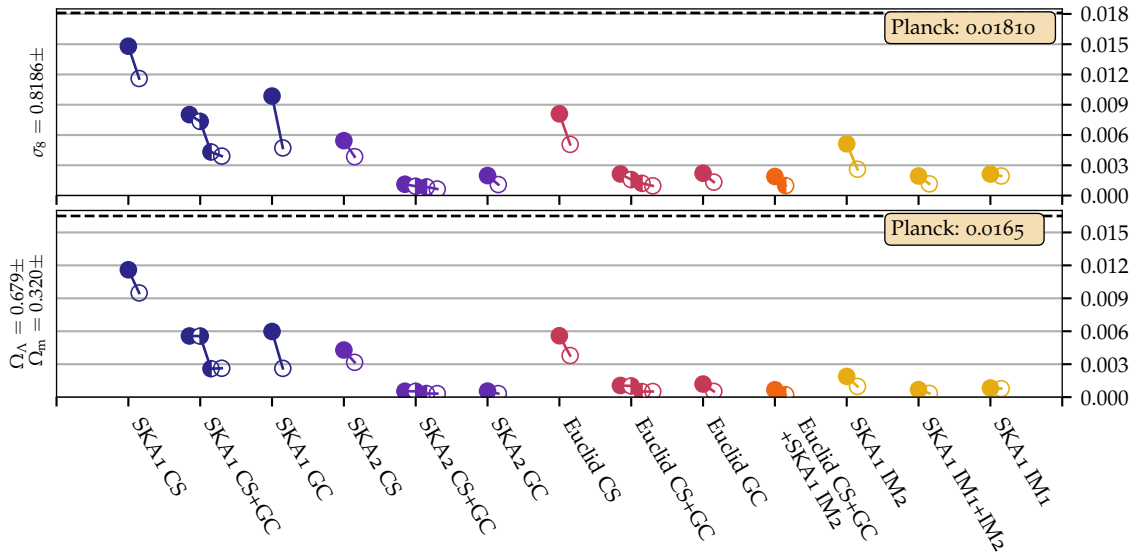

Figure 28: Same as Figure 27 but for the derived parameters Ω_Λ, Ω_m and σ_8.

greatly benefits from the realistic theoretical error configuration (Figure 28). The improvement of the constraints on these parameters with respect to the conservative configuration is caused by low redshift probes breaking the degeneracy intrinsic in CMB data. This mechanism does not work for SKA1 IM2 if a conservative approach to the theoretical error is used, because of the inadequate sensitivity. However, the different correlation of SKA1 IM2 in this part of the parameter space appears if we apply a more optimistic approach to the theoretical error is assumed.

Table 21 and Figure 27 show that both Euclid and SKA will greatly improve the sensitivity to the summed neutrino mass and to the other cosmological parameters with respect to currently available data sets. For a given error prescription, the combination Planck + Euclid (CS + GC) nearly achieves the best possible precision within our baseline ΛCDM+M_ν model, and adding SKA1 (IM2) and SKA2 (CS + GC) seems to lead to a negligible improvement. However, we will see in the next section that SKA1 (IM2) and SKA2 (CS + GC) should improve the sensitivity to the parameters of several extended cosmological models.

Concerning the neutrino mass sum, with a realistic description of the theoretical error, the sensitivity of Planck + Euclid (CS+GC) (Planck + SKA2 (CS + GC)), to a fiducial neutrino mass of 60 meV is 13 meV (12 meV). This fantastic accuracy implies a $5\,\sigma$ detection of a non-zero neutrino mass in the scenario of the minimal mass sum. Using a conservative prescription of the theoretical error degrades the sensitivity to 18 meV (15 meV), which still provides a more than $3\,\sigma$ detection and a factor 5 improvement with respect to Planck-only. Our results on the sensitivity of Euclid and SKA to the neutrino mass sum are consistent with the previous literature [118, 183, 198, 199]. Note that Ref. [199] finds exactly the same uncertainty ($\sigma(M_\nu) = 0.012$ eV) as in our most constraining data combinations (Planck + SKA2 (CS + GC)) by using Planck-polarization + CMB-Stage-IV + BAO-DESI + 21cm-HERA.

Figures 29 and 30 show the marginalized $1\,\sigma$ and $2\,\sigma$ contours and one dimensional posteriors in the $(\omega_b, \omega_{\rm cdm}, A_s, n_s, M_\nu, H_0, \tau_{\rm reio}, \sigma_8, \Omega_m)$ parameter space for the most constraining data combinations with respect to Planck-only. Figure 29 assumes the conservative theoretical error setup, and Figure 30 the realistic one.

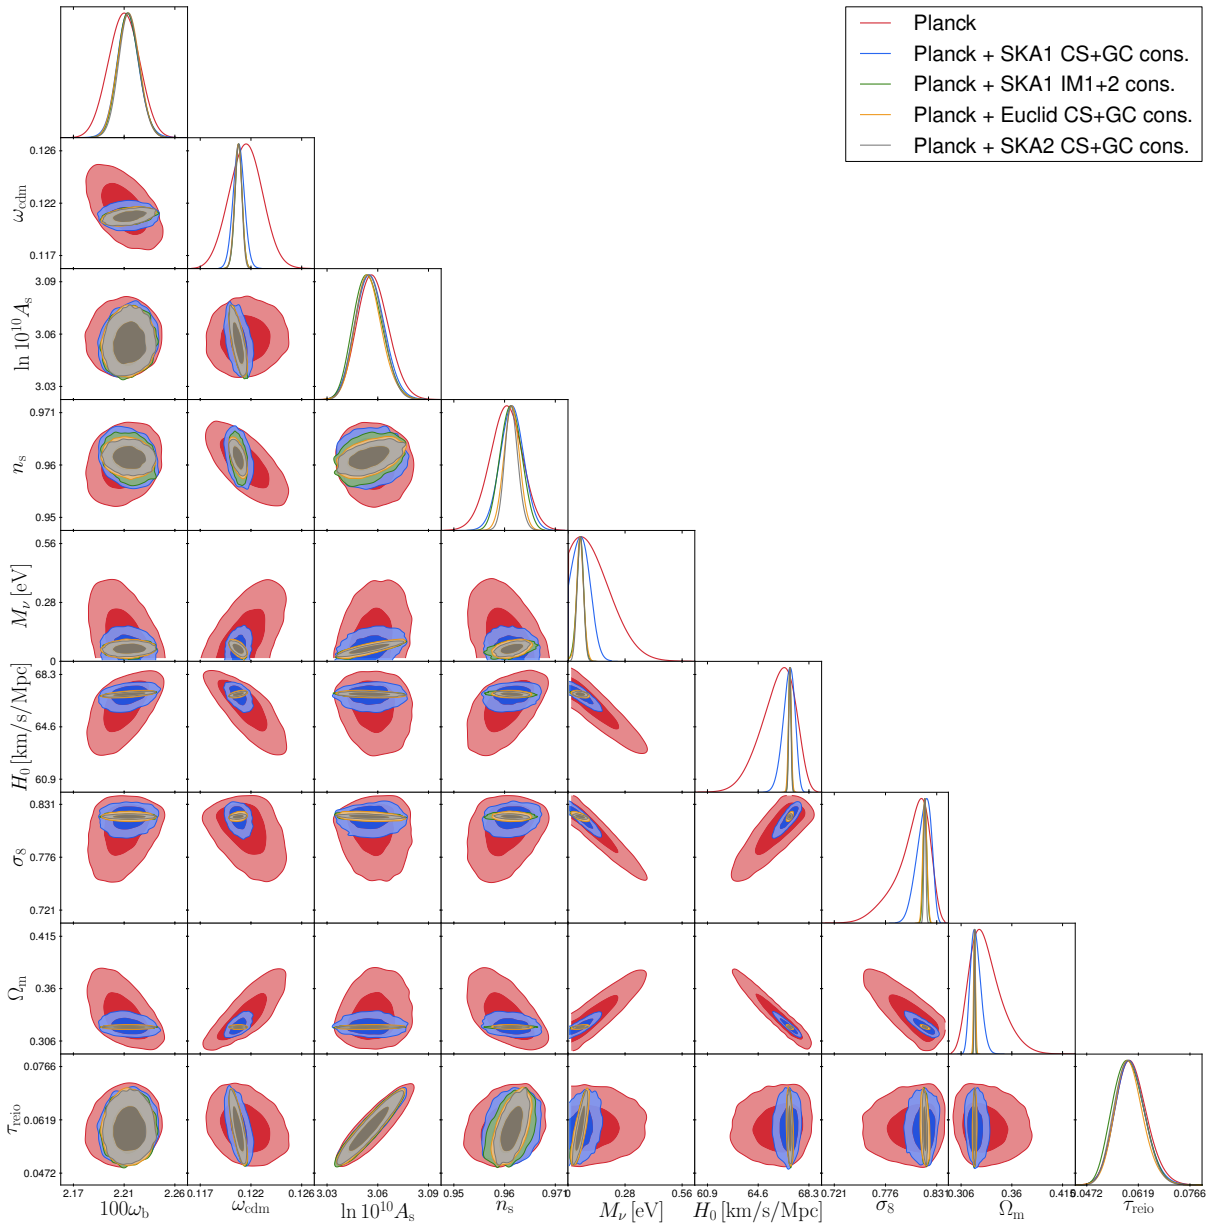

Figure 29: Marginalized 1σ and 2σ contours and one-dimensional posteriors in the $(\omega_b, \omega_{cdm}, A_s, n_s, M_\nu, H_0, \tau_{reio}, \sigma_8, \Omega_m)$ parameter space, showing the expected sensitivity of Planck-only, Planck + Euclid (CS + GC), Planck + SKA1 (CS + GC), Planck+SKA1-IM and Planck + SKA2 (CS + GC). Here the analysis is performed following the conservative approach for the description of the theoretical error.

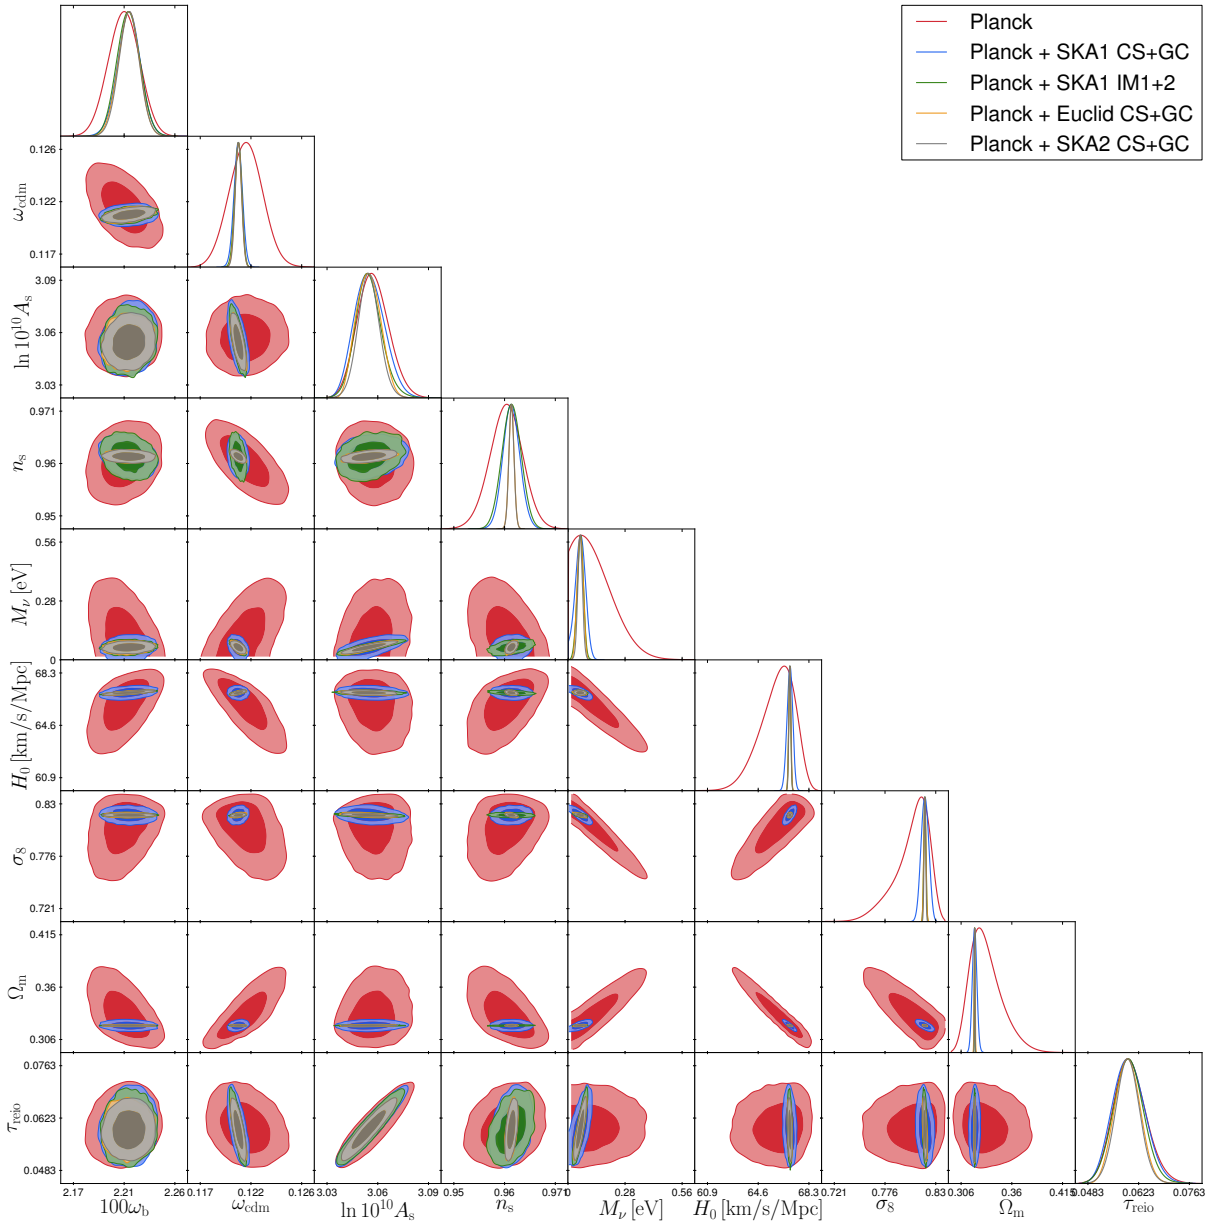

Figure 30: Marginalized 1σ and 2σ contours and one-dimensional posteriors in the $(\omega_b, \omega_{cdm}, A_s, n_s, M_\nu, H_0, \tau_{reio}, \sigma_8, \Omega_m)$ parameter space, showing the expected sensitivity of Planck-only, Planck + Euclid (CS + GC), Planck + SKA1 (CS + GC), Planck + SKA1 IM(1 + 2) and Planck + SKA2 (CS + GC). Here the analysis is performed following the realistic modeling of the theoretical error for GC and IM and the conservative approach for CS. This is done to isolate the effect of a theoretical error approach without a cut-off from contributions of CS.

These figures show that combining high redshift CMB data with low redshift measurements breaks the degenerecies among the cosmological parameters that are present when using only CMB data (red lines and contours). For example, if we look at the contours in the plane $(M_\nu, \omega_{\rm cdm})$ we can see that both Euclid and SKA (SKA1 intensity mapping, or SKA2 (CS + GC) lift the degeneracy and even slightly reverse it (a physical interpretation is given in chapter 5 or Ref. [157]). This propagates to other parameters, causing an overall shrinking of all contours and leading to an increase of the sensitivity. From the triangle plots, we can verify once more that SKA1 (CS + GC) results (blue lines and contours) benefit from a realistic modelling of the theoretical error more than the other experiments. If we look at the SKA1 (CS + GC) one dimensional posterior on H_0, we can see that it is considerably narrower in the realistic configuration. Given the correlation between H_0 and the other cosmological parameters, this effect leads to a remarkable increase of the sensitivity not only to H_0, but also to the other cosmological parameters, e.g. the neutrino mass sum. Indeed, only the realistic approach would allow for a significant detection of the minimal neutrino mass with SKA1 (CS + GC). Interestingly, the combination and Planck and SKA1 intensity mapping (green lines and contours) has a sensitivity which is roughly comparable to that of Planck + Euclid (CS + GC) and Planck + SKA2 (CS + GC) for most parameters, with the notable exception of the primordial power spectrum parameters.

Finally, our analysis shows that, regardless of the theoretical error description, both Euclid and SKA (SKA1 intensity mapping, or SKA2 (CS + GC)) will provide a detection of a non-zero neutrino mass in the next few years. This result appears to be very robust at least as long as we assume a minimal 7-parameter cosmological model. One of the purposes of the next chapter is to assess to which extent this remains true in the presence of additional parameters.

9

SENSITIVITY TO NEUTRINOS IN EXTENDED MODELS

Sections 9.1 to 9.3 are based on Di Valentino, Brinckmann, Gerbino, Poulin et al. [50], section 9.4 is based on Sprenger, Archidiacono, Brinckmann, Clesse & Lesgourgues [102], and section 9.5 is based on Brinckmann, Hooper, Archiadiacono, Lesgourgues & Sprenger (in prep.)

9.1 DEGENERACY BETWEEN NEUTRINO MASS AND OTHER PARAMETERS

In chapter 8, section 8.1, we found a sensitivity of about $\sigma(M_\nu) = 44$ meV for CORE-M5, using any configuration, or 21 meV in combination with future BAOs, and 16 meV with future cosmic shear data. We explained why the sensitivity to M_ν has a very weak dependence on the assumed instrumental settings for CORE. To check how much these predictions depend on the assumed cosmological model, we do several extended forecasts with 8 free parameters instead of 7.

The new parameters studied in this section are the primordial helium fraction, the tensor-to-scalar ratio, the constant Dark Energy equation of state parameter, the primordial scalar tilt running, and the effective density fraction of spatial curvature. Since our focus here is on neutrino masses, we do not investigate the sensitivity to these parameters in as much detail. For instance, we use here a (weak energy principle) prior $w > -1$, while in the Dark Energy section we will also consider phantom Dark Energy or a time-varying w. Also, as in the rest of this paper, we stick to a mock CORE likelihood including only temperature, E-polarisation and lensing data, and not using B-mode information: hence we obtain much worse constraints on r than in the CORE inflation paper [247], in which B modes play an essential role; but at least the present forecast allows to conservatively prove the absence of parameter correlation between M_ν and r at the level of precision of CORE combined with DESI and Euclid.

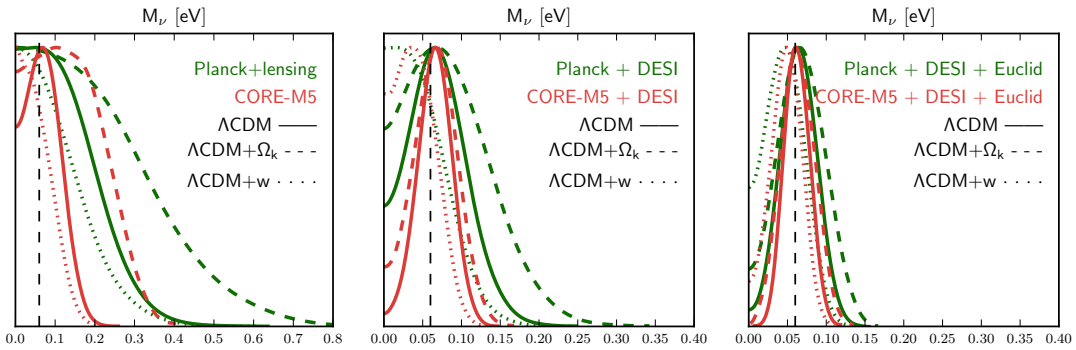

Figure 31: Posterior distribution of the summed neutrino mass in the extended models ΛCDM + M_ν, ΛCDM + M_ν + w and ΛCDM + M_ν + Ω_k. The vertical dashed line shows the fiducial value.

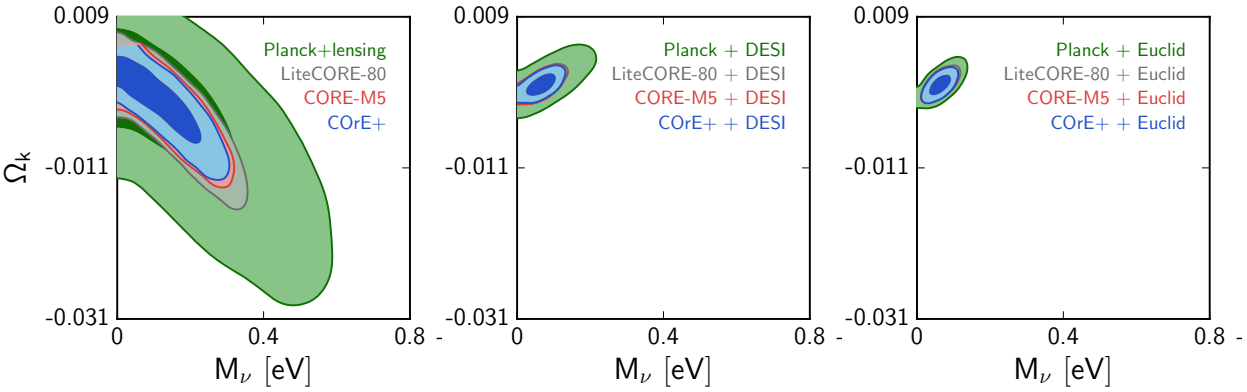

Figure 32: Results for the extended model ΛCDM + M_ν + Ω_k. The (M_ν, Ω_k) degeneracy is removed by adding BAO data.

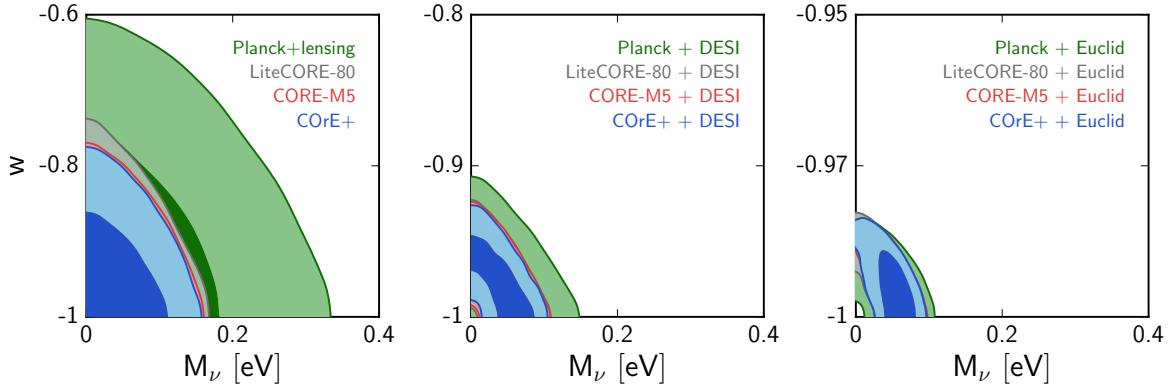

Figure 33: Results for the extended model ΛCDM + M_ν + w (with a prior $w > -1$). The w axis scale changes between plots because of the huge difference of sensitivity between data sets. The (M_ν, w) degeneracy gets partially resolved by adding Euclid cosmic shear data.

Parameter	Planck, TEP	LiteCORE-80, TEP	LiteCORE-120, TEP	CORE-M5, TEP	COrE+, TEP
$Y_{\rm He}$	0.247 ± 0.014	0.2475 ± 0.0042	0.2476 ± 0.0030	0.2476 ± 0.0028	0.2476 ± 0.0025
M_ν (meV)	fixed	77^{+36}_{-60}	74^{+37}_{-53}	74^{+37}_{-52}	73^{+37}_{-51}
r	< 0.055 (68%CL)	< 0.020 (68%CL)	< 0.019 (68%CL)	< 0.020 (68%CL)	< 0.020 (68%CL)
M_ν (meV)	fixed	< 95 (68%CL)	< 97 (68%CL)	68^{+33}_{-52}	66^{+31}_{-51}
$w\ (> -1)$	< -0.74 (68%CL)	< -0.886 (68%CL)	< -0.895 (68%CL)	< -0.900 (68%CL)	< -0.904 (68%CL)
M_ν (meV)	fixed	< 73 (68%CL)	< 72 (68%CL)	< 72 (68%CL)	< 71 (68%CL)
$10^2\, d\ln n_s/d\ln k$	-0.20 ± 0.67	-0.01 ± 0.30	0.00 ± 0.25	-0.00 ± 0.25	0.00 ± 0.23
M_ν (meV)	fixed	< 110 (68%CL)	74^{+37}_{-53}	74^{+37}_{-52}	72^{+37}_{-50}
$10^2\,\Omega_k$	$-0.37^{+0.83}_{-0.69}$	$-0.37^{+0.60}_{-0.33}$	$-0.33^{+0.56}_{-0.31}$	$-0.30^{+0.51}_{-0.29}$	$-0.28^{+0.48}_{-0.30}$
M_ν (meV)	fixed	< 200 (68%CL)	< 197 (68%CL)	< 193 (68%CL)	< 188 (68%CL)

Parameter	Planck, TEP + DESI	LiteCORE-80, TEP + DESI	LiteCORE-120, TEP + DESI	CORE-M5, TEP + DESI	COrE+, TEP + DESI
$Y_{\rm He}$	0.248 ± 0.010	0.2477 ± 0.0041	0.2477 ± 0.0030	0.2477 ± 0.0027	0.2477 ± 0.0024
M_ν (meV)	70^{+33}_{-51}	65^{+25}_{-21}	65^{+23}_{-20}	65 ± 21	66^{+23}_{-19}
r	< 0.062 (68%CL)	< 0.020 (68%CL)	< 0.019 (68%CL)	< 0.020 (68%CL)	< 0.020 (68%CL)
M_ν (meV)	72^{+37}_{-47}	62^{+26}_{-21}	62^{+23}_{-21}	63 ± 22	63^{+23}_{-19}
$w\ (> -1)$	< -0.961 (68%CL)	< -0.965 (68%CL)	< -0.9653 (68%CL)	< -0.965 (68%CL)	< -0.965 (68%)
M_ν (meV)	< 66 (68%CL)	< 56 (68%CL)	< 57 (68%CL)	< 60 (68%CL)	< 57 (68%CL)
$10^2\, d\ln n_s/d\ln k$	-0.01 ± 0.58	0.01 ± 0.31	0.01 ± 0.25	0.01 ± 0.24	-0.00 ± 0.23
M_ν (meV)	71^{+35}_{-48}	66^{+25}_{-22}	66^{+24}_{-21}	66^{+24}_{-20}	61^{+22}_{-19}
$10^2\,\Omega_k$	0.07 ± 0.19	0.01 ± 0.11	0.00 ± 0.10	-0.00 ± 0.10	0.00 ± 0.10
M_ν (meV)	< 122 (68%CL)	66 ± 31	65^{+31}_{-27}	65^{+29}_{-26}	66 ± 28

Parameter	Planck, TEP + DESI + Euclid	LiteCORE-80, TEP + DESI + Euclid	LiteCORE-120, TEP + DESI + Euclid	CORE-M5, TEP + DESI + Euclid	COrE+, TEP + DESI + Euclid
$Y_{\rm He}$	0.2470 ± 0.0091	0.2477 ± 0.0040	0.2477 ± 0.0029	0.2477 ± 0.0027	0.2477 ± 0.0024
M_ν (meV)	63 ± 26	63 ± 17	63 ± 16	62 ± 16	62^{+15}_{-17}
r	< 0.063 (68%CL)	< 0.020 (68%CL)	< 0.020 (68%CL)	< 0.020 (68%CL)	< 0.020 (68%CL)
M_ν (meV)	62 ± 23	60^{+16}_{-18}	60^{+15}_{-17}	60^{+15}_{-17}	60^{+15}_{-17}
$w\ (> -1)$	< -0.9914 (68%CL)	< -0.9920 (68%CL)	< -0.9921 (68%CL)	< -0.9923 (68%)	< -0.9926 (68%)
M_ν (meV)	44^{+23}_{-29}	45^{+23}_{-19}	46^{+22}_{-18}	48^{+22}_{-17}	48^{+21}_{-17}
$10^2\, d\ln n_s/d\ln k$	0.01 ± 0.52	0.01 ± 0.29	0.01 ± 0.24	0.00 ± 0.23	0.01 ± 0.22
M_ν (meV)	63 ± 25	63 ± 17	63 ± 16	63 ± 16	63 ± 16
$10^2\,\Omega_k$	0.03 ± 0.14	0.012 ± 0.100	-0.002 ± 0.094	0.004 ± 0.098	-0.003 ± 0.096
M_ν (meV)	70^{+34}_{-30}	65 ± 22	62 ± 22	63 ± 21	62 ± 21

Table 22: 68% CL constraints on the additional parameters of several extended 8-parameter models, for the different CORE experimental specifications, and with or without external data sets (DESI BAOs, Euclid cosmic shear). For Planck alone, we quote the results from the 2015 data release, *obtained with a fixed mass* $M_\nu = 60$ meV, while for combinations of Planck with future surveys, we fit mock data with a fake Planck likelihood mimicking the sensitivity of the real experiment (although a bit more constraining). In the case with free tensor-to-scalar ratio r, we did not include B-modes in the likelihood, unlike in the CORE inflation paper [247]. In the case with free w we used a (weak energy principle) prior $w > -1$, that will be relaxed in the Dark Energy section of this paper.

Our extended forecast results are summarised in Table 22. When varying the helium fraction, the tensor-to-scalar ratio[1], or the tilt running, we find essentially the same sensitivity to

1 With a free r and in the forecasts based on CMB data alone, the error bar $\sigma(M_\nu)$ can be slightly smaller in the extended model than in the 7-parameter model, which may sound odd. In fact, it comes from a volume effect in Bayesian parameter extraction: models with $r \neq 0$ are less discrepant with the data when M_ν is small, so after marginalising over r the posterior for M_ν is shifted to lower values.

M_ν as in the 7-parameter model. Nonetheless, the cases with free w or Ω_k make the neutrino mass detection more difficult, due to clear parameter degeneracies with M_ν when using CMB data alone (see Figures 31, 32, 33).

We see in Figures 31, 32 that the (M_ν, Ω_k) degeneracy (a particular case of the *geometrical degeneracy* described in [201, 306]) gets broken by the inclusion of BAO data, bringing the error down to $\sigma(M_\nu) \simeq 28$ meV. With additional Euclid cosmic shear data, one would reach $\sigma(M_\nu) \simeq 21$ meV, still guaranteeing a 3σ detection, while Planck+DESI+Euclid could only achieve $\sigma(M_\nu) \simeq 32$ meV for free Ω_k.

In the case with free w (Figures 31, 33), the degeneracy remains problematic even with CORE+BAO data, but ultimately Euclid cosmic shear data could partly differentiate between the physical effects of w and M_ν effects and lead to $\sigma(M_\nu) \simeq 19$ meV under the prior $w > -1$, instead of 26 meV for Planck+DESI+BAO. The error bar would degrade by also allowing for phantom dark energy, but on the other hand, the inclusion of further Large Scale Structure data (e.g. the Euclid galaxy correlation function) would further help to break the degeneracy, since the effect of neutrino masses and w have a different dependence on redshift and scales [183].

9.2 EXTRA RELATIVISTIC RELICS

The minimal cosmological scenario predicts that, at least after the time of nucleosynthesis, the density of relativistic particles is given by the contribution of CMB photons plus that of active neutrino species, until they become non-relativistic due to their small mass. This assumption is summarized by the standard value of the effective neutrino number $N_{\rm eff} = 3.046$ [307] (see Refs. [308, 309] for pioneering work and Ref. [310] for a review of the subject). A more recent calculation beased on the latest data on neutrino physics finds $N_{\rm eff} = 3.045$ [311], but at the precision level of CORE the difference is irrelevant, and we will keep 3.046 as our baseline assumption. However, there are many simple theoretical motivations for relaxing this assumption. We know that the standard model of particle physics is incomplete (e.g. because it does not explain dark matter), and many of its extensions would lead to the existence of extra light or massless particles; depending on their interactions and decoupling time the latter could also contribute to $N_{\rm eff}$. Depending on the context, these extra particles are usually called extra relativistic relics, dark radiation or axion-like particles in more specific cases. In the particular case of particles that were in thermal equilibrium at some point, the enhancement of $N_{\rm eff}$ can be predicted as a function of the decoupling temperature [82]. Even in absence of a significant density of such relics, ordinary neutrinos could have an unexpected density due to non-standard interactions [83], non-thermal production after decoupling [312], or low-temperature reheating [84], leading to a value of $N_{\rm eff}$ larger or smaller than 3.046. There are additional motivations to consider $N_{\rm eff}$ as a free parameter (background of gravitational waves produced by a phase transition, modified gravity, extra dimensions, etc. – see Ref. [313] for a review).

Over the last years the extended ΛCDM + $N_{\rm eff}$ has received a lot of attention within the cosmology community. Assuming $N_{\rm eff} > 3.046$ has the potential to solve tensions in observational data: for instance, internal tensions in pre-Planck CMB data, which have now disappeared ($N_{\rm eff} = 2.99 \pm 0.20$ (68%CL) for Planck 2015 TT,TE,EE+lowP [2]); or tensions between CMB data and direct measurements of H_0 [85] (however, solving this problem by increasing $N_{\rm eff}$ requires a higher value of σ_8, which brings further tensions with other datasets [2]). In any case, the community is particularly eager to measure $N_{\rm eff}$ with better sensitivity in the future,

in order to: (i) test the existence of extra relics and probe extensions of the standard model of particle physics; (ii) get a window on precision neutrino physics (since the contribution of neutrinos to $N_{\rm eff}$ depends on the details of neutrino decoupling); and (iii) check whether the tensions in cosmological data are related to the relativistic density or not.

Since CMB data accurately determines the redshift of equality $z_{\rm eq}$, the impact of $N_{\rm eff}$ on CMB observables is usually discussed at fixed $z_{\rm eq}$ [66, 189, 314]. The time of equality can be kept fixed by simultaneously increasing $N_{\rm eff}$ and the dark matter density $\omega_{\rm cdm}$ (or, depending on the choice of parameter basis, $N_{\rm eff}$ and H_0). The impact on the CMB is then minimal, which explains the well known ($N_{\rm eff}$, $\omega_{\rm cdm}$) or ($N_{\rm eff}$, H_0) degeneracy: the latter is clearly visible with Planck data in Figure 34 (left plot). However, this transformation does not preserve the angular scale of the photon damping scale on the last scattering surface: hence the best probe of $N_{\rm eff}$ comes from accurate measurements of the exponential tail of the temperature and polarisation spectra at high-ℓ. Hence the accuracy with which CMB experiments can measure $N_{\rm eff}$ is directly related to their sensitivity and angular resolution, as confirmed by the following forecasts. Increasing $N_{\rm eff}$ has other effects on the CMB coming from gravitational interactions between photons and neutrinos before decoupling: a smoothing of the acoustic peaks (however, very small, and below the per-cent level for variations of the order of $\Delta N_{\rm eff} \sim 0.1$), and a shift of the peaks towards larger angles caused by the "neutrino drag" effect [66, 189, 314]. This means that in order to keep a fixed CMB peak scale, one should decrease the angular size of the sound horizon θ_s while increasing $N_{\rm eff}$: this implies an anti-correlation between θ_s and $N_{\rm eff}$ that can be observed in Figure 34 (right plot). Therefore, by accurately measuring $N_{\rm eff}$, we could get a more robust and model-independent measurement of the sound horizon scale, which would in turn be very useful for constraining the expansion history with BAO data.

Since the parameter $N_{\rm eff}$ is closely related to neutrino properties, and since we know that neutrinos have a small mass, we forecast the sensitivity of different experimental set-ups to $N_{\rm eff}$ while varying simultaneously the summed neutrino mass M_ν. This leads to more robust predictions than if we had fixed the mass (although a posteriori we find no significant correlation between $N_{\rm eff}$ and M_ν). We investigate the CORE sensitivity to $N_{\rm eff}$ within two distinct models:

- The model "ΛCDM + M_ν +$\Delta N_{\rm eff}^{\rm massless}$" has 3 massive degenerate and thermalised neutrino species, plus extra massless relics contributing as $\Delta N_{\rm eff}^{\rm massless} > 0$. It is motivated by scenarios with standard active neutrinos and extra massless relics (or very light relics with $m \ll 10$ meV).

- The model "ΛCDM + M_ν + $N_{\rm eff}^{\rm massive}$" only has 3 massive degenerate neutrino species, with fixed temperature, but with a rescaled density. During radiation domination they contribute to the effective neutrino number as $N_{\rm eff}^{\rm massive}$, which could be greater or smaller than 3.046. This model provides a rough first-order approximation to specific scenarios in which neutrinos would be either enhanced (e.g. by the decay of other particles) or suppressed (e.g. in case of low-temperature reheating).

Choosing the same fiducial model as in footnote 1, with a summed mass equal to $M_\nu = 60$ meV, we fit the 7-parameter ΛCDM+M_ν model for different CORE settings, alone or in combination with mock DESI BAOs and Euclid cosmic shear data.

The results of our MCMC forecasts are shown in Tables 23, 24, and Figure 34. Since the determination of $N_{\rm eff}$ depends mainly on observations of the exponential tail in the CMB spectra, our results for $\sigma(N_{\rm eff})$ vary a lot with the sensitivity/resolution assumed for CORE,

Parameter	Planck, TEP	LiteCORE-80, TEP	LiteCORE-120, TEP	CORE-M5, TEP	COrE+, TEP
$\Delta N_{\rm eff}^{\rm massless}$	< 0.19 (68%CL)	< 0.062 (68%CL)	< 0.045 (68%CL)	< 0.040 (68%CL)	< 0.036 (68%CL)
M_ν (meV)	< 310 (68%CL)	77^{+37}_{-59}	72^{+34}_{-56}	71^{+34}_{-54}	70^{+35}_{-53}
$\Omega_b h^2$	0.02208 ± 0.00025	0.022305 ± 0.000070	0.022293 ± 0.000052	0.022289 ± 0.000047	0.022284 ± 0.000041
$\Omega_c h^2$	0.1184 ± 0.0030	$0.12056^{+0.00066}_{-0.00096}$	$0.12030^{+0.00057}_{-0.00079}$	$0.12023^{+0.00052}_{-0.00074}$	$0.12015^{+0.00051}_{-0.00071}$
$100\theta_s$	1.04087 ± 0.00046	1.04070 ± 0.00013	1.04070 ± 0.00010	1.040700 ± 0.000094	1.040800 ± 0.000085
τ	0.071 ± 0.018	0.0605 ± 0.0020	0.0606 ± 0.0021	0.0605 ± 0.0021	0.0606 ± 0.0021
n_s	0.9589 ± 0.0095	0.9665 ± 0.0026	0.9663 ± 0.0023	0.9661 ± 0.0023	0.9661 ± 0.0022
$ln(10^{10} A_s)$	3.071 ± 0.037	3.0970 ± 0.0044	3.0964 ± 0.0043	3.0961 ± 0.0042	3.0960 ± 0.0042
H_0 (km/s/Mpc)	64.8 ± 2.3	$67.15^{+0.80}_{-0.58}$	$67.13^{+0.74}_{-0.51}$	$67.12^{+0.71}_{-0.50}$	$67.11^{+0.71}_{-0.47}$
σ_8	$0.778^{+0.038}_{-0.024}$	$0.831^{+0.011}_{-0.006}$	$0.831^{+0.010}_{-0.006}$	$0.8308^{+0.0097}_{-0.0059}$	$0.8307^{+0.0097}_{-0.0055}$
Parameter	Planck, TEP + DESI	LiteCORE-80, TEP + DESI	LiteCORE-120, TEP + DESI	CORE-M5, TEP + DESI	COrE+, TEP + DESI
$\Delta N_{\rm eff}^{\rm massless}$	< 0.15 (68%CL)	< 0.061 (68%CL)	< 0.042 (68%CL)	< 0.038 (68%CL)	< 0.033 (68%CL)
M_ν (meV)	85^{+41}_{-50}	72 ± 24	71^{+23}_{-20}	70^{+23}_{-20}	65^{+22}_{-20}
$\Omega_b h^2$	0.02237 ± 0.00015	0.022310 ± 0.000065	0.022293 ± 0.000050	0.022289 ± 0.000045	0.022279 ± 0.000038
$\Omega_c h^2$	$0.1216^{+0.0012}_{-0.0020}$	$0.12046^{+0.00048}_{-0.00081}$	$0.12023^{+0.00040}_{-0.00059}$	$0.12017^{+0.00036}_{-0.00054}$	$0.12045^{+0.00034}_{-0.00046}$
$100\theta_s$	1.04050 ± 0.00036	1.04070 ± 0.00013	1.04070 ± 0.00010	1.040700 ± 0.000091	1.040700 ± 0.000080
τ	0.0614 ± 0.0046	0.0605 ± 0.0021	0.0606 ± 0.0021	0.0605 ± 0.0021	0.0605 ± 0.0018
n_s	$0.9695^{+0.0037}_{-0.0053}$	$0.9667^{+0.0020}_{-0.0026}$	0.9662 ± 0.0020	0.9661 ± 0.0020	$0.9653^{+0.0016}_{-0.0020}$
$ln(10^{10} A_s)$	3.102 ± 0.010	3.0967 ± 0.0044	3.0962 ± 0.0043	3.0959 ± 0.0040	3.0966 ± 0.0036
H_0 (km/s/Mpc)	$67.56^{+0.42}_{-0.65}$	67.23 ± 0.33	67.15 ± 0.29	67.13 ± 0.29	67.13 ± 0.28
σ_8	0.833 ± 0.011	0.8316 ± 0.0044	0.8311 ± 0.0040	0.8309 ± 0.0038	0.8309 ± 0.0037
Parameter	Planck, TEP + DESI + Euclid	LiteCORE-80, TEP + DESI + Euclid	LiteCORE-120, TEP + DESI + Euclid	CORE-M5, TEP + DESI + Euclid	COrE+, TEP + DESI + Euclid
$\Delta N_{\rm eff}^{\rm massless}$	< 0.111 (68%CL)	< 0.054 (68%CL)	< 0.040 (68%CL)	< 0.038 (68%CL)	< 0.032 (68%CL)
M_ν (meV)	84^{+25}_{-28}	71^{+16}_{-18}	68^{+15}_{-18}	68^{+15}_{-17}	67^{+14}_{-17}
$\Omega_b h^2$	0.02234 ± 0.00013	0.022301 ± 0.000061	0.022290 ± 0.000048	0.022289 ± 0.000045	0.022282 ± 0.000038
$\Omega_c h^2$	$0.1211^{+0.0007}_{-0.0013}$	$0.12043^{+0.00034}_{-0.00065}$	$0.12026^{+0.00029}_{-0.00050}$	$0.12023^{+0.00028}_{-0.00046}$	$0.12017^{+0.00027}_{-0.00040}$
$100\theta_s$	1.04060 ± 0.00034	1.04070 ± 0.00012	1.040700 ± 0.000095	1.040700 ± 0.000089	1.040800 ± 0.000080
τ	0.0611 ± 0.0046	0.0605 ± 0.0021	0.0604 ± 0.0021	0.0605 ± 0.0021	0.0597 ± 0.0020
n_s	$0.9678^{+0.0031}_{-0.0040}$	0.9662 ± 0.0021	0.9660 ± 0.0019	0.9659 ± 0.0018	0.9658 ± 0.0017
$ln(10^{10} A_s)$	$3.100^{+0.008}_{-0.011}$	3.0967 ± 0.0043	3.0960 ± 0.0041	3.0961 ± 0.0041	3.0958 ± 0.0039
H_0 (km/s/Mpc)	$67.37^{+0.28}_{-0.42}$	67.18 ± 0.23	67.14 ± 0.20	67.12 ± 0.19	67.10 ± 0.19
σ_8	$0.8314^{+0.0037}_{-0.0030}$	$0.8319^{+0.0034}_{-0.0026}$	$0.8318^{+0.0032}_{-0.0024}$	$0.8317^{+0.0032}_{-0.0023}$	$0.8318^{+0.0030}_{-0.0022}$

Table 23: 68% CL constraints on cosmological parameters in the ΛCDM + M_ν + $\Delta N_{\rm eff}^{\rm massless}$ model (accounting for standard massive neutrino plus extra massless relics, with $\Delta N_{\rm eff}^{\rm massless} > 0$) from the different CORE experimental specifications and with or without external data sets (DESI BAOs, Euclid cosmic shear). For Planck alone, we quote the results from the 2015 data release, while for combinations of Planck with future surveys, we fit mock data with a fake Planck likelihood mimicking the sensitivity of the real experiment (although a bit more constraining).

and are only marginally affected by the inclusion of extra datasets like BAOs and cosmic shear surveys. The value $l_{\rm max}$ at which the signal-to-noise blows up in the temperature or polarisation spectrum varies a lot between the different experimental settings, as can be seen in Figure 1. Thus there is a dramatic improvement in $\sigma(N_{\rm eff})$ between Planck and LiteCORE-80 (factor 3), and still a substantial one between LiteCORE-80 and COrE+ (factor 1.7). However, stepping back to the design of CORE-M5, one maintains a very good sensitivity, $\sigma(N_{\rm eff}) = 0.041$, only 10% worse than what could be achieved with the better angular resolution of the

Parameter	Planck, TEP	LiteCORE-80, TEP	LiteCORE-120, TEP	CORE-M5, TEP	COrE+, TEP
$N_{\rm eff}^{\rm massive}$	2.93 ± 0.19	3.045 ± 0.063	3.047 ± 0.045	3.045 ± 0.041	3.045 ± 0.036
M_ν (meV)	< 310 (68%CL)	< 110 (68%CL)	73^{+37}_{-53}	73^{+37}_{-52}	72^{+37}_{-49}
$\Omega_b h^2$	0.02208 ± 0.00025	0.022250 ± 0.000089	0.022255 ± 0.000066	0.022254 ± 0.000060	0.022255 ± 0.000051
$\Omega_c h^2$	0.1184 ± 0.0030	0.1198 ± 0.0011	0.11983 ± 0.00082	0.11981 ± 0.00077	0.11979 ± 0.00071
$100\theta_s$	1.04087 ± 0.00046	1.04080 ± 0.00016	1.04080 ± 0.00012	1.04080 ± 0.00011	1.04080 ± 0.00010
τ	0.071 ± 0.018	0.0604 ± 0.0021	0.0604 ± 0.0021	0.0604 ± 0.0021	0.0603 ± 0.0021
n_s	0.9589 ± 0.0095	0.9642 ± 0.0036	0.9644 ± 0.0031	0.9643 ± 0.0030	0.9643 ± 0.0028
$ln(10^{10} A_s)$	3.071 ± 0.037	3.0950 ± 0.0048	3.0950 ± 0.0045	3.0950 ± 0.0045	3.0948 ± 0.0043
H_0 (km/s/Mpc)	64.8 ± 2.3	$66.81^{+0.88}_{-0.71}$	$66.86^{+0.78}_{-0.59}$	$66.85^{+0.76}_{-0.58}$	$66.86^{+0.70}_{-0.55}$
σ_8	$0.778^{+0.038}_{-0.024}$	$0.829^{+0.011}_{-0.007}$	$0.8291^{+0.0098}_{-0.0065}$	$0.8289^{+0.0094}_{-0.0066}$	$0.8291^{+0.0090}_{-0.0063}$
Parameter	Planck, TEP + DESI	LiteCORE-80, TEP + DESI	LiteCORE-120, TEP + DESI	CORE-M5, TEP + DESI	COrE+, TEP + DESI
$N_{\rm eff}^{\rm massive}$	3.07 ± 0.15	3.044 ± 0.061	3.047 ± 0.045	3.046 ± 0.040	3.044 ± 0.035
M_ν (meV)	74^{+35}_{-54}	65 ± 25	66^{+24}_{-22}	65^{+24}_{-21}	61 ± 21
$\Omega_b h^2$	0.02228 ± 0.00018	0.022257 ± 0.000082	0.022258 ± 0.000062	0.022257 ± 0.000057	0.022251 ± 0.000048
$\Omega_c h^2$	0.1200 ± 0.0025	0.1197 ± 0.0010	0.11973 ± 0.00075	0.11970 ± 0.00068	0.12002 ± 0.00059
$100\theta_s$	1.04080 ± 0.00045	1.04080 ± 0.00016	1.04080 ± 0.00012	1.04080 ± 0.00011	1.040800 ± 0.000097
τ	0.0608 ± 0.0045	0.0603 ± 0.0021	0.0603 ± 0.0021	0.0603 ± 0.0021	0.0604 ± 0.0018
n_s	0.9655 ± 0.0065	0.9644 ± 0.0032	0.9646 ± 0.0028	0.9645 ± 0.0026	0.9637 ± 0.0024
$ln(10^{10} A_s)$	3.096 ± 0.012	3.0944 ± 0.0049	3.0944 ± 0.0045	3.0944 ± 0.0044	3.0953 ± 0.0038
H_0 (km/s/Mpc)	67.05 ± 0.82	66.96 ± 0.42	66.97 ± 0.35	66.97 ± 0.33	66.97 ± 0.32
σ_8	0.830 ± 0.012	0.8307 ± 0.0045	0.8305 ± 0.0041	0.8305 ± 0.0039	0.8304 ± 0.0037
Parameter	Planck, TEP + DESI + Euclid	LiteCORE-80, TEP + DESI + Euclid	LiteCORE-120, TEP + DESI + Euclid	CORE-M5, TEP + DESI + Euclid	COrE+, TEP + DESI + Euclid
$N_{\rm eff}^{\rm massive}$	3.05 ± 0.11	3.044 ± 0.057	3.046 ± 0.042	3.046 ± 0.039	3.045 ± 0.034
M_ν (meV)	66^{+31}_{-35}	62 ± 20	62 ± 18	62 ± 17	62^{+15}_{-17}
$\Omega_b h^2$	0.02225 ± 0.00016	0.022253 ± 0.000081	0.022258 ± 0.000062	0.022256 ± 0.000055	0.022253 ± 0.000047
$\Omega_c h^2$	0.1198 ± 0.0017	0.11976 ± 0.00089	0.11978 ± 0.00067	0.11978 ± 0.00062	0.11977 ± 0.00054
$100\theta_s$	1.04080 ± 0.00038	1.04080 ± 0.00016	1.04080 ± 0.00012	1.04080 ± 0.00011	1.040800 ± 0.000092
τ	0.0607 ± 0.0045	0.0602 ± 0.0021	0.0603 ± 0.0021	0.0602 ± 0.0021	0.0595 ± 0.0020
n_s	0.9646 ± 0.0049	0.9644 ± 0.0029	0.9644 ± 0.0025	0.9644 ± 0.0024	0.9644 ± 0.0023
$ln(10^{10} A_s)$	3.095 ± 0.010	3.0944 ± 0.0048	3.0946 ± 0.0044	3.0945 ± 0.0043	3.0944 ± 0.0041
H_0 (km/s/Mpc)	66.97 ± 0.54	66.96 ± 0.32	66.98 ± 0.27	66.98 ± 0.25	66.97 ± 0.23
σ_8	$0.8313^{+0.0039}_{-0.0029}$	$0.8316^{+0.0035}_{-0.0026}$	$0.8315^{+0.0033}_{-0.0026}$	0.8315 ± 0.0028	$0.8315^{+0.0030}_{-0.0024}$

Table 24: Same as previous table, but for the ΛCDM + M_ν + $N_{\rm eff}^{\rm massive}$ model (accounting for non-thermalised active neutrinos degenerate in mass).

COrE+ mission. Instead, LiteCORE-120 would be 25% worse than COrE+. Hence CORE-M5 appears as a good compromise for the purpose of measuring $N_{\rm eff}$.

By achieving $\sigma(N_{\rm eff}) = 0.041$ with CORE-M5 alone, or $\sigma(N_{\rm eff}) = 0.039$ in combination with future BAO data from DESI and/or cosmic shear data from Euclid, we could set very strong bounds on extra relics, neutrino properties, the temperature of reheating, etc., especially compared to Planck + DESI BAOs, which would only yield $\sigma(N_{\rm eff}) = 0.15$. To be more specific, let us consider the case of early decoupled thermal relics, like in Ref. [82]. Assuming that the last-decoupled relics leave thermal equilibrium at a temperature $T_{\rm F}$, and that the subsequent number of relativistic degrees of freedom is entirely accounted for by standard model particles, we notice that there are many well-motivated scenarios predicting a value of $\Delta N_{\rm eff}$ ranging from 0.05 to 0.3, because this corresponds to particles decoupling during the QCD phase transition. In case of a non-detection of extra relics by CORE, the 95% exclusion bound from

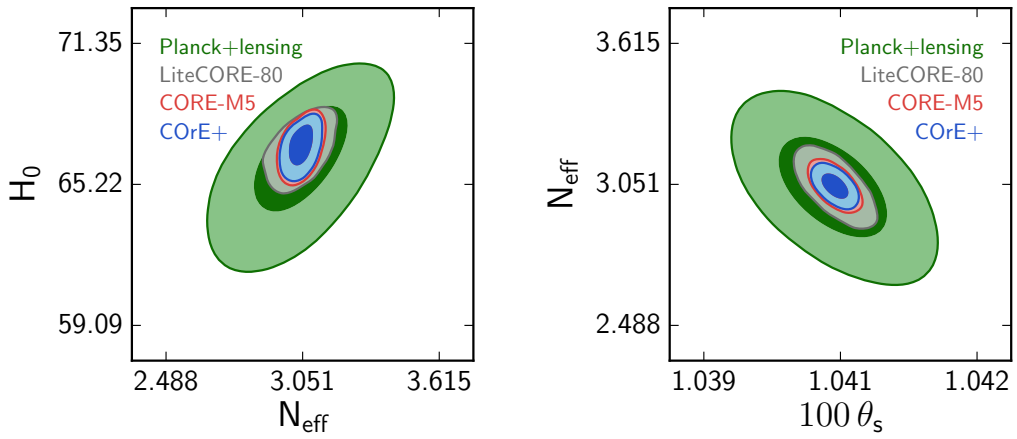

Figure 34: Parameter degeneracy between $N_{\rm eff}$ and H_0 or θ_s, assuming the extended model "DEG+Neff", with three experimental settings for CORE or with a fake Planck likelihood mimicking the sensitivity of the real experiment (always using all CMB information from TT,TE,EE + lensing extraction). The correlations observed in the Planck case are explained in the text. The degeneracy with H_0 is almost entirely resolved by CORE, while that with θ_s is limited to a much smaller range.

CORE + BAOs, $\Delta N_{\rm eff} < 0.076$, would exclude most of this range, while Planck + BAOs would not even touch it.

A sensitivity of $\sigma(N_{\rm eff}) = 0.041$ would also have crucial implications for the determination of other important cosmological parameters, through a considerable reduction of parameter degeneracies. For instance, without making assumptions on $N_{\rm eff}$, Planck + DESI BAOs would measure H_0 with 1.2% uncertainty, and $\omega_{\rm cdm}$ with 2% uncertainty. Figure 34 (left plot) shows that CORE-M5 would almost completely resolve the ($N_{\rm eff}$, H_0) degeneracy, such that CORE + DESI BAOs would pinpoint both H_0 and $\omega_{\rm cdm}$ with 0.5% uncertainty. This would have repercussions on several other parameters, and would allow to fully exploit the synergy between different types of cosmological data. Also, the determination of $N_{\rm eff}$ based on the observation of the CMB damping tails would reduce the uncertainty on the sound horizon angular scale, from $\sigma(\theta_s) = 0.00046$ for Planck to $\sigma(\theta_s) = 0.00011$ for CORE: hence the calibration of the sound horizon scale in future BAO data would be much more accurate, and the scientific impact of these observations (for instance, on Dark Energy models) would be enhanced.

9.3 LIGHT STERILE NEUTRINOS

Right-handed or sterile neutrinos are present in several well-motivated extensions of the standard model of particle physics [78, 315]. If their mass is of the order of a few keV or bigger, they can play the role of warm or cold dark matter, and they are constrained mainly by X-ray and Lyman-alpha observations [315]. If their mass is of the order of the meV or smaller, they will simply behave as extra relativistic relics contributing to $N_{\rm eff}$. There is another interesting range deserving a specific study: that of light sterile neutrinos with a mass in the meV to eV range. Such particles have been extensively discussed over the past years, for the reason that the oscillations between such sterile neutrinos and active neutrinos (or more precisely, between the mass eigenstates formed of active and sterile neutrinos) could explain a number of possible anomalies in short-baseline neutrino oscillation data (see e.g. Ref. [316]).

Sterile neutrinos with large mixing angles would normally acquire a thermal distribution through oscillations with active neutrinos, and their mass would then be very constrained (essentially, as much as that of active neutrinos). However, the explanation of short baseline anomalies requires an $\mathcal{O}(1)$ eV mass in tension with cosmological data. To avoid these bounds, people have discussed several ways to prevent sterile neutrino thermalisation (see e.g. Refs. [78, 79, 317]). In that case, the bounds on the sterile neutrino mass become model-dependent, but a wide category of models can be parametrised in good approximation with two numbers (N_s, $m_s^{\rm eff}$), related to the asymptotic density at early times, given by $\Delta N_{\rm eff} = N_s$, and the asymptotic density at late times, given by the effective mass $m_s^{\rm eff} = 94.1\omega_s$ eV [2, 61], where ω_s is the sterile neutrino density. This covers both the case of light early-decoupled thermal relics, and that of Dodelson-Widrow (i.e. non-resonantly produced) sterile neutrinos. For the later case, the physical mass of the sterile neutrino is given by $m_s = m_s^{\rm eff}/N_s$.

To investigate the sensitivity of CORE to a non-thermal sterile neutrino, we stick to the same fiducial model as in the last subsections (total mass $M_\nu = 60$ meV and $N_{\rm eff} = 3.046$), but we now fit it with an extended model with 9 free parameters, including the summed mass of active neutrinos $M_\nu^{\rm active}$, as well as N_s and $m_s^{\rm eff}$. We impose in our forecasts a top-hat prior $m_s^{\rm eff}/N_s < 5$ eV, designed to eliminate models such that the extra species has a large mass, a very small number density, and behaves like extra cold dark matter.

Parameter	Planck, TEP	LiteCORE-80, TEP	LiteCORE-120, TEP	CORE-M5, TEP	COrE+, TEP
$m_{\rm eff}^{\rm sterile}$ (meV)	< 332	< 56	< 41	< 37	< 34
$\Delta N_{\rm eff}^{\rm sterile}$	< 0.20	< 0.080	< 0.058	< 0.053	< 0.048
$M_\nu^{\rm active}$ (meV)	fixed	< 79	< 76	< 77	< 82
Parameter	Planck, TEP + DESI	LiteCORE-80, TEP + DESI	LiteCORE-120, TEP + DESI	CORE-M5, TEP + DESI	COrE+, TEP + DESI
$m_{\rm eff}^{\rm sterile}$ (meV)	< 123	< 52	< 42	< 39	< 35
$\Delta N_{\rm eff}^{\rm sterile}$	< 0.168	< 0.077	< 0.059	< 0.054	< 0.049
$M_\nu^{\rm active}$ (meV)	< 65	43^{+21}_{-33}	45^{+24}_{-28}	47 ± 26	48 ± 26
Parameter	Planck, TEP + DESI + Euclid	LiteCORE-80, TEP + DESI + Euclid	LiteCORE-120, TEP + DESI + Euclid	CORE-M5, TEP + DESI + Euclid	COrE+, TEP + DESI + Euclid
$m_{\rm eff}^{\rm sterile}$ (meV)	< 60	< 43	< 37	< 35	< 33
$\Delta N_{\rm eff}^{\rm sterile}$	< 0.151	< 0.074	< 0.057	< 0.054	< 0.049
$M_\nu^{\rm active}$ (meV)	< 63	43^{+23}_{-26}	44 ± 23	44 ± 22	44^{+22}_{-20}

Table 25: 68% CL constraints and upper bounds on cosmological parameters in the ΛCDM + $M_\nu^{\rm active}$ + $m_{\rm eff}^{sterile}$ + $\Delta N_{\rm eff}^{sterile}$ model (accounting for massive active neutrinos plus one light and non-thermalised sterile neutrino) from the different CORE experimental specifications and with or without external data sets (DESI BAOs, Euclid cosmic shear). For Planck alone, we quote the results from the 2015 data release, *obtained with a fixed active neutrino mass* $M_\nu^{\rm active} = 60$ meV, while for combinations of Planck with future surveys, we fit mock data with a fake Planck likelihood mimicking the sensitivity of the real experiment (although a bit more constraining). For concision, we only show the bounds for the extended model parameters.

Our results for the parameters $(M_\nu^{\rm active}, N_s, m_s^{\rm eff})$ are given in Table 25, and the probability contours for $(N_s, m_s^{\rm eff})$ are shown in Figure 35. For CORE-M5, the bounds on the sterile sector are impressive:

$$(m_s^{\rm eff}, N_s) < (37\,{\rm meV}, 0.053), \quad ({\rm CORE-M5},\ 68\%{\rm CL})$$

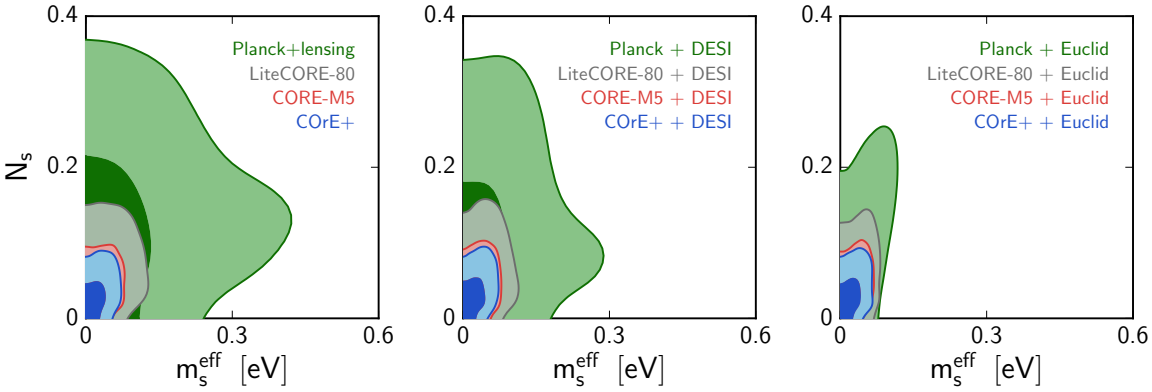

Figure 35: Results for the extended model ΛCDM + M_ν + one light and non-thermalised sterile neutrino with effective mass $m_s^{\rm eff}$, contributing to the effective neutrino number as N_s.

to be compared with

$$(m_{\rm s}^{\rm eff}, N_{\rm s}) < (330\,{\rm meV}, 0.2). \quad ({\rm PlanckTT + lowP + lensing + BAO,\ 68\%CL}).$$

The sensitivity to $(m_{\rm s}^{\rm eff}, N_{\rm s})$ depends heavily on the CORE settings. The error on N_s varies by a factor two between LiteCORE-80 and COrE+. As discussed in section 9.2, this comes mainly from the ability to measure the temperature and polarisation damping tail up to high multipoles when the instrumental sensitivity and resolution are good enough. Besides, the measurement of the CMB lensing potential constrains the density of hot dark matter today, and hence roughly $M_\nu^{\rm active} + m_{\rm s}^{\rm eff}$. If this were the only effect, all CORE configurations would lead essentially to $M_\nu^{\rm active} + m_{\rm s}^{\rm eff} = 60 \pm 44$ meV at one sigma, and to the same constraints on $m_{\rm s}^{\rm eff}$. However, there is some extra sensitivity to $m_{\rm s}^{\rm eff}$ coming from the fact that for small $N_{\rm s}$, the physical mass associated to a given value of $m_{\rm s}^{\rm eff}$ can be large[2], such that the sterile neutrinos have their non-relativistic transition before photon decoupling. In that case, there are additional effects on CMB primary anisotropies[3] that an experiment sensitive to smaller angular scales can constrain better. This explains the gain in sensitivity to $m_{\rm s}^{\rm eff}$ between LiteCORE-80 and COrE+. CORE-M5 appears as a good compromise, more constraining than LiteCORE-80 by 50% for both $N_{\rm s}$ and $m_{\rm s}^{\rm eff}$. In summary, with a sensitivity to $m_{\rm s}^{\rm eff}$ ten times better than Planck, CORE-M5 appears as an ideal instrument for constraining light sterile neutrinos, and the CORE data release will play a key role in the discussion of anomalies in short baseline neutrino oscillations.

Note that with CORE data alone, we find no lower bound on the active neutrino mass M_ν in presence of a sterile neutrino, because the physical effect of the mass $M_\nu = 60$ meV in the fiducial model can be partially endorsed by the sterile neutrino mass. In other words, the data is not able to tell whether the fiducial mass of 60 meV belongs to active neutrinos, or to a mixture of sterile and active neutrinos. By removing degeneracies, BAO data from DESI makes the CMB lensing spectrum more sensitive to $M_\nu^{\rm active} + m_{\rm s}^{\rm eff}$, and given the upper bound on $m_{\rm s}^{\rm eff}$, one now finds a lower bound on $M_\nu^{\rm active}$. Cosmic shear data from Euclid directly probes the free-streaming effect associated with $M_\nu^{\rm active} + m_{\rm s}^{\rm eff}$, which results in a slightly better sensitivity to $M_\nu^{\rm active}$, but the constraints on the sterile neutrino sector remain roughly the same as when considering CORE alone.

2 More precisely, the velocity dispersion given by $\langle p \rangle / m_{\rm s}$ can be much smaller than for active neutrinos

3 Hot Dark Matter particles become non-relativistic before photon decoupling have a direct impact on CMB fluctuations at the level of primary anisotropies: they tend to suppress small-scale fluctuations.

9.4 EXTENDED MODELS WITH IMPROVED LARGE-SCALE STRUCTURE TREATMENT

In this section, we describe the results obtained by fitting our mock data to three physically motivated extensions of the baseline model: (1) ΛCDM+M_ν+$N_{\rm eff}$, (2) ΛCDM+M_ν+w_0, (3) ΛCDM+M_ν+(w_0, w_a). These extensions have been discussed at various points in the past as potentially degenerate to some extent with neutrino masses, so the results of this section also aim at assessing the robustness of a future neutrino mass detection against extended cosmology assumptions. We leave for future work the study of models with a free spatial curvature parameter, which is also likely to degrade the sensitivity to the total neutrino mass (see e.g. [304] for a recent discussion).

In Table 26, we report the expected sensitivity of various combinations of Planck, Euclid and SKA probes to the non-standard cosmological parameters and to those among the standard ones that show a relevant deviation from the baseline case. In Figure 36, we depict the corresponding $1\,\sigma$ uncertainty, to show how the sensitivity degrades in extended models. The decrease in the accuracy is caused by the degeneracies illustrated in the contour plots of Figure 37.

We shall now proceed to discuss the results for each one of the three extended models. Notice that throughout this section, what we call the "realistic" case assumes the realistic theoretical error prescription of section 3.4 for GC and IM, but still the conservative one for CS.

9.4.1 ΛCDM+M_ν+$N_{\rm eff}$

We first promote the effective number of relativistic degrees of freedom $N_{\rm eff}$. This quantity parameterizes the radiation density (ρ_r) of the universe in the early universe beyond the photon density (ρ_γ),

$$\rho_r = \rho_\gamma \left[1 + \frac{7}{8}\left(\frac{4}{11}\right)^{4/3} N_{\rm eff}\right].$$

In the cosmological standard model $N_{\rm eff}$ is equal to 3.045 [307, 311], representing the three active neutrinos. We use this as our fiducial value. A deviation of $N_{\rm eff}$ from 3.045 could be caused, for instance, by several plausible extensions of the Standard Model of particle physics, e.g. with sterile neutrinos [317–319]. Notice that a variation of $N_{\rm eff}$ implies a different expansion rate of the universe, with a profound impact on the cosmological observables at any redshift [66, 313, 314, 320, 321].

From Table 26, we can see that in the realistic theoretical error scenario the uncertainty is $\sigma(N_{\rm eff}) \leq 0.051$ for Planck + Euclid and Planck + SKA2. Such an accuracy allows for a detection of more exotic models, e.g. with new bosons from new broken symmetries, leading to early decoupled or partially thermalized additional degrees of freedom [322].

In Figure 36 the predictions for this model appear as squared markers. We see that the sensitivity to M_ν, $\omega_{\rm cdm}$ and H_0 degrades significantly with respect to the baseline model (circle markers) for every probe combination. In general, this degradation is caused by the strong degeneracies between $N_{\rm eff}$ and the aforementioned parameters. The importance of these correlations is illustrated in the top-left plot of Figure 37. Observables that are more sensitive to the expansion rate of the Universe, such as cluster number counts, can break these degeneracies [216], and would bring back the sensitivity to all the baseline parameters close to the values obtained in the baseline model. Interestingly, the strongest correlations affect $\omega_{\rm cdm}$ and

		Planck+SKA2	Planck+Euclid	Planck+Euclid+SKA1	Planck+SKA1
	σ	(CS+GC)	(CS+GC)	(CS+GC)+(IM2)	(IM1)
ΛCDM	$\omega_{\rm cdm}$	0.00080	0.00083	0.00082	0.00121
		0.00073	0.00073	0.00071	0.00118
$+\,{\rm M}_\nu$	$H_0/[\frac{\rm km}{\rm s\,Mpc}]$	0.237	0.272	0.251	0.350
		0.205	0.207	0.194	0.334
$+\,{\rm N}_{\rm eff}$	$M_\nu/[{\rm meV}]$	0.021	0.022	0.022	0.028
		0.015	0.016	0.014	0.027
	$N_{\rm eff}$	0.048	0.051	0.050	0.076
		0.041	0.041	0.039	0.073
ΛCDM	$\omega_{\rm cdm}$	0.00029	0.00030	0.00030	0.00030
		0.00026	0.00026	0.00025	0.00030
$+\,{\rm M}_\nu$	$H_0/[\frac{\rm km}{\rm s\,Mpc}]$	0.081	0.154	0.088	0.100
		0.072	0.105	0.069	0.100
$+\,{\rm w}_0$	$M_\nu/[{\rm meV}]$	0.019	0.021	0.019	0.020
		0.016	0.017	0.013	0.020
	w_0	0.0033	0.0073	0.0029	0.0046
		0.0023	0.0047	0.0014	0.0046
ΛCDM	$\omega_{\rm cdm}$	0.00029	0.00029	0.00030	0.00034
		0.00026	0.00026	0.00025	0.00034
$+\,{\rm M}_\nu$	$H_0/[\frac{\rm km}{\rm s\,Mpc}]$	0.101	0.232	0.099	0.200
		0.084	0.162	0.072	0.202
$+\,{\rm w}_0$	$M_\nu/[{\rm meV}]$	0.027	0.027	0.025	0.030
		0.024	0.024	0.018	0.031
$+\,{\rm w}_{\rm a}$	w_0	0.0045	0.0157	0.0046	0.0114
		0.0032	0.0098	0.0016	0.0115
	$w_{\rm a}$	0.027	0.050	0.027	0.049
		0.020	0.033	0.013	0.050

Table 26: Expected $1\,\sigma$ sensitivity of Planck, Euclid and SKA to the cosmological parameters relevant in each extended model. For each probe combination the results in the top row are obtained with the conservative theoretical error approach applied to every observable; the results in the bottom row assume the realistic error prescription for GC and IM, while the conservative one is used for CS.

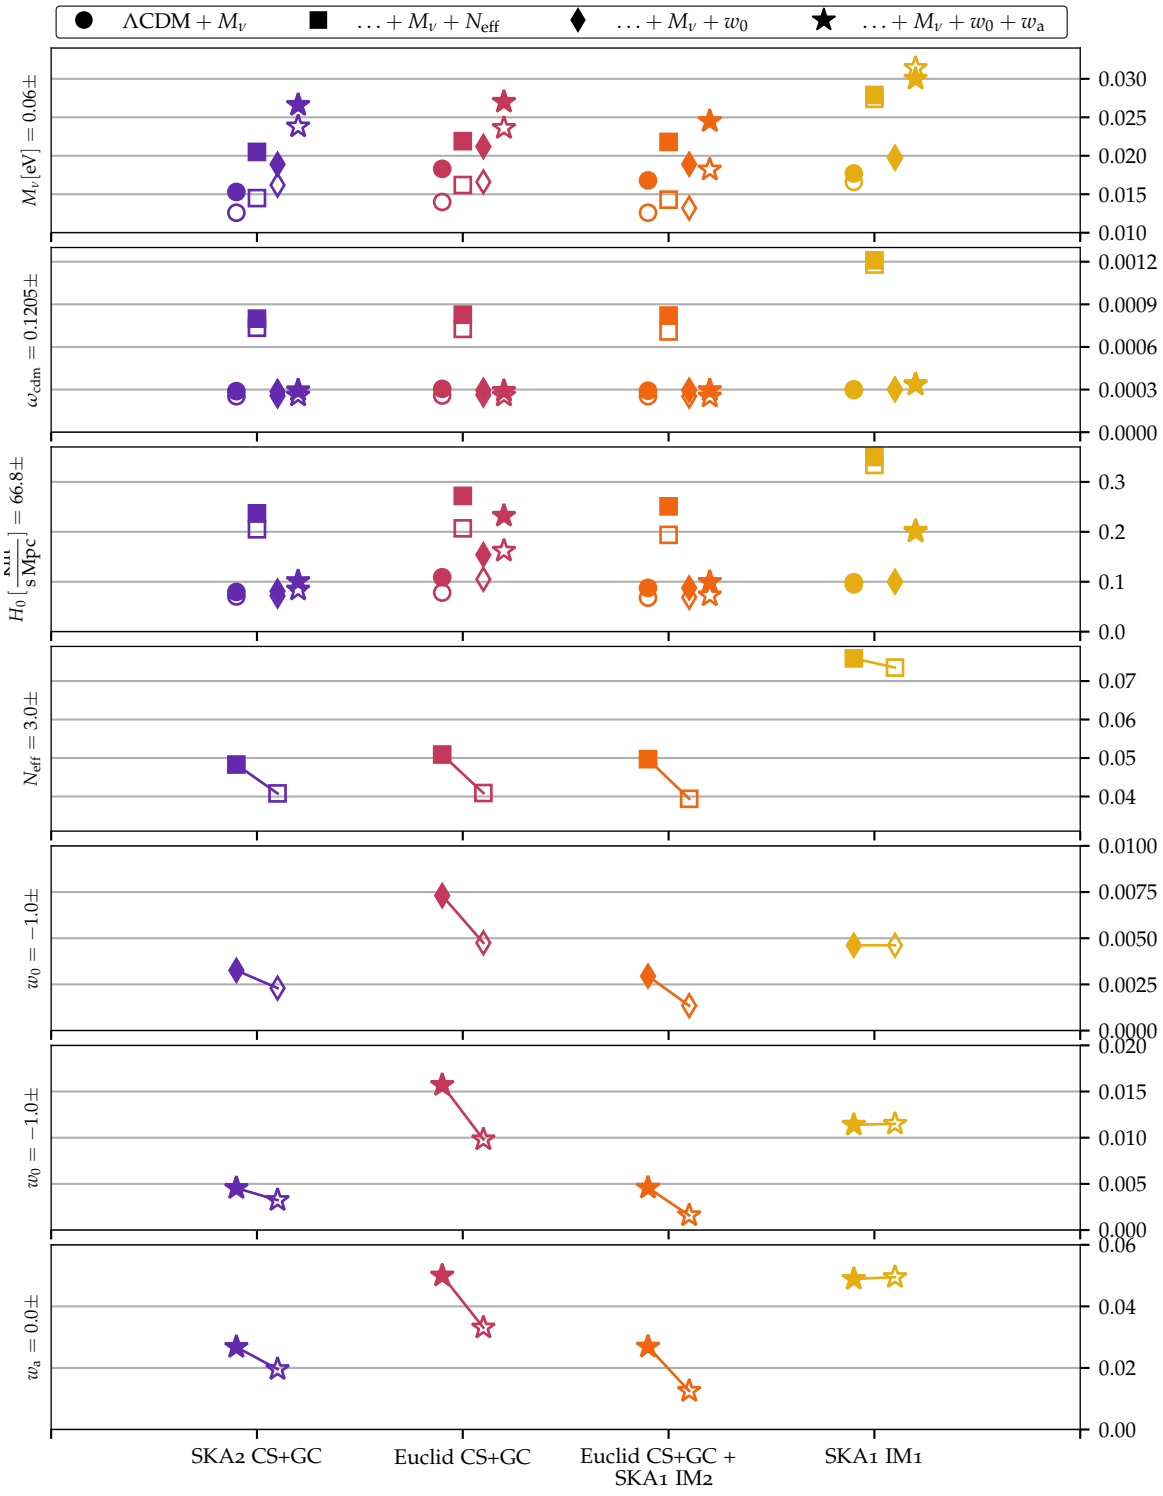

Figure 36: $1\,\sigma$ uncertainty for various combinations of experiments, all including Planck. The different experiments, as well as the probe combinations, are specified on the x axis at the bottom of the plot. The fiducial value is given on the left axis, while the $1\,\sigma$ uncertainty is written on the right axis. Finally, empty markers denote realistic settings, while filled markers denote conservative settings. Different markers denote different extended models, according to the legend at the top of the plot.

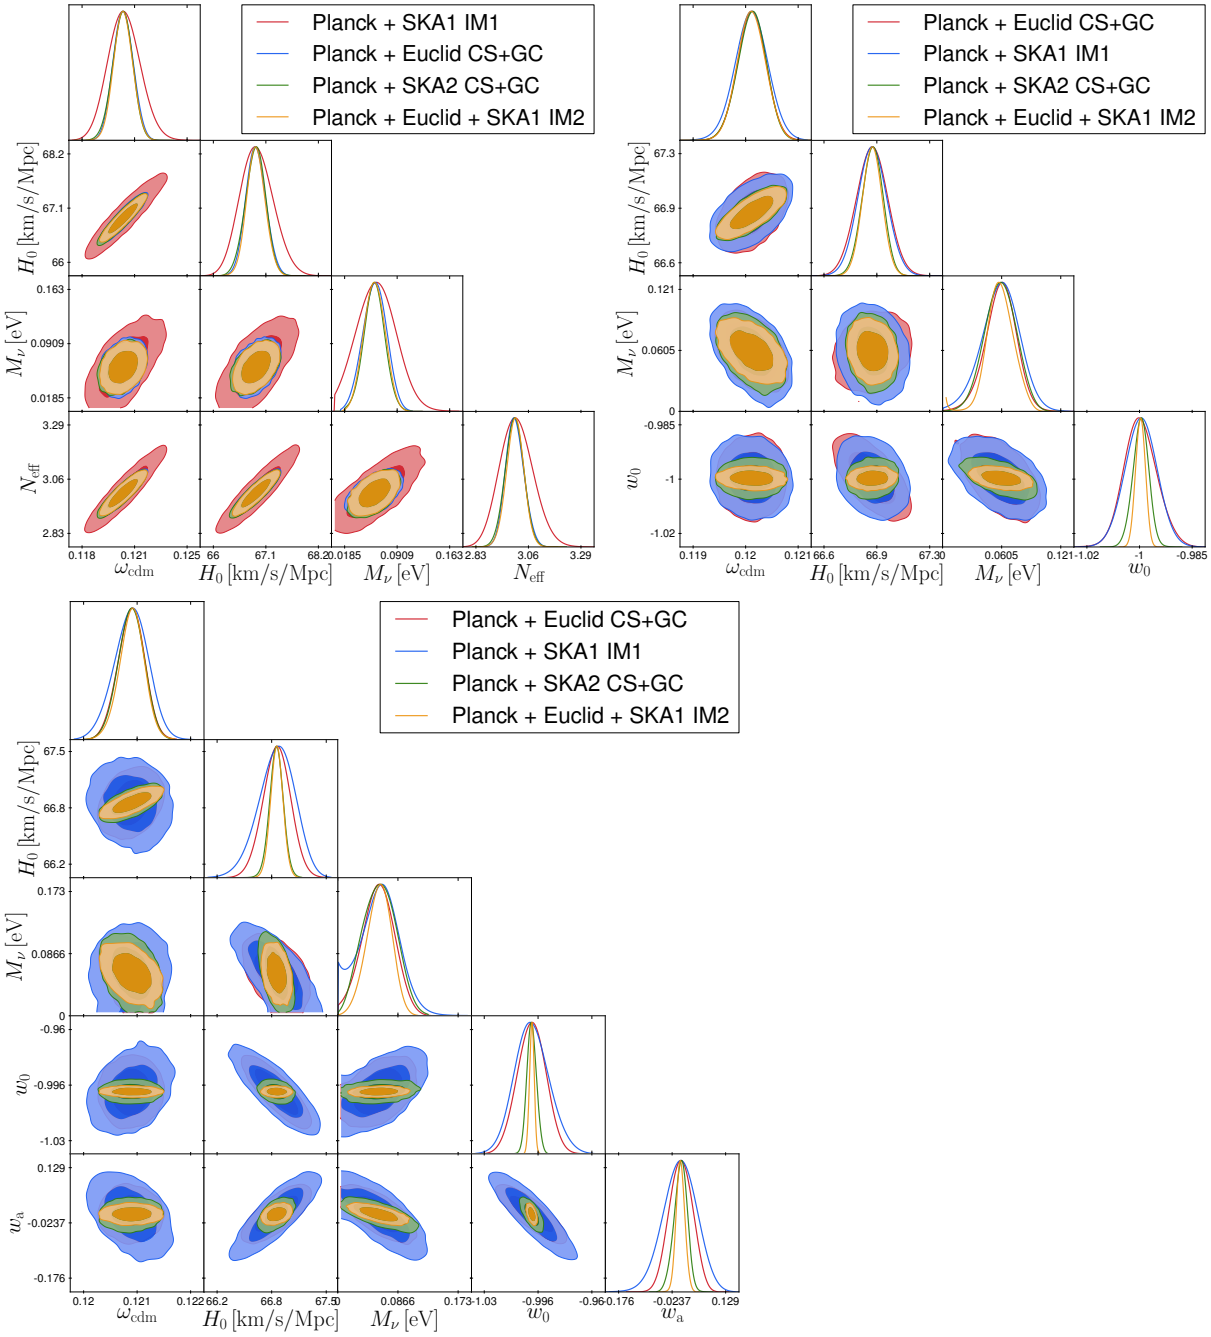

Figure 37: Marginalized one- and two-σ contours and one dimensional posteriors for the three extended models ΛCDM+M_ν+$N_{\rm eff}$ (top left), ΛCDM+M_ν+w_0 (top right) and ΛCDM+M_ν+(w_0, w_a) (bottom left), in the realistic theoretical error scenario.

H_0, for which the reduction in sensitivity is more prominent than in the other dark energy extended models. Instead, concerning M_ν, adding SKA (IM2) to Euclid with a realistic theoretical error leads to $\sigma(M_\nu) = 0.014$ eV, i.e. nearly the same sensitivity as in the baseline model, and a 4.3σ detection of a non-zero absolute neutrino mass. Therefore, the effect of $N_{\rm eff}$ on the cosmological probes can easily be disentangled from the effect of the neutrino mass sum with forthcoming galaxy and hydrogen surveys.

9.4.2 $\Lambda CDM+M_\nu+w_0$

The second extended model includes a constant Dark Energy equation of state parameter w_0, with fiducial value -1 as in the ΛCDM model. From Table 26 and Figure 36, we observe an important degradation (by almost a factor two) of the sensitivity to M_ν and H_0 for Euclid CS+GC and SKA2 CS+GC. The importance of the parameter degeneracies with w_0 can be seen on the posterior distributions, displayed on Figure 37. But we also find that these degeneracies can be reduced, 1) by combining Euclid with SKA1-IM, 2) in the realistic theoretical error scenario, in such a way that the sensitivities to M_ν and H_0 are brought back to the ones of the baseline model.

It is also found that the $1\,\sigma$ sensitivity to w_0 is improved almost by a factor two when considering the realistic theoretical error scenario, with $\sigma(w_0) = 0.0047/0.0023/0.0014$ respectively for Planck+Euclid, Planck+SK2 and Planck+Euclid+SKA1-IM. It is therefore worth noticing that combining Euclid with the low-redshift SKA intensity mapping survey allows a very important improvement of the forecasted sensitivity to w_0. This is due to the intensity mapping of band 2 of SKA1 being sensitive down to very low redshift ($z_{\rm min} = 0.05$).

9.4.3 $\Lambda CDM+M_\nu+(w_0, w_a)$

The third extended model allows for a variation of the Dark Energy equation of state with the CPL parameterization [237]

$$w(a) = w_0 + (1 - a)w_a \ .$$

We adopt the ΛCDM values of these parameters $w_0 = -1$ and $w_a = 0$ in the fiducial model.

In the conservative scenario, our results for Planck+Euclid are well compatible with those of Ref. [4], and globally the sensitivities are significantly degraded with respect to the previous model ΛCDM+M_ν+w_0 model. In the conservative scenario, this is still true for the Planck+SKA1-IM combination. However, as soon as we include CS+GC probes with a realistic error, the results are impressively stable and very mildly affected by the presence of an additional free parameter, as can be seen in Table 26 and Figure 36. In Figure 37 we see the importance of adding CS+GC information (green and yellow contours) in order to break the degeneracies of the Planck+SKA1-IM results (blue contours).

Ultimately, by combining Planck, Euclid CS+GC and SKA1-IM, we can expect sensitivities down to $\sigma(w_0) = 0.0016$ and $\sigma(w_a) = 0.013$, as well as $\sigma(M_\nu) = 0.018\,$eV. By comparing with the results obtained for Planck and Euclid CS+GC only, we see that intensity mapping with SKA1 should lead to a useful increase of the sensitivity to the neutrino mass (allowing potentially for a 3.4σ instead of 2.5σ detection of the minimal mass), and to a very strong improvement in the sensitivity to the two DE parameters.

9.5 FUTURE CMB EXPERIMENTS IN COMBINATION WITH IMPROVED LSS TREATMENT

The complete results of our forecasts are summarised in in tables 27 to 32, tables 27 to 32. However, as the main focus of this thesis is neutrinos, we will restrict our discussion to the results for the neutrino mass sum. A more intuitive and graphical summary of the sensitivity to the neutrino mass sum is presented in the form of the standard deviation of that parameter in figs. 38 and 39.

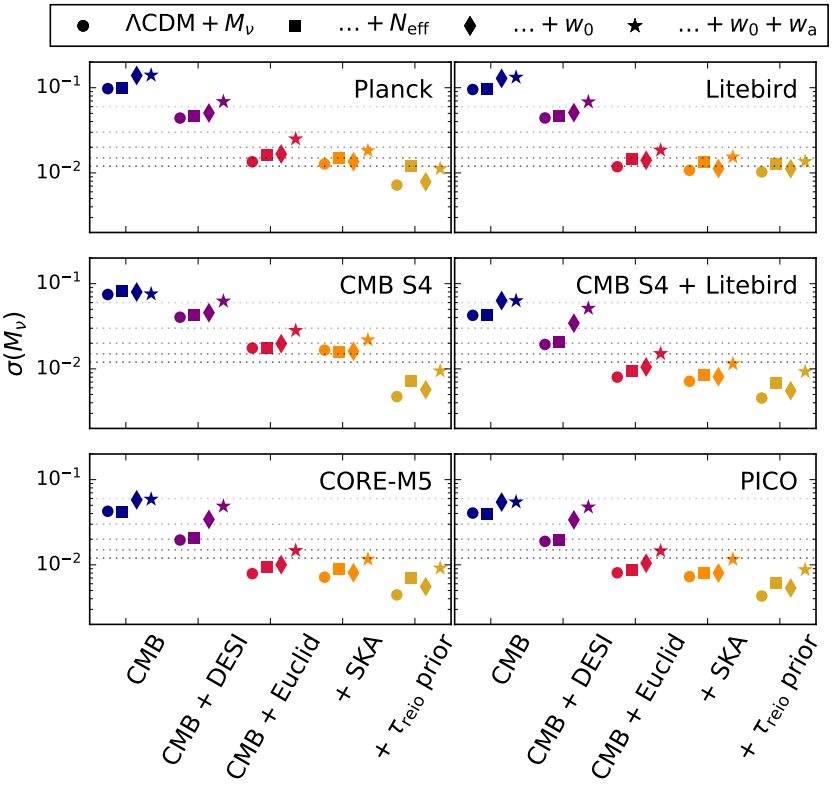

Figure 38: Neutrino mass sensitivity for each CMB experiment, alone and in combination with DESI, Euclid, Euclid + SKA1 IM, Euclid + SKA1 IM + τ_{reio} prior. Each subplot corresponds to one CMB setup (Planck, Litebird, CMB-S4, CMB-S4+Litebird, CORE-M5, or PICO from top left to bottom right) and relevant combinations with large-scale structure surveys. The five vertical dotted lines indicate 1 to 5-σ significance towards a detection of a non-zero neutrino mass of around the minimum normal hierarchy mass $M_\nu \approx 60$ meV. See text for discussion.

The pair of figures actually contain redundant information, but fig. 38 is ordered in a way to highlight the impact of each LSS dataset in combination with a given CMB experiment, whereas fig. 39 show the importance of using more precise CMB datasets in combination with a given LSS experiment.

First let us consider the minimal cosmology scenario (ΛCDM + M_ν). Figures 38 and 39 show that, as long as a minimal cosmology is assumed, for a minimal normal hierarchy scenario with $M_\nu = 0.06$ eV we find:

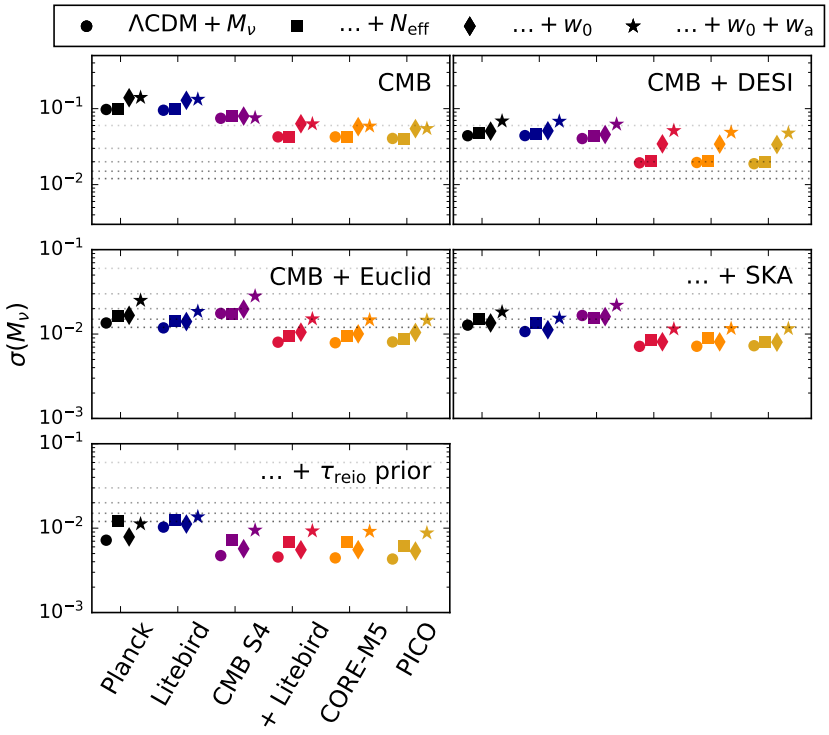

Figure 39: Neutrino mass sensitivity for each CMB experiment, alone and in combination with DESI, Euclid, Euclid + SKA1 IM, Euclid + SKA1 IM + τ_{reio} prior. Each subplot corresponds to one CMB setup plus large-scale structure survey combination (CMB only, CMB+DESI, CMB+Euclid, CMB+Euclid+SKA1 IM, CMB+Euclid+SKA1 IM+$\tau_{\rm reio}$ prior from top left to bottom right). The vertical dotted lines indicate 1 to 5-σ significance towards a detection of a non-zero neutrino mass of around the minimum normal hierarchy mass $M_\nu \approx 60$ meV. See text for discussion.

- CORE and PICO are so sensitive that they would only need to be combined with the BAO scale data from DESI for a 3-σ detection,
- more than a 4-σ detection could be achieved already by Planck or LiteBird when combined with Euclid,
- LiteBird in combination with Euclid and SKA1 intensity mapping reaches the 5-σ threshold, which is nearly true for Planck as well,
- a 7-σ measurement could be achieved by CORE or PICO in combination with Euclid,
- a staggering 13-σ detection once SKA1 intensity mapping and the $\tau_{\rm reio}$ prior is added, illustrating the enormous benefit towards a precise neutrino mass detection from having a very accurate independent determination of $\tau_{\rm reio}$, e.g. from surveys focused on reionization and the dark ages,
- even when combining with Euclid and SKA1, only a 3-σ detection could be achieved by CMB-S4 when viewed in isolation. However, it is important to keep in mind that adding

information from low-ℓ polarization data strongly constrains $\tau_{\rm reio}$, which leads to greatly improved sensitivity on M_ν, and therefore CMB-S4 provides much better sensitivity once LiteBird or the $\tau_{\rm reio}$ prior is included (similarly, low-ℓ Planck data would already help in this regard). This is the effect we discussed in chapter 5 and the effect is illustrated in fig. 40.

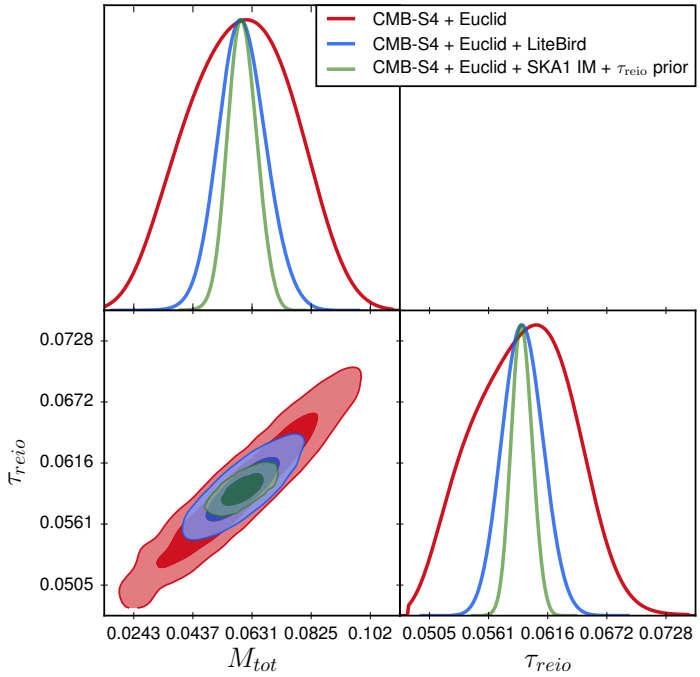

Figure 40: Optical depth to reionization $\tau_{\rm reio}$ vs neutrino mass for the three configurations CMB-S4 + Euclid, CMB-S4 + Euclid + LiteBird, CMB-S4 + Euclid + SKA1 IM + $\tau_{\rm reio}$ prior in the minimal 7 parameter ΛCDM+M_ν model. See text for discussion.

For extended models the neutrino mass sum sensitivity degrades, as expected. We make the following observations about the sensitivity of our surveys to the neutrino mass sum in extended models, comparing to a minimal normal hierarchy scenario with $M_\nu = 0.06$ eV:

- the sensitivity only degrades slightly when also varying the number of extra relativistic degrees of freedom, $N_{\rm eff}$, as expected for current and future surveys. The exception is for forecasts involving the $\tau_{\rm reio}$ prior: if we could precisely determine the optical depth to reionization, we would be able to strongly constrain the neutrino mass sum (at better than 5-σ, even with current CMB data in combination with Euclid and SKA1 intensity mapping), but the addition of $N_{\rm eff}$, which is strongly degenerate with the CDM energy density and slightly so with the spectral index of the primordial power spectrum and the neutrino mass sum, and through these with the amplitude of the primordial power spectrum, A_s, which in turn means accurately determining the optical depth to reionization (and thereby A_s) is a little less helpful in precisely measuring the neutrino mass sum than for cases without varying $N_{\rm eff}$,
- the worst sensitivity is always obtained when including dynamical dark energy, i.e. in order to obtain a reasonable level of significance for the detection CMB plus DESI BAO is never enough and we need at least the combination of LiteBird and Euclid (3-σ), CORE or PICO along with Euclid (4-σ), or Planck with both Euclid and SKA1 IM (3-σ),

- in order to reach a 5-σ neutrino mass sum detection, we would require a very sensitive CMB experiment (i.e. CORE, PICO, or a combination of CMB-S4 and LiteBird) combined with Euclid and SKA1 intensity mapping, when including dynamical dark energy,
- however, if we are able to make an accurate independent measurement of the optical depth to reionization, even Planck in combination with Euclid and SKA1 intensity mapping would be enough for a 5-σ neutrino mass detection in any of the extended models considered, showing that *the expectation of a neutrino mass detection from cosmology in the next decade is robust to different choices of cosmological model*.

		CMB only	CMB + DESI	CMB + Euclid	+ SKA	+ $\tau_{\rm reio}$ prior
	$\sigma\left(100*\omega_{\rm b}\right)$	0.016	0.013	0.011	0.011	0.011
	$\sigma\left(\omega_{\rm cdm}\right)$	0.0014	0.00076	0.00026	0.00026	0.00017
ΛCDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	1.4	0.25	0.08	0.073	0.069
	$\sigma\left(\ln 10^{10} A_s\right)$	0.0089	0.0089	0.0063	0.0063	0.0023
+ M_ν	$\sigma\left(n_s\right)$	0.004	0.003	0.00056	0.0005	0.00047
	$\sigma\left(\tau_{\rm reio}\right)$	0.0045	0.0045	0.0034	0.0033	0.00097
	$\sigma\left(M_\nu\right)/[eV]$	0.097	0.044	0.014	0.013	**0.0072**
	$\sigma\left(100*\omega_{\rm b}\right)$	0.025	0.019	0.013	0.012	0.012
	$\sigma\left(\omega_{\rm cdm}\right)$	0.003	0.003	0.00076	0.00075	0.00071
ΛCDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	2.1	0.91	0.21	0.21	0.2
	$\sigma\left(\ln 10^{10} A_s\right)$	0.013	0.012	0.0065	0.0065	0.0035
+ M_ν	$\sigma\left(n_s\right)$	0.0091	0.0069	0.00068	0.00066	0.00059
	$\sigma\left(\tau_{\rm reio}\right)$	0.0046	0.0045	0.0033	0.0033	0.00095
+ $N_{\rm eff}$	$\sigma\left(M_\nu\right)/[eV]$	0.099	0.047	0.016	0.015	**0.012**
	$\sigma\left(N_{\rm eff}\right)$	0.19	0.17	0.044	0.043	0.044
	$\sigma\left(100*\omega_{\rm b}\right)$	0.017	0.013	0.011	0.011	0.011
	$\sigma\left(\omega_{\rm cdm}\right)$	0.0016	0.00082	0.00027	0.00026	0.00018
CDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	21	0.84	0.11	0.073	0.069
	$\sigma\left(\ln 10^{10} A_s\right)$	0.0089	0.0089	0.0065	0.0063	0.0023
+ M_ν	$\sigma\left(n_s\right)$	0.0043	0.0031	0.00065	0.0005	0.00046
	$\sigma\left(\tau_{\rm reio}\right)$	0.0047	0.0045	0.0035	0.0033	0.00099
+ w_0	$\sigma\left(M_\nu\right)/[eV]$	0.14	0.051	0.017	0.014	**0.0079**
	$\sigma\left(w_0\right)$	0.48	0.038	0.0051	0.0014	0.0014
	$\sigma\left(100*\omega_{\rm b}\right)$	0.018	0.013	0.011	0.012	0.011
CDM	$\sigma\left(\omega_{\rm cdm}\right)$	0.0016	0.00092	0.00027	0.00027	0.00019
	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	20	1.7	0.17	0.079	0.072
+ M_ν	$\sigma\left(\ln 10^{10} A_s\right)$	0.009	0.0087	0.0075	0.0067	0.0023
	$\sigma\left(n_s\right)$	0.0043	0.0032	0.00066	0.00055	0.00046
+ w_0	$\sigma\left(\tau_{\rm reio}\right)$	0.0046	0.0045	0.0039	0.0035	0.00096
	$\sigma\left(M_\nu\right)/[eV]$	0.14	0.069	0.025	0.018	**0.011**
+ w_a	$\sigma\left(w_0\right)$	0.84	0.2	0.01	0.0016	0.0016
	$\sigma\left(w_a\right)$	2.2	0.52	0.034	0.013	0.012

Table 27: Expected $1\,\sigma$ sensitivity of for Planck alone and in combination with DESI, Euclid, Euclid + SKA1 IM, Euclid + SKA1 IM + τ_{reio} prior. Combinations reaching the 3-σ (5-σ) threshold towards a neutrino mass sum detection for a minimal normal hierarchy neutrino mass sum $M_\nu = 0.06$ are highlighted in green (bolded dark green), see text for discussion.

		CMB only	CMB + DESI	CMB + Euclid	+ SKA	+ $\tau_{\rm reio}$ prior
	$\sigma\left(100*\omega_{\rm b}\right)$	0.018	0.014	0.011	0.011	0.011
	$\sigma\left(\omega_{\rm cdm}\right)$	0.0011	0.00072	0.00024	0.00023	0.00024
ΛCDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	1.3	0.25	0.09	0.082	0.082
	$\sigma\left(\ln 10^{10} A_s\right)$	0.0051	0.0052	0.0042	0.004	0.0039
+ M_ν	$\sigma\left(n_s\right)$	0.0045	0.004	0.00056	0.00052	0.00051
	$\sigma\left(\tau_{\rm reio}\right)$	0.0022	0.0022	0.002	0.0019	0.0019
	$\sigma\left(M_\nu\right)/[eV]$	0.095	0.044	**0.012**	**0.011**	**0.01**
	$\sigma\left(100*\omega_{\rm b}\right)$	0.025	0.02	0.011	0.011	0.011
	$\sigma\left(\omega_{\rm cdm}\right)$	0.0046	0.0045	0.00095	0.00093	0.00092
ΛCDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	1.9	1.3	0.25	0.24	0.24
	$\sigma\left(\ln 10^{10} A_s\right)$	0.012	0.012	0.0043	0.0042	0.0031
+ M_ν	$\sigma\left(n_s\right)$	0.008	0.0075	0.00082	0.00081	0.00077
	$\sigma\left(\tau_{\rm reio}\right)$	0.0022	0.0021	0.0019	0.0019	0.00091
+ $N_{\rm eff}$	$\sigma\left(M_\nu\right)/[eV]$	0.098	0.047	0.014	0.013	0.013
	$\sigma\left(N_{\rm eff}\right)$	0.27	0.26	0.059	0.057	0.058
	$\sigma\left(100*\omega_{\rm b}\right)$	0.018	0.015	0.011	0.011	0.01
	$\sigma\left(\omega_{\rm cdm}\right)$	0.0012	0.00078	0.00024	0.00024	0.00024
CDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	21	0.85	0.12	0.082	0.08
	$\sigma\left(\ln 10^{10} A_s\right)$	0.0052	0.0051	0.0041	0.004	0.0041
+ M_ν	$\sigma\left(n_s\right)$	0.0046	0.004	0.00061	0.00052	0.00053
	$\sigma\left(\tau_{\rm reio}\right)$	0.0022	0.0021	0.002	0.0019	0.0019
+ w_0	$\sigma\left(M_\nu\right)/[eV]$	0.13	0.051	0.014	**0.011**	**0.011**
	$\sigma\left(w_0\right)$	0.48	0.038	0.0046	0.0013	0.0013
	$\sigma\left(100*\omega_{\rm b}\right)$	0.019	0.015	0.011	0.011	0.011
CDM	$\sigma\left(\omega_{\rm cdm}\right)$	0.0012	0.00083	0.00026	0.00025	0.00023
	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	20	1.7	0.16	0.082	0.081
+ M_ν	$\sigma\left(\ln 10^{10} A_s\right)$	0.0052	0.0051	0.0044	0.0042	0.0028
	$\sigma\left(n_s\right)$	0.0047	0.0039	0.00062	0.00053	0.00052
+ w_0	$\sigma\left(\tau_{\rm reio}\right)$	0.0023	0.0022	0.002	0.002	0.00092
	$\sigma\left(M_\nu\right)/[eV]$	0.13	0.068	0.019	0.015	0.014
+ w_a	$\sigma\left(w_0\right)$	0.81	0.19	0.0095	0.0016	0.0016
	$\sigma\left(w_a\right)$	2.1	0.5	0.031	0.012	0.012

Table 28: Expected $1\,\sigma$ sensitivity of for Litebird alone and in combination with DESI, Euclid, Euclid + SKA1 IM, Euclid + SKA1 IM + τ_{reio} prior. Combinations reaching the 3-σ (5-σ) threshold towards a neutrino mass sum detection for a minimal normal hierarchy neutrino mass sum $M_\nu = 0.06$ are highlighted in green (bolded dark green), see text for discussion.

		CMB only	CMB + DESI	CMB + Euclid	+ SKA	+ $\tau_{\rm reio}$ prior
	$\sigma\left(100*\omega_{\rm b}\right)$	0.0035	0.0034	0.0026	0.0026	0.0025
	$\sigma\left(\omega_{\rm cdm}\right)$	0.00079	0.00064	0.0003	0.0003	8.6e-05
ΛCDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	0.77	0.24	0.056	0.041	0.025
	$\sigma\left(\ln 10^{10} A_s\right)$	0.022	0.017	0.0087	0.0088	0.002
+ M_ν	$\sigma\left(n_s\right)$	0.0024	0.0022	0.00058	0.00054	0.00042
	$\sigma\left(\tau_{\rm reio}\right)$	0.011	0.0096	0.0047	0.0047	0.00097
	$\sigma\left(M_\nu\right)/[eV]$	0.075	0.04	0.018	0.017	**0.0047**
	$\sigma\left(100*\omega_{\rm b}\right)$	0.0051	0.005	0.0046	0.0046	0.0044
	$\sigma\left(\omega_{\rm cdm}\right)$	0.00092	0.00078	0.00047	0.00047	0.00032
ΛCDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	0.85	0.29	0.11	0.11	0.096
	$\sigma\left(\ln 10^{10} A_s\right)$	0.024	0.018	0.0086	0.0082	0.0025
+ M_ν	$\sigma\left(n_s\right)$	0.0039	0.0034	0.00063	0.00056	0.00044
	$\sigma\left(\tau_{\rm reio}\right)$	0.012	0.0099	0.0048	0.0045	0.001
+ $N_{\rm eff}$	$\sigma\left(M_\nu\right)/[eV]$	0.081	0.043	0.017	0.016	**0.0073**
	$\sigma\left(N_{\rm eff}\right)$	0.042	0.039	0.02	0.021	0.02
	$\sigma\left(100*\omega_{\rm b}\right)$	0.0034	0.0035	0.0026	0.0026	0.0026
	$\sigma\left(\omega_{\rm cdm}\right)$	0.00086	0.00071	0.00029	0.00027	9.6e-05
CDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	10	0.83	0.083	0.038	0.025
	$\sigma\left(\ln 10^{10} A_s\right)$	0.039	0.018	0.0087	0.0081	0.002
+ M_ν	$\sigma\left(n_s\right)$	0.0024	0.0022	0.00066	0.00052	0.00042
	$\sigma\left(\tau_{\rm reio}\right)$	0.02	0.01	0.0047	0.0044	0.00097
+ w_0	$\sigma\left(M_\nu\right)/[eV]$	0.08	0.046	0.02	0.016	**0.0057**
	$\sigma\left(w_0\right)$	0.18	0.036	0.0045	0.0014	0.0013
	$\sigma\left(100*\omega_{\rm b}\right)$	0.0034	0.0035	0.0025	0.0026	0.0026
CDM	$\sigma\left(\omega_{\rm cdm}\right)$	0.00086	0.00077	0.0003	0.00028	0.00011
	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	12	1.6	0.16	0.051	0.031
+ M_ν	$\sigma\left(\ln 10^{10} A_s\right)$	0.036	0.02	0.01	0.0087	0.002
	$\sigma\left(n_s\right)$	0.0025	0.0024	0.0007	0.00059	0.00043
+ w_0	$\sigma\left(\tau_{\rm reio}\right)$	0.019	0.011	0.0055	0.0047	0.001
	$\sigma\left(M_\nu\right)/[eV]$	0.076	0.062	0.028	0.022	**0.0095**
+ w_a	$\sigma\left(w_0\right)$	0.37	0.18	0.01	0.0016	0.0017
	$\sigma\left(w_a\right)$	2.1	0.38	0.032	0.013	0.012

Table 29: Expected $1\,\sigma$ sensitivity of for CMB S4 alone and in combination with DESI, Euclid, Euclid + SKA1 IM, Euclid + SKA1 IM + τ_{reio} prior. Combinations reaching the 3-σ (5-σ) threshold towards a neutrino mass sum detection for a minimal normal hierarchy neutrino mass sum $M_\nu = 0.06$ are highlighted in green (bolded dark green), see text for discussion.

		CMB only	CMB + DESI	CMB + Euclid	+ SKA	+ $\tau_{\rm reio}$ prior
	$\sigma\left(100*\omega_{\rm b}\right)$	0.0033	0.0032	0.0024	0.0024	0.0025
	$\sigma\left(\omega_{\rm cdm}\right)$	0.00053	0.00025	0.00014	0.00013	8.2e-05
ΛCDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	0.54	0.23	0.044	0.027	0.024
	$\sigma\left(\ln 10^{10} A_s\right)$	0.004	0.0039	0.0036	0.0035	0.0018
+ M_ν	$\sigma\left(n_s\right)$	0.0018	0.0015	0.00046	0.00043	0.00041
	$\sigma\left(\tau_{\rm reio}\right)$	0.002	0.002	0.0019	0.0018	0.0009
	$\sigma\left(M_\nu\right)/[eV]$	0.042	0.019	**0.008**	**0.0071**	**0.0045**
	$\sigma\left(100*\omega_{\rm b}\right)$	0.005	0.0048	0.0043	0.0042	0.0044
	$\sigma\left(\omega_{\rm cdm}\right)$	0.00072	0.00061	0.00033	0.00034	0.00031
ΛCDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	0.63	0.29	0.098	0.092	0.091
	$\sigma\left(\ln 10^{10} A_s\right)$	0.0044	0.0043	0.0038	0.0037	0.0023
+ M_ν	$\sigma\left(n_s\right)$	0.0029	0.0026	0.00047	0.00045	0.00044
	$\sigma\left(\tau_{\rm reio}\right)$	0.0021	0.0021	0.0019	0.0019	0.00089
+ $N_{\rm eff}$	$\sigma\left(M_\nu\right)/[eV]$	0.042	0.021	**0.0096**	**0.0086**	**0.0068**
	$\sigma\left(N_{\rm eff}\right)$	0.038	0.037	0.019	0.019	0.019
	$\sigma\left(100*\omega_{\rm b}\right)$	0.0033	0.0033	0.0025	0.0024	0.0024
	$\sigma\left(\omega_{\rm cdm}\right)$	0.00061	0.00032	0.00014	0.00013	8.8e-05
CDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	6.2	0.82	0.081	0.027	0.025
	$\sigma\left(\ln 10^{10} A_s\right)$	0.0042	0.0039	0.0036	0.0035	0.0018
+ M_ν	$\sigma\left(n_s\right)$	0.002	0.0016	0.00052	0.00044	0.00041
	$\sigma\left(\tau_{\rm reio}\right)$	0.0021	0.0021	0.002	0.0019	0.00089
+ w_0	$\sigma\left(M_\nu\right)/[eV]$	0.063	0.034	**0.011**	**0.0081**	**0.0055**
	$\sigma\left(w_0\right)$	0.11	0.032	0.0045	0.0014	0.0013
	$\sigma\left(100*\omega_{\rm b}\right)$	0.0033	0.0033	0.0025	0.0025	0.0025
CDM	$\sigma\left(\omega_{\rm cdm}\right)$	0.00063	0.00039	0.00015	0.00014	0.0001
	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	8.9	1.6	0.14	0.033	0.031
+ M_ν	$\sigma\left(\ln 10^{10} A_s\right)$	0.0043	0.0039	0.0038	0.0036	0.0018
	$\sigma\left(n_s\right)$	0.002	0.0016	0.00052	0.00046	0.00043
+ w_0	$\sigma\left(\tau_{\rm reio}\right)$	0.002	0.0021	0.002	0.0019	0.00091
	$\sigma\left(M_\nu\right)/[eV]$	0.063	0.051	0.015	**0.011**	**0.0093**
+ w_a	$\sigma\left(w_0\right)$	0.31	0.18	0.0095	0.0016	0.0015
	$\sigma\left(w_a\right)$	1.8	0.45	0.031	0.012	0.012

Table 30: Expected 1 σ sensitivity of for CMB S4 + Litebird alone and in combination with DESI, Euclid, Euclid + SKA1 IM, Euclid + SKA1 IM + τ_{reio} prior. Combinations reaching the 3-σ (5-σ) threshold towards a neutrino mass sum detection for a minimal normal hierarchy neutrino mass sum $M_\nu = 0.06$ are highlighted in green (bolded dark green), see text for discussion.

		CMB only	CMB + DESI	CMB + Euclid	+ SKA	+ $\tau_{\rm reio}$ prior
	$\sigma\left(100*\omega_{\rm b}\right)$	0.0039	0.0039	0.0029	0.0029	0.0027
	$\sigma\left(\omega_{\rm cdm}\right)$	0.00053	0.00025	0.00014	0.00013	7.9e-05
ΛCDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	0.54	0.22	0.045	0.027	0.024
	$\sigma\left(\ln 10^{10} A_s\right)$	0.004	0.0039	0.0035	0.0035	0.0017
+ M_ν	$\sigma\left(n_s\right)$	0.0017	0.0015	0.00046	0.00044	0.00042
	$\sigma\left(\tau_{\rm reio}\right)$	0.0021	0.002	0.0019	0.0019	0.00089
	$\sigma\left(M_\nu\right)/[eV]$	0.042	0.02	**0.0079**	**0.0072**	**0.0044**
	$\sigma\left(100*\omega_{\rm b}\right)$	0.0059	0.0057	0.0049	0.0052	0.0049
	$\sigma\left(\omega_{\rm cdm}\right)$	0.00073	0.00067	0.00035	0.00036	0.00032
ΛCDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	0.64	0.31	0.1	0.1	0.095
	$\sigma\left(\ln 10^{10} A_s\right)$	0.0044	0.0043	0.0038	0.0039	0.0023
+ M_ν	$\sigma\left(n_s\right)$	0.003	0.0026	0.00047	0.00045	0.00043
	$\sigma\left(\tau_{\rm reio}\right)$	0.0021	0.0021	0.0019	0.0019	0.00091
+ $N_{\rm eff}$	$\sigma\left(M_\nu\right)/[eV]$	0.042	0.021	**0.0094**	**0.0089**	**0.0069**
	$\sigma\left(N_{\rm eff}\right)$	0.041	0.04	0.019	0.02	0.019
	$\sigma\left(100*\omega_{\rm b}\right)$	0.0039	0.0037	0.0028	0.0028	0.0028
	$\sigma\left(\omega_{\rm cdm}\right)$	0.00058	0.00031	0.00014	0.00013	8.5e-05
CDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	5.2	0.8	0.079	0.028	0.025
	$\sigma\left(\ln 10^{10} A_s\right)$	0.0041	0.0038	0.0035	0.0036	0.0017
+ M_ν	$\sigma\left(n_s\right)$	0.0018	0.0015	0.00051	0.00044	0.00041
	$\sigma\left(\tau_{\rm reio}\right)$	0.002	0.0021	0.0019	0.0019	0.00089
+ w_0	$\sigma\left(M_\nu\right)/[eV]$	0.058	0.034	**0.01**	**0.0081**	**0.0055**
	$\sigma\left(w_0\right)$	0.12	0.032	0.0044	0.0014	0.0013
	$\sigma\left(100*\omega_{\rm b}\right)$	0.004	0.0038	0.0029	0.0028	0.0029
CDM	$\sigma\left(\omega_{\rm cdm}\right)$	0.00061	0.00037	0.00015	0.00013	0.0001
	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	8.2	1.6	0.14	0.033	0.031
+ M_ν	$\sigma\left(\ln 10^{10} A_s\right)$	0.0042	0.0039	0.0037	0.0036	0.0018
	$\sigma\left(n_s\right)$	0.0019	0.0015	0.00051	0.00045	0.00043
+ w_0	$\sigma\left(\tau_{\rm reio}\right)$	0.002	0.0021	0.0019	0.0019	0.00091
	$\sigma\left(M_\nu\right)/[eV]$	0.059	0.049	0.015	**0.012**	**0.0092**
+ w_a	$\sigma\left(w_0\right)$	0.3	0.18	0.0089	0.0017	0.0016
	$\sigma\left(w_a\right)$	1.9	0.43	0.03	0.012	0.012

Table 31: Expected $1\,\sigma$ sensitivity of for CORE-M5 alone and in combination with DESI, Euclid, Euclid + SKA1 IM, Euclid + SKA1 IM + τ_{reio} prior. Combinations reaching the 3-σ (5-σ) threshold towards a neutrino mass sum detection for a minimal normal hierarchy neutrino mass sum $M_\nu = 0.06$ are highlighted in green (bolded dark green), see text for discussion.

		CMB only	CMB + DESI	CMB + Euclid	+ SKA	+ $\tau_{\rm reio}$ prior
	$\sigma\left(100*\omega_{\rm b}\right)$	0.0029	0.0029	0.0022	0.0021	0.0022
	$\sigma\left(\omega_{\rm cdm}\right)$	0.00049	0.00024	0.00014	0.00013	7.5e-05
ΛCDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	0.5	0.22	0.043	0.026	0.022
	$\sigma\left(\ln 10^{10} A_s\right)$	0.0041	0.0038	0.0036	0.0035	0.0017
+ M_ν	$\sigma\left(n_s\right)$	0.0017	0.0014	0.00046	0.00044	0.0004
	$\sigma\left(\tau_{\rm reio}\right)$	0.0021	0.002	0.0019	0.0019	0.00089
	$\sigma\left(M_\nu\right)/[eV]$	0.041	0.019	**0.0081**	**0.0073**	**0.0043**
	$\sigma\left(100*\omega_{\rm b}\right)$	0.0043	0.0041	0.0038	0.0038	0.0037
	$\sigma\left(\omega_{\rm cdm}\right)$	0.00062	0.00052	0.00031	0.00029	0.00027
ΛCDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	0.58	0.27	0.089	0.081	0.079
	$\sigma\left(\ln 10^{10} A_s\right)$	0.0043	0.0041	0.0036	0.0034	0.0021
+ M_ν	$\sigma\left(n_s\right)$	0.0027	0.0022	0.00047	0.00044	0.00041
	$\sigma\left(\tau_{\rm reio}\right)$	0.0022	0.002	0.0019	0.0018	0.00089
+ $N_{\rm eff}$	$\sigma\left(M_\nu\right)/[eV]$	0.04	0.02	**0.0087**	**0.008**	**0.0061**
	$\sigma\left(N_{\rm eff}\right)$	0.032	0.031	0.016	0.016	0.016
	$\sigma\left(100*\omega_{\rm b}\right)$	0.003	0.0028	0.0022	0.0022	0.0022
	$\sigma\left(\omega_{\rm cdm}\right)$	0.00053	0.00029	0.00014	0.00013	8.7e-05
CDM	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	4.3	0.79	0.081	0.026	0.023
	$\sigma\left(\ln 10^{10} A_s\right)$	0.0042	0.0038	0.0037	0.0035	0.0017
+ M_ν	$\sigma\left(n_s\right)$	0.0018	0.0014	0.00051	0.00043	0.00041
	$\sigma\left(\tau_{\rm reio}\right)$	0.0021	0.002	0.002	0.0019	0.0009
+ w_0	$\sigma\left(M_\nu\right)/[eV]$	0.055	0.034	**0.01**	**0.008**	**0.0054**
	$\sigma\left(w_0\right)$	0.11	0.03	0.0045	0.0013	0.0014
	$\sigma\left(100*\omega_{\rm b}\right)$	0.0029	0.0029	0.0023	0.0021	0.0022
CDM	$\sigma\left(\omega_{\rm cdm}\right)$	0.00059	0.00035	0.00015	0.00013	9.8e-05
	$\sigma\left(H_0\right)/[\frac{\rm km}{\rm s\,Mpc}]$	7.1	1.6	0.14	0.032	0.028
+ M_ν	$\sigma\left(\ln 10^{10} A_s\right)$	0.0042	0.0038	0.0036	0.0036	0.0017
	$\sigma\left(n_s\right)$	0.0018	0.0015	0.00049	0.00045	0.00042
+ w_0	$\sigma\left(\tau_{\rm reio}\right)$	0.002	0.002	0.0019	0.0019	0.00088
	$\sigma\left(M_\nu\right)/[eV]$	0.055	0.048	0.015	**0.012**	**0.0088**
+ w_a	$\sigma\left(w_0\right)$	0.28	0.17	0.0097	0.0015	0.0017
	$\sigma\left(w_a\right)$	1.8	0.41	0.029	0.012	0.011

Table 32: Expected $1\,\sigma$ sensitivity of for PICO alone and in combination with DESI, Euclid, Euclid + SKA1 IM, Euclid + SKA1 IM + τ_{reio} prior. Combinations reaching the 3-σ (5-σ) threshold towards a neutrino mass sum detection for a minimal normal hierarchy neutrino mass sum $M_\nu = 0.06$ are highlighted in green (bolded dark green), see text for discussion.

10

SUMMARY OF RESULTS AND CONCLUSIONS

10.1 NEUTRINOS IN COSMOLOGY

Chapter 5 and sections 8.2 and 8.3 are based on
Archidiacono, Brinckmann, Lesgourgues & Poulin [157]

The foundations of a new era in precision cosmology are based on two cornerstones: the high sensitivity of future CMB and galaxy survey experiments, and a deep understanding of the physics governing the processes of recombination and structure formation. The extreme accuracy of future data will offer the opportunity to constrain particle physics with cosmology, exceeding in many cases the precision of laboratory experiments. However, in order to exploit the new data, cosmologists will need an accurate enough theoretical model taking into account the underlying physics.

Neutrinos provide an excellent example of how the sensitivity of future cosmological surveys may lead to such an important result as the summed neutrino mass detection, even when uncertainties on the details of the cosmological model are marginalised over.

In chapter 5, we have provided a careful discussion of the physical effects induced by massive neutrinos and their impact on cosmological observables, as they will appear in the data analysis of the next generation of cosmological experiments. We have shown how the unique nature of light neutrinos, being relativistic until very late times and behaving as a matter component after the non relativistic transition, makes possible to identify different signatures at different epochs of the cosmic history. Therefore, the correlation between the summed neutrino mass and the other cosmological parameters changes, depending on the redshift range probed by the various data sets.

The results relating to this study were presented in sections 8.2 and 8.3 and focused on the sensitivity of future CMB-CORE and BAO-DESI experiments to the summed neutrino mass are consistent with the literature (see Refs. [198, 199]). Moreover, the results of our forecasts including a Euclid-like survey prove the importance of cosmic shear and galaxy clustering as complementary probes. We pointed out that the results of our Euclid cosmic shear + galaxy correlation forecasts depend very much on the choice of the theoretical error introduced to account for the systematics coming from the deep non-linear regime. Nevertheless, they are again compatible with previously published results. For instance, Ref. [183] found $\sigma(M_\nu) = 11$ meV for Planck + Euclid cosmic shear / galaxy correlation, but with a different treatment of the uncertainty on non linear corrections. Ref. [174] found a larger error, close to 20 meV, but for Planck + Euclid cosmic shear or Planck + Euclid galaxy correlation, not trying to combine the two LSS probes together and without CORE data. Ref. [198] found $\sigma(M_\nu) = 12$ meV for Planck + CMB-Stage-IV + BAO-DESI + 21cm-HERA, identical to our estimate for CORE + Euclid + 21cm-$\tau_{\rm reio}$-prior.

Anyhow, the main goal of this study was not to present a new set of forecasts, but to discuss the details of physical effects and parameter degeneracies involving neutrino masses. In particular, we clarified the reason for which an unexpected degeneracy between the neutrino mass sum and the optical depth at reionization will appear in the analysis of future high precision galaxy surveys, as already pointed out e.g. in [198, 199]. We showed that this degeneracy is not present in a CMB-only analysis, because the neutrino mass effects on CMB lensing can be compensated by playing with h and $\omega_{\rm cdm}$ in a better way than by adjusting $(A_s, \tau_{\rm reio})$. However, the former degeneracy is lifted once BAO and LSS low redshift measurements are taken into account. Moreover, we demonstrated that the LSS data introduce a strong correlation between M_ν and A_s, which finally leads to a clear $(M_\nu, \tau_{\rm reio})$ degeneracy in the combined CMB+LSS analysis.

These conclusions clarify why further independent measurements of the optical depth will benefit a neutrino mass determination, as previously noticed by the authors of [198]. For instance, the results from the HERA or SKA 21cm surveys will provide an independent constraint on $\tau_{\rm reio}$, thus breaking this degeneracy. Our results indicate that this could reduce the error on M_ν with respect to the CMB+LSS case, leading to a robust detection of the summed neutrino mass at more than 5 σ for CORE+Euclid+HERA or SKA. In principle, it would be possible to do even better if H_0 could be measured in an independent and robust way with an error below $\sigma(H_0) \sim 0.17$km/s/Mpc.

In conclusion, the remarkable complementarity of future different cosmological data will lead to extremely accurate constraints on the neutrino mass sum and, possibly, on other neutrino properties, answering some of the still open questions of modern physics.

10.2 FUTURE COSMIC MICROWAVE BACKGROUND EXPERIMENTS

Section 10.2 is based on Di Valentino, Brinckmann, Gerbino, Poulin et al. [50]

In sections 7.1, 8.1 and 9.1 to 9.3, we forecasted the constraints on several cosmological parameters that can be achieved by the CORE-M5 satellite proposal. Table 33 provides a summary of our main results. Assuming ΛCDM, the improvement with respect to Planck is extremely significant: CORE-M5 can simultaneously improve constraints on key parameters by a factor ~ 8 (σ_8), ~ 5.5 (H_0, $\Omega_{cdm}h^2$), 4.5 ($\Omega_b h^2$, τ), and 3 (n_s).

Some of the parameters such as σ_8, H_0, and $\Omega_b h^2$ can be measured or derived independently by galaxy surveys or luminosity distance measurements. Future comparisons with the CORE-M5 results will therefore provide a crucial test for cosmology and the ΛCDM scenario and its extensions. The interest of such measurements by several means is exemplified by the current tensions between the Planck dataset and the local determination of the Hubble constant from [85] or measurements of weak lensing cosmic shear from surveys as CFHTLenS and KiDS-450 [291, 292]. These tensions may reveal either previously unknown systematic effects, or new physics. While these current tensions will likely be resolved by the time CORE flies, the large improvement brought by CORE on so many parameters will surely bring new opportunities for revealing tensions with whatever precision datasets will be available by then. These are opportunites for fundamental breakthoughs.

In sections 8.1 and 9.1 to 9.3, we have considered several possible extensions to the basic six parameters ΛCDM model. The forecasted constraints on these extra parameters are summarized in the second section of Table 33. As we can see, also on these extensions CORE-M5 can provide significantly more stringent constraints than the current ones, with a factor of n

Parameter	Description	Planck 2015 + Lensing	CORE-M5 expected uncertainties
ΛCDM			
$\Omega_b h^2$	Baryon Density	$\Omega_b h^2 = 0.02226 \pm 0.00016$ (68 % CL) [2]	$\sigma(\Omega_b h^2) = \mathbf{0.000037}$ {4.3}
$\Omega_c h^2$	Cold Dark Matter Density	$\Omega_c h^2 = 0.1193 \pm 0.0014$ (68 % CL) [2]	$\sigma(\Omega_c h^2) = \mathbf{0.00026}$ {5.4}
n_s	Scalar Spectral Index	$n_s = 0.9653 \pm 0.0048$ (68 % CL) [2]	$\sigma(n_s) = \mathbf{0.0014}$ {3.4}
τ	Reionization Optical Depth	0.063 ± 0.014 (68 % CL) [2]	$\sigma(\tau) = \mathbf{0.002}$ {7.0}
H_0 [km/s/Mpc]	Hubble Constant	$H_0 = 67.51 \pm 0.64$ (68 % CL) [2]	$\sigma(H_0) = \mathbf{0.11}$ {5.8}
σ_8	r.m.s. mass fluctuations	$\sigma_8 = 0.8150 \pm 0.0087$ (68 % CL) [2]	$\sigma(\sigma_8) = \mathbf{0.0011}$ {7.9}
Extensions			
M_ν	Total Neutrino Mass	$M_\nu < 0.315$ eV (68 % CL) [2]	$\sigma(M_\nu) = \mathbf{0.043}$ eV {7.3}
$N_{\rm eff}$	Relativistic Degrees of Freedom	$N_{\rm eff} = 2.94 \pm 0.20$ (68 % CL) [2]	$\sigma(N_{\rm eff}) = \mathbf{0.040}$ {5.0}
(m_s^{eff}, N_s)	Sterile Neutrino Parameters	$(m_s^{eff} < 0.33eV, N_s < 3.24)$ (68 % CL) [2]	$\sigma(m_s^{eff}, N_s) = (\mathbf{0.037}$ eV$, \mathbf{0.053})$ {8.9, 4.5}

Table 33: Current limits from Planck 2015 and forecasted CORE-M5 uncertainties. The first 6 rows assume a ΛCDM scenario and the following rows give the constraints on parameter extensions, always while also varying the neutrino mass sum. In the fourth column, numbers in curly brackets {...} give the improvement in the parameter constraint when moving from Planck 2015 to CORE-M5, defined as the ratio of the uncertainties $\sigma^{Planck}/\sigma^{CORE}$.

improvement that ranges from 4 up to more than 6, clearly opening the window to new tests or discoveries for physics beyond the standard model.

In particular, we found that:

- CORE-M5 alone could detect neutrino masses with an uncertainty of $\sigma(M_\nu) = 0.043$ eV, enough to rule out the inverted mass hierarchy at more than 95% c.l.. When combined with future galaxy clustering data as expected from surveys as DESI or EUCLID, CORE-M5 will provide a guaranteed discovery for a neutrino mass. Other cosmological information from CORE-M5, as cluster number counts (see the CORE cluster science paper [323]) could further reduce these uncertainties.

- CORE-M5 could also provide extremely stringent constraints on the neutrino effective number $N_{\rm eff}$ with $\sigma(N_{\rm eff}) = 0.040$. This uncertainty, that can be further reduced by combining the CORE-M5 data with cluster number counts data from CORE-M5 itself and/or complementary galaxy surveys, will test the presence of extra light particles at recombination and the process of neutrino decoupling from the primordial plasma at redshift $z \sim 10^9$.

- Large angular scale polarization will also be measured by CORE, providing new constraints on the reionization process. It is here worthwhile to note that the ability of CORE-M5 to measure polarization over a wide range of angular scales will provide a crucial test for the cosmological scenario. The constraints on the optical depth τ from large angular scales, for example, can be only validated by a measurement of small angular scale polarization with results consistent with the overall ΛCDM scenario.

It is also interesting to summarize the constraints from different experimental configurations and to compare them. We do this in table 34, where we report the ratio of the 1-σ forecasted error of a certain experimental configuration over the expected 1 σ error from the proposed CORE-M5 setup. For generality, we also compare the constraints with those expected from the JAXA Litebird proposal [204] that at the time of this work was in conceptual design phase (called ISAS Phase-A1). Litebird presents a significantly different experimental design with respect to the CORE configurations studied in this thesis, with, for example, a smaller primary mirror of 60 cm.

Parameter	CORE-M5 vs Litebird	CORE-M5 vs LiteCORE-80	CORE-M5 vs LiteCORE-120	CORE-M5 vs COrE+
ΛCDM				
$\Omega_b h^2$	3.5	1.4	1.1	0.9
$\Omega_c h^2$	2.3	1.3	1.2	1.0
$100\theta_{MC}$	5.8	1.3	1.1	0.9
τ	1.0	1.0	1.0	1.0
n_s	2.6	1.1	1.1	1.0
$ln(10^{10}A_s)$	1.2	1.1	1.0	1.0
H_0[km/s/Mpc]	3.0	1.3	1.1	0.9
σ_8	2.5	1.3	1.1	0.9
Extensions				
$N_{\rm eff}$	4.8	1.5	1.1	0.9
M_ν	1.6	1.1	1.0	1.0

Table 34: Improvements from CORE-M5 on cosmological parameters with respect to several proposed configuration defined as the ratio of the forecasted 1 σ constraints, $\sigma/\sigma_{CORE-M5}$.

As we can see from the results in table 34, any CORE configuration is expected to constrain cosmological parameters with an improvement that ranges from a factor 2 to 5 respect to Litebird. CORE-M5, for example, will constrain the neutrino effective number with a precision about 5 times better than Litebird. It is clear from the results presented in the Table that CORE will have the possibility to probe new physics that will not be accessible by Litebird alone. However, constraints on the reionization optical depth will be comparable, since the imprint of reionization is mainly on large scale polarization that can be equally measured by Litebird and CORE.

Also from table 34, we see that CORE-M5 could produce constraints that are up to 50% better than those expected from the cheaper LiteCORE-80 configuration. A significantly higher precision is indeed expected on key parameters as the baryon abundance, the Hubble constant, and the neutrino effective number. On the other hand, the differences between CORE-M5 and LiteCORE-120 and COrE+ are expected to be of the order of $\sim 10\%$.

From one side we can then consider the forecasts presented here for CORE-M5 as conservative: if the experimental sensitivity will be for some reason degraded to LiteCORE-120 we expect no significant variations in the constraints presented in this thesis. On the other hand, the more expensive COrE+ configuration would only slightly improve the main parameter constraints and would not present a decisive improvement in the specific scientific aspect of parameters recovery and model testing. Indeed the scientific driver for higher angular resolution is not the improvement in parameters accuracy.

To conclude, we have presented in sections 7.1, 8.1 and 9.1 to 9.3 a large number of forecasts on cosmological parameters for the CORE-M5 proposed mission. The expected improved constraints, presented in Table 33 clearly calls for of a next CMB satellite mission as CORE. CORE-M5 can probe new physics with unprecedented precision. We have compared the constraints with different experimental configurations and found that the expected constraints are stable under a degradation of the experimental configuration to LiteCORE-120 that has a significantly smaller number of detectors. Assuming the ΛCDM cosmological scenario, we also found that the CORE-M5 setup can produce constraints that are almost identical (at worst a $\sim 10\%$ degradation) to the ones achievable by the larger aperture COrE+ configuration.

10.3 IMPROVED LARGE-SCALE STRUCTURE TREATMENT

The production of robust and accurate forecasts on the sensitivity to cosmological parameters of future surveys like Euclid and the Square Kilometer Array is an important task in the context of the preparation phase of these surveys. A major difficulty comes from the theoretical uncertainties arising on mildly non-linear scale, induced by the complexity of physical processes at play, such as non-linear clustering and baryonic feedback. In chapter 3, we take these uncertainties into account with a method that has lots of similarities with previous attempts in Refs. [118, 119]. We discuss a compromise between the fully uncorrelated theoretical error of Ref. [118], which is arguably too conservative, and the approach of Ref. [119] based on an exponentially-decaying error correlation, which is too expensive numerically for MCMC forecasts. Our method is at the same time realistic, firmly rooted on physical results from various astrophysical studies, and computationally tractable. It relies on an ansatz for the error amplitude and correlation length on non-linear scales up to $k_{\rm max} = 10\ h\ {\rm Mpc}^{-1}$.

For the first time, we present, in sections 8.4 and 9.4, forecasts based on MCMC simulations for four cosmological scenarios, two different modelings of the non-linear theoretical error (one conservative and one realistic) and 14 experimental configurations and combinations: overall more than 140 MCMC simulations. Additionally, in section 9.5, we present a grid of 120 MCMC simulations for the same four cosmological models, considering six different experimental CMB configurations: we consider the current Planck sensitivity as a baseline and discuss the impact of four mission or survey projects: LiteBird, CORE, CMB-S4 and PICO, including a combination of CMB-S4 and Litebird. Finally, for LSS surveys, we restrict ourselves to future BAO, galaxy redshift surveys, cosmic shear surveys and intensity mapping surveys from DESI, Euclid and SKA, where the latter two are a subset of the array of experimental configurations and non-linear theoretical error modelling used in sections 8.4 and 9.4. We also discuss the impact of a τ prior, motivated by future 21cm measurements. This work is therefore the most exhaustive analysis of this kind released so far, both for Euclid and the Square Kilometer Array, and the combination of them.

It also paves the way towards a realistic implementation of the non-linear theoretical errors, going beyond the usual cut-off scale method that is i) too pessimistic on non-linear scales because it cannot exploit all the power of the survey, ii) at the same time too optimistic because it does not account for the theoretical uncertainties on the power spectrum below the mildly non-linear cut-off scale. Three observational probes have been considered: galaxy clustering and weak lensing (for Euclid, SKA1 and SKA2), and HI intensity mapping (for SKA1) at low redshift (probing the already reionized universe). Compared to the usual Fisher approach, a bayesian MCMC method allows to probe non-Gaussian posteriors and is immune to the sometimes critical numerical stability issues that are linked to the choice of step size for numerical derivatives.

Our main findings arise from the impact of the non-linear theoretical uncertainty on the power spectrum, and are summarized thereafter:

1. Despite the fact that the theoretical error removes a lot of information from large wavenumbers, we find that for galaxy clustering, when considering scales up to $k_{\rm max} = 10\ h\ {\rm Mpc}^{-1}$ with the theoretical error, we increase the lever arm to constrain n_s and improve the sensitivity to this parameter by about a factor two.

2. Because of the degeneracy breaking, the former point also leads to an improvement for H_0 (by about 50% for Planck+SKA1 and 25% for Planck+Euclid) and for the total

neutrino mass M_ν (by about 70% for Planck+SKA1 and 50% for Planck+Euclid). We also observe a factor two improvement on $\omega_{\rm cdm}$ for SKA1-GC.

3. The sensitivity to the derived parameters σ_8, $\Omega_{\rm m}$ and Ω_Λ significantly improves for nearly all the experimental configurations.

4. Concerning the extended cosmological models, there is a remarkable improvement for the constraints on the dark energy equation of state parameters w_0 and w_a.

5. With either Planck+Euclid or Planck+SKA2, the neutrino mass could be constrained to $\sigma(M_\nu) = 0.012\,\mathrm{eV}$ with the baseline model and assuming a realistic theoretical error, leading to a 5σ-detection.

6. The degeneracy of several parameters (including M_ν) with $N_{\rm eff}$, w_0 and $w_{\rm a}$ is not as severe as previous thought. CS+GC data with the realistic theoretical error assumption can break the degeneracies. The final sensitivity to e.g. the total neutrino mass is stable at least against these simple extensions of the minimal cosmological model.

It is worth noticing, that any source of error on the power spectrum can be incorporated in our non-linear treatment, eventually at the price of modifying the overall shape of the theoretical error and the correlation length, if this source is dominant. The assumed error on the spectrum has strong impact on parameter sensitivities. Therefore, we recommend to use an efficient implementation of the theoretical error, as the one illustrated here, in the analysis of forthcoming surveys.

We have studied the constraining power of each individual probes and the advantage of combining them. In particular, we considered the combination of the Euclid survey with the SKA HI intensity mapping survey. The sensitivity to various parameters is reported in Table 21 and 26. Our main conclusions are:

1. Overall, SKA1-IM is more constraining than SKA1 (CS + GC), whereas Euclid and SKA2 perform similarly and better than SKA1-IM.

2. The forecasts on w_0 and w_a improve by up to a factor five when including non-linear scales with a realistic modelling of uncertainties and when combining Euclid with SKA1-IM. In the most constraining scenario we obtain $\sigma(w_0) = 0.0016$ and $\sigma(w_a) = 0.013$, i.e. a factor forty of improvement compared to Planck alone.

These results emphasize the importance of combining Euclid with a survey extending the information down to very low redshift, such as SKA1 intensity mapping. This combination, together with the modelling of non-linear uncertainties, could make the difference between a 4-σ strong indication and a 5-σ detection of the total neutrino mass when the DE is modelled with two free parameters. It is even more crucial for constraining the DE parameters themselves.

Our method to deal with the non-linear uncertainties is only a first proxy that could be made more accurate and updated with the results of future N-body simulations, with a better understanding of the baryonic feedback, or with analytical progress on any other source of error. The suggested implementation can nevertheless be considered as a realistic target, given that it is based on the current understanding of those processes and on conservative assumptions about the expected precision of future N-body simulations. For intensity mapping, our analysis could be refined by using a more precise foreground modelling.

Realistic forecasts could be produced for other cosmological scenarios, e.g., specific dark-energy/modified-gravity parameterizations, or for assessing the accuracy of bayesian selection of inflationary models, see Refs. [324, 325].

Obviously, extending the parameter space leads to more pronounced parameter degeneracies and call for additional data. Extra constraining power might come from independent probes (e.g., the 21cm intensity mapping from reionization provided by SKA), and also from the cross-correlations between different probes (e.g., between galaxy shear and clustering), left for further study.

Our main findings for the MCMC grid with an array of CMB experiments are, for a minimal 7-parameter ΛCDM + M_ν model:

1. If we are able to obtain a precise independent measurement of the optical depth to reionization, Planck in combination with Euclid and SKA1 intensity mapping would be able to reach a 5-σ detection of the neutrino mass sum, even for a minimal normal hierarchy neutrino mass sum of $M_\nu \approx 0.06$.

2. LiteBird does not require an independent $\tau_{\rm reio}$ measurement, the combination with Euclid and SKA1 is enough to reach the 5-σ level.

3. CORE, PICO, or a combination of LiteBird and CMB-S4: reach the 3-σ level when only adding BAO data from DESI; safely exceed the 5-σ level when data from Euclid is added instead.

Additionally, we find that the expectation of a neutrino mass sum detection from cosmology is robust to the choice of model, i.e. we find in the case of the extended model that degrades the neutrino mass sum sensitivity the most (of those under consideration, which here is dynamical dark energy):

1. LiteBird + Euclid as well as Planck + Euclid + SKA1 intensity mapping reach the 3-σ level towards a neutrino mass sum detection.

2. To reach the 4-σ level, we need at least Euclid plus CORE, PICO, or a combination of LiteBird and CMB-S4.

3. We would expect to obtain at least a 5-σ detection of a non-zero neutrino mass sum if we combine a highly sensitive CMB experiment (CORE, PICO, or a combination of LiteBird and CMB-S4) with Euclid and SKA1 intensity mapping.

The conclusion of this thesis is, therefore, *we are all but guaranteed a non-zero neutrino mass sum detection in the next decade*, or at least have a very strong indication that physics beyond the Standard Model of Particle Physics or Cosmology needs to be considered.

BIBLIOGRAPHY

[1] S. Vagnozzi, T. Brinckmann, M. Archidiacono, K. Freese, M. Gerbino, J. Lesgourgues et al., *Bias due to neutrinos must not uncorrect'd go*, 1807.04672.

[2] Planck collaboration, P. A. R. Ade et al., *Planck 2015 results. XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13, [1502.01589].

[3] EUCLID collaboration, R. Laureijs et al., *Euclid Definition Study Report*, 1110.3193.

[4] Euclid Theory Working Group collaboration, L. Amendola et al., *Cosmology and fundamental physics with the Euclid satellite*, *Living Rev. Rel.* **16** (2013) 6, [1206.1225].

[5] SKA Cosmology SWG collaboration, R. Maartens, F. B. Abdalla, M. Jarvis and M. G. Santos, *Overview of Cosmology with the SKA*, *PoS* **AASKA14** (2015) 016, [1501.04076].

[6] M. G. Santos et al., *Cosmology with a SKA HI intensity mapping survey*, 1501.03989.

[7] EoR/CD-SWG, Cosmology-SWG collaboration, J. Pritchard et al., *Cosmology from EoR/Cosmic Dawn with the SKA*, *PoS* **AASKA14** (2015) 012, [1501.04291].

[8] M. C. Gonzalez-Garcia, M. Maltoni and T. Schwetz, *Global Analyses of Neutrino Oscillation Experiments*, *Nucl. Phys.* **B908** (2016) 199–217, [1512.06856].

[9] F. Capozzi, E. Lisi, A. Marrone, D. Montanino and A. Palazzo, *Neutrino masses and mixings: Status of known and unknown 3ν parameters*, *Nucl. Phys.* **B908** (2016) 218–234, [1601.07777].

[10] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler and T. Schwetz, *Updated fit to three neutrino mixing: exploring the accelerator-reactor complementarity*, *JHEP* **01** (2017) 087, [1611.01514].

[11] F. Capozzi, E. Di Valentino, E. Lisi, A. Marrone, A. Melchiorri and A. Palazzo, *Global constraints on absolute neutrino masses and their ordering*, *Phys. Rev.* **D95** (2017) 096014, [1703.04471].

[12] P. F. de Salas, D. V. Forero, C. A. Ternes, M. Tortola and J. W. F. Valle, *Status of neutrino oscillations 2018: 3σ hint for normal mass ordering and improved CP sensitivity*, *Phys. Lett.* **B782** (2018) 633–640, [1708.01186].

[13] N. Palanque-Delabrouille et al., *Neutrino masses and cosmology with Lyman-alpha forest power spectrum*, *JCAP* **1511** (2015) 011, [1506.05976].

[14] E. Di Valentino, E. Giusarma, M. Lattanzi, O. Mena, A. Melchiorri and J. Silk, *Cosmological Axion and neutrino mass constraints from Planck 2015 temperature and polarization data*, *Phys. Lett.* **B752** (2016) 182–185, [1507.08665].

[15] A. J. Cuesta, V. Niro and L. Verde, *Neutrino mass limits: robust information from the power spectrum of galaxy surveys*, *Phys. Dark Univ.* **13** (2016) 77–86, [1511.05983].

[16] Q.-G. Huang, K. Wang and S. Wang, *Constraints on the neutrino mass and mass hierarchy from cosmological observations*, *Eur. Phys. J.* **C76** (2016) 489, [1512.05899].

[17] M. Moresco, R. Jimenez, L. Verde, A. Cimatti, L. Pozzetti, C. Maraston et al., *Constraining the time evolution of dark energy, curvature and neutrino properties with cosmic chronometers*, *JCAP* **1612** (2016) 039, [1604.00183].

[18] E. Giusarma, M. Gerbino, O. Mena, S. Vagnozzi, S. Ho and K. Freese, *Improvement of cosmological neutrino mass bounds*, *Phys. Rev.* **D94** (2016) 083522, [1605.04320].

[19] BOSS collaboration, S. Alam et al., *The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample*, *Mon. Not. Roy. Astron. Soc.* **470** (2017) 2617–2652, [1607.03155].

[20] S. Vagnozzi, E. Giusarma, O. Mena, K. Freese, M. Gerbino, S. Ho et al., *Unveiling ν secrets with cosmological data: neutrino masses and mass hierarchy*, *Phys. Rev.* **D96** (2017) 123503, [1701.08172].

[21] F. Couchot, S. Henrot-Versillé, O. Perdereau, S. Plaszczynski, B. Rouillé d'Orfeuil, M. Spinelli et al., *Cosmological constraints on the neutrino mass including systematic uncertainties*, *Astron. Astrophys.* **606** (2017) A104, [1703.10829].

[22] A. Caldwell, A. Merle, O. Schulz and M. Totzauer, *Global Bayesian analysis of neutrino mass data*, *Phys. Rev.* **D96** (2017) 073001, [1705.01945].

[23] C. Doux, M. Penna-Lima, S. D. P. Vitenti, J. Tréguer, E. Aubourg and K. Ganga, *Cosmological constraints from a joint analysis of cosmic microwave background and large-scale structure*, 1706.04583.

[24] S. Wang, Y.-F. Wang and D.-M. Xia, *Constraints on the sum of neutrino masses using cosmological data including the latest extended Baryon Oscillation Spectroscopic Survey DR14 quasar sample*, *Chin. Phys.* **C42** (2018) 065103, [1707.00588].

[25] L. Chen, Q.-G. Huang and K. Wang, *New cosmological constraints with extended-Baryon Oscillation Spectroscopic Survey DR14 quasar sample*, *Eur. Phys. J.* **C77** (2017) 762, [1707.02742].

[26] A. Upadhye, *Neutrino mass and dark energy constraints from redshift-space distortions*, 1707.09354.

[27] L. Salvati, M. Douspis and N. Aghanim, *Constraints from thermal Sunyaev-Zeldovich cluster counts and power spectrum combined with CMB*, *Astron. Astrophys.* **614** (2018) A13, [1708.00697].

[28] R. C. Nunes and A. Bonilla, *Probing the properties of relic neutrinos using the cosmic microwave background, the Hubble Space Telescope and galaxy clusters*, *Mon. Not. Roy. Astron. Soc.* **473** (2018) 4404–4409, [1710.10264].

[29] M. Zennaro, J. Bel, J. Dossett, C. Carbone and L. Guzzo, *Cosmological constraints from galaxy clustering in the presence of massive neutrinos*, *Mon. Not. Roy. Astron. Soc.* **477** (2018) 491–506, [1712.02886].

[30] L.-F. Wang, X.-N. Zhang, J.-F. Zhang and X. Zhang, *Impacts of gravitational-wave standard siren observation of the Einstein Telescope on weighing neutrinos in cosmology*, *Phys. Lett.* **B782** (2018) 87–93, [1802.04720].

[31] S. R. Choudhury and S. Choubey, *Updated Bounds on Sum of Neutrino Masses in Various Cosmological Scenarios*, 1806.10832.

[32] S. R. Choudhury and A. Naskar, *Bounds on Sum of Neutrino Masses in a 12 Parameter Extended Scenario with Non-Phantom Dynamical Dark Energy ($w(z) \geq -1$)*, 1807.02860.

[33] S. Hannestad and T. Schwetz, *Cosmology and the neutrino mass ordering*, *JCAP* **1611** (2016) 035, [1606.04691].

[34] L. Xu and Q.-G. Huang, *Detecting the Neutrinos Mass Hierarchy from Cosmological Data*, *Sci. China Phys. Mech. Astron.* **61** (2018) 039521, [1611.05178].

[35] M. Gerbino, M. Lattanzi, O. Mena and K. Freese, *A novel approach to quantifying the sensitivity of current and future cosmological datasets to the neutrino mass ordering through Bayesian hierarchical modeling*, *Phys. Lett.* **B775** (2017) 239–250, [1611.07847].

[36] F. Simpson, R. Jimenez, C. Pena-Garay and L. Verde, *Strong Bayesian Evidence for the Normal Neutrino Hierarchy*, *JCAP* **1706** (2017) 029, [1703.03425].

[37] T. Schwetz, K. Freese, M. Gerbino, E. Giusarma, S. Hannestad, M. Lattanzi et al., *Comment on "Strong Evidence for the Normal Neutrino Hierarchy"*, 1703.04585.

[38] S. Hannestad and T. Tram, *Optimal prior for Bayesian inference in a constrained parameter space*, 1710.08899.

[39] A. J. Long, M. Raveri, W. Hu and S. Dodelson, *Neutrino Mass Priors for Cosmology from Random Matrices*, *Phys. Rev.* **D97** (2018) 043510, [1711.08434].

[40] S. Gariazzo, M. Archidiacono, P. F. de Salas, O. Mena, C. A. Ternes and M. Tórtola, *Neutrino masses and their ordering: Global Data, Priors and Models*, *JCAP* **1803** (2018) 011, [1801.04946].

[41] A. F. Heavens and E. Sellentin, *Objective Bayesian analysis of neutrino masses and hierarchy*, *JCAP* **1804** (2018) 047, [1802.09450].

[42] W. Handley and M. Millea, *Maximum entropy priors with derived parameters in a specified distribution*, 1804.08143.

[43] P. F. de Salas, S. Gariazzo, O. Mena, C. A. Ternes and M. Tórtola, *Neutrino Mass Ordering in 2018: Global Status*, 1806.11051.

[44] S. Hannestad, *Neutrino masses and the dark energy equation of state - Relaxing the cosmological neutrino mass bound*, *Phys. Rev. Lett.* **95** (2005) 221301, [astro-ph/0505551].

[45] S. Joudaki, *Constraints on Neutrino Mass and Light Degrees of Freedom in Extended Cosmological Parameter Spaces*, *Phys. Rev.* **D87** (2013) 083523, [1202.0005].

[46] W. Yang, R. C. Nunes, S. Pan and D. F. Mota, *Effects of neutrino mass hierarchies on dynamical dark energy models*, *Phys. Rev.* **D95** (2017) 103522, [1703.02556].

[47] C. S. Lorenz, E. Calabrese and D. Alonso, *Distinguishing between Neutrinos and time-varying Dark Energy through Cosmic Time*, *Phys. Rev.* **D96** (2017) 043510, [1706.00730].

[48] W. Sutherland, *The CMB neutrino mass / vacuum energy degeneracy: a simple derivation of the degeneracy slopes*, *Mon. Not. Roy. Astron. Soc.* **477** (2018) 1913–1920, [1803.02298].

[49] M. Sahlén, *Cluster-Void Degeneracy Breaking: Neutrino Properties and Dark Energy*, 1807.02470.

[50] CORE collaboration, E. Di Valentino et al., *Exploring cosmic origins with CORE: Cosmological parameters*, *JCAP* **1804** (2018) 017, [1612.00021].

[51] COBE collaboration, G. F. Smoot et al., *Structure in the COBE differential microwave radiometer first year maps*, *Astrophys. J.* **396** (1992) L1–L5.

[52] A. D. Miller, R. Caldwell, M. J. Devlin, W. B. Dorwart, T. Herbig, M. R. Nolta et al., *A measurement of the angular power spectrum of the cmb from l = 100 to 400*, *Astrophys. J.* **524** (1999) L1–L4, [astro-ph/9906421].

[53] N. W. Halverson et al., *DASI first results: A Measurement of the cosmic microwave background angular power spectrum*, *Astrophys. J.* **568** (2002) 38–45, [astro-ph/0104489].

[54] ACBAR collaboration, C.-l. Kuo et al., *High resolution observations of the CMB power spectrum with ACBAR*, *Astrophys. J.* **600** (2004) 32–51, [astro-ph/0212289].

[55] Boomerang collaboration, P. de Bernardis et al., *A Flat universe from high resolution maps of the cosmic microwave background radiation*, *Nature* **404** (2000) 955–959, [astro-ph/0004404].

[56] Boomerang collaboration, C. B. Netterfield et al., *A measurement by Boomerang of multiple peaks in the angular power spectrum of the cosmic microwave background*, *Astrophys. J.* **571** (2002) 604–614, [astro-ph/0104460].

[57] S. Hanany et al., *MAXIMA-1: A Measurement of the cosmic microwave background anisotropy on angular scales of 10 arcminutes to 5 degrees*, *Astrophys. J.* **545** (2000) L5, [astro-ph/0005123].

[58] Archeops collaboration, A. Benoit et al., *Cosmological constraints from Archeops*, *Astron. Astrophys.* **399** (2003) L25–L30, [astro-ph/0210306].

[59] WMAP collaboration, D. N. Spergel et al., *First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Determination of cosmological parameters*, *Astrophys. J. Suppl.* **148** (2003) 175–194, [astro-ph/0302209].

[60] WMAP collaboration, G. Hinshaw et al., *Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results*, *Astrophys. J. Suppl.* **208** (2013) 19, [1212.5226].

[61] Planck collaboration, P. A. R. Ade et al., *Planck 2013 results. XVI. Cosmological parameters*, *Astron. Astrophys.* **571** (2014) A16, [1303.5076].

[62] R. J. Cooke, M. Pettini, K. M. Nollett and R. Jorgenson, *The primordial deuterium abundance of the most metal-poor damped Lyα system*, *Astrophys. J.* **830** (2016) 148, [1607.03900].

[63] K. M. Nollett and G. P. Holder, *An analysis of constraints on relativistic species from primordial nucleosynthesis and the cosmic microwave background*, 1112.2683.

[64] E. Di Valentino, C. Gustavino, J. Lesgourgues, G. Mangano, A. Melchiorri, G. Miele et al., *Probing nuclear rates with Planck and BICEP2*, *Phys. Rev.* **D90** (2014) 023543, [1404.7848].

[65] A. Lewis and A. Challinor, *Weak gravitational lensing of the cmb*, *Phys. Rept.* **429** (2006) 1–65, [astro-ph/0601594].

[66] J. Lesgourgues, G. Mangano, G. Miele and S. Pastor, *Neutrino cosmology*. Cambridge Univ. Press, Cambridge, 2013.

[67] J. Lesgourgues and S. Pastor, *Massive neutrinos and cosmology*, *Phys. Rept.* **429** (2006) 307–379, [astro-ph/0603494].

[68] F. Capozzi, G. L. Fogli, E. Lisi, A. Marrone, D. Montanino and A. Palazzo, *Status of three-neutrino oscillation parameters, circa 2013*, *Phys. Rev.* **D89** (2014) 093018, [1312.2878].

[69] M. C. Gonzalez-Garcia, M. Maltoni and T. Schwetz, *Updated fit to three neutrino mixing: status of leptonic CP violation*, *JHEP* **11** (2014) 052, [1409.5439].

[70] D. V. Forero, M. Tortola and J. W. F. Valle, *Neutrino oscillations refitted*, *Phys. Rev.* **D90** (2014) 093006, [1405.7540].

[71] Planck collaboration, N. Aghanim et al., *Planck intermediate results. XLVI. Reduction of large-scale systematic effects in HFI polarization maps and estimation of the reionization optical depth*, *Astron. Astrophys.* **596** (2016) A107, [1605.02985].

[72] M. Gerbino, M. Lattanzi and A. Melchiorri, *ν generation: Present and future constraints on neutrino masses from global analysis of cosmology and laboratory experiments*, *Phys. Rev.* **D93** (2016) 033001, [1507.08614].

[73] E. Di Valentino, E. Giusarma, O. Mena, A. Melchiorri and J. Silk, *Cosmological limits on neutrino unknowns versus low redshift priors*, *Phys. Rev.* **D93** (2016) 083527, [1511.00975].

[74] R. Jimenez, T. Kitching, C. Pena-Garay and L. Verde, *Can we measure the neutrino mass hierarchy in the sky?*, *JCAP* **1005** (2010) 035, [1003.5918].

[75] J. F. Beacom, N. F. Bell and S. Dodelson, *Neutrinoless universe*, *Phys. Rev. Lett.* **93** (2004) 121302, [astro-ph/0404585].

[76] S. Hannestad, A. Mirizzi, G. G. Raffelt and Y. Y. Y. Wong, *Cosmological constraints on neutrino plus axion hot dark matter*, *JCAP* **0708** (2007) 015, [0706.4198].

[77] M. Archidiacono, S. Hannestad, A. Mirizzi, G. Raffelt and Y. Y. Y. Wong, *Axion hot dark matter bounds after Planck*, *JCAP* **1310** (2013) 020, [1307.0615].

[78] K. N. Abazajian et al., *Light Sterile Neutrinos: A White Paper*, 1204.5379.

[79] M. Archidiacono, N. Fornengo, C. Giunti, S. Hannestad and A. Melchiorri, *Sterile neutrinos: Cosmology versus short-baseline experiments*, *Phys. Rev.* **D87** (2013) 125034, [1302.6720].

[80] E. Di Valentino, A. Melchiorri and O. Mena, *Dark radiation sterile neutrino candidates after Planck data*, *JCAP* **1311** (2013) 018, [1304.5981].

[81] M. Blennow, E. Fernandez-Martinez, O. Mena, J. Redondo and P. Serra, *Asymmetric Dark Matter and Dark Radiation*, *JCAP* **1207** (2012) 022, [1203.5803].

[82] D. Baumann, D. Green and B. Wallisch, *New Target for Cosmic Axion Searches*, *Phys. Rev. Lett.* **117** (2016) 171301, [1604.08614].

[83] G. Mangano, G. Miele, S. Pastor, T. Pinto, O. Pisanti and P. D. Serpico, *Effects of non-standard neutrino-electron interactions on relic neutrino decoupling*, *Nucl. Phys.* **B756** (2006) 100–116, [hep-ph/0607267].

[84] P. F. de Salas, M. Lattanzi, G. Mangano, G. Miele, S. Pastor and O. Pisanti, *Bounds on very low reheating scenarios after Planck*, *Phys. Rev.* **D92** (2015) 123534, [1511.00672].

[85] A. G. Riess et al., *A 2.4% Determination of the Local Value of the Hubble Constant*, *Astrophys. J.* **826** (2016) 56, [1604.01424].

[86] E. Di Valentino, A. Melchiorri and J. Silk, *Reconciling Planck with the local value of H_0 in extended parameter space*, *Phys. Lett.* **B761** (2016) 242–246, [1606.00634].

[87] CMB-S4 collaboration, K. N. Abazajian et al., *CMB-S4 Science Book, First Edition*, 1610.02743.

[88] BICEP2 collaboration, P. A. R. Ade et al., *Detection of B-Mode Polarization at Degree Angular Scales by BICEP2*, *Phys. Rev. Lett.* **112** (2014) 241101, [1403.3985].

[89] J. M. Bardeen, *Gauge Invariant Cosmological Perturbations*, *Phys. Rev.* **D22** (1980) 1882–1905.

[90] R. K. Sachs and A. M. Wolfe, *Perturbations of a cosmological model and angular variations of the microwave background*, *Astrophys. J.* **147** (1967) 73–90.

[91] E. Komatsu and D. N. Spergel, *Acoustic signatures in the primary microwave background bispectrum*, *Phys. Rev.* **D63** (2001) 063002, [astro-ph/0005036].

[92] J. M. Maldacena, *Non-Gaussian features of primordial fluctuations in single field inflationary models*, *JHEP* **05** (2003) 013, [astro-ph/0210603].

[93] N. Bartolo, E. Komatsu, S. Matarrese and A. Riotto, *Non-Gaussianity from inflation: Theory and observations*, *Phys. Rept.* **402** (2004) 103–266, [astro-ph/0406398].

[94] Planck collaboration, P. A. R. Ade et al., *Planck 2015 results. XX. Constraints on inflation*, *Astron. Astrophys.* **594** (2016) A20, [1502.02114].

[95] A. Suzuki et al., *The LiteBIRD Satellite Mission - Sub-Kelvin Instrument*, in *17th International Workshop on Low Temperature Detectors (LTD 17) Kurume City, Japan, July 17-21, 2017*, 2018. 1801.06987.

[96] L. Perotto, J. Lesgourgues, S. Hannestad, H. Tu and Y. Y. Y. Wong, *Probing cosmological parameters with the CMB: Forecasts from full Monte Carlo simulations*, *JCAP* **0610** (2006) 013, [astro-ph/0606227].

[97] T. Okamoto and W. Hu, *CMB lensing reconstruction on the full sky*, *Phys. Rev.* **D67** (2003) 083002, [astro-ph/0301031].

[98] J. Errard, S. M. Feeney, H. V. Peiris and A. H. Jaffe, *Robust forecasts on fundamental physics from the foreground-obscured, gravitationally-lensed CMB polarization*, *JCAP* **1603** (2016) 052, [1509.06770].

[99] J. Lesgourgues, L. Perotto, S. Pastor and M. Piat, *Probing neutrino masses with cmb lensing extraction*, *Phys. Rev.* **D73** (2006) 045021, [astro-ph/0511735].

[100] A. Benoit-Levy, K. M. Smith and W. Hu, *Non-Gaussian structure of the lensed CMB power spectra covariance matrix*, *Phys. Rev.* **D86** (2012) 123008, [1205.0474].

[101] M. M. Schmittfull, A. Challinor, D. Hanson and A. Lewis, *Joint analysis of CMB temperature and lensing-reconstruction power spectra*, *Phys. Rev.* **D88** (2013) 063012, [1308.0286].

[102] T. Sprenger, M. Archidiacono, T. Brinckmann, S. Clesse and J. Lesgourgues, *Cosmology in the era of Euclid and the Square Kilometre Array*, 1801.08331.

[103] V. Tansella, C. Bonvin, R. Durrer, B. Ghosh and E. Sellentin, *The full-sky relativistic correlation function and power spectrum of galaxy number counts. Part I: theoretical aspects*, *JCAP* **1803** (2018) 019, [1708.00492].

[104] A. Schneider, R. Teyssier, D. Potter, J. Stadel, J. Onions, D. S. Reed et al., *Matter power spectrum and the challenge of percent accuracy*, *JCAP* **1604** (2016) 047, [1503.05920].

[105] R. Takahashi, M. Sato, T. Nishimichi, A. Taruya and M. Oguri, *Revising the Halofit Model for the Nonlinear Matter Power Spectrum*, *Astrophys. J.* **761** (2012) 152, [1208.2701].

[106] A. Mead, J. Peacock, C. Heymans, S. Joudaki and A. Heavens, *An accurate halo model for fitting non-linear cosmological power spectra and baryonic feedback models*, *Mon. Not. Roy. Astron. Soc.* **454** (2015) 1958–1975, [1505.07833].

[107] L. Casarini, S. A. Bonometto, E. Tessarotto and P. S. Corasaniti, *Extending the Coyote emulator to dark energy models with standard w_0-w_a parametrization of the equation of state*, *JCAP* **1608** (2016) 008, [1601.07230].

[108] E. Jennings, R. H. Wechsler, S. W. Skillman and M. S. Warren, *Disentangling redshift-space distortions and non-linear bias using the 2D power spectrum*, *Mon. Not. Roy. Astron. Soc.* **457** (2016) 1076–1088, [1508.01803].

[109] M. P. van Daalen, J. Schaye, C. M. Booth and C. D. Vecchia, *The effects of galaxy formation on the matter power spectrum: A challenge for precision cosmology*, *Mon. Not. Roy. Astron. Soc.* **415** (2011) 3649–3665, [1104.1174].

[110] A. Schneider and R. Teyssier, *A new method to quantify the effects of baryons on the matter power spectrum*, *JCAP* **1512** (2015) 049, [1510.06034].

[111] M. Rabold and R. Teyssier, *Precision cosmology with baryons: non-radiative hydrodynamics of galaxy groups*, *Mon. Not. Roy. Astron. Soc.* **467** (2017) 3188–3211, [1701.05337].

[112] S. Hilbert, D. Xu, P. Schneider, V. Springel, M. Vogelsberger and L. Hernquist, *Intrinsic Alignments of Galaxies in the Illustris Simulation*, *Mon. Not. Roy. Astron. Soc.* **468** (2017) 790–823, [1606.03216].

[113] P. Brax, A.-C. Davis, B. Li, H. A. Winther and G.-B. Zhao, *Systematic simulations of modified gravity: chameleon models*, *JCAP* **1304** (2013) 029, [1303.0007].

[114] P. Brax, A.-C. Davis, B. Li, H. A. Winther and G.-B. Zhao, *Systematic Simulations of Modified Gravity: Symmetron and Dilaton Models*, *JCAP* **1210** (2012) 002, [1206.3568].

[115] H. A. Winther et al., *Modified Gravity N-body Code Comparison Project*, *Mon. Not. Roy. Astron. Soc.* **454** (2015) 4208–4234, [1506.06384].

[116] B. S. Wright, H. A. Winther and K. Koyama, *COLA with massive neutrinos*, *JCAP* **1710** (2017) 054, [1705.08165].

[117] T. Basse, J. Hamann, S. Hannestad and Y. Y. Y. Wong, *Getting leverage on inflation with a large photometric redshift survey*, *JCAP* **1506** (2015) 042, [1409.3469].

[118] B. Audren, J. Lesgourgues, S. Bird, M. G. Haehnelt and M. Viel, *Neutrino masses and cosmological parameters from a Euclid-like survey: Markov Chain Monte Carlo forecasts including theoretical errors*, *JCAP* **1301** (2013) 026, [1210.2194].

[119] T. Baldauf, M. Mirbabayi, M. Simonović and M. Zaldarriaga, *LSS constraints with controlled theoretical uncertainties*, 1602.00674.

[120] E. Castorina, E. Sefusatti, R. K. Sheth, F. Villaescusa-Navarro and M. Viel, *Cosmology with massive neutrinos II: on the universality of the halo mass function and bias*, *JCAP* **1402** (2014) 049, [1311.1212].

[121] E. Castorina, C. Carbone, J. Bel, E. Sefusatti and K. Dolag, *DEMNUni: The clustering of large-scale structures in the presence of massive neutrinos*, *JCAP* **1507** (2015) 043, [1505.07148].

[122] P. Lemos, A. Challinor and G. Efstathiou, *The effect of Limber and flat-sky approximations on galaxy weak lensing*, *JCAP* **1705** (2017) 014, [1704.01054].

[123] M. Asgari, A. Taylor, B. Joachimi and T. D. Kitching, *Flat-Sky Pseudo-Cls Analysis for Weak Gravitational Lensing*, 1612.04664.

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

[125] J. C. Jackson, *Fingers of God: A critique of Rees' theory of primoridal gravitational radiation*, *Mon. Not. Roy. Astron. Soc.* **156** (1972) 1P–5P, [0810.3908].

[126] P. Bull, P. G. Ferreira, P. Patel and M. G. Santos, *Late-time cosmology with 21cm intensity mapping experiments*, *Astrophys. J.* **803** (2015) 21, [1405.1452].

[127] A. Raccanelli, L. Verde and F. Villaescusa-Navarro, *Biases from neutrino bias: to worry or not to worry?*, 1704.07837.

[128] M. LoVerde, *Halo bias in mixed dark matter cosmologies*, *Phys. Rev.* **D90** (2014) 083530, [1405.4855].

[129] J. E. Geach et al., *Empirical H-alpha emitter count predictions for dark energy surveys*, *Mon. Not. Roy. Astron. Soc.* **402** (2010) 1330, [0911.0686].

[130] E. Giusarma, S. Vagnozzi, S. Ho, S. Ferraro, K. Freese, R. Kamen-Rubio et al., *Scale-dependent galaxy bias, CMB lensing-galaxy cross-correlation, and neutrino masses*, 1802.08694.

[131] P. Dewdney, T. Stevenson and A. M. McPherson, *Ska1 system baseline design v2*, 2016.

[132] S. Yahya, P. Bull, M. G. Santos, M. Silva, R. Maartens, P. Okouma et al., *Cosmological performance of SKA HI galaxy surveys*, *Mon. Not. Roy. Astron. Soc.* **450** (2015) 2251–2260, [1412.4700].

[133] P. Bull, *Extending cosmological tests of General Relativity with the Square Kilometre Array*, *Astrophys. J.* **817** (2016) 26, [1509.07562].

[134] H. A. Feldman, N. Kaiser and J. A. Peacock, *Power spectrum analysis of three-dimensional redshift surveys*, *Astrophys. J.* **426** (1994) 23–37, [astro-ph/9304022].

[135] R. E. Smith and L. Marian, *Towards optimal estimation of the galaxy power spectrum*, *Mon. Not. Roy. Astron. Soc.* **454** (2015) 1266–1289, [1503.06830].

[136] I. Harrison, S. Camera, J. Zuntz and M. L. Brown, *SKA weak lensing – I. Cosmological forecasts and the power of radio-optical cross-correlations*, *Mon. Not. Roy. Astron. Soc.* **463** (2016) 3674–3685, [1601.03947].

[137] S. Furlanetto, S. P. Oh and F. Briggs, *Cosmology at Low Frequencies: The 21 cm Transition and the High-Redshift Universe*, *Phys. Rept.* **433** (2006) 181–301, [astro-ph/0608032].

[138] Y. Mao, M. Tegmark, M. McQuinn, M. Zaldarriaga and O. Zahn, *How accurately can 21 cm tomography constrain cosmology?*, *Phys. Rev.* **D78** (2008) 023529, [0802.1710].

[139] A. Lewis and A. Challinor, *The 21cm angular-power spectrum from the dark ages*, *Phys. Rev.* **D76** (2007) 083005, [astro-ph/0702600].

[140] S. Clesse, L. Lopez-Honorez, C. Ringeval, H. Tashiro and M. H. G. Tytgat, *Background reionization history from omniscopes*, *Phys. Rev.* **D86** (2012) 123506, [1208.4277].

[141] P. Brax, S. Clesse and A.-C. Davis, *Signatures of Modified Gravity on the 21-cm Power Spectrum at Reionisation*, *JCAP* **1301** (2013) 003, [1207.1273].

[142] R. A. Battye, I. W. A. Browne, C. Dickinson, G. Heron, B. Maffei and A. Pourtsidou, *HI intensity mapping : a single dish approach*, *Mon. Not. Roy. Astron. Soc.* **434** (2013) 1239–1256, [1209.0343].

[143] A. Hall, C. Bonvin and A. Challinor, *Testing General Relativity with 21-cm intensity mapping*, *Phys. Rev.* **D87** (2013) 064026, [1212.0728].

[144] F. Villaescusa-Navarro, D. Alonso and M. Viel, *Baryonic acoustic oscillations from 21 cm intensity mapping: the Square Kilometre Array case*, *Mon. Not. Roy. Astron. Soc.* **466** (2017) 2736–2751, [1609.00019].

[145] L. C. Olivari, C. Dickinson, R. A. Battye, Y.-Z. Ma, A. A. Costa, M. Remazeilles et al., *Cosmological parameter forecasts for HI intensity mapping experiments using the angular power spectrum*, *Mon. Not. Roy. Astron. Soc.* **473** (2018) 4242–4256, [1707.07647].

[146] D. Alonso, P. Bull, P. G. Ferreira and M. G. Santos, *Blind foreground subtraction for intensity mapping experiments*, *Mon. Not. Roy. Astron. Soc.* **447** (2015) 400, [1409.8667].

[147] VIRGO Consortium collaboration, R. E. Smith, J. A. Peacock, A. Jenkins, S. D. M. White, C. S. Frenk, F. R. Pearce et al., *Stable clustering, the halo model and nonlinear cosmological power spectra*, *Mon. Not. Roy. Astron. Soc.* **341** (2003) 1311, [astro-ph/0207664].

[148] F. Köhlinger et al., *KiDS-450: The tomographic weak lensing power spectrum and constraints on cosmological parameters*, *Mon. Not. Roy. Astron. Soc.* **471** (2017) 4412–4435, [1706.02892].

[149] A. F. Heavens, S. Matarrese and L. Verde, *The Nonlinear redshift-space power spectrum of galaxies*, *Mon. Not. Roy. Astron. Soc.* **301** (1998) 797–808, [astro-ph/9808016].

[150] D. Bianchi, M. Chiesa and L. Guzzo, *Improving the modelling of redshift-space distortions – I. A bivariate Gaussian description for the galaxy pairwise velocity...*, *Mon. Not. Roy. Astron. Soc.* **446** (2015) 75–84, [1407.4753].

[151] Y.-S. Song, Y. Zheng, A. Taruya and M. Oh, *Hybrid modeling of redshift space distortions*, 1801.04950.

[152] S. Bird, M. Viel and M. G. Haehnelt, *Massive Neutrinos and the Non-linear Matter Power Spectrum*, *Mon. Not. Roy. Astron. Soc.* **420** (2012) 2551–2561, [1109.4416].

[153] R. Teyssier, *Cosmological hydrodynamics with adaptive mesh refinement: a new high resolution code called ramses*, *Astron. Astrophys.* **385** (2002) 337–364, [astro-ph/0111367].

[154] D. Potter, J. Stadel and R. Teyssier, *PKDGRAV3: Beyond Trillion Particle Cosmological Simulations for the Next Era of Galaxy Surveys*, 1609.08621.

[155] V. Springel, *The Cosmological simulation code GADGET-2*, *Mon. Not. Roy. Astron. Soc.* **364** (2005) 1105–1134, [astro-ph/0505010].

[156] M. Shimon, N. J. Miller, C. T. Kishimoto, C. J. Smith, G. M. Fuller and B. G. Keating, *Using Big Bang Nucleosynthesis to Extend CMB Probes of Neutrino Physics*, *JCAP* **1005** (2010) 037, [1001.5088].

[157] M. Archidiacono, T. Brinckmann, J. Lesgourgues and V. Poulin, *Physical effects involved in the measurements of neutrino masses with future cosmological data*, *JCAP* **1702** (2017) 052, [1610.09852].

[158] V. Desjacques, D. Jeong and F. Schmidt, *Large-Scale Galaxy Bias*, *Phys. Rept.* **733** (2018) 1–193, [1611.09787].

[159] F. Villaescusa-Navarro, F. Marulli, M. Viel, E. Branchini, E. Castorina, E. Sefusatti et al., *Cosmology with massive neutrinos I: towards a realistic modeling of the relation between matter, haloes and galaxies*, *JCAP* **1403** (2014) 011, [1311.0866].

[160] M. Costanzi, F. Villaescusa-Navarro, M. Viel, J.-Q. Xia, S. Borgani, E. Castorina et al., *Cosmology with massive neutrinos III: the halo mass function andan application to galaxy clusters*, *JCAP* **1312** (2013) 012, [1311.1514].

[161] M. Biagetti, V. Desjacques, A. Kehagias and A. Riotto, *Nonlocal halo bias with and without massive neutrinos*, *Phys. Rev.* **D90** (2014) 045022, [1405.1435].

[162] U. Seljak, *Redshift space bias and beta from the halo model*, *Mon. Not. Roy. Astron. Soc.* **325** (2001) 1359, [astro-ph/0009016].

[163] L. Amendola, E. Menegoni, C. Di Porto, M. Corsi and E. Branchini, *Constraints on a scale-dependent bias from galaxy clustering*, *Phys. Rev.* **D95** (2017) 023505, [1502.03994].

[164] M. LoVerde, *Spherical collapse in $\nu\Lambda$CDM*, *Phys. Rev.* **D90** (2014) 083518, [1405.4858].

[165] M. LoVerde, *Neutrino mass without cosmic variance*, *Phys. Rev.* **D93** (2016) 103526, [1602.08108].

[166] C.-T. Chiang, W. Hu, Y. Li and M. Loverde, *Scale-dependent bias and bispectrum in neutrino separate universe simulations*, *Phys. Rev.* **D97** (2018) 123526, [1710.01310].

[167] J. B. Muñoz and C. Dvorkin, *Efficient Computation of Galaxy Bias with Neutrinos and Other Relics*, 1805.11623.

[168] L. Amendola et al., *Cosmology and fundamental physics with the Euclid satellite*, *Living Rev. Rel.* **21** (2018) 2, [1606.00180].

[169] L. Samushia, W. J. Percival and A. Raccanelli, *Interpreting large-scale redshift-space distortion measurements*, *Mon. Not. Roy. Astron. Soc.* **420** (2012) 2102–2119, [1102.1014].

[170] F. Villaescusa-Navarro, A. Banerjee, N. Dalal, E. Castorina, R. Scoccimarro, R. Angulo et al., *The imprint of neutrinos on clustering in redshift-space*, 1708.01154.

[171] D. Blas, J. Lesgourgues and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS) II: Approximation schemes*, *JCAP* **1107** (2011) 034, [1104.2933].

[172] T. Matsubara, *Resumming Cosmological Perturbations via the Lagrangian Picture: One-loop Results in Real Space and in Redshift Space*, *Phys. Rev.* **D77** (2008) 063530, [0711.2521].

[173] M. Sato and T. Matsubara, *Nonlinear Biasing and Redshift-Space Distortions in Lagrangian Resummation Theory and N-body Simulations*, *Phys. Rev.* **D84** (2011) 043501, [1105.5007].

[174] B. Audren, J. Lesgourgues, K. Benabed and S. Prunet, *Conservative Constraints on Early Cosmology: an illustration of the Monte Python cosmological parameter inference code*, *JCAP* **1302** (2013) 001, [1210.7183].

[175] T. Brinckmann and J. Lesgourgues, *MontePython 3: boosted MCMC sampler and other features*, 1804.07261.

[176] J. Lesgourgues, S. Pastor and L. Perotto, *Probing neutrino masses with future galaxy redshift surveys*, *Phys. Rev.* **D70** (2004) 045016, [hep-ph/0403296].

[177] F. De Bernardis, T. D. Kitching, A. Heavens and A. Melchiorri, *Determining the Neutrino Mass Hierarchy with Cosmology*, *Phys. Rev.* **D80** (2009) 123509, [0907.1917].

[178] C. Carbone, L. Verde, Y. Wang and A. Cimatti, *Neutrino constraints from future nearly all-sky spectroscopic galaxy surveys*, *JCAP* **1103** (2011) 030, [1012.2868].

[179] A. C. Hall and A. Challinor, *Probing the neutrino mass hierarchy with CMB weak lensing*, *Mon. Not. Roy. Astron. Soc.* **425** (2012) 1170–1184, [1205.6172].

[180] M. Gerbino, K. Freese, S. Vagnozzi, M. Lattanzi, O. Mena, E. Giusarma et al., *Impact of neutrino properties on the estimation of inflationary parameters from current and future observations*, *Phys. Rev.* **D95** (2017) 043512, [1610.08830].

[181] A. Gelman and D. B. Rubin, *Inference from Iterative Simulation Using Multiple Sequences*, *Statist. Sci.* **7** (1992) 457–472.

[182] W. Hu, D. J. Eisenstein and M. Tegmark, *Weighing neutrinos with galaxy surveys*, *Phys. Rev. Lett.* **80** (1998) 5255–5258, [astro-ph/9712057].

[183] J. Hamann, S. Hannestad and Y. Y. Y. Wong, *Measuring neutrino masses with a future galaxy survey*, *JCAP* **1211** (2012) 052, [1209.1043].

[184] Topical Conveners: K.N. Abazajian, J.E. Carlstrom, A.T. Lee collaboration, K. N. Abazajian et al., *Neutrino Physics from the Cosmic Microwave Background and Large Scale Structure*, *Astropart. Phys.* **63** (2015) 66–80, [1309.5383].

[185] J. Lesgourgues and S. Pastor, *Neutrino mass from Cosmology*, *Adv. High Energy Phys.* **2012** (2012) 608515, [1212.6154].

[186] Z. Hou et al., *Constraints on Cosmology from the Cosmic Microwave Background Power Spectrum of the 2500 deg^2 SPT-SZ Survey*, *Astrophys. J.* **782** (2014) 74, [1212.6267].

[187] J. R. Bond, G. Efstathiou and J. Silk, *Massive Neutrinos and the Large Scale Structure of the Universe*, *Phys. Rev. Lett.* **45** (1980) 1980–1984.

[188] M. Kaplinghat, L. Knox and Y.-S. Song, *Determining neutrino mass from the CMB alone*, *Phys. Rev. Lett.* **91** (2003) 241301, [astro-ph/0303344].

[189] S. Bashinsky and U. Seljak, *Neutrino perturbations in CMB anisotropy and matter clustering*, *Phys. Rev.* **D69** (2004) 083002, [astro-ph/0310198].

[190] S. Hannestad, *Neutrino physics from precision cosmology*, *Prog. Part. Nucl. Phys.* **65** (2010) 185–208, [1007.0658].

[191] J. Brandbyge, S. Hannestad, T. Haugbølle and B. Thomsen, *The Effect of Thermal Neutrino Motion on the Non-linear Cosmological Matter Power Spectrum*, *JCAP* **0808** (2008) 020, [0802.3700].

[192] J. Brandbyge, S. Hannestad, T. Haugbølle and Y. Y. Y. Wong, *Neutrinos in Non-linear Structure Formation - The Effect on Halo Properties*, *JCAP* **1009** (2010) 014, [1004.4105].

[193] Y. Ali-Haimoud and S. Bird, *An efficient implementation of massive neutrinos in non-linear structure formation simulations*, *Mon. Not. Roy. Astron. Soc.* **428** (2012) 3375–3389, [1209.0461].

[194] M. Archidiacono and S. Hannestad, *Efficient calculation of cosmological neutrino clustering in the non-linear regime*, *JCAP* **1606** (2016) 018, [1510.02907].

[195] C. Carbone, M. Petkova and K. Dolag, *DEMNUni: ISW, Rees-Sciama, and weak-lensing in the presence of massive neutrinos*, *JCAP* **1607** (2016) 034, [1605.02024].

[196] H. Dupuy and F. Bernardeau, *On the importance of nonlinear couplings in large-scale neutrino streams*, *JCAP* **1508** (2015) 053, [1503.05707].

[197] F. Fürer and Y. Y. Y. Wong, *Higher-order massive neutrino perturbations in large-scale structure*, *JCAP* **1503** (2015) 046, [1412.2764].

[198] A. Liu, J. R. Pritchard, R. Allison, A. R. Parsons, U. Seljak and B. D. Sherwin, *Eliminating the optical depth nuisance from the CMB with 21 cm cosmology*, *Phys. Rev.* **D93** (2016) 043013, [1509.08463].

[199] R. Allison, P. Caucal, E. Calabrese, J. Dunkley and T. Louis, *Towards a cosmological neutrino mass detection*, *Phys. Rev.* **D92** (2015) 123535, [1509.07471].

[200] W. T. Hu, *Wandering in the Background: A CMB Explorer*. PhD thesis, UC, Berkeley, 1995. astro-ph/9508126.

[201] C. Howlett, A. Lewis, A. Hall and A. Challinor, *CMB power spectrum parameter degeneracies in the era of precision cosmology*, *JCAP* **1204** (2012) 027, [1201.3654].

[202] CORE collaboration, G. De Zotti et al., *Exploring cosmic origins with CORE: Extragalactic sources in cosmic microwave background maps*, *JCAP* **1804** (2018) 020, [1609.07263].

[203] CORE collaboration, J. Delabrouille et al., *Exploring cosmic origins with CORE: Survey requirements and mission design*, *JCAP* **1804** (2018) 014, [1706.04516].

[204] T. Matsumura et al., *LiteBIRD: Mission Overview and Focal Plane Layout*, *J. Low. Temp. Phys.* **184** (2016) 824–831.

[205] J. Lesgourgues, *The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview*, 1104.2932.

[206] J. Lesgourgues and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS) IV: efficient implementation of non-cold relics*, *JCAP* **1109** (2011) 032, [1104.2935].

[207] Z. Pan, L. Knox and M. White, *Dependence of the Cosmic Microwave Background Lensing Power Spectrum on the Matter Density*, *Mon. Not. Roy. Astron. Soc.* **445** (2014) 2941–2945, [1406.5459].

[208] Planck collaboration, P. A. R. Ade et al., *Planck 2015 results. XV. Gravitational lensing*, *Astron. Astrophys.* **594** (2016) A15, [1502.01591].

[209] F. Beutler, C. Blake, M. Colless, D. H. Jones, L. Staveley-Smith, L. Campbell et al., *The 6dF Galaxy Survey: Baryon Acoustic Oscillations and the Local Hubble Constant*, *Mon. Not. Roy. Astron. Soc.* **416** (2011) 3017–3032, [1106.3366].

[210] A. J. Ross, L. Samushia, C. Howlett, W. J. Percival, A. Burden and M. Manera, *The clustering of the SDSS DR7 main Galaxy sample – I. A 4 per cent distance measure at* $z = 0.15$, *Mon. Not. Roy. Astron. Soc.* **449** (2015) 835–847, [1409.3242].

[211] BOSS collaboration, L. Anderson et al., *The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: baryon acoustic oscillations in the Data Releases 10 and 11 Galaxy samples*, *Mon. Not. Roy. Astron. Soc.* **441** (2014) 24–62, [1312.4877].

[212] A. Font-Ribera, P. McDonald, N. Mostek, B. A. Reid, H.-J. Seo and A. Slosar, *DESI and other dark energy experiments in the era of neutrino mass measurements*, *JCAP* **1405** (2014) 023, [1308.4164].

[213] T. Namikawa, S. Saito and A. Taruya, *Probing dark energy and neutrino mass from upcoming lensing experiments of CMB and galaxies*, *JCAP* **1012** (2010) 027, [1009.3204].

[214] E. Di Dio, F. Montanari, J. Lesgourgues and R. Durrer, *The CLASSgal code for Relativistic Cosmological Large Scale Structure*, *JCAP* **1311** (2013) 044, [1307.1459].

[215] E. Di Dio, F. Montanari, R. Durrer and J. Lesgourgues, *Cosmological Parameter Estimation with Large Scale Structure Observations*, *JCAP* **1401** (2014) 042, [1308.6186].

[216] T. Basse, O. E. Bjaelde, J. Hamann, S. Hannestad and Y. Y. Y. Wong, *Dark energy properties from large future galaxy surveys*, *JCAP* **1405** (2014) 021, [1304.2321].

[217] M. Archidiacono, T. Basse, J. Hamann, S. Hannestad, G. Raffelt and Y. Y. Y. Wong, *Future cosmological sensitivity for hot dark matter axions*, *JCAP* **1505** (2015) 050, [1502.03325].

[218] A. Rassat, A. Amara, L. Amendola, F. J. Castander, T. Kitching, M. Kunz et al., *Deconstructing Baryon Acoustic Oscillations: A Comparison of Methods*, 0810.0003.

[219] J. Hamann, S. Hannestad, J. Lesgourgues, C. Rampf and Y. Y. Y. Wong, *Cosmological parameters from large scale structure - geometric versus shape information*, *JCAP* **1007** (2010) 022, [1003.3999].

[220] V. Poulin, P. D. Serpico and J. Lesgourgues, *A fresh look at linear cosmological constraints on a decaying dark matter component*, *JCAP* **1608** (2016) 036, [1606.02073].

[221] A. Lewis and S. Bridle, *Cosmological parameters from CMB and other data: A Monte Carlo approach*, *Phys. Rev.* **D66** (2002) 103511, [astro-ph/0205436].

[222] A. Lewis, *Efficient sampling of fast and slow cosmological parameters*, *Phys. Rev.* **D87** (2013) 103529, [1304.4473].

[223] J. Zuntz, M. Paterno, E. Jennings, D. Rudd, A. Manzotti, S. Dodelson et al., *CosmoSIS: Modular Cosmological Parameter Estimation*, *Astron. Comput.* **12** (2015) 45–59, [1409.3409].

[224] J. Lesgourgues, *The Cosmic Linear Anisotropy Solving System (CLASS) III: Comparision with CAMB for LambdaCDM*, 1104.2934.

[225] M. Zumalacárregui, E. Bellini, I. Sawicki, J. Lesgourgues and P. G. Ferreira, *hi_class: Horndeski in the Cosmic Linear Anisotropy Solving System*, *JCAP* **1708** (2017) 019, [1605.06102].

[226] G. W. Pettinari, C. Fidler, R. Crittenden, K. Koyama and D. Wands, *The intrinsic bispectrum of the Cosmic Microwave Background*, *JCAP* **1304** (2013) 003, [1302.0832].

[227] P. Stöcker, M. Krämer, J. Lesgourgues and V. Poulin, *Exotic energy injection with ExoCLASS: Application to the Higgs portal model and evaporating black holes*, *JCAP* **1803** (2018) 018, [1801.01871].

[228] B. Bolliet, B. Comis, E. Komatsu and J. F. Macías-Pérez, *Dark Energy from the Thermal Sunyaev Zeldovich Power Spectrum*, *Mon. Not. Roy. Astron. Soc.* (2017) , [1712.00788].

[229] Y. Dirian, S. Foffa, M. Kunz, M. Maggiore and V. Pettorino, *Non-local gravity and comparison with observational datasets. II. Updated results and Bayesian model comparison with ΛCDM*, *JCAP* **1605** (2016) 068, [1602.03558].

[230] A. Lewis, A. Challinor and A. Lasenby, *Efficient computation of CMB anisotropies in closed FRW models*, *Astrophys. J.* **538** (2000) 473–476, [astro-ph/9911177].

[231] A. Refregier, L. Gamper, A. Amara and L. Heisenberg, *PyCosmo: An Integrated Cosmological Boltzmann Solver*, 1708.05177.

[232] J. Dunkley, M. Bucher, P. G. Ferreira, K. Moodley and C. Skordis, *Fast and reliable mcmc for cosmological parameter estimation*, *Mon. Not. Roy. Astron. Soc.* **356** (2005) 925–936, [astro-ph/0405462].

[233] H. Haario, E. Saksman and J. Tamminen, *An adaptive metropolis algorithm*, *Bernoulli* **7** (04, 2001) 223–242.

[234] S. Henrot-Versillé, O. Perdereau, S. Plaszczynski, B. R. d'Orfeuil, M. Spinelli and M. Tristram, *Agnostic cosmology in the CAMEL framework*, 1607.02964.

[235] B. Schroer, *Forecasts on the sensitivity of future cosmological experiments using a new algorithm for bayesian parameter extraction*, *Bachelor thesis, TTK, RWTH Aachen University* (2016) .

[236] D. Coe, *Fisher Matrices and Confidence Ellipses: A Quick-Start Guide and Software*, 0906.4123.

[237] M. Chevallier and D. Polarski, *Accelerating universes with scaling dark matter*, *Int. J. Mod. Phys.* **D10** (2001) 213–224, [gr-qc/0009008].

[238] E. V. Linder, *Exploring the expansion history of the universe*, *Phys. Rev. Lett.* **90** (2003) 091301, [astro-ph/0208512].

[239] W. Fang, W. Hu and A. Lewis, *Crossing the Phantom Divide with Parameterized Post-Friedmann Dark Energy*, *Phys. Rev.* **D78** (2008) 087303, [0808.3125].

[240] F. Feroz and M. P. Hobson, *Multimodal nested sampling: an efficient and robust alternative to MCMC methods for astronomical data analysis*, *Mon. Not. Roy. Astron. Soc.* **384** (2008) 449, [0704.3704].

[241] F. Feroz, M. P. Hobson and M. Bridges, *MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics*, *Mon. Not. Roy. Astron. Soc.* **398** (2009) 1601–1614, [0809.3437].

[242] F. Feroz, M. P. Hobson, E. Cameron and A. N. Pettitt, *Importance Nested Sampling and the MultiNest Algorithm*, 1306.2144.

[243] J. Buchner, A. Georgakakis, K. Nandra, L. Hsu, C. Rangel, M. Brightman et al., *X-ray spectral modelling of the AGN obscuring region in the CDFS: Bayesian model selection and catalogue*, *Astron. Astrophys.* **564** (2014) A125, [1402.0004].

[244] J. Goodman and J. Weare, *Ensemble samplers with affine invariance*, *Communications in Applied Mathematics and Computational Science, Vol. 5, No. 1, p. 65-80, 2010* **5** (2010) 65–80.

[245] D. Foreman-Mackey, D. W. Hogg, D. Lang and J. Goodman, *emcee: The MCMC Hammer*, *Publ. Astron. Soc. Pac.* **125** (2013) 306–312, [1202.3665].

[246] J. Akeret, S. Seehars, A. Amara, A. Refregier and A. Csillaghy, *CosmoHammer: Cosmological parameter estimation with the MCMC Hammer*, 1212.1721.

[247] CORE collaboration, F. Finelli et al., *Exploring cosmic origins with CORE: Inflation*, *JCAP* **1804** (2018) 016, [1612.08270].

[248] Planck collaboration, J. Tauber, M. Bersanelli, J. M. Lamarre, G. Efstathiou, C. Lawrence, F. Bouchet et al., *The Scientific programme of Planck*, astro-ph/0604069.

[249] C.-L. Kuo et al., *Improved Measurements of the CMB Power Spectrum with ACBAR*, *Astrophys. J.* **664** (2007) 687–701, [astro-ph/0611198].

[250] L. Anderson et al., *The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations in the Data Release 9 Spectroscopic Galaxy Sample*, *Mon. Not. Roy. Astron. Soc.* **427** (2013) 3435–3467, [1203.6594].

[251] B. Audren, *Separate Constraints on Early and Late Cosmology*, *Mon. Not. Roy. Astron. Soc.* **444** (2014) 827–832, [1312.5696].

[252] E. Sanchez, A. Carnero, J. Garcia-Bellido, E. Gaztanaga, F. de Simoni, M. Crocce et al., *Tracing The Sound Horizon Scale With Photometric Redshift Surveys*, *Mon. Not. Roy. Astron. Soc.* **411** (2011) 277–288, [1006.3226].

[253] G. C. Carvalho, A. Bernui, M. Benetti, J. C. Carvalho and J. S. Alcaniz, *Baryon Acoustic Oscillations from the SDSS DR10 galaxies angular correlation function*, *Phys. Rev.* **D93** (2016) 023530, [1507.08972].

[254] J. S. Alcaniz, G. C. Carvalho, A. Bernui, J. C. Carvalho and M. Benetti, *Measuring baryon acoustic oscillations with angular two-point correlation function*, *Fundam. Theor. Phys.* **187** (2017) 11–19, [1611.08458].

[255] E. de Carvalho, A. Bernui, G. C. Carvalho, C. P. Novaes and H. S. Xavier, *Angular Baryon Acoustic Oscillation measure at $z = 2.225$ from the SDSS quasar survey*, *JCAP* **1804** (2018) 064, [1709.00113].

[256] G. C. Carvalho, A. Bernui, M. Benetti, J. C. Carvalho and J. S. Alcaniz, *Measuring the transverse baryonic acoustic scale from the SDSS DR11 galaxies*, 1709.00271.

[257] M. A. Buen-Abad, M. Schmaltz, J. Lesgourgues and T. Brinckmann, *Interacting Dark Sector and Precision Cosmology*, *JCAP* **1801** (2018) 008, [1708.09406].

[258] H. C. Chiang et al., *Measurement of CMB Polarization Power Spectra from Two Years of BICEP Data*, *Astrophys. J.* **711** (2010) 1123–1140, [0906.1181].

[259] BICEP2, Keck Array collaboration, P. A. R. Ade et al., *Improved Constraints on Cosmology and Foregrounds from BICEP2 and Keck Array Cosmic Microwave Background Data with Inclusion of 95 GHz Band*, *Phys. Rev. Lett.* **116** (2016) 031302, [1510.09217].

[260] W. C. Jones et al., *A Measurement of the angular power spectrum of the CMB temperature anisotropy from the 2003 flight of BOOMERANG*, *Astrophys. J.* **647** (2006) 823–832, [astro-ph/0507494].

[261] J. L. Sievers et al., *Implications of the cosmic background imager polarization data*, *Astrophys. J.* **660** (2007) 976–987, [astro-ph/0509203].

[262] C. Heymans et al., *CFHTLenS tomographic weak lensing cosmological parameter constraints: Mitigating the impact of intrinsic galaxy alignments*, *Mon. Not. Roy. Astron. Soc.* **432** (2013) 2433, [1303.1808].

[263] D. Larson et al., *Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Power Spectra and WMAP-Derived Parameters*, *Astrophys. J. Suppl.* **192** (2011) 16, [1001.4635].

[264] M. Moresco, L. Pozzetti, A. Cimatti, R. Jimenez, C. Maraston, L. Verde et al., *A 6% measurement of the Hubble parameter at z~0.45: direct evidence of the epoch of cosmic re-acceleration*, *JCAP* **1605** (2016) 014, [1601.01701].

[265] M. Moresco et al., *Improved constraints on the expansion rate of the Universe up to z~1.1 from the spectroscopic evolution of cosmic chronometers*, *JCAP* **1208** (2012) 006, [1201.3609].

[266] J. Simon, L. Verde and R. Jimenez, *Constraints on the redshift dependence of the dark energy potential*, *Phys. Rev.* **D71** (2005) 123001, [astro-ph/0412269].

[267] D. Stern, R. Jimenez, L. Verde, M. Kamionkowski and S. A. Stanford, *Cosmic Chronometers: Constraining the Equation of State of Dark Energy. I: H(z) Measurements*, *JCAP* **1002** (2010) 008, [0907.3149].

[268] X.-H. Fan, C. L. Carilli and B. G. Keating, *Observational constraints on cosmic reionization*, *Ann. Rev. Astron. Astrophys.* **44** (2006) 415–462, [astro-ph/0602375].

[269] J. Schaye, T. Theuns, M. Rauch, G. Efstathiou and W. L. W. Sargent, *The Thermal history of the intergalactic medium*, *Mon. Not. Roy. Astron. Soc.* **318** (2000) 817, [astro-ph/9912432].

[270] B. Stölzner, A. Cuoco, J. Lesgourgues and M. Bilicki, *Updated tomographic analysis of the integrated Sachs-Wolfe effect and implications for dark energy*, *Phys. Rev.* **D97** (2018) 063506, [1710.03238].

[271] SDSS collaboration, M. Betoule et al., *Improved cosmological constraints from a joint analysis of the SDSS-II and SNLS supernova samples*, *Astron. Astrophys.* **568** (2014) A22, [1401.4064].

[272] Planck collaboration, P. A. R. Ade et al., *Planck 2013 results. XV. CMB power spectra and likelihood*, *Astron. Astrophys.* **571** (2014) A15, [1303.5075].

[273] J. Dunkley et al., *The Atacama Cosmology Telescope: likelihood for small-scale CMB data*, *JCAP* **1307** (2013) 025, [1301.0776].

[274] R. Keisler et al., *A Measurement of the Damping Tail of the Cosmic Microwave Background Power Spectrum with the South Pole Telescope*, *Astrophys. J.* **743** (2011) 28, [1105.3182].

[275] Planck collaboration, N. Aghanim et al., *Planck 2015 results. XI. CMB power spectra, likelihoods, and robustness of parameters*, *Astron. Astrophys.* **594** (2016) A11, [1507.02704].

[276] Planck collaboration, P. A. R. Ade et al., *Planck 2015 results. XXIV. Cosmology from Sunyaev-Zeldovich cluster counts*, *Astron. Astrophys.* **594** (2016) A24, [1502.01597].

[277] POLARBEAR collaboration, P. A. R. Ade et al., *A Measurement of the Cosmic Microwave Background B-Mode Polarization Power Spectrum at Sub-Degree Scales with POLARBEAR*, *Astrophys. J.* **794** (2014) 171, [1403.2369].

[278] QUaD collaboration, M. L. Brown et al., *Improved measurements of the temperature and polarization of the CMB from QUaD*, *Astrophys. J.* **705** (2009) 978–999, [0906.1003].

[279] SDSS collaboration, M. Tegmark et al., *Cosmological Constraints from the SDSS Luminous Red Galaxies*, *Phys. Rev.* **D74** (2006) 123507, [astro-ph/0608632].

[280] B. A. Reid et al., *Cosmological Constraints from the Clustering of the Sloan Digital Sky Survey DR7 Luminous Red Galaxies*, *Mon. Not. Roy. Astron. Soc.* **404** (2010) 60–85, [0907.1659].

[281] Planck collaboration, R. Adam et al., *Planck intermediate results. XLVII. Planck constraints on reionization history*, *Astron. Astrophys.* **596** (2016) A108, [1605.03507].

[282] R. Amanullah et al., *Spectra and Light Curves of Six Type Ia Supernovae at $0.511 < z < 1.12$ and the Union2 Compilation*, *Astrophys. J.* **716** (2010) 712–738, [1004.1711].

[283] K. K. Schaffer et al., *The First Public Release of South Pole Telescope Data: Maps of a 95-square-degree Field from 2008 Observations*, *Astrophys. J.* **743** (2011) 90, [1111.7245].

[284] S. H. Suyu et al., *Two accurate time-delay distances from strong lensing: Implications for cosmology*, *Astrophys. J.* **766** (2013) 70, [1208.6010].

[285] D. Parkinson et al., *The WiggleZ Dark Energy Survey: Final data release and cosmological results*, *Phys. Rev.* **D86** (2012) 103518, [1210.2130].

[286] E. A. Kazin et al., *The WiggleZ Dark Energy Survey: improved distance measurements to $z = 1$ with reconstruction of the baryonic acoustic feature*, *Mon. Not. Roy. Astron. Soc.* **441** (2014) 3524–3542, [1401.0358].

[287] WMAP collaboration, C. L. Bennett et al., *Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results*, *Astrophys. J. Suppl.* **208** (2013) 20, [1212.5225].

[288] S. Grandis, D. Rapetti, A. Saro, J. J. Mohr and J. P. Dietrich, *Quantifying tensions between CMB and distance data sets in models with free curvature or lensing amplitude*, *Mon. Not. Roy. Astron. Soc.* **463** (2016) 1416–1430, [1604.06463].

[289] J. L. Bernal, L. Verde and A. G. Riess, *The trouble with H_0*, *JCAP* **1610** (2016) 019, [1607.05617].

[290] V. Bonvin et al., *HoLiCOW – V. New COSMOGRAIL time delays of HE 0435-1223: H_0 to 3.8 per cent precision from strong lensing in a flat ΛCDM model*, *Mon. Not. Roy. Astron. Soc.* **465** (2017) 4914–4930, [1607.01790].

[291] S. Joudaki et al., *CFHTLenS revisited: assessing concordance with Planck including astrophysical systematics*, *Mon. Not. Roy. Astron. Soc.* **465** (2017) 2033–2052, [1601.05786].

[292] H. Hildebrandt et al., *KiDS-450: Cosmological parameter constraints from tomographic weak gravitational lensing*, *Mon. Not. Roy. Astron. Soc.* **465** (2017) 1454, [1606.05338].

[293] Q.-G. Huang and K. Wang, *How the dark energy can reconcile Planck with local determination of the Hubble constant*, *Eur. Phys. J.* **C76** (2016) 506, [1606.05965].

[294] P. Ko and Y. Tang, *Light dark photon and fermionic dark radiation for the Hubble constant and the structure formation*, *Phys. Lett.* **B762** (2016) 462–466, [1608.01083].

[295] T. Karwal and M. Kamionkowski, *Dark energy at early times, the Hubble parameter, and the string axiverse*, *Phys. Rev.* **D94** (2016) 103523, [1608.01309].

[296] S. Kumar and R. C. Nunes, *Probing the interaction between dark matter and dark energy in the presence of massive neutrinos*, *Phys. Rev.* **D94** (2016) 123511, [1608.02454].

[297] P. Ko and Y. Tang, *Residual Non-Abelian Dark Matter and Dark Radiation*, *Phys. Lett.* **B768** (2017) 12–17, [1609.02307].

[298] V. Prilepina and Y. Tsai, *Reconciling Large And Small-Scale Structure In Twin Higgs Models*, *JHEP* **09** (2017) 033, [1611.05879].

[299] F. Villaescusa-Navarro, P. Bull and M. Viel, *Weighing neutrinos with cosmic neutral hydrogen*, *Astrophys. J.* **814** (2015) 146, [1507.05102].

[300] G. Mellema et al., *Reionization and the Cosmic Dawn with the Square Kilometre Array*, *Exper. Astron.* **36** (2013) 235–318, [1210.0197].

[301] Y. Oyama, A. Shimizu and K. Kohri, *Determination of neutrino mass hierarchy by 21 cm line and CMB B-mode polarization observations*, *Phys. Lett.* **B718** (2013) 1186–1193, [1205.5223].

[302] Y. Oyama, K. Kohri and M. Hazumi, *Constraints on the neutrino parameters by future cosmological 21 cm line and precise CMB polarization observations*, *JCAP* **1602** (2016) 008, [1510.03806].

[303] J. Bonn, K. Eitel, F. Gluck, D. Sevilla-Sanchez and N. Titov, *The KATRIN sensitivity to the neutrino mass and to right-handed currents in beta decay*, *Phys. Lett.* **B703** (2011) 310–312, [0704.3930].

[304] A. Boyle and E. Komatsu, *Deconstructing the neutrino mass constraint from galaxy redshift surveys*, *JCAP* **1803** (2018) 035, [1712.01857].

[305] A. Obuljen, E. Castorina, F. Villaescusa-Navarro and M. Viel, *High-redshift post-reionization cosmology with 21cm intensity mapping*, *JCAP* **1805** (2018) 004, [1709.07893].

[306] K. M. Smith, W. Hu and M. Kaplinghat, *Cosmological Information from Lensed CMB Power Spectra*, *Phys. Rev.* **D74** (2006) 123002, [astro-ph/0607315].

[307] G. Mangano, G. Miele, S. Pastor, T. Pinto, O. Pisanti and P. D. Serpico, *Relic neutrino decoupling including flavor oscillations*, *Nucl. Phys.* **B729** (2005) 221–234, [hep-ph/0506164].

[308] A. D. Dolgov, S. H. Hansen and D. V. Semikoz, *Nonequilibrium corrections to the spectra of massless neutrinos in the early universe: Addendum*, *Nucl. Phys.* **B543** (1999) 269–274, [hep-ph/9805467].

[309] R. E. Lopez, S. Dodelson, A. Heckler and M. S. Turner, *Precision detection of the cosmic neutrino background*, *Phys. Rev. Lett.* **82** (1999) 3952–3955, [astro-ph/9803095].

[310] A. D. Dolgov, *Neutrinos in cosmology*, *Phys. Rept.* **370** (2002) 333–535, [hep-ph/0202122].

[311] P. F. de Salas and S. Pastor, *Relic neutrino decoupling with flavour oscillations revisited*, *JCAP* **1607** (2016) 051, [1606.06986].

[312] R. Tomas, H. Pas and J. W. F. Valle, *Generalized bounds on Majoron - neutrino couplings*, *Phys. Rev.* **D64** (2001) 095005, [hep-ph/0103017].

[313] M. Archidiacono, E. Calabrese and A. Melchiorri, *The Case for Dark Radiation*, *Phys. Rev.* **D84** (2011) 123008, [1109.2767].

[314] Z. Hou, R. Keisler, L. Knox, M. Millea and C. Reichardt, *How Massless Neutrinos Affect the Cosmic Microwave Background Damping Tail*, *Phys. Rev.* **D87** (2013) 083008, [1104.2333].

[315] M. Drewes et al., *A White Paper on keV Sterile Neutrino Dark Matter*, *JCAP* **1701** (2017) 025, [1602.04816].

[316] C. Giunti, *Light Sterile Neutrinos: Status and Perspectives*, *Nucl. Phys.* **B908** (2016) 336–353, [1512.04758].

[317] M. Archidiacono, S. Gariazzo, C. Giunti, S. Hannestad, R. Hansen, M. Laveder et al., *Pseudoscalar—sterile neutrino interactions: reconciling the cosmos with neutrino oscillations*, *JCAP* **1608** (2016) 067, [1606.07673].

[318] S. Bridle, J. Elvin-Poole, J. Evans, S. Fernandez, P. Guzowski and S. Soldner-Rembold, *A Combined View of Sterile-Neutrino Constraints from CMB and Neutrino Oscillation Measurements*, *Phys. Lett.* **B764** (2017) 322–327, [1607.00032].

[319] M. Archidiacono, E. Giusarma, A. Melchiorri and O. Mena, *Dark Radiation in extended cosmological scenarios*, *Phys. Rev.* **D86** (2012) 043509, [1206.0109].

[320] P. Crotty, J. Lesgourgues and S. Pastor, *Current cosmological bounds on neutrino masses and relativistic relics*, *Phys. Rev.* **D69** (2004) 123007, [hep-ph/0402049].

[321] M. Archidiacono, E. Giusarma, S. Hannestad and O. Mena, *Cosmic dark radiation and neutrinos*, *Adv. High Energy Phys.* **2013** (2013) 191047, [1307.0637].

[322] D. Baumann, D. Green and B. Wallisch, *Searching for Light Relics with Large-Scale Structure*, 1712.08067.

[323] CORE collaboration, J. B. Melin et al., *Exploring Cosmic Origins with CORE: Cluster Science*, *JCAP* **1804** (2018) 019, [1703.10456].

[324] J. Martin, C. Ringeval, R. Trotta and V. Vennin, *The Best Inflationary Models After Planck*, *JCAP* **1403** (2014) 039, [1312.3529].

[325] J. Martin, C. Ringeval and V. Vennin, *How Well Can Future CMB Missions Constrain Cosmic Inflation?*, *JCAP* **1410** (2014) 038, [1407.4034].