

This article has been accepted for publication in Monthly Notices of the Royal Astronomical Society ©: 2021 The Authors. Published by Oxford University Press on behalf of the Royal Astronomical Society. All rights reserved.

NuW CDM cosmology from the weak-lensing convergence PDF

Aoife Boyle ¹★, Cora Uhlemann ², Oliver Friedrich,^{3,4} Alexandre Barthelemy,¹ Sandrine Codis,^{1,5} Francis Bernardeau,^{1,5} Carlo Giocoli,^{6,7} and Marco Baldi ^{6,7,8}

¹CNRS & Sorbonne Université, UMR 7095, Institut d'Astrophysique de Paris, F-75014 Paris, France

²School of Mathematics, Statistics and Physics, Newcastle University, Herschel Building, Newcastle-upon-Tyne NE1 7RU, UK

³Kavli Institute for Cosmology, University of Cambridge, Cambridge CB3 0HA, UK

⁴Churchill College, University of Cambridge, Cambridge CB3 0DS, UK

⁵IPhT, DRF-INP, UMR 3680, CEA, Orme des Merisiers Bat 774, F-91191 Gif-sur-Yvette, France

⁶INAF – Osservatorio di Astrofisica e Scienza dello Spazio di Bologna, via Piero Gobetti 93/3, I-40129 Bologna, Italy

⁷INFN – Sezione di Bologna, viale Berti Pichat 6/2, I-40127 Bologna, Italy

⁸Dipartimento di Fisica e Astronomia, Alma Mater Studiorum Università di Bologna, via Piero Gobetti 93/2, I-40129 Bologna, Italy

Accepted 2021 May 4. Received 2021 April 16; in original form 2020 December 14

ABSTRACT

Pinning down the total neutrino mass and the dark energy equation of state is a key aim for upcoming galaxy surveys. Weak lensing is a unique probe of the total matter distribution whose non-Gaussian statistics can be quantified by the one-point probability distribution function (PDF) of the lensing convergence. We calculate the convergence PDF on mildly nonlinear scales from first principles using large-deviation statistics, accounting for dark energy and the total neutrino mass. For the first time, we comprehensively validate the cosmology dependence of the convergence PDF model against large suites of simulated lensing maps, demonstrating its per cent level precision and accuracy. We show that fast simulation codes can provide highly accurate covariance matrices, which can be combined with the theoretical PDF model to perform forecasts and eliminate the need for relying on expensive N -body simulations. Our theoretical model allows us to perform the first forecast for the convergence PDF that varies the full set of Λ CDM parameters. Our Fisher forecasts establish that the constraining power of the convergence PDF compares favourably to the two-point correlation function for a Euclid-like survey area at a single source redshift. When combined with a cosmic microwave background prior from Planck, the PDF constrains both the neutrino mass M_ν and the dark energy equation of state w_0 more strongly than the two-point correlation function.

Key words: methods: analytical – methods: numerical – cosmology: theory – large-scale structure of Universe.

1 INTRODUCTION

The path of light travelling from distant background galaxies is perturbed by matter density fluctuations along the line of sight. Distortions in the images of these background galaxies caused by weak gravitational lensing (reviewed by Kilbinger 2015) can be used to infer the convergence field and hence directly probe the projected matter density between the source galaxies and the observer. The subtle effects of weak gravitational lensing create a need for large statistics in order to constrain cosmological parameters, such as those acquired by the Dark Energy Survey (DES; DES Collaboration 2018), the Kilo-Degree Survey (KiDS; Heymans et al. 2021), and the Hyper Suprime Cam (HSC; Hikage et al. 2019). The future promises even larger and more powerful weak-lensing surveys such as Euclid (Laureijs et al. 2011; Amendola et al. 2018) and the Rubin Observatory (Ivezić et al. 2019).

The non-Gaussian evolution of the Universe at low redshifts and small scales implies that two-point statistics of the weak-lensing field cannot capture all of the information available. Several previous works have demonstrated that the information content of convergence

maps can be better captured when higher order statistics are included. Since it is difficult to find probes of higher order statistics whose cosmological dependence can be accurately derived, the bulk of analysis has relied on numerical simulations.

Petri et al. (2013) showed that adding Minkowski functionals or higher order moments to the weak-lensing convergence power spectrum can provide meaningful improvements in the constraints on cosmological parameters. Petri, May & Haiman (2016) showed that the non-Gaussian information contained in moments of the convergence field and convergence peak counts could also complement constraints from weak-lensing tomography. Vicinanza et al. (2018) confirmed that moments of the convergence field can break some cosmological degeneracies originating in the shear power spectrum. Peel et al. (2018) proposed the use of higher order weak-lensing statistics to break degeneracies between massive neutrinos and modified gravity $f(R)$ models in the weak-lensing power spectrum. Davies et al. (2020) recently examined the constraining power of weak-lensing voids alone in a Rubin-Observatory-like survey, and showed that these statistics outperform the shear-shear correlation function as a probe. Martinet et al. (2021) investigated the constraining power of weak-lensing aperture mass statistics, such as peaks, voids and the probability distribution function (PDF), also showing significant benefits over the shear two-point

* E-mail: boyle@iap.fr

correlation function. Zürcher et al. (2020) examined the benefits to constraints of peak counts, minimum counts or Minkowski functionals compared to the angular cosmic shear power spectrum, and confirmed that the non-Gaussian statistics could significantly enhance constraints in the Ω_m - σ_8 plane. A very recent study by Harnois-Déraps et al. (2020) found one of the tightest constraints on $S_8 = \sigma_8\sqrt{\Omega_m/0.3}$ in DES year-1 data based on a simulation-based combined peak count and correlation function analysis. Munshi et al. (2020a) recently demonstrated how the Minkowski functionals of convergence maps can be inferred using a fitting function to the bispectrum.

The full convergence PDF includes more information than can be accessed through a small set of moments, with the additional benefit that it can reliably modelled theoretically. Patton et al. (2017) performed a Fisher forecast for the one-point convergence PDF based on the L-PICOLA perturbative mocks for a DES-sized survey. They showed that although the degeneracy direction between Ω_m and σ_8 was very similar to the cosmic shear two-point correlation function, adding information from the PDF improved constraints by about a factor of two compared to two-point statistics alone. Additionally, the presence of measurement systematics was shown to tilt the degeneracy directions between the PDF and the power spectrum. Liu & Madhavacheril (2019) focused on the ability of the weak-lensing convergence PDF to help constrain the total neutrino mass (M_ν), performing a Markov chain Monte Carlo (MCMC) forecast for a Rubin-Observatory-like survey based on the MassiveNuS simulations. Their analysis included an examination of the benefits of using a range of source redshifts to perform weak-lensing tomography. They showed that the PDF alone could provide a 20 per cent stronger constraint on M_ν than the power spectrum, and that combining the two measurements could improve the constraint by 35 per cent over that of the power spectrum alone.

The one-point PDF of the convergence field therefore provides a natural complement to two-point measurements for extracting cosmological information in the mildly nonlinear regime. The statistics of weak-lensing fields can be difficult to model because of the mixing of different scales (with different degrees of nonlinearity) along the line-of-sight. It is well known as an empirical result that the weak-lensing PDF can be approximated as lognormal (Taruya et al. 2002; Hilbert, Hartlap & Schneider 2011), or by integrating over a lognormally distributed matter field (Xavier, Abdalla & Joachimi 2016). However, to extract cosmological information a more physically principled model of parametrizing the full cosmology dependence of the PDF is required. Having an accurate theory for predicting cosmological statistics presents many advantages. It allows for easier intuitive understanding of the cosmological dependence of the statistic, removes the need for analysis with expensive simulations, and greatly simplifies the performance of cosmological survey forecasts. In this work, we present the first comprehensive forecast for the lensing convergence PDF simultaneously varying all Λ CDM parameters, as well as individually considering the extensions of a free total neutrino mass M_ν and dark energy equation of state w . Note that beyond galaxy weak lensing, the cosmic microwave background (CMB) lensing PDF (Liu et al. 2016; Barthelemy, Codis & Bernardeau 2020b) and moments of the thermal Sunyaev-Zeldovich field (Hill & Sherwin 2013) have also shown promise in cosmological forecasts. Recently, the convergence skew-spectrum has also been computed for the weak lensing of galaxies and the CMB in Munshi et al. (2020b).

Large-deviation theory (LDT; for a general review see Touchette 2011) has successfully been applied to many cosmic large-scale structure fields in recent years. Bernardeau & Reimberg (2016)

clarified the applicability of the theory to the cosmological density field, demonstrating its connection to earlier works on the calculation of cumulants and modelling the matter PDF using perturbation theory and spherical collapse dynamics (Valageas 2002; Bernardeau, Pichon & Codis 2014a). In Uhlemann et al. (2016), this approach to calculating the PDF was greatly simplified with an analytical approximation and Codis et al. (2016) demonstrated its potential for constraining cosmology. Recently, the power of the matter PDF to constrain the neutrino mass or primordial non-Gaussianities was explored in Uhlemann et al. (2020) and Friedrich et al. (2020a), while Ivanov, Kaurov & Sibiryakov (2019) investigated a renormalization procedure for nonperturbative effects. Although the matter PDF is not a direct observable, it is the foundation of observable quantities like luminous tracer statistics (Repp & Szapudi 2020) and weak-lensing fields, with combined approaches like density-split statistics (Brouwer et al. 2018; Friedrich et al. 2018; Gruen et al. 2018) recently showing particular promise. Uhlemann et al. (2018b) applied the principles of LDT to densities in cylinders, and pointed out that LDT provides a natural basis for the joint modelling of the matter and weak-lensing fields, overcoming a fundamental limitation of the lognormal approximation.

Barthelemy et al. (2020a) presented an LDT-based prediction of the top-hat-filtered weak-lensing convergence PDF on mildly nonlinear scales building on early work by Bernardeau & Valageas (2000). The light-cone is divided into thin redshift slices, and the cumulant generating function (CGF) of the projected matter density in the resulting thin cylinders is related to that of the 3D matter density field in each slice and summed over. For this work, we adopt this approach (Barthelemy et al. 2020a), which is compatible with density-split statistics from Friedrich et al. (2018). A nulling method based on the BNT transform (Bernardeau, Nishimichi & Taruya 2014b) was implemented to remove contributions from very nonlinear small scales to significantly improve the accuracy of the theory. Note that Thiele, Hill & Smith (2020) recently presented an alternative analytic model for the weak-lensing convergence PDF using a halo-model-based approach. Their model targets a smaller scale, more non-Gaussian regime and is not accurate enough for the mildly nonlinear scales considered here, but their Fisher forecast shows promising results for constraining the neutrino mass in a Rubin-Observatory-like survey.

The purpose of this work is to demonstrate how well fundamental cosmological parameters can be constrained through measurements of the weak-lensing PDF. We validate the cosmology dependence of our model against convergence maps obtained from ray-tracing through the DUSTGRAIN-*cosmo* simulations (Giocoli, Baldi & Moscardini 2018) and the MassiveNuS simulations (Liu et al. 2018). After establishing the accuracy and precision of our model, we use it to perform a Fisher forecast for stage 3 cosmic shear surveys such as Euclid. Our forecast is the first for the convergence PDF to simultaneously vary the full set of Λ CDM parameters. We demonstrate that the convergence PDF outperforms the two-point correlation function in constraining cosmological parameters. In addition, we compare the information contained in the full PDF with that provided by a finite number of cumulants of the field. We focus on deploying the theoretical model for weak galaxy lensing at a single source redshift, and leave a tomographic analysis for future work.

Our paper is structured as follows: in Section 2, we describe our theoretical model for the smoothed weak-lensing convergence PDF, its sensitivity to cosmology, and how it can be extended to account for massive neutrinos. In Section 3, we validate our convergence PDF model using weak-lensing convergence maps

extracted from ray-traced N -body simulations. In Section 4 we perform a Fisher analysis for a Euclid-scale survey to quantify the potential cosmological constraints available from the weak-lensing convergence PDF, establishing its complementarity to the two-point correlation function and CMB data. We conclude and summarize potential avenues for future work in Section 5.

2 THEORETICAL MODEL FOR THE WEAK-LENSING CONVERGENCE PDF

In this section, we briefly summarize the key equations required for the calculation of the weak-lensing convergence PDF from large deviation theory following the formalism developed in Barthelemy et al. (2020a). A more complete overview of the derivation of these equations is given in Appendix A.

The lensing convergence κ can be calculated from a projection of the matter density contrast between the source and the observer along the line-of-sight. Using the Born approximation and neglecting lens–lens coupling, it can be written in terms of a redshift integral (Mellier 1999):

$$\kappa(\hat{\mathbf{n}}) = \int_0^{z_s} dz R'(z) \omega(z, z_s) \delta(z, R(z)\hat{\mathbf{n}}), \quad (1)$$

where $R(z)$ is the comoving distance to redshift z (equal to the comoving angular diameter distance $d_A(z)$ in a flat universe, as assumed throughout this work) and R_s is the comoving distance to the source redshift, z_s . The weight function ω is defined as

$$\omega(z, z_s) = \frac{3 \Omega_m H_0^2}{2 c^2} \frac{d_A(z) d_A(z_s, z)}{d_A(z_s)} (1+z), \quad (2)$$

and the comoving distance is a function of the Hubble parameter

$$R(z) = \int_0^z \frac{c}{H(z)} dz, \quad R'(z) = \frac{c}{H(z)}, \quad (3)$$

with the Hubble parameter being sensitive to cosmological parameters through the Friedmann equation. Note that for a flat universe, the lensing weight from equation (2) combined with $R'(z)$ is independent of the Hubble parameter today.

2.1 Cumulants and the PDF

The lensing convergence in a cone with a given opening angle θ , κ_θ , for sources at redshift z_s , can be viewed as a weighted integral of the density contrast δ_{cyl} in cylindrical slices of the cone with radii $d_A(z)\theta$ and depth L . In the small angle approximation it is possible to assume¹ that L is much larger than $d_A(z)\theta$. For standard cosmological power spectra,² the n -order cumulants of the density $k_n^{\text{cyl}}(d_A\theta, L, z)$ in such cylinders will scale like $1/L^{n-1}$. This is true in particular for the cylindrical variance $\sigma^2(d_A\theta, L, z)$, which reads

$$\sigma^2(d_A\theta, L, z) = \frac{1}{L} \int \frac{d^2 \mathbf{k}_\perp}{(2\pi)^2} P(k_\perp, z) W_{2D}^2(d_A\theta k_\perp), \quad (4)$$

where $W_{2D}(l) = 2J_1(l)/l$ is the Fourier transform of the cylindrical top-hat kernel, J_1 is the Bessel function of first order, \mathbf{k}_\perp are the

¹This assumption is valid as long as the scale of the radial variation of the lensing kernel function and the statistical properties of the density field is much larger than the transverse size of the beam. This is a generalization of the Limber approximation.

²In this context the cumulants are indeed dominated by modes whose wavelength is smaller or about the transverse size of the cylinder.

wave vector modes perpendicular to the line-of-sight and $P(k, z)$ is the matter power spectrum at redshift z .

More generally, the cumulant generating function (CGF) of the lensing convergence, $\phi_{\kappa, \theta}$, can be straightforwardly related to the CGF of the density in such cylindrical slices, $\phi_{\delta, \text{cyl}}$, as (Bernardeau & Valageas 2000)

$$\phi_{\kappa, \theta}(y) = \int_0^{z_s} dz R'(z) \phi_{\delta, \text{cyl}}(\omega(z, z_s) y, d_A(z)\theta, z), \quad (5)$$

where $\omega(z, z_s)$ is given in equation (2) and $\phi_{\delta, \text{cyl}}$ is evaluated for cylindrical slices of size $d_A(z)\theta$. A prescription for the CGF of the density in cylinders is provided by large deviation theory (LDT; for a more detailed outline see Appendix A), which describes the rate function Ψ_{cyl} (i.e. the leading behaviour of the logarithm of the PDF) as some driving parameter (here the variance) goes to zero. Formally, the CGF and the rate function are Legendre transformations of one another.

The determination of the rate function Ψ in the case of a cylindrical density relies on a key property of LDT known as the contraction principle. In this particular context, the contraction principle means that the appropriate rate function for the late-time density field can be computed from the initial conditions using only the most likely mapping between the initial and final densities $\rho = 1 + \delta$ in cylinders, which we assume to be the cylindrical collapse based on symmetry arguments (Valageas 2002; Barthelemy et al. 2020a). With this, the rate function for the late-time density in cylinders (assuming Gaussian initial conditions³) takes the form

$$\Psi_{\text{cyl}}(\rho) = \frac{\tau^2(\rho)}{2\sigma_l^2(r_{\text{ini}}, z)}, \quad (6)$$

where σ_l^2 is the linear cylindrical variance in the Lagrangian radius r_{ini} (through mass conservation, $r_{\text{ini}} = d_A\theta\rho^{1/2}$), and τ is the most likely linear density contrast corresponding to the final density ρ . The relevant equations for cylindrical collapse are summarized in appendix A of Friedrich et al. (2018); a robust approximation was provided by Bernardeau (1995):

$$\zeta(\tau) = \left(1 - \frac{\tau}{\nu}\right)^{-\nu} \Leftrightarrow \tau(\rho) = \nu \left(1 - \rho^{-\frac{1}{\nu}}\right). \quad (7)$$

The value of ν is set to 1.4, so as to reproduce the value of the tree-order skewness for cylindrical symmetry that is computed from perturbation theory (Uhlemann et al. 2018b). The linear cylindrical variance can be calculated from equation (4) when the power spectrum is replaced by the linear power spectrum $P_l(k, z)$, which can be obtained using Boltzmann codes like CLASS (Blas, Lesgourgues & Tram 2011).

Although the LDT results become exact only in the limit $\sigma \rightarrow 0$, previous work has shown that the formalism also performs well when extrapolated to non-zero values of the variance. This phenomenological extrapolation, beyond LDT, has to be performed in principle for each random variable – i.e. each redshift slice – and its corresponding driving parameter, namely the nonlinear variance at that redshift. This means that a modelling of the evolving nonlinear variance is needed as an external input to our formalism. In practice, a rescaling of the CGF has to be performed to ensure a correct nonlinear variance. We calculate the nonlinear κ variance, $\sigma_{\kappa, l}^2$, using equation (4) with the nonlinear power spectrum from Halofit (Bird, Viel & Haehnelt 2012; Takahashi et al. 2012). In Barthelemy et al.

³Primordial non-Gaussianities can also be straightforwardly accounted for in this formalism as shown by Uhlemann et al. (2018a) and Friedrich et al. (2020a).

(2020a), the correction for the nonlinear variance was performed by rescaling the CGF with the correct projected nonlinear κ variance, either from Halofit or as measured from simulations. In the present case, we modify this approach slightly and rescale the CGF in each redshift slice by the nonlinear variance at the scale $d_A(z)\theta$

$$\phi_{\delta, \text{cyl, nl}}(y) = \frac{\sigma_1^2(d_A(z)\theta)}{\sigma_{\text{nl}}^2(d_A(z)\theta)} \phi_{\delta, \text{cyl}} \left(\frac{\sigma_{\text{nl}}^2(d_A(z)\theta)}{\sigma_1^2(d_A(z)\theta)} y \right). \quad (8)$$

This is the same approach as adopted in Friedrich et al. (2018) and implemented in the publicly available code CosMomentum,⁴ with which we cross-check some of our results.

Once the CGF has been constructed, the convergence PDF is obtained using an inverse Laplace transform

$$P(\kappa) = \int_{-i\infty}^{+i\infty} \frac{dy}{2\pi i} \exp(-y\kappa + \phi_{\kappa, \theta}(y)). \quad (9)$$

The inverse Laplace transform requires an analytical expression for the CGF so as to be able to perform an analytic continuation in the complex plane. In Barthelemy et al. (2020a) this was achieved through the fitting of a double-power-law approximation to the linear variance. In this work, we instead move to a method first suggested in Bernardeau & Valageas (2000), which involves reframing the derived CGF in terms of an effective mapping $\zeta_{\text{eff}}(\tau)$ between an initial unsmoothed field and the final smoothed convergence field (for a detailed description see Appendix C or Barthelemy, Codis & Bernardeau 2021). A polynomial fit of the effective mapping is used to perform the inverse Laplace transform numerically.

2.2 Cosmology dependence

To get preliminary insights into the cosmology dependence of the PDF derived as above, one can examine the dependence of the cumulants of cylindrical densities and the lensing convergence on the cosmological parameters. In this section, we consider the cold dark matter fraction Ω_{cdm} , the power spectrum amplitude σ_8 and the dark energy equation of state w_0 . Specifically, we focus here on understanding the cosmology dependence of the PDF through the effects of the cosmological parameters on its second and third cumulants (the variance and κ_3), extending to the dark energy equation of state what was pioneered in Bernardeau, van Waerbeke & Mellier (1997). From equation (6), it is clear that the density PDF in the context of LDT is intimately related to both the scale-dependence of the linear variance and the dynamics of the cylindrical collapse. More precisely, different densities ρ scan the linear variance at scales $r\rho^{1/2}$ for a range of values around the radius r . This is encapsulated in the behaviour of the tree-order perturbation theory prediction for the reduced skewness, S_3 , of the density at scale r . In an EdS universe, this quantity is determined by the first logarithmic derivative of the linear variance (Bernardeau 1994, 1995)

$$S_3^{2D}(r) = \frac{\langle \delta^3(r) \rangle}{\langle \delta^2(r) \rangle^2} = \frac{36}{7} + \frac{3}{2} \frac{d \log \sigma_l^2(r)}{d \log r}. \quad (10)$$

In principle, the first term of this expression depends on the expansion rate of the Universe and therefore on the cosmological parameters. This dependence has been found to be very mild in practice (Bernardeau 1992), and we have verified that it makes no difference to our results by comparing with output of the CosMomentum code, which implements exact spherical collapse solutions. In the context




Figure 1. Solid lines: The response of the linear cylindrical variance σ_L^2 at $z = 0$ to different changes in cosmology (denoted by colours) from the fiducial cosmology. The increments used are those specified in Table 1 (using the smaller increment for Ω_{cdm}). Dashed lines: The corresponding differences in the reduced skewness S_3 as calculated using equation (10).

of weak-lensing observations, the main dependence on the expansion rate of the Universe will be encoded in the projection effects.

Note that $S_3^{2D}(r)$ does not depend on the overall amplitude of the density fluctuation σ_8 as its dependence cancels out in the logarithmic derivative. It depends however on the scale-dependence of the variance. In Fig. 1, we show how Ω_{cdm} impacts the linear variance at $z = 0$ and hence the reduced skewness. Note that the shape of the full PDF is also sensitive to a combination of the reduced kurtosis S_4 and higher order cumulants, which depend on higher order logarithmic derivatives.

As mentioned before, the main source of dependence on the expansion rate is encoded in the projection effects, as shown in equation (12) at the level of the CGF, or at the level of the cumulants as shown hereafter. In general, the cumulants of κ can indeed be related to the cumulants in line-of-sight cylinders. Viewing the projected density as the sum of densities in such cylinders, it can be shown that

$$k_n^\kappa(\theta, z_s) = \int_0^{z_s} dz R'(z) \omega^n(z, z_s) k_n^{\text{cyl}}(d_A(z)\theta, L, z) L^{n-1}, \quad (11)$$

where $k_n^{\text{cyl}}(d_A(z)\theta, L, z) = \langle \delta_n^{\text{cyl}}(d_A(z)\theta, L, z) \rangle_c$ are the cumulants of the 3D density contrast filtered in a cylinder of transverse size $d_A\theta$ and depth L , with a well-defined limit $L \rightarrow \infty$. We can rewrite the cumulants for cylinders in terms of their reduced counterparts, S_n^{cyl} (see also equation A3), which are then independent of L . At tree order, they are furthermore expected to be redshift-independent and given by their 2D expression (equation 10) so that

$$\begin{aligned} k_n^{\text{cyl}}(d_A\theta, L, z) &= S_n^{\text{cyl}}(d_A\theta, z) \sigma_{\text{nl}}^{2(n-1)}(d_A\theta, L, z) \\ &\simeq S_n^{2D}(d_A\theta) \sigma_l^{2(n-1)}(d_A\theta, L) D^{2(n-1)}(z). \end{aligned} \quad (12)$$

It is then possible to give the explicit dependence of k_3^κ on any cosmological parameter. More specifically, the upper panel of Fig. 2 shows the integrand for k_3^κ for various cosmologies. The lensing kernel for a source redshift $z_s = 2$ peaks around $z = 0.5$. The middle panel shows the fractional changes in the integrand induced by varying Ω_{cdm} . The plot dissects the different contributions for Ω_{cdm} : the contribution from the Ω_{cdm} factor in the lensing weight (equation 2; sparsely dotted line), which is counteracted by the change in the linear growth factor $D(z)$ (see equation 12; dot-dashed line) and the comoving distance $R(z)$ (see equation 3; tightly dotted line). The total change for Ω_{cdm} depends on the smoothing scale due to the impact of

⁴<https://github.com/OliverFHD/CosMomentum>




Figure 2. Upper panel: The integrand used to calculate the third cumulant of κ in equation (12) for various cosmologies for $\theta = 10'$. The solid lines correspond to increases in the corresponding cosmological parameters (by the smaller increments in Table 1) and the dashed lines to decreases. Middle panel: Fractional changes in the integrand induced by changing Ω_{cdm} for the two smoothing scales considered for our Fisher forecast. The solid and dashed lines correspond to increases and decreases in the cosmological parameters once again. The additional lines represent the various contributions to the overall differences (see the main text). Lower panel: Fractional changes induced in the integrand by increasing or decreasing w_0 , and by changing M_ν from 0 to 0.15 eV. The changes with w_0 are independent of smoothing scale (so only one set of lines are shown) and the changes with M_ν vary little with smoothing scale (compare the dark green and yellow lines). The M_ν integrand is almost constant with redshift. The dotted dark green line shows the effect of changing only the Ω_{m} factor in the integrand.

the cylindrical reduced cumulant S_3 and the scale-dependent variance σ_{cyl}^2 . The lower panel of Fig. 2 shows the fractional changes in the integrand that result from varying w_0 and M_ν . In this case, we see that the change induced by w_0 is independent of the smoothing scale. This is a shared property of all the moments, which allows us to disentangle the Ω_{cdm} and w_0 contributions from each other even at a single source redshift. Note that while changes in Ω_{cdm} and w_0 lead to a redshift-dependent change in the integrand of the moments, a change in σ_8 causes a redshift-independent enhancement or diminution. We also emphasize that while we have used the

linear variance here for illustrative purposes, our model consistently includes the nonlinear variance.

2.3 Extending the model to include massive neutrinos

Including massive neutrino cosmologies requires a few adjustments in the implementation of the theory. Massive neutrinos affect the PDF in multiple ways.

First, the total matter density fraction Ω_{m} now contains a contribution from Ω_ν . Ω_{m} enters into the lensing kernel through equation (2). This is by far the dominant effect, as can be seen in Fig. 2 by comparing the solid dark green line (the overall effect of M_ν on the k_3 integrand) and the sparsely dotted dark green line (the effect of changing Ω_{m} alone). Because the effect of changing Ω_{m} dominates, the effect of changing M_ν in Fig. 2 is relatively insensitive to redshift and smoothing scale. This creates a natural degeneracy with σ_8 .

Next, the scaling of the comoving distance $R(z)$ with redshift changes because the Hubble parameter $H(z)$ changes. The scaling of $H(z)$ is now more complicated, as massive neutrinos contribute with a radiative equation of state at early times and gradually transition into forming a contribution like cold dark matter at late times. For an analytic solution to the evolution of the energy density of neutrinos that can be inserted into the Friedmann equation, see Slepian & Portillo (2018). The simplest solution in this case is to use the Boltzmann code CLASS to extract the appropriate evolution of $H(z)$ and $R(z)$ with redshift.

The unique evolution of massive neutrinos from ultrarelativistic to non-relativistic particles over the course of the history of the Universe has a distinctive effect on the growth of structure. At early times, neutrinos free-stream out of gravitational perturbations and therefore do not contribute to the growth of structure, but they do cluster at late times. This leads to a distinctive scale-dependence in the formation of structures, even on linear scales, with the growth of structure on smaller scales being relatively suppressed. Ultimately, this means that the linear variance $\sigma_l^2(R, z)$ can no longer be evolved with redshift through multiplication by a scale-independent growth factor $D^2(z)$, as the growth factor becomes mildly scale-dependent. For practical implementation, the scale-dependence of the variance must now be independently measured in each redshift slice.

Finally, massive neutrinos make a small contribution to the cylindrical collapse mechanism (which we parametrize as in equation 7). LoVerde (2014) showed that for neutrino masses within realistic current bounds, the main effect of massive neutrinos on spherical collapse is to increase the collapse threshold by at most 1 per cent. The impact of such a change can be estimated by varying the parameter $\nu \propto \delta_c$, and Uhlemann et al. (2020) showed that such a change resulted in a change in the bulk of the 3D matter PDF of less than 1 per cent. For cylindrical collapse, we assume that the effect is also negligible, and find this approximation sufficient to achieve good accuracy for the bulk of the convergence PDF in the next section.

3 VALIDATING THE MODEL WITH SIMULATIONS

Before proceeding with a forecast on cosmological parameters, we perform an extensive validation of our theoretical model with simulations. We use sets of weak-lensing convergence maps from two suites of ray-traced N -body simulations to determine the accuracy of our theoretical PDFs and their Fisher derivatives: DUSTGRAIN-*cosmo* (see Appendix B1 and Giocoli et al. 2018) and MassiveNuS (see Appendix B2 and Liu et al. 2018). Both of those simulations produce convergence maps of relatively small size (25 deg² for

DUSTGRAIN-*cosmo* and 12.25 deg^2 for MassiveNuS). To complement the small patch maps, we also make use of a set of full sky lensing maps from Takahashi et al. (2017) at a fixed cosmology.

3.1 The fiducial model PDF

We smooth the DUSTGRAIN-*cosmo* κ maps with top-hat filters of size 50 and 70 pixels, corresponding to 7.32 arcmin and 10.25 arcmin, respectively. For MassiveNuS, the same scales correspond to exactly 25 and approximately 18 pixels. The smoothing is implemented as a direct convolution of the map with the smoothing kernel in the DUSTGRAIN-*cosmo* case, and using a harmonic-space convolution method for MassiveNuS, with the difference being purely practical based on the different storage formats used for the maps. After smoothing, we discard all pixels within a smoothing radius of the edge of a given patch. We measure the histogram of the smoothed κ maps in 101 linearly spaced bins for the range $\kappa \in [-0.01, 0.01]$. In principle, one might prefer a finer binning, that for example allows to accurately infer the moments from the binned PDF. However, one should consider that when extracting the Fisher information the Kaufman–Hartlap factor given by equation (20) below essentially includes a penalty for too many bins due to the finite number of realizations (here 256–500) from which the covariance is estimated.

For demonstrating the accuracy of the theory for realistic surveys, it is best to do a comparison of this type with full-sky simulations to avoid additional contributions to the errors from small patch sizes (see Appendix B5). We therefore compare our theoretical result with that from the simulations of Takahashi et al. (2017) in Fig. 3. We show a comparison of the mean PDF over 108 full-sky realizations (long dashed lines) and the theoretical PDF derived from our model using matching cosmological parameters (solid lines) for the two scales 7.32 (blue) and 10.25 (red) arcmin. Overall, the prediction and measurement are in very good agreement with the residuals being at a few percent level in the 2σ region around the mean. As expected, the theoretical prediction improves with increasing smoothing scale and the observed departures exhibit a clear signature of a correction to the predicted (tree order) skewness for the smaller scale.⁵

As demonstrated in Barthelemy et al. (2020a), the mixing of scales involved in integrating the density field along the line-of-sight limits the theoretical accuracy that can be expected. This scale-mixing problem can be eliminated by considering linear combinations of convergence maps at different source redshifts that allow for the ‘nulling’ of small scales and hence facilitate excellent agreement between theory and simulation. The PDF of the nulled convergence κ_{null} can be computed using the same formalism as above when replacing the weak-lensing weight from equation (2) by a linear combination of lensing kernels from three different source redshifts z_s^i (called the BNT transform)

$$\omega_{\text{null}}(z, \{z_s^i\}) = \sum_{i=1}^3 p_i \omega(z, z_s^i) \Theta(d_A(z_s^i, z)), \quad (13)$$

where Θ is the Heaviside step function and the weights p_i are chosen to cancel the contribution from the smallest redshifts and scales by ensuring $\omega_{\text{null}}(z < z_s^1) = 0$. We demonstrate this procedure on full-sky simulations from Takahashi et al. (2017) in Fig. 4, where the maps

⁵In an Edgeworth expansion of the PDF, the first non-Gaussian correction that appears is proportional to the rms fluctuation times the skewness times an Hermite function of the field of order 3. Any correction to the skewness hence introduces an H_3 modulation in the residuals, which is consistent with our measurements.




Figure 3. Upper panel: The lensing convergence PDF for the two smoothing scales used in our forecasts for source redshift $z_s = 2$ as measured from the mean over 108 realizations of full-sky simulated maps from Takahashi et al. (2017), compared to the corresponding theoretical prediction from large deviation theory (shown both with and without shape noise). Lower panel: Residuals between the theoretical and simulated noiseless PDFs. The error bars indicate the standard error on the mean computed from the 108 simulations. The vertical coloured lines define the bin ranges we use for our Fisher analysis. The shaded areas in the lower panel highlight 2 and 5 per cent ranges.

have been smoothed with the harmonic-space convolution method. As expected, the residual on the bottom panel is very small in this case, below the percent in the 2σ range around the mean, despite the PDF being quite non-Gaussian (upper panel). This illustrates the power of LDT to capture the non-Gaussian shape of the convergence PDF with high accuracy in the mildly nonlinear regime probed here thanks to the BNT transform. When applying exactly the same procedure to the maps from DUSTGRAIN-*cosmo* in Fig. 4 we find a residual disagreement of a few per cent. We suspect that this hints at potential small discrepancies between the different simulations and map-making approaches, whose further study is beyond the scope of this work. As for resolution effects, it is known that the 3D matter PDF at small smoothing scales is impacted at a few per cent level by a finite mesh resolution and a finite particle number (see e.g. fig. 5 in Uhlemann et al. 2020), which are expected to exacerbate the situation for projected lensing quantities that mix small and large scales. See Appendix B for details of the simulations and a discussion on the effects of small patch size maps. While Takahashi et al. (2017) relies on a full-sky healpix map generated from several runs of nested boxes with varying mass resolution (see fig. 1 in Takahashi et al. 2017), DUSTGRAIN-*cosmo* and MassiveNuS maps are created from a single simulation box using randomization and replication procedures to




Figure 4. PDFs of the nulled lensing convergence from 256 patches in the DUSTGRAIN-*cosmo* simulations, 1000 patches in the MassiveNuS simulations, and 108 full-sky simulations from Takahashi et al. (2017). Error bars are neglected because the question of the number of effective independent realizations for the DUSTGRAIN-*cosmo* simulations is unclear (see Appendix B1 for a discussion of how the simulations are run). It is clear that the simulations of Takahashi et al. and MassiveNuS perform well, but there may be some complications for the DUSTGRAIN-*cosmo* case that fall outside the scope of this work. The nulling source redshifts in each case are given in the captions of the lower panel.

cover a small light-cone patch. Nonetheless, we see the procedure works well when applied to the convergence maps from MassiveNuS.

Since, observationally, the weak-lensing convergence map is obtained from cosmic shear measurements and galaxies themselves are intrinsically elliptical, the observed shear contains a contribution from this intrinsic signal. Shape noise is caused by the variance of the intrinsic ellipticity, which is the dominant source of noise in shear measurements and impacts the convergence PDF (and its moments and cumulants) as if it was convolved with a Gaussian centred at zero with variance σ_{SN}^2 (Clerkin et al. 2016)

$$\mathcal{P}_\kappa^{SN}(\kappa) = \frac{1}{\sqrt{2\pi}\sigma_{SN}} \int_{\kappa_{\min}}^{\infty} dk' \exp\left(-\frac{(\kappa - \kappa')^2}{2\sigma_{SN}^2}\right) \mathcal{P}(\kappa'). \quad (14)$$

For the variance describing the shape noise we assume

$$\sigma_{SN} = \frac{\sigma_\epsilon}{\sqrt{n_g \cdot \Omega_\theta}} \propto \theta^{-1}, \quad (15)$$

where $\sigma_\epsilon = 0.30$ is the shape-noise parameter, n_g the galaxy density (here assumed to be 30 arcmin^{-2}) and Ω_θ the solid angle (in arcmin^2). For the simulated maps, we include shape noise by adding a white noise Gaussian random map to the ‘raw’ simulated convergence map before smoothing, keeping the seed fixed for a given realization when varying cosmologies. For the theoretical predictions, we convolve the

Table 1. The fiducial cosmology and parameter increments chosen for our calculations, largely chosen for consistency with the DUSTGRAIN-*cosmo* simulation suite. Those simulations (see Appendix B1 for more details) were performed with two step sizes for the three main parameters that were varied (Ω_{cdm} , σ_8 , and w) as shown. We do not have simulations that vary the parameters below the horizontal line (h , n_s , Ω_b), but calculate their effects theoretically in Section 4 using the step sizes shown.

Parameter	Fiducial value	Small Increment	Large Increment
Ω_{cdm}	0.26436	0.0125	0.1
σ_8	0.842	0.034	0.135
w	−1	0.04	0.16
M_ν (eV)	0	0.15	–
h	0.6731	0.005	–
n_s	0.9658	0.01	–
Ω_b	0.0491	0.0022	–

theoretical ‘raw’ convergence PDF at the desired smoothing scale with a Gaussian of the appropriate width. Since the area scales quadratically with the radius, the variance due to shape noise is inversely proportional to the smoothing scale. The standard deviation of the noiseless smoothed convergence scales approximately as $\sigma_\kappa \propto \theta^{-1/3}$ for a smoothing scale θ . For a source redshift of $z_s = 2$ and smoothing scales of $\theta_1 = 7.32 \text{ arcmin}$ and $\theta_2 = 10.25 \text{ arcmin}$, the cosmological signal dominates over the shape noise contribution. Note that due to the symmetry of this convolution, the third moment κ_3 is only very mildly affected.

In Fig. 3, we also demonstrate the impact of shape noise on the smoothed weak-lensing convergence PDF for the scales of interest here. As expected, the shape noise broadens the PDF compared to the noiseless case. Note that the addition of shape noise simultaneously decreases the PDF values around the peak (signal) and the PDF covariance (error).

3.2 The cosmology-dependence of the PDF

In this section, we compare the predicted response of the PDF to cosmological parameter changes with measurements from simulations. The PDF response enters the Fisher forecast in Section 4 in terms of derivatives with respect to cosmological parameters. The calculation of these theoretical derivatives requires several ingredients. For each cosmology (see Table 1), we generate linear and nonlinear (Halofit) matter power spectra from CLASS (Blas et al. 2011) at a range of redshifts between $z = 0$ and our source redshift $z_s = 2$. We use these to calculate the linear and nonlinear variance values as a function of redshift following equation (4). CLASS is also used to output $H(z)$ and $R(z)$ for the relevant redshifts. The density CGF is calculated for redshift slices between $z = 0$ and $z = 2$ in steps of 0.05 (see Appendix A for details of the calculation of $\phi_{\delta, \text{cyl}}$ from LDT). These CGFs are then rescaled with the nonlinear variance as outlined in equation (8), before being combined as in equation (5) and undergoing an inverted Laplace transform to produce the PDF.

The simulations with we compare our results in this section have produced very small-area maps (5 deg on a side for DUSTGRAIN-*cosmo* and 3.5° on a side for MassiveNuS). The variance with which we rescale our CGF must take this into account. In these cases, instead of rescaling by $\sigma_{\text{nl}}^2(d_A(z)\theta)$ at each redshift, we rescale with a projected variance derived from Halofit convergence angular power spectra. This allows us to straightforwardly apply a cut in l to exclude the large modes beyond the size of the maps with which we compare. This slightly rougher method is sufficient to validate




Figure 5. Derivatives of the lensing convergence PDF with respect to Ω_{cdm} , σ_8 , w_0 and M_ν computed for large/small steps in the simulation (dashed/dotted) and the theory for large steps (solid) at two smoothing scales, $\theta = 7.32$ arcmin and $\theta = 10.25$ arcmin. In the M_ν case, only a theoretical line is shown, and this is validated using the MassiveNuS simulations instead in Fig. 6. One can see that the small parameter steps in the simulation suffer from numerical noise, which is why we rely on the large steps to validate our theoretical predictions. When calculating the theory derivatives for validation purposes, we include a large-scale cut in l when calculating the variance at $l = 36$, corresponding to the patch size of DUSTGRAIN-cosmo. The vertical dotted lines correspond to the locations of the peaks in the fiducial PDFs. The coloured crosses indicate the extrema of the derivatives.

our derivatives in this section, as we will see. In our Fisher matrix calculations, no cut in l is implemented and the rescaling with the Halofit variance is carried out in each redshift slice individually, but we have validated that whether we rescale with the projected variance or the individual variance in each redshift slice makes no difference to the results, although the latter approach gives slightly better residuals (see Fig. 3).

In Fig. 5, we compare the derivatives of the PDF with respect to Ω_{cdm} , σ_8 and w_0 as obtained using the theory and the DUSTGRAIN-cosmo simulations at the two relevant scales. We also show a theoretical derivative for M_ν , although we cannot directly validate it with DUSTGRAIN-cosmo (we validate it by comparison with MassiveNuS instead in Fig. 6). We find the results to be in very good agreement. We note that with the number of realizations we have, the smaller step sizes provided for the simulations (see Table 1) result in significant numerical noise (see the dotted lines in Fig. 5). This kind of noise can result in exaggerated constraints by artificially breaking degeneracies between parameters. The larger step sizes provide much smoother derivatives, but we find that even those contain some small amount of noise that can artificially enhance their constraining power by a significant amount when inserted into a




Figure 6. Derivative of the lensing convergence PDF with respect to M_ν computed as a single-sided derivative using the difference between the fiducial cosmology from MassiveNuS (which has a total neutrino mass M_ν of 0.1 eV) and their corresponding cosmology without massive neutrinos. This particular comparison is shown only for validation purposes. We calculate the actual theoretical derivatives used in the Fisher matrix calculations in Section 4 for the cosmology given in Table 1. Here we include a cut in l at $l = 60$, corresponding to the MassiveNuS patch sizes. The smaller smoothing scale here is set to 7.38 arcmin instead of 7.34 arcmin simply because it corresponds to a whole number of pixels for the MassiveNuS resolution. The direction of the derivative is reversed compared to that in Fig. 5 because σ_8 is not fixed here and adding massive neutrinos lowers σ_8 . The dotted ‘ σ_8 matched’ lines show how the derivative looks if the PDF with massive neutrinos in the theoretical derivative is replaced by a PDF for a cosmology without massive neutrinos but with equal σ_8 ($\sigma_8 = 0.8295$ from Liu et al. 2018). The difference between the solid and dotted lines therefore corresponds to the orange line in Fig. 5. The vertical dotted lines correspond to the locations of the peaks in the fiducial PDFs (the leftmost is for the smaller scale).

Fisher matrix. This presents an obvious advantage to having a theory with which to model observables in such calculations. In addition, our validated complete cosmological model for the PDF allows us to extend our forecasts to parameters beyond those varied in particular simulations (see Section 4.3). Note that while the general shape of derivatives might look similar at a first glance, there are two hints for potential degeneracy breaking. If derivatives have either different zero crossings or different extrema locations, this means they are not proportional to each other. The first can be seen for example in the positive κ region for the red Ω_{cdm} and the blue σ_8 line. The second can be observed from the green line for w whose peak is at a different location from the rest of the dips.

In Fig. 6, we use the MassiveNuS simulations as a basis for determining how well our theory captures the effects of massive neutrinos on the PDF. The fiducial MassiveNuS cosmology (referred to as model 1b in Liu et al. 2018) has cosmological parameters $\Omega_{\text{m}} = 0.3$, $\sigma_8 = 0.8295$ and $M_\nu = 0.1$ eV. This is used with a model with the same Ω_{m} but without massive neutrinos, which has $\sigma_8 = 0.8523$ (model 1a) to calculate the difference shown in Fig. 6. All relevant cosmological parameters are provided in Table 1 of Liu et al. (2018) and its caption. We see from Fig. 6 that the effects of adding massive neutrinos to the cosmology are well captured by the theory.

3.3 The two-point correlation function

We measure the two-point correlation function from the unsmoothed kappa maps using the TreeCorr Python package (Jarvis, Bernstein & Jain 2004). We use 50 logarithmically spaced bins from 0.1 arcmin to 400 arcmin. To replicate the small patch areas of our two sets of maps, we cut low multipoles $l < l_{\text{cut}} = \pi/\theta_{\text{sim}}$. For our analysis, we discard

scales below $\theta_{\min} = 5$ arcmin to mitigate small-scale inaccuracies and artefacts resulting from our cut-off $l_{\max} = 4096$ when generating FLASK maps Appendix B4. Note that this restriction of scales is similar to the one adopted in Patton et al. (2017), where they performed an l -cut to equate the number of independent modes in the PDF and the two-point correlation function. In their case, the PDF smoothing scale of $\theta \simeq 14$ arcmin ($N_{\text{side}} = 256$) related to $l_{\max} = 886$ and a corresponding $\theta_{\min} \simeq \pi/l_{\max} \simeq 12$ arcmin. The cosmological derivatives for the two-point correlation function match those generated using Halofit very well and we use the Halofit derivatives as the input for our Fisher forecasts to avoid simulation noise and ensure compatibility between the PDF and two-point correlation function.

4 FISHER FORECAST FOR A EUCLID-LIKE SURVEY

In this section, we quantify the information content of the lensing convergence PDF on some key w CDM cosmological parameters using the Fisher matrix formalism. After having validated our theoretical predictions for the PDF and the Halofit two-point correlation function against simulation measurements, we perform our forecast using purely those inputs. This is important because, as mentioned in Section 3.2, we find that apparently negligible amounts of numerical noise in the PDF derivatives extracted from simulations can lead to artificially strong constraints.

The Fisher matrix for a set of cosmological parameters, θ , given a (combination of) statistics s is defined as

$$F_{ij} = \sum_{\alpha, \beta} \frac{\partial s_{\alpha}}{\partial \theta_i} C_{\alpha\beta}^{-1} \frac{\partial s_{\beta}}{\partial \theta_j}, \quad (16)$$

where s_i is element i of the statistic s and $C_{\alpha, \beta}$ are the elements of the covariance matrix of measurements \hat{s} of s , defined as

$$C_{\alpha\beta} = \langle (\hat{s}_{\alpha} - \langle \hat{s}_{\alpha} \rangle)(\hat{s}_{\beta} - \langle \hat{s}_{\beta} \rangle) \rangle. \quad (17)$$

where $\langle \cdot \rangle$ denotes an ensemble average.

The Fisher matrix allows us to determine the error contours on a set of cosmological parameters under the assumption that the likelihood is Gaussian. The inverse of the Fisher matrix gives the parameter covariance. The error on the parameter θ_i , marginalized over all other parameters, is given by

$$\delta\theta_i \geq \sqrt{(F^{-1})_{ii}}. \quad (18)$$

A Fisher analysis has a number of limitations: it only yields realistic error bars if measurements of the considered data vectors have Gaussian noise and if the responses of these data vectors to changing cosmological parameters are close to linear. We demonstrate that the distribution of individual bins of the PDFs measured in the simulations are sufficiently close to a Gaussian distribution in Fig. 7. We focus on the lowest and highest two κ bins we consider for the Fisher analysis (see the following section) and make a histogram of the fluctuations of the individual realizations around the mean in units of the measured variance. From this figure, we can expect a small impact of residual non-Gaussianity on the total width of our forecasted contours. Note that, in contrast, the moments of the lensing convergence can have strongly non-Gaussian distributions around the mean, in particular in small simulation patches as demonstrated in our discussion of moments-based constraints in Appendix A.

Fully realistic data analyses might have to account for systematic effects by marginalizing over additional nuisance parameters. However, the main focus of this study is to explore the complementarity




Figure 7. The distribution of realizations around the mean PDF in the lowest (red) and highest (blue) bins for the smaller smoothing scale (solid) and the larger scale (dashed). The black line shows a Gaussian of zero mean and unit variance, which is the expectation for Gaussian noise and a good approximation to all distributions.

between the two-point correlation function and the one-point lensing convergence PDF as cosmological probes. Hence, we expect these limitations to have limited impact on our findings.

4.1 Covariance matrices

In the following, the data vector s entering equation (16) will either consist of histogram bins of measurements of the convergence PDF, of measurements of the two-point correlation function of convergence in a set of angular bins or of a combination of both. In the cases where s includes PDF measurements we estimate the covariance matrix that enters equation (16) from a set of N_{sim} simulated measurements \hat{s}_i , $i = 1, \dots, N_{\text{sim}}$, as

$$\hat{C}_{\alpha, \beta} = \frac{1}{N_{\text{sim}} - 1} \sum_i (\hat{s}_{i, \alpha} - \bar{s}_{\alpha})(\hat{s}_{i, \beta} - \bar{s}_{\beta}), \quad (19)$$

where $\bar{s} = \sum_i \hat{s}_i / N_{\text{sim}}$ is the mean of all measurements. To compute the Fisher matrix with equation (16), we need to estimate the inverse covariance matrix or precision matrix. Since matrix inversion is a nonlinear operation, the noise in the above estimate of the covariance elements $\hat{C}_{\alpha, \beta}$ will lead to a bias in the elements of the precision matrix. Through a miracle of nature this bias is just a factor multiplying the entire matrix, the so-called Kaufman–Hartlap factor (Kaufman 1967; Hartlap, Simon & Schneider 2006),

$$h = (N_{\text{sim}} - 2 - N_s) / (N_{\text{sim}} - 1), \quad (20)$$

where N_s is the number of data points in s . Please note that the Kaufman–Hartlap factor (but also the more advanced treatment of Sellentin & Heavens 2016) only correct the width of parameter contours for bias from covariance estimation noise. This is sufficient for a Fisher analysis like the one presented here, but it is not sufficient for parameter estimation based on an actual measurement of s . In the latter case, covariance noise will not only impact the width of parameter contours but also introduce additional scatter to the contour location. This needs to be taken into account independently of the Kaufman–Hartlap factor, as e.g. described in Dodelson & Schneider (2013) or even more completely in Percival et al. (2014) (see also fig. 1 of Friedrich & Eifer 2018, for a visualization of this effect).

In Fig. 8, we show the correlation matrix

$$\text{Corr}_{ij} = \frac{C_{ij}}{\sqrt{C_{ii}C_{jj}}}, \quad (21)$$




Figure 8. Upper panel: The reduced cross-correlation matrix (equation 21) of the weak-lensing convergence PDF at radii $\theta_1 = 7.32$ arcmin and $\theta_2 = 10.25$ arcmin and source redshift $z_s = 2$ without shape noise. We display bins $[-0.024, 0.026]$ for θ_1 and $[-0.022, 0.022]$ for θ_2 , which corresponds to cutting [2 per cent, 5 per cent] probability in low/high κ . The black line indicates the diagonal. The lower triangle shows the result from the 256 κ maps created from one DUSTGRAIN-*cosmo* *N*-body simulation with randomization and the upper triangle shows the result from 500 independent full-sky maps created from FLASK.

for measurements of the convergence PDF at our two smoothing scales. This matrix visualizes the amount of correlation throughout the data vector elements. We compare the correlation matrices obtained from 256 maps from the DUSTGRAIN-*cosmo* simulations and from 500 lognormal maps generated using the publicly available FLASK tool (Xavier et al. 2016) in Fig. 8. In practice, we implement the latter in our Fisher matrix calculations, but we see that both methods are in reasonable agreement. Naturally, in both cases, the measured covariances are rescaled to the area of a Euclid-sized survey, assuming that the covariance is proportional to survey area. In the two blocks on the diagonal in Fig. 8 we see that, as expected, neighbouring PDF bins are positively correlated, while intermediate underdense and overdense bins are anticorrelated with each other. Note that the tails of the PDF, which are excluded in the plot, are strongly correlated with each other and anticorrelated with the peak. Additionally, we observe that PDFs at different scales are strongly but not perfectly correlated with each other. This is expected, as the two smoothing apertures partially (but not perfectly) overlap. The cross-correlations between PDF bins of different scales look very similar to the bin correlations for the individual PDFs, because the matter clustering changes mildly with radius (see fig. 5 in Uhlemann et al. 2017).

In all contours that only consider the convergence two-point function we use an analytic covariance model in order to circumvent noise (and potential small-scale numerical artefacts) associated with covariance estimation from FLASK. Our covariance model is based on the lognormal covariance model by Hilbert et al. (2011) but uses the procedure of Friedrich et al. (2018) to fix the parameters of the lognormal distribution as well as the analytic treatment of sky curvature and bin averaging that is detailed in Friedrich et al. (2020b). We find that our analytic covariance agrees well with the corresponding estimate from FLASK for the range of scales (bigger than a few arcmin) for which FLASK can be reasonably applied, see Fig. B3.




Figure 9. Fisher forecast constraints on Ω_{cdm} and σ_8 for a Euclid-like survey from the weak-lensing convergence PDF at $z_s = 2$. Contours are shown for the PDF measured for a single smoothing scale of $\theta_1 = 7.32$ arcmin (cyan), a single scale of $\theta_2 = 10.25$ arcmin (blue), and the two scales combined (purple). It is clear that the use of two smoothing scales significantly improves the constraints. The combined PDF contour is oriented along the line of constant $\Sigma_8 = \sigma_8 \sqrt{\Omega_m}$ (black dotted line).

4.2 Forecasted constraints for Ω_m , σ_8 , w_0 and M_ν

We now analyse the cosmology-constraining power of the weak-lensing convergence PDF and its complementarity to the two-point correlation function. Fig. 9 shows constraints from the convergence PDF for two individual smoothing scales and their combination. It is clear that combining two smoothing scales helps reduce the degeneracy between Ω_{cdm} and σ_8 because the Ω_{cdm} derivative is a function of the smoothing scale (see Fig. 2). We see that the degeneracy direction aligns roughly with the $\Sigma_8 = \sigma_8 \sqrt{\Omega_m}$ direction.

Fig. 10 compares the PDF constraints on Ω_{cdm} and σ_8 at two scales to other probes: the two-point correlation function with $\theta_{\text{min}} = 5$ arcmin and CMB-based constraints from Planck 2018 (Planck Collaboration VI 2020). The Planck constraints are extracted from the TTTEEE_lowl_lowE chains in the Planck legacy archive for the base cosmology. This corresponds to data from the Planck TT, TE and EE spectra at $l > 30$ plus low- l TT and EE spectra. In later figures, we use instead the base_w data set for the w CDM cosmology or base_mnu for the Λ CDM+ M_ν cosmology. The covariance matrices are generated specifically using Ω_{cdm}/h^2 as the Ω_{cdm} value for each run (i.e. not using the values provided for Ω_m). It is important to note that all the Planck constraints here come marginalized over the standard Λ CDM parameters as a minimum (including h , n_s , Ω_b) while our PDF and two-point correlation function Fisher matrices are marginalized only over the parameters shown in each contour plot.

We see that the constraints in Fig. 10 compare favourably with those from the two-point correlation function, although the degeneracy direction in both cases is similar. This is due to the fact that we have fixed all other cosmological parameters here. The degeneracy directions can differ as soon as one varies more parameters. We demonstrate this effect in Fig. 11, which extends the calculation to include constraints on w_0 and looks more promising. Here we

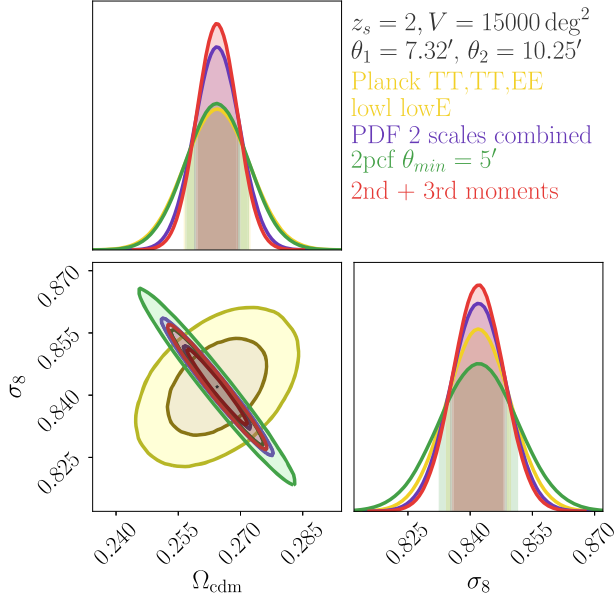


Figure 10. As Fig. 9, but presenting a comparison of the constraints from the PDF with two smoothing scales (purple) to constraints from other probes. Contours for the two-point correlation function with $\theta_{\min} = 5$ arcmin are shown in green. CMB-based constraints from the Planck 2018 TT,EE,EE+lowl+lowE data set are shown in yellow (see the text for details) and show a clearly complementary degeneracy direction. Constraints from the second and third moments alone (pink) slightly outperform the PDF. We note this is not a contradiction as we use a truncated PDF. Some details on the use of moments for constraints are given in Appendix A.




Figure 11. As Fig. 10, but extended to include the dark energy equation of state w_0 as an additional free parameter.

show that adding information from the PDF has the potential to break degeneracies between w_0 and both Ω_{cdm} and σ_8 , potentially tightening constraints considerably. We see something similar in Fig. 12, which shows simultaneous constraints on Ω_{cdm} , σ_8 and M_ν . It is not possible to extract meaningful constraints on M_ν using either CMB or large-scale structure information alone. But the PDF information content compares well with that from the two-




Figure 12. As Fig. 10, but extended to include the total mass of neutrinos M_ν as an additional free parameter.

point correlation function, and the degeneracy directions are shifted somewhat, allowing for potential degeneracy breaking by combining the two probes.

We can also use the Fisher matrix forecasting method to analyse the robustness of our results. For instance, we find that implementing the full differential equations for cylindrical collapse instead of the approximation given in equation (7) makes no difference to the results.

4.3 Extended parameter set constraints

The PDF derivatives for the parameters presented in Figs 9–12 have all been validated using simulations (see Section 3.2). It is reasonable to assume that it is therefore safe to extend the application of the model to other cosmological parameters to produce more quantitative results relevant for a realistic future surveys. While an in-depth analysis of the relevant degeneracies and a realistic redshift distribution is beyond the scope of this work, we analyse the key effects of varying the full set of cosmological parameters. In Table 2, we present forecasted cosmological constraints from the PDF marginalized over the standard Λ CDM parameters, and combined with the CMB prior from Planck described in the previous subsection.

As we aim for these results to represent more realistic constraints, we also choose to be more conservative with our shape noise parameters, adopting n_{gal} of 8 arcmin $^{-2}$ instead of 30 arcmin $^{-2}$, which approximately doubles the shape noise standard deviation. We do this because in practice the total number of galaxies would not be located at a single source redshift, but instead distributed across different redshift bins according to the source galaxy distribution $n_s(z_s)$. This can be included in our formalism by replacing the weight function from equation (2) by

$$\omega_{n_s}(z) = \frac{3 \Omega_m H_0^2}{2 c^2} \int dz_s \frac{d_A(z) d_A(z_s, z)}{d_A(z_s)} (1+z) n_s(z_s). \quad (22)$$

Overall, the constraints in Table 2 are very promising. The constraints in the Λ CDM case are comparable, but the PDF significantly outperforms the two-point correlation function once a free neutrino

Table 2. Upper panel: Percentage improvements in the 1σ constraints from Planck (TTTEEE_lowl_lowE) when adding the convergence PDF at two combined smoothing scales at $z_s = 2$ from a Euclid-like survey. Lower panel: Percentage improvements in the constraints from Planck combined with the PDF over constraints from Planck and the two-point correlation function. Note that we take a more conservative approach to shape noise here compared to in Figs 9–12 (see the text).

Planck+PDF versus Planck	Λ CDM	Λ CDM+ M_ν	w CDM
Ω_{cdm}	+55%	+43%	+65%
σ_8	+43%	+40%	79%
Ω_b	+52%	+37%	+67%
n_s	+25%	+24%	+20%
h	+53%	+39%	+74%
M_ν (eV)	–	+31%	–
w_0	–	–	+78%
Planck+PDF versus Planck+2pcf	Λ CDM	Λ CDM+ M_ν	w CDM
Ω_{cdm}	0%	+23%	+36%
σ_8	+8%	+32%	+35%
Ω_b	+0%	+26%	+37%
n_s	+0%	+2%	+4%
h	+0%	+25%	+36%
M_ν (eV)	–	+27%	–
w_0	–	–	+40%

mass or dark energy equation of state is introduced. We generally find that the resulting contours have quite similar (though not precisely identical) degeneracy directions, as was also seen by Patton et al. (2017), for example. In the Λ CDM case, strong constraints from Planck mean the advantage of the PDF over the two-point correlation function seen in Fig. 10 becomes less relevant. When additional free parameters that are poorly constrained by the Planck covariance matrix are introduced, this is no longer the case. The PDF can constrain M_ν over 25 per cent better and w_0 about 40 per cent better than the two-point correlation function. An examination of the individual Fisher matrices reveals that the advantage of the PDF lies primarily in its ability to better constrain σ_8 than the two-point correlation function. Of course, both M_ν and w_0 are anticorrelated with σ_8 , resulting in the improved constraints we see in Table 2.

Finally, we note that combining Planck with both the PDF and two-point correlation, we expect a modest improvement of 15 per cent on σ_8 constraints for base LCDM. For the extended models we see up to a 50 per cent improvement for most parameters in the w_0 CDM and ν LCDM cosmologies, except for n_s , which is still strongly constrained by Planck. In particular we noticed a 40–45 per cent improvement on σ_8 , M_ν , w_0 and h . Note that the improvement of the PDF over the two-point correlation tends to be reduced by the presence of shape noise. We expect that when repeating a similar analysis at lower source redshifts $z_s \sim 1$ with a larger smoothing scale ~ 20 arcmin (thus keeping the degree of nonlinearity roughly the same as for our current set-up), the impact of shape noise will be smaller and the PDF potentially even more powerful.

4.4 Comparison with previous work

There are several works in the literature that provide us with useful points of comparison for our results. The results of Patton et al. (2017) are perhaps the best suited to comparison with this work. They provide a simulation-based Fisher analysis of the PDF using the fast L-PICOLA perturbative mocks. Although the mocks are based on Lagrangian perturbation theory and not expected to fully capture the non-Gaussian convergence PDF, they confirm that competitive

constraints can be achieved from the convergence PDF compared to the cosmic shear power spectrum including an equivalent number of modes in each measurement. They extract their PDFs from a convergence field smoothed on a scale of approximately 13.7 arcmin, which is a little larger than the scales considered for our Fisher forecast. For their Fisher forecast, they vary the parameters $\{\Omega_m, H_0, \sigma_8\}$ and include a CMB prior on $\Omega_m = \Omega_m h^2$. They find that combining the PDF with the power spectrum improves the $\Omega_m - \sigma_8$ constraints by a factor of two despite the quite similar degeneracy directions between these two parameters for the two probes. They also highlight that limiting the convergence PDF to values of $\kappa > 0$ results in significant degradations in the constraints, and that the convergence PDF provides significant information that is not available in the peaks of the convergence field alone.

Liu & Madhavacheril (2019) performed a joint simulation-based analysis of the constraining power of the weak-lensing power spectrum and PDF for an Rubin Observatory-like survey, focusing particularly on constraining the neutrino mass. Their simulation-based MCMC forecast (using the same MassiveNuS simulation) also combines information from five source redshift bins. The authors focus on the PDF smoothed in Fourier space, on a scale roughly corresponding to 2 arcmin, which represents a more nonlinear regime than that considered in this work, and also give results focused on a single source redshift for $z_s = 1$. They also find that the PDF compares well with the power spectrum as a cosmological probe, finding that the PDF provides generally stronger constraints. They find that both combining the PDF and power spectrum and making use of a tomographic analysis can significantly improve constraints. For now we keep these extensions for future work.

A recent paper by Thiele et al. (2020) provides a nice complement to this work. The authors provide an analytical model for the weak-lensing convergence PDF based on the halo model, which is suited for smaller smoothing scales but not accurate enough on the mildly nonlinear scales considered in this work. The authors provide results from a Fisher forecast for a Rubin Observatory-like survey and show promising results for constraints on the neutrino mass when a CMB prior on A_s is included, and find reasonable agreement with the results of Liu & Madhavacheril (2019).

Friedrich et al. (2018) and Gruen et al. (2018) indirectly analysed the joint PDF of line-of-sight projected galaxy density and lensing in year-1 data of the Dark Energy Survey, using the technique of density split statistics. In this approach, the PDF of foreground galaxy counts is split into a set of density quantiles, and the strength of tangential gravitational shear around these quantiles is used to probe the underlying projected matter density PDF. Using the modelling of Friedrich et al. (2018) (which is based on principles related to the LDT ansatz presented here) Gruen et al. (2018) were able to simultaneously measure σ_8 , Ω_m , linear galaxy bias and galaxy-matter stochasticity and at the same time test for deviations from the Λ CDM predictions for the skewness of the matter density field. Their success with observational data is a proof of concept, indicating also that the results presented here carry over to real data analysis.

5 CONCLUSION

In this work, we have extended the theoretical model for the weak-lensing convergence PDF from Barthelemy et al. (2020a) to include massive neutrinos and the dark energy equation of state. We showed that its per cent level accuracy allows for robust calculations of Fisher derivatives for a range of cosmological parameters through validation with simulated lensing maps, see Figs 5 and 6. We found that Halofit (Bird et al. 2012; Takahashi et al. 2012) provides sufficiently

accurate nonlinear variances for the calculation of such derivatives. Furthermore, accurate covariance matrices for Fisher forecasts can be derived from fast simulation codes like FLASK. Therefore, theory-based forecasts as provided in this paper can easily be carried out without the need for expensive N -body/ray-tracing simulations. We have focused on the central region of the PDF, which has a Gaussian likelihood, making the Fisher matrix a natural tool for quantifying the cosmological information contained within it. This is to be contrasted with the first few moments of the convergence field, for which the likelihood exhibits significant non-Gaussian features (see Fig. A1).

Using our Fisher matrix forecasts, we have demonstrated the value of the weak-lensing convergence PDF as a cosmological probe, and its complementarity to the two-point correlation function and CMB data. We have shown that even with a single source redshift, the convergence PDF can provide useful information for constraining cosmology that outperforms that from the convergence two-point correlation function. Our theoretical model has allowed us to perform the first forecast for the weak-lensing convergence PDF that varies a complete set of Λ CDM parameters (see Table 2). Our comparison could be seen as quite conservative, as we focus on a single source redshift, although it is at quite high redshift ($z_s = 2$) due to limitations of the smoothing scales accessible in small aperture simulated maps. While lower source redshifts should produce more non-Gaussian information, they also push the model further into the nonlinear regime, which can be compensated by increasing the smoothing scale accordingly. The particular power of the PDF over the two-point correlation comes from its superior capacity to constrain σ_8 – this leads directly to better constraints on both M_v and w_0 , both of which are degenerate with σ_8 .

There are many natural extensions to this work. An obvious place to start would be with introducing multiple source redshifts to examine the constraining power of a tomographic analysis, which Liu et al. (2018) found produced significant improvements in forecasted constraints on M_v . In Barthelemy et al. (2020b), the model for the PDF outlined in this work was applied to the CMB lensing PDF, and the cosmological constraints achievable with such a formalism would be another interesting avenue for exploration, although the non-Gaussian information content is smaller so the PDF might be a relatively weaker measurement in this case.

Finally, we would like to comment on our choice to use an indirect observable, the convergence, rather than shear components for instance, with the risk of introducing biases in the conversion. The main advantage of basing a theoretical analysis like this one on weak-lensing mass maps is that they allow for easier intuitive understanding of the connection between results for the projected matter density and lensing. Nonetheless, we note that Patton et al. (2017) allowed for both multiplicative and additive shear biases in their analysis, and that they found qualitatively similar results with and without marginalization over these systematics, highlighting that the cosmological information from the PDF seems less sensitive to these biases than the power spectrum. Furthermore, we note that an analysis of one-point statistics for the DES has been performed based on moments of the convergence (Gatti et al. 2020) and shear profiles around density quantiles (Friedrich et al. 2018; Gruen et al. 2018). We also note that the tangential shear profile around the line of sight can be directly inferred from the convergence profile as in equation (II.6) of Friedrich et al. (2018). Pires et al. (2020) presented significant improvements on the Kaiser–Squires method with a novel mass-inversion method (KS+) that corrects for systematic effects, and pointed out that convergence maps could play an important role in the analysis for future surveys (highlighting non-Gaussianity studies as a specific example) because the lensing signal is more compressed

in convergence maps, leading to less computational expense. We also emphasize once again that the mean (which cannot be reconstructed) has been subtracted in the simulated mass maps we have used. As shown by Barthelemy et al. (2020a) it is possible to extend large deviation theory and the nulling principle to the aperture mass, which is much more tightly linked to measured shear data and we believe that the main conclusion of our investigations will not be changed with such more specific observables.

ACKNOWLEDGEMENTS

This work is partially supported by the SPHERES grant ANR-18-CE31-0009 of the French Agence Nationale de la Recherche and by Fondation Mercat. ABA is supported by a fellowship from CNES. CG acknowledges the grants ASI n. I/023/12/0, ASI n. 2018-23-HH.0, and PRIN 2015 Cosmology and Fundamental Physics: illuminating the ‘Dark Universe with Euclid’. CG is also supported by PRIN-MIUR 2017 WSCC32 ‘Zooming into dark matter and proto-galaxies with massive lensing clusters’. MB also acknowledges support by the project ‘Combining Cosmic Microwave Background and Large Scale Structure data: an Integrated Approach for Addressing Fundamental Questions in Cosmology’, funded by the PRIN-MIUR 2017 grant 2017YJYZAH.

This work has made use of the Horizon Cluster hosted by Institut d’Astrophysique de Paris. We thank Stéphane Rouberol for running smoothly this cluster for us. ABA and ABO thank the Institut de Physique Théorique for hosting several visits during the completion of this work. We thank the authors of Takahashi et al. (2017) and Liu et al. (2018) for making their lensing maps publicly available and providing help for analysing their maps. We also thank the Euclid higher-order weak-lensing statistics team for useful discussions.

DATA AVAILABILITY

There is no new data associated with this article.

REFERENCES

- Amendola L. et al., 2018, *Living Rev. Relativ.*, 21, 2
 Barthelemy A., Codis S., Uhlemann C., Bernardeau F., Gavazzi R., 2020a, *MNRAS*, 492, 3420
 Barthelemy A., Codis S., Bernardeau F., 2020b, *MNRAS*, 494, 3368
 Barthelemy A., Codis S., Bernardeau F., 2021, *MNRAS*, 503, 5204
 Bernardeau F., 1992, *ApJ*, 392, 1
 Bernardeau F., 1994, *A&A*, 291, 697
 Bernardeau F., 1995, *A&A*, 301, 309
 Bernardeau F., Reimberg P., 2016, *Phys. Rev. D*, 94, 063520
 Bernardeau F., Valageas P., 2000, *A&A*, 364, 1
 Bernardeau F., van Waerbeke L., Mellier Y., 1997, *A&A*, 322, 1
 Bernardeau F., Pichon C., Codis S., 2014a, *Phys. Rev. D*, 90, 103519
 Bernardeau F., Nishimichi T., Taruya A., 2014b, *MNRAS*, 445, 1526
 Bird S., Viel M., Haehnelt M. G., 2012, *MNRAS*, 420, 2551
 Blas D., Lesgourgues J., Tram T., 2011, *J. Cosmol. Astropart. Phys.*, 2011, 034
 Brouwer M. M. et al., 2018, *MNRAS*, 481, 5189
 Clerkin L. et al., 2016, *MNRAS*, 466, 1444
 Codis S., Pichon C., Bernardeau F., Uhlemann C., Prunet S., 2016, *MNRAS*, 460, 1549
 Davies C. T., Cautun M., Giblin B., Li B., Harnois-Déraps J., Cai Y.-C., 2020, *MNRAS*, preprint ([arXiv:2010.11954](https://arxiv.org/abs/2010.11954))
 DES Collaboration, 2018, *Phys. Rev. D*, 98, 043526
 Dodelson S., Schneider M. D., 2013, *Phys. Rev. D*, 88, 063537
 Friedrich O., Eifler T., 2018, *MNRAS*, 473, 4150
 Friedrich O. et al., 2018, *Phys. Rev. D*, 98, 023508

- Friedrich O., Uhlemann C., Villaescusa-Navarro F., Baldauf T., Manera M., Nishimichi T., 2020a, *MNRAS*, 498, 464
- Friedrich O. et al., 2020b, preprint (arXiv:2012.08568)
- Gatti M. et al., 2020, *MNRAS*, 498, 4060
- Giocoli C., Meneghetti M., Metcalf R. B., Ettori S., Moscardini L., 2014, *MNRAS*, 440, 1899
- Giocoli C., Baldi M., Moscardini L., 2018, *MNRAS*, 481, 2813
- Gruen D. et al., 2018, *Phys. Rev. D*, 98, 023507
- Harnois-Déraps J., Martinet N., Castro T., Dolag K., Giblin B., Heymans C., Hildebrandt H., Xia Q., 2020, *MNRAS*, preprint (arXiv:2012.02777)
- Hartlap J., Simon P., Schneider P., 2006, *A&A*, 464, 399
- Heymans C. et al., 2021, *A&A*, 646, A140
- Hikage C. et al., 2019, *PASJ*, 71, 43
- Hilbert S., Hartlap J., Schneider P., 2011, *A&A*, 536, A85
- Hilbert S. et al., 2020, *MNRAS*, 493, 305
- Hill J. C., Sherwin B. D., 2013, *Phys. Rev. D*, 87, 023527
- Ivanov M. M., Kurov A. A., Sibiryakov S., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 009
- Ivezic Ž. et al., 2019, *ApJ*, 873, 111
- Jarvis M., Bernstein G., Jain B., 2004, *MNRAS*, 352, 338
- Kaufman G. M., 1967, Report No. 6710, Center for Operations Research and Econometrics. Catholic University of Louvain, Heverlee, Belgium
- Kilbinger M., 2015, *Rep. Prog. Phys.*, 78, 086901
- Laureijs R. et al., 2011, preprint (arXiv:1110.3193)
- Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 538, 473
- Liu J., Madhavacheril M. S., 2019, *Phys. Rev. D*, 99, 083508
- Liu J., Hill J. C., Sherwin B. D., Petri A., Böhm V., Haiman Z., 2016, *Phys. Rev. D*, 94, 103501
- Liu J., Bird S., Matilla J. M. Z., Hill J. C., Haiman Z., Madhavacheril M. S., Petri A., Spergel D. N., 2018, *J. Cosmol. Astropart. Phys.*, 2018, 049
- LoVerde M., 2014, *Phys. Rev. D*, 90, 083518
- Martinet N., Harnois-Déraps J., Jullo E., Schneider P., 2021, *A&A*, 646, A62
- Mellier Y., 1999, *ARA&A*, 37, 127
- Munshi D., Namikawa T., McEwen J. D., Kitching T. D., Bouchet F. R., 2020a, preprint (arXiv:2010.05669)
- Munshi D., Namikawa T., Kitching T. D., McEwen J. D., Bouchet F. R., 2020b, *MNRAS*, 498, 6057
- Patton K., Blazek J., Honscheid K., Huff E., Melchior P., Ross A. J., Suchyta E., 2017, *MNRAS*, 472, 439
- Peel A., Pettorino V., Giocoli C., Starck J.-L., Baldi M., 2018, *A&A*, 619, A38
- Percival W. J. et al., 2014, *MNRAS*, 439, 2531
- Petri A., Haiman Z., Hui L., May M., Kratochvil J. M., 2013, *Phys. Rev. D*, 88, 123002
- Petri A., May M., Haiman Z., 2016, *Phys. Rev. D*, 94, 063534
- Pires S. et al., 2020, *A&A*, 638, A141
- Planck Collaboration VI, 2020, *A&A*, 641, A6
- Puchwein E., Baldi M., Springel V., 2013, *MNRAS*, 436, 348
- Repp A., Szapudi I., 2020, *MNRAS*, 498, L125
- Roncarelli M., Moscardini L., Borgani S., Dolag K., 2007, *MNRAS*, 378, 1259
- Sellentin E., Heavens A. F., 2016, *MNRAS*, 456, L132
- Slepian Z., Portillo S. K. N., 2018, *MNRAS*, 478, 516
- Szyszkowicz S. S., Yanikomeroğlu H., 2009, *IEEE Trans. Commun.*, 57, 3538
- Takahashi R., Sato M., Nishimichi T., Taruya A., Oguri M., 2012, *ApJ*, 761, 152
- Takahashi R., Hamana T., Shirasaki M., Namikawa T., Nishimichi T., Osato K., Shiroshima K., 2017, *ApJ*, 850, 24
- Taruya A., Takada M., Hamana T., Kayo I., Futamase T., 2002, *ApJ*, 571, 638
- Thiele L., Hill J. C., Smith K. M., 2020, *Phys. Rev. D*, 102, 123545
- Touchette H., 2011, preprint (arXiv:1106.4146)
- Uhlemann C., Codis S., Pichon C., Bernardeau F., Reimberg P., 2016, *MNRAS*, 460, 1529
- Uhlemann C., Codis S., Kim J., Pichon C., Bernardeau F., Pogosyan D., Park C., L'Huillier B., 2017, *MNRAS*, 466, 2067
- Uhlemann C., Pajer E., Pichon C., Nishimichi T., Codis S., Bernardeau F., 2018a, *MNRAS*, 474, 2853
- Uhlemann C., Pichon C., Codis S., L'Huillier B., Kim J., Bernardeau F., Park C., Prunet S., 2018b, *MNRAS*, 477, 2772
- Uhlemann C., Friedrich O., Villaescusa-Navarro F., Banerjee A., Codis S., 2020, *MNRAS*, 495, 4006
- Valageas P., 2002, *A&A*, 382, 412
- Vicinanza M., Cardone V. F., Maoli R., Scaramella R., Er X., 2018, *Phys. Rev. D*, 97, 043534
- Xavier H. S., Abdalla F. B., Joachimi B., 2016, *MNRAS*, 459, 3693
- Zürcher D., Fluri J., Sgier R., Kacprzak T., Refregier A., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 028

APPENDIX A: LARGE DEVIATION THEORY AND MOMENT-BASED CONSTRAINTS

A1 Large-deviation theory in a nutshell

In this appendix, we provide more background equations relevant to the derivation of the theoretical model presented in Section 2 and more specific details on the calculation of the matter density CGF in cylinders. For a more elaborate derivation we refer the reader to Barthelemy et al. (2020a).

In general terms, the cumulant generating function (CGF) ϕ_X of a continuous random variable X (such as the density contrast) is the logarithm of the moment generating function, which is the Laplace transform of the associated PDF

$$\phi_X(y) = \log(M_X(y)) = \sum_{n=1}^{\infty} k_n \frac{y^n}{n!}. \quad (\text{A1})$$

k_n are the cumulants (connected moments) of the distribution. In our theory, the mean, k_1 , is assumed to be zero. The scaled cumulant generating function (SCGF), φ_X ,

$$\varphi_X(y) = \lim_{k_2 \rightarrow 0} \sum_{n=0}^{\infty} S_n \frac{y^n}{n!} \quad (\text{A2})$$

is a related quantity that is key to large deviation theory and that defines the reduced cumulants S_n :

$$S_n = \frac{k_n}{k_2^{n-1}}. \quad (\text{A3})$$

The CGF and SCGF in the zero variance limit are straightforwardly related as $\varphi_X(y) = \lim_{k_2 \rightarrow 0} k_2 \phi_X(y/k_2)$. For random variables satisfying a large deviation principle, the SCGF is given through Varadhan's theorem as the Legendre–Fenchel transform of Ψ_X , which is called the rate function

$$\varphi_X(y) = \sup_X [yX - \Psi_X(X)]. \quad (\text{A4})$$

When Ψ is convex, this reduces to a simple Legendre transform and

$$\varphi_X(y) = yX - \Psi_X(X), \quad (\text{A5})$$

with the stationary condition defining the relationship between X and y

$$y = \frac{\partial \Psi_X}{\partial X}. \quad (\text{A6})$$

The crucial element of LDT is the contraction principle, which gives that the rate function for any variable that is a function of X can be determined assuming only the most likely mapping between the two variables. This is what allows us to define the rate function using only cylindrical collapse in Section 2. The final SCGF of the




Figure A1. The distributions of the third and fourth central moments taken from 256 25 deg^2 patches from FLASK-generated maps.

cylindrically-filtered density is given by

$$\varphi_{\text{cyl}}(y) = \sup_{\delta} \left[y\delta - \frac{\sigma_l^2(D\theta, L, z)}{2\sigma_l^2(D\theta\sqrt{1+\delta}, L, z)} \tau^2(1+\delta) \right]. \quad (\text{A7})$$

This can easily be re-written in terms of $\phi_{\text{cyl}}(y)$, which is the form that is implemented in our code. This also connects it to the overview given in Section 2, where we deal only in terms of ϕ for simplicity.

A2 Moments

In Fig. 10, we show that a large part of the information contained in the PDF comes from the second and third moments alone. This is to be treated with some caution, however, as we remove information from the PDF by truncating the tails, while the full distribution of values is included in our measurements of the moments. The moments-only constraints also benefit from a smaller Kaufman–Hartlap factor (see equation 20). In addition, constraints with a finite number of moments can only perform well when the total number of observables (in the case of Fig. 10, two scales and two moments, so four) is greater than the number of parameters being constrained.

Furthermore, we can highlight a particular advantage of performing Fisher matrix calculations with the bulk of the PDF instead of individual moments: the individual moments show very non-Gaussian distributions when measured from simulations with small patch sizes. Fig. A1 shows the distributions of the third and fourth moments measured at two smoothing scales measured from patches of FLASK-generated maps with side length 5° . This is not the case for the bulk of the PDF as used in this work (see Fig. 7).

APPENDIX B: SIMULATIONS AND NUMERICAL TOOLS

B1 DUSTGRAIN-*cosmo*

The DUSTGRAIN-*pathfinder* simulations (Giocoli et al. 2018) are a suite of cosmological simulations originally designed to sample several different combinations of modified gravity and massive neutrinos cosmologies – performed with the MG-GADGET N -body code (Puchwein, Baldi & Springel 2013) and subsequently post-processed with the weak-lensing map making code MapSim (Giocoli et al. 2014). The original DUSTGRAIN-*pathfinder* suite described in Giocoli et al. (2018) has been further extended to include a set of standard Λ CDM cosmologies with specific deviations of individual cosmological parameters from their fiducial values, as summarized

in Table 1. We call this new set of simulations the DUSTGRAIN-*cosmo* suite. The simulations follow the evolution of 768^3 dark matter particles in a periodic cosmological volume of $750 \text{ Mpc } h^{-1}$ per side, starting from initial conditions generated at $z = 99$ from a random realization of an initial matter power spectrum computed through the CAMB Boltzmann solver (Lewis, Challinor & Lasenby 2000) within the Zel’dovich approximation.

The light-cone is built using the MAPSIM routine (Giocoli et al. 2014) following the geometry illustrated in Fig. B1. The approach is based on using several snapshots from a single realization N -body simulation to build a light-cone up to $z = 4$. For DUSTGRAIN-*cosmo* there are 21 snapshots available. Given the box length of $750 \text{ Mpc } h^{-1}$, roughly seven (five) boxes are needed to cover the comoving distance of about 5 (3.6) $\text{Gpc } h^{-1}$ to $z_s = 4$ ($z_s = 2$). To obtain better redshift sampling, each simulation volume necessary for constructing the light-cone is divided along the line of sight into multiple redshift slices obtained from the individual snapshots. If the redshift slice reaches beyond the boundary of a single box, two lens planes are constructed from a single snapshot. The total number of lens planes up to $z_s = 4$ ($z_s = 2$) is 27 (19). To avoid replicating the same structure along the line of sight, the seven boxes needed to cover the light-cone are randomized. This randomization procedure allows to extract multiple realizations from a single simulation. Randomization is achieved by using seeds that act on the simulation boxes based on

- (i) changing the location of the observer, typically placed on the centre of one of the faces of the box,
- (ii) redefining the centre of the box (taking advantage of periodic boundary conditions), and
- (iii) changing the signs of the box axes.

Note that the same randomization process is applied to all redshift slices belonging to the same box to avoid spatial discontinuities (Roncarelli et al. 2007). A comparison of key statistics of the lensing field obtained from different routines based on replicating the original box volume along the line of sight can be found in Hilbert et al. (2020). Note that in the end the mean of the individual weak-lensing convergence maps is subtracted.

B2 MassiveNuS

The MassiveNuS simulation suite is described in Liu et al. (2018). It is based on N -body simulations for dark matter particles with a linear treatment of neutrinos in a box of side length $512 \text{ Mpc } h^{-1}$. Similarly to the case for DUSTGRAIN-*cosmo*, for a given cosmology, the light-cone is built from a single N -body realization with ray-tracing as implemented by LENS-TOOLS. The slices used to construct the light-cone have a constant thickness of $126 \text{ Mpc } h^{-1}$ and are taken from different randomizations of the single box at the given snapshot. This is different from the MapSim procedure adopted for DUSTGRAIN-*cosmo*, where the same randomizations are kept across different slices fitting in a single box (indicated as a common colour in Fig. B1). Large sets of maps (10 000 each) usable for validating Fisher derivatives are provided only for two cosmologies, with equal Ω_m and A_s but different values of M_ν and therefore of σ_8 . Note that the individual κ maps in MassiveNuS have non-zero means, which is one difference between these maps and those of DUSTGRAIN-*cosmo*.

B3 Takahashi full-sky lensing simulations

In Section 3.1, we use a PDF measured from the 108 full-sky lensing convergence maps provided by Takahashi et al. (2017) to validate




Figure B1. Illustration explaining how the DUSTGRAIN-*pathfinder* light-cones up to source redshift $z_s = 2$ are built from a single realization N -body simulation of box size $750 \text{ Mpc } h^{-1}$. To reach the source redshift of $z_s = 2$, the box is replicated 5 times along the light-cone up to $z_s = 2$. A randomization procedure is applied to (a) avoid replicating the same structure along the line of sight (same colour indicates same randomization) and (b) generate 256 realizations.

our theoretical model. These simulations are ideal for validating our model because they avoid any errors introduced by small patch sizes, but they are only available for a single fiducial cosmology, and so cannot be used to validate our cosmological derivatives. The method used to construct these lensing maps differs slightly from that of DUSTGRAIN-*cosmo* and MassiveNuS. The ray-tracing procedure is performed through a set of nested cubic simulation boxes around the observer. Each box has 2048^3 DM particles, so the resolution is higher in boxes closer to the observer. The final maps are generated with a resolution of 0.86 arcmin .

B4 FLASK

We use FLASK (Xavier et al. 2016) to generate a large number of lensing maps with which to compute the covariance matrix for our PDF Fisher matrix calculations. FLASK can generate lensing convergence maps using two possible approaches. The first is by assuming the convergence field can be well fit by a shifted lognormal distribution. The second approach instead simulates the density field as a function of redshift as a lognormal distribution, and obtains the convergence map by integrating along the line of sight. Here we make use of the first approach which is sufficient to obtain good agreement with the simulation measurements. FLASK is built to generate lensing convergence (or other) maps whose one-point distribution is close to a shifted lognormal (Clerkin et al. 2016; Hilbert et al. 2011)

$$\mathcal{P}_{\text{LN}}(\kappa) = \frac{\Theta(\kappa - \kappa_0)}{(\kappa - \kappa_0)\sqrt{2\pi}\sigma} \exp\left[-\frac{(\ln(\kappa - \kappa_0) - \mu)^2}{2\sigma^2}\right], \quad (\text{B1})$$

where the mean is $\mu = \ln(-\kappa_0 + \langle\kappa\rangle) - \sigma^2/2$. The variance of κ is related to the parameters as $\sigma_\kappa^2 = \exp(2\mu + \sigma^2)(\exp(\sigma^2) - 1)$.

The primary input required by FLASK is a set of lensing C_l generated using a Boltzmann solver like CLASS or CAMB, which we generate up to $l = 10000$. The shift parameter, λ , is then set to replicate the desired skewness (Xavier et al. 2016):

$$\lambda = \frac{\sigma}{\tilde{\mu}_3} \left(1 + y(\tilde{\mu}_3)^{-1} + y(\tilde{\mu}_3)\right) - \langle\kappa\rangle, \quad (\text{B2})$$

$$y(\tilde{\mu}_3) = \sqrt[3]{\frac{2 + \tilde{\mu}_3^2 + \tilde{\mu}_3\sqrt{4 + \tilde{\mu}_3^2}}{2}},$$

where $\langle\kappa\rangle$ is the desired mean, σ the target variance, and $\tilde{\mu}_3$ the skewness. We use FLASK to generate maps of HEALPIX resolution $N_{\text{side}} = 4096$, which corresponds to a pixel size of approximately

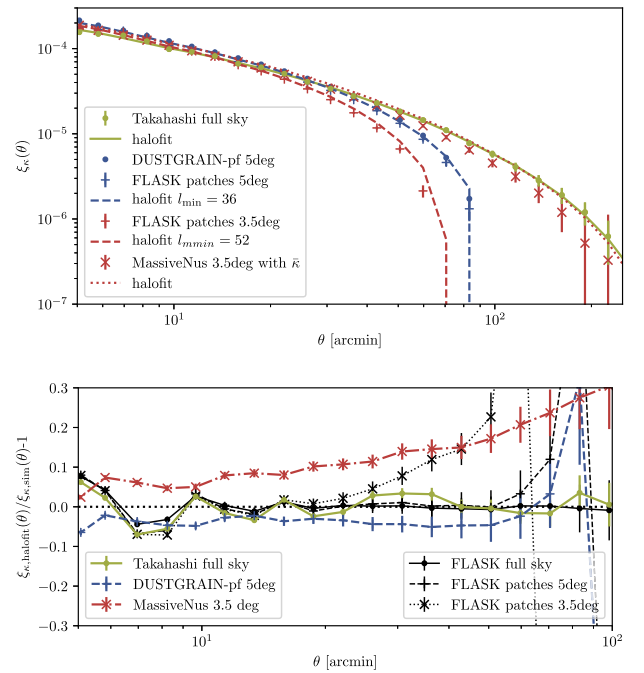


Figure B2. Upper panel: Logarithmic view of the κ correlation functions measured from different simulations (data points). The lines show the halofit predictions for full sky (solid and dotted line) and small patches of 5° and 3.5° with subtracted mean (dashed lines). To show the cut-off that is imposed by using small patches with subtracted mean $\bar{\kappa}$, we plot measurements from small patches cut from the FLASK full sky maps which have been generated with the halofit power spectrum. MassiveNuS maps have been obtained from maps including the mean $\bar{\kappa}$ instead which restores large-scale power. Lower panel: Residuals between the curves shown above.

0.86 arcmin , and then smooth to our required scales. In the end, we find that the mean FLASK PDF agrees with the theoretical one to within 1.25 per cent for all bins used in our Fisher analysis for both smoothing scales. Note that our procedure relies on a theorem by Szyszkowicz & Yanikomeroglu (2009). They show that the average of correlated lognormal random variables yields a random variable well described by a lognormal PDF with the same value of $\tilde{\mu}_3$. This allows us to impose our desired value of the skewness on the grid scale and obtain the same value also on larger smoothing scales.

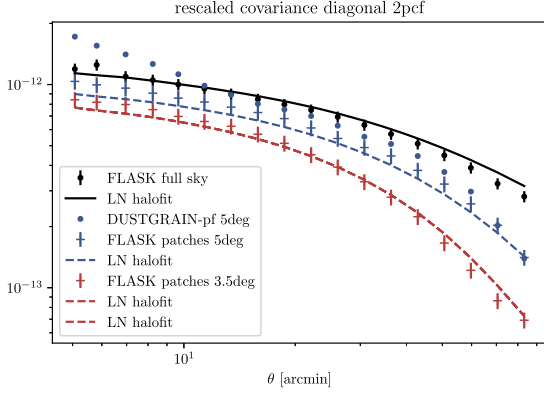


Figure B3. Variance of the κ two-point correlation function bins rescaled to a common sky area of $15\,000\text{ deg}^2$. We validate the theoretical prediction from the lognormal model (black line) against the measurements from full sky FLASK maps (black data points). We also show that considering small patches with mean subtraction (blue and red crosses) underestimate the covariance and are equivalent to imposing an l_{\min} cut appropriate for the different patch sizes in the theory (dashed lines). Those predictions are in broad agreement with the measurements from small patch simulations of DUSTGRAIN-*pathfinder* (blue points), with deviations driven by residuals of the two-point correlation with respect to Halofit (see Fig. B2).

B5 Limitations of measurements in small simulated maps

In Fig. B2, we show a comparison of the measured two-point correlation function in the Takahashi full-sky maps (Takahashi et al. 2017), as well as the small patches from the DUSTGRAIN-*pathfinder* simulations (Giocoli et al. 2018) and MassiveNuS (Liu & Madhavacheril 2019) to the Halofit prediction (Takahashi et al. 2012; Bird et al. 2012) as implemented in CLASS. Note that the convergence maps from the DUSTGRAIN-*pathfinder* simulations have been constructed to remove the mean $\bar{\kappa}$, while this mean is included in the MassiveNuS maps. The upper panel shows an overall good agreement between the measurements and the predictions when taking into account that a small map of side length L with mean subtraction requires to perform a cut in the l -range with roughly $l_{\min} \simeq 180^\circ/L$. The lower panel focuses on residuals and shows an exquisite agreement between Halofit and the full-sky maps from (Takahashi et al. 2017), but a few-percent discrepancies between Halofit and the two-point correlation functions measured from small maps of DUSTGRAIN-*pathfinder* and MassiveNuS. To establish a baseline and isolate potential errors coming from the finite range of $l \in [l_{\min}, l_{\max} = 20\,000]$, we also show residuals for κ maps from FLASK (generated from the Halofit C_l s) when considering a full sky or smaller patches (black lines). Note that the residuals seen in Fig. B2 are very consistent across all DUSTGRAIN-*cosmo* cosmologies, which ultimately means that they cancel out in the derivatives and we find very good agreement between simulation and theory.

In Fig. B3, we demonstrate that the covariance of the two-point correlation measured from small patches significantly underestimates the covariance from full-sky maps. This can artificially improve the constraining power of the two-point correlation function (or equivalently the power spectrum C_l) compared to higher-order statistics, which can be less affected from small patch sizes. As we demonstrated in the main text, the covariance of the κ PDF bins are largely independent of the patch size. Hence we caution that analyses based on covariances from small patches can underestimate the constraining power of the convergence PDF compared to the two-point correlation.

APPENDIX C: PERFORMING THE LAPLACE TRANSFORM USING AN EFFECTIVE MAPPING

To perform the inverse Laplace transform in equation (9), one can apply a method first described in Bernardeau & Valageas (2000). We assume the difference between the unsmoothed Gaussian initial field and the evolved, smoothed field can be combined into a single effective mapping, $\zeta(\tau_{\text{eff}})$. In this case, the SCGF (see equation A5) can be written as

$$\varphi(y) = y\zeta(\tau_{\text{eff}}) - \frac{1}{2}\tau_{\text{eff}}^2. \quad (\text{C1})$$

The stationary condition (equation A6) then becomes

$$y = \frac{d}{d\zeta} \frac{\tau_{\text{eff}}^2}{2} = \tau_{\text{eff}} \left(\frac{d\zeta(\tau_{\text{eff}})}{d\tau_{\text{eff}}} \right)^{-1}. \quad (\text{C2})$$

The effective mapping $\zeta(\tau_{\text{eff}})$ can now be written in the form of a vertex generating function

$$\zeta(\tau_{\text{eff}}) = \sum_{k=0}^n \frac{\mu_k}{k!} \tau_{\text{eff}}^k, \quad (\text{C3})$$

where $\mu_0 = 0$, $\mu_1 = 1$. The remaining coefficients can be fitted to the derived SCGF we want to transform, and this fit can then easily be continued into the complex plane. To fit the values of μ , we must express both $\zeta(\tau_{\text{eff}})$ and τ_{eff} in terms of the variable y . This can be achieved using equation (A5):

$$\frac{d\varphi(y)}{dy} = \zeta(\tau_{\text{eff}}), \quad (\text{C4})$$

and from equation (C1) therefore

$$\frac{1}{2}\tau_{\text{eff}}^2 = y \frac{d\varphi(y)}{dy} - \varphi(y). \quad (\text{C5})$$

We find that the fitting of μ using a table of τ_{eff} and $\zeta(\tau_{\text{eff}})$ values performs well with a polynomial of order 9. In practice one can also perform the fit for the CGF $\phi(y)$ instead, with the μ values being linked as $\mu_k^{\text{CGF}} = \sigma^k \mu^k$. Ultimately, this allows for straightforward numerical computation of the inverse Laplace transform.

Technically the convergence itself does not satisfy a large deviation principle, unlike the projected density in slices used to build up the model used here. However, this effective mapping approach still works well for the convergence because the form of the bulk of the PDF is dominated by the first few cumulants and its unimodality, while the assumption of LDT only affects the PDF around the critical point quite far into the tails. Barthelemy et al. (2021) tried applying the effective mapping approach to individual redshift slices, but found the result was not significantly different to that obtained by applying the effective mapping approach directly to the convergence CGF.