Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **553,** 1–20 (2026)
Advance Access publication 2026 October 8

https://doi.org/10.1093/mnras/stag1905

# Cosmological constraining power of the redshifts, heights, and angular clustering of weak gravitational lensing peaks

Jeger C. Broxterman [1,2★] Matthieu Schaller [1,2] Ian G. McCarthy [3] Willem Elbers [4] John Helly,[4] Henk Hoekstra [2] Konrad Kuijken,[2] Jaime Salcido [3] Joop Schaye [2] Naomi Schutte[2] and Elena Sellentin[2,5]

[1]*Lorentz Institute for Theoretical Physics, Leiden University, PO Box 9506, NL-2300 RA Leiden, the Netherlands*
[2]*Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands*
[3]*Astrophysics Research Institute, Liverpool John Moores University, Liverpool L3 5RF, UK*
[4]*Department of Physics, Institute for Computational Cosmology, University of Durham, South Road, Durham DH1 3LE, UK*
[5]*Mathematical Institute, Leiden University, PO Box 9512, NL-2300 RA Leiden, the Netherlands*

Accepted 2026 October 5. Received 2026 September 15; in original form 2026 July 31

## ABSTRACT

Weak gravitational lensing (WL) peaks probe non-Gaussian information of the large-scale distribution of matter that is not captured by two-point lensing statistics. We study the cosmological potential of the height distribution, redshift distribution, and angular clustering of high-valued WL peaks using a Bayesian inference approach that mimics a *Euclid* analysis. We use a forthcoming dark-matter-only hypercube, which varies cosmology in a ten-dimensional space, including evolving dark energy, neutrino mass, decaying dark matter, and the running of the scalar spectral index. We find the individual WL peak statistics to be complementary, as the redshift distribution best constrains the matter density, $\Omega_\mathrm{m}$, and the dark energy equation-of-state parameters, $w_0$, and $w_\mathrm{a}$; the height distribution and angular clustering are most sensitive to the amplitude of the primordial power spectrum, $\ln(10^{10}A_\mathrm{s})$; while combining the three statistics allows us to also probe the baryon density, $\Omega_\mathrm{b}$, and the Hubble parameter, $h$. A comparison to the shear two-point correlation function demonstrates that WL peaks alone outperform the commonly used statistics, while even tighter constraints are obtained when combining all. Considering the two-dimensional $\Omega_\mathrm{m}-\ln(10^{10}A_\mathrm{s})$ and $w_0-w_\mathrm{a}$ parameter planes, we find that the redshift distribution outperforms the angular clustering and peak-height distributions. We study the impact of the smoothing scale and find that, typically, the smallest scales yield the best results and the figure of merit improves by a factor of $\approx 2$ when combining multiple scales.

**Key words:** gravitational lensing: weak – methods: numerical – large-scale structure of Universe – cosmology: theory.

## 1 INTRODUCTION

The concordance model of cosmology, Lambda cold dark matter ($\Lambda$CDM), is able to describe a plethora of observations with only a very limited number of free parameters and main model ingredients, namely dark energy (DE; $\Lambda$ if its energy density is constant with cosmic time), CDM, baryons, and inflation (for a recent review, see e.g. O. Lahav & A. R. Liddle 2024). However, as current surveys become more precise, hints of deviations from this model are emerging. Specifically, recent Dark Energy Spectroscopic Instrument (DESI) results, which constrain cosmology through baryonic acoustic oscillations (BAOs) and the clustering of galaxies, quasars, and the Lyman-alpha forest, provide hints for evolving DE (M. Abdul Karim et al. 2025; A. G. Adame et al. 2025).

Similarly, in its recent Year 6 inference, the Dark Energy Survey (DES), following the Chevallier–Polarski–Linder (CPL) $w_0-w_\mathrm{a}$ parametrization (M. Chevallier & D. Polarski 2001; E. V. Linder 2003), finds that the best-fitting equation of state of DE slightly deviates from the $\Lambda$CDM expectation. Although their best-fitting measurement from the combination of weak gravitational lensing (WL) and galaxy clustering ($3\times2$ pt) is still consistent with no evolving DE, when they add their BAO and supernova (SN) measurements, they find a $2.2\sigma$ deviation from a cosmological constant (DES Collaboration 2026b).

One of the main cosmological probes for studying the time evolution of DE is cosmic shear, the coherent distortion of the images of background galaxies caused by gravitational lensing by the large-scale structure of the Universe (see e.g. H. Hoekstra & B. Jain 2008). As the Stage III lensing surveys, the Kilo-degree Survey (KiDS; J. T. A. de Jong et al. 2013), DES (The Dark Energy Survey Collaboration 2005), and Hyper Suprime-Cam (S. Miyazaki et al. 2012), are in the final stages of their (extended) cosmological analyses (e.g. R. Reischke et al. 2026),

★ E-mail: broxterman@lorentz.leidenuniv.nl

© The Author(s) 2026.
Published by Oxford University Press on behalf of Royal Astronomical Society. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

the next generation is already underway as *Euclid* (Euclid Collaboration: Mellier et al. 2025) is carrying out its first cosmological inference, *Rubin*'s Legacy Survey of Space and Time has started (LSST Science Collaboration 2009), and *Roman* has launched (D. Spergel et al. 2015). Combined, these instruments are expected to exhaust the sky best suited to cosmic shear, while combining deep observations with high-quality imaging.

Stage III surveys have consistently measured lower $S_8 \equiv \sigma_8(\Omega_m/0.3)^{1/2}$ values than inferred by *Planck*. Although the extent of these differences is smaller than previously reported (for a recent review, see I. Pantos & L. Perivolaropoulos 2026), next-generation surveys are needed to conclusively address these discrepancies and identify their origins. *Euclid* alone is expected to provide per cent-level constraints on the equation-of-state parameters of DE (Euclid Collaboration: Cañas-Herrera et al. 2025). Comparing the results of these different probes and surveys should help us understand the minor discrepancies currently observed. Extracting all cosmological information from the observed data will be key in resolving these differences.

The common statistics used by cosmic shear surveys to constrain cosmology are two-point projections of the matter power spectrum (see e.g. A. H. Wright et al. 2025; DES Collaboration 2026a), which quantify the variance of the matter density contrast field as a function of scale. However, more precise constraints may be obtained by including additional observables that probe non-Gaussian information not captured by traditional two-point measurements (e.g. Euclid Collaboration: Ajani et al. 2023, Euclid Collaboration: Vinciguerra et al. 2026). Examples of these statistics are the WL scattering transform (e.g. S. Cheng et al. 2020), Minkowski functionals (e.g. G. A. Marques et al. 2019), Betti numbers (e.g. J. Feldbrugge et al. 2019), wavelet phase harmonics (E. Allys et al. 2020), persistent homology (e.g. S. Heydenreich, B. Brück & J. Harnois-Déraps 2021), or three-point statistics (e.g. L. Linke et al. 2026).

This paper focuses on another higher-order weak lensing statistic, WL peaks, the local maxima in WL convergence fields (e.g. B. Jain & L. Van Waerbeke 2000; S. Wang, Z. Haiman & M. May 2009), which typically correspond to the most massive structures in the Universe that have experienced the strongest non-linear collapse (T. Hamana, M. Takada & N. Yoshida 2004; X. Yang et al. 2011). It has been shown that these peaks contain cosmological information in their height distribution (J. P. Dietrich & J. Hartlap 2010; J. M. Kratochvil, Z. Haiman & M. May 2010; J. C. Broxterman et al. 2024), steepness (Z. Li, X. Liu & Z. Fan 2023), angular clustering (L. Marian et al. 2013; C. T. Davies, M. Cautun & B. Li 2019), and redshift distribution (J. C. Broxterman et al. 2025). At the same time, the measurement is only first-order in the observed shear and could therefore respond differently to systematics than two-point statistics. Combined, the different WL peak properties could probe information similar to that from traditional galaxy cluster cosmological inferences, as they are sensitive to the mass and location of massive structures. One major advantage is that the peaks are WL- and thus mass-selected and therefore do not suffer from selection effects that require assumptions about the dynamical state of massive structures, or the form and scatter of mass-to-observable scaling relations, which typically limit X-ray- or thermal-Sunyaev Zel'dovich-selected cluster abundance inferences (see e.g. R. Kugel et al. 2024; I. Marini et al. 2024; R. Kugel et al. 2025).

In a cosmological forecast, D. Zürcher et al. (2021) showed that for a stage-III-like WL survey, including the peak height distribution can already increase the figure of merit (FoM) in the $\Omega_m - \sigma_8$ plane by a factor of 4 compared to the angular power spectrum alone. Similarly, C. T. Davies et al. (2022) find that the angular clustering of WL peaks is significantly more sensitive to the cosmological parameters $h$ and $w_0$ than the peak number density distribution, and when the two properties are combined, the constraining power on $\Omega_m$, $S_8$, $h$, and $w_0$ more than doubles. In this paper, we aim to extend this work by conducting a cosmological forecast for high-signal-to-noise ratio (SNR) WL peaks for a Stage IV WL cosmological inference, combining different WL peak properties to study their complementarity. We do this by using a new set of simulations in a ten-dimensional (10D) cosmological parameter space that includes evolving DE, decaying dark matter (DDM), and a running scalar spectral index in $(700\ \mathrm{cMpc})^3$ volumes (McCarthy et al., in preparation). We include the number densities, redshift distribution, and angular clustering of high-valued WL peaks.

In Section 2, we introduce the relevant WL theory to infer WL signals from a general matter distribution. In Section 3, we introduce the hypercube simulations. In Section 4, we demonstrate how WL signals are inferred from the simulation, which peak signals are selected, and present the details of the Bayesian likelihood forecast. The results and discussion are given in Section 5, and the final conclusions in Section 6.

## 2 THEORY

In this section, we summarize the main equations to estimate the WL signal from a matter density distribution. For extended reviews of WL and cosmic shear, see M. Bartelmann & P. Schneider (2001) and M. Kilbinger (2015). We assume a spatially flat universe such that the comoving angular diameter distance, $f_K(\chi)$, reduces to the comoving distance, $\chi$. In WL, photon deflections are described via the lensing equation, which relates the unlensed angular position, $\boldsymbol{\beta}$, to the observed angular position $\boldsymbol{\theta}$ as

$$\boldsymbol{\beta} = \boldsymbol{\theta} - \nabla_{\boldsymbol{\theta}}\,\psi(\boldsymbol{\theta}, \chi), \tag{1}$$

where $\psi$ is the lensing potential. We assume the Born approximation, i.e. we evaluate the lensing quantities on unperturbed photon paths, such that $\psi$ is given by

$$\psi(\boldsymbol{\theta}, \chi) = \frac{2}{c^2}\int_0^{\chi} \mathrm{d}\chi'\,\frac{\chi - \chi'}{\chi\chi'}\,\phi(\boldsymbol{\theta}\,\chi', \chi'), \tag{2}$$

where $c$ is the speed of light in vacuum and $\phi$ is the Newtonian gravitational potential. Conventionally, the lens equation is linearized to first order in the lensing potential, and the effect is quantified in terms of the magnification matrix, $A$, as

$$\frac{\partial\beta_i}{\partial\theta_j} \equiv A_{ij} = \delta_{ij} - \partial_i\partial_j\psi, \tag{3}$$

where the magnification matrix is decomposed as

$$A = \begin{pmatrix} 1-\kappa-\gamma_1 & -\gamma_2 \\ -\gamma_2 & 1-\kappa+\gamma_1 \end{pmatrix}, \tag{4}$$

where $\gamma = \gamma_1 + i\gamma_2$ is the WL shear and $\kappa$ is the WL convergence, which is used here to quantify the cosmic shear strength. Equations (2)–(4) can be combined to express the WL convergence as

$$\kappa(\boldsymbol{\theta}, \chi) = \frac{1}{c^2}\int_0^{\chi} \mathrm{d}\chi'\,\frac{\chi - \chi'}{\chi'\chi}\nabla^2_{\boldsymbol{\theta}}\phi(\chi'\boldsymbol{\theta}, \chi'). \tag{5}$$

We use the three-dimensional (3D) Poisson equation to estimate the deflection strength

$$\nabla^2\phi = 4\pi\overline{\rho_m}a^2G\delta \approx \chi^{-2}\nabla^2_{\boldsymbol{\theta}}\phi, \tag{6}$$



Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

where $G$ is the gravitational constant, $\overline{\rho_{\rm m}}$ is the mean matter density, $a$ the scale factor, and $\delta$ the matter density contrast, and, in the final step, we assume that the contributions along the line of sight cancel, i.e. the Limber approximation (D. N. Limber 1953). To account for a population of source galaxies, we integrate equation (5) as

$$\kappa(\boldsymbol{\theta}) = \int_0^{z_{\rm hor}} {\rm d}z\, n_{\rm s}(z)\, \kappa[\boldsymbol{\theta}, \chi(z)], \tag{7}$$

where the normalized source redshift distribution, $n_{\rm s}$, extends to the edge of the survey, $z_{\rm hor} = 3$ in our case. We assume a *Euclid*-like source redshift distribution given by (Euclid Collaboration: Blanchard et al. 2020)

$$n(z) \propto \left(\frac{z}{z_0}\right)^2 \exp\left[-\left(\frac{z}{z_0}\right)^{3/2}\right], \tag{8}$$

with $z_0 = 0.9/\sqrt{2}$, such that the median redshift is 0.9.

Our analysis ignores several higher-order WL effects, contaminants, and the possibly non-negligible impact of approximations, such as the observable being shear rather than reduced shear (e.g. C. Seitz & P. Schneider 1995; M. Bradač et al. 2005), source clustering (e.g. M. Gatti et al. 2024), intrinsic alignments (IAs; e.g. M. E. Lee et al. 2026; C. Vedder et al. 2026), and the Born approximation (e.g. T. Lu & Z. Haiman 2021). We also do not marginalize over nuisance parameters such as magnification bias (e.g. J. Liu et al. 2014) or shifts in the mean or shape of the redshift distribution (e.g. H. Hildebrandt et al. 2020). However, we assume these effects do not affect the relative differences we study in our idealized inference. Nevertheless, we stress that any actual inference should consider these effects in greater detail. While a quantitative assessment of these systematics requires a more elaborate study, we can draw some qualitative expectations from their known dependencies. Baryonic feedback suppresses the amplitude and abundance of the highest convergence WL peaks, since these are sourced by the most massive collapsed structures. Because the redshift and height distributions in this work are built from the same underlying peak population, they are likely to respond to baryonic suppression in a correlated way, as a systematic reduction in the number and amplitude of the most massive haloes would suppress both statistics coherently. However, as shown in J. C. Broxterman et al. (2025), baryonic feedback primarily changes the amplitude of the redshift distribution and not its shape, which is affected by cosmology. Similarly, angular clustering at fixed WL peak number density is primarily sensitive to the spatial distribution of peaks rather than their amplitude and is thus naively expected to be less sensitive to baryonic feedback. Because we work directly with the convergence field from mass maps, IA, which arises from correlated intrinsic source galaxy shapes, does not enter our analysis. For the height distribution, T. Zhang et al. (2022) show that IA impacts the height distribution in a non-negligible way. Given that our redshift assignment distribution is based on properties from the observed KiDS bright sample (as explained in more detail in Section 4.2), which includes IA, the redshift distribution is expected to be less impacted by this systematic.

## 3 SIMULATIONS

In our analysis, we use a new set of simulations, which will be presented in detail in McCarthy et al. (in preparation). We use two sets of 50 spatially flat simulations in 10D and nine-dimensional (9D) cosmological parameter spaces. The initial conditions have identical phases, and the amplitudes of the large-scale modes are fixed to the mean variance. The parameters that are varied are detailed in Table 1, and include the present-day matter density parameter in the limit of no DDM, $\Omega_{\rm m,noDDM}$, i.e. the value that $\Omega_{\rm m}$ would have been if the simulation had no DDM; the cosmic baryon fraction given by the ratio of the redshift zero baryon density to matter density parameters, $f_{\rm b}$; the Hubble parameter, $h$; the scalar spectral index, $n_{\rm s}$; the amplitude, $\ln(10^{10}A_{\rm s})$, and the running of the scalar spectral index, $\alpha_{\rm s} \equiv {\rm d}n_{\rm s}/{\rm d}\ln k$, of the primordial power spectrum; the DE equation-of-state parameters following the CPL parametrization, $w_0$ and $w_{\rm a}$; the sum of three massive neutrino species, $\sum m_\nu c^2$ (eV); and the decay rate of decaying CDM, $\Gamma_{\rm dcdm}$ (km s$^{-1}$ Mpc$^{-1}$), following the model introduced in W. Elbers et al. (2025b). In the DDM extension, a fraction of the CDM is treated as an unstable species with mean lifetime $\tau$, decaying into a relativistic dark radiation component. $\Gamma_{\rm dcdm} = \tau^{-1}$ is the decay rate. The decay is implemented by adjusting the particle masses in proportion to the fraction of DM that has decayed.

**Table 1.** Varied parameters and their ranges in the hypercube. From top to bottom, the present-day total matter density parameter in the limit of no DDM, $\Omega_{\rm m,noDDM}$; the baryon fraction, $f_{\rm b}$; the dimensionless Hubble parameter, $h$; the scalar spectral index, $n_{\rm s}$, and the amplitude, $\ln(10^{10}A_{\rm s})$, of the primordial power spectrum; the CPL equation-of-state parameters of DE, $w_0$ and $w_{\rm a}$; the sum of neutrino masses, $\sum m_\nu c^2$ (eV); the running of the spectral index, $\alpha_{\rm s}$, of the primordial power spectrum; and the DDM decay rate, $\Gamma_{\rm dcdm}$ (km s$^{-1}$ Mpc$^{-1}$).

| Parameter | Hypercube range |
|---|---|
| $\Omega_{\rm m,noDDM}$ | $[0.2, 0.4]$ |
| $f_{\rm b}$ | $[0.14, 0.175]$ |
| $h$ | $[0.6, 0.8]$ |
| $n_{\rm s}$ | $[0.9, 1.0]$ |
| $\ln(10^{10}A_{\rm s})$ | $[2.965, 3.115]$ |
| $w_0$ | $[-1.2, -0.3]$ |
| $w_{\rm a}$ | $[-3.7(w_0+1.0), -3.7(w_0+1.0)+0.5]$ |
| $\sum m_\nu c^2$ (eV) | $[0, 0.45]$ |
| $\alpha_{\rm s}$ | $[-0.03, 0.03]$ |
| $\Gamma_{\rm dcdm}$ (km s$^{-1}$ Mpc$^{-1}$) | $[0, 20]$ |

The sampling within the 10D hypercube follows a Latin hypercube design, which is approximately uniform in the marginal distribution of each parameter individually, with the exception of $w_0$ and $w_{\rm a}$. $w_0$ and $w_{\rm a}$ are not sampled uniformly independently, as large parts of the $w_0$–$w_{\rm a}$ parameter plane are ruled out by various observations. Instead, they are sampled along a degeneracy direction that is approximately parallel to the geometric degeneracy from the cosmic microwave background (CMB), DESI BAO, and SNe (M. Abdul Karim et al. 2025). The displacement along the degeneracy direction is varied by $\pm 0.5$, which brackets all current constraints from the individual probes. The 9D hypercube can be viewed as a projection of the 10D hypercube on to the subspace where the DDM rate is fixed to zero, while the remaining nine parameter values are identical in both hypercubes. Combining the two hypercubes results in a joint 100-node training set in which the $\Gamma_{\rm dcdm} = 0$ hyperplane is sampled more densely than any other value of $\Gamma_{\rm dcdm}$. In our forecasts, we always jointly use these two hypercube sets to train the emulators.

Each hypercube node corresponds to a single simulation realization, with no repeated realizations at fixed cosmology to sample cosmic variance or shape noise. While we find this is sufficient to obtain unbiased constraints in our simulation forecast

Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

using mode fixing and inferring the cosmology from a simulation run within the same framework, it may not be sufficient for an inference on real data.

The hypercube simulations were run using the same set-up as the original FLAMINGO simulations (R. Kugel et al. 2023; J. Schaye et al. 2023), only in smaller volumes. The simulations were run using SWIFT (M. Schaller et al. 2024) and each contains $1260^3$ DM particles and $700^3$ massive neutrino particles in a box with a side length of 700 cMpc, which results in a typical initial DM particle mass of a few $\times\, 10^9\,\mathrm{M}_\odot$, depending on the exact values of the cosmological parameters. The primordial spectrum, depending on $\ln(10^{10}A_\mathrm{s})$, $n_\mathrm{s}$, and $\alpha_\mathrm{s}$ together with the transfer functions computed with CLASS (J. Lesgourgues 2011), sets the linear matter power spectrum used to generate the initial conditions using MONOFONIC (O. Hahn et al. 2020; M. Michaux et al. 2021). Each simulation has a virtual observer with a spherical full-sky light-cone from which we use the matter shells.

We ignore the impact of baryonic physics, i.e. the redistribution of matter due to black hole feedback, SN explosions, or other physical processes pertaining to baryons that impact the distribution of matter on (mildly) non-linear scales, even though this is one of the dominant systematics in WL inferences (see e.g. E. Semboloni et al. 2011; E. Semboloni, H. Hoekstra & J. Schaye 2013). In J. C. Broxterman et al. (2024, 2025), we found that the WL peak height distribution and redshift distribution depend differently on cosmology and baryonic physics, and combining the two statistics helps to simultaneously constrain cosmology and calibrate baryonic physics. In this forecast, we also include the angular clustering of these peaks, which will depend in its own way on baryonic physics, and therefore may improve the baryonic feedback calibration. At the same time, we stress that this should be tested using hydrodynamical simulations and that we do not expect this to affect the relative differences between our results for the constraining power of the WL peak properties.

## 4 METHODS

In this section, we present the pipeline for deriving full-sky WL convergence maps from maps of the matter density, and the subsequent measurements of the WL peak statistics and WL shear two-point correlation functions (2PCFs) from these maps. Afterwards, we present the inference pipeline consisting of the emulator, likelihood, and covariance matrix, which will be used to infer cosmological constraints from an unseen simulation run.

The construction of the WL convergence map closely follows the method described in J. C. Broxterman et al. (2025). We use the concentric HEALPIX (K. M. Górski et al. 2005) light-cones at $N_\mathrm{side} = 8192$ that are centred around a virtual observer at the same position in each simulation. The shells have a thickness of $\Delta z = 0.05$. Full-sky WL convergence maps are constructed from spherical matter shells by discretizing equation (7) as

$$\kappa(\boldsymbol{\theta}, \chi) = \frac{4\pi G}{c^2} \sum_i [1+z(\chi_i)]^{-2} \Delta\chi_i\, \delta_i(\boldsymbol{\theta})\, \Omega_\mathrm{m}(\chi_i)\, \rho_\mathrm{crit}(\chi_i) W(\chi_i), \tag{9}$$

where $\rho_\mathrm{crit}$ is the critical energy density of the universe, $\Omega_\mathrm{m}(\chi_i) = \rho_\mathrm{m}(\chi_i)/\rho_\mathrm{crit}(\chi_i)$ is the value of the matter density parameter at $\chi_i$, whose time evolution is different from $w_0w_a$CDM as the simulations include DDM, and $\Delta\chi_i$ is the comoving thickness of light-cone shell $i$. We evaluate the matter density contrast, $\delta$, directly from the total matter shells. Each time the light-cone diameter exceeds an integer multiple of the box size, we apply a random rotation to all shells within the next integer multiple using the same rotation angles, thereby avoiding spurious discontinuities at the box boundary while still suppressing repeated large-scale structure along the line of sight. The WL kernel is given by

$$W(\chi_i) = \chi_i \sum_{j=i}^{N_\mathrm{shells}} n_\mathrm{s}(\chi_j) \left(1 - \frac{\chi_i}{\chi_j}\right) \Delta\chi_j, \tag{10}$$

with $n_\mathrm{s}(\chi)\mathrm{d}\chi = n_\mathrm{s}(z)\mathrm{d}z$ and which runs until $N_\mathrm{shells} = 60$. We mimic shape noise by assigning a randomly drawn value from a Gaussian with mean $\mu$ and standard deviation $\sigma$

$$\mathcal{N}\left\{\mu = 0, \sigma = \frac{\sigma_\epsilon}{\sqrt{2 n_\mathrm{gal} A_\mathrm{pix}}}\right\}, \tag{11}$$

where $A_\mathrm{pix}$ is the area of one pixel (0.18 arcmin$^2$ for $N_\mathrm{side} = 8192$), $n_\mathrm{gal}$ is the source galaxy number density, and $\sigma_\epsilon$ is the root mean square total intrinsic ellipticity of source galaxies. To mimic a *Euclid* survey, we take $n_\mathrm{gal} = 30$ arcmin$^{-2}$, and $\sigma_\epsilon = 0.26\sqrt{2}$ (Euclid Collaboration: Mellier et al. 2025).

### 4.1 Smoothing

We explore the impact and complementarity of combining different smoothing scales. In our fiducial set-up, similarly to J. C. Broxterman et al. (2024, 2025), we smooth the integrated WL convergence map using a Gaussian kernel with a full width at half maximum of 1 arcmin, corresponding to 2.3 pixels. This was identified by J. Liu et al. (2015) as close to the optimal smoothing scale for their cosmological inference from CFHTLenS data, based solely on the WL peak number density. However, they already highlight that combining multiple smoothing scales may improve constraining power, as different smoothing scales correspond to different spatial scales. As we, instead, mimic a *Euclid*-like analysis, with a different resolving power, and we combine multiple peak statistics, we re-evaluate the impact of varying the smoothing scale on the constraining power. When selections are made based on the SNR, we recompute the noise standard deviation from the smoothed, randomly generated noise realization for each smoothing scale. In Section 5.5, we compare the fiducial results at 1 arcmin smoothing to those obtained with a smoothing scale that is a factor of 2 larger or smaller, or when combining the three smoothing scales.

### 4.2 Peak statistics

We use three different WL peak statistics in our forecast, the redshift distribution, $\mathrm{d}n(z)/\mathrm{d}z$; height distribution, $\mathrm{d}n(\kappa)/\mathrm{d}\kappa$; and angular clustering, $\omega(\theta)$. More details on the binning are given in Appendix A. The peaks are selected from the integrated WL convergence maps after noise and smoothing have been applied. Peaks are defined as pixels with values greater than those of their eight nearest neighbours, as commonly used in WL peak studies (e.g. A. J. Weiss et al. 2019; D. Grandón et al. 2024). For each peak, we determine its WL convergence value, subsequently referred to as the 'peak height'. For SNR $> 3$ peaks, we identify the primary redshift that generates each peak, following the procedure of J. C. Broxterman et al. (2025). In the simulations, using the mass maps, the redshift of the overdensity that contributes most to the WL convergence at each WL peak, i.e. the dominant contribution to equation (9), is assigned to that peak.



Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

The redshifts are distorted, mimicking the observational uncertainty in photometric redshift estimation for massive galaxy clusters. As the high-value WL peaks generally correspond to massive clusters, the redshifts of these objects are expected to be measured with much greater accuracy than those of the average *Euclid* galaxy. Whereas the *Euclid* requirement on the precision of the photo-$z$ estimates is $\sigma_z < 0.05(1+z)$ with catastrophic failures rate of less than 10 per cent (Euclid Collaboration: Mellier et al. 2025), these objects are expected to be measured more accurately using regular photo-$z$ estimation procedures or red-sequence cluster finders (see e.g. E. S. Rykoff et al. 2014; I.-N. Chiu et al. 2024; K.-F. Chen et al. 2025). Based on the properties of the KiDS-1000 bright sample (M. Bilicki et al. 2021), which were estimated using a similar number of photometric bands to those that will be available for *Euclid*, the simulation redshifts are distorted by, for each peak, drawing a random redshift from a normal distribution centred on the true redshift, $z_{\rm true}$, and with a standard deviation of $\sigma_z = 0.02(1+z_{\rm true})$, without any catastrophic outliers. The WL peak redshift distribution thus aims to exploit these precise redshift measurements of these massive objects for cosmological purposes.

As a third statistic, for the SNR $> 3$ peaks, we estimate the WL peak angular 2PCF. We use the Landy–Szalay estimator (S. D. Landy & A. S. Szalay 1993)

$$\omega(\theta) = \frac{\mathrm{DD}(\theta) - 2\mathrm{DR}(\theta) + \mathrm{RR}(\theta)}{\mathrm{RR}(\theta)}, \tag{12}$$

to estimate the clustering signal. Here, $\theta$ is the angular separation between two WL peaks, and DD, RR, and DR represent the number of pairs of peaks separated by $\theta$ in the data sample (DD), a uniformly randomly generated sample (RR), and their cross-catalogue (DR). In practice, we use the package CORRFUNC (M. Sinha & L. H. Garrison 2020) to estimate the angular correlation function. The random set is chosen to be 8 times larger than the data catalogue in order to suppress Poisson noise originating from the random catalogue.

Although we probe the angular clustering and redshift distribution of the high-SNR peaks, these could, in theory, also be combined into a single statistic measuring the 3D clustering. This is, however, beyond the scope of this work.

### 4.3 Two-point correlation functions

To compare our forecast with those similar to a typical WL inference, we additionally estimate the WL angular 2PCFs, $\xi_{+/-}$. We do so by computing the WL convergence angular correlation function, $\mathcal{C}^{\kappa\kappa}$, from the full-sky WL convergence maps, which in the absence of $B$ modes relates to $\xi_{+/-}$ through a Hankel transform as (P. Schneider, L. van Waerbeke & Y. Mellier 2002)

$$\xi_{+/-}(\theta) = \int_0^\infty \frac{\mathrm{d}\ell\,\ell}{2\pi} \mathcal{C}^{\kappa\kappa}(\ell) J_{0/4}(\ell\theta), \tag{13}$$

where $J_{0/4}$ are Bessel functions of the first kind and $\ell$ is the multipole moment. We measure the 2PCFs in nine logarithmic bins between 0.1 and 3 deg, which roughly corresponds to the ranges used by current cosmic shear inferences (see e.g. DES Collaboration 2026a), and is the same angular range as used for the WL peak angular clustering.

### 4.4 Inference forecast

In this section, we provide details of the forecasting pipeline we use to generate our forecasts. As explained in Section 3, the WL peak statistics are extracted through forward modelling the hypercube simulations. We then train an emulator to predict the peak statistics as a function of the 10 cosmological parameters that are varied in the hypercube (Section 4.4.1). Finally, we estimate a covariance matrix from additional simulations (Section 4.4.2) and perform a Bayesian likelihood inference (Section 4.4.3).

#### *4.4.1 Emulation*

To predict the WL peak statistics across the 10D parameter space spanned by the hypercube, we use Gaussian process (GP) emulators trained on statistics computed from the $2 \times 50$ simulations. All emulators are implemented with the `GaussianProcessRegressor` class from SCIKIT-LEARN (F. Pedregosa et al. 2011). We train separate emulators for the different statistics and smoothing scales, and concatenate them to produce the combined forecasts. The details of the emulator construction as well as a leave-one-out (LOO) test are presented in Appendix A. In general, we find sub-per-cent-level relative emulator accuracy across the ranges on which the statistics are used, which is sufficient for our idealized forecast.

#### *4.4.2 Covariance matrix*

To estimate the covariance, we run an additional suite of flat ΛCDM simulations. This set uses the same framework as the hypercube, but assumes the fiducial DES year 3 '3 × 2 pt + All Ext.' ΛCDM cosmology (D3A) of J. Schaye et al. (2023), i.e. a flat ΛCDM ($w_0 = -1, w_a = 0$) without DDM ($\Gamma_{\rm dcdm} = 0\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$), no running of the scalar spectral index ($\alpha_s = 0$), only a single massive neutrino species with $m_\nu c^2 = 0.06$ eV, and the same cosmology as the reference cosmology we are trying to infer. We assume that the covariance over the entire cosmological parameter space can be approximated by the covariance generated from these ΛCDM simulations. This assumption is expected to be a subdominant effect for cosmic shear 2 pt inferences (see e.g. D. Kodwani, D. Alonso & P. Ferreira 2019), but has not yet been studied for the non-Gaussian statistics we consider. For each simulation, we use a different random seed to generate the initial conditions with MONOFONIC (O. Hahn et al. 2020; M. Michaux et al. 2021), and we do not fix any modes, unlike the simulations that are part of the hypercube. As such, our covariance matrix should reflect the full cosmic variance that is inherent to the inference.

For each forecast and for realization $i$, we compute the relevant WL peak statistics, $\mathbf{x}_i$. The sample covariance, $\hat{\mathbf{C}}$, is then computed as

$$\hat{\mathbf{C}} = \frac{1}{N_{\rm cov} - 1} \sum_{i=1}^{N_{\rm cov}} (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^\top, \tag{14}$$

where $\bar{\mathbf{x}} = N_{\rm cov}^{-1} \sum_i \mathbf{x}_i$ is the sample mean. Since our analysis mimics a *Euclid*-like survey covering $f_{\rm sky} = 15\,000/41\,253$ of the sky, the covariance is rescaled as

$$\mathbf{C} = \frac{\hat{\mathbf{C}}}{f_{\rm sky}}. \tag{15}$$

Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

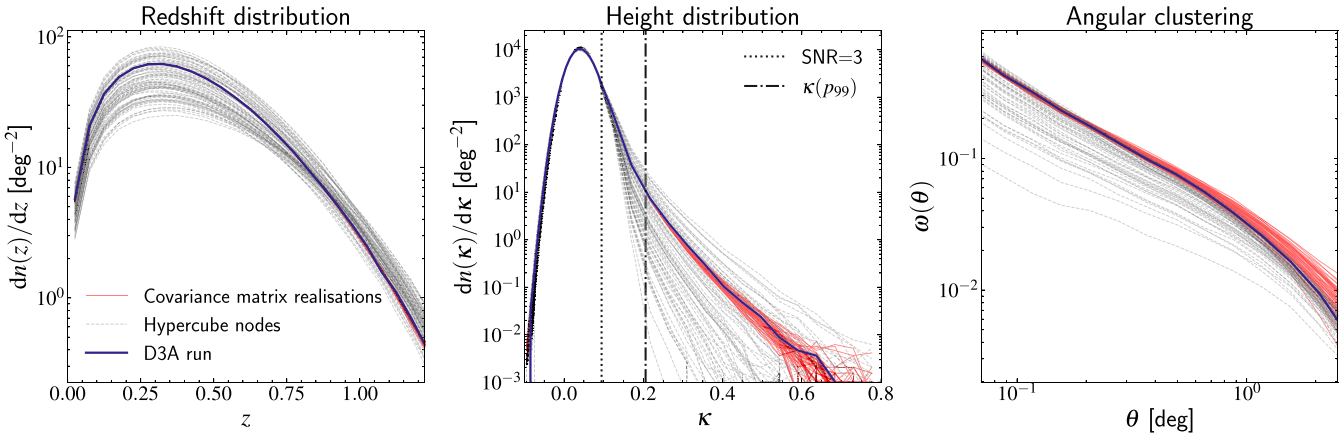



**Figure 1.** WL peak statistics used in the fiducial forecast with 1 arcmin smoothing. Left panel: Redshift distribution of SNR > 3 peaks. Middle panel: Height distribution. Right panel: Angular correlation function of SNR > 3 peaks. The measurements from the hypercube simulations are shown as dashed light-grey curves, the covariance matrix realizations as red lines, and the measurements from the additional simulation at the FLAMINGO D3A cosmology, which is inferred in this paper, as a purple curve. The dotted and dash–dotted lines in the central panel indicate the SNR = 3 threshold used for all statistics and the $\kappa$ value corresponding to the 99th percentile peak, which is used as an upper bound for the peak-height distribution.

When multiple statistics or smoothing scales are combined into a joint data vector, the covariance matrix is estimated from the same realizations, thereby capturing all cross-correlations among them. In Appendix B, we show the correlation matrix for our fiducial combined WL peak forecast, i.e. the height distribution, redshift distribution, and angular clustering at 1 arcmin smoothing, illustrating that the angular correlation function exhibits strong correlations between different angular bins, whereas the cross-correlations between the different statistics and the correlations between individual bins of the height distribution and the redshift distribution are weaker.

### 4.4.3 Likelihood

For our forecast, we assume an explicit likelihood. Whereas simulation-based inferences that learn the relationship between cosmological parameters and observables directly could, in theory, provide more precise constraints (S. Vinciguerra, N. Martinet & M. Gatti 2026; J. Williamson et al. 2026), we use an explicit likelihood because it allows us to gain more physical insight into the constraining properties of different statistics. We assume the likelihood follows a multivariate $t$-distribution to account for the fact that the covariance matrix $\mathbf{C}$ is estimated from $N_{\rm cov} = 80$ simulations rather than being known exactly (E. Sellentin & A. F. Heavens 2016). The likelihood is given by

$$\mathcal{L}(\boldsymbol{x}|\boldsymbol{\theta}) \propto \left\{1 + \frac{[\boldsymbol{x} - M(\boldsymbol{\theta})]^\top \mathbf{C}^{-1} [\boldsymbol{x} - M(\boldsymbol{\theta})]}{N_{\rm cov} - 1}\right\}^{-N_{\rm cov}/2}, \tag{16}$$

where the data vector $\mathbf{x}$ consists of (a combination of) the different WL peak statistics, $\mathbf{C}^{-1}$ is the inverse sample covariance matrix, and $M(\boldsymbol{\theta})$ is the emulator prediction. In practice, we evaluate the log-likelihood, which is given by

$$\begin{aligned}\ln \mathcal{L} = &\ln \Gamma\left(\frac{N_{\rm cov}}{2}\right) - \ln \Gamma\left(\frac{N_{\rm cov} - N_{\rm dv}}{2}\right) - \frac{N_{\rm dv}}{2} \ln[\pi(N_{\rm cov} - 1)] \\ &- \frac{1}{2} \ln |\mathbf{C}| - \frac{N_{\rm cov}}{2} \ln\left(1 + \frac{[\boldsymbol{x} - M(\boldsymbol{\theta})]^\top \mathbf{C}^{-1} [\boldsymbol{x} - M(\boldsymbol{\theta})]}{N_{\rm cov} - 1}\right),\end{aligned} \tag{17}$$

where $\Gamma$ is the gamma function and $N_{\rm dv}$ is the dimension of the data vector $\boldsymbol{x}$. In the analysis, we infer the cosmology of the reference data vector, $\boldsymbol{x}$, for which we run an additional simulation using the hypercube infrastructure but at the fiducial cosmology of the FLAMINGO simulations (J. Schaye et al. 2023), which is the best-fitting D3A (T. M. C. Abbott et al. 2022) with a single massive neutrino species, from which we extract the same statistics through forward modelling. We infer the best-fitting parameters by sampling the log-likelihood with the Markov chain Monte Carlo sampler EMCEE (D. Foreman-Mackey et al. 2013), and our results are plotted with GETDIST (A. Lewis 2025).

For most parameters, we apply flat priors. For the two beyond $w_0 w_{\rm a}$ CDM parameters, $\alpha_{\rm s}$ and $\Gamma_{\rm dcdm}$, we apply, respectively, a Gaussian and half-Gaussian prior centred on the truth with widths 0.02 and 1, to penalize strong deviations from these less common extensions.

We quantify the FoM for a two-dimensional (2D) cosmological parameter combination of $\theta_1$ and $\theta_2$ as

$$\text{FoM} = \frac{1}{\sqrt{|C(\theta_1, \theta_2)|}}, \tag{18}$$

where $C$ is the $2 \times 2$ parameter covariance matrix of $\theta_1$ and $\theta_2$. Similarly, we quantify the figure of bias (FoB) as

$$\text{FoB} = \sqrt{(\theta_{\rm truth} - \theta_{\rm inferred})^\top C(\theta_{\rm inferred})^{-1} (\theta_{\rm truth} - \theta_{\rm inferred})}, \tag{19}$$

where $\theta_{\rm truth}$ is the value of the cosmological parameters in the D3A cosmology and $\theta_{\rm inferred}$ is the best-fitting emulator prediction.

## 5 RESULTS AND DISCUSSION

In Fig. 1, we illustrate the measurements of the three different WL peak statistics across the different parts of the forecast pipeline. The left panel shows the redshift distribution of the SNR > 3 peaks, the central panel shows the height distribution, and the right panel shows the angular 2PCF of the SNR > 3 peaks. These measurements correspond to our fiducial smoothing scale of 1 arcmin. In each panel, the solid red curves indicate the measurements from the covariance-matrix realizations, the solid purple curve indicates the target distribution measured from the D3A simulation, and the dashed grey curves show the measurements



Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

in the 100 hypercube nodes. The hypercube nodes span a wide range of predictions, demonstrating sensitivity across the full 10D parameter space. The central panel also includes the SNR = 3 cut as a vertical dotted line.

For each statistic, the D3A measurement falls within the range of the covariance realizations. The spread of the covariance-matrix realizations relative to the D3A target varies across the three statistics, reflecting the distinct cosmic-variance properties of each observable. For the redshift distribution, the covariance realizations show very little scatter around the D3A measurement across the entire redshift range, reflecting the small sample variance on the shape and amplitude of the redshift distribution of high-SNR WL peaks expected for the forecasted *Euclid* measurements. For the height distribution, the spread of the covariance realizations increases for $\kappa$ values above the SNR cut. This is expected, as the number density of the peaks drops steeply with increasing $\kappa$, and the Poisson noise from the rare high-mass haloes that generate these peaks thus increases. Additionally, for the highest mass objects, the relative location of the (virtual) observer also matters, as it determines the exact peak height via the lensing kernel. As discussed in more detail in Appendix A, we do not use the highest 1 per cent $\kappa$ peaks during emulator training and inference for the height distribution, as we found the emulator performed better without them, likely due to increased stochasticity. For the angular clustering, the covariance realizations show a scatter of a factor of a few in the signal across the entire angular range, with the relative spread increasing towards larger angular separations.

### 5.1 The impact of cosmology

Before examining the actual cosmological forecast, we first discuss the possible ways in which the changes in cosmology may impact our WL peak forecast. The 10 cosmological parameters varied in the hypercube can affect the WL peak statistics in three distinct physical ways.

The first possibility is through the direct impact on the physical matter density. From equation (9), $\kappa$ is directly proportional to $\Omega_{\rm m}(\chi_i)\rho_{\rm crit}(\chi_i)$, i.e. the physical total matter density over time. Within the cosmologies in the hypercube, $\rho_{\rm crit}(z) = 3H(z)^2/(8\pi G)$ and the physical matter density thus scales as $\propto \Omega_{\rm m}h^2$, illustrating that $\Omega_{\rm m}$ and $h$ directly impact the WL signals through this channel. Additionally, $\Gamma_{\rm dcdm}$ impacts the total matter density, and therefore directly rescales the amplitude of each peak, albeit with a small magnitude as the DDM rates are small.

The second way in which the WL peaks can be impacted by cosmology is through a change in the growth of structures, which sets the matter density contrast $\delta_i(\boldsymbol{\theta})$ in equation (9). Structure growth controls the masses of the structures that generate the WL peaks, thereby also controlling the number density and height of the peaks. Parameters that impact the structure growth directly are the amplitude of the primordial power spectrum, $\ln(10^{10}A_{\rm s})$, which sets the overall strength of initial density fluctuations, the sum of neutrino masses, $\sum m_\nu c^2$, whose free-streaming suppresses structure growth, but also DE ($w_0$, $w_{\rm a}$), as it may suppress the growth of matter perturbations through its impact on the growth factor.

Finally, the cosmological parameters may impact the WL peak properties through changes in the geometry, i.e. by changing the comoving distance–redshift relation, $\chi(z)$. Therefore, any parameter that alters the expansion history through the Friedmann equation changes the WL kernel (equation 10) and the conversion between physical separations and angles on the sky, which is relevant for the angular clustering signal.

Although it is not possible to cleanly separate the impact of these different channels on the parameter posteriors, in the next sections, we discuss possibilities for how individual parameters may explain the differences in constraining power we observe.

### 5.2 Cosmological constraints forecast

Next, we forecast the cosmological constraints from our Bayesian analysis for our fiducial set-up for each of the peak statistics individually, and for the forecast that combines the three statistics. In Fig. 2, we show the marginalized one-dimensional (1D) posterior distributions for all 10 cosmological parameters for each of the three individual statistics and their combination, at the fiducial smoothing scale of 1 arcmin. The forecast using the height distribution, redshift distribution, angular clustering, or the combined approach yields the red, orange, blue, and black posteriors, respectively. The shapes of the priors are indicated by the grey dotted line and grey shaded area. In Appendix C, we show the full corner plot corresponding to forecasts using the three WL peak statistics together, the redshift distribution individually, and the angular clustering and height distributions together. This comparison shows the value of just adding the redshift distribution over the other two statistics.

We first assess whether the inference correctly recovers the fiducial D3A cosmology. The combined inference of the three peak statistics yields strong constraints on $\Omega_{\rm m}$, $\Omega_{\rm b}$, $\ln(10^{10}A_{\rm s})$, $w_0$, $w_{\rm a}$, and $h$, whereas the parameters $\sum m_\nu c^2$, $n_{\rm s}$, $\alpha_{\rm s}$, and $\Gamma_{\rm dcdm}$ are either weakly constrained or recover the prior. The combined inference (black posteriors) recovers the input parameter values well for those parameters that are well constrained, as the marginalized posteriors are consistent with the true values (dashed black curves) within the 68 per cent credible intervals, as shown in Fig. 2. The average FoB of the maximum a posteriori of the parameters that are well constrained, $\Omega_{\rm m}$, $\Omega_{\rm b}$, $h$, $w_0$, $w_{\rm a}$, and $\ln(10^{10}A_{\rm s})$ is $0.59\sigma$. The correspondence demonstrates that the entire pipeline of emulator, likelihood, and covariance matrix together form a self-consistent and proper framework for our cosmological forecasts, and that the added noise does not severely bias the parameter constraints.

Although we expect our idealized inference to recover all parameters well, minor offsets from the truth can be due to (a combination of) noise in the likelihood arising from the finite number of covariance realizations, emulation prediction error, the covariance matrix not varying with cosmology, the added noise components, or because of cosmic variance in the D3A run whose cosmology is inferred. We find some minor offsets from the truth, particularly when a parameter is not well constrained, as is visible in the multimodal blue angular clustering posteriors for $\Omega_{\rm m}$, $h$, and $n_{\rm s}$.

Next, we discuss the constraining power on each cosmological parameter and the complementarity between the different WL peak statistics in probing different parts of the cosmological parameter space.

#### *5.2.1 The matter density and amplitude of the primordial power spectrum*

The matter density parameter, $\Omega_{\rm m}$, and the amplitude of the primordial power spectrum, $\ln(10^{10}A_{\rm s})$, are among the best-

Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

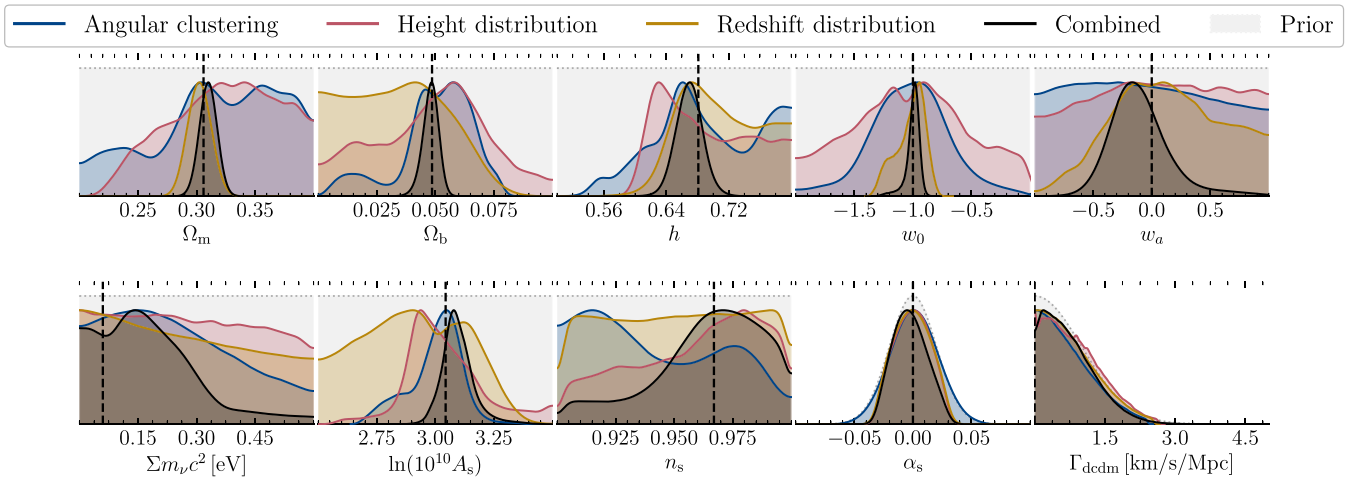



**Figure 2.** 1D posteriors of the fiducial forecast using WL peak angular clustering (blue), height distribution (red), redshift distribution (orange), or the three statistics combined (black) for SNR > 3 WL peaks with 1 arcmin smoothing. The prior is indicated by the dotted grey curve and the shaded grey area, which uniformly spans the parameter range for all panels except the final two. The true values are the D3A cosmology indicated by the dashed black lines. For the combined inference, informative unbiased constraints are recovered for the matter and baryon densities, the Hubble parameter, the DE equation of state, and the amplitude of the primordial power spectrum, whereas the neutrino mass, the running and amplitude of the scalar spectral index, and the DDM rate are either poorly constrained or largely recover the prior.

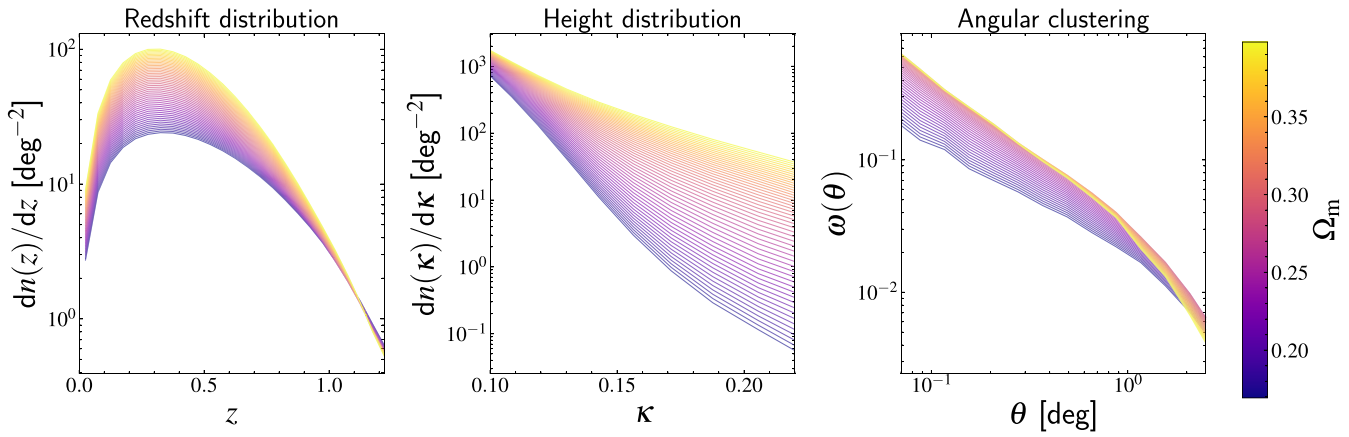



**Figure 3.** Emulator predictions for the redshift distribution (left panel), height distribution (middle panel), and angular clustering (right panel) of SNR > 3 WL peaks at 1 arcmin smoothing for varying $\Omega_m$, as indicated by the colour bar, while keeping all other parameters fixed to their mean values within the hypercube. All three statistics are impacted. Larger values of $\Omega_m$ increase the overall number of peaks at fixed WL convergence and give rise to peaks preferentially corresponding to redshifts where the WL kernel peaks. Below 1 deg the angular clustering signal shows a monotonic trend of larger $\Omega_m$ resulting in a stronger clustering signal, but at larger angular scales this trend reverses.

constrained parameters in our forecast, which is expected for $\Omega_m$ since it impacts the inference through each channel as discussed before, whereas $\ln(10^{10}A_s)$ only acts directly through the second channel. For the combined inference, we find $\Omega_m = 0.310 \pm 0.006$ and $\ln(10^{10}A_s) = 3.09 \pm 0.06$, while the input values were $\Omega_m = 0.306$ and $\ln(10^{10}A_s) = 3.04$. The corresponding best-fitting $S_8$ is $0.819 \pm 0.011$, compared to the true value of 0.815. Since we do not marginalize over any systematics, we limit our interpretation to checking that the inferred values are consistent with the truth, rather than drawing conclusions about the absolute constraining power.

The strong constraining power for $\Omega_m$ and $\ln(10^{10}A_s)$ is consistent with the results of cosmic shear surveys that typically report their cosmology constraints in terms of the $S_8$ parameter (e.g. T. M. C. Abbott et al. 2022; A. H. Wright et al. 2025), since this is the combination in the $\Omega_m$–$\sigma_8$ plane that current cosmic shear surveys are, approximately, most sensitive to. In our forecast, we find that the constraint on $\Omega_m$ is driven by the redshift distribution, whereas the angular clustering and height distribution both contribute significantly to the constraint on $\ln(10^{10}A_s)$.

The sensitivity of the three WL peak statistics to $\Omega_m$ is illustrated in Fig. 3, which shows the emulator predictions for each statistic as $\Omega_m$ is varied individually while all other parameters are held fixed at their mean values. The colour indicates the values of the cosmological parameters, which range from the minimum to the maximum of the hypercube, as also indicated by the colour bar. In Appendix D, we show similar figures for $h$, $w_0$, $w_a$, and $\Omega_b$. Fig. 3 illustrates different dependencies on



Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

$\Omega_m$. The emulator predictions across different $\Omega_m$ values result in changes of nearly two orders of magnitude in the amplitudes of the redshift and height distributions, and a factor of a few in the angular clustering amplitude, with the curves spanning the full dynamic range of the axes in all three panels. Out of the 10 cosmological parameters, within the hypercube range, $\Omega_m$ shows the largest deviations across each of the three statistics. The amplitude of the redshift distribution is set by the number of peaks above the SNR $> 3$ threshold that is applied to the height distributions. Additionally, the redshift distribution shows a secondary dependence, as $\Omega_m$ also changes its shape. Although the distribution always peaks at the same redshift, which is close to the redshift at which the lensing kernel peaks given the chosen source redshift distribution, smaller values of $\Omega_m$ preferentially lead to a smaller fraction of lower redshift peaks, which could be due to the fact that as $\Omega_m$ decreases, $\Omega_\Lambda$ increases to maintain a spatially flat universe. Therefore, late-time density perturbations are increasingly suppressed, as in the limit of DE domination, the linear growth factor asymptotically approaches a constant, halting further structure growth.

For low values of $\Omega_m$, the angular clustering signal increases monotonically with $\Omega_m$. However, as $\Omega_m$ is increased beyond 0.3, the clustering signal shows smaller relative changes than the other two statistics and the signal at $\theta > 0.3$ deg even starts to decrease, while the smaller scales still show the same monotonic behaviour. This non-monotonic behaviour could explain why the clustering signal is not able to provide precise constraints on $\Omega_m$ with the 1D posterior even being multimodal. However, because the clustering constraint reaches the imposed prior boundary on $\Omega_m$, these constraints should be interpreted with caution.

The amplitude of the primordial power spectrum, $\ln(10^{10}A_s)$, sets the overall normalization of the matter power spectrum, directly controlling the abundance and masses of the structures that generate the high-SNR WL peaks. A larger value of $A_s$ increases the amplitude of density fluctuations, leading to more massive haloes and a higher number density of high-SNR peaks. The constraint on $\ln(10^{10}A_s)$ is driven by the angular clustering and peak height distribution, as a higher $\ln(10^{10}A_s)$ produces more WL peaks at fixed $\kappa$ and a more strongly clustered large-scale structure, leading to a stronger two-point signal and more high-SNR peaks.

#### 5.2.2 Dark energy equation of state

In our forecast, the CPL equation-of-state parameters of DE, $w_0$ and $w_a$, affect the inference through their effects on the expansion history, which enter the lensing kernel, and through their impact on the WL convergence through changes in geometric and growth rates. The DE sets the late-time expansion history, where a less negative $w_0$ corresponds to stronger DE dominance, which both suppresses the growth of structure and alters the angular diameter distances to the lens and source planes. Because the simulations do not model DE perturbations, the effect of evolving DE parameters on the growth rate is purely through the expansion rate. Together, these effects modify the WL convergence value of a WL peak generated by a halo of fixed mass. In particular, the amplitude of the integrated convergence (equation 7) along the line of sight scales with the product of the matter density contrast, the comoving distances, and the kernel $W(\chi_i)$, all of which depend on $w_0$ and $w_a$.

As, among the three WL peak statistics, the redshift distribution is the only statistic that directly probes a signal over cosmic time, it is unsurprising that it dominates the constraining power on $w_0$ and $w_a$ in our forecast. The forecast yields $w_0 = -0.98 \pm 0.06$ and $w_a = -0.15 \pm 0.24$. This represents a 6 per cent relative precision on $w_0$ while $w_a$ remains poorly constrained with $\sigma_{w_a} = 0.24$. The parameter $w_a$ controls the time evolution of the equation of state and is generally less well constrained than $w_0$. This is because the constraining signal for $w_a$ requires resolving the differential evolution of the DE contribution across redshift, which remains more uncertain in our forecast.

There is a strong degeneracy direction across all $w_0$–$w_a$ constraints, as the two parameters affect the forecast in a very similar way. This is also reflected by the figures in Appendix D showing the impact of varying the two parameters individually on the WL peak statistics.

#### 5.2.3 Hubble parameter

The Hubble parameter, $h$, enters the WL peak statistics most directly through the physical matter density, which scales as $\Omega_m h^2$, and sets the overall amplitude of the WL convergence (equation 9). As a consequence, $h$ and $\Omega_m$ show a negative correlation that is visible in the 2D contour for the height distribution and angular clustering in Fig. C1, as a larger value of $h$ can compensate for a lower value of $\Omega_m$ to generate a peak of the same WL convergence value. This anticorrelation is primarily present in the posteriors corresponding to the height distribution and angular clustering, which therefore struggle to constrain $h$ independently of $\Omega_m$. The redshift distribution is able to break this degeneracy as it provides a tight estimate on $\Omega_m$, such that the combination is also able to probe $h$, possibly through the impact of $\Omega_m$ on the geometry that is being probed by the redshift distribution. The qualitative impact of $h$ on the angular clustering and height distribution is very similar to that of $\Omega_m$, as shown in Fig. D1, albeit with smaller absolute changes. Regarding the redshift distribution, $h$ primarily impacts its amplitude, whereas $\Omega_m$ also affects the shape, illustrating how including the redshift distribution helps break the $\Omega_m$–$h$ degeneracy.

#### 5.2.4 Baryon density

In the hypercube, $\Omega_m$ is varied individually and $\Omega_b$ is varied through changes in the ratio $\Omega_b/\Omega_m$. $\Omega_b$ at fixed $\Omega_b/\Omega_m$ then impacts the forecast through its impact on the expansion history. At fixed $\Omega_m$, the expansion history is not affected, and changes in $\Omega_b/\Omega_m$, instead, likely affect the inference through changes in the primordial power spectrum. Individually, none of the three statistics provides a tight constraint on $\Omega_b$. However, when combining all three statistics, the constraint on $\Omega_b$ improves drastically. The physical origin of this improvement is not straightforward to identify, but it reflects the breaking of a degeneracy in the full 10D parameter space that no single statistic can resolve on its own. We note that $\Omega_b$ is, in any case, already well constrained independently through big bang nucleosynthesis (e.g. B. D. Fields et al. 2020), and a strong external prior on this parameter could help break the residual degeneracies even further, potentially tightening the constraints on the remaining cosmological parameters in the process.

Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

#### 5.2.5 Neutrino mass, scalar spectral index, and beyond $w_0w_aCDM$ parameters

Several parameters in the hypercube are only weakly constrained by our fiducial forecast, and the marginalized posteriors recover much of the prior range. The sum of neutrino masses $\sum m_\nu c^2$ is not well constrained in our combined forecast compared to competitive results from measurements of, for example, the CMB combined with DESI BAO, which report constraints on the neutrino mass with a marginalized error of 0.078 eV for a $w_0w_a$ CDM + $\Sigma m_\nu$ inference (W. Elbers et al. 2025a). The constraining power, however, is still competitive with CMB-independent constraints from DESI. Using the same simulations as us, J. Wang et al. (2026) showed that persistence strips, a formalism that segments Betti curves by topological persistence, can reach a neutrino mass constraint with an uncertainty down to 0.05 eV.

The scalar spectral index $n_s$ and its running $\alpha_s$ affect the shape of the primordial power spectrum and, hence, the matter power spectrum on small scales, which ultimately influence the mass function of haloes and, in turn, the high-SNR peaks. However, given that the range over which these parameters are varied in the hypercube is not very large relative to existing CMB constraints, and the constraining power of the WL peak statistics is limited, we fail to constrain these parameters within our fiducial forecast.

Similarly, the DDM decay rate $\Gamma_{\rm dcdm}$ is only weakly constrained, and the marginalized posterior recovers the half-Gaussian prior we impose. This is somewhat surprising, as in J. C. Broxterman et al. (2025) we found that the amplitude of the FLAMINGO variations with DDM rates of 0.12 and 0.24 $H_0/h$ reduced the amplitude of the redshift distribution by factors of $\approx 2$ and $\approx 4$, respectively. Although the latter is outside the range of the hypercube used in this work, naively, given the change in the number of peaks above the SNR = 3 threshold in the 0.12 $H_0/h$ run, we would still expect an additional constraint beyond the prior in this work. The two analyses differ in their peak-height cut and noise properties, although the exact reason we do not recover stronger constraints on $\Gamma_{\rm dcdm}$ is unclear, it might be the fact that we already impose a strong prior on the value of $\Gamma_{\rm dcdm}$ and the small variations within this prior do not impact the analysis much. The prior on $\Gamma_{\rm dcdm}$ is more than 10 times stronger than the weakest DCDM scenario considered in the FLAMINGO suite. For $\Gamma_{\rm dcdm} = 1\,{\rm km\,s^{-1}\,Mpc^{-1}}$, which is the width of the half-Gaussian prior we apply, the effect on the linear power spectrum is less than 1 per cent and the effect on the expansion rate is less than 0.1 per cent. An alternative explanation is that our emulators emulate $\Omega_{\rm m}$, i.e. the real present-day value of $\Omega_{\rm b} + \Omega_{\rm cdm} + \Omega_\nu$ and not $\Omega_{\rm m,noDDM}$, which is the parameter that is varied in the hypercube. Some of the sensitivity to $\Gamma_{\rm dcdm}$ may therefore be absorbed by $\Omega_{\rm m}$.

### 5.3 Complementarity of WL peak statistics

One of the key results of this work is the complementarity of the three different WL peak statistics. Fig. 2, and in more detail, the 2D corner plot in Appendix C, demonstrates that the combination of the WL peak height distribution, redshift distribution, and angular clustering significantly outperforms each individual statistic for the majority of parameters. Particularly the constraints on $\Omega_{\rm b}$, $h$, $w_a$, $\sum m_\nu c^2$, and $\ln(10^{10}A_s)$ improve significantly compared to any individual statistic.

The height distribution is directly sensitive to the number density of massive haloes as a function of their mass. Individually, the height distribution does not impose strong constraints on any cosmological parameters. Since the height distribution measures the WL convergence values by integrating over the entire source redshift distribution, it averages over the complete growth history, which limits its ability to constrain parameters that affect the growth history in different ways but have similar effects on the amplitude of the WL signal.

Adding the redshift distribution allows us to break some of these degeneracies by probing the evolution of structure formation. Individually, the redshift distribution is only able to constrain well $\Omega_{\rm m}-w_0$. However, the corner plot shows that combining the statistics helps break degeneracies. Individually, none of the statistics can probe the neutrino mass, but together they can place a clear upper bound, albeit still not competitive with other cosmological probes. Angular clustering provides a qualitatively different source of information by probing the spatial distribution of peaks on the sky. Consequently, it provides clear constraints on $\ln(10^{10}A_s)$.

C. T. Davies et al. (2022) found that the clustering of WL peaks is a factor $\approx 3$ more sensitive to $h$ and $w_0$ than the peak height distribution is. For $w_0$, we see a qualitatively similar improvement between the red and blue 1D posteriors. However, the redshift distribution outcompetes the height distribution and angular clustering for these parameters. For $h$, we do not see the same improvement. If anything, the height distribution seems to provide slightly better results. The comparison is complicated by the fact that, in our 10D inference, we cannot constrain these parameters within the hypercube bounds, so the prior bounds are therefore important. Also, C. T. Davies et al. (2022) combine the clustering of peaks selected at different heights, which they show significantly improves the constraints compared to clustering of peaks at a single cut-off value.

### 5.4 Comparison to shear two-point correlation functions

As our analysis is idealized, since we do not include marginalization over nuisance parameters, it is difficult to quantitatively compare our constraints to those of other works, let alone to the precision of current inferences or requirements for *Euclid*. Therefore, to more straightforwardly assess the benefits and downsides of our WL peak forecast, we compare our results to those of a similarly idealized forecast using WL shear 2PCFs, $\xi_\pm$, which are the standard statistic used by current cosmic shear surveys. The 2PCFs are computed from the same simulations and WL convergence maps and use the same smoothing, priors, and covariance matrix realizations as the fiducial set-up. Fig. 4 shows the marginalized 2D posterior distributions for all 10 cosmological parameters for the $\xi_\pm$ 2PCFs (blue), the combined WL peak statistics (black), and the combination of all four statistics (orange), i.e. the shear 2 pt functions plus the three peak statistics. In general, the combination of WL peak statistics outperforms the $\xi_\pm$ 2PCFs. In reality, the cosmic shear signal is measured across different tomographic bins, meaning our forecast on the constraining power of the 2PCFs is a conservative estimate. Additionally, the cosmic shear measurements are usually combined with galaxy clustering and galaxy–galaxy lensing in a 3 × 2 pt analysis, yielding better constraints than the cosmic shear two-point signal alone. A full comparison to the 3 × 2 pt, or even 6 × 2 pt, forecast, however, is beyond the scope of this work.



Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

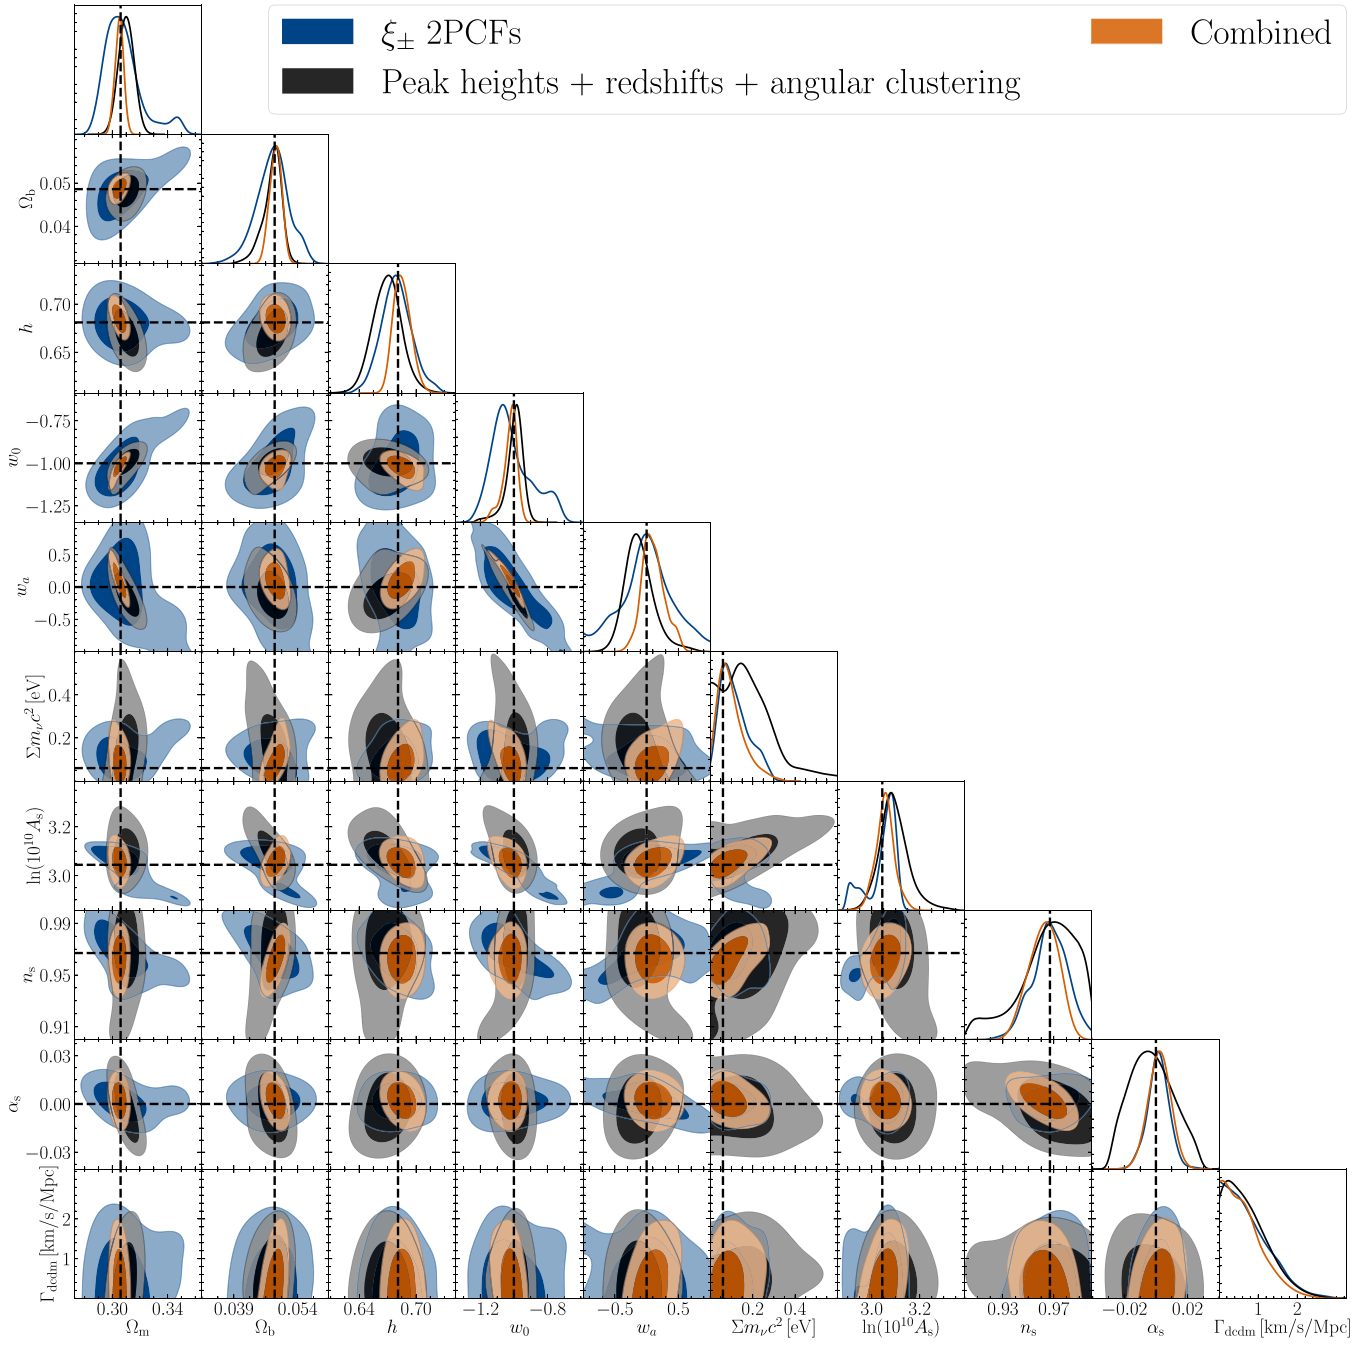



**Figure 4.** Corner plot of the cosmological parameter forecast using the WL 2PCFs (blue) or angular clustering, redshifts and heights of WL peaks (black) of SNR > 3 peaks, or a combined inference using all four statistics (orange) at 1 arcmin smoothing. The contours indicate the $1\sigma$ and $2\sigma$ credible intervals. The combination of WL peak statistics outperforms the $\xi_{+/-}$ 2PCFs for $\Omega_m$, $\Omega_b$, $w_0$, and $w_a$, while the constraints on $h$, $\ln(10^{10}A_s)$, and $\Gamma_{dcdm}$ are similar and the remaining parameters, $\Sigma m_\nu c^2$, $n_s$, and $\alpha_s$ are probed best by the $\xi_{+/-}$ 2PCFs.

For the parameters $\Omega_m$, $\Omega_b$, $w_0$, and $w_a$, the combination of WL peak statistics provides a tighter constraint than the $\xi_{+/-}$, while the constraints on $h$, $\ln(10^{10}A_s)$, and $\Gamma_{dcdm}$ are of a similar level. Only for the parameters $\Sigma m_\nu c^2$, $n_s$, and $\alpha_s$, which were poorly constrained in the WL peak forecast, do the $\xi_{+/-}$ outperform the WL peaks. The constraints on these parameters, however, remain worse than those obtained from other cosmological observables, such as the CMB (see e.g. Planck Collaboration VI 2020). Whereas the high-valued WL peaks are sensitive to the most massive haloes and structures, the 2PCFs are instead sensitive to the shape of the matter power spectrum across a wide range of scales, which is why the $\xi_{+/-}$ constrains these parameters better (see e.g. T. Hamana et al. 2004; K. Heitmann et al. 2009; J. C. Broxterman et al. 2025; Y. Yang, S. Bird & M.-F. Ho 2025).

The joint forecast tightens the contours even further, particularly for $\Omega_m$, where some parameter degeneracy seems to be broken as the WL peaks and $\xi_{+/-}$ show different degeneracies in, for example, the $\Omega_m - h$ and $\Omega_m - \ln(10^{10}A_s)$ subpanels.

Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

The results are qualitatively consistent with C. T. Davies et al. (2022), who also find that WL peaks can probe $S_8$ and $w_0$ more precisely than the shear 2PCFs. However, they find a similar level of constraining power on $\Omega_m$ and that the shear 2PCFs are able to probe $h$ better than the peak statistics. This could be explained by the fact that they do not include the redshift distribution of WL peaks, which drives our constraints on $\Omega_m$. As expected, they generally find that combining the statistics leads to even tighter constraints.

Overall, these results suggest that WL peaks and 2PCFs are complementary probes, as the peaks carry non-Gaussian information that helps break degeneracies in the 2PCF posterior, while the 2PCFs still carry constraints on parameters to which the peak statistics are less sensitive. Their combination, therefore, yields a more powerful cosmological inference than either statistic alone, motivating a joint analysis for future *Euclid*-like surveys, which can be further improved by including additional non-Gaussian statistics (see e.g. Euclid Collaboration: Ajani et al. 2023).

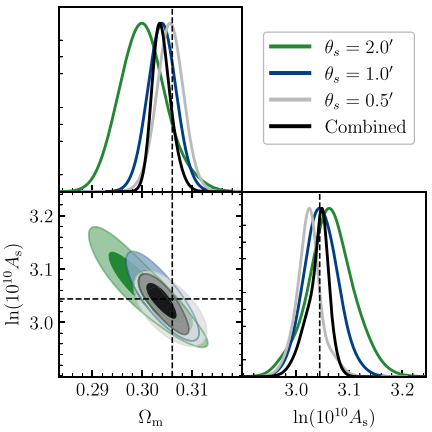
### 5.5 Impact of the smoothing scale

In this section, we explore the impact of the smoothing scale on the cosmological constraints. We chose our fiducial smoothing scale based on the findings of J. Liu et al. (2015), who found it to be the most constraining for their inference of the WL peak-height distribution. However, they already showed that combining scales may improve the results, leaving a more dedicated analysis to future work.

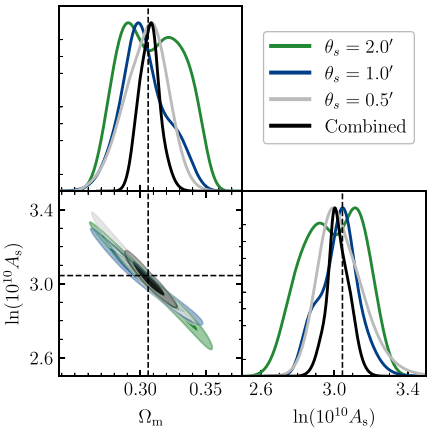
Physically, a different smoothing scale corresponds to a different physical size at each lens redshift. In order to measure an object with the highest significance, the physical aperture corresponding to the smoothing scale should be similar to the object's size. The smoothing will thus affect which WL peaks are detected and at what significance level. By combining scales that are better suited for extracting smaller high-redshift or larger lower-redshift objects, more information could possibly be probed. Additionally, for the smallest scales, extra information may be obtained from different radial profile parts of haloes, as a single low-redshift object may generate multiple peaks.

We fix eight cosmological parameters to their true values and infer only the remaining two. Consequently, the absolute constraining power may differ from the previous figures. We focus on the $\Omega_m - \ln(10^{10}A_s)$ and $w_0 - w_a$ planes because these are the parameter combinations that upcoming WL inferences will likely be most interested in. At the same time, because we did not re-optimize our emulators for different smoothing scales or for this 2D parameter forecast, we found the results to be less stable to small variations in the emulation set-up. We therefore decided to add the emulator uncertainty as an additional noise component to the covariance matrix $\mathbf{C}$ (equation 15). Following the same LOO approach described in Appendix A, we estimate an emulator error covariance matrix, $\Sigma_{emu}$, from the residuals between the emulator predictions and the true measurements of the withheld hypercube nodes. The total covariance used in this section is then computed following the combined likelihood function for a multivariate $t$ and Gaussian distribution from A. Heavens, L. Whiteway & E. Sellentin (2026). The additional term increases the uncertainty, preventing the 2D forecast from becoming artificially overconfident due to emulator imperfections.

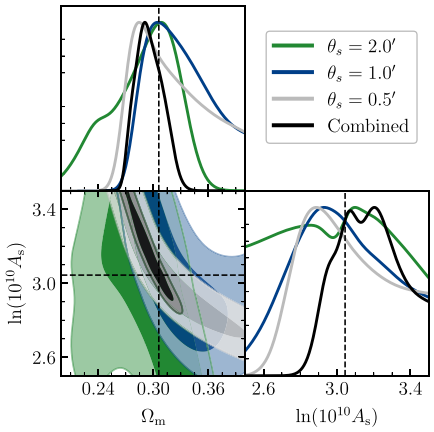
From top to bottom, Figs 5 and 6 show the marginalized 2D posterior distributions for the redshift distribution, height distribution, and angular clustering, respectively, for smoothing scales




(a) Redshift distribution



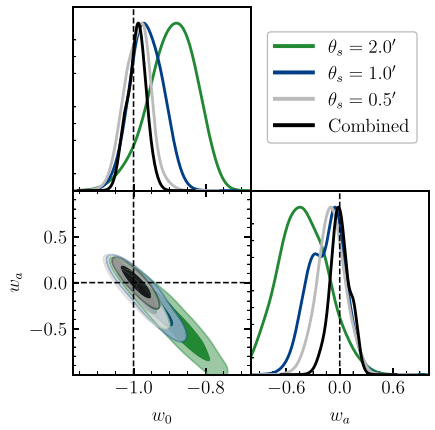

(b) Height distribution



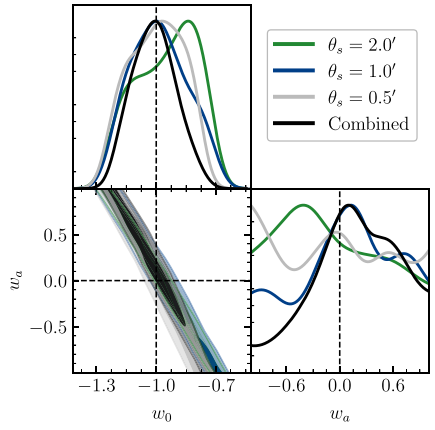

(c) Angular clustering

**Figure 5.** Posteriors for the 2D forecast for $\Omega_m - \ln(10^{10}A_s)$ using (a) the redshift distribution (top panel), (b) the height distribution (middle panel), or (c) the angular clustering (bottom panel) for 0.5 arcmin (grey), 1.0 arcmin (blue), 2.0 arcmin (green), or the combination of all three (black) smoothing scales. The redshift distribution shows the clearest improvement when combining multiple smoothing scales. The height distribution shows a weaker improvement, whereas the angular clustering shows the characteristic WL banana shape.

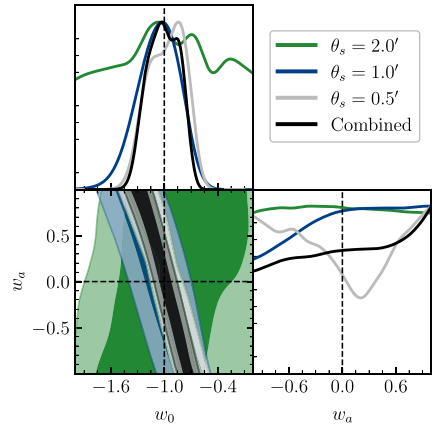
Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026




(a) Redshift distribution




(b) Height distribution




(c) Angular clustering

**Figure 6.** Same as Fig. 5 but for $w_0$−$w_a$ posteriors. For the redshift distribution, a smaller smoothing scale yields tighter constraints, while combining smoothing scales yields significantly tighter constraints. The height distribution shows a weaker improvement, and even in the combined case the angular clustering remains insensitive to the time evolution of DE.

of $\theta_s = 0.5$ (grey), 1.0 (blue), and 2.0 (green) arcmin, as well as the forecast using the three scales combined (black). The results are further quantified in Table 2. There, for each statistic at the different smoothing scales, as indicated by the first column, the FoM in the $\Omega_m - \ln(10^{10}A_s)$ and $w_0$−$w_a$ planes are listed in the second and fourth columns, respectively. We only interpret the relative changes in the FoMs due to our idealized forecasts. Similarly, the corresponding FoB values are listed in the third and fifth columns.

The FoB at the 68th percentile confidence level for this 2D inference is 1.52. Forecasts with a bias below this threshold are therefore less biased than the commonly used 1σ value to quantify the bias for a single parameter. The FoB values are below the 1.52 threshold in almost every case. Only for the $w_0$−$w_a$ forecast using the redshift distribution at $\theta_s = 2.0$ arcmin the FoB is 2.35.

The three statistics show similar improvements when combining the three smoothing scales. For the redshift distribution, combining all three scales improves the FoM by a factor of 2.15 and 2.66 for $\Omega_m - \ln(10^{10}A_s)$ and $w_0$−$w_a$, respectively, compared to only using our fiducial 1 arcmin scale. The absolute improvement is the largest of the three statistics. The height distribution shows a relatively similar improvement for $\Omega_m - \ln(10^{10}A_s)$ of 2.33, while for $w_0$−$w_a$ the combined FoM improves by a factor of 1.41 relative to the fiducial scale. The relative gain is similar to that of the redshift distribution, even though the absolute constraining power of the height distribution on $w_0$−$w_a$ remains weaker than that of the redshift distribution at every individual scale. The angular clustering shows the largest relative improvement of all three statistics, with the combined FoM exceeding the fiducial-scale value by a factor of 8.76 for $\Omega_m - \ln(10^{10}A_s)$ and 2.46 for $w_0$−$w_a$. As mentioned before, C. T. Davies et al. (2022) found that combining peaks selected at multiple SNR thresholds substantially improved angular clustering constraints. Instead, we compared peaks selected from the same SNR threshold but from different smoothing scales, which qualitatively probe the same information. We also find that the results improve significantly when combining multiple, differently selected angular clustering distributions. However, in terms of absolute constraining power, the angular clustering remains the poorest-performing statistic in both parameter planes.

Comparing the different statistics at the fiducial 1 arcmin smoothing scale, the redshift distribution dominates the constraining power in both parameter planes, exceeding the height distribution by a factor of 6 and the angular clustering by a factor of 116 in FoM for $\Omega_m - \ln(10^{10}A_s)$, and by factors of 4 and 24, respectively, for $w_0$−$w_a$.

Regarding the sensitivity to the smoothing scale itself, each of the three statistics shows a monotonic trend in which the smallest scale yields the tightest constraints in both parameter planes. Given the number of covariance matrix realizations, it is not possible to include an extra scale in the forecast as the number of data vector bins for the redshift distribution would exceed the number of covariance matrix realizations, even though this might further improve the constraints. However, forecasts of each of the three statistics in both parameter planes, with 5 arcmin smoothing, show that constraints weaken further (not shown), consistent with the trend that smaller smoothing scales yield better constraints.

Compared to the findings of J. Liu et al. (2015), who identified ≈1 arcmin as close to the optimal smoothing scale for a CFHTLenS-depth inference based solely on the peak number density, our results suggest that, for a *Euclid* inference using SNR > 3 peaks, a slightly smaller scale is preferred, likely due to (a combination of) the better angular resolution, lower noise levels, and higher source number density. The smallest smoothing scale we consider corresponds to only 1.16 pixels for the $N_{side} = 8192$ maps that are used. As HEALPIX's functionality, which we need to smooth and rotate our mass maps, is limited by 32-bit indexing,

Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

**Table 2.** FoM and FoB for the $\Omega_m - \ln(10^{10}A_s)$ and $w_0 - w_a$ parameter planes when all other cosmological parameters are fixed while varying the smoothing scale for a forecast using the redshift distribution, the height distribution, or angular clustering, at smoothing scales of $\theta_s = 0.5$, 1.0, or 2.0 arcmin, and for the combination of all three scales. The FoM and FoB corresponding to contours that hit the prior bounds, as shown in Figs 5 and 6, should be treated carefully, as the prior truncation may bias the inferred covariance used to compute the FoM and FoB. The redshift distribution provides the tightest constraint on both parameter combinations. The FoM of all statistics and parameter combinations typically increases by a factor of $\approx 2$ when combining the different scales.

| Statistic and smoothing scale | $\Omega_m - \ln(10^{10}A_s)$ | | $w_0 - w_a$ | |
|---|---|---|---|---|
| | Figure of merit | Figure of bias | Figure of merit | Figure of bias |
| Redshift distribution at $\theta_s = 0.5$ arcmin | 18780.5 | 0.87 | 347.2 | 0.90 |
| Redshift distribution at $\theta_s = 1$ arcmin | 17850.2 | 0.75 | 243.8 | 0.73 |
| Redshift distribution at $\theta_s = 2$ arcmin | 10744.9 | 1.31 | 151.9 | 2.35 |
| Redshift distribution at $\theta_s = 0.5$ arcmin + $\theta_s = 1$ arcmin + $\theta_s = 2$ arcmin | 38363.5 | 1.14 | 649.0 | 0.48 |
| Height distribution at $\theta_s = 0.5$ arcmin | 2868.8 | 0.26 | 59.6 | 1.24 |
| Height distribution at $\theta_s = 1$ arcmin | 2842.6 | 0.28 | 57.5 | 0.21 |
| Height distribution at $\theta_s = 2$ arcmin | 2359.0 | 0.07 | 39.8 | 1.05 |
| Height distribution at $\theta_s = 0.5$ arcmin + $\theta_s = 1$ arcmin + $\theta_s = 2$ arcmin | 6623.2 | 0.75 | 80.9 | 0.28 |
| Angular clustering at $\theta_s = 0.5$ arcmin | 251.2 | 0.50 | 16.4 | 0.14 |
| Angular clustering at $\theta_s = 1$ arcmin | 153.5 | 0.43 | 10.2 | 0.31 |
| Angular clustering at $\theta_s = 2$ arcmin | 105.2 | 0.26 | 3.2 | 0.41 |
| Angular clustering at $\theta_s = 0.5$ arcmin + $\theta_s = 1$ arcmin + $\theta_s = 2$ arcmin | 1344.7 | 0.95 | 25.1 | 0.11 |

we did not explore smaller scales, even though they may provide even stronger constraints.

The statistics exhibit qualitatively different directions of degeneracy. The redshift and height distributions both show negatively correlated degeneracies between $\Omega_m$ and $\ln(10^{10}A_s)$ at every smoothing scale, whereas for the angular clustering forecast, the $\theta_s = 0.5$ and 1 arcmin $\Omega_m - \ln(10^{10}A_s)$ posteriors show a banana-shaped degeneracy that is not apparent in the fiducial 10D forecast, which is similar to the banana-shaped posteriors typically found in WL two-point cosmic shear inferences in the $\sigma_8 - \Omega_m$ plane. In our inference, the curvature is a consequence of the non-monotonic $\Omega_m$ dependence of the clustering signal, as shown before in Fig. 3. The clustering amplitude increases with $\Omega_m$ only up to $\Omega_m \approx 0.3$, beyond which the impact decreases and a reversed trend is even observed at large scales. At the same time, within our hypercube range, the dependence of clustering amplitude on $\ln(10^{10}A_s)$ is always monotonic, with larger values of $\ln(10^{10}A_s)$ resulting in a stronger clustering signal (not shown). Despite being noisier, the 2 arcmin scale shows a different degeneracy than the other two smoothing scales. Consequently, when combining the scales, the slightly different degeneracies cause a factor $\approx 9$ increase in the FoM compared to the fiducial smoothing scale. It is unclear whether this represents a different physical behaviour or is instead a result of some part of the pipeline. Similar, somewhat confusing behaviour can be seen in the $w_0 - w_a$ posteriors for the angular clustering forecast at 2 arcmin, which appear to reverse the degeneracy direction, albeit remaining practically unconstrained within the prior ranges.

More generally, our exploration of sensitivity to the smoothing scale is illustrative rather than exhaustive. Although we typically find that the smaller smoothing scale yields the tightest constraint, we have considered only three discrete scales and only in two restricted parameter planes, and we have ignored baryonic physics, which potentially impacts the smallest scales the strongest. A more complete analysis that combines multiple smoothing scales and multiple statistics across the full 10D forecast is left for future work. Similarly, we did not explore the possibility of mixing different scales for the different statistics.

## 6 CONCLUSIONS

We have investigated the potential of using WL peak properties to constrain cosmology in a *Euclid*-like set-up that indirectly aims to exploit the cosmology dependence of the mass, angular clustering, and redshift distribution of massive structures. We have done so by forward-modelling the final expected signal in a new DMO hypercube, a suite of simulations spanning a 10D cosmological parameter space varying the matter density, evolving DE, neutrino mass, and beyond $w_0 w_a$ CDM parameters that control the running of the scalar spectral index and the decay rate of CDM. We have run an additional suite of 80 $\Lambda$CDM DMO simulations within the same simulation framework to estimate the covariance matrix at the fixed D3A cosmology.

Within these simulations, we generate WL convergence maps from which we extract the local maxima, i.e. WL peaks. We study the cosmological constraining potential of the different WL peak properties, namely (the combination of) their angular clustering, redshift distribution, and height distributions. We find that different WL peak statistics probe different regions of the cosmological space. For our fiducial set-up that uses SNR $> 3$ peaks measured from a map smoothed with 1 arcmin, the WL peak redshift distribution provides the tightest constraints on the matter density, $\Omega_m$, and on the DE equation-of-state parameters, $w_0$ and $w_a$. The angular clustering and peak-height distribution jointly constrain best the amplitude of the primordial power spectrum, $\ln(10^{10}A_s)$. The combination of different statistics can also probe the Hubble parameter, $h$, and the baryon matter density, $\Omega_b$, through the breaking of parameter degeneracies. None of the considered WL peak statistics, even when combined, provide strong constraints on the DM decay rate, the neutrino mass, or the amplitude and running of the scalar spectral index.

We interpret the different cosmological constraining power by considering three pathways through which cosmological forecasts are affected by changes in cosmology, namely through direct impact on the matter density, changes in geometry, or changes in structure formation. Although it is difficult to isolate the effect of a single channel, we find that all channels likely contribute to the constraints, given that the different statistics



Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

exhibit distinct degeneracies and sensitivities to cosmological parameters.

We additionally compare the constraining power to that of the commonly used two-point cosmic-shear correlation functions. We find that for most parameters, the combination of WL peak statistics provides equal or better constraining power than the $\xi_{+/-}$. Only for the neutrino mass and the (running of the) scalar spectral index do we find that the shear 2PCFs outperform the WL peaks, which do not constrain these parameters well. When combining both, the constraints improve further.

We additionally explore the impact of the smoothing scale to explore the optimal scale and complementarity of different scales in the 2D $\Omega_\mathrm{m} - \ln(10^{10} A_\mathrm{s})$ and $w_0 - w_\mathrm{a}$ planes by comparing cosmological forecasts for peaks selected from maps smoothed with 0.5, 1, and 2 arcmin. For both parameter planes, we find the redshift distribution to provide the tightest constraints, with its FoM being a factor $\approx 10-100$ larger than the corresponding FoMs for the angular clustering or height distribution. When combining different smoothing scales, which select partially independent sets of objects with different masses and redshifts, we find the FoMs to typically improve by a factor of $\approx 2$.

The cosmological forecast set-up of this paper is idealized, as it ignores the impact of baryonic physics, IA, source clustering, reduced shear, and nuisance parameters that are expected to significantly impact the results. Although the impact of some of these effects and systematics is expected to be minor or qualitatively different for the different WL peak statistics we considered, future analyses will need to better account for these effects to obtain precise and accurate cosmological constraints from upcoming data from *Euclid*, *Roman*, and *Rubin*.

In addition, we restricted our analysis to WL peaks and did not consider other higher-order WL statistics in detail, although including them can further improve the constraints (see e.g. G. A. Marques et al. 2024, for the inclusion of WL minima).

## ACKNOWLEDGEMENTS

This work was supported by the Science and Technology Facilities Council (grant no. ST/Y002733/1). This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 769130). This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1, and ST/S002502/1, Durham University, and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure. We have used ASTROPY (Astropy Collaboration 2022), HEALPY (A. Zonca et al. 2019), MATPLOTLIB (J. D. Hunter 2007), NUMPY (C. R. Harris et al. 2020), and SWIFTSIMIO (J. Borrow & A. Borrisov 2020) in our analysis.

## DATA AVAILABILITY

The data supporting the figures in this article are available upon reasonable request to the corresponding author. The hypercube will be presented in McCarthy et al. (in preparation) and will eventually be made public.

## REFERENCES

Abbott T. M. C. et al., 2022, Phys. Rev. D, 105, 023520
Abdul Karim M. et al., 2025, Phys. Rev. D, 112, 083515
Adame A. G. et al., 2025, J. Cosmol. Astropart. Phys., 2025, 028
Allys E., Marchand T., Cardoso J.-F., Villaescusa-Navarro F., Ho S., Mallat S., 2020, Phys. Rev. D, 102, 103506
Astropy Collaboration, 2022, ApJ, 935, 167
Bartelmann M., Schneider P., 2001, Phys. Rep., 340, 291
Bilicki M. et al., 2021, A&A, 653, A82
Borrow J., Borrisov A., 2020, J. Open Source Softw., 5, 2430
Bradač M., Schneider P., Lombardi M., Erben T., 2005, A&A, 437, 39
Broxterman J. C. et al., 2024, MNRAS, 529, 2309
Broxterman J. C., Schaller M., Hoekstra H., Schaye J., McGibbon R. J., Forouhar Moreno V. J., Kugel R., Elbers W., 2025, MNRAS, 538, 755
Chen K.-F. et al., 2025, Open J. Astrophys., 8, 2
Cheng S., Ting Y.-S., Ménard B., Bruna J., 2020, MNRAS, 499, 5902
Chevallier M., Polarski D., 2001, Int. J. Mod. Phys. D, 10, 213
Chiu I.-N. et al., 2024, Open J. Astrophys., 7, 90
DES Collaboration, 2026a, preprint (arXiv:2601.14559)
DES Collaboration, 2026b, preprint (arXiv:2605.27221)
Davies C. T., Cautun M., Li B., 2019, MNRAS, 488, 5833
Davies C. T., Cautun M., Giblin B., Li B., Harnois-Déraps J., Cai Y.-C., 2022, MNRAS, 513, 4729
de Jong J. T. A., Verdoes Kleijn G. A., Kuijken K. H., Valentijn E. A., 2013, Exp. Astron., 35, 25
Dietrich J. P., Hartlap J., 2010, MNRAS, 402, 1049
Elbers W. et al., 2025a, Phys. Rev. D, 112, 083513
Elbers W. et al., 2025b, MNRAS, 537, 2160
Euclid Collaboration: Ajani V. et al., 2023, A&A, 675, A120
Euclid Collaboration: Blanchard A. et al., 2020, A&A, 642, A191
Euclid Collaboration: Cañas-Herrera G. et al., 2025, A&A, 713, A188
Euclid Collaboration: Mellier Y. et al., 2025, A&A, 697, A1
Euclid Collaboration: Vinciguerra S. et al., 2026, A&A, 707, A235
Feldbrugge J., van Engelen M., van de Weygaert R., Pranav P., Vegter G., 2019, J. Cosmol. Astropart. Phys., 2019, 052
Fields B. D., Olive K. A., Yeh T.-H., Young C., 2020, J. Cosmol. Astropart. Phys., 2020, 010
Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306
Gatti M. et al., 2024, MNRAS, 527, L115
Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759
Grandón D., Marques G. A., Thiele L., Cheng S., Shirasaki M., Liu J., 2024, Phys. Rev. D, 110, 103539
Hahn O., Michaux M., Rampf C., Uhlemann C., Angulo R. E., 2020, Astrophysics Source Code Library, record ascl:2008.024
Hamana T., Takada M., Yoshida N., 2004, MNRAS, 350, 893
Harris C. R. et al., 2020, Nature, 585, 357
Heavens A., Whiteway L., Sellentin E., 2026, Open J. Astrophys., 9, 65025
Heitmann K., Higdon D., White M., Habib S., Williams B. J., Lawrence E., Wagner C., 2009, ApJ, 705, 156
Heydenreich S., Brück B., Harnois-Déraps J., 2021, A&A, 648, A74
Hildebrandt H. et al., 2020, A&A, 633, A69
Hoekstra H., Jain B., 2008, Annu. Rev. Nucl. Part. Sci., 58, 99
Hunter J. D., 2007, Comput. Sci. Eng., 9, 90
Jain B., Van Waerbeke L., 2000, ApJ, 530, L1
Kilbinger M., 2015, Rep. Prog. Phys., 78, 086901
Kodwani D., Alonso D., Ferreira P., 2019, Open J. Astrophys., 2, 3
Kratochvil J. M., Haiman Z., May M., 2010, Phys. Rev. D, 81, 043519
Kugel R. et al., 2023, MNRAS, 526, 6103
Kugel R., Schaye J., Schaller M., McCarthy I. G., Braspenning J., Helly J. C., Forouhar Moreno V. J., McGibbon R. J., 2024, MNRAS, 534, 2378
Kugel R., Schaye J., Schaller M., Forouhar Moreno V. J., McGibbon R. J., 2025, MNRAS, 537, 2179
Lahav O., Liddle A. R., 2024, preprint (arXiv:2403.15526)
Landy S. D., Szalay A. S., 1993, ApJ, 412, 64
Lee M. E., Haiman Z., Pandey S., Genel S., 2026, ApJ, 996, 36

Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

Lesgourgues J., 2011, preprint (arXiv:1104.2932)
Lewis A., 2025, J. Cosmol. Astropart. Phys., 2025, 025
Li Z., Liu X., Fan Z., 2023, MNRAS, 520, 6382
Limber D. N., 1953, ApJ, 117, 134
Linder E. V., 2003, Phys. Rev. Lett., 90, 091301
Linke L. et al., 2026, preprint (arXiv:2606.12389)
Liu J., Haiman Z., Hui L., Kratochvil J. M., May M., 2014, Phys. Rev. D, 89, 023515
Liu J., Petri A., Haiman Z., Hui L., Kratochvil J. M., May M., 2015, Phys. Rev. D, 91, 063507
LSST Science Collaboration, 2009, preprint (arXiv:0912.0201)
Lu T., Haiman Z., 2021, MNRAS, 506, 3406
Marian L., Smith R. E., Hilbert S., Schneider P., 2013, MNRAS, 432, 1338
Marini I. et al., 2024, A&A, 689, A7
Marques G. A., Liu J., Zorrilla Matilla J. M., Haiman Z., Bernui A., Novaes C. P., 2019, J. Cosmol. Astropart. Phys., 2019, 019
Marques G. A. et al., 2024, MNRAS, 528, 4513
Michaux M., Hahn O., Rampf C., Angulo R. E., 2021, MNRAS, 500, 663
Miyazaki S. et al., 2012, in McLean I. S., Ramsay S. K., Takami H., eds, Proc. SPIE Conf. Ser. Vol. 8446, Ground-based and Airborne Instrumentation for Astronomy IV. SPIE, Bellingham, p. 84460Z
Pantos I., Perivolaropoulos L., 2026, Phys. Dark Universe, 52, 102286
Pedregosa F. et al., 2011, J. Mach. Learn. Res., 12, 2825
Planck Collaboration VI, 2020, A&A, 641, A6
Rasmussen C. E., Williams C. K. I., 2006, Gaussian Processes for Machine Learning
Reischke R. et al., 2026, A&A, 709, A82
Rykoff E. S. et al., 2014, ApJ, 785, 104
Schaller M. et al., 2024, MNRAS, 530, 2378
Schaye J. et al., 2023, MNRAS, 526, 4978
Schneider P., van Waerbeke L., Mellier Y., 2002, A&A, 389, 729
Seitz C., Schneider P., 1995, A&A, 297, 287
Sellentin E., Heavens A. F., 2016, MNRAS, 456, L132
Semboloni E., Hoekstra H., Schaye J., van Daalen M. P., McCarthy I. G., 2011, MNRAS, 417, 2020
Semboloni E., Hoekstra H., Schaye J., 2013, MNRAS, 434, 148
Sinha M., Garrison L. H., 2020, MNRAS, 491, 3022
Spergel D. et al., 2015, preprint (arXiv:1503.03757)
The Dark Energy Survey Collaboration, 2005, preprint (astro-ph/0510346)
Vedder C., Bakx T., Chisari N. E., Hoekstra H., Schaller M., 2026, MNRAS, 550, stag1008
Vinciguerra S., Martinet N., Gatti M., 2026, preprint (arXiv:2607.17942)
Wang J., Elbers W., Frenk C. S., Cole S., Yang X., McCarthy I. G., van de Weygaert R., 2026, MNRAS, 549, stag937
Wang S., Haiman Z., May M., 2009, ApJ, 691, 547
Weiss A. J., Schneider A., Sgier R., Kacprzak T., Amara A., Refregier A., 2019, J. Cosmol. Astropart. Phys., 2019, 011
Williamson J. et al., 2026, preprint (arXiv:2606.11309)
Wright A. H. et al., 2025, A&A, 703, A158
Yang X., Kratochvil J. M., Wang S., Lim E. A., Haiman Z., May M., 2011, Phys. Rev. D, 84, 043529
Yang Y., Bird S., Ho M.-F., 2025, Phys. Rev. D, 111, 083529
Zhang T., Liu X., Wei C., Li G., Luo Y., Kang X., Fan Z., 2022, ApJ, 940, 96
Zonca A., Singer L., Lenz D., Reinecke M., Rosset C., Hivon E., Gorski K., 2019, J. Open Source Softw., 4, 1298
Zürcher D., Fluri J., Sgier R., Kacprzak T., Refregier A., 2021, J. Cosmol. Astropart. Phys., 2021, 028

## APPENDIX A: EMULATOR CONSTRUCTION AND VALIDATION

In this appendix, we provide additional details on the construction and accuracy of the emulators used for the forecasting. We use Gaussian process (GP) emulators that are built using SCIKIT-LEARN's `GaussianProcessRegressor` class (F. Pedregosa et al. 2011). All emulators use a Matérn-$\frac{3}{2}$ kernel with a trainable amplitude (see e.g. C. E. Rasmussen & C. K. I. Williams 2006). We add a white-noise kernel to absorb stochasticity, as can be seen in, for example, the high-$\kappa$ tail of the peak height distributions in Fig. 1.

The kernel hyperparameters are optimized by maximizing the GP log-marginal likelihood, with the optimization repeated 10 times with random restarts to avoid local maxima. The different WL peak summary statistics are emulated separately at each smoothing scale, and the resulting emulators are concatenated before forecasting. The different statistics and their corresponding emulator architectures are described below.

### A1 Redshift distribution

SNR $> 3$ peaks are selected and binned in redshift using 25 linearly spaced bins over $z \in [0,\ 1.25]$, with bin widths normalized to yield the density $\mathrm{d}N/\mathrm{d}z$. We chose this SNR cut because in J. C. Broxterman et al. (2025), we showed WL peaks above the corresponding $\kappa$-cut can be robustly assigned to a single massive object. Also, lower-valued peaks are more sensitive to the

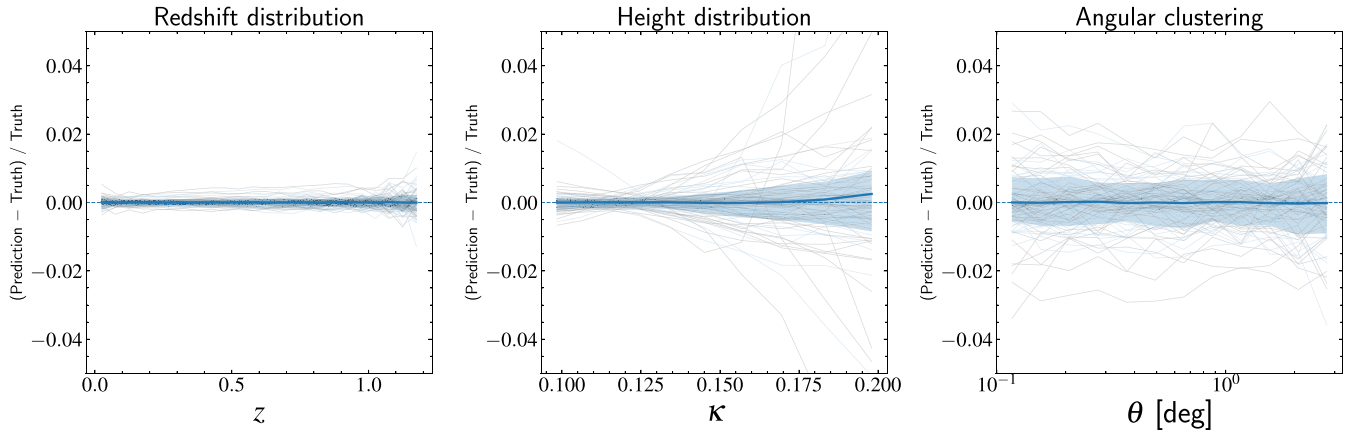



**Figure A1.** LOO emulator test showing the fractional accuracy of the predicted statistics by training on all but one hypercube node and predicting the statistics at the cosmology of the final hypercube node for the SNR $> 3$ WL peak redshift distribution (left panel), height distribution (middle panel), and angular clustering (right panel), corresponding to the statistics shown in Fig. 1. The emulators show per-cent-level accuracy across the complete ranges for which the statistics are used in the forecast.



Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

noise and instrument systematics, as discussed in appendix B of J. C. Broxterman et al. (2024). In this way, we aim at a potentially robust cosmological inference. Note, however, that we apply stronger noise, so the SNR differs. The binning corresponds to the FLAMINGO light-cone redshift resolution and is similar to the expected redshift uncertainty on the massive objects that generate the high-valued WL peaks. For additional details, see the discussion in section 3.5 of J. C. Broxterman et al. (2025). Before emulating the distributions, we apply a logarithmic transform $\ln(\mathrm{d}N/\mathrm{d}z + \epsilon)$ (with $\epsilon = 10^{-12}$) to stabilize the variance at low counts. The data vectors are then emulated using principal component analysis (PCA) separately for the shape of the redshift distribution and the normalization, i.e. the total number of peaks above the SNR $> 3$ threshold. We emulate the first 15 PCA components. When predicting the redshift distribution, the shape is reconstructed from an inverse PCA transform and multiplied by the amplitude prediction to yield the full data vector.

### A2 Height distribution

Peaks above the SNR threshold are binned into 10 logarithmically spaced SNR bins, with the upper bin edge set at the 99th percentile of the $\kappa$ distribution across all training simulations. We do this because preliminary tests showed it yielded better predictions than including all high-end $\kappa$ tail peaks, likely due to cosmic variance. At the same time, we do not expect these peaks to carry much of the cosmological information, given the relative increase in cosmic variance, as shown in the spread among the covariance matrix realizations in Fig. 1. Nevertheless, this may be improved in future work. We again normalize the counts to a density $n(\nu) = \mathrm{d}N/\mathrm{d}\kappa$, and apply a logarithmic transform $\ln[n(\nu) + \epsilon]$ before emulation. For the height distribution, we use the same 15-component PCA + amplitude GP emulator as described for the peak redshift distribution.

### A3 Angular clustering

The angular clustering of the peaks is measured in 12 logarithmically spaced bins between $\theta = 0.1{-}3.0$ deg. We chose this range because it corresponds to the angular range we also use for our $\xi_{+/-}$ 2PCFs. To reduce the dynamic range and improve the emulation accuracy, we transform the correlation function distributions as $\tilde{\omega}(\theta) = \ln[\theta\,\omega(\theta)]$. For $\tilde{\omega}(\theta)$, we adopt a similar PCA GP emulator.

### A4 Leave-one-out test

We assess the emulator accuracy by performing a LOO test. Each of the $2 \times 50$ training simulations is withheld in turn, and an emulator is trained on the remaining 99 nodes. We then evaluate the prediction at the cosmology of the remaining node and compare it to the actual measurements. For the fiducial forecast set-up considered in this work, Fig. A1 shows the emulator's fractional accuracy on the log-transformed quantities. The left panel shows the redshift distribution, the central panel the height distribution, and the right panel the angular clustering. For each statistic, averaged over the entire emulated range, the emulator typically has an accuracy of within 1 per cent. The mean $1\sigma$ accuracy averaged across all bins is 0.09, 0.40, and 0.58 per cent for the redshift distribution, height distribution, and angular clustering, respectively. Additionally, for the clustering, although the simulations are predicted with per-cent-level accuracy on average, there are individual simulations in which the angular clustering is mispredicted by up to $\approx 3$ per cent across all angular scales.

We stress that we designed the emulators to perform well in the fiducial forecast set-up, i.e. with 1 arcmin smoothing while inferring all 10 cosmological parameters simultaneously. We did not change the architecture for the 0.5 or 2.0 arcmin smoothing tests. Inspection of the emulator LOO accuracy tests shows that the emulators do not perform much worse at those smoothing scales.

## APPENDIX B: CORRELATION MATRIX

In this appendix, we present the correlation matrix for our fiducial forecast set-up. Fig. B1 shows the Pearson correlation coefficient values, $\rho_{i,j} = C_{ij}/(\sigma_i\sigma_j)$, where $C_{ij}$ is the $i$, $j$th component of the covariance matrix and $\sigma_i$ and $\sigma_j$ denote the standard deviation of the $i$th and $j$th component, respectively, for the forecast corresponding to the black contours in Figs 2 and C1. The auto- and cross-correlation blocks corresponding to separate (combinations of) statistics are separated by solid black lines. From bottom left to top right, the figure shows the clustering, $\omega(\theta)$, height distribution, $\mathrm{d}n(\kappa)/\mathrm{d}\kappa$, and redshift distribution, $\mathrm{d}n(z)/\mathrm{d}z$. The $x$-axis indicates the angular, WL convergence, or redshift bin values. The clustering shows strong correlations, particularly between individual low-angular-separation bins and between individual high-separation bins, whereas the other statistics and cross-correlations show milder correlations, though there is still typically some moderate correlation.

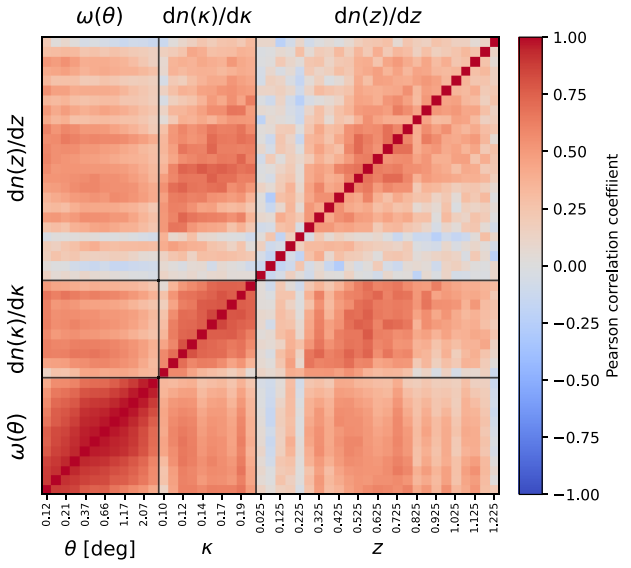



**Figure B1.** Matrix of Pearson correlation coefficients for the inference for the fiducial inference set-up. The solid black curves separate the bins corresponding to the different WL peak statistics. Bottom-left panel: Angular clustering of SNR $> 3$ peaks. Centre-middle panel: Height distribution. Top-right panel: Redshift distribution of SNR $> 3$ WL peaks. The different clustering bins are strongly correlated, whereas the other statistics and cross-correlations show weaker correlations.

Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

## APPENDIX C: COSMOLOGICAL CONSTRAINTS FROM WL PEAKS

In this appendix, we present the full corner plot of the 2D posteriors for the analysis of the three different WL peak statistics jointly, the redshift distribution individually, and the angular clustering and height distribution combined, with 1 arcmin smoothing, whose individual 1D parameter posteriors were already shown in Fig. 2. Fig. C1 shows the posteriors corresponding to the forecast using the redshift distribution (orange), height distribution and angular clustering (purple), or all three distributions (black).

Individually, the WL peak statistics cannot provide strong constraints over the entire 10D parameter space. $\Omega_m$, $w_0$, and $w_a$ are most tightly constrained by the redshift distribution. The peak height distribution and angular clustering constrain $\ln(10^{10}A_s)$ best, whereas $\Omega_b$ and $h$ are primarily constrained due to breaking

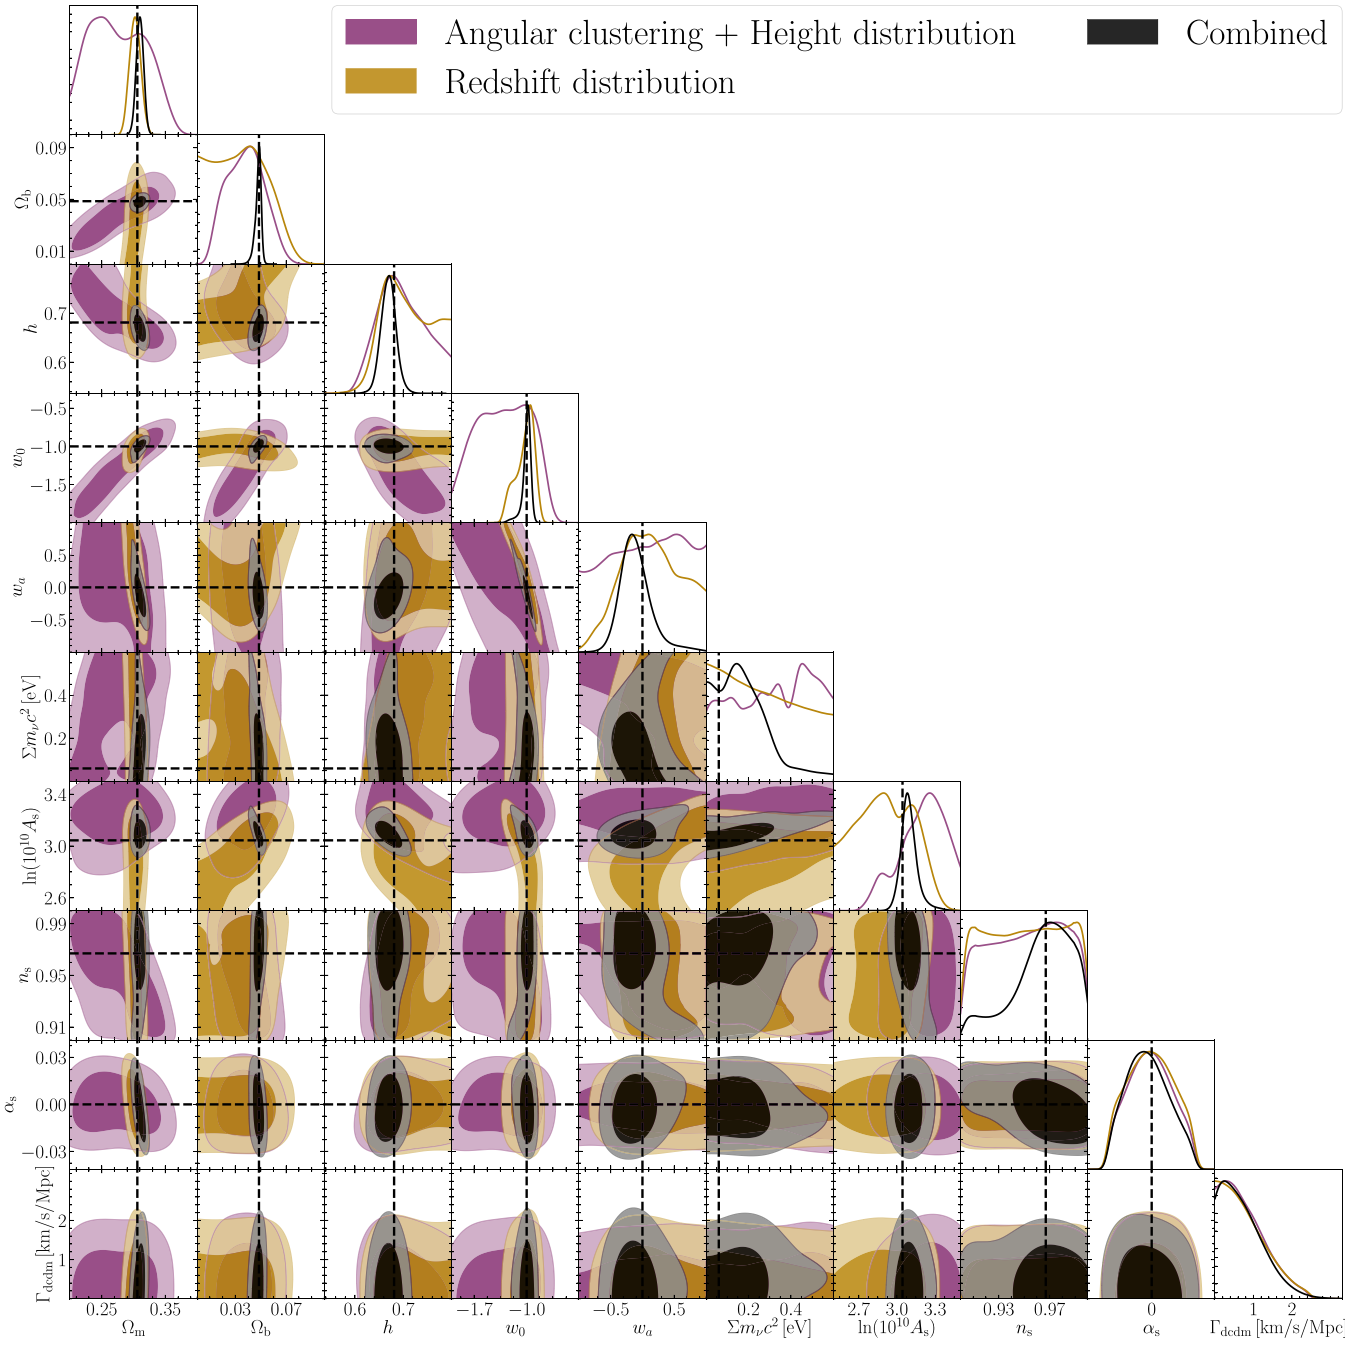



**Figure C1.** Corner plot of the cosmological parameter forecast using the angular clustering and height distributions (purple), redshift distribution (orange), or the combination of the three statistics (black) of WL peaks. The contours indicate the $1\sigma$ and $2\sigma$ credible intervals. Different statistics are sensitive to different parameters, with the redshift distribution being most sensitive to $\Omega_m$, $w_0$, $w_a$, and the angular clustering and height distribution jointly probing $\ln(10^{10}A_s)$. Additionally, due to the breaking of degeneracies, the combination of statistics is also able to probe $\Omega_b$ and $h$, whereas $\sum m_\nu c^2$, $n_s$, $\alpha_s$, and $\Gamma_{dcdm}$ remain unconstrained.



Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

of degeneracies in the combined forecast, as can be seen in the $\Omega_m - \Omega_b$ and $\Omega_m - h$ panels. $\Sigma m_\nu c^2$, $n_s$, $\alpha_s$, and $\Gamma_{dcdm}$ remain unconstrained or recover the imposed prior even in the combined case.

The corner plot illustrates the complementarity of combining the different statistics to break degeneracies and obtain tighter constraints. In particular, in the $\Omega_m$–$\Omega_b$ plane, no individual statistic can simultaneously constrain both parameters, whereas their combination yields a clearly localized posterior, which seems to result from the combination of the angular clustering and the redshift distribution.

Some of the 2D contours, particularly when parameters are not well constrained, show multimodal posteriors, which may indicate an insufficiently regular emulator. This particularly holds for the angular clustering and height distribution combination, as for instance seen in the $\Omega_m - \Sigma m_\nu c^2$, where Appendix A already showed that individual nodes can be misestimated on average by several per cent. This behaviour could indicate that improvements in emulation or additional hypercube nodes are necessary to obtain accurate predictions across the entire 10D cosmological parameter space.

## APPENDIX D: IMPACT OF INDIVIDUAL COSMOLOGICAL PARAMETERS ON THE WL PEAK STATISTICS

In this appendix, we show the impact of additional individual cosmological parameters on the WL peak statistics using emulator predictions, analogous to the distributions shown for varying $\Omega_m$ in Fig. 3. Nine out of the ten parameters that are input to the emulator are set to their mean value, and the 10th parameter is varied within its range within the hypercube. Figs D1, D2, D3, and D4, respectively, show the redshift distribution (left), height distribution (centre), and angular clustering (right) for varying $h$, $w_0$, $w_a$, and $\Omega_b$ as indicated by the colour bar. Individually varying all remaining hypercube parameters results in smaller variations than seen for these parameters.

Varying $h$ affects each of the WL peak statistics. The amplitude of the redshift distribution has a positive correlation with $h$. The height distribution is more sensitive to $h$, as larger values of $h$ systematically increase the abundance of peaks at fixed $\kappa$, with larger $\kappa$ values being preferred for larger $h$ values. For the angular clustering distributions, larger $h$ increases the signal for $\theta < 1$ deg.

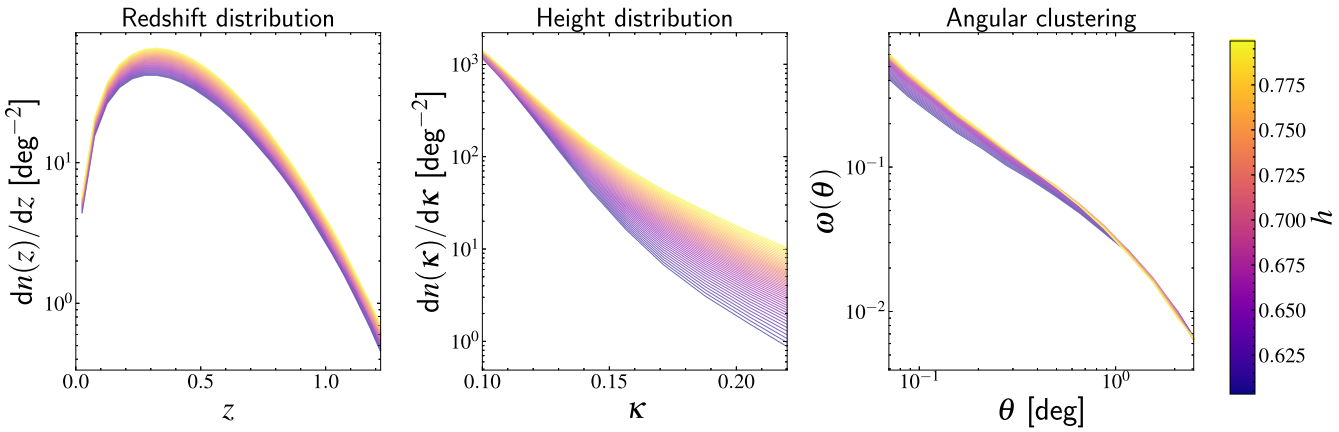



**Figure D1.** Same as Fig. 3 but varying $h$. Larger values of $h$ generally enhance the amplitude of each of the three signals. Additionally, the height distribution is enhanced preferentially at large $\kappa$ and the clustering signal at small $\theta$.

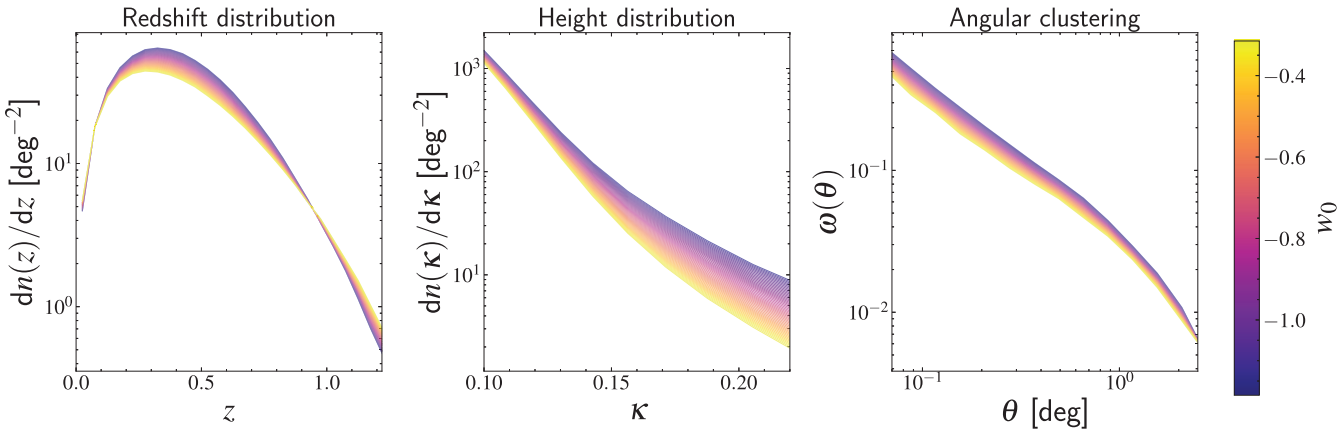



**Figure D2.** Same as Fig. 3 but varying $w_0$. Stronger evolving DE, corresponding to less negative values of $w_0$, suppresses each of the considered statistics, with the suppression most pronounced at low redshift, high $\kappa$, and small $\theta$.

Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

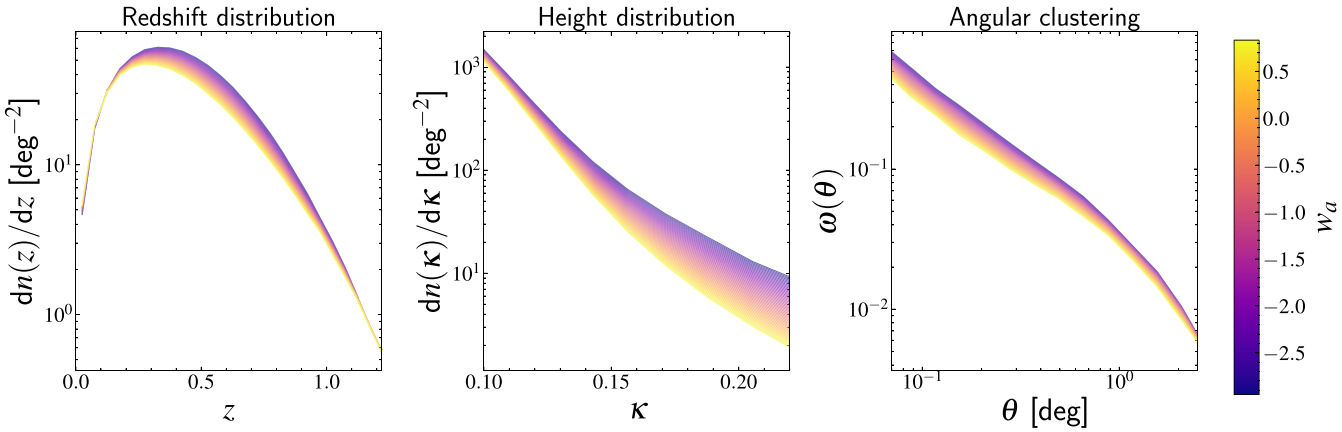



**Figure D3.** Same as Fig. 3 but varying $w_a$. The impact of varying $w_a$ is very similar to that of varying $w_0$ as shown in Fig. D2.

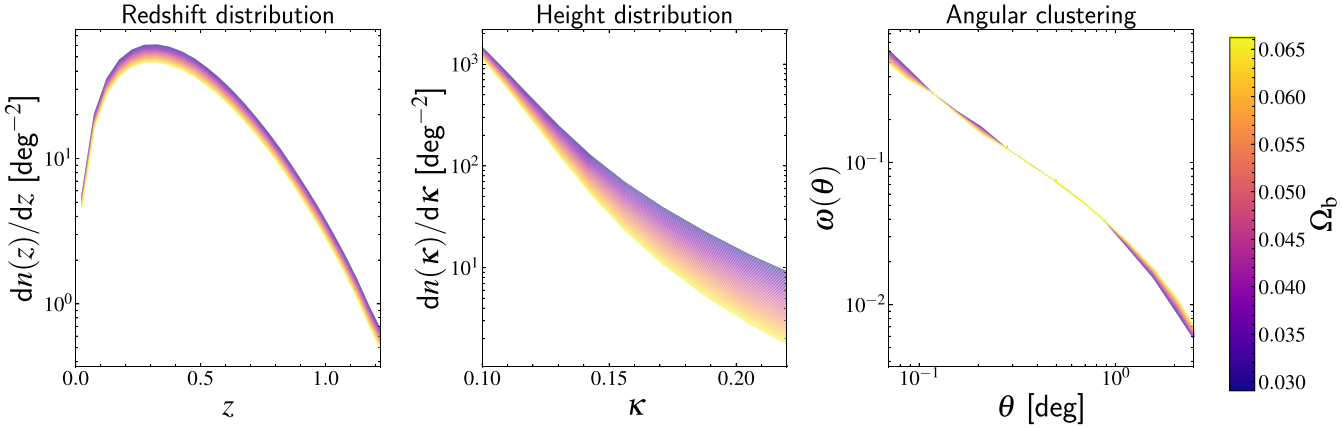



**Figure D4.** Same as Fig. 3 but varying $\Omega_b$. The impact of varying $\Omega_b$ is most distinct on the height distribution, where lower values of $\Omega_b$ lead to more higher-valued WL peaks at fixed $\kappa$.

Stronger evolving DE, with larger values of $w_0$ or $w_a$, leads to a suppression of each of the three WL peak statistics. The behaviour of $w_0$ and $w_a$ is very similar across the three statistics. Only above $z \approx 1$ do larger values of $w_0$ lead to a small enhancement of the number of peaks, whereas for the $w_a$ values considered, there is only a suppression.

Varying $\Omega_b$ only has a minor impact on the redshift distribution and angular clustering, whereas the height distribution shows a stronger dependence, with higher values of $\Omega_b$ leading to a suppression of the number of high-valued peaks.

This paper has been typeset from a TEX/LATEX file prepared by the author.

Downloaded from academic.oup.com/mnras/article/553/1/stag1905/8885810 by Sarah Dakin user on 09 November 2026

© The Author(s) 2026.
Published by Oxford University Press on behalf of Royal Astronomical Society. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.