

Cosmological forecasts with the clustering of weak lensing peaks

Christopher T. Davies,¹★ Marius Cautun², Benjamin Giblin³, Baojiu Li⁴,
Joachim Harnois-Déraps^{5,6} and Yan-Chuan Cai³

¹Faculty of Physics, Ludwig-Maximilians-Universität, Scheinerstr. 1, D-81679 Munich, Germany

²Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, The Netherlands

³Scottish Universities Physics Alliance, Institute for Astronomy, University of Edinburgh, Blackford Hill, Edinburgh EH8 9YL, UK

⁴Institute for Computational Cosmology, Department of Physics, Durham University, South Road, Durham DH1 3LE, UK

⁵School of Mathematics, Statistics and Physics, Newcastle University, Herschel Building, Newcastle upon Tyne NE1 7RU, UK

⁶Astrophysics Research Institute, Liverpool John Moores University, 146 Brownlow Hill, Liverpool L3 5RF, UK

Accepted 2022 April 11. Received 2022 March 28; in original form 2021 October 19

ABSTRACT

Maximizing the information that can be extracted from weak lensing (WL) measurements is a key goal for upcoming stage IV surveys. This is typically achieved through statistics that are complementary to the cosmic shear two-point correlation function, the most well established of which is the WL peak abundance. In this work, we study the clustering of WL peaks, and present parameter constraint forecasts for an LSST-like survey. We use the *cosmo*-SLICS w CDM simulations to measure the peak two-point correlation function for a range of cosmological parameters, and use the simulation data to train a Gaussian process regression emulator that is applied to generate likelihood contours and provide parameter constraint forecasts from mock observations. We investigate the dependence of the peak two-point correlation function on the peak height, and find that the clustering of low-amplitude peaks is complementary to that of high-amplitude peaks. Consequently, their combination gives significantly tighter constraints than the clustering of high peaks alone. The peak two-point correlation function is significantly more sensitive to the cosmological parameters h and w_0 than the peak abundance, and when the probes are combined, constraints on Ω_m , S_8 , h , and w_0 improve by at least a factor of 2, relative to the peak abundance alone. Finally, we compare the forecasts for WL peaks and voids, and show that the two are also complementary; both probes can offer better constraints on S_8 and w_0 than the shear correlation function by roughly a factor of 2.

Key words: gravitational lensing: weak – methods: data analysis – cosmology: theory – large-scale structure of universe.

1 INTRODUCTION

The standard cosmological model, Λ CDM, consists of matter dominated by cold dark matter (CDM), and a late-time accelerated expansion that is driven by a positive cosmological constant Λ . This model is highly successful at describing a number of independent observations, which constrain the Λ CDM parameters with a large degree of concordance. Notably, measurements of fluctuations in the cosmic microwave background (CMB; Planck Collaboration VI 2018) have been used to constrain cosmological parameters including the present-day expansion rate of the Universe, H_0 , the matter density parameter, Ω_m , and the matter fluctuation amplitude σ_8 , defined as the root-mean-squared matter density perturbations smoothed on $8h^{-1}$ Mpc scales.

Gravitational lensing is another promising observational probe that can be used to constrain many cosmological parameters, where the light of distant source images is distorted by the gravitational potentials of the foreground matter. In the weak lensing (WL) regime, light is deflected by the large-scale structure (LSS) of the Universe, and the WL signal is measured through the correlations in distortions

of many source galaxies (Bacon, Refregier & Ellis 2000; Kaiser, Wilson & Luppino 2000; Van Waerbeke et al. 2000; Wittman et al. 2000). This allows us to probe the total matter distribution of the Universe on the largest scales (see Bartelmann & Schneider 2001; Kilbinger 2015, for reviews), and offers a powerful method to study the properties of dark matter and dark energy.

Recent WL observations that supplement the CMB parameter measurements include the Dark Energy Survey (DES; DES Collaboration et al. 2021),¹ Hyper Supreme-Cam (HSC; Hikage et al. 2019),² and the Kilo-Degree Survey (KiDS; Asgari et al. 2020)³ WL surveys. All of these surveys measure lower values of σ_8 compared to Planck, with a statistically significant disagreement arising between the Planck and KiDS constraints. However, the more recent results from DES are statistically compatible with Planck. This is one example where different observations point to slightly different values of certain cosmological parameters, which, in more extreme cases, may imply the presence of either unaccounted for systematics or new physics that is not modelled. An example of a larger discrepancy, leading to a parameter tension, is the H_0 tension, where multiple

¹<https://www.darkenergysurvey.org/>.

²<https://hsc.mtk.nao.ac.jp/ssp/>.

³<http://kids.strw.leidenuniv.nl/>.

* E-mail: christopherdavies1234@googlemail.com

observations find that measurements from the early Universe are broadly inconsistent with those of the late Universe (Verde, Treu & Riess 2019), particularly the distance scale measurement of H_0 based on Cepheids by the SHOES collaboration (Riess et al. 2019).

In order to address these parameter tensions, and more deeply probe the nature of the Universe, it is important to measure cosmological parameters as precisely as possible. This can be achieved by maximizing the information that can be extracted from a given survey. The standard approach for WL surveys is to measure Λ CDM parameters with two-point statistics, such as the shear–shear correlation function (Schneider et al. 2002; Hoekstra et al. 2006; Semboloni et al. 2006; Fu et al. 2008; Heymans et al. 2012; Kilbinger et al. 2013; Hildebrandt et al. 2017; Troxel et al. 2018; Aihara et al. 2019; Hikage et al. 2019; Asgari et al. 2020; DES Collaboration et al. 2021). However, the shear two-point correlation function (2PCF) does not capture non-Gaussian information, and, due to the non-linear evolution of the Universe, WL data are highly non-Gaussian. To fully exploit the data, many non-Gaussian statistics have been developed, which encapsulate information beyond two-point statistics. A well-established example is the abundance of WL peaks (local maxima in the convergence field), which has been shown to be complementary to the shear two-point function and helps break the Ω_m – σ_8 parameter degeneracy (Jain & Van Waerbeke 2000; Pen et al. 2003; Dietrich & Hartlap 2010). Peaks are also shown to outperform the standard methods for constraining the sum of neutrino mass (Li et al. 2019) and w_0 (Martinet et al. 2020), and can be used to provide constraints on modified gravity theories (Liu et al. 2016). When used in conjunction with the shear two-point correlation function, WL peaks have been used to provide the tightest constraints on S_8 from DES-Y1 WL data (Harnois-Déraps et al. 2020). WL peaks also offer utility for other non-Gaussian statistics, such as WL voids (Davies, Cautun & Li 2018; Davies et al. 2020b), where the peaks can be used as tracers to identify the voids. By including WL peaks as a complementary statistic, the measurement errors on cosmological parameters can be reduced, which will help inform the statistical significance of any parameter tensions between multiple observations.

When used to constrain cosmological parameters, peak analyses typically focus on the WL peak abundance, which is the number density of WL peaks as a function of their lensing amplitude. Studies have shown that the WL peaks with the highest amplitudes tend to correspond to large haloes along the line of sight (Hamana, Yoshida & Takada 2004; Liu & Haiman 2016; Wei et al. 2018). For this reason, WL peaks identified in surveys such as HSC can be used to search for galaxy clusters (e.g. Hamana, Shirasaki & Lin 2020). Furthermore, shear 2PCF measurements are typically combined with measurements of galaxy clustering (and galaxy–galaxy lensing) to further tighten the cosmological constraints (e.g. DES Collaboration et al. 2021). So, if WL peaks correspond to massive haloes, and hence massive galaxies, and the clustering of these galaxies is known to contain complementary information, then studying and exploiting the clustering of WL peaks is a natural next step in maximizing the utility of WL peaks.

Previously, Marian et al. (2013) have shown that the 2PCF of WL peaks with high lensing amplitudes does not contain much complementary information to the peak abundance alone. In Davies, Cautun & Li (2019a), we presented some simple scaling relations for the WL peak 2PCF, and found that the clustering of low-amplitude WL peaks also appears to be sensitive to the cosmological parameters Ω_m and σ_8 . In this work, with a more detailed analysis, we show that the clustering of low-amplitude peaks contains significant complementary information to the clustering of high peaks, and that

when the clustering of multiple peak height ranges are combined, the WL peak 2PCF offers similar constraining power to the peak abundance alone, where the two probes have different degeneracy directions. Therefore, we also show that, when their abundance and 2PCF are combined, the total cosmological information that is extracted from the WL peaks is significantly improved.

We use the numerical simulation suite, *cosmo*-SLICS (Harnois-Déraps, Giblin & Joachimi 2019) to measure the peak abundance and 2PCF, for a range of cosmological parameters. These data are then used to train a Gaussian process regression emulator, which, combined with Markov chain Monte Carlo (MCMC), allows us to generate likelihood contours and provide forecast parameter constraints for an LSST-like survey.

The layout of this paper is as follows. In Section 2, we outline the relevant theory for our WL peak analysis. In Section 3, we describe how we generate our mock observational data, and our emulation and likelihood analysis pipeline. In Section 4, we present the WL peak statistics used in the analysis, first from the mock data, and then from the emulator, in order to understand how these statistics depend on the cosmological parameters Ω_m , S_8 , h , and w_0 . In Section 5, we present the parameter constraint forecasts for the WL peak 2PCF and peak abundance. Finally, our conclusions are presented in Section 6. We also have two appendices where we present the covariance matrix used in our analysis, and study the accuracy of our emulator.

2 THEORY

The lens equation for a gravitationally lensed image, relating the deflection angle α to the true position of the source β and the observed position of the image θ , is

$$\alpha = \beta - \theta. \quad (1)$$

Neglecting second-order effects, the deflection angle is the gradient of a 2D lensing potential ψ ,

$$\alpha = \nabla\psi, \quad (2)$$

where ψ is given by

$$\psi(\theta, \chi) = \frac{2}{c^2} \int_0^\chi \frac{f_K(\chi - \chi')}{f_K(\chi)f_K(\chi')} \Phi(\chi'\theta, \theta) d\chi'. \quad (3)$$

Here, χ is the comoving distance from the observer to the source and χ' is the comoving distance from the observer to the continuously distributed lenses. $f_K(\chi)$ is the comoving angular distance, with the spatial curvature of the universe denoted by K . Note that for a flat universe with $K = 0$ (as used in this work) $f_K(\chi) = \chi$. Φ is the 3D lensing potential of the lens, and c is the speed of light. Φ is given by the Poisson equation

$$\nabla^2\Phi = 4\pi G a^2 \bar{\rho} \delta, \quad (4)$$

where ρ is the matter density of the Universe (with the mean denoted by a bar), $\delta \equiv \rho/\bar{\rho} - 1$, a is the scalefactor, and G is the gravitational constant.

The convergence κ and shear $\gamma = \gamma_1 + i\gamma_2$ can be related to the lensing potential via

$$\kappa \equiv \frac{1}{2} \nabla_\theta^2 \psi, \quad (5)$$

and

$$\gamma_1 \equiv \frac{1}{2} (\nabla_{\theta_1} \nabla_{\theta_1} - \nabla_{\theta_2} \nabla_{\theta_2}) \psi, \quad \gamma_2 \equiv \nabla_{\theta_1} \nabla_{\theta_2} \psi, \quad (6)$$

where $\nabla_{\theta} \equiv (\chi)^{-1} \nabla$. Equations (3), (4), and (5) allow us to express the convergence in terms of the matter overdensity

$$\kappa(\theta, \chi) = \frac{3H_0^2 \Omega_m}{2c^2} \int_0^\chi \frac{f_K(\chi - \chi')}{f_K(\chi)} f_K(\chi') \frac{\delta(\chi' \theta, \chi')}{a(\chi')} d\chi'. \quad (7)$$

The above derivation uses a fixed source plane. However, in real WL observations, the catalogue of source galaxies has a probability distribution $n(\chi)$ that spans over a range of χ values, and equation (7) is then weighted by this distribution to obtain $\kappa(\theta)$ (see, e.g. Kilbinger 2015, for details):

$$\kappa(\theta) = \int_0^\chi n(\chi') \kappa(\theta, \chi') d\chi'. \quad (8)$$

WL observations rely on accurately measuring the shapes of galaxies, and cross-correlating the shapes of neighbouring galaxies. However, measurements of the galaxy shapes used to extract the lensing signal are dominated by the random galaxy shapes and orientations, which is referred to as galaxy shape noise (GSN). Since the lensing signal is weak, when identifying WL peaks (local maxima in the convergence field $\kappa(\theta)$) it is convenient to express the convergence relative to the standard deviation of the corresponding GSN component of the field, σ_{GSN} . This is given by the following definition of signal-to-noise ratio,

$$v = \frac{\kappa}{\sigma_{\text{GSN}}}, \quad (9)$$

where σ_{GSN} is the standard deviation of the contributions to the signal from galaxy shape noise. The σ_{GSN} term can be calculated by generating mock GSN maps using the prescription from Jain & Van Waerbeke (2000) and Van Waerbeke et al. (2000) and applying to them any transformations also applied to the convergence maps, such as smoothing. Mock GSN maps are generated by assigning to pixels random convergence values from a Gaussian distribution with standard deviation

$$\sigma_{\text{pix}}^2 = \frac{\sigma_{\text{int}}^2}{2\theta_{\text{pix}}^2 n_{\text{gal}}}, \quad (10)$$

where θ_{pix} is the width of each pixel, σ_{int} is the intrinsic ellipticity dispersion of the source galaxies, and n_{gal} is the measured source galaxy number density. In this work we use $\sigma_{\text{int}} = 0.28$ and $n_{\text{gal}} = 20 \text{ arcmin}^{-2}$ as will be discussed in Section 3.1.

As mentioned in the introduction, WL peaks are closely related to the dark matter haloes along the line of sight. In cosmology, both the abundance and large-scale clustering of haloes encode useful information about the underlying cosmological model and parameter values. Therefore, as well as studying the abundance of WL peaks, we will also study their clustering. The extent to which objects are clustered can be measured through the two-point correlation function (2PCF) that is defined as the excess probability, relative to a random distribution, of finding a pair of objects at a given separation θ . Formally, this is written as

$$dP_{ij}(\theta) = \bar{n}^2 (1 + \xi(\theta)) dA_i dA_j, \quad (11)$$

where \bar{n} is the expected tracer number density, dA_i and dA_j are two sky area elements that are separated by a displacement θ with amplitude θ , and $\xi(\theta)$ is the 2PCF. We have $\xi(\theta) = \xi(\theta)$, thanks to statistical isotropy. In practice, the 2PCF can be measured through the Landy–Szalay estimator (Landy & Szalay 1993) that requires the generation of matching catalogues containing randomly distributed points and is given by

$$\xi_{\text{LS}}(\theta) = 1 + \left(\frac{N_{\text{R}}}{N_{\text{D}}} \right)^2 \frac{DD(\theta)}{RR(\theta)} - \left(\frac{N_{\text{R}}}{N_{\text{D}}} \right) \frac{DR(\theta)}{RR(\theta)}. \quad (12)$$




Figure 1. The four-dimensional parameter space ($[\Omega_m, S_8, h, w_0]$) sampled by the *cosmo-SLICS* simulation suite. The fiducial cosmology is indicated by a star with parameter values $[0.29, 0.82, 0.69, -1.00]$.

In the above, N_{D} and N_{R} are the numbers of data and random points, and DD , DR , and RR are the numbers of data–data, data–random, and random–random pairs in bins $\theta \pm \delta\theta$, respectively. See Davies et al. (2019a) for more details about the measurement of the peak 2PCF, which are important for small lensing maps.

3 METHODOLOGY

In this section, we describe the methodology followed in this work, including the simulations, mock lensing data, emulation, and likelihood analysis.

3.1 Mock data

In this work, we use the SLICS and *cosmo-SLICS* (Harnois-Déraps & van Waerbeke 2015; Harnois-Déraps et al. 2018, 2019) mock WL convergence maps, which we briefly outline in this subsection.

The *cosmo-SLICS* are a suite of high-resolution N -body simulations that were run for 26 sets of $[\Omega_m, S_8, h, w_0]$ cosmological parameters. Here $S_8 \equiv \sigma_8(\Omega_m/0.3)^{0.5}$, $h = H_0/100 \text{ km s}^{-1} \text{ Mpc}^{-1}$ is the reduced Hubble constant, and w_0 the dark energy equation-of-state parameter, which is assumed to be a constant that is allowed to deviate from -1 (cosmological constant).

The four-dimensional parameter space is sampled using a Latin hypercube, which samples the parameter space comprehensively with a low node count. The exact cosmological parameter space that is probed by the *cosmo-SLICS* is shown in Fig. 1. Each simulation volume is a cube with length $L = 505 h^{-1} \text{ Mpc}$, with $N = 1536^3$ dark matter particles. To reduce the impact of cosmic variance, two simulations are run for each cosmology, starting from different (paired) initial conditions. For each set of cosmological parameters, 50 pseudo-independent light-cones are constructed by resampling projected mass sheets, which are then ray-traced under the Born approximation to construct lensing maps and catalogues

(see Harnois-Déraps et al. 2019, for full details about the light-cone and catalogue construction).

We use the *Generic cosmo-SLICS* source catalogue selected over the range $z_s = [0.6, 1.4]$ to match LSST specifications, which gives a conservative source galaxy number density of 20 arcmin^{-2} . From this, we generate 50 WL convergence maps for each of the nodes, with a sky coverage of $10 \times 10 \text{ deg}^2$ each and 3600^2 pixels, following the method described in Giblin et al. (2018). These maps are then smoothed with a Gaussian filter with smoothing scale $\theta_s = 1 \text{ arcmin}$. Davies et al. (2019a) contains a study of the impact of different smoothing scales on the WL peak 2PCF.

For estimates of the covariance matrices, we use the SLICS suite to produce 615 WL convergence maps at the fiducial cosmology, which match the properties of the *cosmo-SLICS* maps. These are obtained from fully independent N -body realizations carried out at the same cosmology,⁴ but with different seeds in their initial conditions, allowing to accurately capture the sample variance. The larger number of SLICS realizations relative to *cosmo-SLICS* allows us to calculate robust covariance matrices and to use large data vectors in the likelihood analysis below when combining probes.

3.2 Emulation and likelihood analysis

In this subsection, we outline the procedure used to test the sensitivity of WL peak statistics to the cosmological parameters Ω_m , S_8 , h , and w_0 .

First, we measure the WL peak statistics from the 50 convergence maps for each of the nodes shown in Fig. 1. Then, in order to predict the WL peak statistics at arbitrary points in this parameter space, we use the Gaussian process (GP) regression emulator from SCIKIT-LEARN (Pedregosa et al. 2011) to interpolate the peak statistics between nodes. GP regression is a non-parametric Bayesian machine learning algorithm used to make probabilistic predictions that are consistent with the training data (see, e.g. Habib et al. 2007; Schneider et al. 2008, for some of its early applications in cosmology). The accuracy of the GP emulator trained on *cosmo-SLICS* has been tested extensively and shown to reach a few per cent level in predictions of WL shear two-point correlation functions (Harnois-Déraps et al. 2019), density split statistics (Burger et al. 2020), persistent homology statistics (Heydenreich, Brück & Harnois-Déraps 2020), aperture mass statistics (Martinet et al. 2020) and WL void statistics (Davies et al. 2020b). In this work, the average peak statistics and their standard errors at each node are used as the training data. We present results of the accuracy of the emulator for the peak statistics in Appendix A.

Finally, once the emulator has been trained and tested, we use MCMC to estimate the posteriors of the parameters for the entire parameter space and produce likelihood contours. We use the EMCEE PYTHON package (Foreman-Mackey et al. 2013) to conduct the MCMC analysis in this work sampling the 4D parameter space as follows. We employ a Bayesian formalism, in which the likelihood of the set of cosmological parameters $\mathbf{p} = [\Omega_m, S_8, h, w_0]$, given a data set \mathbf{d} , is given by

$$P(\mathbf{p}|\mathbf{d}) = \frac{P(\mathbf{p})P(\mathbf{d}|\mathbf{p})}{P(\mathbf{d})}, \quad (13)$$

where $P(\mathbf{p})$ is the prior, $P(\mathbf{d}|\mathbf{p})$ is the likelihood of the data conditional on the parameters, and $P(\mathbf{d})$ is the normalization. In this

⁴The SLICS cosmology has the following parameter values: $[\Omega_m, \sigma_8, h, w_0, n_s, \Omega_b] = [0.2905, 0.826, 0.6898, -1.0, 0.969, 0.0474]$.

work we use flat priors with the following upper and lower limits, respectively, for Ω_m : $[0.10, 0.55]$, S_8 : $[0.61, 0.89]$, h : $[0.60, 0.81]$, w_0 : $[-1.99, -0.52]$, which matches the parameter space sampled by *cosmo-SLICS*. The log likelihood can be expressed as

$$\log(P(\mathbf{d}|\mathbf{p})) = -\frac{1}{2} [\mathbf{d} - \mu(\mathbf{p})] C^{-1} [\mathbf{d} - \mu(\mathbf{p})], \quad (14)$$

where $\mu(\mathbf{p})$ is the prediction generated by the emulator for a set of parameters \mathbf{p} , and C^{-1} is the inverse of the covariance matrix. We use the emulator's prediction of a statistic at the fiducial cosmology as the data \mathbf{d} . This choice is for simplicity and presentation purposes, which ensures that the confidence intervals are always centred on the true values of the cosmological parameters allowing for easier comparisons between multiple probes.

The likelihood returns a 4D probability distribution that indicates how well different regions of the parameter space match the input data \mathbf{d} . Note that equation (14) assumes that the covariance matrix does not depend on the cosmological parameters.

We use the 615 SLICS WL map realizations (which match the fiducial cosmology) to calculate the covariance matrices, and then divide it by a factor of 180 to rescale the covariance matrix from a 100 deg^2 area to the LSST survey area, which we take as $18\,000 \text{ deg}^2$. The joint covariance matrix for the peak probes studied in this work is presented in Appendix B. We also multiply the inverse covariance matrix by a debiasing factor α , which accounts for the bias introduced when inverting a noisy covariance matrix (Anderson 2003; Hartlap, Simon & Schneider 2007), and is given by

$$\alpha = \frac{N - N_{\text{bin}} - 2}{N - 1}. \quad (15)$$

Here $N = 615$ is the number of WL maps used to calculate the covariance matrix and N_{bin} is the number of bins used to measure the statistic.

4 WL PEAK STATISTICS

In this section, we present the WL peak statistics studied in this work, which include the peak abundance and the peak 2PCF. For each statistic, we first show their measurements from the *cosmo-SLICS* nodes that are used as the training data for the emulator. We then present emulations of the statistic by varying one cosmological parameter at a time, to exemplify its sensitivity to different cosmological parameters, which will aid the interpretation of the forecast cosmological constraints in Section 5.

4.1 WL peak abundance

Fig. 2 shows the differential WL peak abundance (number density) measured in each of the 26 nodes in Fig. 1. The abundance of the fiducial cosmology is shown by the blue curve, with the rest of the cosmologies plotted in grey.

First, the figure shows that there are an appreciable number of peaks with amplitudes below $\nu = 0$, which correspond to local maxima in regions that are underdense. In this analysis we do not use peaks with $\nu < 0$ for our forecast constraints for the following two reasons. First, Martinet et al. (2018) have shown that peaks with $\nu < 0$ correlate very strongly with peaks with $\nu > 0$, and so there is little gain in parameter constraints when these peaks are included. Second, due to their low amplitude, these peaks are also more likely to be affected by GSN.

The maxima of the peak abundance occur just above $\nu = 1$ for all cosmologies. Due to the low signal-to-noise ratio, this indicates that




Figure 2. (Colour online) The differential WL peak number density (abundance) as a function of peak height ν . The grey curves correspond to the 26 *cosmo*-SLICS nodes in Fig. 1, with the fiducial cosmology plotted as a blue thick curve. The orange shaded regions shows the 1σ standard error measured from the 50 fiducial *cosmo*SLICS realizations, multiplied by a factor of 10 to increase visibility.

a large fraction of the total number of peaks correspond to spurious local maxima induced by galaxy shape noise, rather than a physical signal induced by matter overdensities along the line of sight, while more peaks at the high- ν end are produced by a true physical signal.

As the peak abundance approaches higher ν values, the spread in the peak abundances between different cosmologies increases significantly, where the abundance of peaks at $\nu = 6$ can differ by an order of magnitude between the most extreme cosmologies. This is due to two factors. First, because the peaks at this amplitude are not dominated by noise, differences in the physical signal are more visible. Secondly, because the high-mass end of the halo mass function varies more significantly as a function of cosmological parameters, so does the peak abundance, since the largest peaks are created by the largest haloes (Liu & Haiman 2016; Wei et al. 2018).

Whilst the high- ν end of the abundance exhibits the greatest variation amongst the different cosmological parameters, this region also has the highest sample variance, since high peaks are orders of magnitude less abundant than low peaks. Therefore, as ν increases, the increased spread between cosmologies is in direct competition

with the increased statistical uncertainty. For this reason, it is important to consider the abundance of peaks over a wide ν range in our forecasts. We also note that large ν peaks are more affected by uncertainties such as intrinsic alignments and baryonic physics, as shown in Harnois-Déraps et al. (2020).

Next, in order to aid the physical interpretation of how the WL peak abundances (Fig. 2) depend on the four parameters, Ω_m , S_8 , h , and w_0 , we use the *cosmo*-SLICS data to train a GP emulator as discussed in Section 3.2 and present the emulated peak abundances in the *cosmo*-SLICS parameter space. In Fig. 3, we present these emulator predictions while varying one parameter at a time.

The emulated peak abundances plotted in Fig. 3 sample the signal around the test cosmology with parameters $[\Omega_m, S_8, h, w_0] = [0.3, 0.8, 0.7, -1.0]$. Each subpanel contains curves where one parameter is varied above and below these values. The bottom row of subpanels shows the ratio of the curves relative to the test cosmology. 1σ standard errors measured from the 50 *cosmo*-SLICS realizations are included on the test cosmology, shown by the shaded blue region. Finally, the ν range plotted here is slightly narrower than that presented in Fig. 2, since we are now showing the ν range that will be used to forecast the peak abundance constraints, which is $\nu \in [0, 6]$. We do not use peaks with $\nu > 6$ for two reasons. First, due to a low number density, the sample variance in this regime is very high. Secondly, as previously mentioned, this regime is significantly affected by uncertainties such as intrinsic alignments and baryonic physics. We have also performed tests using a higher upper limit of $\nu < 8$, and find that it has nearly no impact on our results. Therefore, using $\nu < 6$ allows us to stay within the regime that is less affected by the afore-mentioned uncertainties, whilst still encapsulating maximal cosmological information.

Panel (a) shows how the peak abundance depends on Ω_m . Note that since $S_8 = \sigma_8(\Omega_m/0.3)^{0.5}$, σ_8 increases when Ω_m is reduced (and vice versa), in order for S_8 to remain constant. Increasing Ω_m (with S_8 and the other parameters held constant) reduces the abundance of WL peaks with amplitudes $\nu < 2$ but increases the abundance of peaks with amplitudes $\nu > 2$, relative to the fiducial case. This is because a higher Ω_m increases the matter content of the universe, which allows dark matter haloes to grow more massive, increasing their lensing signal and the resulting peak amplitudes. The opposite behaviour can be seen when Ω_m is reduced relative to the fiducial case, with more peaks below $\nu = 2$ and fewer above. There is an

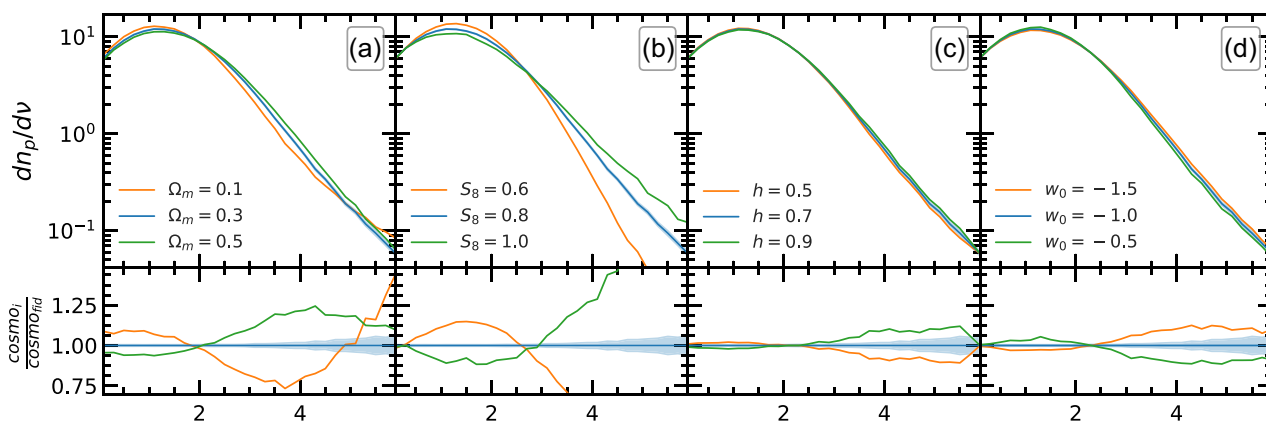


Figure 3. (Colour online) Top row: the emulated peak abundance. The curves correspond to the cosmological parameters $[\Omega_m, S_8, h, w_0]$ with values $[0.3, 0.8, 0.7, -1.0]$, unless otherwise stated in the subpanel legends. Each subpanel corresponds to varying one cosmological parameter at a time, denoted in the legend. Bottom row: the curves from the top row, divided by the fiducial cosmology (blue curve in the top row). The 1σ standard errors measured from the 50 *cosmo*-SLICS realizations are included on the fiducial (blue) curves, and they are barely visible in the upper subpanels.

upturn in the peak abundance for $\Omega_m = 0.1$ at $\nu \approx 5$, which is due to the fact that S_8 is held constant, rather than σ_8 .

Panel (b) shows the peak abundance for different S_8 values. The results presented in this subpanel vary more strongly compared to all other subpanels, verifying that the peak abundance is the most sensitive to S_8 of all the parameters studied here. Similar to the behaviour seen for Ω_m , increasing S_8 reduces the number of small peaks below $\nu \approx 2.7$, but increases the number of large peaks above this point. The opposite behaviour is seen for decreasing S_8 . Increasing S_8 leads to greater clustering of matter, which will place more haloes closer together. The increased mass along an overdense line of sight translates into a greater lensing signal, which produces more peaks of higher amplitudes. This also reduces the number of small peaks since fewer haloes are in isolation that would produce small peaks.

Panel (c) shows how the peak abundance changes with h . Peaks with amplitudes $\nu < 3$ are mostly unaffected, however there is a small amount of sensitivity to h at the high ν end, where increasing h slightly increases the number of high peaks and vice versa. This result is not entirely surprising: From equation (7), we can see that the dependences of h , or equivalently H_0 , cancel out, because the comoving distance can be written as

$$\chi(z) = \frac{c}{H_0} \int \frac{1}{E(z')} dz', \quad (16)$$

where $E_{\Lambda\text{CDM}}(z)$ is defined as

$$E_{\Lambda\text{CDM}}(z) \equiv \frac{H_{\Lambda\text{CDM}}(z)}{H_0} = \Omega_m(1+z)^3 + 1 - \Omega_m, \quad (17)$$

for a flat ΛCDM cosmology (as is the case of our fiducial cosmology), and is independent of h . This means that the H_0 factors in the pre-factor, and the χ and $d\chi$ terms of equation (7) cancel out, so that the only dependence on H_0 in κ would come through the matter density contrast δ . In the linear-perturbation regime, the evolution of δ can be expressed in the linear growth factor D_+ , which, for a flat ΛCDM cosmology, is given by the following solution:

$$D_+(z) = E_{\Lambda\text{CDM}}(z) \left[\int_0^\infty \frac{(1+z')}{E_{\Lambda\text{CDM}}^3(z')} dz' \right]^{-1} \int_z^\infty \frac{(1+z')}{E_{\Lambda\text{CDM}}^3(z')} dz', \quad (18)$$

where the term in the bracket offers the normalization to ensure that $D_+(z=0) = 1$ as per the usual convention – this suggests that for flat ΛCDM models with the same Ω_m , σ_8 and D_+ , κ is independent of h . However, we remark that the above argument only applies to a universe with strictly no radiation. In practice, increasing h with Ω_m fixed would mean that the physical matter density today, $\Omega_m h^2$, increases, which brings the matter-radiation equality to higher redshift. Since the growth of matter perturbations is slower during radiation domination but faster during matter domination, this means that *small-scale* matter perturbations experience a stronger growth in the case of a larger h , and thus it requires a *lower* value of A_s (the amplitude of the primordial power spectrum) in order to reach the desired σ_8 . Consequently, the matter clustering on *large scales* – e.g. at k smaller than $\simeq 0.01 h\text{Mpc}^{-1}$, which corresponds to the horizon scale at matter-radiation equality – will indeed be *weaker*. Actually, a more detailed calculation shows that, when comparing the cases of $h = 0.9$ and 0.7 (with Ω_m and σ_8 fixed), the late-time matter power spectrum $P(k)$ is higher (lower) in the latter than in the former for $k \gtrsim 0.1 h\text{Mpc}^{-1}$ ($k \lesssim 0.1 h\text{Mpc}^{-1}$). This will have non-trivial implications for the peak 2PCFs as we shall see shortly. Nevertheless, for the peak abundance, the most relevant scales are

$k \simeq 0.1\text{--}1 h\text{Mpc}^{-1}$, where the cases of $h = 0.7$ and 0.9 have similar matter clustering amplitudes, which is slightly higher for larger h : \bar{A} s this k range corresponds to the sizes of halo-forming regions, this is consistent with the high- ν behaviour of the middle right-hand panel.

Panel (d) shows the peak abundances with varying w_0 . Similar to h , the peak abundance does not appear to be very sensitive to changes in w_0 , but increasing w_0 does indeed create slightly more low- ν peaks and fewer high- ν peaks compared to the fiducial case, and vice versa. A different dark energy equation of state can change the expansion rate, therefore affecting the comoving distances, the lensing kernel in equation (7), and the growth rate of matter perturbation δ . The physics underlying the qualitative behaviours shown in these panels is actually complicated and quite interesting. Usually, a more negative w_0 , e.g. $w_0 < -1.0$, implies an increase of the dark energy density with time and therefore (for the same matter density) a faster transition from the phase of decelerated expansion to an accelerated one dominated by dark energy, compared to standard ΛCDM . But given that we fix h and therefore H_0 , at $z > 0$ the expansion rate is actually slower than the fiducial ΛCDM model, because the density of dark energy in this case decreases with redshift, and so at $z > 0$ the total density of matter and dark energy is smaller than in ΛCDM . More explicitly, we have $E_{w_0}(z) \leq E_{\Lambda\text{CDM}}(z)$ for $w_0 < -1$, where

$$E_{w_0}(z) = \frac{H_{w_0}(z)}{H_0} = \Omega_m(1+z)^3 + [1 - \Omega_m](1+z)^{3(1+w_0)}, \quad (19)$$

which reduces to equation (17) when $w_0 = -1$. Because the dark energy in our simulations is assumed to be non-clustering, the only effect of varying w_0 is to modify the background expansion history, which leads to a *scale-independent* change in the linear matter clustering, $P(k)$. It may seem that, since $w_0 = -1.5$ leads to a slower expansion, it will increase matter clustering. While this is true, we have to note that in this comparison we have fixed S_8 (and equivalently σ_8) today in all three cases, and so this means that, in order to have the same σ_8 at the present day, the primordial power spectrum must be lower in the $w_0 = -1.5$ case (we have explicitly checked this using CAMB). At an initial thought, this seems to suggest that this model predicts less structure formation than ΛCDM (until $z = 0$), which is against the results of Fig. 3. However, recall that another effect of having a slower expansion is that the Universe becomes older at $z = 0$, and distances to the same redshift become larger; the latter, in particular, means that in between the observer and the source(s) there would be more volume, and more structures such as large dark matter haloes. Since these haloes produce the high- ν peaks, the net effect can be a larger abundance of such peaks. Fig. 3 indeed confirms that the two competing effects – the decrease of LSS due to the lowered primordial power spectrum and the larger volume between the observer and a fixed source redshift – give rise to a higher peak number count at $\nu \gtrsim 2.3$. For low peaks, the total effect is less clear-cut, since a peak is defined as a pixel in the convergence map with a higher κ value than all its neighbouring pixels, and both the central and the neighbouring pixels could be affected by chance alignments of small dark matter structures; Fig. 3 shows that at $\nu \lesssim 2.3$ the peak abundance is smaller in $w_0 = -1.5$ than in ΛCDM . The behaviour of the $w_0 = -0.5$ model can be similarly explained.

Figs 2 and 3 show that the peak abundance is mostly sensitive to changes in Ω_m and S_8 , and less so to h and w_0 , with sensitivity to cosmology coming from both the high- and low-amplitude peaks.




Figure 4. (Colour online) The WL peak 2PCF, where each subpanel corresponds to the 2PCF of a different peak catalogue. The various peak catalogues (and hence their 2PCF) only contain peaks with amplitudes $\nu > 1$ (top left-hand panel), $\nu > 2$ (top right-hand panel), $\nu > 3$ (bottom left-hand panel), and $\nu > 4$ (bottom right-hand panel). The curves in each subpanel correspond to the 26 *cosmo*-SLICS nodes in Fig. 1, with the fiducial cosmology plotted in blue. The shaded orange regions show the 1σ standard errors measure from the 50 *cosmo*SLICS realizations. Note the change of y-axis range between the upper and lower subpanels.

4.2 WL peak two-point correlation function

In order to measure the WL peak 2PCF, we first remove peaks with amplitudes below a given ν threshold, and then repeat this step with different ν thresholds in order to create multiple peak catalogues. This procedure is motivated by the following factors.

First, as discussed in Section 4.1, a significant fraction of the WL peak population is noise-dominated (low ν), suggesting that their spatial distribution (to which the 2PCF is sensitive) may not contain useful cosmological information. We have tested this assertion, and found that the WL peak 2PCF measured using peaks of all heights has a very low amplitude and exhibits only small variations between the 26 *cosmo*-SLICS nodes. This indicates that when no distinction is made based on peak heights, the average clustering of WL peaks is close to that of a randomly distributed sample. Therefore, in order to extract useful information on the clustering of WL peaks, we must first remove the low-amplitude noise-dominated peaks. We have tested a range of ν thresholds, and use two criteria to determine the best threshold. First, the amplitude of the 2PCF measurement must be sufficiently high so that the clustering signal is not noise dominated. Secondly, the variance between 2PCFs for different cosmological parameters must be larger than the variance with which the fiducial 2PCF is measured. This ensures that any cosmological information will not be lost to noisy measurements. We found that both these criteria can be met with a threshold as low as $\nu = 1$.

Secondly, varying the ν threshold and using multiple WL peak catalogues produces multiple WL peak 2PCF measurements. The change in the 2PCF as the ν threshold changes is also sensitive to the underlying cosmology, so we expect that the different 2PCF measurements will contain complementary information to each other, so that combining these measurements will yield tighter cosmological parameter constraints.

Fig. 4 shows the 2PCFs for four WL peak catalogues with $\nu > 1$ (top left-hand panel), $\nu > 2$ (top right-hand panel), $\nu > 3$ (bottom

left-hand panel), and $\nu > 4$ (bottom right-hand panel). The 2PCFs for the 26 *cosmo*-SLICS nodes are plotted, with the fiducial cosmology plotted in blue, and all other cosmologies plotted in grey.

The 2PCF measurements for the $\nu > 1$ catalogue have the lowest amplitude. As the ν threshold increases, so does the amplitude of the 2PCF for all cosmologies, indicating that the high ν peaks are more clustered than the low ν peaks. Both the gradient and the amplitude of the 2PCF change as the ν threshold increases, however the changes in amplitude appear to be the most dominant feature. This can be explained by the relationship between WL peaks and dark matter haloes – more massive haloes are known to be more strongly biased and clustered, because they form from higher density peaks of the primordial density field.

Similar to Section 4.1, we use the *cosmo*-SLICS data from Fig. 4 to train a GP emulator as discussed in Section 3.2, and present emulated peak 2PCFs in the *cosmo*-SLICS parameter space by varying one parameter at a time. The results are plotted in Fig. 5. The bottom row in each section shows the ratio relative to the test cosmology. The 1σ standard errors measured from the 50 *cosmo*-SLICS realizations are included for the fiducial cosmology and are shown by the shaded blue region. The top and bottom sections of Fig. 5 shows results for the 2PCF of peaks with $\nu > 2$ and $\nu > 4$, respectively. We choose to show results for $\nu > 2$ rather than $\nu > 1$ since, as we will see in Fig. 6, $\nu > 2$ gives stronger parameter constraints than the $\nu > 1$ case.

Panel (a) shows the emulated 2PCF for $\nu > 2$ varying only Ω_m . Increasing Ω_m has less of an impact on the 2PCF for small θ compared to large θ , effectively steepening the curve relative to the fiducial case by a small amount. When decreasing Ω_m , the above behaviour is mirrored, except the overall magnitude of the change is larger compared to the case where Ω_m is increased. This shows that Ω_m dictates the gradient of the 2PCF, and appears to be more sensitive to low Ω_m values. It might seem counter-intuitive that a




Figure 5. (Colour online) The emulated peak 2PCF for two different peak catalogues, $\nu > 2$ (top section) and $\nu > 4$ (bottom section). The curves correspond to the cosmological parameters $[\Omega_m, S_8, h, w_0]$ with values $[0.3, 0.8, 0.7, -1.0]$, unless otherwise stated in the subpanel legends. Each subpanel corresponds to varying one cosmological parameter at a time, as specified in the legend. The bottom rows in each section show the ratio of the curves relative to the fiducial cosmology. The 1σ standard errors measured from the 50 *cosmo*-SLICS realizations are shown for the fiducial cosmology by the shaded blue region.

model with smaller Ω_m would predict a stronger clustering for WL peaks, but we note again that here S_8 has been fixed when Ω_m is being varied, so that a smaller Ω_m corresponds to a larger σ_8 , and the latter means there is more matter clustering.

Panel (e) shows the peak 2PCFs for $\nu > 4$ for the same Ω_m values. Similar to the $\nu > 2$ case, increasing Ω_m decreases the 2PCF amplitude and vice versa, and the behaviour relative to the fiducial case is asymmetric, where changing Ω_m by a fixed amount in either direction has a larger impact on the amplitude when Ω_m is decreased, suggesting that the $\nu > 4$ catalogue is also more sensitive to small Ω_m . However, compared to the $\nu > 2$ case, there appears to be slightly less change to the overall slope of the 2PCF as Ω_m is varied.

Panel (b) is the same as the previous panels except S_8 is varied in this case. The figure shows that changes to S_8 affect the amplitude of the 2PCF, where lowering S_8 lowers the 2PCF amplitude since it corresponds to a smaller σ_8 (remember that Ω_m is held constant here) and therefore less clustering of matter. Increasing S_8 by the same amount increases the amplitude relative to the fiducial case, but the magnitude of the change is slightly smaller compared to the decreased S_8 case. Panel (f) shows the $\nu > 4$ 2PCF for the same three S_8 values. The overall trend here is the opposite to the $\nu > 2$ case. Initially, it seems counter-intuitive that higher S_8 values would lead to a lower clustering amplitude; however, as shown by Fig. 3, the abundance of peaks is also larger for this catalogue. Therefore, when S_8 increases, the number of peaks with $\nu > 4$ increases, meaning that smaller maxima in the primordial density field – which are less

biased and hence less clustered tracers of the matter density field – end up contributing to this peak catalogue, and so the clustering of the peaks decreases and vice versa.

Panel (c) shows how the 2PCF for the $\nu > 2$ catalogue depends on h . The 2PCF appears to be sensitive to changes in h , where increasing h decreases its amplitude, and vice versa. This observation is actually consistent with the discussion above about the physical impact of varying h – with Ω_m and σ_8 fixed – on matter clustering: increasing h from 0.7 to 0.9 weakens the late-time matter clustering at $k \lesssim 0.1 h \text{Mpc}^{-1}$, and these are the scales most relevant for the peak clustering (which is expected to trace the dark matter clustering) as well. Unlike the behaviour seen for Ω_m and S_8 , changing h by a fixed amount in either direction appears to change the 2PCF amplitude by an equal amount. Panel (g) shows the $\nu > 4$ 2PCF for the same three h values. As for the case of $\nu > 2$, the 2PCF amplitude increases when h decreases and vice versa. Indeed, the effects of varying h are similar for both the $\nu > 2$ and the $\nu > 4$ catalogues, except the $\nu > 2$ case appears to be more sensitive to h at large θ .

Panel (d) shows the $\nu > 2$ 2PCFs for different values of w_0 . Increasing w_0 decreases the amplitude and vice versa, with no apparent changes to the gradient. This behaviour also appears to be symmetric relative to the fiducial cosmology, similar to that seen for h , and unlike Ω_m and S_8 . Panel (h) is the same but shows the $\nu > 4$ 2PCF, which appears to have little sensitivity to changes in w_0 . In both catalogues, we think that the physical reason underlying

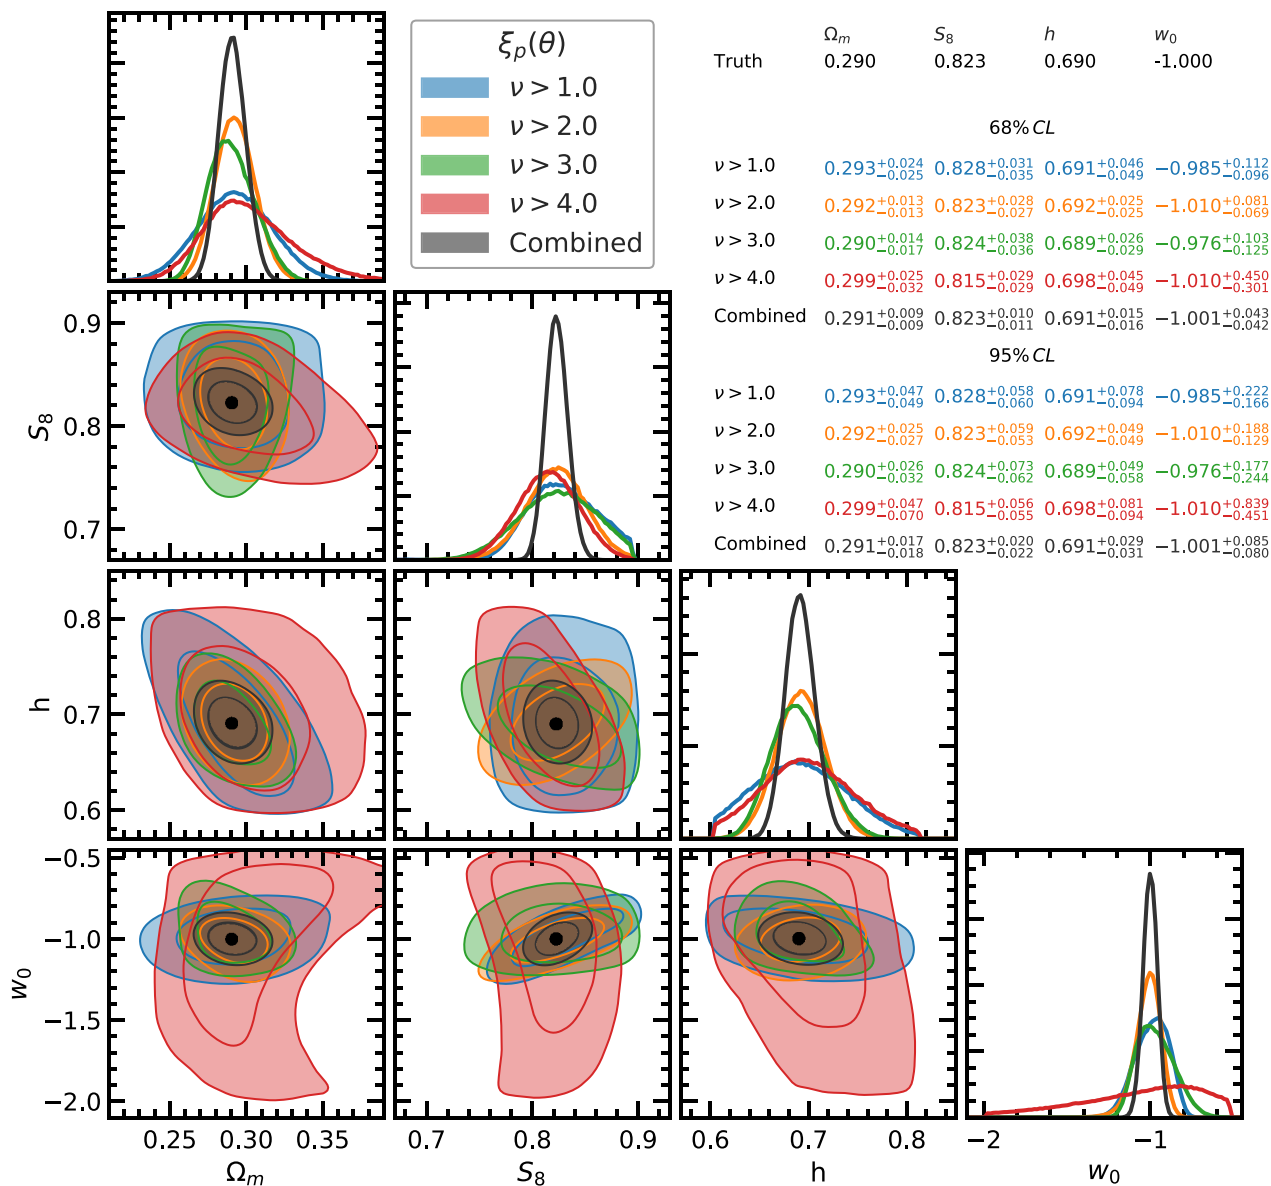


Figure 6. (Colour online) Constraint forecasts on cosmological parameters measured from the WL peak 2PCF. Contours are shown for 2PCFs measured from WL peak catalogues with $\nu > 1$ (blue), $\nu > 2$ (orange), $\nu > 3$ (green), and $\nu > 4$ (red) and the combination of all four catalogues (black). The true cosmological parameter values used to generate the data are indicated by the black point. The diagonal panels show the 1D marginalized probability distribution, and the remaining panels show the marginalized 2D probability contours enclosing the 68 and 95 per cent confidence intervals. The table in the top right shows true parameter values (top section) and the inferred parameter values for the different peak catalogues with 68 (upper section) and 95 per cent (lower section) confidence limits.

the w_0 dependence is the same as in the case of peak abundance. Although, naively, it seems that the case of $w_0 = -1.5$ – which has faster growth of structures at late times and hence requires a lower initial power spectrum amplitude to achieve the same σ_8 or S_8 at $z = 0$ – should predict less matter clustering during the entire lensing kernel and so lead to a lower amplitude of the peak 2PCF, we note that this model also covers a larger volume for the same redshift range and therefore receives contribution from a greater number of massive haloes. Also, the peak 2PCF is a projection effect, and the projection depth is larger for the case of $w_0 = -1.5$, which leads to a larger line-of-sight integration. These different effects compete with each other and can have cancellations, which may explain why

for the $\nu > 4$ catalogue there is almost no dependence on w_0 (also note that the same $\nu > 4$ peak height threshold can lead to different peak populations for the different models, which could also have an impact on the peak correlation).

Comparing the bottom row ($\nu > 4$) to the top row ($\nu > 2$), we see that the amplitudes of the 2PCFs are all higher. Given that fewer tracers are used for the $\nu > 4$ measurements, the errors on these curves should be larger. This will be in direct competition with any increased sensitivity to the cosmological parameters relative to the $\nu > 2$ case. None the less, separate catalogues still contain complementary information to each other regardless of which factors wins out, as we will show later.

5 PARAMETER CONSTRAINTS FORECAST

In this section, we present parameter constraint forecasts for the statistics studied in Section 4.

Fig. 6 shows the parameter constraint forecasts for an LSST-like survey for the WL peak 2PCF. We present constraints for the four WL peak 2PCFs measured from WL peak catalogues with four heights $\nu > 1, 2, 3,$ and $4,$ and the combination of all four 2PCFs. The true cosmological parameter values used to generate the data are indicated by the black point. The diagonal panels show the 1D marginalized probability distributions, while the remaining panels show the marginalized 2D probability contours enclosing the 68 and 95 per cent confidence intervals. All confidence intervals, along with the true parameter values, are explicitly stated in the table in the top right of the figure.

In general, as the ν threshold increases, the contour sizes start off large ($\nu > 1$), begin to shrink ($\nu > 2$ and $\nu > 3$), and become large again ($\nu > 4$). The shape and orientation of the contours also change significantly as the ν threshold increases. For example, in the Ω_m - S_8 plane, the $\nu > 3$ contour is smaller than the $\nu > 4$ contour; however, the two are orthogonal to each other. This behaviour shows that the constraining power of the WL peak 2PCF can be significantly improved when the 2PCFs of multiple peak catalogues are combined. Even in the case of very large contours that fully enclose the contours from lower ν thresholds, the presence of complementary information between the different 2PCFs is not ruled out. This is because it depends not only on the size, shape and orientation of the contours, but also on the correlation between the contours. This is discussed in detail in Appendix B. The benefit to combining multiple peak catalogues is shown by the grey contours, which are significantly smaller than any individual contour in all cases.

We find that the $\nu > 2$ and $\nu > 3$ peak 2PCFs give the tightest constraints on Ω_m , the $\nu > 2$ and 4 2PCFs both give the similar constraint on S_8 , $\nu > 2$ and 3 are tightest on h and $\nu > 2$ gives the best constraint on w_0 . It is interesting to note that the constraints on w_0 are roughly nine times smaller for the combination of all catalogues compared to $\nu > 4$ alone, indicating that a significant amount of cosmological information is contained in the clustering of low-amplitude peaks.

Fig. 7 shows parameter constraint forecasts for the combination of the peak 2PCFs similar to the black contours of Fig. 6, but now using a finer selection from eight peak catalogues with $\nu > 1.0, 1.5, \dots, 4.5$ in orange. The constraints from the peak abundance for peaks with heights $0 < \nu < 6$ are shown in blue, and the combination of the abundance and 2PCFs are shown in green. We note that when multiple probes are combined, it is important to account for any duplicate information between the probes through the covariance matrix of the data vector, including the cross-correlation between the multiple probes. The covariance matrix of all probes studied in this work is presented in Appendix B and discussed in detail therein. The orange contour shows how the parameter constraints are improved when the 2PCFs of many more WL peak catalogues are combined: constraints from the combined 2PCFs are much smaller than the best constraints from any individual catalogue (cf. Fig. 6). Increasing the number of catalogues used in the combined case from four to eight, improves the constraints on Ω_m and S_8 by roughly 30 and 20 per cent, respectively, there is only a small improvement for h , and the w_0 constraints improve by nearly a factor of 2. Potentially, one could use a very large number of ν thresholds, by reducing the ν increments further, such as $\nu > 1.0, 1.25, \dots, 4.5$. However at some point the 2PCFs from adjacent catalogues become so correlated that there is no gain in extra information, and this significantly increases

the length of our data vector. Our choice is a compromise between having several ν bins that span a range of peak heights while keeping a small enough data vector to calculate accurate covariance matrices.

The complementarity between 2PCFs with different thresholds, including the potential for degeneracy breaking, is the main factor that contributes to the strong peak 2PCF constraints, when many catalogues are used. In the following paragraphs, we provide a physical interpretation for this behaviour.

There have been several studies on the structures that produce WL peaks (Yang et al. 2011; Liu & Haiman 2016; Wei et al. 2018). These studies show that in general, medium-height peaks correspond to chance alignments of multiple haloes along the line of sight, and high-amplitude peaks are often associated with single clusters along the line of sight. Peaks created by LSS projections will contain different cosmological information to those created by single clusters. For the peak abundance, this can be seen in the covariance matrix (Fig. B1, also studied in Martinet et al. 2018) where there is anticorrelation between low and high peaks. There is a similar behaviour in the peak 2PCF covariance matrix in Fig. B2, where there is very little correlation between the $\nu > 1$ and $\nu > 4$ catalogues. However this feature is less pronounced, due to adjacent catalogues sharing a large fraction of the same peaks.

Because peaks of different amplitudes are produced by different types of dark matter structures, the 2PCF of each WL peak catalogue probes slightly different regimes of the dark matter distribution. By combining multiple peak catalogues with different peak height thresholds, we essentially add additional information about the peak height to the 2PCF.

As shown by the table in Fig. 7, the peak abundance and peak 2PCF provide similar constraints on Ω_m ; however, the constraints on S_8 are twice as strong for the peak abundance compared to the peak 2PCF: In the Ω_m - S_8 plane, the peak abundance contour is significantly tighter than the peak 2PCF contour in the S_8 direction. When the two probes are combined, there is an overall improvement on the Ω_m and S_8 constraints by a factor of 2, relative to the peak abundance alone. This leads to a good overall improvement in the Ω_m - S_8 plane when the peak abundance and 2PCF are combined, as shown by the green contour.

The peak 2PCF is able to constrain both h and w_0 with greater accuracy than the peak abundance. There also appears to be some orthogonality between the abundance and 2PCF constraints in the h - S_8 and w_0 - Ω_m planes. In the w_0 - h plane, the parameter constraints are dominated by the peak 2PCF contours, while the peak abundance contours are significantly larger than the former. This indicates that the peak 2PCF offers a great deal of complementary information to the peak abundance, and combining the two probes can significantly improve constraints in the four-dimensional parameter space studied here. The behaviour of the constraints from peak abundance and 2PCF, especially those on h and w_0 , are consistent with the observations we made above for Figs 3 and 5.

In Fig. 8, we also introduce the parameter constraint forecasts for WL voids from Davies et al. (2020b), which are measured with the same methodology and specifications used in this work, to compare the constraining power of these two different probes. This is important since the voids studied in Davies et al. (2020b), which were found to be a promising void definition (Davies et al. 2020a), are identified as underdense regions in the distribution of WL peaks. This means that the properties of WL voids are likely correlated with the number and clustering of peaks, and we need a joint analysis to reveal the amount of complementary information contained in the two probes. The forecasts from the WL voids (blue), making use of both their abundances and tangential shear profiles,

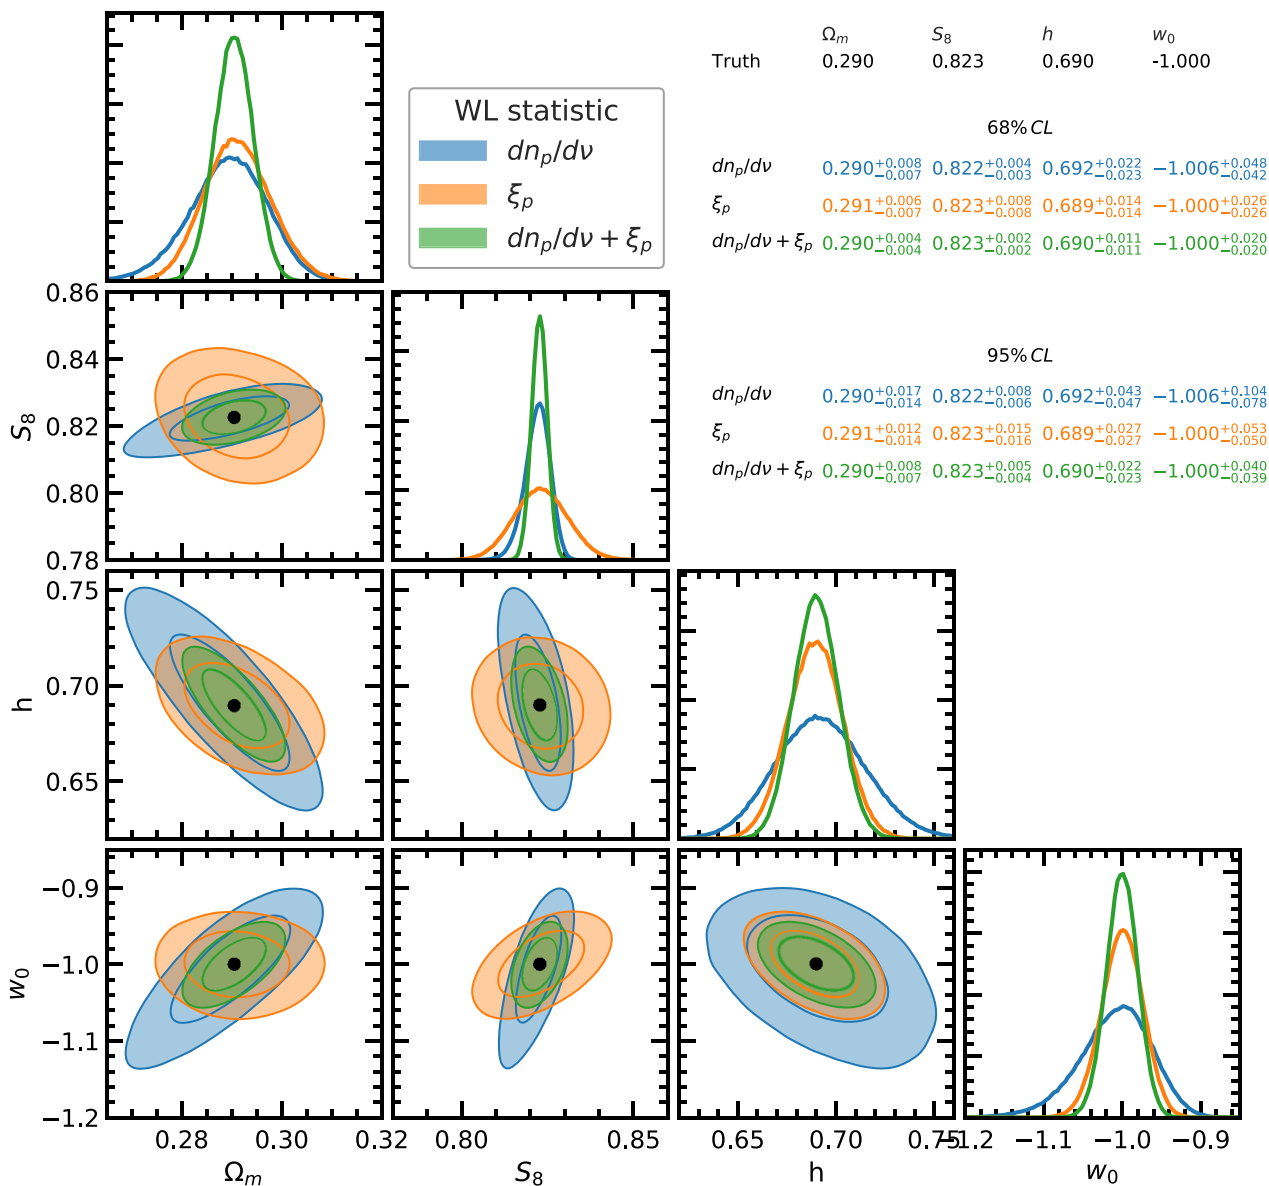


Figure 7. (Colour online) The same as Fig. 6 but for the combination of eight peak 2PCFs from peak catalogues with $\nu > 1.0, 1.5, \dots, 4.5$ (orange), the peak abundance (blue), and the combination of the two (green).

are compared to the WL peak forecasts (orange), which combine the peak abundance and peak 2PCF. We note that both the void and peak contours presented here are for the combination of the $\nu > 1, 2, 3$, and 4 peak catalogues (excluding the peak abundance that does not combine multiple catalogues). This is to provide a fair comparison between the voids and the peaks. In principle, the void contours could be measured for the eight catalogues used for the orange peak 2PCF contours in Fig. 7, however this would cause our data vector to become too large, even for our high number of SLICS realizations.

Overall, both the peaks and voids are able to measure the four cosmological parameters with similar accuracy. The voids provide notably tighter measurements of h and w_0 . The void contours are smaller than the peak contours, and follow similar degeneracy directions for all combinations of parameters. The void and peak contours are most similar in the S_8 - w_0 plane, and most distinct in the w_0 - h plane. When the peaks and voids are combined (green contours), there is a small improvement on the Ω_m , S_8 , and h

measurements, and, there is also a reasonable improvement on w_0 , indicating that the WL peak and void statistics are complementary to each other.

As a comparison, we also include the forecast contours using the standard cosmic shear 2PCFs (ξ_+ and ξ_- combined) in grey (for details on how these are measured, see Davies et al. 2020b). For fair comparisons, the cosmological model dependence and covariance matrix for these were both obtained using the same simulation data as used for the peak and void analyses throughout this paper. For Ω_m and h , WL peaks or WL voids (or both of them) give similar constraints as the shear 2PCFs; however, for S_8 and w_0 , the former probes actually can place tighter constraints (for this survey specification) by roughly a factor of 2, indicating again the benefit of exploring beyond-two-point WL statistics to help maximize the information that can be extracted. In some parameter planes, such as S_8 - h and h - w_0 , there is a clear orthogonality between the degeneracy directions of the peak/void statistics and the shear 2PCFs.

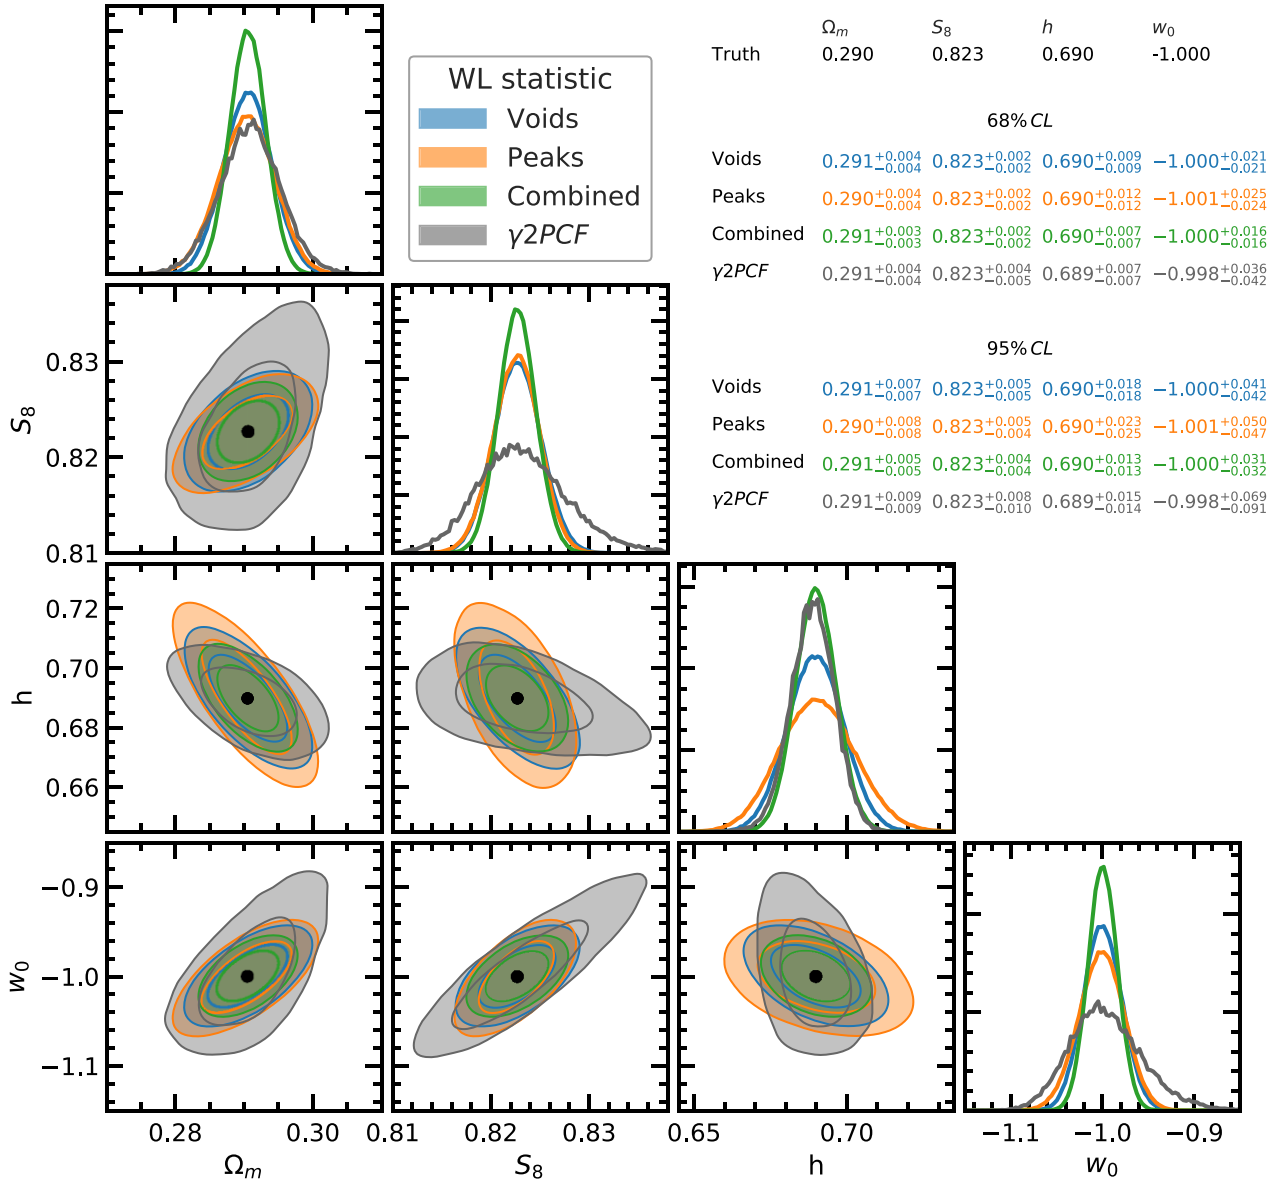


Figure 8. (Colour online) The same as Fig. 6, but for void statistics presented in (Davies et al. 2020b) (blue), peak statistics (orange), and the combination of peak and void statistics (green). Both the peak and void statistics use the combination of the $\nu > 1, 2, 3,$ and 4 catalogues. Shear 2PCF forecasts are shown in grey.

6 DISCUSSION AND CONCLUSIONS

We have tested the sensitivity of the WL peak statistics to the cosmological parameters Ω_m , S_8 , h , and w_0 and compared the peak 2PCF to the peak abundance. In order to achieve this, we have trained a Gaussian Process emulator with 26 cosmologies sampled in the 4D parameter space using a Latin hypercube, which we used to predict the peak statistics for arbitrary cosmologies (within the range spanned by the training cosmologies). We have run MCMC samplings from our mock WL data to forecast the accuracy's at which these four parameters can be constrained by a future, LSST-like, lensing survey, using the above WL peak statistics.

Using the emulators, we have studied the behaviour of the WL peak 2PCF in detail, and made connections to the well-established peak abundance. A main feature of our peak 2PCF analysis is that we generate a WL peak catalogue from the entire peak population

by introducing a peak height (ν) threshold, below which all peaks are removed, and then vary this threshold to generate multiple catalogues. We then study the behaviour of the WL peak 2PCF of these catalogues as this ν threshold changes.

In Marian et al. (2013), it has been shown that the WL peak 2PCF of high-amplitude peaks provides little complementary information to the peak abundance. In this work, we have presented some additional steps that are able to further push the utility of the WL peak 2PCF. These additional steps significantly improve the overall constraining power of WL peaks when the abundance and 2PCFs are combined. First, we study the 2PCF of low-amplitude (ν) peaks, and find that it contains significant cosmological information compared to the 2PCF of high-amplitude peaks. For example, in Fig. 6, the constraints on w_0 are roughly four times stronger for the $\nu > 2$ catalogue compared to the $\nu > 4$ catalogue. Secondly, we find that the 2PCFs of multiple catalogues are complementary to each other, and when combined,

the peak 2PCF can constrain Ω_m with tighter accuracy than the peak abundance, and that it is able to constrain both h and w_0 with significantly greater accuracy than the peak abundance alone. We also find that the peak abundance provides constraints that are twice as tight on S_8 than the combined peak 2PCF, indicating that in order to fully exploit the cosmological information contained in WL peaks, both their abundance and their clustering should be measured and combined. This is illustrated by the green contours in Fig. 7, which show the total constraints from WL peaks in which the abundance is combined with the combined 2PCF from different ν catalogues. Here, the abundance plus the clustering forecasts are roughly twice as strong as those for either of the individual cases (orange for 2PCF and blue for abundance). When we compare the constraints from the peak abundance plus the peak 2PCF to those from the shear 2PCF, we find that the peaks are able to constrain Ω_m , S_8 , and w_0 with greater precision than the shear 2PCF, the most significant improvement is for S_8 and w_0 , which improve by roughly a factor of 2. Finally, the information required to measure the peak 2PCF is already present when the peak abundance is measured. Therefore, the addition of the peak 2PCF to any pre-existing peak abundance analysis pipeline will require minimal modifications, making the peak 2PCF a very promising probe.

We also include a comparison of the forecasts from WL peaks to the WL voids studied in Davies et al. (2020b), and find that the combination of the two can improve the constraints on Ω_m , S_8 , and h , and can provide significant improvements on the w_0 measurements. The WL voids are sensitive to the N -point correlation function of peaks (White 1979), and the improved constraints resulting from combining WL peaks and voids shows that three and higher-order correlation functions in the peak distribution contain complementary cosmological information. WL voids are one simple way to access the information contained in the higher order correlation functions of peaks.

The improved parameter constraints from the combination of all of the peak and void probes presented here can also be explained through their full covariance matrix, which includes cross-correlations between probes, present in Appendix B. From the covariance matrix, it is clear that many of the probes studied here have a large degree of statistical independence. This leads to complementary information between the probes which yields tighter constraints when the different probes are combined.

We highlight that the work carried out here applies to the 4D parameter space in Fig. 1, and may change if additional parameters in the Λ CDM model, such as the spectral index, are included. Our results may also be sensitive to changes in curvature, massive neutrinos or other sources of additional physics. For example, in Davies, Cautun & Li (2019b), we found that the peak abundance is sensitive to the nDGP modified gravity model, and Liu et al. (2016) have used the WL peak abundance to constrain $f(R)$ gravity.

We also note that the ray-tracing method used to obtain our WL maps employs some approximations, the most important being the Born approximation. We do not expect this simplified framework to have affected our results. For example, Hilbert et al. (2020) have compared multiple ray-tracing methods, such as our approach versus full techniques run on the fly at the same time as the numerical simulation, and found very good agreement in the WL convergence and shear power spectrum, as well as in the abundance of peaks with heights $\nu < 6$. Hilbert et al. have found some discrepancies between methods in the peak abundance for $\nu > 6$, which should not be surprising since such peaks correspond to massive clusters; however, such differences are unlikely to affect our results since peaks with such high amplitudes constitute only a very small fraction of the

population even for our $\nu > 4$ peak catalogue (e.g. see Fig. 2). The agreement of different ray-tracing methods when predicting the peak 2PCF remains to be studied. In this work the peak 2PCF is measured on scales larger than $0^{\circ}.1$, so any approximations we employed, such as the Born approximation, need to fail on similar scales to have an impact on our measurements.

In addition, the simulations used to construct the emulators for the different WL statistics analysed here are limited in their number of nodes sampled with the Latin hypercube. As the results of this paper suggest, future WL observations can place competitive constraints on the various cosmological parameters, with significantly smaller contours than the current status. As the contours keep shrinking around the best-fitting model, improved emulators that can more accurately capture the small effects induced by small variations of parameters will be needed. In the future, it will be necessary to simulate cosmological models sampled using a nested Latin hypercube, or nested Latin hypercubes, to refine the emulators used in this work.

The results presented here may be further improved with the inclusion of tomography. This is the standard approach when using the shear two-point correlation function to measure cosmological parameters, and typically this significantly improves the constraints on w_0 . Therefore, in a future work, we will test how cosmological parameter constraints can be improved when using WL voids with tomography.

Finally, in order to use the WL peak 2PCF in observations, it will be important to understand the impact of baryonic physics, the intrinsic alignment of galaxies, and the uncertainty associated with photometric redshifts and shear calibration. It is already established that the WL peak abundance is altered in presence of baryons (Osato, Shirasaki & Yoshida 2015; Coulton et al. 2019; Fong et al. 2019; Weiss et al. 2019) and intrinsic galaxy alignments (Harnois-Déraps, Martinet & Reischke 2021), so it will also be necessary to test how the WL peak 2PCF is affected by these. The behaviour of these systematics will be an important factor in determining which peak catalogues can be used to reliably measure the peak 2PCF in observations. When identifying WL peaks in observational data, previous studies (e.g. Kacprzak et al. 2016; Martinet et al. 2018; Harnois-Déraps et al. 2020) have discarded peaks with $\nu < 0$ and $\nu > 4$, in order to mitigate the impact of these systematics on the peak abundance. We would expect a similar approach to be valid for the peak 2PCF, though that remains to be tested. Additionally, the removal of the high-amplitude peaks that are most affected by these systematics may not change our posterior forecasts by a significant amount, since they also contribute the least to the parameter constraints.

ACKNOWLEDGEMENTS

We thank Ludovic Van Waerbeke for sharing the code used in the WL convergence map reconstruction.

CTD is funded by a UK Science and Technology Facilities Council (STFC) PhD studentship through grant ST/R504725/1. MC is supported by the EU Horizon 2020 research and innovation programme under a Marie Skłodowska-Curie grant agreement 794474 (DancingGalaxies). JHD acknowledges support from an STFC Ernest Rutherford Fellowship (project reference ST/S004858/1). BG acknowledges the support of the Royal Society through an Enhancement Award (RGF/EA/181006). YC acknowledges the support of the Royal Society through a University Research Fellowship and an Enhancement Award. BL is supported by an ERC Starting Grant,

ERC-StG-PUNCA-716532, and additionally supported by the STFC Consolidated Grants [ST/P000541/1, ST/T000244/1].

The SLICS numerical simulations can be found at <http://slics.roe.ac.uk/>, while the *cosmo*-SLICS can be made available upon request. 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.

DATA AVAILABILITY

The data used in this work are available upon request. See Harnois-Déraps et al. (2019) for more details.

REFERENCES

- Aihara H. et al., 2019, *PASJ*, 71, 114
 Anderson T. W., 2003, *An Introduction to Multivariate Statistical Analysis*. Wiley-Interscience, New Jersey
 Asgari M. et al., 2020, *A&A*, 645, 31
 Bacon D. J., Refregier A. R., Ellis R. S., 2000, *MNRAS*, 318, 625
 Bartelmann M., Schneider P., 2001, *Phys. Rep.*, 340, 291
 Burger P., Schneider P., Demchenko V., Harnois-Déraps J., Heymans C., Hildebrandt H., Unruh S., 2020, *A&A*, 642, A161
 Coulton W. R., Liu J., McCarthy I. G., Osato K., 2019, *MNRAS*, 495, 2531
 DES Collaboration et al., 2021, *Phys. Rev. D*, 105, 023520
 Davies C. T., Cautun M., Li B., 2018, *MNRAS*, 480, L101
 Davies C. T., Cautun M., Li B., 2019a, *MNRAS*, 488, 5833
 Davies C. T., Cautun M., Li B., 2019b, *MNRAS*, 490, 4907
 Davies C. T., Paillas E., Cautun M., Li B., 2020a, *MNRAS*, 500, 2417
 Davies C. T., Cautun M., Giblin B., Li B., Harnois-Déraps J., Cai Y.-C., 2020b, *MNRAS*, 507, 2267
 Dietrich J. P., Hartlap J., 2010, *MNRAS*, 402, 1049
 Fong M., Choi M., Catlett V., Lee B., Peel A., Bowyer R., King L. J., McCarthy I. G., 2019, *MNRAS*, 488, 3340
 Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
 Fu L. et al., 2008, *A&A*, 479, 9
 Giblin B. et al., 2018, *MNRAS*, 480, 5529
 Habib S., Heitmann K., Higdon D., Nakhleh C., Williams B., 2007, *Phys. Rev. D*, 76, 083503
 Hamana T., Yoshida N., Takada M., 2004, *MNRAS*, 350, 893
 Hamana T., Shirasaki M., Lin Y.-T., 2020, *PASJ*, 72, 78
 Harnois-Déraps J., van Waerbeke L., 2015, *MNRAS*, 450, 2857
 Harnois-Déraps J. et al., 2018, *MNRAS*, 481, 1337
 Harnois-Déraps J., Giblin B., Joachimi B., 2019, *A&A*, 631, A160
 Harnois-Déraps J., Martinet N., Castro T., Dolag K., Giblin B., Heymans C., Hildebrandt H., Xia Q., 2020, *MNRAS*, 506, 1623
 Harnois-Déraps J., Martinet N., Reischke R., 2021, *MNRAS*, 509, 3868
 Hartlap J., Simon P., Schneider P., 2007, *A&A*, 464, 399
 Heydenreich S., Brück B., Harnois-Déraps J., 2020, *A&A*, 648, 18
 Heymans C. et al., 2012, *MNRAS*, 427, 146
 Hikage C. et al., 2019, *PASJ*, 71, 43
 Hilbert S. et al., 2020, *MNRAS*, 493, 305
 Hildebrandt H. et al., 2017, *MNRAS*, 465, 1454
 Hoekstra H. et al., 2006, *ApJ*, 647, 116
 Jain B., Van Waerbeke L., 2000, *ApJ*, 530, L1
 Kacprzak T. et al., 2016, *MNRAS*, 463, 3653
 Kaiser N., Wilson G., Luppino G. A., 2000, preprint ([arXiv:astro-ph/0003338](https://arxiv.org/abs/astro-ph/0003338))
 Kilbinger M., 2015, *Rep. Prog. Phys.*, 78, 086901
 Kilbinger M. et al., 2013, *MNRAS*, 430, 2200
 Landy S. D., Szalay A. S., 1993, *ApJ*, 412, 64

- Li Z., Liu J., Matilla J. M. Z., Coulton W. R., 2019, *Phys. Rev. D*, 99, 063527
 Liu J., Haiman Z., 2016, *Phys. Rev. D*, 94, 043533
 Liu X. et al., 2016, *Phys. Rev. Lett.*, 117, 051101
 Marian L., Smith R. E., Hilbert S., Schneider P., 2013, *MNRAS*, 432, 1338
 Martinet N. et al., 2018, *MNRAS*, 474, 712
 Martinet N., Harnois-Déraps J., Jullo E., Schneider P., 2020, *A&A*, 646, 16
 Osato K., Shirasaki M., Yoshida N., 2015, *ApJ*, 806, 186
 Pedregosa F. et al., 2011, *J. Machine Learning Res.*, 12, 2825
 Pen U.-L., Zhang T., van Waerbeke L., Mellier Y., Zhang P., Dubinski J., 2003, *ApJ*, 592, 664
 Planck Collaboration VI, 2018, *A&A*, 641, 77
 Riess A. G., Casertano S., Yuan W., Macri L. M., Scolnic D., 2019, *ApJ*, 876, 85
 Schneider P., van Waerbeke L., Kilbinger M., Mellier Y., 2002, *A&A*, 396, 1
 Schneider M. D., Knox L., Habib S., Heitmann K., Higdon D., Nakhleh C., 2008, *Phys. Rev. D*, 78, 063529
 Semboloni E. et al., 2006, *A&A*, 452, 51
 Troxel M. A. et al., 2018, *Phys. Rev. D*, 98, 043528
 Van Waerbeke L. et al., 2000, *A&A*, 358, 30
 Verde L., Treu T., Riess A. G., 2019, *Nat. Astron.*, 3, 891
 Wei C., Li G., Kang X., Liu X., Fan Z., Yuan S., Pan C., 2018, *MNRAS*, 478, 2987
 Weiss A. J., Schneider A., Sgier R., Kacprzak T., Amara A., Refregier A., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 011
 White S. D. M., 1979, *MNRAS*, 186, 145
 Wittman D. M., Tyson J. A., Kirkman D., Dell'Antonio I., Bernstein G., 2000, *Nature*, 405, 143
 Yang X., Kratochvil J. M., Wang S., Lim E. A., Haiman Z., May M., 2011, *Phys. Rev. D*, 84, 043529

APPENDIX A: EMULATOR ACCURACY

In this section, we present the accuracy of the peak abundance and peak 2PCF emulator used for our cosmological forecasts.

In order to test the accuracy of the emulator, we employ a cross-validation test, which is outlined as follows. First, one node from the training data (simulated data) is removed, and the emulator is then trained with the remaining 25 nodes, for a given statistics. The emulator prediction for the removed node is then calculated, and this




Figure A1. (Colour online) The cross-validation of the WL peak abundance emulator accuracy. One node is removed from the training set, and the difference between the emulation (Em) and simulation (Sim) predictions of the removed node are compared relative to the simulation standard error (σ). This process is repeated for each of the 26 nodes, giving an upper limit on the emulator accuracy. The iteration where the node for the fiducial cosmology is removed is shown by the blue line. Dashed lines are added at the 1σ level to help guide the reader.




Figure A2. (Colour online) The same as Fig. A1, but showing the percentage accuracy relative to the simulated predictions. The dashed lines enclose the 1 per cent region.

result is compared to the simulated version, by taking the difference between the two and dividing it by the standard error measured in the simulated data for that node. The above steps are repeated 25 more times, removing a different node from the training data at each iteration. This results in measurements of the emulator accuracy at each node, which is an upper limit, since the accuracy increases as more training data are used, and the cross-validation measurements uses training data with one less node than the training data used in the main analysis.

Fig. A1 shows the cross-validation test for WL peak abundance. The cross-validation for the fiducial cosmology is shown in blue, and the remaining nodes are shown in grey. The fiducial cosmology is the node of most interest, as all posterior contours presented in this work reside close to this region. The dashed lines delineate the region where the accuracy of the emulator is within the standard error of the simulated data. The blue curve shows that, for the fiducial cosmology, the emulator is accurate to within 1σ , as roughly 68 per cent of the bins are within the 1σ region. The grey curves show that the accuracy is lower for the other nodes, and we find that the accuracy decreases as we approach the edges of the *cosmo*-SLICS parameter space. This is to be expected, as the emulator has less data to train from for these regions. Fig. A2 is similar to Fig. A1, but shows the percentage accuracy of the emulator for the cross-validation test. We can observe that the accuracy is within roughly 1 per cent for $\nu \lesssim 3$, increasing to up to 4 per cent at $\nu \gtrsim 5$ due to the more noisy measurement for the high- ν peaks.

Fig. A3 is the same as Fig. A1, but shows the cross-validation test for the WL peak 2PCF, for peak catalogues with heights $\nu > 1$ (top left-hand panel), $\nu > 2$ (top right-hand panel), $\nu > 3$ (bottom left-hand panel), and $\nu > 4$ (bottom right-hand panel). The figure shows that, similar to the WL peak abundance, the emulator is accurate to within 1σ at the fiducial cosmology for the peak 2PCF for all four catalogues. Fig. A4 is the same as Fig. A3, but shows the percentage accuracy of the peak 2PCF emulator applied in the cross-validation test. For the $\nu > 1$ catalogue, the accuracy is mostly within 10 per cent, with a few bin at 20 per cent. The accuracy is within 10% for the $\nu > 2$ and 3 peak catalogues, and for $\nu > 4$, the accuracy is within 10 per cent except for the final bin.




Figure A3. The same as Fig. A1 but for the WL peak 2PCFs. The four panels correspond to the WL peak 2PCFs of WL peaks with heights $\nu > 1$ (top left-hand panel), $\nu > 2$ (top right-hand panel), $\nu > 3$ (bottom left-hand panel), and $\nu > 4$ (bottom right-hand panel).




Figure A4. The same as Fig. A3 but for the percentage accuracy relative to the simulated predictions. The dashed lines enclose the 10 per cent region.

APPENDIX B: COVARIANCE

As shown by equation (14), we require the (inverted) covariance matrix of the data vector in order to produce our forecasts. Within the matrix, the diagonal elements correspond to the variance of each of the data vector bins and the off-diagonal elements give the covariance between all possible pairs of bins. When multiple probes are combined into a single data vector, any correlated or duplicate information between the probes is accounted for by the cross-covariance within the matrix.

In Fig. B1, we present the total correlation matrix for the all of the probes studied in this work. This corresponds to the matrix that is used to produce the green likelihood contour in Fig. 8. We present the correlation matrix instead of the covariance matrix, as it allows for easier visual interpretation, which is expressed in terms of the covariance matrix as follows

$$R_{ij} = \frac{\text{cov}(i, j)}{\sigma_i \sigma_j}, \quad (\text{B1})$$

where R is the correlation matrix, cov is the covariance matrix, and σ is the standard deviation for a given bin.

Starting from the bottom left-hand panel of the figure, the diagonal tiles enclosed by the black lines show the correlation for the following statistics (which are labelled with the range of peak heights used in their identification): peak abundance ($0 < \nu < 6$), peak 2PCF (four catalogues with thresholds $\nu > 1, 2, 3, 4$), and WL void abundance and WL void tangential shear profiles where the voids are identified using the same four peak catalogues. The remaining off-diagonal terms show the cross-covariances between all possible combinations of the probes.

For the peak abundance, the figure shows that the low-amplitude peaks are somewhat correlated with other low-amplitude peaks, and a similar behaviour is present for the high-amplitude peaks, as shown by the green regions in the bottom left and top right of the peak

abundance correlation tile. There also appears to be a small amount of anticorrelation between low- and high-amplitude peaks, as shown by the dark regions in the top left- and bottom right-hand corner of the tile.

For the diagonal peak 2PCF tiles, each bin in the 2PCF appears to be correlated with all of the other bins. For the off-diagonal tiles between the different peak 2PCFs, there is also a high amount of correlation, which is again expected, as the main difference between the 2PCFs is simply a change in amplitude, and all catalogues have some fraction of their tracer population in common.

For the tiles representing the correlation between the peak 2PCFs and the WL void abundances, we see some correlation between the peak 2PCFs and the small radii WL voids (especially for high ν thresholds). This is also to be expected since the WL voids are identified from a Delaunay triangulation of the peaks, which will be sensitive to the peak clustering. It is interesting to see that this correlation drops off as the void size increases, which may indicate that higher-order clustering such as the three-point correlation function of WL peaks dictates the abundance of large voids.

In Fig. B2, we show the correlation matrix for the peak abundance combined with the eight peak 2PCFs with $\nu > 1.0, 1.5, \dots, 4.5$. The figure shows that, for the peak 2PCF, adjacent catalogues (similar ν thresholds) are highly correlated. This is to be expected as the tracer populations are very similar for adjacent catalogues. The correlation reduces significantly as the difference between the ν thresholds increases, which is again expected as this is where the tracer populations will differ the most. The low correlation between the peak 2PCFs with very different ν thresholds is a strong contribution to the improved constraining power from the combination of multiple peak 2PCFs, alongside any complementary parameter degeneracy directions.




Figure B1. (Colour online) The correlation matrix for all probes studied in this work, which are as follows (from the left to right): peak abundance (dn/dk), peak 2PCF (ξ_p), void abundance (dn/dR_v), and void tangential shear profiles (γ_t).




Figure B2. (Colour online) The correlation matrix for the combination of the peak abundance (dn_p/dk) and the peak 2PCF (ξ_p , for eight catalogues with $\nu > 1.0, 1.5, \dots, 4.5$).

This paper has been typeset from a \TeX/L\AA\TeX file prepared by the author.