Version September 29, 2020
Preprint typeset using LaTeX style openjournal v. 09/06/15

arXiv:2007.07253v2 [astro-ph.CO] 28 Sep 2020

THE IMPACT OF SIGNAL-TO-NOISE, REDSHIFT, AND ANGULAR RANGE ON THE BIAS OF WEAK LENSING 2-POINT FUNCTIONS

Amy J. Louca and Elena Sellentin
Leiden Observatory, Leiden University, Huygens Laboratory, Niels Bohrweg 2, NL-2333 CA Leiden, The Netherlands.
Version September 29, 2020

ABSTRACT

Weak lensing data follow a naturally skewed distribution, implying the data vector most likely yielded from a survey will systematically fall below its mean. Although this effect is qualitatively known from CMB-analyses, correctly accounting for it in weak lensing is challenging, as a direct transfer of the CMB results is quantitatively incorrect. While a previous study (Sellentin et al. 2018) focused on the magnitude of this bias, we here focus on the frequency of this bias, its scaling with redshift, and its impact on the signal-to-noise of a survey. Filtering weak lensing data with COSEBIs, we show that weak lensing likelihoods are skewed up until $\ell \approx 100$, whereas CMB-likelihoods Gaussianize already at $\ell \approx 20$. While COSEBI-compressed data on KiDS- and DES-like redshift- and angular ranges follow Gaussian distributions, we detect skewness at 6σ significance for half of a Euclid- or LSST-like data set, caused by the wider coverage and deeper reach of these surveys. Computing the signal-to-noise ratio per data point, we show that precisely the data points of highest signal-to-noise are the most biased. Over all redshifts, this bias affects at least 10% of a survey's total signal-to-noise, at high redshifts up to 25%. The bias is accordingly expected to impact parameter inference. The bias can be handled by developing non-Gaussian likelihoods. Otherwise, it could be reduced by removing the data points of highest signal-to-noise.

1. INTRODUCTION

Weak lensing is a central observable for contemporary cosmology. Arising from the bending of light rays in gravitational fields, it allows not only to study gravity beyond its Newtonian limit, but also to map dark and luminous matter alike. When jointly analyzed with galaxy clustering, it can constrain the gravitational slip parameter (Daniel et al. 2008) and thereby discriminate between different types of gravitational theories. Considerable observational efforts are thus invested into weak lensing, including the Kilo Degree Survey (Hildebrandt et al. 2020), Dark Energy Survey (Abbott et al. 2019) and the Hyper Suprime-Cam (Mandelbaum et al. 2018).

Weak lensing is sensitive to the cosmological parameter σ_8, which describes the 'clumpiness' of matter in the Universe in the sense that σ_8 sets the amplitude of the matter power spectrum. Measurements of σ_8 derived from weak lensing experiments are mildly discrepant with measurements of σ_8 from the cosmic microwave background (CMB), see Planck Collaboration (2014, 2016, 2018). Since many years, weak lensing yields lower values of σ_8, and it has been repeatedly proposed that this might be indicative of an unsolved bias in contemporary weak lensing constraints.

The term 'bias' usually carries a negative connotation. However, in a previous study Sellentin et al. (2018) had shown that although the summary statistics derived from weak lensing observations are indeed 'biased', this bias has here to be understood in the literal sense. An estimator is said to be biased if the value it yields most likely differs from the one it yields on average. An experiment's average outcome then deviates from its most likely outcome. This occurs if the estimator is drawn from a skewed distribution.

sellentin@strw.leidenuniv.nl

Sellentin et al. (2018) had shown that all weak lensing 2-point function estimators will be affected to various degrees by such a literal bias, as they follow a skewed distribution. In consequence, a sound weak lensing data vector will likely display an amplitude that is below average.

If such weak lensing data were analyzed with the matching skewed likelihood, then parameter inference could proceed self-consistently. However, the current standard weak lensing analysis employs a symmetric Gaussian likelihood instead. When used to analyze a data vector that is biased low, the Gaussian likelihood will force the fit to center on this low amplitude. As σ_8 scales the amplitude of weak lensing observables, it can thus be expected that a systematically lowered value for σ_8 is found in the standard analysis.

Quantifying by how much σ_8 will be biased low – and removing this bias – is a highly non-trivial challenge: solving both questions implies the *correct* weak lensing likelihood must be found, rather than a heuristic non-Gaussian substitute.

Accordingly, a skewed baseline likelihood to handle weak lensing observations alongside galaxy clustering and their cross-correlation has been presented in Manrique-Yus & Sellentin (2020). It takes the form of a Bayesian Hierarchical Model and at its core lies the same Wishart distribution that also describes the skewed distribution of the low multipoles of the cosmic microwave background (CMB), see Hamimeche & Lewis (2008, 2009); Upham et al. (2020). In comparison to CMB data, weak lensing data follow a decidedly more skewed distribution, as was revealed by simulations (Sellentin & Heavens 2018; Sellentin et al. 2018; Harnois-Déraps & van Waerbeke 2015) and subsequently by Diaz Rivero & Dvorkin (2020). This additional skewness can be implemented in the skewed baseline likeli-

hood through a free parameter $g_{\rm eff}$, but to date $g_{\rm eff}$ can only be measured from simulations, and computing it from first principles is still unsolved.

In this paper, we thus return to the origin of skewness in weak lensing likelihoods. In particular, we study whether the problematic skewness can be suppressed without losing precision. Although the skewness could be suppressed by invoking the Central Limit Theorem, this is highly suboptimal: inconsiderate averaging will remove information. We therefore opt for the more refined approach by filtering, and study whether skewness can be filtered out. Should such a suppression succeed, then developing non-Gaussian likelihoods could be put aside. We here opt for a set of filters known as COSEBIs, as these are constructed to remove further weak lensing systematics. If also able to suppress skewness, a COSEBI-analysis would thus present a highly controllable weak lensing analysis.

The setup of this paper is as follows. In Sect. 2 we introduce weak lensing. In Sect. 3 we describe why it can be hoped that filtering may suppress skewness and ensuing biases. In Sect. 4 we introduce the essentials of COSEBI filters, and in Sect. 5 we create 819 COSEBI samples to estimate their sampling distribution. Our strategy to detect deviations from a Gaussian distribution is described in Sect. 6. From Sect. 7 we present the main results of this paper: Sect. 7.1 describes how sampling distribution skewness evolves with redshift, and Sect. 7.2 describes the excess skewness of weak lensing measurements in comparison to the CMB case. Sect. 7.3 contrasts Euclid- and LSST-like angular ranges with KiDS- and DES-like angular ranges. Sect. 8 then studies the excess probability of finding a 'low' weak lensing outcome, if a Gaussian likelihood is insisted upon. Our conclusions are presented in Sect. 9, and the appendix describes the construction of COSEBIs and answers questions concerning gravity-induced skewness often encountered.

2. WEAK LENSING SIGNAL AND NOISE

Matter between the observer and far-away sources distorts the image of the sources via gravitational lensing. As light from neighbouring sources propagates through similar foreground fields, weak lensing causes the images of neighbouring sources to have a correlated alignment. This correlation can be extracted from observations, see Bartelmann & Schneider (2001); Kilbinger (2015).

Such weak lensing signals carry two types of noise, 'shape noise' due to the diversity of galaxy shapes and orientations, and 'cosmic variance' due to the stochastic nature of the matter fields transversed by the propagating light.

A weak lensing analysis begins from a *weak lensing galaxy catalogue* that stores for N galaxies their celestial positions $\boldsymbol{x}$ together with the two measured ellipticity components ϵ_1 and ϵ_2 and the redshift. From such a catalogue, the standard analysis of a weak lensing data set computes binned 2-point statistics, such as the correlation functions $\xi_\pm$. If θ denotes the angular separation on the sky, then the weak lensing correlation functions $\xi_\pm$ can be extracted via the estimator

$$\hat{\xi}_\pm(\theta) = \frac{\sum_{ij} w_i w_j \left[\epsilon_t(\boldsymbol{x}_i)\epsilon_t(\boldsymbol{x}_j) \pm \epsilon_\times(\boldsymbol{x}_i)\epsilon_\times(\boldsymbol{x}_j)\right]}{\sum_{ij} w_i w_j}, \quad (1)$$

where $w_{i,j}$ are the weights arising from the shape measurement process on realistic galaxies, and the subscript $\pm$ is inherited from the plus-minus on the right-hand side. We use hats to denote noisy quantities.

The angular separation θ and the positions $\boldsymbol{x}$ of galaxies in the sky are still two-dimensional only. As our Universe evolves our time, it is advantageous to split a galaxy catalogue radially into subsets of galaxies. This splitting is known as redshift tomography. Denoting the redshift by z, the ith redshift bin is $n^i(z)$, which is the normalized distribution of the galaxies assigned to the bin i.

These redshift bins are then re-expressed as a function of the comoving distance χ,

$$g^i(\chi) = \int_\chi^{\chi_{\rm h}} {\rm d}\chi' n^i_\chi(\chi') \frac{\chi' - \chi}{\chi'}, \quad (2)$$

known as the weak lensing kernel corresponding to the ith galaxy population.

The light emitted from galaxies per bin will propagate through foreground fields towards the observer. Of these, cosmological theory can predict the matter power spectrum P_δ. The continuous propagation of light through these foreground fields causes that the weak lensing power spectrum $\mathrm{P}^{ij}_\gamma(\ell)$ integrates over the matter power spectrum, weighted by the lensing kernels

$$\mathrm{P}^{ij}_\gamma(\ell) = \frac{9H_0^4\Omega_m^2}{4c^4} \int_0^{\chi_{\rm h}} {\rm d}\chi \frac{g^i(\chi)g^j(\chi)}{a^2} \mathrm{P}_\delta\left(\frac{\ell+1/2}{\chi}, \chi\right), \quad (3)$$

where the Limber approximation has been used. Furthermore, ℓ is a spherical harmonic multipole, H_0 is the Hubble constant, Ω_m is the matter density parameter, c is the speed of light. Both the matter power spectrum P_δ and the comoving distance $\chi(z)$ as a function of redshift depend on a broad range of cosmological parameters, which can accordingly be inferred from weak lensing data.

3. FILTERING TO SUPPRESS SKEWNESS?

Comparing Eq. (1) and Eq. (3), an inherent disparity between weak lensing observations and its theoretical predictions becomes apparent: while the weak lensing signal is easiest to extract in real space θ, the theoretical predictions are easier to compute in harmonic space ℓ. The translation from harmonic space to real space occurs via a filter $F(\ell\theta)$ such that a general weak lensing 2-point function can be written as

$$\xi^{ij}_{\rm F}(\theta) = \frac{1}{2\pi} \int {\rm d}\ell\ \ell F(\ell\theta) P^{ij}_\gamma(\ell), \quad (4)$$

where the filter F is to be specified. The correlation function $\xi_+(\theta)$ employs the filter $J_0(\ell\theta)$ and $\xi_-(\theta)$ employs $J_4(\ell\theta)$, which are the zeroth and fourth order Bessel functions respectively. In the manner described in Bartelmann & Schneider (2001); Kilbinger (2015), $\xi_\pm$ are the two correlation functions associated to the spin-2 shear field of weak lensing.

Within the context of determining whether a Gaussian likelihood function is adequate, the filter in Eq. (4) bears a special significance: as shown by Sellentin et al. (2018), it is the low ℓ-modes of the transversed matter fields that skew the distribution from which weak lensing data are drawn. Thus, if we succeed in constructing a filter $F(\ell\theta)$

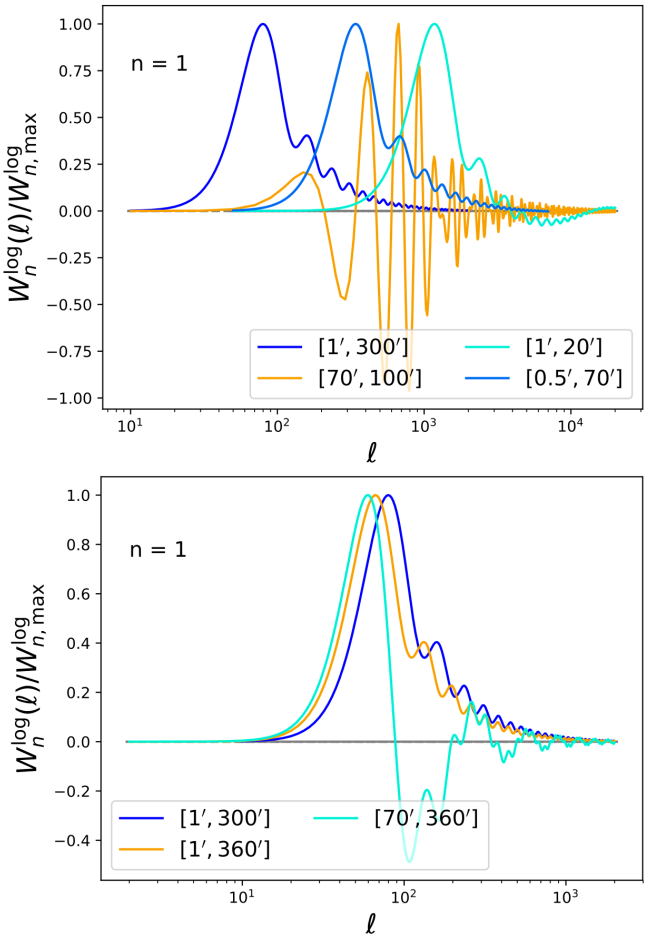

FIG. 1.— The $W_n^{\log}(\ell)$ filters for logarithmic COSEBIs, as a function of mode index $n = 1$ and multipole ℓ. The angular range $[\theta_{\min}, \theta_{\max}]$ is indicated in arc-minutes. It is known that low ℓ-modes cause non-Gaussianities in 2-point function likelihoods. As low n-modes assign large weight to low ℓ-modes, it can be expected that COSEBIs with low n-mode will require non-Gaussian likelihoods.

that downweighs the contribution of low ℓ-modes to a weak lensing observable, then this filter will Gaussianize the distribution.

The downside to filtering out non-Gaussianity is that if it is done naively (e.g. by setting the filter to zero in the low ℓ-regime), then constraining power is lost. We thus have to find a filter that is both efficient in retaining cosmological information, *and* in suppressing non-Gaussianity. In this context, the so-called COSEBI filters proposed by Schneider & Kilbinger (2007); Schneider et al. (2010) are of interest. Originally developed to split weak lensing E-modes from B-modes, and applied as such to the data of the CFHTLenS (Asgari et al. 2017), KiDS and DES survey (Asgari et al. 2019), COSEBIs are not only efficient in retaining cosmological information, but might as well suppress non-Gaussianity.

4. ESSENTIALS OF LOGARITHMIC COSEBIS

Complete Orthogonal Sets of E-/B-mode Integrals, or COSEBIs for short, were originally introduced to split weak lensing E-modes from B-modes (Schneider & Kilbinger 2007; Schneider et al. 2010). The physics of weak lensing is expected to produce E-modes only, and the detection of B-modes would thus directly indicate systematic effects. By discarding B-mode COSEBIs and only analyzing their E-mode counterparts, a clean data vector can be gained and used for parameter inference. However, this clean E-mode data vector still requires the correct likelihood to be known, otherwise the use of the incorrect likelihood becomes a systematic itself.

With the aim of determining the likelihood shape of COSEBIs, we thus study whether the COSEBI filter $W_n^{\log}(\ell)$ succeeds in suppressing non-Gaussianity that skews the likelihood for parameter inference.

We detail the derivation of COSEBIs in appendix A. For the study of the COSEBI likelihood shape, it is sufficient to recognize that the COSEBI filters will act according to Eq. (4): given a power spectrum, the filter lets certain ℓ-modes pass, while others are suppressed, depending on the amplitude that the filter assigns to these modes.

As it is known that the subclass of logarithmic COSEBIs is more efficient in retaining cosmological information, we here specialize to *logarithmic COSEBIs*, which we denote as E_n. COSEBIs are discrete modes, enumerated by their index n, and we expand until $n = 7$, which has been shown to retain the majority of cosmological information (Asgari et al. 2012).

In case of logarithmic COSEBIs, the filter F from Eq. (4) is usually referred to as $W_n^{\log}(\ell)$. Per index n, the filter will output the corresponding COSEBI E_n, implying the filter's shape will change when varying n.

As COSEBIs are constructed to only use data on given angular ranges (see appendix A), the filters $W_n^{\log}(\ell)$ implicitly also depend on $\theta_{\min}$ and $\theta_{\max}$, the minimum and maximum angular range of data that the filter lets pass.

The exclusion of angular ranges becomes more evident when transforming the $W_n^{\log}(\ell)$ filters to real space, where they filter $\xi_\pm(\theta)$ instead of the shear power spectrum $P_\gamma^{ij}(\ell)$. The real space analogues of the $W_n^{\log}(\ell)$ filters are here denoted as $T_{+,n}^{\log}(\theta)$ and $T_{-,n}^{\log}(\theta)$, depending on whether they filter ξ_+ or ξ_-.

We plot harmonic space representations $W_n^{\log}(\ell)$ of selected COSEBI filters in Fig. 1. Fig. 2 displays examples of $T_\pm$ filters for $n \in [1, 7]$ and an angular range of 0.5 to 100 arc-minutes. The angular ranges $[\theta_{\min}, \theta_{\max}]$ that each COSEBI filter integrates over are indicated in the legends of our plots.

From the harmonic space representation $W_n^{\log}(\ell)$ in Fig. 1 it is evident that COSEBI filters let a selected range of ℓ-modes pass. As Sellentin et al. (2018) identified low ℓ-modes to be the cause of likelihood skewness, COSEBIs might thus be able to exclude these problematic modes. In practice, the filter functions displayed in Fig. 2 are applied to correlation functions $\xi_\pm$ estimated from observed galaxy catalogues. In Sect. 5, we will use the $T_\pm$ filters on correlation functions estimated from simulations, with the aim to gain a representative sample of simulated COSEBIs to study the shape of their sampling distribution.

5. CREATING COSEBI SAMPLES

In order to study whether COSEBIs suppress non-Gaussianity in a weak lensing data set, we have to es-

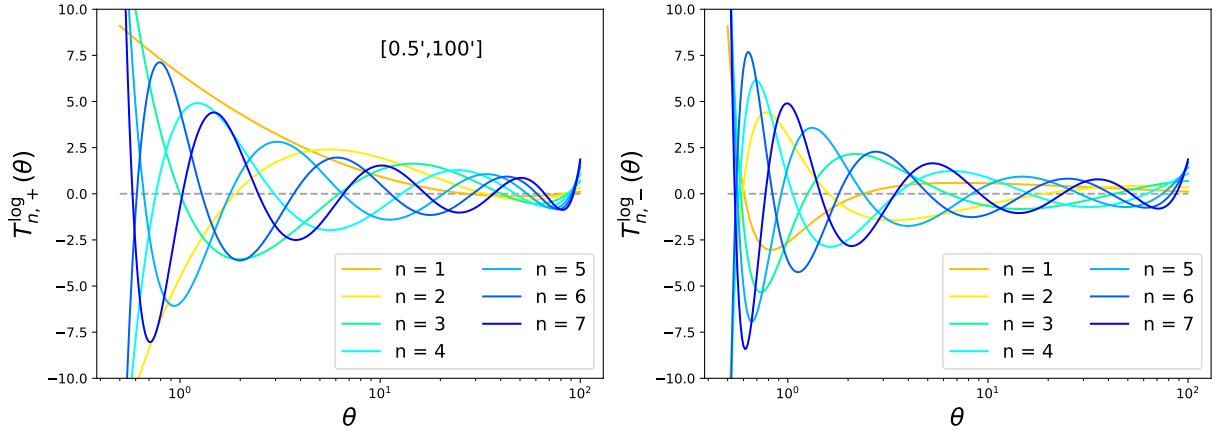

FIG. 2.— The logarithmic filter functions $T^{\log}_{\pm}(\theta)$ as a function of angle for $\theta_{\rm min} = 0.5'$ and $\theta_{\rm max} = 100'$. Filters as these are used to create the first 7 COSEBI modes, E_n and B_n, from shear 2-point correlation functions $\xi_{\pm}(\theta)$.

timate the distribution from which such COSEBIs are drawn. This distribution is called the 'sampling distribution' of the data. In parameter inference the sampling distribution reoccurs as the likelihood, when it will be evaluated conditionally on the parameters. To estimate the sampling distribution, we compute 819 simulated COSEBI data vectors from the SLICS weak lensing simulations (Harnois-Déraps & van Waerbeke 2015) in a Euclid- or LSST-like setup of redshift binning.

The SLICS simulations provide square celestial patches of the cosmic large scale structure of size $A_{\rm sky} = 100$ deg^2 at a mass resolution of $2.88 \cdot 10^9 M_\odot/h$. The simulation suit contains 930 independent realizations of a gravitationally evolved dark matter field projected onto lens planes in order to compute the light deflection and therefrom the weak lensing signal per simulation.

Of these simulations, 819 were post-processed to imitate a 10-bin tomographic weak lensing survey, with $\bar{n} = 2.6/\text{arcmin}^2$ galaxies per bin. This corresponds to a total of 30 galaxies per square-arcminute as expected for a Euclid-like survey.

To compute COSEBIs, we first run the TreeCorr implementation (Jarvis et al. 2004) of the $\hat{\xi}_{\pm}$-estimator given in Eq. (1) on the SLICS weak lensing catalogues. When computing accurate COSEBIs an unusually high angular resolution is of the essence. Following the recommendation of Asgari et al. (2017), we thus use 10000 linear θ-bins, of width $\Delta\theta = 0.00995'$. We then filter the resulting $\hat{\xi}_{\pm}(\theta)$ estimates with the $T^{(n)}_{\pm}(\theta)$ filters according to Eq. (A3). This results in 819 estimates of E- and B-mode COSEBIs E_n and B_n. We retain the E-modes for further statistical analysis, but discard the B-modes as these would not contain physical information in an analysis of real data.

In our computation of simulated data vectors $\hat{\xi}_{\pm}$, we observe the known loss of power in $\hat{\xi}_{\pm}$ as reported in Fig. 6 of Harnois-Déraps et al. (2018). This loss of power in turn affects the amplitude to the COSEBIs computed from $\hat{\xi}_{\pm}$. Fortunately, this is a deterministic rather than a stochastic effect, and is caused by the finite mass resolution and box size of the simulations. As Harnois-Déraps et al. (2018), we hence rescale the extracted $\hat{\xi}_{\pm}$ such that they average to the best-fitting parameters of the original simulation run.

To study shape noise and cosmic variance individually, we estimate pure cosmic-variance $\hat{\xi}_{\pm}(\theta)$ correlation functions from the SLICS simulations. As shape noise is understood to be Gaussianly distributed, we process the pure cosmic-variance $\hat{\xi}_{\pm}$ functions further by adding shape noise directly to $\hat{\xi}_{\pm}$ rather than per galaxy.

The shape noise variance per galaxy splits as $\sigma^2_\epsilon = \sigma^2_{\epsilon_1} + \sigma^2_{\epsilon_2}$ where the indices 1, 2 indicate the two components of the estimated ellipticities. We use the typical standard deviation per component of $\sigma_{\epsilon_i} = 0.29$. Due to $\hat{\xi}_{\pm}$ being computed in angular bins, the shape noise variance $\mathsf{C}_{\rm s}$ of $\hat{\xi}_{\pm}$ is given by

$$\mathsf{C}_{\rm s}(\theta) = \frac{\sigma^2_\epsilon}{A_{\rm sky}\bar{n}^2 2\pi\theta\Delta\theta}, \tag{5}$$

where $A_{\rm sky} = 100$ deg^2 is the area of the simulated celestial patch and $\bar{n}$ the number density of about 2.6 galaxies per arcmin^2 per tomographic bin.

The correlation functions with shape noise then follow to be

$$\hat{\xi}^{\rm SN}_{\pm}(\theta) = \hat{\xi}_{\pm}(\theta) + \hat{s}(\theta), \tag{6}$$

where the random variables $\hat{s}$ are drawn from Gaussian distributions of mean zero and variance $\mathsf{C}_{\rm s}(\theta)$,

$$\hat{s}(\theta) \sim \mathcal{G}\left(0,\ \mathsf{C}_{\rm s}(\theta)\right). \tag{7}$$

Note, that Eq. (6) and Eq. (7) imply that shape noise does not change the signal, but only enhances the scatter around the signal: the quantity $\hat{s}(\theta)$ is a Gaussian random variable, taking with equal probability values above and below its mean.

Both the correlation functions with and without shape noise are then filtered according to Eq. (A3) in order to yield 819 E-mode COSEBIs for each n and each redshift bin combination. With 10 redshift bins, there will be 55 redshift bin combinations. As we compute 7 COSEBIs per redshift bin combination, we thus create 819 samples of a 385 dimensional data vector.

As an example, Fig. 3 depicts all two-dimensional marginal distributions of the first seven COSEBI E-

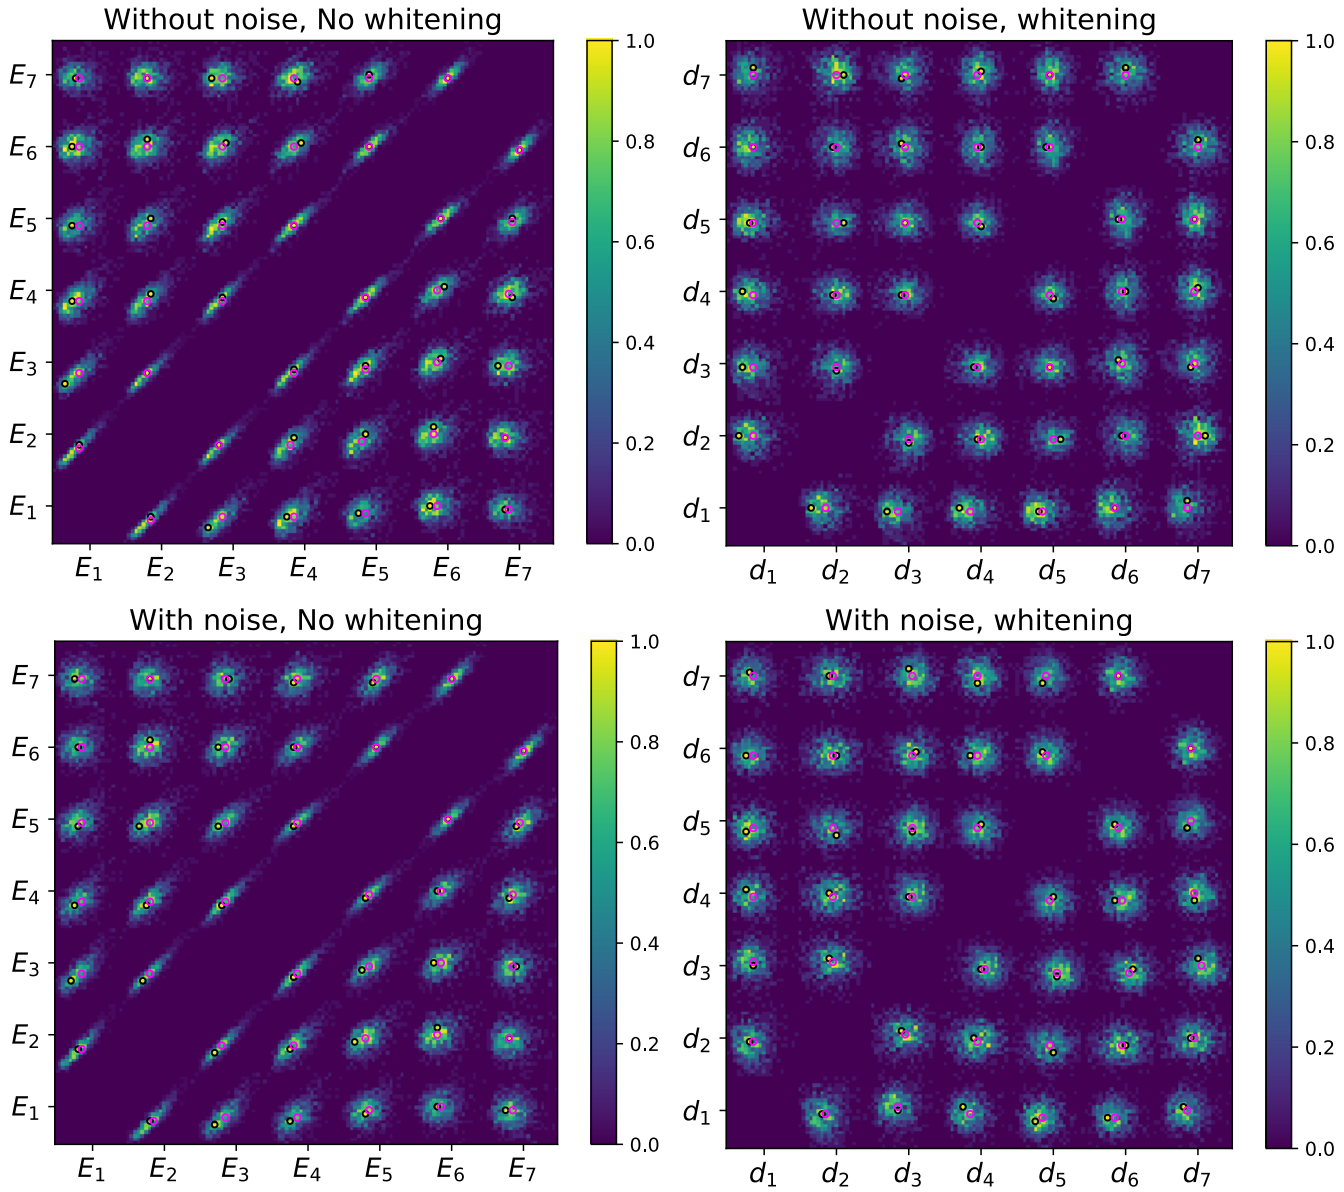

FIG. 3.— Two-dimensional marginal distributions of the COSEBIs in redshift bin 8. The seven E_n modes filtered the angular scale [1',300']. Per marginal, the two E_n modes are indicated on the axes. The black circles mark the peak of the marginals which are often off-center, due to the asymmetry of the distributions; the pink circles mark the means. When Cholesky-whitening the left column, the right column ensues, where each d_i is a linear superposition of the E_n modes. Cholesky-whitening destroys all Gaussian correlations between data points. This is easily seen close to the diagonal where the strongly ellipsoidal (i.e. correlated) COSEBI distributions were transformed to more isotropic, roundish distributions. Non-Gaussian dependencies between random variables survive Cholesky-whitening, which is why the right column still displays asymmetric non-Gaussian features. These residual non-Gaussianities are quantified in Sect. 6.

modes of the redshift bin 8 correlated with itself. This figure can be repeated for the remaining 54 redshift bin combinations, which provides the basis for the non-Gaussianity detection in Sect. 6.

From Fig. 3, we can already deduce by eye that the assumption of a Gaussian distribution for COSEBIs there used will be approximate. This is evidenced by the marginals deviating from two-dimensional Gaussian distributions. It can be seen that the distributions' peaks are often off-center, which affects in particular the COSEBIs E_2 and E_3. Additionally, the marginals including the COSEBI E_6 display a sharp drop in intensity, causing an edge in the depicted histograms. These first visual impressions of non-Gaussianity will be quantified in Sect. 6.

6. NON-GAUSSIANITY DETECTION STRATEGY

The purpose of this section is to quantify the level of non-Gaussianity in the high-dimensional data set of seven COSEBIs per 55 redshift bin combinations of a Euclid- or LSST-like survey. For detected non-Gaussianities, a follow-up study of the ensuing biases is presented in Sect. 8.

A covariance matrix is sufficient if there exist only Gaussian correlations in pairs of data points. The components of the covariance matrix are estimated via the usual sample estimator

$$C_{mn} = \frac{1}{N-1}\sum_{i=1}^{N}(E_m^i - \bar{E}_m)(E_n^i - \bar{E}_n), \qquad (8)$$

which uses the sample mean

$$\bar{E}_n = \frac{1}{N}\sum_{i=1}^{N} E_n^i. \tag{9}$$

If the likelihood truly is Gaussian, then computing an accurate covariance matrix is the final stepping stone for an unbiased analysis. In contrast, if the likelihood is non-Gaussian, then even an arbitrarily precise covariance matrix is insufficient to parameterize the actual likelihood. The purpose of a *trans-covariance matrix*, as developed by Sellentin & Heavens (2018), is to flag where in a data set computing a covariance matrix is insufficient. If the likelihood truly is Gaussian, then the trans-covariance matrix will vanish. Otherwise, if non-Gaussian dependencies affect the data, then the trans-covariance matrix will flag pairs of data points for which more than their covariance must be understood to yield the correct likelihood.

The logic of a trans-covariance matrix is to first destroy the Gaussian correlations in pairs of data points, and then measure whether any further non-Gaussian dependencies between the two data points remain. If $\boldsymbol{d} = (d_a, d_b)$ is a pair of two data points taken from a usually much larger data vector $\boldsymbol{x}$, then the covariance matrix of this pair is given by

$$\mathsf{C} = \begin{pmatrix} \mathrm{var}(d_a) & \mathrm{cov}(d_a, d_b) \\ \mathrm{cov}(d_a, d_b) & \mathrm{var}(d_b) \end{pmatrix}. \tag{10}$$

This covariance matrix is Cholesky decomposed as $\mathsf{C} = \mathsf{L}\mathsf{L}^T$, and the pair $\boldsymbol{d}$ is whitened via

$$\tilde{\boldsymbol{d}} = \mathsf{L}^T \boldsymbol{d}. \tag{11}$$

If the covariance matrix of Eq. (10) fully captures the statistical dependence between d_a and d_b, then the whitened vector $\tilde{\boldsymbol{d}}$ will now consist of two independent normally distributed random variables. This is testable. If the test is failed, non-Gaussian dependencies between the pair exist.

The test for non-Gaussian dependencies must be repeated for every pair of data points. This repetition builds up a matrix: just like the covariance matrix computes the covariance between a pair, the trans-covariance matrix flags the non-Gaussian dependencies in a pair. Plotting the two matrices side by side directly reveals for which pairs the covariance matrix is an all-encompassing statistical description, and which pairs require an advanced statistical treatment.

Our test for non-Gaussian dependencies in a pair proceeds as follows. Assume $\tilde{d}_a, \tilde{d}_b$ indeed were whitened normal distributed variables. Then

$$\tilde{d}_a \sim \mathcal{G}(0,1), \quad \tilde{d}_b \sim \mathcal{G}(0,1), \tag{12}$$

would hold. A direct consequence of this assumption is that the sum of $\tilde{d}_a, \tilde{d}_b$ follows

$$\tilde{d}_a + \tilde{d}_b \sim \mathcal{G}(0,2). \tag{13}$$

If there exist non-Gaussian dependencies between d_a and d_b, then the assumption that $\tilde{d}_a, \tilde{d}_b$ are independently drawn from the two normal distributions in Eq. (12) breaks down, and Eq. (13) will be incorrect as a consequence. From our simulated COSEBIs, we can create a histogram of the sum in Eq. (13). If this histogram tends to a Gaussian with variance two, then the pair d_a and d_b is fully described by its covariance matrix. If the histogram deviates from a Gaussian of variance two, then d_a and d_b depend on each other in ways not captured by a Gaussian likelihood.

Denoting the histogram as $\mathcal{H}(\tilde{d}_a + \tilde{d}_b)$, its deviance from $\mathcal{G}(0,2)$ is computed by integrating over the difference between the histogram and the Gaussian. If the histogram consists of W rectangular bins, then this integral is simply given by the sum

$$S_{ab} = \frac{1}{W}\sum_{w=1}^{W}\left[\mathcal{H}_w\left(\tilde{d}_a + \tilde{d}_b\right) - \mathcal{G}\left(\tilde{d}_a + \tilde{d}_b|0,2\right)\right]^2, \tag{14}$$

where $\mathcal{H}_w$ is the wth histogram bin. The sum S_{ab} provides the (a,b) component of the trans-covariance matrix. The corresponding components of the covariance matrix are given in Eq. (8).

Finally, we have to mitigate the impact of having only 819 simulations, which will cause noise in the trans-covariance matrices. We thus compute trans-covariance matrices for 819 Gaussian data sets, and measure the mean and standard deviation of their matrix elements. We then mean-subtract the trans-covariance matrices of the actual COSEBI simulations, and divide by the standard deviation. This yields a trans-covariance matrix in units of a standard deviation, such that we can speak of a '5σ' detection, etc. The units σ are thus assigned to the colour bars of all plotted trans-covariance matrices.

7. RESULTS: BIASES AND SKEWNESS IN COSEBIS

In this section we present the main findings of this paper, which include the scaling of biases from skewness with redshift, ℓ-mode, and angular range $\theta_{\min}, \theta_{\max}$ in a weak lensing data set.

7.1. *Increase of skewness with redshift*

In Fig. 4 we depict the correlation matrices and the trans-covariance matrices for a Euclid- or LSST-like survey with 10 tomographic redshift bins. The redshift bins 1-4 of this simulated survey fall below a redshift of unity. The redshift bin 5 marks the transition to a redshift of unity, and the redshift bins 6-10 finally include galaxies above a unit redshift. For a contemporary survey such as KiDS and DES, the vast majority of their observed galaxies used for weak lensing fall below a unit redshift, see e.g. van Uitert et al. (2018); Hildebrandt et al. (2020); Abbott et al. (2019). Accordingly, these surveys can roughly be recognized in the first four redshift bins of Fig. 4.

Fig. 4 displays the COSEBI modes E_1 to E_7 per redshift bin combination. The angular range of these COSEBIs is $\theta_{\min} = 1'$ and $\theta_{\max} = 300'$.

From the correlation matrices with and without shape noise we recognize the familiar trend that weak lensing observations in different redshift bins are strongly correlated. The trans-covariance matrix without shape noise (lower left) reveals a 14σ to 21σ detection of a non-Gaussian sampling distribution for COSEBIs at low redshift. In the case at hand, this non-Gaussianity takes the form of a skewed sampling distribution, as is expected for any 2-point function, see Sellentin et al. (2018) and refer-

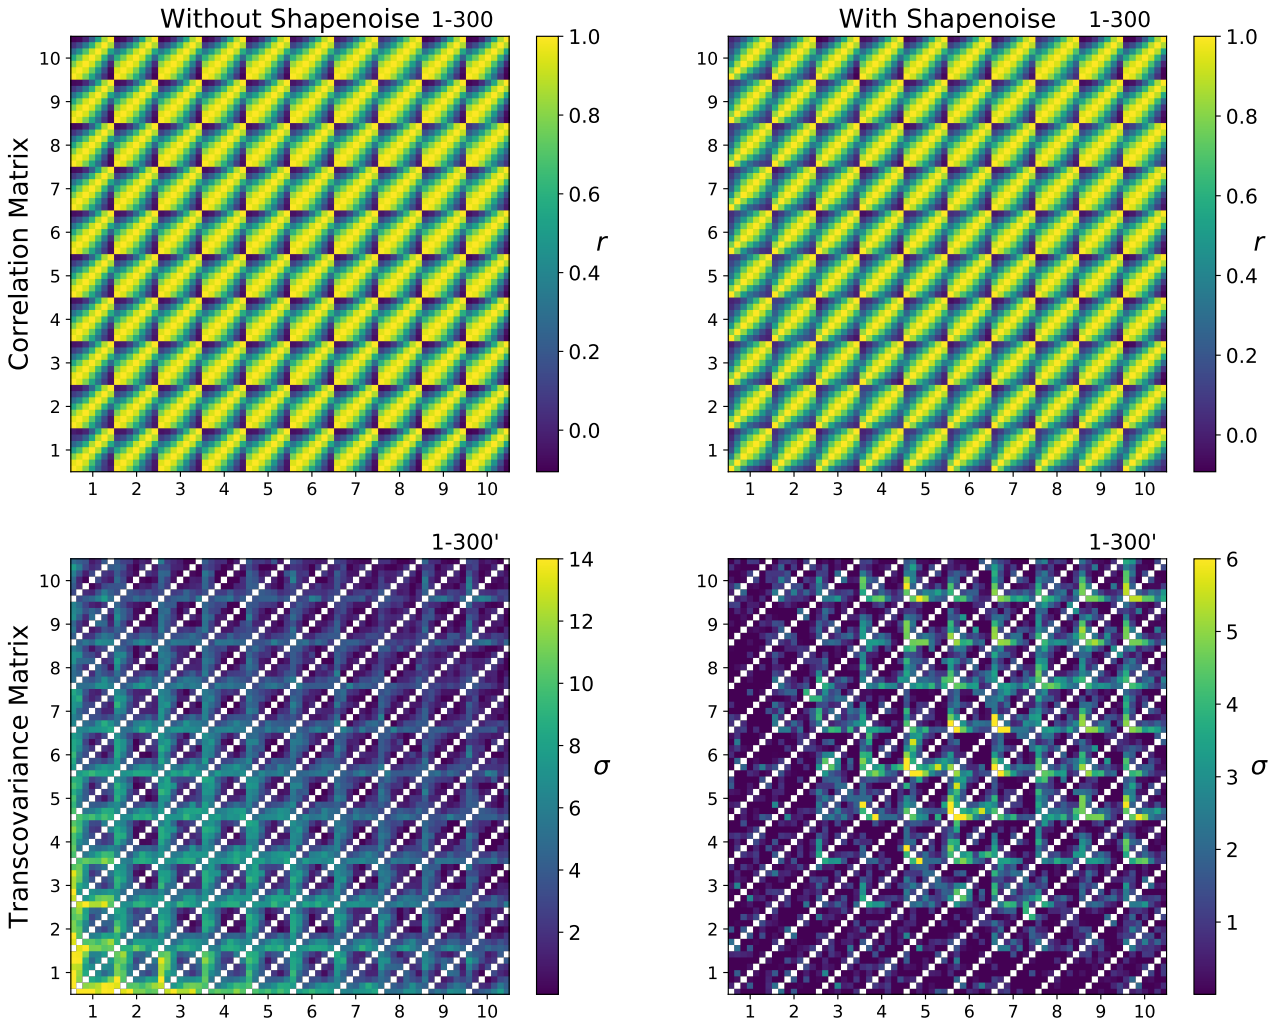

FIG. 4.— **Description:** The correlation- and trans-covariance matrices of the seven E_n COSEBIs for all 55 redshift bin combinations of a 10-bin survey. These COSEBIs were computed by filtering $\hat{\xi}_{\pm}(\theta)$ binned in 250 angular bins. The axis labels indicate the redshift bins; within each redshift bin the mode index n increases from 1 (left) to 7 (right). The left column singles out cosmic variance; the right column also adds shape noise. The angular ranges $[\theta_{\rm min}, \theta_{\rm max}]$ of the COSEBIs are indicated. **Interpretation:** While all COSEBIs are strongly correlated as seen in the upper row, only selected COSEBIs show additional non-Gaussian dependencies. Without shape noise, non-Gaussianity is detected in the low redshift bins with up to 21σ significance. For visibility, the colour bar is cut at 14σ. With shape noise, a redshift trend for non-Gaussianity emerges: While the redshift bins 1-4 are unaffected, higher bins show a 6σ detection of non-Gaussianity. The redshift bins 1-4 all fall below a redshift of unity, bin 5 marks the transition, and from bin 6 onwards the binned galaxies reside at redshifts $z > 1$. The exact redshift bin boundaries here used are inherited from the SLICS simulations and are stated in Fig. A1 of (Harnois-Déraps et al. 2018). They imitate the standard 10 bin setup of a Euclid-like survey. While the bulk of galaxies of contemporary surveys is located below a redshift of unity, half of the galaxies in a Euclid- or LSST-like survey will reside at redshifts beyond unity. This implies the next generation of surveys will be affected by the found skewness.

ences therein. This shape-noise-free panel is included to show the upper bound on non-Gaussianity, as it would occur in a survey of many more galaxies than detectable by Euclid or LSST.

The lower right trans-covariance matrix adds shape noise for a Euclid- or LSST-like survey. Shape noise follows a Gaussian distribution and convolves the originally strongly skewed distributions that gave rise to the 14 to 21σ detection on the left. The convolution results in a suppression of non-Gaussianity, and accordingly the trans-covariance matrix with shape noise flags a 6σ detection of residual skewness, which is still considerable. All pairs of data points that are flagged as subject to non-Gaussian dependencies will require a non-Gaussian likelihood for bias-free inference.

The perhaps most important result of Fig. 4 is the evidence that a COSEBI-based KiDS and DES analysis on the angular scale $[1', 300']$ is free from non-Gaussian biases. This conclusion is evidenced by the first four redshift bins not displaying any prominent non-Gaussianity detection. In contrast, for a Euclid- or LSST-like survey, half of the galaxies reside at redshifts above unity, and will thus be affected by the skewness of sampling distribution and likelihood here detected.

The computation of Fig. 4 was only feasible by binning the $\hat{\xi}_{\pm}(\theta)$ functions in 250 angular bins before applying the COSEBI filters $T_{\pm}(\theta)$. While 250 angular bins in $\hat{\xi}_{\pm}$ are insufficient to yield a percent-level accuracy of the *mean* COSEBI signal (see Asgari et al. (2017)), we find that the COSEBI trans-covariance matrix con-

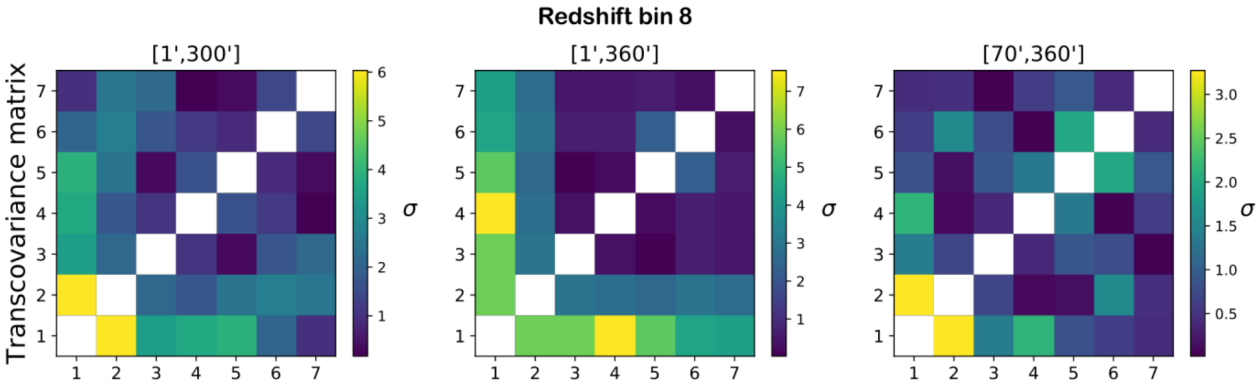

FIG. 5.— **Description:** Three magnifications of redshift bin combination 8-8, recomputed from $\hat{\xi}_{\pm}(\theta)$ with 10000 θ-bins to show that the results with 250 θ-bins are accurate. The axis labels enumerate the seven COSEBI modes. The angular ranges $[\theta_{\rm min}, \theta_{\rm max}]$ of the COSEBIs are indicated per panel. **Interpretation:** The strength of non-Gaussianity indeed depends on the ℓ-range the W-filter assigns large weight. Comparing to Fig. 1, the [1',300']-filter amplifies ℓ-modes around $\ell \approx 100$ for which non-Gaussian sampling distributions are identified with 6σ significance (left). The [1',360']-filter emphasizes slightly lower ℓ-modes causing a non-Gaussianity detection significance of more than 7σ (middle). The [70',100']-filter changes sign, thereby balancing positive and negative skewness against each other; this results in a suppression of non-Gaussianity that is consequently only detected at slightly more than 3σ (right).

verges much more rapidly than the mean signal. To evidence that 250 angular bins were sufficient, we display in Fig. 5 a recomputation of the trans-covariance matrix of redshift bin 8, but now with 10000 angular bins. This yields precisely the same detection significance of the found non-Gaussianities. Further recomputations of the first three redshift bins at angular resolution of 10000 angular bins are presented in Sect. 7.3.

7.2. *ℓ-ranges: Excess skewness in comparison to the CMB*

The sampling distribution (and in consequence the likelihood) of any 2-point function is a priori expected to be skewed. This arises due to 2-point functions being quadratic in the underlying field values that they correlate, and in consequence even affects the power spectrum estimates C_ℓ of the cosmic microwave background (CMB), although the microwave background itself is compatible with a Gaussian random field (Hamimeche & Lewis 2008, 2009).

Known from analyses of the CMB is that the sampling distribution of C_ℓ is skewed until $\ell \approx 20$, see e.g. the discussion in the Planck likelihood of Planck Collaboration et al. (2019). As COSEBI filters allow us to single out certain ℓ-ranges with more precision than could be achieved with a $\xi_{\pm}(\theta)$ analysis, we are able to compare likelihood skewness in weak lensing surveys to likelihood skewness in CMB analysis. This study is depicted in Fig. 5.

All three panels of Fig. 5 depict the trans-covariance matrix of the COSEBI modes E_1 to E_7 for redshift bin 8. These COSEBIs were computed by filtering $\xi_{\pm}$ binned in 10000 angular bins, as described in appendix A. From left to right, the angular range of the COSEBIs was varied, to the effect of singling out different ℓ-ranges. The left trans-covariance matrix reveals a 6σ detection of skewness on the angular range $[1', 300']$. The corresponding $W_n^{\log}(\ell)$ filter has been displayed in blue in Fig. 1. Evidently, this filter emphasizes ℓ-modes in the range $\ell \approx 80 - 150$, and continues to assign significant amplitude to ℓ-modes at $\ell \approx 200$. Weak lensing is therefore subject to likelihood skewness over an ℓ-range roughly ten times larger than for the CMB.

The middle panel of Fig. 5 depicts the trans-covariance matrix for $\theta_{\rm min} = 1'$ and $\theta_{\rm max} = 360'$. The corresponding $W_n^{\log}(\ell)$ filter is depicted in yellow in the lower panel of Fig. 1. This filter emphasizes ℓ-modes slightly below those of the COSEBIs in the left panel. Accordingly, non-Gaussianities are now detected with more than 7σ significance, rather than 6σ as in the case of the previously discussed $[1', 300']$ COSEBIs.

The right-most trans-covariance matrix of Fig. 5 employs the angular range $[70', 360']$, whose $W_n^{\log}(\ell)$ filter is depicted in green in Fig. 1. This filter also assigns a large amplitude on ℓ-modes around $\ell \approx 100$, but additionally shows a reversal of sign. As neighbouring $\mathrm{P}_\gamma(\ell)$-estimates are highly correlated in weak lensing, this filter will thus sum over nearly identical $\mathrm{P}_\gamma(\ell)$ values which it weights with positive and negative signs. As the skewness of a distribution is reversed when the sign of the random variable is reversed, this leads to a cancellation of the skewness during the summation. Accordingly, the right-most trans-covariance matrix of Fig. 5 only displays a 3.2σ detection of residual skewness.

7.3. *Scaling with angular range*

Euclid and LSST will not only compile galaxy catalogues that reach deeper into the cosmic past than those of KiDS and DES, but will additionally cover a larger celestial area. Accordingly, it is natural that KiDS and DES limit their analysis of weak lensing correlation functions to smaller angular ranges than we discussed in Sect. 7.1.

In Fig. 6 we therefore depict trans-covariance matrices for KiDS- and DES-like redshift ranges and angular ranges. The depicted COSEBIs are E_1 to E_7, computed for the first three redshift bins of Fig. 4. The upper limit of these redshift bins is $z \approx 1$. We depict three sets of COSEBIs, with those of the left column spanning the angular range $[1', 20']$, the middle column using $[0.5', 70']$ and the right-most column $[70', 100']$. All these COSEBIs were computed from $\hat{\xi}_{\pm}(\theta)$ binned in 10000 angular bins.

The upper row of Fig. 6 reveals that these low redshifts

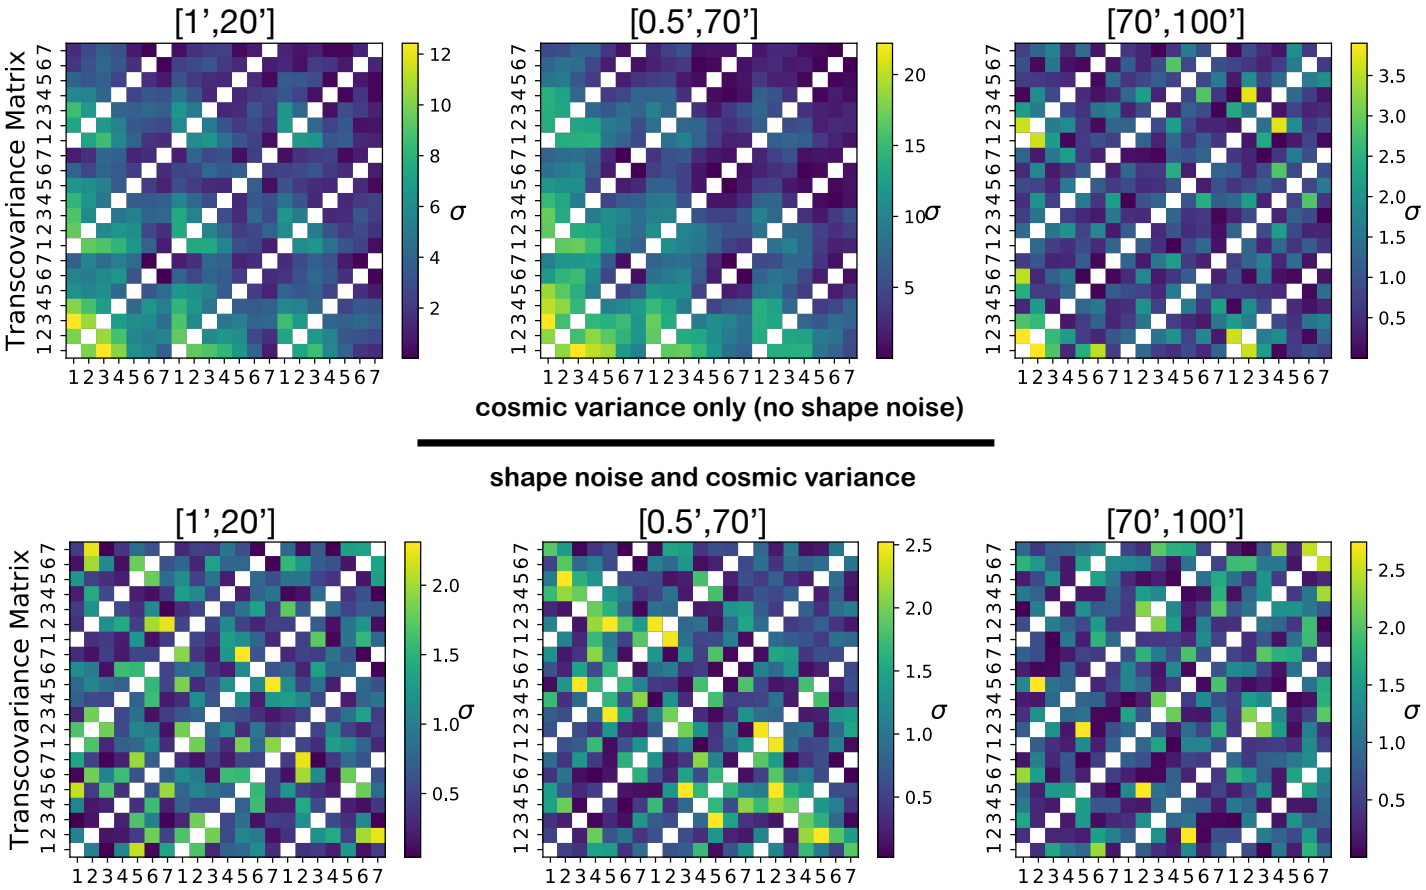

FIG. 6.— The trans-covariance matrices of the COSEBI E-modes on an angular scale of 0.5' - 20' (left), 0.5'-70' (middle), 70'-100' (right) for redshift bins 1 to 3 (and their cross-correlations) of a Euclid- or LSST-like survey. These three redshift bins extend until $z \approx 1$ and cover the bulk of KiDS- and DES-like galaxy catalogues. The top row depicts E-modes without shape noise, the lower row includes it. The comparison reveals that although the KiDS- and DES-like redshift ranges are affected by non-Gaussianity from cosmic structure formation, it is the addition of shape noise that renders this skewness negligible. This is evidenced by the up to 18σ detections of skewness without shape noise (middle top panel), which is suppressed as soon as shape noise is added. Accordingly, the lower row shows no significant detection of skewness throughout the depicted redshift ranges.

are subject to strong non-Gaussianity when cosmic variance alone is regarded. The skewness of the sampling distribution of COSEBIs with $\theta_{\rm min} = 1'$ and $\theta_{\rm max} = 20'$ is detected with more than 12σ significance. For the angular range $[0.5', 70']$ the skewness of the COSEBI distribution is even detected with 21σ significance. On angular ranges $[70', 100']$ small clusters of 3.5σ significance appear, but otherwise the sampling distribution of these COSEBIs is compatible with a Gaussian.

The lower row of Fig. 6 reveals that the addition of shape noise suppresses the skewness due to cosmic variance efficiently: with shape noise, none of of the COSEBIs display any significant detection of skewness anymore.

8. EXCESS PROBABILITY OF FINDING 'LOW' WEAK LENSING MEASUREMENTS

As known from Sellentin et al. (2018), data points for $\xi_+(\theta)$ on angular scales of 100 arcminutes and beyond are expected to be biased low by 0.3 standard deviations in the low redshift bins of a Euclid- or LSST-like survey. This bias was shown to exacerbate with redshift. ξ_- is similarly biased low, but to a lesser degree.

COSEBIs form linear combinations of $\xi_+(\theta)$ and $\xi_-(\theta)$. If the COSEBI filter is positive everywhere, any COSEBI computed from low $\xi_\pm$ will in turn be biased low itself. Only when the COSEBI filter changes sign can it be expected that biases are suppressed, see the discussion of Fig. 5 in Sect. 7.2. With Euclid and LSST set up to target large angular scales, we therefore investigate the biases of COSEBIs on the [1',300'] scale, and set these biases into perspective with the signal-to-noise contributed by each COSEBI mode.

In Fig. 7 we plot the excess probability ratio of COSEBI modes taking values below their mean. This 'excess' is rated in comparison to a Gaussian. If COSEBI data points were drawn from a Gaussian distribution, they are equally likely to take values below or above their mean.

To quantify the excess probability ratio of COSEBIs falling below the mean, we analyze all our COSEBI samples directly, without binning in a histogram. For each COSEBI mode E_n, we compute the fraction of samples that falls below and above the mean $\bar{E}_n$. This estimates the two probabilities $\mathcal{P}(E_n < \bar{E}_n)$ and $\mathcal{P}(E_n > \bar{E}_n)$. For a Gaussian, the ratio of these probabilities will equal unity, with some small scatter due to estimating the probability from 819 samples.

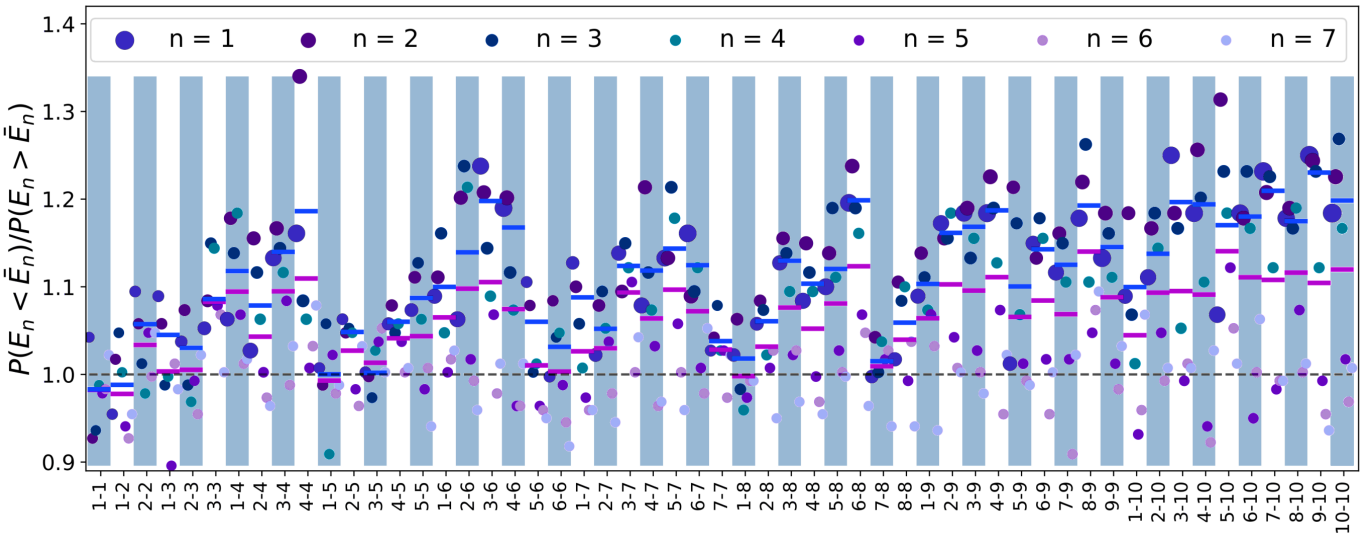

FIG. 7.— Excess probability of drawing a 'low' weak lensing data vector in a 10-bin tomographic survey when insisting on a Gaussian likelihood. All COSEBIs here used span the angular range [1',300'], and include cosmic variance and shape noise. Plotted is the probability ratio of a COSEBI mode falling below or above its mean. For a symmetric Gaussian distribution, this ratio equals unity. The grey and white bands mark the different redshift bin combinations as indicated on the x-axis. The colourful points per band mark the probability ratio of the individual E_n modes per bin. As can be seen, in particular E_1 often has an excess probability of 20-25% of falling below its mean which is unexpected by a Gaussian likelihood. Moreover, the marker size of each mode is proportional to its signal-to-noise ratio (SNR), revealing that the COSEBIs contributing the highest SNR are most prominently biased low. The blue (purple) bars indicate the SNR-weighted (unweighted) average probability ratio of all modes within one redshift bin being biased low. This probability of being biased low increases whenever at least one redshift bin lies at large redshift, causing the visible sinusoidal pattern. Typical unweighted probabilities of being biased low amount to 10%. Weighting by the SNR, it can be seen that in 20-25% of all cases, the signal will be systematically lower than expected from a Gaussian likelihood. Roughly a quarter of a Euclid- or LSST-sized data set will thus be affected.

We define the ratio of these probabilities as

$$e_n = \frac{\mathcal{P}(E_n < \bar{E}_n)}{\mathcal{P}(E_n > \bar{E}_n)}, \tag{15}$$

which is the 'excess' of COSEBIs taking values below their mean.

Fig. 7 reveals that this excess probability for the COSEBIs deviates from unity. We observe that in particular the low COSEBI modes E_1 to E_4 are at least 10% more likely to fall below the mean than expected from a Gaussian. However, in particular for high redshifts, this excess can amount to 20-25%, in rare cases even exceeding 30%.

Assuming all COSEBI modes contributed with equal weight to the signal in a redshift bin, we compute the arithmetic mean per redshift bin

$$\bar{e}_n = \frac{1}{7}\sum_{i=1}^{7} e_n. \tag{16}$$

This average excess probability ratio of all COSEBIs per bin falling below their mean is indicated as purple horizontal bars in Fig. 7. A sinusoidal pattern of this averaged excess can be seen, which originates from the ordering of redshift bin combinations: combinations with at least one bin at high redshift are systematically more likely to yield low values for COSEBIs. In comparison to a Gaussian, Fig. 7 indicates typically a 10% excess of yielding a low weak lensing data vector, if all seven COSEBIs are assigned equal weights while averaging.

In reality however, the different COSEBI modes contribute a different signal-to-noise to the entire measurement. Computing an unweighted mean therefore neglects that certain COSEBIs are more influential for the final inference than others. To refine the study, we therefore proceed by computing the signal-to-noise per COSEBI, and then compute the signal-to-noise weighted expectation for low COSEBI amplitudes.

Fig. 8 depicts the signal-to-noise ratio for the different COSEBI modes in a 10-bin tomographic survey. In Fig. 7, the marker size of each COSEBI mode is proportional to its signal-to-noise. We define the signal-to-noise ratio as

$$\mathrm{SNR}(n) = \frac{\bar{E}_n}{\sigma(E_n)}, \tag{17}$$

where $\bar{E}_n$ is the mean of a COSEBI mode, and $\sigma(E_n)$ its standard deviation. In Fig. 8, the seven modes used in this study are indicated along the horizontal axes, whereas the redshift bin combinations are indicated on the vertical axis. As weak lensing integrates along the line of sight, the signal-to-noise increases with redshift. Also visible is that the information content of COSEBIs saturates with increasing n, implying that also within a single redshift bin combination the different COSEBI modes will have unequal impact on the total signal per bin.

To account for this inequality of modes, we compute the weighted average excess probability

$$\bar{e}_n^{\mathrm{w}} = \frac{\sum_{i=1}^{7}\mathrm{SNR}(n)e_n}{\sum_{i=1}^{7}\mathrm{SNR}(n)}, \tag{18}$$

which weighs modes according to their signal-to-noise. These weighted excess probability ratios for the weak lensing signal per bin to fall below its mean are marked by blue bars in Fig. 7. As the modes with highest sig-

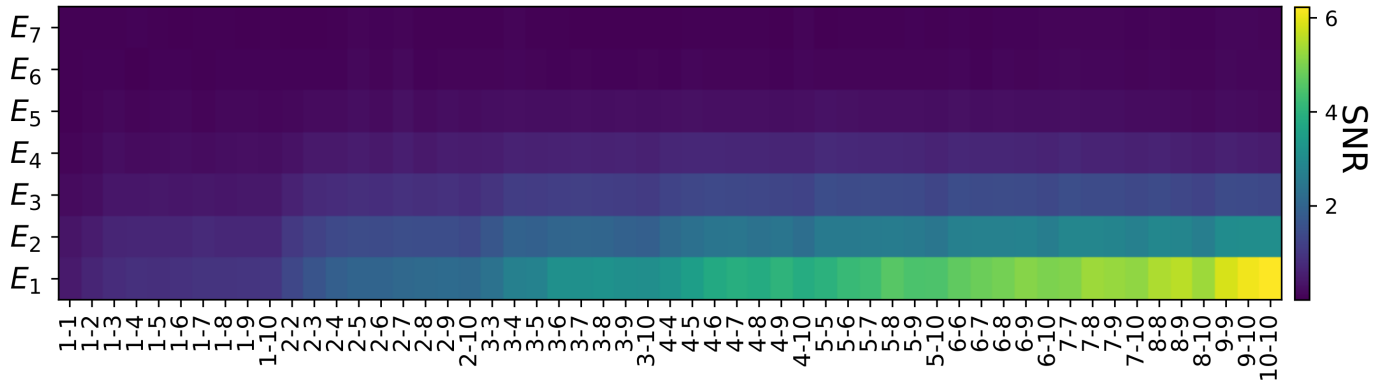

FIG. 8.— The signal-to-noise that each E_n COSEBI contributes in a 10-bin tomographic survey with shape noise. Indicated along the horizontal axis are the redshift bin combinations, along the vertical axis the index n of the COSEBI E-modes. As expected from weak lensing being an integral effect, the highest signal-to-noise is carried by the highest redshift bin combinations. With increasing n, the COSEBI modes contribute ever less signal, motivating the truncation at $n = 7$.

nal to noise are precisely those that are the most biased, the weighted average for all COSEBIs taking values below their mean is larger than its unweighted equivalent: Fig. 7 reveals a 15-20% excess probability ratio of all modes taking 'surprisingly low' values, if regarded with a Gaussian expectation. This trend exacerbates for high redshifts.

In total, the implications for a Euclid- or LSST-like survey are clear: given Fig. 7 there exists the choice between removing the modes of highest signal-to-noise from the data set in order to debias — or a non-Gaussian likelihood is required for a fully bias-free inference. As both surveys are conceived to yield a precision measurement of weak lensing, the removal of the highest signal-to-noise modes does not seem to be the natural solution.

9. CONCLUSIONS

The sampling distribution — and in consequence the likelihood — of any 2-point function is a priori expected to be skewed. This has been implemented in likelihoods of the CMB, see e.g. Planck Collaboration et al. (2019), and while Manrique-Yus & Sellentin (2020) developed a skewed core-likelihood for weak lensing, galaxy clustering and their cross-correlation, the standard in weak lensing analyses still employs a symmetric Gaussian likelihood. Skewness in a sampling distribution implies the measurement gained most likely will systematically differ from the average measurement. As already shown in Sellentin et al. (2018), this affects weak lensing, and both the magnitude and direction of this bias are similar to the lowness of σ_8 as in comparison to Planck's measurements.

From a somewhat misplaced appeal to the Central Limit Theorem, one could hope that data compression reduces the bias arising from skewness. The hope therein is that large and small values of data points cancel during compression, thereby yielding a compressed data vector that scatters more tightly around its mean. This hope is somewhat misplaced, as the Central Limit Theorem implies that (potentially weighted) sums over random variables will tend towards Gaussianly distributed variables. This is still consistent with the statement that if all random variables happen to be 'low', then their sum will equally be 'low'.

Cosmological data compression often takes the form of filtering, where the scalar product with the filter causes the weighted averaging. Accordingly, if the original signal is low, the filtered signal will also be low, unless the filter has specifically been constructed to cancel the low values. This argument holds for all filter-based compression techniques (Heavens et al. 2000, 2017, 2020; Alsing & Wandelt 2018) and in weak lensing also for COSEBI filters (Schneider & Kilbinger 2007; Schneider et al. 2010).

COSEBIs were originally designed to split weak lensing E- and B-modes. A priori, they are thus not made to suppress biases from skewness, but could coincidentally do so. Accordingly, this paper studied whether COSEBIs suppress biases from sampling distribution skewness.

In Sect. 7.1 we focussed on a Euclid- or LSST-like 10 bin tomoraphic weak lensing survey and showed that the skewness of a given COSEBI mode's sampling distribution increases with redshift. Up to a unit redshift, which comprises the majority of all galaxies observed by KiDS and DES, no significant skewness is detected, implying a COSEBI analysis of the current surveys may use a Gaussian likelihood. In contrast, roughly half of the galaxies observed with LSST or Euclid will fall into redshift ranges where the skewness of the COSEBI distribution is detected at 6σ significance.

In Sect. 7.2 we provide evidence that weak lensing observables are drawn from a much more skewed distribution than CMB observations of the same multipole. While the CMB likelihood for spherical harmonics power spectra C_ℓ Gaussianizes at $\ell \approx 20$, the weak lensing likelihood remains skewed until $\ell \approx 100$.

The skewness of 2-point function sampling distributions implies that if the data are analyzed with a symmetric Gaussian likelihood, then one will often find 'surprisingly low' amplitudes. This arises since the Gaussian misestimates how frequently the data will take values below their mean. In Sect. 8 we showed that if a Gaussian likelihood is insisted upon for a Euclid- or LSST-like survey, then in 10% of all cases the data points will appear to be systematically low. Particularly affected are those COSEBI modes which contribute the highest signal-to-noise, such that in the high signal-to-noise regime systematically low amplitudes are to be expected in up to 25% of the cases. For modes affected by high shape noise, this can be understood intuitively: as shape noise adds Gaussian noise, a mode whose SNR is suppressed

by shape noise is automatically a mode whose scatter is dominated by Gaussian noise. Hence these modes will show little bias, as the Gaussian shape noise will cause dominating symmetric noise. In total, we conclude that the next generation of surveys has the choice between lowering the signal-to-noise by discarding the modes which are affected by skewness-induced biases, or to retain these modes and develop a non-Gaussian likelihood for their analysis.

The skewness here discussed is known from analyses of the cosmic microwave background; as such, it is fundamentally different from gravity-induced non-Gaussianity. The interested reader may find a further discussion of gravitationally induced non-Gaussianity in weak lensing in appendix B.

APPENDIX

COMPUTATION OF COSEBIS

In this appendix we detail the construction of COSEBIs, following Schneider et al. (2010); Schneider & Kilbinger (2007).

COSEBIs were initially introduced in order to split E- and B-modes in a weak lensing survey. Their filters are polynomials that are linear or logarithmic in the angular scale θ. As linear COSEBIs have been shown to be significantly less efficient in retaining cosmological information (Schneider et al. 2010), we directly used logarithmic COSEBIs in this paper.

Unlike a correlation function that is a continuous function of angular separation θ, COSEBIs are set of discrete modes, labelled by an index n. This index is accordingly indicated in all relevant figures. We denote modes that only contain E-type power as E_n, and COSEBI B-modes by B_n. If a cosmology of cold dark matter (CDM), and a non-standard dark energy equation of state w is to be studied, then expanding COSEBI modes until $n = 7$ is sufficient (Asgari et al. 2012).

To derive COSEBIs via their property of splitting E- and B-modes, we introduce an E-mode power spectrum $P_E^{ij}(\ell)$, and a B-mode power spectrum $P_B^{ij}(\ell)$, where ij determines the tomographic bin combination. Weak lensing cannot induce a B-mode power spectrum, and it is likewise hoped that no systematics contribute to the E-mode power spectrum. In this case it then follows that $P_E(\ell)$ equals Eq. (3).

In terms of both $P_E^{ij}(\ell)$ and a putative $P_B^{ij}(\ell)$, the weak lensing correlations functions $\xi_\pm$ are given by

$$\xi_+^{ij}(\theta) = \int_0^\infty \frac{\mathrm{d}\ell\ell}{2\pi} \mathrm{J}_0(\theta\ell) \left[\mathrm{P}_E^{ij}(\ell) + \mathrm{P}_B^{ij}(\ell)\right], \tag{A1}$$

and

$$\xi_-^{ij}(\theta) = \int_0^\infty \frac{\mathrm{d}\ell\ell}{2\pi} \mathrm{J}_4(\theta\ell) \left[\mathrm{P}_E^{ij}(\ell) - \mathrm{P}_B^{ij}(\ell)\right], \tag{A2}$$

implying they will mix E- and B-modes, if B-modes are present (Schneider & Kilbinger 2007). To create a linear filter that splits E- from B-modes, the ansatz

$$\begin{aligned} E_n^{(ij)} &= \frac{1}{2}\int_0^\infty \mathrm{d}\theta\theta\Big[T_+^{(n)}(\theta)\xi_+^{(ij)}(\theta) + T_-^{(n)}(\theta)\xi_-^{(ij)}(\theta)\Big], \\ B_n^{(ij)} &= \frac{1}{2}\int_0^\infty \mathrm{d}\theta\theta\Big[T_+^{(n)}(\theta)\xi_+^{(ij)}(\theta) - T_-^{(n)}(\theta)\xi_-^{(ij)}(\theta)\Big], \end{aligned} \tag{A3}$$

is chosen. The filter functions $T_\pm^{(n)}(\theta)$ thus depend on the index n which determines the order of the COSEBI mode.

To construct a filter which splits E- from B-modes, we transform the filters $T_\pm^{(n)}(\theta)$ to harmonic space

$$\begin{aligned} W_+(\ell) &= \int_0^\infty \mathrm{d}\theta\ \theta T_+^{(n)}(\theta) J_0(\ell\theta), \\ W_-(\ell) &= \int_0^\infty \mathrm{d}\theta\ \theta T_-^{(n)}(\theta) J_4(\ell\theta), \end{aligned} \tag{A4}$$

as the corresponding E- and B modes in harmonic space are then simply given by

$$\begin{aligned} E_n &= \int_0^\infty \frac{\mathrm{d}\ell\ \ell}{2\pi}\Big(P_E(\ell)[W_+^{(n)}(\ell) + W_-^{(n)}(\ell)] \\ &\qquad + P_B(\ell)[W_+^{(n)}(\ell) - W_-^{(n)}(\ell)]\Big), \\ B_n &= \int_0^\infty \frac{\mathrm{d}\ell\ \ell}{2\pi}\Big(P_E(\ell)[W_+^{(n)}(\ell) - W_-^{(n)}(\ell)] \\ &\qquad + P_B(\ell)[W_+^{(n)}(\ell) + W_-^{(n)}(\ell)]\Big). \end{aligned} \tag{A5}$$

This allows us to read off the defining property of a filter that splits E- and B-modes, namely the equality

$$W_+^{(n)}(\ell) = W_-^{(n)}(\ell), \tag{A6}$$

such that the angular brackets in Eq. (A5) vanish, thereby enforcing the split.

Fortunately, infinitely many solutions for the constraint Eq. (A6) exist. This enables the specification of further advantageous side constraints on the filter. COSEBIs use this freedom to impose additionally that only a finite range in θ be used. The integrals in Eq. (A3) will then only run over the interval $[\theta_{\rm min}, \theta_{\rm max}]$ rather than from zero to infinity. By setting $[\theta_{\rm min}, \theta_{\rm max}]$ to within the bounds of the observation, we can thus enforce that only observed ranges are used.

To exclude angular ranges, the filter $T_+^{(n)}(\theta) \equiv 0$ outside $[\theta_{\rm min}, \theta_{\rm max}]$, implying the filter's zero amplitude does not let pass data from outside these angular ranges. As $T_+^{(n)}$ and $T_-^{(n)}$ are related to each other via the identity of their transforms Eq. (A6), the filter $T_+^{(n)}$ must then additionally fulfil

$$\int_{\theta_{\rm min}}^{\theta_{\rm max}} {\rm d}\theta\ \theta T_+^{(n)}(\theta) = \int_{\theta_{\rm min}}^{\theta_{\rm max}} {\rm d}\theta\ \theta^3 T_+^{(n)}(\theta) = 0. \tag{A7}$$

such that also the $T_-^{(n)}$ does not let pass data outside the chosen angular range.

The two constraints Eq. (A6) and Eq. (A7) are of direct importance, as they enforce the E/B-mode split and the usage of observed angular ranges only. Schneider et al. (2010) furthermore impose the constraint that the sequence of filters be orthonormal, which is convenient, but not necessary. Importantly however, Schneider et al. (2010) show that cosmological information is better retained if the T-filters are logarithmic in θ, thus motivating the last transformation

$$T_+^{(n)}(\theta) = t_+^{(n)}(z), \quad z = \ln(\theta/\theta_{\rm min}). \tag{A8}$$

At this stage, Eq. (A6) till Eq. (A8) fully specified the scientific aims to be met with COSEBIs. The constraints following therefrom can now be met, implying the COSEBI-filters can now be explicitly constructed. To do so, we follow Eq. (28) – Eq. (39) of Schneider et al. (2010), such that we compute our polynomial COSEBIs via their roots. This yields the filters shown in Fig. 1 and Fig. 2 which are used to gain the results on skewness in this paper.

DISCUSSION OF GRAVITATIONAL SKEWNESS

This paper has focused on a *quantification* of skewness in the sampling distribution and likelihood of 2-point functions derived from weak lensing observations. This quantification is required to understand biases during parameter inference. Nonetheless, it transpired that much interest exists to also understand the role of gravity in this skewness.

This interest is in particular caused by the excess skewness of weak lensing distributions (until $\ell \approx 100$), in comparison to the likelihood of the CMB power spectrum, which is only noticeably skewed until $\ell \approx 20$. While the impact of shape noise and cosmic variance on the skewness has already been factored apart in Sellentin et al. (2018), gravity is often put forward as a potential additional cause, and we shall thus discuss it here.

We here observed the trend that excess skewness increases with redshift, which is also consistent with previous studies on gravitationally caused non-Gaussianity: Schmit et al. (2019) show that the gravitational bispectrum is dominated by triangles of the squeezed and elongated type, consistent with our observed redshift trend, and Munshi et al. (2020) show that the weak lensing bispectrum is also dominated by squeezed configurations. The amplitude of the weak lensing bispectrum equally grows with redshift, again consistent with the findings of this paper. We limit ourselves to describing these as 'consistencies' and refrain from attributing the excess skewness to gravity alone.

ACKNOWLEDGEMENTS

ES is funded by the Oort fellowship programme of Leiden Observatory. This research has profited from academic discussions with Koen Kuijken, Henk Hoekstra and Tom Kitching. We thank Henk Hoekstra, Simon Portegies Zwart and Leiden's IT group for sharing their clusters andwelcoming this paper's greedy CPU demands. We thank Joachim Harnois-Deraps for making public the SLICS mock data, which can be found at http://slics.roe.ac.uk/.

REFERENCES

Abbott T. M. C., et al., 2019, Phys. Rev. D, 100, 023541
Alsing J., Wandelt B., 2018, MNRAS, 476, L60
Asgari M., Schneider P., Simon P., 2012, A&A, 542, A122
Asgari M., Heymans C., Blake C., Harnois-Deraps J., Schneider P., Van Waerbeke L., 2017, MNRAS, 464, 1676
Asgari M., et al., 2019, A&A, 624, A134
Bartelmann M., Schneider P., 2001, Phys. Rep., 340, 291
Daniel S. F., Caldwell R. R., Cooray A., Melchiorri A., 2008, Phys. Rev. D, 77, 103513
Diaz Rivero A., Dvorkin C., 2020, arXiv e-prints, p. arXiv:2007.05535
Hamimeche S., Lewis A., 2008, Phys. Rev. D, 77, 103013
Hamimeche S., Lewis A., 2009, Phys. Rev. D, 79, 083012
Harnois-Déraps J., van Waerbeke L., 2015, MNRAS, 450, 2857
Harnois-Déraps J., et al., 2018, MNRAS, 481, 1337
Heavens A. F., Jimenez R., Lahav O., 2000, MNRAS, 317, 965
Heavens A. F., Sellentin E., de Mijolla D., Vianello A., 2017, MNRAS, 472, 4244
Heavens A., Sellentin E., Jaffe A., 2020, arXiv e-prints, p. arXiv:2006.06706
Hildebrandt H., et al., 2020, A&A, 633, A69
Jarvis M., Bernstein G., Jain B., 2004, MNRAS, 352, 338
Kilbinger M., 2015, Reports on Progress in Physics, 78, 086901
Mandelbaum R., et al., 2018, PASJ, 70, S25
Manrique-Yus A., Sellentin E., 2020, MNRAS, 491, 2655

Munshi D., Namikawa T., Kitching T. D., McEwen J. D., Takahashi R., Bouchet F. R., Taruya A., Bose B., 2020, MNRAS, 493, 3985
Planck Collaboration 2014, A&A, 571, A16
Planck Collaboration 2016, A&A, 594, A13
Planck Collaboration 2018, arXiv e-prints, p. arXiv:1807.06209
Planck Collaboration et al., 2019, arXiv e-prints, p. arXiv:1907.12875
Schmit C. J., Heavens A. F., Pritchard J. R., 2019, MNRAS, 483, 4259
Schneider P., Kilbinger M., 2007, A&A, 462, 841
Schneider P., Eifler T., Krause E., 2010, A&A, 520, A116
Sellentin E., Heavens A. F., 2018, MNRAS, 473, 2355
Sellentin E., Heymans C., Harnois-Déraps J., 2018, MNRAS, 477, 4879
Upham R. E., Whittaker L., Brown M. L., 2020, MNRAS, 491, 3165
van Uitert E., et al., 2018, MNRAS, 476, 4662

This paper was built using the Open Journal of Astrophysics LaTeX template. The OJA is a journal which provides fast and easy peer review for new papers in the `astro-ph` section of the arXiv, making the reviewing process simpler for authors and referees alike. Learn more at `http://astro.theoj.org`.