

Enhancing the Cosmic Shear Power Spectrum

Fergus Simpson,^{1*} Joachim Harnois-Déraps,² Catherine Heymans,³ Raul Jimenez,^{1,4,5}
Benjamin Joachimi⁶ and Licia Verde^{1,4,7}

¹ *Instituto de Ciencias del Cosmos, University of Barcelona, UB-IEEC, Martí i Franques 1, E08028, Barcelona, Spain.*

² *Department of Physics and Astronomy, The University of British Columbia, 6224 Agricultural Road, Vancouver, B.C., V6T 1Z1, Canada.*

³ *Institute for Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh EH9 3HJ, UK.*

⁴ *ICREA (Institució Catalana de Recerca i Estudis Avançat).*

⁵ *Institute for Applied Computational Science, Harvard University, MA 02138, USA.*

⁶ *Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK.*

⁷ *Institute of Theoretical Astrophysics, University of Oslo, Oslo 0315, Norway.*

11 November 2015

ABSTRACT

Applying a transformation to a non-Gaussian field can enhance the information content of the resulting power spectrum, by reducing the correlations between Fourier modes. In the context of weak gravitational lensing, it has been shown that this gain in information content is significantly compromised by the presence of shape noise. We apply clipping to mock convergence fields, a technique which is known to be robust in the presence of noise and has been successfully applied to galaxy number density fields. When analysed in isolation the resulting convergence power spectrum returns degraded constraints on cosmological parameters. However substantial gains can be achieved by performing a combined analysis of the power spectra derived from both the original and transformed fields. Even in the presence of realistic levels of shape noise, we demonstrate that this approach is capable of reducing the area of likelihood contours within the $\Omega_m - \sigma_8$ plane by more than a factor of three.

1 INTRODUCTION

The extraction of useful information from cosmological fields (such as galaxy number densities or weak lensing) is often heavily compromised by two limiting factors. Firstly the smallest physical scales are too challenging to model theoretically and are therefore excluded from the analysis. This is particularly problematic for more evolved fields where the amplitude of fluctuations are sufficiently large that the particle dynamics behave in a highly non-linear manner. Secondly, for any non-Gaussian field, the power spectrum fails to retain all of the statistical information captured by the field’s configuration. The amplitude of the field’s Fourier modes are correlated, thereby compromising the information content of the power spectrum. These correlations can be interpreted as originating from the largest dark matter halos, which form spikes in the density field. Each spike in real space translates to a uniform signal in Fourier space, making a coherent contribution to the power spectrum across all wavelengths. Both of these issues, the modelling of small scale clustering and the correlation of Fourier modes, are predominantly associated with the highest density regions of the field.

Manipulating the observed field prior to evaluating the power spectrum can help alleviate these problems. Suppressing fluctuations within the highest density regions in the field improves the performance of perturbation theory, allowing a greater number of Fourier modes to be included in a likelihood analysis. This has been demonstrated with simulated galaxy fields (Simpson et al. 2013) and applied to data from the Galaxy And Mass Assembly (GAMA) galaxy redshift survey (Simpson et al. 2015). The suppression of non-Gaussian peaks also leads to a reduction in the correlation be-

tween Fourier modes, thereby enhancing the amount of information retained by the power spectrum. While previous applications of clipping involved three-dimensional fields, in this work we turn our attention to weak gravitational lensing which involves a projection of the cosmological density field onto two dimensions.

The apparent alignment of galaxies on the sky arises from the optical distortion imposed by the intervening distribution of inhomogeneous matter. When considering the two-point statistics associated with this weak gravitational lensing signal, a prominent degeneracy emerges between two cosmological parameters: the amplitude of linear density perturbations σ_8 and the cosmic matter density Ω_m . For example, the cosmic shear correlation function from the Canada-France-Hawaii Telescope Lensing Survey (CFHTLenS) has been found to measure the combination $\sigma_8 (\Omega_m/0.27)^{0.6} = 0.79 \pm 0.03$ within the context of a standard flat Λ CDM cosmology (Kilbinger et al. 2013). There are a number of possible routes to improve this measurement. Using either tomography (Heymans et al. 2013) or the full three-dimensional information (Kitching et al. 2014) could lead to significant gains. Alternatively, further information can be extracted from the field by analysing the three-point statistics, as demonstrated by Fu et al. (2014), or the abundance of peaks (Cardone et al. 2015; White et al. 2002; Bergé et al. 2010; Maturi et al. 2010; Dietrich & Hartlap 2010).

Seo et al. (2011) found that by invoking a logarithmic transformation, considerable improvements could be made to the information content of the convergence power spectrum, while Yu et al. (2012) demonstrated that further enhancements were achievable by performing a wavelet decomposition of the convergence field.

Joachimi et al. (2011) present a comprehensive analysis of applying Box-Cox transformations (among which the logarithmic transforms form a subset) to the convergence field. They also reported that significant enhancements to the power spectra were achievable under idealised conditions. However each study concluded that realistic levels of shape noise substantially reduce these gains.

Our work aims to build on these previous studies in two ways. First of all we shall focus on the clipping transformation which is well suited to noisy fields. Secondly, we will investigate the potential benefits of combining the information from both spectra, before and after the transformation, by taking into account their cross-covariance.

In §2 we briefly review how local transformations may be applied to weak lensing fields. Methods for estimating the covariance matrices via numerical simulations are detailed in §3, followed by our main results which are presented in §4. Our conclusions are summarised in §5.

2 TRANSFORMING WEAK LENSING FIELDS

Clipping is a local transformation characterised by the application of a saturation value δ_0 to a scalar field $\delta(\mathbf{x})$ such that

$$\begin{aligned} \delta_c(\mathbf{x}) &= \delta_0 & (\delta(\mathbf{x}) > \delta_0) \\ \delta_c(\mathbf{x}) &= \delta(\mathbf{x}) & (\delta(\mathbf{x}) \leq \delta_0) \end{aligned} \quad (1)$$

yielding the clipped field $\delta_c(\mathbf{x})$. The motivation for applying clipping to cosmological fields is twofold: (1) the potential improvement in modelling smaller scale clustering, and (2) the introduction of new information to the power spectrum. In this work we shall focus on the latter.

This transformation has the unique property of either eliminating or unperturbing the existing noise field. Consider a noise field n , defined as the difference between the original field $\delta(\mathbf{x})$ and observed field $\delta_n(\mathbf{x})$

$$n(\mathbf{x}) = \delta_n(\mathbf{x}) - \delta(\mathbf{x}), \quad (2)$$

then the transformed noise field $n'(\mathbf{x})$ is the difference in the transformed fields

$$\begin{aligned} n'(\mathbf{x}) &= f(\delta_n) - f(\delta), \\ &= n(\mathbf{x})f'(\delta), \end{aligned} \quad (3)$$

where f' is the functional derivative of f . For clipping, as specified by (1), $f'(\delta)$ can only take two values: zero or unity. Therefore, the noise contribution is either removed (in regions above the threshold) or left untouched (in regions below the threshold).

The benefits of decorrelating the Fourier modes within the galaxy power spectrum are limited since its amplitude is dictated by the uncertain manner in which galaxies trace the underlying matter distribution. Conversely, enhancing the weak lensing power spectrum would be of greater consequence as it acts as a direct measure of the amplitude of matter fluctuations.

A possible disadvantage of applying clipping to lensing stems from the fact that lensing involves a projection of the three dimensional density field onto two dimensions. This averaging process will reduce the level of non-Gaussianity in the field, and so a greater proportion of the field must be subject to clipping in order to achieve the desired result.

While the direct observable from weak lensing surveys are the shapes of galaxies, it is the derived convergence field which directly relates to the projected mass density. The convergence field $\kappa(\boldsymbol{\theta})$ is

therefore the more natural quantity to work with. It may be thought of as a weighted integral of the matter perturbations along the line of sight

$$\kappa(\boldsymbol{\theta}) \simeq \frac{3}{2} \left(\frac{H_0}{c} \right)^2 \Omega_m \int g(\chi) f_K(\chi) \frac{\delta(\boldsymbol{\theta}, \chi)}{a(\chi)} d\chi, \quad (4)$$

where χ denotes the radial coordinate distance. The lensing efficiency $g(\chi)$ is determined by the radial distribution of source galaxies $n_g(\chi)$ and the comoving angular diameter distance $f_K(\chi)$

$$g(\chi) = \int_{\chi}^{\infty} d\chi' n_g(\chi') \frac{f_K(\chi' - \chi)}{f_K(\chi')}. \quad (5)$$

If we decompose the matter field into linear and non-linear components, $\delta = \delta_G + \delta_x$, where δ_G is the underlying Gaussian random field and the residual term δ_x incorporates all departures from Gaussianity, then from Eq. (4) the convergence is separable in the same manner, yielding $\kappa(\boldsymbol{\theta}) = \kappa_G(\boldsymbol{\theta}) + \kappa_x(\boldsymbol{\theta})$. In the limit that the non-Gaussian component $\kappa_x(\boldsymbol{\theta})$ only contributes at field values larger than the chosen threshold, the application of clipping as defined by Eq. (1) ensures that only the Gaussian component of the density field $\delta_G(\boldsymbol{\theta})$ contributes to the resulting convergence power spectrum P_κ .

For comparison we shall also assess the performance of the logarithmic transformation (Seo et al. 2011; Joachimi et al. 2011)

$$\bar{\kappa}(\boldsymbol{\theta}) = \kappa_0 \ln [1 + \kappa(\boldsymbol{\theta})/\kappa_0], \quad (6)$$

where the parameter κ_0 can be adjusted to compensate for the amplitude of the convergence field, serving a similar purpose to the clipping threshold δ_0 . We follow the prescription used in Seo et al. (2011) and set $\kappa_0 = |\kappa_{\min}| + 0.001$, where κ_{\min} is the most negative value of the convergence. This small offset ensures that the transformed field remains finite.

3 METHODOLOGY

In order to quantify the benefits of clipping the convergence field, we perform a Fisher matrix analysis to determine the error forecast in the $\Omega_m - \sigma_8$ plane for a fiducial survey, and compare the size of the expected likelihood contours before and after the convergence field is subject to a local transformation.

It is well known that in circumstances where the posterior is not well described by a multivariate Gaussian, the Fisher matrix acts as a poor estimator (see for example Sellentin et al. (2014); Sellentin & Schäfer (2015)). However when considering forecasts for weak lensing surveys in the $\Omega_m - \sigma_8$ plane, then provided the likelihood contours are sufficiently small, comparisons between Fisher predictions and Markov Chain Monte Carlo (MCMC) outputs are found to be highly consistent (see for example Figure 3 of Taylor & Kitching 2010 and Figure 6 of Wolz et al. 2012).

3.1 Simulations

We make use of mock convergence fields from simulations described in Harnois-Déraps & van Waerbeke (2015). We apply a variety of clipping thresholds to the mock convergence fields associated with 70 independent lines of sight, each of which spans 60 square degrees. Sources are placed at a redshift of $z = 0.8$ and the fiducial amplitude of linear density fluctuations is taken to be

$\sigma_8 = 0.826$. These lines of sight serve as our estimator for the covariance of the angular power spectra $P_\kappa(\ell)$. In addition to these, we utilise a set of five simulations which are seeded with an initial particle distribution with matching phases, but with small displacements in the values of Ω_m and σ_8 . This allows for a clean extraction of the partial derivative of the observables with respect to these parameters, as required for estimating the Fisher Matrices, without being contaminated by cosmic variance. Other cosmological parameters such as the spectral index n_s and global curvature Ω_k are taken to be fixed, as they can be measured with a much greater precision from other observations such as the Cosmic Microwave Background.

For each line of sight the convergence field is assigned to a 300×300 grid, such that each grid cell corresponds to approximately 1.5 arcminutes, and the convergence power spectra are evaluated in eight equally spaced bins for multipoles $\ell < 1,500$. Our results are largely insensitive to the choice of smoothing length. However if very fine grids are used then clipping predominantly acts on noise spikes rather than true peaks in the projected density field.

When generating each simulated convergence field, the source redshift plane ought to remain fixed. However the ray-tracing takes place across an integer number of simulation boxes, and since the redshift-distance relation is altered when perturbing Ω_m , there is inevitably a small unwanted shift in the redshift of the sources. We compensate for this by rescaling the amplitude of the convergence field, which accounts for the change in lensing efficiency $g(\chi)$ due to the small difference between the desired and simulated distances to the source plane. The magnitude of this correction is approximately 1% and so would not change our results by more than a few percent if it were neglected.

3.2 Shape Noise

Starting from the true convergence field, we use the inverted Kaiser Squires method (Kaiser & Squires 1993) to recover the shear field at each cell location. Then noise is superposed by drawing two random numbers at each cell location, from a Gaussian distribution with mean of zero and standard deviation $\bar{\sigma}_e/\sqrt{2}$, which constitute the real and imaginary parts of the mean intrinsic ellipticity of the galaxies within that cell. The cell averaged ellipticity $\bar{\sigma}_e$ is given by σ_e/\sqrt{n} where n is the number of source galaxies per cell, and we take the rms intrinsic ellipticity $\sigma_e = 0.3$ throughout this work. The noisy shear field is then given as a sum of the true and intrinsic shear fields. The noisy convergence field is recovered by applying the Kaiser-Squires transformation (Kaiser & Squires 1993) to the noisy shear field.

The intrinsic shape noise introduces an additive white noise contribution which includes a distortion associated with the window function of the grid

$$P_n(k) = \frac{\sigma_e^2}{2\bar{n}} \left\langle \frac{1}{\text{sinc}(k_x/2)^2} \frac{1}{\text{sinc}(k_y/2)^2} \right\rangle, \quad (7)$$

where k_x and k_y specify the elements of the wavevector in the x and y dimensions, in units of the Nyquist frequency, and the averaging takes place within the annulus associated with the wavevector bin. Clipping reduces the noise contribution to the power spectrum, since the regions of the field lying above the threshold are smoothed. This leads to an estimate of the clipped shot noise contribution $P_{cn}(k)$, derived from equation (5) of Simpson et al. (2013), where the leading order term may be expressed as




Figure 1. The ratio of the convergence power spectrum from the clipped convergence field, to that of the original field, for three different values of σ_8 . In this example a clipping threshold is chosen such that on the largest angular scales the power is reduced by a factor of two, which typically provides a good balance between the removal of non-linear power while maintaining a high level of signal to noise.

$$P_{cn}(k) \simeq (1 - f_c)^2 P_n(k), \quad (8)$$

where f_c is the proportion of the field which lies above the clipping threshold. For the fields considered in this work, the value of f_c required to reduce the large scale power by a factor of two is around 20%. By comparison, the proportion of a three-dimensional galaxy field subject to a comparable clipping strength is typically 2% (Simpson et al. 2015).

3.3 Covariance Estimation

Since estimating the covariance matrix is a critical step in this analysis, and one which can be highly susceptible to the influence of noise, we compare two different methods for estimating the covariance matrix. First we use the Ledoit-Wolf shrinkage estimator (Ledoit & Wolf 2004) to evaluate the covariance matrix \hat{C} associated with our binned power spectra.

$$\hat{C} = \delta^* F + (1 - \delta^*) S, \quad (9)$$

We use publicly available code¹ to estimate the shrinkage target F and the shrinkage constant δ^* . The form of the shrinkage target is based upon the samples being independent and identically distributed. The sample covariance matrix S is defined as the ensemble average over 70 lines of sight

$$S_{ij} = \langle \Delta P_\kappa^m(i) \Delta P_\kappa^n(j) \rangle, \quad (10)$$

where $\Delta P_\kappa^m(i)$ represents the deviation in power from the sample mean, and m denotes the particular transformation(s) under consideration. The total number of columns in the covariance matrix is given by the product of the number of angular bins in $P_\kappa(\ell)$ with the number of transformations included in the analysis, in order to accommodate their cross-covariance.

The finite number of fields is a source of sampling noise which

¹ <http://www.econ.uzh.ch/faculty/wolf/publications/covCor.m.zip>




Figure 2. *Left:* Constraints in the $\Omega_m - \sigma_8$ plane using the angular power spectra ($\ell < 1,500$) of the original convergence field (dashed), the clipped field (thin solid) and the log field (dotted). The thick set of contours relate to the combined analysis of power spectra from the original and clipped fields. These constraints correspond to a survey of 10,000 square degrees. *Right:* The same format as the left panel, but we now include the effects of shape noise, with a galaxy number density of 20 arcmin $^{-2}$ and an rms intrinsic ellipticity $\sigma_e = 0.3$.

will inevitably propagate through to the estimation of our likelihood contours. In order to confirm that the uncertainty in our likelihood contours is negligible, we employ a second algorithm to estimate the covariance matrix. Instead of using a shrinkage estimator, we simply invert the sample covariance matrix S before applying a rescaling factor which accounts for the noise-induced bias in the amplitude (Hartlap et al. 2007)

$$\hat{C}^{-1} = \frac{n-p-2}{n-1} S^{-1}, \quad (11)$$

where p represents the length of the data vector (either 8 or 16), and n is the number of samples. When all 70 fields are used to construct the sample covariance S , these two algorithms produce highly consistent results, the derived likelihood contours differ in area by less than 5%.

A further validation test for convergence of the covariance matrix involves repeating this procedure with a small subset of the available samples. This also proves useful in order to identify which algorithm produces more consistent results. For example, when using only 20 samples the rescaling method defined by Eq. (11) systematically overestimates the area of the likelihood contour by more than 40%. By contrast, the shrinkage estimator recovers likelihood contours which are consistently within 40% of the area derived from the full sample. This improved stability leads us to conclude that the shrinkage algorithm likely offers a superior performance. It also predicts slightly more conservative gains when combining the clipped and unclipped spectra. For these two reasons, the shrinkage algorithm defined by Eq. (9) is the one we adopt when presenting our main results.

3.4 Constructing the Fisher Matrix

The precision with which a parameter may be determined is governed by the rate at which the observable responds to a change in the parameter. The key observable to consider is the power spectrum of the convergence field $P_\kappa(\ell)$. This may be expressed as

$$P_\kappa(\ell) = \frac{9}{4} \Omega_m^2 \left(\frac{H_0}{c} \right)^4 \int_0^\infty \frac{g^2(\chi)}{a^2(\chi)} P_\delta \left(\frac{\ell}{f_K(\chi)}, \chi \right) d\chi, \quad (12)$$

where the lensing efficiency $g(\chi)$ is defined in Eq. (5).

Predicting the power spectrum of a clipped convergence field poses a significantly more challenging problem. As is apparent from Eq. (12), the convergence field can be thought of as arising from the weighted sum of intervening lens planes. Applying clipping to a convergence field derived from a single lens plane would possess a power spectrum in line with the model presented in Simpson et al. (2013). However each of the stacked lens planes is at a slightly different stage of gravitational evolution. We therefore resort to using numerical simulations in order to determine the expected behaviour. In order to define a consistent clipping strength, independent of smoothing length, we specify the fractional drop in amplitude of the large scale power. We perform an iterative procedure on the threshold value δ_0 to converge upon the desired clipping strength. Note that this iteration is not performed for each field - it is performed once, then this threshold value is fixed for all remaining fields within the ensemble. Furthermore, prior to applying any transformation, each convergence field has its mean subtracted. This reflects the uncertain normalisation of a given convergence field derived from the measurement of a shear field, although this is less problematic for fields spanning a large area.

Our fiducial survey spans 10,000 square degrees, and involves a galaxy number density of 20 arcmin $^{-2}$. For simplicity we only consider a single redshift bin, rather than using tomography, and the mock survey is idealised in that there are no holes or masks.

The Fisher information matrix, formally the expectation of the Hessian matrix of the log likelihood, may be expressed as

$$F = D^T C^{-1} D, \quad (13)$$

where C is the covariance matrix corresponding to the various bins in $P_\kappa(\ell)$ and is estimated following the methodology in §3.3. Where a combined analysis of two spectra is performed, the covariance matrix incorporates the cross-covariance between the two spectra. The Jacobian matrix D is comprised of elements D_{ij} which

	Unclipped	Logarithm	Clip 40%	Clip 60%	Clip 80%
Unclipped	1	1.4	3.1	3.3	2.9
Logarithm		0.73	2.6	2.5	1.8
Clip 40%			0.59	0.87	1.5
Clip 60%				0.59	1.2
Clip 80%					0.72

Table 1. The improvement in the Figure of Merit as defined in Eq. (15), for various combinations of power spectra from transformed convergence fields, relative to that of the original power spectrum. The four local transformations - three clipping thresholds and the logarithmic transform - all prove detrimental when considered alone, yet yield substantial gains when coupled with the power spectrum of the original field.

hold the differential of the i th multipole bin with respect to the j th parameter,

$$D_{ij} = \frac{\partial P_{\kappa}(i)}{\partial p_j}. \quad (14)$$

We expect transformed fields will be susceptible to shape noise and so D must be re-evaluated for each transformation when noise is introduced, as discussed in further detail in Seo et al. (2012). We take the average power derived from twenty different noise realisations in order to ensure that the above derivatives are stable.

There are two distinct analyses we wish to perform in the $\Omega_m - \sigma_8$ plane. The first is the constraint which one would achieve from a conventional interpretation of the weak lensing power spectrum, without clipping. In this case the determination of the differential matrix D can be performed analytically. This is not the case for the power spectra derived from the clipped convergence fields. In this case, to determine the differential matrix D we utilise the simulations which were seeded with matching phases in the initial density fluctuations. Due to the nonlinear nature of structure formation, the smallest density fluctuations exhibit a stochastic behaviour across the different simulations, but on the angular scales of interest we find the derivatives to be numerically stable.

4 RESULTS FROM SIMULATIONS

Using the simulations outlined in the previous section, we explore the consequences of applying local transformations to the convergence field, and construct likelihood contours derived from the resulting convergence power spectra.

4.1 Power Spectra

Figure 1 shows the change in the shape of the convergence power spectrum for three different cosmologies, after applying a clipping threshold chosen such that the largest angular bin in C_ℓ is reduced in amplitude by a factor of two. Previous studies of clipping cosmological fields conclude that reducing the linear power by between 40–60% strikes a good balance by suppressing higher order terms, while maintaining a high level of signal to noise. This therefore drives our choice of the threshold value, which for the central value ($\sigma_8 = 0.831$) corresponds to $\kappa_0 = 0.026$. The three different values of σ_8 generate significantly different responses, illustrating how clipping can reveal the degree of non-linearity present in the field. For more evolved density fields, a greater suppression of power occurs on small angular scales. This response can be understood by considering the original field to be a superposition of its linear and nonlinear components. The clipped power is then predominantly comprised of the clipped linear power spectrum and the clipped

nonlinear power (Simpson et al. 2013). However owing to the positive skewness of the nonlinear pdf, the nonlinear power suffers a much stronger loss of power than the Gaussian pdf. Thus, fields which are more evolved and host a stronger nonlinear component suffer a greater loss of power on small scales. This behavior has previously been observed for the case of the three-dimensional simulated density fields (Simpson et al. 2013). Applying stronger clipping with a lower threshold value leads to a more linear power spectrum with less correlated Fourier modes. This improvement saturates once the amplitude of the linear power spectrum is reduced by approximately a factor of two, so there is little to be gained by adopting a lower threshold.

4.2 Parameter Constraints

The left hand panel of Figure 2 illustrates the 68 and 95% joint likelihood contours derived from the Fisher matrix of the convergence power spectrum ($\ell < 1500$). The dashed lines are derived from the standard convergence field, while the dotted lines are derived from the log transform as defined by Eq. (6). The outer set of solid contours make use of the clipped convergence field as defined by Eq. (1), with a threshold selected such that the large scale power is reduced to 50% of its original value. These contours are somewhat elongated relative to the original spectrum, but also exhibit a small clockwise rotation of their degeneracy direction. The inner set of solid contours result from combining both the original and clipped power spectra. The minor axes of the combined contours are not noticeably reduced from those of the original spectrum, owing to the strong cross-covariance between the amplitudes of the two spectra. However the major axes shrink considerably, suggestive of a powerful complementarity between the information held by the two power spectra.

Correlations between data points are often perceived to have a detrimental impact on the dataset’s overall information content. But there are many circumstances in which this is not the case. To see why, consider the simple case of two data points, to which we wish to fit a straight line. The only two free parameters are the gradient and the amplitude. If the errors on these two data points are perfectly correlated, the constraint on the amplitude weakens by a factor of $\sqrt{2}$ relative to the uncorrelated case. However the gradients uncertainty vanishes - it is now determined to an arbitrarily high precision.

This toy model is intimately linked to the case we consider here. For weak lensing contours in the $\Omega_m - \sigma_8$ plane, the size of the minor axis relates to the uncertainty in the amplitude of the weak lensing power spectrum. This is often quoted as a measurement of the parameter combination $\sigma_8 \Omega_m^\alpha$. The fact that the clipped and unclipped spectra are highly correlated means that when a joint analysis is performed, any improvement in the measurement of this amplitude is negligible. Hence the minor axis does not shrink for the case of the ‘Combined’ contours in Figure 2. Meanwhile the ex-

tent of the contour’s major axis is governed by how well the shape of the power spectrum is measured. More evolved fields with higher values of σ_8 will possess more power on smaller angular scales, and this helps lift the degeneracy within the $\Omega_m - \sigma_8$ plane. When considered separately, the precision with which the shape of each spectra can be determined is ultimately limited by cosmic variance. It might be that the large scale power is low by chance, causing a spurious upward tilt in the lensing power spectrum, thereby mimicking a higher value of σ_8 . However this deception can be revealed with a joint analysis of both spectra. A higher value of σ_8 leads to a stronger decrement of small scale power in the clipped spectrum relative to the unclipped. This relative change in power is largely insensitive to the particular cosmic variance realisation. The contour’s major axis therefore shrinks further than would have been possible if the two spectra were uncorrelated. This is explored further in §4.3 and Figure 3 below.

Note that this is not the only cosmological example where correlations prove to be beneficial. Performing an analysis of redshift space distortions with multiple tracers allow cosmic variance fluctuations to be effectively nulled, permitting an enhanced measurement of the growth of structure.

These results are consistent with the findings of Joachimi et al. (2011) who noted that upon performing a Box-Cox transformation the likelihood contours demonstrate a small clockwise rotation of the degeneracy direction, and become considerably elongated. This elongation of the contours for transformed fields occurs because information regarding the shape of the spectra (as induced by non-linear structure associated with higher values of σ_8) is erased. Indeed the full extent of this elongation effect is likely lost in our case due to the limitations of the Fisher formalism. Nonetheless the addition of the original information from the unclipped field truncates the longer contours. The combined contours therefore offer a faithful representation of the true likelihood.

Our findings differ from Seo et al. (2012) in that even in the noise-free case, we do not observe significant gains from applying a logarithmic transform alone. However there are a number of differences in our methodologies. They use multiple redshift bins and explore a broader parameter space - for example we do not marginalise over the dark energy parameters w_0 and w_a . Furthermore the efficacy of the log transform is strongly governed by the choice of pixel size.

The right hand panel of Figure 2 takes the same format as the left panel, but here we include the impact of shape noise. There is only a modest degradation of the unclipped contours in this case. The degeneracy direction of the transformed fields experiences a slightly reduced rotation, but there is still sufficient complementarity to yield substantial gains when performing a combined analysis. The logarithmic contours also show little degradation, which may appear surprising given the information loss found by Seo et al. (2011). However the impact of shape noise is highly sensitive to the choice of smoothing length used to define the field, particularly since the amplitude of the convergence fundamentally alters the transformation defined in (6), and here we are using a considerably coarser pixel size.

The extent to which the benefits of applying clipping and logarithmic transformations are degraded by the presence of shape noise also depends on the specific choice of transformation chosen. Higher clipping thresholds lead to a higher power spectrum and thus a lower susceptibility to shape noise. Yet they are also less efficient at recovering information to the power spectrum in the low noise limit.

In order to quantify these gains we can readily compute the




Figure 3. Simulated constraints for a 200 square degree convergence field, before (red dashed contours) and after (blue solid contours) applying a logarithmic transformation (95% C.L.). The inner black solid contour represents a combined analysis of both spectra. The consequence of ignoring the cross-covariance between the original and transformed spectra is illustrated by the dotted black contour. Shape noise is not included in the convergence field. The black marker signifies the fiducial cosmology of the simulation.

change in the figure of merit such as those previously defined in Albrecht et al. (2009) and Simpson & Peacock (2010)

$$\text{FoM} = \sqrt{F_{\Omega\Omega}F_{\sigma\sigma} - F_{\Omega\sigma}^2}, \quad (15)$$

where the notation $F_{\theta\theta}$ denotes the elements of the Fisher matrix associated with the parameter θ . This quantity is inversely proportional to the area enclosed by the likelihood contours. The absolute value is not of great interest, since it will strongly depend on the configuration and area of the survey in question, so in Table 1 we quote in each case the fractional change in the Figure of Merit relative to the original power spectrum analysis. We show results for three different clipping thresholds, corresponding to those which reduce the large scale power to 40, 60 and 80% of its original value. In addition we include the original convergence field, and the log transform as defined by Eq. (6). The diagonal elements demonstrate that for each transformation under consideration, the areas enclosed by the ellipses have enlarged. The off-diagonal elements in Table 1 quantify the significant gains which can be made by performing a combined analysis, particularly when combining the original spectrum with that from a clipped field, where the gain is more than a factor of three, despite the presence of shape noise. We have also evaluated intermediate clipping strengths, finding that the gain derived from combining two clipped power spectra deteriorates rapidly as the difference between their clipping strengths is reduced.

Thus while the findings of previous related studies (Seo et al. 2011; Joachimi et al. 2011) are entirely consistent with the results presented here, our conclusions differ markedly in that we find substantial gains can be achieved from a local transformation of the convergence field even in the presence of realistic levels of shape noise. The source of this difference stems from performing a combined analysis with the original power spectrum, which holds com-

plementary information, rather than considering each power spectrum in isolation.

4.3 Full likelihood model

The posterior in the parameters Ω_m and σ_8 as produced by contemporary cosmic shear surveys is not well approximated by a multivariate Gaussian. In order to assess whether our conclusions are affected by the limitations of the Fisher matrix formalism, we consider the case of a full mock likelihood analysis following the methodology presented in Joachimi et al. (2011). In this work analytical models for power spectra of convergence fields transformed via Box-Cox transformations, which include the logarithm, were derived using a perturbative expansion and recalibration of higher-order contributions via the third to fifth moments of the convergence measured from N-body simulations. Joachimi et al. (2011) used a suite of 100 convergence fields at a single source redshift $z = 1$ simulated with the SUNGLASS pipeline (Kiessling et al. 2011) for signal and covariance estimation via Eq. (11). We adopt these measurements for convergence field without shape noise and smoothed with a Gaussian of width 1.5 arcmin. Moreover, we rescale the covariance to correspond to a 200 deg² survey. This toy model setup is deliberately chosen to yield wide contours with a strong non-linear degeneracy, which is least well represented by a Gaussian approximation to the posterior. We use the value for the free parameter in the logarithmic transformation, equivalent to κ_0 in Eq. (5), as determined by Joachimi et al. (2011) in the Box-Cox optimisation.

The resulting parameter constraints for a power spectrum analysis of the original convergence fields, of the log-transformed fields, and their combined analysis are shown in Figure 3. We have used 10 angular frequency bins logarithmically spaced in the range 150 to 1500. The contours for the log-transformed fields are slightly biased high, which is caused by the difficulty of analytically modelling the transformed power spectrum at high ℓ where higher-order correlations become important (see Joachimi et al. 2011 for a more detailed discussion). However, this does not affect our conclusion that there is excellent qualitative agreement between the full likelihood analysis and the Fisher forecasts. The dotted contours represent a joint analysis in which the cross-covariance has been set to zero, as would arise if the two spectra were taken from fields in different parts of the sky. This greatly underestimates the achievable constraints, demonstrating the importance of the cross-correlations in suppressing noise contributions.

This independent analysis offers striking confirmation of the considerable gains which are achievable by accounting for the strong cross-covariance which exists between the power spectra of a locally transformed field and its original form. In terms of the areas enclosed by the contours in Figure 3, the combined analysis offers a gain in performance of more than a factor of five compared to the transformed field alone.

5 CONCLUSIONS

In this work we have demonstrated how applying a transformation such as clipping to a convergence field can not only introduce new information to the resulting power spectrum, but that this new information helps lift the degeneracy in the $\Omega_m - \sigma_8$ plane. When considering moderately large angular scales $\ell < 1500$, parameter constraints can be improved by more than a factor of two, even when realistic levels of shape noise are included. The loss of performance

reported in previous studies for a single transformed power spectrum is largely attributed to the elongation of the contours along the degeneracy direction, but this elongation is readily truncated by including the original power spectrum which exhibits a different degeneracy direction.

A second major advantage of performing a local transformation such as clipping is that it may also permit smaller angular scales to be studied, since the resulting power spectrum is more straightforward to model. We may therefore have significantly underestimated the gains achievable when clipping convergence fields. The use of tomography - dividing the source galaxies into discrete redshift bins - in conjunction with clipping will further enhance the level of precision attainable, but we leave this topic for the subject of future work. A suite of N-body cosmological simulations would be required to adequately model the clipped convergence power spectra. Many simulations will in any case be required for modelling the conventional power spectrum, or alternatively rescaling methods could prove effective for exploring a greater volume of parameter space (Angulo & White 2010; Mead & Peacock 2014).

As found when applying clipping to large scale structure, the choice of threshold is fairly flexible. Provided the threshold is sufficiently strong to reduce the large scale amplitude of the power spectrum by over 30%, there is little change in performance. The optimal choice ultimately rests upon the level of noise within the particular field under investigation. For the convergence fields considered in this work we find that in order for clipping to reduce the large scale power by a factor of two, a threshold value of $\kappa_0 = 0.026$ is required, corresponding to approximately 20% of the field being clipped.

The greatest obstacle to applying a local transformation to weak lensing data is that the observable field is not a direct tracer of the matter density, so we must convert the shear field to a convergence field in order to find a suitable proxy for the projected matter density. The conversion process is particularly challenging in the presence of noise and masks, but significant progress is being made (Van Waerbeke et al. 2013; Chang et al. 2015). We therefore envision that a combined analysis of clipped and unclipped lensing observations will help extract more robust and precise information from future and forthcoming surveys.

ACKNOWLEDGEMENTS

FS acknowledges support by the European Research Council under the European Community's Seventh Framework Programme FP7-IDEAS-Phys.LSS 240117. JHD is supported by the NSERC of Canada. LV acknowledges support from the European Research Council (grant FP7-IDEAS-Phys.LSS 240117) and a Mineco grant FPA2011- 29678-C02-02. RJ acknowledges support from Mineco grant FPA2011- 29678-C02-02. CH acknowledges support from the European Research Council under the EC FP7 grant number 240185. BJ acknowledges support by an STFC Ernest Rutherford Fellowship, grant reference ST/J004421/1.

Computations for the N-body simulations were performed on the GPC supercomputer at the SciNet HPC Consortium. SciNet is funded by: the Canada Foundation for Innovation under the auspices of Compute Canada; the Government of Ontario; Ontario Research Fund - Research Excellence; and the University of Toronto.

REFERENCES

- Albrecht A. et al., 2009, ArXiv e-prints: 0901.0721
Angulo R. E., White S. D. M., 2010, MNRAS, 405, 143
Bergé J., Amara A., Réfrégier A., 2010, ApJ, 712, 992
Cardone V. F., Camera S., Sereno M., Covone G., Maoli R., Scaramella R., 2015, Astron. & Astrophys., 574, A141
Chang C. et al., 2015, ArXiv e-prints: 1505.01871
Dietrich J. P., Hartlap J., 2010, MNRAS, 402, 1049
Fu L. et al., 2014, MNRAS, 441, 2725
Harnois-Déraps J., van Waerbeke L., 2015, MNRAS, 450, 2857
Hartlap J., Simon P., Schneider P., 2007, Astron. & Astrophys., 464, 399
Heymans C. et al., 2013, MNRAS, 432, 2433
Joachimi B., Taylor A. N., Kiessling A., 2011, MNRAS, 418, 145
Kaiser N., Squires G., 1993, ApJ, 404, 441
Kiessling A., Heavens A. F., Taylor A. N., Joachimi B., 2011, MNRAS, 414, 2235
Kilbinger M. et al., 2013, MNRAS, 430, 2200
Kitching T. D. et al., 2014, MNRAS, 442, 1326
Ledoit O., Wolf M., 2004, Journal of multivariate analysis, 88, 365
Maturi M., Angrick C., Pace F., Bartelmann M., 2010, Astron. & Astrophys., 519, A23
Mead A. J., Peacock J. A., 2014, MNRAS, 440, 1233
Sellentin E., Quartin M., Amendola L., 2014, MNRAS, 441, 1831
Sellentin E., Schäfer B. M., 2015, ArXiv e-prints
Seo H.-J., Sato M., Dodelson S., Jain B., Takada M., 2011, Astrophys. J. Lett., 729, L11
Seo H.-J., Sato M., Takada M., Dodelson S., 2012, ApJ, 748, 57
Simpson F. et al., 2015, ArXiv e-prints: 1505.03865
Simpson F., Heavens A. F., Heymans C., 2013, PRD, 88, 083510
Simpson F., Peacock J. A., 2010, PRD, 81, 043512
Taylor A. N., Kitching T. D., 2010, MNRAS, 408, 865
Van Waerbeke L. et al., 2013, MNRAS, 433, 3373
White M., van Waerbeke L., Mackey J., 2002, ApJ, 575, 640
Wolz L., Kilbinger M., Weller J., Giannantonio T., 2012, JCAP, 9, 9
Yu H.-R., Harnois-Déraps J., Zhang T.-J., Pen U.-L., 2012, MNRAS, 421, 832