



Universiteit
Leiden
The Netherlands

Means of confusion: how pixel noise affects shear estimates for weak gravitational lensing

Melchior, P.; Viola, M.

Citation

Melchior, P., & Viola, M. (2012). Means of confusion: how pixel noise affects shear estimates for weak gravitational lensing. *Monthly Notices Of The Royal Astronomical Society*, 424(4), 2757-2769. doi:10.1111/j.1365-2966.2012.21381.x

Version: Not Applicable (or Unknown)

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/46825>

Note: To cite this publication please use the final published version (if applicable).

Means of confusion: how pixel noise affects shear estimates for weak gravitational lensing

P. Melchior^{1,2★} and M. Viola³

¹Center for Cosmology and Astro-Particle Physics, The Ohio State University, 191 W. Woodruff Ave., Columbus, OH 43210, USA

²Department of Physics, The Ohio State University, 191 W. Woodruff Ave., Columbus, OH 43210, USA

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

Accepted 2012 May 23. Received 2012 May 18; in original form 2012 April 23

ABSTRACT

Weak-lensing shear estimates show a troublesome dependence on the apparent brightness of the galaxies used to measure the ellipticity: in several studies, the amplitude of the inferred shear falls sharply with decreasing source significance. This dependence limits the overall ability of upcoming large weak-lensing surveys to constrain cosmological parameters.

We seek to provide a concise overview of the impact of pixel noise on weak-lensing measurements, covering the entire path from noisy images to shear estimates. We show that there are at least three distinct layers, where pixel noise not only obscures but also biases the outcome of the measurements: (1) the propagation of pixel noise to the non-linear observable ellipticity; (2) the response of the shape-measurement methods to limited amount of information extractable from noisy images and (3) the reaction of shear estimation statistics to the presence of noise and outliers in the measured ellipticities.

We identify and discuss several fundamental problems and show that each of them is able to introduce *biases in the range of a few tens to a few per cent* for galaxies with typical significance levels. Furthermore, all of these biases do not only depend on the brightness of galaxies but also depend on their ellipticity, with more elliptical galaxies often being harder to measure correctly. We also discuss existing possibilities to mitigate and novel ideas to avoid the biases induced by pixel noise. We present a new shear estimator that shows a more robust performance for noisy ellipticity samples. Finally, we release the open-source PYTHON code to predict and efficiently sample from the noisy ellipticity distribution and the shear estimators used in this work at <https://github.com/pmelchior/epsnoise>.

Key words: gravitational lensing; weak – techniques: image processing.

1 INTRODUCTION

Current and in particular upcoming wide-field imaging surveys such as the Dark Energy Survey,¹ the Kilo Degree Survey,² the Hyper Suprime Camera Survey, Euclid³ and Wide-Field Infrared Survey Telescope (WFIRST)⁴ require highly accurate shape-measurement methods to reach the forecasted accuracy of cosmological parameter constraints, e.g. for the energy density of matter Ω_m , the normalization of the matter power spectrum σ_8 , the dark energy equation-of-state parameter w and its variation with time. The currently most demanding lensing method, the cosmic shear two-point correlation

function, allows for multiplicative errors (defined as the deviation of actual shear from the measurement by a factor m), with $|m|$ not larger than some per mille (e.g. Huterer et al. 2006; Amara & Refregier 2008). But with increasing survey volumes also traditionally less demanding techniques, such as stacked cluster lensing, will require $|m|$ of the order of 1 per cent (Weinberg et al. 2012). Shear measurement methods known so far can reach these requirements in certain regimes, for instance for well-resolved and bright galaxies. However, they commonly struggle with small and especially with faint galaxies (Massey et al. 2007; Bridle et al. 2010; Kitching et al. 2012). Many suggestions have been brought forward as to why prominent pixel noise hampers shape measurements, but often it was difficult to disentangle the causes from their consequences.

Bernstein & Jarvis (2002) computed that this so-called *noise rectification bias* generally scales inversely with the second power of the object's significance. Hirata et al. (2004) obtained an analytic description of the dependence of this bias on the size of galaxies, but this derivation only applies to their adaptive moment-based

★E-mail: melchior.12@osu.edu

¹ <http://www.darkenergysurvey.org/>

² <http://kids.strw.leidenuniv.nl/>

³ <http://sci.esa.int/euclid>

⁴ <http://wfIRST.gsfc.nasa.gov/>

measurement method. Recently, Refregier et al. (2012) showed that biases in the maximum-likelihood estimators of model-fitting approaches are a direct consequence of the presence of non-linear fit parameters. They also provide an analytic expression for the bias of the ellipticity estimate, valid for Gaussian-shaped galaxies and point spread functions (PSFs).

We seek to generalize and extend these previous findings. In particular, we choose an approach, which is as method independent as possible, offering insights into the several ways in which pixel noise obscures the estimation of weak gravitational shear.

1.1 Approach

Throughout this work, we aim to identify conceptual problems \mathcal{P} for the shape estimation task, which give rise to deviations of the measured source ellipticity $\epsilon'(\mathcal{P})$ from the true ellipticity ϵ . We assume a probabilistic approach, where we inspect the probability distribution of \mathcal{P} under noise, $p(\mathcal{P}|\nu)$, with ν denoting some suitable characterization of the significance of the measurement, and its consequence, the probability distribution of the measured ellipticity caused by \mathcal{P} ,

$$p_{\mathcal{P}}(\epsilon'|\nu) = \int d\mathcal{P} \epsilon'(\mathcal{P}) p(\mathcal{P}|\nu). \quad (1)$$

The ellipticity distribution is thus given by the impact \mathcal{P} has on ϵ' , weighted by the probability that \mathcal{P} actually occurs in a measurement with significance ν . The key assumption here is that we could measure the function $\epsilon'(\mathcal{P})$ perfectly, i.e. without pixel noise, while the action of the noise is entirely contained in the width and shape of the probability distribution of \mathcal{P} .

In practice, both $\epsilon'(\mathcal{P})$ and $p(\mathcal{P}|\nu)$ additionally depend on the apparent shape of the source, i.e. its intrinsic shape and the effects of the convolution with the PSF. We therefore introduce a parametrization of the apparent source morphology with a parameter vector θ .

Although the true ellipticity ϵ could be regarded as one of these source parameters, we choose to make the dependency of ϵ' on ϵ explicit, such that our most general form of the ellipticity distribution caused by \mathcal{P} reads as

$$p_{\mathcal{P}}(\epsilon'|\epsilon, \theta, \nu) = \int d\mathcal{P} \epsilon'(\mathcal{P}|\epsilon, \theta) p(\mathcal{P}|\epsilon, \theta, \nu). \quad (2)$$

We can now define when we consider an ellipticity measurement to be *unbiased* by \mathcal{P} , namely if

$$\langle p_{\mathcal{P}}(\epsilon'|\epsilon, \theta, \nu) \rangle = \epsilon, \quad (3)$$

where the average is taken over independent noise realizations of images of identical galaxies, parametrized by (ϵ, θ) . In Sections 2 and 3 we are going to inspect cases, where, for different reasons, biases occur for fixed ϵ , whereas in Section 4 we are going to discuss the application of statistics to samples of noisy ellipticity estimates, whose true values are drawn from an underlying distribution $p(\epsilon)$. We conclude in Section 5.

2 NON-LINEAR ERROR PROPAGATION

An object's ellipticity is necessarily a non-linear quantity since any definition needs to invoke a ratio of the two key parameters of the geometric ellipse, its semi-major and semi-minor axes. As with any parameter that depends non-linearly on the data, even a symmetric distribution of noise for each data points translates into much more complicated, in general asymmetric and skewed, distribution of the ellipticity. Refregier et al. (2012) describe this generic problem

specifically for the case of model-dependent galaxy shape measurements and work out the bias on size and ellipticity of the best-fitting model that occurs even if the functional form of the galactic shape is perfectly known. We extend their finding by providing a theoretical form of the distribution of noisy ellipticity estimates that works for model- and moment-based approaches. The key ingredient of the derivation is the understanding that any attempt to measure the ellipticity is affected by the spurious ellipticity pattern of the noise realization recorded in the image, which is entirely describable by its moments, whereas spatial models for arbitrary noise configurations are not meaningful.

The second-order moments of the light distribution $I(\mathbf{x})$, centred at $\bar{\mathbf{x}}$,

$$Q_{ij} \equiv \int d^2x I(\mathbf{x})(x_i - \bar{x}_i)(x_j - \bar{x}_j), \quad (4)$$

give rise to a number of ellipticity estimators, two of which are in widespread use,

$$\chi \equiv \frac{Q_{11} - Q_{22} + 2iQ_{12}}{Q_{11} + Q_{22}} \quad \text{and} \quad (5a)$$

$$\epsilon \equiv \frac{Q_{11} - Q_{22} + 2iQ_{12}}{Q_{11} + Q_{22} + 2\sqrt{Q_{11}Q_{22} - Q_{12}^2}}. \quad (5b)$$

We review their advantages and disadvantages in Appendix A. In summary, ϵ is, at least theoretically, an unbiased estimator of the shear, while χ has a much more favourable distribution under noise, but requires higher-order corrections when used as shear estimator. Since we seek to obtain a distribution of ellipticities under noise, we choose χ for the derivation, but continue to use ϵ when referring to the proper ellipticity of an object.

In the presence of noise equation (4) needs to be modified,

$$Q'_{ij} = \int d^2x W(\mathbf{x} - \bar{\mathbf{x}})[I(\mathbf{x}) + n(\mathbf{x})](x_i - \bar{x}_i)(x_j - \bar{x}_j), \quad (6)$$

where W is a centred weight function and n is the (uncorrelated) noise term, which, in the background-dominated limit of faint objects, can be assumed to be drawn from an uncorrelated Gaussian distribution with variance σ_n^2 ,

$$n \sim \mathcal{N}(0, \sigma_n^2), \quad \langle n(\mathbf{x})n(\mathbf{x}') \rangle = \sigma_n^2 \delta(\mathbf{x} - \mathbf{x}'). \quad (7)$$

Even though weight functions are only explicitly present in moment-based approaches, model-based approaches effectively weigh pixels according to their distance and the shape of the employed model, so they too make use of a weight function.

Since moments are linear in the image data, they inherit the Gaussian error distribution from n . By the same token, the algebraic sum of moments defining the numerator and the denominator in equation (5a) are Gaussian distributed. But what about their ratio? It is basic knowledge in statistics that the ratio t of two uncorrelated Gaussian variates with mean of zero and unit variance, $\mathcal{N}(0, 1)$, is distributed according to the Cauchy distribution,

$$\mathcal{C}(t) = \frac{1}{\pi(1+t^2)}. \quad (8)$$

But for the ellipticity-measurement problem, the combination of moments defining χ do not have zero mean nor are they uncorrelated. It is obvious that at least the denominator of χ does not vanish for a source with non-negative brightness distribution. Also, the origin of the correlation quickly becomes apparent, when we choose a frame such that $|\chi| = \chi_1 > 0$, which can always be realized by a suitable rotation. We can write the mapping of (Q_{11}, Q_{22}) on to new

variables (w, z) , the denominator and numerator of χ_1 , as a linear operator

$$\mathbf{M} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \quad (9)$$

The covariance matrix of (w, z) is then given by

$$\begin{aligned} \mathbf{S}_{w,z} &= \mathbf{M}\mathbf{S}_{11,22}\mathbf{M}^T = \begin{pmatrix} \sigma_{11}^2 + \sigma_{22}^2 & \sigma_{11}^2 - \sigma_{22}^2 \\ \sigma_{11}^2 - \sigma_{22}^2 & \sigma_{11}^2 + \sigma_{22}^2 \end{pmatrix} \\ &\equiv \begin{pmatrix} \sigma_w^2 & \rho\sigma_w\sigma_z \\ \rho\sigma_w\sigma_z & \sigma_z^2 \end{pmatrix}, \end{aligned} \quad (10)$$

where $\mathbf{S}_{11,22} = \text{Diag}(\sigma_{11}^2, \sigma_{22}^2)$ denotes the diagonal covariance matrix of the moments Q_{ii} . Thus, the correlation between w and z ,

$$\rho = \frac{\sigma_{11}^2 - \sigma_{22}^2}{\sigma_w\sigma_z}, \quad (11)$$

only vanishes if the variances of the two moments Q_{11} and Q_{22} are equal, i.e. for circular galaxies. However, if the weight function W is adjusted to match the apparent shape of the galaxy, then the variances will generally be different. In our case with $\chi_1 > 0$, the pixels along the 1-direction (the semi-major axis) have a larger weight than those perpendicular to it. It follows that $\sigma_{11} > \sigma_{22}$. Thus, the variance of both w and z are driven by σ_{11} , and they become more and more correlated the larger χ_1 gets.

But this picture is not yet entirely complete because for equation (10) we assumed the moments Q_{11} and Q_{22} to be uncorrelated, but for any given image they are determined by the same noise realization. Imagine a perfectly noise-free image, where we add a positive noise fluctuation in only one pixel. According to equation (6), both Q'_{11} and Q'_{22} will then be larger than their noise-free counterparts since their distance-weighting factors $(x_i - \bar{x}_i)^2$ are positive or zero. Effectively, the errors of these two moments are not independent. As we are unaware of a way to compute the expected correlation theoretically, we measured it off images that contained only noise, given by the Gaussian noise model of equation (7). The result is shown in Fig. 1. We can see that there is a non-vanishing correlation $\rho_n \approx 0.325$, which, surprisingly, does not depend on the ellipticity of the weight function. Additional tests showed that this correlation is largely unaffected by changes in the size and radial




Figure 1. Correlation coefficient ρ_n between noisy moments Q'_{11} and Q'_{22} (black line) and between the moment combinations $z = Q'_{11} - Q'_{22}$ and $w = Q'_{11} + Q'_{22}$ (red line) as a function of the ellipticity of the Gaussian weighting function. Due to $\rho_n \approx 0.325$, the errors of w and z are different (their ratio: blue line). Dots indicate the measurement of these quantities in numerical tests with 10 000 noise realizations for each galaxy image. Variation in the size and radial profile of the weight function amounted to sub-per-cent changes in the quantities shown above.

profile of the weight function. The consequence of this correlation is a modification of equation (10),

$$\mathbf{S}_{w,z} = \begin{pmatrix} \sigma_{11}^2 + \sigma_{22}^2 + 2\rho_n\sigma_{11}\sigma_{22} & \sigma_{11}^2 - \sigma_{22}^2 \\ \sigma_{11}^2 - \sigma_{22}^2 & \sigma_{11}^2 + \sigma_{22}^2 - 2\rho_n\sigma_{11}\sigma_{22} \end{pmatrix}, \quad (12)$$

which does not alter ρ but the variances of w and z . The ratio σ_z/σ_w is also shown in Fig. 1, together with ρ . It is remarkable how strongly correlated the two moment combinations become even at modest ellipticities. Again, this result shows only marginal changes under variation of weighting function size or radial profile.

We are now equipped with variances and correlation of w and z and want to know the distribution of the ratio $t = \frac{z}{w}$. Marsaglia (1965, 2006) proved that the ratio of correlated Gaussian variates $w \sim \mathcal{N}(\mu_w, \sigma_w)$ and $z \sim \mathcal{N}(\mu_z, \sigma_z)$ with correlation coefficient ρ is given by

$$p_M(t) = r f(r(s - t)), \quad (13)$$

with constants defined as

$$r = \frac{\sigma_w}{\sigma_z\sqrt{1-\rho^2}} \quad \text{and} \quad s = \rho\frac{\sigma_z}{\sigma_w}. \quad (14)$$

The function $f(\tau)$ describes the probability distribution of $\frac{a+x}{b+y}$, with $x, y \sim \mathcal{N}(0, 1)$:

$$f(\tau) = \frac{e^{-\frac{1}{2}(a^2+b^2)}}{\pi(1+\tau^2)} \left(1 + \frac{\pi}{2} q e^{\frac{1}{2}q^2} \text{Erf} \left(\sqrt{\frac{1}{2}q} \right) \right), \quad (15)$$

where

$$q = \frac{b+a\tau}{\sqrt{1+\tau^2}}, \quad a = \frac{\mu_z/\sigma_z - \rho\mu_w/\sigma_w}{\sqrt{1-\rho^2}}, \quad \text{and} \quad b = \frac{\mu_w}{\sigma_w}. \quad (16)$$

We will refer to equation (13) as the *Marsaglia distribution*.⁵ For the ellipticity distribution we only have to substitute

$$\mu_w = Q_{11} + Q_{22}, \quad \mu_z = Q_{11} - Q_{22} \quad (17)$$

and to take variances and correlation from equation (12).

A non-vanishing correlation $\rho > 0$ has important consequences, foremost $s > 0$ and thus a shift of the peak of the ellipticity distribution towards higher values of χ . In Fig. 2, we show the distribution for a fixed ellipticity, i.e. fixed ρ , as a function of the image significance ν , varied in steps of 1 mag.⁶ We can see that even for fairly bright galaxies, a substantial shift occurs, which grows with decreasing ν . Less elliptical galaxies exhibit a weaker but still noticeable shift of the peak.

There are two other remarkable features of the Marsaglia distribution. First, it is not only shifted but in general also skewed (this is true even for $\rho = 0$) with a long tail towards lower values of $|\chi|$. It is thus not obvious whether the expectation value $\langle p_M(\chi|\epsilon, \theta, \nu) \rangle$ deviates from χ , constituting a bias according to equation (3). The top panel of Fig. 3 shows the integral over the entire distribution minus the true value of χ as a function of ϵ and different values of ν (solid lines). As expected, there is almost no bias for bright objects (black solid line), even up to large ellipticities. But with increasing noise level, the bias becomes initially negative (long tail dominates) before it turns positive (shift of the peak dominates). That means, the simple fact that the ellipticity is related to the image data in a non-linear way constitutes a bias of remarkable amplitude and non-trivial behaviour.

⁵ Unsurprisingly, in the case of uncorrelated variates $w, z \sim \mathcal{N}(0, 1)$, the first term of equation (15) recovers the Cauchy distribution of equation (8).

⁶ For the entire paper, we use the definition of ν from Erben et al. (2001).




Figure 2. The Marsaglia distribution (equation 13) of the ellipticity χ of a galaxy with true ellipticity $\chi = 0.88$ (equivalent to $\epsilon = 0.6$, indicated by the vertical dotted line) as a function of image significance ν . The correlation strength $\rho = 0.993$ is characteristic of this value of ϵ (cf. equation 11 and Fig. 1). The hatched region indicates the unphysical but possible range of outliers with $|\chi| \geq 1$.




Figure 3. Top: ellipticity bias $b(\chi)$, defined as the difference between the mean of the Marsaglia distribution of equation (13) and the ellipticity χ , as a function of the proper ellipticity ϵ for different significance levels (solid lines). The restriction of the integral to within $|\chi| = 1$ leads to stronger biases (dashed lines). Bottom: the fraction of measurements with $|\chi| > 1$.

Similar findings, namely a shift of the peak and a skewed distribution, are reported by Refregier et al. (2012) and Kacprzak et al. (2012), but our derivation does not require us to adopt a model-fitting approach or Gaussian likelihoods. It is thus not restricted to a particular shape of the galaxy and can also deal with convolutions with arbitrary PSF shapes as long as the process can be described in terms of moments (cf. Melchior et al. 2011). It is, however, not entirely obvious how to quantitatively compare to their distribution of maximum-likelihood estimators of ϵ .

Secondly, the Marsaglia distribution it is not bound by $|\chi| \leq 1$. In fact, it inherits the wide Cauchy-type wings and is thus capable of generating outliers with unbounded errors. Still, most errors are comparatively small such that outliers most often originate from galaxies with large initial ellipticities, rendering the outlier fraction strongly ellipticity dependent (bottom panel Fig. 3). With increasing noise level, ever smaller ellipticities become possible outliers. The exclusion of outliers from the integral over the distribution $p_M(\chi)$

leads to much stronger and mostly negative biases (dashed lines in the top panel) as the positive impact of the shifted peak becomes limited. How measurement codes and shear statistics respond to such outliers will be a reoccurring topic in the remainder of this work.

3 MEASUREMENT-RELATED ELLIPTICITY BIASES

As we laid out above, an ellipticity measurement cannot provide an unbiased result, simply because the error distribution is shifted and skewed for any ellipticity $|\epsilon| > 0$. Instead of requiring a shape-measurement method to yield unbiased ellipticity estimates, we should rather require that it reproduces the theoretically expected noisy distribution. In this section, we argue that in general not even that is possible because the attempt to measure a shape necessarily contributes its own uncertainties and, for non-linear shape parameters such as the galaxy size, its own biases (Refregier et al. 2012).

Even though the biases arising from the limited amount of information extractable from a noisy image are highly specific to the measurement method employed, there are general problems affecting each method in a similar way. In this section, we highlight the most relevant of these problems and seek to describe their level of systematic contamination of the ellipticity estimates. The tests are all carried out with the moment-based method DEIMOS (Melchior et al. 2011), but great care has been taken to ensure that our approach and conclusions we draw from the tests can be generalized to other methods.

The test galaxy images follow the Sérsic radial profile (Sersic 1968) with intrinsic parameters, including the ellipticity, taken from model fits to galaxies in the GEMS *Hubble Space Telescope* survey (Häussler et al. 2007, see Fig. 4). The galaxies are convolved with circular PSFs of Moffat-type (Moffat 1969). We choose two different PSF widths to mimic ground- and space-based conditions. The pixel noise follows equation (7) with two levels of pixel noise, the first one corresponds to optimistic weak-lensing conditions of $\nu = 35$, defined as in Erben et al. (2001, equation 16 therein), while the second one is 1 mag fainter ($\nu = 15$). In both cases, the images




Figure 4. Ellipticity distribution of GEMS galaxies from Häussler et al. (2007). Galaxies in the specified magnitude range are only considered if GALFIT was able to fit a single-Sérsic model. For moderate ellipticities, the distribution is well described by the Rayleigh function (solid line), the expected distribution for the absolute value of the two-dimensional, or complex, ellipticity if each ellipticity component followed a common Gaussian distribution (here: $\sigma_e = 0.33$). The measured distribution lacks highly elliptical galaxies, either because galaxies with large ellipticity are not as abundant in nature as predicted by the Rayleigh distribution or because their measurement is more difficult (Section 4.1 gives a possible explanation).

exhibit prominent pixel noise, but the galaxies remain clearly detectable. The ground-based PSF has a Moffat-index of 9, and the galaxy resolution factor $R_2 \approx 0.4$ (defined in Hirata et al. 2004, equation 8). The space-based PSF has a Moffat-index of 3, and the resolution is $R_2 \approx 0.7$.

3.1 Centroid errors

Determining the accurate centroid $\bar{\mathbf{x}}$ of an object is crucial for any further shape analysis. This is obvious for moment-based approaches, for which the ellipticity definitions given in equation (5) are sensible only if the centroid is chosen such that the dipole moment

$$D_i \equiv \int d^2x I(\mathbf{x})(x_i - \bar{x}_i) = 0 \quad \text{for } i \in \{1, 2\}, \quad (18)$$

which therefore needs to be enforced by such methods.

Analogously, model-fitting approaches for galaxies or stars rely on models with peaked light distributions and finite support. Due to the limited amount of information in the image data, such models are often derived from a radial profile $p(r)$, whose radial coordinate undergoes an ellipticity transformation, rendering the light distribution axisymmetric. The radial coordinate is expressed relative to the centroid or the peak position of the light distribution,

$$r = \left| \begin{pmatrix} 1 - \epsilon_1 & -\epsilon_2 \\ -\epsilon_2 & 1 + \epsilon_1 \end{pmatrix} (\mathbf{x} - \bar{\mathbf{x}}) \right|, \quad (19)$$

such that the estimation of the model parameters ϵ_i explicitly depends on the estimation of $\bar{\mathbf{x}}$.

With non-vanishing pixel noise, any measured centroid position is to some degree inaccurate, giving rise to an error $\Delta_{\bar{\mathbf{x}}}$ in the measured location. It is useful to limit the discussion to perfectly elliptical shapes, i.e. galaxies whose isophotes have the same centre, orientation and ellipticity. By doing so, we can introduce polar coordinates, $\Delta_{\bar{\mathbf{x}}} \rightarrow (r_c, \phi_c)$, and rotate again into a frame with only one non-vanishing ellipticity component (cf. Fig. 5). Then, we can identify the centroid error with \mathcal{P} in equation (2),

$$p_c(\epsilon' | \epsilon, \theta, \nu) = \int_0^\infty dr_c \int_{-\pi/2}^{\pi/2} d\phi_c \epsilon'(r_c, \phi_c | \epsilon, \theta) p(r_c, \phi_c | \epsilon, \theta, \nu). \quad (20)$$

The top panel of Fig. 6 shows the distribution of centroid errors under noise. In general, centroid errors for an elliptical light distribution show a similarly elliptical distribution, i.e. errors along the semi-major axis are larger than along the semi-minor axis where




Figure 5. Sketch of a perfectly elliptical ‘galaxy’ affected by a positive noise fluctuation (red area), causing a centroid offset of length r_c at an angle ϕ_c from the semi-major axis.

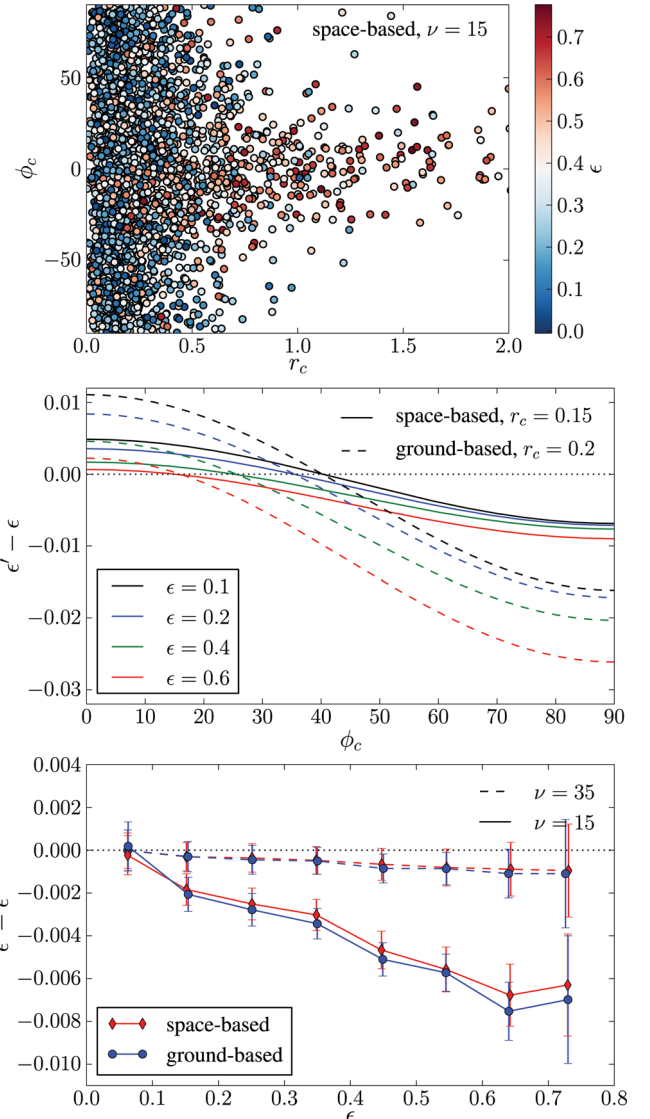


Figure 6. Top: distribution of centroid offsets in polar (r_c, ϕ_c) coordinates for well-resolved galaxies with morphologies from the GEMS survey in images with high levels of pixel noise. Middle: ellipticity error as a function of the miscentering angle ϕ_c for a noise-free galaxy image. The offsets r_c were set to the median values of the high-noise ($\nu = 15$) simulations. Bottom: ellipticity error as a function of intrinsic ellipticity for the same ensemble used in the top panel, seen with two different noise levels (dashed and solid lines) and with space-based or ground-based resolution (red and blue lines). Error bars denote 1σ errors of the mean ellipticity in each bin. The scatter at the high-ellipticity end is mainly driven by the lower number of simulated galaxies (cf. Fig. 4).

the light distribution has a larger gradient (e.g. Lewis 2009). In the polar coordinate frame this translates into

$$p(\phi_c, r_c) \propto r_c \exp\left(-\frac{r_c^2}{a^2} \cos^2(\phi_c) - \frac{r_c^2}{b^2} \sin^2(\phi_c)\right), \quad (21)$$

where we assumed an elliptical Gaussian distribution for $p(\Delta_{\bar{\mathbf{x}}})$ with semi-major axis a and semi-minor axis b . Remarkable about this distribution is the enhanced probability of finding centroid errors with small ϕ_c , and that this preferential alignment with the semi-major axis is most prominent for galaxies with large ellipticities. For this plot we chose the well-resolved galaxies with the high

noise setting such as to emphasize the alignment of the centroid errors with the semi-major axis. For a ground-based instrument the distribution $p(\phi_c)$ would be much less peaked around zero for any r_c , and larger values of r_c would be more likely. At lower noise, the distribution would be shifted towards smaller values of r_c .

To assess the impact of miscentring, we artificially shift the centroid position assumed in our shape-measurement code, and record the resulting deconvolved ellipticity in absence of any pixel noise. The middle panel of Fig. 6 illustrates the impact of miscentring on a typical disc-like galaxy as a function of the angle ϕ_c . Shown are the ellipticity estimates for four different intrinsic ellipticities, with either space-based resolution (solid curves) or ground-based resolution (dashed curves). Several aspects of this plot are worth mentioning: centroid shifts with small angles ϕ_c , i.e. aligned with the semi-major axis, shift the flux of the peak to non-vanishing distances without altering the orientation, thus leading to an increase in the ellipticity. On the opposite end, centroid shifts with large ϕ_c change the perceived orientation away from the actual one, thus lowering the inferred ellipticity. Furthermore, while over- and underestimate are fairly balanced for small ϵ , galaxies with large ellipticity suffer much more strongly from miscentring. Finally, the effects are stronger for the ground-based case since the PSF deconvolution amplifies any ellipticity signal, both the true and the spurious one. This is additionally enhanced by the larger offsets r_c encountered for the ground-based resolution at fixed ν .

The shape of the miscentring curves can be approximately described by

$$|\epsilon'| \approx [|\epsilon| + \epsilon_0 - \epsilon_{90}] \cos^2(\phi_c) + \epsilon_{90}, \quad (22)$$

where the parameters $\epsilon_{0,90} \geq 0$ depend on r_c (which in turn depends on ν) and ϵ , as well as other source parameters θ , most notably the width of the PSF, and the method employed.

In the bottom panel of Fig. 6 we show the ellipticity error induced by miscentring. We can see that the overall effect of the miscentring is a small, but consistent underestimation of the inferred ellipticity, which scales linearly with ϵ . This is not surprising, since large ellipticities suffer more strongly from miscentring. This effect is partially compensated by the preferential alignment of the centroid with the semi-major axis such that the centroid errors are not entirely isotropic and the average error bias is smaller than what would naively be expected when only considering the middle panel of Fig. 6. Qualitatively, this result is in good agreement with the derivation by Bernstein & Jarvis (2002, see their section 8.2 and equation 8.1) in that the bias depends linearly on ϵ and scales approximately as ν^{-2} . However, we only find a mild dependence on the image resolution R .

3.2 Misalignment

In noisy images, methods that employ either an elliptical model or – in the case of moment-based approaches – an elliptical weight function are subject to random errors in the determination of the orientation. Like before, we choose a coordinate system aligned with the semi-major axis of the elliptical source and introduce the misalignment angle ϕ_m (cf. Fig. 7). Identifying misalignment as the conceptual problem in equation (2) leads to

$$p_m(\epsilon'|\epsilon, \theta, \nu) = \int d\phi_m \epsilon'(\phi_m|\epsilon, \theta) p(\phi_m|\epsilon, \theta, \nu), \quad (23)$$

where $\epsilon'(\phi_m|\epsilon, \theta)$ denotes the ellipticity measurement performed with a potentially erroneous orientation. As indicated in Fig. 7, we expect the measured shape of the object to be biased towards




Figure 7. Sketch of a perfectly elliptical ‘galaxy’ (solid) affected by a positive noise fluctuation (red area), which gives rise to a misalignment of an elliptical model or weight function with respect to the true orientation by an angle ϕ_m (dotted).

smaller sizes and ellipticities. Estimates with smaller sizes were indeed reported by Kacprzak et al. (2012, fig. 1 therein) for model-fitting approaches.

In Fig. 8 we show – from top to bottom – the measured distribution of misalignment angles ϕ_m from noisy images, the impact of misalignment on the ellipticity estimates in noise-free images and the net effect of misalignment on the inferred ellipticity as a function of the true ellipticity. For these tests, ϕ_m denotes the orientation of the weight function used for the measurement, but since the PSF is circular, the orientation is identical to the actual orientation of the galaxy. That means $\epsilon'(\phi_m)$ is obtained by rotating the correctly matched elliptical weight function by the angle ϕ_m .⁷ Two essential findings can be made with this test. Pixel noise leads to an error on the orientation inferred from the image data, which becomes smaller with increasing ellipticity. This is not surprising since a circular source has a uniform distribution of orientation angles, while increasing the ellipticity renders the determination of the orientation easier, even in noisy images. The second finding is that for any $|\phi_m| > 0$ an underestimation of ϵ is observed with a characteristic shape that approximately follows the relation

$$|\epsilon'| \approx (|\epsilon| - \epsilon_{90}) \cos^2(\phi_m) + \epsilon_{90}, \quad (24)$$

with some ϵ_{90} that depends on the parameters θ of the apparent galactic shape. Combining both findings, we obtain a low bias induced in noisy images caused by misalignment. Looking at the bottom panel of Fig. 8, we can see that the effects of misalignment are more severe for modestly elliptical galaxies, where we observe an underestimation also in the moderate noise case with $\nu = 35$. This is an immediate consequence of the poor alignment constraints for galaxies with small $|\epsilon|$, and as misalignment is dependent on the apparent, i.e. convolved, ellipticity, ground-based imaging is more prone to this kind of bias.

3.3 Method-related limitations

So far, we have dealt with the problem of pixel noise assuming that a method for estimating the ellipticity is employed, whose validity does not become questionable in presence of noise. Unfortunately, we cannot expect shape-measurement methods to be entirely indifferent to increasing amounts of noise.

⁷ In order to single out the effect of misalignment, we do not alter the size or the total ellipticity of the weight function, even though these changes would occur in practice.

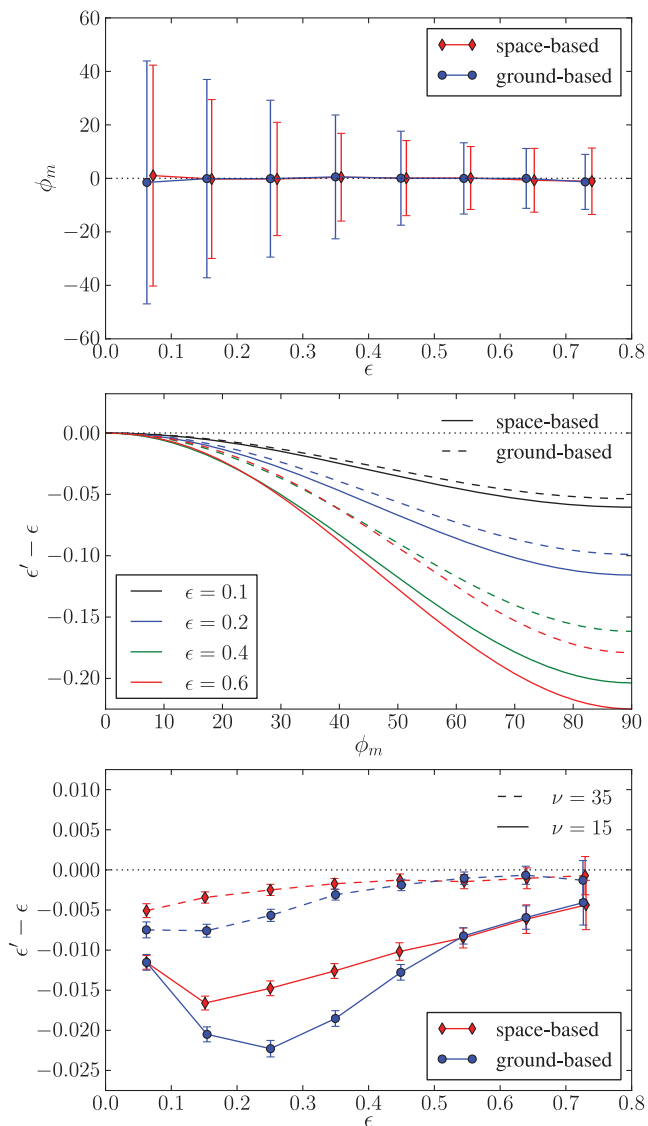


Figure 8. Top: width of the orientation error angle ϕ_m for the high noise level. For better visibility, the space-based data have been shifted horizontally by 0.01. Error bars denote 1σ intervals. Middle: error of the ellipticity when the weight function is rotated by the angle ϕ_m for four different galactic ellipticities. Bottom: error of the ellipticity in noisy images caused by misalignment as a function of galactic ellipticity. Error bars denote the 1σ errors of the mean in each ellipticity bin.

Moment-based methods, such as KSB (Kaiser, Squires & Broadhurst 1995), `HOLICS` (Okura & Futamase 2009) or `FDNT` (Bernstein 2010), need to apply a weight function to the data to suppress the impact of pixel noise at large distance from the source. The application of the weight function leaves an imprint on the measured moments, which can be corrected, albeit only approximately. With increasing amplitude of the noise, the only way to limit the variance of the ellipticity estimates is to shrink the weight function. Then, the approximation for the weight function correction becomes increasingly inaccurate, leading to errors in the deweighted moments. The direction and amplitude of the errors are a function of the apparent galaxy shape, in particular its slope, and properties of the PSF (e.g. Melchior et al. 2011, fig. 1 therein).

Model-fitting approaches do not need an artificial weight function as they make use of the compactness of their galaxy model,

which is often related to the Sérsic radial profile. One problematic aspect is the validity of this model – or model family – to faithfully describe the morphologies of all galaxies present in the observation. If data with a higher significance and possibly also a higher spatial resolution are available, the model assumptions can be verified, but this is often unfeasible. Bayesian approaches in model-fitting (e.g. `LENSEFIT`, Miller et al. 2007; Kitching et al. 2008) additionally employ priors on some parameters of the model, which are hard to estimate from the data alone. These priors are themselves derived from data of higher quality, and do not necessarily apply to the data at hand. But even if galaxies were purely elliptical – as many models assume – unbiased shear estimates require the radial profile to be accurately matched to the observed galaxies (Voigt & Bridle 2010), which cannot be guaranteed with images of severely limited significance.

Model assumptions could be avoided by using a decomposition into complete basis function sets, such as shapelets (Refregier 2003) or Sérsiclets (Ngan et al. 2009). However, the pixel noise limits the number of modes used in the fit such that the resulting model becomes dominated by the shape of the zeroth order. In the case of the circular shapelet basis function set, this introduces a bias towards circular objects (Melchior et al. 2010). For Sérsiclets, which can be considered a generalization of shapelets, using a finite number of modes leads to a relation between the slope of the radial profile and the spatial scale of higher-order fluctuations, which is not necessarily obeyed by observed galaxies (Andrae, Melchior & Jahnke 2011).

In summary, any method that deals with severely degraded data invokes additional assumptions about the data, which may turn out to be wrong once an adequate assessment can be devised, e.g. with data of higher quality. It needs to be shown that the methods employed are able to meet requirements demanded to reach the scientific goals of the project at hand. Therefore, simulations with simplified galaxy models (e.g. Heymans et al. 2006; Bridle et al. 2010; Kitching et al. 2012) provide a clean way of comparing several methods, but should be complemented by simulations with realistic galaxy morphologies (such as Massey et al. 2004; Meneghetti et al. 2008; Mandelbaum et al. 2012).

Additional problems arise from the occurrence of outliers, i.e. $|\epsilon'| \geq 1$. From the discussion in Section 2 and by looking at Fig. 3, we know that outliers must occur if the combination of apparent ellipticity and pixel noise exceeds some value. Moment-based methods will just let these outliers pass (unless they try to avoid them by shrinking the weight function) and expect a subsequent lensing analysis to deal with them. We revisit this problem in Section 4.1. On the other hand, by construction model-fitting methods cannot have such outliers, but they will encounter catastrophic events, such as convergence failures, and flag these objects, which again excludes them from the analysis later on.

Alternatively, model-fitting approaches might invoke a prior on the ellipticity, which could in principle completely prevent the occurrence of outliers or modelling failures. However, this comes at the price of essentially recovering the prior for data with very little constraining power on the ellipticity. As we said above, this prior does not necessarily describe the actual data accurately. Moreover, even if the prior is an accurate description of the ensemble ellipticity distribution, the fact that at fixed noise level outliers preferentially occur for galaxies with large ellipticity leads to a stronger impact of the prior on those galaxies. Since large ellipticities are less abundant than those around $|\epsilon| = 0.3$ (cf. Fig. 4), we expect a bias towards the prior which grows with galaxy ellipticity.

For years, attempts have been made to calibrate the bias away based on simulated images, e.g. by introducing global ‘fudge’

factors that boost the measured ellipticities (e.g. Heymans et al. 2006, table A1 therein). More recently, sophisticated supervised learning methods have been employed to correct for the bias as a function of several input parameters (Gruen et al. 2010; Kacprzak et al. 2012; Tewes et al. 2012). Although they acknowledge and account for the presence of biases in the measurements, irrespective of their origin, they hinge on quality, size and representative nature of the training set. We thus regard these methods to work well for galaxies that are abundant in the training data (cf. Kitching et al. 2012) and whose properties accurately resemble those of actually observed galaxies.

4 INADEQUATE SHEAR STATISTICS

We now turn to a different type of systematic problem caused by the pixel noise, which does not occur at the level of individual ellipticity estimates but later on, namely when statistics of the measured ellipticity distribution are calculated. Pixel noise can render these statistics difficult to interpret or entirely inappropriate by violating their fundamental assumptions. One such assumption in virtually every lensing analysis is that a proper definition of ellipticity provides an unbiased estimator of the (reduced) shear,

$$\langle \epsilon \rangle \equiv \int d^2\epsilon \epsilon p(\epsilon) = g. \quad (25)$$

For a noise-free measurement of the ellipticity in the form of equation (5b), this is in fact the case (Seitz & Schneider 1997; Bartelmann & Schneider 2001).

For what follows, we are going to assume the optimistic scenario, in which the measurement methods provide ellipticity estimates, which follow directly the Marsaglia distribution of equation (13). That means all complications discussed in Section 3 are eliminated. In practice, this can be realized by using the sampling method outlined in Appendix B, which provides a realistic ellipticity distribution of Gaussian-shaped galaxies under noise but is not affected by shape-measurement biases. We recall the two main distinctions between the two popular moment-based ellipticity estimators from Appendix A: ϵ has a problematic distribution under noise (cf. Fig. A1), while χ depends on the shear in a non-linear way. Both features will prove to be problematic.

4.1 Outliers and their rejection

As we have stressed several times now, the pixel noise can lead to ellipticity outliers with $|\epsilon'| \geq 1$ (cf. Fig. 3) and thus gives rise to a second population of ellipticity measurements, namely those on the unit circle in ϵ -space. If the resulting sample is naively inserted in equation (25), the presence of these outliers, whose position on the ring is still loosely correlated with their noise-free location, should lead to an overestimation of the inferred shear, simply because they all have unit ellipticity and thus large impact on the mean of the distribution. Weirdly, this is *not* observed in the top-left panel of Fig. 9, where we show the mean ellipticity as a function of shear and noise level. If the shear only has one non-vanishing component, the population on the unit circle almost perfectly compensates the underestimation we expect from the mean of the Marsaglia distribution (cf. Fig. 3). This balance is delicate, and it is quite possible that it is not maintained for galaxies with a different radial profile, where the means and errors of the measured moments deviate from our Gaussian calculation. But even if the balance persisted, statistics other than the mean, e.g. the two-point correlation function, will in general pick up the presence of outliers

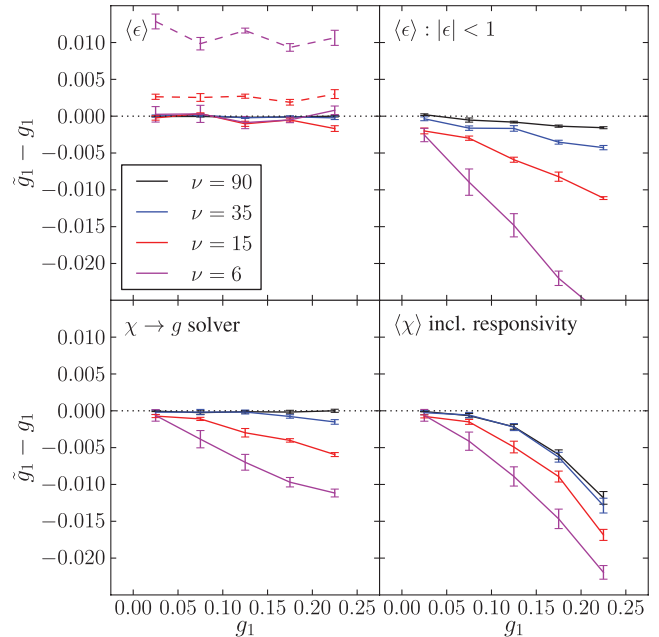


Figure 9. Performance of shear statistics under pixel noise. Top left: average of the entire ϵ distribution, following equation (25). Top right: same as in the left-hand panel, but after rejection of outliers with $|\epsilon| \geq 1$. Bottom left: non-linear solver for the shear from measurements of χ , based on equation (27). Bottom right: linearized relation between χ and the shear, following equation (A2). For all panels, shear was only applied on one-direction (solid lines; dashed lines in the top-left panel: $g_2 = 0.1$), and the ellipticity noise was of Marsaglia-type, simulated with the algorithm described in Appendix B. Means and errors are taken from 10 independent noise realizations at each value of the shear, each realization comprised 10000 samples.

and perform erratically. Another concerning aspect of $\langle \epsilon \rangle$ as shear estimator becomes apparent when the other shear component comes into play. Because the noise affects the outliers by altering their phase on the unit circle, the estimate \tilde{g}_1 becomes dependent on g_2 (dashed lines). The cross-talk between the two shear components is relevant for the faintest two noise settings, where the outliers constitute a significant portion of the entire sample.⁸

On the other hand, excluding these outliers would be absolutely justified since any ellipticity definition needs to be bound by 1, otherwise the ratio of semi-minor to semi-major axis is non-sensical. Unfortunately, this commonly adopted approach also leads to biases because we now sample from

$$p_o(\epsilon' | \epsilon, \theta, \nu) = \int_{|\epsilon+n|<1} d^2n (\epsilon + n) p(n | \epsilon, \theta, \nu). \quad (26)$$

Even if $p(n)$ were isotropic and had zero mean, the truncation of the integration range at the ellipticity unit circle poses a bias for all galaxies with sufficiently large ellipticity. If $p(n)$ is monotonically decreasing with increasing $|n|$ as it is common, this bias will be negative because the outlier-excluded distribution is more compact than the actual noisy distribution. In other words, this negative bias increases with ellipticity or, equivalently, with the probability of obtaining a noise contribution that can push the measured ellipticity

⁸ The other estimators shown in Fig. 9 do not suffer significantly from this sort of cross-talk since they either do not have outliers (top right panel) or the outlier population has a well-behaved shape (both estimators based on χ). Hence, we only show the effect for $\langle \epsilon \rangle$.

beyond the unit circle. Consequently, sampling from $p_o(\epsilon'|\epsilon)$ rather than from $p(\epsilon)$ in equation (25) results in a low bias on \mathbf{g} since any coherent distortion of the ellipticity distribution will increase the probability of falling outside of the unit circle for all galaxies, whose intrinsic ellipticity is aligned with the additional distortion. Hence, this distortion becomes suppressed in the outlier-rejected sample. This is shown in the top-right panel of Fig. 9. The bias becomes more prominent with increasing shear since large ellipticities become more abundant and the entire distribution thus more likely to create outliers at any non-vanishing noise level.

It is important to note that even though we have discussed the outlier problem solely in terms of moment-based measurements, shear estimates from model-fitting methods are susceptible to the same bias: any filtering on catastrophic modelling failures is equivalent to outlier rejection, and if such failures become more prominent with increasing source ellipticity, the shape of the bias curves will follow the one in the top-right panel of Fig. 9.

4.2 Non-linear statistics

Using χ rather than ϵ provides the advantage of a much simpler noise distribution with a closed description: the Marsaglia distribution of equation (13). However, the relation between χ and \mathbf{g} is non-linear. As before, we are thus faced with error propagation in a non-linear system and have to expect biased shear estimates even from a perfect measurement.

The theoretically correct way of obtaining shear estimates from the measurements of χ is solving for the shear \mathbf{g} that nulls the mean of the source-plane ellipticity (Schneider & Seitz 1995),

$$\chi^s = \frac{\chi - 2\mathbf{g} + \mathbf{g}^2 \chi^*}{1 + |\mathbf{g}|^2 - 2\Re(\mathbf{g} \chi^*)} \quad (27)$$

as the distribution of unlensed galaxy ellipticities is assumed to be isotropic, i.e. with zero mean. To our knowledge, the bottom-left panel of Fig. 9 shows the first application of this estimator. While unbiased for all shears at zero noise, and being still unbiased for $\nu = 35$ and $g_1 < 0.15$, the solver exhibits negative bias for the two faintest noise settings.

In practice, equation (27) is approximated to first order in the shear, which leads to equation (A2) and the usage of the so-called responsivity correction. Unsurprisingly, this simplified relation introduces its own bias that grows with the shear, which is shown in the bottom-right panel of Fig. 9. Moreover, this simplified estimator does not perform any better than the fully non-linear solver in the left panel. In fact, our novel estimator appears to cope well with outliers in the noisy ellipticity distribution and to perform more reliably than the other three. It has the highest computational demands as it involves the minimization of an objective function, the modulus of χ^s in equation (27), but can be implemented efficiently.

4.3 Ellipticity weights

Often, ellipticity estimates do not directly enter equation (25), but get weighted before. Such a weighted average is tempting for two reasons. First, one can reduce the noise-induced variance by penalizing faint galaxies. Secondly, one might even be able to reduce systematic biases by down-weighting galaxies from regimes, where the employed method yields consistently wrong ellipticity estimates. For instance, as the outlier problem increases with apparent ellipticity (cf. Section 4.1) while the responsivity to shear becomes weaker, applying larger weights for less elliptical galaxies seems logical.

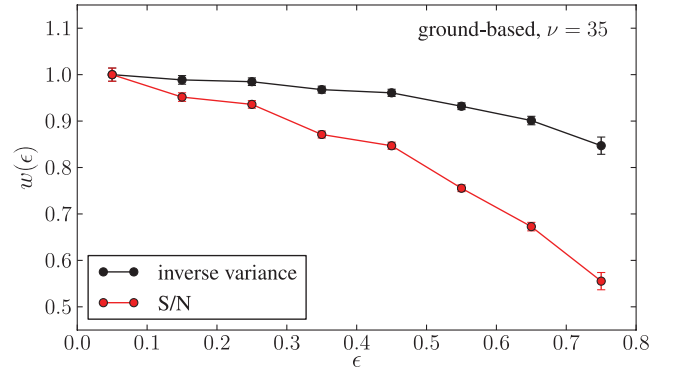


Figure 10. Ellipticity weights: *inverse variance* denotes the scheme that attempts to minimize the mean square error, considering the ellipticity-dependent measurement noise σ_n and the intrinsic scatter σ_e , according to $w \propto [\sigma_n^2 + \sigma_e^2]^{-1}$ (Hoekstra, Franx & Kuijken 2000). *S/N* denotes a scheme where the weight is directly proportional to the measurement significance, which often favours circular objects. Measurement errors and *S/N* are obtained from the DEIMOS method, and the intrinsic dispersion was assumed to be $\sigma_e = 0.3$.

But a weighting scheme can introduce a bias on its own, even for perfectly unbiased, noise-free ellipticity estimates, namely in this seemingly beneficial case of ellipticity-dependent weights. If we replace the theoretically unbiased average $\langle \epsilon \rangle$ from equation (25) by its weighted and properly normalized equivalent

$$\langle w\epsilon \rangle = \frac{\int d^2\epsilon w\epsilon p(\epsilon)}{\int d^2\epsilon w p(\epsilon)}, \quad (28)$$

we can immediately see that the result is the same, i.e. unbiased, only if w does not depend on ϵ . If the weight decreases with increasing ellipticity, we are confronted with an altered and more compact distribution

$$p_w(\epsilon') = w(\epsilon')p(\epsilon'), \quad (29)$$

similar to the case of the outlier-rejected distribution. Even if we do not explicitly want the weighting scheme to penalize large ellipticities, the measurement method might report statistics of the effective signal-to-noise ratio level ν that are turned into weights. Due to the more difficult task to measure large ellipticities, these statistics tend to implicitly depend on the source ellipticity (cf. Fig. 10 for two plausible weighting schemes). For the simple case of an approximately linear relation of the weight on ϵ , described by the intercept w_0 and the slope c , we show in Appendix C that the resulting bias on the shear estimate is to the first order purely multiplicative,

$$\langle w\epsilon \rangle - \langle \epsilon \rangle \approx \mathbf{g} \sqrt{\frac{2}{\pi}} \frac{c}{w_0} \sigma_e, \quad (30)$$

where σ_e denotes the dispersion of the ellipticity estimates. With reasonable values of $\frac{c}{w_0} \approx -0.1$, we obtain a bias of -2.5 per cent.

Also the two-point correlation function of (post-lensing) ellipticities

$$\langle \epsilon_i \epsilon_j \rangle(\theta) = \xi_g(\theta), \quad (31)$$

which averages over pairs of galaxies with separation $\theta = |\mathbf{x}_i - \mathbf{x}_j|$, is expected to be sensitive to an ellipticity-dependent weighting scheme. Relevant for the estimation of cosmological parameters are the amplitude and shape of the shear correlation function $\xi_g(\theta)$. A straightforward calculation (outlined in Appendix C1) shows that the normalized correlation function is given by

$$\frac{\langle w_i \epsilon_i w_j \epsilon_j \rangle}{\langle w_i w_j \rangle} \approx \frac{c^2}{w_0^2} [\sigma_e^2 + \sigma_g^2]^2 + \xi_g(\theta) + \frac{c^2}{w_0^2} \xi_g^2(\theta) \quad (32)$$

Table 1. Overview of the ellipticity and shear estimation biases described in this work.

| Description | Section | Effect | References |
|---------------------------------|-------------|-----------------------------------------------------------------------------------------------------------------------|---------------------------|
| Error propagation | Section 2 | Noise distribution shifted and skewed, average biased low, strongly ellipticity-dependent | This work |
| Centroiding error | Section 3.1 | Bias $\propto \epsilon/\nu^2$ | Bernstein & Jarvis (2002) |
| Misalignment | Section 3.2 | Bias most problematic for intermediate ellipticities | This work |
| Various shape estimation issues | Section 3.3 | Deviation of measured ellipticity distribution from its expectation | See Section 3.3 |
| Ellipticity outliers | Section 4.1 | Cross-talk between shear components (with outliers) or multiplicative shear underestimation (after outlier rejection) | This work |
| Non-linear shear statistics | Section 4.2 | Multiplicative shear underestimation | This work |
| Ellipticity-dependent weighting | Section 4.3 | Multiplication shear underestimation, negligible in correlation function ξ_g | This work |

where σ_g^2 denotes the variance of the shear field. Since the constant terms in the above equation can be determined by looking at separations θ where we expect the cosmological lensing signal to vanish, the correlation function is, surprisingly, largely unaffected by the weighting scheme. Only at very small scales, it is steepened due to the last term in equation (32), but we expect this effect to be clearly subdominant compared to the influence of baryonic physics at these small scales.

In contrast to the outlier rejection bias – which corresponds to a binary weight: either 1 or 0 – this bias is less severe for large ellipticities, but it generally affects all galaxies, including those with smaller ellipticities, for which the creation of outliers is not a significant issue. Consequently, whenever weighted ellipticities are inserted into statistics that have been derived for unweighted ellipticities, correction terms (such as equation 30) need to be applied, which requires accurate knowledge of the ellipticity-dependence of the weighting scheme.

5 SUMMARY, CONCLUSIONS AND OUTLOOK

We compiled theoretical and practical evidence that pixel noise biases shear estimates at three subsequent levels: (1) the propagation of the pixel noise into the non-linear quantity ellipticity; (2) additional uncertainties and biases introduced by ellipticity-measurement methods and (3) the application of statistics to infer the gravitational shear from a sample of ellipticity measurements, which are either unaware of the presence of pixel noise or themselves non-linear and thus biased.

We summarize our findings in Table 1. In practice, lensing analyses are affected by a combination of the aforementioned biases, most of which are negative. Considering these findings, it is not surprising that every investigation of any shear estimation methodology the authors are aware of shows a weakening response to shear with increasing noise level. It is inevitable. Even though the details differ, the biases we revealed or reinvestigated generally not only depend on the object significance but also depend on its ellipticity. We therefore recommend to extend shear accuracy test programs to also inspect trends with ellipticity.

That biases occur at several stages of the analysis pipeline leads to an unfortunate interdependence of the idiosyncrasies of the image data (foremost the galaxy ellipticity and signal-to-noise distributions), shape-measurement method and shear statistic. This means that a setup which has been found to work well in one situation does not necessarily perform so well in others. Given the scope and fundamental nature of the majority of these biases, we do not believe that the needs of upcoming lensing surveys in terms of accuracy

and reliability can be met without a substantial effort in correcting for or avoiding biased ellipticity and shear estimates.

Method-dependent biases can be studied with simulated images. Special attention should be paid to cases where the measured ellipticity distribution deviates from the expected Marsaglia distribution. For instance, a pile-up of samples shortly before the unit circle is a clear indication of a bias introduced by the shape-measurement code in order to prevent unphysical outliers.

To correct biases in an actual measurement, one needs to be able to identify those parameters that determine the bias. As we argued, this is mainly the significance and the ellipticity, but other factors, such as the radial profile or changes of the ellipticity with radius (Bernstein 2010), will also play a role. Therefore, correction schemes need to fully consider the performance of the shape-measurement method as a function of all relevant parameters as well as the fact that these parameters themselves can only be inferred with a certain precision from the image data (Kacprzak et al. 2012). This is computationally and practically challenging.

Instead of correcting ellipticity estimates, which also does not completely eliminate the problem of noise for the shear statistic, we should seek solutions that avoid the biases, for instance by staying linear in the data for as long as possible (exemplified by stacking methods in Bridle et al. 2010). We conclude by sketching out an alternative idea. Since we worked out the theoretical form of the noisy ellipticity distribution, we are able to predict the (biased) outcome of a measurement given an assumed galaxy ellipticity and noise level. As we showed in Section 3, we can also model the errors shape-measurement methods exhibit once we assume to know the underlying galaxy parameters. We can thus ask the question: how likely is a certain ellipticity given a measured one and its measurement errors. The result will be an ellipticity likelihood for each galaxy that incorporates the non-linearities of the measurement process. Combining full likelihoods should finally lead to unbiased shear estimates. Whether this idea works in practice remains to be seen.

To facilitate the review of our findings and to support forthcoming development, we make the code used in this work public. The PYTHON implementation comprises the computation of the Marsaglia distribution with efficient sample generation as well as all shear estimators of Section 4 and is available at <https://github.com/pmelchior/epsnoise>.

ACKNOWLEDGMENTS

PM is supported by the U.S. Department of Energy under Contract No. DE-FG02-91ER40690. MV is supported by a STFC Rolling Grant RA0888.

REFERENCES

- Amara A., Refregier A., 2008, MNRAS, 391, 228
 Andrae R., Melchior P., Jahnke K., 2011, MNRAS, 417, 2465
 Bartelmann M., Schneider P., 2001, Phys. Rep., 340, 291
 Bernstein G. M., 2010, MNRAS, 406, 2793
 Bernstein G. M., Jarvis M., 2002, AJ, 123, 583
 Bridle S. et al., 2010, MNRAS, 405, 2044
 Erben T., Van Waerbeke L., Bertin E., Mellier Y., Schneider P., 2001, A&A, 366, 717
 Gruen D., Seitz S., Koppenhoefer J., Riffeser A., 2010, ApJ, 720, 639
 Häussler B. et al., 2007, ApJS, 172, 615
 Heymans C. et al., 2006, MNRAS, 368, 1323
 Hirata C. M. et al., 2004, MNRAS, 353, 529
 Hoekstra H., Franx M., Kuijken K., 2000, ApJ, 532, 88
 Huterer D., Takada M., Bernstein G., Jain B., 2006, MNRAS, 366, 101
 Kacprzak T., Zuntz J., Rowe B., Bridle S., Refregier A., Amara A., Voigt L., Hirsch M., 2012, preprint (arXiv:1203.5049)
 Kaiser N., Squires G., Broadhurst T., 1995, ApJ, 449, 460
 Kitching T. D., Miller L., Heymans C. E., van Waerbeke L., Heavens A. F., 2008, MNRAS, 390, 149
 Kitching T. D. et al., 2012, preprint (arXiv:1202.5254)
 Lewis A., 2009, MNRAS, 398, 471
 Mandelbaum R., Hirata C. M., Leauthaud A., Massey R. J., Rhodes J., 2012, MNRAS, 420, 1518
 Marsaglia G., 1965, J. Am. Stat. Assoc., 60, 193
 Marsaglia G., 2006, J. Stat. Software, 16, 1
 Massey R., Refregier A., Conselice C. J., Bacon D. J., 2004, MNRAS, 348, 214
 Massey R. et al., 2007, MNRAS, 376, 13
 Melchior P., Böhnert A., Lombardi M., Bartelmann M., 2010, A&A, 510, A75
 Melchior P., Viola M., Schäfer B. M., Bartelmann M., 2011, MNRAS, 412, 1552
 Meneghetti M. et al., 2008, A&A, 482, 403
 Miller L., Kitching T. D., Heymans C., Heavens A. F., van Waerbeke L., 2007, MNRAS, 382, 315
 Moffat A. F. J., 1969, A&A, 3, 455
 Ngan W., van Waerbeke L., Mahdavi A., Heymans C., Hoekstra H., 2009, MNRAS, 396, 1211
 Okura Y., Futamase T., 2009, ApJ, 699, 143
 Refregier A., 2003, MNRAS, 338, 35
 Refregier A., Kacprzak T., Amara A., Bridle S., Rowe B., 2012, preprint (arXiv:1203.5050)
 Schneider P., Seitz C., 1995, A&A, 294, 411
 Seitz C., Schneider P., 1997, A&A, 318, 687
 Sersic J. L., 1968, Atlas de galaxias australes. Observatorio Astronomico, Cordoba, Argentina
 Tewes M., Cantale N., Courbin F., Kitching T. D., Meylan G., 2012, preprint (arXiv:1203.4429)
 Viola M., Melchior P., Bartelmann M., 2011, MNRAS, 410, 2156
 Voigt L. M., Bridle S. L., 2010, MNRAS, 404, 458
 Weinberg D. H., Mortonson M. J., Eisenstein D. J., Hirata C., Riess A. G., Rozo E., 2012, preprint (arXiv:1201.2434)

APPENDIX A: MOMENT-BASED ELLIPTICITY DEFINITIONS

In equation (5) we have introduced two particular forms of the complex ellipticity, which are related according to (Bartelmann & Schneider 2001)

$$\chi = \frac{2\epsilon}{1 + |\epsilon|^2} \quad \text{or} \quad \epsilon = \frac{\chi}{1 + \sqrt{1 - |\chi|^2}}. \quad (\text{A1})$$

Since the only significant difference between the two definitions occurs in the denominator, they share the same complex phase, but have different amplitude. There are two important distinctions

between these definitions. First, the relation between gravitational shear \mathbf{g} and mean ellipticity is

$$\mathbf{g} = \langle \epsilon \rangle \quad \text{or} \quad \mathbf{g} = \frac{\langle \chi \rangle}{2 + \sigma_\chi^2} + \mathcal{O}(\langle \chi^3 \rangle), \quad (\text{A2})$$

where the averages are taken over galaxies affected by a constant shear \mathbf{g} . That means, ϵ is an unbiased estimator of the shear (Seitz & Schneider 1997), while χ needs a so-called *responsivity* correction (the denominator in the equation above), and even then it remains an approximate estimator of \mathbf{g} (see Viola, Melchior & Bartelmann 2011 for a discussion of higher-order corrections to this relation).

The reason for the wide-spread use of χ as shear estimator rather than ϵ stems from their second distinction, the distribution under pixel noise. If we assume the pixel noise to be uncorrelated and Gaussian, the second-order moments share this property due to their linearity in the data. However, both ϵ and χ are ratios of combinations of second-order moments. In the case of χ , numerator and denominator are linear combination of moments and thus still follow Gaussian distributions such that their ratio follows the Marsaglia distribution of equation (13), a generalization of the Cauchy distribution (equation 8). The Cauchy distribution is known for its diverging variance, which allows arbitrarily large errors with finite probability, namely when $Q_{11} + Q_{22} \rightarrow 0$. In contrast, due to the non-linear term $\sqrt{Q_{11}Q_{22} - Q_{12}^2}$ in the denominator of equation (5b), ϵ has a much more complicated distribution. Instead of allowing infinite errors, the additional term in the denominator can lead to a complex phase, which alters the orientation of ϵ in the cases where χ would come to lie outside the unit circle of viable ellipticities. This effectively couples the two components of ϵ and leads to two separate distributions, the ordinary one within the unit circle and the one on the unit circle, which renders a theoretical description much more difficult.

We show the different distributions of ϵ and χ under noise in Fig. A1. While χ' shows a continuous distribution across the edge of the unit circle, outlier ellipticities (red points) are located right on the unit circle in ϵ' -space. Performing a projection on to one component or computing the absolute value of ϵ' would thus lead to an enhanced probability of measurements close to the unit circle. Moreover, global statistics such as the mean or the variance of the ϵ' distribution are significantly altered by the presence of the ring at $|\epsilon'| = 1$.

In summary, while ϵ is theoretically an unbiased estimator of the shear, its distribution under noise effectively undermines this property. On the other hand, the distribution of χ is much easier to describe and statistics thereof are fairly robust against noise, but the interpretation of these statistics is hampered by the non-linear relation to shear. One therefore needs to choose the ellipticity estimator based on the application at hand, depending on what kind of drawback can most effectively be dealt with.

APPENDIX B: SAMPLING THE NOISY ELLIPTICITY DISTRIBUTION

We present the recipe to simulate a realistic ellipticity distribution, considering the effects of non-linear error propagation discussed in Section 2. Instead of sampling from the complicated Marsaglia distribution of equation (13), we follow the path of the pixel noise, i.e. we compute the means and errors for the moment combinations $w = Q_{11} + Q_{22}$ and $z = Q_{11} - Q_{22}$ and then their ratio $\chi = \frac{z}{w}$.

For a fast and analytic moment calculation, we assume an elliptical Gaussian shape for both galaxy and/or weight function and rotate into a frame, such that its semi-major axis is aligned with the




Figure A1. Ellipticity distributions of measured χ' (top) and measured ϵ' (bottom) from a synthetic catalogue of 5000 objects with intrinsic dispersion $\sigma_e = 0.3$, constant shear $\mathbf{g} = (0.15, 0.05)$, indicated by red cross markers, and pixel noise equivalent to $\nu = 15$. The circle in the top plot corresponds to the limit of intrinsic $|\epsilon| < 1$, and the red points in the bottom plot indicate galaxies with $|\epsilon'| \geq 1$. The distributions are created with the algorithm outlined in Appendix B.

one-direction,

$$W(\mathbf{x}) = \exp \left[\frac{-(1-\epsilon)^2 x_1^2 - (1+\epsilon)^2 x_2^2}{2s^2} \right]. \quad (\text{B1})$$

The size of the galaxy is then defined by s and its flux $F = \int d^2x W(\mathbf{x})$. Errors of the moments, measured with weight function W , are given by (Melchior et al. 2011)

$$\sigma_{i,j}^2 = \sigma_n^2 \int d^2x W^2(\mathbf{x}) x_1^i x_2^j, \quad (\text{B2})$$

where σ_n^2 denotes the pixel noise variance and i, j describe a moment $\{W\}_{i,j} = \int d^2x W(\mathbf{x}) x_1^i x_2^j$. Note that this notation differs from the notation used throughout the rest of this work. In the case of a Gaussian-shaped W , these errors can be analytically evaluated. The algorithm can trivially be extended to work on moments from arbitrarily shaped galaxies, measured from noise-free images. With

the convolution relation in moment-space (Melchior et al. 2011, equation 9 therein), also the effect of the PSF convolution can be taken into account.

(1) For a source with ellipticity ϵ , we compute the modulus $\epsilon = |\epsilon|$, its equivalent χ from equation (A1), and the phase $\phi = \arg(\epsilon)$.

(2) For a Gaussian shape, the scale s defines the sum of second moments, $w = Fs^2$. Then, $z = \chi w$.

(3) We adopt the definition of significance from Erben et al. (2001, equation 16 therein) and insert the error of the flux F (the $\{W\}_{0,0}$ moment in the notation used above): $\nu = \frac{F}{\sigma_n \sqrt{\pi s}} \sqrt{(1+\epsilon)(1-\epsilon)}$. Thus, if we specify ν we get the pixel noise level σ_n .

(4) The Gaussian errors σ_{ij} for the second moments Q_{ij} are then given by

$$\begin{aligned} \sigma_{11}^2 &= \sigma_{2,0}^2 = \sigma_n^2 \frac{3\pi}{4} \frac{s^6}{(1-\epsilon)^5(1+\epsilon)} \\ \sigma_{12}^2 &= \sigma_{1,1}^2 = \sigma_n^2 \frac{\pi}{4} \frac{s^6}{(1-\epsilon)^3(1+\epsilon)^3} \\ \sigma_{22}^2 &= \sigma_{0,2}^2 = \sigma_n^2 \frac{3\pi}{4} \frac{s^6}{(1-\epsilon)(1+\epsilon)^5}. \end{aligned} \quad (\text{B3})$$

(5) For the errors ΔQ_{11} and ΔQ_{22} , we sample from a correlated bi-variate Gaussian, whose covariance matrix is given by

$$\mathbf{S}_{11,22} = \begin{pmatrix} \sigma_{11}^2 & \rho_n \sigma_{11} \sigma_{22} \\ \rho_n \sigma_{11} \sigma_{22} & \sigma_{22}^2 \end{pmatrix} \quad (\text{B4})$$

[cf. discussion that led to equation (12)]. This can be realized by applying the Cholesky decomposition

$$\mathbf{S}_{11,22} = \mathbf{A}\mathbf{A}^T \rightarrow \mathbf{A} = \begin{pmatrix} \sigma_{11} & 0 \\ \rho_n \sigma_{22} & \sigma_{22} \sqrt{1 - \rho_n^2} \end{pmatrix} \quad (\text{B5})$$

to a vector of two $\mathcal{N}(0, 1)$ variates.

(6) Then, noisy samples $w' = w + \Delta Q_{11} + \Delta Q_{22}$ and $z' = z + \Delta Q_{11} - \Delta Q_{22}$ exhibit the correct variances and correlation given by equation (12), and their ratio $\chi'_1 = \frac{z'}{w'}$ is distributed according to the Marsaglia distribution of equation (13).

(7) The second component of χ has an independent numerator $\Delta Q_{12} \sim \mathcal{N}(0, \sigma_{12}^2)$, but the same denominator: $\chi'_2 = \frac{\Delta Q_{12}}{w'}$.

(8) The original orientation is recovered by $\chi' \rightarrow \chi' e^{i\phi}$. Then, ϵ' can be obtained from equation (A1).

Since the only terms entering the Marsaglia distribution are ratios of moments and errors, and both scale as s^2 , we can set $s = 1$. Similarly, the significance ν depends on the ratio of flux F and noise dispersion σ_n , so that we can also set $F = 1$ without changing the results. Effectively, one only has to specify the ellipticity ϵ and the significance ν to uniquely describe the effects of noise on the ellipticity. All information about the galactic shape, which determines the moments of the brightness distribution and their errors, is then internally computed (assuming a Gaussian radial profile).

Realistic distributions, such as the ones shown in Fig. A1, additionally need to start from a decent intrinsic distributions of ellipticities ϵ' . With the common assumption of each ellipticity component being drawn from an independent Gaussian distribution of some dispersion $\sigma_e \approx 0.3$, their modulus is drawn from the Rayleigh distribution $f(\epsilon) = \frac{\epsilon}{\sigma_e^2} e^{-\epsilon^2/2\sigma_e^2}$ (cf. Fig. 4 for an observed ellipticity distribution), and the orientation is uniform, $\mathcal{U}(0, \pi)$. Finally, a sheared ellipticity distribution can be obtained from (Seitz &

Schneider 1997)

$$\epsilon = \frac{\epsilon^s + \mathbf{g}}{1 + \epsilon^s \mathbf{g}^*}. \quad (\text{B6})$$

We implemented the entire sampling procedure (as well as the theoretical form of the Marsaglia distribution and the shear estimators described in Section 4) in PYTHON. The code is open-source and available at <https://github.com/pmelchior/epsnoise>.

APPENDIX C: WEIGHTING-SCHEME INDUCED BIAS

We seek an approximate analytical description of the bias an ellipticity-dependent weighting scheme introduces on globally acting statistics such as the average and the two-point correlation function. Therefore, we choose a coordinate frame, which is aligned with the semi-major axis of each galaxy, such that we only have to consider the one component of ϵ . We start by Taylor-expanding the ellipticity dependence of the weighting scheme to the first order,

$$w(\epsilon') = w_0 + \left. \frac{\partial w(\epsilon)}{\partial \epsilon} \right|_{\epsilon'} \epsilon' + \mathcal{O}(\epsilon'^2). \quad (\text{C1})$$

Assuming the most simple form of a positive offset $w_0 > 0$ and small, constant slope $\partial w(|\epsilon|)/\partial |\epsilon| = c$, we can simplify the previous equation to

$$w(\epsilon) = w_0 \begin{cases} +c\epsilon' & \text{if } \epsilon > 0 \\ -c\epsilon' & \text{if } \epsilon < 0. \end{cases} \quad (\text{C2})$$

Now we split the integrals in equation (28) into the lower and upper half (which allows us to apply the reduced weights also to negative ϵ):

$$\begin{aligned} \int d\epsilon w(\epsilon)\epsilon p(\epsilon) &= w_0 \int_{-\infty}^0 d\epsilon \epsilon p(\epsilon) - c \int_{-\infty}^0 d\epsilon \epsilon^2 p(\epsilon) \\ &\quad + w_0 \int_0^{\infty} d\epsilon \epsilon p(\epsilon) + c \int_0^{\infty} d\epsilon \epsilon^2 p(\epsilon) \end{aligned} \quad (\text{C3})$$

and likewise for the denominator. Since there is no sign-flip in the terms with w_0 , we can combine these integrals again and exploit $\int d\epsilon p(\epsilon) = 1$ and $\int d\epsilon \epsilon p(\epsilon) = g$. For the c -terms, we need to carry out the integration over the ellipticity distribution explicitly, simply

because the presence of the shear shifts and skews the distribution such that the c -terms in the equation above do not exactly cancel. We therefore assume the pre-lensing distribution to be Gaussian with dispersion σ_ϵ and the shear to only shift the entire distribution without changing its shape: $p(\epsilon) \rightarrow \mathcal{N}(\epsilon - g, \sigma_\epsilon)$. This assumption will restrict our derivation to small shears. We thus linearize the resulting integrals to the first order in the shear and obtain

$$\begin{aligned} \int_{-\infty}^0 d\epsilon \epsilon \mathcal{N}(\epsilon - g, \sigma_\epsilon) - \int_0^{\infty} d\epsilon \epsilon \mathcal{N}(\epsilon - g, \sigma_\epsilon) &\approx -\sqrt{\frac{2}{\pi}} \sigma_\epsilon \\ \int_{-\infty}^0 d\epsilon \epsilon^2 \mathcal{N}(\epsilon - g, \sigma_\epsilon) - \int_0^{\infty} d\epsilon \epsilon^2 \mathcal{N}(\epsilon - g, \sigma_\epsilon) &\approx -2\sqrt{\frac{2}{\pi}} \sigma_\epsilon g. \end{aligned} \quad (\text{C4})$$

Inserting all terms into equation (28) yields equation (30). The result is accurate to the first order in $w(\epsilon)$, g and $\frac{c}{w_0}$.

C1 Shear correlation function

In the limit of small shear, we can simplify equation (B6) between observed and pre-lensing ellipticities, $\epsilon \rightarrow \epsilon^s + \mathbf{g}$, which corresponds to the above case: the shear only shifts the ellipticity distribution, but does not skew it. When considering the weighting scheme from equation (C2), we get the correlation functions of the weights

$$\langle w_i w_j \rangle(\theta) = \langle (w_0 + c\epsilon_i)(w_0 + c\epsilon_j) \rangle = w_0^2 + c^2 \langle \epsilon_i \epsilon_j \rangle(\theta) \quad (\text{C5})$$

and of the weighted ellipticities

$$\langle w_i \epsilon_i w_j \epsilon_j \rangle(\theta) = w_0^2 \langle \epsilon_i \epsilon_j \rangle(\theta) + c^2 \langle \epsilon_i^2 \epsilon_j^2 \rangle(\theta), \quad (\text{C6})$$

where $\langle \epsilon_i \epsilon_j \rangle$ is given by equation (31) and

$$\langle \epsilon_i^2 \epsilon_j^2 \rangle(\theta) = \sigma_\epsilon^4 + 2\sigma_\epsilon^2 \sigma_g^2 + \sigma_g^4 + 2\xi_g^2(\theta) \quad (\text{C7})$$

with $\sigma_g^2 \equiv \langle \mathbf{g}_i^2 \rangle$ being the variance of the shear field. In this derivation we assumed the intrinsic ellipticity field to be uncorrelated for separation $\theta > 0$ and the shear field to be uncorrelated with the intrinsic ellipticity field, such that terms like $\langle \epsilon_i \mathbf{g}_i \rangle$ and $\langle \epsilon_i \mathbf{g}_j \rangle$ vanish. Expanding the ratio of equations (C5) and (C6) to the first order in $\frac{c^2}{w_0^2}$ yields equation (32).

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.