Astronomy & Astrophysics

Problems using ratios of galaxy shape moments in requirements for weak lensing surveys

H. Israel[1,2,3], T. D. Kitching[4], R. Massey[1,2], and M. Cropper[4]

[1] Institute for Computational Cosmology, Durham University, South Road, Durham DH1 3LE, UK
e-mail: `holger.israel@durham.ac.uk`
[2] Centre for Extragalactic Astronomy, Durham University, South Road, Durham DH1 3LE, UK
[3] Universitätssternwarte München, Fakultät für Physik, Ludwig-Maximilians-Universität München, Scheinerstraße 1,
81679 München, Germany
[4] Mullard Space Science Laboratory, University College London, Holmbury St Mary, Dorking, Surrey RH5 6NT, UK

ABSTRACT

Context. The shapes of galaxies are typically quantified by ratios of their quadrupole moments. Knowledge of these ratios (i.e. their measured standard deviation) is commonly used to assess the efficiency of weak gravitational lensing surveys. For faint galaxies, observational noise can make the denominator close to zero, so the ratios become ill-defined.
Aims. Since the requirements cannot be formally tested for faint galaxies, we explore two complementary mitigation strategies. In many weak lensing contexts, the most problematic sources can be removed by a cut in measured size. This first technique is applied frequently. As our second strategy, we propose requirements directly on the quadrupole moments rather than their ratio.
Methods. As an example of the first strategy, we have investigated how a size cut affects the required precision of the charge transfer inefficiency model for two shape measurement settings. For the second strategy, we analysed the joint likelihood distribution of the image quadrupole moments measured from simulated galaxies, and propagate their (correlated) uncertainties into ellipticities.
Results. Using a size cut, we find slightly wider tolerance margins for the charge transfer inefficiency parameters compared to the full size distribution. However, subtle biases in the data analysis chain may be introduced. These can be avoided using the second strategy. To optimally exploit a Stage-IV dark energy survey, we find that the mean and standard deviation of a population of galaxies' quadrupole moments must to be known to better than 1.4×10^{-3} arcsec^2, or the Stokes parameters to 1.9×10^{-3} arcsec^2.
Conclusions. Cuts in measured size remove sources that otherwise make ellipticity statistics of weak lensing galaxy samples diverge. However, size cuts bias the source population non-trivially. Assessing weak lensing data quality directly on the quadrupole moments instead mitigates the need for size cuts. Such testable requirements can form the basis for performance validation of future instruments.

Key words. methods: data analysis – techniques: image processing – gravitational lensing: weak – large-scale structure of Universe

1. Introduction

Weak gravitational lensing is the deflection of light in the presence of a weak gravitational field along the line of sight between source and observer, because of the relativistic distortion of space-time. Galaxy images become sheared and amplified by a small amount given by second-order derivatives of the local projected gravitational potential. Shear information is encoded in the shape of a galaxy image, namely its ellipticity. While extremely noisy for the individual source, weak lensing becomes powerful with statistics, when wide or deep imaging data are available to measure the shapes of large numbers of galaxies to a high degree of accuracy and infer their ellipticities. Cosmic shear, i.e. weak lensing by the large-scale structure, provides a reliable probe of cosmology via model constraints on the two-point correlation function or power spectrum of the ellipticity field. However, to be unbiased the ellipticity measurements also need to be unbiased (e.g. Antonik et al. 2013; Chang et al. 2013a; Huterer & White 2002; Kitching et al. 2009; Cropper et al. 2013; Jarvis et al. 2016).

The requirements on the average bias in ellipticity and size measurements, over the ensemble of galaxies used in a weak lensing analysis, have been discussed by Paulin-Henriksson et al. (2008), Amara & Réfrégier (2008), Kitching et al. (2009),

Massey et al. (2013). In these papers, the parent requirement on the ellipticity measurements was broken-down into contributions from various instrumental and telescope effects, and into requirements on the measurement of the size of galaxies and stars.

Ellipticity is usually defined as the ratio of linear combinations of measured quadrupole moments of an image. Any observational measurements are noisy, and, through the central limit theorem, these moments follow Gaussian error distributions. However, ratios of Gaussian distributed variables do not have a simple distribution. The probability distribution of two correlated random variables with non-zero means is defined in Hinkley (1969), where it is also shown that the moments of this distribution diverge so that mean and variance are not defined.

Given certain conditions on numerator and denominator, a parameter transformation can be applied such that meaningful first and second moments can be computed (Marsaglia 2006, cf. Appendix B). However, these conditions are typically not fulfilled when measuring ellipticities, as the distribution in the denominator approaches or even crosses zero.

In this paper, we extend Israel et al. (2015, I15)'s analysis of the precision with which charge transfer inefficiency (CTI) can be corrected. The problem of divergent ratios makes it hard to set requirements on how high that precision must be, and hard to determine how high a precision was achieved. We consider two

possible methods to circumvent the divergent ratios. The overall performance of a CTI correction is not affected by recasting the requirements; rather it clarifies the way in which performance is measured, and the conditions that an experiment needs to meet.

Charge transfer inefficiency comprises the effects of temporary trapping and release of photoelectrons during CCD readout by defects in the detector material caused by radiation in space. CTI manifests itself by charge trails in the direction of the CCD readout, stemming from deferred electrons that after a duration specific to the lattice defect have been released from the charge traps created by irradiation. Image processing software can correct for CTI trails by iterative modelling of the capture and release processes, reshuffling signal counts to the pixel where they would have been registered in the absence of CTI (Massey et al. 2010, 2014; I15).

A sensitivity study was performed by I15, analysing the combined effects of readout noise and errors in the determination of the trap parameters (e.g. their densities and release timescales) on the accuracy with which measured galaxy parameters can be recovered, compared to the case of no CTI. Based on the scientific requirements of the *Euclid*[1] (Laureijs et al. 2011) survey, I15 derived tolerance margins for the uncertainties in the trap parameters.

One main purpose of this article is to address a complication that I15 did not take into account. So-called size cuts, that is a (lower) limit on the measured size of the objects to be further considered as likely weakly lensed sources are widely included in weak lensing analysis pipelines, in particular in moment based approaches, for example KSB (Kaiser et al. 1995; Erben et al. 2001), as implemented in, for example, Hoekstra (2007), Okabe & Umetsu (2008), Israel et al. (2010), Applegate et al. (2014), but also in model fitting algorithms (e.g., Zuntz et al. 2013; Jarvis et al. 2016).

Motivated mainly by the need to distinguish galaxies from stars, size cuts removing the smallest sources effectively mitigate the vanishing denominators problem, as they bar the most problematic objects from the analysis. This is the first of the two solutions we consider here. In terms of cosmic shear, because weak lensing statistics (see e.g. Bartelmann & Schneider 2001) are not sensitive to the galaxy selection function this is a good approach. A similar approach can be taken for stellar objects, as it is not necessary for all stars to be used in PSF modelling; only a sufficient number.

However, size cuts are not always viable. Problematic scenarios include the validation of telescope (or pipeline) performance on a set of galaxies whose selection is predefined, or the measurement of the distribution of the changes of sizes in objects caused by CTI itself (cf. I15). In the absence of any radiation damage to CCDs the size change caused by CTI is by definition zero. To solve the problem in this case, we propose setting requirements on (and analysing) the ensemble mean and error of the quadrupole moments rather than their ratio.

This is a critical issue in the design of future weak lensing experiments. The requirements for both the *Euclid* and Large Synoptic Survey Telescope (LSST, Chang et al. 2013b) weak lensing surveys are currently set on the measurement of galaxy ellipticities (and hence through ratios of quantities). In this paper, we show why these are not verifiable, due to their divergent moments. Instead we avoid this issue by suggesting that "top-level" requirements be recast on the quadrupole moments themselves. These requirements can then be propagated or "flown down" in a similar way to the process that has been followed for ellipticity

[1] http://euclid-ec.org

and size variables (Cropper et al. 2013). In fact the propagation and proportioning of these requirements into various components should be more straightforward as the effect on quadrupole moments is typically linear for both PSF and detector effects (Melchior et al. 2011). This second approach is conceptually similar to but distinct from the corrections for PSF ellipticity the KSB and RRG (Rhodes et al. 2001) algorithms performs on the moments.

We note that the ratio statistics problem has been known in the weak lensing literature (e.g. Refregier et al. 2012) – although it has been seldom addressed directly – and shape measurement algorithms avoiding the evaluation of ratios of stochastic quantities exist, in particular model fitting methods (cf. Miller et al. 2007, 2013; Sheldon 2014; Jarvis et al. 2016). However, the phrasing (or definition) of survey requirements must also be updated to reflect this knowledge.

In Sect. 2 we formally state the problem and show why the ellipticity denominator can be measured negative. In Sect. 3 we describe the methodologies used to study both the size cut and the recasting of requirements. Section 4 explores the impact of size cuts using the example of CTI correction. We then present the recasted requirements based on quadrupole moments. After discussing the findings in the context of *Euclid* CTI correction in Sect. 5, we conclude in Sect. 6.

2. Statement of the problem

2.1. Divergent terms in the requirement flowdown

The weak lensing requirement derivations for *Euclid* and LSST start with the measured quadrupole moments of a galaxy or stellar image $I(x_1, x_2)$,

$$Q_{ij} = \frac{\int dx_1 dx_2 (x_i - \bar{x}_i)(x_j - \bar{x}_j) I(x_1, x_2) W_\omega(x_1, x_2)}{\int dx_1 dx_2 I(x_1, x_2) W_\omega(x_1, x_2)} \tag{1}$$

where i and $j = \{1, 2\}$, and (x_1, x_2) is a Cartesian coordinate system. $W_\omega(x_1, x_2)$ is a weight function that is typically assumed to a be a multivariate Gaussian of scale ω. There are three quadrupole moments $\boldsymbol{Q} = \{Q_{11}, Q_{22}, Q_{12}\}$ that are therefore defined, and these can be related to the ellipticity of the object in question by

$$e = \chi = \chi_1 + i\chi_2 = \frac{Q_{11} - Q_{22} + 2iQ_{12}}{Q_{11} + Q_{22}} \tag{2}$$

the third eccentricity. Its denominator,

$$A = Q_{11} + Q_{22} \tag{3}$$

measures the size of an object (referred to as R^2 in other works). As Fig. 1 showing how negative values of A can be measured in real data illustrates, the R^2 nomenclature, which sounds positive definite, appears to be more appropriate for the alternative estimator $A' = (Q_{11}'^2 + Q_{22}'^2)/F^2$, with Q_{ij}' the numerator of Eq. (1), and F its denominator. However we note the oddity that A' is in units of angle to the fourth power. While A, in units of angle squared, can be understood as the solid angle subtended by the object, there is no similarly straightforward interpretation for A'.

Although other definitions, for example, $\sqrt{(Q_{11}Q_{22} - Q_{12}^2)}$ yield positive definite size estimators, we restrict our discussion to $Q_{11} + Q_{22}$, because of its relevance as the denominator of χ, as we will demonstrate. In the presence of a point spread function

(PSF) and detector effects the observed size and ellipticity transform as (see Appendix A and Bartelmann & Schneider 2001)

$$e_{\rm obs} \equiv e_{\rm gal} + \frac{A_{\rm PSF}}{A_{\rm gal} + A_{\rm PSF}}(e_{\rm PSF} - e_{\rm gal}) + e_{\rm NC} \tag{4}$$

and

$$A_{\rm obs} \equiv A_{\rm gal} + A_{\rm PSF} + A_{\rm NC} \tag{5}$$

where "PSF" refers to any convolutive effect and "NC" refers to any non-convolutive effect (for example due to CTI), "obs" refers to the observed quantity and "gal" denotes the (true) galaxy quantity that would be observed given no additional effects. Equations (4) and (5) correct Eqs. (31) and (32) of Massey et al. (2013, M13), and we detail the subsequent changes to the requirement flowdown model in Appendix A.

Cosmic shear analyses typically involve calculating correlations between galaxy shapes by taking the two-point correlation function of Eq. (4), and in taking the ensemble average (mean) of the derived expressions, requirements can be determined on each the convolutive and non-convolutive elements of an experiment design. This leads to expressions like $\langle \delta A/A \rangle$, where δ refers to a measurement uncertainty, upon which there is a requirement. This is particularly important for the non-convolutive effects where there are requirements that depend on quantities such as $\langle \delta A_{\rm NC}/A_{\rm obs} \rangle$ (M13; Cropper et al. 2013).

Because the quadrupole moments are sums over pixels, they are expected (through the central limit theorem) to be Gaussian distributed. Derived quantities that are linear combinations of the quadrupole moments, such as numerator and denominator (the size A) in Eq. (2), are also Gaussian distributed. However it is by taking the ratio in Eq. (2) that χ follows a distribution whose mean and variance formally and practically diverge. We note that χ will diverge for a general distribution of the quadrupole moments, independent of the choice of W_ω.

In the cosmic shear requirement flowdown (M13), it is the denominator in terms such as $\langle \delta A_{\rm NC}/A_{\rm NC} \rangle$ that causes the problem: if the distribution of the denominator crosses zero then the distribution of the ratio diverges. Because the mean of the ratio of two correlated variables is undefined it is therefore not formally possible to verify if quantities such as $\langle \delta A_{\rm NC}/A_{\rm obs} \rangle$ are being measured correctly – one can make the approximation $\langle \delta A_{\rm NC} \rangle/\langle A_{\rm obs} \rangle$ but then it is not possible to verify that this is a sufficient approximation.

2.2. Why objects with negative sizes A exist

Let us revisit the size measurement of simulated galaxies in I15, and analyse the role and impact of very small objects. Galaxies of negative measured size $A_{\rm obs}$ are problematic for two related reasons: as we have just seen, they make terms with the size in the denominator diverge (recall that $A_{\rm obs}$ itself is Gaussian distributed). Moreover, because the distribution of $A_{\rm obs}$ extends to negative values the Marsaglia (2006) mitigation technique (cf. Appendix B) cannot be applied to measure ellipticity statistics. From where do $A_{\rm obs} \leq 0$ sources arise?

In general, a measured Q_{ij} can be negative if an image is noisy. Consider the case that there is a local background $M(x_1, x_2)$, with mean $\langle M \rangle$ and noise about this mean, then we measure in the presence of this background

$$\tilde{Q}_{11} = \int [I(x_1, x_2) + M(x_1, x_2)]\, W_\omega(x_1, x_2)\, x^2\, {\rm d}x_1 {\rm d}x_2, \tag{6}$$
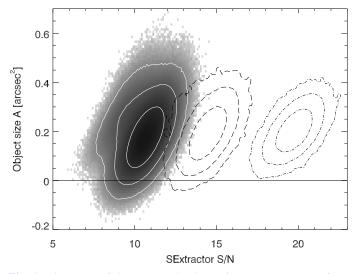

Fig. 1. Histogram of the measured values of $A_{\rm obs} = Q_{11} + Q_{22}$ from simulations of the CTI effect, as a function of the object signal-to-noise ratio. The logarithmic grey scale and white contours (enclosing 68.3%, 95.45%, and 99.73% of samples, respectively) show the 10^7 exponential disk galaxies analysed in I15. Input simulations containing Poisson distributed sky noise were subjected to CTI and Gaussian read-out noise. Then, the CTI was removed using the correct trap model. Dashed and dot-dashed contours (both smoothed) enclose the same density levels measured from 10^5 simulations with ~35% (~85%) higher signal-to-noise. In all three cases, the noise causes a non-negligible fraction of the galaxies to be measured with $A_{\rm obs} < 0$.

where $I(x_1, x_2)$ is the galaxy intensity. To obtain the galaxy's shape moments we would have to subtract the mean local moment

$$Q_{11} = \tilde{Q}_{11} - \int M(x_1, x_2)\, x^2\, {\rm d}x_1 {\rm d}x_2, \tag{7}$$

which can be zero or less then zero depending on how noisy the image and background is. In Eq. (7), the weight function $W_\omega(x_1, x_2)$ drops out because $\langle M \rangle$ is constant over object scales. In practice the mean background is usually subtracted leaving "negative" pixels in the data in a process referred to as "background subtraction". Therefore $A = Q_{11} + Q_{22}$ can be negative in practice.

Figure 1 shows an example from I15 of the distribution of measured sizes $A_{\rm obs}$, as a function of SExtractor (Bertin & Arnouts 1996) signal-to-noise (S/N) ratio. The Rhodes et al. (2001, RRG) shape-measurement algorithm used here subtracts a mean background before calculating image moments as just described. I15 applied the ArCTIC (Massey et al. 2010, 2014) algorithm for CTI correction in CTI-addition mode to each of the 10^7 exponential disk galaxy images used in Fig. 1, and then iteratively corrected the CTI trails using the same software and trap model. For a detailed account of our input simulation, we refer to I15.

Because ArCTIC restores the input simulation (sources convolved with a *Euclid* visual instrument PSF plus Poisson distributed sky noise) perfectly except for read-out noise added during the emulated CCD read-out, the distribution in Fig. 1 looks very similar to that of the input simulations. That means we would have recovered a similar distribution of size and signal to noise even in the absence of CTI. But we include it here for realism (including slightly correlated background noise due to the CTI correction). We choose the CTI-corrected images because the represent what can be measured from real observations.

While the measured size $A_{\rm obs}$ and S/N are correlated, the negative size objects in Fig. 1 do not represent the very low S/N end of the distribution in either a relative or an absolute sense. Indeed, in the I15 analysis pipeline, all of them are bona fide `SExtractor` detections, accounting for one in 128 (0.78%) of these galaxies sampling the faintest population to be included in the *Euclid* cosmic shear experiments. Increasing the S/N of the input simulations by ∼15% (∼35%), reduces the fraction of $A_{\rm obs} \leq 0$ galaxies to 0.50% (0.19%), but this is a slow drop-off (cf. dashed and dot-dashed contours in Fig. 1). Even in a simulation with mean $S/N \approx 20$, we still observe negative sizes for eight out of 10^5 samples, a tiny, but not negligible fraction in a *Euclid*-like survey. The tail on the low-S/N side of the distribution is an artifact of the RRG algorithm's choice of scale ω of the weight function W in Eq. (1).

We note that the results shown in Fig. 1 have been calculated with a specific choice of `SExtractor` centroids and a Gaussian weight function ($\langle\omega\rangle = 0''34$). We expect the results to depend quantitatively, but not qualitatively on these choices. In particular, the small value of ω compared to usual choices in the weak lensing literature already minimises the effect of noise on the measured moments. However, with only a few signal pixels, these sources are intrinsically vulnerable to noise.

We also note that while Fig. 1 uses the `SExtractor` S/N instead of a measurement based on the same moments used to compute A, this will not affect our conclusions: any two viable S/N estimators can be transformed into each other via a bijective mapping, such that, locally for the range of Fig. 1, the abscissa can be rescaled linearly into the reader's preferred S/N estimate.

Finally, one might argue that many shape measurement algorithms employ inverse variance weights to the galaxies in their shear catalogues. Because of the large uncertainties in the resulting shear estimates, $A < 0$ objects will preferentially get assigned smaller weights than the average galaxy. However, unless the $A < 0$ galaxies are given zero weight, this will not mitigate the problem. Our simulations take into account the ensemble scatter of the estimated $A = Q_{11} + Q_{22}$.

3. Methodology

Here, we present the two mitigation strategies we discuss. Sect. 3.1 describes the effect of size cuts on CTI correction as an example of the first strategy. Section 3.2 details how requirements can be recast in terms of the normally distributed quadrupole moments.

3.1. Removing negative size sources from CTI simulations

Although the $A_{\rm obs} \leq 0$ sources are legitimate objects for shape measurement, removing them from the catalogue by means of a size cut can solve some of the mathematical problems arising from the vanishing denominator in Eq. (2). In fact, most existing shear measurement algorithms impose a size cut at or above the size scale measured from observed PSF tracing stars, for practical purposes. However, the I15 sensitivity analysis did not consider such size cuts when translating the requirements on observables like $e_{\rm obs}$ into requirements on the accuracy and precision to which the parameters of ArCTIC, the CTI model, need to be determined by calibration. Instead, I15 maximised their sample statistics by taking into account the full distribution in $A_{\rm obs}$.

We repeat the analysis of I15 with an increasingly selective size cut on $A_{\rm obs}$ (that is on the size of objects after CTI has been first applied and then corrected), mimicking realistic conditions.

Because the Gaussian read-out noise that is added just after CTI has been applied, and is uncorrelated to the Poisson distributed sky noise in the input simulation, and

$$A_{\rm obs} = A_{\rm gal} + A_{\rm PSF} + A_{\rm sky\_noise} + A_{\rm NC} + A_{\rm read\_noise}, \qquad (8)$$

the sources $A_{\rm obs} \leq 0$ only rarely coincide with the sources $A_{\rm gal} + A_{\rm PSF} + A_{\rm sky\_noise} \leq 0$ in the input simulations. Indeed, our simulations allow us to trace that in the I15 setup and sample $A_{\rm NC} + A_{\rm read\_noise}$ are well fit by a Gaussian distribution of mean $0''013$ and standard deviation of $0''112$ (the I15 simulations used here and in Sect. 4.1 have a pixel scale of $0''1$/pixel).

3.2. How to link requirements to moments

As we will see in Sect. 4.1, making size cuts in weak lensing analyses can make requirements assessment more robust. However there are also negative impacts, most notably (1) the reduction in the number of galaxies; and (2) the introduction of a non-trivial relationship between the size cut, CTI correction, signal-to-noise, and the shape measurement method employed (Sect. 4.1.2). Thus we explore an alternative mitigation strategy.

In this approach we propose that instead of setting requirements on the ellipticity and size, requirements need to be set in the quadrupole moment space. We still propose that ellipticities are used for shear inference (using the quadrupole moments themselves is explored in Viola et al. 2014), and that only the requirements are set in the moment space.

The requirements we will set are on the accuracy with which the true distribution of moments needs to be known (that would have been observed in the absence of any systematic biases, in other words, the prior distribution of the quadrupole moments). We therefore start by measuring this distribution, about which perturbations can be made. To get a realistic fiducial baseline we measure this from data using the `GalSim` (Mandelbaum et al. 2012; Rowe et al. 2015) deconvolved sample of galaxies (we use all galaxies in this sample; for a full description of the magnitude range and other properties we refer to the `GalSim` papers), where we use a weight function for the moments that is an multivariate Gaussian with a FWHM of 20 pixels (the pixel scale is 0.2 arcsec); we find that the results are independent of the exact choice of this width since we take perturbations about the fiducial distribution. Throughout units of the quadrupole moments are in arcseconds squared unless otherwise stated.

Viola et al. (2014) show how a given measured set of quadrupole moment values $\mathbf{Q}$ and their measurement errors $\sigma_{\mathbf{Q}}$ relate to a probability distribution for ellipticity p_χ. This is known as the "Marsaglia-Tin" distribution (Marsaglia 1965, 2006; Tin 1965) and is the multivariate correlated case of the ratio distribution. The mean and maximum likelihood of the ellipticity probability distributions are biased, in a way that is dependent on the measured error on the quadrupole moments (or signal-to-noise of the observed galaxy image).

We define a measured quadrupole distribution $p_i(\mathbf{Q})$ for a galaxy i. In a Bayesian setting the distribution of the true quadrupole moments can be considered as a prior $\Pi(\mathbf{Q})$ from which the galaxy is drawn. That means the probability of measuring a value $\mathbf{Q}$ given some data D can be written like $p_i(\mathbf{Q}|D) \propto p(D|\mathbf{Q})\Pi(\mathbf{Q})$. The distribution $p_i(\mathbf{Q})$ can then be mapped into ellipticity via the Marsaglia-Tin distribution, arriving at the biases in the mean of the ellipticity components.
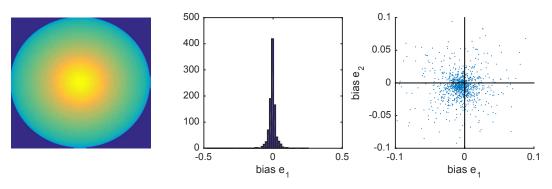
Fig. 2. Example of Monte-Carlo sampling of $\Pi(\boldsymbol{Q})$, with no perturbations (the fiducial case) propagated into p_ϵ. The *left-hand panel* shows the inferred prior in ϵ_1 (*x*-axis) and ϵ_2 (*y*-axis) calculated by taking the product of the individual ellipticity distributions. The *middle panel* shows the biases in ϵ_1 and the *right-hand panel* shows the distribution of biases in ϵ_1 and ϵ_2. Biases are in units of arcsec^2.

Doing so, it is advantageous to use here the third flatterning, an alternative ellipticity estimator

$$\epsilon_1 + i\epsilon_2 = \frac{Q_{11} - Q_{22} + 2iQ_{12}}{Q_{11} + Q_{22} + 2\sqrt{Q_{11}Q_{22} - Q_{11}^2}}, \qquad (9)$$

that relates to χ as $\chi = 2\epsilon/(1 + |\epsilon|^2)$. We choose to transform to ϵ here instead of χ, because the range of ϵ is the unit disk, while χ has an infinite range (cf. Fig. 1 of Viola et al. 2014). Thus, a ϵ allows us to sample extremely elliptical objects in an unbiased way, leading to well-defined moments for p_ϵ. Moreover, we explicitly take into account the correlation between the moments, in particular also between numerator and denominator of χ.

We will place requirements on how well the mean and error of the moment distribution needs to be known in order to ensure small biases on ellipticity. The procedure we take is as follows. We perform Monte-Carlo sampling from the $\Pi(\boldsymbol{Q})$ distribution, for each sampled value we assign a measurement error equal to the value given above to define a $p_i(\boldsymbol{Q})$ ($\sigma(Q_{11}) = 0.020$, $\sigma(Q_{22}) = 0.019$, $\sigma(Q_{12}) = 0.012$, all in arcsec^2). We then transform this to $p_{\epsilon,i}$ and compute the bias in the mean of this distribution away from the values computed by using the mean values $\mu(Q_{ij})$ in Eq. (9); the bias is the difference between the two. This results in a distribution of biases in ϵ_1 and ϵ_2 from which a mean bias $\langle$ bias $\epsilon_i\rangle$, and error on the bias $\sigma(\epsilon_i)$, can be computed. We perform this for the fiducial distribution and then repeat the process for distributions $\Pi(\boldsymbol{Q})$ for which the mean and error have been perturbed. We can then compute the relative change in the biases caused by the perturbations $b_e(\{\boldsymbol{Q}_F\}) - b_e(\{\boldsymbol{Q}_F + \delta\boldsymbol{Q}\})$ (where b_e are biases in ellipticity that are a function of a fiducial set of moments $\{\boldsymbol{Q}_F\}$ and perturbations about these $\{\boldsymbol{Q}_F + \delta\boldsymbol{Q}\}$), and therefore relate the knowledge of this distribution to biases in ellipticity. We assume that given a well-defined measurement of the moments the fiducial bias can be corrected for using the analytic results from the Marsaglia-Tin distribution.

We note that the process allows one to place requirements on the quadrupole moments, that ensure unbiased ellipticity measurements, with avoiding the need to estimate any ratios of variables. The ellipticities thus derived are therefore usable in an unbiased way for cosmological parameter inference following the formalisms followed in Amara & Réfrégier (2008), Kitching et al. (2009) and M13.

In Fig. 2 we show an example of the process. We present 1000 Monte-Carlo realisations of the fiducial $\Pi(\boldsymbol{Q})$ distribution propagated into p_ϵ. We show the product of the ellipticity distributions $\prod p_{\epsilon,i}$, which returns the inferred prior (the intrinsic ellipticity) distribution (as shown in Miller et al. 2007) and the
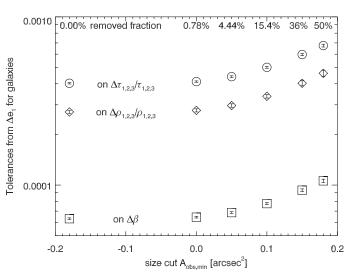

Fig. 3. Relative tolerances for changes in ArCTIC trap model parameters based on the *Euclid* requirement for Δe_1 in galaxies, as a function of the minimum observed size $A_{\mathrm{obs,min}}$ included in the analysis. Larger tolerances mean less strict calibration requirements. We show results for the well fill power β, and the densities $\rho_{1,2,3}$ and release time-scales $\tau_{1,2,3}$ of all trap species considered by I15 (cf. their Table 3).

distribution of biases in ϵ_1 and ϵ_2 as a result of taking realisations from the $\Pi(\boldsymbol{Q})$ distribution.

4. Results

4.1. Results for the size cut method

4.1.1. Size cuts and CTI correction sensitivity

Re-creating the results of I15 with different size cuts $A_{\mathrm{obs,min}}$ in place, we find the effect of a size cut on the CTI correction sensitivity, that is the required precision to which the ArCTIC parameters need to calibrated, to be reassuringly small. To first order, and especially so for size cuts $A_{\mathrm{obs,min}}$ affecting only the extreme tail of the distribution, CTI correction works independent of a size cut. After all, we deal with a pixel-level correction before the extraction of a catalogue.

Figure 3 presents the tolerances (bias margins) for the parameters I15 found to yield the tightest margins given the *Euclid* requirement on CTI-induced ellipticity bias $\Delta e_1 = \langle e_{1,\mathrm{corrected}}\rangle - \langle e_{1,\mathrm{input}}\rangle$, as a function of a size cut $A_{\mathrm{obs}} > A_{\mathrm{obs,min}}$. We only consider CTI along the x_1 axis. We probe the densities ρ and release time-scales τ of the trap species. The same relative biases
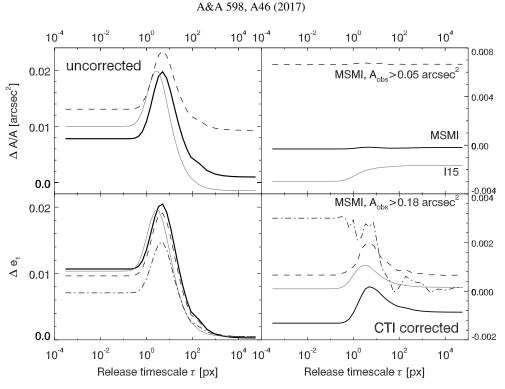
Fig. 4. CTI-induced relative size bias $\Delta A/A$ (*upper panels*) and ellipticity bias Δe_1 (*lower panels*) caused by a single trap species of time scale τ (in units of pixels clocked charge has travelled) and unit density. Measurements before (*left panels*) and after (*right panels*) CTI correction are shown for the I15 faint galaxy sample for four shape measurement choices: the I15 default (iterative centroiding and adaptive weight function width ω and $\langle\omega\rangle = 0\overset{''}{.}34$; grey solid lines; also cf. Fig. 2 of I15); the most shape-measurement independent (MSMI) choice of `SExtractor` centroids and fixed $\omega = 0\overset{''}{.}34$ (black solid lines); MSMI with a size cut $A_{\rm obs} > 0\overset{''}{.}05$ (dashed lines) and MSMI with a size cut $A_{\rm obs} > 0\overset{''}{.}18$ (dot-dashed lines, off scale for $\Delta A/A$). Uncertainties are smaller than the line widths. The results in this figure correspond to a trap density of 24% of the *Euclid* end-of-life estimate.

in ρ_i and τ_i are tested simultaneously for all three trap species in the I15 baseline model. We also probe the well fill power β describing the volume growth of a charge cloud inside a pixel as a function of number of electrons. We derive tolerances on the deviation of the parameters from the fiducial values by fitting to the measured $\Delta e_1(\Delta\beta)$ and the other parameters in the same way as I15. Thus, without a size cut (leftmost points in Fig. 3), we reproduce the tolerances given in Table 3 of I15.

We observe a significant change in tolerances only for size cuts removing at least several per cent of the catalogue sources. Moreover, the size cuts act in a way that render Δe_1 more robust to biases in the trap parameters. This is what we expect when removing objects with a denominator in Eq. (2) close to zero, while slightly increasing the average S/N. A size cut at the scale of the PSF size would have the welcome side-effect of widening the margins in the crucial ArCTIC parameters by ~70%.

4.1.2. Disentangling CTI correction and shape measurement

Any size cut, however, affects what I15 termed the CTI correction zeropoint, that is the residual bias due to overcorrected read-out noise that is present even if the ArCTIC parameters are perfectly known and correct. (Consistent with I15, we tacitly assumed this effect to be absent in the analysis leading to Fig. 3. Indeed, the zeropoint can be calibrated out of real data at the catalogue stage.) This is obvious for the zeropoint in $\Delta A = \langle A_{\rm corrected}\rangle - \langle A_{\rm input}\rangle \equiv \langle A_{\rm obs}\rangle - \langle A_{\rm input}\rangle$. Removing objects below a threshold in $A_{\rm obs}$ increases $\langle A_{\rm obs}\rangle$, but leaves $\langle A_{\rm input}\rangle$ unchanged, as we saw in Sect. 4.1.1.

While we observe the expected monotonic increase in the ΔA zeropoint, the zeropoint CTI bias in e_1 (with A in the

denominator) shows a more complicated, non-monotonic behaviour as a function of $A_{\rm obs,min}$.

Moreover, analysing simulations with a variable amount of read-out noise, we also find a non-monotonic dependency of the Δe_1 zeropoint, instead of the increasing (in absolute terms) CTI correction residuals illustrated in Fig. 3 of I15. Because adding more noise between applying CTI and correcting cannot lead to a better reconstruction of the true, underlying pre-CTI image, these findings are best explained by an artifact of the (simple) shear measurement pipeline we are using. `SExtractor` catalogues are fed into the RRG algorithm which iteratively determines a centroid and, in the I15 setup, calculates the size (standard deviation) of the Gaussian weight function in Eq. (1) as $\omega = 2\sqrt{\Omega/\pi}$, with Ω the `SExtractor` area.

These steps are susceptible to the same noise fluctuations of the local background that can make A negative. Because our goal is to allocate uncertainty margins to each element of the *Euclid* cosmic shear experiment, and validate algorithms against these requirements, we seek to disentangle effects of shape measurement and CTI correction. We thus propose a "most shape measurement independent" (MSMI) measurement of the CTI-induced biases in galaxy morphometry, and compare to I15 to gauge the magnitude of pipeline effects on CTI correction.

Our MSMI setting directly uses `SExtractor` centroids for the galaxies, switching off the iterative refinement. We fix the weight function size to a fixed, small value of $\omega = 0\overset{''}{.}34$, to minimise the effect of outlying sky pixels. Our value matches the sample average of ω I15 recorded for the same galaxies.

Figure 4 shows the CTI-induced $\Delta A/A$ (upper panels) and Δe_1 (lower panels) arising from a single trap species of time scale τ before (left panels) and after (right panels) CTI correction.

Table 1. Parametric description of CTI-induced bias charge traps of different species cause in the measured sizes d and ellipticities e_1 of faint galaxies (black lines in Fig. 4).

	C	D_a	D_p	D_w	G_a	G_p	G_w
MSMI-measured galaxies: degraded images							
$\Delta d/d_{\text{true}}$	0.04350 ± 0.00034	-0.0235 ± 0.0005	1.92 ± 0.02	0.39 ± 0.08	0.1236 ± 0.0038	0.713 ± 0.011	0.449 ± 0.017
Δe_1	0.05356 ± 0.00022	-0.0360 ± 0.0004	1.99 ± 0.02	0.38 ± 0.04	0.1058 ± 0.0019	0.666 ± 0.004	0.434 ± 0.008
MSMI-measured galaxies: CTI-corrected images							
$\Delta d/d_{\text{true}}$	-0.00249 ± 0.00030	0.0005 ± 0.0005	1.83 ± 1.27	0.82 ± 3.91	0.0013 ± 0.0009	0.675 ± 0.246	0.432 ± 0.372
Δe_1	-0.00107 ± 0.00012	0.0002 ± 0.0001	1.40 ± 0.50	0.10 ± 0.00	0.0015 ± 0.0003	0.792 ± 0.145	0.517 ± 0.141

Notes. The measurements assume a density of one trap per pixel, and the astrophysical measurement is fitted as a function of the charge trap's characteristic release time τ as $C + D_a \operatorname{atan}((\log \tau - D_p)/D_w) + G_a \exp\left(-(\log \tau - G_p)^2/2G_w^2\right)$. The parameters in this table correspond to a trap density of 24% of the *Euclid* end-of-life estimate.

Qualitatively, the MSMI (black lines) and I15 (grey lines; see also their Fig. 2) are in broad agreement, with the traps causing the strongest biases slightly shifting towards longer release times τ. Possibly this is due to more objects being slightly off-centred in the more simplistic MSMI setup, and thus more sensitive to their electrons being dragged out of the aperture.

After CTI correction, the MSMI measurements return a significantly smaller residual $\Delta A/A$ than the I15 settings over the whole eight decades in release time τ we tested. Curiously, the I15 pipeline performs better in residual Δe_1 for a random τ, but our reducing the influence of the shape-measurement pipeline nulls away the zeropoint bias of the most effective charge traps at the peak of the curve.

We also introduce size cuts in the MSMI measurements (dashed and dot-dashed lines in Fig. 4). These size cuts lower the bias Δe_1 for all traps, likely by removing the some of the most biases sources. However, by the mechanism described in Sect. 4.1.1, size cuts introduce an additional bias in $\Delta A/A$ that the CTI correction cannot account for (but which could be removed by calibration).

A complete understanding of the interaction between the ArCTIC CTI correction, shape measurement algorithms and source selection by size cuts exceeds the scope of this paper. We conclude that they should be disentangled as far as possible and provide empirical fits to the results of Fig. 4, updating Table 1 of I15 as a baseline for further research. As I15 demonstrated, the combined effect of several trap species are linear combinations of the biases caused by their component trap species.

4.2. Results for requirements recast on quadrupole moments

Figure 5 shows the measured distribution of moments $\Pi(\boldsymbol{Q})$ measured in `GalSim` galaxies. We find that indeed the quadrupole moments are consistent with a Gaussian distribution with a mean for each component of $\mu(Q_{11}) = 0.042$, $\mu(Q_{22}) = 0.039$, $\mu(Q_{12}) = 7 \times 10^{-4}$ and an error on each component of $\sigma(Q_{11}) = 0.020$, $\sigma(Q_{22}) = 0.019$, $\sigma(Q_{12}) = 0.012$; reported in units of arcseconds squared throughout.

We also measured these values from COSMOS data (Leauthaud et al. 2007) and found values (in arcsec²) of $\sigma(Q_{11}) = 0.09$, $\mu(Q_{11}) \sim 0.02$, $\sigma(Q_{12}) = 0.040$, $\mu(Q_{12}) \sim 0$ before PSF correction. We repeat the analysis in Sect. 4.2 with these values, and recover the same requirements on ellipticity to four decimal places, or $\sim 5\%$ of the requirement.

The error on Q_{12} is expected to be approximately $\sigma(Q_{12}) \approx 0.5(\sigma^2(Q_{11}) + \sigma^2(Q_{22}))^{1/2}$. Through expected symmetries we assume in our analysis that $\mu(Q_{12}) = 0$, $\mu(Q_{11}) = \mu(Q_{22})$ and $\sigma(Q_{11}) = \sigma(Q_{22})$. The moments are correlated as shown in
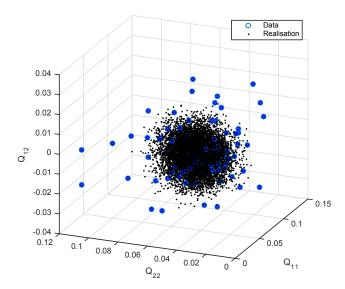

Fig. 5. Distribution of the measured quadrupole moments measured in `GalSim` galaxies. We show only 100 points here, in blue. The black points show a set of 1000 values sampled from a multivariate Gaussian derived from the measured values. All axes are in units of arcsec².

Viola et al. (2014), and using these correlation coefficients we can then make random realisations of this distribution that we also show in Fig. 5.

About the fiducial quadrupole moment distribution $\Pi(\boldsymbol{Q})$ we make a 1000 perturbations where we vary the mean and the errors, each a uniform random number between $r = [-0.01, 0.01]$, for example for the mean $\mu(Q_{11}) \rightarrow \mu(Q_{11}) + r$, and similarly for the other variables each with an independent value. In this way we are searching the six-dimensional "requirement space" (the means and standard deviations of each moment direction) in a random way – a more sophisticated implementation could use Markov-chain optimisation for example. For each realisation of $\Pi(\boldsymbol{Q})$ we sample 10^4 points from this distribution and then follow the procedure outlined in Sect. 3.2.

In Fig. 6 we show the dependency of the mean and error of the ellipticity biases, as a function of perturbations in the mean and error on the quadrupole moment distribution. The mean and error of the bias for each ellipticity component depends on the all the six varied parameters, however there is a primary direction in this parameter space along which the strongest dependency occurs, for example the mean bias in ϵ_1 is most sensitive to changes in the mean of the Q_{11} as is expected, so we only show these strongest dependencies. We now derive our main requirements taking only into account these dominant dependencies.
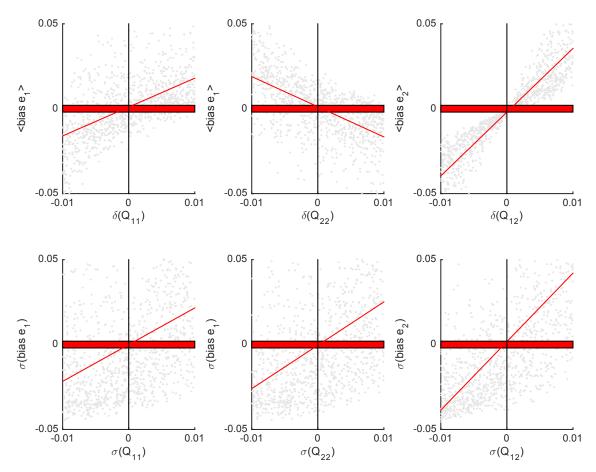
Fig. 6. Dependency of the mean and error of ellipticity biases as a function of changes in the mean and error of the ensemble quadrupole moment distribution (in arcsec^2). The red lines show linear fits to the distributions (grey points), and the red bar shows $\pm 2 \times 10^{-3}$ arcsec^2.

Table 2. The derived requirements on the quadrupole moments.

$\delta(Q_{11})$	0.0017
$\delta(Q_{22})$	0.0017
$\delta(Q_{12})$	0.0014
$\sigma(Q_{11})$	0.0016
$\sigma(Q_{22})$	0.0016
$\sigma(Q_{12})$	0.0017

Notes. Values need to be less than these quantities, in units of arcsec^2.

We find that there is an approximately linear dependency between the mean bias and the mean of the quadrupole moments, and between the error on the bias and the error of the quadrupole moment distribution. By taking the requirement in biases on ellipticity (e.g. from Kitching et al. 2009) of 2×10^{-3} arcsec^2 we can therefore set a requirement on the moments of quadrupole moment distribution – where the bias in ellipticity exceeds the absolute value of this. We show the knowledge of mean and error of each quadrupole moment show in Table 2. As a rule of thumb we find that the error on the standard deviation of each component needs to be smaller than 1.4×10^{-3} arcsec^2. Alternatively the requirement on Stokes parameters $Q_{11} + Q_{22}$ and $2Q_{12}$ is 1.9×10^{-3} arcsec^2.

5. Implications for *Euclid* CTI correction

The wider uncertainty margins on CTI trap parameters found in Sect. 4 – up to 70% – underline the positive role that size cuts to the galaxy catalogue can have in shear experiments. However,

we have to caution that this new result does not mean efforts on CTI calibration have become less crucial for the *Euclid* mission, for three reasons.

First, there will be a size cut in *Euclid* shear catalogue anyway, very likely at some value slightly larger than the PSF size. What has improved is our modelling of an experiment's sensitivity to CTI, addressing an aspect that was missing in I15.

Second, the 70% improvement specifically applies to our baseline trap model, the mix of trap species defined in I15, which is informed by laboratory experiments on *Euclid* VIS CCDs, but should be seen as a realistic example, not a prediction of the actual *Euclid* CTI.

Third, we observe the CTI correction, the (simple) shape measurement pipeline we used, and the size cut on the source catalogue to interact non-trivially. We simplified the shape measurement algorithm further and find the residual size and ellipticity biases after CTI correction to decrease for many relevant charge trap species. Our results provide a new baseline for further research.

While size cuts further prevent biases that would arise if a galaxy of measured size smaller than the PSF was corrected for PSF effects, no PSF correction has been performed in either I15 or this work. This approach is justified because *Euclid* requirements on shape measurement are kept deliberately distinct from those on PSF knowledge and correction. Setting separate requirements on the shear numerator and denominator is also consistent with the separation of numerator and denominator (for PSF and weight function correction) in the RRG method.

When deriving these individual measurements of the surface brightness moments, we do not investigate the requirements on the covariances of the moments, taking instead the theoretical covariance estimated from Viola et al. (2014). In future work, however, these could also be investigated, and are likely to depend on the weight function used in the analysis (which could be even optimised for this purpose).

Moreover, we need to caution that while the most important quantities for which a weak lensing survey needs to set requirements can be expressed via surface brightness moments, there will be some degree of information loss inherent to the compression of the lensing map (from a two-dimensional input to a two-dimensional output) using only low-order surface brightness moments. This limitation is common to ellipticity-based and moment-based requirements.

6. Conclusions

Weak lensing is a potentially powerful probe of cosmology, but the experiments and algorithms to measure this phenomenon need to be carefully designed. To this end a series of requirements on instrumental and detector systematic effects has been previously derived based on the propagation of measured changes in ellipticity and size into changes on cosmological parameter inference. However in doing this, the relationship between ellipticity and size leads to requirements being placed on the mean of two measured random variables. Such moments are not defined in general and therefore cannot be measured, so that such requirements cannot be tested.

Especially problematic are those sources whose measured size is close to zero or negative, because estimation techniques like the Marsaglia (2006) re-parametrisation cannot be applied. However galaxies of negative measured size represent legitimate samples from the size distribution of faint, small objects in the presence of noise that still occur with non-negligible frequently at $S/N \gtrsim 15$.

Removing the smallest galaxies from source catalogues by means of a size cut is a viable strategy in many, but not all weak lensing contexts. Sampling from a clipped source distribution may introduce unwanted biases in the complex data analysis chains. We extend the I15 sensitivity study of the charge transfer inefficiency (CTI) correction to include a size cut. We find the tolerance margins in CTI correction parameters to show a moderate dependence on removing the smallest sources. In fact, uncertainty margins are found to be wider by up to ~70% in the tested set-up.

As a more robust long-term solution, we present a formalism that allows requirements on ellipticity to be set in the space of quadrupole moments, that are linear functions of the data, where no ratios need to be computed, using the probability distribution for ellipticity derived in Viola et al. (2014). We find that the mean and the error of the distribution of quadrupole moments over the ensemble of galaxies used in a Stage-IV weak lensing experiment needs known to better than 1.4×10^{-3} arcsec^2 in each component for the ellipticity measurements to be unbiased at the level of 2×10^{-3} arcsec^2.

This requirement can now serve as a basis from which a breakdown and proportion into individual requirement on PSF and detector effects can be made as is done in Cropper et al. (2013). This should be straightforward given the formulae provided in Melchior et al. (2010) for example. We do not perform this breakdown here as the proportioning is flexible and should be done with instrument-specific knowledge: for example one may have a very stable PSF and wish to proportion more flexibility to instrument effects or vice versa. In this study we propagate only requirements to ellipticity, and marginalise over size; however if one wishes to use weak lensing magnification as an additional cosmological probe then this could be used to set joint ellipticity and size requirements. The setting of these requirements can now serve as a firm statistical basis from which weak lensing experimental design can proceed.

Finally, we point out that because recasting the requirements does not affect the galaxy selection as the size cuts do, there is no need to update the CTI trap model sensitivity results of I15 for this case. Rather, these two elements work together: we have updated the CTI results to a more realistic model, while identifying a more appropriate way to define the requirements setting the context for the CTI correction in *Euclid*.

References

Amara, A., & Réfrégier, A. 2008, MNRAS, 391, 228
Antonik, M. L., Bacon, D. J., Bridle, S., et al. 2013, MNRAS, 431, 3291
Applegate, D. E., von der Linden, A., Kelly, P. L., et al. 2014, MNRAS, 439, 48
Bartelmann, M., & Schneider, P. 2001, Phys. Rep., 340, 291
Bertin, E., & Arnouts, S. 1996, A&AS, 117, 393
Chang, C., Jarvis, M., Jain, B., et al. 2013a, MNRAS, 434, 2121
Chang, C., Kahn, S. M., Jernigan, J. G., et al. 2013b, MNRAS, 428, 2695
Cropper, M., Hoekstra, H., Kitching, T., et al. 2013, MNRAS, 431, 3103
Erben, T., Van Waerbeke, L., Bertin, E., Mellier, Y., & Schneider, P. 2001, A&A, 366, 717
Hinkley, D. V. 1969, Biometrika, 56, 635
Hoekstra, H. 2007, MNRAS, 379, 317
Huterer, D., & White, M. 2002, ApJ, 578, L95
Israel, H., Erben, T., Reiprich, T. H., et al. 2010, A&A, 520, A58
Israel, H., Massey, R., Prod'homme, T., et al. 2015, MNRAS, 453, 561
Jarvis, M., Sheldon, E., Zuntz, J., et al. 2016, MNRAS, 460, 2245
Kaiser, N., Squires, G., & Broadhurst, T. 1995, ApJ, 449, 460
Kitching, T. D., Amara, A., Abdalla, F. B., Joachimi, B., & Refregier, A. 2009, MNRAS, 399, 2107
Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, ArXiv e-print [arXiv:1110.3193]
Leauthaud, A., Massey, R., Kneib, J.-P., et al. 2007, ApJS, 172, 219
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., Stoughton, C., Leauthaud, A., et al. 2010, MNRAS, 401, 371
Massey, R., Hoekstra, H., Kitching, T., et al. 2013, MNRAS, 429, 661
Massey, R., Schrabback, T., Cordes, O., et al. 2014, MNRAS, 439, 887
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
Miller, L., Kitching, T. D., Heymans, C., Heavens, A. F., & van Waerbeke, L. 2007, MNRAS, 382, 315
Miller, L., Heymans, C., Kitching, T. D., et al. 2013, MNRAS, 429, 2858
Okabe, N., & Umetsu, K. 2008, PASJ, 60, 345
Paulin-Henriksson, S., Amara, A., Voigt, L., Refregier, A., & Bridle, S. L. 2008, A&A, 484, 67
Refregier, A., Kacprzak, T., Amara, A., Bridle, S., & Rowe, B. 2012, MNRAS, 425, 1951
Rhodes, J., Réfrégier, A., & Groth, E. J. 2001, ApJ, 552, L85
Rowe, B. T. P., Jarvis, M., Mandelbaum, R., et al. 2015, Astron. Comput., 10, 121
Sheldon, E. S. 2014, MNRAS, 444, L25
Tin, M. 1965, J. Am. Stat. Assoc., 60, 294
Viola, M., Kitching, T. D., & Joachimi, B. 2014, MNRAS, 439, 1909
Zuntz, J., Kacprzak, T., Voigt, L., et al. 2013, MNRAS, 434, 1604

Appendix A: Update of M13 CTI requirements derivation

Since M13, we have found it more physical to consider non-convolutive effects like CTI as linearly affecting an object size

$$A_{\rm obs} = A_{\rm gal} + A_{\rm NC} + A_{\rm PSF} \tag{A.1}$$

and ellipticity

$$e_{\rm obs} = e_{\rm gal} + e_{\rm NC} + \frac{A_{\rm PSF}}{A_{\rm gal} + A_{\rm PSF}} e_{\rm PSF}. \tag{A.2}$$

We now derive equivalent versions of Eqs. (34) to (37) in M13, the Taylor expansion detailing the sensitivity of observables to biases in ellipticity e_1 and size A. We ignore noise terms in the following and note that Eqs. (A.1) and (A.2) only hold true to linear order in moments. Not including PSF, we have

$$e_{\rm obs} = \frac{Q_{11} - Q_{22} + (\delta Q_{11} - \delta Q_{22})}{Q_{11} + Q_{22} + (\delta Q_{11} + \delta Q_{22})} \tag{A.3}$$

where δQ are the change in moments caused by CTI. This is approximately $e_{\rm obs} = e_{\rm gal} + e_{\rm NC}$ if $\delta Q \ll Q$ to linear order. We note this is a further reason for using Q requirements directly. This means that:

$$
\begin{aligned}
e_{\rm gal} &= \frac{(e_{\rm obs} - e_{\rm NC})(A_{\rm obs} - A_{\rm NC}) - e_{\rm PSF} A_{\rm PSF}}{(A_{\rm obs} - A_{\rm NC})} \\
&= (e_{\rm obs} - e_{\rm NC}) - \frac{e_{\rm PSF} A_{\rm PSF}}{(A_{\rm obs} - A_{\rm NC})},
\end{aligned} \tag{A.4}
$$

and consequently

$$\frac{\partial e_{\rm gal}}{\partial A_{\rm NC}} = -\frac{e_{\rm PSF} A_{\rm PSF}}{(A_{\rm obs} - A_{\rm NC})^2}, \tag{A.5}$$

and

$$\frac{\partial e_{\rm gal}}{\partial e_{\rm NC}} = -1 \tag{A.6}$$

so that the $\delta R_{\rm NC}$ term in the M13 Taylor expansion (their Eq. (36)) should be replaced by

$$\frac{A_{\rm PSF}\, e_{\rm PSF}}{(A_{\rm gal} + A_{\rm PSF})^2} \delta A_{\rm NC}, \tag{A.7}$$

and the $\delta e_{\rm NC}$ term (Eq. (37) in M13) is simply

$$-\delta e_{\rm NC} \tag{A.8}$$

this leads to a term like

$$\left\langle \left(\frac{A_{\rm PSF}}{(A_{\rm gal} + A_{\rm PSF})^2} \right)^2 \right\rangle, \tag{A.9}$$

in the ellipticity two-point correlation function or power spectrum, which needs to be evaluated in order to check the requirement level contribution. Problems will occur in taking the ratio in the regime that $A_{\rm gal} + A_{\rm PSF} = A_{\rm obs} - A_{\rm NC} \rightarrow 0$.

These terms make intuitive sense in that if the PSF ellipticity is zero, then any ellipticity-independent size change will not cause a bias in ellipticity. Whereas a change in ellipticity caused by CTI, that adds linearly, should cause a linear bias.

Appendix B: The Marsaglia mitigation technique

Marsaglia (2006) demonstrate that the ratio z/w of two normally distributed random variables z and w with respective means μ_z, μ_w, variances σ_z, σ_w and correlation ρ follows the same distribution as $\frac{a+x}{b+y}$ where x and y are uncorrelated standard normal random variables. The constants a and b are given by

$$b = \frac{\mu_w}{\sigma_w}, \quad a = \pm \frac{\mu_z/\sigma_z - \rho \mu_w/\sigma_w}{\sqrt{1-\rho^2}} > 0. \tag{B.1}$$

Although no formal proof is given, Marsaglia (2006) show that, for practical purposes, the mean and variance of the ratio exist and are approximated by

$$\mu = \frac{a}{1.01b - 0.2713}, \quad \sigma = \frac{a^2 + 1}{b^2 + 0.108b - 3.795 - \mu^2}, \tag{B.2}$$

as long as the conditions $b > 4$ and $a < 2.5$ are satisfied. In other words, negative samples must only occur very rarely (as 4σ events or scarcer) for the Marsaglia mitigation technique to be applicable. Figure 1 shows that this condition is not met for typical faint weak lensing test galaxies.