

Weak-lensing shear estimates with general adaptive moments, and studies of bias by pixellation, PSF distortions, and noise

Patrick Simon and Peter Schneider

Argelander-Institut für Astronomie, Universität Bonn, Auf dem Hügel 71, 53121 Bonn, Germany
e-mail: psimon@astro.uni-bonn.de

Received May 19, 2017

ABSTRACT

In weak gravitational lensing, weighted quadrupole moments of the brightness profile in galaxy images are a common way to estimate gravitational shear. We employ general adaptive moments (GLAM) to study causes of shear bias on a fundamental level and for a practical definition of an image ellipticity. The GLAM ellipticity has useful properties for any chosen weight profile: the weighted ellipticity is identical to that of isophotes of elliptical images, and in absence of noise and pixellation it is always an unbiased estimator of reduced shear. We show that moment-based techniques, adaptive or unweighted, are similar to a model-based approach in the sense that they can be seen as imperfect fit of an elliptical profile to the image. Due to residuals in the fit, moment-based estimates of ellipticities are prone to underfitting bias when inferred from observed images. The estimation is fundamentally limited mainly by pixellation which destroys information on the original, pre-seeing image. We give an optimized estimator for the pre-seeing GLAM ellipticity and quantify its bias for noise-free images. To deal with images where pixel noise is prominent, we consider a Bayesian approach to infer GLAM ellipticity where, similar to the noise-free case, the ellipticity posterior can be inconsistent with the true ellipticity if we do not properly account for our ignorance about fit residuals. This underfitting bias, quantified in the paper, does not vary with the overall noise level but changes with the pre-seeing brightness profile and the correlation or heterogeneity of pixel noise over the image. Furthermore, when inferring a constant ellipticity or, more relevantly, constant shear from a source sample with a distribution of intrinsic properties (sizes, centroid positions, intrinsic shapes), an additional, now noise-dependent bias arises towards low signal-to-noise if incorrect prior densities for the intrinsic properties are used. We discuss the origin of this prior bias. With regard to a fully-Bayesian lensing analysis, we point out that passing tests with source samples subject to constant shear may not be sufficient for an analysis of sources with varying shear.

Key words. Gravitational lensing: weak – Methods: data analysis – Methods: statistical

1. Introduction

Over the past decade measurements of distortions of galaxy images by the gravitational lensing effect have developed into an important, independent tool for cosmologists to study the large-scale distribution of matter in the Universe and its expansion history (recent reviews: Schneider 2006; Munshi et al. 2008; Hoekstra & Jain 2008; Massey et al. 2010; Kilbinger 2015). These studies exploit the magnification and shear of galaxy light bundles by the tidal gravitational field of intervening matter. The shear gives rise to a detectable coherent distortion pattern in the observed galaxy shapes. The distortions are usually weak, only of order of a few per cent of the unlensed shape of a typical galaxy image. Therefore, the key to successfully devising gravitational shear as cosmological tool are accurate measurements of the shapes of many, mostly faint and hardly resolved galaxy images.

There has been a boost of interest in methods of shape measurements in anticipation of the gravitational lensing analysis of the upcoming next generation of wide-field imaging surveys (e.g., Euclid; Laureijs et al. 2011). Despite being sufficient for contemporary surveys, current methodologies are not quite at the required level of accuracy to fully do justice to the amount of cosmological information in future lensing surveys (Heymans et al. 2006; Massey et al. 2007). The challenge all methodologies face is that observable, noisy (post-seeing) galaxy images

are modifications of the actual (pre-seeing) image owing to instrumental and possible atmospheric effects. Post-seeing galaxy images are subject to pixellation as well as instrumental noise, sky noise, photon noise, and random overlapping with very faint objects (Kitching et al. 2012; Hoekstra et al. 2015). In addition, galaxies are not intrinsically circular such that their ellipticities are noisy estimators of the cosmic distortion. Current theoretical work consequently focuses on sources of bias in shape measurements, such as pixel-noise bias, shape-noise bias, underfitting bias, colour gradients, or several selection biases (Hirata & Seljak 2003; Hirata et al. 2004; Mandelbaum et al. 2005; Melchior et al. 2010; Viola et al. 2011; Melchior & Viola 2012; Kacprzak et al. 2012; Massey et al. 2013; Semboloni et al. 2013).

One major source of bias is the pixel-noise bias or simply noise bias hereafter. This bias can at least partly be blamed on the usage of point estimates of galaxy shapes in a statistical analysis, i.e., single-value estimators of galaxy ellipticities (Refregier et al. 2012). This begs the question whether it is feasible to eradicate noise bias by means of a more careful treatment of the statistical uncertainties in the measurement of galaxy ellipticities within a fully Bayesian framework. Indeed recent advances in image processing for weak gravitational lensing strongly support this idea, at least for the inference of constant shear (Sheldon 2014; Bernstein & Armstrong 2014, BA14 hereafter; Bernstein et al. 2016, BAKM16 hereafter). In contrast, the contemporary philosophy with point estimates is to perform elaborate,

time-consuming calibrations of biased estimators by means of simulated images; the calibration accuracy is, additionally, only as good as the realism of simulated images (e.g., Hoekstra et al. 2015). To be fair, code implementations of non-Bayesian techniques are typically an order of magnitude or more faster than Bayesian codes which could be a decisive factor for upcoming surveys.

We take here a new look into possible causes of bias in shear measurements on a fundamental level. To this end, we examine, step-by-step with increasing complexity, a fully-Bayesian lensing analysis based on weighted brightness moments of galaxy images (Gelman et al. 2013; MacKay 2003). While the method in BA14 and BAKM16 is set in Fourier space, we work with moments in angular space which has benefits in the case of correlated noise or missing pixels in realistic images. Moment-based methods as ours are non-parametric; this means they are free from assumptions about the galaxy brightness profile. They hence appear to be advantageous for reducing bias, but nonetheless the specific choice of the adaptive weight for the moments is known to produce bias (Viola et al. 2014; Voigt & Bridle 2010). The origin of this problem, which principally also affects unweighted moments, becomes obvious in our formalism. We define as practical measure of galaxy shape a generalization of the impractical ellipticity ϵ expressed in terms of unweighted moments (Kaiser et al. 1995; Seitz & Schneider 1997, SS97 hereafter). Being Bayesian, our measurement of ellipticity results in a Monte-Carlo sample of the probability distribution function (PDF) of ϵ which should be propagated in a fully Bayesian analysis. That is: we do not devise point estimators in order to ideally stay clear of noise bias. This overall approach of general adaptive moments, GLAM hereafter, is inspired by and comparable to Bernstein & Jarvis (2002) apart from the Bayesian framework and some technical differences: (i) for any adaptive weight, the perfectly measured GLAM ellipticity is an unbiased estimator of gravitational shear unaffected by shape-noise bias; (ii) the adaptive weight may have a non-Gaussian radial profile; (iii) our inference of the pre-seeing ellipticity is realised as forward-fitting of elliptical profiles (so-called templates), that is we do not determine the brightness moments of the post-seeing image and correct them to estimate the pre-seeing moments (cf. Hirata & Seljak 2003; Mandelbaum et al. 2005).

As a disclaimer, the GLAM methodology outlined here is prone to bias, even where a fully Bayesian analysis can be realised, and, at this stage, its performance is behind that of other techniques. The aim of this paper is to elucidate causes of bias, instead of proposing a new technique that is competitive to existing techniques. However, these findings are also relevant for other methodologies because model-based or moment-based approaches are linked to GLAM.

For this paper, we exclude bias from practically relevant factors: the insufficient knowledge of the PSF or noise properties, blending of images, and the selection of source galaxies (see e.g., Heymans et al. 2006; Hartlap et al. 2011; Dawson et al. 2014). We focus on the core of the problem of shape measurements which is the inference of pre-seeing ellipticities from images whose full information on the brightness profile have been lost by instrumental limitations.

The paper is laid out as follows. In Sect. 2, we introduce the formalism of GLAM for a practical definition of ellipticity with convenient transformation properties under the action of gravitational shear. We also analytically investigate the limits of measuring the pre-seeing ellipticity from a noise-free but both PSF-convolved and pixellated image. In Sect. 3, we construct a statistical model for the GLAM ellipticity of noisy post-seeing

images. We then study the impact of inconsistencies in the posterior model of ellipticity in three, increasingly complex scenarios. First, we analyse with independent exposures of the same pre-seeing image the ellipticity bias due to a misspecified likelihood in the posterior (underfitting bias). Second, we consider samples of noisy images with the same ellipticity but distributions of intrinsic properties such as sizes or centroid positions. Here a new contribution to the ellipticity bias emerges if the prior densities of intrinsic properties are incorrectly specified (prior bias). Third in Sect. 4, we perform numerical experiments with samples of noisy galaxy images of random intrinsic shapes that are subject to constant shear. With these samples we study the impact of inconsistent ellipticity posteriors on shear constraints (shear bias). We also outline details on our technique to Monte-Carlo sample ellipticity or shear posteriors. We discuss our results and possible improvements of the GLAM approach in Sect. 5.

2. General adaptive moments

2.1. Definition

Let $I(\mathbf{x})$ be the light distribution in a galaxy image of infinite resolution and without noise where \mathbf{x} is the position on the sky. A common way of defining a (complex) ellipticity of a galaxy image uses the quadrupole moments

$$Q_{ij} = \frac{\int d^2x (x_i - X_{0,i})(x_j - X_{0,j}) I(\mathbf{x})}{\int d^2x I(\mathbf{x})} \quad (1)$$

of $I(\mathbf{x})$ relative to a centroid position

$$\mathbf{X}_0 = \frac{\int d^2x \mathbf{x} I(\mathbf{x})}{\int d^2x I(\mathbf{x})} \quad (2)$$

of the image (e.g., Bartelmann & Schneider 2001). For real applications, the quadrupole moments are soundly defined only if they involve a weight, decreasing with separation from the centroid position \mathbf{X}_0 , because galaxies are not isolated so that the normalisation $\int d^2x I(\mathbf{x})$ and the brightness moments diverge. Hirata & Seljak (2003, H03 hereafter) address the divergence problems by realising an adaptive weighting scheme by minimising the error functional, sometimes dubbed the energy functional,

$$E(\mathbf{p}|I) = \frac{1}{2} \int d^2x [I(\mathbf{x}) - A f(\rho)]^2, \quad (3)$$

with the quadratic form

$$\rho := (\mathbf{x} - \mathbf{x}_0)^T \mathbf{M}^{-1} (\mathbf{x} - \mathbf{x}_0) \quad (4)$$

and the second-order tensor

$$\mathbf{M} = \frac{T}{2} \begin{pmatrix} 1 + e_1 & e_2 \\ e_2 & 1 - e_1 \end{pmatrix}. \quad (5)$$

The tensor \mathbf{M} is expressed in terms of the complex ellipticity $e = e_1 + ie_2$ and the size T of the image; $f(\rho)$ is a weight function that HS03 chose to be a Gaussian weight $f(\rho) = e^{-\rho/2}$. In comparison to HS03, we have slightly changed the definition of ρ for convenience: here we use ρ instead of ρ^2 . The set $\mathbf{p} = (A, \mathbf{x}_0, \mathbf{M})$, comprises a set of six parameters on which the functional E depends for a given galaxy image I .

Frequently another definition of complex ellipticity, the ϵ -ellipticity, is more convenient (sometimes also known as the third flattening). It arises if we write ρ in the form

$$\rho = |V^{-1}(\mathbf{x} - \mathbf{x}_0)|^2 = (\mathbf{x} - \mathbf{x}_0)^T V^{-2}(\mathbf{x} - \mathbf{x}_0), \quad (6)$$

where V is symmetric. Obviously, we have $V^2 = M$ or $V = \sqrt{M}$. By writing V in the form

$$V = \frac{t}{2} \begin{pmatrix} 1 + \epsilon_1 & \epsilon_2 \\ \epsilon_2 & 1 - \epsilon_1 \end{pmatrix} \quad (7)$$

we see that $V^2 = M$ implies $2T = t^2(1 + |\epsilon|^2)$, and

$$e = \frac{2\epsilon}{1 + |\epsilon|^2}; \quad \epsilon = \epsilon_1 + i\epsilon_2 = \frac{e}{1 + \sqrt{1 - |e|^2}}. \quad (8)$$

We henceforth use ϵ as parametrisation of M because ϵ is an unbiased estimator of reduced shear in the absence of a PSF and pixellation (SS97). Conversely, the ellipticity e has to be calibrated with the distribution of unsheared ellipticities which poses another possible source of bias in a lensing analysis (H03).

As generally derived in Appendix A, the parameters \mathbf{p} at the minimum of (3) are:

$$\mathbf{x}_0 = \frac{\int d^2x \mathbf{x} I(\mathbf{x}) f'(\rho)}{\int d^2x I(\mathbf{x}) f'(\rho)}; \quad A = \frac{\int d^2x I(\mathbf{x}) f''(\rho)}{\int d^2x f''(\rho)} \quad (9)$$

and

$$M = \beta \frac{\int d^2x (\mathbf{x} - \mathbf{x}_0)(\mathbf{x} - \mathbf{x}_0)^T I(\mathbf{x}) f'(\rho)}{\int d^2x I(\mathbf{x}) f'(\rho)}, \quad (10)$$

where $\beta := -2 \int_0^\infty ds f^2(s)/f^2(0)$ is a constant. These equations are derived by HS03 for a Gaussian $f(\rho)$, for which we have $f'(\rho) := df/d\rho = -f/2$. This shows that the best-fitting $f(\rho)$ has the same centroid and, up to a scalar factor, the same second moment $M = V^2$ as the with $f'(\rho)$ adaptively weighted image $I(\mathbf{x})$. This is basically also noted in Lewis (2009) where it is argued that the least-square fit of *any* sheared model to a pre-seeing image $I(\mathbf{x})$ provides an unbiased estimate of the shear, even if it fits poorly.

In this system of equations, the centroid \mathbf{x}_0 and tensor M are implicitly defined because both $f(\rho)$ and $f'(\rho)$ on the right-hand-side are functions of the unknowns \mathbf{x}_0 and M : the weights *adapt* to the position, size, and shape of the image. The Eqs. (9) and (10) therefore need to be solved iteratively. The iteration should be started at a point which is close to the final solution. Such a starting point could be obtained by using the image position, determined by the image detection software, as initial value for \mathbf{x}_0 , and tensor M determined from a circular weight function with the same functional form as f . Nevertheless, there is no guarantee that the solution of this set of equations is unique. In fact, for images with two brightness peaks one might suspect that there are multiple local minima in \mathbf{p} of the functional E . This may occur, for instance, in the case of blended images. One standard solution to this particular problem is to identify blends and to remove these images from the shear catalogue. Alternatively we could in principle try to fit two template profiles to the observed image, i.e., by adding a second profile $E_2(\mathbf{p}_2|I)$ to the functional (3) and by minimising the new functional with respect to the parameter sets \mathbf{p} and \mathbf{p}_2 of both profiles simultaneously.

2.2. Interpretation

If an image $I(\mathbf{x})$ has confocal elliptical isophotes, with the same ellipticity for all isophotes, one can define the ellipticity of the image uniquely by the ellipticity of the isophotes. In this case, the ellipticity ϵ defined by the minimum of (3) coincides with the ellipticity of the isophotes for any weight f . We show this property in the following.

Assume that the brightness profile $I(\mathbf{x})$ is constant on confocal ellipses so that we can write $I(\mathbf{x}) = S(\zeta)$ where $\zeta := (\mathbf{x} - \mathbf{x}_c)^T B^{-2}(\mathbf{x} - \mathbf{x}_c)$. Here \mathbf{x}_c denotes the centre of the image, and the matrix elements of B describe the size and the shape of the image, in the same way as we discussed for the matrix V before. The function $S(\zeta)$ describes the radial brightness profile of the image. We start by writing (9) in the form

$$0 = \int d^2x S(\zeta) f'(\rho) (\mathbf{x} - \mathbf{x}_0), \quad (11)$$

and introduce the transformed position vector $\mathbf{z} = B^{-1}(\mathbf{x} - \mathbf{x}_c)$, or $\mathbf{x} = B\mathbf{z} + \mathbf{x}_c$. Then the previous equation becomes

$$0 = \int d^2z S(|z|^2) f'(\rho) (B\mathbf{z} + \mathbf{x}_c - \mathbf{x}_0), \quad (12)$$

where in terms of \mathbf{z} the quadratic form ρ is

$$\rho = (B\mathbf{z} + \mathbf{x}_c - \mathbf{x}_0)^T V^{-2}(B\mathbf{z} + \mathbf{x}_c - \mathbf{x}_0). \quad (13)$$

From these equations, we can see that $\mathbf{x}_0 = \mathbf{x}_c$ is the solution of (12) since then ρ is an even function of \mathbf{z} , S is an even function of \mathbf{z} , whereas the term in the parenthesis of (12) is odd, and the integral vanishes due to symmetry reasons. Thus we found that our adaptive moments approach yields the correct centre of the image.

Next, we rewrite (10) in the form

$$\begin{aligned} V^2 \int d^2x I(\mathbf{x}) f(\rho) &= \beta \int d^2x (\mathbf{x} - \mathbf{x}_0)(\mathbf{x} - \mathbf{x}_0)^T I(\mathbf{x}) f'(\rho) \\ &= \beta \det B \int d^2z S(|z|^2) f'(\rho) B\mathbf{z}\mathbf{z}^T B, \end{aligned} \quad (14)$$

where β , see Eq. (10), is a constant factor (see Appendix A.3). We again used the transformation from \mathbf{x} to $\mathbf{z} = B^{-1}(\mathbf{x} - \mathbf{x}_c)$ and employed the fact that $\mathbf{x}_0 = \mathbf{x}_c$. Accordingly, we have $\rho = \mathbf{z}^T B V^{-2} B \mathbf{z}$. We now show that the solution of (14) is given by $V = \lambda B$ with λ being a scalar factor. Using this Ansatz, we get $\rho = \lambda^{-2} |z|^2$, and (14) can be written, after multiplying from the left and from the right by B^{-1} , as

$$B^{-1} V^2 B^{-1} = \lambda^2 \mathbf{1} = \beta \frac{\int d^2z S(|z|^2) f'(|z|^2/\lambda^2) \mathbf{z}\mathbf{z}^T}{\int d^2z S(|z|^2) f(|z|^2/\lambda^2)}. \quad (15)$$

Since both S and f' in the numerator depend solely on $|z|^2$, the tensor on the right hand side is proportional to the unit tensor $\mathbf{1}$, and (15) becomes a scalar equation for the scalar λ ,

$$\lambda^2 = \beta \frac{\int ds s S(s) f'(s/\lambda^2)}{\int ds S(s) f(s/\lambda^2)}, \quad (16)$$

whose solution depends on the brightness profile S and the chosen weight function f . However, the fact that B differs from V only by the scalar factor λ implies that the derived ellipticity ϵ of V is the same as that of the elliptical image. Therefore we have shown that the approach of adapted moments recovers the true

ellipticity with elliptical isophotes for *any* radial weight function f .

For a general brightness profile of the image, the interpretation of the GLAM ellipticity ϵ is less clear, and the ellipticity generally depends on the weight f . Nevertheless, ϵ is uniquely defined as long as a minimum of the functional (3) can be found.

More importantly, for *any* weight f the GLAM ellipticity obeys the same simple transformation law under the action of gravitational shear, as shown in the following section.

2.3. Transformation under shear

We now consider the effect of a shear $\gamma = \gamma_1 + i\gamma_2$ and convergence κ on the GLAM ellipticity ϵ (Bartelmann & Schneider 2001).

[REDACTED]

For general images, let $\epsilon_s = \epsilon_{s,1} + i\epsilon_{s,2}$ be the complex ellipticity of the image in the source plane and $\epsilon = \epsilon_1 + i\epsilon_2$ its complex ellipticity in the lens plane. We show now that for any brightness profile and template $f(\rho)$, GLAM ellipticities have the extremely useful property to transform under the action of a reduced shear according to

$$\epsilon(g, \epsilon_s) = \begin{cases} \frac{\epsilon_s + g}{1 + g^* \epsilon_s} & , |g| \leq 1, \\ \frac{1 + \epsilon_s^* g}{\epsilon_s^* + g^*} & , |g| > 1 \end{cases} . \quad (17)$$

This is exactly the well-known transformation obtained from unweighted moments (SS97). To show this generally for GLAM, let $I_s(\mathbf{y})$ be the surface brightness of an image in the source plane with source plane coordinates \mathbf{y} . The centroid \mathbf{y}_0 and moment tensor \mathbf{M}_s of I_s are defined by the minimum of (3) or, alternatively, by the analog of Eqs. (9) and (10) through

$$\int d^2\mathbf{y} I_s(\mathbf{y}) f(\rho_s) (\mathbf{y} - \mathbf{y}_0) = 0 \quad (18)$$

and

$$\mathbf{M}_s = \beta \frac{\int d^2\mathbf{y} (\mathbf{y} - \mathbf{y}_0)(\mathbf{y} - \mathbf{y}_0)^T I_s(\mathbf{y}) f'(\rho_s)}{\int d^2\mathbf{y} I_s(\mathbf{y}) f(\rho_s)} \quad (19)$$

with

$$\rho_s = (\mathbf{y} - \mathbf{y}_0)^T \mathbf{M}_s^{-1} (\mathbf{y} - \mathbf{y}_0) . \quad (20)$$

The ellipticity ϵ_s is given by

$$\mathbf{M}_s = \mathbf{V}_s^2 ; \mathbf{V}_s = \frac{t_s}{2} \begin{pmatrix} 1 + \epsilon_{s,1} & \epsilon_{s,2} \\ \epsilon_{s,2} & 1 - \epsilon_{s,1} \end{pmatrix} . \quad (21)$$

We now shear the image I_s . The shear and the convergence are assumed to be constant over the extent of the image, i.e., lens plane positions \mathbf{x} are linearly mapped onto source plane positions \mathbf{y} by virtue of $\mathbf{y} - \mathbf{y}_c = \mathcal{A}(\mathbf{x} - \mathbf{x}_c)$ where

$$\mathcal{A} = \mathcal{A}^T = (1 - \kappa) \begin{pmatrix} 1 - g_1 & -g_2 \\ -g_2 & 1 + g_1 \end{pmatrix} ; \quad (22)$$

\mathbf{x}_c and \mathbf{y}_c are such that the point \mathbf{x}_c is mapped onto the point \mathbf{y}_c by the lens equation, and both are chosen to be the central points around which the lens equation is linearised. We then find for the centroid \mathbf{y}_0 in the source plane

$$\mathbf{y} - \mathbf{y}_0 = \mathbf{y} - \mathbf{y}_c + \mathbf{y}_c - \mathbf{y}_0 = \mathcal{A}(\mathbf{x} - \mathbf{x}_1) , \quad (23)$$

where $\mathbf{x}_1 - \mathbf{x}_c = \mathcal{A}^{-1}(\mathbf{y}_0 - \mathbf{y}_c)$. This then yields for (20)

$$\rho_s = (\mathbf{x} - \mathbf{x}_1)^T \mathcal{A} \mathbf{M}_s^{-1} \mathcal{A} (\mathbf{x} - \mathbf{x}_1) . \quad (24)$$

In the next step, the expression (18) for the source centre can be rewritten by transforming to the image coordinates and using the conservation of surface brightness $I_s(\mathbf{y}(\mathbf{x})) = I(\mathbf{x})$:

$$\begin{aligned} 0 &= \int d^2\mathbf{y} I_s(\mathbf{y}) f(\rho_s) (\mathbf{y} - \mathbf{y}_0) \\ &= \det(\mathcal{A}) \mathcal{A} \int d^2\mathbf{x} I(\mathbf{x}) f(\rho_s) (\mathbf{x} - \mathbf{x}_1) . \end{aligned} \quad (25)$$

With the same transformation, we rewrite the moment tensor (19) as

$$\mathbf{M}_s = \beta \mathcal{A} \frac{\int d^2\mathbf{x} (\mathbf{x} - \mathbf{x}_1)(\mathbf{x} - \mathbf{x}_1)^T I(\mathbf{x}) f'(\rho_s)}{\int d^2\mathbf{x} I(\mathbf{x}) f(\rho_s)} \mathcal{A} . \quad (26)$$

On the other hand, minimising the functional (3) for the surface brightness $I(\mathbf{x})$ in the lens plane yields the expressions (9) and (10) for \mathbf{x}_0 and \mathbf{M} , respectively. We then see that the Eqs. (25), (26) and (9), (10) agree with each other, if we set

$$\mathbf{M}_s = \mathcal{A} \mathbf{M} \mathcal{A} ; \mathbf{x}_1 = \mathbf{x}_0 , \quad (27)$$

for which $\rho_s = \rho$. In particular, this shows that the centre \mathbf{x}_0 of the image is mapped onto the centre \mathbf{y}_0 of the source.

The relation

$$\mathbf{M} = \mathcal{A}^{-1} \mathbf{M}_s \mathcal{A}^{-1} \quad (28)$$

can be rewritten in terms of the square roots of the matrix \mathbf{M} as

$$\mathbf{M}_s \equiv \mathbf{V}_s^2 = \mathcal{A} \mathbf{V}^2 \mathcal{A} \quad (29)$$

where $\mathbf{V}_s = \sqrt{\mathbf{M}_s}$ is uniquely defined by requiring that for the symmetric, positive-definite matrix \mathbf{M}_s , \mathbf{V}_s is symmetric and positive-definite. Although both \mathbf{V} and \mathcal{A} are symmetric, $\mathcal{A} \mathbf{V}$ is in general not. Therefore \mathbf{V}_s cannot be readily read off from (29). Instead, we use a rotation matrix

$$\mathbf{R}(\varphi) = \begin{pmatrix} \cos \varphi & -\sin \varphi \\ +\sin \varphi & \cos \varphi \end{pmatrix} ; \mathbf{R}^T(\varphi) = \mathbf{R}^{-1}(\varphi) = \mathbf{R}(-\varphi) , \quad (30)$$

to write

$$\mathbf{V}_s^2 = \mathcal{A} \mathbf{V} \mathbf{R}^{-1}(\varphi) \mathbf{R}(\varphi) \mathbf{V} \mathcal{A} = [\mathbf{R}(\varphi) \mathbf{V} \mathcal{A}]^T [\mathbf{R}(\varphi) \mathbf{V} \mathcal{A}] . \quad (31)$$

If we now choose φ to be such that $\mathbf{R}(\varphi) \mathbf{V} \mathcal{A}$ is symmetric, then $\mathbf{V}_s = \mathbf{R}(\varphi) \mathbf{V} \mathcal{A}$. After a bit of algebra, we find that the rotation angle φ is given through

$$e^{-i\varphi} = \frac{1 - \epsilon g^*}{|1 - \epsilon g^*|} , \quad (32)$$

and we obtain as final result \mathbf{V}_s as in (21) with

$$t_s = |1 - \epsilon g^*| t ; \epsilon^s = \frac{\epsilon - g}{1 - \epsilon g^*} . \quad (33)$$

The inverse of (33) is given by

$$= \frac{s + g}{1 + sg} : \quad (34)$$

We recover for the GLAM ellipticity exactly the transformation law of unweighted moments (SS97). The GLAM ellipticity is therefore an unbiased estimator of the reduced shear along the line-of-sight of the galaxy, and there is no need to determine unweighted moments.

As a side remark, the transformation between s and s is a linear conformal mapping from the unit circle onto the unit circle, and from the origin of the plane onto the point g in the s -plane. If $|g| > 1$, then g has to be replaced by $1/g$ in Eq. (34), but we shall not be concerned here with this situation in the strong lensing regime.

2.4. Point spread function and pixellation

We have defined the GLAM ellipticity of an image (x) relative to an adaptive weight (w) . This definition is idealized in the sense that it assumes an infinite angular resolution and absence of any atmospheric or instrumental distortion or image blur. Equally important, it ignores pixel noise. In this section we move one step further to discuss the recovery of the original image after it has been convolved with a PSF and pixelated. The problem of properly dealing with noise in the image is discussed subsequently.

Let $I_{\text{pre}}(x)$ be the original image prior to a PSF convolution and pixellation. This we call the 'pre-seeing' image. Likewise, by the vector I_{post} of N_{pix} values we denote the 'post-seeing' image that has been subject to a convolution with a PSF and pixellation. For mathematical convenience, we further assume that $I_{\text{pre}}(x)$ is binned on a fine auxiliary grid with N_{pre} pixels of solid angle Ω_{pre} . We list these pixel values as vector I_{pre} . The approximation of $I_{\text{pre}}(x)$ by the vector I_{pre} becomes arbitrarily accurate for $N_{\text{pre}} \gg 1$. Therefore we express the post-seeing image $I_{\text{post}} = LL^T I_{\text{pre}}$ by the linear transformation matrix L applied to the pre-seeing image I_{pre} . The matrix L with $N_{\text{post}} \times N_{\text{pre}}$ elements combines the effect of a (linear) PSF convolution and pixellation. Similarly, we bin the template $t(x)$ in the pre-seeing frame to the grid of I_{pre} , and we denote the binned template by the vector t_{pre} as usual, the quadratic form is a function of the variables $(x_0; \theta; t)$, Eq. (6). The GLAM parameters p_{pre} of the pre-seeing image are given by the minimum of $E_{\text{pre}}(p_{\text{pre}})$, or approximately by

$$E_{\text{pre}}(p_{\text{pre}}) := (I_{\text{pre}} - A p_{\text{pre}})^T (I_{\text{pre}} - A p_{\text{pre}}) \approx 2^{-1} E(p_{\text{pre}}) : \quad (35)$$

For the recovery of the pre-seeing ellipticity, the practical challenge is to derive the pre-seeing parameters from the observed image I_{post} in the post-seeing frame. For this task, we assume that the transformation is exactly known. Note that a linear mapping L is an approximation here; we ignore the nonlinear effects in the detector (Plazas et al. 2014; Gruen et al. 2015; Niemi et al. 2015; Melchior et al. 2015).

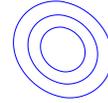
For a start, imagine a trivial case where no information is lost by going from I_{pre} to I_{post} . We express this case by a transformation L that can be inverted, i.e., we have $N_{\text{post}} = N_{\text{pre}}$ and L is regular. We then obtain p_{pre} by minimising $E_{\text{pre}}(p_{\text{pre}} | I_{\text{post}})$ with respect to p_{pre} : we map I_{post} to the pre-seeing frame and analyse $I_{\text{pre}} = L^{-1} I_{\text{post}}$ there. This is equivalent to minimising the form $(I_{\text{post}} - A L p_{\text{pre}})^T (L L^T)^{-1} (I_{\text{post}} - A L p_{\text{pre}})$ in the post-seeing frame.

For realistic problems where L^{-1} does not exist, because $N_{\text{post}} < N_{\text{pre}}$, this trivial case at least suggests to determine the

high-res template



high-res template with PSF



pixelated template

Fig. 1. Examples of GLAM templates in the pre-seeing frame (top left), and the post-seeing frame (other panels); the templates are Gaussian radial profiles with $t(r) = e^{-r^2}$. The bottom left panel simulates only pixellation, whereas the right column also shows the impact of a PSF, indicated in the top right corner, without (top) and with (bottom) pixellation.

pixelated template with PSF

minimum p_{post} of the new functional

$$E_{\text{post}}(p_{\text{post}}) := (I_{\text{post}} - A L p_{\text{post}})^T U (I_{\text{post}} - A L p_{\text{post}}) \quad (36)$$

as estimator of p_{pre} . This way we are setting up an estimator by forward-fitting the template t to the image in the post-seeing frame with the matrix U being a metric for the goodness of the fit. Clearly, should L^{-1} exist we recover (35) only by adopting $U = (L L^T)^{-1}$. So we could equivalently obtain p_{pre} without bias, by fitting $L p_{\text{pre}}$ to the observed image I_{post} in this case. However, realistically L is singular: the recovery of p_{pre} from (36) can only be done approximately. Then we could at least find an optimal metric to minimise the bias. We return to this point shortly. In any case, the metric has to be positive-definite and symmetric such that always $E_{\text{post}} \geq 0$. Note that the moments at the minimum of (36) are related but not identical to the adaptive moments in the post-seeing frame. To obtain the latter we would need a pixelated t with $U = 1$ to I_{post} . The bottom and top right images in Fig. 1 display examples of post-seeing templates that are fitted to a post-seeing image to estimate p_{pre} with the functional (36).

For singular L , the minimum of the functional yields an unbiased p_{pre} for any U if

1. $I_{\text{pre}}(x)$ has confocal elliptical isophotes with the radial brightness profile $S(x)$;
2. and if we choose $t(x) = S(x)$ as GLAM template;
3. and if E_{post} has only one minimum (non-degenerate).

To explain, due to 1. and 2. we find a vanishing residual

$$R_{\text{pre}} = I_{\text{pre}} - A p_{\text{pre}} = 0 ; \text{ for } p_{\text{pre}} = p_{\text{pre}} ; \quad (37)$$

at the minimum of E_{pre} and consequently $E_{\text{pre}}(p_{\text{pre}} | I_{\text{pre}}) = 0$. At the same time for any metric U , we also have $E_{\text{post}}(p_{\text{pre}} | I_{\text{post}}) = 0$

because $\mathbf{L}^T \mathbf{U} \mathbf{L} = \mathbf{L}^T \mathbf{R}_{pre} = 0$ for $\mathbf{p} = \mathbf{p}_{pre}$. Because of the lower bound $E_{post} \geq 0$, these parameters \mathbf{p}_{pre} have to coincide with a minimum of E_{post} and hence indeed $\mathbf{p}_{post} = \mathbf{p}_{pre}$. Note that the previous argument already holds for the weaker condition $\mathbf{L}^T \mathbf{R}_{pre} = 0$ so that a mismatch between \mathbf{p}_{pre} and the template at \mathbf{p}_{pre} produces no bias if the mapped residuals vanish in the post-seeing frame. In addition, if this is the only minimum of E_{post} then the estimator \mathbf{p}_{post} is also uniquely defined (condition 3). An extreme example of a violation of condition 3 is the degenerate case $N_{pix} = 1$: the observed image consists of only one pixel. Then every parameter set \mathbf{p} produces $E_{post} = 0$, if \mathbf{A} is chosen correspondingly. But this should be a rare case since it is not expected to occur for $N_{pix} > 6$, thus for images that span over more pixels than GLAM parameters.

A realistic pre-seeing image is neither elliptical nor is our chosen template $\tilde{f}(\mathbf{r})$ likely to perfectly fit the radial light profile of the image, even if it were elliptical. This mismatch produces a bias in \mathbf{p} only if \mathbf{L} is singular, and the magnitude of the bias scales with the residual of the template \tilde{f} in the pre-seeing frame. To see this, let $\tilde{\mathbf{f}}_{pre}$ be the best-fitting template \tilde{f} with parameters \mathbf{p}_{pre} . The residual of the template \tilde{f} in the pre-seeing frame is $\mathbf{R}_{pre} = \mathbf{I}_{pre} - \tilde{\mathbf{f}}_{pre}$. In the vicinity of \mathbf{p}_{pre} , we express the linear change of $\mathbf{A}\tilde{f}$ for small $\delta\mathbf{p}$ by its gradient \mathbf{G} at \mathbf{p}_{pre} ,

$$\mathbf{G} = (\mathbf{G}_1; \dots; \mathbf{G}_6) = \mathbf{r}_p(\mathbf{A}\tilde{f}) ; \quad (38)$$

where

$$\mathbf{G}_i = \frac{\partial(\mathbf{A}\tilde{f})}{\partial p_i} \Big|_{\mathbf{p}=\mathbf{p}_{pre}} ; \text{ for } i = 1; \dots; 6 ; \quad (39)$$

Each column \mathbf{G}_i of the matrix \mathbf{G} denotes the change $\Delta\tilde{f}$ with respect to p_i . Therefore, close to \mathbf{p}_{pre} we find the Taylor expansion

$$\begin{aligned} \mathbf{A}\tilde{f} &= \tilde{\mathbf{f}}_{pre} + \sum_{i=1}^6 \mathbf{G}_i \delta p_i + O(\delta p_i \delta p_j) \\ &= \tilde{\mathbf{f}}_{pre} + \mathbf{G} \delta \mathbf{p} + O(\delta p_i \delta p_j) ; \end{aligned} \quad (40)$$

Furthermore, since \mathbf{p}_{pre} is a local minimum of E_{pre} , Eq. (35), we find at the minimum the necessary condition

$$\mathbf{r}_p E_{pre}(\mathbf{p} | \mathbf{p}_{pre}) = 0 \Rightarrow \mathbf{G}^T \mathbf{R}_{pre} = 0 ; \quad (41)$$

This means that the residual \mathbf{R}_{pre} is orthogonal to every \mathbf{G}_i . Now let $\mathbf{R} = \mathbf{L} \mathbf{R}_{pre} = \mathbf{I}_{post} - \mathbf{L} \tilde{\mathbf{f}}_{pre}$ be the residual mapped to the post-seeing frame. Then (36) in the vicinity \mathbf{p}_{pre} is approximately

$$E_{post}(\mathbf{p}_{pre} + \delta \mathbf{p} | \mathbf{p}_{post}) = (\mathbf{R} - \mathbf{L} \mathbf{G} \delta \mathbf{p})^T \mathbf{U} (\mathbf{R} - \mathbf{L} \mathbf{G} \delta \mathbf{p}) ; \quad (42)$$

which we obtain by plugging (40) into (36). This approximation is good if the bias is small, i.e., if the minimum \mathbf{p}_{post} of $E_{post}(\mathbf{p} | \mathbf{p}_{post})$ is close to \mathbf{p}_{pre} . As shown in Aitken (1934) in the context of minimum-variance estimators, the first-order bias $\mathbf{p}_{min} = \mathbf{p}_{post} - \mathbf{p}_{pre}$ that minimises (42) is then given by

$$\mathbf{p}_{min} = \mathbf{G}^T \mathbf{L}^T \mathbf{U} \mathbf{L} \mathbf{G}^{-1} \mathbf{G}^T \mathbf{L}^T \mathbf{U} \mathbf{R} = \mathbf{G}^T \mathbf{U} \mathbf{L} \mathbf{G}^{-1} \mathbf{G}^T \mathbf{U} \mathbf{L} \mathbf{R}_{pre} ; \quad (43)$$

with $\mathbf{U}_L := \mathbf{L}^T \mathbf{U} \mathbf{L}$. This reiterates that the bias \mathbf{p}_{min} always vanishes either for vanishing residuals $\mathbf{R}_{pre} = 0$, or if \mathbf{L}^{-1} exists and we choose $\mathbf{U} = (\mathbf{L} \mathbf{L}^T)^{-1}$ as metric. The latter follows from Eq. (43) with $\mathbf{U}_L = \mathbf{1}$ and Eq. (41). Note that \mathbf{p}_{min} does not change if we multiply \mathbf{U} by a scalar $\gamma > 0$. Thus any metric \mathbf{U} generates as much bias as

With regard to an optimal metric \mathbf{U} , we conclude from Eq. (43) and $\mathbf{G}^T \mathbf{R}_{pre} = 0$ that we can minimise the bias by the choice of \mathbf{U} for which $\mathbf{U}_L = \mathbf{1}$, or if $\|\mathbf{L}^T \mathbf{U} \mathbf{L} - \mathbf{1}\|_F^2$ is minimal with respect to the Frobenius norm $\|\mathbf{Q}\|_F^2 = \text{tr}(\mathbf{Q} \mathbf{Q}^T)$. This optimised choice of \mathbf{U} corresponds to the so-called pseudo-inverse $(\mathbf{L} \mathbf{L}^T)^+$ of $\mathbf{L} \mathbf{L}^T$ (Press et al. 1992). The pseudo-inverse is the normal inverse of $\mathbf{L} \mathbf{L}^T$ if the latter is regular.

A practical computation of $\mathbf{U} = (\mathbf{L} \mathbf{L}^T)^+$ could be attained by choosing a set of orthonormal basis functions in the pre-seeing frame, or approximately a finite number N_{base} of basis functions that sufficiently describes images in the pre-seeing frame. For every basis function, one then computes the image and the matrix

$$\mathbf{L} \mathbf{L}^T = \sum_{i=1}^{N_{base}} \mathbf{b}_i \mathbf{b}_i^T \mathbf{A}^T = \sum_{i=1}^{N_{base}} (\mathbf{L} \mathbf{b}_i) (\mathbf{L} \mathbf{b}_i)^T ; \quad (44)$$

The pseudo-inverse of this matrix is the metric in the post-seeing frame.

Specifically, for images that are only pixelated the matrix $\mathbf{L} \mathbf{L}^T$ is diagonal. To see this, consider pixelations that map \mathbf{e}_i in the pre-seeing frame to a single pixel \mathbf{e}_i^0 in the post-seeing frame, or $\mathbf{e}_i = \mathbf{e}^0(\mathbf{e}_i)$; both \mathbf{e}_i and $\mathbf{e}^0(\mathbf{e}_i)$ are unit vectors from the standard bases in the two frames. According to (44), matrix $\mathbf{L} \mathbf{L}^T = \sum_i \mathbf{e}^0(\mathbf{e}_i) [\mathbf{e}^0(\mathbf{e}_i)]^T$ is then always diagonal, typically proportional to the unit matrix, and easily inverted to obtain the optimal metric \mathbf{U} . Unfortunately, this \mathbf{U} only makes the linear-order bias (43) vanish while we still can have a higher-order bias because of the singular

2.5. Similarity between model-based and moment-based techniques

Through the formalism in the previous sections it becomes evident that there is no fundamental difference between model-based techniques and those involving adaptive moments. Model-based techniques perform fits of model profiles to the observed image, whereas moment-based ellipticities with the adaptive weight $f^0(\mathbf{r})$ are equivalently obtained from the image by fitting the elliptical profile $f(\mathbf{r})$ to the image. This requires, however, the existence of a unique minimum of the function $E(\mathbf{p} | \mathbf{I})$ which we assume throughout the paper. We note that the similarity also extends to the special case of unweighted moments in Eqs. (1) and (2) which are in principle obtained by fitting $f(\mathbf{r}) = \text{since } f^0(\mathbf{r}) = 1$ in this case.

Nonetheless one crucial difference to a model-fitting technique is that a fit of $f(\mathbf{r})$ does not assume a perfect match to $\mathbf{I}(\mathbf{x})$: the functional $E(\mathbf{p} | \mathbf{I})$ needs to have a minimum, but the fit is allowed to have residuals \mathbf{R}_{pre} , i.e., $E(\mathbf{p} | \mathbf{I}) > 0$ at the minimum. As shown, for the estimator based on (36) this may cause bias, Eq. (43), but only when analysing post-seeing images for $L < 1$. In practice, different methodologies to estimate the pre-seeing moment-based ellipticity certainly used different approaches. Our choice of a forward-fitting estimator (36) is very specific but is optimal in the sense that it is always unbiased and $\mathbf{L} \mathbf{R}_{pre} = 0$ or regular \mathbf{L} . Yet other options are conceivable. For instance, we could estimate moments in the post-seeing frame first and try to map those to the pre-seeing frame (e.g., H03 or Melchior et al. 2011 which aim at unweighted moments). It is then unclear which weight is effectively applied in the pre-seeing frame. Therefore the expression Eq. (43) for the bias is strictly applicable only to our estimator and adaptive moments. But it seems plausible that the bias of estimators of pre-seeing moments generally depends on the residual \mathbf{R}_{pre} since the brightness

moments $x_{0,i}$ and M_{ij} are the solutions to a best-fit of elliptical templates.

In the literature the problem of bias due to residuals in models is known as model bias or underfitting bias (Zuntz et al. 2010; Bernstein 2010). Consequently, moment-based techniques are prone to underfitting bias as model-based methodologies.

3. Statistical inference of ellipticity

Realistic galaxy images are superimposed by instrumental noise I . Therefore the pre-seeing GLAM ellipticity can only be inferred statistically with uncertainties, and it is, according to the foregoing discussion, subject to underfitting bias. For a statistical model of the ellipticity, we exploit the previous conclusions according to which the ellipticity of (x) for the adaptive weight $f^0(\cdot)$ is equivalent to of the best-fitting template $f(\cdot)$. This renders the inference of a standard forward-fit of a model $f(\cdot)$ to I .

We consider post-seeing images with Gaussian noise, i.e., $I = I_{\text{post}} + I$. The covariance of the noise is $\Sigma = h | I^T I$, while $I_{\text{post}} = L I_{\text{pre}}$ is the noise-free image in the post-seeing frame. A Gaussian noise model is a fair assumption for faint galaxies in the sky-limited regime (Miller et al. 2007). Possible sources of noise are: read-out noise, sky noise, photon noise, faint objects that blend with the galaxy image. If an appropriate Gaussian model is not applicable, the following model of the likelihood has to be modified accordingly.

The statistical model of noise are given by the likelihood $L(I|p)$ of an image $I = L I_{\text{pre}} + I$ given the GLAM parameters p . We aim at a Bayesian analysis for which we additionally quantify our prior knowledge on parameters by the PDF $P(p)$. We combine likelihood and prior to produce the marginal posterior

$$P(p|I) = \int dA dt d^2x_0 L(I|p) P(p) \quad (45)$$

of ellipticity by integrating out the nuisance parameters $(A; t)$; the constant normalisation of the posterior is irrelevant in this paper but we assume that the posterior is proper (it can be normalised). Our choice for the numerical experiments in this study is a uniform prior $P_p(p)$ for positive sizes and amplitudes A , ellipticities $j < 1$, and centroid positions x_0 inside the thumbnail image. As known from previous Bayesian approaches to shear analyses, the choice of the prior affects the consistency of the ellipticity posteriors (see, e.g., BA14). The origin of prior-dependence will become clear in Sect. 3.3.

With regard to notation, we occasionally have to draw random numbers or vectors of random numbers from a PDF $P(x)$ or a conditional density $P(x|y)$. We denote this by the shorthand $x \sim P(x)$ and $x \sim P(x|y)$, respectively. As common in statistical notation, distinct conditional probability functions may use the same symbol, as for instance the symbols $P(x|y)$ and $P(y|x)$.

3.1. Caveat of point estimates

The bias in a lensing analysis is not only affected by how we statistically infer galaxy shapes but also how we process this statistical information later on. To demonstrate in this context the disadvantage of point estimators in comparison to a fully Bayesian treatment, we consider here a simplistic nonlinear toy model. This model has one parameter and one single observable $y = x^3 + n$ that is subject to noise n . By $n \sim N(0; \sigma)$ we draw random noise from a Gaussian distribution $N(0; \sigma)$ with mean zero

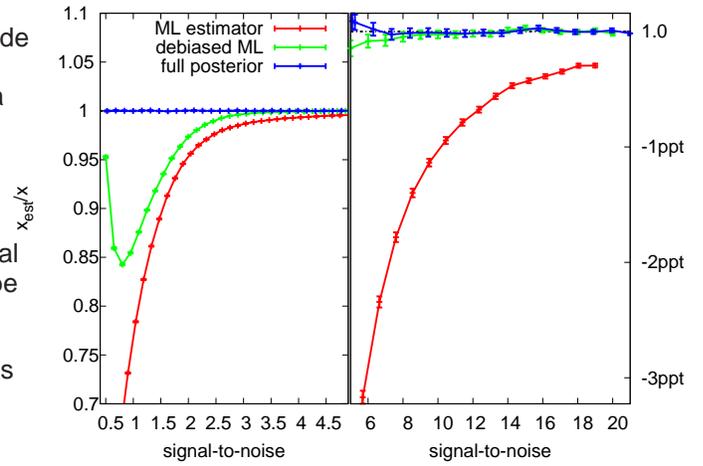


Fig. 2. Toy-model demonstration of a maximum-likelihood estimator (red), a maximum-likelihood estimator with first-order bias correction (green), and an estimator exploiting the full posterior (blue). Data points display the estimator average over 10 data points at varying signal-to-noise levels (x-axis). The true value to be estimated is 1. The panels show different signal-to-noise regimes: 1 ppt denotes $\sigma = 1$, $n = 10^3$.

and variance σ . From the data y , we statistically infer the original value of x . Towards this goal we consider the (log-)likelihood of y given x which is $2 \ln L(y|x) = (y - x^3)^2 / \sigma^2 + \text{const}$.

A maximum likelihood estimator of x is given by $x_{\text{est}} = y^{1/3}$, the maximum of $L(y|x)$. We determine the bias of x_{est} as function of signal-to-noise ratio (SN) $x = y/\sigma$ by averaging the estimates of $N_{\text{real}} = 10^6$ independent realisations of y . The averages and the standard errors are plotted as red line in Fig. 2. Clearly, x_{est} is increasingly biased low towards lower SN levels. In the context of lensing, this would be noise bias. As an improvement we then correct the bias by employing the first-order correction of Refregier et al. (2012) for each realisation of A as seen in the figure, this correction indeed reduces the systematic error, but nevertheless breaks down for $\text{SN} < 3$.

On the other hand in a fully Bayesian analysis, we obtain constraints on x that are consistent with the true value for any SN. For this purpose, we make N_{real} independent identically distributed realisations (i.i.d.) y_i and combine their posterior densities $P_{\text{post}}(x|y_i) / L(y_i|x) P_{\text{prior}}(x)$ by multiplying the likelihoods; we adopt a uniform (improper) prior $P_{\text{prior}}(x) = \text{const}$. This gives us for $y = (y_1; \dots; y_{N_{\text{real}}})$, up to a normalisation constant, the combined posterior

$$\ln P_{\text{post}}(x|y) + \text{const} = \ln L(y|x) = \frac{1}{2} \sum_{i=1}^{N_{\text{real}}} (y_i - x^3)^2 / \sigma^2 \quad (46)$$

As expected due to the asymptotic normality of posteriors (under regularity conditions), for i.i.d. experiments the product density is well approximated by a Gaussian $N(x; \hat{x}, \hat{\sigma}^2)$ and is consistent with the true value x (van der Vaart 1998). We plot values of x_{est} and \hat{x} in Fig. 2 as blue data points.

In conclusion, keeping the full statistical information $P_{\text{post}}(x|y)$ in the inference of x yields consistent constraints over the entire SN range probed: the noise bias vanishes. Also note that the Bayesian approach has not substantially increased the error $\hat{\sigma}_x$ compared to the error of the point estimator (relative sizes of error bars); both approaches have similar efficiency.

Inspired by the foregoing Bayesian toy model that is free of bias, we set up a Bayesian approach for GLAM ellipticities. To construct a likelihood for a GLAM t in the pre-seeing frame, we first consider, similar to Sect. 2.4, the trivial case where L is regular. This is straightforward since we can map the image $I = L^{-1}I'$ back to the pre-seeing frame and determine the noise residual for given pre-seeing

$$I_{\text{pre}}(p) = L^{-1}I' - R_{\text{pre}} - Af : \quad (47)$$

The inverse noise covariance in the pre-seeing frame is $L^{-1}L^{-T}$. The logarithmic likelihood of $I_{\text{pre}}(p)$ in the Gaussian case is thus

$$\begin{aligned} 2 \ln L_{\text{pre}}(I'jp) + \text{const} &= I_{\text{pre}}(p)^T L^{-1} L^{-T} L^{-1} I_{\text{pre}}(p) \\ &= (I' - ALf - LR_{\text{pre}})^T N^{-1} (I' - ALf - LR_{\text{pre}}) \\ &=: k(I' - ALf - LR_{\text{pre}})^2_N ; \end{aligned} \quad (48)$$

where 'const' expresses the normalisation of the likelihood. Therefore, we can equivalently write the pre-seeing t in terms of available post-seeing quantities. Note here that the scalar LR_{pre} of the (unknown) pre-seeing residual is for singular L and $R_{\text{pre}} \neq 0$ not equal to the post-seeing residual R_{post} which we defined in a least-square fit of I' to I_{post} (see Sect. 2.4).

In reality, L is singular so the previous steps cannot be applied. Nevertheless, Eq. (48) is the correct specification of a forward t if the model $m(p) := Af + R_{\text{pre}}$ perfectly describes the brightness profile I_{pre} for some parameters p_{true} . Therefore we expect no inconsistencies for singular L as long as the correct R_{pre} can be given. Since R_{pre} is unknown, however, we investigate in the following the impact of a misspecified likelihood that does not properly account for our ignorance in the pre-seeing residuals. We do this by assuming $R_{\text{pre}} = 0$ and employing

$$2 \ln L(I'jp) + \text{const} = k(I' - ALf)^2_N : \quad (49)$$

Obviously, this is a reasonable approximation of (48) if $kLR_{\text{pre}} \ll k(I' - ALf)$ so that residuals are only relevant if they are non-small when compared to the image in the post-seeing frame. It also follows from Eq. (48) that for $R_{\text{pre}} = 0$ the likelihood is correctly specified even if $R_{\text{pre}} \neq 0$ and L being singular, similar to the noise-free case. More generally we discuss later on how we could modify the likelihood $L(I'jp)$ to factor in our insufficient knowledge about R_{pre} . Until then the approximation (49) introduces under tting bias into the likelihood model, making it inconsistent with the true

To show the inconsistency, we proceed as in the foregoing section on the toy model. We consider a series of d. realisations I_i of the same image I_{post} and combine their likelihoods $L(I_ijp)$ for a given p into the joint (product) likelihood

$$2 \ln \prod_{i=1}^n L(I_ijp) + \text{const} = \sum_{i=1}^n k(I_i - ALf)^2_N ; \quad (50)$$

and we work out its limit for $n \rightarrow \infty$. The joint likelihood can be written as

$$\begin{aligned} 2 \ln \prod_{i=1}^n L(I_ijp) + \text{const} &= \sum_{i=1}^n \left[-\frac{1}{2} (I_i - ALf)^T N^{-1} (I_i - ALf) \right] \\ &= -\frac{1}{2} \sum_{i=1}^n (I_i^T N^{-1} I_i - 2 I_i^T N^{-1} ALf + (ALf)^T N^{-1} ALf) \\ &= -\frac{1}{2} \left[\sum_{i=1}^n I_i^T N^{-1} I_i - 2 \sum_{i=1}^n I_i^T N^{-1} ALf + n (ALf)^T N^{-1} ALf \right] \end{aligned} \quad (51)$$

Here we have made use of the properties of the trace of matrices, namely its linearity $\text{tr}(A+B) = \text{tr}(A) + \text{tr}(B)$ and that $\text{tr}(I_i^T A I_i) = \text{tr}(I_i^T A I_i) = \text{tr}(A I_i I_i^T)$. For large n , we can employ the asymptotic expressions

$$\sum_{i=1}^n I_i^T N^{-1} I_i \approx n I_{\text{post}}^T N^{-1} I_{\text{post}} + \sum_{i=1}^n I_i I_i^T N^{-1} \quad (52)$$

to the following effect:

$$\begin{aligned} 2 \ln \prod_{i=1}^n L(I_ijp) + \text{const} &= -\frac{1}{2} \left[n \text{tr}(I_{\text{post}}^T N^{-1} I_{\text{post}}) + n \text{tr}(I_{\text{post}}^T N^{-1} ALf) \right. \\ &\quad \left. + n \text{tr}(ALf^T N^{-1} ALf) - 2 \sum_{i=1}^n I_i^T N^{-1} ALf \right] \\ &= -\frac{1}{2} \left[n \text{tr}(I_{\text{post}}^T N^{-1} I_{\text{post}}) + n \text{tr}(ALf^T N^{-1} ALf) \right] \end{aligned} \quad (53)$$

Thus the joint likelihood peaks for $n \rightarrow \infty$ at the minimum of $k(I_{\text{post}} - ALf)^2_N$. This is equivalent to the location of the minimum of $E_{\text{post}}(p|I_{\text{post}})$, Eq. (36), with metric $\mathcal{U} = N^{-1}$. Consequently, as in the noise-free case, we find an inconsistent likelihood if the residual R_{pre} is non-vanishing and if $L^{-1}L^{-T}$ is not proportional to the unity matrix. We additionally expect a smaller bias $p = p_{\text{post}} - p_{\text{pre}}$ for smaller levels of residuals.

With regards to the noise dependence of under tting bias, we find that p does not change if we increase the overall level of noise in the image. This can be seen by scaling the noise covariance $N \rightarrow \gamma N$ with a scalar γ . Any value $\gamma > 0$ results in the same minimum location for $E_{\text{post}}(p|I_{\text{post}})$ so that p is independent of γ : there is no noise bias.

Moreover, the bias p of a misspecified likelihood seems to depend on our specific assumption of a Gaussian model for the likelihood. It can be argued on the basis of general theorems on consistency and asymptotic normality of posteriors, that for $n \rightarrow \infty$ we obtain the same results for other likelihood models under certain regularity conditions (van der Vaart & Tibshirani, Appendix B in Gelman et al. 2013). The latter requires continuous likelihoods that are identifiable, hence $L(I'jp_1) \neq L(I'jp_2)$ for $p_1 \neq p_2$, and that the true p_{true} is not at the boundary of the domain of all parameters. This is stricter than our previous assumptions for the Gaussian model where we needed only a unique global maximum of the likelihood $L(I_{\text{post}}jp)$. It could therefore be that a unique maximum is not sufficient for non-Gaussian likelihoods.

¹ Incidentally, the bias p in the likelihood and the under tting bias in the noise-free case are equal to $(L^{-1}L^{-T})^{-1} / (L^{-1}L^{-T})^+$, which for regular L is equivalent to homogeneous, uncorrelated noise in the pre-seeing frame, which means we have $L^{-1}L^{-T} = I$.

The foregoing section discusses the consistency of the likelihood of a single image. We can interpret the analysis also in a different way: if we actually had independent exposures of the same pre-seeing image, then combining the information of all exposures results in a posterior $P(p|I) / \prod_{i=1}^n L(I_{ij}|p) P_p(p)$ that is consistent with p_{post} for a uniform prior $P_p(p) = 1$. As discussed in van der Vaart (1998), the more general Bernstein-von Mises theorem additionally shows that under regularity conditions the choice of the prior is even irrelevant provided the prior mass around p_{post} (Cromwell's rule). This might suggest that for a correctly specified likelihood $L(I_{ij}|p)$, a fully Bayesian approach for the consistent measurement might be found that is independent of the specifics of the prior and has no prior bias in the sense of Sect. 3.2. This is wrong as shown in the following.

Namely, in contrast to the previous simplistic scenario, sources in a lensing survey have varying values: they are intrinsically different. For more realism, we therefore assume now $i = 1 : : n$ pre-seeing images that, on the one hand, shall have different centroid positions $x_{0,i}$, sizes t_i , amplitudes A_i but, on the other hand, have identical ellipticities. Our goal in this experiment is to infer from independent image realisations $I_i = I_{\text{post}i} + I_i$ by marginalizing over the nuisance parameters $q_i = (x_{0,i}; t_i; A_i)$. This experiment is similar to the standard test for shear measurements where a set of pre-seeing images is considered whose realisations are subject to the same amount of shear (e.g., Bridle et al. 2010). As a matter of fact, the inference of constant shear from an ensemble of images would just result in additional nuisance parameters for the intrinsic shapes with essentially the same following equations.

Let $P_q(q_i)$ be the prior for the four nuisance parameters of the i th image and $P(\cdot) = 1$ a uniform prior for \cdot . We combine the GLAM parameters $\text{inp}_i := (q_i; \cdot)$, and we assume that all images have the same prior density and that the noise covariance N applies to all images. The marginal posterior is then the integral

$$N P(j|I_1; \dots; I_n) = \int \prod_{i=1}^n d^4 q_i L(I_{ij}|p) P_q(q_i) P(\cdot) \\ = \int \prod_{i=1}^n d^4 q_i \dots \prod_{i=1}^n L(I_{ij}|p) \prod_{i=1}^n P_q(q_i); \quad (54)$$

with N being a normalization constant.

The product of the likelihood densities inside the integral is given by

$$2 \ln \prod_{i=1}^n L(I_{ij}|p) + \text{const} \\ = n \text{tr} \left[\frac{1}{n} \sum_{i=1}^n I_i I_i^T \right] - 2 \text{tr} \left[\sum_{i=1}^n I_i^T (A_i L f_{\cdot}; i) \right] \\ + \sum_{i=1}^n k A_i L f_{\cdot}; i k_N^2; \quad (55)$$

Here we have taken into account that the GLAM parameters are partly different, indicated by the additional index A_i and $f_{\cdot}; i$. This is different in Eq. (53) where we take the product of full likelihoods in p -space without marginalization. In the limit of $n \rightarrow \infty$,

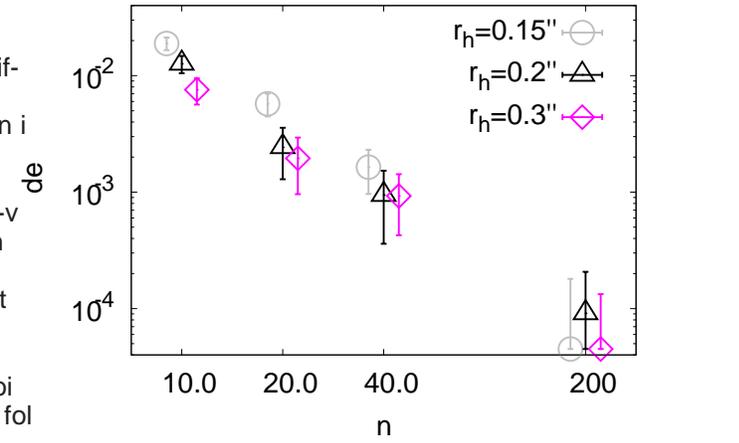


Fig. 3. Prior bias in the marginal posterior $P(j|I_1; \dots; I_n)$ as function of SN for different galaxy sizes, (in arcsec). The posterior assumes a uniform prior. Shown is the error $\sigma_j = |j - j_{\text{true}}|$ of the inferred for a true $j_{\text{true}} = 0.3$ as obtained by combining the marginal posteriors of $5 \cdot 10^3$ exposures of the same galaxy with random centroid positions. The pixel size 0.1 arcsec equals the PSF size (M0 Galaxy profiles and GLAM templates have a Sérsic profile with $n = 2$: there is no underfitting).

we find in addition to the relations (52) that

$$\text{tr} \left[\frac{1}{n} \sum_{i=1}^n I_i^T (A_i L f_{\cdot}; i) \right] = \text{tr} \left[\frac{1}{n} \sum_{i=1}^n (I_{\text{post}i} + I_i)^T (A_i L f_{\cdot}; i) \right] \\ = \text{tr} \left[\frac{1}{n} \sum_{i=1}^n I_{\text{post}i}^T A_i L f_{\cdot}; i \right] + \text{tr} \left[\frac{1}{n} \sum_{i=1}^n I_i^T A_i L f_{\cdot}; i \right] \quad (56)$$

because I_i is uncorrelated to $A_i L f_{\cdot}; i$ so that $I_i^T (A_i L f_{\cdot}; i)$ vanishes on average for many I_i . Therefore, for the asymptotic statistic we can replace all I_i by $I_{\text{post}i}$ in Eq. (55) to obtain

$$\prod_{i=1}^n L(I_{ij}|p) = \prod_{i=1}^n L(I_{\text{post}i}|p); \quad (57)$$

and, as a result, for Eq. (54)

$$N P(j|I_1; \dots; I_n) = \int \prod_{i=1}^n d^4 q_i L(I_{\text{post}i}|p) P_q(q_i) =: P(j|I_{\text{post}}); \quad (58)$$

where $P(j|I_{\text{post}})$ is the marginal ellipticity posterior for this noise-free image.

The limit (58) has the interesting consequence that the consistency with p_{post} of the marginal posterior depends on the specific choice of the prior density $P_q(q_i)$. To show this, consider one particular case in which, for simplicity, all pre-seeing images are identical such that $I_{\text{post}i} = I_{\text{post}}$. Then, according to (58), the ellipticity posterior converges in distribution to $\mathcal{P}(j|I_{\text{post}})^n$ which for $n \gg 1$ peaks at the global maximum of $\mathcal{P}(j|I_{\text{post}})$. It is then easy to see that we can always change the positions of the maximum by varying the prior density $\mathcal{P}(j|I_{\text{post}})$. In particular, even if the likelihoods are correctly specified, we generally find an inconsistent marginal posterior depending on the prior $\mathcal{P}(j|I_{\text{post}})$. A similar argument can be made for $I_{\text{post}i}, I_{\text{post}j}$ for i, j .

For a concrete example, we perform for Fig. 3 a simulated analysis of $5 \cdot 10^3$ noisy images with $j_{\text{true}} = 0.3$. All pre-seeing

images are identical to one particular template (Sérsic profile with $n = 2$). Therefore we have a correctly specified likelihood model and no underfitting. We adopt a uniform prior $P(q)$ (and $P(\gamma)$). The details on the simulated images and their analysis are given in Sect. 4. For each data point, we plot the mean and variance of the marginal posterior $P(\gamma | j_1, \dots, j_n)$ relative to the true ellipticity γ_{true} as function of SN and for different image sizes n . Evidently, the bias increases for smaller n , thereby producing a noise-dependent bias which is not present when analysing the consistency of individual images as in Sect. 3.2.

The sensitivity of the posterior to the prior implies that consistency (for a correctly specified likelihood) could be improved by choosing an appropriate prior density $P(q)$. On the other hand, the dependence on the prior runs against the common wisdom that the prior should become asymptotically irrelevant for $n \gg 1$ as the joint likelihood starts dominating the information on inferred parameters. Indeed, general theorems show under certain regularity conditions (see, e.g., discussion in Chapter 4 of Gelman et al. 2013). These conditions are, however, given here because the total number of model parameters is fixed but rather increases linearly with n due to a new set of nuisance parameters for every new galaxy image. The observed breakdown of consistency is the result. A trivial (but practically useless) prior to fix the problem is one that puts all prior mass at the true values of the nuisance parameters which essentially leaves only as free model parameter. Likewise, an analysis with sources that knowingly have the same values for the nuisance parameters also yields consistent constraints; this is exactly what is done in Sect. 3.2. A non-trivial solution, as reported by BAIT, is to use 'correct priors' for the nuisance parameters that equal to the actual distribution of q in the sample. On the downside, this raises the practical problem of obtaining correct priors from observational data.

In summary, for an incorrect prior $P(q)$, such as our uniform prior, we find a noise-dependent bias in the inferred ellipticity despite a fully Bayesian approach and a correctly specified likelihood. We henceforth call this bias 'prior bias'. This emphasizes the difference to the noise bias found in the context of point estimators.

4. Simulated shear bias

We have classified two possible sources of bias in our Bayesian implementation of adaptive moments: underfitting bias and prior bias. In this section, we study their impact on the inference of reduced shear by using samples of mock images with varying SN, galaxy sizes, and galaxy brightness profiles. In contrast to our previous experiments, the images in each sample have random intrinsic shapes but are all subject to the same amount of shear, and we perform numerical experiments to quantify the bias. Moreover, this section outlines practical details of a sampling technique for the posteriors of GLAM ellipticities at the reduced-shear, which is of interest for future applications (Sect. 4.4 and Sect. 4.5).

The PSF shall be exactly known for these experiments. For a large shear bias, we choose a relatively small PSF size for galaxy images that are not much larger than a few pixels. According to the discussion in Sect. 2.4, underfitting bias is a result of pixellation which mathematically cannot be fixed. Note that a larger PSF size would reduce the underfitting bias since it spreads out images over more image pixels. Overall, the values for shear bias presented here are larger than what is

typically found in realistic surveys (e.g., Zuntz et al. 2013). We start this section with a summary of our simulation specifications

4.1. Point-spread function

If not stated otherwise, our simulated postage stamps consist of 20×20 pixels (squares), of which one pixel has a size of $1''$. The simulated galaxies are relatively small with a half-light radius r_{hl} of only a few times the pixel size ($0.5 < r_{\text{hl}} < 0.3$). Additionally, we adopt an isotropic Moffat PSF,

$$I_{\text{psf}}(x) \propto \frac{1}{1 + \frac{|x|^2}{2}}^{-1}; \quad \sigma := \frac{\text{FWHM}}{2 \sqrt{2^{1/n} - 1}}; \quad (59)$$

with the full width half maximum (FWHM) of $\text{FWHM} = 0.1''$ and $n = 5$ (Mo et al 1969).

4.2. GLAM template profiles

As GLAM templates we employ truncated Sérsic-like profiles with index n ,

$$f(r) = \exp\left(-\frac{r}{r_0}\right) \frac{1}{\Gamma(n)} \left(\frac{r}{r_0}\right)^{n-1}; \quad h(x) := \frac{1}{e^{5(x/3)^3} + 1} \quad (60)$$

$$f(r) := (1.992n - 0.3271)^{2n} \quad (61)$$

(Capaccioli 1989). To avoid a numerical bias due to aliasing at the edges of the grid during the Fourier transformations step of the sampling code, we have introduced the auxiliary function $h(x)$ that smoothly cuts off the Sérsic profile beyond $r = 9$. If $f(r)$ were the radial light profile of an elliptical galaxy, the truncation would be located at about three half-light radii.

We use template profiles with $n = 2$ throughout. This index falls between the values of used for the model galaxies and thus is a good compromise to minimise the fit residual and therefore underfitting bias.

4.3. Mock galaxy images

We generate postage stamps of mock galaxy images with varying half-light radii r_{hl} , radial light profiles, and signal-to-noise ratios. To this end, we utilise the code that computes the post-seeing GLAM templates (Appendix B). We consider pre-seeing images of galaxies with elliptical isophotes of three kinds and light profiles, Eq. (60): (1) exponential profiles with Sérsic index $n = 1$ (EXP; exponential), (2) de-Vaucouleur profiles with $n = 4$ (DEV; bulge-like) and (3) galaxies with profiles $n = 2$ that match the profile of the GLAM template (TMP; template-like). TMP galaxies hence cannot produce underfitting bias.

We devise uncorrelated Gaussian noise in the simulation of postage stamps. To determine the RMS variance of the pixel noise for a given γ , let f_p be the flux inside image pixels that is free of noise and $f_{\text{tot}} = \sum_i f_i$ the total flux. From this, we compute a half-light flux threshold f_{th} defined such that the integrated flux $f_{\text{hl}} = \sum_{f_i > f_{\text{th}}} f_i$ above the threshold is $f_{\text{hl}} = f_{\text{tot}}/2$ or as close as possible to this value. The pixels with $f_i > f_{\text{th}}$ are defined to be within the half-light radius of the image; their number is $N_{\text{hl}} := \sum_{f_i > f_{\text{th}}} 1$; the integrated noise variance within the half-light radius is $\sigma_{\text{hl}}^2 = \sum_{f_i > f_{\text{th}}} \sigma^2$. The signal-to-noise ratio within the half-light radius is therefore $\text{SN}_{\text{hl}} = f_{\text{hl}} / \sigma_{\text{hl}} = \frac{f_{\text{hl}}}{\sigma \sqrt{N_{\text{hl}}}} \approx \frac{f_{\text{hl}}}{\sigma \sqrt{N_{\text{hl}}^{1-2}}} \frac{1}{\sigma_{\text{rms}}}$, or

$$\text{SN}_{\text{hl}} = \frac{f_{\text{hl}}}{\sigma \sqrt{N_{\text{hl}}}}; \quad (62)$$

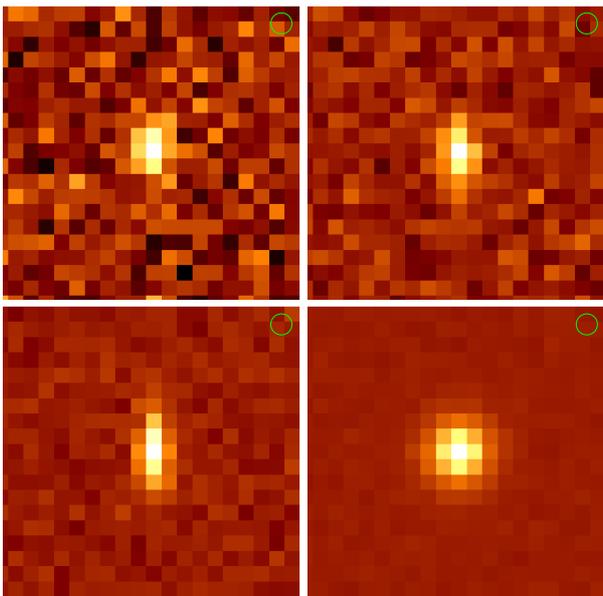


Fig. 4. Examples of simulated images of galaxies with random ellipticities. Signal-to-noise ratios from left to right and top to bottom: $\nu=10$, 20, 40, and 200. The radial light profile of the pre-seeing images is exponential with $r_h = 0.2''$; the pixel size is $0.1''$. The FWHM of the PSF is indicated by the circle in the top right corners.

Figure 4 depicts four examples with added noise for different ν .

To simulate galaxy images with intrinsic ellipticities that are sheared by g , we make a random realisation of an intrinsic ellipticity ϵ_s and compute the pre-seeing ellipticity with (17). As PDF for the intrinsic ellipticities we assume a bivariate Gaussian with a variance of $\sigma_\epsilon = 0.3$ for each ellipticity component; we truncate the PDF beyond $|\epsilon_s| \geq 1$.

In realistic applications, centroid positions of galaxy images are random within an image pixel, so we average all our following results for bias over subpixel offsets within the quadratic solid angle of one pixel at the centre of the image. This means: in an sample of mock images, every image has a different subpixel position which is chosen randomly from a uniform distribution. In this averaging process, care has to be taken to perform an even sampling of subpixel offsets. To ensure this with a finite number of random points in all our following tests, we employ a subrandom Sobol sequence in two dimensions for the random centroid offsets (Press et al. 1992).

4.4. Monte-Carlo sampling of ellipticity posterior

For practical applications of a Bayesian GLAM analysis, we produce a Monte Carlo sample of the posterior $P_\epsilon(\epsilon|\mathbf{I})$, Eq. (45) with the approximate likelihood $\mathcal{L}(\mathbf{I}|\mathbf{p})$ in Eq. (49). This sample consists of a set (ϵ_i, w_i) of $1 \leq i \leq N_{\text{real}}$ pairs of values ϵ_i and w_i , which determine the sampling position ϵ_i and a sampling weight w_i . For this paper, we use $N_{\text{real}} = 50$ sampling points. In contrast to a lensing analysis with single-valued point estimators of ellipticity, a Bayesian analysis employs the sample (ϵ_i, w_i) of each galaxy. To attain the sample (ϵ_i, w_i) we invoke the importance sampling technique (e.g., Marshall 1956; Kilbinger et al. 2010).

For this technique, we define an approximation $Q(\mathbf{p})$ of $P_p(\mathbf{p}|\mathbf{I})$, the so-called importance distribution function, from which we draw a random sample \mathbf{p}_i . The ellipticity component ϵ_i of \mathbf{p}_i is then assigned the weight $w_i = P_p(\mathbf{p}_i|\mathbf{I}) Q(\mathbf{p}_i)^{-1}$ which

we normalise to $\sum_i w_i = 1$ afterwards. As importance function, we use a multivariate Gaussian with mean at the maximum \mathbf{p}_{ml} of the likelihood $\mathcal{L}(\mathbf{I}|\mathbf{p})$ and a covariance defined by the inverse Fisher matrix \mathbf{F}^{-1} at the maximum (Fisher 1935; Tegmark et al. 1997). More implementation details are given in Appendix B.

If $Q(\mathbf{p})$ is too different from $P_p(\mathbf{p}|\mathbf{I})$, the sample will be dominated by few points with large weights, usually due to sampling points in the tail of the posterior for which the importance function predicts a too low probability, i.e., w_i becomes large. This is indicated by a small effective number $N_{\text{eff}} = (\sum_i w_i^2)^{-1}$ of points compared to N_{real} . This can produce a poor convergence and thus extreme outliers in the analysis that merely appear to have tight constraints on ϵ . Typically affected by this are images that are both small compared to the pixel size and low in signal-to-noise. These images tend to exhibit posteriors that allow parameter solutions with small sizes t compared to the true t or high eccentricities $|\epsilon| \gtrsim 0.7$, giving the posterior a complex, distinctly non-Gaussian shape. In future applications, this problem can be addressed by finding a better model of the importance function.

For the scope of this paper, we address this problem at the expense of computation time by increasing the number of sampling points and by resampling. This means: for Monte-Carlo samples with $N_{\text{eff}} < N_{\text{real}}/2$, we draw more samples from the importance distribution until $N_{\text{eff}} \geq N_{\text{real}}/2$ in the expanded sample. By an additional rule, we stop this process if the expanded sample size becomes too large and reaches $10 \times N_{\text{real}}$. This case is indicative of a failure of the importance sampling. If this happens, we switch to a time-consuming standard Metropolis algorithm to sample the posterior with 10^3 points after 10^2 preliminary burn-in steps (Metropolis et al. 1953). This technique does not assume a particular shape for the posterior but performs a random walk through the \mathbf{p} space, thereby producing a correlated sample along a Monte-Carlo Markov chain. As symmetric proposal distribution for this algorithm, we adopt a multivariate Gaussian with covariance $1.5 \times \mathbf{F}^{-1}$; all points of the chain have equal weight w_i ; the starting point of the chain is \mathbf{p}_{ml} . Finally, in the resampling phase, we bootstrap the expanded or Metropolis sample by randomly drawing N_{real} points from it with probability w_i (with replacement). All selected data points are given the equal weights $w_i = N_{\text{real}}^{-1}$ in the final catalogue.

4.5. Posterior density of reduced shear

For the inference of g , we convert the ellipticity posterior $P_\epsilon(\epsilon|\mathbf{I})$ into a posterior $P_g(g|\mathbf{I})$ of shear. To this end, let us first determine the $P_g(g|\epsilon)$ for an exactly known ϵ . We express our uncertainty on the intrinsic ellipticity ϵ_s by the prior $P_s(\epsilon_s)$, and $P_g(g)$ is our a-priori information on g . The values of g and ϵ_s shall be statistically independent, i.e., $P_{\text{sg}}(\epsilon_s, g) = P_s(\epsilon_s) P_g(g)$.² Applying a marginalization over ϵ_s and then Bayes' rule, we find

$$P_g(g|\epsilon) = \int d^2\epsilon_s P_{\text{sg}}(\epsilon_s, g|\epsilon) = \int d^2\epsilon_s \frac{P_\epsilon(\epsilon|g, \epsilon_s) P_{\text{sg}}(\epsilon_s, g)}{N(\epsilon)}, \quad (63)$$

² This assumption would be false if shear and intrinsic ellipticity were correlated, for instance due to selection effects related to the shape of a sheared image (Hirata & Seljak 2003; Heymans et al. 2006). Thus correlations between intrinsic shapes and shear may affect Bayesian methodologies already at this level.

or, equivalently,

$$\begin{aligned} N(\cdot) P_g(g_j) &= \int d^2 s P(g; s) P_g(g) P_s(s) \\ &= \int d^2 s_D (g; s) P_g(g) P_s(s) \\ &= P_g(g) P_s(s; g) \frac{d^2 s}{d^2} : \end{aligned} \quad (64)$$

By $N(\cdot)$ we denote the normalisation $\int d^2 g P_g(g)$,

$$N(\cdot) := \int d^2 g P_g(g) P_s(s; g) \frac{d^2 s}{d^2} : \quad (65)$$

The normalisation $N(\cdot)$ only depends on the modulus of g for isotropic priors for symmetry reasons. The integration over the Dirac delta function $\delta(x)$ uses Eq. (17). The determinant $|\frac{d^2 s}{d^2 g}| = |\frac{d^2 s}{d^2 g}|$, the Jacobian of the linear conformal mapping $(g; s)$, is a function of g and s . For the weak lensing regime $|g| \ll 1$ of interest, this is

$$\frac{d^2 s}{d^2} = \frac{(1 - |g|^2)^2}{|1 - g|^4} \quad (67)$$

(Geiger & Schneider 1998). To now account for the measurement error of g in the shear posterior, we marginalise $P_g(g)$, Eq. (65), over $P(g|j)$ as error distribution,

$$\begin{aligned} P_g(g|j) &= \int d^2 P_g(g) P(j|g) \\ &= \int d^2 P_g(g) (1 - |g|^2)^2 \int \frac{P_s(s; g) P(j|g)}{N(\cdot) |1 - g|^4} : \end{aligned} \quad (68)$$

In our Monte Carlo scheme, we sample the ellipticity posterior of every galaxy through $(i_j; w_{ij}) P(j|i_j)$ by $1 - |g|$ N_{real} values i_j of ellipticity and statistical weight w_{ij} . To convert this sample to an approximation of the posterior, we replace $P(j|i_j)$ inside the integral (68) by the point distribution $\sum_j w_{ij} \delta(g - i_j)$,

$$P_g(g|i_j) = P_g(g) (1 - |g|^2)^2 \sum_{j=1}^{N_{\text{real}}} \frac{w_{ij} P_s(s; g; i_j)}{N(i_j) |1 - i_j|^4} : \quad (69)$$

For the following, we adopt a uniform prior for g , i.e., $P_g(g) / H(1 - |g|)$. For reasons discussed in Sect. 3.3, the prior on the nuisance s presumably has to be equal to the distribution of intrinsic shapes for a consistent shear posterior. To investigate the sensitivity of shear bias as a function of s we therefore use two types of priors: the correct prior, which is the true intrinsic-shape distribution in Sect. 4.3, or a uniform prior $P_s(s)$. For each prior $P_s(s)$, we numerically compute $N(\cdot)$ for different s and interpolate between them later on in (69). We finally combine the posteriors $P_g(g|i_j)$ of all images i_j in the sample by multiplying posterior values on a g -grid. For this, we apply (69) to every galaxy in the sample independently.

4.6. Results

For each experiment, we consider a sample of statistically independent galaxy images $\hat{c}_i = \{f_{i1}, \dots, f_{in}\}$ that are subject to the same reduced shear but have random intrinsic ellipticities. To infer g from \hat{c} in a fully Bayesian manner, we compute the

Table 1. Multiplicative bias in per cent for three profiles TMP, EXP, DEV, and sizes $r_h = 0:2^{00}$ only. The PSF size is $0:2^{00}$ which is larger than in Fig. 5. Values are quoted in per cent for different α and for a correct prior $P_s(s)$. The 1- σ posterior errors are given only for TMP galaxies but are similar for the other galaxy models.

	TMP		EXP	DEV
10	22:1	5:0	12:8	17:8
20	3:4	4:4	+1:6	11:6
40	1:7	4:2	+5:3	6:8
200	+0:2	4:1	+5:0	4:8

posteriors $P_g(g|i_i)$ of g for every image i_i separately with Eq. (69) and then combine all posteriors by the product

$$P_g(g|\hat{c}) = \prod_{i=1}^{N_i} P_g(g|i_i) : \quad (70)$$

This test is related, but not identical, to the ring test in Kijima & Bernstein (2007): we do not fix the ellipticity magnitude $|g|$ during one test run but instead use a random sample of intrinsic ellipticities. Each galaxy sample consists of $n = 5 \cdot 10^4$ simulated noisy galaxy images to typically attain a precision of the order 10^3 . This number applies for the non-uniform prior $P_s(s)$. The precision is lower for the less informative uniform prior. Moreover, for each galaxy with intrinsic g we also use s in another image belonging to the sample. By doing so the mean of intrinsic ellipticities in a sample is always zero, which substantially reduces the numerical noise in the result. For every set of image parameters, we vary the input shear between the three values $g_{\text{true}} \in \{2/0; 0/05; 0/1\}$. For these three measurements, let $\langle g|g_{\text{true}} \rangle$ be the mean of the posterior (70) for a given g_{true} of the sample. We then determine the multiplicative bias and the additive bias in

$$\langle g|g_{\text{true}} \rangle = (1 + m) g_{\text{true}} + c \quad (71)$$

by best-fitting m and c to the mean and dispersion of (70) for the three input values of g_{true} ; m shall be a real-valued number.

In Fig. 5, we plot m (filled data points) juxtaposed with the corresponding values based on a uniform prior (open data points); same data points indicate the same size while colours, varying by row, indicate the light profiles of the images. The numbers in the table and the data points in the figure, errors bar quantify a 68% posterior interval of the measurement around the quoted mean. The size of error bars account for pixel noise and shape noise. For the range of α studied here, the impact of pixel noise is subdominant: the errors increase only between 10% a 20% from $\alpha = 200$ to $\alpha = 10$; the increase is larger for small images. Values for the additive bias (not shown) for all profiles are small within a few 10^4 , except at $\alpha = 10$ for DEV with $r_h = 0:15^{00}, 0:2^{00}$ and for TMP at $0:15^{00}$ there can reach negative amplitudes of the order of 10^4 . This indicates that an additive bias could be relevant for image sizes close to the p size.

In contrast, a multiplicative bias is typically significant. At high $S/N = 200$, this bias in Fig. 5 is mainly underfitting bias which is consequently absent for TMP galaxies since the likelihood is correctly specified. The shear of bulge-like galaxies DEV, with steeper slope compared to $\alpha = 10$, is underestimated by $m = 9:6 \pm 0:5\%$; the shear of exponential galaxies EXP, which have a shallower slope, is overestimated by $+7:7 \pm 0:5\%$.

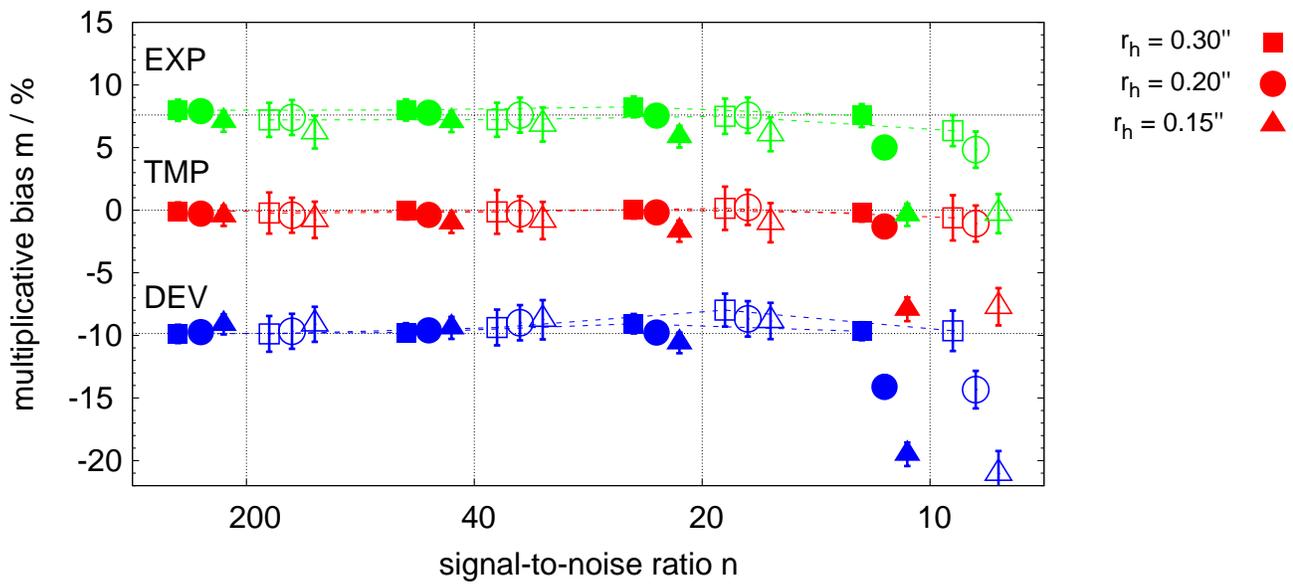


Fig. 5. Plots of the multiplicative bias for simulated images, based on Eq. (69), for two different priors $P_s(s)$ (filled data points: correct prior; open data points: uniform prior). Different styles for the data points indicate different image sizes s_h , see key inside figure, while colors vary with galaxy type: GLAM template (TMP; red); exponential (EXP; green); bulge-like (DEV; blue). Data points for $r_h = 0.30''$, same galaxy type, and same prior are connected by dotted lines to guide the eye. The bias for galaxy sizes, amplitudes A , and centroids x_0 is uniform giving rise to noise-dependent prior bias; the constant bias of EXP and DEV is due to underfitting. A square image pixelates the size $1''$, which also equals the PSF size. Results for a larger PSF size can be found in Table 1.

These numbers are the mean and standard error for all galaxies; (ii) under the influence of shear they behave exactly like unbiased estimators of reduced shear. (i)-(iii) assume ideal conditions with no pixelation, PSF convolution, or pixel noise. These results are fundamentally limited by our ability to determine the ellipticity of galaxies from noisy images (see Sect. 3.3). The noise-dependence of prior bias thus becomes only relevant here at significantly low SN. The particular value of $s_h = 2''$ is specific to the PSF size. This can be seen in Table 1 which lists estimates for galaxies, in a slightly smaller sample $n = 2 \times 10^4$, with size $r_h = 0.20''$ but for a modified, roughly twice as large PSF size $2''$. Now the noise-dependence of prior bias is already visible below $SN = 40$. On the other hand, the underfitting bias, to be read at high $SN = 200$, is roughly halved since the images are spread over more pixels.

Compared to the underfitting bias at $SN = 200$, the bias significantly drops if $SN < 20$ for all galaxies, most prominently at $r_h = 0.15''$, but stays roughly constant within a couple of per cent otherwise. We attribute the noise-dependence of prior bias owing to our choice of an (incorrect) uniform prior for the nuisance parameters of the images (see Sect. 3.3). The prior bias thus becomes only relevant here at significantly low SN. The particular value of $s_h = 2''$ is specific to the PSF size. This can be seen in Table 1 which lists estimates for galaxies, in a slightly smaller sample $n = 2 \times 10^4$, with size $r_h = 0.20''$ but for a modified, roughly twice as large PSF size $2''$. Now the noise-dependence of prior bias is already visible below $SN = 40$. On the other hand, the underfitting bias, to be read at high $SN = 200$, is roughly halved since the images are spread over more pixels.

The choice of a uniform prior for the intrinsic ellipticity is distinctly different from the true distribution of e in the sample, has a negligible effect on shear bias here for both PSF sizes and at same in Fig. 5. Nevertheless, the resulting statistical errors of shear, as opposed to its bias, are larger by roughly a factor 1.7 with the uniform prior, which reflects the larger uncertainty on s_h . This implies that a shear bias due to an incorrect prior $P_s(s)$ is negligible in this experiment.

5. Discussion

Unlike galaxy ellipticities defined in terms of unweighted moments, GLAM provide a practical definition of ellipticity with useful properties for any chosen adaptive weight function: they (i) are identical with the ellipticity of isophotes of elliptical galaxy

images; (ii) under the influence of shear they behave exactly like unbiased estimators of reduced shear. (i)-(iii) assume ideal conditions with no pixelation, PSF convolution, or pixel noise. These results are fundamentally limited by our ability to determine the ellipticity of galaxies from noisy images (see Sect. 3.3). The noise-dependence of prior bias thus becomes only relevant here at significantly low SN. The particular value of $s_h = 2''$ is specific to the PSF size. This can be seen in Table 1 which lists estimates for galaxies, in a slightly smaller sample $n = 2 \times 10^4$, with size $r_h = 0.20''$ but for a modified, roughly twice as large PSF size $2''$. Now the noise-dependence of prior bias is already visible below $SN = 40$. On the other hand, the underfitting bias, to be read at high $SN = 200$, is roughly halved since the images are spread over more pixels.

If the minimum of the functional (3) uniquely exists, the moments $(k_0; M_1)$ of the best-fitting template (\hat{e}) will, for a constant scalar α , be the adaptively weighted moments $(\hat{k}_0; \hat{M}_1)$ of the galaxy light profile $I(x)$ for the weight $f_0(\hat{e})$ (Appendix A). This means that the radial profile of the weight at separation r from the centroid is $w(r) / r^{-1} df_r(r) = dr$ for a template profile $f_r(r) := f(r^2)$. We exploited this relation between adaptive moments and moments of a best-fitting elliptical profile to adaptively moment-based methods under the loss of image information.

Namely, for realistic images subject to pixelation and PSF effects, underfitting bias can be present for GLAM if estimated from the post-seeing image. It potentially arises because of a loss of information on the pre-seeing galaxy image, fundamentally due to pixelation. To explore the limitations, we assume noise-free images and express in Sect. 2.4 the mapping between a pre-seeing and post-seeing image by the linear operator

pre-seeing adaptive moments are equivalent to a least-squares fit of an elliptical profile $f(\theta)$ to the pre-seeing image $I(x)$. If estimated from the post-seeing image whilst ignoring residuals, the inferred ellipticity is biased and a singular. The latter indicates a loss of information on the pre-seeing image. For noise-free images and the estimator the bias is, to linear order, given by Eq. (43). In (36) the bias is optimally reduced through the metric $\tilde{L} = (LL^T)^+$. With this metric the estimator is unbiased for regular noise in principle also for an invertible convolution with an anisotropic PSF; in the convolutions have kernels $K(x)$ with $\tilde{K}(\theta)$, $\theta = 0$ in the Fourier domain. Practically, however, a singular mapping is always present due to pixellation so that the quadratic estimator (36) does not fully remove underfitting bias.

The inference of GLAM ellipticity from noisy images with a likelihood that ignores residuals produces underfitting bias; the underfitting bias depends on the correlation of pixel noise and its homogeneity over the image. To deal with pixel noise in images, we have put GLAM in a Bayesian setting in Sect. 3 by employing an approximate model (49) for the likelihood that ignores residuals and thus equals an imperfect model fit with an elliptical profile. We use this to explore the impact of a misspecified likelihood that, as a result, can be subject to underfitting bias. The bias is revealed in the combination of the likelihoods of N independent exposures of the same image (see Sect. 3.2). To lowest order, we find the bias to vanish if $L^T N^{-1} L$ is proportional to the unit matrix. The underfitting bias is unchanged for a rescaled noise covariance: there hence is no noise bias from a misspecified likelihood. Additionally, the magnitude of bias depends on the details of L and can also emerge for $\text{tr}(L^T N^{-1} L) = 1$ if N is not proportional to the unit matrix, which is the case for heterogeneous or correlated noise. Due to the close relation of GLAM ellipticity to other definitions of ellipticity we expect similar dependencies of bias on the noise covariance there, which may impact future calibrations in lensing applications because they apparently should not ignore cases of correlated or heterogeneous noise. Note that several steps in the reduction of lensing data can produce correlated pixel noise.

The underfitting bias from the likelihood can presumably be addressed at the expense of an increased statistical uncertainty, as by means of a follow-up high-resolution survey that gathers statistical information on the residuals R_{pre} . First of all, one obvious way to reduce underfitting bias is to choose a template that provides a good fit to galaxy images to reduce the residuals. The optimal template has the profile $I_{\text{pre}}(r) / S_r(r)$ for a pre-seeing galaxy profile $S_r(r)$; this also optimises the signal-to-noise ratio of the ellipticity (Bernstein & Jarvis 2002). Even so, elliptical galaxies can never be perfectly fit by an elliptical model and are always leaving residuals R_{pre} . To deal with these residuals, we can imagine a more elaborate Bayesian scheme where we add additional nuisance parameters that we marginalize over by using an empirical distribution $P(R_{\text{pre}}|p)$ of R_{pre} given p (prior), i.e.,

$$L(I|p) = \int dR_{\text{pre}} L(I|p; R_{\text{pre}}) P(R_{\text{pre}}|p); \quad (72)$$

where $L(I|p; R_{\text{pre}})$ is the likelihood of a post-seeing image, correctly specified by the post-seeing model $L(I|p) = L(Af + R_{\text{pre}})$ and the noise distribution. The marginalization would then increase the statistical uncertainty in the posterior $P(p|I)$ in Eq. (45). It is conceivable to measure the prior in a dedicated survey by assessing the distribution of residuals in templates

high resolution, high SN images. This approach would be similar to that of BA14 where a (correct) prior on galaxy-specific descriptors is needed. In comparison to BA14, it is noteworthy here that the residuals in the pre-seeing frame are those of sheared images for which we infer no assumptions on the prior. This, however, still has to be tested, and it is not clear how it can be properly implemented and if this marginalization is the source of another bias.

A consistent likelihood of single images is not sufficient for a lensing analysis as new inconsistencies can arise in samples of intrinsically different sources through incorrect priors for intrinsic source parameters, such as sizes or orientations (prior bias). Take for example our TMP galaxies in Set 4 which have a correctly specified likelihood by definition; they are free of underfitting bias. For the two experiments presented in Fig. 3 and Fig. 5, the source sample consists of pre-seeing images with an unknown distribution of intrinsic parameters. The first experiment considers samples with constant ellipticity. The second samples with constant reduced shear. Despite the tabs underfitting, noise bias now emerges in both cases if we apply our uniform prior for the nuisance parameters in the marginal posterior of either ellipticity or reduced shear; see TMP data points at $\gamma = 20$. We argue in Sect. 3.3 that the emerging noise bias is related to the specific choice of priors and can only be avoided by a special family of correct priors related to the maximum of the posterior of noise-free images. According to previous work, correct priors are the distributions of nuisance parameters of the sources in the sample which have to be defined externally (e.g., BA14, Miller et al. 2013, Kitching et al. 2018). A conceivable alternative to external priors, potentially worthwhile investigating in future studies, might be to specify priors down to a family of distributions, so-called hyperparameters (Gelman et al. 2013). A hierarchical approach would then jointly determine the posterior density of θ and the hyperparameters from all I_i alone in a fully Bayesian way. Marginalizing over the hyperparameters accounts for the uncertainty in the prior. Note that the incorrect uniform prior for θ in Fig. 5 has no significant impact on the shear bias (compare open to filled data points). Clearly, the importance of a correct prior depends on the type of nuisance parameter.

That prior bias in Fig. 5 or Table 1 becomes relevant only at low SN can be understood qualitatively. For a sufficiently high SN, the mass of an (identifiable) likelihood $L(I_{\text{post}}|p)$ concentrates around p_{post} such that the prior $P_p(p)$ details have no strong impact on the marginal posterior of one source. We also expect for high SN the likelihood to be well approximated by a normal Gaussian with maximum p_{post} such that the marginalization over $q = (A; t; x_0)$ approximately yields a Gaussian posterior $P(p|I_{\text{post}})$ with maximum near p_{post} which is the true pre-seeing ellipticity for a correctly specified likelihood. Note that the marginalization over q of a multivariate Gaussian with maximum at $p_0 = (q_0; q_0)$ produces another Gaussian density with maximum at p_0 . In addition, should the likelihood be misspecified, p_{post} will be prone to underfitting bias which is clearly visible by the offsets in Fig. 5.

Finally, we point out that while marginal posteriors $P(p|I_j)$, obtained on an image-by-image basis, are sufficient to infer a constant ellipticity (or shear) it is not obvious how to incorporate them into a fully Bayesian analysis of sources with varying shear. For example, in a fully Bayesian inference of ellipticity correlations $\# = h_{i,j}$ for images j at separation $\#$, we would

compute the likelihood $\mathcal{L}(\hat{\mathbf{I}}|\xi_\theta)$ for given values of ξ_θ and a range of θ , which principally requires the repeated shape measurements for the entire sample of images $\hat{\mathbf{I}}$ for every new ξ_θ . This poses a practical, computational problem. Proposed solutions to this problem leave the path of a fully Bayesian analysis and commonly use unbiased point estimators of ϵ_i on a image-by-image basis as input for a statistical analysis (e.g., Miller et al. 2013). This technique has essentially been applied to all lensing analyses so far, but requires a calibration of the ellipticity estimates, especially for the noise bias. The recently proposed hierarchical Bayesian technique by Alsing et al. (2016), applied to data in Alsing et al. (2017), is closest to a fully Bayesian analysis in the above sense but also uses as input unbiased estimators of source ellipticities with a multivariate Gaussian probability density for the statistical errors of the estimates. Yet, it is surely conceivable to extend this technique by directly modelling the likelihood of source images \mathbf{I}_i . To reverse-propagate our $P_\epsilon(\epsilon|\mathbf{I}_i)$ to a posterior of ξ_θ , we might be tempted here to Monte-Carlo a posterior PDF of ξ_θ by independently drawing ellipticities $\epsilon_{im} \sim P_\epsilon(\epsilon|\mathbf{I}_i)$ from the marginal posteriors to produce realisations of a joint vector $\epsilon_m = (\epsilon_{1m}, \dots, \epsilon_{nm})$, which could be used with an estimator of ξ_θ . Unfortunately, this falsely assumes a separable joint posterior $P_\epsilon(\epsilon|\hat{\mathbf{I}}) = P_\epsilon(\epsilon_1|\mathbf{I}_1) \times \dots \times P_\epsilon(\epsilon_n|\mathbf{I}_n)$ with statistically independent ellipticities, potentially biasing (low) constraints on ξ_θ . As alternative we speculate about the following hierarchical Bayesian scheme. We assume a parametric model $P_\epsilon(\epsilon|\hat{\mathbf{I}}, \theta)$ for the joint posterior density of ellipticities that is determined by (i) the measured marginal ellipticities $P_\epsilon(\epsilon|\mathbf{I}_i)$ and (ii) a set of unknown hyperparameters $\theta = (\theta_1, \theta_2, \dots)$ that quantify the correlations between the source ellipticities. A convenient choice in this respect seem to be copula models that are specified exactly that way (e.g., Sect. 2 in Takeuchi 2010). We further express our ignorance about θ_i by a (sufficiently uninformative) hyperprior $P_\theta(\theta)$. We then produce realisations ϵ_m by first randomly drawing $\theta_m \sim P_\theta(\theta)$ followed by $\epsilon_m \sim P_\epsilon(\epsilon|\hat{\mathbf{I}}, \theta_m)$. A sample $\{\epsilon_1, \epsilon_2, \dots\}$ of joint ellipticities is then utilized to propagate uncertainties on source ellipticities and their mutual correlations through the lensing analysis, based on an estimator for ξ_θ given a joint sample ϵ . Note that in the limit $\nu \rightarrow \infty$ the hyperprior for θ becomes irrelevant because all marginal posteriors $P_\epsilon(\epsilon|\mathbf{I}_i)$ will be sharply peaked at the true ellipticity (if correctly specified).

Acknowledgments

The authors are particularly grateful to Gary Bernstein who read and commented on an earlier manuscript of this paper. The comments lead to a substantial improvement. We would also like to thank Tim Schrabback for helpful discussions and comments, and our colleague Reiko Nakajima for organising a weekly seminar on shape measurements in our institute that triggered the research presented in this paper. We also acknowledge the comments by the referee of this paper which helped to clarify and revise certain points, in particular the section on the prior bias. This work has been supported by the Deutsche Forschungsgemeinschaft in the framework of the Collaborative Research Center TR33 ‘The Dark Universe’. Patrick Simon also acknowledges support from the German Federal Ministry for Economic Affairs and Energy (BMWi) provided via DLR under project no. 50QE1103.

References

Aitken, A. 1934, On least squares and linear combination of observations, Vol. 55, 42–48

- Alsing, J., Heavens, A., & Jaffe, A. H. 2017, MNRAS, 466, 3272
 Alsing, J., Heavens, A., Jaffe, A. H., et al. 2016, MNRAS, 455, 4452
 Bartelmann, M. & Schneider, P. 2001, Phys. Rep., 340, 291
 Bernstein, G. M. 2010, MNRAS, 406, 2793
 Bernstein, G. M. & Armstrong, R. 2014, MNRAS, 438, 1880
 Bernstein, G. M., Armstrong, R., Krawiec, C., & March, M. C. 2016, MNRAS, 459, 4467
 Bernstein, G. M. & Jarvis, M. 2002, AJ, 123, 583
 Bertin, E. & Arnouts, S. 1996, A&AS, 117, 393
 Bezdek, J., Hathaway, R., Howard, R., Wilson, C., & Windham, M. 1987, Journal of Optimization Theory and Applications, 54, 471
 Bridle, S., Balan, S. T., Bethge, M., et al. 2010, MNRAS, 405, 2044
 Capaccioli, M. 1989, in World of Galaxies (Le Monde des Galaxies), ed. H. G. Corwin, Jr. & L. Bottinelli, 208–227
 Dawson, W. A., Schneider, M. D., Tyson, J. A., & Jee, M. J. 2014, arXiv:1406.1506
 Fisher, R. A. 1935, J. Roy. Stat. Soc., 98, 39
 Frigo, M. & Johnson, S. G. 2005, Proceedings of the IEEE, 93, 216, special issue on ‘‘Program Generation, Optimization, and Platform Adaptation’’
 Geiger, B. & Schneider, P. 1998, MNRAS, 295, 497
 Gelman, A., Carlin, J., Stern, H., et al. 2013, Bayesian Data Analysis, Third Edition, Chapman & Hall/CRC Texts in Statistical Science (Taylor & Francis)
 Gruen, D., Bernstein, G. M., Jarvis, M., et al. 2015, arXiv:1501.02802
 Hartlap, J., Hilbert, S., Schneider, P., & Hildebrandt, H. 2011, A&A, 528, A51
 Heymans, C., Van Waerbeke, L., Bacon, D., et al. 2006, MNRAS, 368, 1323
 Hirata, C. & Seljak, U. 2003, MNRAS, 343, 459
 Hirata, C. M., Mandelbaum, R., Seljak, U., et al. 2004, MNRAS, 353, 529
 Hoekstra, H., Herbommet, R., Muzzin, A., et al. 2015, MNRAS, 449, 685
 Hoekstra, H. & Jain, B. 2008, Annual Review of Nuclear and Particle Science, 58, 99
 Kacprzak, T., Zuntz, J., Rowe, B., et al. 2012, MNRAS, 427, 2711
 Kaiser, N., Squires, G., & Broadhurst, T. 1995, ApJ, 449, 460
 Kilbinger, M. 2015, Reports on Progress in Physics, 78, 086901
 Kilbinger, M., Wraith, D., Robert, C. P., et al. 2010, MNRAS, 405, 2381
 Kitching, T. D., Balan, S. T., Bridle, S., et al. 2012, MNRAS, 423, 3163
 Kitching, T. D., Miller, L., Heymans, C. E., van Waerbeke, L., & Heavens, A. F. 2008, MNRAS, 390, 149
 Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, arXiv:1110.3193
 Lewis, A. 2009, MNRAS, 398, 471
 MacKay, J. D. 2003, Information Theory, Inference, and Learning Algorithms (Cambridge University Press, Cambridge, UK)
 Mandelbaum, R., Hirata, C. M., Seljak, U., et al. 2005, MNRAS, 361, 1287
 Marquardt, D. 1963, Journal of the Society for Industrial and Applied Mathematics, 11, 431
 Marshall, A. 1956, in In M. Meyer (Ed.), Symposium on Monte Carlo methods (Wiley), 123–140
 Massey, R., Heymans, C., Bergé, J., et al. 2007, MNRAS, 376, 13
 Massey, R., Hoekstra, H., Kitching, T., et al. 2013, MNRAS, 429, 661
 Massey, R., Kitching, T., & Richard, J. 2010, Reports on Progress in Physics, 73, 086901
 Melchior, P., Böhnert, A., Lombardi, M., & Bartelmann, M. 2010, A&A, 510, A75
 Melchior, P., Suchyta, E., Huff, E., et al. 2015, MNRAS, 449, 2219
 Melchior, P. & Viola, M. 2012, MNRAS, 424, 2757
 Melchior, P., Viola, M., Schäfer, B. M., & Bartelmann, M. 2011, MNRAS, 412, 1552
 Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. 1953, J. Chem. Phys., 21, 1087
 Miller, L., Heymans, C., Kitching, T. D., et al. 2013, MNRAS, 429, 2858
 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
 Munshi, D., Valageas, P., van Waerbeke, L., & Heavens, A. 2008, Phys. Rep., 462, 67
 Nakajima, R. & Bernstein, G. 2007, AJ, 133, 1763
 Niemi, S.-M., Cropper, M., Szafraniec, M., & Kitching, T. 2015, Experimental Astronomy
 Plazas, A. A., Bernstein, G. M., & Sheldon, E. S. 2014, Journal of Instrumentation, 9, C4001
 Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, Numerical recipes in C. The art of scientific computing (Cambridge: University Press, |c1992, 2nd ed.)
 Refregier, A., Kacprzak, T., Amara, A., Bridle, S., & Rowe, B. 2012, MNRAS, 425, 1951
 Schneider, P. 2006, in Saas-Fee Advanced Course 33: Gravitational Lensing: Strong, Weak and Micro, ed. G. Meylan, P. Jetzer, P. North, P. Schneider, C. S. Kochanek, & J. Wambsganss, 269–451
 Seitz, C. & Schneider, P. 1997, A&A, 318, 687
 Semboloni, E., Hoekstra, H., Huang, Z., et al. 2013, MNRAS, 432, 2385
 Sheldon, E. S. 2014, MNRAS, 444, L25
 Takeuchi, T. T. 2010, MNRAS, 406, 1830
 Tegmark, M., Taylor, A. N., & Heavens, A. F. 1997, ApJ, 480, 22
 van der Vaart, A. W. 1998, Asymptotic statistics, Cambridge Series in Statistical and Probabilistic Mathematics (Cambridge University Press)
 Viola, M., Kitching, T. D., & Joachimi, B. 2014, MNRAS, 439, 1909
 Viola, M., Melchior, P., & Bartelmann, M. 2011, MNRAS, 410, 2156
 Voigt, L. M. & Bridle, S. L. 2010, MNRAS, 404, 458
 Zuntz, J., Kacprzak, T., Voigt, L., et al. 2013, MNRAS, 434, 1604

Appendix A: Least-square conditions

In this section, we derive general characteristics of GLAM parameters $\mathbf{p} = (A, \mathbf{x}_0, \mathbf{M})$ at a local minimum of the error functional $E(\mathbf{p}|I)$, Eq. (3). For convenience, we introduce the bilinear scalar product

$$\langle g(\mathbf{x}), h(\mathbf{x}) \rangle := \frac{1}{2} \int d^2x g(\mathbf{x})h(\mathbf{x}) \quad (\text{A.1})$$

to write the functional as

$$E(\mathbf{p}|I) = \langle I(\mathbf{x}) - Af(\rho), I(\mathbf{x}) - Af(\rho) \rangle. \quad (\text{A.2})$$

Notice that $\rho := (\mathbf{x} - \mathbf{x}_0)^T \mathbf{M}^{-1} (\mathbf{x} - \mathbf{x}_0)$ is a function of \mathbf{p} . The scalar product has the properties

$$\begin{aligned} \langle g(\mathbf{x}), h(\mathbf{x}) \rangle &= \langle h(\mathbf{x}), g(\mathbf{x}) \rangle; \\ \langle a(\mathbf{p})g(\mathbf{x}) + b(\mathbf{p})f(\mathbf{x}), h(\mathbf{x}) \rangle &= \\ & a(\mathbf{p})\langle g(\mathbf{x}), h(\mathbf{x}) \rangle + b(\mathbf{p})\langle f(\mathbf{x}), h(\mathbf{x}) \rangle; \\ \langle a(\mathbf{p})g(\mathbf{x}), h(\mathbf{x}) \rangle &= \langle g(\mathbf{x}), a(\mathbf{p})h(\mathbf{x}) \rangle; \end{aligned}$$

of which we make heavily use in the following.

Starting from here, we find as necessary condition for the least-square solution of any parameter p_i :

$$\begin{aligned} \frac{\partial E(\mathbf{p}|I)}{\partial p_i} &= 0 \quad (\text{A.3}) \\ &= \left\langle -\frac{\partial[Af(\rho)]}{\partial p_i}, I(\mathbf{x}) - Af(\rho) \right\rangle + \left\langle I(\mathbf{x}) - Af(\rho), -\frac{\partial[Af(\rho)]}{\partial p_i} \right\rangle \\ &= -2 \left\langle \frac{\partial[Af(\rho)]}{\partial p_i}, I(\mathbf{x}) - Af(\rho) \right\rangle \\ &= -2 \left\langle \frac{\partial[Af(\rho)]}{\partial p_i}, I(\mathbf{x}) \right\rangle + 2A \left\langle \frac{\partial[Af(\rho)]}{\partial p_i}, f(\rho) \right\rangle, \end{aligned}$$

or simply

$$\left\langle \frac{\partial[Af(\rho)]}{\partial p_i}, A^{-1}I(\mathbf{x}) \right\rangle = \left\langle \frac{\partial[Af(\rho)]}{\partial p_i}, f(\rho) \right\rangle \quad (\text{A.4})$$

that is fulfilled for all p_i .

Appendix A.1: Template normalisation

From (A.4) follows for the normalisation $p_i = A$ at the minimum of $E(\mathbf{p}|I)$ that

$$A = \frac{\langle f(\rho), I(\mathbf{x}) \rangle}{\langle f(\rho), f(\rho) \rangle}. \quad (\text{A.5})$$

Appendix A.2: Image centroid position

For the least-square centroid position $x_{0,i}$ with $i = 1, 2$ at the minimum of $E(\mathbf{p}|I)$, we find $\partial[Af(\rho)]/\partial x_{0,i} = Af'(\rho) \partial\rho/\partial x_{0,i}$ using $f'(\rho) := df(\rho)/d\rho$. This can be expanded further by considering

$$\begin{aligned} \frac{\partial\rho}{\partial x_{0,i}} &= \frac{\partial}{\partial x_{0,i}} \sum_{k,l} (x_k - x_{0,k}) [M^{-1}]_{kl} (x_l - x_{0,l}) \quad (\text{A.6}) \\ &= - \sum_{k,l} \left(\delta_{ki}^K [M^{-1}]_{il} (x_l - x_{0,l}) + \delta_{li}^K [M^{-1}]_{ki} (x_k - x_{0,k}) \right) \\ &= -2 \sum_k [M^{-1}]_{ki} (x_k - x_{0,k}) = -2 [M^{-1}(\mathbf{x} - \mathbf{x}_0)]_i. \end{aligned}$$

Here we have exploited the symmetry $[M^{-1}]_{ij} = [M^{-1}]_{ji}$ and introduced the Kronecker symbol δ_{ij}^K . Therefore we find for the i th component of the centroid position

$$\langle Af'(\rho) [M^{-1}(\mathbf{x} - \mathbf{x}_0)]_i, f(\rho) - A^{-1}I(\mathbf{x}) \rangle = 0 \quad (\text{A.7})$$

or, equally, for both centroid components combined

$$AM^{-1} \langle f'(\rho)(\mathbf{x} - \mathbf{x}_0), f(\rho) \rangle = M^{-1} \langle f'(\rho)(\mathbf{x} - \mathbf{x}_0), I(\mathbf{x}) \rangle. \quad (\text{A.8})$$

The scalar product on the left-hand-side has to vanish due to symmetry reasons. Therefore using the right-hand-side, we find at the minimum

$$\mathbf{x}_0 = \frac{\langle f'(\rho)\mathbf{x}, I(\mathbf{x}) \rangle}{\langle f'(\rho), I(\mathbf{x}) \rangle}. \quad (\text{A.9})$$

This shows that the centroid \mathbf{x}_0 at the least-square point is the first-order moment \mathbf{x} of the image $I(\mathbf{x})$ weighted with the kernel $f'(\rho)$. The weight is thus the derivative $f'(\rho)$ not $f(\rho)$.

Appendix A.3: Image second-order moments

For the components of \mathbf{M} , we consider the least-square values of the inverse $W_{ij} = [M^{-1}]_{ij}$, hence values W_{ij} for which $\partial[Af(\rho)]/\partial W_{ij} = Af'(\rho) \partial\rho/\partial W_{ij}$ vanishes. The function ρ can be expressed as

$$\rho = \text{tr}((\mathbf{x} - \mathbf{x}_0)^T \mathbf{W} (\mathbf{x} - \mathbf{x}_0)) = \text{tr}(\mathbf{W}\mathbf{X}), \quad (\text{A.10})$$

using $\mathbf{X} := (\mathbf{x} - \mathbf{x}_0)(\mathbf{x} - \mathbf{x}_0)^T = \mathbf{X}^T$ and the relation $\text{tr}(\mathbf{ABC}) = \text{tr}(\mathbf{BCA})$ for traces of matrices. From this follows that $\partial\rho/\partial W_{ij} = [\mathbf{X}]_{ij} = X_{ij}$, or

$$\begin{aligned} A \langle f'(\rho) X_{ij}, A^{-1}I(\mathbf{x}) \rangle &= A \langle f'(\rho) X_{ij}, f(\rho) \rangle \quad (\text{A.11}) \\ \iff \langle f'(\rho) X_{ij}, I(\mathbf{x}) \rangle &= \langle f'(\rho) X_{ij}, Af(\rho) \rangle \end{aligned}$$

by employing Eq. (A.4). This set of (four) equations can compactly be written as

$$\langle f'(\rho)\mathbf{X}, Af(\rho) \rangle = \langle f'(\rho)\mathbf{X}, I(\mathbf{x}) \rangle \quad (\text{A.12})$$

or utilising Eq. (A.5):

$$\frac{\langle f'(\rho)\mathbf{X}, f(\rho) \rangle}{\langle f(\rho), f(\rho) \rangle} = \frac{\langle f'(\rho)\mathbf{X}, I(\mathbf{x}) \rangle}{\langle f(\rho), I(\mathbf{x}) \rangle}. \quad (\text{A.13})$$

This last relation shows that the best-fitting template $f(\rho)$ has, up to a constant scalar α , the same central second-order moments \mathbf{M}_I as the image,

$$\mathbf{M} = \frac{\langle f'(\rho)\mathbf{X}, f(\rho) \rangle}{\langle f'(\rho), f(\rho) \rangle} \quad (\text{A.14})$$

$$= \frac{\langle f(\rho), f(\rho) \rangle \langle f'(\rho), I(\mathbf{x}) \rangle \langle f'(\rho)\mathbf{X}, I(\mathbf{x}) \rangle}{\langle f'(\rho), f(\rho) \rangle \langle f(\rho), I(\mathbf{x}) \rangle \langle f'(\rho), I(\mathbf{x}) \rangle} \quad (\text{A.15})$$

$$=: \alpha \frac{\langle f'(\rho)\mathbf{X}, I(\mathbf{x}) \rangle}{\langle f'(\rho), I(\mathbf{x}) \rangle} = \alpha \mathbf{M}_I. \quad (\text{A.16})$$

In particular, the best-fitting template $Af(\rho)$ has, for a given profile $f(s)$, the same ellipticity ϵ as the image $I(\mathbf{x})$. Moreover, we derive from (A.13) the relation

$$\mathbf{M} = \frac{\langle f(\rho), f(\rho) \rangle \langle f'(\rho)\mathbf{X}, I(\mathbf{x}) \rangle}{\langle f'(\rho), f(\rho) \rangle \langle f(\rho), I(\mathbf{x}) \rangle} =: \beta \frac{\langle f'(\rho)\mathbf{X}, I(\mathbf{x}) \rangle}{\langle f(\rho), I(\mathbf{x}) \rangle}. \quad (\text{A.17})$$

The scalar prefactor equals $\beta = -2$ for $f(\rho) = e^{-\rho/2}$. Note that we have on the right-hand-side $f(\rho)$ in the denominator rather than $f'(\rho)$.

We evolve the expression of β a bit further by changing the variables $\Delta \mathbf{x} = \mathbf{M}^{1/2} \mathbf{y}$ and hence $d^2x = \sqrt{\det \mathbf{M}} d^2y$ for $\rho = \Delta \mathbf{x}^T \mathbf{M}^{-1} \Delta \mathbf{x}$ in the integrals of the denominator and nominator of β ,

$$\beta = \frac{\int d^2y f^2(\mathbf{y}^T \mathbf{y})}{\int d^2y f'(\mathbf{y}^T \mathbf{y}) f(\mathbf{y}^T \mathbf{y})} = \frac{\int_0^\infty dy y f^2(y^2)}{\int_0^\infty dy y f'(y^2) f(y^2)} \quad (\text{A.18})$$

or after another change of variables $y = \sqrt{s}$ and $dy = ds/(2\sqrt{s})$,

$$\beta = \frac{\langle f(\rho), f(\rho) \rangle}{\langle f'(\rho), f(\rho) \rangle} = \frac{\int_0^\infty \frac{ds}{2} f^2(s)}{\int_0^\infty ds \frac{d}{4ds} f^2(s)} = -2 \frac{\int_0^\infty ds f^2(s)}{f^2(0)}, \quad (\text{A.19})$$

where the last step assumes that the profile $f(s)$ vanishes for $s \rightarrow \infty$.

Appendix B: Importance sampling of posterior distribution

Algorithm 1: Sampling of the ellipticity posterior $P_\epsilon(\epsilon|\mathbf{I})$ for a given post-seeing galaxy image \mathbf{I} . Note that GLAM parameters are $\mathbf{p} = (A, \mathbf{x}_0, t, \epsilon)$; A : amplitude; \mathbf{x}_0 : centroid; t : size; ϵ : ellipticity; assumption for residuals: $\mathbf{R}_{\text{pre}} \equiv 0$.

Input :

\mathbf{I} : pixellated, noisy galaxy image;
 \mathbf{N} : noise covariance of image;
 $A\mathbf{L}\mathbf{f}_\rho$: post-seeing GLAM templates with parameters \mathbf{p} ;
 $P_p(\mathbf{p})$: prior of GLAM parameters;

Output:

$1 \leq i \leq N_{\text{real}}$ realisations (ϵ_i, w_i) with weights w_i ; N_{eff} : effective number of sampling points;

```

1 begin
2    $\mathbf{r}(\mathbf{p}) := \mathbf{I} - A\mathbf{L}\mathbf{f}_\rho$ ;
3    $-2 \ln \mathcal{L}(\mathbf{I}|\mathbf{p}) := \mathbf{r}^T(\mathbf{p})\mathbf{N}^{-1}\mathbf{r}(\mathbf{p})$ ;
4    $\mathbf{p}_{\text{ml}} \leftarrow -2 \ln \mathcal{L}(\mathbf{I}|\mathbf{p}_{\text{ml}}) = \min.$  (ML point);
5    $F_{ij} \leftarrow -\left\langle \frac{\partial^2 \ln \mathcal{L}(\mathbf{I}|\mathbf{p})}{\partial p_i \partial p_j} \right\rangle_{\mathbf{p}_{\text{ml}}}$  (Fisher matrix; Eq. B.2);
6   for  $i \leftarrow 1$  to  $N_{\text{real}}$  do
7     repeat
8        $\forall 1 \leq j \leq 6$   $x_{ij} \sim N(0, 1)$ ;
9        $\mathbf{p}_i \leftarrow \mathbf{p}_{\text{ml}} + \sqrt{\mathbf{F}}^{-1} \mathbf{x}_i$ ;
10       $\epsilon_i \leftarrow \mathbf{p}_i = (A_i, \mathbf{x}_{0,i}, t_i, \epsilon_i)$ ;
11     until  $|\epsilon_i| < 1$  and  $P_p(\mathbf{p}_i) > 0$ ;
12      $\Delta \mathbf{p}_i \leftarrow \mathbf{p}_i - \mathbf{p}_{\text{ml}}$ ;
13      $Q_i \leftarrow \exp(-\Delta \mathbf{p}_i^T \mathbf{F} \Delta \mathbf{p}_i / 2)$ ;
14      $\pi_i \leftarrow \mathcal{L}(\mathbf{I}|\mathbf{p}_i) P_p(\mathbf{p}_i)$ ;
15      $w'_i \leftarrow \frac{\pi_i}{Q_i}$ ;
16      $\forall i$   $w_i \leftarrow \frac{w'_i}{\sum_{i=1}^{N_{\text{real}}} w'_i}$ ;
17      $N_{\text{eff}} \leftarrow \frac{1}{\sum_{i=1}^{N_{\text{real}}} w_i^2}$ ;

```

An overview of our algorithm for sampling the posterior $P_p(\mathbf{p}|\mathbf{I})$ is given by Algorithm 1, more details follow below.

For the importance sampling of the posterior GLAM variables \mathbf{p} , we initially construct an analytical approximation $Q(\mathbf{p})$ of the (full) posterior $P_p(\mathbf{p}|\mathbf{I}) := \mathcal{L}(\mathbf{I}|\mathbf{p}) P_p(\mathbf{p})$ in Eq. (45), the so-called importance function. Our importance function is a local Gaussian approximation of $\mathcal{L}(\mathbf{I}|\mathbf{p})$ that is truncated for $|\epsilon| \geq 1$; $Q(\mathbf{p})$ is hence defined by its mean and its covariance. As previously discussed, we ignore residuals and set $\mathbf{R}_{\text{pre}} \equiv 0$ in Eq. (48) for the following.

For the mean \mathbf{p}_{ml} of the Gaussian $Q(\mathbf{p})$, we numerically determine the maximum-likelihood (ML) point of $\mathcal{L}(\mathbf{I}|\mathbf{p})$ by minimising (48) with respect to \mathbf{p} . This is the most time consuming step in our algorithm, which is done in the following three steps.

1. We quickly guess a starting point \mathbf{p}_0 of the centroid position \mathbf{x}_0 and the galaxy size t by measuring the unweighted first and second-order moments of the post-seeing postage stamp. The initial ellipticity ϵ is set to zero. Herein we ignore the PSF. Alternatively, one could use estimates from object detection software such as SExtractor if available (Bertin & Arnouts 1996).
2. For an improved initial guess of the size t and the centroid position \mathbf{x}_0 , starting from \mathbf{p}_0 we perform a series of one-dimensional (1D) searches wherein we vary one parameter p_i while all other parameters p_j with $i \neq j$ are fixed (coordinate descent; Bezdek et al. 1987). For each 1D search, we compute $-2 \ln \mathcal{L}(\mathbf{I}|\mathbf{p})$ at five points over the range of values of p_i that are allowed according to the prior $P_p(\mathbf{p})$. We spline interpolate between the points to estimate the global minimum p_i^{min} . Then we update p_i to the value p_i^{min} and move on to the 1D search of the next parameter. After every update we replace A by its most likely value

$$A_{\text{ml}} = \frac{\mathbf{I}^T \mathbf{N}^{-1} \mathbf{L} \mathbf{f}_\rho}{[\mathbf{L} \mathbf{f}_\rho]^T \mathbf{N}^{-1} \mathbf{L} \mathbf{f}_\rho} \quad (\text{B.1})$$

(Eq. A.5) given the current values of $(\mathbf{x}_0, \epsilon, t)$. We carry out the searches first for the size t and then for the two components of the centroid position \mathbf{x}_0 .

3. Starting from \mathbf{p} in step 2, we apply the Levenberg-Marquardt algorithm (LMA) to iteratively approach \mathbf{p}_{ml} (Marquardt 1963). This algorithm devises a linear approximation of $\mathbf{L} \mathbf{f}_\rho$ at \mathbf{p} based on the partial derivatives of $\mathbf{L} \mathbf{f}_\rho$. Our implementation of the LMA searches only in the directions of $(\mathbf{x}_0, \epsilon, t)$ and again updates A to A_{ml} after each iteration as in step 2. We find a slow convergence for large ellipticities, i.e., $|\epsilon| \gtrsim 0.3$. This is improved, however, by adding an extra 1D search for the size t , ϵ_1 and ϵ_2 after each LMA iteration with an interval width of $\pm 0.1t$ (10 spline points) and ± 0.1 (5 points) around the latest values of t and ϵ_i , respectively. For $|\epsilon_n| \geq 0.6$, we use 10 instead of 5 spline points which increases the accuracy for very elliptical objects. We repeat the LMA until iteration $n > 5$ for which $|\epsilon_{n-1} - \epsilon_n| \leq 10^{-4}$; we change this threshold to 2×10^{-5} once $|\epsilon_n| \geq 0.6$. The latter is needed because the change of ϵ_n at each iteration becomes small for high ϵ .

For the covariance of $Q(\mathbf{p})$, we then numerically compute the elements F_{ij} of the (expected) Fisher matrix \mathbf{F} around the ML point \mathbf{p}_{ml} ,

$$F_{ij} = \left(\frac{\partial [A\mathbf{L}\mathbf{f}_\rho]}{\partial p_i} \right)^T \times \mathbf{N}^{-1} \times \left(\frac{\partial [A\mathbf{L}\mathbf{f}_\rho]}{\partial p_j} \right), \quad (\text{B.2})$$

by considering partial derivatives of the best-fitting GLAM templates (e.g., Tegmark et al. 1997). This gives us

$$Q(\mathbf{p}) \propto \exp\left(-\frac{1}{2} [\mathbf{p} - \mathbf{p}_{\text{ml}}]^T \mathbf{F} [\mathbf{p} - \mathbf{p}_{\text{ml}}]\right) H(P_p(\mathbf{p})), \quad (\text{B.3})$$

where $H(x)$ denotes the Heaviside step function. To quickly construct the Fisher matrix, we employ the partial derivatives of $\mathbf{L}\mathbf{f}_\rho$ in the last iteration of the previous LMA; the derivative with respect to A is simply $\mathbf{L}\mathbf{f}_\rho$.

For the main loop of the importance sampling, we randomly draw N_{real} points $\mathbf{p}_i \sim Q(\mathbf{p})$. Each realisation \mathbf{p}_i receives an importance weight $w_i = P_p(\mathbf{p}_i|\mathbf{I}) Q(\mathbf{p}_i)^{-1}$ for which we have to evaluate the posterior $P_p(\mathbf{p}_i|\mathbf{I})$; neither $P_p(\mathbf{p}|\mathbf{I})$ nor $Q(\mathbf{p})$ need to be normalised here. In the end, we normalise all weights to unity for convenience, i.e., $\sum_i w_i \equiv 1$. Moreover, for later usage, we keep only the ellipticities ϵ_i of the parameter sets \mathbf{p}_i . This Monte Carlo marginalizes over the uncertainties of all other parameters.

As another technical detail, the above algorithm has to repeatedly produce pixellated and PSF convolved versions of GLAM templates $\mathbf{L}\mathbf{f}_\rho$ for a given set of parameters \mathbf{p} . To implement a fast computation in this regard, we construct the template images on a high resolution (high-res) 256×256 grid in Fourier space. To reduce boundary effects, we make the angular extent of the high-res grid about 20 per cent larger than the extent of the post-seeing galaxy image. Let $\tilde{c}_{\text{psf}}(\boldsymbol{\ell})$ be the Fourier transform of the PSF and

$$\tilde{f}(\boldsymbol{\ell}) = 2\pi \int_0^\infty ds f(s) J_0(s\ell) \quad (\text{B.4})$$

the transform of the radial template profile $f(\rho)$. Then the Fourier coefficients of the PSF-convolved template are

$$\tilde{f}_{\text{psf}}(\boldsymbol{\ell}) = (\det V) e^{i\boldsymbol{\ell} \cdot \mathbf{x}_0} \tilde{c}_{\text{psf}}(\boldsymbol{\ell}) \tilde{f}(\boldsymbol{\ell}^T V^2 \boldsymbol{\ell}). \quad (\text{B.5})$$

After setting up the grid of values in Fourier space, we back-transform to coordinate space by means of the FFTW package (Frigo & Johnson 2005). Note that the Fourier coefficient at $\boldsymbol{\ell} = 0$ has to be set to $\tilde{f}_{\text{psf}}(0) = \det V \tilde{c}_{\text{psf}}(0) \tilde{f}(0)$ to avoid a constant offset of the template. After the FFT, we apply pixellation to produce the vector $\mathbf{L}\mathbf{f}_\rho$ of post-seeing pixel values by averaging pixels of the high-res grid within the extent of the post-seeing image. We choose the angular extent of the high-res grid such that an integer number of high-res pixels is enclosed by one post-seeing pixel. For example, for our 20×20 post-seeing images with pixel size $r_{\text{post}} = 0.103125 \approx 0.1$ arcsec, we use a high-res grid with extent of 2.4 arcsec along each axis, or a high-res pixel size of $r_{\text{hres}} = 2.4/256 = 0.009375$ arcsec (square). Hence exactly $(r_{\text{post}}/r_{\text{hres}})^2 = 121$ high-res pixel correspond to one post-seeing pixel.

Furthermore, to quickly compute the partial derivatives of $\partial(\mathbf{L}\mathbf{f}_\rho)/\partial p_i$, needed for the LMA and the Fisher matrix, we carry out the foregoing procedure, but we insert the derivatives $\partial \tilde{f}_{\text{psf}}(\boldsymbol{\ell})/\partial p_i$ as Fourier coefficients instead. These can be computed analytically from Eq. (B.5).