

Reconstructing spatially varying multiplicative bias for Stage IV weak lensing galaxy surveys with a quadratic estimator

Konstantinos Tanidis,^{1,2★} David Alonso,¹ Lance Miller¹ and Joachim Harnois-Déraps³ 

¹Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK

²Center for Astrophysics and Cosmology, University of Nova Gorica, 1280 Nova Gorica, Slovenia

³School of Mathematics, Statistics and Physics, Newcastle University, Newcastle-upon-Tyne, NE1 7RU, UK

Accepted 2026 March 17. Received 2026 March 6; in original form 2026 January 5

ABSTRACT

We present a quadratic estimator that detects and reconstructs spatially varying multiplicative (m -) bias in weak lensing shear measurements, by exploiting the EB mode coupling that it generates. The method combines E and B modes with inverse-variance weights, to yield an unbiased reconstruction of $m(\theta)$ to first order. We study the ability of future Stage IV surveys to obtain an unbiased reconstruction of the m -bias in differing scenarios, considering differing bias morphologies, and characteristic scales, as well as differing metrics to quantify the signal-to-noise ratio of the reconstructed map. We consider an m pattern repeating on $\sim 1^\circ \times 1^\circ$ sky patches, as might be the case for an m field caused by focal-plane systematics. With a *Euclid*-like redshift distribution, we find that ~ 5 per cent root mean square (rms) variations in m -bias may be detected at the 20σ level, after stacking between ~ 400 and ~ 1000 patches (rising to between ~ 2800 and ~ 7600 for 1 per cent rms variations, data volumes that are becoming available with upcoming surveys), depending on the morphology of the m pattern. We show that these results are robust against the cosmological model assumed in the reconstruction, as well as the presence of intrinsic alignments or baryonic effects, and that the method shows no spurious response to additive (c -) bias. These results demonstrate that percent-level, spatially varying m -bias can be detected at high significance, enabling diagnosis and mitigation in the Stage IV weak lensing era.

Key words: gravitational lensing: weak – large-scale structure of Universe – cosmology: observations.

1 INTRODUCTION

Stage IV weak lensing surveys have the potential to make significant improvements in our understanding of the cosmological model (A. Albrecht et al. 2006), a prospect which is now being realized with large-area, high-accuracy surveys such as *Euclid* (Euclid Collaboration 2025) and the Vera Rubin Observatory’s Legacy Survey of Space and Time (LSST; The LSST Dark Energy Science Collaboration 2018). However, it is widely recognized that the achievement of accurate cosmological inference from weak lensing surveys requires understanding and mitigation of a wide range of systematic effects (R. Mandelbaum 2018).

In the weak lensing regime, where the shear due to gravitational lensing is small ($\lesssim 0.01$), it is common to express the effects of systematic biases to leading order as linear multiplicative (m -) and additive (c -) biases (C. Heymans et al. 2006; D. Huterer et al. 2006),

$$\gamma^{\text{biased}} \simeq (1 + m)\gamma^{\text{cosm}} + c, \quad (1)$$

for biased and true cosmic shear, γ^{biased} and γ^{cosm} , respectively.¹ There may be multiple effects contributing to these biases, including both astrophysical and measurement systematics. Among the set of measurement systematics, the errors in the modelling and correction for the imaging point spread function (PSF) may cause both m - and c -biases (S. Paulin-Henriksson et al. 2008; R. Massey et al. 2013), and these biases may be further decomposed into correlations with parameters of the measurement, such as ‘PSF leakage’ (M. Jarvis et al. 2016).

Among the additive biases, methods have been developed to detect these in data. One set of methods is based around cross-correlation with expected sources of error, such as the ‘Rho’ statistics of B. Rowe (2010) and M. Jarvis et al. (2016), and ‘Tau’ statistics of M. Gatti et al. (2021), with cross-correlation between measured galaxy shapes and both PSF ellipticity and size and PSF modelling error. Alternatively, additive bias that has some repeating pattern, for example due to time-invariant errors on the scale of the field of view (FOV) due to PSF modelling error, may be revealed in the data by stacking of galaxy ellipticity measures

¹We should note that here we neglect the distinction between reduced shear g and shear γ , as it does not make a difference for the formalism we develop in this work. Therefore, from now on, we treat γ^{biased} as the observed shear field.

* E-mail: konstantinos.tanidis@physics.ox.ac.uk

from a wide sky area in FOV coordinates (e.g. E. Uiter & P. Schneider 2016; H. Hildebrandt et al. 2020). Finally, template deprojection techniques may be employed to mitigate the impact of additive bias from sources with a known spatial pattern (D. Alonso, J. Sanchez & A. Slosar 2019; T. Cornish et al. 2026).

PSF modelling errors may also generate multiplicative bias: specifically, it is expected that errors in the knowledge of PSF size lead to joint m - and c -bias (S. Paulin-Henriksson et al. 2008; R. Massey et al. 2013). The statistical measurement of galaxy shear generally also introduces m -bias (P. Melchior & M. Viola 2012; A. Refregier et al. 2012), which must be corrected either by realistic simulations (e.g. S.-S. Li et al. 2023) or by self-calibration (e.g. ‘metacalibration’; E. Huff & R. Mandelbaum 2017; E. S. Sheldon & E. M. Huff 2017). Further sources of measurement m -bias include the effects of galaxy neighbours, blends and redshift biases (S. Samuroff et al. 2018; N. MacCrann et al. 2022). However, the direct detection of m -bias in weak lensing data is problematic: the shear-ratio test of B. Jain & A. Taylor (2003) and A. N. Taylor et al. (2007) may be used to detect redshift-dependent m -bias (P. Schneider 2016), but only with limited sensitivity. For this reason, obtaining accurate knowledge of the total m -bias in Stage IV weak lensing surveys is considered to be one of the most challenging aspects of the analysis (R. Mandelbaum 2018).

When calibrating m -bias in Stage III surveys, from simulations, it has generally been assumed that the bias may be dependent on galaxy properties, such as galaxy size, and hence may also be redshift-dependent (S.-S. Li et al. 2023). However, it has been assumed that the bias is otherwise invariant with sky position. It had been thought that the effects of spatially varying m -bias would be negligible, with accurate knowledge only required of the mean m -bias (T. D. Kitching et al. 2019). More recently, however, it has been recognized that spatially dependent m -bias may have a significant convolutional effect on measured shear power spectra: for Stage IV surveys, percent-level root mean square (rms) variations in m -bias on angular scales of degrees or larger may introduce a significant error in Stage IV cosmology analyses (C. Cragg et al. 2023).

A spatially varying m -bias effectively leads to a spatial modulation of the true shear field, inducing additional statistical couplings between different independent Fourier modes and shear components. This is reminiscent of the impact of gravitational lensing on the cosmic microwave background (CMB) temperature and polarization anisotropies, in the form of a spatial modulation proportional to the gradient of these anisotropies. The statistical coupling induced by this modulation can then be used to reconstruct the field causing it (the lensing potential or the reionization in the case of CMB lensing, or the m -bias in the case of cosmic shear), by considering correlations between unequal Fourier modes, through the so-called ‘quadratic estimator’ approach (A. Lewis & A. Challinor 2006; C. Dvorkin & K. M. Smith 2009). Other applications of the quadratic estimator include the removal of CMB systematics (J. Williams et al. 2021). Here we present a quadratic estimator for the reconstruction and detection of a spatially dependent m -bias in cosmic shear data.

This paper is structured as follows. Section 2 introduces the method. We discuss how the method may be applied either to cross-correlation with templates of potential systematic effects, or in a direct-detection method where we stack on the scale of the FOV – in either case somewhat analogously to what may already be achieved for additive bias detection and measurement. Section 3 describes how the method is demonstrated in this paper,

using some simplified models, with results and conclusions in Sections 4 and 5.

2 METHOD

2.1 Introduction to the method

The m -bias detection method acts on the harmonic space two-point functions, E and B , which we start by defining here. For the demonstration in this paper, we assume the flat-sky approximation, so that angular coordinates θ can be treated as two-dimensional vectors on a plane.² The spin-2 cosmic shear field is described by two components $\gamma_1(\theta)$ and $\gamma_2(\theta)$, which can be combined into a complex field $\gamma(\theta) = \gamma_1(\theta) + i\gamma_2(\theta)$. This field can be decomposed into E -mode and B -mode components via the Fourier transform of γ :

$$\gamma(\theta) = \int \frac{d^2\ell}{(2\pi)^2} \gamma(\ell) e^{i\ell\cdot\theta}, \quad (2)$$

where $\gamma(\ell)$ is the shear in Fourier-space. In terms of the Fourier-space γ_1 and γ_2 , one can obtain the E and B modes by a rotation in Fourier space (e.g. W. Hu 2000; P. Schneider 2006):

$$E(\ell) = \gamma_1(\ell) \cos(2\varphi_\ell) + \gamma_2(\ell) \sin(2\varphi_\ell), \quad (3)$$

$$B(\ell) = -\gamma_1(\ell) \sin(2\varphi_\ell) + \gamma_2(\ell) \cos(2\varphi_\ell), \quad (4)$$

with φ_ℓ the polar angle of the wavevector ℓ . By construction, $E(\ell)$ and $B(\ell)$ are real-valued fields describing the parity-even and parity-odd components of the shear field, respectively. At linear order, and in the absence of systematics, the cosmological shear signal γ^{cosm} , due to gravitational lensing by large-scale structure, is expected to produce entirely E -mode power. We therefore consider the cosmological B -mode to be zero, $B^{\text{cosm}}(\ell) = 0$, and that all cosmological shear information resides in $E^{\text{cosm}}(\ell)$.³ Assuming statistical isotropy, the E -mode angular power spectrum C_ℓ^{EE} is defined by

$$\langle E^{\text{cosm}}(\ell) E^{\text{cosm}*}(\ell') \rangle = (2\pi)^2 \delta^D(\ell - \ell') C_\ell^{EE}, \quad (5)$$

where δ^D is the two-dimensional Dirac delta function. Thus, any measured B -modes in the observed shear field γ^{biased} could therefore be an indicator of systematic effects or noise. In the current work, we focus on building an estimator which detects the mode coupling between E and B , as well as between Fourier modes with different wave vectors, as induced by a spatially varying multiplicative bias field contaminating the cosmic shear signal.

2.2 Spatially varying shear bias and mode coupling

We model a real (not imaginary) spatially varying multiplicative shear bias $m(\theta)$ as a dimensionless scalar field that modulates the true cosmological shear. Equation (1) becomes

$$\gamma^{\text{biased}}(\theta) \simeq [1 + m(\theta)] \gamma^{\text{cosm}}(\theta) + c(\theta), \quad (6)$$

where $c(\theta)$ is also a spatially varying c -bias contamination. In practice, $m \ll 1$ and $c \ll 1$ for weak systematics. In Fourier space,

²For analysis on large angular scales, a curved-sky methodology could be adopted (e.g. C. Dvorkin & K. M. Smith 2009).

³Although this is a good approximation in linear order, there are also non-linear lensing effects producing small contributions to B -modes (P. Schneider & M. Kilbinger 2006; S. Saga, D. Yamauchi & K. Ichiki 2015). It is also important to note that intrinsic alignments can also produce B -modes (J. A. Blazek et al. 2019).

the product in equation (6) becomes a convolution. Let $\gamma^{\text{biased}}(\boldsymbol{\ell})$ and $\gamma^{\text{cosm}}(\boldsymbol{\ell})$ be the Fourier transforms of $\gamma^{\text{biased}}(\boldsymbol{\theta})$ and $\gamma^{\text{cosm}}(\boldsymbol{\theta})$, respectively, and $m(\mathbf{L})$ (with $\hat{m}^*(\mathbf{L}) = \hat{m}(-\mathbf{L})$ since m is assumed to be real in real space) and $c(\boldsymbol{\ell})$ the Fourier transforms of the m -bias and the c -bias fields. Then, we have (for a pedagogical derivation see also A. N. Taylor & T. D. Kitching 2018):

$$\gamma^{\text{biased}}(\boldsymbol{\ell}) \simeq \gamma^{\text{cosm}}(\boldsymbol{\ell}) + \int \frac{d^2\mathbf{L}}{(2\pi)^2} m(\mathbf{L}) \gamma^{\text{cosm}}(\boldsymbol{\ell} - \mathbf{L}) + c(\boldsymbol{\ell}). \quad (7)$$

The first term is the cosmic shear at wavevector $\boldsymbol{\ell}$, the second term is the contribution from the coupling between the m -bias and the cosmological shear field and the third term is the c -bias. The coupling term encodes how $m(\boldsymbol{\theta})$ redistributes power in the shear field: it mixes Fourier modes separated by wavevector \mathbf{L} . In particular, the m -bias field produces off-diagonal correlations in Fourier space. Decomposing equation (7) into E and B components using equations (3) and (4), one finds that a non-zero m -bias field mixes E and B modes. Intuitively, a spatially varying multiplicative factor introduces spurious ‘curl’ into the shear pattern, sourcing B -modes from an originally pure E -mode field.

More concretely, inserting equation (7) into the definitions of E^{biased} and B^{biased} , and assuming here $B^{\text{cosm}} \neq 0$ for generality, we can derive the first-order effect of m -bias on the observed E and B fields. The observed E -mode receives a direct contribution from the cosmological E plus a convolution with m -bias and also a c -bias term:

$$\begin{aligned} E^{\text{biased}}(\boldsymbol{\ell}) &= E^{\text{cosm}}(\boldsymbol{\ell}) \\ &+ \int \frac{d^2\mathbf{L}}{(2\pi)^2} m(\mathbf{L}) \{ E^{\text{cosm}}(\boldsymbol{\ell} - \mathbf{L}) \cos[2(\varphi_{\boldsymbol{\ell}-\mathbf{L}} - \varphi_{\boldsymbol{\ell}})] \\ &- B^{\text{cosm}}(\boldsymbol{\ell} - \mathbf{L}) \sin[2(\varphi_{\boldsymbol{\ell}-\mathbf{L}} - \varphi_{\boldsymbol{\ell}})] \} + E_c(\boldsymbol{\ell}), \end{aligned} \quad (8)$$

and the observed B -mode is similarly given by

$$\begin{aligned} B^{\text{biased}}(\boldsymbol{\ell}) &= B^{\text{cosm}}(\boldsymbol{\ell}) \\ &+ \int \frac{d^2\mathbf{L}}{(2\pi)^2} m(\mathbf{L}) \{ E^{\text{cosm}}(\boldsymbol{\ell} - \mathbf{L}) \sin[2(\varphi_{\boldsymbol{\ell}-\mathbf{L}} - \varphi_{\boldsymbol{\ell}})] \\ &+ B^{\text{cosm}}(\boldsymbol{\ell} - \mathbf{L}) \cos[2(\varphi_{\boldsymbol{\ell}-\mathbf{L}} - \varphi_{\boldsymbol{\ell}})] \} + B_c(\boldsymbol{\ell}). \end{aligned} \quad (9)$$

Here, the $\cos(2\Delta\varphi)$ and $\sin(2\Delta\varphi)$ factors arise from the spin-2 nature of the shear field. An important consequence of equations (8) and (9) is that E and B modes become statistically correlated in the presence of m -bias, even if the c -bias is zero. In the absence of m -bias and c -bias, one would have $\langle E^{\text{biased}}(\boldsymbol{\ell}) B^{\text{biased}*}(\boldsymbol{\ell}') \rangle = 0$ (since B^{biased} would be pure noise or very small non-linear weak lensing cosmological signals, uncorrelated with E^{biased}). With $m \neq 0$, the coupling induces a characteristic off-diagonal correlation between E and B . To first order in m , using equation (7) or directly from equations (8) and (9), one finds for $\boldsymbol{\ell} - \boldsymbol{\ell}' = \mathbf{L}$ that

$$\begin{aligned} \langle E^{\text{biased}}(\boldsymbol{\ell}) B^{\text{biased}*}(\boldsymbol{\ell} - \mathbf{L}) \rangle &= \\ \left[\left(C_{\boldsymbol{\ell}}^{EE,s} + C_{|\boldsymbol{\ell}-\mathbf{L}|}^{BB,s} \right) \sin 2(\varphi_{\boldsymbol{\ell}} - \varphi_{\boldsymbol{\ell}-\mathbf{L}}) \right] m(\mathbf{L}) &+ E_c(\boldsymbol{\ell}) B_c^*(\boldsymbol{\ell} - \mathbf{L}). \end{aligned} \quad (10)$$

To linear order, we treat m and c as fixed external fields and average over the cosmic shear as a statistically isotropic field, while we neglect additional cross and higher-order terms. Equation (10) is the key statistical signature of a spatially varying m -bias: The first term shows that an E -mode at wavevector $\boldsymbol{\ell}$ couples to a B -mode at wavevector $\boldsymbol{\ell} - \mathbf{L}$ such that their correlation is proportional to the bias field mode $m(\mathbf{L})$. The second term acts as a mean field c -bias contamination to the m -bias response, if this is

present. At this point, we should note that the spectra in equation (10) contain only the signal. A fully consistent treatment would replace these by total spectra that include the shape noise as m -bias scales the measured shape. We verified that this choice does not change the results as we discuss in Section 2.3.

The methodology we develop here is directly analogous to the case of CMB lensing, where lensing by a potential ϕ induces correlations between E and B polarization modes of the CMB. In fact, equation (10) is the weak lensing shear counterpart of the CMB lensing relation given in A. Lewis & A. Challinor (2006) for EB mode coupling by the lensing potential (see their section 7.2).

2.3 Quadratic estimator for the m -bias field

Now, we construct a quadratic estimator for m in Fourier space, denoted $\hat{m}(\mathbf{L})$, formed from the quadratic combination of observed E and B fields. Following the formalism of A. Lewis & A. Challinor (2006), we can write the Fourier-space estimator for the m -bias field as

$$\hat{m}(\mathbf{L}) = \mathcal{N}(\mathbf{L}) \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} E^{\text{biased}}(\boldsymbol{\ell}) B^{\text{biased}*}(\boldsymbol{\ell} - \mathbf{L}) W(\boldsymbol{\ell}, \boldsymbol{\ell} - \mathbf{L}), \quad (11)$$

where $W(\boldsymbol{\ell}, \boldsymbol{\ell} - \mathbf{L})$ is a weighting function to be specified, and $\mathcal{N}(\mathbf{L})$ is an overall normalization factor. Equation (10) shows, if c -bias is zero, that the quadratic product $X_{\boldsymbol{\ell},\mathbf{L}} = E^{\text{biased}}(\boldsymbol{\ell}) B^{\text{biased}*}(\boldsymbol{\ell} - \mathbf{L})$ has mean

$$\langle X_{\boldsymbol{\ell},\mathbf{L}} \rangle = f(\boldsymbol{\ell}, \boldsymbol{\ell} - \mathbf{L}) m(\mathbf{L}), \quad (12)$$

where

$$f(\boldsymbol{\ell}, \boldsymbol{\ell} - \mathbf{L}) = (C_{\boldsymbol{\ell}}^{EE,s} + C_{|\boldsymbol{\ell}-\mathbf{L}|}^{BB,s}) \sin 2(\varphi_{\boldsymbol{\ell}} - \varphi_{\boldsymbol{\ell}-\mathbf{L}}). \quad (13)$$

Therefore, we estimate $m(\mathbf{L})$ by filtering $X_{\boldsymbol{\ell},\mathbf{L}}$ against the response f with inverse-variance weights for W , and use N to recover an unbiased, minimum-variance estimator. We note again here that, if there is a mean field EB c -bias contamination in equation (10), this also enters the m -bias reconstruction in equation (11). The integration runs over all Fourier modes $\boldsymbol{\ell}$, pairing an E -mode at $\boldsymbol{\ell}$ with a B -mode at $\boldsymbol{\ell} - \mathbf{L}$ so that the pair corresponds to the m -bias mode \mathbf{L} . Because the estimator is quadratic in the observed shear field (proportional to $E \times B$), it constitutes a quadratic estimator for m .

The role of $W(\boldsymbol{\ell}, \boldsymbol{\ell} - \mathbf{L})$ is to optimally weight each (E, B) mode pair in the sum. Intuitively, modes with a higher signal-to-noise, either due to larger cosmological signal or lower noise, should receive higher weight. We derive the form of W by demanding that \hat{m} is the minimum-variance unbiased estimator of m . First, the unbiasedness condition means that the expectation value of \hat{m} should equal the true m . However, until now, we have assumed that there is a bias term from the c mean field. Thus, using equation (10), we require

$$\langle \hat{m}(\mathbf{L}) \rangle = m(\mathbf{L}) + \mathcal{N}(\mathbf{L}) \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} E_c(\boldsymbol{\ell}) B_c^*(\boldsymbol{\ell} - \mathbf{L}) W(\boldsymbol{\ell}, \boldsymbol{\ell} - \mathbf{L}). \quad (14)$$

At this point, we should note that the c -bias, as we have already discussed in Section 1, can be detected and, in principle, corrected using null tests from the data itself (e.g. E. Uiter & P. Schneider 2016). Assuming that the data we treat here from now on are c -bias corrected unless otherwise stated, equation (14) becomes $\langle \hat{m}(\mathbf{L}) \rangle = m(\mathbf{L})$. Then, we can substitute the definition (11) and use $\langle E^{\text{biased}}(\boldsymbol{\ell}) B^{\text{biased}*}(\boldsymbol{\ell} - \mathbf{L}) \rangle$ from equation (10), so that the

condition becomes

$$\mathcal{N}(\mathbf{L}) \int \frac{d^2\ell}{(2\pi)^2} (C_\ell^{EE,s} + C_{|\ell-L|}^{BB,s}) \sin[2(\varphi_\ell - \varphi_{\ell-L})] W(\ell, \ell - \mathbf{L}) = 1, \quad (15)$$

to linear order in m . Equation (15) fixes the normalization factor $\mathcal{N}(\mathbf{L})$ in terms of the chosen weight function W . In practice, we choose W (up to an overall normalization) to minimize the variance of $\hat{m}(\mathbf{L})$. The variance $\text{Var}[\hat{m}(\mathbf{L})]$ can be derived by considering $\langle \hat{m}(\mathbf{L}) \hat{m}^*(\mathbf{L}') \rangle$ and assuming Gaussian statistics for the cosmic shear field $E^{\text{cosm}}(\ell)$. In this limit, we can use G. C. Wick (1950)'s theorem to express the resulting four-point correlators of E^{cosm} into sums of products of two-point functions. The dominant contributions to the estimator variance come from terms where the E and B fields in equation (11) contract with themselves, leading to noise bias. One finds that the variance is minimized if W is chosen proportional to the signal coupling kernel from equation (10), divided by the total power in each of the E and B modes (signal + noise) (c.f. A. Lewis & A. Challinor 2006). In other words, the optimal weight is

$$W(\ell, \ell - \mathbf{L}) = \frac{\sin[2(\varphi_\ell - \varphi_{\ell-L})]}{C_\ell^{EE,\text{tot}} C_{|\ell-L|}^{BB,\text{tot}}} (C_\ell^{EE,s} + C_{|\ell-L|}^{BB,s}), \quad (16)$$

where $C_\ell^{EE,\text{tot}} = C_\ell^{EE,s} + N_\ell^{EE}$ and $C_\ell^{BB,\text{tot}} = C_\ell^{BB,s} + N_\ell^{BB}$ denote the total E - and B -mode angular power spectra including noise.⁴ Here N_ℓ^{EE} and N_ℓ^{BB} are the noise power spectra of the E and B fields from intrinsic shape noise in the shear catalogs (note that shape noise typically contributes equally to E and B modes, on average). The factor $(C_\ell^{EE,s} + C_{|\ell-L|}^{BB,s})$ in the numerator of equation (16) arises from the fact that either leg of the estimator pair could carry the signal coupling. The $\sin[2(\varphi_\ell - \varphi_{\ell-L})]$ factor encodes the geometric dependence of the mode coupling (as in equation 10), ensuring that only the appropriate cross-orientation of E and B modes contributes to the signal. With this choice, modes are weighted by the inverse of their variance (the product of their total power spectra), which is the standard optimal weighting for a quadratic estimator.

Given the weight function (equation 16), we can now specify the normalization $\mathcal{N}(\mathbf{L})$ using equation (15):

$$\mathcal{N}(\mathbf{L})^{-1} = \int \frac{d^2\ell}{(2\pi)^2} \frac{(C_\ell^{EE,s} + C_{|\ell-L|}^{BB,s})^2}{C_\ell^{EE,\text{tot}} C_{|\ell-L|}^{BB,\text{tot}}} \sin^2[2(\varphi_\ell - \varphi_{\ell-L})]. \quad (17)$$

At this point, we should mention that in the case of cosmic shear, and contrary to the CMB lensing case, the m -bias affects the measured shear which already includes the shape noise. Thus, it is worth checking the inclusion of the shape noise also in the numerators of equations (16) and (17), despite using theory predictions for the spectra. We confirmed that the performance of the quadratic estimator is not impacted, and decided to retain the original formalism, with the shape noise only in the denominator.

In practice, one can compute $\mathcal{N}(\mathbf{L})$ given a fiducial E -mode angular power spectrum, assuming a cosmological model and known noise levels.

In principle, a more general estimator may be derived, using also the EE and BB correlations of the observed shear field, similar to the global minimum variance quadratic estimator used in CMB lensing reconstruction (A. S. Maniyar et al. 2021). Nevertheless, here, we restrict ourselves to the EB pair, which likely dominates the total signal-to-noise ratio (SNR) of the reconstruction,

⁴We note here that if c -bias is present, it contributes to the variance as $C_\ell^{XY,\text{tot}} = C_\ell^{XY,s} + N_\ell^{XY} + C_\ell^{XY,c}$.

since EB correlations should be zero in the absence of systematics. Its detection is thus a clearer indication of the m -bias field.

Finally, after computing $\hat{m}(\mathbf{L})$ in Fourier space, one can transform it back to real space, to obtain a two-dimensional map $\hat{m}(\boldsymbol{\theta})$ of the m -bias across the survey region. The resulting \hat{m} map can be used to diagnose m -bias in cosmic shear survey data. We note here that the method cannot constrain a constant m -bias (the $L = 0$ mode). In the next sections, we test this estimator on simulated Stage-IV-like shear maps, with injected artificial m -bias patterns, to validate its performance.

3 SURVEY SIMULATIONS AND m -BIAS REALIZATIONS

3.1 The testing methodology

We aim to first test the method on simplified realizations of weak lensing surveys, quantifying the ability of the method to recover some injected patterns of m -bias variation, noting the point that the method relies on detection of mode coupling, and hence is sensitive only to spatially dependent bias.

In weak lensing surveys, the cosmic shear signal is buried within the ‘shape noise’ caused by the broad distribution of intrinsic galaxy ellipticity (see e.g. S.-S. Li et al. 2023). We show below that, for realistic galaxy number density, it is not possible to recover percent-level m -biases at arbitrary positions within a survey. Any m -bias we measure has much larger statistical noise than the signatures that we aim to detect in Stage IV surveys.

We can overcome this limitation by analogous methods to those adopted for c -bias detection. First, if we have position-dependent templates or maps of systematic effects that are possible sources of bias, we may compare those maps with the m -bias map and assess the significance of any detection. Such systematics templates might be maps of PSF error, or perhaps other astronomical effects such as Galactic reddening, that might impact the shear measurement.

The second method is just to stack the m -bias estimator outputs themselves on some length-scale, analogously to the c -bias stacking methods. This approach is particularly appropriate for the case where m -bias arises from instrumental effects within the FOV – for example, some spatially varying PSF error which is invariant in telescope coordinates, leading to a periodic bias variation across the survey, dependent on the set of instrument pointings that make up the survey.

We shall investigate both approaches in the following sections. In order to reproduce the effect of a replicating pattern, we generate multiple realizations of patches of survey data, with randomly distributed galaxy shear values across the entire survey, but with an m -bias pattern that is invariant between the realizations.

3.2 Galaxy survey sample

The cosmic shear signal power spectrum $C_\ell^{EE,s}$ is related to the matter power spectrum via

$$C_\ell^{EE,s} = \int_0^{\chi_h} \frac{d\chi}{\chi^2} [W^E(\chi)]^2 P_{mm} \left(k = \frac{\ell + 1/2}{\chi}, \chi \right), \quad (18)$$

where $\chi(z)$ is the comoving distance for flat cosmologies, χ_h the comoving distance at the horizon and $P_{mm}(k, \chi)$ is the matter power spectrum as a function of wavenumber k and comoving distance (which we use here as a proxy for redshift, or cosmic time, in the lightcone). The weak lensing kernel is




Figure 1. *Left:* Galaxy redshift distribution $p(z)$ for a *Euclid*-like photometric sample (see equation 20). *Middle:* a blob-like m -bias given by equation (22). *Right:* a strip-like m -bias model given by equation (23). Both m -bias patterns have an rms of 5 per cent and are 20×20 pixel grids.

$$W^E(\chi) = \frac{3\Omega_{m,0}H_0^2}{2c^2} \chi [1 + z(\chi)] \int_{z(\chi)}^{\infty} dz p(z) \frac{\chi(z) - \chi}{\chi(z)}, \quad (19)$$

with $\Omega_{m,0}$ the fraction of the matter density of the Universe at present, H_0 the Hubble constant, c the speed of light and $p(z)$ the galaxy sample’s redshift distribution. The expressions above make use of the Limber approximation (D. N. Limber 1953; N. Kaiser 1992), appropriate for the broad weak lensing kernel.

To calculate C_ℓ^{EE} we use the Core Cosmology Library (CCL, N. E. Chisari et al. 2019b), with the matter power spectrum given by the HALOFIT parametrization of R. E. Smith et al. (2003) and R. Takahashi et al. (2012). We consider a *Euclid*-like (Euclid Collaboration 2025) redshift distribution describing the photometric/cosmic shear sample with the form

$$p(z) \propto z^\alpha \exp \left[- \left(\frac{z}{z_0} \right)^\beta \right], \quad (20)$$

with $\alpha = 2$, $\beta = 1.5$ and $z_0 = 0.637$ in $0 < z \leq 3$ (R. Laureijs et al. 2011). We assume a total number density of $\bar{p}_{\text{gal}} = 30 \text{ arcmin}^{-2}$ (see left panel of Fig. 1 for $p(z)$ normalized to unit integral). Most cosmological analyses will make use of tomography, dividing the sample into different redshift bins to enhance the resulting cosmological constraints. Instead here we simply assume a single redshift bin containing the whole sample. On the one hand, this is the simplest approach to showcase the potential of the method as a proof of concept. On the other, combining the full sample reduces the level of shape noise in the data, which likely improves the sensitivity to a redshift-independent (but spatially varying) m -bias.

The contribution to the total observed power spectrum receives a contribution from shape noise of the form

$$N_\ell^{EE} = N_\ell^{BB} = \sigma_e^2 / \bar{p}_{\text{gal}}, \quad (21)$$

with $\sigma_e = 0.28$ the per-component intrinsic ellipticity dispersion for a *Euclid*-like galaxy sample (Euclid Collaboration 2025).

Besides the contribution from gravitational lensing, the correlated galaxy shapes also receive a contribution from the intrinsic alignments (IAs) of galaxies due to the tidal interactions and other long-range astrophysical effects taking place during galaxy formation (B. Joachimi et al. 2015). In the simplest parametrization, the so-called non-linear alignment model (S. Bridle & L. King 2007), the IA contribution is a pure E mode, and its presence therefore does not affect the performance of the quadratic estimator to reconstruct the m -bias map. Furthermore, omitting the IA contribution to C_ℓ^{EE} in the quadratic estimator does not

significantly affect our results. We verified this by including this contribution, choosing an IA amplitude $A_{\text{IA}} = 2$, both in C_ℓ^{EE} and in the simulated maps from the Gaussian realizations (for more details, see Section 3.4) used in the tests described below, or including it only in one of the two. In both cases, the reconstructed m -bias maps were indistinguishable from the maps reconstructed in the absence of IAs.

In addition to IAs, weak lensing measurements are also affected by baryonic feedback, which systematically suppresses the matter power spectrum at small scales (N. E. Chisari et al. 2019a). For this, we performed the same series of tests using the HMCODE (A. J. Mead et al. 2021), with the baryonic feedback efficiency parameter $\log_{10}(T_{\text{AGN}}/\text{K})$ set to 7.8, which is the fit to the $\nu\Lambda\text{CDM}$ BAHAMAS simulations, and also found that it did not affect the results. Thus, we decided not to include IA or baryonic effects in the main analysis in Section 4.

3.3 m -bias patterns

We introduce two ‘toy models’ describing flat-sky spatially varying m -bias patterns $m(\theta)$, inspired by the models applied in T. D. Kitching et al. (2019), which could be induced by residual effects in the PSF (see their Sec. 3). We define a square patch of size $L_x = L_y$ (in radians) sampled on an $N_x \times N_y$ pixel grid. Let $\theta = (\theta_x, \theta_y)$ denote angular coordinates in radians, with grid arrays X, Y covering $[-L_x/2, L_x/2] \times [-L_y/2, L_y/2]$. We introduce a ‘blob’-like pattern by parametrizing $m(\theta)$ as

$$m_{\text{blob}}(\theta) = A \sin(k_b |\theta_x|) \sin(k_b |\theta_y|), \quad (22)$$

with A an arbitrary constant, and k_b is the angular frequency characterizing the size of the blobs set to 50. An example can be seen in the middle panel of Fig. 1. As an alternative example, we introduce a ‘strip’-like pattern, given by

$$m_{\text{strip}}(\theta) = -A \sin(\pi - k_s \theta_x), \quad (23)$$

where k_s characterizes the size of the strips along the x axis set to 100 (see right panel of Fig. 1).

In both cases we fix the amplitude parameter A to ensure a desired rms variation m_{rms}

$$m_{\text{rms}} = \sqrt{\langle m_{\text{pattern}}^2 \rangle_{\text{pix}}}, \quad (24)$$

where $\langle \cdot \rangle_{\text{pix}}$ denotes the pixel average over the $N_x \times N_y$ grid. We consider two rms values, $m_{\text{rms}} = 0.05$ (5 per cent) and 0.01 (1 per cent), motivated by the values identified by C. Cragg et al.

(2023) as being potentially of significance for Stage IV weak lensing surveys.

3.4 Gaussian realizations

The m -bias modulates the cosmic shear signal as well as the intrinsic shape noise, both of which we also simulate. As a first simple model, we generate Gaussian realizations of both components.

We generate flat-sky signal realizations following the E -mode power spectrum in equation (18), using the `NaMaster` package (D. Alonso et al. 2019), including angular scales in the range $[\ell_{\min}, \ell_{\max}]$, with $\ell_{\min} = 2\pi/L_{\text{side}}$ and $\ell_{\max} = \sqrt{2\pi}N_{\text{side}}/L_{\text{side}}$. The input power spectrum was generated for a fiducial Λ CDM cosmology with parameters

$$(\Omega_c, \Omega_b, h, \sigma_8, n_s) = (0.2432, 0.0473, 0.6898, 0.8364, 0.969), \quad (25)$$

where Ω_c is the present cold dark matter density fraction, Ω_b is the present baryon density fraction, $h = H_0/100 \text{ km s}^{-1} \text{ Mpc}^{-1}$ is the dimensionless Hubble constant, σ_8 is the rms of the linear matter density fluctuations in spheres of $8h^{-1} \text{ Mpc}$, and n_s is the scalar spectral index.

We assign shape noise in each pixel as follows. For simplicity we assume the same number of galaxies in each pixel, given by $N_p = \bar{p}_{\text{gal}}L_xL_y$. We then generate values for the shape noise in each pixel and shear component as uncorrelated Gaussian numbers with zero mean and a standard deviation $\sigma_p = \sigma_e/\sqrt{N_p}$.

Finally, the observed shear field is generated by adding the signal and noise components, and multiplying them by the m -bias patterns described in the previous section. These are then used in equation (11) to reconstruct the input bias pattern.

3.5 N-body simulations

Apart from using Gaussian realizations, we also test our quadratic estimator with N-body simulations from the *cosmo-SLICS* suite (J. Harnois-Déraps, B. Giblin & B. Joachimi 2019). The output of these N-body runs is convergence κ and shear γ_1, γ_2 maps derived with 26 w CDM cosmology models and two random seeds each. In addition, for each model, there is a light cone spanning from $z_{\min} = 0$ to $z_{\max} = 3$ with fixed redshift snapshots, spanning from 15 to 28 depending on the cosmology model. There are 50 light cones per cosmology model, in total. We choose their fiducial Λ CDM cosmology model maps (see Section 3.4 for the exact values). The shear maps γ_1, γ_2 are grids of 7745×7745 pixels spanning an area of 100 deg^2 (10 deg per grid side). Based on these, we construct a ray-tracing of each light cone given the *Euclid*-like $p(z)$ redshift distribution (see equation 20) by estimating the weighted sum of the shear maps $\sum_i w_i \gamma_1^i, \sum_i w_i \gamma_2^i$, where $w_i = p(z_i)\Delta z_i$, with i the redshift snapshot and Δz_i their half intervals. Then, we assign the shape noise and the m -bias pattern in the pixel grids that enter the quadratic estimator, as we similarly described for the Gaussian realizations in Section 3.4. Although the N-body simulations have a limit to their available volume, the Gaussian realizations contain only two-point information, so they cannot capture the non-linear, non-Gaussian structure of the shear field that arises from late-time structure formation. The *cosmo-SLICS* N-body light cones deliver the non-Gaussian lensing maps with realistic small-scale power and super-sample coupling. Therefore, these simulations are more accurate descriptions of the observed cosmic shear field. In addition, non-linearity can lead to coupling between ℓ modes,

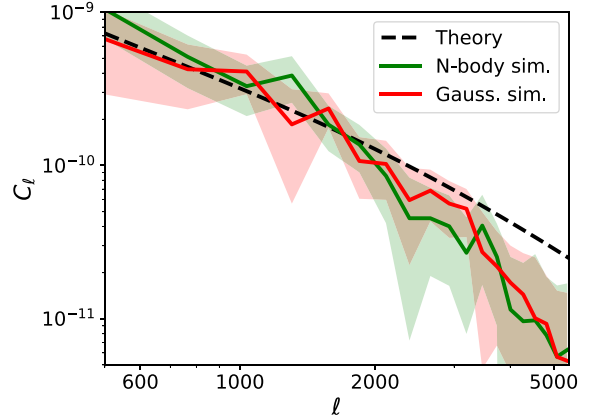


Figure 2. The angular power spectra from an N-body (green) and a Gaussian (red) simulation at an area of 1 deg^2 with their 1σ uncertainties and the input theory prediction (dashed black). The high- ℓ suppression visible in the figure is due to finite-resolution effects.

which would act as an additional source of noise for the reconstructed m -bias.

Despite the lack of non-linear structures in the Gaussian realizations compared to the N-body simulations at small scales below 1 deg^2 , the difference between their angular power spectra is expected to be small (see fig. 4 of J. Harnois-Déraps et al. 2019). In Fig. 2, we make a quick sanity check by comparing one N-body simulation with one Gaussian realization. We compute the flat-sky angular power spectra along with their 1σ uncertainty bounds using the `pymaster` package for one random pair of γ_1 and γ_2 pixel maps (31×31 downsampled pixels) from one N-body simulation and a Gaussian realization (green and red lines). We downsample (more details on the downsampling are described in Sections 4.3 and 4.4) the Gaussian realization here the same way we do for the N-body simulations, in order to have the same resolution and pixel window size effects. We also plot the theory prediction from `HALOFIT`. Indeed, we found that the two spectra are in agreement within their 1σ uncertainties. However, we see that the simulated spectra start to deviate from the theory prediction at $\ell \gtrsim 2000$, due to resolution and pixel window effects (see J. Harnois-Déraps et al. 2019). We have validated that, assuming a theory power spectrum in the estimator of equation (11), described by a simple `HALOFIT` at up to $\ell_{\max} \lesssim 1000$, in order to avoid such deviations, we are able to reconstruct accurately the non-Gaussian m -bias patterns without any loss of information.

4 RESULTS

4.1 Evaluation of results

As discussed in Section 3.1, an average over a number of realizations of the integral of equation (11), equivalent to averaging over multiple sky patches in a weak lensing survey, is needed in order to suppress the shape noise.

Although, in reality, the m -bias pattern and its rms amplitude are generally unknown, in cases where there are available templates describing certain types of systematic effects such as the PSF model, these could be compared to the data in a χ^2 analysis. In this exercise, we have created m -bias toy models that we assume as templates (see again equations 22 and 23). Therefore, we

can quantify the confidence with which the quadratic estimator in equation (11) is able to reconstruct the input *m*-bias pattern in terms of its spatial variation and rms, by defining an SNR. This way, we can come up with an estimate of the number of realizations or sky-patch stacks one would need in order to detect an *m*-bias pattern of a certain rms. An estimate of the SNR even in the absence of a template is also possible by simply measuring the detection significance of the signal itself compared to noise.

4.2 Signal-to-noise estimates

To assess the performance of the quadratic estimator, we must quantify the significance of the reconstructed *m*-bias map. Let $\hat{m}_i(\boldsymbol{\theta})$ be the reconstructed multiplicative bias map in the *i*-th simulated realization, where each realization corresponds to a sky patch of a large-area survey. We estimate the local variance of a given realization as

$$s^2(\boldsymbol{\theta}) = \frac{1}{N-1} \sum_{i=1}^N [\hat{m}_i(\boldsymbol{\theta}) - \bar{m}_N(\boldsymbol{\theta})]^2, \quad (26)$$

where *N* is the total number of realizations, and \bar{m}_N is the mean reconstructed map

$$\bar{m}_N(\boldsymbol{\theta}) = \frac{1}{N} \sum_{i=1}^N \hat{m}_i(\boldsymbol{\theta}). \quad (27)$$

Additionally, since all realizations are independent, the standard deviation of \bar{m}_N is $\sigma_{\bar{m}}(\boldsymbol{\theta}) \equiv s(\boldsymbol{\theta})/\sqrt{N}$.

We use three different SNR metrics to quantify the significance of the reconstructed *m*-bias map:

(i) **Template fitting SNR:** Let us assume a linear model for the reconstructed *m*-bias map of the form $\mathbf{d} = \alpha \mathbf{t} + \mathbf{n}$, where \mathbf{d} is a vector containing the reconstructed map in each pixel, and \mathbf{t} is a known template map (e.g. the focal-plane map of a given systematic known to be a potential source of multiplicative bias), with α an unknown constant, and \mathbf{n} is the reconstruction noise component. The significance of the detection in this case may be estimated as

$$\text{SNR}_{\text{temp}} = \bar{\alpha} / \sigma_{\alpha}, \quad (28)$$

where $\bar{\alpha}$ and σ_{α} are the best-fitting value of α and its statistical uncertainty, given by

$$\bar{\alpha} = \frac{\mathbf{t}^T C^{-1} \mathbf{d}}{\mathbf{t}^T C^{-1} \mathbf{t}}, \quad \sigma_{\alpha} = \frac{1}{\sqrt{\mathbf{t}^T C^{-1} \mathbf{t}}}. \quad (29)$$

Here *C* is the per-pixel covariance matrix, which we approximate as $C = \text{diag}(\sigma_m^2)$. We acknowledge that this approximation is optimistic, since it underestimates the SNR, as we discuss in Section 4.3. Despite that, the scaling with patch number or the relative conclusions of the paper do not change.

(ii) **Data-driven SNR:** One can also derive an SNR estimate from the data itself. The null χ^2 is defined as

$$\chi_{\text{null}}^2 = \mathbf{d}^T C^{-1} \mathbf{d}, \quad (30)$$

and by considering the degrees of freedom, which in our case is the number of pixels N_{pix} , we can derive a σ -detection definition which is

$$\text{SNR}_{\text{data}} = \sqrt{\chi_{\text{null}}^2 - N_{\text{pix}}}. \quad (31)$$

In case of no *m*-bias in the data, $\chi_{\text{null}}^2 \rightarrow N_{\text{pix}}$, so that $\text{SNR}_{\text{data}} \rightarrow 0$

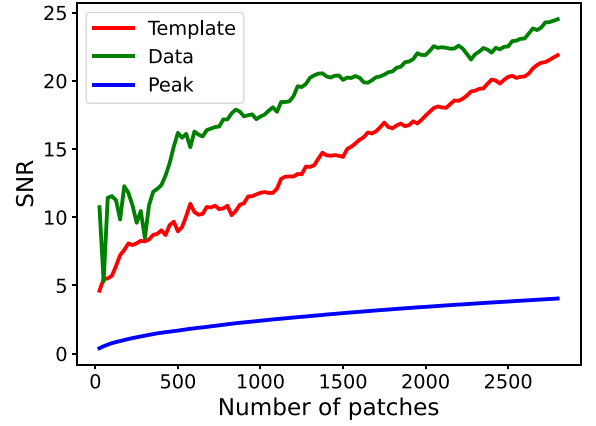


Figure 3. The ‘template’, ‘data’, and ‘peak’ SNR definitions as a function of the number of patches. We consider patches of $10 \times 10 \text{ deg}^2$ and a 5 per cent rms blob-like *m*-bias.

(iii) **Peak SNR:** Finally, for an alternative SNR definition, one can define a ‘peak mask’ \mathcal{S} , selecting the brightest parts of the injected pattern, specifically those pixels where the absolute value of the true *m*-bias map is above half of its maximum value

$$\mathcal{S} \equiv \left\{ \boldsymbol{\theta} : |m(\boldsymbol{\theta})| > \frac{1}{2} \max_{\boldsymbol{\theta}'} |m(\boldsymbol{\theta}')| \right\}. \quad (32)$$

We then average the signal amplitude of the true map and the noise from the realizations over this mask and form a local SNR dubbed as ‘peak’:

$$\text{SNR}_{\text{peak}} = \frac{\langle |m(\boldsymbol{\theta})| \rangle_{\boldsymbol{\theta} \in \mathcal{S}}}{\langle \sigma_{\hat{m}}(\boldsymbol{\theta}) \rangle_{\boldsymbol{\theta} \in \mathcal{S}}}. \quad (33)$$

This SNR definition gives a sense of the significance with which the largest deviations in the *m*-bias pattern may be identified in a single pixel.

We established the number of sky patches required to reach a target $\text{SNR}_{\text{def}}^{\text{req}}$ by increasing the number of patches until $\text{SNR}_{\text{def}} \geq \text{SNR}_{\text{def}}^{\text{req}}$.

4.3 *m*-bias reconstruction on large scales

We aim to test the *m*-bias quadratic estimator at large scales. We start by constructing Gaussian realizations of $10 \times 10 \text{ deg}^2$ patches. We choose a $N_x \times N_y = 20 \times 20$ pixel grid for our maps and the number of galaxies in each patch is $\sim 10^7$. In principle, one could choose a finer pixel grid, but the reason not to opt for this is two-fold. First, the spatially varying *m*-bias toy models we explore here are characterized by large enough scale variations corresponding to 5 deg (blob-like) and 2 deg (strip-like) in a 100 deg^2 area, which are captured with such a pixel grid size (analogously, the same is true for the 1 deg^2 patches we investigate in Section 4.4). Secondly, it is a considerable speed-up for the code as well, since the computational cost scales with the FFT done on the number of pixels in the grid, allowing us to explore large numbers of realizations.

Within this setup, and using a blob-like *m*-bias template with 5 per cent rms fluctuations, Fig. 3 shows the number of stacked realizations needed to achieve a given SNR values with the ‘template’, ‘data’, and ‘peak’ definitions. The ‘peak’ definition is naturally stricter, since it quantifies the SNR in a single bright pixel.

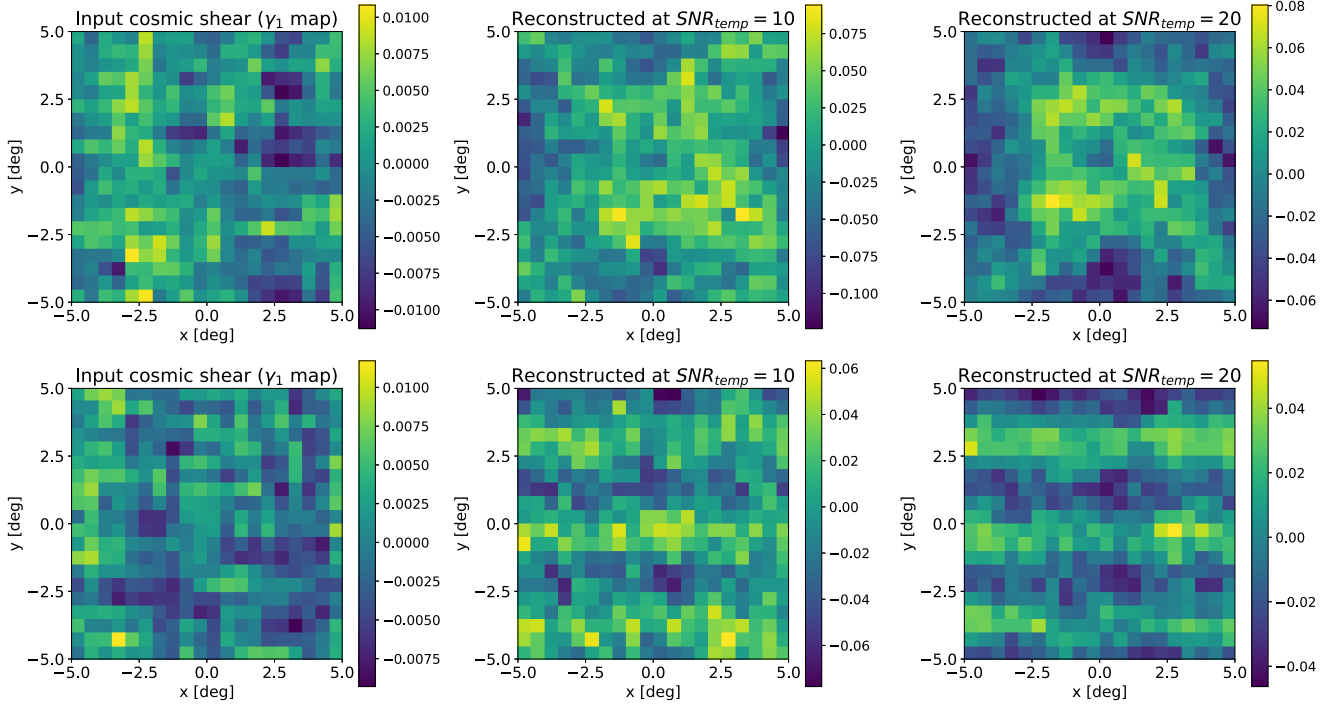


Figure 4. *Top left:* one input γ_1 component of the cosmic shear from Gaussian realizations. *Top middle:* reconstruction of the 5 per cent rms blob-like m -bias field at $\text{SNR}_{\text{temp}} = 10$ with 600 realizations. *Top right:* reconstruction of the same m -bias field at $\text{SNR}_{\text{temp}} = 20$ with 2800 realizations. The grid area is 100 deg^2 with 20×20 pixels. Similarly, the bottom panels show the 5 per cent rms strip-like m -bias model reconstructed with 2000 (bottom middle) and 6700 realizations (bottom right).

It reaches a value of $\text{SNR}_{\text{peak}} = 4$ for 2800 realizations, at which point the template and data-based definitions reach $\text{SNR}_{\text{temp}} \sim 20$ and $\text{SNR}_{\text{data}} \sim 24$. Fig. 4 shows one of the Gaussian realizations used for this exercise, as well as the reconstructed m -bias map after stacking enough realizations to achieve a total $\text{SNR}_{\text{temp}} = 10$ (corresponding to $\text{SNR}_{\text{peak}} = 2$), where the bias pattern starts becoming visually evident. We can also appreciate that the reconstruction looks less noisy when we reach $\text{SNR}_{\text{temp}} = 20$. At this point we verify that $\bar{\alpha}$ and its uncertainty σ_{α} in equation (29) is 1.003 ± 0.056 at $\text{SNR}_{\text{temp}} = 20$, indicating no bias in the estimator. Additionally, as already hinted in Section 4.2, neglecting off-diagonal terms in the pixel-by-pixel covariance matrix, makes SNR optimistic. We quantified this by estimating the template amplitude $\bar{\alpha}$ (see again equation 29) on a realization-by-realization basis and computing its scatter directly from simulations. The Monte Carlo estimate of σ_{α} fully includes all the pixel-by-pixel correlations in the reconstructed maps. For a 5 per cent blob-like m -bias pattern at $\text{SNR}_{\text{temp}} \sim 20$, we find that the diagonal covariance approximation underestimates σ_{α} by a factor of ~ 1.7 (factor scaling unchanged with amplitude and pattern of m -bias), corresponding to a Monte Carlo calibrated $\text{SNR}_{\text{temp}} \sim 12$. Likewise, the diagonal covariance approximation underestimates SNR_{data} by a factor of 1.2. Given the numerical instability of inverting the full covariance for our setup, we retain the diagonal approximation in this work and treat SNR as optimistic detection metrics, while basing our conclusions on the relative scaling with patch number and m -bias amplitude, which are unaffected by this choice. For a statistical detection of the m -bias through template fitting, corresponding to $\text{SNR}_{\text{temp}} = 5$ (i.e. a ‘ 5σ ’ detection), we require 50 realizations, corresponding to a total sky area of 5000 deg^2 (even though the m -bias pattern is not visually evident at

that point). The quadratic estimator approach proposed here may therefore be used to identify a spatially varying multiplicative bias as long as an informed template for its spatial pattern can be constructed. This may be constructed, for example, as sky maps of scalar survey conditions known to be highly correlated with shape measurement errors (e.g. co-added PSF size, or Galactic reddening, which may impact shear measurement through its systematic effect on galaxy observed spectral energy distributions). This procedure is reminiscent of the techniques used to identify additive systematics in cosmic shear and galaxy clustering (F. Berlfein et al. 2024), but now applicable to multiplicative bias effects.

The bottom panels in Fig. 4 show the same results for a strip-like m -bias, where the same conclusions drawn for the blob-like case hold qualitatively (6700 realizations are now needed). In what follows, we report the number of realizations/patches needed to achieve an $\text{SNR}_{\text{temp}} = 20$ unless otherwise stated, where the m -bias pattern can be visually identified (corresponding to $\text{SNR}_{\text{peak}} \sim 5$ and $\text{SNR}_{\text{data}} \sim 24$ or much higher, depending on the case).

For independent realizations the reconstruction noise is the same and scales with the number of realizations N as $\propto 1/\sqrt{N}$. Therefore, a desired SNR is reached at $\propto \text{rms}\sqrt{N}$ and going from 5 per cent to 1 per cent rms values requires ~ 25 times the number of realizations to achieve the same SNR. Indeed, we recover the input 1 per cent rms m -bias at $\text{SNR}_{\text{temp}} = 20$, but now with an increased number of realizations by that factor, corresponding to 80 750 and 160 875 for the blob and the strip patterns, respectively. We acknowledge that these numbers of realizations or equivalently multiple sky patches needed imply large sky areas which are not available in a real data analysis.




Figure 5. *Top left:* one input γ_1 component of the cosmic shear from N-body realizations. *Top right:* reconstruction of the 20 per cent rms *m*-bias field at $\text{SNR}_{\text{temp}} = 81$ with 50 simulations. The grid area is 100 deg^2 , with 100×100 pixels. Bottom panels are for a 20 per cent rms strip-like *m*-bias, reconstructed with $\text{SNR}_{\text{temp}} = 81$, using also 50 simulations.

We only choose these numbers here as a reference in order to obtain the good visual reconstruction at $\text{SNR}_{\text{temp}} = 20$. In fact, we could still detect the *m*-bias of both patterns at $\text{SNR}_{\text{temp}} = 3$ with roughly 100 realizations, numbers certainly available in the upcoming data since they correspond to a total sky area of $10\,000 \text{ deg}^2$. The results for $\text{SNR}_{\text{temp}} = 20$ visually look very similar to the recovered patterns of Fig. 4, and therefore we do not show them. It is worth noting here that the increased number of realizations needed is expected, given that we need to suppress the shear noise further in order to detect an *m*-bias with lower rms values.

We now turn our attention to the shear maps from the N-body simulations. As discussed in Section 3.5, the N-body simulations include realistic non-Gaussianities that can affect the performance of the quadratic estimator, by introducing additional noise or biases in the *m*-bias reconstruction. The output maps are pixel grids of size 7745×7745 . We assign the shape noise and then downsample these maps with a Gaussian filtering, with a width corresponding to the new pixel Nyquist frequency and also with a cubic interpolation scheme, into grids of size 100×100 . Afterwards, we multiply with the *m*-bias toy models as we did for the Gaussian realizations. The total number of realizations we consider here is 50 (25 light cones generated with two random seeds). Given that small number of simulations, we examine an exaggerated *m*-bias scenario in order to have a good visual emergence of the *m*-bias field. Thus, by setting the rms to 20 per cent just for this case, we evaluate the performance of the

quadratic estimator using all the 50 N-body simulation shear maps covering 100 deg^2 patches. In Fig. 5, we see that with all the simulations available we manage to capture very well visually the spatial-variation of the input *m*-bias patterns at $\text{SNR}_{\text{temp}} \sim 81$ for both the blob and strip models as seen in the top and the bottom panels of Fig. 5, respectively ($\text{SNR}_{\text{temp}} \sim 20$ is reached with only 5). At this point, we cannot safely conclude yet that our quadratic estimator is not affected by the non-linearities in the N-body simulations since we use an exaggerated 20 per cent rms *m*-bias, but this is a first test verifying that the quadratic estimator works using the N-body simulations. We will see in detail in Section 4.4 that indeed our estimator is also not impacted by these non-Gaussianities at small scales using realistic rms *m*-bias values.

4.4 *m*-bias reconstruction on small scales

There are good reasons to test the quadratic estimator at a smaller area of the order of $\sim 1 \times 1 \text{ deg}^2$: forthcoming weak lensing galaxy surveys, such as *Euclid* (Euclid Collaboration 2025) and LSST (LSST Dark Energy Science Collaboration 2012), have \sim -degree-sized fields of view (diameter $D_{\text{FOV}} \sim 0.7^\circ$ and $\sim 3.5^\circ$, respectively). Many shape measurement systematic effects such as PSF modelling errors will be approximately constant in telescope coordinates, and therefore will leave a repeated (or at least highly correlated) imprint across different exposures on




Figure 6. *Top left:* one γ_1 component patch of the cosmic shear from N-body realizations. *Top right:* reconstruction of the m -bias field with 5 per cent amplitude at $\text{SNR}_{\text{temp}} = 20$ with 400 patches. Bottom panels same as top but for a strip-like m -bias model with 1000 patches. The reconstructed grid area is $0.68 \times 0.68 \text{ deg}^2$ with 21×21 pixels.

scales smaller than D_{FOV} . The presence of these systematics could thus be detected by stacking maps centred at those exposures in telescope coordinates using the quadratic estimator described here. The total sky area covered by these surveys, of order $\sim 1.5 \times 10^4 \text{ deg}^2$, would allow for $\sim 1.5 \times 10^4$ degree-sized patches. Here we will show that these would be more than enough to detect a 5 per cent rms spatially varying m -bias patterns at a high significance (e.g. $\text{SNR}_{\text{temp}} = 20$) with a realistic small-scale cosmic shear signal from N-body simulations. Moreover, with a *Euclid*-like redshift distribution, it would be possible to detect 1 per cent rms m -bias at the same significance with a few thousand patches.

For this test we consider the cosmic shear realizations from the N-body simulations, and assume the same *Euclid*-like galaxy density number used above (see Section 3.2). The corresponding number of galaxies residing in a typical 1 deg^2 patch is roughly $\sim 10^5$ sources. In each of the 50 simulations, we split the 7745×7745 pixel grid spanning an area of 100 deg^2 into 100 sub-grids of size 774×774 , cropping 5 pixels at the outer edges to form a 10×10 tiling, each spanning an area now of $\sim 1 \text{ deg}^2$. Similarly to the procedure in Section 4.3, we include Gaussian shape noise and downsample these grid maps with a Gaussian filter and a cubic interpolation scheme into 31×31 pixel grids. We imprint the multiplicative bias pattern in the central 21×21 pixels of each patch, leaving a buffer with a width of 5 pixels where $m(\theta) = 0$. We do this to avoid edge effects caused by the implicit assumption of periodic boundary conditions in the FFTs. This assumption was valid in the previous sections, where we analysed

the full simulation boxes, but it breaks down when we work with smaller shear patches: the shear field is discontinuous at the patch boundaries, so large-scale modes are poorly represented and generate artefacts. We therefore introduce a buffer region around each patch before convolving with the usual m -bias pattern and applying the quadratic estimator in equation (11). We have verified that, without this buffer, spurious features appear in the reconstructed m -bias field near the patch edges, whereas they disappear when the buffer is included.

In Fig. 6, we present the 5 per cent rms for blob-like and strip-like m -bias patterns reconstructed at $\text{SNR}_{\text{temp}} = 20$ by averaging over 400 and 1000 patches. Since the appropriate templates at sub-FOV scales may be difficult to construct, we also report the data-driven SNR_{data} values which are 30 and 39 using the 400 and 1000 patches, respectively. For an 1 per cent rms m -bias the $\text{SNR}_{\text{temp}} = 20$ is achieved now with 2800 and 7600 patches for a blob-like and strip-like m -bias, respectively, corresponding to SNR_{data} of 54 and 86. These constitute the largest detection significance reported here, and show that spatially varying m -biases on sub-FOV scales with ~ 1 per cent amplitudes can be detected at high significance using the quadratic estimator approach.

4.5 Cosmology independence

The quadratic estimator presented here requires an estimate of the cosmic shear power spectrum C_ℓ^{EE} (see equation 16). Since this is *a priori* unknown, it is interesting to quantify whether using an incorrect C_ℓ^{EE} would lead to a significant error in the

reconstructed *m*-bias map. To do this, we repeat our analysis of the N-body simulation maps, which use the fiducial cosmology Λ CDM, but assuming a significantly different w CDM cosmological model to calculate the input C_{ℓ}^{EE} , with parameters

$$(w, \Omega_c, \Omega_b, h, \sigma_8, n_s) = (-1.9127, 0.5009, 0.0473, 0.7645, 0.4716, 0.969). \quad (34)$$

This cosmological model was shown in J. Harnois-Déraps et al. (2019), to exhibit strong differences in the shape of C_{ℓ}^{EE} , due to the strong changes in the growth history and distance-redshift relation caused by a high Ω_c and an extreme phantom-like value of w . We use 1 deg² patches, as in Section 4.4.

Despite the inconsistent cosmology models between the simulations in the shear maps and the theory spectra, we manage to recover the *m*-bias field successfully. The reconstruction for a 1 per cent rms blob-like and strip-like *m*-bias with N-body simulations with $\text{SNR}_{\text{temp}} = 20$ is achieved with ~ 2800 and ~ 7800 patches, respectively. The reconstructions for the 5 per cent *m*-bias patterns are similar to the findings in Section 4.4 as well. This means that the quadratic estimator formalism in equation (11) we employ here can be largely insensitive to the cosmology model we assume in the theory angular power spectrum entering the quadratic estimator. This is not entirely surprising: the estimator recovers the *m*-bias map by correlating *EB* Fourier modes with different wavenumbers. Using an incorrect C_{ℓ}^{EE} only leads to a sub-optimal weighting of the different mode pairs, which may affect the noise in the reconstructed *m*-bias map, but not the ability of the method to recover it.

4.6 Additive bias contamination

We also test whether the presence of an additive *c*-bias affects our ability to reconstruct the *m*-bias. The *c*-bias is expected to be one order of magnitude smaller than the *m*-bias in Stage IV galaxy surveys (T. D. Kitching et al. 2019). Following equation (1), we consider the case of an rms strip-like *c*-bias pattern amplitude of 0.1 per cent and a 1 per cent rms blob-like *m*-bias pattern (similar results are found when swapping the strip-like and blob-like patterns). We find that the presence of a *c*-bias contaminant has no discernible impact in the reconstructed *m*-bias map, and does not change any of our conclusions regarding its detectability.

We further confirm that our quadratic estimator definition in equation (11) is only sensitive to the *m*-bias, and not to the *c*-bias, by setting the *m*-bias pattern to zero and keeping only the *c*-bias patterns. For this, we test *c*-bias maps with an rms amplitude of 0.1 per cent, as well as two unrealistic cases with much larger amplitudes of 1 per cent and 5 per cent. In all cases, we fail to reconstruct the input *c*-bias patterns, and the reconstructed *m*-bias is compatible with zero. This finding confirms that the quadratic estimator is insensitive to contaminating *c*-bias.

5 CONCLUSIONS

We have presented and validated a quadratic estimator that directly reconstructs spatially varying *m*-shear bias, by exploiting the *EB* off-diagonal mode couplings that such a field induces in cosmic shear maps. Building on this, our estimator combines quadratic combinations of *E* and *B* with optimal normalized weights to reconstruct a map of *m*-bias variations to first order, in close analogy with standard CMB lensing reconstruction methods (A. Lewis & A. Challinor 2006). This statistic is by construc-

tion insensitive to a constant *m*-bias, and it does not respond to *c*-bias.

To quantify the method’s ability to detect a given *m*-bias pattern, we used three SNR definitions: ‘template’, ‘data’, and ‘peak’ (where the pattern appears visually), with the ‘template’ definition being the most relevant when an external template for the spatial pattern is available. We have explored two different scenarios in which such an estimator could be deployed.

First, multiple systematic effects are naturally described in telescope coordinates (e.g. PSF variations across the focal plane), and will imprint a repeating pattern across different exposures on scales smaller than the instrument’s FOV. A spatially varying *m*-bias may then be reconstructed by applying our quadratic estimator to FOV-sized patches and then stacking the result over multiple telescope pointings. In this case, we have shown that assuming a *Euclid*-like survey and using non-Gaussian N-body shear field maps from the *cosmo*-SLICS suite, we would be able to detect a blob-like (strip-like) *m*-bias pattern with a 5 per cent amplitude with an $\text{SNR}_{\text{temp}} = 20$ after stacking over only 400 (1000) patches. In turn, 1 per cent variations could be detected at the same significance with 2800 (7600) patches, still within the ‘few-thousand-patch’ regime, feasible for a wide-area Stage IV experiment. Similar conclusions hold if using the ‘data’ SNR definition instead.

Secondly, we explored the possibility of reconstructing variations on larger scales, as could be caused by varying observing conditions throughout the experiment’s observing campaign. In such a scenario, a template search for spatially varying *m*-bias may still be feasible, by building template maps of the most likely sources of shape measurement systematics from per-exposure metadata. On 10×10 deg² patches, Gaussian shear realizations show that a 5 per cent rms *m*-bias pattern is recovered at $\text{SNR}_{\text{temp}} = 20$ with 2800 and 6700 co-added sky patches, for ‘blob-like’ and ‘strip-like’ patterns, respectively. Although visual morphology requires a high SNR, template-amplitude fitting achieves 5σ detections with as few as ~ 50 patches. For 1 per cent rms, reaching the same $\text{SNR}_{\text{temp}} = 20$ significance demands of the order of 100 000 patches, for both patterns, consistent with the need to suppress noise much further at the percent level but unfeasible given the galaxy survey areas. Despite that, a 3σ detection is still feasible with ~ 100 patches. Using again the N-body simulations in the same 100 deg² geometry, we were able to test the estimator for a large input bias of 20 per cent and found that this could be detected at high $\text{SNR}_{\text{temp}} \simeq 81$. We were not able to test the method on larger areas, owing to the limited simulation volume available, but a naive scaling by the square root of the observed area predicts that $\text{SNR}_{\text{temp}} \simeq 5$ could be achieved for 1 per cent bias with a sky area of order 8000 deg², broadly consistent with the results from the Gaussian simulations.

A set of robustness tests supports the method’s applicability to real analyses. First, changing the theory spectrum used in the estimator weights from the fiducial Λ CDM to a deliberately extreme w CDM model leaves the reconstruction performance essentially unchanged, indicating that the estimator’s response is largely insensitive to moderate cosmology mis-specification in the filtering. Secondly, injecting *c*-bias at realistic levels does not bias the *m*-bias reconstruction nor increase the number of stacks required, and, in the absence of *m*, the quadratic estimator does not yield any false detection arising from contaminating additive bias. Thirdly, including intrinsic alignments or baryonic feedback either in the Gaussian realization maps, the theory spectra, or both produced indistinguishable results from the IA-free and baryon

feedback-free case at the scales and noise levels investigated. Finally, the flat-sky implementation used here is adequate for the patch sizes analysed, while a curved-sky generalization with spin harmonics and Wigner-3j couplings should be straightforward in a future development.

The results in this paper demonstrate that spatially varying m -bias at the few-percent to percent level is detectable at high significance with realistic survey characteristics, provided one stacks a sufficient number of FOV-scale patches or, when available, leverages physically motivated templates (e.g. PSF size/error proxies, detector coordinates). For the *Euclid*/LSST surveys, the available number of independent FOVs comfortably supports the stack sizes implied by our 1 deg^2 forecasts. There is also another natural extension apart from the curved-sky formalism. The method can include masks, and also allow for a forward-modelling of the reconstructed m -bias maps at the power spectrum level, via a weight map coupling, in order to correct the data from the spatially varying m -bias pattern.

In summary, the estimator introduced here provides a practical and robust route to detecting spatially varying m shear bias at the angular scales and amplitudes most relevant for Stage IV weak lensing cosmology. Its demonstrated performance on both Gaussian and fully non-Gaussian shear fields from N-body simulations, its weak dependence on cosmology assumptions, and its limited susceptibility to c -bias contamination, make it immediately useful for *Euclid*- and LSST-era systematics control, both as a survey-scale diagnostic and as a quantitative ingredient in end-to-end mitigation pipelines.

ACKNOWLEDGEMENTS

The authors thank the anonymous referee for their invaluable advice on the improvement of the presentation of the current work. KT was supported by the STFC grant ST/W000903/1 and by the European Structural and Investment Fund. KT further acknowledges support by the European Union's Horizon Europe research and innovation program under the Marie Skłodowska-Curie CO-FUND Postdoctoral Programme grant agreement No.101081355-SMASH and from the Republic of Slovenia and the European Union from the European Regional Development Fund. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. DA acknowledges support from the Beecroft Trust. We made extensive use of computational resources at the University of Oxford Department of Physics, funded by the John Fell Oxford University Press Research Fund.

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

Albrecht A. et al., 2006, preprint (astro-ph/0609591)
 Alonso D., Sanchez J., Slosar A., 2019, *MNRAS*, 484, 4127
 Berlfein F., Mandelbaum R., Dodelson S., Schafer C., 2024, *MNRAS*, 531, 4954

Blazek J. A., MacCrann N., Troxel M., Fang X., 2019, *Phys. Rev. D*, 100, 103506
 Bridle S., King L., 2007, *New J. Phys.*, 9, 444
 Chisari N. E. et al., 2019a, *Open J. Astrophys.*, 2
 Chisari N. E. et al., 2019b, *ApJS*, 242, 2
 Cornish T., Alonso D., Leistedt B., Wolz K., 2026, *MNRAS*, 547, 2026
 Cragg C., Duncan C. A. J., Miller L., Alonso D., 2023, *MNRAS*, 518, 4909
 Dvorkin C., Smith K. M., 2009, *Phys. Rev. D*, 79, 043003
 Euclid Collaboration, 2025, *A&A*, 697, A1
 Gatti M. et al., 2021, *MNRAS*, 504, 4312
 Harnois-Déraps J., Giblin B., Joachimi B., 2019, *A&A*, 631, A160
 Heymans C. et al., 2006, *MNRAS*, 368, 1323
 Hildebrandt H. et al., 2020, *A&A*, 633, A69
 Hu W., 2000, *Phys. Rev. D*, 62, 043007
 Huff E., Mandelbaum R., 2017, preprint (arXiv:1702.02600)
 Huterer D., Takada M., Bernstein G., Jain B., 2006, *MNRAS*, 366, 101
 Jain B., Taylor A., 2003, *Phys. Rev. Lett.*, 91, 141302
 Jarvis M. et al., 2016, *MNRAS*, 460, 2245
 Joachimi B. et al., 2015, *Space Sci. Rev.*, 193, 1
 Kaiser N., 1992, *ApJ*, 388, 272
 Kitching T. D., Paykari P., Hoekstra H., Cropper M., 2019, *Open J. Astrophys.*, 2, 5
 LSST Dark Energy Science Collaboration, 2012, preprint (arXiv:1211.0310)
 Laureijs R. et al., 2011, Euclid Definition Study Report, preprint (arXiv:1110.3193)
 Lewis A., Challinor A., 2006, *Phys. Rep.*, 429, 1
 Li S.-S. et al., 2023, *A&A*, 670, A100
 Limber D. N., 1953, *ApJ*, 117, 134
 MacCrann N. et al., 2022, *MNRAS*, 509, 3371
 Mandelbaum R., 2018, *ARA&A*, 56, 393
 Maniyar A. S., Ali-Haïmoud Y., Carron J., Lewis A., Madhavacheril M. S., 2021, *Phys. Rev. D*, 103, 083524
 Massey R. et al., 2013, *MNRAS*, 429, 661
 Mead A. J., Brieden S., Tröster T., Heymans C., 2021, *MNRAS*, 502, 1401
 Melchior P., Viola M., 2012, *MNRAS*, 424, 2757
 Paulin-Henriksson S., Amara A., Voigt L., Refregier A., Bridle S. L., 2008, *A&A*, 484, 67
 Refregier A., Kacprzak T., Amara A., Bridle S., Rowe B., 2012, *MNRAS*, 425, 1951
 Rowe B., 2010, *MNRAS*, 404, 350
 Saga S., Yamauchi D., Ichiki K., 2015, *Phys. Rev. D*, 92, 063533
 Samuroff S. et al., 2018, *MNRAS*, 475, 4524
 Schneider P., 2006, *Weak Gravitational Lensing*. Springer, Heidelberg, Berlin, p. 269
 Schneider P., 2016, *A&A*, 592, L6
 Schneider P., Kilbinger M., 2006, *A&A*, 462, 841
 Sheldon E. S., Huff E. M., 2017, *ApJ*, 841, 24
 Smith R. E. et al., 2003, *MNRAS*, 341, 1311
 Takahashi R., Sato M., Nishimichi T., Taruya A., Oguri M., 2012, *ApJ*, 761, 152
 Taylor A. N., Kitching T. D., 2018, *MNRAS*, 477, 3397
 Taylor A. N., Kitching T. D., Bacon D. J., Heavens A. F., 2007, *MNRAS*, 374, 1377
 The LSST Dark Energy Science Collaboration, 2018, preprint (arXiv:1809.01669)
 van Uitert E., Schneider P., 2016, *A&A*, 595, A93
 Wick G. C., 1950, *Phys. Rev.*, 80, 268
 Williams J., McCallum N., Rotti A., Thomas D. B., Battye R., Brown M. L., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 016

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