

A blinding solution for inference from astronomical data

Elena Sellentin[★]

Leiden Observatory, Leiden University, Huygens Laboratory, Niels Bohrweg 2, NL-2333 CA Leiden, the Netherlands

Accepted 2020 January 1. Received 2019 December 19; in original form 2019 October 22

ABSTRACT

This paper presents a joint blinding and deblinding strategy for inference of physical laws from astronomical data. The strategy allows for up to three blinding stages, where the data may be blinded, the computations of theoretical physics may be blinded, and –assuming Gaussianly distributed data – the covariance matrix may be blinded. We found covariance blinding to be particularly effective, as it enables the blinder to determine close to exactly where the blinded posterior will peak. Accordingly, we present an algorithm which induces posterior shifts in predetermined directions by hiding untraceable biases in a covariance matrix. The associated deblinding takes the form of a numerically lightweight post-processing step, where the blinded posterior is multiplied with deblinding weights. We illustrate the blinding strategy for cosmic shear from KiDS-450, and show that even though there is no direct evidence of the KiDS-450 covariance matrix being biased, the famous cosmic shear tension with Planck could easily be induced by a mischaracterization of correlations between ξ_{-} at the highest redshift and all lower redshifts. The blinding algorithm illustrates the increasing importance of accurate uncertainty assessment in astronomical inferences, as otherwise involuntary blinding through biases occurs.

Key words: methods: data analysis – methods: statistical – cosmology: observations.

1 INTRODUCTION

Astronomy provides many data sets that enable us to infer physical laws on energy-, size-, and time-scales that are inaccessible to Earth-bound laboratories. Of particular interest to inference are astronomical observations which are often so rare that only a few – or even no – comparable observations are expected in a human lifetime. Amongst such unique data sets rank e.g. observations of our Milky Way (Lindegren et al. 2016; Gaia Collaboration 2018) or across the entire cosmos (Laureijs et al. 2011; Jain et al. 2015; Planck Collaboration VI 2018), where by definition no second data set like the original will ever exist.

Many astronomers therefore remember e.g. peculiar stars in the Milky Way, the cold spot in the cosmic microwave background, or other directly visible features in the data, such as the ‘Great Wall’ (Einasto et al. 2011) in galaxy surveys. Inference, on the other hand, is the attempt to go beyond the directly visible, and assigns a credibility to a hidden, unobservable quantity, such as a model of theoretical physics, or the values of free physical parameters.

Due to remembering data features on the one hand, but also due to iterative data cleaning to handle unexpected systematics, the physics inferred from astronomical data is sometimes regarded with skepticism. Post-dictions and retrospectively adapted models rank amongst frequently encountered points of critique. Iterative and

often subconscious tampering with the analysis constitute further elements of concern, as they might lead to the confirmation of prior held beliefs (Croft & Dailey 2011; Seehars et al. 2016). The wish to avoid that such biases impact the inferred physics is therefore becoming widespread, and can be addressed by conducting analyses blindly.

Blinding strategies extend an analysis such that it becomes impossible to predict which physics will be discovered from it. Blinding can be extremely difficult to achieve. Indeed, most blinding techniques operate exclusively on the data, and therefore either interfere with the inevitable necessity to make informed decisions when cleaning astronomical data – or are easily spotted as counterfeits.

Thus, the aim of this paper is to establish a blinding technique which meets the needs of astronomical inference. The raw data are left untouched, but the transition to science-ready data may be blinded if this does not cause substantial costs. Additionally, the computations of physical models may also be blinded. Crucially though, we find blinding of a likelihood by biasing a covariance matrix provides a very powerful third tier. It enables the blinder to specify nearly perfectly where the blinded posterior peaks, whilst causing negligible numerical costs.

In Section 2, we detail why astronomical inference both requires and enables special blinding strategies. The up to three-stage blinding algorithm is described in Section 2.1. In Sections 3, 4, and 5 we develop the algorithm for blinding a covariance matrix. The associated deblinding is described in Sections 6 and 7. Throughout

[★] E-mail: elena.sellentin@posteo.de

the paper, we demonstrate the algorithm on the data of Hildebrandt et al. (2017), which is a cosmic shear analysis of the Kilo-Degree Survey (KiDS) further described in Kuijken et al. (2015), Fenech Conti et al. (2017). The KiDS data¹ are processed by THELI (Erben et al. 2013) and Astro-WISE (Begeman et al. 2013; de Jong et al. 2015). Shears were measured with LENSFIT (Miller et al. 2007), and photometric redshifts were obtained from PSF-matched photometry and calibrated using external overlapping spectroscopic surveys (see Hildebrandt et al. 2017). The essentials of cosmic shear are summarized in Appendix A. Appendix B summarizes the implications of this paper’s findings for the KiDS-450 survey and its reported mild tension with Planck (Planck Collaboration VI 2018).

2 WHY ASTRONOMY HAS SPECIAL BLINDING NEEDS

We briefly review blinding techniques used in neighbouring fields, and contrast them with astronomy.

Popular blinding techniques in particle physics include masking of a ‘signal region’, as e.g. carried out by CMS and ATLAS during the discovery of the Higgs boson (Aad et al. 2012; Chatrchyan et al. 2012). Also the ANTARES neutrino telescope blinded a spatial signal region when studying excessive neutrino flux from the galactic ridge (Adrián-Martínez et al. 2016). The Large Underground Xenon experiment LUX (Akerib et al. 2017) and the gravitational wave facility LIGO (Abbott et al. 2016) instead injected artificial signals into the detector, an approach known as ‘salting’.

Common to these blinding techniques is their direct operation on the raw data, either by masking or by imitating signals. This is effective when new physics leaves visible imprints in a particle physics experiment: Novel particles may cause unconventional tracks in a detector, and decaying new particles will cause visible peaks above the detector’s background, called ‘resonances’. Salting and blinding of a signal region effectively masks the presence of such features.

In contrast, the majority of astronomical data sets do not exhibit a split into a signal- and a background region, thereby having nothing whose hiding would be of any advantage. The lack of a signal region is e.g. illustrated by Fig. 1, which depicts the cosmic shear data set from Hildebrandt et al. (2017), from which the best-fitting parameters for the cosmic dark matter density Ω_m was determined to be $\Omega_m = 0.2$, and the normalization of the matter power spectrum, σ_8 , which essentially measures how clumpily matter is distributed, was determined to be $\sigma_8 = 0.838$, keeping all other cosmological parameters fixed (Sellentin & Starck 2019). Fig. 1 illustrates that there is no clear signal region from which the value of these parameters could have been read off. It simply depicts two correlation functions, ξ_+ and ξ_- , as a function of angular separation ϑ .

This illustrates aptly how indirect astronomical constraints of physics often are. In fact, they are often true ‘inference problems’: Given the data, one wishes to infer by definition *unobservable* quantities, namely the physical parameters. Only upon adoption of a likelihood can information on the parameters be distilled from the data.

In contrast to their inferences being extremely indirect, most astronomical raw data are easy to visualize and memorize. Salting,

i.e. adding artificial signals to the data, is hence correctly disregarded in many astronomical disciplines due to artificial additions or omissions being easily spotted.

Astronomical and particle physics measurements are therefore close to opposites of each other, and a astronomical blinding technique needs to reflect this. We therefore describe an up to three-stage blinding strategy which may blind the likelihood, the theoretical predictions, and one of the last stages in the long and weary transition from raw data to science-ready data.

An example of this strategy is seen in Fig. 2, which uses the same data as Fig. 1, only that the covariance matrix contained biases to shift the posterior. Overplotted is the blinded best-fit of $\Omega_m = 0.274$, $\sigma_8 = 0.754$, which differs by 1.5 posterior standard deviations from the actual best fit. This corresponds to a shift towards higher $S_8 = \sigma_8 \sqrt{\Omega_m/0.3}$, in a direction perpendicular to the degeneracy between Ω_m and σ_8 . Telling Figs 2 and 1 apart is extremely difficult, which underlines the power of the blinding algorithm now to be presented.

2.1 Three-stage blinding setup

Although the below arguments are easily extended, we now specialize to Gaussianly distributed data. The likelihood is then Gaussian if the covariance matrix is known. If the covariance matrix is estimated from simulations, then the likelihood is the t-distribution of Sellentin & Heavens (2016), Sellentin & Heavens (2017) instead. These two cases are the most encountered likelihoods.

Parameter inference then requires three ingredients: the correct science-ready data \mathbf{x} , the correct covariance matrix Σ (either estimated or analytically computed), and the capability to compute the theoretical mean $\mu(\theta)$ at full accuracy, where θ are the parameters of interest. It is thus natural to introduce three blinders, and we denote (potentially) blinded quantities with a breve, \check{x} , $\check{\Sigma}$, $\check{\mu}(\theta)$.

We assume that to the general researcher, all three quantities \check{x} , $\check{\Sigma}$, $\check{\mu}(\theta)$ appear blinded. To each of the three blinders, two out of three quantities appear blinded, and the third is the one whose blinding is their task.

The data-blinder requires access to the process of distilling science-ready data \mathbf{x} from the raw data. The theory-blinder requires access to the software computing $\mu(\theta)$. Essentially any analysis at some point uses a look-up table, an emulator, or fixed nuisance parameters, all of which can be biased. The theory-blinder may decide to restrict the other blinders’ capabilities to compute $\check{\mu}(\theta)$ – this may be advantageous if blinding is partially assigned to external researchers.

Finally, we assume the likelihood-blinder has exclusive access to the code which computes the covariance matrix. This code is of no direct use as long as there is no data vector yet. The likelihood-blinder is then given \check{x} , and computes its true covariance matrix. Additionally, the likelihood-blinder is given access to the potentially restricted computational facilities to generate $\check{\mu}(\theta)$. This blinder then uses the below algorithm to generate a biased covariance matrix, which is used in the evaluation of the blinded posterior.

Three-stage blinding also has to consider, whether the covariance matrix can be directly estimated from the data (e.g. through bootstrapping), or whether the covariance matrix is independently modelled (e.g. analytically). In the rare case that the covariance matrix can be estimated from the astronomical data, then blinding \check{x} automatically also blinds $\check{\Sigma}$, as it is a quadratic function of the blinded data. In the usual case of data and covariance matrix being independently prepared, both can also independently be blinded.

¹The public data products are available at <http://kids.strw.leidenuniv.nl/cosmicshear2016.php>

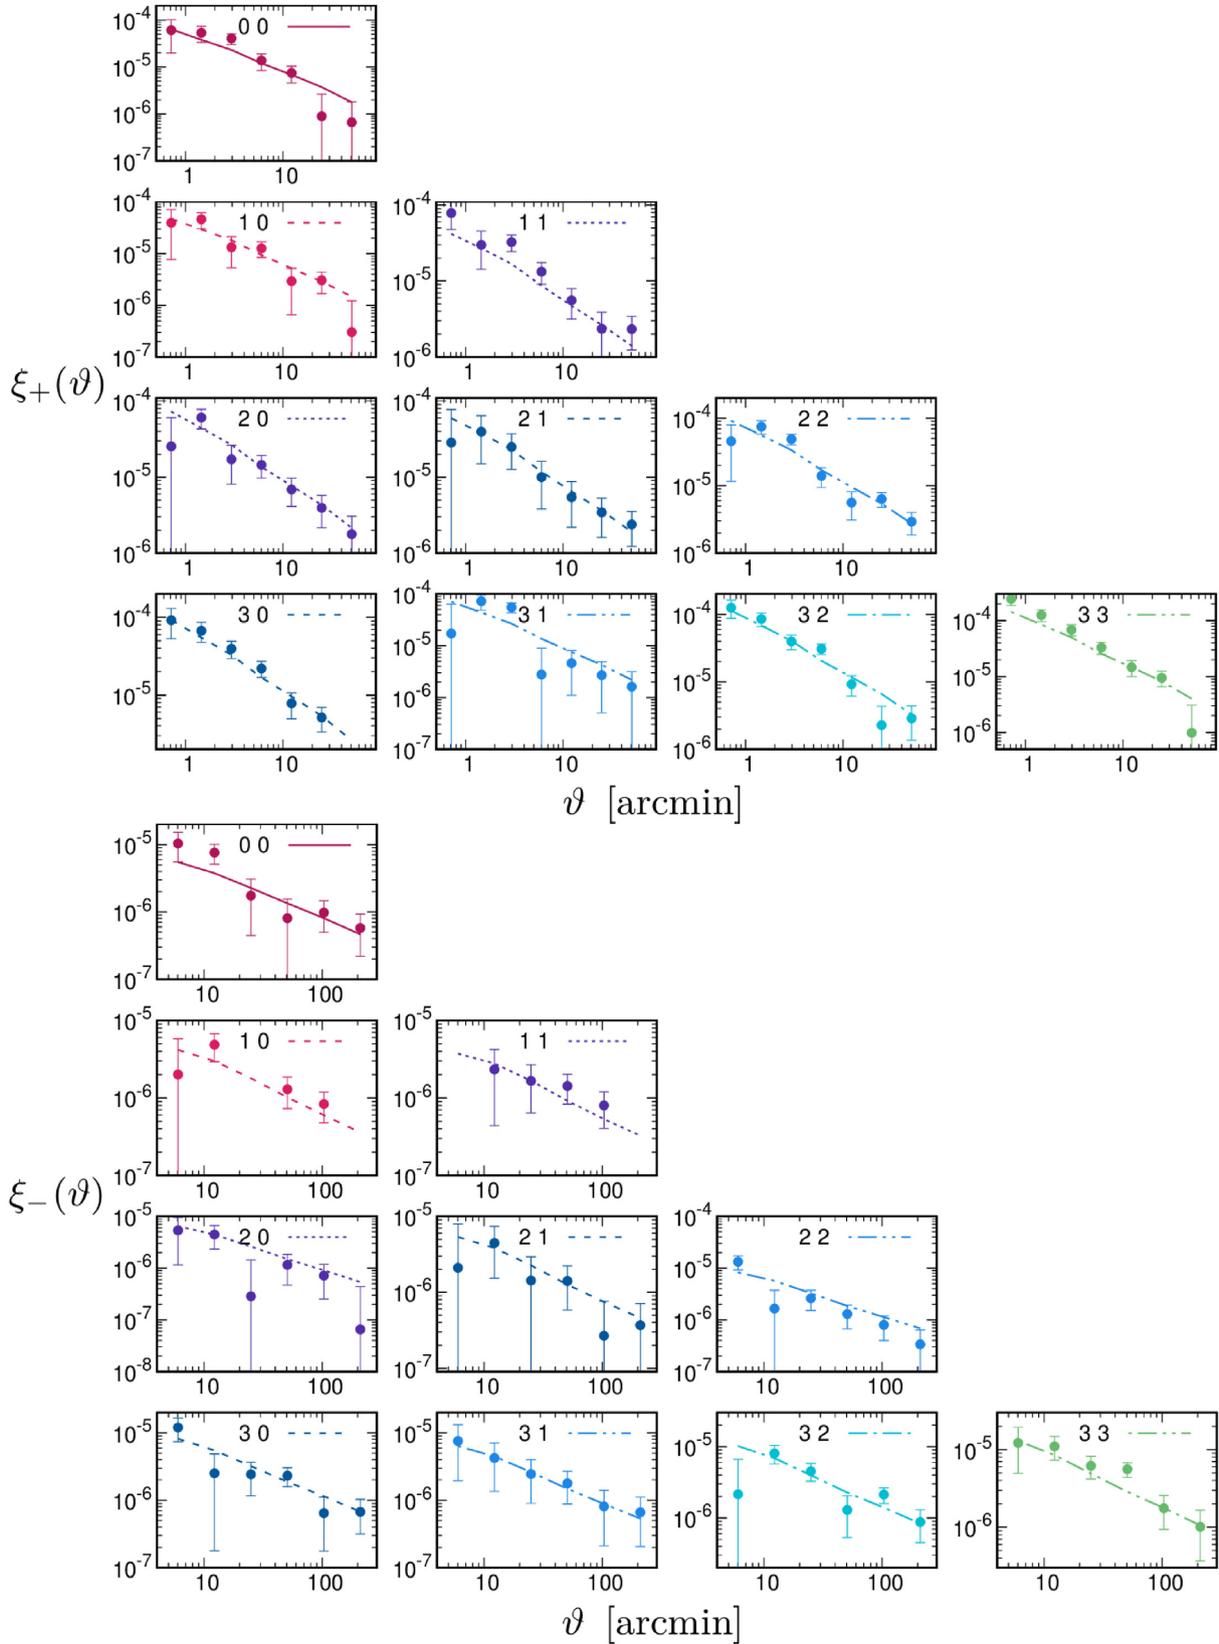


Figure 1. Example of an unblinded analysis: Joint plot of the original KiDS-450 data (Hildebrandt et al. 2017), the original error bars, and the original best-fitting theory curve, computed with $\Omega_m = 0.2$ and $\sigma_8 = 0.838$. The upper triangle depicts the cosmic shear correlation function $\xi_+(\vartheta)$ over all redshift bin combinations as labelled; the lower triangle depicts the correlation function $\xi_-(\vartheta)$.

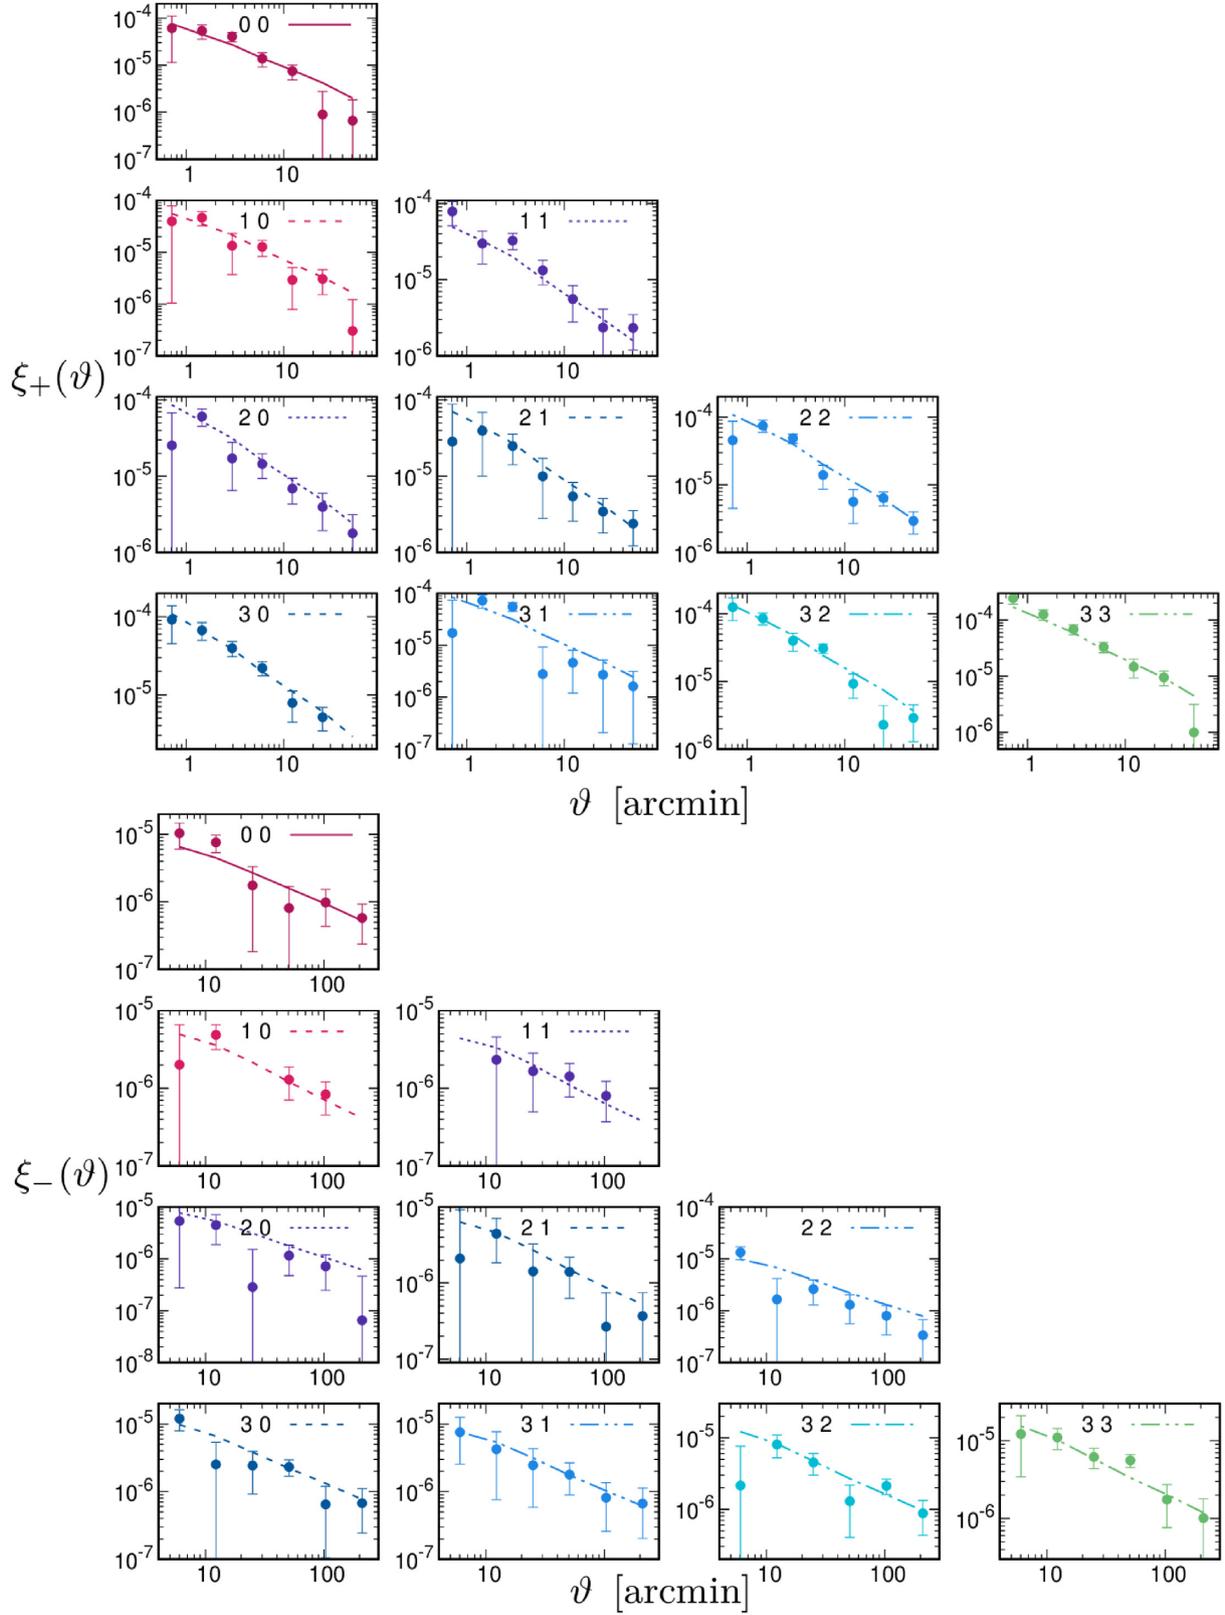


Figure 2. Example of a blinded analysis: The data are the same as in Fig. 1, but the error bars differ, as the analysis was blinded by biasing the covariance matrix to enforce a shift of the posterior by about 1.5 posterior standard deviations in S_8 . The now plotted theory curve is the peak of the blinded posterior, which lies at $\Omega_m = 0.2748$, $\sigma_8 = 0.7548$. Comparison with Fig. 1 illustrates that ‘fitting by eye’ is impossible and that it is extremely difficult to tell the blinded and unblinded analysis apart.

The blinding of \check{x} and $\check{\mu}(\theta)$ will be highly field specific. Many disciplines may even chose not to blind x . Therefore, this paper now focuses on an algorithm to blind Σ . The (potentially) blinded data and theory will then be re-encountered when jointly deblinding in Section 7.

3 BLINDING PREPARATIONS

3.1 Why blinded covariance matrices shift posteriors

Having specialized to a Gaussian or t -distribution likelihood, blinding the likelihood is akin to blinding a covariance matrix. We thus have to explain why biases in a covariance matrix shift posteriors. This was partially discussed in Sellentin & Starck (2019), which we here extend by showing how best-fitting parameters of a Gaussian likelihood directly depend on the covariance matrix.

We adopt a d -dimensional data vector x , and indicate expectation values by angular brackets $\langle \cdot \rangle$. The mean is $\langle x \rangle = \mu = (\mu_1, \dots, \mu_d)$, which is a function of a p -dimensional parameter vector $\theta = (\theta_1, \dots, \theta_p)$. The parameters are of physical significance and shall be inferred.

In general, $\mu(\theta)$ will depend non-linearly on its parameters, but to illustrate the biasing effect of covariance matrices we Taylor expand around an initial parameter point θ_1 . To this end, we introduce a non-square matrix X whose j th column shall be the derivative of the mean with respect to the j th parameter,

$$X = \begin{pmatrix} \frac{\partial \mu}{\partial \theta_1} & \dots & \frac{\partial \mu}{\partial \theta_p} \end{pmatrix}. \quad (1)$$

The matrix X is thus $p \times d$ dimensional and its ij th element is given by

$$X_{ij} = \frac{\partial \mu_i}{\partial \theta_j}. \quad (2)$$

This matrix can be computed with any conventional Fisher matrix forecasting code (Tegmark, Taylor & Heavens 1997; Lesgourgues 2011; Sellentin, Quartin & Amendola 2014; Sellentin 2015). The linearized mean is then

$$\mu_{\text{lin}}(\theta) = \mu(\theta_1) + X(\theta - \theta_1), \quad (3)$$

which replaces the non-linear dependence on θ by a linear dependence. In order for equation (3) to hold, the derivatives in X must be evaluated at θ_1 . The aim is to determine the best-fitting parameters $\hat{\theta}$, which maximize the log-likelihood. We adopt a Gaussian posterior with parameter-independent covariance matrix Σ

$$\mathcal{P}(\theta|x, \Sigma) \propto \exp\left(-\frac{1}{2} [x - \mu(\theta)]^\top \Sigma^{-1} [x - \mu(\theta)]\right), \quad (4)$$

where the superscript \top denotes transposition. Replacing $\mu(\theta)$ with $\mu_{\text{lin}}(\theta)$ it follows that the linearized best-fitting parameters must solve

$$\frac{\partial}{\partial \theta} [x - \mu_{\text{lin}}(\theta)]^\top \Sigma^{-1} [x - \mu_{\text{lin}}(\theta)] = 0. \quad (5)$$

Using the relation

$$\frac{\partial}{\partial s} (r - As)^\top \Omega (r - As) = -2A^\top \Omega (r - As), \quad (6)$$

for vectors s, r , and matrices A, Ω of matching dimensions, the solution to equation (5) yields the linearized best-fitting parameters

$$\hat{\theta}_{\text{lin}} = (X^\top \Sigma^{-1} X)^{-1} X^\top \Sigma^{-1} [x - \mu(\theta_1) + X\theta_1]. \quad (7)$$

This illustrates that the position of the best-fit depends on the covariance matrix in a dual manner: the term

$$\Sigma^{-1} [x - \mu(\theta_1) + X\theta_1], \quad (8)$$

inverse-variance weights the distance between the data and the mean, thereby preferring means which match the data in units of the covariance. The term

$$(X^\top \Sigma^{-1} X)^{-1}, \quad (9)$$

describes a compression, due to X being non-square. This term compresses the information from fitting in data space into the lower-dimensional parameter space.

As the best-fitting parameters depend in this dual manner on the covariance matrix, it directly follows that a bias in the covariance matrix will translate into a shift of the best-fitting parameters – this opens the possibility to blind by changing the covariance matrix.

In fact, for the purpose of blinding, the compression term of equation (9) has a further appeal: due to the compression, many different covariance matrices will lead to the same shift in parameters, due to this being a highly underdetermined system. This allows us to set side-constraints, for example that the biased covariance matrix not only induces a specified shift from best-fitting parameters $\hat{\theta}$ to blinded best-fitting parameters $\check{\theta}$, but at the same time also maintains e.g. all its original variances, and (as a further example) its determinant, or the sign of all its correlation coefficients.

The upcoming sections therefore describe blinding by constructing one (or multiple) blinded covariance matrices. To linear order, the blinded best-fit will then lie at

$$\check{\theta}_{\text{lin}} = (X^\top \check{\Sigma}^{-1} X)^{-1} X^\top \check{\Sigma}^{-1} [x - \mu(\theta_1) + X\theta_1], \quad (10)$$

which will be extended to a fully non-linear inference with sampling in Section 6.

3.2 Stages of the covariance blinding algorithm

The upcoming blinding algorithm passes through different stages. At first, the algorithm assists the blinder in determining sensible magnitudes for the parameter shift (Section 3.3). This is non-trivial, as the posterior width is not yet known. The algorithm then transforms the data from their physical units on to a representation which is natural for statistical manipulations (Section 3.4). Subsequently, Section 4 describes how to adapt the covariance matrix to induce the intended posterior shift.

Theoretically, the algorithm could then stop. However, wilful deblinding might at that stage still be possible, and the entirety of Section 5 is thus devoted to making wilful deblinding impossible by using an encryption algorithm and allowing the specification of side constraints to be met. Section 5.6 computes the shift of parameters to determine whether blinding succeeded. Finally, the physical units are restored, and the output is a purposefully biased covariance matrix which leads to the requested bias in physical parameters, with the origin of the bias being untraceable due to the intermediate encryption and side constraints.

3.3 Finding sensible target parameters for blinding

At the outset of the blinding strategy, the blinder has to pick an origin θ_0 of the posterior shift, and target parameters θ_1 which are preferred after shifting. Typically, shifting the yet unknown posterior by two to three standard deviations is the sought aim. Having never computed the full posterior, its standard deviations

Table 1. $\Delta\chi^2$ -values where the 90 percent credibility contour lies above the minimum χ^2 , as a function of number of parameters. If the blinder chooses target parameters θ_i which differ from the initial parameters θ_o by more than $\Delta\chi^2$ shown here, then a posterior shift of more than three posterior standard deviations will ensue. This is not advised.

Parameters	6	8	10	12	14	16	18	20	22	24	26	28	30	40	50	100
$\Delta\chi^2$ (90 per cent)	10.6	13.3	15.9	18.5	21.0	23.5	25.9	28.4	30.8	33.2	35.5	37.9	40.2	51.8	63.1	118.5

are however not known yet, so we effectively wish to shift by the multiple of a yet unknown quantity. We thus require an abstract prediction of the shift's magnitude.

To this aim, we exploit scaling relations: For a multivariate Gaussian posterior, the best-fit will occur at the minimal chisquared, χ_{\min}^2 . The associated 90 per cent credibility contour can then be chosen to run along an isocontour which lies by $\Delta\chi^2$ above this minimum. The more parameters are estimated, the larger $\Delta\chi^2$ has to be, since more parameters will widen up the joint posterior. The values of $\Delta\chi^2$ where the 90 per cent credibility contour lies above the minimum are known (they are integrals of χ^2 -distributions), and we tabulate them as a function of free parameters in Table 1.

The blinder thus has to know how many parameters are to be inferred, and has to pick initial parameters θ_o and target parameters θ_i which differ approximately by $\Delta\chi^2$ -values as given in Table 1.

If the blinder has no physical intuition for which parameter values θ_o are likely to be preferred by the data, then the linearized best-fitting estimator of equation (7) can be evaluated, which will yield parameter values $\hat{\theta}_{\text{lin}}$ close to the full non-linear best fit $\hat{\theta}$.

We found that the values θ_o do not need to be picked with too much care, as long as they are within about two posterior standard deviation from the true best-fit $\hat{\theta}$. We provide a public code² for the algorithm, which assists the blinder via equation (7) and Table 1 in finding sensible parameters θ_o for the shift origin, and for the shift destination θ_i .

3.4 Data transformation

The general astronomical data set will come in ‘natural’ units, which might depend on estimator choice or the adopted (physical) units of measurements, such as parsec or mega-parsec. To gain generality, we thus standardize the data set.

Let Σ_{ii} be the i th diagonal element of the covariance matrix, then the conditional standard deviations σ_i are

$$\sigma_i = \sqrt{\Sigma_{ii}}. \quad (11)$$

The standard deviations have the same dimensions as the data points themselves. We therefore introduce the standardized variables

$$u_i = \frac{x_i}{\sigma_i}, \quad v_i = \frac{\mu_i}{\sigma_i}, \quad (12)$$

such that \mathbf{u} is the data vector expressed as multiples of its former standard deviation, and \mathbf{v} is the standardized theoretical mean, which still depends on parameters θ . Dividing out the standard deviations at this point allows us to restore them post-blinding. This is required for situations where variances can be easily remembered, and must thus remain unchanged during blinding.

The covariance matrix of the standardized data is then the correlation matrix

$$\langle (\mathbf{u} - \langle \mathbf{u} \rangle)(\mathbf{u} - \langle \mathbf{u} \rangle)^\top \rangle = \mathbf{C}. \quad (13)$$

The posterior of the parameters given the standardized data is then a Gaussian, with the inverse correlation matrix as precision matrix

$$\mathcal{P}(\theta|\mathbf{u}) \propto \exp\left(-\frac{1}{2}[\mathbf{u} - \mathbf{v}(\theta)]^\top \mathbf{C}^{-1}[\mathbf{u} - \mathbf{v}(\theta)]\right). \quad (14)$$

Equations (14) and (4) are exactly the same posterior, only once expressed in units natural for statistics (equation 14), and once in units natural for the astronomer (equation 4).

Finally, we Cholesky decompose the inverse correlation matrix

$$\mathbf{C}^{-1} = \mathbf{L}\mathbf{L}^\top, \quad (15)$$

where \mathbf{L} is a lower triangular matrix, and its transpose is upper-triangular.

Technically, $\mathbf{L}^\top \mathbf{u}$ is a whitening transform, with the astronomical implication being that \mathbf{L} causes the often rich structure of astronomical data. Factorizing it out at this point enables us to multiply it back in later, whereupon deceptively naturally looking covariance matrices are restored.

4 MAIN BLINDING ALGORITHM

At this stage, the data have been transformed into a easily manageable representation, and the blinder has decided which shift of the posterior shall be induced. We therefore now lay out the mathematics of how to change a covariance matrix such that it shifts the posterior in a wanted direction.

To clarify the aim, we depict in Fig. 3 three correlation matrices as produced with the upcoming algorithm. One of these is the original correlation matrix of Hildebrandt et al. (2017), and two are matrices biased by our algorithm. When used in a Gaussian likelihood, the three matrices lead to the three posteriors of Fig. 4. Obviously, maximally one of the three posteriors can be correct. To underline the point of how easily such shifts are hidden in a covariance matrix, we refrain from revealing which of the three matrices in Fig. 3 is the correct correlation matrix.³ The algorithm for blinding the covariance matrix is as follows.

4.1 Biasing the covariance matrix

The algorithm begins by translating the posterior, and side constraints are enforced later.

To achieve a translation, we have to adapt the χ^2 -surface. Expressed by the Cholesky decomposition, the chisquare surface is

$$\chi^2(\theta, \mathbf{u}, \mathbf{L}) = [\mathbf{u} - \mathbf{v}(\theta)]^\top \mathbf{L}\mathbf{L}^\top [\mathbf{u} - \mathbf{v}(\theta)]. \quad (16)$$

We now introduce the biased inverse correlation matrix $\check{\mathbf{C}}^{-1}$, for which we chose the ansatz

$$\check{\mathbf{C}}^{-1} = \mathbf{L}\mathbf{B}\mathbf{B}^\top \mathbf{L}^\top, \quad (17)$$

³In the spirit of reproducible research, all plots in this paper can be reproduced with our public code and the equally public likelihood of Hildebrandt et al. (2017).

²<https://github.com/elenasellentini/StellarBlind>

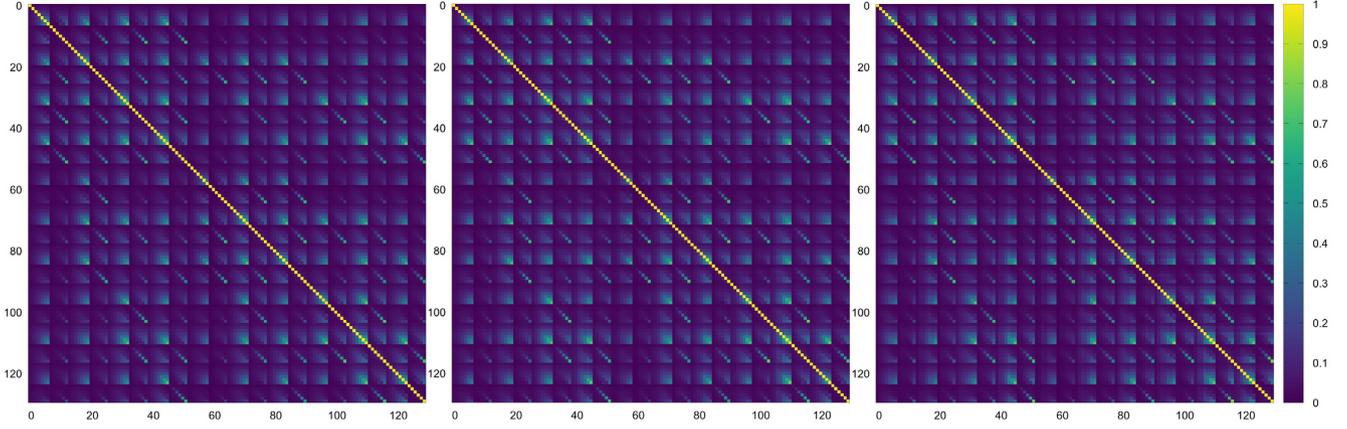


Figure 3. Three correlation matrices, of which one is the original correlation matrix of Hildebrandt et al. (2017), and two were constructed such that the posterior peaks at pre-defined parameter values. These matrices result in the three posteriors of Fig. 4. The extreme difficulty of identifying the correct correlation matrix illustrates the power of the here presented blinding algorithm.

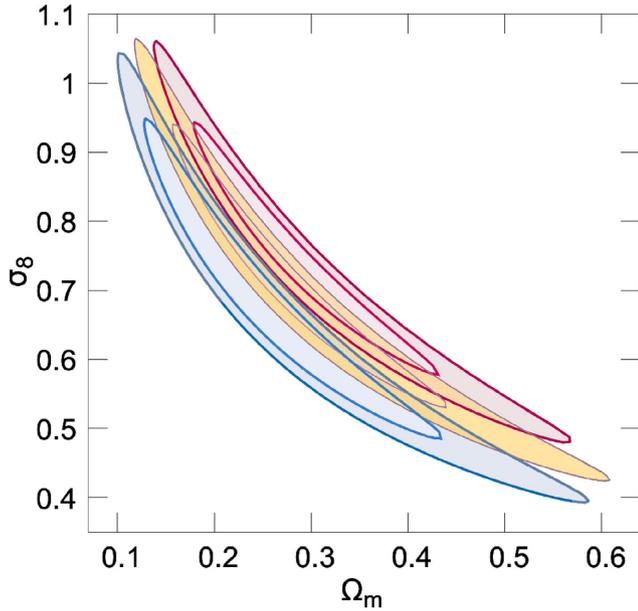


Figure 4. Blinded analysis of the KiDS-450 data vector, once with the true covariance matrix (orange), and once with two blinded covariance matrices, whose corresponding correlation matrices are depicted in Fig. 3.

where $\mathbf{B} = \text{diag}(b_{11}, b_{22}, \dots, b_{dd})$ is a diagonal matrix which causes the translation. The matrix \mathbf{B} can be interpreted as artificial signals, which are hidden in the correlation matrix. equation (17) also implies that this blinding technique requires a dense correlation matrix: If the original correlation matrix were diagonal, then \mathbf{L} were diagonal as well, and the blinding would then easily be discovered.

To compute \mathbf{B} such that the posterior shifts from θ_o to θ_t , we follow the arguments of Sellentin & Starck (2019) and demand $\chi^2(\theta_o, \mathbf{u}, \mathbf{C}) = \chi^2(\theta_t, \mathbf{u}, \tilde{\mathbf{C}})$ which requires

$$\begin{aligned} [\mathbf{u} - \mathbf{v}(\theta_o)]^\top \mathbf{L} \mathbf{L}^\top [\mathbf{u} - \mathbf{v}(\theta_o)] \\ = [\mathbf{u} - \mathbf{v}(\theta_t)]^\top \mathbf{L} \mathbf{B} \mathbf{B}^\top \mathbf{L}^\top [\mathbf{u} - \mathbf{v}(\theta_t)]. \end{aligned} \quad (18)$$

This constraint expresses that the χ^2 -value prior to blinding at the origin shall equal the χ^2 -value post-blinding at the parameters targeted by blinding. As \mathbf{B} is diagonal, it can now easily be

computed. We introduce the vectors

$$\begin{aligned} \mathbf{L}^\top [\mathbf{u} - \mathbf{v}(\theta_o)] &= \mathbf{e}, \\ \mathbf{L}^\top [\mathbf{u} - \mathbf{v}(\theta_t)] &= \tilde{\mathbf{e}}, \end{aligned} \quad (19)$$

and the diagonal elements of \mathbf{B} are then

$$b_{ii} = \frac{e_i}{\tilde{e}_i}. \quad (20)$$

Theoretically, the blinding algorithm could stop here: one could now directly compute $\tilde{\mathbf{C}} = (\mathbf{L} \mathbf{B} \mathbf{B}^\top \mathbf{L}^\top)^{-1}$, which is the blinded correlation matrix. The variances which were divided out in equation (12) would need to be multiplied back in, and the result would be a blinded covariance matrix which shifts the posterior.

However, at this point human deblinding might still be possible as the above is a deterministic calculation. There might exist situations where enough intuition about the true covariance matrix can be gained in order to reverse-engineer which blinding parameters θ_t the blinder chose. The blinding could then be undone.

Additionally, because the determinant $|\mathbf{B}|$ was not enforced to be unity, we will have changed the determinant of the covariance matrix. This will change the size of the posterior.

We therefore regard the ansatz equation (17) with a diagonal \mathbf{B} only as a convenient starting point for the algorithm, and we will now exploit the fact that the constraint of equation (18) is strongly underdetermined: it sets a single constraint to solve for the elements of a (in general dense) $d \times d$ matrix. There thus exist infinitely many matrices \mathbf{B} to induce the wanted posterior translation, and in the following we use this freedom to adapt the blinded correlation matrix.

5 DISABLING ACCIDENTAL DEBLINDING

Section 4 biased a covariance matrix such that the posterior prefers a chosen set of parameters θ_t . The entirety of this section is devoted to making the blinding untraceable. As this requires us to employ random manipulations, this section ends by controlling whether the requested shift is still achieved, despite the manipulations.

5.1 Disabling recovery of the blinding parameters

We improve the quality of the algorithm by making it impossible to reconstruct the parameters chosen by the blinder. This requires us

to change the values b_{ii} , subject to still inducing the wanted shift. To this aim, we note that the magnitude of the b_{ii} will in general be too large, and it is preferable if the b_{ii} are as close to unity as only possible during blinding.

To compute how close to unity we can push the b_{ii} , we compute the average $\Delta\chi^2$ that the blinding must induce. Using the vectors from equation (19), the χ^2 at the target parameters is

$$\chi^2(\boldsymbol{\theta}_t, \mathbf{u}, \mathbf{C}) = \check{\boldsymbol{\epsilon}}^\top \check{\boldsymbol{\epsilon}}, \quad (21)$$

before blinding. This will be larger than the χ^2 after blinding, which is

$$\chi^2(\boldsymbol{\theta}_t, \mathbf{u}, \check{\mathbf{C}}) = \check{\boldsymbol{\epsilon}}^\top \mathbf{B} \mathbf{B}^\top \check{\boldsymbol{\epsilon}}. \quad (22)$$

The $\Delta\chi^2$ bridged during blinding is thus

$$\begin{aligned} \Delta\chi^2 &= \chi^2(\boldsymbol{\theta}_t, \mathbf{u}, \mathbf{C}) - \chi^2(\boldsymbol{\theta}_t, \mathbf{u}, \check{\mathbf{C}}) \\ &= \sum_{i=1}^d \check{\epsilon}_i^2 [1 - b_{ii}^2]. \end{aligned} \quad (23)$$

It is senseless to fit perfectly to one realization $\check{\boldsymbol{\epsilon}}$ of the data, and we thus average over multiple realizations. On average, we will have $\langle \check{\epsilon}_i^2 \rangle \approx 1$ because $\check{\boldsymbol{\epsilon}}$ is approximately a white vector, where the approximation is that blinding conditions on an incorrect mean, see equation (19).

As a consequence of $\check{\boldsymbol{\epsilon}}$ being approximately white it follows that all b_{ii} are approximately the same. Then, inserting $\langle \check{\epsilon}_i^2 \rangle = 1$ into equation (23), we have

$$\langle \Delta\chi^2 \rangle \approx d(1 - b_{ii}^2). \quad (24)$$

Solving for the elements b_{ii} , we have

$$|b_{ii}| = \sqrt{1 - \frac{\langle \Delta\chi^2 \rangle}{d}}. \quad (25)$$

This result has a highly intuitive interpretation: if blinding bridges a larger gap $\Delta\chi^2$, then the magnitude of the b_{ii} increases, implying larger biases are needed. On the other hand, if the dimension d of the data increases, then smaller elements b_{ii} suffice to still shift the posterior. For $d \rightarrow \infty$, we have $b_{ii} \rightarrow 1$, meaning even the tiniest changes in the covariance matrix will suffice to induce major shifts of the posterior. Accordingly, it is easier to blind large data sets.

For our blinding algorithm, we hence set

$$\forall i : \sqrt{1 - K} < b_{ii} < \sqrt{1 + K}. \quad (26)$$

These thresholds will deteriorate the goodness of fit, which we will compensate for later. The scalar K is

$$K = W \frac{\chi^2(\boldsymbol{\theta}_t, \mathbf{u}, \mathbf{C}) - \chi^2(\boldsymbol{\theta}_t, \mathbf{u}, \check{\mathbf{C}})}{d}, \quad (27)$$

where W has to be positive and acts as a tolerance: if $W < 1$, then blinding will not succeed, as the b_{ii} can vary insufficiently to induce the requested change in chisquared. If $W > 1$, then the b_{ii} can induce even larger changes in $\Delta\chi^2$ than is needed to shift the posterior. The latter allows further constraints to be enforced.

We now additionally demand the determinant $|\mathbf{B}|$ be unity. We thus rescale all elements $b_{ii} \rightarrow b_{ii} |\mathbf{B}|^{-1/d}$, where $|\mathbf{B}|$ is the determinant of \mathbf{B} . After rescaling, the determinant of \mathbf{B} is unity, which will not alter the determinant of the blinded correlation matrix. This will in turn leave the size of the posterior unchanged.

The elements b_{ii} have now been changed twice: first their magnitude has been subjected to upper and lower bounds, and then all elements were jointly changed multiplicatively. Both will induce degeneracies in attempts of recovering the original b_{ii} and

the blinder's $\boldsymbol{\theta}_t$, thereby making successful reverse engineering very unlikely. To disguise the biases even further, we shall now hide their presence through a series of random changes and matrix inversions.

5.2 Encrypting the correlation matrix

At this stage, we multiplied a biasing matrix to the Cholesky decomposition of the inverse correlation matrix

$$\check{\mathbf{L}}^\top = \mathbf{B}^\top \mathbf{L}^\top, \quad (28)$$

where \mathbf{B} has already been pre-optimized. The biased inverse correlation matrix would then be $\check{\mathbf{C}}^{-1} = \check{\mathbf{L}} \check{\mathbf{L}}^\top$, but in equation (27), we left margin to optimize further ($W > 1$), which we now use.

We now transit to random manipulations, such that as long as the random seeds of these random manipulations are unknown, they cannot be reversed. The below algorithm can be seen as an encryption, where the random seeds are the decryption keys.

We introduce the ‘symmetrized mean absolute percentage error’ (SMAPE) between two values v, v' , which is given by

$$\text{SMAPE}(v, v') = \frac{|v - v'|}{|v| + |v'|}. \quad (29)$$

This symmetrized per cental error yields zero if $v' = v$. Its upper bound of unity is reached in the limits $v \gg v'$ and $v' \gg v$. For $v' = 2v$ it yields 1/3, so it scales differently from the usual per cental error. The latter is a desired benefit: during blinding, elements of either the blinded or unblinded matrices can become zero, which the symmetrized error handles well, whereas the usual per cental error would return misleading zeros or infinite values.

To disguise the presence of the biases, we now shrink $\check{\mathbf{L}}$ element wise randomly towards the unbiased \mathbf{L} . The diagonal elements are left unchanged, in order to not change the determinant. Thus,

$$\begin{aligned} \forall i < j : & \text{ if } \text{SMAPE}(\check{L}_{ij}, L_{ij}) > S_{\text{inv}}, \\ & \text{ then } \check{L}_{ij} \sim \text{Uniform}(\check{L}_{ij}, L_{ij}). \end{aligned} \quad (30)$$

Here, $S_{\text{inv}} \in [0, 1]$ is a user defined threshold for the SMAPE. If the element wise SMAPE is exceeded, then the element of \check{L}_{ij} is reset to resemble the unbiased element L_{ij} more closely, by randomly drawing from a uniform distribution with upper and lower bound given by the two matrix elements. For this random draw, a seed has to be specified.

Given the element wise edited matrix $\check{\mathbf{L}}$, we compute $\check{\mathbf{C}}^{-1} = \check{\mathbf{L}} \check{\mathbf{L}}^\top$, and invert it to yield $\check{\mathbf{C}}$. This inversion redistributes the random changes in a manner impossible to predict by humans. To edit even further, we now Cholesky decompose the blinded and unblinded correlation matrices

$$\begin{aligned} \mathbf{C} &= \mathbf{R} \mathbf{R}^\top, \\ \check{\mathbf{C}} &= \check{\mathbf{R}} \check{\mathbf{R}}^\top, \end{aligned} \quad (31)$$

and we repeat the shrinking towards the original Cholesky decomposition

$$\begin{aligned} \forall i < j : & \text{ if } \text{SMAPE}(\check{R}_{ij}, R_{ij}) > S_{\text{corr}}, \\ & \text{ then } \check{R}_{ij} \sim \text{Uniform}(\check{R}_{ij}, R_{ij}). \end{aligned} \quad (32)$$

Here, S_{corr} is a threshold for the SMAPE which specifies the maximal changes the blinder tolerates in elements of the Cholesky decomposition of the correlation matrix.

The random resetting here conducted partially erases the desired biases, which is why it was important to leave room in equation (27). Should the partial erasing cause the blinding control of Section 5.6

to fail, then W in equation (27), S_{inv} and S_{corr} have to be adjusted. Crucially though, it is by this stage impossible to reconstruct how the biases entered the correlation matrix – two random editing processes with an intermediate inversion lie in the way.

5.3 Optimization of further side-constraints

The editing process of the biased correlation matrix has meanwhile progressed so far that the biased correlation matrix will strongly resemble the unblinded matrix. This makes it easy to optimize for final constraints which may be necessitated by the specifics of the research field – this step is highly important, as it is the only one to enforce that all physical constraints are met. Should logical inconsistencies remain and be discovered during the blinded analysis, then fractions of the blinding might become reversible.

For example, one might wish that the variances do not change during blinding, or that the correlation coefficients do not change sign, if there is a physical reason for positive or negative correlation between data points. Other fields might require the eigenvalues to be unchanged. If independent experiments are combined, then a block-diagonal structure of the covariance matrix ensues, which also ought to be preserved during blinding. Almost certainly one might wish that the formerly preferred parameter point θ_o is now disfavoured with at least a certain $\Delta\chi^2$ with respect to the target parameters θ_t .

If any such constraints has to be perfectly fulfilled, then the blinder should enforce it directly, e.g. by resetting the variances, and transit to controlling the success of the blinding algorithm in Section 5.6. But in general, enforcing additional constraints without major care may corrupt the correlation matrix. For example, resetting variances may lead to a non-positive definite matrix. Instead of implementing any constraints by brute-force, we rather advocate the following stochastic gradient descent algorithm, which operates on Cholesky decompositions instead.

We define a loss function F which is the sum over all constraints, such as

$$F = \text{SMAPE}(\check{L}, L) + \text{SMAPE}(\check{R}, R) + \left[\chi^2(\theta_o, \mathbf{u}, \check{C}) - \chi_{\text{req}}^2(\theta_o) \right]^2 + \left[\chi^2(\theta_t, \mathbf{u}, \check{C}) - \chi_{\text{req}}^2(\theta_t) \right]^2 + \sum_i (\check{C}_{ii} - C_{ii})^2, \quad (33)$$

where we define the SMAPE of a matrix to be taken element wise, χ^2 -values without subscript are those achieved when using the current iteration's matrix, and χ_{req}^2 with subscript are numbers which are the blinder's requested values at these parameter points. Omitting or adding further constraints to the loss function is possible until an overconstrained system is reached. This also implies that if the data set is small, and the covariance matrix is additionally known to be diagonal, then blinding is likely to leave visible traces, as little freedom exists before the system is overdetermined.

The loss function F now has to be minimized. We found a particularly efficient minimization alternately changes random elements of \check{L} and \check{R} on the few percent level. If the loss F decreased, the random change is accepted and the iteration proceeds to changing new matrix elements.

If the loss did not decrease, the random change is discarded without updating the current matrices \check{L} and \check{R} , and a new iteration is begun.

The minimization of F can be stopped when the blinder's targets are reached, or if F begins to asymptote to the minimal loss achievable under the set constraints. In practice, the loss function must include constraints on disfavoured the old parameters θ_o with respect to the target parameters θ_t , otherwise minimizing the loss function will reproduce the unblinded correlation matrix.

5.4 Setting the target chisquare

Former experience with blinding of Hildebrandt et al. (2017) revealed that out of a set of blinded posteriors, many researchers suggested the one with the smallest minimum- χ^2 represents the true posterior. This is an incorrect assumption, and indeed turned out to be wrong for Hildebrandt et al. (2017). To counter this pre-conception, we suggest the blinder also enforce a χ^2 -value of their choice.

This is easily achieved, e.g. through the loss function of Section 5.3. Another possibility, which will change the determinant, is to reset the eigenvalues of the biased correlation matrix. In this case, the biased correlation matrix is to be spectrally decomposed $\check{C} = \mathbf{Q}\Lambda\mathbf{Q}^{-1}$, where \mathbf{Q} is an orthogonal $d \times d$ matrix satisfying $\mathbf{Q}^T = \mathbf{Q}^{-1}$. The matrix Λ is the diagonal matrix of eigenvalues. The chisquared value at the target parameter point is thus

$$\chi^2(\theta_t) = [\mathbf{u} - \mathbf{v}(\theta_t)]^T [\mathbf{Q}\Lambda\mathbf{Q}^{-1}]^{-1} [\mathbf{u} - \mathbf{v}(\theta_t)] = \sum_i \frac{1}{\lambda_i} (\mathbf{q}_i^T [\mathbf{u} - \mathbf{v}(\theta_t)])^2. \quad (34)$$

Here, λ_i is the i th eigenvalue and the vectors \mathbf{q}_i are the i th row of the matrix \mathbf{Q}^{-1} .

From equation (34) we thus see that the eigenvalues weight the contribution of each summand to the total χ^2 . The blinder may thus reset either a single, or multiple eigenvalues to enforce the χ^2 of their choice at target parameters θ_t .

5.5 Finalization of the algorithm and output

At this stage, the blinding algorithm has nearly completed, with the current output being the prototype \check{C} of the biased correlation matrix. Importantly though, the algorithm does not enforce the diagonal elements of the correlation matrix to be unity during blinding. This is fully acceptable and simply corresponds to a rescaling of the variances. The final step is thus to transform back to the natural units of the astronomical data, which yields the final biased covariance matrix and the final biased correlation matrix.

We therefore multiply back in the variances that were factored out in equation (11)

$$\check{\Sigma}_{ij} = \check{C}_{ij} \sigma_i \sigma_j. \quad (35)$$

This creates the final blinded covariance matrix $\check{\Sigma}$. Thus, if any of the $\check{C}_{ii} \neq 1$ at this stage, then this simply rescales the variance $\check{\Sigma}_{ii}$. The final blinded correlation matrix will then none the less have unit diagonal elements. It results from the computation

$$\check{C}_{ij} = \frac{\check{\Sigma}_{ij}}{\sqrt{\check{\Sigma}_{ii}} \sqrt{\check{\Sigma}_{jj}}}. \quad (36)$$

At this stage, the blinding algorithm is completed, with $\check{\Sigma}$ and \check{C} being the final blinded covariance and correlation matrix.

An example of the relative differences between the unblinded and final blinded covariance matrix is seen in Fig. 5, from which it can indeed be seen that the variances (diagonal elements) changed

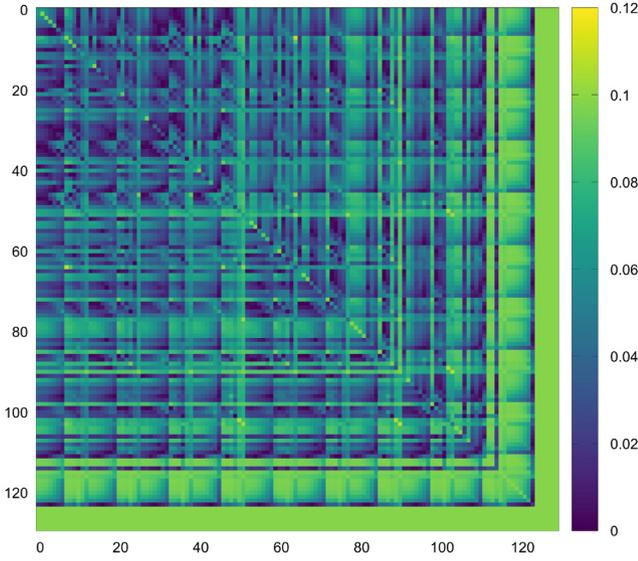


Figure 5. Plot of the relative changes each matrix element of the covariance matrix underwent whilst blinding the KiDS-450 analysis. The colour bar measures the ‘symmetrized mean absolute percentage error’ (SMAPE), defined in equation (29). The here plotted relative difference matrix causes the upwards shift from orange to red posterior in Fig. 4. The green bars of SMAPE values around 0.1 for matrix rows 120–130 affect the highest redshift bin of KiDS-450, and there preferentially ξ_{-} . This implies that the upwards shift of the KiDS-450 posterior can be caused by mischaracterizing the correlation of the data from the highest redshift bin with all lower redshift bins. Appendix B provides further cosmic shear specific context for this figure.

somewhat, but that the posterior shift is mostly induced by having changed off-diagonal covariance elements. This also explains why spotting the blinding from the joint plot of data and theory curves in Figs 1 and 2 is essentially impossible: the covariances do not show up when overplotting data and theory predictions, and thus go unnoticed when attempting to fit by eye.

Examples of final biased correlation matrices are seen in Fig. 3, where it is difficult to spot the unbiased correlation matrix amongst the two biased correlation matrices. The corresponding biased and unbiased covariance matrices are equally difficult to tell apart, but due to the disadvantageous scaling over many order of magnitudes (see y-axes of the data in Fig. 1) this is difficult to visualize in a colour plot.

5.6 Blinding control

Although highly reliable, the algorithm does contain free algorithmic parameters W , S_{inv} , S_{corr} , an adaptable loss function F , and random seeds. It may thus sometimes fail, either due to a user error or due to chance. Before using the blinded covariance matrix $\check{\Sigma}$ in a posterior, the blinder has to control the sanity of the matrices, and that the intended posterior shift was achieved.

The biased covariance matrix will be mathematically sound, if all variances are positive, and if the final biased correlation matrix is positive definite, and its elements take values on the interval $[-1, 1]$. As the algorithm used Cholesky decompositions and spectral decompositions, positive definiteness should be guaranteed.

Whether the blinded covariance matrix will shift the posterior as intended can be evaluated in multiple ways. To linear order, it can be checked whether the best-fitting estimator equation (10) indeed yields parameters close to the targeted θ_t . To non-linear order, it

should be ensured that χ^2 at the shift’s origin has increased during blinding, i.e.

$$\chi_{\text{final}}^2(\theta_o, \mathbf{u}, \check{\mathbf{C}}) > \chi^2(\theta_o, \mathbf{u}, \mathbf{C}), \quad (37)$$

which expresses that the blinded covariance matrix disfavours the former parameters. Simultaneously, χ^2 at the target parameters should have decreased due to blinding

$$\chi_{\text{final}}^2(\theta_t, \mathbf{u}, \check{\mathbf{C}}) < \chi^2(\theta_t, \mathbf{u}, \mathbf{C}), \quad (38)$$

expressing that the formerly disfavoured parameters are now a better fit than before.

To check whether the posterior shifted by the intended number of standard deviations, the final check is to compute the final $\Delta\chi^2$ between formerly preferred and target parameters

$$\Delta\chi_{\text{final}}^2 = \chi_{\text{final}}^2(\theta_o, \mathbf{u}, \check{\mathbf{C}}) - \chi_{\text{final}}^2(\theta_t, \mathbf{u}, \check{\mathbf{C}}). \quad (39)$$

This $\Delta\chi_{\text{final}}^2$ has to be compared against Table 1, in order to control whether its magnitude shifts by sufficiently many standard deviations for the total number of parameters to be fitted. Should this $\Delta\chi_{\text{final}}^2$ be negative, then the old parameters were still preferred.

If any of these tests is failed, then the blinding algorithm has to be run with adapted algorithmic parameters. In the public code accompanying this paper, the code reports whether a test is failed, and recommends improved settings of the algorithmic parameters. Otherwise, if all tests are passed, then the yielded biased covariance matrix $\check{\Sigma}$ is qualified for use in a blinded posterior evaluation.

6 COMPUTING THE BLINDED POSTERIOR

The above algorithm allows the construction of biased covariance matrices $\check{\Sigma}$ which shift the posterior in a requested direction. As the algorithm has free parameters, including S_{inv} , S_{corr} , θ_t and also random seeds, many such covariance matrices can be computed. If multiple covariance matrices are computed, care should be taken that the entire set of covariance matrices does not allow joint debinding, for example by averaging.

Any such biased covariance matrix would then be used to compute the blinded Gaussian posterior

$$\check{\mathcal{P}}(\theta|\check{\mathbf{x}}, \check{\Sigma}) \propto \exp\left(-\frac{1}{2}[\check{\mathbf{x}} - \mu(\theta)]^\top \check{\Sigma}^{-1}[\check{\mathbf{x}} - \mu(\theta)]\right), \quad (40)$$

or in the blinded t -distribution of Sellentin & Heavens (2016). This posterior’s peak position will be jointly influenced by the data-blinder’s target parameters, and the likelihood-blinder’s target parameters. An example of such shifted posteriors is seen in Fig. 4, where the original data \mathbf{x} of KiDS-450 were used.

7 DEBLINDING

The essential step at the end of any blinded analysis is of course to deblind.⁴ In our setup, up to three quantities were blinded, the data, the theory computations, and the covariance matrix. The theory blinder is strongly recommended to deblind directly after having received the blinded data and the blinded covariance. This means the theory blinder is the only one who is recommended to deblind *prior* to computing the blinded posterior.

⁴We distinguish between unblinded and debinded. An unblinded quantity never was blinded, and a debinded quantity was intermittently blinded but the blinding is then undone.

This recommendation has a utilitarian aim: the numerical costs of posterior computations are usually dominated by the theoretical predictions $\mu(\theta)$. We therefore recommend storing all computed values of $\mu(\theta)$ when computing the blinded posterior. Deblinding is then achieved with the following numerically lightweight post-processing of the blinded posterior.

7.1 Deblinding by posterior post-processing

During the blinded analysis, a blinded posterior $\check{\mathcal{P}}(\theta|\check{\mathbf{x}}, \check{\Sigma})$ was computed. The aim is now to compute the unblinded posterior $\mathcal{P}(\theta|\mathbf{x}, \Sigma)$ whilst minimizing computational overload. This is achieved by multiplying the posterior with deblinding weights, which corresponds to importance sampling.

We relate the blinded and deblinded posterior by

$$\mathcal{P}(\theta|\mathbf{x}, \Sigma) = w(\theta)\check{\mathcal{P}}(\theta|\check{\mathbf{x}}, \check{\Sigma}). \quad (41)$$

The blinded posterior will have been computed at N discrete points θ_i , with $i \in [1, N]$. If the posterior was computed on a grid, then the θ_i are regularly spaced; if the posterior was sampled with a Monte Carlo Markov Chain (MCMC) technique, then the index i enumerates the samples of the chain. The aim is now to avoid that a new chain must be run after deblinding, as this may be extremely costly.

We therefore use that the general MCMC sampler will have produced a chain which will equilibrate to the *blinded* posterior, but each sample will additionally have a weight which depends on the exact algorithm used. This weight can be multiplied with deblinding weights, in order to produce chains which equilibrate to the deblinded posterior. Each sample of the deblinded chain will then contribute with more or less weight to the deblinded posterior.

For all samples, the deblinding weights are

$$w(\theta_i) = \frac{\mathcal{P}(\theta_i|\mathbf{x}, \Sigma)}{\check{\mathcal{P}}(\theta_i|\check{\mathbf{x}}, \check{\Sigma})}. \quad (42)$$

As the means $\mu(\theta_i)$ were stored, the weights $w(\theta_i)$ are quickly evaluated. For a grid-based posterior computation, the blinded posterior is directly multiplied with the weights, according to equation (41). For the MCMC sampled case, the former weights of each sample are multiplied by the deblinding weights. Deblinding by posterior weighting is depicted in Fig. 6.

We caution that MCMC convergence after deblinding should be carefully assessed. It may be advisable to hide one unblinded analysis amongst other blinds, to enforce a high sample density in the region of importance.

8 DISCUSSION

This publication established a numerically lightweight blinding and deblinding algorithm which ties in with astronomy's special circumstance of having often unique, irrepeatable, and easily recognizable data sets.

It is often said blinding avoids iterative or subconscious biasing of an analysis in order to fall in line with the status quo of a field. The positive flip-side of this view is that blinding eases the possibility to convincingly disprove a status quo, and thereby avoid stagnation of the field. Blinding also motivates increased model-independent testing of the analysis, thereby strengthening the understanding of the astronomical data prior to inferring physics. Blinding has therefore only positive aspects, if numerically cheap, as is the case for the here presented strategy.

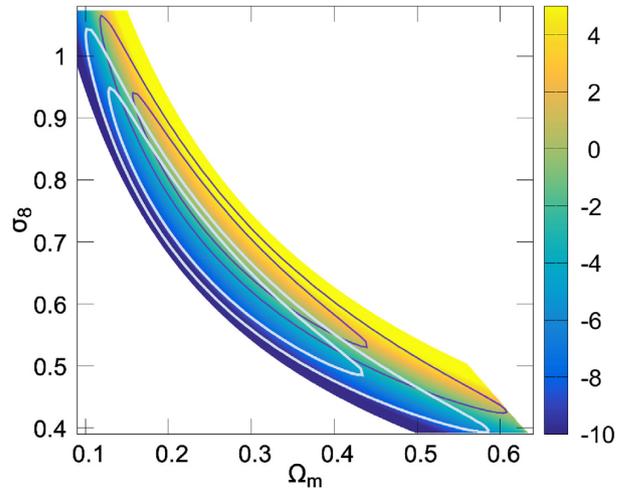


Figure 6. Illustration of deblinding: the colourbar refers to the logarithm of the deblinding weights from equation (42). Multiplication of the bright-blue blinded posterior with the deblinding-weights results in the deblinded dark-purple posterior. Where the blinded and true posterior overlap, the deblinding weights take values around units (zero on the plotted log-scale). Positive weights (yellow) indicate an increase of posterior probability during deblinding. Negative weights (blue and purple) indicate downweighted parameter probabilities during deblinding.

Our blinding strategy allows limited control over the inference to be assigned to external researchers, thereby addressing potential concerns about the details of distilling science-ready data out of astronomical raw data. In total, the presented strategy enables up to three-stage blinding, where especially the covariance-blinder has – up to parameter degeneracies – close to perfect control over determining where the blinded posterior shall peak.

Once sampled or computed on a grid, the posterior can be deblinded by multiplying with deblinding weights, thereby revealing the parameters actually preferred by the data.

ACKNOWLEDGEMENTS

ES is supported by Leiden's Oort-Fellowship programme and thanks Hendrik Hildebrandt, Catherine Heymans, Koen Kuijken, and Henk Hoekstra. This research is based on data products from observations made with ESO Telescopes at the La Silla Paranal Observatory under programme IDs 177.A-3016, 177.A-3017, and 177.A-3018.

REFERENCES

- Aad G. et al., 2012, *Phys. Lett. B*, 716, 1
 Abbott B. P. et al., 2016, *Phys. Rev. X*, 6, 041015
 Adrián-Martínez S. et al., 2016, *Phys. Lett. B*, 760, 143
 Akerib D. S. et al., 2017, *Phys. Rev. Lett.*, 118, 021303
 Begeman K., Belikov A. N., Boxhoorn D. R., Valentijn E. A., 2013, *Exp. Astron.*, 35, 1
 Chatchyan S. et al., 2012, *Phys. Lett. B*, 716, 30
 Croft R. A. C., Dailey M., 2011, preprint ([arXiv:1112.3108](https://arxiv.org/abs/1112.3108))
 de Jong J. T. A. et al., 2015, *A&A*, 582, A62
 Einasto M. et al., 2011, *ApJ*, 736, 51
 Erben T. et al., 2013, *MNRAS*, 433, 2545
 Fenech Conti I., Herbonnet R., Hoekstra H., Merten J., Miller L., Viola M., 2017, *MNRAS*, 467, 1627
 Gaia Collaboration, 2018, *A&A*, 616, A1
 Hildebrandt H. et al., 2017, *MNRAS*, 465, 1454

- Jain B. et al., 2015, preprint (arXiv:1501.07897)
 Kuijken K. et al., 2015, *MNRAS*, 454, 3500
 Laureijs R. et al., 2011, preprint (arXiv:1110.3193)
 Lesgourgues J., 2011, preprint (arXiv:1104.2932)
 Lindegren L. et al., 2016, *A&A*, 595, A4
 Miller L., Kitching T. D., Heymans C., Heavens A. F., van Waerbeke L., 2007, *MNRAS*, 382, 315
 Planck Collaboration, 2018, preprint (arXiv:1807.06209)
 Seehars S., Grandis S., Amara A., Refregier A., 2016, *Phys. Rev. D*, 93, 103507
 Sellentin E., 2015, *MNRAS*, 453, 893
 Sellentin E., Heavens A. F., 2016, *MNRAS*, 456, L132
 Sellentin E., Heavens A. F., 2017, *MNRAS*, 464, 4658
 Sellentin E., Starck J.-L., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 021
 Sellentin E., Quartin M., Amendola L., 2014, *MNRAS*, 441, 1831
 Tegmark M., Taylor A., Heavens A., 1997, *ApJ*, 480, 22

APPENDIX A: ESSENTIALS OF COSMIC SHEAR

Cosmic shear is a cosmological observation technique for which we showcased the blinding algorithm. Cosmic shear measures the shearing of galaxy images on the sky. This effect arises from general relativity, according to which light follows null-geodesics which adapt to the presence of matter. This leads to light being deflected by massive objects.

As the cosmic matter fields are perturbed, the deflection of light traversing them imprints similar perturbation patterns on galaxy images. Cosmic shear measures these distortions over significant fractions of the sky, and computes two correlation functions from it. In this article, these two correlation functions are denoted as ξ_+ and ξ_- , and are measured as a function of angular separation ϑ expressed in arcminutes.

Both $\xi_+(\vartheta)$ and $\xi_-(\vartheta)$ are simultaneously caused by cosmic shear – the presented data set in Fig. 1 therefore includes the upper (ξ_+) as well as the lower panels (ξ_-).

The triangular arrangement of the data set results from having partitioned all observed galaxies into ‘bins’, where each bin is identified by its mean redshift. Approximately, the galaxies closest to us are assigned to redshift bin ‘0’ (see plot labels), and the galaxies furthest from us are assigned to bin ‘3’. Of all bins, the autocorrelation and cross-correlation functions are measured, which leads to the labels ‘0 0’, ‘0 1’..., ‘3 3’ in the subpanels of the triangular plots.

In total, all subpanels of Fig. 1 display one joint 130-dimensional data set. For further detail we refer the interested reader to Hildebrandt et al. (2017) and references therein.

APPENDIX B: IMPLICATIONS FOR KIDS-450

This paper developed a general blinding technique and illustrated it on the KiDS-450 data which were first analysed in Hildebrandt et al. (2017). This appendix embeds our findings in the larger context of cosmic shear research.

Cosmic shear is highly sensitive to the dark matter density Ω_m and the power spectrum amplitude σ_8 via the combination $S_8 = \sigma_8 \sqrt{\Omega_m/0.3}$. A mild tension between cosmic shear constraints for S_8 and Planck constraints on S_8 (Planck Collaboration VI 2018) has persisted for multiple years now, with cosmic shear returning lower values of S_8 than Planck. Our shift of the original KiDS-450 posterior (yellow in Fig. 4) towards higher S_8 (red in Fig. 4) should therefore be put into context.

Fig. 5 illustrates that SMAPE errors of at most 0.12 (defined in equation 29 and similar to percentual deviations) for covariance matrix elements suffice to allow shifts of the KiDS posterior into Planck compatible regions. As cosmic shear covariance matrices are either analytical approximations or numerical estimates, they will indeed be biased to a certain degree, but currently no evidence exists that the very specific bias required for the posterior shift affects the KiDS-450 covariance matrix.

Nonetheless, it is surprising that all data points of the highest redshift bin in KiDS-450 light up *consistently* in Fig. 5. Usually, the method here presented will affect all data points to a low degree without any preference of physically meaningful subgroups in the data.

It is thus unclear why Fig. 5 consistently impacts the highest redshifts. The safest interpretation of Fig. 5 is that the S_8 tension correlates with the total uncertainty at high redshifts – whether this correlation implies a causal connection is not known, but it illustrates that an agnostic route towards understanding the origin of the tension between Planck and cosmic shear has to include a detailed understanding of cosmic shear uncertainties at high redshifts.

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