

Cosmological Inference with Cosmic Shear



Peter Llewelyn Taylor
Mullard Space Science Laboratory,
Department of Space and Climate Physics
University College London

A thesis submitted for the degree of
Doctor of Philosophy
July 2019

Statement of Originality

I confirm that the work presented in this thesis is my own except where explicitly stated otherwise. The three research chapters are based on my papers:

- Peter L. Taylor, Francis Bernardeau, and Thomas D. Kitching.
k-cut cosmic shear: Tuneable power spectrum sensitivity to test gravity. *Phys. Rev. D*, 98(8):083514, 2018.
- Peter L. Taylor, Thomas D. Kitching, and Jason D. McEwen.
Nonparametric cosmology with cosmic shear. *Phys. Rev. D*, 99:043532, 2019.
- Peter L. Taylor, Thomas D. Kitching, Justin Alsing, Benjamin D. Wandelt, Stephen M. Feeney, and Jason D. McEwen. Cosmic Shear: Inference from Forward Models. *arXiv e-prints*, arXiv:1904.05364, 2019. (*Phys. Rev. D* accepted)

closely following the exposition therein. I have also contributed to:

- Peter L. Taylor, Thomas D. Kitching, Jason D. McEwen, and Thomas Tram. Testing the cosmic shear spatially-flat universe approximation with generalized lensing and shear spectra. *Phys. Rev. D*, 98(2):023522, 2018.
- Peter L. Taylor, Thomas D. Kitching, and Jason D. McEwen. Preparing for the cosmic shear data flood: Optimal data extraction and simulation requirements for stage iv dark energy experiments. *Phys. Rev. D*, 98:043532, 2018.

- Alessio Spurio Mancini, Peter L. Taylor, R Reischke, T. Kitching, V. Pettorino, B. M. Schafer, B. Zieser, and P. M. Merkel. 3d cosmic shear: Numerical challenges, 3d lensing random fields generation, and minkowski functionals for cosmological inference. *Phys. Rev. D*, 98(10):103507, 2018.
- Thomas D. Kitching, Peter L. Taylor, Peter Capak, Daniel Masters, and Henk Hoekstra. Rainbow cosmic shear: Optimization of tomographic bins. *Phys. Rev. D*, 99(6):063536, 2019.

which will be cited like any other papers. I have contributed to two further publications over the funding period on a separate topic:

- Peter Taylor, Richard Massey, Mathilde Jauzac, Frederic Courbin, David Harvey, Remy Joseph, and Andrew Robertson. A test for skewed distributions of dark matter, and a possible detection in galaxy cluster abell 3827. *Monthly Notices of the Royal Astronomical Society*, 468(4):50045013, 2017.
- Richard Massey, David Harvey, Jori Liesenborgs, Johan Richard, Stuart Stach, Mark Swinbank, Peter Taylor, Liliya Williams, Douglas Clowe, Frederic Courbin, et al. Dark matter dynamics in abell 3827: new data consistent with standard cold dark matter. *Monthly Notices of the Royal Astronomical Society*, 477(1):669677, 2018.

which will not be discussed.

Signature: _____

Abstract

Over the next decade, data from large Stage IV survey telescopes including Euclid, LSST and WFIRST will provide some of the tightest cosmological constraints.

To extract information from these surveys we take advantage of gravitational lensing, an effect predicted by Einstein’s general theory of relativity. Gravitational lensing simply refers to the bending of light rays around massive bodies. This causes small changes in the observed ellipticity of galaxies, which is called weak gravitational lensing or — on the largest scales — cosmic shear. By examining these shape distortions over millions, or even billions of galaxies, we can distinguish between alternative cosmological models and measure the fundamental cosmological parameters precisely.

While the constraining power of these upcoming data sets will improve by more than an order of magnitude, our statistical methods are not keeping pace. In this thesis I develop three new techniques to take full advantage of next generation surveys.

The first of these is a method called k -cut cosmic shear. It allows us to efficiently remove sensitivity to small scales that are too difficult to model accurately due to complicated baryonic physics and nonlinear structure formation.

Next I present a method called non-parametric cosmology with cosmic shear. I show how to extract information about the growth of structure and the background expansion of the Universe with no *a priori*

assumption about the underlying cosmological model. This can be used to search for failures of the Lambda-Cold Dark Matter (LCDM) model.

Finally I show how to perform inference with full forward models of the cosmic shear data. This approach allows us to seamlessly propagate all astrophysical, theoretical and instrumental systematics into the final parameter constraints, sidestepping complicated issues including the deconvolution of the survey mask and an assumption about the functional form of the likelihood.

Impact Statement

This thesis presents new methods for analysing cosmic shear data from next generation billion dollar class missions. Ultimately this may play a role in determining the true nature of dark energy and measuring the parameters which govern the behaviour of the Universe on the largest scales.

To my parents and Alex

Acknowledgements

I could not have asked for better supervisors. Tom has suggested some excellent avenues of enquiry and I know that I have benefited immeasurably from his almost encyclopaedic knowledge of everything weak lensing. Jason has always been around to help, particularly when there were some tricky numerics involved. My time at MSSL would not have been nearly as enjoyable without my friends: Ahlam, Aisha, Anurag, Brian, Chris, Choong Ling, Dennis, Ellis, James, Jennifer, Jo, Luke, Matt, Monica, Monu, Richard, Sam and Tom. I will miss our trips to the pub and seemingly unending coffee breaks. All the staff and postdocs have made MSSL everything an astronomy department should be – a truly friendly, welcoming and inquisitive place.

I can also never express enough gratitude to my masters supervisors in Durham, Richard Massey and Mathilde Jauzac, and my office mate Andrew Robertson. When I started as a research student I was extremely naïve about what research entailed. I had no formal training in astronomy and never programmed a computer. It was only because of their patience and encouragement that I have made it this far.

I would also like to thank my viva examiners Alkistis Poutsis and Daisuke Kawata for taking the time to read this thesis and my hosts and collaborators: Björn Malte Schäfer, Alessio Spurio Mancini, Robert Reischke, Justin Alsing, Benjamin Wandelt, Stephen Feeney and Francis Bernardeau in Heidelberg, the Flatiron Institute and the IAP. I am indebted to Joe Zuntz and Thomas Tram for taking time to help me use their codes and to everyone else who helped me along

the way at some point in my research career including: Joel Primack,
Pedro Ferreira and Eric Jullo.

Contents

1	Introduction	23
1.1	General Relativity	26
1.1.1	General Relativity	26
1.1.2	Deflection of Light Due to Gravity	27
1.2	Geometry and Distances in the Universe	28
1.2.1	The Expanding Universe and Accelerating Universe	28
1.2.2	The Cosmological Principle and the Friedmann-Robertson-Walker Metric	29
1.2.3	Cosmological Redshift	30
1.2.4	Distance Measures	30
1.3	Contents of the Universe	31
1.3.1	The Conservation Equation and The Equation of State	31
1.3.2	Components of the Universe	32
1.4	Kinematics of the Universe	32
1.4.1	The Friedmann Equation	32
1.4.2	Background Expansion of the Universe	33
1.4.3	Kinematics for a Generalised Dark Energy Equation of State	34
1.5	Structure Growth	35
1.5.1	Linear Structure growth	35
1.5.2	Linear Growth of Structure: A Case Study	36
1.6	Statistics of the Cosmological Density Field	36
1.6.1	Random Fields and Power Spectra	36

1.6.2	Initial Conditions	37
1.6.3	Computing the Matter Power Spectrum	38
1.7	Cosmic Shear	40
1.7.1	Weak Lensing	40
1.7.2	The Shear Field	42
1.7.3	Shape and Distance Measurements	45
1.7.4	The Generalised Lensing Spectrum	46
1.7.5	The Tomographic Lensing Spectrum	48
1.7.6	Shot-Noise	48
1.7.7	Intrinsic Alignments	49
1.7.8	E/B-Modes	52
1.7.9	Current and Future Data	54
2	<i>k</i>-cut Cosmic Shear	55
2.1	Motivation	55
2.2	The Shell Universe Thought Experiment	57
2.3	The Three-plane Solution	58
2.4	Constructing the Bernardeau-Nishimichi-Taruya (BNT) Transform .	59
2.5	<i>k</i> -cut Cosmic Shear	60
2.6	Fisher Matrix Formalism	62
2.7	Results	63
2.8	Outlook and Future Prospects	65
3	Non-parametric Cosmology with Cosmic Shear	72
3.1	Motivation	72
3.2	Correlation Functions	73
3.3	CFHTLenS Data	75
3.4	The Non-parametric Likelihood	75
3.5	Non-parametric Information Extraction	77
3.6	Adaptive Grid Compression	78

3.7	Principle Component Compression	80
3.8	Cosmological Parameter Inference from Non-parametric Information	81
3.9	Adaptive Grid Reconstruction Results	83
3.10	PCA Reconstruction Results	86
3.11	Cosmological Inference from Adaptive Grid Reconstruction	88
3.12	Cosmological Inference from PCA Reconstruction	88
3.13	Conclusion and Future Prospects	91
4	Inference from Forward Models	95
4.1	Motivation	95
4.2	Cosmic Shear Model Choices and the Lognormal Field Approximation	99
4.2.1	Cosmic Shear Modelling Choices	99
4.2.2	The Lognormal Field Approximation	99
4.2.3	Band-limit Bias from the Lognormal Field	103
4.3	Cosmological Parameter Inference	104
4.3.1	The Potential Insufficiency of the Gaussian likelihood Approximation in Cosmic Shear	104
4.3.2	Density-estimation Likelihood-free Overview	105
4.3.3	Density-estimation Likelihood-free Compression	105
4.3.4	Neural Density Estimators (NDEs) in DELFI	106
4.3.5	Active Sequential Learning	107
4.4	The Full Forward Model	108
4.4.1	The Mask	108
4.4.2	Shot-Noise Model	108
4.4.3	Forward Modelling the Mask	109
4.4.4	Mimicking a Stage IV Experiment	112
4.5	Prospects for Stage III Data	113
4.5.1	Validating the Lognormal Simulations	113
4.5.2	Model Improvements	114
4.6	Testing the Gaussian Likelihood Approximation	115

4.6.1	Modelling Choices In Pipeline II	115
4.6.2	Impact of the Gaussian Likelihood Approximation	117
4.7	Future Prospects	120
4.8	Conclusion	122
5	Conclusion	123
A	Calculation of Light Deflection in General Relativity	125
B	Deflection angle in the observer's frame	128
C	Computation of the Lensing Spectrum	130
D	Derivation of the Shot-noise	133
E	The Reduced Shear Correction From Forward Models	134

List of Figures

1.1	The linear and nonlinear power spectra at redshifts $z = 0$ and $z = 3$. As structure collapses through cosmic time, the amplitude of the power spectrum increases. At small scales, above $k \sim 10^{-1}h \text{ Mpc}^{-1}$, nonlinear corrections become relevant. By $k \sim 2h \text{ Mpc}^{-1}$, the linear power spectrum underpredicts the nonlinear power spectrum by an order of magnitude. Cosmic shear is primarily sensitive to structure in the range $k \in [10^{-1}h \text{ Mpc}^{-1}, 10h \text{ Mpc}^{-1}]$ [124]. The linear power spectrum was generated using <code>Camb</code> [84] and the nonlinear corrections were generated using <code>HALOFIT</code> [120] inside the modular framework <code>Cosmosis</code> [136].	39
1.2	The combined effect of shear and convergence. Convergence changes the radius of a circular source while the shear changes the ellipticity. This figure originally appears in [98].	43
1.3	The ellipticity components γ_1 and γ_2 . As the ellipticity rotates around π radians, the orientation angle, ψ , rotates around 2π radians. This figure originally appears in [70].	44
1.4	The green ellipsoid represents a large dark matter halo. Nearby galaxies (red) become tidally aligned with the dark matter halo while the observed ellipticity of background galaxies (blue) is gravitationally sheared. The intrinsic alignments of galaxies around dark matter halos is orthogonal to the shearing effect which suppresses the cosmic shear signal. This figure originally appears in [65].	49

3.5 The posterior distribution on the measured PC amplitudes. The constraints on the first two PCs are in good agreement with those expected from the Fisher matrix prediction, but the constraints on the other PCs are up to twice as wide as expected. I also plot the Gaussian distribution that we have fit to the chains to form the likelihood $\ln \mathcal{L}_1(p)$, in grey.	89
3.6 The LCDM posteriors found using the standard cosmic shear likelihood analysis (blue) and from the non-parametric information (red), using adaptive grid compression. The two techniques are in good agreement, but the posteriors in the later case are broader. This is unsurprising since information is lost in the adaptive grid compression step.	90
3.7 LCDM posteriors derived from the standard cosmic shear likelihood analysis (blue) and those derived from the non-parametric information (green), using PCA compression. The two techniques are in good agreement but there is $\sim 1\sigma$ tension in the $\sigma_8 - \Omega_m$ plane. Because the PCs do not form a complete set, they do not capture all the variance leading us to underestimate error bars, and could also cause a shift in the parameter constraints. I do not advocate the PCA data compression method for this reason.	92
4.1 A schematic of the two forward model pipelines presented in this chapter. In Pipeline I I develop a forward model pipeline of cosmic shear data that takes full advantage of the forward model paradigm. For example, there is no need to deconvolve the mask or the pixel window function. In Pipeline II, I use a Gaussian field, do not use a mask, subtract off the shot-noise or deconvolve the pixel window function. These choices allow me to make a direct comparison between DELFI and a Gaussian likelihood analysis and test the impact of the Gaussian likelihood approximation.	98

4.2 A single masked data realisation of the convergence field, κ , and the two observable shear components: γ_1 and γ_2 (including shape noise) for a typical Stage IV experiment. I show the lowest redshift bin of six, where the effect of non-Gaussianity is largest. The non-Gaussianity is clearly visible in the κ -map, where the majority of pixels are very slightly negative with a small number of pixels taking very large (positive) κ -values. Note that the colour scale in the κ -map is not symmetric about zero, to make the non-Gaussianity as clearly visible as possible. The mask cuts all pixels lying within 22.5 deg of the galactic and ecliptic planes, mimicking a Euclid-like survey.	100
4.3 68% and 95% credible region parameter constraints found with DELFI analysis DA1 after 1000 simulations, for a mock Stage IV experiment. I confirm that the input cosmology is recovered within statistical errors. The convergence is plotted in Figure 4.4. In a realistic situation there may be a larger number of nuisance parameters. One should not expect this to dramatically slow convergence because we could ‘nuisance harden’ the data compression step, and only learn the posterior for the parameters of interest. For more details see [11].	110
4.4 The negative loss function defined in equation (4.18) for the training and validation sets as a function of the number of simulations. This suggests that $\mathcal{O}(1000)$ simulations are needed for a Stage IV experiment. As this is similar to the number found in [7], which only considered a simple Gaussian field forward model with no intrinsic alignments, the convergence rate should be fairly insensitive to model improvements in the future in preparation for Stage IV experiments.	111

4.5 The coloured lines show the absolute value of the difference between the average recovered cross and intra-bin spectra from 100 lognormal and 100 Gaussian realisations (4 tomographic bins, $N_{\text{side}} = 512$ and $\ell \in [10, 1535]$). The difference is due to the band-limit and regularisation bias in the lognormal field. The combined bias is safely below 1% for nearly all data points.	113
4.6 The absolute value of the bias, $ b $, due to imperfect pixel-window deconvolution and noise subtraction relative to the statistical error, σ , from 500 Pipeline II simulations. The bias is so small that the comparison between DELFI and the likelihood analysis presented in Section 4.6.2 will be unaffected.	116
4.7 The 68% and 95% credible region parameter constraints for three random data realisations found using a Gaussian likelihood analysis and DELFI. The mock data input cosmology is labeled by black dotted lines. Only in the first realisation, does the input cosmology lie outside the 68% credible region – but statistically this is to be expected for a small number of realisations. The contours found using the two different analyses are very similar for all three data realisations suggesting that the Gaussian likelihood approximation has negligible impact. This is confirmed in Figure 4.8.	118
4.8 The 68% and 95% credible region of the MLE distribution, assuming a Gaussian likelihood. The value of the input cosmology is indicated by the black dotted lines, and lies at the center of the contours. This implies that the Gaussian likelihood approximation does not lead to any measurable bias in our setup.	119

Chapter 1

Introduction

Far from slowing down under the attractive force of gravity, the Universe’s expansion is actually accelerating. The leading explanation is that 70% of the energy is in the form of a cosmological constant – a mysterious form of energy that has the same density everywhere. Meanwhile, to explain why galaxies do not fly apart under their own rotation, it is conjectured that 25% of the remaining rest energy in the Universe is in the form of dark matter – a substance that interacts only through gravity. These assumptions form the Lambda-Cold Dark (LCDM) paradigm. While LCDM passes nearly all current observational tests, there is no widely accepted physical model that explains the existence of dark energy and dark matter. Over the next decade, data from large Stage IV survey telescopes including Euclid¹ [80], the Wide Field Infrared Survey Telescope² (WFIRST) and the Large Synoptic Survey Telescope³ (LSST), will help us test the LCDM paradigm and in particular determine if the Universe’s accelerated expansion is driven by the cosmological constant – or something else. Potential alternatives include the addition of extra fields or modified gravity. See [38] for a review.

To extract the cosmological information from Stage IV surveys we take advantage of gravitational lensing, an effect predicted by Einstein’s general theory of relativity. Gravitational lensing simply refers to the bending of light rays around

¹<http://euclid-ec.org>

²<https://www.nasa.gov/wfirst>

³<https://www.lsst.org>

massive bodies. This causes small changes in the observed ellipticity and size of galaxies, which is referred to as weak gravitational lensing. By examining the shape distortions over millions, or even billions of galaxies, we can distinguish between cosmological models and measure cosmological parameters precisely.

A Gaussian-likelihood analysis of the two-point statistics is by far the most popular way to extract cosmological information from large photometric survey shear data [51, 56, 130]. The first step is to measure the ellipticity, angular position and photometric redshift (a proxy for distance) of every galaxy in the survey. From this, one computes the correlations in ellipticity between pairs of galaxies in redshift bins i and j as a function of angular separation on the sky, θ . These are called the cosmic shear two-point correlation functions and they are written as $\xi_+^{ij}(\theta)$ and $\xi_-^{ij}(\theta)$. The plus and minus indicate whether we take the sum or the difference of the so called ‘tangential’ and ‘cross’ ellipticities in the frame joining the two galaxies. Alternatively we can compute power spectra C_ℓ^{ij} , which expresses the correlations between pairs of galaxy ellipticities in harmonic space.

Assuming a cosmological model with free-parameters p we infer the posterior distribution on these parameters by Markov Chain Monte Carlo (MCMC) sampling from a likelihood function, which we assume to be Gaussian. It is written as

$$\ln \mathcal{L}(p) = -\frac{1}{2} \sum_{a,b} [D_a - T_a(p)] C_{ab}^{-1} [D_b - T_b(p)], \quad (1.1)$$

where D_a is the data vector composed of the observed correlation functions, $T_a(p)$ is the expected value of the correlation functions given the cosmological parameters p , C_{ab}^{-1} is the inverse of the covariance matrix and the sum is over all data points pairs (a, b) . This procedure encapsulates the way that most shear analyses are done today.

Relative to today’s current surveys, the number of measured galaxies in Stage IV surveys will increase by two orders of magnitude [5], yet the statistical framework used to extract the cosmological information is not keeping pace. With Stage IV experiments coming online in just five years, we urgently need to address serious

issues with this framework. These can be divided into three broad categories:

- **Issue 1 Current techniques are prone to bias:** This broader problem breaks down in two sub-issues:

A. Despite probing the large scale structure evolution of the Universe, cosmic shear is also sensitive to the physics of small scale high-density regions. The nonlinear structure growth and baryonic (non-gravitational) physics in this regime are very difficult to model. For example, the impact of winds from the massive black holes at the centre of galaxies – which suppresses the formation of structure in a process called AGN feedback – is very uncertain [60]. If not handled correctly, modelling uncertainties at small scales will lead to significant bias in the dark energy measurement [33, 60, 97, 115, 124].

B. To extract information about the cosmological model parameters, M , from the shear data, D , we use Bayes' theorem to write $P(M|D) = P(D|M)P(M)$ and sample from the posterior distribution, $P(M|D)$. The prior, $P(M)$, is computed from other cosmological probes. Meanwhile in current weak lensing studies the likelihood, $P(D|M)$, is assumed to be a Gaussian distribution. It was recently claimed that this could lead to bias in Stage IV experiments [114], so the Gaussian likelihood approximation must be tested or avoided altogether

- **Issue 2 We are not ready for an unexpected or paradigm shifting result:** At the moment, we assume a cosmological model and then infer the value of just a few parameters. With the increased statistical power of Stage IV experiments it is conceivable we will make a measurement that is inconsistent with the cosmological constant, or in tension with other cosmological probes. The few model parameters we measure today may not encapsulate why Λ CDM has failed, making it difficult to distinguish between exciting new physics and systematics.

- **Issue 3 We do not perform inference on full forward models of the data:**

Inside the current likelihood framework we must rapidly compute the theoretical lensing prediction for a large number of points in cosmological parameter space to infer the underlying cosmological parameters. Instrumental systematics (e.g the blurring of galaxies due to the telescope’s point spread function) must be deconvolved, and this may lead to bias. Meanwhile ‘theoretical systematics’ (e.g the reduced shear correction [32], which is known to cause bias for next generation data) are not considered as they are computationally expensive to compute. While it is difficult to compute the *expected* impact of instrumental and theoretical, it is easy to forward model realisations of the data which include these effects. Ideally, we need a framework to perform cosmological inference on a small number of forward model realisations of the data.

I propose a set of techniques which mitigates all these issues in the next three research chapters. In Chapter 1 I present a method which overcomes Issue 1A, in Chapter 2 I address Issue 2 and in Chapter 3 I tackle Issue 3 and investigate Issue 1B. In the remainder of this chapter I review all the relevant general relativity, cosmology and weak lensing formalism needed to understand the remainder of the thesis.

1.1 General Relativity

In this section I aim to concisely summarise the general relativity relevant to this thesis.

1.1.1 General Relativity

Of the four known fundamental forces of nature, it is gravity that dominates on cosmological scales. In general relativity, gravity is described by the Einstein field

equations:

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}, \quad (1.2)$$

where G is Newton's constant and c is the speed of light in vacuum. The left hand side is the Einstein tensor, which is a geometric quantity that gives the curvature of spacetime, while on the right, $T_{\mu\nu}$ gives the energy-momentum tensor. In a sense, the matter and energy on the right, tells spacetime how to curve, on the left. It is possible to generalise the Einstein field equations to:

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}, \quad (1.3)$$

where $g_{\mu\nu}$ gives the metric, and Λ is the infamous cosmological constant – which can be thought of as the energy of the vacuum if moved to the right hand side. Finally we have the condition that energy-momentum is preserved:

$$\nabla_\mu T^{\mu\nu} = 0, \quad (1.4)$$

where ∇_μ is the covariant derivative.

1.1.2 Deflection of Light Due to Gravity

The total travel time of a light ray between two points is given by:

$$t = \int dt = \frac{1}{c} \int dt \frac{c}{v[\mathbf{x}(\ell)]} \frac{d\ell}{dt} = \frac{1}{c} \int d\ell n[\mathbf{x}(\ell)], \quad (1.5)$$

where $v[\mathbf{x}(\ell)]$ is the velocity of light along the path and $n[\mathbf{x}(\ell)] = \frac{c}{v[\mathbf{x}(\ell)]}$ is the refractive index along the path. Using Fermat's Principle and extremising the travel time along the path in a static gravitational potential, it is possible to show that the total deflection angle is given by:

$$\hat{\boldsymbol{\alpha}} = \frac{2}{c^2} \int d\lambda \nabla_\perp \Phi, \quad (1.6)$$

where Φ is the gravitational potential, λ parameterises the path and ∇_{\perp} is the gradient perpendicular to the path. Since this result is central to this thesis a full derivation is given in Appendix A.

We can also apply the *Born Approximation* and integrate over the path that the light ray would have followed if it was undeflected. Assuming the light ray initially travels in the z -direction, the deflection angle is then given by:

$$\hat{\boldsymbol{\alpha}} = \frac{2}{c^2} \int dz \nabla_{\perp} \Phi. \quad (1.7)$$

It is shown in [30] that this approximation will have a negligible impact on cosmological parameter constraints from the two-point statistics used in this thesis.

1.2 Geometry and Distances in the Universe

In this section I give a concise overview of the cosmology relevant to this thesis. This material is standard, so unless otherwise stated I will closely follow the review discussions presented in [18, 26].

1.2.1 The Expanding Universe and Accelerating Universe

Edwin Hubble and Vesto Slipher discovered that the Universe is expanding – not static. More recently it has been discovered that the expansion of the Universe is actually accelerating [104, 109]. We describe this expansion in terms of a dimensionless scale factor, $a(t)$, which increases with time, t . The physical separation, x_{phys} , between two points separated by some comoving coordinate distance, r , is then:

$$x_{\text{phys}} = a(t)r. \quad (1.8)$$

In an expanding universe two points separated by some comoving coordinate distance, r , will move apart at a rate, v_{phys} . It is customary to normalise $a(t) = 1$ today and I will implicitly use this convention for the remainder of this thesis.

Differentiating both sides with respect to t , and assuming that the comoving separation is invariant, implies:

$$v_{\text{phys}} = H(t)x_{\text{phys}}, \quad (1.9)$$

where $H(t) = \dot{a}/a$. The value of $H(t)$ today, with $t = 0$, is called the *Hubble constant* and it is written as H_0 .

The precise value of the Hubble constant is hotly debated, but current accepted values lie in the range $65 - 75 \text{ (km/s) Mpc}^{-1}$ [2, 19, 108]. For historical reasons, we sometimes write $H_0 = h_0 100 \text{ (km/s) Mpc}^{-1}$, where h_0 is a dimensionless quantity. The subscript is also usually omitted in this case.

1.2.2 The Cosmological Principle and the Friedmann-Robertson-Walker Metric

Assuming the *Cosmological Principle* – that the Universe is spatially homogenous and isotropic on scales larger than $\sim 100 \text{ Mpc}$ – the metric for the Universe can be written in terms of the normal spherical comoving coordinates (r, θ, ϕ) , time coordinate t and an evolving expansion factor, $a(t)$, as:

$$ds^2 = -c^2 dt^2 + a^2(t) [dr^2 + f_k(r)^2 (d\theta^2 + \sin^2\theta d\phi^2)], \quad (1.10)$$

where

$$f_k(r) \equiv \begin{cases} k^{-1/2} \sin(k^{1/2}r) & \text{if } k > 0 \\ r & \text{if } k = 0 \\ (-k)^{-1/2} \sinh((-k)^{1/2}r) & \text{if } k < 0, \end{cases} \quad (1.11)$$

and $k = -1, 0, 1$ corresponds to negative, flat and positive spatial curvature respectively. This is called the *Friedmann-Robertson-Walker metric*.

1.2.3 Cosmological Redshift

The expansion of the Universe will increase the observed wavelength of light between the emitter and the observer at different cosmological times according to:

$$\lambda_0 = \frac{a(t_0)}{a(t_e)} \lambda_e. \quad (1.12)$$

where λ_0 is the wavelength observed from a source emitting at wavelength λ_e and $a(t_0)$ and $a(t_e)$ are the corresponding scale factors at the location of the observer and emitter respectively. The redshift, z , is defined as:

$$z = \frac{\lambda_0 - \lambda_e}{\lambda_e}. \quad (1.13)$$

Combining equations (1.12) and (1.13) yields:

$$1 + z = \frac{1}{a(t_e)}. \quad (1.14)$$

1.2.4 Distance Measures

The comoving distance, r , is not directly observable, but we can express it in terms of the redshift which is directly measurable. On a radial null geodesic, putting $c = 1$,

$$0 = -dt^2 + a^2(t) dr^2. \quad (1.15)$$

Rearranging, integrating and using equation (1.14) implies:

$$r(z) = \int_0^t \frac{dt'}{a(t')} = \int_0^z \frac{dz'}{H(z')} \quad (1.16)$$

The distance measure relevant to this thesis is the angular-diameter distance. It is defined as the ratio of an object's physical size relative to the angle made on the sky:

$$d_A = \frac{x_{\text{phys}}}{\theta} = \frac{f_k(r(z)) a(z) \theta}{\theta} = f_k(r) a(z), \quad (1.17)$$

so that for a flat cosmology the comoving angular diameter distance is just the comoving distance $r(z)$. For this reason I will refer to the comoving distance and comoving angular distance, $d_A/a(z)$, interchangeably for the remainder of the thesis.

1.3 Contents of the Universe

1.3.1 The Conservation Equation and The Equation of State

In general relativity the energy-momentum tensor for a homogenous and isotropic fluid is:

$$T_{\mu\nu} = \left(\rho + \frac{P}{c^2} \right) U_\mu U_\nu - P g_{\mu\nu}, \quad (1.18)$$

where ρ is the energy density and P is the pressure. Substituting into equation (1.4) yields the *conservation equation*:

$$\frac{\partial \rho}{\partial t} + 3H(\rho + P) = 0. \quad (1.19)$$

The density, ρ , pressure, P , and time-evolving Hubble parameter, H , are all implicitly functions of the redshift, z , but I will use this compact notation for the remainder of the thesis. It is now convenient to define an equation of state:

$$P = \omega \rho, \quad (1.20)$$

where we have set $c = 1$. Plugging this into equation (1.19) and solving gives:

$$\rho \propto a^{-3(1+\omega)}. \quad (1.21)$$

1.3.2 Components of the Universe

The Universe has three components: matter, radiation and vacuum energy⁴:

- **Matter:** On cosmological scales we can treat matter as a cloud of massive collisionless particles. In this approximation the pressure, P , is zero. Hence, the equation of state is also zero. It follows from equation (1.21) that the density of matter in the Universe goes as: $\rho \propto a^{-3}$.
- **Radiation:** The equation of state is: $P = -\rho/3$ [26]. Following the same argument as for matter gives: $\rho \propto a^{-4}$. Since the density of radiation falls off faster than for matter, accounting for radiation is only relevant in the very early Universe.
- **Vacuum energy:** The vacuum energy density due to the cosmological constant is time invariant i.e $\partial\rho/\partial t = 0$. Plugging this into equation (1.19) implies that $P = -\rho$ and hence that $w = -1$. Thus $\rho \propto a^0$ and the cosmological constant becomes the dominant component in the late Universe.

1.4 Kinematics of the Universe

1.4.1 The Friedmann Equation

Substituting the metric defined in equation (1.10) and the energy momentum-tensor for a perfect fluid defined in equation (1.18) into the field equations (1.2) gives the *first Friedmann equation* (see [18] for a full derivation):

$$H^2 = \frac{8\pi G}{3}\rho - \frac{k}{a^2}. \quad (1.22)$$

Observations place the value of the curvature k extremely close to zero [2]. I will set it to this value for the remainder of the thesis. For this to happen, equation (1.22)

⁴Neutrinos are also important. This is a topic in its own right so I refer the reader to the review article [83].

implies that the density of the Universe must equal the critical density:

$$\rho_{\text{crit}} = \frac{3H_0^2}{8\pi G}. \quad (1.23)$$

Dimensionless density parameters can then be defined in terms of the critical density: $\Omega_m \equiv \rho_m/\rho_{\text{crit}}$, $\Omega_r \equiv \rho_r/\rho_{\text{crit}}$ and $\Omega_\Lambda \equiv \rho_\Lambda/\rho_{\text{crit}}$. Then the Friedmann equation can be rewritten:

$$H^2 = H_0^2 [\Omega_r a^{-4} + \Omega_m a^{-3} + \Omega_\Lambda]. \quad (1.24)$$

1.4.2 Background Expansion of the Universe

For a single component, equation (1.24) can be solved analytically. In the case $\Omega_m = 1$ and $\Omega_i = 0$ for all other i , the growing mode solution is:

$$a \propto t^{2/3}, \quad (1.25)$$

giving an expanding, but not accelerating universe. For $\Omega_\Lambda = 1$, $\Lambda > 0$ and $\Omega_i = 0$ for all other components, yields the growing mode solution:

$$a \propto e^{H_0 t}, \quad (1.26)$$

corresponding to accelerating expansion, that is $\ddot{a} > 0$.

In practice the background expansion is driven by a combination of dark energy, matter and radiation so we need to use a code such as `Class` [21] or `Camb` [84] to generate a realistic expansion history, $a(t)$. These codes automatically plug this into equation (1.16) calculating all the cosmological distances relevant to this thesis.

1.4.3 Kinematics for a Generalised Dark Energy Equation of State

To search for deviations from a simple cosmological constant, it is customary to take the Chevallier-Polarski-Linder (CPL) parameterisation:

$$w(a) = w_0 + w_a(1 - a), \quad (1.27)$$

for the dark energy equation of state proposed in [86]. This particular parameterisation is not theoretically motivated and it is just a Taylor expansion about $a = 1$. This is part of the motivation for the method presented in Chapter 3. Nevertheless, either $w_a \neq 0$ or $w_0 \neq -1$ would imply that the accelerated expansion of the Universe is not caused by a cosmological constant.

If the dark energy equation of state is not constant, then the first Friedmann equation, and hence the expansion history of the universe, will also change. For a general dark energy equation of state, equation (1.19) can be re-written:

$$\frac{d \log(\rho_{\text{DE}})}{d \log(a)} = -3(1 + w_{\text{DE}}[a]). \quad (1.28)$$

This has the solution:

$$\rho_{\text{DE}}(a) \propto \exp \left[-3 \int_1^a d(\log a') (1 + w[a']) \right]. \quad (1.29)$$

Thus for a general parameterisation of the dark energy equation of state, the first Friedmann equation becomes:

$$H^2 = H_0^2 \left(\Omega_r a^{-4} + \Omega_m a^{-3} + \Omega_\Lambda \exp \left[-3 \int_1^a d(\log a') (1 + w[a']) \right] \right), \quad (1.30)$$

and from this **Camb** and **Class** compute the background expansion as before.

1.5 Structure Growth

1.5.1 Linear Structure growth

Structure growth also depends on the underlying cosmology. To gain intuition I reproduce the linear Newtonian theory below. We start with the continuity and Euler equations for an inviscid fluid:

$$\begin{aligned}\partial_t \rho + \nabla_{\mathbf{x}} \cdot (\rho \mathbf{u}) &= 0 \\ (\partial_t + \mathbf{u} \cdot \nabla_{\mathbf{x}}) \mathbf{u} &= - \frac{\nabla_{\mathbf{x}} P}{\rho} - \nabla_{\mathbf{x}} \Phi,\end{aligned}\tag{1.31}$$

and a Poisson equation for gravity:

$$\nabla_{\mathbf{x}}^2 \Phi = 4\pi G \rho,\tag{1.32}$$

where Φ is the Newtonian potential and all spatial components are in physical coordinates denoted by \mathbf{x} . Defining the fractional overdensity as:

$$\delta \equiv \frac{\rho - \bar{\rho}}{\bar{\rho}},\tag{1.33}$$

we expand equations (1.31) and (1.32) linearly about the mean values in \mathbf{u} , ρ and P . Re-expressing in comoving coordinates, and combining yields:

$$\ddot{\delta} + 2H\delta - \frac{c_s^2}{a^2} \nabla^2 \delta = 4\pi G \bar{\rho} \delta,\tag{1.34}$$

where c_s is the sound speed with the implicit assumption of adiabatic expansion, so that $c_s^2 = \frac{\delta P}{\delta \rho}$.

1.5.2 Linear Growth of Structure: A Case Study

For a dark matter dominated universe, $c_s = 0$ (since there is no pressure) and equation (1.34) offers the growing mode solution:

$$\delta \propto a. \quad (1.35)$$

Meanwhile in a Λ -dominated universe, the Laplacian of the overdensity is zero since there is no spatial clustering. Also dropping the source term, since the expansion term will dominate, this universe yields the growing mode solution:

$$\delta \propto \text{const.} \quad (1.36)$$

Together, equations (1.35) and (1.36) characterise the growth of large linear scales for most of the history of the Universe. While in the matter dominated phase, density perturbations will grow as the scale factor, but once in the Λ -dominated phase, perturbation growth will be ‘frozen out’.

1.6 Statistics of the Cosmological Density Field

1.6.1 Random Fields and Power Spectra

Fields are of fundamental importance in cosmology. It is not the field itself that tells us about the underlying cosmology, but rather the statistics of the field. This is because given a cosmology, the field we observe is one among many equally likely configurations.

Given a random field $f(\mathbf{x})$, we define the ensemble average as a functional integral over all possible field configurations:

$$\langle f(\mathbf{x}) \rangle \equiv \int \mathcal{D}f \mathbb{P}[f] f(\mathbf{x}), \quad (1.37)$$

where $\mathbb{P}[f]$ is the probability distribution function for fields f . Very often this does

not contain any information (for example, $\langle \delta \rangle = 0$) and we must go to higher-order statistics. The two-point correlation function is defined as:

$$\langle f(\mathbf{x}) f(\mathbf{x} + \mathbf{y}) \rangle \equiv \int \mathcal{D}f \mathbb{P}[f] f(\mathbf{x}) f(\mathbf{x} + \mathbf{y}), \quad (1.38)$$

which by homogeneity and isotropy is a function of $|\mathbf{y}|$ only. It is sometimes desirable to work in harmonic space since this allows us to explicitly exclude small scales which can not be theoretically modelled. As a concrete example, in normal $3D$ Euclidean space, we can take the Fourier transform of the field, given by $f(\mathbf{k})$, and compute the quantity:

$$\begin{aligned} \langle f(\mathbf{k}) f^*(\mathbf{k}') \rangle &= \left\langle \int d^3x f(\mathbf{x}) e^{i\mathbf{k}\cdot\mathbf{x}} \int d^3x' f(\mathbf{x}') e^{-i\mathbf{k}'\cdot\mathbf{x}'} \right\rangle \\ &= \int d^3x e^{i(\mathbf{k}-\mathbf{k}')\cdot\mathbf{x}} \int d^3y e^{-i\mathbf{k}'\cdot\mathbf{y}} \langle f(\mathbf{x}) f(\mathbf{x} + \mathbf{y}) \rangle \\ &= (2\pi)^3 \delta^D(\mathbf{k} - \mathbf{k}') P(k), \end{aligned} \quad (1.39)$$

where we have made the substitution $\mathbf{x}' = \mathbf{x} + \mathbf{y}$, $*$ denotes complex conjugation, δ^D denotes the $3D$ Dirac delta function and the power spectrum, $P(k)$, is defined as the Fourier transform of the two-point correlation function. This evolves over cosmic time so we write the $3D$ matter power spectrum as $P(k, z)$, where z labels the redshift.

1.6.2 Initial Conditions

To compute the power spectrum in the late Universe we must assume a set of initial conditions. Measurements from Planck [3] confirm that to a very good approximation the primordial density field is Gaussian with a power law spectrum:

$$P(k) = A_s k^{n_s}. \quad (1.40)$$

Best estimates place the value of the spectral index $n_s = 0.965 \pm 0.004$ [2]. In late Universe observational cosmology the primordial amplitude, A_s , is not commonly

used, and instead we introduce a new parameter, σ_8 , to parameterise the amplitude of the power spectrum. It is defined by:

$$\sigma_8^2 = \int dk P(k)W(k, R), \quad (1.41)$$

where

$$W(k, R) = \frac{3k^2}{2\pi^2(kR)^3} [\sin(kR) - kR \cos(kR)], \quad (1.42)$$

with $R = 8h^{-1}$ Mpc. This can be interpreted as the variance of the overdensity field, δ , on scales of $R = 8h^{-1}$ Mpc. Recent measurements place $\sigma_8 \approx 0.8$, but there is a mild tension between CMB [2] and local Universe measurements [1, 57].

1.6.3 Computing the Matter Power Spectrum

The matter power spectrum, plotted in Figure 1.1, on large linear scales can be computed using a relativistic reformulation of the Newtonian theory encapsulated in equation (1.34)⁵. In practice the public codes `Camb` and `Class` are used to compute the power spectrum in this linear regime.

In high-density regions ($k \gtrsim 10^{-1}h$ Mpc $^{-1}$), linear perturbation theory breaks down and the *linear matter power spectrum* under-predicts the true *nonlinear matter power spectrum*. At $z = 0$ and $k \gtrsim 10h$ Mpc $^{-1}$, linear theory under-predicts the nonlinear power spectrum by nearly two orders of magnitude.

On scales $k \gtrsim 1h$ Mpc $^{-1}$, baryonic physics has a large impact on the matter power spectrum [60]. For $k \lesssim 10h$ Mpc $^{-1}$, feedback from galactic nuclei (AGN) and supernovae suppresses the matter power spectrum while for $k \gtrsim 10h$ Mpc $^{-1}$ radiative cooling enhances power [60, 116].

To compute the true matter power spectrum we must resort to emulators [50, 120] or halo model [97] codes trained on N -body simulations. Still, modelling uncertainties due to nonlinear structure growth and baryonic physics will bias weak lensing parameter constraints [5, 15, 81, 118] in upcoming surveys. Mitigating this

⁵This is beyond the scope of this thesis but more details can be found in [18]

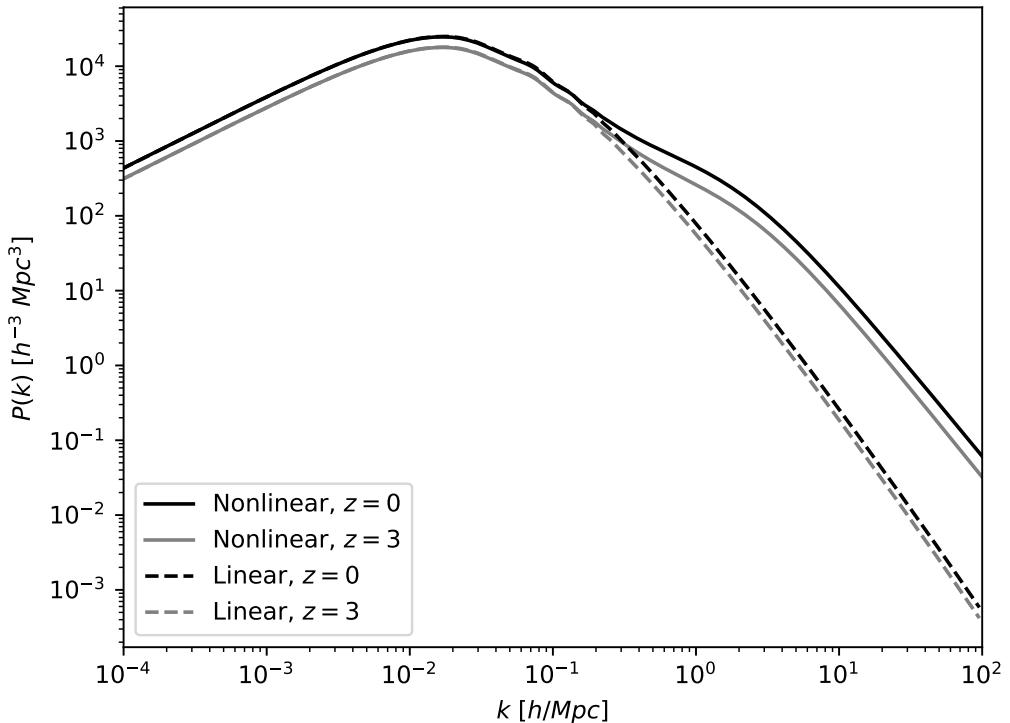


Figure 1.1: The linear and nonlinear power spectra at redshifts $z = 0$ and $z = 3$. As structure collapses through cosmic time, the amplitude of the power spectrum increases. At small scales, above $k \sim 10^{-1} h \text{ Mpc}^{-1}$, nonlinear corrections become relevant. By $k \sim 2h \text{ Mpc}^{-1}$, the linear power spectrum underpredicts the nonlinear power spectrum by an order of magnitude. Cosmic shear is primarily sensitive to structure in the range $k \in [10^{-1} h \text{ Mpc}^{-1}, 10h \text{ Mpc}^{-1}]$ [124]. The linear power spectrum was generated using `Camb` [84] and the nonlinear corrections were generated using `HALOFIT` [120] inside the modular framework `Cosmosis` [136].

bias is the focus of Chapter 2.

1.7 Cosmic Shear

In this section I outline the theory of cosmic shear. Most of the discussion is standard so I follow the discussion in the review articles [17, 70, 98] unless otherwise stated.

1.7.1 Weak Lensing

In what follows we work in a spatially flat universe. To generalise to a spatially curved universe, just replace, r , with the general comoving angular diameter distance, $f_k(r)$ [126].

Imagine that light travels through a gravitational potential, Φ , between a source and an observer. Given a source at comoving angular diameter, r , the total deflection angle in the observer's frame is:

$$\boldsymbol{\alpha} = \frac{2}{c^2} \int_0^r dr' \frac{(r - r')}{r} \nabla_{\perp} \Phi(\mathbf{x}, r'), \quad (1.43)$$

where \mathbf{x} gives the spatial coordinate along the path, assuming the Born approximation. This is a central result so a full derivation is given in Appendix B. If $\boldsymbol{\theta}$ gives the observed location of the source on the sky, spotting that ∇_{\perp} is equivalent to $1/r' \nabla_{\boldsymbol{\theta}}$, we write the components of $\boldsymbol{\alpha}$ as:

$$\alpha_i = \frac{\partial \phi}{\partial \theta_i}, \quad (1.44)$$

where the *lensing potential* is defined as:

$$\phi(\mathbf{r}) = \frac{2}{c^2} \int_0^r dr' F(r, r') \Phi(r, \boldsymbol{\theta}), \quad (1.45)$$

and the *lensing kernel* is:

$$F(r, r') = \frac{r - r'}{rr'}. \quad (1.46)$$

This is maximally efficient when the lens is halfway between the observer and the source. Writing β for the angle to the source along a straight line, we have the *lens equation*:

$$\beta = \theta - \alpha. \quad (1.47)$$

Thus the Jacobian transformation matrix is:

$$\begin{aligned} A_{ij} &\equiv \frac{\partial \beta_i}{\partial \theta_j} = \delta_{ij}^K - \frac{\partial \alpha_i}{\partial \beta_j} \\ &= \delta_{ij}^K - \frac{\partial^2 \phi}{\partial \theta_i \partial \theta_j} \end{aligned} \quad (1.48)$$

where δ_{ij}^K is the Kronecker delta. This can be rewritten as:

$$A = \begin{pmatrix} 1 - \kappa - \gamma_1 & -\gamma_2 \\ -\gamma_2 & 1 - \kappa + \gamma_1 \end{pmatrix}, \quad (1.49)$$

where the *convergence* is defined as

$$\kappa = \frac{1}{2} (\partial_1 \partial_1 + \partial_2 \partial_2) \phi, \quad (1.50)$$

and the two *shear* components are defined by:

$$\begin{aligned} \gamma_1 &= \frac{1}{2} (\partial_1 \partial_1 - \partial_2 \partial_2) \phi \\ \gamma_2 &= \partial_1 \partial_2 \phi, \end{aligned} \quad (1.51)$$

where the partial derivatives are with respect to the components of θ . We say that we are in the *weak lensing regime* if $\kappa < 1$ and in most cases $\kappa \ll 1$.

The inverse of the matrix A gives the local mapping from the source to the

image plane. It is readily found that:

$$A^{-1} = \mu(1 - \kappa) \left[\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + g \begin{pmatrix} \cos 2\psi & \sin 2\psi \\ \sin 2\psi & -\cos 2\psi \end{pmatrix} \right], \quad (1.52)$$

where the magnification, μ , is just the ratio of the size in the image plane relative to the source plane:

$$\mu = \frac{1}{\det A} = \frac{1}{(1 - \kappa)^2 - (\gamma_1^2 + \gamma_2^2)}, \quad (1.53)$$

the *reduced shear*, g , is:

$$g = \frac{\gamma}{1 - \kappa}, \quad (1.54)$$

$\gamma_1 = \gamma \cos 2\psi$ and $\gamma_2 = \gamma \sin 2\psi$, for some angle ψ . Equation (1.52) gives us a good intuition of what the shear and convergence mean physically. If $\kappa = 0$, then a circular source will be sheared into an ellipse with ellipticity (third-flattening) $e = \frac{a-b}{a+b} = |\gamma|$, where a and b are the semi-major and semi-minor axes lengths respectively. Meanwhile if $\gamma = 0$, then a circle with radius R would be mapped to a circle with radius $R/(1 - \kappa)$. The combined effect of the shear and convergence is illustrated in Figure 1.2.

1.7.2 The Shear Field

Like the fractional overdensity, $\delta(\mathbf{r})$, the shear, $\gamma(\boldsymbol{\theta}, z)$, and convergence, $\kappa(\boldsymbol{\theta}, z)$, are random fields which are sensitive to the underlying cosmology. Measuring their statistics provides cosmological constraints. Of course it is impossible to find the field values at every point in space, but it is possible to measure a galaxy's size, orientation and ellipticity. Since the intrinsic shape and size of a galaxy is unknown, this only gives a noisy estimate of the shear and convergence fields. For the remainder of this thesis, I focus on the shear signal, because it has higher signal-to-noise. For more information about cosmological inference from the convergence field see [10, 61].

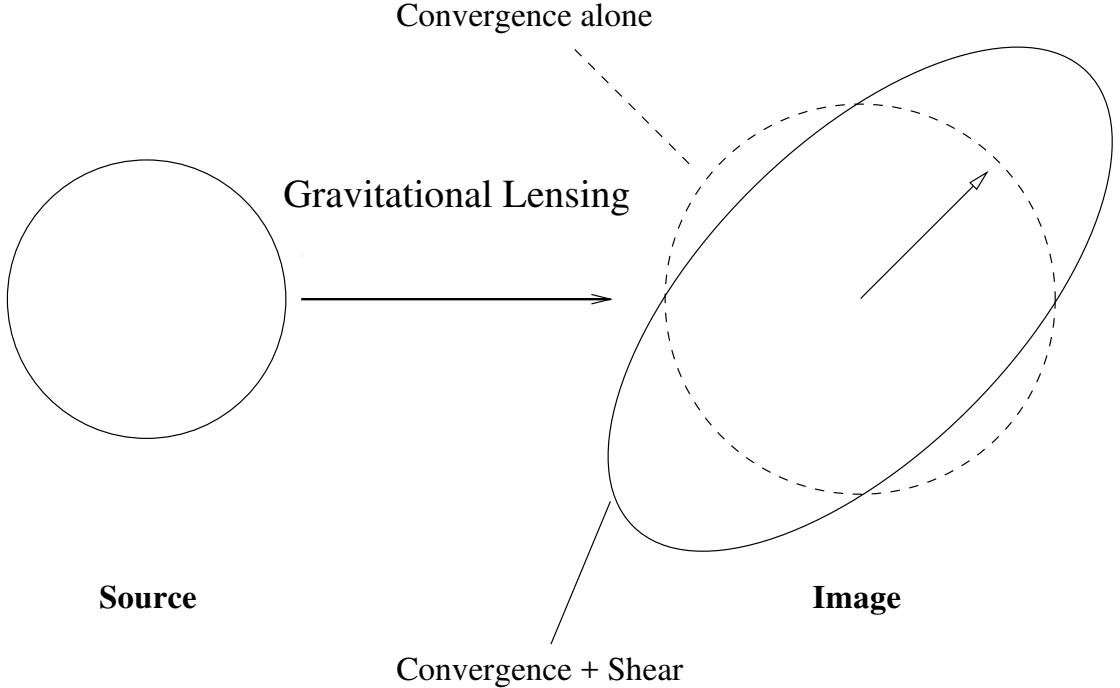


Figure 1.2: The combined effect of shear and convergence. Convergence changes the radius of a circular source while the shear changes the ellipticity. This figure originally appears in [98].

For mathematical convenience, the shear is normally written as a complex number

$$\gamma \equiv \gamma_1 + i\gamma_2. \quad (1.55)$$

We say that shear is a *spin two* field, which means that it is invariant under local rotations by π radians, since an ellipse has this symmetry. We write:

$$\gamma = e^{is\psi} |\gamma|, \quad (1.56)$$

for $s = 2$.

Defining the ellipticity of a galaxy as:

$$\epsilon = \left(\frac{a - b}{a + b} \right) e^{2i\psi}, \quad (1.57)$$

where a and b give the length of the semi-major and semi-minor axes and ψ gives the angle of orientation (see Figure 1.3), one can write the observed ellipticity, ϵ^o ,

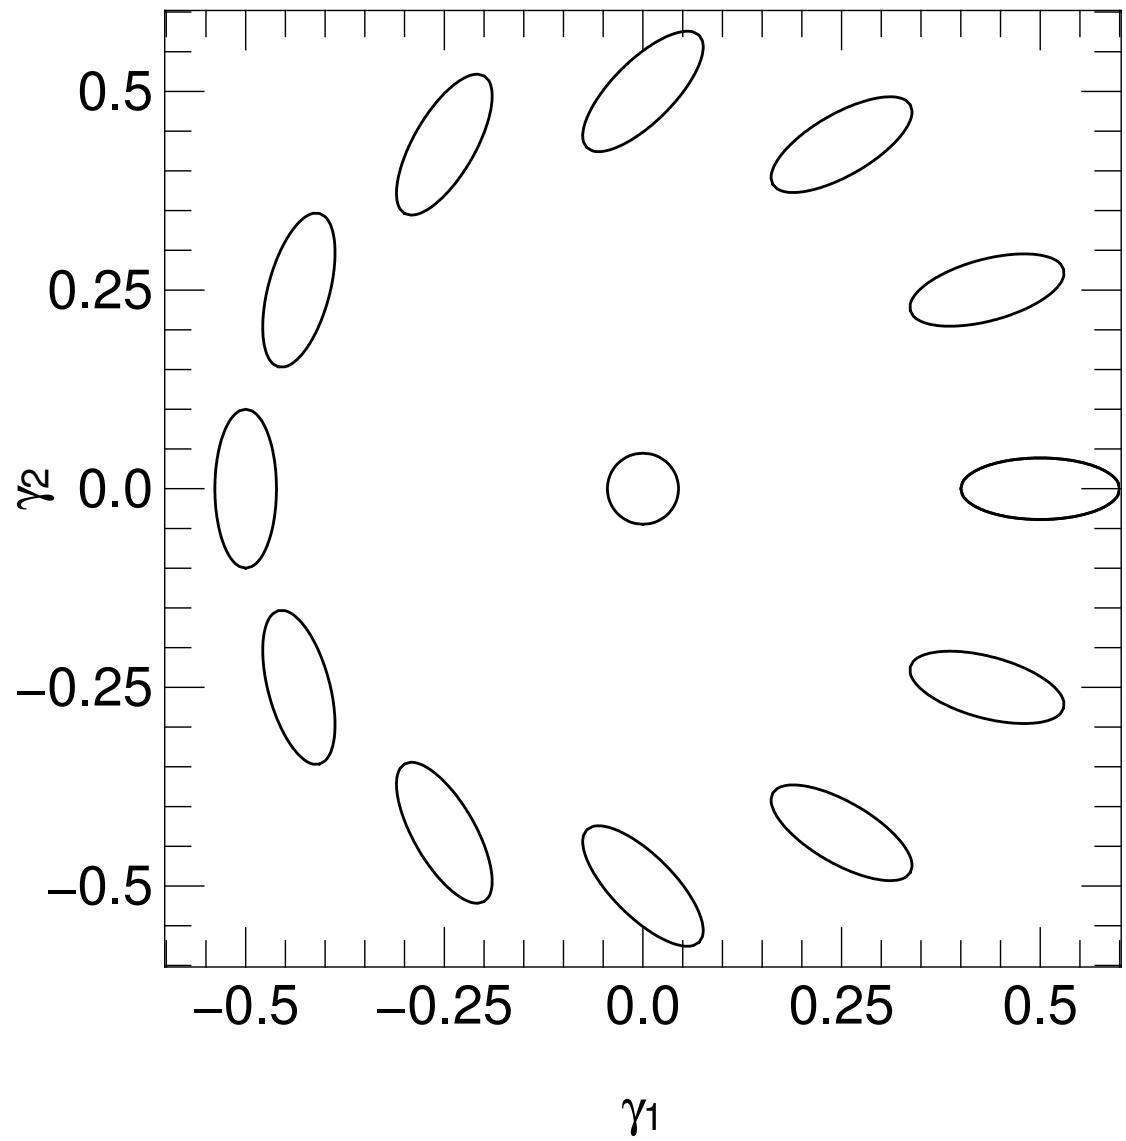


Figure 1.3: The ellipticity components γ_1 and γ_2 . As the ellipticity rotates around π radians, the orientation angle, ψ , rotates around 2π radians. This figure originally appears in [70].

in terms of the ellipticity of the source, ϵ^s , by:

$$\epsilon^o = \frac{\epsilon^s + g}{1 + g^* \epsilon^s}, \quad (1.58)$$

using equation (1.52), where $*$ denotes complex conjugation. Making the anzatz that $\kappa \ll 1$ and Taylor expanding implies:

$$\epsilon^o \approx \epsilon^s + \gamma. \quad (1.59)$$

In the absence of intrinsic alignments (see Section 1.7.7), $\langle \epsilon^s \rangle = 0$ so that the observed ellipticity is an unbiased estimator for the shear: $\langle \epsilon^o \rangle = \gamma$. In practice the typical shear induced by large scale structure is approximately 100 times smaller than the intrinsic shape dispersion, thus weak lensing surveys must contain a large number of galaxies to extract cosmological information.

1.7.3 Shape and Distance Measurements

To perform inference with real photometric data, we must generate a catalog containing the ellipticities and redshifts of all the galaxies in the survey. Measuring the photometric redshift (photo-z) and shape of a galaxy are topics in their own right, so I aim to only give a brief overview of these topics in this section.

Shape measurement algorithms fall into two broad categories: model fitting and free-form algorithms. In the model fitting case, parametric profiles are fit to the light profiles of each galaxy. Meanwhile in the free-form case we can define the second moments of the brightness as:

$$Q_{ij} = \frac{\int d^2\boldsymbol{\theta} w [I(\boldsymbol{\theta})] (\boldsymbol{\theta}_i - \bar{\boldsymbol{\theta}}_i)(\boldsymbol{\theta}_j - \bar{\boldsymbol{\theta}}_j)}{\int d^2\boldsymbol{\theta} w [I(\boldsymbol{\theta})]}, \quad (1.60)$$

where w is a weight, $I(\boldsymbol{\theta})$ is the surface brightness at $\boldsymbol{\theta}$, $\bar{\boldsymbol{\theta}}$ is the centre of brightness and $i, j \in \{1, 2\}$. The two components of the observed ellipticity are then given

by:

$$\begin{pmatrix} \epsilon_1^o \\ \epsilon_2^o \end{pmatrix} = \frac{1}{Q_{11} + Q_{22} + 2(Q_{11}Q_{22} - Q_{12}^2)^{1/2}} \begin{pmatrix} Q_{11} - Q_{22} \\ 2Q_{12} \end{pmatrix}. \quad (1.61)$$

The point spread function (PSF) of the telescope and charge transfer inefficiency (CTI) in the charge-coupled devices (CCDs) are the dominant systematics in shape estimation [95]. Thorough reviews of shape measurement techniques are given in [52, 75, 94, 100].

It is not possible to image the entire sky with a high-resolution spectrographic instrument, so we must estimate the redshift of galaxies using a relatively small number of photometric bands. Methods for photo-z estimation and requirements for Stage IV surveys are discussed in [81, 90, 96, 101] and the references therein. Acquiring a representative large spectroscopic sample of galaxies to train the photo-z estimation algorithm is one of the primary challenges [96]. While in practice each galaxy will have a different posterior redshift distribution, in Chapter 2 and Chapter 4, I account for the photometric redshift uncertainty in each tomographic bin by smoothing by the Gaussian kernel:

$$p(z|z_p) \equiv \frac{1}{2\pi\sigma_z(z_p)} e^{-\frac{(z-c_{\text{cal}}z_p+z_{\text{bias}})^2}{2\sigma_{z_p}^2}}, \quad (1.62)$$

with $c_{\text{cal}} = 1$, $z_{\text{bias}} = 0$ and $\sigma_{z_p} = A(1+z_p)$ with $A = 0.05$ [62].

1.7.4 The Generalised Lensing Spectrum

One can define the lensing spectrum, which contains the two-point information of the observed lensing shear field in harmonic space, in much the same way that the power spectrum does for the overdensity field $\delta(\mathbf{r})$. The generalised lensing spectrum, $C_\ell(\eta, \eta')$, is defined as:

$$C_\ell(\eta, \eta') = \langle \gamma_{\ell m}(\eta) \gamma_{\ell' m'}^*(\eta') \rangle \delta_{\ell \ell'}^K \delta_{m m'}^K, \quad (1.63)$$

where δ^K denotes the Kronecker-delta and in [124] I defined the *generalised-spherical harmonic coefficients*, $\gamma_{\ell m}$, as:

$$\gamma_{\ell m}(\eta) = \sqrt{\frac{2}{\pi}} \sum_g \gamma_g(r_g, \boldsymbol{\theta}_g) W_\ell(\eta, r_g) {}_2Y_{\ell m}(\boldsymbol{\theta}_g), \quad (1.64)$$

where the sum is over all galaxies g with angular coordinate $\boldsymbol{\theta}_g$ and radial coordinate r_g , W_ℓ is a weight, η is an as yet arbitrary label and ${}_2Y_{\ell m}$ are the spin-2 spherical harmonics. It can be shown (see Appendix C) that:

$$C_\ell^{\gamma\gamma}(\eta_1, \eta_2) = \frac{9\Omega_m^2 H_0^4}{16\pi^4 c^4} \frac{(\ell+2)!}{(\ell-2)!} \int \frac{dk}{k^2} G_\ell^\gamma(\eta_1, k) G_\ell^\gamma(\eta_2, k), \quad (1.65)$$

where the G-matrix is:

$$G_\ell^\gamma(\eta, k) \equiv \int dz_p dz' n(z_p) p(z'|z_p) W_\ell(\eta, r[z']) U_\ell(r[z'], k), \quad (1.66)$$

and the U -matrix, which contains the cosmological information, is:

$$U_\ell(r[z], k) \equiv \int_0^r dr' \frac{F_K(r, r')}{a(r')} j_\ell(kr') P^{1/2}(k; r'), \quad (1.67)$$

where $n(z)$ is the radial distribution function, $p(z|z_p)$, defined in equation (1.62), gives the photometric error and $j_\ell(kr)$ are the spherical Bessel functions.

We can make the Limber approximation derived in [88], which amounts to making the replacement:

$$j_\ell(x) \rightarrow \sqrt{\frac{\pi}{2x}} \delta(x - (\ell + 1/2)). \quad (1.68)$$

In this approximation the U -matrix becomes:

$$U_\ell(r, k) = \frac{F_k(r, \nu(k))}{ka(\nu(k))} \sqrt{\frac{\pi}{2(\ell + 1/2)}} P^{1/2}(k, \nu(k)), \quad (1.69)$$

where $\nu(k) \equiv \frac{\ell+1/2}{k}$. This is a good approximation for $\ell > 100$ [76].

1.7.5 The Tomographic Lensing Spectrum

There are many possible choices for the weight function, $W_\ell(\eta, r)$. Taking it to be a spherical Bessel function, $j_\ell(kr)$, yields the spherical Bessel cosmic shear spectrum [48]. More recently [77] considered taking the weights in photometric colour space rather than the usual redshift space. However, by far the most popular choice to extract the radial information of the shear field is to choose the weight to be a top hat function in redshift space:

$$W_i(z) \equiv \begin{cases} 1 & \text{if } z \in [z_i, z_{i+1}] \\ 0 & \text{otherwise,} \end{cases} \quad (1.70)$$

where z_i and z_{i+1} give the boundaries of the tomographic bin i . This was first proposed in [59] and is referred to as *tomographic cosmic shear*. Multiplying the prefactor in equation (1.65) by $8\pi^3$ (this is valid provided the shot-noise, defined in the next section, is appropriately rescaled), making the substitution $k = \ell/r$ [76] and mapping $l + 1/2 \rightarrow \ell$ in the Limber approximation yields the usual expression for the tomographic lensing spectrum:

$$C_\ell^{ij} = \int_0^{r_H} dr \frac{q_i(r)q_j(r)}{r^2} P\left(\frac{\ell}{r}, r\right), \quad (1.71)$$

where the lensing efficiency kernel, q_i is defined as:

$$q_i(r) = \frac{3H_0^2\Omega_m}{2c^2} \frac{r}{a(r)} \int_r^{r_H} dr' n_i(r') \frac{r' - r}{r'}, \quad (1.72)$$

and the indices i and j denote the tomographic bin numbers.

1.7.6 Shot-Noise

The intrinsic dispersion of galaxies' ellipticities, σ_e^2 , adds an extra *shot-noise* term to the lensing spectrum. In the case of generalised-spherical cosmic shear, the

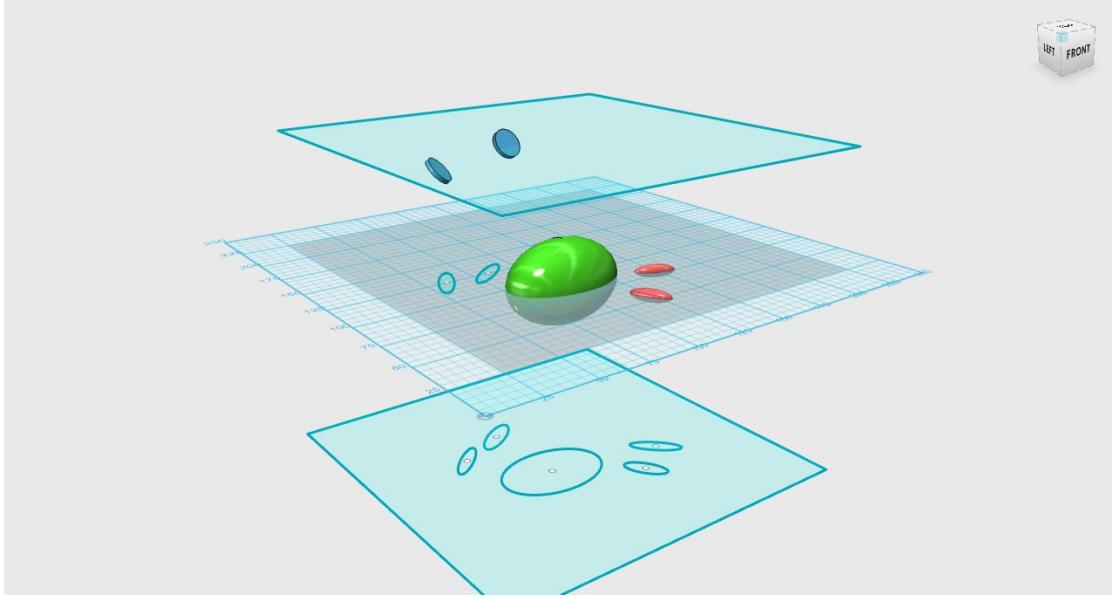


Figure 1.4: The green ellipsoid represents a large dark matter halo. Nearby galaxies (red) become tidally aligned with the dark matter halo while the observed ellipticity of background galaxies (blue) is gravitationally sheared. The intrinsic alignments of galaxies around dark matter halos is orthogonal to the shearing effect which suppresses the cosmic shear signal. This figure originally appears in [65].

shot-noise is given by:

$$N_\ell(\eta, \eta') = \frac{\sigma_e^2}{2\pi^2 A n_{\text{eff}}} \int dz n(z) W_\ell(\eta, r[z]) W_\ell(\eta', r[z]), \quad (1.73)$$

where A gives the area of the survey and n_{eff} gives the number density of galaxies in the survey. I take $\sigma_e = 0.3$ for the remainder of the thesis [23]. I derive this result in Appendix D. In practice it is easier to just compute the shot-noise from noise-only maps as in [53]. This seamlessly handles any pixelisation effects with real data.

1.7.7 Intrinsic Alignments

Galaxies that form near large dark matter halos will tend to be tidally aligned radially with the halo. This is shown in Figure 1.4. The large green ellipsoid represents a massive dark matter halo and two nearby red galaxies are intrinsically aligned with the massive halo. Meanwhile the two background blue galaxies are

gravitationally sheared perpendicularly to the green halo. *Intrinsic alignments* (IA) are an important effect in weak lensing studies and since the effect is in a perpendicular direction to cosmic shear, intrinsic alignments dampen the cosmic shear signal. To account for this, two additional terms must be added to the cosmic shear power spectrum. A ‘GI’ term accounts for the cross-correlation between cosmic shear and intrinsic alignments and an ‘II’ term accounts for the auto-correlation of intrinsic alignments.

In the simple alignment model presented in [58], which assumes that galaxies are distorted by the local gravitational field [28], the IA signal is:

$$\gamma^I = -\frac{\mathcal{A}_I C_1}{4\pi G} \left(\frac{1+z}{1+z_0} \right)^\eta (\partial_1^2 - \partial_2^2, 2\partial_1\partial_2) \Phi_p, \quad (1.74)$$

where Φ_p denotes the gravitational potential at the time of galaxy formation (when the intrinsic alignments were imprinted), $C_1 = 5 \times 10^{-14} h^{-2} M_\odot^{-1} \text{Mpc}^3$ is chosen so that the value of the dimensionless intrinsic alignment amplitude, A_I , is approximately unity [25]. The redshift amplitude term introduced in [91] is added to account for the fact that the strength of the intrinsic alignment may have evolved overtime. It is customary [130] to choose the pivot redshift, z_0 , to be the mean redshift of the survey so that the redshift correction is near unity. I will take this choice throughout the remainder of the thesis. From the Poisson equation we have:

$$\nabla_x^2 \Phi_p = 4\pi G \rho, \quad (1.75)$$

where x denotes physical coordinates. Using the relation $\nabla_x = a^{-1} \nabla_r$ and Fourier transforming gives:

$$\begin{aligned} \Phi_p &= -4\pi G \rho(k) k^{-2} a^2 \\ &= -4\pi G \frac{\rho_{\text{crit}} \Omega_m}{D(z)} k^{-2} a^2 \delta_{\text{lin}}(k, z=0), \end{aligned} \quad (1.76)$$

where we have used the definition of the linear overdensity field noting that the average density of the Universe, $\bar{\rho}$, is related to the critical density, ρ_{crit} , by a factor

of Ω_m , and the dimensionless growth factor normalised to unity today, $D(z)$, is defined by:

$$P_{\text{lin}}(k, z) = D^2(z) P_{\text{lin}}(k, z=0), \quad (1.77)$$

where P_{lin} is the linear matter power spectrum. Replacing the linear overdensity, δ_{lin} , with the true nonlinear overdensity field, δ , plugging equation (1.76) into equation (1.74) and following the derivation of the shear spectrum, noting in particular the similarity between equations (1.76) and (C.9), it follows that if we define the II matter power spectrum as:

$$P_{\text{II}}(k, z) = F^2(z) P(k, z), \quad (1.78)$$

where

$$F(z) = -A_I C_1 \rho_{\text{crit}} \frac{\Omega_m}{D(z)} \left(\frac{1+z(r)}{1+z_0} \right)^\eta, \quad (1.79)$$

we can write the II spectrum as:

$$C_{\ell, \text{II}}^{ij} = \int_0^{r_{\text{H}}} dr \frac{n_i(r) n_j(r)}{r^2} P_{\text{II}} \left(\frac{\ell}{r}, r \right). \quad (1.80)$$

Similarly if we define the GI matter power spectrum as:

$$P_{\text{GI}}(k, z) = F(z) P(k, z), \quad (1.81)$$

the GI spectrum is⁶:

$$C_{\ell, \text{GI}}^{ij} = \int_0^{r_{\text{H}}} dr \frac{q_i(r) n_j(r) + n_i(r) q_j(r)}{r^2} P_{\text{GI}} \left(\frac{\ell}{r}, r \right). \quad (1.82)$$

We see that in the future we can use the rule that every time there is an II contribution we use the kernels $n_i(r) n_j(r)$ and every time there is a GI contribution we use the kernels $q_i(r) n_j(r) + n_i(r) q_j(r)$. Finally, the theoretical lensing spectrum,

⁶This implicitly includes the ‘IG’ term

$C_\ell^{T,ij}$, is given by the sum over the three contributions:

$$C_\ell^T = C_{\ell,GG}^{ij} + C_{\ell,GI}^{ij} + C_{\ell,II}^{ij}. \quad (1.83)$$

I use **Cosmosis** [51, 136] to calculate all intrinsic alignments in this thesis.

Much like the study of photometric redshift estimation and shape measurements, intrinsic alignments are a subject in their own right. The model presented here was successfully used in [130], but for Stage IV experiments a more sophisticated model which takes into account different galaxy types [111] and higher order effects will need to be developed. See [65] for a review.

1.7.8 E/B-Modes

Since the shear components γ_1 and γ_2 , defined in equation (1.51), originate from an underlying potential ϕ , they are not independent – and not all field configurations are possible. Defining a new vector field:

$$\mathbf{u} \equiv \nabla \kappa, \quad (1.84)$$

for the convergence κ , it immediately follows that the curl of \mathbf{u} is zero, i.e $\nabla \wedge \mathbf{u} = 0$. In combination with the definitions in equations (1.50) and (1.51), this implies the constraint:

$$(\partial_1 - \partial_2)\gamma_2 = 0 = \partial_1 \partial_2 \gamma_1. \quad (1.85)$$

We say that gravitational lensing only induces a curl-free or *E-mode* field. Meanwhile the curl components of the observed field are called *B-modes*. Full expressions for the *E* and *B*-modes on the sphere are given in [74]. Checking for *B*-modes is a useful test to check for systematics in the data – although by no means does their absence imply that the data is bias free. Both types of modes are illustrated in Figure 1.5.

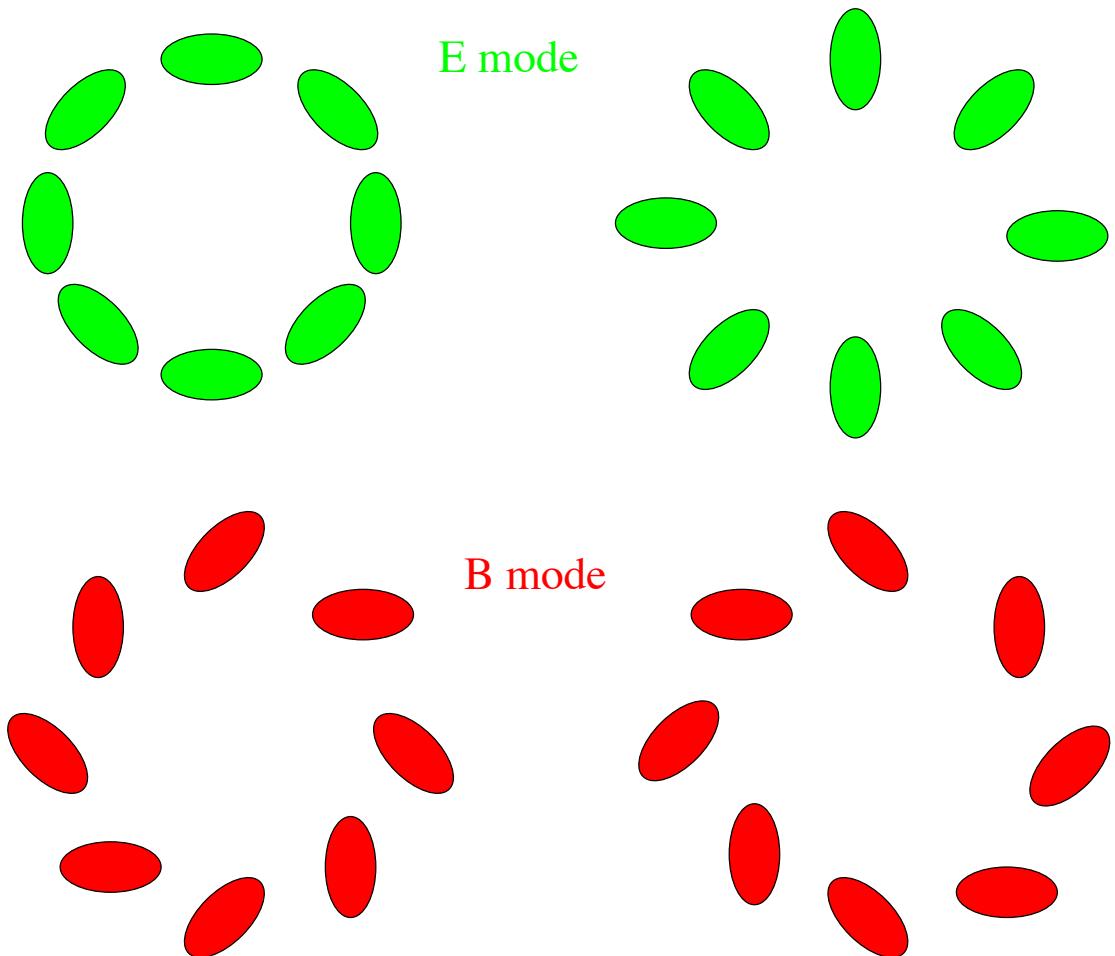


Figure 1.5: Gravitational lensing only induces E -modes like those in the top left corner. Meanwhile the intrinsic tidal alignment of nearby galaxies (see Section 1.7.7) cause E -modes like those in the top right corner. Measuring B -modes, as shown in red, would suggest the presence of a systematic in the data. This figure originally appears in [47].

1.7.9 Current and Future Data

Since the shot-noise in weak lensing surveys goes as $1/N_g$, where N_g is the number of galaxies in the survey, the ‘constraining power’ of these surveys naïvely goes as $N_g^{1/2}$. The quality of the shape measurements are also important, and in general the quality of ground based images are degraded by the atmosphere.

From the first detections of cosmic shear in 2000-2001 [16, 107, 134], studies have entered the realm of precision cosmology [51, 53, 57, 73, 130]. Some of the work in this paper uses Canada-France-Hawaii Telescope Lensing Survey (CFHTLenS) data. This was a Stage II experiment that imaged ~ 4 million galaxies over ~ 120 deg 2 [51]. Today’s leading surveys are: the Dark Energy Survey (DES) [130], the Kilo-Degree Survey (KiDS) [56] and the Subaru Hyper Suprime-Cam Lensing Survey [53]. These are referred to as *Stage III surveys* and they all use ground based imaging. Today’s largest public data set, the DES Year 1 Release, covers $\sim 1,500$ deg 2 with ~ 25 million galaxies. DES will eventually cover 5,000 deg 2 , with an effective number density of 10 galaxies per arcmin 2 , for a total of ~ 180 million galaxies.

The next generation of surveys becoming operational from the early to mid 2020s are called *Stage IV* experiments. These include: Euclid [81], Wide-Field Infrared Space Telescope (WFIRST) [118] and the Large Synoptic Survey Telescope (LSST) [15]. Euclid, a European Space Agency Mission, with a launch date of 2022 will measure the shapes of ~ 1.5 billion galaxies over $\sim 15,000$ deg 2 in its wide-field survey. LSST, a ground based telescope expected to be operational in 2021, will measure the shape of 4 billion galaxies over $\sim 20,000$ deg 2 . WFIRST is a NASA mission expected to launch by the mid-2020s. The High Latitude Imaging Survey will measure the shapes of 500 million galaxies over $\sim 2,000$ deg 2 . Developing the tools to take advantage of these next generation data sets is the focus for the remainder of this thesis.

Chapter 2

k-cut Cosmic Shear

In this chapter I outline the *k*-cut cosmic shear technique, which I presented in [122], closely following the presentation therein. The original idea for *k*-cut cosmic shear is attributable to Francis Bernardeau, but I have realised the first implementation and suggested using the mean redshift of the bin for the cut rather than the minimum redshift. This means that we can still get information from the lowest redshift bin (see below for more details). The code used to compute the Bernardeau-Nishimichi-Taruya (BNT) Transform in this chapter was written by Francis Bernardeau, but the remainder of the work is my own. For the remainder of this chapter I assume the radial distribution function:

$$n(z) = (z/z_e)^2 \exp \left[-(z/z_e)^{3/2} \right], \quad (2.1)$$

with $z_e = 0.9/\sqrt{2}$ and 30 galaxies per arcmin² over 15,000 deg², mimicking the Euclid survey. I will also drop the tomographic index labels and just write C_ℓ for C_ℓ^{ij} .

2.1 Motivation

Cosmic shear is highly sensitive to the small scale behaviour of the matter power spectrum – down to $k = 7 h \text{ Mpc}^{-1}$ [124]. This is both a blessing and a curse.

Sensitivity to these scales gives us information which can not be gleamed from other cosmological probes. However if left unchecked, poor theoretical models for these scales — due to baryonic physics (primarily feedback from Active galactic Nuclei (AGN) and the radiative cooling of gas [116]) and nonlinear structure formation — will bias parameter constraints in upcoming surveys. I refer to this as the *small scale sensitivity problem*.

In today’s experiments the matter power spectrum is usually computed with an emulator [82, 121] or halo model code [97] calibrated on N -body simulations. However, current state-of-the-art emulator codes are not sufficiently accurate for Stage IV lensing surveys [33, 60, 97, 115, 124]. To make matters worse, the power spectra from high resolution N -body simulations are in disagreement (refer to Figure 1 in [60]). If these differences were propagated through to the cosmological parameter constraints, the bias would be similar to the size of the statistical standard deviation (refer to Figure 4 in [60]).

Two possible methods are compared in [60] which deal with this issue:

- marginalising out principal components that encapsulate the difference between the full power spectrum and the dark matter only spectrum,
- marginalising over free halo parameters inside a halo model code.

Both methods fail to achieve unbiased results in some circumstances [60] and in the former case it is difficult to know how many principle components to exclude without knowing the true underlying matter power spectrum of the Universe.

Furthermore, even if a large suite of extremely accurate N -body simulations could be produced, it is computationally infeasible to run the large number required to test *all* theories of gravity, without using the untested assumption that nonlinear and baryonic feedback is cosmology and model independent. Indeed [93] used the halo model code [97] trained on Λ CDM N -body simulations to generate theoretical lensing spectra when testing the Horndeski class of gravity with shear data.

I propose a cleaner geometric solution to the small scale sensitivity problem that cuts the lensing spectrum’s sensitivity to small scale structure (corresponding to high- k in the matter power spectrum). In linear theory each k -mode evolves independently, but nonlinear and baryonic corrections couple k -modes smearing modelling errors across a wide range in k . This is why the accuracy of leading matter power spectrum emulators and halo model codes only vary slowly across a large range in k . For example, the stated accuracy of **HALOFIT** [120] is 5% for $k \leq 1h \text{ Mpc}^{-1}$ and 10% for $k \leq 10h \text{ Mpc}^{-1}$. Meanwhile **COSMIC EMU** [50] report 4% accuracy for $k \in [0.1h \text{ Mpc}^{-1}, 10h \text{ Mpc}^{-1}]$ and **HMCode** [97] report 5% accuracy for $k \in [0.1h \text{ Mpc}^{-1}, 10h \text{ Mpc}^{-1}]$. Nevertheless it is generally the case that small k -modes are modelled less accurately than large- k and in the absence of an accurate physical model, it is these scales that we should aim to cut from the shear analysis.

I refer to the procedure presented in this chapter as *k -cut cosmic shear*. It works by taking a redshift-dependent ℓ -cut after applying the Bernardeau-Nishimichi-Taruya (BNT) nulling scheme [20] which reorganises the information in the lensing spectrum to make the relationship between the angular scale, ℓ , and the structure scale, k , precise. Together this allows us to remove sensitivity to large- k . In the remainder of this chapter I develop the theory and quantify the method’s effectiveness at removing sensitivity to small scales before computing the Fisher error on the dark energy equation of state, w_0 , for different k -cuts.

2.2 The Shell Universe Thought Experiment

Suppose we lived in a ‘shell universe’ – where all the matter lay at a comoving distance r . Further imagine we want to remove contributions from structure smaller than some scale (denoted by a k -mode) from the projected lensing spectrum, C_ℓ . The Limber relation (see equation (1.69)) tells us that we could simply cut angular scales $\ell > kr$. Unfortunately in the real Universe, the lensing kernel is broad

(see top plot in Figure 2.1), so lenses across a wide range of distances and scales contribute power to the same ℓ -mode [124]. This means such a strategy will not work by itself. The information in the lensing spectrum must first be reorganized, before taking an ℓ -cut.

2.3 The Three-plane Solution

I now review the three-plane BNT [20] transform which re-weights the standard tomographic C_ℓ so that each tomographic bin contains information only about the *lenses* – rather than the *sources* – inside a small redshift range. This is generalised to an arbitrary number of source planes in the next section. From there it is a simple extension to apply the Limber argument in each bin to cut sensitivity to large- k , as outlined in the previous section.

Suppose there are a discrete number of source planes at radial distances r_i . The weighted convergence, $\tilde{\kappa}$, is defined:

$$\tilde{\kappa} = \frac{3\Omega_m H_0^2}{2c^2} \int_0^{r_i} dr \frac{\delta(r)}{a(r)} w(r), \quad (2.2)$$

where $\delta(r)$ and $a(r)$ are the local matter overdensity and scale factor of the infinitesimal lens at the radial distance r , respectively. The kernel is defined as:

$$w(r) = \sum_{i, r_i > r} p_i \frac{r_i - r}{r_i}, \quad (2.3)$$

where $\{p_i\}$ are a set of weights [20].

Assuming that there are just three discrete source planes: $r_1 < r_2 < r_3$, the key step in the BNT nulling scheme is to construct constant weights p_i so that $w(r) = 0$ for $r < r_1$. Clearly ‘lenses’ at $r' > r_3$ do not contribute to $\tilde{\kappa}$. Thus the weighted convergence is only sensitive to lenses that lie in the radial range $r \in [r_1, r_3]$. In detail, spotting that the expression for the weight, $w(r)$, can be

rewritten:

$$w(r) = r^2 \left(\frac{1}{r} \sum_{i, r_i > r} p_i - \sum_{i, r_i > r} \frac{p_i}{r_i} \right), \quad (2.4)$$

we require:

$$\sum_i p_i = 0, \text{ and } \sum_i \frac{p_i}{r_i} = 0. \quad (2.5)$$

This has solution:

$$p_2/p_1 = c(2, 3, 1)/c(1, 2, 3), \text{ and } p_3/p_1 = c(3, 1, 2)/c(1, 2, 3), \quad (2.6)$$

where

$$c(i, j, k) = r_i(r_j - r_k). \quad (2.7)$$

I refer to this as *the three-plane solution*.

2.4 Constructing the Bernardeau-Nishimichi-Taruya (BNT) Transform

For a set of N source planes the solution is no longer unique, but we can nevertheless find a solution. Given source planes at $r_1 < r_2 \dots < r_N$, which we take to be at the mean redshift of each tomographic bins, and imposing the conditions:

$$p_1^1 = 1, \quad p_2^1 = -1, \quad \text{and} \quad p_2^2 = 1, \quad (2.8)$$

constructing the weights iteratively using the three plane solution, yields the solution:

$$p_a^a = 1$$

$$p_a^{a-1} = c(a-1, a-2, a)/c(a-2, a, a-1)$$

$$p_a^{a-2} = c(a-2, a-1, a)/c(a-2, a, a-1).$$

It is worth noting that the first tomographic bin is not transformed with this choice. The BNT weighted lensing efficiency kernel defined as:

$$q_a(r) = \sum_i p_a^i q_i(r) \quad (2.9)$$

is shown in Figure 2.1, for the 10 tomographic bin case.

Defining the constant transformation matrix, M , by $M_{ia} = p_a^i$, and taking it to be the weight in the generalise lensing spectrum (see equation (1.65)) and the shot-noise (see equation (1.73)), we find that the BNT lensing spectrum is: $\tilde{C}_\ell = M C_\ell M^T$ ¹ and the shot-noise is: $\tilde{N}_\ell = M N_\ell M^T$, where C_ℓ and N_ℓ are the lensing and shot-noise spectra for normal tomographic cosmic shear. The matrix, M , has $\det(M) = 1$, and thus is a volume preserving transform implying that the total signal-to-noise remains unchanged. I compute the transformation matrix using a code written by Francis Bernardeau.

2.5 *k*-cut Cosmic Shear

In the analysis in this thesis I compute the BNT weight matrix, M , for 10 tomographic bins with the same number of galaxies per bin. Then for each BNT reweighed bin, i , I remove sensitivity to the matter power spectrum above some k_{cut} by cutting $\ell > k_{\text{cut}} r^i$, where r^i is the mean distance to each reweighted bin. Although this means that I do not entirely cut sensitivity above a given k_{cut} , it means that I do not have to cut the first bin entirely. I must also assume a fiducial cosmology to go from the redshift, z , to comoving distance $r(z)$. I assume a w_0 CDM cosmology and take the fiducial parameter values as: $(\Omega_m, \Omega_k, w_0, \Omega_b, h_0, n_s, A_s, \tau) = (0.32, 0.0, -1.0, 0.04, 0.67, 0.96, 2.1 \times 10^9, 0.08)$. Finally in the cross-correlation between bins I take whichever ℓ -cut is smaller. I refer to the joint procedure of BNT weighting and applying a lens-redshift dependent angular scale cut as *k-cut*

¹This applies to both the shear and convergence spectrum because the two full-sky spectra are related by $C_\ell^{\kappa\kappa} = \frac{\ell^2(\ell+1)^2}{(\ell+2)(\ell+1)\ell(\ell-1)} C_\ell^{\gamma\gamma}$. See [27].

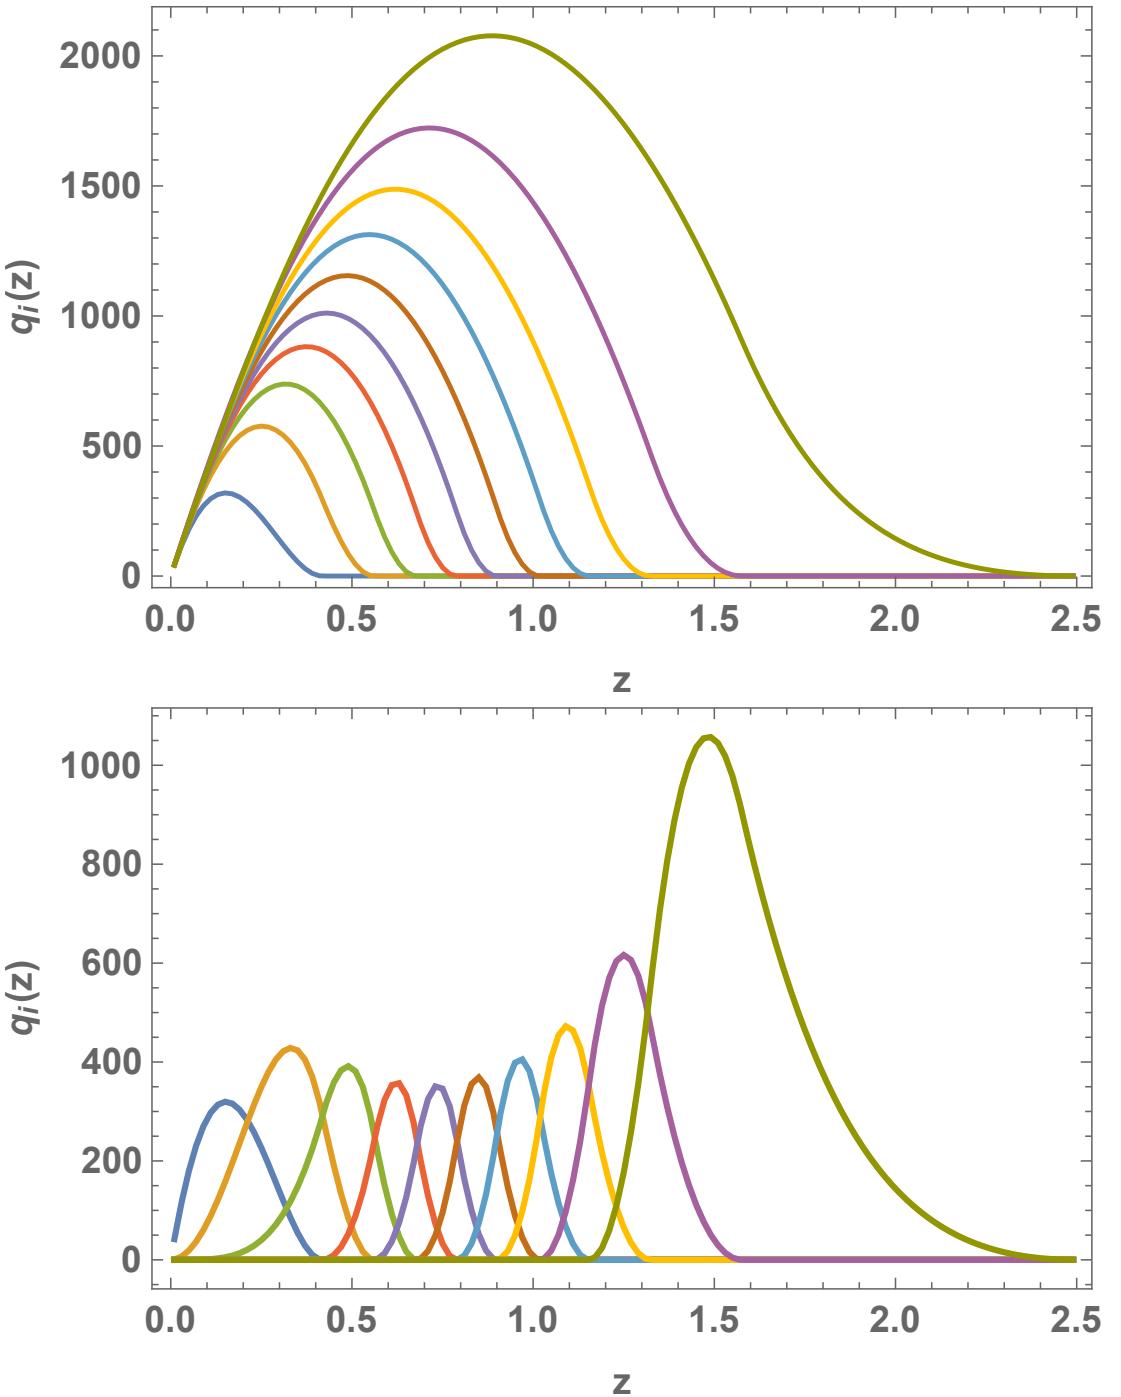


Figure 2.1: **Top:** The lensing efficiency kernels, $q_i(r)$, defined in (1.72) for standard 10-bin tomography. These are broad in z which physically means that lensing is sensitive to lenses across a wide range of redshifts. **Bottom:** The same kernels after BNT re-weighting (see equation 2.9). The kernels are much narrower in z , and bounded below, which means that it is possible to find a much more precise relationship between the angular wave-mode, ℓ , and structure wave-mode, k . This figure was produced using a script originally written by Francis Bernardeau that I have modified slightly.

cosmic shear.

2.6 Fisher Matrix Formalism

I now review the Fisher matrix formalism that I will use to evaluate the sensitivity of the standard C_ℓ analysis, BNT reweighted cosmic shear, and k -cut cosmic shear to regions of the matter power spectrum in k - z space and compare constraints on the dark energy equation of state, w_0 .

For a set of parameters, p , the Fisher matrix formalism allows us to determine how tightly the parameters will be constrained *before* the experiment is conducted [127]. Assuming a Gaussian likelihood (this is a good assumption for cosmic shear, see Chapter 4) centred about a fiducial cosmology $\{p_*$, the Fisher matrix is defined as:

$$F_{\alpha\beta} = - \left\langle \frac{\partial^2 \mathcal{L}_*}{\partial p_\alpha \partial p_\beta} \right\rangle, \quad (2.10)$$

where \mathcal{L}_* is the log-likelihood centred about the fiducial cosmology. Since this is a measure of how quickly the likelihood falls off around the maximum likelihood point, its inverse, F^{-1} , is a good estimate for the parameter covariance matrix [127], C , and we have:

$$C = F^{-1}, \quad (2.11)$$

so that the anticipated error on a parameter p marginalised over all other parameters is given by:

$$\sigma(p) = \sqrt{(F^{-1})_{pp}}. \quad (2.12)$$

Assuming an isotropic Gaussian field, the expression for the Fisher matrix in a single redshift bin is found in [127]. Taking into account the cross-correlation between redshift bins and defining S_ℓ as the sum of the signal and the noise, $S_\ell = C_\ell + N_\ell$, the Fisher matrix for cosmic shear is:

$$F_{\alpha\beta} = \sum_\ell \frac{2\ell + 1}{2} \text{Tr} [S_\ell^{-1} C_{\ell,\alpha} S_\ell^{-1} C_{\ell,\beta}], \quad (2.13)$$

where $C_{\ell,\alpha}$ denotes the derivative with respect to parameter p_α . I compute the lensing spectra using my code **GLaSS** [126] which is integrated into **Cosmosis** [136], and I choose not to include intrinsic alignments for now.

To measure the sensitivity of cosmic shear to the matter power spectrum, I follow the analysis in my paper [124] which I now review (see [39] for an alternative approach that suffers from interpolation errors). First, I divide the matter power spectrum $P(k, z)$, into logarithmically and linearly spaced grid cells in k and z , respectively. I compute the fractional amplitude change inside each grid cell g :

$$P_g(k, z, \mathcal{A}) \equiv \begin{cases} (1 + \mathcal{A}) P(k, z) & \text{if } (k, z) \text{ in cell } g \\ P(k, z) & \text{otherwise,} \end{cases} \quad (2.14)$$

where \mathcal{A} is a fixed small amplitude change that I take to be 5%. The two sided derivative is:

$$C_{\ell,g} = \frac{C_\ell[P_g(k, z, \mathcal{A})] - C_\ell[P_g(k, z, -\mathcal{A})]}{2\mathcal{A}}, \quad (2.15)$$

where $, g$ denotes the derivative with respect to amplitude of cell g . Substituting into (2.13) gives the Fisher matrix, F , for the matter power spectrum grid cells. Then the sensitivity to power spectrum cell g is given by the inverse error, $\sigma^{-1}(\mathcal{A}_g)$, which is:

$$\sigma^{-1}(\mathcal{A}_g) = \frac{1}{\sqrt{(F^{-1})_{gg}}}. \quad (2.16)$$

I compute the error on Ω_m , τ , Ω_b , H_0 , σ_8 and w_0 in a similar fashion². I do not compute the constraints on w_a because I have found that this can be sensitive to exactly how the derivative is defined.

2.7 Results

Using the formalism presented in the previous section I determine the sensitivity of different analyses to regions of the matter power spectrum and compare the

² τ is the optical depth to reionization. It has a very small impact on weak lensing studies but is included for completeness.

constraints on the dark energy equation of state parameter w_0 . I consider:

- the standard cosmic shear C_ℓ approach with a large constant $\ell_{\max} = 3000$;
- a BNT re-weighed C_ℓ analysis with no ℓ -cuts;
- k -cut cosmic shear for target k_{cuts} of the form $k_{\text{cut}} = A_{\text{cut}}$ (redshift independent) and $k_{\text{cut}} = A_{\text{cut}}(1+z)^2$. In the former case I use $A_{\text{cut}} \in [0.64, 1.94, 3.38]$, while in the later case I consider $A_{\text{cut}} \in [0.2, 0.6, 2]$, which roughly follows the redshift evolution of the highest k -mode in: the linear regime, a k -value in the quasi-linear regime and a k -value in the fully nonlinear regime.

In Figure 2.2 I show the inverse error on the amplitude of power spectrum cells for these six different analyses. Cosmic shear is most sensitive to dark blue regions. Using the standard C_ℓ analysis (top left) $\sim 50\%$ of the signal comes from hard to model scales, above $k_{\text{cut}} = 1 \text{ h Mpc}^{-1}$. The top right panel shows the case where I have applied *BNT* re-weighting with no angular scale cuts. This has no effect on the sensitivity compared to the standard case, as expected. Finally in the last two rows I plot the sensitivity of k -cut cosmic shear for different target k -cuts. Sensitivity to regions above the target cut is dramatically reduced to essentially zero sensitivity, for all the cuts considered. This is true even when photometric redshift errors defined in equation (1.62) are included in the C_ℓ computation but not in the calculation of the BNT weight, as is the case in this analysis.

In Figure 2.3 I summarise the reduction in sensitivity to small scales plotting the fraction of the matter power spectrum information that comes from scales above the cut. I define this as the sum of the inverse errors (see equation (2.16)) on the cells above the cut relative to sum over all cells. Never more than 5% of the information comes from scales above the target cut with k -cut cosmic shear for all the cuts I have considered. This contrasts with the standard C_ℓ approach where in some cases $> 50\%$ of the structure information comes from scales above the cut.

I plot the fraction of the power spectrum information retained using different

k -cuts, relative to the standard approach in Figure 2.4 . I define the information as the sum of the inverse errors on the power spectrum cells. When I take a very conservative choice, $k_{\text{cut}} = 1.94 h \text{ Mpc}^{-1}$, $\sim 50\%$ of the power spectrum information is lost. In fact for all cuts I have considered, $> 35\%$ of the power spectrum information is lost. However, most of the constraining power on the dark energy equation of state, w_0 is retained (see Figure 2.5), because information about this parameter comes predominantly from large scales (small- k) in the power spectrum [31] and from the background geometry [124].

In Figure 2.5 I plot the k -cut cosmic shear Fisher constraints on w_0 , relative to the standard C_ℓ approach. Scale cutting results in some loss of constraining power. However, in all but the most extreme case that I considered, this never degrades the w_0 constraint by more than a factor of 2. For example, cutting scales above $k_{\text{cut}} = 1.94 h \text{ Mpc}^{-1}$ results in a 31% increase on the size of the error. Meanwhile, with the most aggressive cut that was considered – where I remove sensitivity to all nonlinear scales taking $k_{\text{cut}} = 0.2(1+z)^2$ – the size of the error increases by a factor of 2.8.

2.8 Outlook and Future Prospects

In this chapter I have demonstrated that k -cut cosmic shear is a clean and efficient way to remove sensitivity to small scales, which are not yet modelled accurately enough for next generation surveys.

Ultimately we should strive to develop accurate models of the matter power spectrum; there is information to be had here, which is not accessible with other cosmological probes. But the reality is that, with data arriving from LSST and Euclid in less than three years, we will likely need to cut scales – and k -cut cosmic shear does this optimally. The cut-scale, k , can be gradually pushed upwards as models of the matter power spectrum improve, or, given multiple favoured models, the cut scale can be taken to correspond to the scale where the results derived

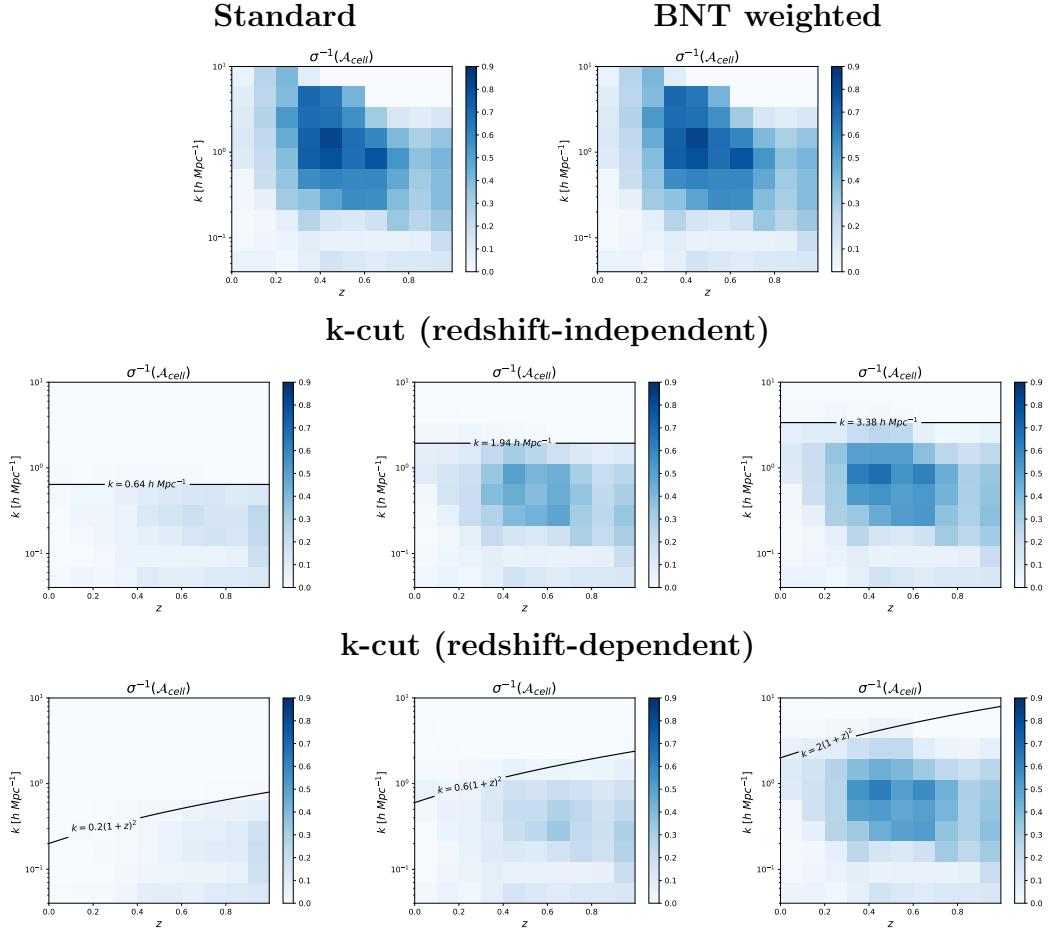


Figure 2.2: The inverse error on the measured amplitude of each power spectrum cell, $\sigma^{-1}(\mathcal{A})$, using different techniques with the Fisher matrix analysis. The technique is sensitive to regions where $\sigma^{-1}(\mathcal{A})$ is high (darker blue). **Top left:** The standard C_ℓ approach. **Top right:** BNT weighting with no ℓ -cut. This should not change the total sensitivity and there is at most a 0.02% fractional in any cell relative to the standard C_ℓ approach due to numerical imprecisions in my implementation. **Centre row:** k -cut lensing with a target k_{cut} of the form $k_{cut} = A_{cut}$. **Bottom row:** k -cut lensing with a target k_{cut} of the form $k_{cut} = A_{cut}(1+z)^2$. In the bottom two rows, k -cut cosmic shear efficiently removes sensitivity to the power spectrum above the desired k .

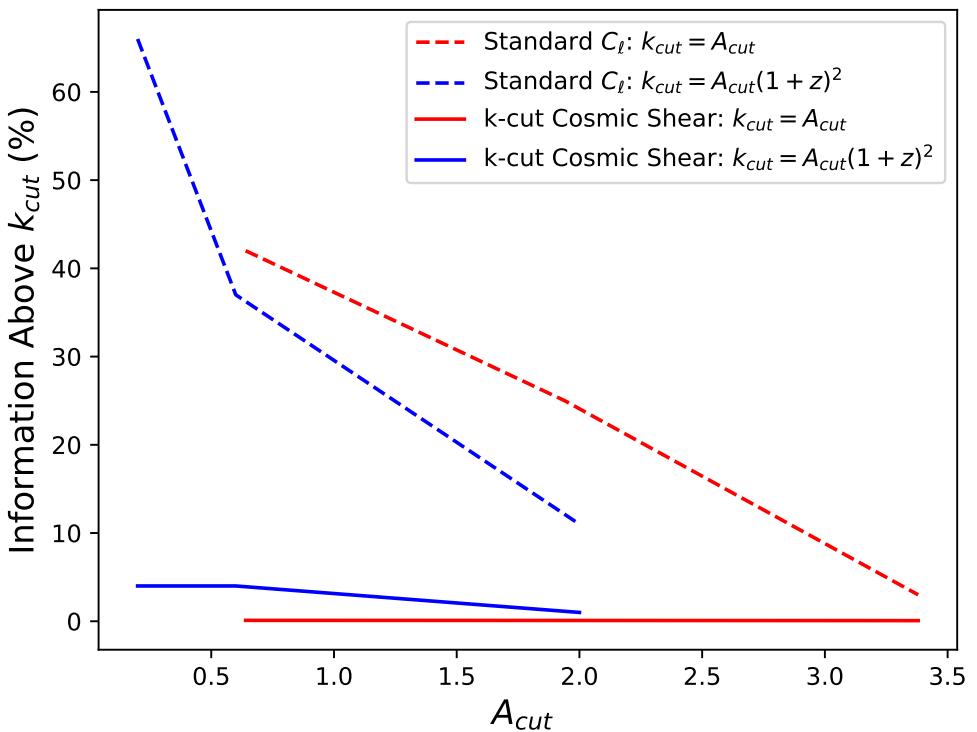


Figure 2.3: Fraction of the information (defined as the sum of the inverse errors (see equation (2.16)) on the cells above the cut relative to sum over all cells) coming from above the desired cut scale using k -cut cosmic shear and the standard C_ℓ approach. In the standard analysis a significant fraction of the signal comes from above the cuts, while k -cut cosmic shear removes nearly all sensitivity to small scales.

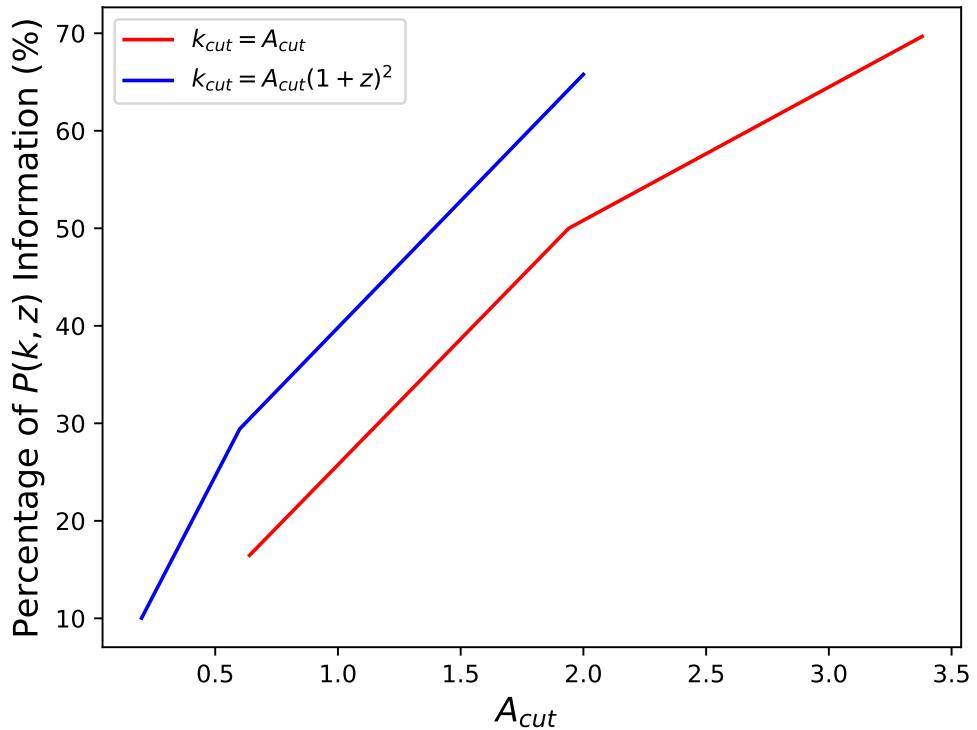


Figure 2.4: Fraction of the power spectrum information (defined as the sum of the inverse errors (see equation (2.16)) on the cells shown in Figure 2.2) captured by k -cut cosmic shear relative to the standard C_ℓ approach. A large share of the power spectrum information is lost using k_{cut} cosmic shear but by comparing with Figure 2.5, we find that most of the information about the dark energy equation of state, w_0 , is retained. For example, when we take $k_{cut} = 1.94 h \text{ Mpc}^{-1}$, the size of error on w_0 increases by $\sim 30\%$, even though $\sim 50\%$ of the power spectrum information is lost.

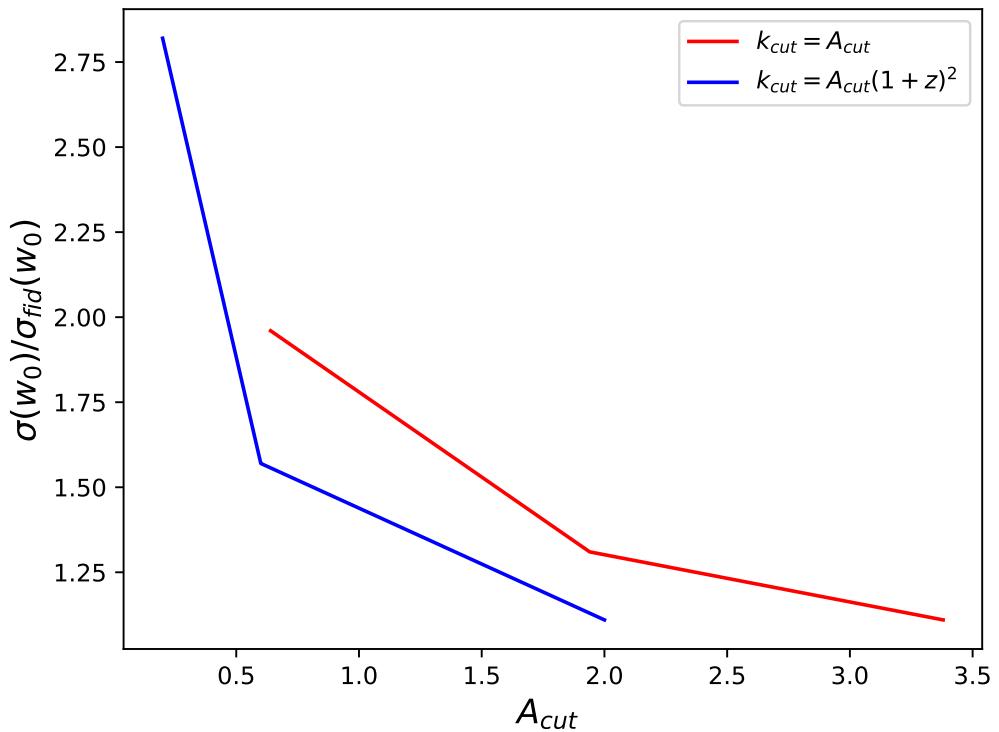


Figure 2.5: Size of 1σ marginalized k -cut constraints on the dark energy equation of state, w_0 , relative to the standard approach. Except in the most extreme case, where a very aggressive k -cut is used, most of the sensitivity to small scales can be cut without degrading the w_0 constraint by over a factor of 2. Applying the BNT transformation with no k -cut does not result in a loss of information, as expected.

from different power spectrum models start to diverge.

The next step is to apply the k -cut cosmic shear to data. This can be done by applying the BNT transformation to the C_ℓ computed from the data, and dealing with masked regions of the sky either with a forward model approach or a pseudo- C_ℓ approach [6, 24, 133]; see Chapter 4 for more details. However, there are a few additional considerations which must be addressed:

- **Intrinsic Alignments (IA):** Since the IA signal is generated with a different kernel from the gravitational shear signal, k -cut cosmic shear does not remove small scales from the IA contribution. This is not a major concern because the IA contribution is (i) already very sub-dominant to the shear signal and (ii) primarily sensitive to large scales through tidal distortions induced by massive halos. Making these statements precise is left to a future work.
- **Cosmological Dependence:** To make the BNT transform we must assume a fiducial cosmology to make the mapping from redshift, z , to comoving distance $r(z)$. This is investigated briefly in [20] and they find this should not be an issue, particularly since Baryonic Acoustic Oscillation and Type IA supernovae measurements already place tight constraints on the background expansion [4, 110].
- **Covariance Matrix:** It is necessary to generate a valid data covariance matrix for the k -cut cosmic shear C_ℓ 's to perform inference on real data. With current data this can be achieved by adapting the pipeline that I present in Chapter 4. Alternatively we could use a likelihood-free approach to perform the inference step, which is also presented in Chapter 4. This is favourable because we would not have to worry about the impact of the mask.
- **Verification:** In light of the issues above, we must verify that k -cut cosmic shear really does remove sensitivity to the power spectrum above the desired

scale. The easiest way to do this is to adapt the non-parametric technique presented in the next chapter to check that the data does not constrain the power spectrum above the desired k -cut.

Since in the short term k -cut cosmic shear is an efficient way to deal with small scales that are difficult to model and because it might be the only computationally feasible way to deal with these scales to test a large number of theories of gravity, overcoming these remaining issues should be a priority.

Chapter 3

Non-parametric Cosmology with Cosmic Shear

In this Chapter I present non-parametric cosmology with cosmic shear which I introduced in [125], closely following the discussion therein.

3.1 Motivation

The leading cosmological model, Lambda-Cold Dark Matter (Λ CDM) is phenomenological. With no widely accepted physical mechanism that explains the existence of dark matter or the accelerated expansion of the Universe, in addition to measuring the Λ CDM parameters to ever great precision, we must test – rather than assume – the Λ CDM paradigm.

I use data from the Canada-France-Hawaii Lensing Survey (CFHTLenS) to reconstruct the lensing amplitude \mathcal{A}_G , the linear intrinsic alignments amplitude, \mathcal{A}_{IA} , the comoving distance, $r(z)$ and the matter power spectrum, $P(k, z)$. I will refer to this as non-parametric cosmology with cosmic shear since this information is always measurable without ever needing to assume a cosmological model parametrised in terms of a small number of physical parameters. I will refer to the information contained in these amplitudes as the non-parametric information.

By pinpointing the exact scales where discrepancies occur, this information

will help us look for physics beyond Λ CDM and more importantly help localise any discrepancies with Λ CDM to regions of the matter power spectrum or a given redshift in the expansion history.

It is worth noting that although cosmic shear extracts both distance and structure growth information, measurements of Baryonic Acoustic Oscillations (BAO) [4] and Type Ia supernovae (SNe Ia) [110] already constrain the cosmic distance to within a few percent at low redshifts. A disagreement in the inferred expansion history between a non-parametric cosmic shear reconstruction and two other two measurements would indicate the presence of systematics effects in one or more of the experiments.

Another motivation for this study is that once the non-parametric information is extracted, it is possible to test *any* cosmological model, without having to recompute lensing observables. Although they are not too difficult to compute at the level of precision necessary for today’s data, in the future we may have to take into account higher order effects and detector systematics (see Chapter 4), greatly increasing computation time. Inferring cosmological parameters from the non-parametric information can also be used to verify the fidelity of my non-parametric reconstruction. I infer the Λ CDM parameters directly from the non-parametric information and compare the results to those found using the standard cosmic shear likelihood analysis later in the chapter.

3.2 Correlation Functions

The original CFHTLenS tomographic cosmic shear analysis [51] (hereafter H13) used correlation functions rather than the lensing spectrum. This statistic contains the same information as the lensing spectrum, but it is easier to deal with the mask, although some care must be taken [129]. Since all the relevant correlation function data products have already been computed, for convenience I choose to use this statistic for the remainder of this chapter. However a similar analysis

could equally well be performed with the lensing spectrum.

For pairs of galaxies the tangential and cross-ellipticity $\epsilon_{+,\times}$ are defined in the frame joining a pair of galaxies $\{\alpha, \beta\}$. The correlation function for angular separation, θ , is computed from the data by:

$$\hat{\xi}_{\pm}^{ij}(\theta) = \frac{\sum w_{\alpha} w_{\beta} [\epsilon_{+}^i(x_{\alpha}) \epsilon_{+}^j(x_{\beta}) \pm \epsilon_{\times}^i(x_{\alpha}) \epsilon_{\times}^j(x_{\beta})]}{\sum w_{\alpha} w_{\beta}}, \quad (3.1)$$

where w_{α} and w_{β} are weights and i and j denote the tomographic bin number, and the sums are over all galaxy pairs. In the limit of high- ℓ [76], the theoretical correlation function is related to the convergence power spectrum, $C_{\ell, \text{GG}}^{ij}$, by:

$$\xi_{\pm, \text{GG}}^{ij}(\theta) = \frac{1}{2\pi} \int d\ell \ell C_{\ell, \text{GG}}^{ij} J_{\pm}(\ell\theta), \quad (3.2)$$

where $J_{+}(\ell\theta)$ is the zeroth order Bessel function of the first kind and $J_{-}(\ell\theta)$ is the fourth order Bessel function of the first kind. This result is derived in [67] (see also [76]). Since correlation functions are formed from the lensing spectrum filtered by a Bessel function, physical scales do not separate as naturally as for the lensing spectrum and not nearly as well as the BNT weighted lensing spectrum presented in Chapter 2. Analogous expressions can be found for the ‘IG’ and ‘II’ contributions written, $\xi_{\pm, \text{GI}}^{ij}(\theta)$ and $\xi_{\pm, \text{II}}^{ij}(\theta)$, by replacing $C_{\ell, \text{GG}}^{ij}$ with $C_{\ell, \text{GI}}^{ij}$ and $C_{\ell, \text{II}}^{ij}$, in the expression above. The theoretical correlation functions are found as a sum of the contributions:

$$\xi_{\pm}^{ij}(\theta) = \xi_{\pm, \text{II}}^{ij}(\theta) + \xi_{\pm, \text{GI}}^{ij}(\theta) + \xi_{\pm, \text{GG}}^{ij}(\theta). \quad (3.3)$$

As in H13 I assume that the intrinsic alignment amplitude has no redshift dependence by setting $\eta = 0$ in equation (1.79), so that the intrinsic alignment contribution is given by a single amplitude, \mathcal{A}_{IA} . Throughout the remainder of this chapter I use **Nicea** [71] integrated into **Cosmosis** to compute the correlation functions. I use my own code to compute the lensing spectrum, defined in equa-

tion (1.71), after making the transformation $\ell = kr$. This means that $r(z)$ can be perturbed independently without changing the matter power spectrum label.

3.3 CFHTLenS Data

The Canada-France-Hawaii Lensing Survey covers 154 deg^2 in five optical bands. The observed median redshift of galaxies in the survey is $z_m = 0.70$ with an effective weighted number density of $n_{\text{eff}} = 11$ galaxies per square arcmin. Catalogs were produced by combining weak lensing data processing from **THELI** [36], shear measurements from **Lensfit** [99] and photometric redshift estimates were found using the method presented in [55].

In what follows, I use the same angular bins and the same 6 tomographic bins for galaxies in the redshift range $0.2 < z < 1.30$ as in H13. Radial bins were defined by dividing galaxies into the redshift ranges: $z_1 \in [0.2, 0.39]$, $z_2 \in [0.39, 0.58]$, $z_3 \in [0.58, 0.72]$, $z_4 \in [0.72, 0.86]$, $z_5 \in [0.86, 1.02]$, and $z_6 \in [1.02, 1.30]$. Bins are smoothed by Gaussian kernel with dispersion $\sigma_z = 0.04(1+z)$ to account for the photometric redshift uncertainty.

3.4 The Non-parametric Likelihood

To extract the cosmological information from the shear catalog, I assume a Gaussian likelihood:

$$\ln \mathcal{L}_1(p) = -\frac{1}{2} \sum_{a,b} [D_a - T_a(p)] C_{ab}^{-1} [D_b - T_b(p)], \quad (3.4)$$

where D_a is the data vector composed of the observed $\hat{\xi}_{\pm}^{ij}$ and $T_a(p)$ is formed from the theoretical prediction of ξ_{\pm}^{ij} given parameters p , and C_{ab}^{-1} is the inverse of the covariance matrix; see Chapter 4 for a discussion of the Gaussian likelihood approximation and ways to sidestep this approximation altogether.

The data and theory vectors are taken to be the correlation functions defined

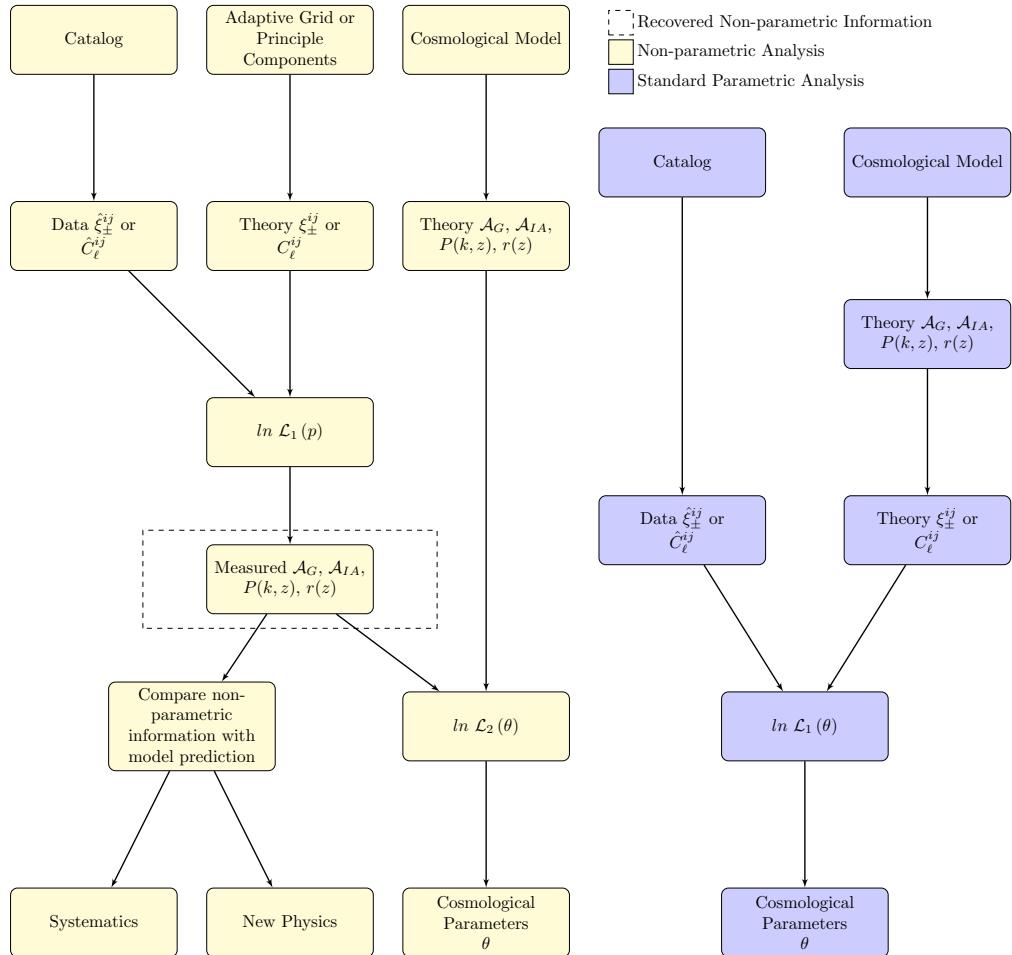


Figure 3.1: The main steps of the non-parametric reconstruction (yellow) and the standard parametric inference (purple). It is useful for the reader to refer back to this figure throughout the remainder of this chapter. I discuss these techniques in detail in Sections 3.5-3.8. The non-parametric reconstruction has a number of desirable features. 1. The expansion and structure growth history is recovered, with no need to assume a cosmological model. 2. Once the non-parametric information is recovered, any cosmological model can be tested without needing to re-compute lensing observables. 3. Comparing the non-parametric and parametric reconstructions pinpoints the precise redshifts and scales where the cosmological model fails – if any. This could help narrow the search for previously unidentified systematic effects (see the second paragraph of Section 3.13). After a thorough search – if the discrepancies are believed to be physical – this would indicate precisely how the Universe deviates from Λ CDM in a fully non-parametric way.

in equations (3.1) and (3.3). In the standard cosmic shear likelihood analysis, the parameters p are taken to be the cosmological model parameters and a set of nuisance parameters (e.g. the amplitude of the intrinsic alignments \mathcal{A}_{IA}). In the non-parametric analysis which follows, I take the parameters $\{p\}$ to be a set of amplitudes that encode information about the lensing amplitude, the intrinsic alignment amplitude, the comoving distance and the power spectrum (see Sections 3.6 and 3.7 for more details).

I use the publicly available covariance matrix from the CFHTLenS survey¹ integrated into `Cosmosis 2pt` module (see [51] for more details). This covariance matrix is generated from the N-body lensing simulations of [45]. I apply the Anderson-Hartlap correction when inverting the covariance [46, 14], since the covariance is generated from noisy realisations. I use the Markov Chain Monte Carlo (MCMC) sampler `emcee` [40] to sample the likelihood and perform the inference.

3.5 Non-parametric Information Extraction

Correlation functions (alternatively the lensing spectrum) are only sensitive to the cosmology of the Universe through the matter power spectrum, $P(k, z)$, the comoving distance, $r(z)$ and a set of lensing amplitudes. With my choice of intrinsic alignment model, the lensing amplitudes are: an overall shear amplitude, \mathcal{A}_G , and an intrinsic alignment amplitude, \mathcal{A}_{IA} . From equations (1.65) and (1.79) we see $\mathcal{A}_G \propto \Omega_m H_0^2$ and $\mathcal{A}_{IA} \propto A_I \Omega_m$.

The primary aim of this chapter is to divide the power spectrum and comoving distance into cells/bins and simultaneously measure the amplitudes of these cells and the lensing amplitudes. In detail, I generate a fiducial² power spectrum, comoving distance and lensing amplitudes using the CFHTLenS best fit cosmology from H13. I divide the power spectrum into a set of cells in the (k, z) plane

¹This covariance matrix is available for download from <http://www.cfhtlens.org/astronomers/cosmological-data-products>

²I have found that all the results presented in this chapter are insensitive to changing the fiducial cosmology parameters by up to 15%, one at a time.

$\{P_i\}$ and the comoving distance into a set of cells $\{r_i\}$. Perturbing the lensing amplitudes, the power spectrum cells and comoving distance cells, I define a vector of amplitudes $p = (\mathcal{A}(\mathcal{A}_G), \mathcal{A}(\mathcal{A}_{IA}), \{\mathcal{A}(P_i)\}, \{\mathcal{A}(r_j)\})$ where the amplitudes are defined relative to the fiducial cosmology, so that:

$$\begin{aligned}\mathcal{A}(\mathcal{A}_G) &= \mathcal{A}_G / \mathcal{A}_G^{\text{fid}} \\ \mathcal{A}(\mathcal{A}_{IA}) &= \mathcal{A}_{IA} / \mathcal{A}_{IA}^{\text{fid}} \\ \mathcal{A}(P_i) &= P_i / P_i^{\text{fid}} \\ \mathcal{A}(r_j) &= r_j / r_j^{\text{fid}},\end{aligned}\tag{3.5}$$

where P_i is the power spectrum in cell i , P_i^{fid} is the fiducial power spectrum inside cell i , r_j is the comoving distance in cell j , and r_j^{fid} is the fiducial comoving inside cell j . New GG, GI and II correlation functions are defined as functions of the amplitude vector, p , and written as $\xi_{\text{GG}}^{ij}(\theta, p)$, $\xi_{\text{GI}}^{ij}(\theta, p)$ and $\xi_{\text{II}}^{ij}(\theta, p)$. Substituting into equation (3.3) forms the theory vector in the Gaussian likelihood defined in equation (3.4) and I infer the posterior distribution on the amplitude vector, p .

I start by dividing the power spectrum into 100 cells on a 10×10 grid in k and z . The comoving distance is divided into 10 cells in z . With the lensing and intrinsic alignment amplitude, this leaves 112 amplitudes to measure. To perform the inference we first compress the amplitude vector, p , down to a more manageable size using two different data compression regimes, which are discussed in the next two sections.

3.6 Adaptive Grid Compression

To perform the compression I employ the Fisher matrix formalism (see Section 2.6) to first assess cosmic shear's sensitivity to each amplitude. Since the covariance matrix is already computed, it follows from the definition of the Fisher matrix

(2.10) and the Gaussian likelihood (3.4) that the Fisher matrix is given by:

$$F_{ij} = \sum_{a,b} \frac{\partial T_a}{\partial p_i} C_{ab}^{-1} \frac{\partial T_b}{\partial p_j}, \quad (3.6)$$

where T_a is the theory vector formed from the correlation functions and C_{ab}^{-1} is the covariance, while p_i and p_j are amplitudes in the vector p . Here I have assumed that the covariance matrix is independent of the cosmology, which is a good approximation in future surveys provided the likelihood is Gaussian [78]. It is now convenient to define the information content contained in a set of amplitudes $\{p_i\}$ as:

$$I = \sum_{i \in \{p_i\}} 1/F_{ii}. \quad (3.7)$$

In my adaptive grid regime, I combine adjacent power spectrum and comoving distance cells so that we are left with a much smaller set of cells that each contain roughly the same amount of information, I .

Specifically, I divide the power spectrum four cells, which each contain roughly 25% of the remaining information. The cell boundaries are plotted in Figure 3.2. Meanwhile, I divide the comoving distance into three cells in z , shown in Figure 3.3. Each contain roughly a third of the comoving distance information. Altogether the compressed amplitude vector, p , is formed of two lensing amplitudes, four power spectrum cells and three comoving cells:

$$p = (\mathcal{A}(\mathcal{A}_G), \mathcal{A}(\mathcal{A}_{IA}), \{\mathcal{A}(P_i)\}, \{\mathcal{A}(r_j)\}), \quad (3.8)$$

where $i \in [1, 4]$ and $j \in [1, 3]$.

Since this is just a proof of concept at this stage, this compression is in no way optimal and several arbitrary choices (e.g. the number of power spectrum cells) have been made. I leave optimising this procedure to a future work with new data.

3.7 Principle Component Compression

As an alternative to the adaptive grid compression, I also employ the principle component analysis (PCA) compression technique. Given a high-dimensional data set, PCA compression works by finding an orthogonal set of vectors that contain the majority of the variance in the data and only saves this information. In this case, I want to solve the opposite problem and minimise the variance by finding the linear combinations of amplitudes $p_i \in p$ that are the most tightly constrained by the shear data.

The predicted variance is encoded in the Fisher matrix, F . Specifically the covariance between amplitudes is $C_p = F^{-1}$. Rotating into the eigenbasis yields:

$$C_p = P^T D P, \quad (3.9)$$

where D is a diagonal matrix of eigenvalues, $\{\lambda_i\}$, and P is a rotation matrix with columns formed from the associated eigenvectors.

Arranging the eigenvalues in ascending order, the corresponding eigenvectors are called the principle components (PCs). I take the first 10 PCs to form the compressed amplitude vector:

$$p = (\{PC_i\}), \quad (3.10)$$

where $i \in [1, 10]$.

If the lensing likelihood was exactly Gaussian, so that the Fisher matrix formalism applies exactly, then these components would contain 77% of the total inverse variance. I stress again that this is an arbitrary choice which will need fine-tuning in the future.

3.8 Cosmological Parameter Inference from Non-parametric Information

Normally cosmological parameters, p , are found directly from the shear data by sampling from $\ln \mathcal{L}_1(p)$, defined in equation (3.4). I now discuss how to extract the cosmological parameters directly from the measured non-parametric amplitudes instead. I use this to validate the non-parametric reconstruction, but in the future this technique could be used to test a large number of cosmological models quickly and consistently *without re-computing lensing observables*.

Using the MCMC chains from the non-parametric reconstruction as data (after removing the burn-in phase), I form the Gaussian likelihood:

$$\ln \mathcal{L}_2(\theta) = -\frac{1}{2} \sum_{a,b} [\hat{p}_a - T_a(\theta)] \hat{C}_{ab}^{-1} [\hat{p}_b - T_b(\theta)], \quad (3.11)$$

where \hat{p} is the mean of amplitude vector over all samples in the chain, and the covariance, \hat{C} , between amplitudes is given by:

$$\hat{C}_{ab} = \langle (p_a - \hat{p}_a) (p_b - \hat{p}_b) \rangle, \quad (3.12)$$

and the average is taken over all samples in the chain. It is useful for the reader to refer to the left hand side of Figure 3.1.

The theory vector, $T_a(\theta)$, depends on which compression regime was used. In the adaptive grid case the theoretical lensing amplitudes scalars are:

$$\begin{aligned} \mathcal{A}^{\text{Th}}(\mathcal{A}_G) &= \frac{\Omega_m H_0^2 \sigma_8}{\Omega_m^{\text{fid}} H_0^{\text{fid}} {}^2 \sigma_8^{\text{fid}}} \\ \mathcal{A}^{\text{Th}}(\mathcal{A}_{\text{IA}}) &= \frac{\Omega_m A_I \sigma_8}{\Omega_m^{\text{fid}} A_I^{\text{fid}} \sigma_8^{\text{fid}}}. \end{aligned} \quad (3.13)$$

I have pulled out an overall scaling amplitude of the power spectrum, σ_8 , so the theoretical power spectrum amplitude inside cell i must be appropriately rescaled.³

³I have found that if I do not do pull out an overall scaling factor, I do not accurately recover the tails of the posterior on p , since a Gaussian does not completely describe the distribution of

It is:

$$\mathcal{A}^{\text{Th}}(P_i) = \left(\frac{\sigma_8^{\text{fid}}}{\sigma_8} \right)^2 \langle P_i(\theta) / P_i^{\text{fid}} \rangle, \quad (3.14)$$

where I take the average over all the points sampled in the cell. Meanwhile, the theoretical comoving distance amplitude inside cell j is given by:

$$\mathcal{A}^{\text{Th}}(r_j) = \langle r_i(\theta) / r_j^{\text{fid}} \rangle, \quad (3.15)$$

where the average is again over all points in the cell. In summary, the theory and data vectors for the adaptive grid technique are given by:

$$\begin{aligned} T &= (\mathcal{A}^{\text{Th}}(\mathcal{A}_G), \mathcal{A}^{\text{Th}}(\mathcal{A}_{\text{IA}}), \{\mathcal{A}^{\text{Th}}(P_i)\}, \{\mathcal{A}^{\text{Th}}(r_j)\}) \\ \hat{p} &= (\langle \mathcal{A}(\mathcal{A}_G) \rangle, \langle \mathcal{A}(\mathcal{A}_{\text{IA}}) \rangle, \{\langle \mathcal{A}(P_i) \rangle\}, \{\langle \mathcal{A}(r_j) \rangle\}), \end{aligned} \quad (3.16)$$

where $i \in [1, 4]$ and $j \in [1, 3]$ and the averages are taken over all samples in the reconstructed amplitude chain, which is found from sampling from $\ln \mathcal{L}_1(p)$.

In the PCA compression case, I take the theoretical amplitudes defined in equations (3.13) - (3.15), and rotate these into PCA space using the rotation matrix P , defined in equation (3.9). That is, I compute

$$\mathcal{A}^{\text{PC,Th}} = I + R(\mathcal{A}^{\text{Th}} - I). \quad (3.17)$$

The first 10 rows of the rotation matrix, R , are the same as P while the remaining rows are set to zero since I assume no contribution from the remaining PCs. The theoretical amplitude vector, \mathcal{A}^{Th} , appearing in equation (3.17), is defined in terms of the amplitudes written in equations (3.13)-(3.15) and is given by:

$$\mathcal{A}^{\text{Th}} = (\mathcal{A}^{\text{Th}}(\mathcal{A}_G), \mathcal{A}^{\text{Th}}(\mathcal{A}_{\text{IA}}), \{\mathcal{A}^{\text{Th}}(P_i)\}, \{\mathcal{A}^{\text{Th}}(r_j)\}) \quad (3.18)$$

where $i \in [1, 100]$ and $j \in [1, 10]$ run over the original cells. Finally, I is a dimension 112 vector with all entries equal to unity. This is subtracted before

the chains.

rotation in equation (3.17) because the PCA amplitudes are defined relative to unity. To summarise, the theory and data vectors for the PCA technique are:

$$T = \left(\{ \mathcal{A}_i^{\text{PC,Th}} \} \right) \quad (3.19)$$

$$\hat{p} = (\{ \langle \mathcal{A}_i \rangle \}) ,$$

for $i \in [1, 10]$ and the average is defined over all samples in the chain, just like in the adaptive grid case.

I can now sample from likelihood defined in equation (3.11), to compute the posterior distribution on the cosmological parameters, using both compression techniques. This is schematically represented in Figure 3.1.

3.9 Adaptive Grid Reconstruction Results

I sample from the likelihood $\ln \mathcal{L}_1(p)$, to measure 2 lensing amplitudes, 4 power spectrum amplitudes and 3 comoving distance amplitudes, taking a broad flat prior for each. The recovered posterior distribution of the amplitudes is plotted in Figure 3.4. Only the lensing amplitude and the comoving distance amplitudes are tightly constrained, while the amplitude of individual matter power spectrum cells are hardly constrained at all.

In Figure 3.2 I show the best fit non-parametric power spectrum. For comparison, a parametric power spectrum generated using the best fit LCDM cosmological parameters in H13 is shown. The two are in good agreement, particularly since the size of the error bars on the non-parametric reconstruction are so large (see Figure 3.4).

In Figure 3.3 I plot the non-parametric reconstruction of the comoving distance $r(z)$. The LCDM prediction generated with the best fit parameters from both H13 and the Planck 2018 combined analysis (including BAO) [2] are shown for comparison. It might appear that there is an internal inconsistency between the non-parametric reconstruction and the parametric analysis of H13, but this is just

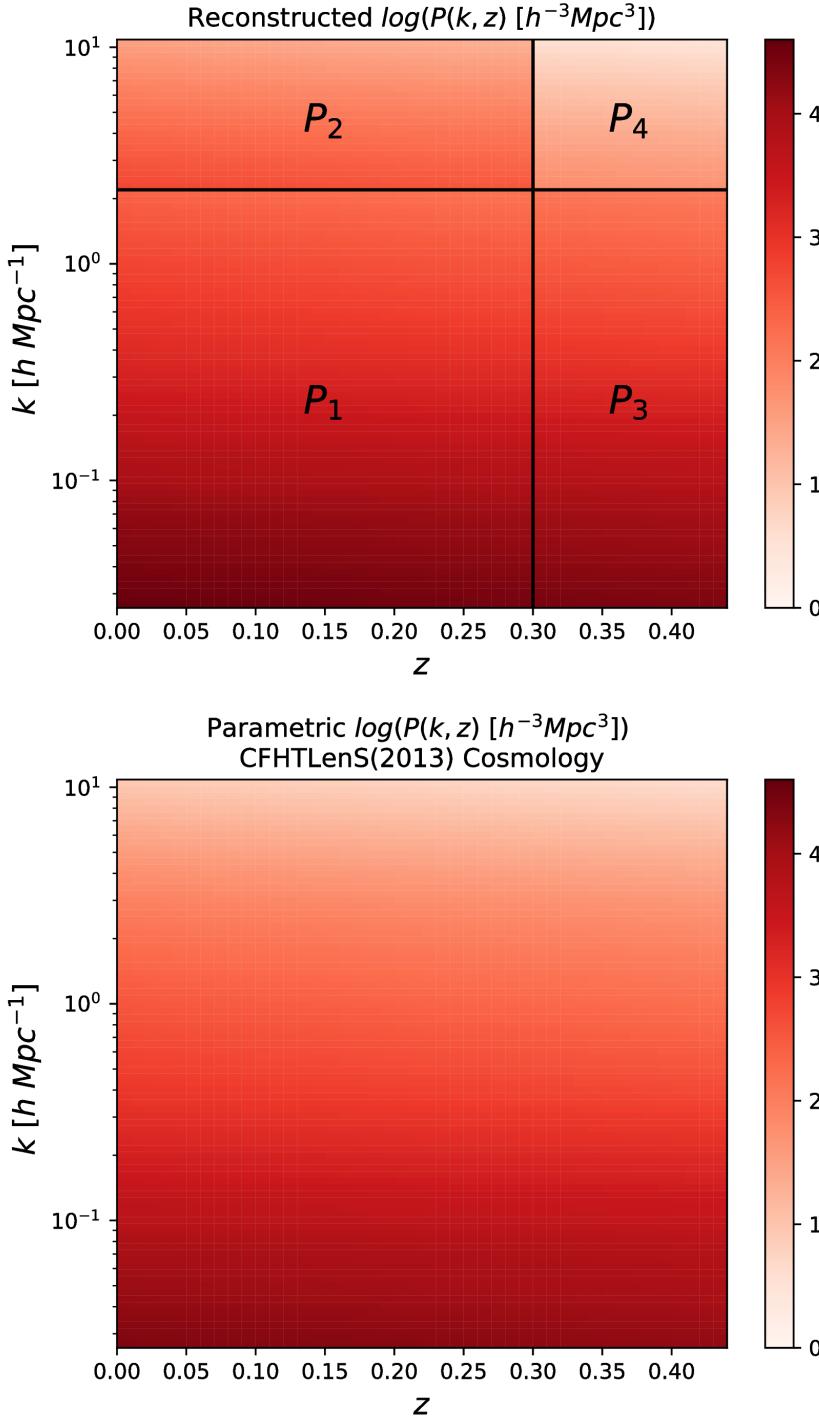


Figure 3.2: **Top:** Best fit non-parametric reconstructed power spectrum. This is the first non-parametric reconstruction of the time evolving matter power spectrum from shear data. The errors on this reconstruction are very large (see Figure 3.4), but these will shrink by a factor of $\sim 20 - 25$ with a Stage IV experiment. **Bottom:** Power spectrum generated by `CAMB` [84] and `HALOFIT` [120] using H13 best fit LCDM parameters. The non-parametric and parametric LCDM power spectrum are in good agreement.

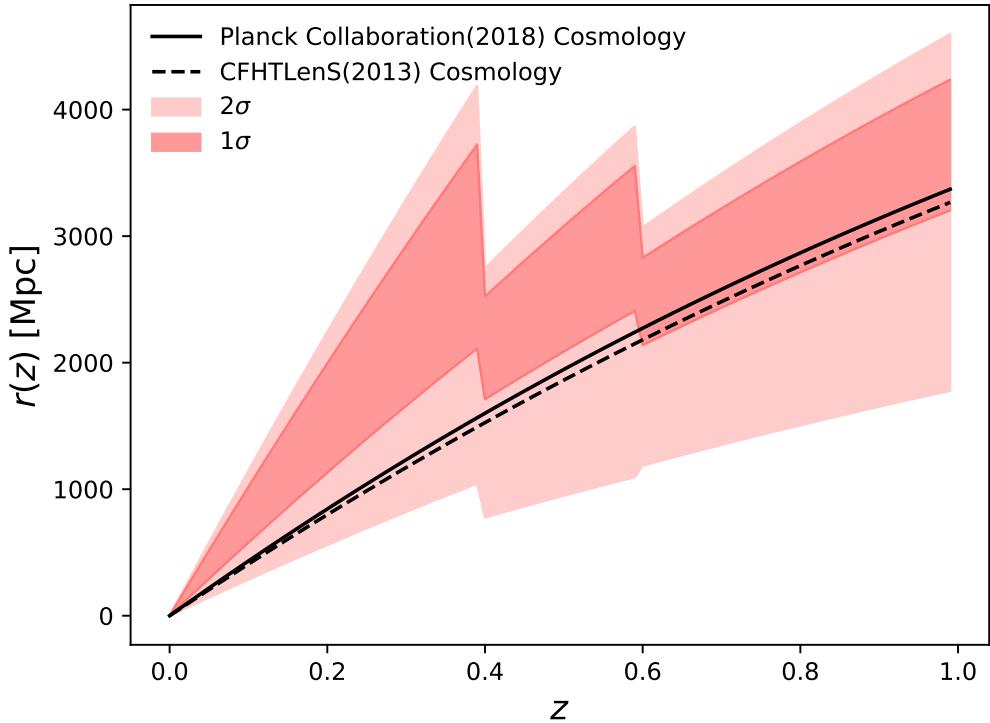


Figure 3.3: The 1σ and 2σ constraints on the reconstructed non-parametric $r(z)$. The jumps in the constraints are due to binning. This function is fairly well constrained with CFHTLenS data, unlike the power spectrum. I also plot the parametric $r(z)$ using a Λ CDM cosmology with the best fit parameters from H13 and the Planck 2018 combined analysis (including BAO) [2]. In H13 the Hubble parameter, h_0 , is given a tight prior (see Section 3.11). Repeating the H13 analysis with a flat prior on h_0 , I find it is only constrained in the range $(0.4, 1.2)$ at the 1σ level. Since $r(z)$ is proportional to h_0^{-1} , there is no internal tensions between the non-parametric distance measurement and the parametric analysis. However, below $z = 0.4$ the non-parametric reconstruction is in mild $\sim 1.5\sigma$ tension with the Planck combined cosmology $r(z)$.

due to the choice of prior on h_0 in the analysis presented in H13 (see the discussion in the caption of Figure 3.3).

For $z < 0.4$, the non-parametric reconstruction of $r(z)$ is in $\sim 1.5\sigma$ tension with the Planck LCDM predictions. In the range $0.4 < z < 0.6$, this drops to a $\sim 1\sigma$ tension, while for $z > 0.6$ there is no tension at all.

The discrepancy between the non-parametric $r(z)$ and the Planck LCDM reconstruction is unlikely to be caused by poor photometric redshift error estimation because I would expect these to get worse at higher redshifts, not lower redshifts where the tension occurs. It is also pointed out in [29], that positive values for the intrinsic alignment parameter, \mathcal{A}_{IA} , are favoured by CFHTLenS. This is the opposite sign to what is expected by theory and could point to lingering systematic effects in the shear catalog. Whatever the cause of the comoving distance tension, I intend to investigate this further with data from other surveys.

3.10 PCA Reconstruction Results

Sampling from the likelihood $\ln \mathcal{L}_1(p)$, I also infer the first 10 PC amplitudes. I show the posterior distribution on the amplitudes in Figure 3.5.

Possibly because the lensing likelihood is not exactly Gaussian [114], the fact that there is numerical noise when inverting the Fisher matrix, F , and the fiducial cosmology may not exactly correspond to the maximum likelihood point in the data, the posteriors do not agree with the Fisher expectation. Degeneracies between a few PC amplitudes (e.g between \mathcal{A}_1 and \mathcal{A}_2) are present even though degeneracies between principle components should be absent, by construction.

Although it is possible to rotate the PC amplitudes back into the (k, z) plane to reproduce Figures 3.2-3.3, I do not advocate this approach. The PCs are not a spanning set, and they will not capture the variance in the unmeasured components, leading us to underestimate the size of the error bars. To summarise: the adaptive grid method is a complete set for $r(z)$ and $P(k, z)$ but it is potentially

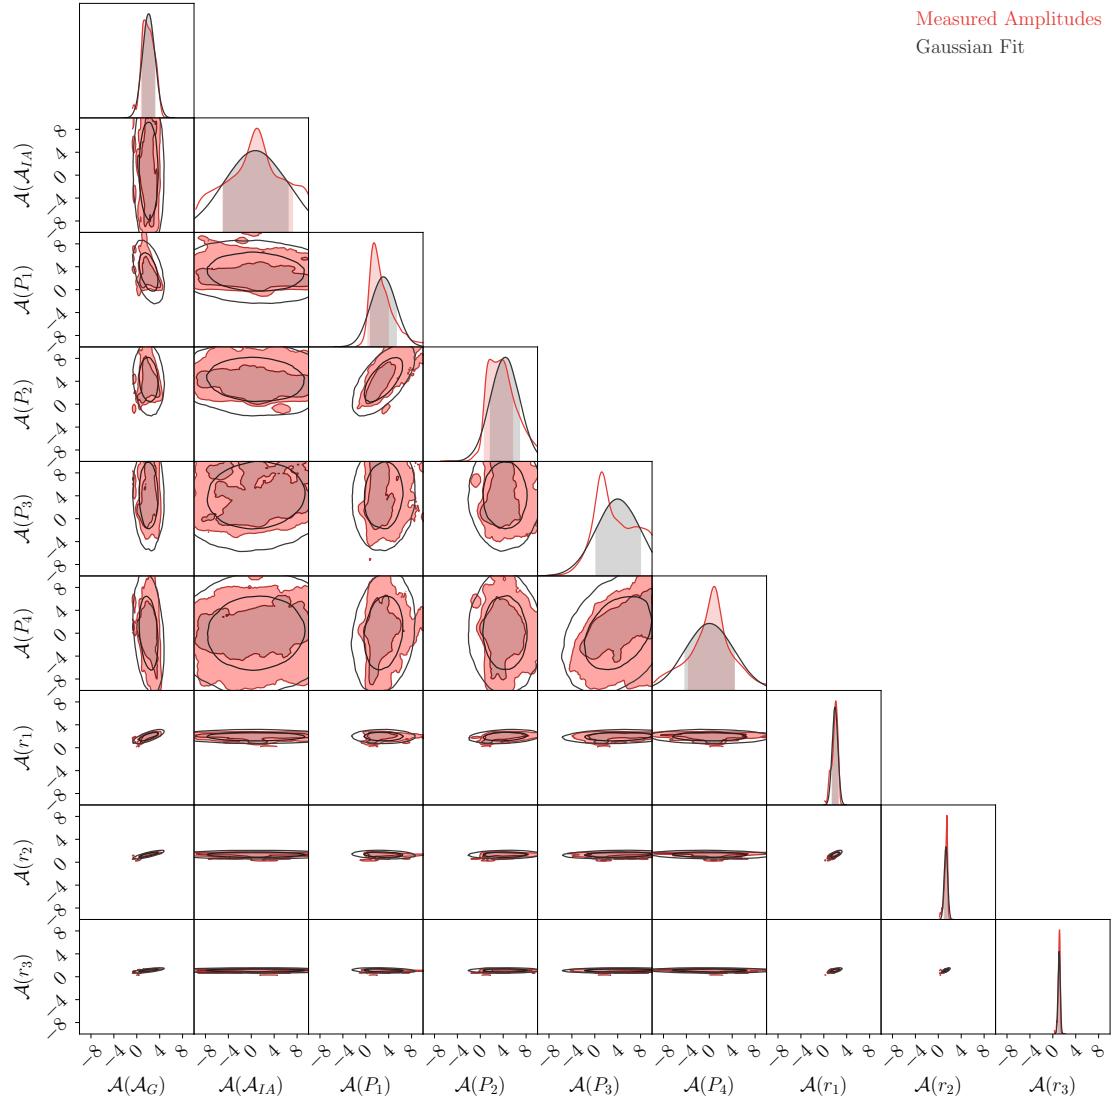


Figure 3.4: The reconstructed amplitude posteriors using the adaptive grid technique. I have inferred the amplitude of the lensing signal \mathcal{A}_G , the intrinsic alignment amplitude \mathcal{A}_{IA} , the amplitude of four power spectrum bins in k and z , and the amplitude of three comoving distances bins in z . The bin boundaries for the power spectrum and comoving distance are illustrated in Figures 3.2 and 3.3. I intentionally plot all amplitudes on the same axes to show how tightly the lensing amplitude and comoving distances are constrained relative to the power spectrum cells. Only the amplitude of the lensing signal and the comoving distance are well constrained. There is a degeneracy between the lensing amplitude and the comoving distance amplitudes because both are strongly dependent on Ω_m and h_0 . I plot the Gaussian distribution that I have fit to the chains to form the likelihood $\ln \mathcal{L}_1(p)$, in grey.

sub-optimal (leading to large error bars). The PCA approach is not a complete set (leading to potential biases), but is perhaps more optimal (smaller error bars). I recommend a conservative approach favouring the adaptive grid formulation, since it is, in general, better to be unbiased but have larger error bars than the other way round.

3.11 Cosmological Inference from Adaptive Grid Reconstruction

I sample from the likelihood $\ln \mathcal{L}_2(\theta)$ to place constraints on the LCDM parameters, θ , directly from the non-parametric information found using the adaptive grid compression. This is compared to the results of the usual cosmic shear likelihood analysis found by sampling from $\ln \mathcal{L}_1(\theta)$. In both cases, following the analysis of H13, I place a Gaussian prior on the Hubble constant: $h_0 = 0.73 \pm 0.024$. The resulting parameter constraints are shown in Figure 3.6.

Parameter constraints from the standard likelihood analysis are shown in blue, while constraints from the non-parametric information are in red. The non-parametric posteriors are wider as is expected, since information is lost in the adaptive grid compression step. Otherwise the two techniques are in good agreement. Optimising the compression step is left for a future work.

3.12 Cosmological Inference from PCA Reconstruction

Sampling from $\ln \mathcal{L}_1(\theta)$, I place constraints on the LCDM parameters directly from the non-parametric information derived using PCA compression. I show the resulting constraints in Figure 3.7, overlaying the posteriors found using the standard technique for comparison.

The constraints from the PCA compression are much tighter than those found

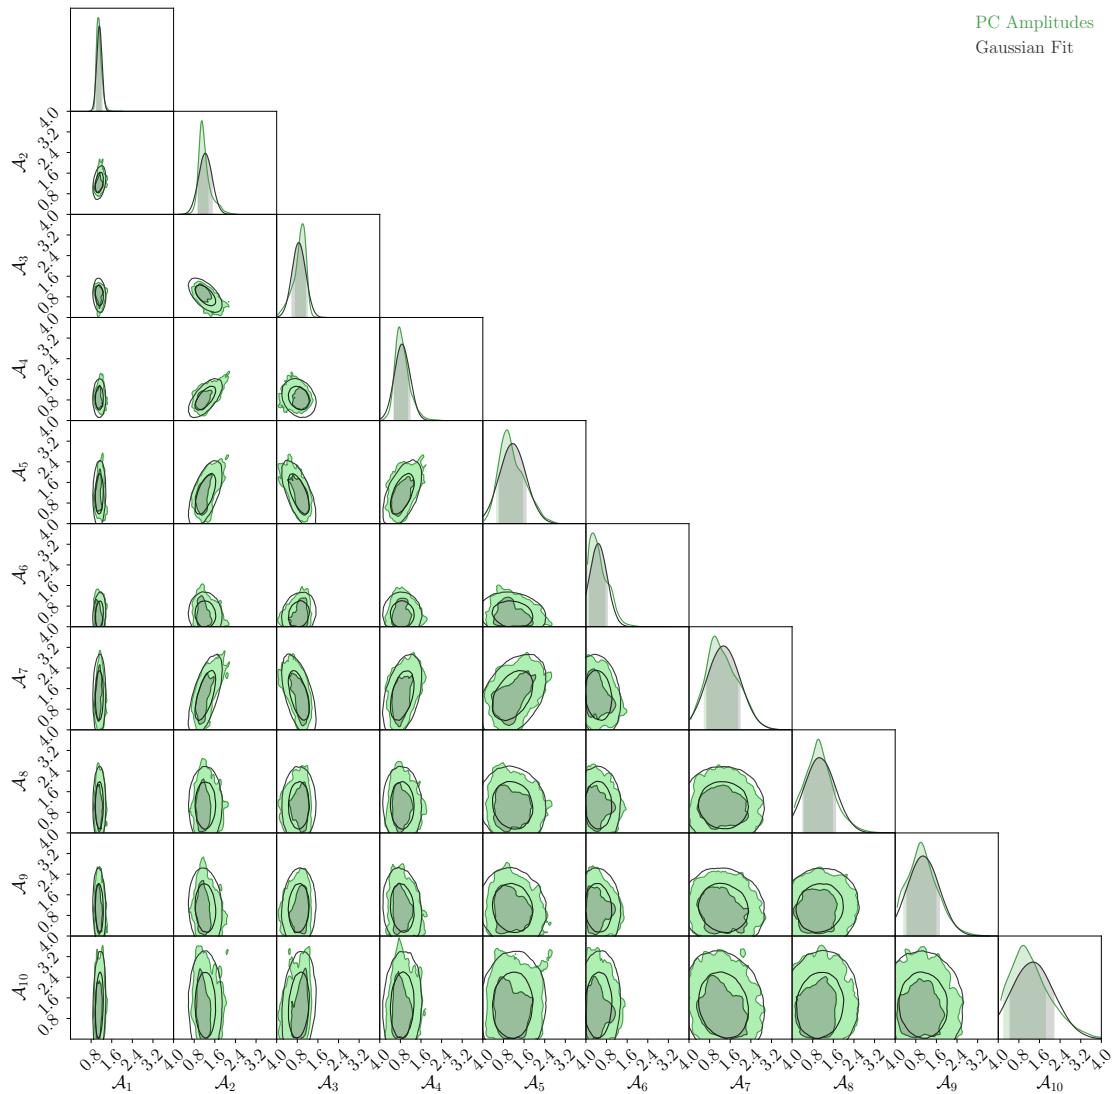


Figure 3.5: The posterior distribution on the measured PC amplitudes. The constraints on the first two PCs are in good agreement with those expected from the Fisher matrix prediction, but the constraints on the other PCs are up to twice as wide as expected. I also plot the Gaussian distribution that we have fit to the chains to form the likelihood $\ln \mathcal{L}_1(p)$, in grey.

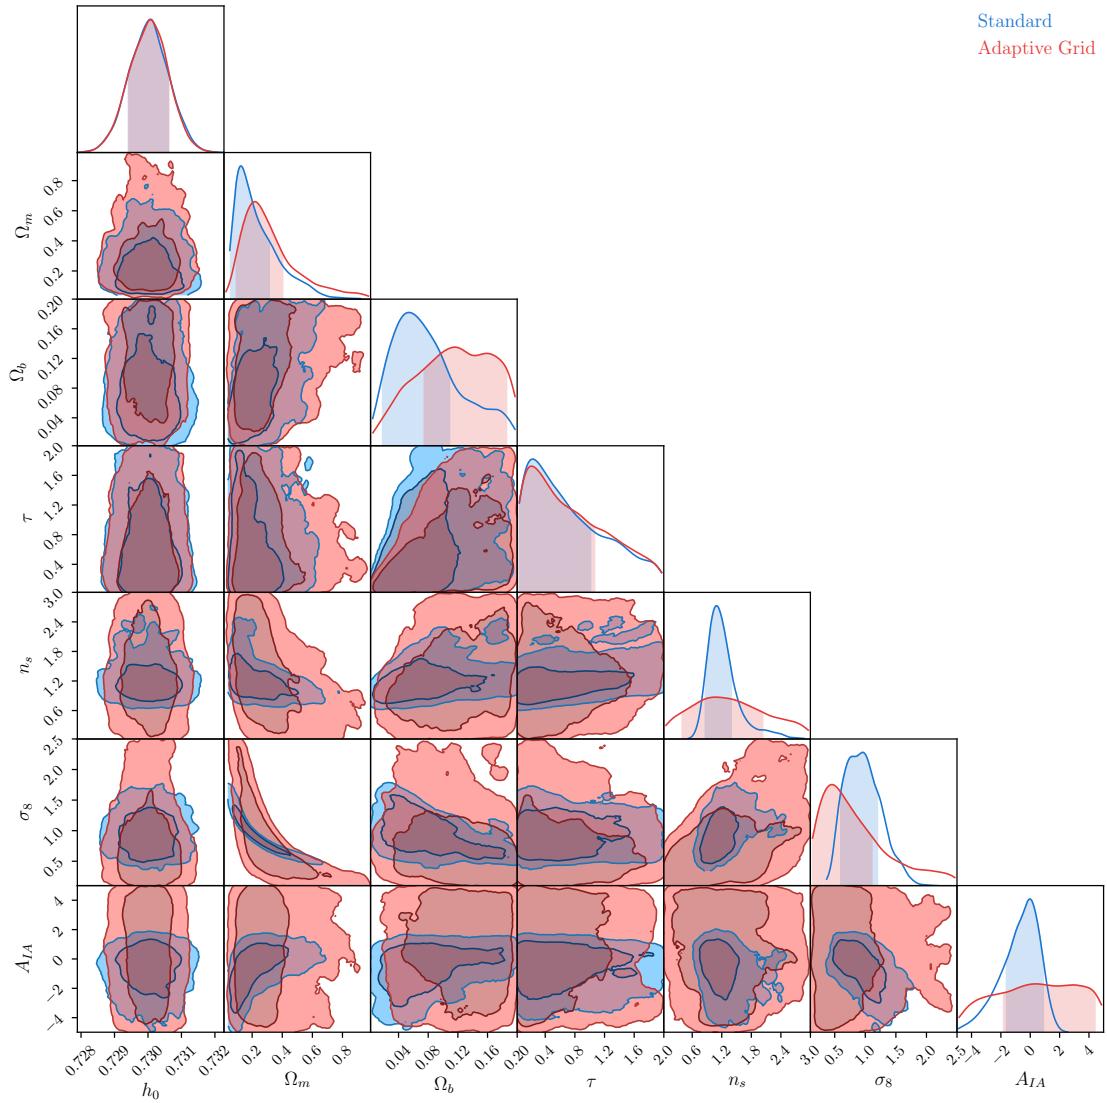


Figure 3.6: The LCDM posteriors found using the standard cosmic shear likelihood analysis (blue) and from the non-parametric information (red), using adaptive grid compression. The two techniques are in good agreement, but the posteriors in the later case are broader. This is unsurprising since information is lost in the adaptive grid compression step.

using the adaptive grid compression. They generally agree with the posteriors from the standard cosmic shear likelihood analysis, but there is $\sim 1\sigma$ tension in the $\sigma_8 - \Omega_m$ plane. Discrepancies are expected in the PCA method. As previously discussed, the PCs do not form a complete set, so we do not capture all the variance in the unmeasured PCs. This leads us to underestimate error bars, and could also cause a shift in the parameter constraints, which is why I do not advocate for the use of the PCA data compression method in this case.

3.13 Conclusion and Future Prospects

In this chapter, I have shown how to reconstruct the non-parametric information from shear data. I have applied this method to CFHTLenS shear data finding that the majority of the information is contained in a single lensing amplitude and the background expansion. To reduce the dimensionality of the reconstruction problem I have employed two different data compression regimes. Although the PCA technique is efficient, the PCs do not form a spanning set, so I conclude that the adaptive grid is preferred. Optimising the adaptive grid compression regime is left to a future work. The non-parametric method has a bright future. Since the error bars shrink roughly as the square root of both the number density and the survey area, the non-parametric constraints should shrink by a factor of ~ 2 using Dark Energy Survey Year 1 (DES Y1) data [130], and by a factor of $\sim 20 - 25$ with data from a Stage IV experiment [81].

Comparing the non-parametric and parametric reconstructions can help us identify systematic effects and search for new physics – pinpointing the precise scale and redshift where they occur. Discrepant power spectrum measurements might point to un-modelled baryonic physics, while a discrepant comoving distance measurement could point to redshift dependent systematic effects such as photometric redshift errors or colour dependent shear estimates – or even a non-constant dark energy equation of state.

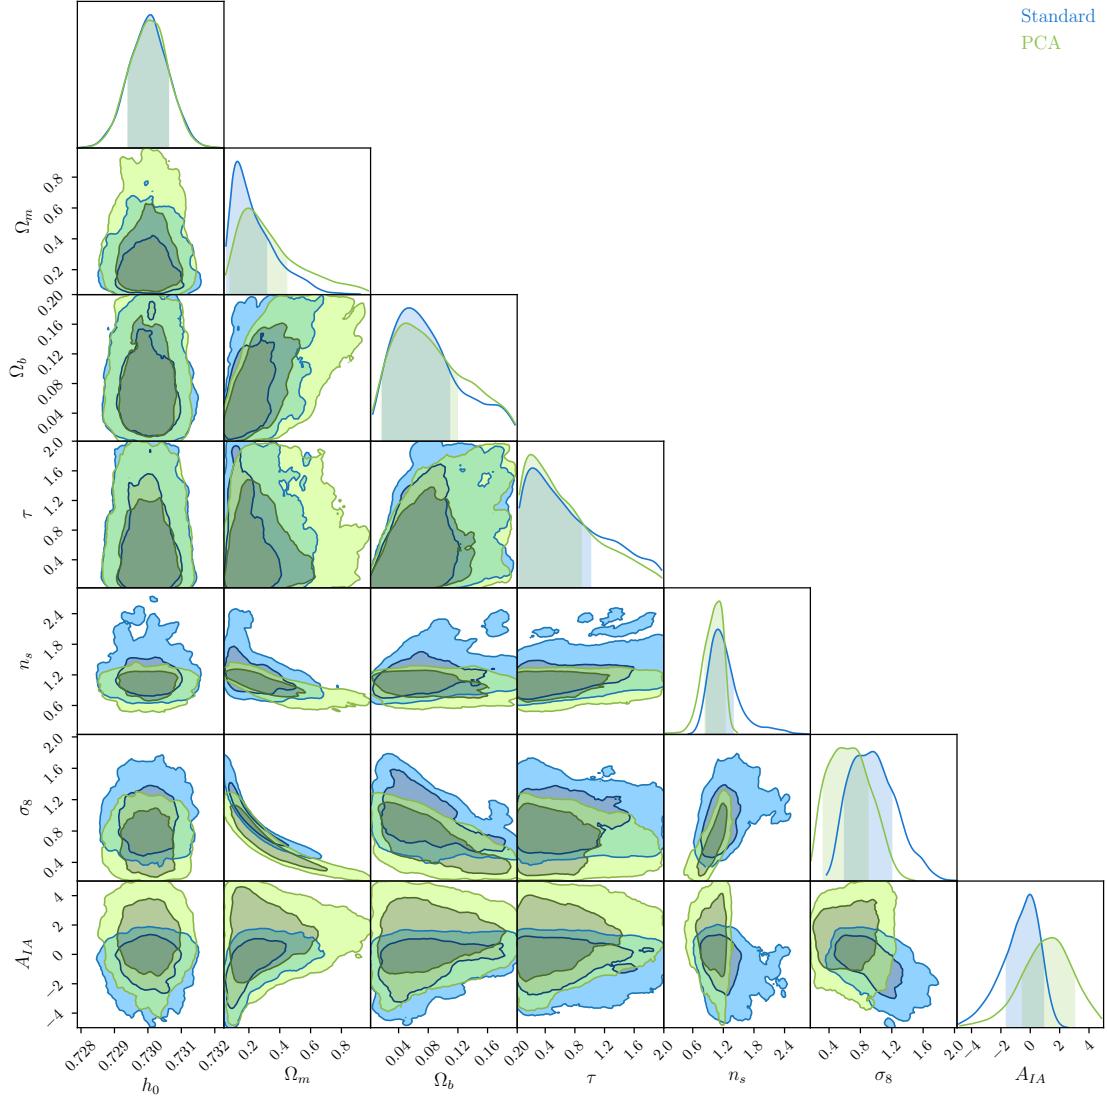


Figure 3.7: LCDM posteriors derived from the standard cosmic shear likelihood analysis (blue) and those derived from the non-parametric information (green), using PCA compression. The two techniques are in good agreement but there is $\sim 1\sigma$ tension in the $\sigma_8 - \Omega_m$ plane. Because the PCs do not form a complete set, they do not capture all the variance leading us to underestimate error bars, and could also cause a shift in the parameter constraints. I do not advocate the PCA data compression method for this reason.

In the reconstruction I have presented above, I have already identified a tension. Below $z = 0.4$, the reconstruction of the comoving distance, $r(z)$, is larger than the Planck cosmology Λ CDM prediction which could be the first sign of new physics, down to unaccounted for systematic effects or (since the tension is mild) just statistical variance. Since distance measurements from BAO and SNe Ia constrain the growth to within a few percent, unless there are large systematic effects in these other two probes, the last two explanations are more likely. In general, because BAO and SNe Ia already constrain the distance measurements so well, ensuring the cosmic shear non-parametric distance reconstruction agrees will become a useful zeroth order test to check for any systematics in a shear catalog.

The next step is to repeat the analysis presented here on the DESY1 data to see if this discrepancy persists. Folding in galaxy clustering [35] and galaxy-galaxy lensing [105] into the framework would also be informative. With a different set of systematic effects, these techniques will serve as a useful cross-check, as well as tightening constraints. Using the non-parametric method will also be a good check to ensure that the k -cut technique (see Chapter 2) works as intended with real data, by ensuring that power spectrum amplitudes above the desired cut are poorly constrained.

In the future, large experiments could repeat the non-parametric analysis and release the non-parametric reconstructions as a final data product. Then, anyone could consistently test new physical theories without having to repeat the cosmic shear analysis or model lensing observables.

The general philosophy of this chapter was to completely separate the inference of the non-parametric information – which is valid regardless of the underlying cosmological model – from the inference of the cosmological parameters. This allowed me to validate the non-parametric reconstruction by recovering the Λ CDM parameters. A fully non-parametric reconstruction also enables the inference of parameters from any cosmological model, without having to repeat the lensing analysis. However, if the sole purpose is to search for systematics or a failure of

the Λ CDM model, then it is more efficient to simultaneously infer the cosmological parameters, θ , and a set of perturbing amplitudes $\{\mathcal{A}(\theta)\}$. After marginalising out the cosmological parameters, any amplitude which is not consistent with unity is a red flag for the presence of unidentified systematics or new physics. This simultaneous approach would allow us to get away from assuming a fiducial cosmology, which may not be valid as error bars shrink in future surveys.

Chapter 4

Inference from Forward Models

In this chapter I discuss how to perform inference with full forward models of the shear data, using density-estimation likelihood-free inference (DELFI). This work was first presented in my paper [123]. It is the first realistic cosmic shear DELFI implementation which takes into account the non-Gaussianities of the shear field and intrinsic alignments.

4.1 Motivation

Except for [8, 9], all existing studies of the shear two-point statistics [51, 53, 57, 73, 130] use a Gaussian likelihood analysis to infer the cosmological parameters. This approach is far from perfect. With the improved statistical precision of next generation data, the impact of complicated ‘theoretical systematics’ (e.g. reduced shear [32]) and detector effects [95] will need to be propagated into the final parameter constraints. It is much easier to produce forward model realisations than to compute the expected impact of these effects, which is required for a likelihood analysis. Recently it was also claimed that since the true lensing likelihood is left-skewed, not Gaussian, parameter constraints from correlation functions are biased low in the $\sigma_8 - \Omega_m$ plane [113, 114]. This will be discussed in Section 4.3.

A new method called density-estimation likelihood-free inference (DELFI) [13, 22, 37, 7, 89, 102, 103] offers a way to overcome these issues. DELFI is used

to estimate the posterior distribution with summary statistics (the C_ℓ statistic, in this case) generated from full forward models of the data at different points in cosmological parameter space. This allows us to seamlessly propagate all the effects in a forward model into the final parameter constraints without needing to assume a likelihood. Because this approach uses summary statistics we avoid fitting to potentially aphysical relics in our forward models, in contrast to ‘blind’ machine learning approaches [43].

Performing inference on realisations of the data may seem computationally challenging, but using efficient data compression [12, 13], DELFI only requires $\mathcal{O}(10^3)$ simulations [7]. This is less than the number of simulations already required to produce a valid estimate of the inverse covariance matrix from noisy simulations [49] in a Stage IV likelihood analysis and DELFI is also highly parallelisable. For these reasons DELFI is more, not less, computationally tractable than a traditional likelihood analysis for upcoming surveys.

In light of these advantages, I present two cosmic shear forward model pipelines for DELFI, using the publicly available `pydelfi`¹ implementation, in this chapter. These are summarised in Figure 4.1. Pipeline I takes advantage of all the benefits of forward modelling and is intended for application to real data-sets, while Pipeline II is intended only for comparison with the standard likelihood analysis. I also consider 3 different analyses summarised in Table 4.1. DA1 is a DELFI analysis using shear Pipeline I. Meanwhile I compare the DELFI analysis, DA2, to the likelihood analysis, LA, to test the impact of the Gaussian likelihood approximation. It will be useful for the reader to refer back to Table 4.1 and Figure 4.1 throughout the remainder of the chapter.

¹<https://github.com/justinlalsing/pydelfi/commits/master>

	DA1	DA2	LA
Inference	DELFI	DELFI	Guassian likelihood
Pipeline	Pipeline I	Pipeline II	N/A
Number of galaxies	1.56×10^9	1.56×10^9	1.56×10^9
Number of tomographic bins	6	2	2
Number of ℓ -bins	15 with $\ell \in [10, 1000]$	15 with $\ell \in [10, 1000]$	15 with $\ell \in [10, 1000]$
Field type	Lognormal	Gaussian	N/A
Deconvolve Pixel Window	No	Yes	N/A
Mask	Yes	No	N/A
Subtract shot-noise	No	Yes	N/A

Table 4.1: The three analyses in this chapter. In DA1 I use DELFI to infer the cosmological parameters. Since I perform inference with forward models, there is no need to deconvolve the pixel window function, deconvolve the mask or subtract off the shot noise. This analysis is applied to mock Stage IV data in Section 4.4. By comparing the DELFI analysis, DA2, with the likelihood analysis, LA, I test the impact of the Gaussian likelihood approximation. In DA2 the modelling choices are governed by the constraint that I must match the Gaussian likelihood analysis as closely as possible. Some of the map-level choices are not applicable to the likelihood analysis, LA.

Pipeline I Pipeline II

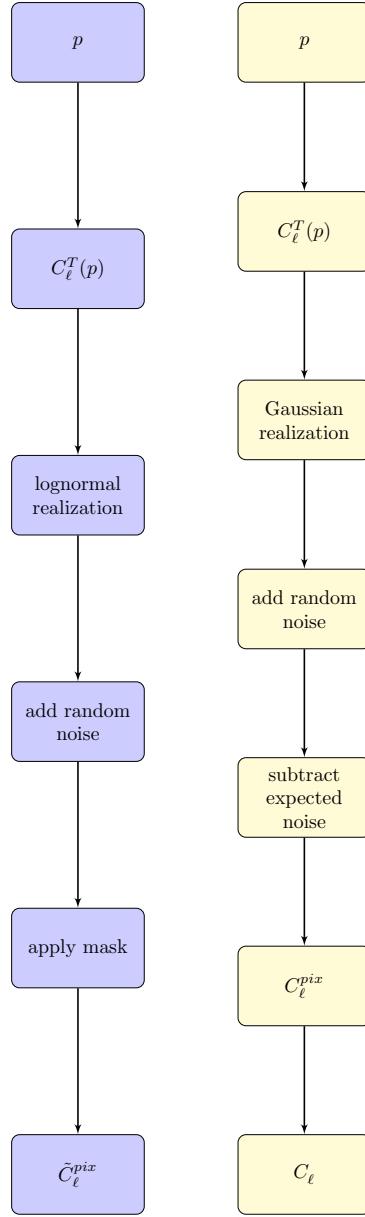


Figure 4.1: A schematic of the two forward model pipelines presented in this chapter. In Pipeline I I develop a forward model pipeline of cosmic shear data that takes full advantage of the forward model paradigm. For example, there is no need to deconvolve the mask or the pixel window function. In Pipeline II, I use a Gaussian field, do not use a mask, subtract off the shot-noise or deconvolve the pixel window function. These choices allow me to make a direct comparison between DELFI and a Gaussian likelihood analysis and test the impact of the Gaussian likelihood approximation.

4.2 Cosmic Shear Model Choices and the Lognormal Field Approximation

4.2.1 Cosmic Shear Modelling Choices

For the remainder of this chapter I take the same assumptions as the lensing spectrum given by equation (1.71), where I generate the N_{tomo} tomographic bins, $n_i(z)$, by dividing the radial distribution function:

$$n(z) \propto \frac{a_1}{c_1} \exp \left[-\frac{(z - 0.7)^2}{b_1^2} \right] + \exp \left[-\frac{(z - 1.2)^2}{d_1^2} \right], \quad (4.1)$$

with $(a_1/c_1, b_1, d_1) = (1.5/0.2, 0.32, 0.46)$ [131] into bins with an equal number of galaxies per bin. I assume a simple model for the photometric redshift error given in equation (1.62). For the intrinsic alignment contribution I use the nonlinear alignment model reviewed in Chapter 1 and to match the analysis in current Stage III experiments [66, 130], I allow the intrinsic amplitude, $A(z)$, to vary as a function of redshift so that $A(z) = [(1 + z_0)/(1 + z)]^\eta$ [91], where z_0 is the mean redshift of the survey. This is $z_0 = 0.76$ for the $n(z)$ given in equation (4.1).

4.2.2 The Lognormal Field Approximation

Lognormal convergence fields [54] are computationally inexpensive to generate and they capture the impact of nonlinear structure growth more accurately than Gaussian realisations. This is because the density contrast, δ , defined in equation (1.33) is bounded below by minus one, but has no upper bound, and this asymmetry can not be captured by a Gaussian field. In the DESY1 [130], lognormal fields were used to compute the covariance matrix from noisy realisations of the data. No differences in parameter constraints were found when the covariance was computed using lognormal fields compared with a halo model approach [79].

I use `Flask` [135] to generate consistent lognormal realisations [54] of the convergence and shear fields – correlated between redshift slices. The procedure is

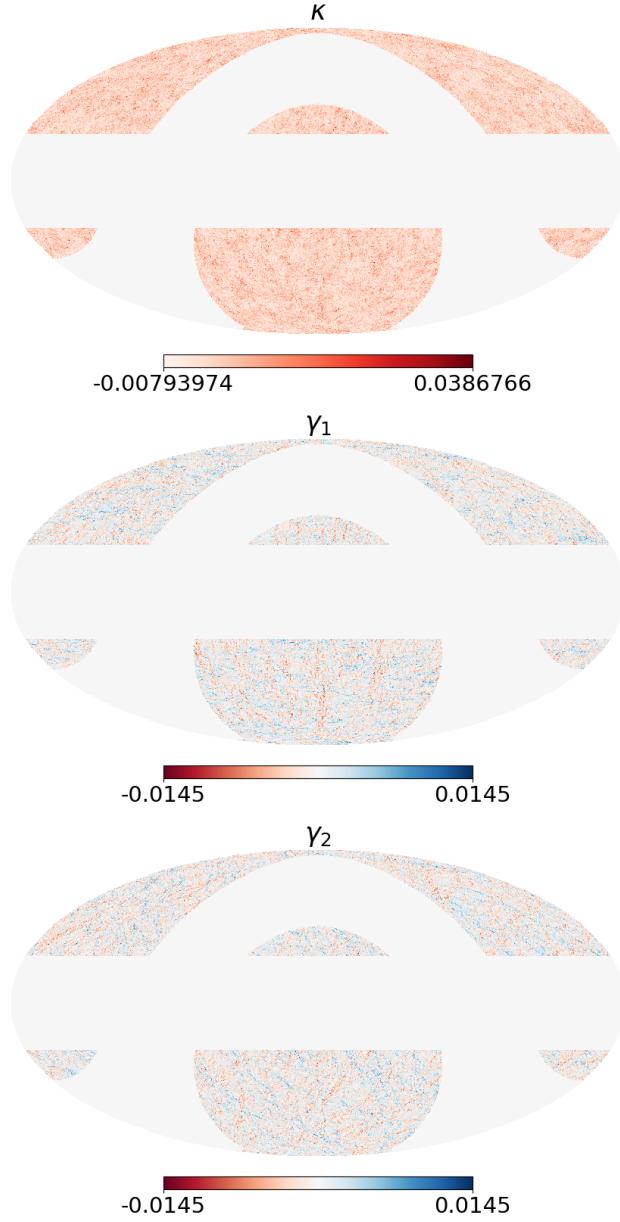


Figure 4.2: A single masked data realisation of the convergence field, κ , and the two observable shear components: γ_1 and γ_2 (including shape noise) for a typical Stage IV experiment. I show the lowest redshift bin of six, where the effect of non-Gaussianity is largest. The non-Gaussianity is clearly visible in the κ -map, where the majority of pixels are very slightly negative with a small number of pixels taking very large (positive) κ -values. Note that the colour scale in the κ -map is not symmetric about zero, to make the non-Gaussianity as clearly visible as possible. The mask cuts all pixels lying within 22.5 deg of the galactic and ecliptic planes, mimicking a Euclid-like survey.

discussed in detail in Section 5.2 of [135] (see also [92]). In particular, for each tomographic bin, i , **Flask** first generates a Gaussian realization by Cholesky decomposing (see [92] for more details) the convergence spectrum, $C_\ell^{T,ij}$, into a products of upper and lower triangular matrices²:

$$C_\ell^{T,ij} = \sum_{k \in \text{bins}} T_\ell^{ik} \left(T_\ell^{kj} \right)^{Tr}, \quad (4.2)$$

where Tr denotes the transpose and the sum is over all tomographic bins. Drawing samples from a standard normal distribution

$$x^k \sim \mathcal{N}(0, 1), \quad (4.3)$$

the Gaussian field in harmonic space, $g_{\ell m}$, is generated by taking

$$g_{\ell m}^i = \sum_{k \in \text{bins}} T_\ell^{ik} x^k, \quad (4.4)$$

for all $m > 0$. Transforming to configuration space yields:

$$g^i(\theta) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} g_{\ell m}^i Y_{\ell m}(\theta), \quad (4.5)$$

where $Y_{\ell m}(\theta)$ are standard spherical harmonics³. Finally, the lognormal field, $\kappa^i(\theta)$, is generated by exponentiating and shifting the Gaussian realisation:

$$\kappa^i(\theta) = \exp [g^i(\theta)] - \kappa_0^i. \quad (4.6)$$

For each tomographic bin, i , I compute the shift parameter, κ_0^i , by taking a

²In the flat sky approximation – which I assume throughout this chapter – the shear and convergence spectrum are the same, but care would be needed to correctly re-scale the input convergence spectrum by the appropriate ℓ -factor if the flat sky approximation was dropped [27, 76].

³spin-weighted spherical-harmonics for spin 0

weighted average of the shift parameter at each redshift:

$$\kappa_0^i = \int dz n_i(z) \kappa_0^i(z), \quad (4.7)$$

using the fitting formula:

$$\kappa_0^i(z) = 0.008z + 0.029z^2 - 0.0079z^3 + 0.0065z^4, \quad (4.8)$$

derived from simulations [54].

While the fitting formula will have some cosmological dependence, the shift parameter does not affect the power spectrum of the field and will only impact cosmological constraints through the covariance matrix. The non-Gaussian corrections to the covariance matrix already have a sub-dominant impact on parameter constraints [34, 112], thus one would expect that the dependence of these corrections on the cosmology is further sub-dominant. For this reason I choose to ignore the cosmological dependence of the shift parameter.

Once the convergence fields are generated, **Flask** estimates the harmonic coefficients of the field using the estimator [42]:

$$\kappa_{\ell m} = \frac{4\pi}{N_{\text{pix}}} \sum_{p=0}^{N_{\text{pix}}} Y_{\ell m}^*(\theta_p) \kappa(\theta_p), \quad (4.9)$$

which is essentially an average over all pixels in the map. Defining

$$a_{E,\ell m} = \sqrt{\frac{(\ell+2)!}{(\ell-2)!}} \frac{\kappa_{\ell m}}{\ell(\ell+1)}, \quad (4.10)$$

we have [27]:

$$\gamma_{1,\ell,m} = i\gamma_{2,\ell,m} = -a_{E,\ell m}, \quad (4.11)$$

where $\gamma_{1,\ell,m}$ and $\gamma_{2,\ell,m}$ are the harmonic coefficients of the shear field. Then the shear field in configuration space is found by taking the inverse spherical-harmonic

transform of the shear coefficients:

$$\gamma(\theta) = \gamma_1 + i\gamma_2 = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} {}_2Y_{\ell m}(\theta) \gamma_{\ell m}, \quad (4.12)$$

where ${}_2Y_{\ell m}(\theta)$ are the spin-weighted spherical harmonics for spin 2.

A valid covariance matrix between data must be positive-definite, but this is not guaranteed for correlations between tomographic lognormal fields as pointed out in [135]. To overcome this issue, `Flask` perturbs the lognormal fields following the regularisation procedure outlined in Section 3.1 of [135]. Provided that the regularisation is applied to a small number of tomographic bins, it is found in [135] that $C_{\ell}^{\text{reg}}/C_{\ell}^{\text{ln}} \ll 1 \times 10^{-5}$, where C_{ℓ}^{reg} is the recovered regularised spectrum and C_{ℓ}^{ln} is the spectrum recovered from the un-regularised map [135]. In Section 4.5 I verify that this will not impact Stage III parameter constraints.

In Figure 4.2 I plot a single lognormal realisation generated with Pipeline I. I show the masked convergence and components of the shear field in the lowest redshift bin, where the non-Gaussianities are most pronounced and clearly visible. In the convergence map the majority of the pixels take small negative values, but there are rare incidences of large positive convergence. These physically correspond to collapsed high-density structures along the line-of-sight.

4.2.3 Band-limit Bias from the Lognormal Field

Unlike Gaussian fields, lognormal realisations are not band-limited in ℓ [135] (see Section 5.2.2 therein). In particular, Taylor expanding the lognormal convergence field, $\kappa^i(\theta)$, in terms of the Gaussian field, $g^i(\theta)$, yields quadratic and higher order terms in g^i . In harmonic space this mixes different ℓ -modes. When a band-limit is imposed, this biases the lensing spectrum recovered from the map. I discuss this issue further in Section 4.5.

4.3 Cosmological Parameter Inference

In this section I discuss why the standard Gaussian likelihood analysis may lead to bias, and review density-estimation likelihood-free inference.

4.3.1 The Potential Insufficiency of the Gaussian likelihood Approximation in Cosmic Shear

In the standard two-point cosmic shear likelihood analysis, I assume the Gaussian likelihood defined in equation (3.4) and I implicitly make the Anderson-Hartlap [14, 46] correction for the inverse covariance matrix because I generate the covariance matrix from noisy simulations. To see why the Gaussian likelihood assumptions can lead to bias, I now summarise the argument given in [114]. Inside a single bin, the unmasked lensing spectrum can be computed from a shear map by taking an average over m -modes, so that:

$$C_\ell = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} |\gamma_{\ell m}|^2. \quad (4.13)$$

Since the harmonic coefficients, $\gamma_{\ell m}$, are computed as a summation over a large number of pixels, they are Gaussian distributed by the central limit theorem. Squaring a Gaussian random variable gives a gamma distribution – which is left-skewed.

Taking a Gaussian rather than a sum of gamma distributions for the likelihood could bias parameter constraints. Since $S_8^2 = \sigma_8^2(\Omega_m/0.3)$ and Ω_m enter into the shear spectrum amplitude, we would expect these parameters to be ones which are most affected – and biased low. Only in the limit of large ℓ – as the C_ℓ itself becomes the sum over a large number of m -modes – does the central limit theorem kick in and the likelihood become Gaussian.

4.3.2 Density-estimation Likelihood-free Overview

The discussion below closely follows the one given in [7]. All the steps, except the data simulation step, outlined in this section are implemented in the publicly available `pydelfi` code that I use throughout the remainder of this chapter. For cosmological parameters, p , and some observed data, d_{obs} , DELFI estimates the posterior distribution $\mathbb{P}(p|d_{\text{obs}})$. DELFI can be broken into three steps: a simulation step, a compression step and an active learning step. In the simulation step we take a set of points in cosmological parameter space $\{p\}$ and forward model the data, with realistic noise, to form a set of parameter-data pairs $\{(p, d)\}$. In the specific case of this chapter, Pipeline I and Pipeline II will be used to simulate the mock data. Since density estimation is most efficient in lower dimensional spaces, we compress the data to a lower dimensional space while trying to minimise the information loss (see Section 4.3.3 for more details). After compression one is left with the compressed observed data vector, t_{obs} , and a set of parameter-compressed-data pairs, $\{(p, t)\}$. DELFI then learns the conditional density $\mathbb{P}(t|p)$, from the parameter-compressed-data pairs $\{(p, t)\}$, using neural density estimation (see Section 4.3.5 for more details). Once this is known we find the likelihood, $\mathbb{P}(t = t_{\text{obs}}|p)$, by taking a hyper-dimensional slice of the conditional density, $\mathbb{P}(t|p)$. Finally the posterior is computed by multiplying the likelihood by the prior, $\mathbb{P}(p)$. I now review these steps in more detail.

4.3.3 Density-estimation Likelihood-free Compression

Through the remainder of this chapter I compress the C_{ℓ} summary statistic following the compression regime suggested in [12]. Specifically the lensing spectra are compressed into a new vector, t , according to:

$$C_{\ell} \rightarrow t = \nabla_p \mathcal{L}_*, \quad (4.14)$$

where p is the set of cosmological parameters that I am inferring, \mathcal{L}_* is a proposal Gaussian likelihood centred at a fiducial set of parameters, which I take to be:

$$(\Omega_m, h_0, \Omega_b, n_s, S_8, A, \eta) = (0.3, 0.72, 0.96, 0.79, 1., 2.8), \quad (4.15)$$

throughout the remainder of the chapter, where $S_8 = \sigma_8(\Omega_m/0.3)^{0.5}$, A and η are the intrinsic alignment parameters defined in Section 1.7.7 and the other parameters take their standard cosmological definitions. The Fisher information is preserved provided the true likelihood is Gaussian [12], but if it is not, some information will be lost. This is investigated in Section 4.6. It is important to note that although the covariance matrix appears in the compression step through the likelihood, it will not be used in the inference step, so we do not have the stringent accuracy requirements on the covariance matrix as in the standard likelihood analysis. For this reason it may be possible to use computationally inexpensive methods to generate the covariance matrix in the future.

4.3.4 Neural Density Estimators (NDEs) in DELFI

Taking parameter-compressed-data pairs $\{(p, t)\}$ as inputs, DELFI uses a neural network parameterised in terms of a set of node weights, w , to learn the conditional distribution $\mathbb{P}(t|p)$. While a full discussion is beyond the scope of this thesis, I highlight the key steps, as outlined in [7], below.

A particularly simple neural network architecture to understand, and one which is used in DELFI, is a mixture density network (MDN). In this case the conditional density is a sum over Gaussians parameterised in terms of the network weights, w , by:

$$\mathbb{P}(t|p, w) = \sum_{i=1}^N A_i(p, w) \mathcal{N}[\mu_i(p, w), C_i(p, w)], \quad (4.16)$$

where N is the number of Gaussian components with amplitudes A_i , means μ_i and covariance matrices C_i . To train the network, we must define a metric for the performance of the network also called a *loss function*. The Kullback-Leibler

divergence (KL-divergence) defined by:

$$KL(\mathbb{P}^*|\mathbb{P}) = \int dt \mathbb{P}^*(t|p) \ln \left(\frac{\mathbb{P}(t|p)}{\mathbb{P}^*(t|p)} \right), \quad (4.17)$$

where $\mathbb{P}^*(t|p)$ is the true conditional density, would be a good choice. However the target distribution $\mathbb{P}^*(t|p)$ is not known, so DELFI instead uses the negative loss function:

$$-\ln U(w) = - \sum_{i=1}^{N_{\text{samples}}} \ln \mathbb{P}(t_i|p_i, w), \quad (4.18)$$

which is a Monte-Carlo estimate of the KL-divergence up to an additive constant [7]. In practice to avoid overfitting and to negate the fact that we choose a particular network architecture, we train an ensemble of NDEs with different network architectures. Then the final estimate for the conditional density is given as a weighted average over all NDEs:

$$\mathbb{P}(t|p, w) = \sum_{i \in \text{networks}} \beta_i \mathbb{P}_i(t|p, w), \quad (4.19)$$

where $\mathbb{P}_i(t|p, w)$ is the estimate of the conditional density from the i^{th} NDE and the weights, β_i , are determined by the performance of each NDE (see [7] for more details).

4.3.5 Active Sequential Learning

DELFI divides the inference task into a set of training steps. As an initial guess for the conditional distribution, DELFI takes the multivariate Gaussian:

$$\mathbb{P}(t|p) = \mathcal{N}(t|p, F^{-1}), \quad (4.20)$$

where F^{-1} is the inverse of the Fisher matrix defined in equation (3.6). At each step thereafter, `pydelfi` trains each NDE on a set of parameter realisation pairs $\{p_i, t_i\}$ drawing samples from the conditional density of the previous step to ensure

that the highest density regions are the most finely sampled. Ten percent of the samples in the first step are saved as a validation set to avoid overfitting, using early stopping. Given a large enough computer, all the simulations in each training step could be run in parallel. In all that follows, I use 20 training steps with 100 simulations per step. Otherwise I use all the default settings in `pydelfi`.

4.4 The Full Forward Model

In this section I present Pipeline I before using it to generate mock Stage IV data. I then run analysis DA1 to recover the input cosmology. This is used to verify that I can recover the input cosmology with DELFI and to estimate the number of simulations needed for a Stage IV experiment.

4.4.1 The Mask

I use a typical Stage IV survey mask shown in Figure 4.2. All pixels lying within 22.5 deg of either the galactic or ecliptic planes are masked. This leaves 14,490 deg² of unmasked pixels which, as a fraction of the full sky, corresponds to $f_{\text{sky}} = 0.35$.

4.4.2 Shot-Noise Model

The magnitude of the noise, γ_p , for each pixel, p , is drawn from a Gaussian distribution [9]:

$$\gamma_p \sim \mathcal{N} \left(0, \frac{\sigma_\epsilon}{\sqrt{\hat{N}_P}} \right) \quad (4.21)$$

where \hat{N}_P is the number of galaxies in each pixel, the orientation is angle is drawn from a uniform distribution, and I take the intrinsic shape dispersion as $\sigma_\epsilon = 0.3$ [25] and use 30 galaxies per arcmin² throughout. Gaussian noise in each pixel is a good approximation since in my simulations there are a large number of galaxies in each pixel, so the central limit theorem applies.

4.4.3 Forward Modelling the Mask

One advantage of performing inference with full forward models of the data is that there is no need to deconvolve the mask using the pseudo- C_ℓ method [6, 24, 133]. This is both computationally simpler and avoids the risk of bias from inaccurate deconvolution that is present in the standard likelihood analysis.

Given two masked shear fields $a(\theta)$ and $b(\theta)$, a naïve estimate of the lensing spectrum is the pixel pseudo- C_ℓ spectrum:

$$\tilde{C}_\ell^{\text{pix},EE} = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} \langle a_{lm}^E b_{lm}^E \rangle, \quad (4.22)$$

where the tilde is used to denote the fact that I have not corrected for the mask, while the ‘pix’ superscript is a reminder that I have not accounted for the pixel window function. Analogous expressions are easily found for the EB , BE and BB spectra.

In an unmasked field, lensing by large scale structure will only induce power in the EE spectra, but to retain information leaked into the EB , BE and BB spectra due to the presence of a mask, in Pipeline I, I use:

$$\tilde{C}_\ell^{\text{pix}} = \tilde{C}_\ell^{\text{pix},EE} + \tilde{C}_\ell^{\text{pix},EB} + \tilde{C}_\ell^{\text{pix},BE} + \tilde{C}_\ell^{\text{pix},BB}, \quad (4.23)$$

as the estimator. I compute this using the public code `HEALpy` [42, 41].

In a future pipeline it may still be desirable to use the pseudo- C_ℓ formalism to avoid mixing between E and B -modes, allowing us to immediately remove B -modes induced by unknown systematics. As long as the data and theory are treated in the same way, the pseudo- C_ℓ formalism will not introduce bias, as it could in the standard likelihood analysis.

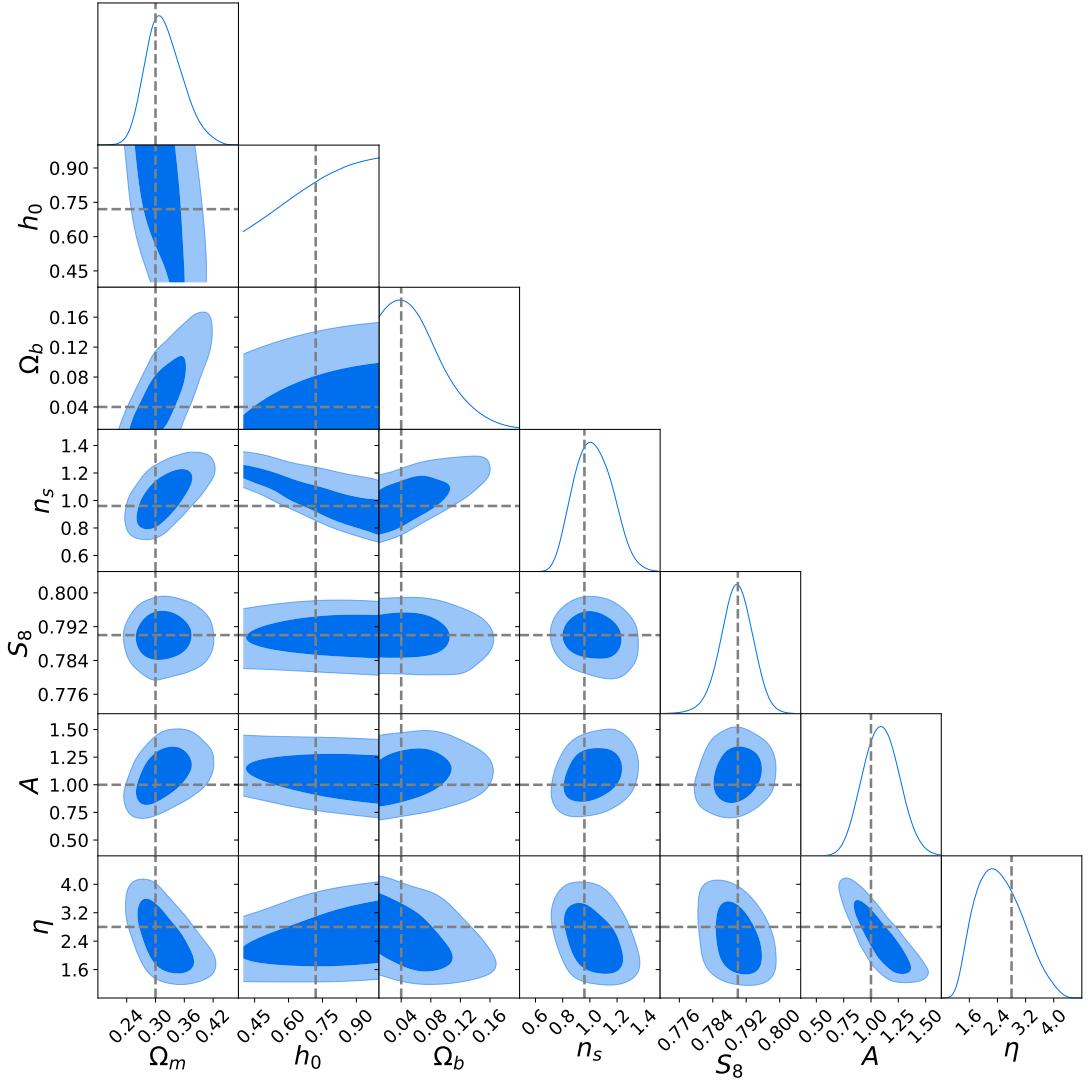


Figure 4.3: 68% and 95% credible region parameter constraints found with DELFI analysis DA1 after 1000 simulations, for a mock Stage IV experiment. I confirm that the input cosmology is recovered within statistical errors. The convergence is plotted in Figure 4.4. In a realistic situation there may be a larger number of nuisance parameters. One should not expect this to dramatically slow convergence because we could ‘nuisance harden’ the data compression step, and only learn the posterior for the parameters of interest. For more details see [11].

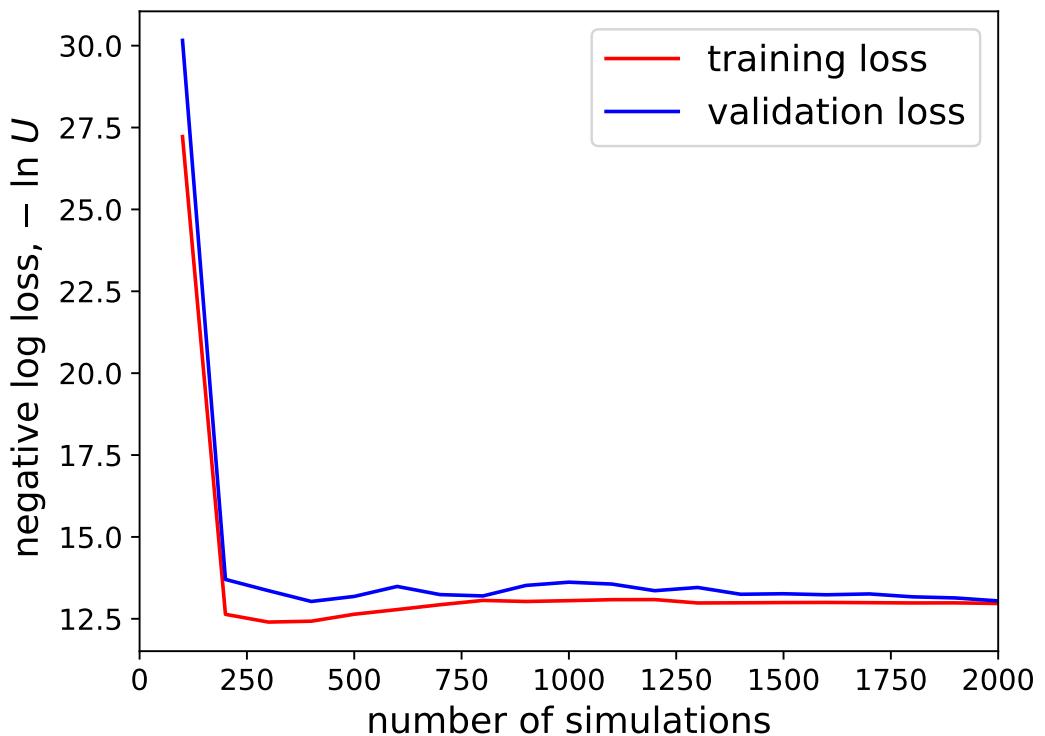


Figure 4.4: The negative loss function defined in equation (4.18) for the training and validation sets as a function of the number of simulations. This suggests that $\mathcal{O}(1000)$ simulations are needed for a Stage IV experiment. As this is similar to the number found in [7], which only considered a simple Gaussian field forward model with no intrinsic alignments, the convergence rate should be fairly insensitive to model improvements in the future in preparation for Stage IV experiments.

4.4.4 Mimicking a Stage IV Experiment

To estimate the number of simulations needed for a Stage IV experiment and ensure that pipeline recovers the input cosmology, I first generate mock data using Pipeline I. I use 6 tomographic bins sampling 15 logarithmicly spaced ℓ -bins in the range $\ell \in [10, 1000]$. I then run `pydelfi` to estimate the posterior distribution of the cosmological parameters with this mock data. The final parameter constraints for a Λ CDM cosmology with two nuisance intrinsic alignment parameters are shown in Figure 4.3, confirming that I recover the input parameters within errors.

In Figure 4.4 I plot the negative loss function, defined in equation (4.18), for the training and validation sets. Both have converged within $\mathcal{O}(1000)$ simulations. This is similar to the number found in the simple Gaussian field pipeline presented in [7], suggesting that the precise details of the forward model, and in particular the field distribution, does not significantly increase the required number of simulations.

When working with real data in the future, one may require a large number of nuisance parameters. This should not dramatically increase the number of simulations needed, since we can always tune the data compression to maximise the information retention of the parameters of interest, following the procedure in [11].

Each simulation takes approximately 33 minutes on a single thread of a 1.8 GHz **Intel Xeon (E5-2650Lv3) Processor**, so if run on 100 threads in parallel, the total simulation time during active learning is only 10 hours. Running on even more threads would further reduce the total run-time because many of the individual modules in the pipeline are multithreaded (e.g `Flask`).

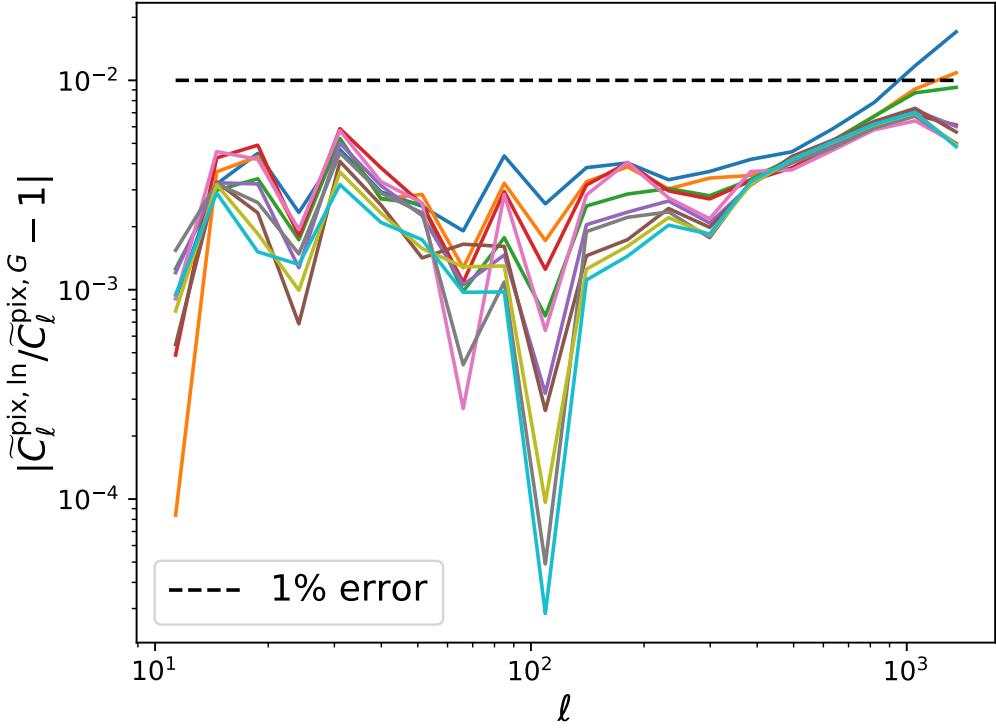


Figure 4.5: The coloured lines show the absolute value of the difference between the average recovered cross and intra-bin spectra from 100 lognormal and 100 Gaussian realisations (4 tomographic bins, $N_{\text{side}} = 512$ and $\ell \in [10, 1535]$). The difference is due to the band-limit and regularisation bias in the lognormal field. The combined bias is safely below 1% for nearly all data points.

4.5 Prospects for Stage III Data

In this section I discuss the viability of applying analysis DA1 to existing Stage III data. I assume a circular mask of 4951 deg^2 , similar to the final coverage of the Dark Energy Survey [128] with 10 galaxies per arcmin^2 , for the remainder of this section. I also use Pipeline I throughout this section – except where modifications are explicitly stated.

4.5.1 Validating the Lognormal Simulations

Lognormal fields were used to generate the covariance matrix in the recent Dark Energy Survey Year 1 analysis [130]. The authors found no difference in parameter constraints between this analysis and one which used a halo model to generate the

covariance matrix – but to verify that Pipeline I is ready for Stage III data, one must also ensure that the C_ℓ recovered from the maps is unbiased.

Provided an accurate input C_ℓ , the only bias in Pipeline I comes from regularising the map (see Section 4.2.2). The band-limit bias of the lognormal field should not be a problem, since imposing a band-limit would affect the data in the same way. However this assumes that the true field is exactly lognormal, so I check to ensure that the combined effect of regularisation and imposing a band-limit is small. This statement is quantified by finding the difference between the average recovered pixelated spectra from 100 Gaussian (where no band-limit bias or regularisation bias is present) and 100 lognormal simulations. Each 4-tomographic bin simulation takes approximately 15 minutes on a single thread and the difference in the recovered spectra is shown in Figure 4.5. The bias is safely below 1% in all but three data points. This confirms that once minor updates have been made (see next subsection), the pipeline will be ready for use on today’s data.

4.5.2 Model Improvements

Only a small number of adjustments must be made to DA1 to apply this analysis to existing Stage III data. These are:

- I must accurately account for baryonic physics. This can be handled using a halo model code [97], in combination with the k -cut cosmic shear approach presented in Chapter 2.
- Several nuisance parameters must be introduced. To match the DESY1 analysis, I must allow for free multiplicative and additive shear biases and photo-z bias parameters in each bin. This will increase the number of nuisance parameters, so to avoid excessive computational costs, we must ‘nuisance harden’ [11] the data compression step. This means that we compress the data in such a way that we only preserve information about the parameters of interest. This will allow us to continue working in a low dimensional

data space, which is key to the DELFI technique.

4.6 Testing the Gaussian Likelihood Approximation

In this section I compare DELFI, which makes no assumption about the underlying likelihood, to the standard Gaussian likelihood analysis by running analyses DA2 and LA, on the same mock Stage IV data. I use Pipeline II to generate the mock data, produce the covariance matrix and generate the forward models in analysis DA2. Because DELFI does not assume any particular likelihood, the differences in the resulting parameter constraints are only due to the Gaussian likelihood assumption in LA. Since it is not possible to forward model everything in analysis LA, care must be taken to ensure that the band-limit bias, deconvolving the mask, deconvolving the pixel window function and subtracting the shot-noise does not lead to additional bias between the two analyses. Controlling for these effects is described in the first subsection.

4.6.1 Modelling Choices In Pipeline II

To avoid the band-limit bias (see Section 4.2.3) I choose to use a Gaussian field, rather than the lognormal field.

I choose not to apply a mask in DA2 as I have found that using the pseudo- C_ℓ method [6, 24] (with the public code **NaMaster** [6]) can bias parameter constraints, with my choice of **HEALpix**⁴ N_{side} and ℓ_{max} by up to 1σ . Instead I use an unmasked full-sky map and adjust the galaxy number density so that total number of galaxies and hence the signal-to-noise remains unchanged, as for the Stage IV experiment considered in Section 4.4.

In LA I take the C_ℓ , with the shot-noise subtracted as the data vector. Thus I subtract off the expected value of the noise in DA2. This is computed by running

⁴<https://sourceforge.net/projects/healpix/>

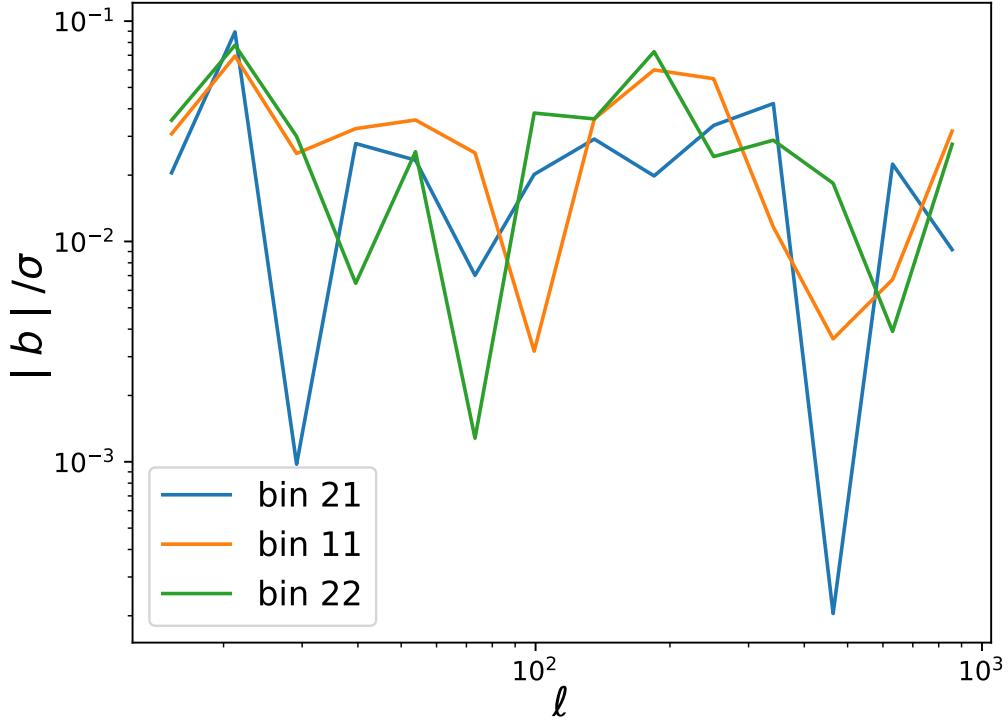


Figure 4.6: The absolute value of the bias, $|b|$, due to imperfect pixel-window deconvolution and noise subtraction relative to the statistical error, σ , from 500 Pipeline II simulations. The bias is so small that the comparison between DELFI and the likelihood analysis presented in Section 4.6.2 will be unaffected.

500 noise-only simulations, as in the analysis of [53].

I must also account for the fact that the shear spectra are computed on pixelised maps – that is, I must deconvolve the pixel window function, w_ℓ , which is defined in [64]. I compute w_ℓ using `HEALpix` which assumes that the scale of the signal is large relative to the pixel scale and that all pixels are the same shape, and also interpolates the function from a lower N_{side} (see the `HEALpix` documentation for more information). The window-corrected spectrum, C_ℓ , is given in terms of the spectrum computed from a pixelized map, C_ℓ^{pix} , by:

$$C_\ell = w_\ell^{-2} C_\ell^{\text{pix}}. \quad (4.24)$$

By running 500 Gaussian field simulations I have confirmed that the combined bias from deconvolving the pixel window function and subtracting the shot-noise

is small, so that I can make a fair comparison between DA2 and LA. This is demonstrated in Figure 4.6 where the absolute value of the bias, $|b|$, is small relative to the statistical error, σ , with $|b|/\sigma < 0.1$ for all data points.

4.6.2 Impact of the Gaussian Likelihood Approximation

To test the impact of the Gaussian likelihood approximation I first generate 1000 mock data realisations using Pipeline II. I take 15 logarithmically spaced ℓ -bins in the range $[10, 1000]$ and I restrict my attention to $S_8 - \Omega_m$ plane, where we would expect to find a bias in the Gaussian likelihood analysis if the true likelihood was skewed. To cut computational time, I use only two tomographic bins. Since the parameters S_8 and Ω_m primarily impact the amplitude of the shear spectrum, not too much constraining power should be lost with this choice [119, 125].

For three random data realisations, I run a DELFI and a Gaussian likelihood analysis, and the resulting posteriors are shown in Figure 4.7. Each subplot corresponds to one of the three realisations.

In all three cases the DELFI and Gaussian likelihood contours are very similar. This suggests the Gaussian likelihood assumption does not bias parameter constraints in the $S_8 - \Omega_m$ plane and furthermore that no information is lost in the compression defined in equation (4.14).

I confirm and quantify this statement by sampling the maximum likelihood estimator (MLE) distribution assuming a Gaussian likelihood, using the 1000 data realisations generated earlier. For each realisation, the MLE is found using the Nelder-Mead algorithm built into `scipy` and wrapped into `Cosmosis` using the default settings. The resulting MLE distribution is shown in Figure 4.8. The input cosmology lies almost exactly at the centre of the 68% credible region which implies that there is no measurable bias from the Gaussian likelihood approximation.

I stress that these conclusions only hold for the C_ℓ analysis presented in this work. In particular, the ℓ -binning strategy matters. By binning ℓ -modes we are taking a sum over random variables so, by the central limit theorem, broader bins

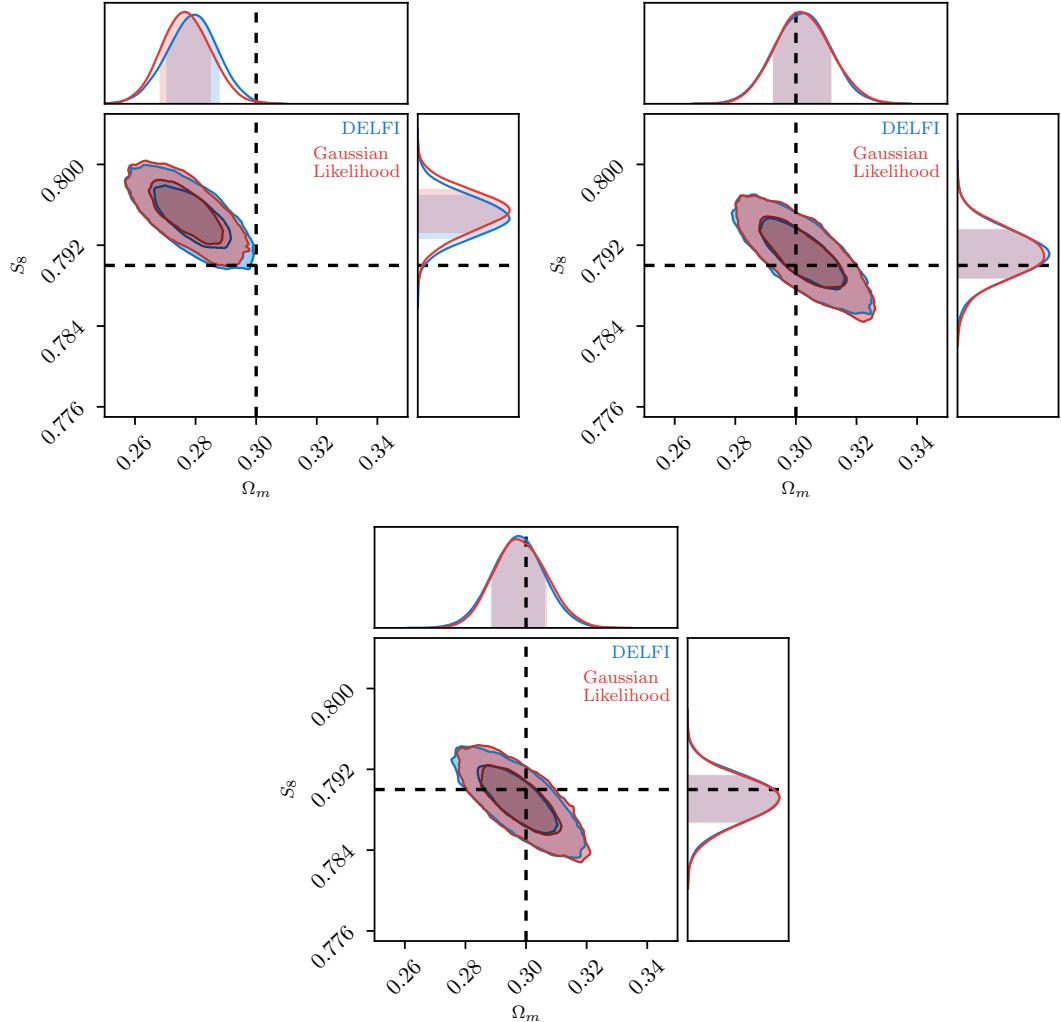


Figure 4.7: The 68% and 95% credible region parameter constraints for three random data realisations found using a Gaussian likelihood analysis and DELFI. The mock data input cosmology is labeled by black dotted lines. Only in the first realisation, does the input cosmology lie outside the 68% credible region – but statistically this is to be expected for a small number of realisations. The contours found using the two different analyses are very similar for all three data realisations suggesting that the Gaussian likelihood approximation has negligible impact. This is confirmed in Figure 4.8.

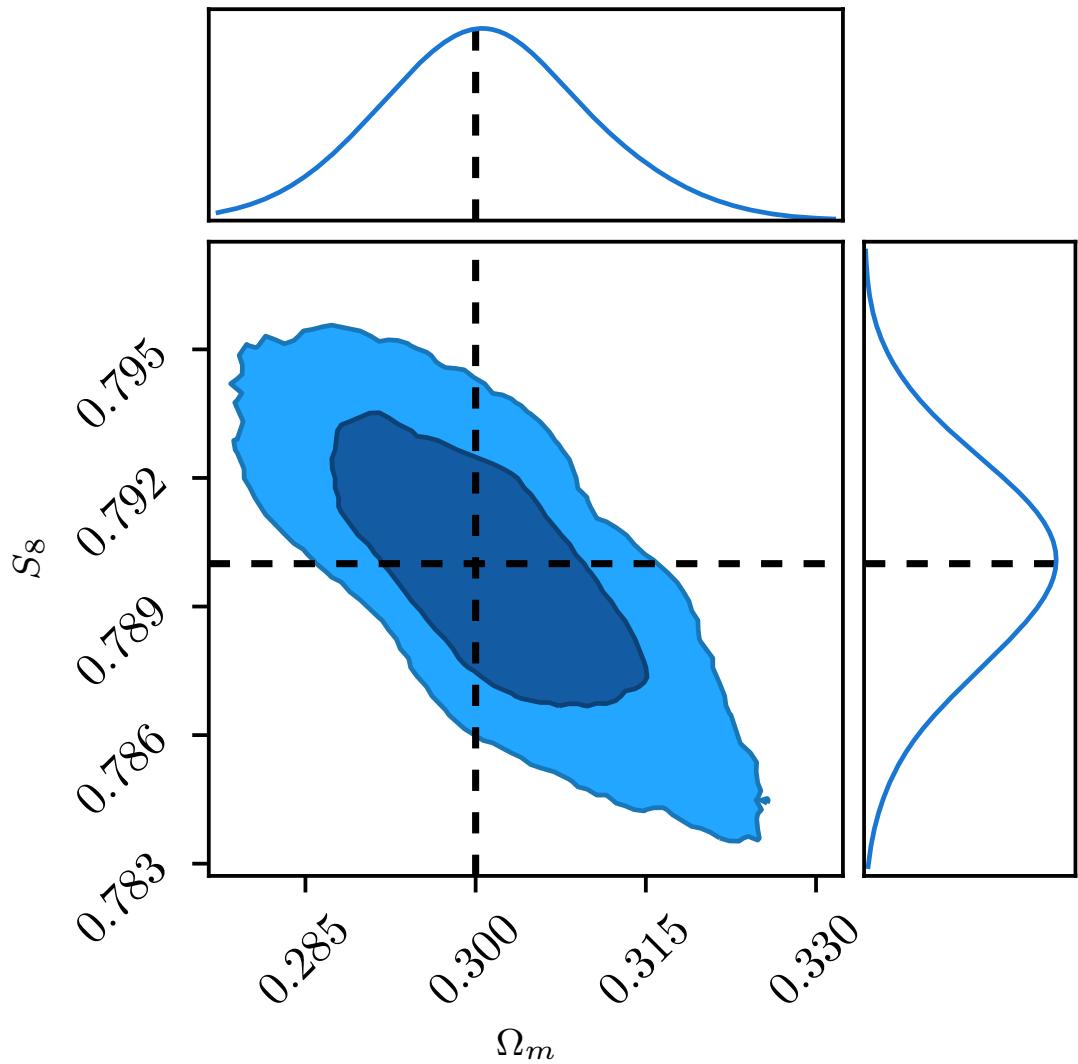


Figure 4.8: The 68% and 95% credible region of the MLE distribution, assuming a Gaussian likelihood. The value of the input cosmology is indicated by the black dotted lines, and lies at the center of the contours. This implies that the Gaussian likelihood approximation does not lead to any measurable bias in our setup.

correspond to more Gaussianised data. The Gaussian likelihood approximation may be important for much narrower ℓ -bins.

Similarly, the Gaussian likelihood approximation might be important in a correlation function analysis (see Section 3.2). However a recent paper, appearing after the paper on which this chapter is based, has suggested that here too this approximation is negligible [85].

In the future, collaborations which continue to use a likelihood analysis must also explicitly check the impact of the Gaussian likelihood approximation on the dark energy parameters w_0 and w_a . These parameters are more sensitive to lower ℓ -modes than σ_8 and Ω_m where non-Gaussianity is important.

4.7 Future Prospects

I now review the main known cosmic shear systematics which must eventually be included in the full forward model. To account for many of these effects one would first need to take the base-model presented in this work to ‘catalog level’. This can be done by first generating a consistent density field – either with `Flask` or by taking the difference between two neighbouring tomographic bins – and then populating the density field with a realistic population of galaxies [100] assuming a biased tracer model (e.g [35]). Cosmic shear systematics break down into four broad categories: data-processing, theoretical, astrophysical and instrumental systematics.

On the data-processing side, accurately measuring the shape and photometric redshift of galaxies is the primary challenge. Rather than using the best fit parameters for each galaxy, we can sample the posterior on each galaxy as in a Bayesian hierarchical model [9] to propagate the measurement uncertainty into the final parameter constraints, as suggested in [7].

Two important theoretical systematics are the reduced shear correction [17, 32] and magnification bias [44, 87]. The former correction accounts for the fact that

we measure the reduced shear $\gamma/(1 - \kappa)$ with a weak lensing experiment. In a likelihood analysis, this can be handled using a perturbative expansion as in [32, 117]. This is slow and requires us to rely on potentially inaccurate fitting functions for the lensing bispectrum. I discuss an alternative approach in the Appendix E. Meanwhile the magnification bias accounts for the fact that galaxies of the same luminosity can fall above/below in regions of high/low lensing magnification. In both cases, these systematics can be easily handled with full forward models of consistent shear and convergence fields.

The two dominant instrumental systematics are the telescope’s point spread function (PSF) [95] and the effect of charge transfer inefficiency (CTI) in the charge-couple devices (CCDs) [95, 106]. Efforts are underway to build pipelines which characterise these effects in upcoming experiments (e.g [132]). Integrating these pipelines into the DELFI pipeline would enable the propagation of instrumental errors through to the final parameter constraints.

On the astrophysical side, the two dominant systematics are the impact of baryons on the density field [116] and the intrinsic alignment of galaxies [58, 68]. For Stage IV data, forward models will likely have to be based on high-resolution N-body lensing simulations [63, 69] to include the effects of baryons. Even with today’s highest resolution simulations the impact of baryons is still uncertain [60], so it will likely be necessary to optimally cut (see Chapter 2) or marginalise out uncertain scales [60]. Meanwhile more sophisticated intrinsic alignment models which account for different alignment behaviour by galaxy type [111] will need to be included.

Eventually higher-order statistics such as peak counts and the shear bispectrum can be added. Since DELFI automatically handles multiple summary statistics in a unified way, the constraints will be tighter than doing the two-point and higher-order statistic analyses separately.

4.8 Conclusion

Even though I have found that the Gaussian likelihood approximation will have a negligible impact on Stage IV parameters constraints, density-estimation likelihood-free inference offers the possibility of performing rapid parallel inference on full forward realisations of the shear data. In the future, this could be used to seamlessly handle astrophysical and detector systematics – at a minimal computational cost.

I have taken the first steps towards this goal by developing a pipeline to rapidly generate realistic non-Gaussian shear data, including the impact of intrinsic alignments. I have integrated this pipeline into `pydelfi`, to confirm that inference with cosmic shear data using DELFI is feasible, requiring only $\mathcal{O}(1000)$ simulations for Stage IV data. I have also verified that the current pipeline is accurate enough for today’s Stage III data. Even though the lognormal field approximation will likely need to be dropped for Stage IV data, as the pipeline is computationally inexpensive, in the future it will be useful for quickly determining which systematics are important as an alternative to a Fisher analysis. However, unlike a Fisher analysis, we will be able to test the impact of systematics which can not be easily modelled analytically.

I conclude that DELFI has a promising future in cosmic shear studies. Since developing fast simulations that fully integrate all relevant astrophysical, detector and modelling effects is the primary hurdle, developing these simulations should be the immediate priority.

Chapter 5

Conclusion

A central theme of this thesis has been that we should not use yesterday’s techniques on tomorrow’s data. In this spirit, I have extended the Gaussian likelihood analysis of the shear two-point statistic in a number of important ways, in preparation for Stage IV surveys.

The k -cut cosmic shear technique is a method which efficiently removes sensitivity to small poorly modelled scales. Part of the advantage of this technique is its simplicity. It is independent of any particular baryonic feedback model and should be easy to implement in all upcoming shear experiments. In the near future we can handle the uncertain baryonic physics simply by taking a k -cut at the scale where competing models start to disagree. The k -cut cosmic shear method is no silver bullet, however. Small scales also contain cosmological information, but this can only be accessed by improving models of baryonic physics. As models of the matter power spectrum improve, the k -cut can gradually be increased until there is no longer a need for the method in w LCDM studies. However, there may never be enough computational resources to compute the impact of baryons on nonlinear structure growth for all modified gravity theories. In these circumstances one may always need to cut scales – and k -cut cosmic shear does this optimally.

The non-parametric technique presented in Chapter 3 offers a unique way to

test LCDM without assuming any particular extension to the model. In the short term the technique should be applied to the Dark Energy Survey Year 1 data to see if the discrepancy in the expansion history at low redshifts between the non-parametric reconstruction and other distance measurements persists, or if this is just a systematic in the CFHTLenS shear catalog. The formalism should also be adapted to simultaneously measure the non-parametric amplitudes and cosmological parameters, as the assumption of a fiducial cosmology may no longer be valid as data improves. One could then marginalise out the cosmological parameters to obtain the non-parametric information. This method could also naturally be extended to photometric galaxy clustering, galaxy-galaxy lensing and CMB lensing.

Finally I have shown that performing inference with full forward models of the data offers a natural way to propagate the effects of systematics into the final parameter constraints. I have confirmed this method is feasible requiring $\mathcal{O}(10^3)$ simulations. After applying this method to real data for the first time, developing fast accurate lensing simulations which fully integrate all known systematics effects is the primary task for the future.

Appendix A

Calculation of Light Deflection in General Relativity

To compute the deflection angle due to gravity I closely follow a particularly clear derivation given in [98] using the result from classical optics that the total travel time of a light ray between two points is given by:

$$t = \frac{1}{c} \int d\ell n [\mathbf{x}(\ell)], \quad (\text{A.1})$$

where $n [\mathbf{x}(\ell)]$ is the refractive index along the path and according to Fermat's Principle the path travelled will extremise the time so that:

$$\delta \frac{1}{c} \int d\ell n [\mathbf{x}(\ell)] = 0. \quad (\text{A.2})$$

Now the metric for a static potential is given by:

$$ds^2 = - \left(1 + 2\frac{\Phi}{c^2}\right) c^2 dt^2 + \left(1 - 2\frac{\Phi}{c^2}\right) d\mathbf{x}^2, \quad (\text{A.3})$$

where Φ is the Newtonian gravitational potential and c is the speed of light in vacuum. Along the light ray $ds^2 = 0$. Hence,

$$v(\mathbf{x}) = \frac{dx}{dt} = c \sqrt{\frac{1 + 2\frac{\Phi}{c^2}}{1 - 2\frac{\Phi}{c^2}}} \approx c \left(1 + 2\frac{\Phi}{c^2}\right), \quad (\text{A.4})$$

where $x = |\mathbf{x}|$. Then the refractive index is:

$$n[\mathbf{x}(\ell)] = \frac{c}{v(\mathbf{x})} = \frac{1}{1 + 2\frac{\Phi}{c^2}} \approx 1 - 2\frac{\Phi}{c^2}. \quad (\text{A.5})$$

Writing the refractive index, n , in terms of a parameter λ , equation (A.2) can be rewritten as:

$$\delta \int d\lambda n[\mathbf{x}(\lambda)] \frac{dx}{d\lambda} = 0. \quad (\text{A.6})$$

The equations of motion are given by the Euler-Lagrange equations:

$$\frac{d}{d\lambda} \left(\frac{\partial \mathcal{L}}{\partial \dot{\mathbf{x}}} \right) - \frac{\partial \mathcal{L}}{\partial \mathbf{x}} = 0, \quad (\text{A.7})$$

where the lagrangian $\mathcal{L} = n(\mathbf{x}(\lambda)) \dot{x}$. Now,

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \dot{\mathbf{x}}} &= n \frac{\dot{\mathbf{x}}}{\dot{x}} \\ \frac{\partial \mathcal{L}}{\partial \mathbf{x}} &= \dot{x} \nabla n. \end{aligned} \quad (\text{A.8})$$

Choosing λ so that $\frac{dx}{d\lambda} = 1$, the tangent vector to the path, \mathbf{e} , is $\mathbf{e} = \dot{\mathbf{x}}$. The Euler-Lagrange Equations imply

$$\frac{d}{d\lambda} (n\mathbf{e}) - \nabla n = 0. \quad (\text{A.9})$$

Expanding, rearranging and applying the chain rule gives:

$$n\dot{\mathbf{e}} = \nabla n - \mathbf{e} \cdot (\nabla n \cdot \mathbf{e}). \quad (\text{A.10})$$

We recognise the second term as the gradient of the refractive index along the path, $\nabla_{\parallel}n$, so that after rearranging we are left with:

$$\dot{\mathbf{e}} = \frac{1}{n} \nabla_{\perp} n = \nabla_{\perp} \log n. \quad (\text{A.11})$$

Assuming a weak field so that $\Phi/c^2 \ll 1$ and Taylor expanding the logarithm to first order yields:

$$\dot{\mathbf{e}} = -\frac{2}{c^2} \nabla_{\perp} \Phi. \quad (\text{A.12})$$

Thus the total deflection angle over the path, $\hat{\boldsymbol{\alpha}}$, is given by:

$$\hat{\boldsymbol{\alpha}} = \frac{2}{c^2} \int d\lambda \nabla_{\perp} \Phi. \quad (\text{A.13})$$

Appendix B

Deflection angle in the observer's frame

We know from equation (A.13) that as viewed by the source, the differential gravitational deflection due to matter at comoving angular diameter distance, r' , is:

$$d\hat{\alpha} = -\frac{2}{c^2} |\nabla_{\perp} \Phi| dr', \quad (B.1)$$

where we choose our coordinate system to lie in the plane of the deflection so that we can drop the vector notation. From the observer's perspective the perpendicular displacement is:

$$dx = (r - r') |\hat{d}\alpha| . \quad (B.2)$$

Meanwhile, assuming the small angle approximation, the deflection angle in observer's frame, $d\alpha$, is:

$$d\alpha \approx \tan d\alpha = \frac{dx}{r - r'} . \quad (B.3)$$

Substituting equation (B.1) into equation (B.2) and then into equation (B.3) and integrating over the path implies:

$$\alpha = \frac{2}{c^2} \int_0^r dr' \frac{(r - r')}{r} \nabla_{\perp} \Phi(\mathbf{x}, r') , \quad (B.4)$$

where \mathbf{x} gives the angular coordinate on the sky and we have reintroduced the vector notation.

Appendix C

Computation of the Lensing Spectrum

The expectation of (1.64) can be rewritten as:

$$\gamma_{\ell m}(\eta) = \int d^3\mathbf{r} \bar{n}(\mathbf{r}) \gamma(\mathbf{r}) W_\ell(\eta, r) {}_2Y_{\ell m}(\boldsymbol{\theta}), \quad (C.1)$$

where $\bar{n}(\mathbf{r})$ gives the mean density of observed galaxies, $\boldsymbol{\theta}$ gives the angle on the sky, r is the radial distance and η parameterises the weight. This can be re-expressed by:

$$\gamma_{\ell m}(\eta) = \frac{1}{4\pi} \int d^2\boldsymbol{\theta} dz dz_p p(z|z_p) n(z_p) \gamma(\mathbf{r}) W_\ell(\eta, r) {}_2Y_{\ell m}(\boldsymbol{\theta}), \quad (C.2)$$

where $n(z_p)$ gives the radially distribution of galaxies in the survey and $p(z|z_p)$ defined in equation (1.62) accounts for the photometric redshift error. The definition of the shear is generalised to the celestial sphere as:

$$\gamma = \frac{1}{2} \vec{\partial} \vec{\partial} \phi(\mathbf{r}) \text{ and } \gamma^* = \frac{1}{2} \vec{\partial} \vec{\partial} \phi(\mathbf{r}), \quad (C.3)$$

where $*$ denotes complex conjugation and \eth is the ‘edth’ operator. See [27] for more details. Hence,

$$\gamma_{\ell m}(\eta) = \frac{1}{8\pi} \int d^2\boldsymbol{\theta} dz dz_p p(z|z_p) n(z_p) [\eth \eth \phi(\mathbf{r})] W_\ell(\eta, r) {}_2Y_{\ell m}(\boldsymbol{\theta}). \quad (\text{C.4})$$

from which it follows that:

$$\gamma_{\ell m}(\eta) = \frac{1}{4\pi c^2} \int d^2\boldsymbol{\theta} dz dz_p p(z|z_p) n(z_p) W_\ell(\eta, r) {}_2Y_{\ell m}(\boldsymbol{\theta}) \int dr' F(r, r') \eth \eth \Phi(\mathbf{r}), \quad (\text{C.5})$$

using equation (1.45). Now we expand the potential, Φ , using a spherical-Bessel transform so that:

$$\Phi(\mathbf{r}) = \sqrt{\frac{2}{\pi}} \int k dk \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \Phi_{\ell m}(k) j_\ell(kr) Y_{\ell m}(\boldsymbol{\theta}), \quad (\text{C.6})$$

where $j_\ell(kr)$ are spherical Bessel functions. In a curved cosmology we would need to replace the Bessel functions with hyperspherical Bessel functions (see [126] for more details). Using the Poisson equation:

$$\nabla_r^2 \Phi(\mathbf{r}; z) = \frac{3\Omega_m H_0^2}{2a(t)} \delta(\mathbf{r}; z), \quad (\text{C.7})$$

where we have assumed a LCDM universe for the prefactor, and noting that spherical harmonics and spherical Bessel functions are eigenfunctions of the Laplace operator, so that:

$$(\nabla_r^2 + k^2) j_\ell(kr) Y_{\ell m}(\theta, \phi) = 0, \quad (\text{C.8})$$

implies the relation:

$$\Phi_{\ell m}(k; z) = -\frac{3\Omega_m H_0^2}{2k^2 a(t)} \delta_{\ell m}(k; z). \quad (\text{C.9})$$

Then using the identities:

$$\bar{\partial} \bar{\partial} Y_{\ell m}(\boldsymbol{\theta}) = \partial \partial Y_{\ell m}(\boldsymbol{\theta}) = \sqrt{\frac{(\ell+2)!}{(\ell-2)!}}^2 Y_{\ell m} \quad (\text{C.10})$$

and

$$\int d^2\boldsymbol{\theta}'_s Y_{\ell' m'}^*(\boldsymbol{\theta})_s Y_{\ell m}(\boldsymbol{\theta}) = \delta_{ss'}^K \delta_{mm'}^K \delta_{\ell\ell'}^K, \quad (\text{C.11})$$

plugging equation (C.6) into equation (C.9) then into equation (C.5) and finally into equation (1.63) and making the equal time approximation (see [72] for more details):

$$P(k; z, z') = [P(k; z)P(k; z')]^{1/2}, \quad (\text{C.12})$$

where in analogy with (1.39) we define the unequal time power spectrum by:

$$\langle \delta(k; z)\delta(k; z') \rangle = (2\pi)^3 \delta(\mathbf{k} - \mathbf{k}') P(k; z, z'), \quad (\text{C.13})$$

we at last we arrive at the expression for the lensing spectrum given in equation (1.65).

Appendix D

Derivation of the Shot-noise

The shot-noise in the cosmic shear power spectrum comes from Poisson sampling the shear field at the location of observed galaxies, with source ellipticity ϵ . Hence it scales as $1/N_g$, where N_g is the number of observed galaxies in the survey. The number of galaxies can be written as $N_g = An_{\text{eff}}$, where A is the area of the survey and n_{eff} is the effective number density. Writing the variance of galaxy ellipticities as $\sigma_\epsilon^2 = \langle \epsilon \epsilon \rangle$ and defining:

$$n_{\ell m}(\eta) = \sqrt{\frac{2}{\pi}} \sum_g \epsilon_{g2} Y_{\ell m}(\boldsymbol{\theta}) W_\ell(\eta, r), \quad (\text{D.1})$$

where the sum is over all galaxies g , and the shot-noise for the generalised spherical lensing spectrum is:

$$\begin{aligned} N_\ell(\eta, \eta') &= \frac{1}{An_{\text{eff}}} \langle n_{\ell m}(\eta) n_{\ell' m'}^*(\eta') \delta_{gg'} \rangle \delta_{\ell\ell'}^K \delta_{mm'}^K \\ &= \frac{2}{\pi An_{\text{eff}}} \int d^3\mathbf{r} d^3\mathbf{r}' \bar{n}(\mathbf{r}) \langle \epsilon \epsilon' \rangle \delta^3(\mathbf{r} - \mathbf{r}') \delta_{\ell\ell'} \delta_{mm'}^K Y_{\ell m}(\boldsymbol{\theta})_2 Y_{\ell m}^*(\boldsymbol{\theta}') W_\ell(\eta, r) W_{\ell'}(\eta', r') \\ &= \frac{2\pi\sigma_\epsilon^2}{\pi An_{\text{eff}}} \int d^3\mathbf{r} \bar{n}(\mathbf{r})_2 Y_{\ell m}(\boldsymbol{\theta})_2 Y_{\ell m}^*(\boldsymbol{\theta}) W_\ell(\eta, r) W_{\ell'}(\eta', r). \end{aligned} \quad (\text{D.2})$$

Using the orthogonality relations for the spherical-harmonics (see equation (C.11)) and writing $d^3\mathbf{r}' \bar{n}(\mathbf{r}) = dz d^2\boldsymbol{\theta} n(z)/4\pi$ yields:

$$N_\ell(\eta, \eta') = \frac{\sigma_\epsilon^2}{2\pi^2 An_{\text{eff}}} \int dz n(z) W_\ell(\eta, r[z]) W_\ell(\eta', r[z]). \quad (\text{D.3})$$

Appendix E

The Reduced Shear Correction From Forward Models

In this Section I show how the forward modelling technique presented in Chapter 4 can be used to naturally account for the reduced shear correction (see Section 1.7.1). This work will appear in [Deshpande et al. in prep]. I outline only my contribution to this paper except where otherwise stated.

Using Pipeline I presented in Chapter 4, I generate 100 shot-noise-free realisations of consistent shear and convergence fields at map-level, from which it is easy to estimate the reduced shear spectrum. I take the same model assumptions as in Chapter 4, but do not include intrinsic alignment contributions to the lensing spectrum. I estimate the reduced shear correction by taking the average over all realisations.

I compare with the results of the perturbative expansion based on the bispectrum, provided by Anurag Deshpande (private communication). The results are shown in Figure E.1. I consistently bin the spectra into logarithmically spaced ℓ -bins in the range $\ell \in [10, 3000]$. The two methods show remarkable qualitative agreement. There is mild disagreement at intermediate ℓ -mode, but the agreement at low- ℓ and in the highest ℓ -bin are striking, particularly given the number of approximations which go into each method. In the future forward models can

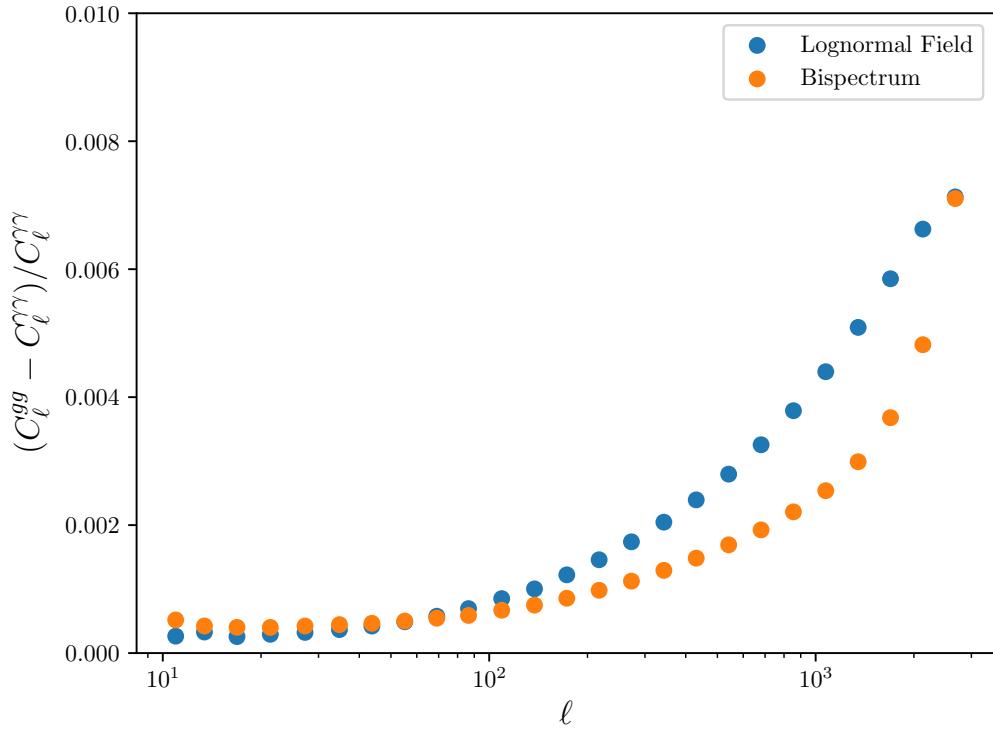


Figure E.1: The reduced shear correction using the bispectrum perturbative approach (see [32, 117]) and using the forward model in the lognormal field approximation as presented in Chapter 4. The data points are plotted at the geometric mean of the ℓ -bin boundaries. There is mild disagreement at intermediate ℓ -mode. This is to be expected given the approximations that go into the bispectrum fitting formula and the lognormal field approximation. Nevertheless the agreement at low- ℓ and in the highest ℓ -bin are striking. In the future, forward models could be used to get around using a fitting formula for the matter bispectrum, which could be biased due to the indeterminate impact of baryons. The bispectrum data was provided by Anurag Deshpande (private communication).

be used to seamlessly account for the reduced shear and other corrections without ever needing to derive complicated corrections.

Bibliography

- [1] TMC Abbott, FB Abdalla, A Alarcon, J Aleksić, S Allam, S Allen, A Amara, J Annis, J Asorey, S Avila, et al. Dark energy survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing. *Physical Review D*, 98(4):043526, 2018.
- [2] N Aghanim, Y Akrami, M Ashdown, J Aumont, C Baccigalupi, M Balldini, AJ Banday, RB Barreiro, N Bartolo, S Basak, et al. Planck 2018 results. vi. cosmological parameters. *arXiv preprint arXiv:1807.06209*, 2018.
- [3] Y Akrami, F Arroja, M Ashdown, J Aumont, C Baccigalupi, M Ballardini, AJ Banday, RB Barreiro, N Bartolo, S Basak, et al. Planck 2018 results. i. overview and the cosmological legacy of planck. *arXiv preprint arXiv:1807.06205*, 2018.
- [4] Shadab Alam, Metin Ata, Stephen Bailey, Florian Beutler, Dmitry Bizyaev, Jonathan A Blazek, Adam S Bolton, Joel R Brownstein, Angela Burden, Chia-Hsun Chuang, et al. The clustering of galaxies in the completed sdss-iii baryon oscillation spectroscopic survey: cosmological analysis of the dr12 galaxy sample. *Monthly Notices of the Royal Astronomical Society*, 470(3):2617–2652, 2017.
- [5] Andreas Albrecht, Gary Bernstein, Robert Cahn, Wendy L Freedman, Jacqueline Hewitt, Wayne Hu, John Huth, Marc Kamionkowski, Edward W Kolb, Lloyd Knox, et al. Report of the dark energy task force. *arXiv preprint astro-ph/0609591*, 2006.

- [6] David Alonso, Javier Sanchez, and Anže Slosar. A unified pseudo-c framework. *Monthly Notices of the Royal Astronomical Society*, 2018.
- [7] Justin Alsing, Tom Charnock, Stephen Feeney, and Benjamin Wandelt. Fast likelihood-free cosmology with neural density estimators and active learning. *arXiv preprint arXiv:1903:00007*, 2019.
- [8] Justin Alsing, Alan Heavens, and Andrew H Jaffe. Cosmological parameters, shear maps and power spectra from cfhtlens using bayesian hierarchical inference. *Monthly Notices of the Royal Astronomical Society*, 466(3):3272–3292, 2016.
- [9] Justin Alsing, Alan Heavens, Andrew H Jaffe, Alina Kiessling, Benjamin Wandelt, and Till Hoffmann. Hierarchical cosmic shear power spectrum inference. *Monthly Notices of the Royal Astronomical Society*, 455(4):4452–4466, 2015.
- [10] Justin Alsing, Donnacha Kirk, Alan Heavens, and Andrew H Jaffe. Weak lensing with sizes, magnitudes and shapes. *Monthly Notices of the Royal Astronomical Society*, 452(2):1202–1216, 2015.
- [11] Justin Alsing and Benjamin Wandelt. Nuisance hardened data compression for fast likelihood-free inference.
- [12] Justin Alsing and Benjamin Wandelt. Generalized massive optimal data compression. *Monthly Notices of the Royal Astronomical Society: Letters*, 476(1):L60–L64, 2018.
- [13] Justin Alsing, Benjamin Wandelt, and Stephen Feeney. Massive optimal data compression and density estimation for scalable, likelihood-free inference in cosmology. *Monthly Notices of the Royal Astronomical Society*, 477(3):2874–2885, 2018.
- [14] Theodore Wilbur Anderson, Theodore Wilbur Anderson, Theodore Wilbur Anderson, Theodore Wilbur Anderson, and Etats-Unis Mathématicien. *An*

introduction to multivariate statistical analysis, volume 2. Wiley New York, 1958.

- [15] J Anthony and LSST Collaboration. Large synoptic survey telescope: Overview. In *Proc. of SPIE Vol*, volume 4836, page 11.
- [16] David J Bacon, Alexandre R Refregier, and Richard S Ellis. Detection of weak gravitational lensing by large-scale structure. *Monthly Notices of the Royal Astronomical Society*, 318(2):625–640, 2000.
- [17] Matthias Bartelmann and Peter Schneider. Weak gravitational lensing. *Physics Reports*, 340(4):291–472, 2001.
- [18] Daniel Baumann. Cosmology: Part iii mathematical tripos.
- [19] Jose Luis Bernal, Licia Verde, and Adam G Riess. The trouble with h0. *Journal of Cosmology and Astroparticle Physics*, 2016(10):019, 2016.
- [20] Francis Bernardeau, Takahiro Nishimichi, and Atsushi Taruya. Cosmic shear full nulling: sorting out dynamics, geometry and systematics. *Monthly Notices of the Royal Astronomical Society*, 445(2):1526–1537, 2014.
- [21] Diego Blas, Julien Lesgourgues, and Thomas Tram. The cosmic linear anisotropy solving system (class). part ii: approximation schemes. *Journal of Cosmology and Astroparticle Physics*, 2011(07):034, 2011.
- [22] Fernando V Bonassi, Lingchong You, and Mike West. Bayesian learning from marginal data in bionetwork models. *Statistical applications in genetics and molecular biology*, 10(1), 2011.
- [23] Michael L Brown, Andy N Taylor, David J Bacon, Meghan E Gray, Simon Dye, Klaus Meisenheimer, and Christian Wolf. The shear power spectrum from the combo-17 survey. *Monthly Notices of the Royal Astronomical Society*, 341(1):100–118, 2003.

- [24] ML Brown, PG Castro, and AN Taylor. Cmb temperature and polarization pseudo-c estimators and covariances. *arXiv preprint astro-ph/0410394*, 2008.
- [25] ML Brown, AN Taylor, NC Hambly, and S Dye. Measurement of intrinsic alignments in galaxy ellipticities. *Monthly Notices of the Royal Astronomical Society*, 333(3):501–509, 2002.
- [26] Sean M Carroll. *Spacetime and geometry. An introduction to general relativity*, volume 1. 2004.
- [27] PG Castro, AF Heavens, and TD Kitching. Weak lensing analysis in three dimensions. *Physical Review D*, 72(2):023516, 2005.
- [28] Paolo Catelan, Marc Kamionkowski, and Roger D Blandford. Intrinsic and extrinsic galaxy alignment. *Monthly Notices of the Royal Astronomical Society*, 320(1):L7–L13, 2001.
- [29] Chihway Chang, Michael Wang, Scott Dodelson, Tim Eifler, Catherine Heymans, Michael Jarvis, M James Jee, Shahab Joudaki, Elisabeth Krause, Alex Malz, et al. A unified analysis of four cosmic shear surveys. *arXiv preprint arXiv:1808.07335*, 2018.
- [30] Asantha Cooray and Wayne Hu. Second-order corrections to weak lensing by large-scale structure. *The Astrophysical Journal*, 574(1):19, 2002.
- [31] David Copeland, Andy Taylor, and Alex Hall. The impact of baryons on the sensitivity of dark energy measurements. *Monthly Notices of the Royal Astronomical Society*, 480(2):2247–2265, 2018.
- [32] Scott Dodelson, Charles Shapiro, and Martin White. Reduced shear power spectrum. *Physical Review D*, 73(2):023009, 2006.
- [33] Tim Eifler, Elisabeth Krause, Scott Dodelson, Andrew R Zentner, Andrew P Hearn, and Nickolay Y Gnedin. Accounting for baryonic effects in cosmic shear tomography: Determining a minimal set of nuisance parameters using

pca. *Monthly Notices of the Royal Astronomical Society*, 454(3):2451–2471, 2015.

[34] Tim Eifler, Elisabeth Krause, Peter Schneider, and Klaus Honscheid. Combining probes of large-scale structure with cosmolike. *Monthly Notices of the Royal Astronomical Society*, 440(2):1379–1390, 2014.

[35] J Elvin-Poole, M Crocce, AJ Ross, T Giannantonio, E Rozo, ES Rykoff, S Avila, N Banik, J Blazek, SL Bridle, et al. Dark energy survey year 1 results: Galaxy clustering for combined probes. *Physical Review D*, 98(4):042006, 2018.

[36] T Erben, H Hildebrandt, L Miller, L van Waerbeke, C Heymans, H Hoekstra, TD Kitching, Y Mellier, J Benjamin, C Blake, et al. Cfhtlens: the canada–france–hawaii telescope lensing survey–imaging data and catalogue products. *Monthly Notices of the Royal Astronomical Society*, 433(3):2545–2563, 2013.

[37] Yanan Fan, David J Nott, and Scott A Sisson. Approximate bayesian computation via regression density estimation. *Stat*, 2(1):34–48, 2013.

[38] Pedro G. Ferreira. Cosmological Tests of Gravity. *arXiv e-prints*, page arXiv:1902.10503, Feb 2019.

[39] Simon Foreman, Matthew R Becker, and Risa H Wechsler. Cosmic shear as a probe of galaxy formation physics. *Monthly Notices of the Royal Astronomical Society*, 463(3):3326–3338, 2016.

[40] Daniel Foreman-Mackey, David W Hogg, Dustin Lang, and Jonathan Goodman. emcee: the mcmc hammer. *Publications of the Astronomical Society of the Pacific*, 125(925):306, 2013.

[41] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. , 622:759–771, April 2005.

- [42] Krzysztof M Gorski, Benjamin D Wandelt, Frode K Hansen, Eric Hivon, and Anthony J Banday. The healpix primer. *arXiv preprint astro-ph/9905275*, 1999.
- [43] Arushi Gupta, José Manuel Zorrilla Matilla, Daniel Hsu, and Zoltán Haiman. Non-gaussian information from weak lensing data via deep learning. *Physical Review D*, 97(10):103515, 2018.
- [44] Takashi Hamana. Lensing magnification effects on the cosmic shear statistics. *Monthly Notices of the Royal Astronomical Society*, 326(1):326–332, 2001.
- [45] Joachim Harnois-Déraps, Sanaz Vafaei, and Ludovic Van Waerbeke. Gravitational lensing simulations—i. covariance matrices and halo catalogues. *Monthly Notices of the Royal Astronomical Society*, 426(2):1262–1279, 2012.
- [46] J Hartlap, Patrick Simon, and P Schneider. Why your model parameter confidences might be too optimistic. unbiased estimation of the inverse covariance matrix. *Astronomy & Astrophysics*, 464(1):399–404, 2007.
- [47] Alan Heavens. Cosmology with gravitational lensing. In Gorini Moschella Matarrese, Colpi, editor, *Dark Matter and Dark Energy*, chapter 3, pages 177–216. Springer, 2011.
- [48] Alan Heavens. 3d weak lensing. *Monthly Notices of the Royal Astronomical Society*, 343(4):1327–1334, 2003.
- [49] Alan F Heavens, Elena Sellentin, Damien de Mijolla, and Alvise Vianello. Massive data compression for parameter-dependent covariance matrices. *Monthly Notices of the Royal Astronomical Society*, 472(4):4244–4250, 2017.
- [50] Katrin Heitmann, Earl Lawrence, Juliana Kwan, Salman Habib, and David Higdon. The coyote universe extended: precision emulation of the matter power spectrum. *The Astrophysical Journal*, 780(1):111, 2013.

[51] Catherine Heymans, Emma Grocott, Alan Heavens, Martin Kilbinger, Thomas D Kitching, Fergus Simpson, Jonathan Benjamin, Thomas Erben, Hendrik Hildebrandt, Henk Hoekstra, et al. Cfhtlens tomographic weak lensing cosmological parameter constraints: Mitigating the impact of intrinsic galaxy alignments. *Monthly Notices of the Royal Astronomical Society*, 432(3):2433–2453, 2013.

[52] Catherine Heymans, Ludovic Van Waerbeke, David Bacon, Joel Berge, Gary Bernstein, Emmanuel Bertin, Sarah Bridle, Michael L Brown, Douglas Clowe, Håkon Dahle, et al. The shear testing programme—i. weak lensing analysis of simulated ground-based observations. *Monthly Notices of the Royal Astronomical Society*, 368(3):1323–1339, 2006.

[53] Chiaki Hikage, Masamune Oguri, Takashi Hamana, Surhud More, Rachel Mandelbaum, Masahiro Takada, Fabian Köhlinger, Hironao Miyatake, Atsushi J Nishizawa, Hiroaki Aihara, et al. Cosmology from cosmic shear power spectra with subaru hyper suprime-cam first-year data. *arXiv preprint arXiv:1809.09148*, 2018.

[54] Stefan Hilbert, Jan Hartlap, and Peter Schneider. Cosmic shear covariance: the log-normal approximation. *Astronomy & Astrophysics*, 536:A85, 2011.

[55] H Hildebrandt, T Erben, K Kuijken, L van Waerbeke, C Heymans, J Coupon, J Benjamin, C Bennett, L Fu, H Hoekstra, et al. Cfhtlens: improving the quality of photometric redshifts with precision photometry. *Monthly Notices of the Royal Astronomical Society*, 421(3):2355–2367, 2012.

[56] H Hildebrandt, M Viola, C Heymans, S Joudaki, K Kuijken, C Blake, T Erben, B Joachimi, D Klaes, L Miller, et al. Kids-450: Cosmological parameter constraints from tomographic weak gravitational lensing. *Monthly Notices of the Royal Astronomical Society*, 465(2):1454–1498, 2016.

[57] H Hildebrandt, M Viola, C Heymans, S Joudaki, K Kuijken, C Blake, T Erben, B Joachimi, D Klaes, L Miller, et al. Kids-450: Cosmological parameter constraints from tomographic weak gravitational lensing. *Monthly Notices of the Royal Astronomical Society*, 465(2):1454–1498, 2016.

ben, B Joachimi, D Klaes, L Miller, et al. Kids-450: Cosmological parameter constraints from tomographic weak gravitational lensing. *Monthly Notices of the Royal Astronomical Society*, 2017.

[58] Christopher M Hirata and Uroš Seljak. Intrinsic alignment-lensing interference as a contaminant of cosmic shear. *Physical Review D*, 70(6):063526, 2004.

[59] Wayne Hu. Power spectrum tomography with weak lensing. *The Astrophysical Journal Letters*, 522(1):L21, 1999.

[60] Hung-Jin Huang, Tim Eifler, Rachel Mandelbaum, and Scott Dodelson. Modeling baryonic physics in future weak lensing surveys. *arXiv preprint arXiv:1809.01146*, 2018.

[61] Eric M Huff and Genevieve J Graves. Magnificent magnification: exploiting the other half of the lensing signal. *The Astrophysical Journal Letters*, 780(2):L16, 2013.

[62] Olivier Ilbert, S Arnouts, HJ McCracken, M Bolzonella, E Bertin, O Le Fevre, Y Mellier, G Zamorani, R Pello, A Iovino, et al. Accurate photometric redshifts for the cfht legacy survey calibrated using the vimos vlt deep survey. *Astronomy & Astrophysics*, 457(3):841–856, 2006.

[63] Albert Izard, Pablo Fosalba, and Martin Crocce. Ice-cola: fast simulations for weak lensing observables. *Monthly Notices of the Royal Astronomical Society*, 473(3):3051–3061, 2017.

[64] Donghui Jeong, Jens Chluba, Liang Dai, Marc Kamionkowski, and Xin Wang. Effect of aberration on partial-sky measurements of the cosmic microwave background temperature power spectrum. *Physical Review D*, 89(2):023003, 2014.

[65] Benjamin Joachimi, Marcello Cacciato, Thomas D Kitching, Adrienne Leonard, Rachel Mandelbaum, Björn Malte Schäfer, Cristóbal Sifón, Henk

Hoekstra, Alina Kiessling, Donnacha Kirk, et al. Galaxy alignments: An overview. *Space Science Reviews*, 193(1-4):1–65, 2015.

[66] Shahab Joudaki, Chris Blake, Andrew Johnson, Alexandra Amon, Marika Asgari, Ami Choi, Thomas Erben, Karl Glazebrook, Joachim Harnois-Déraps, Catherine Heymans, et al. Kids-450+ 2dflens: Cosmological parameter constraints from weak gravitational lensing tomography and overlapping redshift-space galaxy clustering. *Monthly Notices of the Royal Astronomical Society*, 474(4):4894–4924, 2017.

[67] Nick Kaiser. Weak gravitational lensing of distant galaxies. *The Astrophysical Journal*, 388:272–286, 1992.

[68] Alina Kiessling, Marcello Cacciato, Benjamin Joachimi, Donnacha Kirk, Thomas D Kitching, Adrienne Leonard, Rachel Mandelbaum, Björn Malte Schäfer, Cristóbal Sifón, Michael L Brown, et al. Galaxy alignments: Theory, modelling & simulations. *Space Science Reviews*, 193(1-4):67–136, 2015.

[69] Alina Kiessling, AF Heavens, AN Taylor, and B Joachimi. Sunglass: a new weak-lensing simulation pipeline. *Monthly Notices of the Royal Astronomical Society*, 414(3):2235–2245, 2011.

[70] Martin Kilbinger. Cosmology with cosmic shear observations: a review. *Reports on Progress in Physics*, 78(8):086901, 2015.

[71] Martin Kilbinger, Karim Benabed, Julien Guy, Pierre Astier, I Tereno, L Fu, D Wraith, J Coupon, Y Mellier, C Balland, et al. Dark-energy constraints and correlations with systematics from cfhtls weak lensing, snls supernovae ia and wmap5. *Astronomy & Astrophysics*, 497(3):677–688, 2009.

[72] TD Kitching and AF Heavens. Unequal-time correlators for cosmology. *Physical Review D*, 95(6):063522, 2017.

[73] TD Kitching, AF Heavens, J Alsing, T Erben, C Heymans, H Hildebrandt, H Hoekstra, A Jaffe, A Kiessling, Y Mellier, et al. 3d cosmic shear: cosmology from cfhtlens. *Monthly Notices of the Royal Astronomical Society*, 442(2):1326–1349, 2014.

[74] TD Kitching, P Paykari, H Hoekstra, and M Cropper. Propagating residual biases in cosmic shear power spectra. *arXiv preprint arXiv:1904.07173*, 2019.

[75] Thomas Kitching, Adam Amara, Mandeep Gill, Stefan Harmeling, Catherine Heymans, Richard Massey, Barnaby Rowe, Tim Schrabback, Lisa Voigt, Sreekumar Balan, et al. Gravitational lensing accuracy testing 2010 (great10) challenge handbook. *The Annals of Applied Statistics*, pages 2231–2263, 2011.

[76] Thomas D Kitching, Justin Alsing, Alan F Heavens, Raul Jimenez, Jason D McEwen, and Licia Verde. The limits of cosmic shear. *Monthly Notices of the Royal Astronomical Society*, 469(3):2737–2749, 2017.

[77] Thomas D Kitching, Peter L Taylor, Peter Capak, Daniel Masters, and Henk Hoekstra. Rainbow cosmic shear: Optimization of tomographic bins. *Physical Review D*, 99(6):063536, 2019.

[78] Darsh Kodwani, David Alonso, and Pedro Ferreira. The effect on cosmological parameter estimation of a parameter dependent covariance matrix. *arXiv preprint arXiv:1811.11584*, 2018.

[79] Elisabeth Krause and Tim Eifler. Cosmolike–cosmological likelihood analyses for photometric galaxy surveys. *Monthly Notices of the Royal Astronomical Society*, 470(2):2100–2112, 2017.

[80] Rene Laureijs, J Amiaux, S Arduini, J-L Auguera, J Brinchmann, R Cole, M Cropper, C Dabin, L Duvet, A Ealet, et al. Euclid definition study report. *arXiv preprint arXiv:1110.3193*, 2011.

[81] René J Laureijs, Ludovic Duvet, Isabel Escudero Sanz, Philippe Gondoin, David H Lumb, Tim Oosterbroek, and Gonzalo Saavedra Criado. The euclid mission. In *Proc. SPIE*, volume 7731, page 77311H, 2010.

[82] Earl Lawrence, Katrin Heitmann, Martin White, David Higdon, Christian Wagner, Salman Habib, and Brian Williams. The coyote universe. iii. simulation suite and precision emulator for the nonlinear matter power spectrum. *The Astrophysical Journal*, 713(2):1322, 2010.

[83] Julien Lesgourgues and Sergio Pastor. Massive neutrinos and cosmology. *Physics Reports*, 429(6):307–379, 2006.

[84] Antony Lewis and Anthony Challinor. Camb: Code for anisotropies in the microwave background. *Astrophysics Source Code Library*, 2011.

[85] Chien-Hao Lin, Joachim Harnois-Déraps, Tim Eifler, Taylor Pospisil, Rachel Mandelbaum, Ann B. Lee, and Sukhdeep Singh. Non-Gaussianity in the Weak Lensing Correlation Function Likelihood - Implications for Cosmological Parameter Biases. *arXiv e-prints*, page arXiv:1905.03779, May 2019.

[86] Eric V Linder. Cosmic growth history and expansion history. *Physical Review D*, 72(4):043529, 2005.

[87] Jia Liu, Zoltán Haiman, Lam Hui, Jan M Kratochvil, and Morgan May. Impact of magnification and size bias on the weak lensing power spectrum and peak statistics. *Physical Review D*, 89(2):023515, 2014.

[88] Marilena LoVerde and Niayesh Afshordi. Extended limber approximation. *Physical Review D*, 78(12):123506, 2008.

[89] Jan-Matthis Lueckmann, Giacomo Bassetto, Theofanis Karaletsos, and Jakob H Macke. Likelihood-free inference with emulator networks. *arXiv preprint arXiv:1805.09294*, 2018.

[90] Zhaoming Ma, Wayne Hu, and Dragan Huterer. Effects of photometric redshift uncertainties on weak-lensing tomography. *The Astrophysical Journal*, 636(1):21, 2006.

[91] Niall MacCrann, Joe Zuntz, Sarah Bridle, Bhuvnesh Jain, and Matthew R Becker. Cosmic discordance: Are planck cmb and cfhtlens weak lensing measurements out of tune? *Monthly Notices of the Royal Astronomical Society*, 451(3):2877–2888, 2015.

[92] A Spurio Mancini, PL Taylor, R Reischke, T Kitching, V Pettorino, BM Schäfer, B Zieser, and Ph M Merkel. 3d cosmic shear: Numerical challenges, 3d lensing random fields generation, and minkowski functionals for cosmological inference. *Physical Review D*, 98(10):103507, 2018.

[93] Alessio Spurio Mancini, Fabian Köhlinger, Benjamin Joachimi, Valeria Pettorino, Björn Malte Schäfer, Robert Reischke, Samuel Brieden, Maria Archidiacono, and Julien Lesgourgues. Kids+ gama: Constraints on horndeski gravity from combined large-scale structure probes. *arXiv preprint arXiv:1901.03686*, 2019.

[94] Richard Massey, Catherine Heymans, Joel Bergé, Gary Bernstein, Sarah Bridle, Douglas Clowe, Håkon Dahle, Richard Ellis, Thomas Erben, Marco Hetterscheidt, et al. The shear testing programme 2: Factors affecting high-precision weak-lensing analyses. *Monthly Notices of the Royal Astronomical Society*, 376(1):13–38, 2007.

[95] Richard Massey, Henk Hoekstra, Thomas Kitching, Jason Rhodes, Mark Cropper, Jérôme Amiaux, David Harvey, Yannick Mellier, Massimo Meneghetti, Lance Miller, et al. Origins of weak lensing systematics, and requirements on future instrumentation (or knowledge of instrumentation). *Monthly Notices of the Royal Astronomical Society*, 429(1):661–678, 2012.

[96] Daniel Masters, Peter Capak, Daniel Stern, Olivier Ilbert, Mara Salvato,

Samuel Schmidt, Giuseppe Longo, Jason Rhodes, Stephane Paltani, Bahram Mobasher, et al. Mapping the galaxy color–redshift relation: optimal photometric redshift calibration strategies for cosmology surveys. *The Astrophysical Journal*, 813(1):53, 2015.

[97] AJ Mead, JA Peacock, Catherine Heymans, Shahab Joudaki, and AF Heavens. An accurate halo model for fitting non-linear cosmological power spectra and baryonic feedback models. *Monthly Notices of the Royal Astronomical Society*, 454(2):1958–1975, 2015.

[98] Massimo Meneghetti. Introduction to gravitational lensing, lecture notes.

[99] L Miller, C Heymans, TD Kitching, L Van Waerbeke, T Erben, H Hildebrandt, H Hoekstra, Y Mellier, BTP Rowe, J Coupon, et al. Bayesian galaxy shape measurement for weak lensing surveys–iii. application to the canada–france–hawaii telescope lensing survey. *Monthly Notices of the Royal Astronomical Society*, 429(4):2858–2880, 2013.

[100] Lance Miller, TD Kitching, C Heymans, AF Heavens, and L Van Waerbeke. Bayesian galaxy shape measurement for weak lensing surveys–i. methodology and a fast-fitting algorithm. *Monthly Notices of the Royal Astronomical Society*, 382(1):315–324, 2007.

[101] Jeffrey A Newman, Alexandra Abate, Filipe B Abdalla, Sahar Allam, Steven W Allen, Réza Ansari, Stephen Bailey, Wayne A Barkhouse, Timothy C Beers, Michael R Blanton, et al. Spectroscopic needs for imaging dark energy experiments. *Astroparticle Physics*, 63:81–100, 2015.

[102] G Papamakarios, I Murray, and T Pavlakou. Advances in neural information processing systems, 2016.

[103] George Papamakarios, David C Sterratt, and Iain Murray. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. *arXiv preprint arXiv:1805.07226*, 2018.

[104] Saul Perlmutter, G Aldering, G Goldhaber, RA Knop, P Nugent, PG Castro, S Deustua, S Fabbro, A Goobar, DE Groom, et al. Measurements of ω and λ from 42 high-redshift supernovae. *The Astrophysical Journal*, 517(2):565, 1999.

[105] J Prat, C Sánchez, Y Fang, D Gruen, J Elvin-Poole, N Kokron, LF Secco, B Jain, R Miquel, N MacCrann, et al. Dark energy survey year 1 results: Galaxy-galaxy lensing. *Physical Review D*, 98(4):042005, 2018.

[106] Jason Rhodes, Alexie Leauthaud, Chris Stoughton, Richard Massey, Kyle Dawson, William Kolbe, and Natalie Roe. The effects of charge transfer inefficiency (cti) on galaxy shape measurements. *Publications of the Astronomical Society of the Pacific*, 122(890):439, 2010.

[107] Jason Rhodes, Alexandre Refregier, and Edward J Groth. Detection of cosmic shear with the hubble space telescope survey strip. *The Astrophysical Journal Letters*, 552(2):L85, 2001.

[108] Adam G Riess, Stefano Casertano, Wenlong Yuan, Lucas Macri, Beatrice Bucciarelli, Mario G Lattanzi, John W MacKenty, J Bradley Bowers, WeiKang Zheng, Alexei V Filippenko, et al. Milky way cepheid standards for measuring cosmic distances and application to gaia dr2: implications for the hubble constant. *The Astrophysical Journal*, 861(2):126, 2018.

[109] Adam G Riess, Alexei V Filippenko, Peter Challis, Alejandro Clocchiatti, Alan Diercks, Peter M Garnavich, Ron L Gilliland, Craig J Hogan, Saurabh Jha, Robert P Kirshner, et al. Observational evidence from supernovae for an accelerating universe and a cosmological constant. *The Astronomical Journal*, 116(3):1009, 1998.

[110] Adam G Riess, Lucas M Macri, Samantha L Hoffmann, Dan Scolnic, Stefano Casertano, Alexei V Filippenko, Brad E Tucker, Mark J Reid, David O

Jones, Jeffrey M Silverman, et al. A 2.4% determination of the local value of the hubble constant. *The Astrophysical Journal*, 826(1):56, 2016.

[111] S Samuroff, J Blazek, MA Troxel, N MacCrann, E Krause, CD Leonard, J Prat, D Gruen, S Dodelson, TF Eifler, et al. Dark energy survey year 1 results: Constraints on intrinsic alignments and their colour dependence from galaxy clustering and weak lensing. *arXiv preprint arXiv:1811.06989*, 2018.

[112] Masanori Sato and Takahiro Nishimichi. Impact of the non-gaussian covariance of the weak lensing power spectrum and bispectrum on cosmological parameter estimation. *Physical Review D*, 87(12):123538, 2013.

[113] Elena Sellentin and Alan F Heavens. On the insufficiency of arbitrarily precise covariance matrices: non-gaussian weak-lensing likelihoods. *Monthly Notices of the Royal Astronomical Society*, 473(2):2355–2363, 2017.

[114] Elena Sellentin, Catherine Heymans, and Joachim Harnois-Déraps. The skewed weak lensing likelihood: why biases arise, despite data and theory being sound. *Monthly Notices of the Royal Astronomical Society*, 477(4):4879–4895, 2018.

[115] Elisabetta Semboloni, Henk Hoekstra, Joop Schaye, Marcel P van Daalen, and Ian G McCarthy. Quantifying the effect of baryon physics on weak lensing tomography. *Monthly Notices of the Royal Astronomical Society*, 417(3):2020–2035, 2011.

[116] Elisabetta Semboloni, Tim Schrabbach, Ludovic van Waerbeke, Sanaz Vafaei, Jan Hartlap, and Stefan Hilbert. Weak lensing from space: first cosmological constraints from three-point shear statistics. *Monthly Notices of the Royal Astronomical Society*, 410(1):143–160, 2010.

[117] Charles Shapiro. Biased dark energy constraints from neglecting reduced shear in weak-lensing surveys. *The Astrophysical Journal*, 696(1):775, 2009.

- [118] D Spergel, N Gehrels, C Baltay, D Bennett, J Breckinridge, M Donahue, A Dressler, BS Gaudi, T Greene, O Guyon, et al. Wide-field infrared survey telescope-astrophysics focused telescope assets wfirst-afta 2015 report. *arXiv preprint arXiv:1503.03757*, 2015.
- [119] A Spurio Mancini, R Reischke, V Pettorino, BM Schäfer, and M Zumalacárregui. Testing (modified) gravity with 3d and tomographic cosmic shear. *Monthly Notices of the Royal Astronomical Society*, 480(3):3725–3738, 2018.
- [120] Ryuichi Takahashi, Masanori Sato, Takahiro Nishimichi, Atsushi Taruya, and Masamune Oguri. Revising the halofit model for the nonlinear matter power spectrum. *The Astrophysical Journal*, 761(2):152, 2012.
- [121] Ryuichi Takahashi, Masanori Sato, Takahiro Nishimichi, Atsushi Taruya, and Masamune Oguri. Revising the halofit model for the nonlinear matter power spectrum. *The Astrophysical Journal*, 761(2):152, 2012.
- [122] Peter L Taylor, Francis Bernardeau, and Thomas D Kitching. k-cut cosmic shear: Tunable power spectrum sensitivity to test gravity. *Physical Review D*, 98(8):083514, 2018.
- [123] Peter L. Taylor, Thomas D. Kitching, Justin Alsing, Benjamin D. Wandelt, Stephen M. Feeney, and Jason D. McEwen. Cosmic Shear: Inference from Forward Models. *arXiv e-prints*, page arXiv:1904.05364, Apr 2019.
- [124] Peter L. Taylor, Thomas D. Kitching, and Jason D. McEwen. Preparing for the cosmic shear data flood: Optimal data extraction and simulation requirements for stage iv dark energy experiments. *Phys. Rev. D*, 98:043532, Aug 2018.
- [125] Peter L. Taylor, Thomas D. Kitching, and Jason D. McEwen. Nonparametric cosmology with cosmic shear. *Phys. Rev. D*, 99:043532, Feb 2019.

[126] Peter L Taylor, Thomas D Kitching, Jason D McEwen, and Thomas Tram. Testing the cosmic shear spatially-flat universe approximation with generalized lensing and shear spectra. *Physical Review D*, 98(2):023522, 2018.

[127] Max Tegmark, Andy N Taylor, and Alan F Heavens. Karhunen-loeve eigenvalue problems in cosmology: How should we tackle large data sets? *The Astrophysical Journal*, 480(1):22, 1997.

[128] MA Troxel and Mustapha Ishak. The intrinsic alignment of galaxies and its impact on weak gravitational lensing in an era of precision cosmology. *Physics Reports*, 558:1–59, 2015.

[129] MA Troxel, Elizabeth Krause, C Chang, TF Eifler, O Friedrich, D Gruen, N MacCrann, A Chen, C Davis, J DeRose, et al. Survey geometry and the internal consistency of recent cosmic shear measurements. *Monthly Notices of the Royal Astronomical Society*, 479(4):4998–5004, 2018.

[130] MA Troxel, N MacCrann, J Zuntz, TF Eifler, E Krause, S Dodelson, D Gruen, J Blazek, O Friedrich, S Samuroff, et al. Dark energy survey year 1 results: Cosmological constraints from cosmic shear. *Physical Review D*, 98(4):043528, 2018.

[131] Ludovic Van Waerbeke, Jonathan Benjamin, Thomas Erben, Catherine Heymans, Hendrik Hildebrandt, Henk Hoekstra, Thomas D Kitching, Yannick Mellier, Lance Miller, Jean Coupon, et al. Cfhtlens: mapping the large-scale structure with gravitational lensing. *Monthly Notices of the Royal Astronomical Society*, 433(4):3373–3388, 2013.

[132] Roland D Vavrek, René J Laureijs, Jose Lorenzo Alvarez, Jérôme Amiaux, Yannick Mellier, Ruyman Azzollini, Guillermo Buenadicha, Gonzalo Saavedra Criado, Mark Cropper, Christophe Dabin, et al. Mission-level performance verification approach for the euclid space mission. In *Modeling, Sys-*

tems Engineering, and Project Management for Astronomy VI}, volume 9911, page 991105. International Society for Optics and Photonics, 2016.

- [133] Benjamin D Wandelt, Eric Hivon, and Krzysztof M Gorski. Cosmic microwave background anisotropy power spectrum statistics for high precision cosmology. *Physical Review D*, 64(8):083003, 2001.
- [134] David M Wittman, J Anthony Tyson, David Kirkman, Ian Dell’Antonio, and Gary Bernstein. Detection of weak gravitational lensing distortions of distant galaxies by cosmic dark matter at large scales. *Nature*, 405(6783):143, 2000.
- [135] Henrique S Xavier, Filipe B Abdalla, and Benjamin Joachimi. Improving lognormal models for cosmological fields. *Monthly Notices of the Royal Astronomical Society*, 459(4):3693–3710, 2016.
- [136] Joe Zuntz, Marc Paterno, Elise Jennings, Douglas Rudd, Alessandro Manzotti, Scott Dodelson, Sarah Bridle, Saba Sehrish, and James Kowalkowski. Cosmoss: modular cosmological parameter estimation. *Astronomy and Computing*, 12:45–59, 2015.