

Systematic Biases in Weak Lensing Cosmology with the Dark Energy Survey

Simon Samuroff

The University of Manchester
School of Physics and Astronomy

2017

Thesis presented for the degree of Doctor of Philosophy
in the Faculty of Science and Engineering

Contents

1	The Basic Principles of Cosmology	10
1.1	The Standard Model of Cosmology	11
1.1.1	Framework and Parameterisation	11
1.1.2	Components	15
1.1.3	The Cosmic Web & Our Place in the Universe	20
1.1.4	Chronology	21
1.1.5	The Matter Power Spectrum	27
1.2	Observational Probes of the Universe	30
1.2.1	Bayesian Parameter Inference	31
1.2.2	Cosmic Shear	34
1.2.3	Strong Lensing	38
1.2.4	Galaxy Clustering & Galaxy-Galaxy Lensing	40
1.2.5	Supernovae & Supernova Lensing	43
1.2.6	The Cosmic Microwave Background & CMB Lensing	46
1.3	Stage IV and the Future of Cosmology	48
1.3.1	Photometric Surveys	48
1.3.2	Spectroscopic Surveys	50
1.3.3	Non-Conventional Lensing Methods	51
2	The Basic Principles of Weak Gravitational Lensing	53
2.1	Cosmic Shear in Theory: Formalism	53
2.1.1	Point Deflectors, Critical Density & Convergence	53
2.1.2	The Lens Equation	56
2.1.3	Cosmology with Cosmic Shear	58
2.2	Cosmic Shear in Practice: Observables	64
2.2.1	The Point Spread Function	65
2.2.2	Galaxy Shapes	65
2.2.3	Estimators of Cosmic Shear	69
2.2.4	Photometric Redshifts	72
2.2.5	Covariance Matrices	73
2.3	Biases in Cosmic Shear	74
2.3.1	Shear Measurement Bias	75

2.3.2	Intrinsic Alignments	75
2.3.3	Baryons & Modelling Uncertainties	76
2.3.4	Expectation Bias	77
3	Measuring Cosmic Shear in Year 1 of the Dark Energy Survey	78
3.1	Introduction	78
3.2	Data	81
3.2.1	Observing Period and Conditions	81
3.2.2	Fields in DES Y1	81
3.2.3	The GOLD Catalogue & Star/Galaxy Separation	82
3.2.4	Blinding	82
3.2.5	PSF Estimation	83
3.3	The IM3SHAPE Catalogue	84
3.3.1	Overview	84
3.3.2	Galaxy Weights	86
3.4	The HOOPOE Image Simulations	86
3.4.1	Simulating DES Y1: The Image Pipeline	88
3.4.2	Galaxy Sample	89
3.4.3	Comparing Simulations & Data	91
3.4.4	Bias Calibration & Diagnostics	93
3.4.5	Robustness to Tomographic Binning	95
3.4.6	Validating the HOOPOE Simulations	100
3.4.7	Sensitivity to Observable Distributions	102
3.5	Using the Shear Catalogues	102
3.5.1	Mean Shear	102
3.5.2	Catalogue Flags	104
3.5.3	Applying the IM3SHAPE Calibration	106
3.5.4	Number Density	108
3.5.5	Choosing a Prior on Residual Calibration Bias	109
3.6	Summary and Discussion	113
4	The Impact of Image Plane Neighbours on Shear Cosmology	115
4.1	Introduction	115
4.2	The Shear Measurement Problem	118
4.2.1	Shape Measurement with IM3SHAPE	118
4.2.2	Shear Measurement Bias	119
4.3	Toy Model Predictions	120
4.3.1	Single-Galaxy Effects	121
4.3.2	Ensemble Biases	122
4.4	A More Realistic Simulation	127
4.4.1	Parent Data	127
4.4.2	Input Galaxy Selection	127

4.4.3	Neighbour-Free Resimulations	129
4.5	Quantifying Neighbour Bias with HOOPOE	130
4.5.1	Single-Galaxy Effects	131
4.5.2	Neighbour Ensemble Biases	132
4.5.3	Untangling the Knot of Neighbour Bias	133
4.5.4	Isolating the Impact of Subdetection Galaxies	139
4.5.5	Suppressing Neighbour Bias	142
4.6	Cosmological Implications	143
4.6.1	Mean Multiplicative Bias	143
4.6.2	Scale Dependence	146
4.7	Conclusions	149
5	Constraining Redshift Systematics Using Lensing & Galaxy Clustering	152
5.1	Introduction	152
5.2	Methodology & Assumptions	156
5.3	Simultaneous Constraints on Cosmology and Photometric Redshift Bias	157
5.4	Conclusion	160
6	Colour Space Divisions and Intrinsic Alignments in the Dark Energy Survey	163
6.1	Theory & Background	163
6.1.1	Observational Constraints on Intrinsic Alignments	163
6.1.2	Theory Predictions	165
6.2	Data	169
6.2.1	Year One of The Dark Energy Survey	169
6.2.2	Definitions of Red & Blue Galaxies	169
6.2.3	Shape Measurements	173
6.2.4	Photometric Redshifts	174
6.2.5	Two-Point Correlations	176
6.2.6	Covariance Matrix	176
6.3	Results	178
6.3.1	Simultaneous Constraints	180
6.3.2	Redshift and Magnitude Dependence	181
6.3.3	Colour Leakage & Redshift Error	184
6.3.4	Intrinsic Alignment Models	187
6.4	Discussion & Comparison with Previous Results	188
6.5	Conclusion	190
7	Conclusions	192
7.1	Remarks About the Current Work	192
7.2	Prospects for the Future	196
7.2.1	Shear Estimation	196
7.2.2	The Intrinsic Alignment of Galaxies	197

7.2.3 Self-Calibration & Multi-Probe Cosmology	198
A Simulating Galaxies Below the Survey Detection Limit	213
B Quality Checking COSMOS Images Using the DES Scientific Community	219
C Derivation of a Two-Point Modifier for Scale Dependent Bias	222

This thesis (excluding preface material and bibliography) contains a total of 56930 words.

Abstract

This thesis sets out a practical guide to applying shear measurements as a cosmological tool. We first present one of two science-ready galaxy shape catalogues from Year 1 of the Dark Energy Survey (DES Y1), which covers 1500 square degrees in four bands *griz*, with a median redshift of 0.59. We describe the shape measurement process implemented by the DES Y1 IM3SHAPE catalogue, which contains 21.9 million high-quality *r*-band bulge/disc fits. In Chapter 3 a new suite of image simulations, referred to as HOOPOE, are presented. The HOOPOE dataset is tailored to DES Y1 and includes realistic blending, spatial masks and variation in the point spread function. We derive shear corrections, which we show are robust to changes in calibration method, galaxy binning and variance within the simulated dataset. Sources of systematic uncertainty in the simulation-based shear calibration are discussed, leading to a final estimate of the 1σ uncertainties in the residual multiplicative bias after calibration of 0.025.

Chapter 4 describes an extension of the analysis on the HOOPOE simulations into a detailed investigation of the impact of galaxy neighbours on shape measurement and shear cosmology. Four mechanisms by which neighbours can have a non-negligible influence on shear measurement are identified. These effects, if ignored, would contribute a net multiplicative bias of $m \sim 0.03 - 0.09$ in DES Y1, though the precise impact will depend on both the measurement code and the selection cuts applied. We use the cosmological inference pipeline of DES Y1 to explore the cosmological implications of neighbour bias and show that omitting blending from the calibration simulation for DES Y1 would bias the inferred clustering amplitude $S_8 \equiv \sigma_8(\Omega_m/0.3)^{0.5}$ by 1.5σ towards low values. Finally, we use the HOOPOE simulations to test the effect of neighbour-induced spatial correlations in the multiplicative bias. We find the cosmological impact to be subdominant to statistical error at the current level of precision.

Another major uncertainty in shear cosmology is the accuracy of our ensemble redshift distributions. Chapter 5 presents a numerical investigation into the combined constraining power of cosmic shear, galaxy clustering and their cross-correlation in DES Y1, and the potential for internal calibration of redshift errors. Introducing a moderate uniform bias into the redshift distributions used to model the weak lensing (WL) galaxies is shown to produce a $> 2\sigma$ bias in S_8 . We demonstrate that this cosmological bias can be eliminated by marginalising over redshift error nuisance parameters. Strikingly, the cosmological constraint of the combined dataset is largely undiminished by the loss of prior information on the WL distributions. We demonstrate that this implicit self-calibration is the result of complementary degeneracy directions in the combined data.

In Chapter 6 we present the preliminary results of an investigation into galaxy intrinsic alignments. Using the DES Y1 data, we show a clear dependence in alignment amplitude on galaxy type, in agreement with previous results. We subject these findings to a series of initial robustness tests. We conclude with a short overview of the work presented, and discuss prospects for the future.

Declaration

I declare that no portion of the work referred to in the thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning;

This thesis is the result of my own work, unless otherwise stated. Chapter 3 was carried out in collaboration with Joseph Zuntz, and Erin Sheldon, and coordinated by Michael Troxel and Daniel Gruen as part of the Year 1 Dark Energy Survey Weak Lensing analysis group. Chapter 4 was also carried out within DES Y1, with input from our internal reviewers, Tomasz Kacprzak, Gary Bernstein and Eric Huff. The simulations and shape catalogues used in both of these chapters were built using computer resources at the National Center for Science and Energy Research (NERSC), which is maintained by the United States Department of Energy. The analysis for Chapter 5 was performed in collaboration with Michael Troxel and Sarah Bridle.

Parts of this thesis have been published in refereed journals:

- Samuroff S., Troxel M. A., Bridle S. L., Zuntz J., MacCrann N., Krause E., Eifler T., Kirk D., 2017, MNRAS, 465, L20
- Jarvis M. et al, 2016 MNRAS 460 p.2245-228, arXiv:1507.05603

The analyses presented in Chapters 3 and 4 have passed the Dark Energy Survey internal review process, and have been submitted for publication as:

- Samuroff S., Bridle S. L., Zuntz J., Gruen D. et al, 2017, MNRAS, arXiv:1708.01534
- Zuntz J., Sheldon E., Samuroff S., Troxel M.A, Gruen D. et al, 2017, MNRAS, arXiv:1708.01533

Intellectual Property and Copyright

I furthermore make the following declarations, with regards to ownership of copyright and intellectual property rights:

1. The author of this thesis (including any appendices and/or schedules to this thesis) owns certain copyright or related rights in it (the “Copyright”) and s/he has given The University of Manchester certain rights to use such Copyright, including for administrative purposes.
2. Copies of this thesis, either in full or in extracts and whether in hard or electronic copy, may be made only in accordance with the Copyright, Designs and Patents Act 1988 (as amended) and regulations issued under it or, where appropriate, in accordance with licensing agreements which the University has from time to time. This page must form part of any such copies made.
3. The ownership of certain Copyright, patents, designs, trademarks and other intellectual property (the “Intellectual Property”) and any reproductions of copyright works in the thesis, for example graphs and tables (“Reproductions”), which may be described in this thesis, may not be owned by the author and may be owned by third parties. Such Intellectual Property and Reproductions cannot and must not be made available for use without the prior written permission of the owner(s) of the relevant Intellectual Property and/or Reproductions.
4. Further information on the conditions under which disclosure, publication and commercialisation of this thesis, the Copyright and any Intellectual Property and/or Reproductions described in it may take place is available in the University IP Policy (see <http://documents.manchester.ac.uk/DocuInfo.aspx?DocID=24420>), in any relevant Thesis restriction declarations deposited in the University Library, The University Library’s regulations (see <http://www.library.manchester.ac.uk/about/regulations/>) and in The University’s policy on Presentation of Theses

Acknowledgements

First and foremost my thanks go to Sarah Bridle, for everything. Sarah, you introduced me to a world of which I knew almost nothing a handful of years ago. Thoughtful and unfailingly patient in the face of ignorant questions, you taught me how to think about science. You could always see the bigger picture. I will miss our long talks.

I'd also like to thank Joe Zuntz and Michael Troxel, friends and mentors both. You made Manchester an exciting place to work in the early days of my PhD. Joe, you have a better instinct for shape measurement than I will ever have. Thanks for your understanding through the years. And to Troxel, for your support and insight.

Many others are deserving of thanks for guidance and interesting thoughts over the three short years of my PhD. To Donnacha Kirk, my predecessor, thanks for your immense patience. I wish you the best of luck in the real world. I am grateful to my various colleagues within the Dark Energy Survey, particularly to Tim Eifler, Niall MacCrann, Elisabeth Krause, Benjamin Joachimi, Jonathan Blazek and the rest of the Multi Probe Cosmology group. Standing on the shoulders of giants doesn't quite cover it.

Amongst other weak lensers Gary Bernstein, Tomek Kacprzak and Eric Huff are deserving of thanks for agreeing to review our shear measurement paper, and for their various insights on the subject in general.

It would be remiss if I failed to mention our group co-ordinators Daniel Gruen and Troxel (again), for making the Weak Lensing Working Group such a pleasant place to work. I would also like to thank my examiners, Richard Battye and David Bacon, for an interesting discussion and for finding my various typos.

My colleagues at Manchester also more than warrant a few words here. Thanks to Jack, my fellow cosmologist-in-training and office mate, and to Nicolas for being there to offer a friendly eye over everything. Thank you also to Richard Rollins - for helping to make Y1 IM3SHAPE happen.

And through everything, I am indebted to Eileen, Stephen and Alice Samuroff. For sane conversation about anything from politics to books to physics to bees. For listening to me witter. And for so much more besides. It means more that I can say.

“It is a capital mistake to theorise before one has data.
Insensibly one begins to twist facts to suit theories,
instead of theories to suit facts.”
A Scandal in Bohemia, AC Doyle 1891

“The idea is to go
from numbers,
to information,
to understanding.”
Hans Rosling

Chapter 1

The Basic Principles of Cosmology

It has often been said that we have entered an era of precision cosmology. It is a statement which has preceded so many discussions on topics as disparate as the cosmic microwave background, strong lensing modelling and spectroscopy. All too often it is shorthand for an ideal – sunlit uplands always just over the horizon, which we might reach were we only to find a solution to this systematic or that. If usage of the term has acquired a vaguely Damoclean edge, there is good reason for this. Cosmology is now dominated by large collaborations involving sometimes hundreds of people across the world, and sophisticated software, in a way it has never been before. We must trust that each part of the web is functioning (and improving) as it should do. No longer the reserve of the eccentric polymath aristocrat or the clergyman, cosmology has become a community, as reliant on trust as any other.

But the phrase implies something more than this: “precision” implies a measurement. Whereas other fields of science could, from the mid to early Twentieth Century, begin to claim startling accuracy in the predictive power of their theories, those who think about the cosmos professionally have typically been limited to doing just that. Throughout most of the Twentieth Century, measurements on the scale needed to test our theories about the Universe were unimaginable. Beyond our tiny solar system, the Universe was a handful of blurry dots on the sky.

Nonetheless, the rate of progress cosmology has made over just a handful of decades has been remarkable. At the end of the Nineteenth Century, before Einstein and Hubble had made their impact, we had no reason to suspect the Universe was expanding and no real concept of its scale. A hundred years later, our understanding had been reshaped and it was simply a matter of time before we had built the observational machinery left to fill in the blanks. Now, just seventeen years after that, the cosmological community has reduced the Universe into a few relatively well known

numerical parameters. We have a predictive standard model for the entirety of the cosmos, and we can measure its parameters to a few percent.

This model is built on the theoretical foundations laid by Friedmann, Einstein and many others. Beyond the framework, there are free parameters that cannot be deduced *a priori*, and observations have played a significant role in shaping the model. The confirmation of gravitational lensing by Dyson et al. (1920), the discovery of a cosmic microwave background (Penzias, 1965) and the more recent discovery that the late-time expansion is accelerating (Perlmutter et al., 1998) are just a few notable examples of this. In addition, our model relies on a set of basic assumptions, primarily that gravity is well described by Einstein's General Relativity (GR) on all (macroscopic) scales, and that the Universe is isotropic and homogenous on the largest scales (the Cosmological Principle).

In this introductory chapter we aim to furnish the reader with some basic concepts that underpin the research presented in the rest of this thesis. We first provide an outline of the formalism for the standard model, or the Λ CDM Universe. The following paragraphs set out how an isotropic universe is expected to evolve under this theory, and chart a brief history of our own Universe, as it is understood today. This established, we will then consider the observational tools that have been developed to measure it. Since weak lensing is the focus of the bulk of this thesis, we will focus particularly on this area and its most recent developments. We aim to provide a qualitative overview in this chapter. For detail about the practicalities and formalism of lensing we refer the reader to Chapter 2.

1.1 The Standard Model of Cosmology

The following paragraphs introduces the analytical framework and components of Λ CDM. We draw on a number of more comprehensive reviews, to which we refer the reader for more detail (Peacock, 1999; Bartelmann & Schneider, 2001; Dodelson, 2003; Ryden, 2003; Weinberg et al., 2013; Kilbinger, 2015).

1.1.1 Framework and Parameterisation

The Friedmann Lemaître Robertson Walker Universe

The Standard Model is built around GR, which is in essence a formalism for describing a system acting under its own gravity. It makes no specific assumptions about where the mass originates, or how much of it there is. Let's start by defining a spatial coordinate system that is stationary relative

to the large-scale motion of the Universe; we will refer to this as the *comoving* coordinate system. Ideal measurements within this frame are said to be taken by fundamental observers. Assuming the Universe is homogeneous and isotropic throughout its evolution, one can map position coordinates in a fixed local observation frame \mathbf{d} onto the those in the comoving system χ via a time-dependent scale factor $a(t)$:

$$\mathbf{d}(\chi, t) = a(t)\chi. \quad (1.1)$$

The instantaneous radial distance between any two points, as defined in the local frame, d is referred to as the physical distance. The normalisation of the scale factor is arbitrary, and typically it is defined such that the value today is $a_0 = 1$. In GR space-time separations between events are described by a metric. Specifically in the homogenous isotropic case one has the Friedmann Lemaître Robertson Walker (FLRW) metric:

$$ds^2 = c^2 dt^2 - a^2(t) [d\chi^2 + S_K^2(\chi) (d\theta^2 + d\phi^2 \sin^2(\theta))]. \quad (1.2)$$

Here, $(d\chi, d\theta, d\phi)$ define a spatial increment in comoving spherical polar coordinates. The exact functional form of $S_K(\chi)$, which is a distance measure oriented transverse to the line of sight, is dependent on the background curvature of the Universe. One can parameterise the geometry of the 4D surface with a curvature parameter K , which can have three values, corresponding to flat (Euclidean) space ($K = 0$), positive curvature ($K > 0$) and negative curvature ($K < 0$)¹. Specifically it has the form:

$$S_K(\chi) = \begin{cases} \frac{1}{\sqrt{K}} \sinh(\sqrt{K}\chi) & K > 0 \\ \chi & K = 0 \\ \frac{1}{\sqrt{|K|}} \sin(\sqrt{|K|}\chi) & K < 0. \end{cases} \quad (1.3)$$

If one enforces the FLRW metric, Einstein's field equations simplify to two independent expressions (Friedman, 1922), which describe the time evolution of the scale factor in terms of the comoving energy density $\rho(t)c^2$ and pressure $p(t)$ of an ideal fluid:

¹The distinction between the two type of non-zero curvature can be a little bit difficult to visualise. To build a simple picture, imagine a point with a pair of orthogonal lines crossing through it, defining a surface. In the simplest case they're straight - the surface is flat and can be extended infinitely by making the lines longer. Now picture both lines being bent in the same direction at the edges (upwards or downwards), forming a curve with a peak at the crossing point; The area of the surface can no longer be expanded infinitely, since at some point the lines will meet; that is called *positive curvature*. Next imagine one line is curled upwards at both ends and the other downwards. The resulting surface can no longer be closed but instead has a saddle-shaped form. This is called *negative curvature*.

$$\left(\frac{\dot{a}(t)}{a(t)}\right)^2 = \frac{8\pi G}{3c^2}\rho(t)c^2 - \frac{Kc^2}{a^2(t)}, \quad (1.4)$$

and

$$\left(\frac{\ddot{a}(t)}{a(t)}\right) = -\frac{4\pi G}{3c^2} [\rho(t)c^2 + 3p(t)]. \quad (1.5)$$

The dot is used here in the normal sense, indicating a time derivative, and G and c are Newton's constant and the speed of light in a vacuum, which are fundamental constants. This fluid is simply a mathematical description of the large scale mass components of the Universe, and the formalism set out here are an entirely general consequence of gravity. GR does not distinguish between different types of fluid, provided that the pressure and density are equivalent. For this we must build a phenomenological model to fit our Universe, as discussed in Section 1.1.2.

Though it is not required by the above equations, we will assume the Universe is expanding (and so a was less than unity in the past, and will be larger in the future). This is indeed what current evidence implies, a point we will return to later in this chapter.

Cosmological Redshift and Distance Measures

A photon travelling large distances across the Universe will be subject to cosmological redshift effects. The basic principle here is simple and relatively intuitive. First, because it sits within an expanding universe any source of photons will be receding from any observer at the point of emission, and so the emitted light will be redshifted. Second, space is continuously expanding while the photon is in transit; for any arbitrary increment along its path, its original position will be receding from its destination. For a photon following a null-geodesic, equation 1.2 reduces to a simple expression for the comoving distance travelled between emission and observation times, t_{em} and t_{obs} :

$$\chi = c \int_{t_{\text{em}}}^{t_{\text{obs}}} \frac{1}{a(t)} dt. \quad (1.6)$$

In the extreme $t_{\text{em}} \rightarrow 0$, this gives an upper limit on the scale of causal connection. The corresponding physical distance is the *horizon distance*, d_{hor} . For $t_{\text{em}} > 0$, assuming the scale factor is constant between emission of one wavecrest and the next and setting $t_{\text{obs}} = t_0$, one can derive the expression,

$$a(z) = \frac{1}{(1+z)}. \quad (1.7)$$

This relates the cosmological redshift z and the scale factor at the time of emission, a . On a practical level this allows a source's redshift to be used as an observable proxy for cosmic distance (or time). Assuming negligible peculiar velocity, it directly translates into recession speed due to expansion, $v = cz$.

Using equation (1.1), the physical distance of an object is related to its recession speed:

$$v = \left(\frac{\dot{a}(z)}{a(z)} \right) d(z) = H(z) \times d(z). \quad (1.8)$$

The *Hubble parameter*, $H(z) \equiv \dot{a}(z)/a(z)$, describes the expansion rate at z . In the local Universe the Hubble parameter is approximately constant, $H \simeq H_0$, leading to a rough proportionality between z and d . For convenience it is very common to use the dimensionless form $h = H_0/100 \text{ kms}^{-1}\text{Mpc}^{-1}$.

Finally, we introduce the concept of a *standard measure* in cosmology, which can be used to constrain slightly different distance quantities. The underlying concept is this: if one knows *a priori* some intrinsic property of an object, and understands how observations of that property should change per unit of depth between us and it, then it should be possible to infer its distance. Primarily these objects can be classified either as standard candles or as standard rulers, which respectively yield *luminosity distances* and *angular diameter distances*.

To understand this imagine that we observe two point sources with an angular separation on the sky θ , which are known to have a physical separation R . We can make a distance estimate using simple trigonometry as $D_A(\chi) = R/\theta$. We can also see by examination of equation 1.3 that the length of an arc between the two points at equal distance is given by $R = a(z)S_K(\chi)\theta$. Basic algebra then gives us:

$$D_A(\chi) = (1+z)^{-1}S_K(\chi). \quad (1.9)$$

The argument with regard to standard candles follows similar lines. Let's say we measure a small patch of sky and find it has a flux f , and we know that the source of photons has an inherent luminosity L . Defining an arbitrary shell around the source, clearly the measured flux will diminish as the size of the shell expands, which gives us a new distance estimate $D_L(\chi) = (L/4\pi f)^{\frac{1}{2}}$. Flux measurements in an expanding universe are affected both by cosmological redshift, and the fact that at each instant when a photon reaches an imaginary shell at r , the background expansion of

the Universe makes the shell's surface area incrementally larger. Both effects contribute a factor of $(1+z)$ dilution in flux, $f = L(4\pi S_K^2(\chi))^{-1}(1+z)^{-2}$. One can surmise that the luminosity distance should be:

$$D_L(\chi) = (1+z)S_K(\chi) \quad (1.10)$$

with $S_K(\chi)$ defined in the previous section. These quantities, cosmological redshift, angular diameter distance and luminosity distance are the three basic distance metrics in the field of observational cosmology. With these simple observables in hand, we now turn to the question of what the Universe actually contains.

1.1.2 Components

The basic Λ CDM Universe contains three density components: radiation, matter and dark energy. The matter constituent ρ_m can be further decomposed into what we loosely call baryons (the immediately observable stuff of people, planets and stars; clearly this includes electrons, but for our purposes they contribute no mass, and can be safely ignored) and cold dark matter (the other mass we infer must exist, but which neither emits nor absorbs photons). On cosmological scales ($\gtrsim 100$ Mpc) each of these is well approximated as a continuous ideal fluid, which is governed by a linear equation of state relating the pressure p and energy density ρc^2 . The generalised equation of state has the form,

$$p_X = w_X \rho_X c^2, \quad (1.11)$$

with the subscript X denoting the component and a dimensionless parameter of state w_X . Applying equation 1.4 and 1.5, the redshift evolution can be derived as,

$$\rho_X(z) = \rho_{X,0} \times (1+z)^{3(1+w_X)}, \quad (1.12)$$

where $\rho_{X,0} = \rho_X(z=0)$ is the present-day mean density of component X . From this expression one obtains the scaling relations in Table 1.1. As the dark energy equation of state parameter is the only one of these that is not known from first principles, it is common to drop the subscript and refer to it simply as w . We will do the same from here onwards. Finally, equation 1.5 can be cast in a convenient form:

$$H(z) = H_0 \left[\Omega_m(1+z)^{-3} + \Omega_r(1+z)^{-4} + \Omega_K(1+z)^{-2} + f_{de}(z)\Omega_{de} \right]^{\frac{1}{2}}, \quad (1.13)$$

Component	Ω_X	w_X	Evolution
Radiation	$\sim 10^{-4}$	1/3	$(1+z)^4$
Baryonic Matter	0.049	0	$(1+z)^3$
Cold Dark Matter	0.264	0	$(1+z)^3$
Dark Energy	0.683	~ -1	?

Table 1.1: Properties of the Λ CDM energy density components in the present-day Universe. The density parameter is defined in units of the current critical density. The equation of state parameters w are as defined above. Values credited to Planck Collaboration (2015a). The redshift dependence in each component are shown in the right-hand column

where $H(t)$ is defined by equation 1.8, and f_{de} is a dimensionless function describing the evolution of dark energy. For Λ CDM one has $f_{\text{de}}(z) = f_{\text{de},0} = 1$. In a general dark energy model that allows for w to be time dependent, $f_{\text{de}}(z) = \exp \left[3 \int_0^z (1 + w(z')) / (1 + z') dz' \right]$ (Bartelmann, 2010). The density parameter Ω_X is the present-day density of component X in units of the critical density,

$$\rho_c(t) = 3H^2(t)/8\pi G. \quad (1.14)$$

The latter is defined as the threshold density, below which a universe with no dark energy expands forever. Many observable probes are sensitive to physical rather than comoving densities, which mean it is often convenient to cast the density parameter in a form that depends on the expansion rate, $\Omega_X h^2$.

The expansion is initially driven by radiation, with the Universe gradually evolving through matter and finally dark energy dominated eras. Setting $t = t_0$, equation 1.13 yields $1 - \Omega_K = \Omega_m + \Omega_r + \Omega_{\text{de}} \equiv \Omega_0$. Ω_0 is an historically important, if now less favoured cosmological parameter. $\Omega_0 > 1$ implies a spatially finite closed universe, which will recollapse. $\Omega_0 < 1$ indicates an infinite open universe, which will expand forever. $\Omega_0 = 1$ implies an infinite, perpetually expanding universe with flat geometry.

Radiation

The term radiation includes all highly relativistic particles, into which category photons by definition must fall. The photonic energy-density is dominated by the Cosmic Microwave Background (CMB), which behaves to good approximation as a black body at $T_{\text{CMB}} = 2.726$ K and is smooth in temperature to 1 part in 10^4 .

Radiation exerts non-negligible pressure, $p = \rho_r c^2 / 3$, giving an equation of state, $w_r = 1/3$.

The energy density evolves as $(1 + z)^4$ due to a combined effect of the expansion of space, which dilutes the energy density of a bundle of photons, and cosmological redshift.

Baryons

In a cosmological context, the term baryons refers collectively to all forms of mass that emits or absorbs photons. Protons and neutrons naturally dominate the gravity of this component due to the instability of other hadrons and the comparatively low mass of the electron. Even in the present day the bulk of baryonic mass resides in light nuclides formed in nucleosynthesis. Though small quantities of heavier elements do emerge on small scales, in terms of mass these are entirely insignificant.

It is also true that on smaller scales electromagnetic interactions facilitate radiative cooling and sometimes violent feedback effects. Though gravitationally subdominant it is thus possible for baryons to have observable effects on the cosmological mass distribution. At late times baryons are typically non relativistic and pressureless, giving $w_b = 0$.

The current best estimate for the mean baryon density from Black Body Nucleosynthesis (BBN; see Section 1.1.4) is $\Omega_b h^2 = 0.021 \pm 0.002$ (Kirkman et al., 2003), which is consistent with early-time predictions from the CMB (see Section 1.1.4 for a brief discussion of how this is measured).

One longstanding question is thus posed by observations of the X-ray luminous intra-cluster gas and the Inter-Galactic Medium (IGM), which seem to be able to account for at most $\sim 60\%$ of this. The discrepancy is known as the *missing baryons problem*.

A partial answer has been provided in recent years by spectroscopic studies (e.g. Gupta et al. 2012) which imply suggest a large amount of mass may reside in large-scale diffuse halos and filaments, expelled from galaxies by feedback effects. Such matter is too hot for observation at sub X-ray energies but of sufficiently low density to avoid detection of thermal emission at current detector sensitivities.

Dark Matter

As alluded to above, dark matter is a phenomenological term to describe that mass component which must be present but which cannot be accounted for by baryons. Extensive observational evidence has now made the case for a dominant non-luminous mass component very difficult to dispute. The first hint came in the form of a study of the Coma cluster (Zwicky, 1937), which suggested tension between mass estimates derived from X-ray temperature and velocity dispersion




Figure 1.1: Composite multi-wavelength image of the Bullet Cluster, often cited as a signature of dark matter. The colours indicate the plasma temperature, as inferred using X-ray observations from the Chandra satellite. The overlain contours show the mass profile of the structure, as reconstructed using the weak lensing of background galaxies in HST data. **Image credits:** Clowe et al. (2007).

measurements. Archives of galaxy rotation curve data compiled over the past century (Rubin et al., 1978) similarly imply an extended but invisible halo of mass within which the luminous matter of a galaxy resides. More recent imaging of the merging Bullet cluster, shown in Figure 1.1, in optical (i.e. lensing) and X-ray (temperature) provide yet another instance of a cluster scale object with apparently non-identical mass and light distributions (Clowe et al., 2007).

Inferring that an unseen mass constituent exists and understanding what it *is*, however, are two quite different things. Various models have been empirically ruled out, but there is currently no single favoured understanding of dark matter at the microscopic level. For example, the idea that dim baryonic matter (white dwarfs and primordial black holes) can account for the excess mass has been tested and found wanting on all spatial scales; local microlensing studies find no significant dark baryon constituent in galaxy halos (Wambsganss, 2005). On cluster scales, the amplitude of the SZ signal is inconsistent with a dominant baryonic mass contribution (Grego et al., 2001; Myers et al., 2004; Carr et al., 2016). On cosmological scales, the CMB spectrum implies $\Omega_b \ll 1$ (and simultaneously $\Omega_m \sim 0.315$). Hot dark matter models, in which neutrinos make up the dominant mass are likewise disfavoured by large scale structure, which is smoothed out by free streaming particles.

Though we are yet to find a suitable particle candidate, that does not mean that we cannot predict how dark matter should behave on large scales. The Cold Dark Matter in Λ CDM is described as pressureless fluid, defined by $w_{\text{CDM}} = 0$.

It is finally worth noting that here have been attempts to explain the observational signatures

of dark matter using alternative theories of gravity (see Milgrom 1983 or Verlinde 2016 for a more recent attempt). These theories have had some success in reproducing flat galaxy rotation curves, but are notably unable to match other strands of evidence. For at least the present time they are considered rather a fringe activity, and the majority of cosmologists now accept the GR + CDM paradigm.

Dark Energy

One startling discovery, which came at the turn of the Twentieth Century, and for which Riess et al. (1998) and Perlmutter et al. (1998) shared a 2011 Nobel prize, was that the late time Universe is not only expanding, but *accelerating*. In terms of our simple fluid model, this requires an additional (now dominant) component with negative pressure, $p = -\rho_{\text{de}}c^2$. By equation 1.11, one finds an equation of state parameter $w = -1$. Though there is no physical reason it must be, Occam's Razor leads us to include the simplest possible parameterisation, giving us a time-independent unseen energy component in Λ CDM called dark energy.

If the Universe is understood to hold a finite quantity of dark matter, baryons and radiation the density of these components is gradually diluted as it expands. If the fourth component Ω_{de} is held constant, then dark energy eventually becomes the dominant term in the equation 1.13 and gives rise to accelerated expansion at late times. Various models have sought to introduce evolving dark matter; we will not discuss such theories in detail, but merely note that there is currently no strong evidence to support the added complexity. Many modified gravity theories have also sought to explain accelerated expansion without dark energy through perturbations to GR on particular scales, though again it is not our intention here to review these in detail.

The most widely favoured parameterisation for dark energy, introduced by Chevallier & Polarski (2001) uses a series expansion of $w(a)$ with two free parameters:

$$w(a) = w_0 + w_a(1 - a). \quad (1.15)$$

The first evidence for $\Omega_{\text{de}} \neq 0$ came from the lightcurves of distance Type Ia supernova (SNIa; see Section 1.2.5). Further evidence is provided by a combination of other measurements. CMB and BAO data suggest an approximately flat universe ($\Omega_K = 0$), while lensing and the CMB independently imply $\Omega_{\text{m}} \sim 0.3$ and a radiation component $\sim 10^{-4}$. To reconcile these two observations requires the additional component, $\Omega_{\text{de}} \simeq 0.7$. Though there has been much speculation about future measurements, contemporary observations have proved consistent with $w_a = 0$ to within $\sim 1\sigma$. A number of planned and ongoing cosmological surveys, including the Dark Energy Sur-

vey and the Kilo Degree Survey, were designed with the express purpose of shedding light on this question.

1.1.3 The Cosmic Web & Our Place in the Universe

In the above paragraphs we have set out the ingredients from which the Universe is formed and the rules they should follow. Now, with the paint and canvas these provide, we have two options open to us, should we wish to visualise a particular epoch of the Universe. We could simulate a set of initial conditions, and wind our model forwards numerically until we reach the point of interest. Alternatively, we could train our telescopes on a patch of sky, and use redshift measurements to isolate a slice along the line of sight at approximately the desired epoch. Neither approach is without its limitations; they do, however, paint a remarkably consistent picture, of which we provide a brief outline below.

All current observations point to the conclusion that the large scale Universe is very different from the world around us, both in content and in structure. Indeed, the local Universe is highly structured – if it were not, neither you nor this thesis would exist. On scales of ~ 100 Mpc, however, the mass is dominated by dark matter, with only a small fraction accounted for by baryons (matter that absorbs or emits photons and is thus directly visible), mostly in the form of diffuse clouds of cold hydrogen gas. On these scales the Universe is *homogenous* and *isotropic*, but not uniform. That is, though the mass density fluctuates spatially, the statistical properties of those fluctuations appears the same along any arbitrary line of sight and from any random point within its volume. This basic idea, that there is no favoured position in the Universe, is known as the *cosmological principle*.

Neither does this statement imply constancy in time. Look back into the distant Universe, and the amplitude of the fluctuations is smaller than it is today. The large scale mass field is seen to be evolving, with dark matter and baryons accreting gradually onto the peaks in density. The result is a uniformly distributed set of concentrated dark matter cores, known as *halos*. These structures are still growing, merging and making the Universe ever more clustered. Inevitably, as matter clumps and falls inwards the regions between and around the overdensities form underdense patches, which we call *voids*. Separating these largest of structures one finds diffuse filaments of dark matter. The net result is a web-like structure of startling intricacy. We show an example of the cosmic web at four epochs, as recreated using the Millennium simulation in Figure 1.2.

As one moves down in scale, within the largest of halos one finds sub-halos, upon which the surrounding matter collapses to form local peaks. Eventually we reach a scale where the




Figure 1.2: Four snapshots of a patch in the N -body Millennium Simulations presented by Springel et al. (2005). Each represents a $15h^{-1}$ Mpc slice of the simulated universe centred on redshifts of (upper left, upper right, lower left, lower right) 18.3, 5.7, 1.4 and 0.0. The white bar marks a comoving transverse length of $125h^{-1}$ Mpc. The colour indicates the mass density of dark matter, increasing from blue-black to yellow.

infalling matter becomes hot enough for star formation to begin. It is here that galaxies form; first about the central mass kernel of the halo, forming a massive object often known as the brightest cluster galaxy (BCG), and then about sub-halos. In this picture, then, today's Universe is strongly hierarchical. Each galaxy forms within its dark matter halo, which in turn is embedded in a larger halo, host to a cluster of similar objects bound by its gravitational potential.

1.1.4 Chronology

The starting point ($t = 0, z = \infty$) of the Standard Model is the Big Bang. The Universe evolves from a singular point, and proceeds according to gravity. Though it evolves through a series of different phases, in which different energy density constituents dominate its dynamics, the basic axioms of causality and conservation of mass are at all times respected. The current best estimate

for the age of the Universe is ~ 13.8 Gyr (Planck Collaboration 2015a).


Inflation ($t < 10^{-32}$ s)

A period of early rapid expansion, an idea now known as inflation, was theorised in the latter part of the Twentieth Century in response to a number of observations that would be difficult to explain otherwise without fine-tuning the initial conditions. That is, it allows observed late-time properties to evolve from a Gaussian-random mass distribution in the early Universe without reverse engineering (see e.g. Ryden 2003; Dodelson 2003). Specifically, it provides a natural explanation for the apparent flatness of the late-time Universe, and uniformity of the CMB.

Nucleosynthesis ($3 \text{ min} < t < 20 \text{ min}$)

In the hot dense conditions immediately after inflation, a process known as Big Bang Nucleosynthesis ensues. It is this process that generates first stable light nuclides.

At this very early point in the history of the Universe free protons p constitute the only form of baryonic mass. The average pressure is sufficient for inverse β -decay, which generates an equilibrium population of neutrons:



As equations (1.16) and (1.17) involve interaction between leptons and hadrons they occur only as weak interactions. When the Universe expands to become sufficiently cool, neutrinos decouple from baryons and the neutron/proton ratio becomes fixed. Synthesis of deuterium from $p-n$ fusion then competes with the decay of free n . Initially after neutrino freeze-out the mean photon energy is sufficient to ensure immediate photo-disintegration of heavier isotopes. As the probability of fusion directly to elements beyond hydrogen is negligible, this produces a blocking effect called the deuterium bottleneck. The n_n/n_p ratio at the end of this delay and thus the relative yields of BBN elements from subsequent fusion are dependent on the total baryon density. Baryons in well-separated islands of relict primordial gas provide our best direct probe of the cosmological baryon density $\Omega_b h^2$. Current best estimates come from high redshift quasars viewed through intervening gas. Current observations indicate 75% H and 24% ^4He by mass, with the additional 1% made

up by ${}^3\text{He}$, ${}^7\text{Be}$ and ${}^7\text{Li}$, and small quantities of heavier nuclei derived from stellar fusion at later times.

The Linear Growth of Structure

The matter distribution in the very early Universe is not completely smooth; rather it contains inhomogeneities, corresponding to quantum fluctuations, which are magnified by a period of inflation. Most commonly these fluctuations are assumed to be Gaussian random. In general it is useful to describe a local comoving density relative to a spatial average:

$$\delta(\chi, a) = \frac{\rho_m(\chi, a) - \bar{\rho}_m(a)}{\bar{\rho}_m(a)}. \quad (1.18)$$

Here $\rho_m(\chi)$ is the density of matter evaluated at a specific comoving position and the bar represents averaging over all space. In reality the overdensity is a continuous field, with a value at any arbitrary position. Discussing the evolution of “an overdensity”, then, requires some explanation. It is often convenient to deconstruct δ into a set of discrete Fourier modes, each with a characteristic scale, λ_i . In the following, the singular is used in this sense, referring to a particular mode rather than a point-position evaluation of the total field.

At early times matter clumps of high density will grow steadily under their own gravity. While they are small ($\delta \ll 1$), they evolve linearly with time, a process known as *linear growth*. Under these conditions one can use the Newtonian fluid equations to derive an evolution equation (Peacock, 1999; Ryden, 2003; Dodelson, 2003),

$$\ddot{\delta} + 2H\dot{\delta} - 4\pi G\bar{\rho}_m\delta = 0. \quad (1.19)$$

The solutions to this second-order differential equation govern the time evolution of structure in the linear regime. Given the dependence on $H(t)$ and $\bar{\rho}_m$, the exact solutions at a given time are dependent on the relative densities of the different mass constituents. In physical terms, the last part corresponds to the effect of self-gravity, while the middle term codifies the opposing effect of background expansion (called *Hubble drag*). At early times the gravitational term is dominant and overdensity perturbations will grow spontaneously.

One interesting feature of equation 1.19 is that it involves no spatial derivatives or position-dependent coefficients, meaning that overdensities are stationary in the comoving frame. That is, the landscape of large scale structure is fixed early on; though peaks and troughs get deeper they do not spontaneously shift around under gravity. One can thus relate an observed overdensity at a

later time $\delta(\chi)$ to the primordial overdensity from which it evolved δ_p via a position independent growth factor $D(a)$:

$$\delta(\chi, a) = D(a)\delta_p(\chi). \quad (1.20)$$

Early in its history ($t < 0.75$ Myr), the Universe is sufficiently compact that its dynamics are dominated by radiation (photons and neutrinos) ($\bar{\rho}_r \propto a^{-4}$). In this regime baryons are coupled to photons via Thomson scattering. Since the cold dark matter interacts with the other components only through gravity, it is simply dragged along with the radiation as it evolves. Under these conditions, density perturbations on scales larger than the horizon grow² as $D(a) \propto a^2$. For overdensities on scales smaller than the horizon, however, the picture is somewhat different. Since information has time to cross the overdensity's full scale, radiation exerts its own outward pressure, which opposes gravitational collapse. This gives rise to *acoustic oscillations* in the baryon-photon fluid, and the growth of sub-horizon overdensities is suppressed. Note that the horizon also grows over time, and so the suppression scale is continuously expanding during this period. We show a toy model of an evolving overdensity with time and scale factor in Figure 1.3.

Eventually the radiation density will be diluted by expansion to a sufficient extent that (dark) matter becomes dominant. At the point of matter-radiation equality ($\bar{\rho}_r = \bar{\rho}_m$) radiation pressure is no longer a dominant factor and the constraint on sub-horizon growth is lifted. This state of the Universe is well approximated by what is known as Einstein de Sitter cosmology ($\Omega_{de} = 0, \Omega_K = 0, \Omega_m = 1$). Pressureless dark matter drives growth on all scales, evolving as $D(a) \propto a$. Unfortunately it is not always possible to come up with an analytic forms for $D(a)$ in all cosmologies. At late times particularly the simple matter dominated paradigm no longer holds and the exact evolution must be solved numerically. Until a handful of years ago fast fitting formulae based on simulations were routinely used (see e.g. Carroll et al. 1992). In more recent years advances in computing speeds and increasing requirements on accuracy, driven by larger data volumes from new surveys, have seen these become increasingly less common.

As the Universe continues to expand, the mass density $\bar{\rho}_m(a)$ is gradually diluted by the extra volume to the point where photoionisation is not instantaneous. Recombination occurs over $\sim 10^5$ yr, and the Universe eventually becomes to become transparent to photons at the *time of decoupling*. It is possible to define a shell about an observer called the *last scattering surface*, such that photons emitted at that surface during recombination are arriving at the observer's position at the

²This is possible because all points within an overdensity do not need to be in causal contact with all others in order for it to grow. In a physical frame, a volume element's expansion rate is dependent on its local energy-density. Overdense volume elements thus expand at a lower rate than those at $\bar{\rho}_m$.




Figure 1.3: **Left:** Evolution of an overdensity mode of given scale λ as the Universe expands. In the radiation dominated era peaks in density are quadratic in the scale factor until λ is contained by the expanding horizon scale. After this point radiation pressure suppresses further increase in overdensity. After matter-radiation equality ($\Omega_m a^{-3} = \Omega_r a^{-4}$; shown by the vertical dashed line) growth can again proceed, and does so linearly with a . **Right:** Evolution of the scale factor with time. The coloured bands show epochs in which (left to right) radiation, matter and dark energy are the dominant terms in equation 1.13. The hatched section on the far right indicates times beyond the present day. The line was generated numerically using equation 1.13, assuming a flat Λ CDM universe ($\Omega_m = 0.315$, $\Omega_{de} = 0.7$, $\Omega_r = 10^{-5}$).

time of observation t_0 . The photons liberated at this point persist in Universe today as the CMB. The process of decoupling removes radiation pressure as a factor in the behaviour of matter, which has the effect of freezing a set of acoustic oscillations into the large scale mass distribution. The resulting imprinted waves, which also persist today, are called baryon acoustic oscillations (BAO). This same pattern of compressions and rarefactions map onto the CMB temperature.

After decoupling, the baryonic matter is allowed to collapse into the potential wells of the dark matter. It accrues and heats, begins to trigger star formation, and eventually leads to radiative cooling, allowing galaxies and then galaxy clusters to emerge within the most massive clumps of dark matter. This early star formation generates an ultraviolet background, which heats and reionises the gas. In turn, a fraction of CMB photons are re-scattered by newly unbound electrons. Under Λ CDM, large scale structure develops gradually through an upwards hierarchical process of clustering and merging of halos³.

³The reverse process, in which the largest structures form and then fragment, is referred to as top-down structure formation. This model was popular for a while in the 1980s, and there was some political weight behind the argument. Gradually, however, it became clear that the hot dark matter models that would be needed to generate such a paradigm were disfavoured by observations.

Nonlinear Structure Formation

On the smallest scales local gravitational effects from inhomogeneities become significant. As structure grows and the Universe becomes less uniform, the spatial scale of these effects expands with time. Within this regime, the assumption that density perturbations are small ($\delta \ll 1$) is no longer reasonable and the linear approximation breaks down. Unfortunately, the complex nonlinear interactions on these scales cannot be trivially predicted. N-body simulations are commonly required in lieu of analytic predictions to keep track of growth in these regimes. The scale and resolution of a given simulation must be finite, and generating a sizeable volume of mock Universe is unavoidably expensive in computing power. Currently the largest N-body simulations are DEUS and Millennium (see Springel et al. 2005; Alimi et al. 2012; Angulo & Pontzen 2016).

A few basic approximations, however, allow qualitatively correct analytic results, and are thus still useful to consider. A common example, known as the Zel'dovich approximation (Zel'dovich, 1970) simply assumes linear growth extends down to palpably nonlinear scales. The key prediction of this approach is that a perturbation in density will collapse in turn along each of its three axes. A spatial overdensity will tend gradually into a flat 2D disc, then reduce into an extended snake-like structure before finally collapsing into a compact core. In reality, of course, collapse does not eliminate each dimension fully, and one actually finds diffuse filaments and halos. The approximate picture, however, is consistent with numerical simulations and large scale observations.

The Dark Energy Dominated Universe ($t > 9.5$ Gyr)

At late times the Friedmann equation becomes dominated by Ω_{de} . The drag term in equation (1.19) then approaches a constant while the gravitational contribution declines as a^{-3} . The Universe eventually enters a period of exponential expansion $a(t) = \exp[2H_0\Omega_{\text{de}}^{1/2}(t - t_0)]$. Under these conditions the linear approximation yields

$$\delta(\chi, a) = A(\chi) \exp \left[-2H_0\Omega_{\text{de}}^{1/2}(t - t_0) \right] + B(\chi), \quad (1.21)$$

where A and B are time-independent integration constants. During this epoch, growth is frozen out and expansion smooths each point in $\delta(\chi)$ towards an asymptotic value.

1.1.5 The Matter Power Spectrum

The isotropic matter power spectrum is defined as

$$(2\pi)^3 \delta_D(k - k') P_\delta(k) \equiv \left\langle \tilde{\delta}(k) \tilde{\delta}^*(k') \right\rangle, \quad (1.22)$$

where $\tilde{\delta}(k)$ is an overdensity Fourier mode with wavenumber $k \propto \lambda_i^{-1}$. The brackets signify averaging over many realisations of the Universe. Assuming $\delta(\chi)$ is Gaussian, the field is fully characterised by $P(k)$. The power spectrum is a fundamental statistic in cosmology and has a simple physical interpretation. The random field of mass density fluctuations at a fixed time in the history of the Universe can be thought of as a set of Fourier modes, or superposed set of standing waves of different wavelengths. Each k mode in the matter power spectrum corresponds to a sine wave with a particular angular frequency, and the power $P(k)$ corresponds to the square of the amplitude of that wave.

It is commonly assumed that the power spectrum of density perturbations immediately after inflation, $P_{\delta,p}(k)$, follows a power law with an index $n_s \simeq 1$. This linear primordial shape is called the Harrison Zel'dovich Peebles spectrum. More generally, one can write,

$$P_\delta(k, a) = A_s T(k, a) k^{n_s}, \quad (1.23)$$

where A_s is the amplitude of $P_{\delta,p}$ and $T(k, a)$ is the *transfer function*, describing later modifications to the power spectrum. Before matter-radiation equality, only Fourier modes with $\lambda_i > d_{\text{hor}}$ grow. This defines a threshold k_{sup} above which growth is suppressed. As d_{hor} expands, k_{sup} shifts downwards, and larger overdensities are gradually allowed to begin growing. This evolution is shown in Figure 1.4. By the time of equality, differential growth has produced the characteristic shape seen in Figure 1.4. The position of the peak is frozen at this time, $k_{\text{sup}}(a_{\text{eq}}) = k_0$, and subsequent growth distributes power evenly across all k modes. The matter power spectrum today has the approximate form (Bartelmann & Schneider, 2001),

$$P(k) \propto \begin{cases} k & k < k_0 \\ k^{-3} & k > k_0 \end{cases}. \quad (1.24)$$

In addition to those described above, $T(k, a)$ encodes several other effects. First, BAO imprint a sinusoidal pattern at $k \sim k_0$. The signal is weakened by subsequent growth, distorted by baryon infall and exponentially damped at high k by residual photon-baryon interactions after recombination. Second, massive neutrinos alter large-scale structure. This is entirely consistent with the

subdominance of neutrinos as the massive constituent of dark matter, since they can have measurable impact on cold dark matter clumps by free-streaming away at the point the Universe becomes too diffuse for spontaneous inverse beta decay. Finally, nonlinear growth systematically increases power at high k .

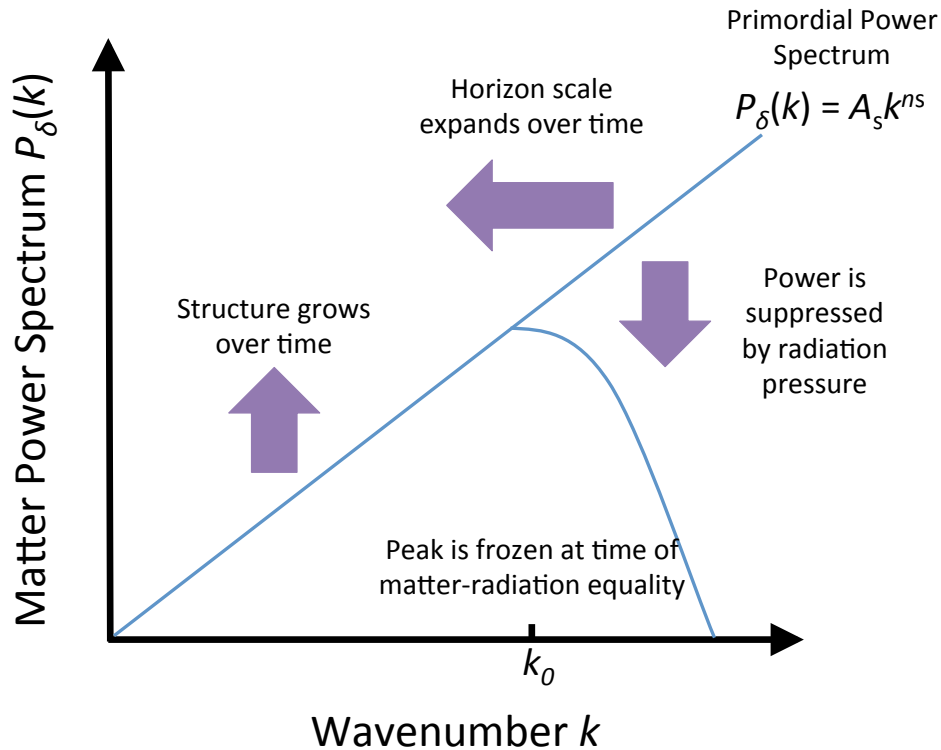


Figure 1.4: Schematic diagram of the evolution of the linear matter power spectrum. The upper arrow denotes the shift in the position of the turning point. The axes are logarithmic.

Power spectra for several cosmologies are shown in Figure 1.5. The shape of $P_\delta(k)$ depends on the primordial spectrum and the peak position indicates d_{hor} at equality, making it sensitive to $\Omega_m h$. The amplitude of the linear spectrum at a redshift of zero is set by a normalisation parameter σ_8 , which by convention is defined as the root mean square (rms) of overdensity fluctuations, averaged in spheres of comoving radius $8h^{-1}$ Mpc. This is a key physical parameter describing the large scale clustering of the late-time Universe, and can be measured directly by any dataset which is sensitive to $P_\delta(k)$. In some models an additional parameter dn_s/dk , the so-called running of the power spectrum, allows scale-dependence in the power law index. Theory does not make *ab initio* predictions of the values of $A_s, \sigma_8, n_s, dn_s/dk$ beyond ruling out unphysical regions of parameter space. Thus, they must be treated as *free parameters* in our model of the Universe, to be

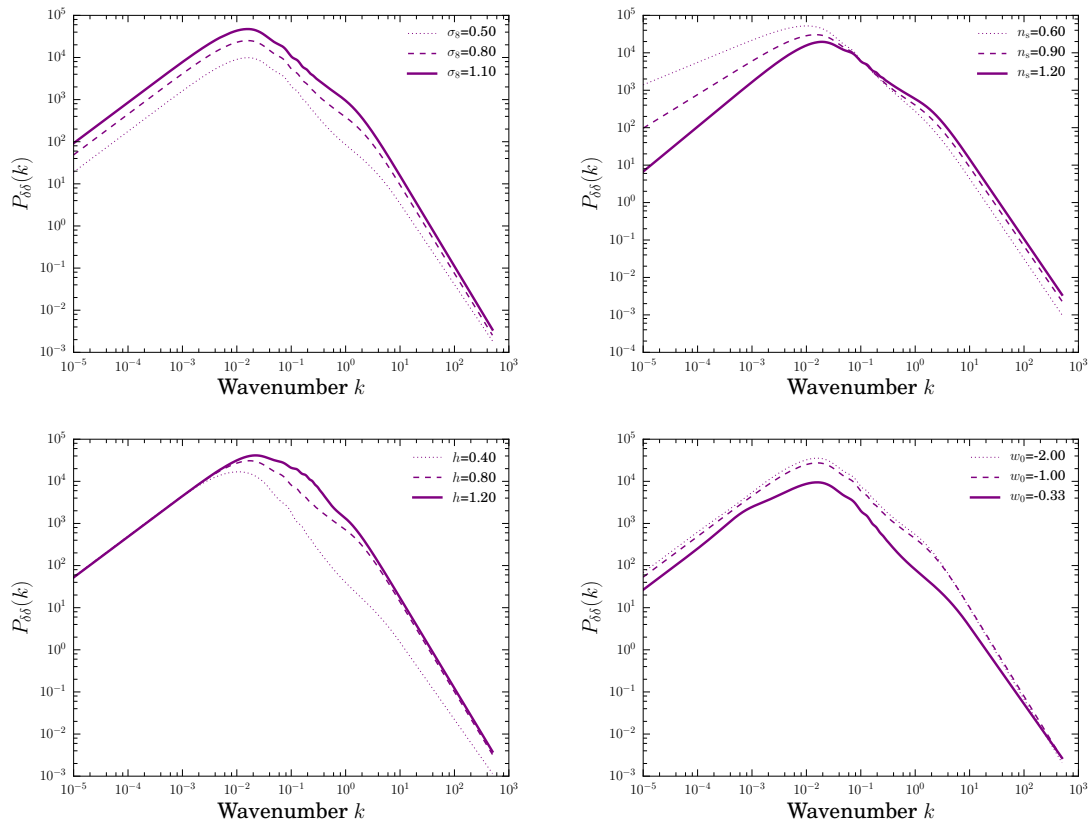


Figure 1.5: Numerical predictions of the linear matter power spectrum under various cosmologies at fixed redshift $z = 0$. In each panel we perturb one of the relevant cosmological parameters about fiducial values. We show the sensitivity to variations in (clockwise from top left) σ_8 , the primordial index n_s , the dark energy equation of state w_0 and the dimensionless Hubble parameter h . Note that σ_8 is a derived quantity, in the sense that it is defined by the matter power spectrum today, and its value is wholly determined by the other parameters. In this case we fix A_s and n_s , and rescale the late-time spectrum to the appropriate σ_8 . For the other parameters the initial amplitude A_s is fixed, and σ_8 is allowed to change.

constrained by observation. The way in which changes in the linear fluctuation amplitude σ_8 affect $P_\delta(k)$ is strongly degenerate with variations in the background average mass density Ω_m , and it is thus common to define best-constrained combinations for particular observational probes (e.g. S_8 in cosmic shear).

Much literature has been devoted to the subject of accurate calculation of $T(k, a)$. The precise solution for a given cosmology requires the Boltzmann time evolution equations to be integrated numerically. Widely used public codes such as CAMB⁴ and CLASS⁵ have been developed for this. As computer-power has become more readily available, this approach is increasingly common. A number of fast fitting-formulae have, however, been developed for tasks with limited time, computing resources or less stringent requirements on accuracy. The basis of this approach is to construct analytic expressions, which are calibrated with N-body simulations. Increasingly sophisticated versions which factor in massive neutrinos (Bond & Efstathiou, 1987; Bardeen et al., 1986; Pogosyan & Starobinsky, 1995; Eisenstein & Hu, 1999) and BAO effects (see Eisenstein & Hu (1998)) have been released. These typically claim accuracy only to the level of $\sim 10\%$.

Calculation of the matter power spectrum typically begins with the linear P_δ for a required cosmology, as a function of wavenumber and redshift. Whereas the transfer function $T(k, a)$ is a template for transforming the primordial power spectrum into the late-time linear power spectrum, in reality the growth of structure does not proceed linearly on all scales. Boltzmann codes can account rigorously for BAO and neutrinos; they do not naturally account for nonlinear growth. The most common approach to this problem is an algorithm known as HALOFIT, which was first presented in the early 2000s by Smith et al. (2003). Based on the halo model of structure formation, it assumes that all matter at t_0 is contained within dark matter halos. Late-time modifications to $P_\delta(k)$ are decomposed into 1- and 2-halo terms, representing intra- and inter-halo interactions and calibrated using simulations. Several improvements to halofit accounting for neutrino effects and refinements in simulation resolution have been published (e.g. Takahashi et al. 2012), and are still in common use today.

1.2 Observational Probes of the Universe

The following section will review the best tools available to the modern cosmologist, with a particular emphasis on gravitational lensing. We first set out a generic methodology for extracting

⁴camb.info

⁵class-code.net

estimates of the underlying model parameters from observations of cosmological data. The following will lay out the basic principles of Bayesian inference, which has become almost ubiquitous in modern cosmology. Note that this is not a discussion of philosophy; we are not concerned here with whether Bayesian probability is the *right* way to think about the Universe, but rather with the utility it provides.

With this generic description in hand, we will then consider a number of specific cosmological observables. For each, theoretical sensitivities will be discussed alongside a short review of past use and recent developments.

1.2.1 Bayesian Parameter Inference

As implied above, the standard methodology for cosmological inference is planted firmly in the framework of Bayesian probability. To most cosmologists today, this is a useful way of thinking and nothing more; there is, however, a self-consistent logic behind this framework, which is the main reason it has been so widely adopted. We will not attempt to review this subject in detail here, but rather outline the basic concepts used later in this thesis. For more comprehensive overviews and discussion see Trotta (2017), Sivia & Skilling (2006), MacKay (2002).

The basic principle behind Bayesian statistics is a particular understanding of what probability means in the physical sense. Under one common school of thought, the probability of a particular result occurring is simply the fraction of times it will show up in the limit of infinitely many repeated measurements. This makes complete sense in, for example, the context of a particle physics experiment, where the measured quantity itself is stochastic. One can simply simulate the experiment many times and measure, say, the number of decays that result in a particle A . It makes less sense when one is considering hypothetical questions like “what is the probability of snow in Stockholm on Friday?”. This way of thinking makes still less sense in the context of cosmological parameters estimates. These are not inherently random variables, but fundamental quantities. We have one Universe, and repetition of a perfectly accurate measurement will yield the same results irrespective of how many times it is repeated.

The more intuitively correct way to think about probability in this context is as a statement of an observer’s degree of confidence in a particular proposal. If we rephrase the earlier question as a definite statement: “there is a one in four chance of snow in Stockholm on Friday”, the first definition would compel us to consider many hypothetical universes, many Stockholms and many configurations of the atoms in the atmosphere. Each universe is a measurement, and the probability is the fraction of universes in which Stockholm sees snow on Friday. Under the Bayesian way of

thinking there need only be one Stockholm, and the probability is just the observer's way to codify her confidence. A probability of 0.25 is simply a statement that there is enough information to moderately disfavour the hypothesis of snow, given our current state of knowledge. This simple idea is deceptively powerful, and motivates much of the framework of modern scientific enquiry.

In completely general terms imagine we have a proposition A and some background information \mathcal{I} . The latter encodes everything we know about the Universe, the laws of physics etc. The statement of equivalence $p(A, \mathcal{I}) = p(\mathcal{I}, A)$ is then trivially correct. By basic logic we can write the statement

$$p(\mathcal{I})p(A|\mathcal{I}) = p(A)p(\mathcal{I}|A). \quad (1.25)$$

This is the most generic form of Bayes' theorem, and underpins much of what follows.

Assume next we have a generic observation (perhaps a two-point correlation function of shear, measured as a function of angular scale), which we will call \mathbf{D} , and is just a column vector of scalar numbers. In general an observation is selected according to the cosmological parameters $\boldsymbol{\theta} = (\theta_1, \theta_2 \dots \theta_m)$ it is sensitive to (under some assumed cosmological model), and also the instrument and type of data available. Given a particular set of parameter values, the model M should be fully predictive of the expected data. Without loss of generality, the aim is then to obtain a conditional probability distribution of $\boldsymbol{\theta}$, given the data and the model. Using Bayes' theorem,

$$P(\boldsymbol{\theta}|\mathbf{D}, M) = \frac{\mathcal{L}(\mathbf{D}|\boldsymbol{\theta}, M)p(\boldsymbol{\theta}|M)}{p(\mathbf{D}|M)}. \quad (1.26)$$

The *prior*, $p(\boldsymbol{\theta}|M)$, encodes restrictions on the regions of parameter space $\boldsymbol{\theta}$ is allowed to occupy, either based on previous observation or on physical arguments. The denominator, $p(\mathbf{D}|M)$, is referred to as the *evidence*. Since it is independent of position in parameter space, it does not affect the shape of the posterior distribution $P(\boldsymbol{\theta}|\mathbf{D}, M)$ and can thus be safely ignored in standard cosmological analyses. It is, however, useful for model comparison (i.e. deciding whether a particular dataset fits better with one cosmological model or another). The first term, $\mathcal{L}(\mathbf{D}|\boldsymbol{\theta}, M)$, is called the *likelihood* and represents the probability of observing data \mathbf{D} for a specific set $\boldsymbol{\theta}$ values. It is generally determined by comparing \mathbf{D} with the prediction given by M , \mathbf{y} . In most cases measurement errors are assumed to be Gaussian random⁶, giving:

⁶If there are multiple error contributions from different sources this is generally a safe assumption by virtue of Central Limit Theorem. Irrespective of the shape of the individual error distributions, if there are enough of them then their sum will tend towards a Gaussian.

$$\mathcal{L}(\mathbf{D}|\boldsymbol{\theta}, M) = \frac{1}{(2\pi)^{m/2} |\mathbf{C}|^{1/2}} \exp \left[-\frac{1}{2} (\mathbf{D} - \mathbf{y}(\boldsymbol{\theta}, M))^T \mathbf{C}^{-1} (\mathbf{D} - \mathbf{y}(\boldsymbol{\theta}, M)) \right], \quad (1.27)$$

where \mathbf{C} is the data covariance matrix. This can be written in a simpler form:

$$-2 \ln \mathcal{L} = \chi^2 = \mathbf{R}^T \mathbf{C}^{-1} \mathbf{R}, \quad (1.28)$$

where we have rewritten the residual difference between data and theory as $\mathbf{R} \equiv \mathbf{D} - \mathbf{y}$. In practice \mathbf{C} can be a very large matrix, and its computation and inversion is often extremely slow (see Section 2.2.5) Once the covariance has been obtained, parameter estimation is then simply a case of repeatedly evaluating equations (1.26) and (1.27) using different trial $\boldsymbol{\theta}$.

For calculations involving few parameters, $P(\boldsymbol{\theta}|\mathbf{D}, M)$ can be evaluated on a grid in parameter space. For high dimensionality, other techniques such as Markov Chain Monte Carlo (MCMC), Importance sampling and adaptive MultiNest or Population Monte Carlo techniques become computationally more efficient. The number of cosmological analyses based partly or entirely on Bayesian statistics has multiplied in the past decade (see, for example Lewis & Bridle 2002; Wraith et al. 2009; Kilbinger et al. 2011; Feroz et al. 2009). Several widely-used packages implementing these techniques, such as EMCEE⁷ and COSMOMC⁸, are available. Each have relative merits and should be chosen according to the dataset size and the parameter space dimensions.

In most cases a confidence region in multidimensional parameter space is more useful than a point statistic such as the maximum likelihood. In practice, however, it is inevitable that a significant fraction of $\boldsymbol{\theta}$ will be made up of relevant unknown but uninteresting parameters, describing astrophysical or measurement systematics. It is therefore very common to contract the parameter space using a process called *marginalisation*. To understand this, consider a simple cosmological experiment, which tells us about two parameters. One of these is the target of the measurement a_1 , but the second a_2 is uninteresting (say an instrumental systematic). If the two parameters are completely independent we can simply discount a_2 , but in general a_1 and a_2 might be degenerate, which means we can only map out a joint distribution $p(a_1, a_2)$. We can pick a point on the y axis and take a 1D slice, but that can only give us a conditional distribution $p(a_1|a_2)$, predicated on somehow knowing a_2 with absolute certainty⁹. If we really don't care about a_2 , and wish to

⁷dan.iel.fm/emcee

⁸cosmologist.info/cosmomc

⁹If this were the case, via external datasets or some other means, the proper way to incorporate this knowledge would be in the prior.

stop thinking about it at this point, we can simply sum over the probability of all the alternative a_2 values at any given a_1 . This effectively compresses the parameter space into the dimension of interest:

$$p(a_1) = \int_{-\infty}^{\infty} p(a_1, a_2) da_2, \quad (1.29)$$

leaving a 1D probability distribution for a_1 . Since a_2 is no longer a factor here, and we have accounted for all of the values it could possibly take, we say we have *marginalised* it.

1.2.2 Cosmic Shear

A standard and well tested prediction of GR is that a concentration of mass will distort the space-time around it, and thus produce a curious phenomenon called gravitational lensing. The most obvious manifestation is about massive galaxy clusters, where background galaxies can be elongated into crescent-shaped arcs. A subtler, but potentially more powerful, consequence of gravitational lensing is that background fluctuations in the density of dark matter induce coherent distortions to a photon's path. This effect is known as cosmic shear, and it was first detected by four groups at around the same time close to two decades ago (Bacon et al., 2000; Van Waerbeke et al., 2000; Kaiser, 2000; Wittman et al., 2000).




Figure 1.6: Graphic illustration of the concept of cosmic shear. Photons emitted from galaxies embedded in the cosmic web of dark matter (shown in red) are continuously deflected as they pass through the surrounding large scale structure. **Image credits:** Stephane Coulombi (IAP).

Cosmic shear is potentially one of the most powerful probes in modern cosmology. The spatial correlations in galaxy shape due to lensing are a direct imprint of the large scale mass distribution of the Universe. Thus they allow one to make inferences about the mass density and level of structure in the low-redshift Universe or to map out the spatial distribution of dark matter on the

sky. Unlike comparable probes it is theoretically very clean, responding directly to the power of dark matter without complications due to our limited ability to model baryonic physics. The angle through which a bundle of photons is deflected is determined by two factors: the geometry of the lens-source system and power of the lens. The latter is very intuitive; a lens becomes more powerful as the total mass enclosed by it (baryonic and dark) increases. As with its classical analogue, the lens has most impact when positioned approximately mid-way between the observer and the source. The standard approach in lensing cosmology, then, is to divide a line of sight into a series of infinitely thin slices of mass, each defined by a surface density.

Though gravitational lensing both *magnifies* and *shears* (see the next chapter for more rigorous definitions), it is the latter effect which is most useful for cosmology. In large part this is simply because it is the more accessible to the observer. If the orientation of the (unlensed) intrinsic galaxy shapes is entirely random, then over a sufficiently large enough population they should average to zero. The cancellation will not be exact in practice, even without systematics, and the numerical residual due to use of a finite number of galaxies is called *shape noise*. In principle, however, this means that given a large enough sample with coherent gravitational shear, one can isolate the cosmological signal by averaging galaxy ellipticities. This is not the case with magnification, for which no convenient cancellation occurs. The basic observable unit in any lensing based cosmology is, then, *galaxy ellipticities*.

Cosmological weak lensing is sensitive to the Hubble parameter, but unfortunately projecting shear measurements into the sky plane mixes physical scales and reduces sensitivity to this parameter. Standard ruler measurements such as the CMB and BAO are more typically considered as the primary geometric probes. The most powerfully constrained quantity is a combination of the mean matter density and σ_8 , and exhibits a characteristic curved degeneracy, roughly given by $S_8 \equiv \sigma_8 (\Omega_m/0.3)^\alpha$, where $\alpha \simeq 0.5$. See Bartelmann & Schneider (2001); Munshi et al. (2008); Massey et al. (2010); Hoekstra (2013); Kilbinger (2015) for extensive reviews of weak lensing theory and methods.

The Dark Energy Survey (DES) is the largest of the current generation of lensing surveys (known as “Stage III”, in the language of a seminal report commissioned by the US Department of Energy from a group of scientists under the umbrella of the Dark Energy Task Force, Albrecht et al. (2006)). Since the first detection of cosmic shear, the field has expanded rapidly. In just under two decades we have gone from detection of weak lensing around clusters, to field measurements using thousands of galaxies, to state-of-the-art lensing surveys containing tens of millions of high-quality galaxy shape measurements. The group of early measurements, labelled retrospectively as Stage I surveys, include VIRMOS-Descart (Van Waerbeke et al., 2005), CTIO (Jarvis et al., 2006),

SDSS (Hirata et al., 2004) and COSMOS (Schrabback et al., 2007). These surveys were typically limited to fields of 10s of square degrees, but they pioneered the first measurements of cosmic shear. Stage II surveys included DLS (Jee et al., 2013), RCSLenS (Hildebrandt et al., 2016). The first truly competitive cosmology constraints with cosmic shear came from the Canada France Hawaii Telescope Lensing Survey (CFHTLenS; Heymans et al. 2012), in a first non-tomographic analysis by Kilbinger et al. (2013), and then a more complete study using shear measurements in redshift bins by Heymans et al. (2013), who jointly fit for cosmology and multiple systematics. These analyses used state-of-the-art shear measurement techniques, sophisticated photo- z codes and non Gaussian covariance estimates, and set a benchmark for modern cosmic shear inference. A few years later came the preliminary Science Verification (SV) results from DES (Jarvis et al. 2015, Dark Energy Survey Collaboration 2016), which followed much the same methodology with comparable results.

We are now well into Stage III, which includes DES, KiDS and HSC (Aihara et al., 2017), all designed with science programmes featuring prominent cosmic shear components. These surveys are currently engaged in the task of collecting photometric images covering fields of 100s to 1000s of square degrees. DES will eventually survey 5000 square degrees, and has currently completed four out of its five (post-SV) seasons of data. To date the most advanced lensing-based cosmology constraints from Stage III come from the analysis of the first year of KiDS data, which covers 450 degrees of the northern sky and is known as the KiDS-450 (Hildebrandt et al., 2016). The analysis of the Year One DES results is also complete (Troxel et al., 2017), and has been submitted for publication. The most recent published constraints from DES SV and KiDS are shown in Figure 1.7.

The next generation of surveys, known as Stage IV will provide a new paradigm for cosmology, with unprecedented late-time constraints not only on the properties of dark matter, but also of dark energy. These planned surveys are currently in the stages of commissioning or construction, and include the satellites Euclid and WFIRST and large ground based experiments, LSST and SKA; these are forecast to afford measurement of the dark energy equation of state to a precision of 1%.

In what has become one of the major questions in modern cosmology, there is a persistent mild tension between the late-time constraints from lensing and those from the Cosmic Microwave Background radiation, which probes a fixed epoch of around z_* . Specifically, both CFHTLenS and KiDS report a best-fitting $S_8 \sim 2.0 - 3.0\sigma$ lower than Planck. This is remarkable given the fact that the two lensing datasets are entirely independent, as are the technical analysis pipelines. The published DES results are agnostic, being consistent with both results, largely due to the relatively low number densities afforded by the SV catalogues. The results of these two flagship cosmology

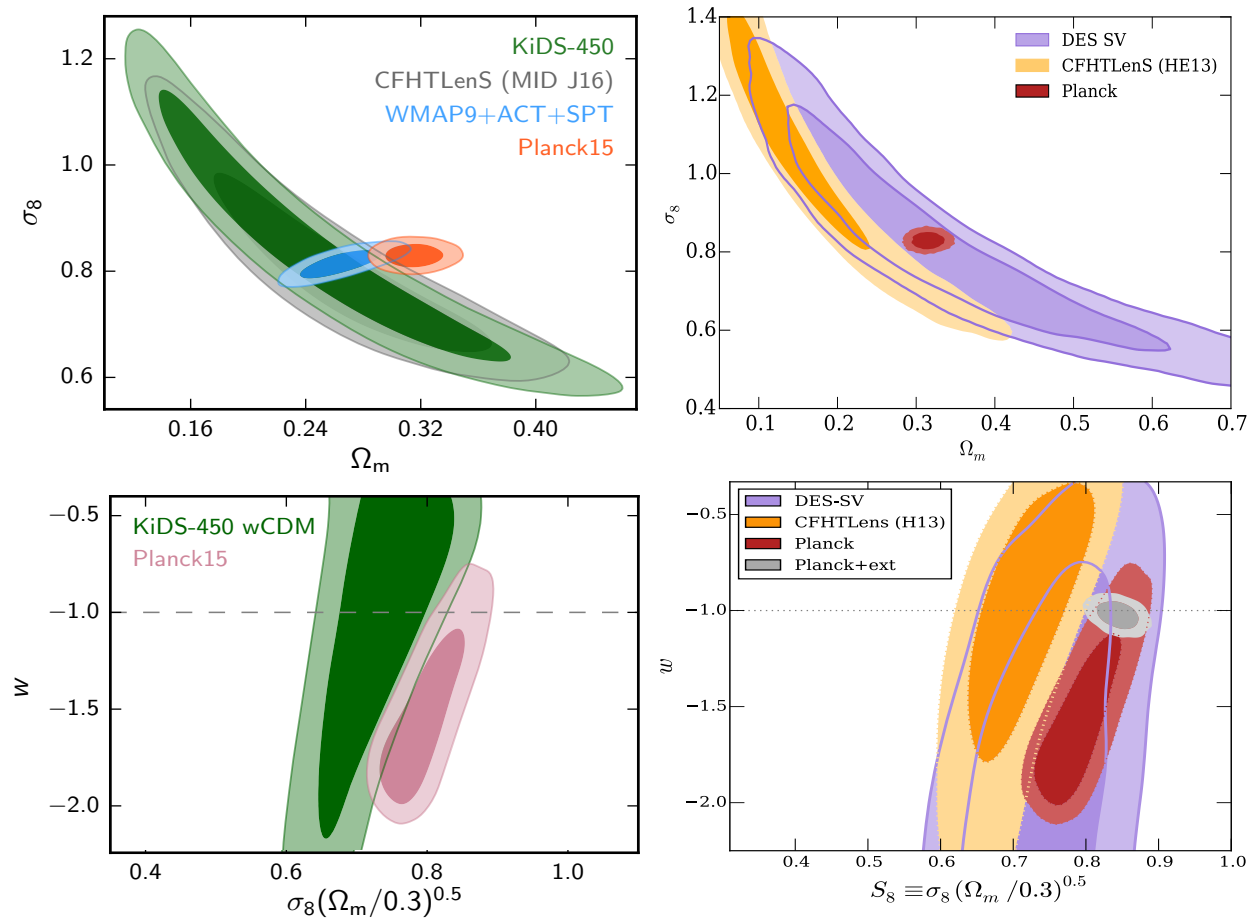


Figure 1.7: A summary of the most recent cosmic shear based cosmology constraints at the time of writing. The upper and lower panels show 1 and 2 σ confidence contours under Λ CDM and w CDM models of cosmology. The left and right-hand panels are the equivalent results from KiDS-450 and DES SV cosmic shear two-point analyses of Hildebrandt et al. (2016) and Dark Energy Survey Collaboration (2016), which are interim studies based on subsets of the final area of each of these experiments.

analyses can be seen in Figure 1.7, alongside the constraints from Planck.

The source of this tension has been the focus of much recent literature, particularly in light of emerging tensions with other low- z probes. Reanalysis of the CFHTLenS data using various conservative systematics models (Joudaki et al., 2017), for example, has been carried out, but finds no convincing single systematic that can account for the disagreement. Choi et al. (2016)'s spectroscopic cross-correlation measurements point to a systematic photo- z bias in CFHTLenS, an idea supported by Kitching et al. (2016b), who claim that the data prefer a combination of (unphysical) negative intrinsic alignment (IA) amplitude in combination with non-zero redshift error.

Later work by Kitching et al. (2016a) asserts that the CFHTLenS tension can be alleviated by higher order (beyond flat sky) extensions to the Limber equation. This claim is, however, contended by Kilbinger et al. (2017), who undergo a comprehensive comparison of such extensions and their expected cosmological impact.

More recently, reanalysis of the KiDS-450 data using an alternative Fourier space shear estimator (Köhlinger et al., 2017) and lensing in combination with galaxy clustering measurements from the GAMA survey (van Uitert et al., 2017) have produced results that are largely consistent with the earlier findings of Hildebrandt et al. (2016).

This is by no means a comprehensive list of such studies, and there has been much discussion, citing everything from redshift error to new physics as plausible explanations for the tension. The current broad consensus is that, as a community, we have not yet ruled out all sources of systematics in one or both datasets, and that robustly doing so would be a necessary precursor to claims about new physics.

1.2.3 Strong Lensing

Strong gravitational lensing is the name given to the distortion of galaxy images in regimes where the weak limit no longer applies. Primarily used as a probe of cluster mass and substructure, strong gravitational lensing is most apparent at intermediate-to-large radii around the most massive structures in the Universe. In recent years, however, it has also been proposed that strong lensing could have a contribution to make as a cosmological probe.

The cosmological power of strong lensing comes in the potential for *cosmography* (see Treu 2010). Most commonly this is exploited using time delays between images of a single background source. If an observer can accurately measure the time delay between image arcs from a single source carried along different geodesics around a massive foreground lens, she can infer the ratio




Figure 1.8: An example of a strongly lensed galaxy cluster as imaged by the HST. The lower panel shows a selection of quadrupole lenses identified as part of the H0LiCOW program (Suyu et al., 2016).

of angular diameter distances within the lens equation. This provides a standard ruler measurement equivalent to the BAO scale in large scale galaxy clustering. In cosmological terms it provides a potentially valuable independent measure of H_0 . Specifically one measures a quantity called the time delay distance:

$$D_{\Delta t}(z) = (1 + z_L) \frac{D_L D_S}{D_{LS}}, \quad (1.30)$$

where the subscripts denote quantities relating to the source and the lens and D_L , D_S and D_{LS} are angular diameter distances. The steps from identifying a strong lens to constraints on H_0 follow as: (a) measure the time delay using the light curve of a variable source with multiple images, (b) predict what the Δt should be under an assumed model for the lens mass distribution, for a given

cosmology, (c) repeat multiple times to map out the likelihood of H_0 .

The power of strong lensing as a way to directly infer the local Hubble parameter was recognised about midway through the Twentieth Century in an essay by the Norwegian astronomer Sjur Refsdal (Refsdal, 1964). In practice, however, it has yet to realise this potential. The latest effort is a project referred to as H0LiCOW (Suyu et al., 2016), which was launched with the stated aim of measuring H_0 to 1% precision using data from the Hubble Space Telescope (HST). A selection of strong lenses identified in the Hubble Deep Field data can be seen in Figure 1.8. Using quasar time delay distances from single lenses, the current best constraints are at just under 4% precision (Wong et al., 2017; Bonvin et al., 2017).

Attempts to use strong lensing as a cosmological probe are currently limited on two fronts. First, there are simply too few strong lenses in the current data of sufficient quality for time delay measurements. This is likely to change in the near- to mid-term future: under some assumptions about the distribution of strongly lensed sources at high z , DES is forecast to reveal 1300 such objects over its lifetime. LSST and Euclid are expected to bring well over 10,000 new strong lens measurements (Collett, 2015). The second limitation, which is less easily bypassed, comes from theoretical uncertainties in the lens models (Treu & Marshall, 2016; Harvey et al., 2016). The biases that can be introduced are currently poorly understood and are without definite mitigation strategies.

1.2.4 Galaxy Clustering & Galaxy-Galaxy Lensing

One alternative to cosmic shear is to map out the positions of galaxies, under the assumption that dark matter collapsing under its own gravity will drag visible matter along with it. Galaxies are believed to form in the centres of dark matter halos, and so one would expect the density of galaxies to trace out the underlying mass fluctuations. In general the correspondence is not exact, due to the complications of late-time effects produced by the collapsing baryons (generically termed *baryon feedback*). One can write a general relation between the power spectra of galaxies and dark matter overdensities,

$$P_{gg}(k, z) = b_g^2 P_\delta(k, z), \quad (1.31)$$

where we have defined $(2\pi)^3 \delta_D(k - k') P_{gg}(k) \equiv \langle \delta_g(k) \delta_g^*(k') \rangle$, and δ_g is the overdensity of galaxies, defined analogously to equation 1.18. The *galaxy bias* b_g is actually the first term in a series expansion of the galaxy overdensity δ_g . A single linear bias is commonly assumed to be sufficient to describe observations on scales above ~ 100 Mpc. In general, though, b_g can be a

function of both scale and redshift. If b_g is known perfectly, galaxy clustering is a powerful probe of cosmology; as things stand, however, cosmological constraints from clustering are significantly degenerate with b_g , and there are considerable gaps in our theoretical understanding of galaxy bias. Clustering is, in general, more commonly considered as a complementary probe, most useful when combined with other measurements.

We will consider four measurements based on mapping out galaxy positions which can be useful, from a cosmological standpoint: baryon acoustic oscillations (BAO), redshift space distortions, the broadband shape of the power spectrum and galaxy-galaxy lensing.

Shape of the Matter Power Spectrum

If one can successfully mitigate the uncertainty of galaxy bias galaxy density measurements provide a direct window onto the dark matter power spectrum, with a peak sensitivity determined by the chosen galaxy sample. This sensitivity is completely analogous to the Limber equation in cosmic shear (see the next chapter), but with the lensing efficiency kernels replaced with the raw redshift distribution of galaxies. That, of course, brings information on all of the cosmological parameters that enter P_δ (σ_8 , A_s , Ω_m , n_s and in principle the neutrino mass).

Baryon Acoustic Oscillations & Redshift Space Distortions

Instead of considering the information from the full power spectrum, one could instead choose to focus on a range of scales where there are characteristic features of particular interest. Specifically, the acoustic oscillations left at the time of matter-radiation equality leave a distinctive sinusoidal pattern at intermediate k scales, which persists at late times.

The length scale associated with BAO r_s thus provides a standard ruler measurement of geometry and expansion. Unlike supernovae, for example, which provide an unambiguous H_0 measurement, BAO probe a degenerate combination of $H_0 r_s$, and thus require an independent measure of r_s in order to meaningfully constrain H_0 . One particularly interesting result to emerge from galaxy clustering measurements in the past decade has been from the combination of BAO and CMB data. Specifically, the combination of WMAP and SDSS (Aubourg et al., 2015) favours a low H_0 , in tension of $\sim 2.5\sigma$ with more direct late-time measurements.

In addition to BAO, line-of-sight peculiar velocities also leave an imprint on the measured distribution of galaxies, distorting distances deduced using equation (1.8). These effects are called *Redshift Space Distortions* (RSD). On scales $\lesssim 1h^{-1}\text{Mpc}$ the peculiar velocities of galaxies are dominated by random intracluster cluster motion, which tends to elongate apparent structure along

the line-of-sight. Conversely, at distances $\gtrsim 10h^{-1}$ Mpc, cluster-scale objects have yet to become virialised and gravitational collapse dominates motion. It is these distant measurements through which RSD become a probe of structure formation.

The RSD datavector commonly incorporates cosmology directly, with each RSD measurement using a sample at some mean redshift providing a single datapoint $f(z)\sigma_8$, where the prefactor is the logarithmic growth rate $f(z) \equiv d \ln D(a)/d \ln a$. (Percival et al., 2011).

Galaxy-Galaxy Lensing




Figure 1.9: A recent example of the constraining power of galaxy-galaxy lensing, as measured from DES SV. The filled blue contours indicate the constraints on cosmological parameters (left) and galaxy bias (right) from galaxy-galaxy lensing combined with galaxy clustering. The left-hand panel also shows the comparable confidence contours provided by cosmic shear alone and the CMB. The solid and dashed lines in the right-hand panel indicate two sets of scale cuts, and were used by the authors to locate the angular scale above which nonlinear bias could be neglected. **Figure credits:** Kwan et al. (2017).

In the above we have discussed the use of galaxy clustering and cosmic shear correlations to independently probe cosmology. The cross-correlation of the two, however, also contains information beyond the sum of the two autocorrelations. This statistic, which correlates the positions of foreground “lenses” with the shapes of background “source” galaxies, probes the galaxy-mass cross spectrum $P_{g\delta}(k)$. Though mapping either $P_{g\delta}(k)$ or $P_{gg}(k)$ onto the matter power spectrum is complicated by an unknown galaxy bias coefficient b_g , they each relate to $P_\delta(k)$ in a slightly different way. Thus the combination of galaxy clustering and galaxy-galaxy lensing can break the

degeneracy between the power spectrum amplitude and b_g . This ability does, however, in general depend on the galaxy bias being linear (i.e. independent of k), which is not the case on small scales. In general one has three options: (a) to cut the data sufficiently such that linear bias holds to good approximation, (b) to explicitly model higher order bias terms (e.g. Kwan et al. 2017) or (c) to devise a statistic which nulls or downweights small scales such that the bias can be treated as linear (Baldauf et al., 2010).

Like galaxy clustering, galaxy-galaxy lensing in isolation is a rather weak probe of cosmology due to uncertainty in galaxy bias. In combination with galaxy clustering, however, it starts to become competitive as a cosmological probe (Mandelbaum et al., 2013; Kwan et al., 2017). We show an example of the constraining power of such data in Figure 1.9, which illustrates the primary value in constraining the linear bias coefficient. Indeed as part of the so-called 3x2pt combination of late-time probes, alongside cosmic shear and galaxy clustering, it has the potential to internally constrain lensing systematics as well as adding statistical constraining power. This was demonstrated in principle by Eifler et al. (2014), Samuroff et al. (2017b) and Krause et al. (2017), and more recently in practice by van Uitert et al. (2017) and DES Collaboration et al. (2017).

1.2.5 Supernovae & Supernova Lensing

A class of exploding star called the Type Ia supernova has been known for some time to produce largely predictable light curves, and so act as a form of standard candle. Measurements of these extreme astrophysical objects provide a direct measurement of the Hubble parameter at specific redshifts through equation (1.8).

Generated by thermonuclear explosions in collapsing carbon-oxygen white dwarfs, every SNIa should have a standard (or at least *standardisable*) duration and peak luminosity. The measurable quantity is the distance modulus or the difference between the apparent and absolute magnitudes, $\mu = m_L - M_L = 5 \ln D_L + 25$. The luminosity distance, D_L , is sensitive to geometry and expansion as,

$$D_L(z) = \begin{cases} \frac{c(1+z)}{H_0\sqrt{\Omega_K}} \sinh\left(\frac{\sqrt{\Omega_K}}{c} H_0\chi\right) & \Omega_K > 0 \\ (1+z)\chi & \Omega_K = 0 \\ \frac{c(1+z)}{H_0\sqrt{|\Omega_K|}} \sin\left(\frac{\sqrt{|\Omega_K|}}{c} H_0\chi\right) & \Omega_K < 0, \end{cases} \quad (1.32)$$

where $\chi = c \int_0^z H^{-1}(z') dz'$ from equation 1.6. In the local Universe geometric effects are small and D_L is approximately linear with z , providing a direct measure of H_0 . At high z equation 1.32

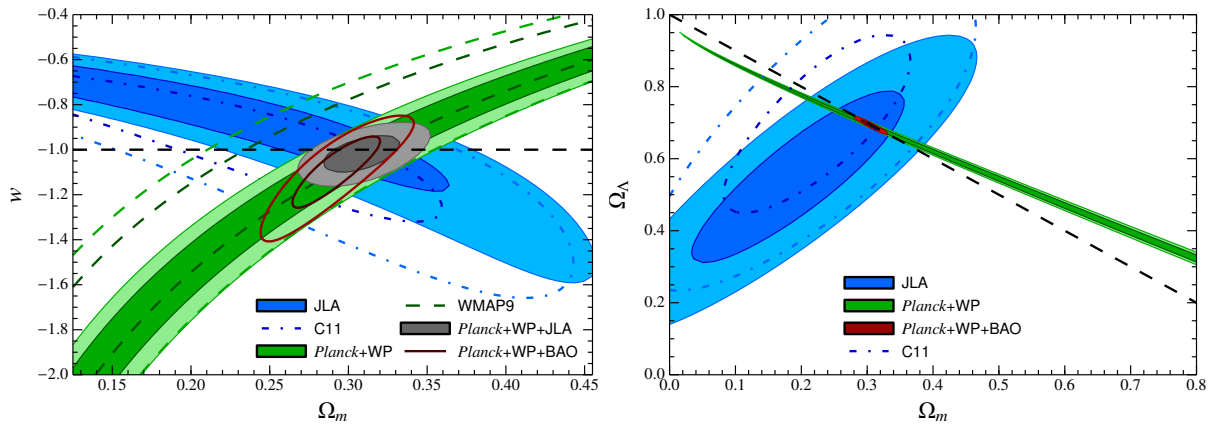


Figure 1.10: An example of the constraining power of type Ia Supernova. Joint constraints on dark energy and dark matter from a joint analysis of low and high redshift supernovae from the SNLS and SDSS datasets. **Figure credits:** Betoule et al. (2014)

can be fitted to constrain Ω_m , Ω_{de} , Ω_K , w_0 and w_a . For general reviews of SNIa cosmology see Goobar & Leibundgut (2011) and Weinberg et al. (2013).

In practice SNe are imperfect standard candles, exhibiting finite astrophysical scatter in M_L . Typically one will make corrections for systematics in the measured shape, colour and peak brightness of the SN light curve. Interstellar reddening, z -evolution, time-dilation, and an observed relation between luminosity and host galaxy mass must also be corrected or marginalised. Due to limitations in the accuracy of theoretical models, SNIa luminosities must be calibrated using external distance estimates. It is common to use variable stars with well-defined periodicities for this (see e.g. Riess et al. (2009); Efstathiou (2014)).

In a discovery which was later awarded a Nobel prize, it was from SNIa lightcurves that Riess et al. (1998) and Perlmutter et al. (1999) derived the first observational evidence for accelerated late-time expansion of the Universe. They found a systematic increase in D_L for given z relative to predictions based on a flat matter dominated (Einstein de Sitter) universe. Their observations implied a flat universe with $\Omega_m = 0.3$, $\Omega_{de} = 0.7$. Significant strides have been made in the size and redshift range of SN catalogues in recent years, thanks in part to improved targetting of host galaxies and new tailored surveys. For example, the Supernova Legacy Survey (SNLS; Guy et al. 2010) and Equation of State: Supernovae trace Cosmic Expansion (ESSENCE; Sollerman et al. 2006) have provided high-time resolution imaging of several hundred SNIa at $z = 0.3 - 1.0$. Others such as the Carnegie Supernova Project (Freedman & Carnegie Supernova Project 2005) have been useful in tracing the local ($z < 0.1$) part of the $D_L - z$ relation. In addition ~ 400 low-

to mid-redshift SNe have been detected by the SDSS-II SN Survey.

The most precise measurements of H_0 and high- z (> 1) SNe currently come from space based instruments. Given the narrow field-of-view, targeted searches of distant clusters greatly aid detection efficiency. The total well-measured SNIa sample contains ~ 700 objects, although the number in the high- z regime useful for probing dark energy is still small. As has been noted elsewhere, low- z uncertainties are dominated by flux calibration and poor understanding of astrophysical systematics, while high- z measurements are currently noise-dominated due to the small samples available.

As alluded to previously there has been some discussion of a mild tension between direct local measurements of H_0 (e.g. using type Ia supernovae) and the early-time standard ruler measurement of the CMB. This has also been noted in BAO-based studies, which provide an entirely independent direct local H_0 constraint.

The SN-CMB tension was highlighted in Riess et al. (2011) and is still present in more recent work by the same authors. Riess et al. (2016) present a $H_0 = 73.24 \pm 1.74 \text{ km s}^{-1} \text{ Mpc}^{-1}$ constraint based upon the HST SHOES Survey, which is discrepant with the Planck-only best-fit by over 3σ . Recent reanalysis by Efstathiou (2014) reduces but does not remove this disagreement.

Given the precision and direct nature of these measurements, it has become standard practice to adopt their constraints as a prior on H_0 in analyses based on other probes. Using high- z SNIa from a range of surveys various attempts have been made to derive constraints on dark energy (Sullivan et al., 2011; Suzuki et al., 2012; Betoule et al., 2014). With the Riess et al. (2016) prior on H_0 , their results are consistent with flat Λ CDM. As shown in Figure 1.10, the parameter combination probed by SNe alone exhibits some degeneracy between matter density and w_0 .

The same is true of the CMB, but the difference in degeneracies makes them complementary cosmological probes. Joint analysis including CMB+BAO measurements from Planck and SDSS alongside SNe, tightens constraints to $w_0 = -1.013_{-0.073}^{+0.068}$ (Betoule et al., 2014). Whereas the Planck+BAO combination weakly favours $w_0 < -1$ (see Figure 1.10), the addition of the SN data appears to shift the joint fit back towards Λ CDM.

Though these authors make the attempt, the current data is insufficient to meaningfully constrain w_a . As pointed out by these authors tomographic analysis, which may aid in this task, is currently hindered the available SN sample size. This situation is unlikely to change at a stroke, and clearly one cannot simply collect supernovae from a survey in anything close to the abundances of galaxies involved in cosmic shear measurement. Larger survey volumes will, however, bring more and deeper supernova samples. By the end of its lifetime, for example, the Dark Energy Survey is projected to provide around 4000 Type Ia SNe (Bernstein et al., 2012).

One final interesting point to note is that photons from distant supernovae will be lensed in the same way as all other light as it propagates through the Universe. This can induce scatter in the measured magnitudes of SNe, detectable either as a change in the moments of the magnitude distribution, or else through the correlations it induces between supernovae. Such signals are in principle detectable but, again, at present time of limited cosmological use due to the size of the available SN samples (Quartin et al., 2014; Scovacricchi et al., 2017).

1.2.6 The Cosmic Microwave Background & CMB Lensing

Temperature fluctuations in the ambient cloud of photons that has pervaded the Universe since the end of the radiation dominated era are a direct imprint of sound waves during that epoch (see Section 1.1.4). On first approach the power spectrum of temperature fluctuations (of which we show an example in Figure 1.11) is a rich source of cosmological information. The flat left-hand region represents super-horizon scales at the time of recombination. As such, its amplitude and gradient reflect the primordial matter power spectrum, (A_s , n_s and potentially dn_s/dk). The height of the primary peak above this plateau measures $\Omega_m h^2$. Its angular scale is sensitive to the product of the angular diameter distance to the CMB last-scattering surface and H_0 . Via the Hubble parameter $H(z)$, it is also sensitive to the density parameters ($\Omega_m, \Omega_K, \Omega_r, \Omega_{de}$). These features are all encoded in the CMB at early times and, assuming foregrounds and instrumental systematics are properly accounted for, provide a window onto the era of radiation domination.

At late times the temperature spectrum is modified by interactions with foreground large scale structure. In addition to being systematics, these alterations can provide useful information about the late-time mass distribution. For example localised gravitational redshifting of the CMB (i.e. the Sunyaev Zel'dovich effect) is commonly used as a method to identify massive galaxy clusters (e.g. Carlstrom et al. 2002; Planck Collaboration 2015b). Lensing by large scale structure also has a net smoothing effect on the CMB temperature fluctuations.

In recent years a number of novel approaches have been suggested for extracting information from the CMB data. Since it passes through the same large scale structure as light from distant galaxies, the CMB undergoes continuous cosmological shear, and so in principle probes the same large scale structure, albeit with a lensing efficiency kernel $g(z)$ which peaks at higher redshift (eg. Planck Collaboration 2016, Kirk et al. 2016).

Though no cosmological observation can be guaranteed to be entirely systematics-free, the widely divergent natures of both the probe and the measurement itself suggest that the systematics entering cosmic shear and the CMB should be different and (hopefully) *uncorrelated*. A number of




Figure 1.11: **Upper panel:** Angular power spectrum for CMB temperature anisotropies, as obtained from the Planck 2015 dataset. Error bars show 68% confidence regions. The blue points are each averaged over ~ 31 modes. The best-fitting Λ CDM prediction is demarked by the red line. **Lower panel:** Data-model residuals, relative to the best fitting Λ CDM case. **Image credits:** Planck Collaboration (2015a).

studies have thus suggested using cross-correlations between the two to constrain systematics in both the late and early time data (see e.g. Schaan et al. 2016, Harnois-Déraps et al. 2017)

The first detection of the CMB (Penzias 1965) was a serendipitous observation, suspected in the first instance as having avian origin. In the decades that followed, however, CMB measurements have become an observational mainstay of the standard Λ CDM model. The satellite-based Cosmic Background Explorer (COBE; Bennett et al. 1996) made the first detection of the temperature anisotropy map at higher order than the dipole. There were a number of interim balloon-based surveys, but the first major cosmology constraints from the CMB came from the the Wilkinson Anisotropy Probe (WMAP; Spergel et al. 2003); the final 9-year data release offered the most stringent constraints at that time (Hinshaw et al. 2013). A new generation of ground-based instruments (e.g. the South Pole Telescope, or the Atacama Cosmology Telescope) have sought to add complementary data to the satellite-based WMAP and Planck, providing high-resolution but narrow-field observations, which are sensitive to the high- ℓ tail of the spectrum.

The strongest single-experiment cosmology constraints came in 2015 with a second data release from Planck (Planck Collaboration, 2015a). Measurement of $n_s < 1$ to 6σ has been confirmed. They have provided the most stringent measurements of large scale geometry, combining Planck CMB data with low redshift galaxy clustering data from SDSS (Blake et al. 2007) and assuming $w_a = 0$, to find $\Omega_K = 0.0005 \pm 0.0033$. In a particularly intriguing element of the recent CMB analyses, mild tension had emerged with local estimates of H_0 based on SN and variable stars (discussed in Sec. 1.2.5).

A number of recent attempts have been made to measure two-point functions of CMB lensing, either with itself or the cross-spectra with conventional galaxy lensing or clustering. For example,

after Planck Collaboration (2016), Giannantonio et al. (2016) present the first *tomographic* CMB lensing results with SPT. As alluded to before, a growing body of work focuses on CMB cross correlations, as a cosmological probe and a means to internally constrain systematics (Hand et al. 2015; Liu & Hill 2015; Kirk et al. 2016; Harnois-Déraps et al. 2017). The main cosmology results of these studies assume a fiducial cosmology, and rescale the lensing power spectrum to match the data. As discussed in Liu & Hill (2015) (see also e.g. Di Valentino et al. 2017) Planck-CFHTLenS cross-correlations seem to favour $A_{\text{lens}} < 1$ (i.e. σ_8 smaller than the Planck-only best fit). Although apparently consistent with tensions between the datasets taken apart, it still not clear at this point that one or more systematic is not responsible for this apparent disagreement.

1.3 Stage IV and the Future of Cosmology

Within the next decade the current generation of Stage III surveys will reach the end of their lifetime. These are the most ambitious projects in mapping the late-time Universe ever undertaken. It will not, however, be long before they are superceded by a new generation of experiments, which are currently in the stages of planning. Unlike their predecessors, these are instruments built primarily for cosmology, and will offer a paradigm shift in statistical power. These experiments are referred to as Stage IV in the language of Albrecht et al. (2006). The following paragraphs will review the field of major cosmology-oriented surveys, planned and extant. We will conclude this chapter with a brief discussion of novel lensing probes which may offer new insights in the years ahead.

1.3.1 Photometric Surveys

For reasons discussed above, the most statistically powerful datasets for cosmic shear necessarily use broad band photometric filters. Table 1.2 summarises the basic characteristics of some recent notable examples in modern cosmology.

The Kilo-Degree Square Survey (KiDS; de Jong et al. 2013), Dark Energy Survey (DES; Flaugher et al. 2015), and Hyper SuprimeCam (HSC; Aihara et al. (2017)) are the largest of the ongoing cosmic shear surveys. Inevitably their survey strategies must compromise depth and area, and clearly the science aims must play a role in survey strategy. Deep data has clear advantages in terms of photo- z accuracy (Benjamin et al., 2013) and resolution. Depth is also useful as it maximises the number of faint, high redshift galaxies in which the cosmological shear is strongest, and may suppress some astrophysical systematics (Krause et al., 2015). It is, however, also susceptible

to PSF instability and cosmic variance. A series of forecasting studies have more or less settled this debate by demonstrating that optimising area for given exposure time even at a cost to depth can in general improve the constraining power of a survey (Amara & Réfrégier, 2007; Kirk et al., 2013).

KiDS and DES favour breadth over depth at 1500 and 5000 square degrees respectively. Notably the former has significant overlap with the VIKING near-infrared field and with the GAMA survey, providing high-quality wide-band imaging. By contrast HSC favours a slightly different balance. The survey is being conducted in two layers, the first covering 1400 square degrees to an r -band depth of ~ 26 th magnitude. The second reaches almost to 27 mag, but over a tiny field of 3.5 square degrees. The ultimate aim of these surveys is to make the first competitive (non-CMB) constraints on w_0 . By the end of their lifetimes they are set to afford percentage level constraints on the dark energy parameters.

For obvious reasons Stage IV will suffer far less from cosmic variance than the current generation. Covering 20,000 square degrees over its 10 year lifespan, the Large Synoptic Survey Telescope¹⁰ (LSST; LSST Science Collaboration (2009)) will provide a step change in statistical power. It is expected that the project will achieve the first competitive (i.e. percentage level) late-time constraints on w_0 and w_a (Albrecht et al., 2006). An analogous weak lensing project at radio frequencies will be undertaken by the Square Kilometre Array¹¹ (Dewdney, 2015). Though the survey strategy is still the subject of much discussion, cosmic shear in a population of faint radio sources has been proposed, both as a means to constrain cosmology and a complementary way to constrain systematics through cross-correlations with optical counterparts (Brown et al., 2015; Patel et al., 2015; Harrison et al., 2016).

It is finally worth mentioning briefly an idea that has been mooted but hitherto not realised: using specialised balloon-borne observations for weak lensing measurement. Balloon experiments have been implemented with some success in the context of the CMB (Mauskopf et al., 1999; Crill et al., 2008). The most recent incarnation of this concept is a NASA-led project called SuperBIT (Romualdez et al., 2016), which in current form will benefit from a wider field of view than extant optical satellites as well as the advantageous seeing conditions brought by altitude. This project is still very much in its interim stages, and the science goals focus on cluster lensing.

Survey (Telescope)	Observing Time	Bands	No. Galaxies ($\times 10^6$)	Sky Coverage (square degrees)
KiDS (VST)	2012-2016	<i>ugri</i>	90	1500
DES (Blanco)	2013-2018	<i>grizY</i>	300	5000
HSC (Subaru)	2014-2019	<i>grizY</i>	400	1400
PAU (WHT)	2013-2016	40 narrow-band	0.03	100-200
LSST (LSST)	2021-2031	<i>ugrizY</i>	4000	20000
Euclid (Euclid*)	2020-2025	<i>R+I+Z,Y,J,H</i>	1500	14700
SKA2 [†] (SKA)	2023+	Radio Bands 1,2,3	1100	30000
WFIRST (WFIRST*)	2024-2030	<i>Y+J+H+F184</i>	500	2000

Table 1.2: Summary of recent, ongoing and planned lensing surveys. All listed include large-scale structure or cosmic shear measurement programmes. Adapted from Table 2 in Kirk et al. (2015). The lines marked with an asterisk are hosted by satellites, where the rest are ground-based. [†] Interim forecasts for for a second-stage SKA galaxy lensing survey taken from Harrison et al. (2016).

1.3.2 Spectroscopic Surveys

Cosmologically speaking, spectroscopy has two main uses. First, it is highly useful for estimation of the source redshift distribution in lensing, both as a means to test photo- z algorithms and also to eliminate systematic errors through cross-correlation (Kirk et al., 2013). Second, it provides highly accurate per-object (rather than ensemble) redshift estimates, which are needed for standard ruler (BAO) measurements. Such detailed single-object information is also necessary for direct measurement of localised astrophysical phenomena such as intrinsic alignments. Dedicated surveys in planning such as SDSS-IV’s eBOSS¹² programme will target massive luminous red galaxies (LRG) and quasars to considerably deeper redshifts than its predecessors. This is designed to cover the epoch of transition to dark energy domination, a region where currently spectroscopic sampling is poor. The planned Dark Energy Spectroscopic Instrument¹³ (DESI; Levi et al. 2013) will be larger and deeper still. In addition to valuable photo- z validation for other surveys, early plans for that experiment list neutrino mass detection amongst its science goals, anticipating a 2σ measurement.

Potentially even more valuable will be a handful of new satellite missions with spectroscopic capabilities, primarily ESA’s Euclid¹⁴ (Refregier & Douspis, 2008) and NASA’s WFIRST¹⁵ (Spergel

¹⁰lsst.org/lstt

¹¹skatelescope.org

¹²sdss.org/surveys/eboss

¹³desi.lbl.gov

¹⁴euclid-ec.org

¹⁵wfirst.gsfc.nasa.gov

Survey (Telescope)	Observing Time	No. Galaxies ($\times 10^6$)	Sky Coverage (square degrees)
eBOSS (APO)	2014-2020	0.6 LRG + 0.07 Lyman- α	7000
DESI (NOAO)	2018-2022	32.0 LRG + 2.0 Lyman- α	18000
PFS (Subaru)	2018-2023	4.0	1400
4MOST (VISTA)	2019-2024	6.0-20.0 bright objects	15000
Euclid (Euclid*)	2020-2025	75.0	14700
WFIRST (WFIRST*)	2024-2030	20.0	2000

Table 1.3: Summary of recent, ongoing and planned spectroscopic surveys. All listed include BAO observation programmes. Adapted from Table 1 in Kirk et al. (2015). As in Table 1.2 the lines marked with an asterik are hosted by satellites and those without are ground-based.

et al., 2015). Euclid’s near-infrared spectrograph affords higher sampling density but narrower sky-coverage, reducing statistical noise within the field with the obvious caveat of sample variance. Deep low noise, high resolution optical data is invaluable for many reasons. Apart from the benefits to photo- z accuracy already mentioned, these missions will provide input for the next generation of shear calibration simulations. At present time virtually all mock galaxy catalogues using real galaxy morphologies (including those presented in this thesis) borrow from the low noise HST COSMOS imaging. For obvious reasons over reliance on a relatively small calibration field to test shape measurements in ever larger/deeper datasets is not ideal (see Kannawadi et al. 2015 for a more detailed discussion of this issue).

1.3.3 Non-Conventional Lensing Methods

Finally in this chapter we will give a very short overview of a number of lensing-based methods that are widely accepted as *potentially* valuable, but for various reasons have yet to realise that potential.

One such novel probe is lensing by voids. Cosmic scale underdensities, like their overdense counterparts, induce deflections to light from distant sources. More accurately thought of as a reduction in the amplitude of the background distortions, this in effect produces radial ellipticity component. This effect has been detected on a handful of occasions (Melchior et al., 2014; Gruen et al., 2016; Sánchez et al., 2017). Cosmological applications largely centre on testing for signs of modified gravity (Cai et al., 2015), though these ideas are yet to be put into practice.

In addition to undergoing a coherent shear, one might expect background sources to be magnified by the continuous background mass distribution. Lensing magnification is significantly harder

to detect than shear, simply because one cannot accurately predict the mean of the unlensed galaxy size distribution. With the aid of galaxy number counts, however, a number of studies have reported detection of a magnification signal (Schmidt et al., 2012; Garcia-Fernandez et al., 2016). Though in principle a probe of large scale structure akin to shear, the veracity of magnification-based inference is largely dependent on accurate fore-knowledge of the galaxy luminosity function.

It has been understood for some time that constructing maps of shear on the sky and searching for local peaks above the noise can be used to identify massive galaxy clusters. There has, however, been more recent recognition that the raw abundance, and correlation functions of shear peaks can also be viable probes of cosmology. In practice this has been shown to be comparable in constraining power to non-tomographic cosmic shear correlations (Kacprzak et al., 2016). There are, however, still significant gaps in our understanding of the systematics, and much less time and effort has been expended into understanding how effects such as intrinsic alignments enter peak counts, compared with two-point cosmic shear statistics.

Chapter 2

The Basic Principles of Weak Gravitational Lensing

2.1 Cosmic Shear in Theory: Formalism

Gravitational lensing is a fundamental consequence of General Relativity, and has been confirmed repeatedly over the course of the Twentieth and Twenty First Centuries. A detailed understanding of GR's formalism is not, however, necessary in order to derive the basic set of results upon which most shear-based analyses rest. This is well-trodden ground, and we refer the reader to the many serviceable reviews of the subject for more detailed derivations.

2.1.1 Point Deflectors, Critical Density & Convergence

To proceed we will perform a simple thought experiment. Imagine one has a lensing system in which a distant source of photons (the “source”) is distorted by foreground mass which is contained by a thin slice of space at a common distance D_L . It isn't important how that arises; just imagine one has a highly compact mass causing lensing and nothing else. Under this thin lens approximation, the path taken by photons between the source, lens and observer follow straight lines, with a single instantaneous point of deflection when they reach the lens plane. We show a sketch diagram of our lens configuration in Figure 2.1.

We'll now make our first and last call upon GR; in the case where the lens has no two dimensional extent (i.e. a point mass in the lens plane) Einstein's relativity predicts that the deflection angle should be given by




Figure 2.1: Schematic diagram of a simple lensing system. The purple (foreground) slice represents an idealised thin lens plane, which contains all of the lens mass. A galaxy observed in the lens plane (green, background) has a true angular position (denoted β) and a projected position after lensing (θ), which are not trivially the same. The purple arrows show the trajectory of a photon from emission at the galaxy's true position to observation. The dotted purple line shows the backwards projection to the apparent position of the galaxy in the image plane. The distances between observer, lens and source (D_L , D_S and D_{LS}) are angular diameter distances as defined in equation 1.9.

$$\hat{\alpha}(\mathbf{R}) = \frac{4Gm}{c^2 R} \hat{\mathbf{R}}. \quad (2.1)$$

As in Figure 2.1, R is the radial position at which the photon hits a lens plane defined about point mass m . The vector quantity $\hat{\mathbf{R}}$ is a unit vector in the direction of that deflection point. Assuming small angles simple geometry then gives us an expression for the deflection angle $\alpha \equiv \theta - \beta = (D_{LS}/D_S)\hat{\alpha}$. One then finds

$$\alpha = \frac{D_{LS}}{D_S D_L} \frac{4Gm}{c^2 \theta^2} \theta, \quad (2.2)$$

since we know the angle of deflection must follow the physical radius R in the lens plane.

This description can be generalised for to describe an arbitrary mass distribution composed of an ensemble of point deflectors. If we call $d\alpha$ the small deflection due to a tiny mass dm , then the total deflection angle is simply the sum of contributions from all of the point masses in the lens. Since we can write the intercept angle as $\mathbf{R} = D_L \boldsymbol{\theta}'$, it follows that a small area of the lens plane at position $\boldsymbol{\theta}'$ will contain a total mass

$$dm(\boldsymbol{\theta}') = D_L^2 \Sigma(\boldsymbol{\theta}') dx' dy'. \quad (2.3)$$

The lens is again considered to be a thin mass sheet with surface density $\Sigma(\boldsymbol{\theta})$ and the coordinates x, y are Cartesian position coordinates in the image plane, $\boldsymbol{\theta} = (x, y)$. Combining equations 2.2 and 2.3 the deflection angle for an ensemble lens distribution is given by the sum over its constituent masses:

$$\begin{aligned}\boldsymbol{\alpha}(\boldsymbol{\theta} - \boldsymbol{\theta}') &= \frac{4G}{c^2} \frac{D_{LS}}{D_S D_L} \sum_i dm(x'_i, y'_i) \frac{\boldsymbol{\theta} - \boldsymbol{\theta}'}{|\boldsymbol{\theta} - \boldsymbol{\theta}'|^2} \\ &\approx \frac{4G}{c^2} \frac{D_{LS} D_L}{D_S} \int \frac{\boldsymbol{\theta} - \boldsymbol{\theta}'}{|\boldsymbol{\theta} - \boldsymbol{\theta}'|^2} \Sigma(\boldsymbol{\theta}') d^2\theta'\end{aligned}\quad (2.4)$$

It is then common to rewrite this in a slightly more elegant form,

$$\boldsymbol{\alpha} = \frac{1}{\pi} \int \frac{\boldsymbol{\theta} - \boldsymbol{\theta}'}{|\boldsymbol{\theta} - \boldsymbol{\theta}'|^2} \kappa(\boldsymbol{\theta}') d^2\theta', \quad (2.5)$$

by defining quantities called *critical density*

$$\Sigma_{\text{cr}} \equiv \frac{D_S}{D_L D_{LS}} \frac{c^2}{4\pi G}, \quad (2.6)$$

and *convergence*

$$\kappa(\boldsymbol{\theta}) \equiv \frac{\Sigma(\boldsymbol{\theta})}{\Sigma_{\text{cr}}}. \quad (2.7)$$

We should note that it is through Σ_{cr} that lensing becomes sensitive to the angular diameter distances within the lens system, and thus the background geometry of the Universe.

At this point it is common to define another scalar quantity called the effective lens potential or the deflection potential as the scalar analogue to the vector field of the deflection angle:

$$\psi(\boldsymbol{\theta}) \equiv \pi^{-1} \int \kappa(\boldsymbol{\theta}') \ln(|\boldsymbol{\theta} - \boldsymbol{\theta}'|) d^2\theta', \quad (2.8)$$

which implies

$$\nabla_{\boldsymbol{\theta}} \psi = \boldsymbol{\alpha}, \quad (2.9)$$

and

$$\nabla_{\boldsymbol{\theta}}^2 \psi = 2\kappa. \quad (2.10)$$

One can think of the three quantities κ , ψ and $\boldsymbol{\alpha}$ as representative of the physical properties of the

lens: a (weighted) mass density, the local gravitational field and the deflection angle imposed on photons, all evaluated at a particular point in the lens plane θ .

2.1.2 The Lens Equation

Now consider a small aperture in the observed image plane containing total pixel intensity $I_{\text{im}}(\theta)$. We can state in general terms that lensing maps photons back onto the source plane¹. If the deflection angle is small the position mapping can be assumed to be linear, and to conserve the total flux in the aperture. One can then write $I_S(\beta) = I_{\text{im}}(\mathbf{A}\theta)$, where I_S and I_{im} are the intensities of points in the source and image planes, or more usefully in terms of a position mapping

$$\delta\beta = \frac{\partial\beta}{\partial\theta}\delta\theta = \mathbf{A}\delta\theta, \quad (2.11)$$

where we've defined a 2×2 matrix \mathbf{A} , which is just the Jacobian of the mapping of a small shift in the image plane back onto the source plane. We will refer to it as the *distortion matrix*.

Considering Figure 2.1 again, more simple geometry provides what is generally labelled the *lens equation*,

$$\beta = \theta - \alpha. \quad (2.12)$$

Using the last two equations, it is then just a matter of algebra to obtain

$$A_{ij} = \delta_{ij} - \frac{\partial\alpha_i}{\partial\theta_j} = \delta_{ij} - \frac{\partial^2\psi}{\partial\theta_i\partial\theta_j}, \quad (2.13)$$

where the indices denote the two spatial coordinate axes and δ_{ij} is the Kronecker delta function (i.e. 1 if $i = j$, 0 otherwise). The second equality relies on the earlier definition in equation 2.9.

Next consider how the lensing distortion alters a straight line between two points in the image plane. In general terms we can decompose the transformation into the sum of two parts. First we have an isotropic (i.e. diagonal) distortion, which changes the length of the line but leaves its shape and direction unchanged. Second, we have a trace-free part, which can extend or compress the line but also rotates it through an angle 2ϕ . That is,

$$\mathbf{A} = \mathbf{A}_{\text{iso}} + \mathbf{A}_{\text{TF}} = \begin{pmatrix} 1 - \kappa & 0 \\ 0 & 1 - \kappa \end{pmatrix} - \begin{pmatrix} \gamma_1 & \gamma_2 \\ \gamma_2 & -\gamma_1 \end{pmatrix}. \quad (2.14)$$

¹i.e. the image plane as it would appear in the absence of lensing.

Here we have introduced a complex quantity called *shear*, $\gamma = \gamma_1 + i\gamma_2 = |\gamma| \exp(-2i\phi)$, where 2ϕ is the rotation angle referred to above which describes the anisotropic part (essentially a rotation matrix with amplitude γ). One can then see that the shear and convergence are given by the second derivatives of ψ :

$$\gamma_1 = \frac{1}{2} \left(\frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial y^2} \right) \psi, \quad \gamma_2 = \frac{\partial^2 \psi}{\partial x \partial y}, \quad \kappa = \frac{1}{2} \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right) \psi. \quad (2.15)$$

So in general lensing will induce both an anisotropic shear and an isotropic magnification.

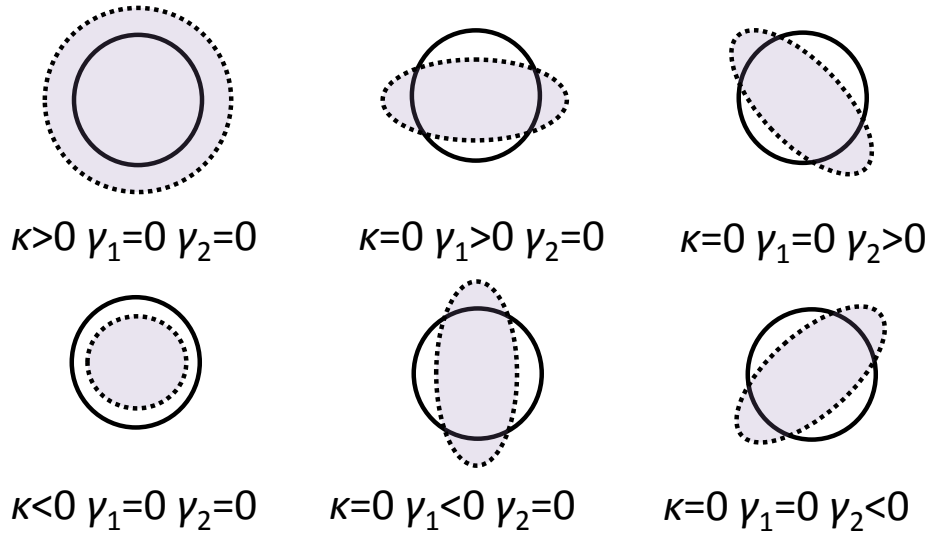


Figure 2.2: Graphical illustration of shear and convergence. The black circles represent an unlensed source, while the shaded ellipses are the same source under various types of distortion.

The schematic diagram in Figure 2.2 gives a qualitative example of how a circular galaxy's image responds under shear and magnification. The factor by which an image is expanded due to the latter effect is given by the inverse determinant of the distortion matrix:

$$\mu = |\mathbf{A}|^{-1} = \frac{1}{(1 - \kappa)^2 - \gamma^2}. \quad (2.16)$$

In general this is very difficult to infer from measurements, since one cannot know *a priori* the inherent size of the galaxy before lensing.

Finally, by evaluating each element in the distortion matrix for a lens comprising two masses

using equation 2.5, and taking the trace-free part, it is possible to derive a general expression for γ . Most commonly this is stated in the form

$$\gamma = \pi^{-1} \int \kappa(\boldsymbol{\theta}') \mathcal{D}(\boldsymbol{\theta} - \boldsymbol{\theta}') d^2\theta', \quad (2.17)$$

with the complex quantity

$$\mathcal{D}(\boldsymbol{\theta}) \equiv \frac{y^2 - x^2 - 2ixy}{|\boldsymbol{\theta}|^4}. \quad (2.18)$$

In practice, the convergence can be recovered from shear measurements via a process called the Kaiser Squires Inversion (Kaiser & Squires, 1993). This method relies on the fact that the convolution here reduces to a product in Fourier space. With this in mind, one can see that reconstructing a spatial convergence map is just a case of taking the inverse Fourier transform of the product $\kappa(\boldsymbol{\theta}') \mathcal{D}(\boldsymbol{\theta} - \boldsymbol{\theta}')^2$.

2.1.3 Cosmology with Cosmic Shear

Effective Convergence

All of the above is valid in the case of a thin lens at fixed distance. In reality, however, the Universe does not contain one foreground lens sheet that is responsible for all observed lensing. Rather, it contains a continuous distribution of mass out to the horizon, which lenses an ensemble of galaxies to different degrees. It is also worth noting here that though the lensing field is just a mathematical quantity tracing the large scale mass distribution, it is only possible to measure it at specific points where galaxies happen to sit. With this in mind we will define a new quantity, which we'll call the *effective convergence* as a weighted sum over the redshift distribution of galaxies observed along the line of sight:

$$\kappa_{\text{eff}}(z_L) = \sum_S \kappa_S = \int_{z_L}^{\infty} \kappa_s p(z_s) dz_s. \quad (2.19)$$

This is the lensing convergence associated with a thin sheet of mass at redshift z_L , and includes contributions from all of the observed sources behind that sheet. Note that we've used the lower

²There is a small practical point we should note here: the integral diverges as $\ell \rightarrow 0$, meaning it is necessary to impose some finite lower bound ℓ_{min} . This leads to an additive term, which in general is unknown. This effect, known as mass sheet degeneracy, is problematic for reconstruction of individual cluster lens masses. In the weak lensing regime for the purpose of constructing mass maps, however, the low ℓ modes simply introduces a unknown flat normalisation.

case s to denote continuous quantities relating to sources in general, and upper case S to indicate one specific source. It follows from the definition of convergence that

$$\kappa_{\text{eff}}(\boldsymbol{\theta}, z_L) = \frac{4\pi G D_L}{c^2} \Sigma(\boldsymbol{\theta}, z_L) \int_{z_L}^{\infty} \frac{D_{LS}(z_s)}{D_S(z_s)} p(z_s) dz_s. \quad (2.20)$$




Figure 2.3: A cartoon diagram of background shear by large scale structure. Light from source galaxies (red ellipses) are lensed a series of effective lens slices along the line of sight. The shaded purple pane shows an example of a single thin slice of the dark matter distribution at redshift z_L , which induces distortions in all galaxies observed at $z_S > z_L$. The left-hand axes (the blue curve) show the observed source density $p(z)$, which falls away to zero at high redshift due to the observer's finite flux sensitivity.

As we said, any line of sight will pass through a spatially fluctuating mass density field; we can think of this as a multitude of thin lenses stacked in series. The total convergence accrued along the line of sight is thus the sum of κ_{eff} contributions. This time there is no need to weight by the galaxy distribution, because the background dark matter is responsible for lensing, not only the visible matter. One finds:




Figure 2.4: An example of a set of lensing efficiency kernels for realistic redshift distributions. The purple (solid) curves show the estimated photometric redshift distributions for the Year 1 Dark Energy Survey IM3SHAPE catalogue of Zuntz et al. (2017) in the four tomographic bins used in the Y1 cosmic shear analysis of Troxel et al. (2017). The blue (dotted lines) in each panel show the corresponding lensing efficiency $g^i(z)$, normalised to unity over the redshift range shown. As can be seen here lensing correlations in each bin probe slightly different redshift ranges, illustrating how cosmic shear can be used to explore the expansion history of the Universe.

$$\kappa_{\text{eff}}(\boldsymbol{\theta}) = \int_0^{\chi_{\text{hor}}} g(\chi)(\delta(\boldsymbol{\chi}) + 1)d\chi, \quad (2.21)$$

with

$$g(\chi) = \frac{3}{2} \frac{H_0^2}{c^2 a(\chi)} \Omega_m \int_{\chi}^{\chi_{\text{hor}}} p(\chi') \frac{S_K(\chi) S_K(\chi' - \chi)}{S_K(\chi')} d\chi' \quad (2.22)$$

where Σ has been rewritten in terms of comoving density, and we make use of equations 1.14 and 1.9. The kernel $g(\chi)$, which is often called the *lensing efficiency*, is defined for a specific set of source galaxies and absorbs the geometry of the lens system. As a function of redshift, it defines a window onto the matter power spectrum and thus the epoch to which lensing is most sensitive. As a rule of thumb $g(\chi)$ peaks approximately midway between the observer and the maximum redshift of the galaxy sample (though this is only exactly true if the observed source distribution is a delta function at χ_{max}). Figure 2.4 illustrates the relationship between redshift distribution and lensing efficiencies for a set of real tomographic redshift measurements from DES Y1.

For clarity the “effective” subscript will be dropped hereafter, and we’ll assume that if a specific lens redshift isn’t given then the convergence is an effective measured convergence over an ensemble of foreground lenses.

The Equivalence of the Convergence & Shear Power Spectra

One non-trivial but key feature of equation 2.17 is that the two-point statistics of the shear and convergence can be treated interchangeably. This can be demonstrated as follows.

We will make a brief diversion here into the general properties of fields, and define some statistics which will be used heavily in the rest of this thesis. First consider a Gaussian random field. It could be scalar or complex (convergence and shear), and it can be two or three dimensional. Call the field α and the n D position coordinate \mathbf{x} . The two-point correlation between any two points is then

$$C(X) = \langle \alpha(\mathbf{x})\alpha^*(\mathbf{x}') \rangle, \quad (2.23)$$

where the angular brackets denote averaging over pairs of points with the same separation $X = |\mathbf{x}^2 - \mathbf{x}'^2|^{\frac{1}{2}}$. The above assumes isotropy in α , which requires any moments of the field to be uniform in space (and so C is dependent only on the separation of the points, not their exact positions or orientation).

If we wished to describe the field in Fourier space we could perform an n D transform on each of the α terms before correlating the result:

$$\langle \tilde{\alpha}(k)\tilde{\alpha}^*(k') \rangle = \int \int e^{i\mathbf{x}\cdot\mathbf{k}} e^{-i\mathbf{x}'\cdot\mathbf{k}'} \langle \alpha(\mathbf{x})\alpha^*(\mathbf{x}') \rangle d^n x d^n x'. \quad (2.24)$$

By using $\mathbf{X} = \mathbf{x} - \mathbf{x}'$, and $d\mathbf{X} = d\mathbf{x}'$ the above becomes

$$\langle \tilde{\alpha}(k)\tilde{\alpha}^*(k') \rangle = \int \int e^{i\mathbf{x}\cdot(\mathbf{k}-\mathbf{k}')} e^{i\mathbf{X}\cdot\mathbf{k}'} \langle \alpha(\mathbf{x})\alpha^*(\mathbf{x}') \rangle d^n x d^n X. \quad (2.25)$$

Each of the integrals reduces to triviality under scrutiny. The first, with respect to x , simply yields a Dirac delta function multiplied by a factor of $(2\pi)^n$, under the assumption that the k modes are uncorrelated on different scales. It then becomes apparent that

$$(2\pi)^n \delta_D(k - k') P(k) = \langle \tilde{\alpha}(k)\tilde{\alpha}^*(k') \rangle, \quad (2.26)$$

where we've defined the *power spectrum* of an arbitrary field α as the Fourier transform of the real (i.e. not harmonic) space correlation $P(k) = \int e^{i\mathbf{X}\cdot\mathbf{k}} C(X) d^n X$.

Now we'll return to the specific case of lensing. Consider equations 2.17 and 2.18 again. We can make two elementary observations here. First, equations 2.17 is effectively a convolution, in which the real space shear field smoothed with a spatial filter $\mathcal{D}(\theta)$. If the operation is mapped into

2D multipole space, the convolution becomes a product,

$$\tilde{\gamma}(\boldsymbol{\ell}) = \frac{1}{\pi} \tilde{\kappa}(\boldsymbol{\ell}) \times \tilde{\mathcal{D}}(\boldsymbol{\ell}), \quad (2.27)$$

and under Fourier transform the filter in equation 2.18 takes the form

$$\tilde{\mathcal{D}}(\boldsymbol{\ell}) = \frac{\ell_1^2 - \ell_2^2 + 2i\ell_1\ell_2}{\ell^2}, \quad (2.28)$$

where $\boldsymbol{\ell} = (\ell_1, \ell_2)$ is the Fourier pair to the spatial coordinates $\boldsymbol{\theta} = (x, y)$. Taking the product at two multipoles one finds

$$\tilde{\mathcal{D}}(\boldsymbol{\ell})\tilde{\mathcal{D}}^*(\boldsymbol{\ell}') = \frac{1}{\ell^2} \frac{1}{\ell'^2} \left[(\ell_1^2 - \ell_2^2)(\ell_1'^2 - \ell_2'^2) + 4\ell_1\ell_1'\ell_2\ell_2' + 2i\ell_1'\ell_2(\ell_1^2 - \ell_2^2) - 2i\ell_1\ell_2(\ell_1'^2 - \ell_2'^2) \right]. \quad (2.29)$$

Now using the definition of the two dimensional shear power spectrum from equation 2.26,

$$(2\pi)^2 \delta_D^2(\boldsymbol{\ell} - \boldsymbol{\ell}') C_{\gamma\gamma}(\boldsymbol{\ell}) = \left\langle \tilde{\kappa}(\boldsymbol{\ell}) \tilde{\mathcal{D}}(\boldsymbol{\ell}) \tilde{\kappa}^*(\boldsymbol{\ell}') \tilde{\mathcal{D}}^*(\boldsymbol{\ell}') \right\rangle, \quad (2.30)$$

into which we can substitute in equation 2.29, whereupon the result simplifies considerably by virtue of the delta function. In fact, the product of $\mathcal{D}(\boldsymbol{\ell})\mathcal{D}^*(\boldsymbol{\ell}')$ goes to unity when we enforce the equality $\boldsymbol{\ell} = \boldsymbol{\ell}'$. The result is

$$(2\pi)^2 \delta_D^2(\boldsymbol{\ell} - \boldsymbol{\ell}') C_{\gamma\gamma}(\boldsymbol{\ell}) = \langle \tilde{\kappa}(\boldsymbol{\ell}) \tilde{\kappa}(\boldsymbol{\ell}') \rangle, \quad (2.31)$$

and so from the definition of the convergence spectrum,

$$C_{\gamma\gamma}(\boldsymbol{\ell}) = C_{\kappa\kappa}(\boldsymbol{\ell}). \quad (2.32)$$

Relating Convergence to the Matter Power Spectrum

For reasons that will become apparent we will now write out explicitly the definition of the Fourier space pair of the convergence in terms of angular multipoles ℓ :

$$\tilde{\kappa}(\boldsymbol{\ell}) = \int e^{i\boldsymbol{\ell}\cdot\boldsymbol{\theta}} \kappa(\boldsymbol{\theta}) d^2\theta, \quad (2.33)$$

where the tilde denotes the two dimensional Fourier transform. The equivalent (inverse) transform

for the overdensity is

$$\delta(\boldsymbol{\chi}) = \frac{1}{(2\pi)^3} \int e^{-i\mathbf{k}\cdot\boldsymbol{\chi}} \tilde{\delta}(\mathbf{k}) d^3k, \quad (2.34)$$

or if we explicitly separate out 2D modes perpendicular to the line of sight and those parallel to it we can write

$$\delta(S_K(\chi)\boldsymbol{\theta}, \chi) = \frac{1}{(2\pi)^3} \int e^{-i\mathbf{k}_\perp \cdot \boldsymbol{\theta} S_K(\chi)} e^{-ik_0\chi} \tilde{\delta}(\mathbf{k}) d^3k. \quad (2.35)$$

We note that the position argument has not changed, but is simply written in terms of parallel and perpendicular components, $\boldsymbol{\chi} = (S_K(\chi)\boldsymbol{\theta}, \chi)$. Next, let us define two populations of source galaxies i and j , each of which can be used to measure a convergence. Substituting equations 2.35 and 2.21 into 2.33 then gives us

$$\begin{aligned} \langle \tilde{\kappa}^i(\boldsymbol{\ell}) \tilde{\kappa}^j(\boldsymbol{\ell}') \rangle &= \int \int \int_0^{\chi_{\text{hor}}} \int_0^{\chi_{\text{hor}}} \int (2\pi)^{-3} \int (2\pi)^{-3} g^i(\chi) g^j(\chi) \langle \tilde{\delta}(\mathbf{k}, \chi) \tilde{\delta}^*(\mathbf{k}, \chi) \rangle \\ &\times e^{-i(\mathbf{k}_\perp S_K(\chi) - \boldsymbol{\ell}) \cdot \boldsymbol{\theta}} e^{-i(\mathbf{k}'_\perp S_K(\chi') - \boldsymbol{\ell}') \cdot \boldsymbol{\theta}'} e^{-i(k_0\chi - k'_0\chi')} d^3k' d^3k d\chi' d\chi d^2\theta' d^2\theta. \end{aligned} \quad (2.36)$$

The definition of the power spectrum then eliminates one of the integrals

$$\begin{aligned} \langle \tilde{\kappa}^i(\boldsymbol{\ell}) \tilde{\kappa}^j(\boldsymbol{\ell}') \rangle &= \int \int \int_0^{\chi_{\text{hor}}} \int_0^{\chi_{\text{hor}}} \int (2\pi)^{-3} g^i(\chi) g^j(\chi) P_\delta(k, \bar{\chi}) \\ &\times e^{-i(\mathbf{k}_\perp S_K(\chi) - \boldsymbol{\ell}) \cdot \boldsymbol{\theta}} e^{-i(\mathbf{k}_\perp S_K(\chi') - \boldsymbol{\ell}') \cdot \boldsymbol{\theta}'} e^{-ik_0(\chi - \chi')} d^3k d\chi' d\chi d^2\theta' d^2\theta, \end{aligned} \quad (2.37)$$

where the correlation is actually between two points at different distances and so the power spectrum is evaluated an intermediate distance $\bar{\chi}$ (the exact value isn't terribly important for reasons we'll come back to). Now the integrals with regard to angular coordinates revert to two dimensional Dirac delta functions of the form $\delta_D^2(S_K(\chi)\mathbf{k}_\perp - \boldsymbol{\ell})$. If we separate the 3D k mode integral into parts which are perpendicular and parallel to the line of sight $d^3k = d^2k_\perp dk_0$, after some delta function manipulation we arrive at the expression

$$\begin{aligned} \langle \tilde{\kappa}^i(\boldsymbol{\ell}) \tilde{\kappa}^j(\boldsymbol{\ell}') \rangle &= 2\pi \int_0^{\chi_{\text{hor}}} \int_0^{\chi_{\text{hor}}} \int g^i(\chi) g^j(\chi) S_K^{-2}(\chi) P_\delta(k, \bar{\chi}) \\ &\quad \times \delta_D^2 \left(\boldsymbol{\ell}' - \boldsymbol{\ell} \frac{S_K(\chi')}{S_K(\chi)} \right) e^{-ik_0(\chi - \chi')} dk_0 d\chi' d\chi. \end{aligned} \quad (2.38)$$

Finally the integral with respect to k_0 yields another Dirac delta function (thanks to the last exponential), and from there the $d\chi'$ integration simply becomes a statement of equality $\chi = \chi'$ (and so $\bar{\chi}$ becomes superfluous). We end up with an expression for the convergence correlation of the form

$$\langle \tilde{\kappa}^i(\boldsymbol{\ell}) \tilde{\kappa}^j(\boldsymbol{\ell}') \rangle = (2\pi)^2 \delta_D^2(\boldsymbol{\ell}' - \boldsymbol{\ell}) \int_0^{\chi_{\text{hor}}} \frac{g^i(\chi) g^j(\chi')}{S_K^2(\chi)} P_\delta \left(\frac{\ell}{S_K(\chi)}, \chi \right) d\chi \quad (2.39)$$

or

$$C_{\gamma\gamma}^{ij}(\ell) = \int_0^{\chi_{\text{hor}}} \frac{g^i(\chi) g^j(\chi')}{S_K^2(\chi)} P_\delta \left(k = \frac{\ell}{S_K(\chi)}, \chi \right) d\chi. \quad (2.40)$$

This simple expression, known as the *Limber equation*, illustrates why cosmic shear is so potent as a cosmological tool; the basic lensing observable (the two-point function of galaxy shear) accesses the power spectrum via a pair of kernels which are sensitive to the redshift distribution of sources and a ratio of angular diameter distances. All of these parts are in principle very well understood (although see Section 2.2.4). Clearly the amplitude of $C_{\gamma\gamma}$ responds directly to σ_8^2 and to Ω_m (or $\Omega_m h^2$ via the efficiency kernel), and so the two exhibit a characteristic degeneracy.

2.2 Cosmic Shear in Practice: Observables

The era of large scale cosmological surveys has necessitated collaboration between experts in various sub-fields. Providing all of the elements for cosmic shear to the required accuracy often needs an in depth understanding of measurement algorithms and their systematics and uncertainties. In general it is no longer feasible for any individual or body of code to handle every aspect. Typically a survey will have multiple (sometimes disjoint) working groups, which each gather the experts and algorithms required. In recent years, then, it has become meaningful to talk of a “shear pipeline” tying together the various ingredients to go from photometric pixel fluxes to constraints on cosmology. The following paragraphs provide a short review of the essential components of such a pipeline.

2.2.1 The Point Spread Function

In a dataset like DES Y1 it is inevitable that there will be degradation in the quality of an image due to the atmosphere and the optics of the telescope. The net effect of stochastic atmospheric distortions is to blur the image, spreading the flux of a point source across several pixels. This is equivalent to a convolution with a kernel known as the *point spread function* (PSF). If the PSF is approximately circular, this tends to make a galaxy appear slightly larger and rounder than it would be otherwise. Unfortunately in ground-based measurements it is often the case that the PSF is comparable in size to the underlying galaxy. It is thus essential that we devise an accurate way to measure and remove the impact if we are to accurately recover the galaxy's shape.

The most common method for PSF estimation is to use a sample of bright local stars, which can be assumed to be point sources to good approximation. This provides a window on to the PSF at particular fixed points in the image, which must then be mapped onto the positions of galaxies. Typically one would use a polynomial basis set on the coordinates of the image to quantify spatial variations. A number of algorithms have been built to perform that task to varying degrees of accuracy (Jarvis et al., 2008; Bertin, 2011; Chang et al., 2012). If one has a PSF reconstruction for a particular galaxy it can then be folded into the shape measurement, either by convolving the model during the forward modelling process to construct mock galaxy images, or else by deconvolving the real image prior to measurement.

Finally it is worth noting here that a selection of measurement algorithms explicitly compensate for anisotropy in the PSF using a process known as “Gaussianisation”, whereby the image is convolved with an additional kernel in order to symmetrise the PSF distortion. We do not use these methods in detail, but refer the reader for Kaiser et al. (1995), Hirata & Seljak (2003), Massey et al. (2007b) and Herbonnet et al. (2017) for examples.

2.2.2 Galaxy Shapes

With the PSF estimates in hand, the next task is to compile a catalogue of reliable galaxy shapes. This is a highly non-trivial technical challenge and, as data volumes steadily grow, can require significant computational resources. The following sections provide a brief overview of the basic observables of shape measurement and their relation to cosmic shear.

Ellipticity Definitions

The derivations above give us the building blocks for cosmic shear: the basic quantities of convergence, deflection angle and shear, which pertain to the physical (but generally unseen) properties of a lens system. We have seen how these quantities are links to the underlying properties of the Universe. The next inevitable question must be how can we go about *measuring* these basic quantities. The data emerging from our telescopes is not a 3D grid of shear or convergence; rather it is a set of noisy, pixelised photometric fluxes.

To begin, imagine we have a cutout galaxy image, consisting of a smooth light profile, which varies with position in the image $I(\boldsymbol{\theta})$. We can define unweighted first and second moments as 2D integrals over position :

$$\mu_i = \frac{\int I(\boldsymbol{\theta})\theta_i d^2\theta}{\int I(\boldsymbol{\theta})d^2\theta}, \quad (2.41)$$

and

$$Q_{ij} = \frac{\int I(\boldsymbol{\theta})(\theta_i - \mu_i)(\theta_j - \mu_j)d^2\theta}{\int I(\boldsymbol{\theta})d^2\theta}, \quad (2.42)$$

where the 2D position $\boldsymbol{\theta} = (x, y)$ is given relative to a centroid position, defined by the first moment. The second moments, or the quadrupole moments of the light profile, form the basic building blocks of galaxy shape measurement. One can construct numerical shape estimators by combining these moments. For example two common quantities, both loosely termed “ellipticity”, are defined as the *distortion*

$$\epsilon \equiv \frac{Q_{xx} - Q_{yy} + 2iQ_{xy}}{Q_{xx} + Q_{yy} + \sqrt{Q_{yy}Q_{xx} - 2Q_{xy}^2}} \quad (2.43)$$

and *polarisation*

$$\chi \equiv \frac{Q_{xx} - Q_{yy} + 2iQ_{xy}}{Q_{xx} + Q_{yy}}. \quad (2.44)$$

Due to their respective mathematical properties each of these quantities has its advantages. In recent years equation 2.43 has become the favoured definition in many circles, for reasons that we will return to below. We show an example of the shear and distortion for a random set of galaxies from DES in Figure 2.5. In the sections of this thesis beyond the present discussion we’ll simply refer to a galaxy’s ellipticity e , meaning the first of these definitions.




Figure 2.5: Illustrative example of the difference between the two ellipticity definitions described in the text. **Top:** the distribution of a random subsample of 10,000 galaxies from the DES Y1 IM3SHAPE catalogue in the e_1, e_2 ellipticity plane as defined by each of the definitions. The black dotted line shows the unit ellipticity circle in each case. **Bottom:** The one dimensional distribution of ellipticity magnitude $e = \sqrt{e_1^2 + e_2^2}$ of the same galaxies using each definition.

Returning to the expression for the distortion matrix in equation 2.14, we note that it can be written,

$$\mathbf{A} = (1 - \kappa) \begin{pmatrix} 1 + g_1 & -g_2 \\ -g_2 & 1 - g_1 \end{pmatrix}, \quad (2.45)$$

where we've defined

$$g \equiv \frac{\gamma}{1 - \kappa}. \quad (2.46)$$

The quantity defined in equation 2.46 is called *reduced shear*. Since in general it is possible to measure a coherent elongation from a population of galaxies but not an analogous change in surface area, in practice we can typically recover \mathbf{A} only up to a multiplicative amplitude. Thus the reduced shear g , and not the shear γ is the observable quantity. In the context of cosmic shear γ and κ are each $\mathcal{O}(10^{-2})$. For all measurements to date it is safe to ignore the factor of $1 - \kappa$ and treat γ and g as interchangeable.

If we define two sets of galaxy properties, in the (unlensed) source plane and in the observed image plane, the positions transform as $\boldsymbol{\theta}^{\text{src}} = \mathbf{A}\boldsymbol{\theta}^{\text{img}}$, and so equation 2.43 gives us (Seitz & Schneider, 1997):

$$\epsilon^{\text{img}} = \frac{\epsilon^{\text{src}} + g}{1 + \epsilon^{\text{src}} g^*}. \quad (2.47)$$

It is generally, then, reasonable to treat the measured shape of a galaxy as the sum of an inherent part ϵ^{src} , which would be measured without lensing, and a cosmological shear g . Under the assumption that the ellipticities of galaxies in the source plane are randomly oriented, clearly the mean ϵ over a large number of galaxies gives an unbiased estimate for the mean shear $\langle g \rangle$. Unfortunately this is not true of the alternative ellipticity definition, and the equivalent calculation yields $\langle \chi^{\text{img}} \rangle = 2/(1 - \langle |\chi^{\text{src}}|^2 \rangle) \times \langle g \rangle$. That is, accurate recovery of the cosmological shear using χ requires prior knowledge of the unlensed ellipticity distribution.

Measurement Algorithms

Modern shape measurement techniques can be divided, broadly, into two classes: *moments-based* and *model fitting* methods. The former were well represented in the earliest generations of shear algorithm (e.g. KSB, Reglens, Ellipto, RRG; see Kaiser et al. 1995; Heymans et al. 2006), and there are a number of successor codes in use today (Bernstein & Armstrong, 2014; Herbonnet

et al., 2017). The key idea is to measure the quadrupole moments of each galaxy image, and then fit these measurements together to construct an estimator for ellipticity. A related technique, referred to as shapelets, involves reconstructing a light profile using a finite sum of polynomial basis functions. In general, although moments-based methods are agnostic to galaxy type in the sense of not requiring a specific analytic model, they do require radial weighting in order for the integrals in equation 2.42 to converge.

Model fitting methods take a more direct route, using galaxy ellipticities as the basic unit of observation. The process typically involves a forward modelling loop, wherein a galaxy model is repeatedly simulated with trial parameter values. The prediction is compared with the data in order to map out the likelihood of the model parameters. Such techniques have become increasingly sophisticated over the decades, and are often now favoured over more traditional moments-based approaches. It is worth pointing out that all of the major shear cosmology studies of the Twenty First Century have adopted such methods, notably IM3SHAPE and NGMIX, which were used in DES SV and DES Y1 (Jarvis et al., 2015; Zuntz et al., 2017) and LENSFIT, which was used in CFHTLenS and KiDS (Miller et al., 2013; Fenech Conti et al., 2016).

2.2.3 Estimators of Cosmic Shear

Two-Point Statistics

Implicit in the discussion of Section 2.1 is that on large scales the one-point shear statistics should approach zero³. In practice it is necessary to use statistics of second order or higher in the shear. One practically useful two-point shear statistic is known as the correlation function:

$$\xi_{\pm}^{ij}(\theta) \equiv \langle \gamma_{+}^i \gamma_{+}^j \rangle_{\theta} \pm \langle \gamma_{\times}^i \gamma_{\times}^j \rangle_{\theta}, \quad (2.48)$$

defined for two populations of galaxies (i, j) , and with the angular brackets denoting averaging over galaxy pairs separated by a particular angular distance θ . The subscripts $(+, \times)$ indicate the orthogonal components of shear rotated into coordinate axes defined by the separation vector between the galaxies:

$$(\gamma_{+}, \gamma_{\times}) = (-(\gamma_1 \cos 2\phi + \gamma_2 \sin 2\phi), -(\gamma_2 \cos 2\phi - \gamma_1 \sin 2\phi)). \quad (2.49)$$

³If this is not immediately obvious consider equation 2.4. If the density of dark matter is a homogeneous Gaussian random field, then each slice along the line of sight is essentially a randomly distributed collection of point sources. Over many positions in the lens plane, then, positive and negative θ values will average to an expectation of zero.

By noting the equivalence of $\gamma = \gamma_+ + i\gamma_\times = \gamma_1 + i\gamma_2$, and writing the shears in terms of their inverse Fourier transforms it can be shown that

$$\xi_{\pm}^{ij}(\theta) = \int_0^{\infty} \frac{\ell}{2\pi} J_{0/4}(\ell\theta) C_{\gamma\gamma}^{ij}(\ell) d\ell, \quad (2.50)$$

where the Bessel function of the first kind,

$$J_n(x) = \frac{1}{i^n \pi} \int_0^{\pi} e^{ix \cos \alpha} \cos(n\alpha) d\alpha, \quad (2.51)$$

serves as an effective window function onto the angular shear spectrum. The two populations referred to here are most commonly galaxies in different photometric redshift bins. By performing coarse binning along the line of sight, one can use the fact that the lensing efficiency kernels in equation 2.40 are sensitive to the redshift distribution of the source galaxies to probe the power spectrum at a particular epoch.

Various alternative two-point statistics have been devised for addressing specific perceived deficiencies in the correlation functions. The Bessel function kernels are relatively wide, which helps to maximise the signal-to-noise of the measurement, but has the drawback of mixing power from different ℓ modes. For example there are various public codes for measuring the angular spectra directly from individual galaxy shapes. Though $C(\ell)$ estimators are closer to the theory and thus avoid the complication of the numerical Hankel transform above, they do not naturally factor in the spatial configuration of the survey masks.

The standard practical estimator for ξ_{\pm} (see Dark Energy Survey Collaboration 2016; Hildebrandt et al. 2016) can be constructed as

$$\xi_{\pm}^{ij}(\theta) = \frac{\sum_a \sum_b w_a^i w_b^j (e_+^i(\theta_a) e_+^j(\theta_b) \pm e_\times^i(\theta_a) e_\times^j(\theta_b))}{\sum_a \sum_b w_a w_b}, \quad (2.52)$$

with the sum being over galaxies within a finite bin of angular separation $\theta_{ab} \in [\theta - \Delta\theta, \theta + \Delta\theta]$. The galaxy weights w_a are typically provided by the shape measurement algorithm, and are intended as a measure of the quality of the ellipticity measurement for each galaxy.

E-Modes & B-Modes

It is often useful to decompose a measured shear in pairs of galaxies into a tangential part, which might be expected from lensing, and a cross part, which would not. A convenient way to do this is to define a vector field $\mathbf{u}(\boldsymbol{\theta})$, which is defined as the 2D gradient of the convergence field $\nabla\kappa$

(Schneider et al., 2002). From 2.10 it is then easy to show the new field in terms of the derivatives of shear

$$\mathbf{u} = \begin{pmatrix} \frac{\partial\gamma_1}{\partial x} + \frac{\partial\gamma_2}{\partial y} \\ \frac{\partial\gamma_2}{\partial x} - \frac{\partial\gamma_1}{\partial y} \end{pmatrix}. \quad (2.53)$$

If κ is a pure scalar field then one finds:

$$\nabla \cdot \mathbf{u} = \nabla^2 \kappa \quad \nabla \times \mathbf{u} = \nabla \times (\nabla \kappa) = 0 \quad (2.54)$$

In practice, however, if one attempts to reconstruct \mathbf{u} using measured shears the result will not always follow the second equality. Astrophysical systematics and foreground instrumental effects can induce a spurious shear component, which in this formalism translates into an imaginary convergence component. In general, then, one can recast the convergence as a complex quantity, consisting of a dominant (real) term κ^E and an imaginary κ^B , such that $\nabla \times \mathbf{u} = \nabla^2 \kappa^B$.

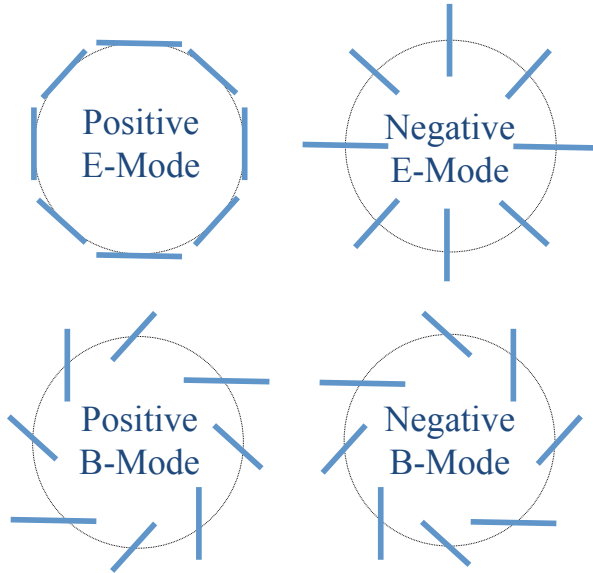


Figure 2.6: Stick diagram of the impact on a round source galaxy of E- and B-mode distortions. Positive E-modes are caused by physical mass overdensities, which lens background galaxies along tangential arcs about the mass centre.

This decomposition $\kappa = \kappa^E + i\kappa^B$, is equivalent to separating the parts of the convergence that cause tangential shear-like (cruciform) and X-like (rotated cruciform) distortions. These two forms of distortion are known as *E-modes* and *B-modes*, and are characted in Figure 2.6.

This additional signal is expected to enter the observed shear correlations in equation 2.50 additively:

$$\xi_{\pm}(\theta) = \int_0^{\infty} \frac{\ell}{2\pi} J_{0/4}(\ell\theta) (C_{EE}(\ell) + C_{BB}(\ell)) d\ell, \quad (2.55)$$

where the two additional terms are the angular power spectra of κ^E and κ^B . In practice attempting to detect a B-mode signal is used as a null test to ensure a shear measurement is free of systematics (see e.g. Becker et al. 2015, Hildebrandt et al. 2016, Jee et al. 2016). Spurious (though generally small) B-modes can be generated by local interactions between galaxies and foreground source clustering, as well as instrumental effects.

2.2.4 Photometric Redshifts

It has been shown elsewhere that dividing galaxies into redshift bins and utilising the extra information in the auto- and cross-correlations is extremely helpful for the cosmological constraining power of a set of shear measurements (Hu, 1999). Unfortunately, however, the photometric image of a galaxy does not trivially demark it in redshift. Instead, an observer must make what inferences she can with the often crude photometry available.

The most accurate method for estimating redshift uses high resolution spectroscopy. If known emission lines can be identified, measurement of the offset from the known rest-frame frequency gives one a precise z measure. Unfortunately spectroscopic redshift estimation requires specialised equipment and long exposure times to reliably identify typically faint spectral lines. This quickly becomes unfeasible for lensing surveys, which must image hundreds of millions of galaxies in a timespan of months to years. Indeed, it is the small faint galaxies which carry the strongest cosmological signal and are thus most valuable for cosmic shear based studies.

A much faster and widely used alternative relies on the use of wide multi-wavelength filters (typically $ugrizY$, or some subset thereof). Pre-calculated spectral templates are fitted to flux in different bands to obtain a z estimate (photo- z). Such methods commonly rely on broad-band features in optical galaxy spectra such as the Lyman- α break. Although *much* faster than spec- z these techniques are often highly unreliable, susceptible to large scatter and systematic errors, and should be treated with caution.

Typically a photo- z code provides a conditional probability distribution, encoding how likely it is that a galaxy with observable properties α is at redshift z . The standard process for reconstructing the ensemble redshift distribution of population i is then to simply add the PDFs for each individual constituent galaxy j :

$$p^i(z) = \sum_j^{catalogue} p(z|\alpha_j) \approx \int p(z|\alpha)p(\alpha)d\alpha. \quad (2.56)$$

This process, known as stacking, effectively marginalises the observable properties of the individual galaxies. In the simplest case i refers to the entire catalogue of galaxy measurements from a survey, but more commonly it is a subset of objects defined by some observable property (e.g. using the mean of $p(z|\alpha)$ to sort galaxies into tomographic bins).

A middle path between photometric and spectroscopic redshift estimation is to use an instrument with superior multi-band photometry to obtain high quality photo- z . Examples of this include the HST GEMS 30-band photometry and the planned 40-band PAU survey (Castander, 2008). Clearly this capability must be built into the instrument during the design stages, but such measurements exist for a number of small fields (e.g. COSMOS), which have become a benchmark for validating the redshift estimates in larger overlapping surveys.

Many analyses, particularly those involving cosmic shear and related lensing measurements, are predicated on an ability to accurately reconstruct galaxy redshifts. Unconstrained redshift error can be a major systematic in our attempts to make inferences about cosmology using photometric data, and can be strongly degenerate with the cosmological parameters we are trying to constrain.

There has, as a result, been much recent focus on identifying and mitigating photo- z error. At the highest level it has become common practice to include some level of redundancy in cosmology analyses, with multiple independent photo- z codes used to verify the robustness of results (Hildebrandt et al., 2016; Dark Energy Survey Collaboration, 2016).

In the long term, there has been some effort to improve machine learning algorithms trained on spectroscopic fields (De Vicente et al., 2016; Hoyle, 2016; Cavuoti et al., 2017).

For more immediate needs there has been some discussion of cross calibration with spectroscopic or high-quality photometric samples as a means to constrain errors (Choi et al., 2016; Samuroff et al., 2017b), bypassing longer running development of photo- z techniques. Cross correlations with galaxy clustering measurements have also been mooted as an alternative means to constrain the redshift distribution of photometric galaxy catalogues (Newman, 2008; Ménard et al., 2013; Johnson et al., 2017).

2.2.5 Covariance Matrices

The final piece required before equation 1.28 can be used to evaluate the likelihood of a trial cosmology is a covariance matrix of the data. Clearly there is some uncertainty in the measured

datavector resulting from the limited statistical precision of the measurement. It is also possible that the data could have become contaminated by residual systematics in any of the previous steps in such a way that mimics lensing (see Section 2.3 below). In some cases we might, to the best of our knowledge, correct for a particular systematic adequately, but concede that there are gaps in our knowledge which *could* be leading us to over- or under-correct the data. The most common way to handle this is to incorporate a systematic uncertainty into the prior when we marginalise over the systematic. It is also, however, possible to include it as an additional contribution to the covariance matrix. Finally, there is the issue of sample variance. The measured shear could be a perfectly accurate representation of the actual lensing signal in that patch of sky, but still result in inaccuracies in the recovered cosmological parameters. Since the measurement encompasses a finite volume of the Universe, there is no guarantee that our line of sight does not, by chance, pass through a particular clumpy patch of the dark matter field. Known as *cosmic variance*, the size of this uncertainty clearly scales with the volume of space probed by the data.

Since we have only one dataset (and indeed one realisation of the Universe) calculating the relevant uncertainties is not simply a case of repeating the measurement and considering the spread of results. Faced with this task there are three broad routes one could take.

The first is to use the data itself by, for example, bootstrap resampling the survey to compute a variance between sub-patches.

The second is to use an analytic approximation. When the constraining power of the data is not great, it is common to assume $\delta(\chi)$ is Gaussian random on all scales, which affords a significant simplification in the required calculation. There are, however, an increasing number of implementations of more sophisticated halo-model calculations, which provide a far more realistic error estimate (e.g. Hildebrandt et al. 2016; Krause et al. 2017).

The last, most computationally expensive, route is to use numerical simulations. By assuming a fixed cosmology one can generate multiple universe realisations (or well separated chunks of a single simulation volume) to evaluate the uncertainty of a measurement. Each method has its limitations, and it is not unprecedented to use multiple methods as a test of the robustness of a set of results.

2.3 Biases in Cosmic Shear

There are a great many ways in which a measurement of cosmic shear can become biased. Given that it will be repeated many times though this thesis, it is worth taking a moment to define “bias”.

In the context of cosmological inference, a bias is any factor (other than parameters of interest) which ultimately influences the outcome of the measurement. Inevitably the data *will* be subject to such effects, and the task of the observer is to prevent them translating into cosmological biases. In the following paragraphs we introduce a number of known biases, known to afflict cosmic shear-based measurements.

2.3.1 Shear Measurement Bias

The most common parameterisation for inaccuracies in the shear recovered from an ensemble of galaxies was introduced by (Heymans et al., 2006) in the context of the STEP challenge. Since the cosmological signal in cosmic shear is small we can expand the residual between the true shear γ^{tr} (which we assume to be constant on a small patch of sky) and the measured shear $\tilde{\gamma}$. The result is:

$$\tilde{\gamma}_i = (1 + m_i)\gamma_i^{\text{tr}} + c_i, \quad (2.57)$$

where γ_i is the $i = (1, 2)$ component shear measured from an ensemble of galaxies, and γ_i^{tr} is the input or “true” value, which is assumed to be the same for all of these galaxies. Note that we assume here that there is no cross-talk between shear components (i.e. the bias in γ_1 is not sensitive to γ_2^{tr}). For completeness m_i and c_i could be replaced by 2×2 matrices, but in general there are no significant off-diagonal terms and most authors favour the simpler index notation.

It is also uncommon for there to be systematic differences between m_1 and m_2 in any region of parameter space, and typically one will assume a simple arithmetic mean $m = m_1 \approx m_2$. The additive terms are often dominated by the PSF-induced effects, whereby the PSF ellipticity at a particular position leaks into measured shear.

There are many ways bias can enter an ensemble shear estimate based on a population of galaxies, and this has been the subject of a decade or more of rigorous simulation-based tests (Heymans et al., 2006; Massey et al., 2007b; Bridle et al., 2010; Kitching et al., 2012; Mandelbaum et al., 2015). There is an extensive canon of literature on the nature and origins of such biases, and we will return to this topic in Chapters 3 and 4.

2.3.2 Intrinsic Alignments

In an idealised case the intrinsic ellipticities in equation 2.47 would average away to zero. While a finite number of galaxies will always produce a finite level of shape noise, this vanishes given a sufficiently large shear catalogue. Unfortunately, however, galaxy shapes are not uncorrelated

samples of the large scale shear field, but complex astrophysical objects embedded in a particular local environment. Specifically galaxies may share a common dark matter halo, the spin and the gravitational tidal field of which will induce local shape correlations. A slightly more pervasive effect arises from the fact that a mass concentration on the line of sight lenses background sources whilst simultaneously tidally distorting foreground galaxies. These astrophysical correlations are called *intrinsic alignments* (IAs). If we measure shape-shape correlations, based on equation 2.47 (assuming we are in the weak lensing regime) we find:

$$\begin{aligned}
\langle e^i(\theta)e(\theta + \theta')^j \rangle &= \langle \gamma^i(\theta)\gamma^j(\theta + \theta') \rangle \\
&+ \langle e^{\text{src},i}(\theta)\gamma^j(\theta + \theta') \rangle + \langle \gamma^i(\theta)e^{\text{src},j}(\theta + \theta') \rangle \\
&+ \langle e^{\text{src},i}(\theta)e^{\text{src},j}(\theta + \theta') \rangle \\
&= GG + GI + II,
\end{aligned} \tag{2.58}$$

where the angle brackets denote averaging over galaxy pairs i, j at positions θ and $\theta + \theta'$. The interpretation of correlations of measured galaxy shapes is no longer a pure measure of the cosmological shear. The additive terms, commonly called II and GI correlations, are typically subdominant to the cosmological GG signal, but sufficiently large that they will bias cosmic shear based inference if we neglect to account for them (Kirk et al., 2012; Krause et al., 2015).

2.3.3 Baryons & Modelling Uncertainties

Though often dominated by shear and photo- z inaccuracies and IA, a great many other effects can influence the result of a shear measurement. Many of the theoretical uncertainties have been shown using numerical forecasts to be acceptable for the current generation of surveys. The Limber approximation, the flat sky approximation and the convergence weighting (magnification bias) all fall under this category. Likewise, the modelling of baryons, which can shift power between scales in the dark matter power spectrum, is somewhat phenomenological and depends significantly on the specific models fed into hydrodynamical simulations. Fortunately many of these effects are also strongly scale dependent, and a large part of the mitigation strategy of past analyses has involved removing scales thought to be worst affected by such uncertainties.

2.3.4 Expectation Bias

There has been a gradual recognition in recent years of another, perhaps less tangible form of cosmological bias, which arises from the observer, rather than the objective measurement they are making. It has long been understood in other fields of science that expectation of a particular result can lead to unconscious bias. In simple terms, measurements which differ from an experimenter's preconceived notion of what the result *should* be will be tested more rigorously than those which seem to agree. This is known as *expectation bias*, or experimenter bias. Definite instances of this are difficult to identify, but there is clear potential, given the nature of observational cosmology, where results are released as a continuous series ever tighter constraints on a handful of basic parameters.

A case in point is the discrepancy between Planck and lensing datasets that have emerged in recent years. Each new analysis published on the topic may be internally robust, but if there is a tendency to search for systematics that drive new lensing results towards Planck (or vice versa) the final cosmology result may be biased. Likewise, analyses (including reusing public datasets) are more likely to be continued to completion and submitted for publication if they resolve known tensions. In recognition of this the most recent lensing cosmology analyses have included some form of blinding at the galaxy catalogue level (Dark Energy Survey Collaboration, 2016; Hildebrandt et al., 2016; Mandelbaum et al., 2017; Troxel et al., 2017). Additionally, and arguably equally important, there is a growing consensus in the lensing community that individual investigations should be blinded at the level of hiding cosmologically sensitive axes until the details of the analysis have been fixed.

Chapter 3

Measuring Cosmic Shear in Year 1 of the Dark Energy Survey

As discussed in the earlier chapters of this thesis, extracting an unbiased measurement of cosmic shear from pixelised galaxy images is a highly challenging task. This chapter presents a summary of work carried out within the weak lensing working group of the Dark Energy Survey with the aim of doing just that. Specifically we detail the process of measuring galaxy shapes using a forward modelling code, IM3SHAPE, and then calibrating and validating those measurements using simulations. The end result is a shape catalogue of sufficient quality for weak lensing cosmology, given the statistical power of DES Y1. The following forms part of the Y1 shear pipeline paper, which has been submitted for publication in Zuntz et al. (2017). In the following we will refer to this longer manuscript as Z17.

3.1 Introduction

Weak lensing, the deflection of light rays by wide-field matter, presents a powerful probe of cosmology and the laws of gravity. As discussed in earlier chapters, the deflection to which a bundle of photons is subjected is dependent on two factors: the geometry of the source-lens system relative to the observer, and the inherent lensing strength of the lens. On cosmological scales the former depends on the expansion history of the Universe via the redshift-distance relation. The latter depends on the level of structure in the Universe (the second moments of the cosmic density field), or alternatively the laws governing large scale gravity. These sensitivities allow lensing to impose limits on the history of the Universe, its underlying cosmological parameters and the equation of

state of dark energy.

The most direct way to measure weak lensing is via the ellipticity of distant galaxies. The net impact of intervening mass is to induce a shear, coherently stretching galaxy images on nearby lines of sight. The magnitude of this effect is only a few percent, (smaller than both the variance of intrinsic galaxy shapes and typical seeing), requiring us to construct vast catalogues of millions of galaxies to maximise statistical power. This chapter reports on weak lensing shape measurements from the first full year of data in a survey designed to meet these requirements.

The Dark Energy Survey (DES) is the largest extant lensing survey, and is part of the current generation of Stage III experiments (Albrecht et al., 2006). As discussed in the opening chapter, Stage III represents the current wave of ongoing surveys, and includes DES, KiDS and HSC. The final footprint of DES will cover 5000 square degrees of the southern sky, and at the time of writing, it has completed four out of its five planned observation seasons (discounting the first science verification phase). The catalogues described in the following chapter are based on images from the first of those years, a dataset covering 1500 square degrees. With a number of Y1 analyses now awaiting publication, work continues apace to extract shape measurements from the much larger Y3 dataset.

Building a catalogue of galaxy ellipticities from image data is a long process, at each stage of which one must pay careful attention to potential biases. For more details of the initial image processing steps see Z17. The first stage involves artefact detection and noise measurement. Composite images are built by coadding single exposures, and stars and galaxies are detected and classified on these so-called coadds. Astrometric transforms and the point spread function are derived from each single-epoch image. We collect single-epoch postage stamps about each detection in a coadd and stack them in a collection of binary tables called a Multi-Epoch Data Structure (MEDS) file (see Jarvis et al. 2015). Finally, we arrive at the shape measurement process itself, which forms the basis of this chapter.

Galaxy shape algorithms naturally fall into two categories. Any such measurement must account for non-cosmological distortions such as convolution by the PSF, which alter the apparent shapes of galaxies. *Forward modelling methods* are currently the most widely used class of technique, and involve generating parametric models of the underlying galaxy, and distorting them in a way designed to mimic instrumental effects. The model predictions are compared with the data in order to obtain a likelihood for the galaxy parameters. The second class are known as inverse methods and typically involve measuring second-order moments directly from the image data. Corrections are applied retrospectively to compensate for the effects of the observing process, and the extracted moments are used to calculate the ellipticity. This category includes many

early methods such as KSB (Kaiser et al., 1995) and Shapelets (Refregier, 2003). Within each class there are a great many qualitatively different algorithms with different assumptions. An advantage of model-fitting codes, which has contributed to their widespread use in practical shear studies, is that they return a likelihood for each fit. This provides a useful metric for assessing the quality of the shape estimate.

Inevitably all of these techniques will be subject to biases, which will be strongly dependent on the specifics of the implementation and the cuts applied. There are various ways in which one can seek to eliminate (or at least mitigate) such estimation biases. The simplest is to accept the existence of biases, and to constrain them by processing mock data with known input properties through the same pipeline as the real data. A number of early studies derived a single global calibration factor (Schraback et al., 2007; Jee et al., 2013) from simulations, but more recent studies have involved constructing per-galaxy corrections as a function of the measured properties (e.g. Jarvis et al. 2015 and Hildebrandt et al. 2016). This is the approach taken by the IM3SHAPE code in DES Y1.

Naturally, simulation-based calibration methods require the mock images to be carefully matched to the properties of the data. Any mismatch could result in inappropriate calibration factors, which would ultimately translate into a cosmological bias. Sparked by the computational and technical challenges of generating large highly accurate simulations, there has been a flurry of interest in recent years in alternative methods that reduce or remove the dependence on simulations. In Fenech Conti et al. (2016), the KiDS collaboration used a method they call self-calibration, in which a parametric copy of each object is generated from the best-fit model parameters. This new version is re-measured and used to correct the initial measurement on the data, a process which removes about half of the noise bias and reduces required simulation volumes. An alternative form of self-calibration is set out by Huff & Mandelbaum (2017) and Sheldon & Huff (2017). This technique, called metacalibration is implemented in the DES Y1 METACALIBRATION pipeline. Estimator biases are calibrated by applying an added shear to the real galaxy images and assessing its impact on galaxy measurement and selection. This has been demonstrated to be highly effective in simulated tests.

The shape measurement methodology in the DES Science Verification (SV) period was exhaustively detailed in Jarvis et al. (2015), hereafter J16. Many aspects of our methodology are identical to the equivalent in SV. We refer the reader to that paper for details of these processes, and focus here on the aspects that have been improved upon since SV.

This chapter is organised as follows: in Section 3.2 the observations upon which this work is based are described. This includes all steps from image reduction to PSF measurement; though a

large and time-consuming effort was expended in these early steps, the involvement of the author was limited. The current chapter is thus limited to a brief overview of these processes. In Section 3.3 we describe the construction of the Y1 IM3SHAPE catalogue, and Section 3.4 outlines the more delicate task of calibrating and testing it using simulations. Section 3.5 outlines the intended usage of the Y1 shape catalogues, including recommended systematic priors based on our current confidence in the calibration method. We conclude in Section 3.6.

3.2 Data

3.2.1 Observing Period and Conditions

The Dark Energy Survey Year One (DES Y1) catalogues are derived from image data from the Dark Energy Camera (DECam) (Flaugher et al., 2015), which is installed on the Blanco telescope at the Cerro Tololo Inter-American Observatory in Coquimbo, northern Chile. The data were collected during the first full season of DES operations, between 29 Aug 2013 and 9 Feb 2014 and targeted regions at $-60^\circ \lesssim \delta \lesssim -40^\circ$. The aim was to obtain 4 “tilings” per filter over this subset of the Y5 region, since covering the full area would reduce the average tiling depth to half of this, which would significantly hinder our ability to mask cosmic rays and per-exposure systematic errors. Due to variable atmospheric conditions the depth of the Y1 area is non-uniform. Figure 3.1 shows the footprint of the Y1 shape catalogues.

Compared to the SV catalogues described by J16, the Y1 shape catalogues represent a significant step forwards in area but with a lower integrated exposure time (up to 4×90 s exposures in *griz* vs 10 in SV). The quality of the Y1 imaging is superior to that taken in SV in several respects (see Z17 for a list of improvements), the cumulative result of which is a smaller and less elliptical PSF, and more uniform depth.

3.2.2 Fields in DES Y1

The DES Y1 footprint is shown in Figure 3.1, and covers a total of 1500 square degrees. Though less than half of the eventual 5000 square degrees in the DES Y5 footprint, this still represents the largest contiguous lensing survey to date. The area is divided into four areas, known as fields. The broad southern area, which accounts for the bulk of Y1 overlaps with the South Pole Telescope (SPT) CMB survey. This region was selected as the main Y1 science sample due to its large continuous area. The Science Verification dataset, which reached the full Y5 depth, was also taken

from ~ 130 square degree patch within the SPT field, known as SPT-East. The elongated strip along the top of the Y1 footprint is known as Stripe 82 due to its significant overlap with the SDSS Stripe 82 region. A collection of 10 supernova fields are also included in Y1. These are significantly deeper than the main survey. Along with the spectroscopic overlap fields, they were used for validation of photometric redshifts, but not included in our final shear catalogues.

3.2.3 The GOLD Catalogue & Star/Galaxy Separation

The initial selection of galaxies on which shape measurement was performed is detailed in Drlica-Wagner et al. (2017), and the selection described therein is referred to as the GOLD catalogue. The steps from raw images through image reduction, photometric calibration, and object detection to the catalogues are described in that paper. The star-galaxy separation process outlined by those authors is applied to the IM3SHAPE catalogue.

Objects detected in the Y1 area are identified as galaxies or stars in the GOLD catalogue using a classifier called MODEST. This quantity is based on the SEXTRACTOR SPREAD_MODEL variable, and is designed to discriminate between point sources and extended object. The IM3SHAPE selection flags used throughout this chapter and the next incorporate a star cut using the the MODEST classifier.

3.2.4 Blinding

In order to counter possible expectation bias, whereby participants of an analysis preferentially accept findings which agree with previous published results to be correct, the Y1 shape catalogues were blinded. The blinding operation transforms the ellipticity magnitude $|e|$ of each galaxy as $|\eta| \equiv 2 \operatorname{arctanh} |e| \rightarrow f|\eta|$, with a factor f assigned an unknown value drawn randomly from a uniform distribution with bounds $0.9 < f < 1.1$. This mapping preserves the fact that the raw measured e values are confined to a unit circle, while rescaling all inferred shears. All Y1 cosmology analyses which made use of the shape measurements were required to freeze their methodology before access to the unblinded catalogues was granted. Additional blinding measures were mandated by the lead authors of individual investigations, as appropriate for the analyses in question.

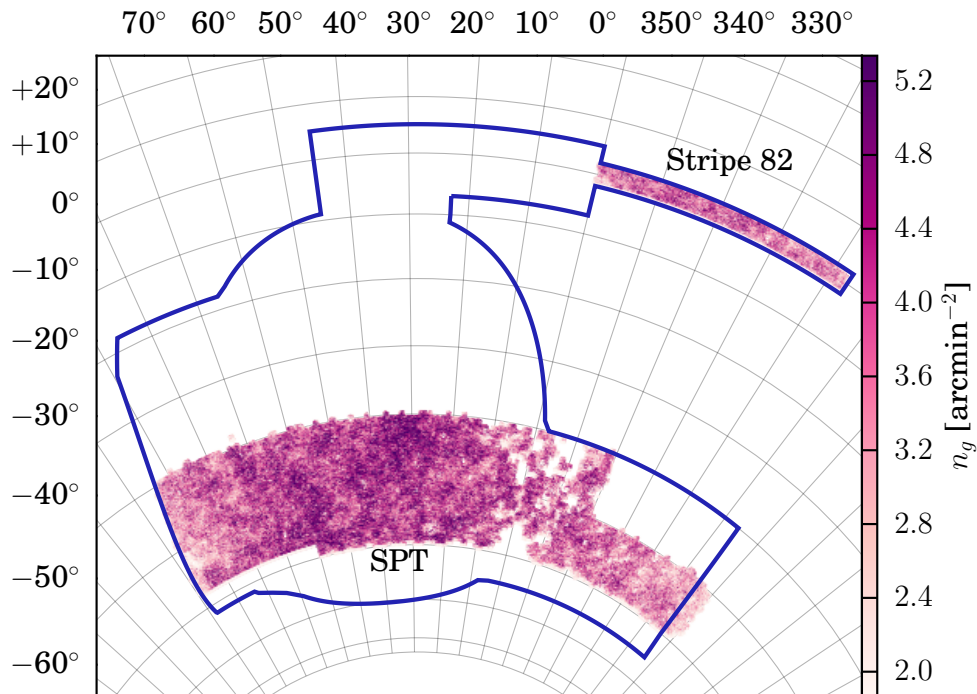


Figure 3.1: The footprint of the DES Y1 IM3SHAPE catalogue. The colour map shows the local raw number density and the solid outline shows the bounds of the Year 5 DES area. METACALIBRATION covers the same area, but at slightly higher number density. The area is divided into fields, as labelled. The largest section at the base of the footprint overlaps with the SPT survey, and is thus referred to as the SPT region. The long equatorial strip to the north is known as Stripe 82. The disjoint supernova fields are not shown here, but fall within the same bounds.

3.2.5 PSF Estimation

One of the most important tasks prior to shape measurement is accurately reconstructing the PSF at each position on the sky. The PSF describes how any point source on the sky is blurred into a two-dimensional profile on the image. The galaxy images we have access to can thus be considered as a convolution of the true surface brightness profile with a PSF kernel.

Since they are essentially point sources, observations of stars give us a direct (albeit noisy) estimate of the PSF at fixed locations. Unfortunately, it varies across, as well as between, single-exposure images and must be interpolated from the positions of stars (where it can be measured) to the locations of galaxies (where it is required).

The process for PSF estimation in Y1 is largely unchanged from the procedure used in J16. We used the PSFEX (Bertin, 2011) software to generate a PSF estimate at the position of each

galaxy in the Y1 shape catalogues. In short, the process involves first identifying stars using a size-magnitude diagram (Z17). The PSF shape and size is measured at the locations of the input stars and interpolated using a second-order polynomial in CCD chip coordinates, which is interpolated separately on each chip. The polynomial is then evaluated at the positions of a subset of stars reserved for validation. We define the PSF size and shape in terms of the second moments of the surface brightness profile (Seitz & Schneider, 1997):

$$T = Q_{xx} + Q_{yy} \quad (3.1)$$

$$e = e_1 + ie_2 = \frac{Q_{xx} - Q_{yy} + 2iQ_{xy}}{Q_{xx} + Q_{yy} + 2\sqrt{Q_{xx}Q_{yy} - Q_{xy}^2}} \quad (3.2)$$

where the moments are defined in equation 2.42. The moments are measured using an adaptive moments based algorithm (HSM; Hirata et al. 2004).

3.3 The IM3SHAPE Catalogue

3.3.1 Overview

As implied above, there are two shape catalogues based on DES Y1, both of which are available for use within the survey’s membership. The author played a major role in generating one of these two science-ready catalogues, and it is upon this that we will focus here. This dataset was generated with the maximum likelihood code IM3SHAPE, a forward modelling algorithm that uses Levenberg-Marquardt minimization of the likelihood of fitted parametric profiles. In its current form IM3SHAPE performs two independent fits using Sérsic models with power-law indices $n = 1$ and $n = 4$. We fit all of the single-exposure data simultaneously with one model, but exclude the coadd. The minimization uses unmasked pixel data from all of the postage stamp cutouts of each galaxy, weighted according to the estimated noise variance. Each galaxy is then identified as either a bulge or a disc, depending on which model returned the superior likelihood. A diagram of the iterative fitting cycle for a given model (i.e. bulge or disc) is shown in Figure 3.2.

The IM3SHAPE code¹ (Zuntz et al., 2015) is largely unchanged from the version used in SV, though the simulations used to calibrate are significantly different. We refer the reader to J16 and the original code release paper Zuntz et al. (2015) for more details about the algorithm itself.

¹<https://bitbucket.org/joezuntz/im3shape-git>

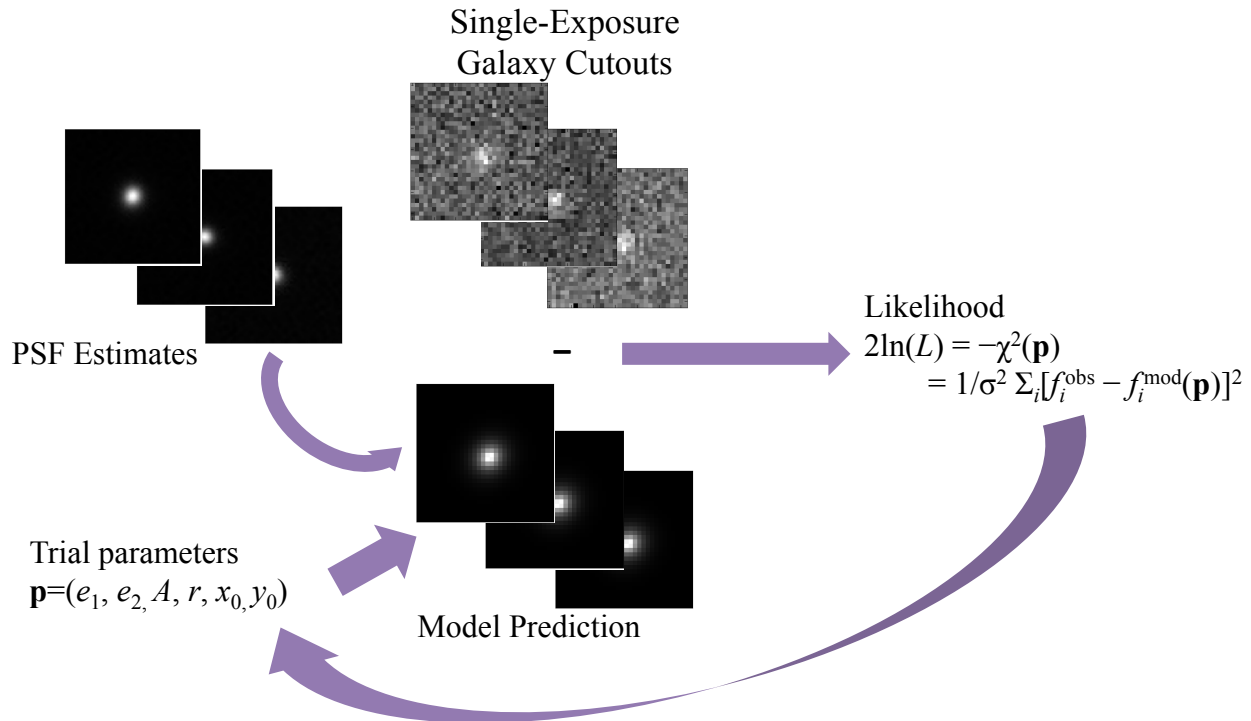


Figure 3.2: Schematic diagram of the fitting process implemented by IM3SHAPE. Trial solutions (bottom left) are repeatedly generated and combined with PSF estimates to create a convolved model prediction. The per-pixel predictions are subtracted from the image, and used to evaluate the likelihood of the model parameters. The process continues iteratively until the likelihood has converged.

The galaxy models used by IM3SHAPE have six free parameters: two ellipticity components (e_1 , e_2), a half light radius r , a centroid offset (x_0, y_0) and an amplitude A . We also compute a signal-to-noise ratio (S/N) for each measurement using the convention of Mandelbaum et al. (2015) and J16, and a PSF-normalised size R_{gp}/R_p . As we point out in J16, this signal-to-noise measure is analogous to a matched filter, favouring maximal agreement between the model fit and the image pixel fluxes. R_{gp}/R_p is defined as the ratio of the full width at half maximum (FWHM) of the convolved galaxy to the PSF FWHM, where the former is measured from a circularised version of the best-fit galaxy profile.

The mean time required by IM3SHAPE to analyse a single galaxy in DES is 1.6 seconds per exposure. In practice many objects converge in less time, and the mean is raised by a small number of faint or otherwise poorly fit objects. The total time required for shape measurements on the full Y1 dataset to complete was approximately 200,000 CPU hours, although this ignores the computation time for running and analysing the necessary suite of calibration simulations.

All IM3SHAPE measurements presented here were carried out at the National Energy and Scientific Research Computing Center² (NERSC) and the GridPP grid computing system³ (Britton et al., 2009). The Y1 HOOPOE simulations were generated entirely using the NERSC facility.

3.3.2 Galaxy Weights

We compute a measurement weight for each galaxy in DES Y1 using a very similar calculation to that outlined by J16. In summary, we first define a 2D grid of signal-to-noise and size, with each cell containing roughly the same number of galaxies. In each cell a zero-centred Gaussian is first fitted to the histogram of the e_1 component ellipticity, and the standard deviation is also calculated directly. This yields two similar but non-identical variance estimates, of which we adopt the maximum. The resulting grid is then interpolated using two dimensional radial basis functions to the position of each galaxy in this parameter grid. The weight allocated to a given galaxy is simply the inverse of the interpolated variance at that position. This process is designed to estimate the total uncertainty of an ellipticity measurement, including both shape noise and measurement uncertainty. The simulated galaxies used to calibrate IM3SHAPE are assigned weights by the same process, and these weights are used in constructing the shear calibration.

3.4 The HOOPOE Image Simulations

In this section we describe an extensive suite of image simulations, tailored to DES Y1 and built specifically for the purpose of constraining shear measurement bias. The resulting cloned survey, the so-called HOOPOE simulations, were used to model the biases, as parameterised in the standard way in equation 2.57. The additive term c_i can be constrained directly from the data, as it does not rely on knowing the underlying shear we are trying to recover at the outset. This is not the case for the multiplicative bias m . No evidence of a systematic difference between m_1 and m_2 has been detected in any region of parameter space, and so our calibration uses the arithmetic mean of the two components $m = (m_1 + m_2)/2$.

The Y1 simulations differ in philosophy from those used in SV. The latter started with postage stamp images of isolated galaxies, to which observational features were added as deemed necessary. For Y1 we start with a set of reduced survey images, and create an object-for-object simulacrum of the entire image plane, preserving as much of the original detail as possible.

²<http://www.nersc.gov/>

³<https://www.gridpp.ac.uk/>

	GREAT-DES (DES SV)	HOOPE (DES Y1)
Multiple Exposures	Coadd Only	Multi-Exposure
Point Spread Function	Discrete, Kolmogorov	PSFEx, Image Plane Variations from Data
Pixel Noise	Gaussian Random per Pixel, Fixed σ_n	Gaussian Random per Pixel, σ_n from Weight Maps
Blending/Neighbours	None (postage stamp simulations)	Simulate Full Image Plane
Galaxies Below Detection Limit	None	Random Positions, Drawn from Faint Cache
Source Detection	None	Rerun SETRACTOR on Simulated Images
Masking	None	Masks & PSF Blacklists from GOLD Catalogue
Input Galaxy Selection	COSMOS (< 23.5 mag)	COSMOS Deep (< 25.2 mag)
Magnitudes/photometry	HST F606W Magnitudes	SDSS (DES) r -band Magnitudes
Stars	None	Point Sources
Number of Galaxies	48M	68.2M
Input Shears	$ g = 0.05, 8$ Discrete Rotations	Continuous Uniform Random $-0.08 < g_{1,2} < 0.08$
Galaxy Morphology	No Morphology Split	Bulges/Discs Calibrated Separately
Interpolation/Fit	Polynomial Fit	Grid Nodes*
PSF Measurements	IM3SHAPE Weighted Moments	Radial Basis Functions Polynomial Fit HSM Adaptive Moments

Table 3.1: Comparison of the IM3SHAPE shear calibration schemes used in DES SV and DES Y1. The upper portion of the table itemises differences in the calibration simulations HOOPE and GREAT-DES. Entries below the dividing line pertain to methodological choices rather than differences between simulations. In the case of the interpolation scheme we compare three methods in this study, as described in the text below. We highlight “grid nodes” as this was adopted as the fiducial scheme for reasons discussed in the text.

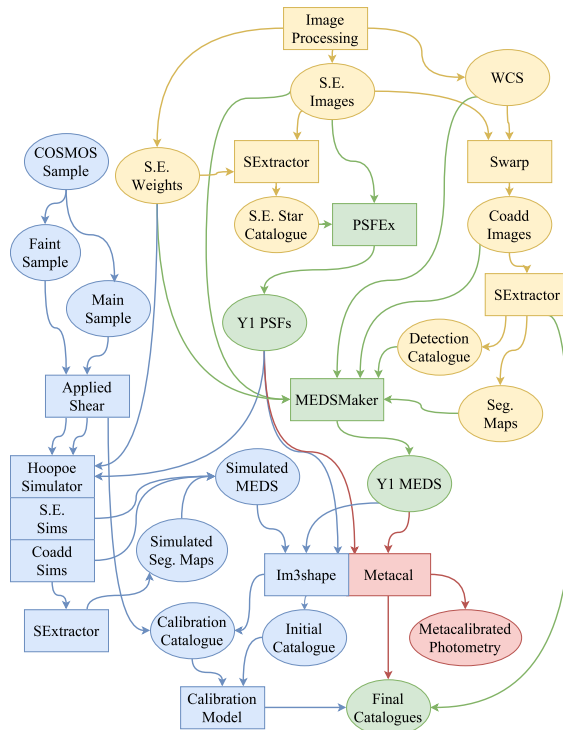


Figure 3.3: Flow-chart illustrating the steps in the DES Year 1 shape analysis, from low-level calibrated data products released by DES Data Management to the final science-ready shape catalogues. Yellow stages are performed by the DES Data Management division prior to releasing the data to the collaboration. Green blocks are tasks conducted within the weak lensing group. Processes unique to each of the two shape pipelines, IM3SHAPE and METACALIBRATION, are shown in blue and red respectively.

3.4.1 Simulating DES Y1: The Image Pipeline

The simulation pipeline for the HOOPOE image simulations is shown in the blue (left-hand) part of Figure 3.3. Subsequent analysis of the simulations was closely matched to the equivalent process used on real data, although we do not repeat the coaddition of single-epoch data or PSF estimation. The galaxy positions, noise levels, and PSFs of each simulated galaxy are taken directly from the real observations. An image mask is constructed from a bad-pixel map, which is imported from the data, and an object segmentation map, which is remade on the simulations.

The HOOPOE image simulator begins by choosing one of the 0.75×0.75 degree coadd tiles output by the DESDM pipeline. Each tile is the result of coaddition of around 70 partially overlapping exposures. In order to simulate each tile we require (a) a source catalogue generated by SExtractor (or similar object detection algorithm), (b) a WCS specifying the image bounds and the transformation between pixel and world coordinates per exposure, (c) a model describing the PSF variation across the image plane, (d) a noise variance weight map per exposure.

With these basic inputs the simulation then proceeds as follows:

1. A set of noise images is generated from the SExtractor weight maps, matched to the bounds of each real input image for this tile. A simulated coadd-image is also generated in the same way.

2. The true detection catalogue for the region is imported, and a population of fainter undetected galaxies is added (see Section 3.4.2).
3. It iterates through galaxy positions, selecting a random COSMOS profile from a rolling cache, set to give an expectation value of the number of unique profiles per coadd tile of 2000. Profiles are simulated with an additional shear and rotation angle drawn from $g_{1,2} \in [-0.08, 0.08]$, and $\theta \in [-\pi, \pi]$.
4. The profile is convolved with the PSF at the appropriate position in each image. The convolved profile is then drawn into each image (including the coadd).
5. If a faint galaxy is associated with a particular position (see Section 3.4.2) then another profile is drawn from a secondary cache of COSMOS faint profiles. It is placed at a random point in the region encompassed by the overlap of all the exposures that contain the current galaxy. This is intended to ensure it is placed in approximately the same part of the image bounds, but will not guarantee that it is a close neighbour. It is convolved, sheared and rotated as in the previous steps, using its own independent random values.
6. Once the full image has been simulated, SExtractor is run on it, generating a new detection catalogue and segmentation mask.
7. It iterates through the detection positions a second time, building the SExtractor mask for each and extracting a postage stamp cutout. In the version of the simulations the stamp size was not recomputed for each object, but was imported from the parent data. Later code versions modified this, but the change was not included for the Y1 simulation runs. As discussed in the final section of this chapter, this has implications for the priors we place on post-calibration residual biases.
8. Our code then stores and stacks the galaxy cutouts in the MEDS format (J16).
9. Finally IM3SHAPE is run on the HOOPOE MEDS files. Deblending flags are propagated into the selection flags, and the measured shears are blinded using the prescription described in Section 3.2.4 with the same factor f as is applied to the data.

3.4.2 Galaxy Sample

To capture the range of morphologies found in a photometric survey like DES, the Y1 HOOPOE simulations use real galaxy profiles rather than analytic constructions. Ideally the selection of galaxies

used for calibration should extend to at least the same depth as the data and have sufficiently low levels of noise and seeing to allow them to be degraded to match DES precisely. These requirements somewhat limit our options. The galaxy catalogue used by HOOPOE comes from Advanced Camera for Surveys (ACS) imaging of HST’s COSMOS field (Scoville et al., 2007). The catalogue of $\sim 73,000$ objects has been “whitened” (correlated noise removal; see Rowe et al. 2015), and is a deeper superset of the galaxies used in the GREAT3 challenge⁴. It extends significantly beyond the Y1 detection limit of $M_{r,\text{lim}} = 24.1$, reaching ~ 25.2 mag in the HST F606W filter and ~ 27.9 mag in the DES r -band. The publicly available HST data is limited to wide band photometry in the optical F606W filter. In order to obtain the desired magnitudes, we match the sky position of each of these galaxies to the COSMOS mock catalogue of Jouvel et al. (2009), which includes photometry specific to the transmission curves of the DES filters.

The input sample for a tile is then generated by splitting the COSMOS catalogue about the DES $M_{r,\text{lim}}$ and excluding objects too faint for detection. Each of these galaxies is simulated at its original COSMOS magnitude, rescaled to the zero-point of the DES images.

Simulated Stars

The mock images also contain stars, simulated at the positions of objects classified as stars in the real data. Stars are rendered as point sources and account for around 10% of simulated objects. This should capture any effect they may have as a source of neighbour bias, including changes they induce in the galaxy selection. Each simulated star is drawn at the measured magnitude of the real star whose position it inherits, and so should the distribution of simulated stellar magnitudes should match the data by construction. We do not re-run star/galaxy separation in the simulations, and so do not account for any bias from mis-classification. The cuts to the IM3SHAPE catalogue in size and S/N , however, will remove the majority of the ambiguous objects, so we expect the impact of this decision be small (see Section 3.5.5).

Galaxies Below the Detection Limit

In addition to simulating objects detected in parent data we also wish to simulate a population of fainter objects that escaped detection. These objects do not correspond to detection positions, but nonetheless appear in the real images. In reality such objects will contribute flux to the images and may affect both the measurements and the selection function of the detected galaxies. The

⁴<http://great3.jb.man.ac.uk>

code chooses a number of faint galaxies for each tile N_{faint} by integrating the full distribution of COSMOS magnitudes,

$$N_{\text{faint}} = \frac{f_{\text{faint}}}{(1 - f_{\text{faint}})} \times N_{\text{det}}, \quad (3.3)$$

where $f_{\text{faint}} \equiv \int_{M_{r,\text{lim}}}^{\infty} p(M_r) dM_r$ is the fraction of the weight of the normalised magnitude distribution $p(M_r)$ above the nominal DES detection limit, and M_r is the aperture magnitude. Each of these extra objects is randomly assigned a companion from the detections within its coadd tile. The faint object is placed at a random positions in the same set of exposures as its detected companion. It does not replace it, nor are their properties linked in any other way, but its position is constrained to the overlapping bounds of the same subset of single-exposure images⁵.

In the real data the flux from these galaxies enters the images prior to reduction, and would affect the background subtraction. For computational reasons we do not to not simulate thermal sky emission and rerun the background subtraction. To gauge the impact of the extra background flux, a small subset of the simulations was rerun with the same random seed settings, but omitting to draw the faint galaxies. The background estimation algorithm was then applied to the two sets of images, which were identical apart from the flux of the faint objects. In summary we find that the first order impact of the sub-detection galaxies is a uniform shift in the mean of the estimated sky background. To correct for this effect we apply a uniform additive flux correction prior to shape measurement. This correction subtracts off the mean per-pixel flux of the faint objects drawn into the specific tile, which is roughly equivalent to their impact on the background flux level. The numerical experiment outlined above is described in more detail in Appendix A. The impact of subdetection galaxies on shear measurement is also explored in more detail in Section 5.4 of Samuroff et al. (2017a) and in the next chapter of this thesis. In the tests presented therein we find a net contribution to the bias budget $m \sim -0.01$.

3.4.3 Comparing Simulations & Data

Given the sensitivities of measurement bias to the observable parameters of an image, most notably signal-to-noise, size, and ellipticity, it is important that the simulations should cover the same parameter space as the data. Whereas the calibration explicitly bins in S/N and R_{gp}/R_p , and so exactly matching the shape of the distributions is only of secondary importance, the same is not

⁵We note that the use of random positions neglects small-scale correlations, which could lead us to underestimate the level of blending in the data. The significance of this effect is beyond the scope of the current analysis, but is a useful subject for further investigation.

true of ellipticity or PSF properties.

In general we find good agreement between observables in the simulations and data. Figure 3.4 compares the measured distributions for a number of salient quantities. Unlike in previous studies, we are convolving simulated galaxies with the measured PSF at each position on the sky. This avoids quantising the PSF, which typically requires a large representative pool of distortion kernels in order to accurately mimic the measured distributions of these properties. It also accurately captures spatial correlations that would be lost if we were to draw randomly from a pool of analytic PSFs. Even using the measured PSF as input, after quality cuts it is non-trivial that the PSF properties still match well to the data. It is thus reassuring that the blue and purple curves in the upper-most row of Figure 3.4 agree as well as they do. Any significant difference in PSF properties *after* these cuts would imply a different selection behaviour with respect to PSF in simulations compared to the data. Unlike in real data the simulations do not include PSF errors, though we do test this and recognise it as a contribution to the final prior on multiplicative bias m . (see Section 3.5.5).

The distribution of *input* simulated ellipticities in Figure 3.4 is notably narrower than the *measured* distributions in both simulations and data. This is not unexpected as a consequence of both pixel noise and blending. It is, however, interesting to compare these results with shape measurements on a set of spin-off simulations described in Chapter 4 (see also Samuroff et al. 2017a). In those simulations, which are identical to HOOPOE, other than the removal of neighbour flux, we find a histogram of measured ellipticities that is much narrower. The qualitative match to the input distribution suggests that blending is a non-trivial factor in determining the shape of the measured ellipticity distribution.

There are still noticeable differences in the measured R_{gp}/R_p distributions, increasing in small galaxies despite the distributions of PSF size and best-fit radius individually matching the data well. A similar discrepancy was apparent between in SV (see J16’s Figure 12), though it was largely eliminated by the more stringent size cut we adopted there. This is thought to be a result of the (partially overlapping) COSMOS sample used in both simulations, or the similar PSF estimation methodology employed in Y1 and SV.

Finally in the lower panel we compare the input and output magnitudes from the simulations. The green shaded (dotted) histogram represents the full COSMOS catalogue, which is complete down to a 25.2 mag in the HST F606W filter. Clearly drawing randomly from this distribution would leave us with many galaxies well below the DES detection limit. The dashed and solid purple lines show respectively the input COSMOS magnitudes, selected to use in the simulation, and the values remeasured from the simulated images. No significant bias in the remeasured mag-

nitudes is detected, nor any serious disagreement with the data.

3.4.4 Bias Calibration & Diagnostics

Multiplicative Bias Scheme

With the simulations generated and shape measurements complete, this section will now define a scheme to correct for the multiplicative bias measured in the simulations. This task by necessity involves interpolating among noisy measurements. Both on theoretical grounds for noise bias (Reffregier et al., 2012) and as has been found in practice (e.g. J16 and Fenech Conti et al. 2016), shear biases are primarily sensitive to galaxy size and the S/N of the measurement. As is now standard in weak lensing analyses the Y1 calibration model is built as a function of these parameters.

The first step in this process the involves sorting the simulated HOOPOE data into a 16×16 grid according to the measured S/N and R_{gp}/R_p , allowing the bin width to vary such that each grid cell contains roughly the same number of galaxies. A multiplicative bias is derived within each cell by subdividing the galaxies into bins of g^{tr} and fitting a linear function to the bin-averaged shear response $\langle e_i \rangle - \langle g_i^{\text{tr}} \rangle$ (see equation 2.57). The resulting bias surface m^{ij} is shown in Figure 3.5.

It is important here to define a well motivated gridding scheme in terms of binning along each axis; too coarse a grid will result in real structure in this parameter space being washed out, while an overly fine sampling will inflate the statistical variance on our grid nodes. We have verified that varying our fiducial 16×16 grid between 6×6 and 20×20 does not lead to a significant change in the results.

The next task is to interpolate this grid onto the actual galaxies in this parameter space. To ensure the results presented here are robust, we compare three methods for performing this mapping. In the first scheme, we follow Fenech Conti et al. (2016), and compute a fine grid in m . If a galaxy falls within cell ij , the mean m in that cell is taken as the bias estimate. The accuracy of such an approach will clearly depend on the resolution of the grid.

In the second scheme we interpolate between grid nodes using radial basis functions. The bias at a given point is calculated as a linear combination of radial basis functions, each of which is centred on one of the nodes:

$$m(x, y) = \frac{\sum_i m_i f((x - x_i)^2 + (y - y_i)^2)}{\sum_i f((x - x_i)^2 + (y - y_i)^2)} \quad (3.4)$$

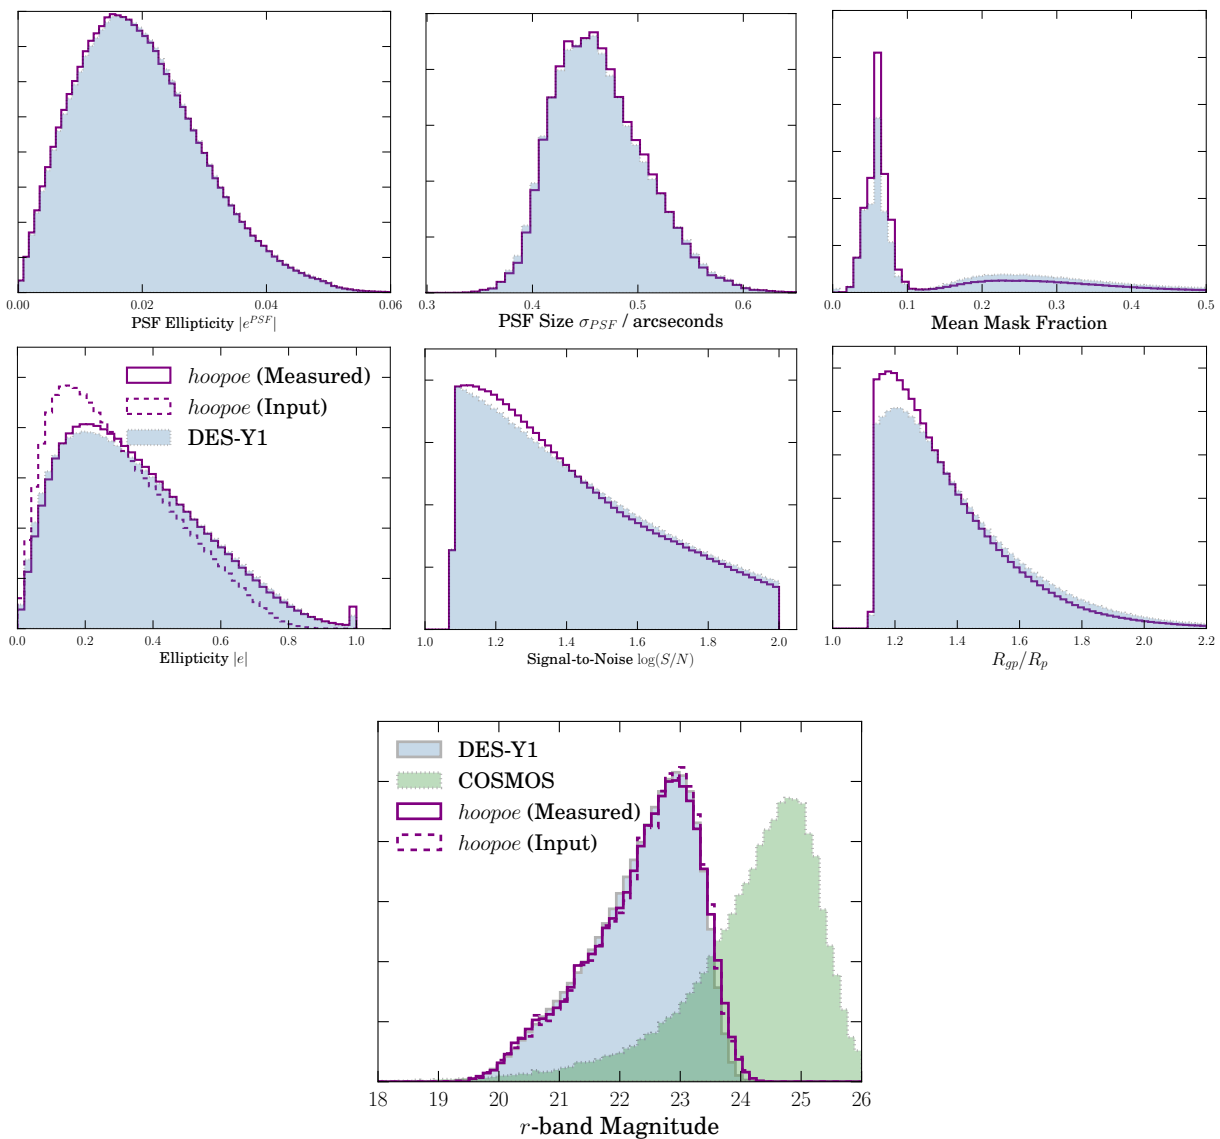


Figure 3.4: A comparison of normalised distributions in the Y1 simulations used for IM3SHAPE calibration (purple) and data (blue). The upper panels show (clockwise from upper left) PSF ellipticity; PSF size, as measured using HSM; the fraction of pixels masked out, averaged across each object’s exposures; IM3SHAPE’s measure of galaxy size relative to the PSF, R_{gp}/R_p ; flux signal-to-noise; and total galaxy ellipticity. In the latter we show both the input and remeasured distributions to the simulations as dashed and dot-dash lines respectively. The lower-most panel shows the distribution of measured and input magnitudes from the simulation, in addition to the data. The shaded green (dotted) curve shows the equivalent r -band magnitudes for the full COSMOS catalogue from which we draw our input sample.

Method	$\Delta m^{(1)}$	$\Delta m^{(2)}$	$\Delta m^{(3)}$	$\Delta m^{(4)}$
Uncalibrated	-0.0886	-0.0981	-0.1200	-0.1547
Grid	0.0069	-0.0014	-0.0074	0.0013
Radial Basis	0.0056	-0.0024	-0.0082	-0.0022
Polynomial	0.0049	-0.0028	-0.0078	-0.0000

Table 3.2: Residual multiplicative bias in the IM3SHAPE calibration simulations, after calibration using different methods for interpolating m^{ij} nodes onto individual galaxies. The calibration is derived globally, and the residuals are computed for the redshift bins used in the cosmic shear analysis in Troxel et al. (2017).

where

$$f(r^2) = (r^2/\epsilon^2 + 1)^{-\frac{1}{2}} \quad (3.5)$$

and the (x, y) coordinates are S/N and R_{gp}/R_p suitably weighted to give the two dimensions parity, ϵ is a fixed smoothing parameter, and the sums are over the grid nodes.

Finally, we fit the polynomial basis used in J16. It is not particularly useful to not write out the entire functional form here, but we note that it consists of a linear combination of 18 terms of the form $(S/N)^{-\alpha}(R_{gp}/R_p)^{-\beta}$, where the indices $\alpha, \beta \in (1.25, 1.5, 1.75, 2, 2.5, 3, 4)$. We will refer to these three methods respectively as *grid*, *RBF* and *polynomial* calibration schemes. Owing to slightly better performance in diagnostic tests the grid scheme is taken as our fiducial choice.

The relative performance of the three schemes can be assessed visually from Figure 3.6, where we show the residual bias after calibration as a function of signal-to-noise and galaxy size. The residual bias after applying the grid-based calibration might naively be expected to be zero, when binning on the same quantities used to calculate the calibration in the first place; this is only true if the same binning is applied, which is not the case in Figure 3.6⁶.

3.4.5 Robustness to Tomographic Binning

A simulation-based calibration of the sort presented here may be valid for the full dataset, and yield residual biases within tolerance, but it does not trivially follow that this is true of all sub-divisions of the data. It is perfectly possible that there are competing sources of biases, which by chance

⁶Imagine taking a set of galaxies with roughly the same S/N , dividing them into two bins about the median S/N , and calculating the mean m in each half. If we apply those corrections to the same halves and recompute m , then indeed, by construction the residual bias will be zero. Instead we can define a new box, placed across the interface of the two bins, such that it contains galaxies from both. It is no longer trivial that fitting for m in this new selection will give the mean of the values computed from the two bins




Figure 3.5: **Top:** Multiplicative bias estimates for Y1 IM3SHAPE, using the HOOPOE image simulations for objects fitted using bulge profiles (right) and disc profiles (left). The colored circles represent the grid of directly evaluated m described in the text. The underlying colour map is generated using radial basis functions to interpolate between nodes, and is for illustrative purposes only. **Bottom:** Bulge fraction as a function of galaxy signal-to-noise and size. The bulge fraction is calculated on a 16×16 grid and interpolated to generate the smooth map shown. The circles represent the grid cell positions, and are drawn at a size proportional to the total IM3SHAPE lensing weight of the galaxies contained.

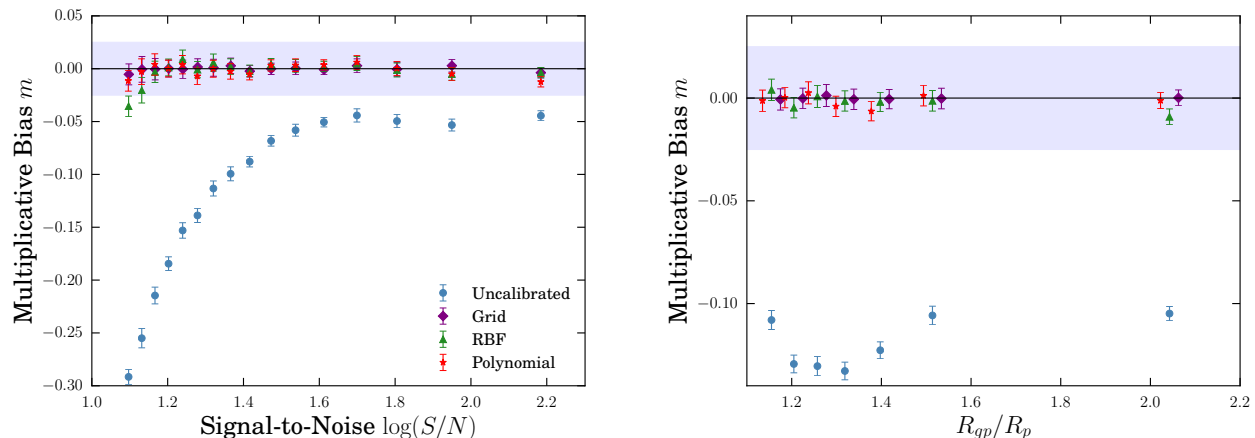


Figure 3.6: Multiplicative bias for IM3SHAPE measured from the full Y1 simulations, as a function of galaxy signal-to-noise ratio and size. The blue circles in both panels are the measured biases prior to calibration. The other points, labelled grid, RBF and polynomial are the result of correction using the three methods described in the text. The shaded blue band marks the $\pm 1\sigma$ Gaussian width of the recommended m prior for the Y1 IM3SHAPE catalogue.

cancel to zero. It is also possible to induce biases by introducing extra post-calibration selections based on quantities which correlate with galaxy shape. We show some explicit examples of this in Section 3.5.3.

Many science applications of the Y1 shape catalogues require a calibration that is robust to selection in bins of redshift and angular scale. Unlike the equivalent simulations in previous WL studies, HOOPOE generates images in sky coordinates, allowing us to test both of these. In this section we focus on the tomographic selection; we refer the reader to Samuroff et al. (2017a) and Section 4.6.2 of this thesis for discussion of scale-dependent selection effects.

The redshift information we have for each COSMOS galaxy has the form of single point-estimate photo- z , as estimated using the ACS 30-band photometry. In the following we assume this measurement is of sufficient quality to allow us to treat it as an input “true” redshift z_{tr} . Unfortunately large ground-based surveys typically have only a handful of broad filters, affording somewhat less reliable point redshifts. The best photometric redshift estimates in DES Y1 are unlikely to match their true redshifts exactly.

We build two sets of redshift bins for the calibration. For the first we use only the best estimate of z_{tr} directly for each object; this corresponds to ideal tophat redshift bins $n(z)$ (how the galaxies would be binned in the absence of redshift errors). In the second set we mimic the scatter in photometric redshift: each HOOPOE galaxy is stochastically allocated to one of the four

Y1 redshift bins. First, we construct a realistic set of DES Y1 redshift estimates using the Y1 IM3SHAPE catalogue. The redshift PDFs obtained from the BPZ code are stacked in four bins $z = [(0.2 - 0.43), (0.43 - 0.63), (0.63 - 0.9), (0.9 - 1.3)]$, resulting in four normalised distributions $n^i(z)$. We assign each galaxy with true redshift z_{tr} to a bin i , with probability $n^i(z_{\text{tr}}) / [\sum_{j=1}^4 n^j(z_{\text{tr}})]$. The resulting histograms of z_{tr} in each bin cover the full range $z \in [0.2 - 1.3]$, and approximately match the measured $n(z)$ in each bin. This random assignment of redshifts does not simulate systematic correlation between photometric redshifts and shear. It does, however, address the smearing out of the estimated redshift information due to noise, which we expect to be the dominant effect.

To test the redshift dependence of our calibration, we measure the residual bias after splitting into these bins. The results are shown in Figure 3.7 and Table 3.2 (the latter also includes values comparing alternative interpolation methods). The tophat results show larger scatter than the ones in slightly more realistic bins, which blur out the bias slightly. This worsens the residual m , particularly in the upper bin, although the residual is still within the 1σ width of our final m prior.

The residuals in our tophat redshift bins demonstrate an important limitation of our current calibration procedure: namely that galaxy morphology (and thus measurement bias) varies with redshift. Our calibration assumes that S/N and size are a sufficient proxy for change, which will be true only to some level of accuracy. The results on DES-like photometric bins suggest this has less impact on the real data, but as the requirements on m become ever tighter and we find new ways to improve the accuracy of our redshift estimates, this effect is something of which lensing community must be aware.

It may be argued that the tests outlined above are somewhat circular; we derive a calibration from the simulation, then use the same simulation to test its accuracy. As a further test we split the simulated dataset sample into halves, and use one half to generate a calibration model for the other. We perform this test twice, using different splitting methods. First, dividing the catalogue that each part contains an equal number of HOOPOE galaxies. The second time we define a split such that roughly equalises the number the unique COSMOS profiles in each half. Since measurement biases depend on both the distribution of galaxy morphologies and the specific observing conditions, neither test is redundant to the other. The results are shown in Table 3.3.

Though subdominant to the other forms of systematic bias discussed in this paper, the residual bias in the third redshift bin is statistically significant. Some residual biases might be expected, given that we are using a rigid two parameter grid to describe complex morphology-dependent biases. Unfortunately, it is not possible to predict the magnitude or sign of these residuals, which

Split Type	$\Delta m^{(1)}$	$\Delta m^{(2)}$	$\Delta m^{(3)}$	$\Delta m^{(4)}$
None	0.0069 ± 0.0044	-0.0014 ± 0.0046	-0.0074 ± 0.0030	0.0013 ± 0.0034
At random	0.0021 ± 0.0046	-0.0018 ± 0.0039	-0.0095 ± 0.0039	-0.0027 ± 0.0054
By COSMOS profile	0.0034 ± 0.0062	-0.0006 ± 0.0060	-0.0048 ± 0.0037	0.0073 ± 0.0039

Table 3.3: Residual multiplicative bias in the HOOPOE simulations under various divisions. For reference the top line shows the result of applying the fiducial calibration to the whole catalogue, and is identical to the “grid” line in Table 3.2 and the purple diamonds in Figure 3.6. The other lines show the remeasured biases when using disjoint calibration and validation subsets of the simulation. We split first at random, such that there are equal numbers of HOOPOE galaxies in each subset, and then such that there are equal numbers of COSMOS profiles in each.




Figure 3.7: Residual multiplicative bias for IM3SHAPE measured from the full HOOPOE catalogue in four tomographic bins after fiducial calibration. For the “tophat” points objects are binned by their COSMOS redshifts, and for the “DES” bins they are assigned to match DES Y1 redshift distributions, partially simulated photometric redshift errors. As above the shaded band shows the 1σ width of our Gaussian prior on m in the Y1 IM3SHAPE catalogue, and the horizontal dotted lines the redshift bin boundaries.

depend on the details of the COSMOS sample and how they are distributed between redshift bins. It is thus not guaranteed that the measured residual m in the third redshift bin implies an equivalent bias in the data.

To account for this uncertainty we widen our prior on the residual m in the calibrated catalogue. The maximum amplitude of the residual bias in Figure 3.7 is $|\Delta m^{(i)}| = 0.0080$. We include this amplitude rounded up to $\sigma_m = 0.01$ as a systematic contribution to the prior on residual bias in the IM3SHAPE catalogue (see Section 3.5.5). In order to ensure we are not underpredicting the level of uncertainty we also widen the m prior to account for the fact that these residual biases will be correlated between redshift bins.

3.4.6 Validating the HOOPOE Simulations

The calibration tests presented above are promising, but make use only of the simulations themselves. As will inevitably be the case with simulations, however, HOOPOE is not a perfect copy of the data, and systematic differences are present that may affect the multiplicative bias calibration.

The first limitation comes from the fact that a finite selection of COSMOS galaxies is used to simulate a much larger sample of DES galaxies. The cache of input profiles, though continuously updated, is relatively small. This results in particular COSMOS galaxies appearing repeatedly within particular regions of the simulated images. Such effects could conceivably lead to biases, either additive (since the drawing process is directional) or multiplicative, if the frequency of repetition is sufficiently high. To test this we divide the HOOPOE galaxies according to COSMOS identifier. For each unique profile we construct a k -d tree data structure on the coadd pixel grid. This is repeatedly queried to locate the nearest instance of the same COSMOS profile. The distance to this self-neighbour is referred to as the recurrence scale.

We find a mean recurrence scale of ~ 150 pixels or 40.5 arcseconds, though there is a significant asymmetry in the distribution of distances with a heavy tail out to 1000 pixels and higher. The fraction of galaxies with a relatively close self-neighbour is, however, also non-vanishing. We thus perform the following test. HOOPOE galaxies are first assigned to four tophat redshift bins, as described in Section 3.4.5. In each bin we fit for multiplicative and additive biases (a) using all galaxies and (b) using only galaxies with no instance of the same profile within a radius of 100 pixels. The raw number removed by the cut is relatively small, but it could conceivably favour small round objects. To ensure we are measuring the true impact of self-neighbours, and not a selection effect from the cut devised to remove them, we reweight the surviving galaxies. Weights are assigned based on S/N and R_{gp}/R_p , such that, when applied, the 2D histogram $p(S/N, R_{gp}/R_p)$ matches the data. We find no significant change in multiplicative nor additive bias in any of the redshift bins ($\Delta m \sim 10^{-4}$, $\Delta c_i \sim 10^{-5}$).

A second limitation concerns the nature of the input COSMOS profiles themselves. The simulations make use of an early release of the deep COSMOS catalogue. Due to masking errors and deblending failures a (small) fraction of this input catalogue are visibly defective. Given that these represent unphysical characteristics, it is very difficult to gauge their impact on shape measurement. Though one might expect the majority to be flagged by IM3SHAPE’s quality cuts, we cannot assume that they have no net impact. To test this, an internal collaboration crowdsourcing exercise was devised, the details of which can be found in Appendix B. The aim was to categorise the COSMOS galaxies into groups according to their visual characteristics. In total the cut resulting

from this exercise to excise the corrupted profiles removes 0.51M/18M objects from the simulated shape catalogue. Using a similar nearest neighbour search as described above, we estimate mean distance to the nearest “bad” COSMOS profile to be ~ 90 pixels. We recompute the biases m and c_i under three scenarios: (a) using all galaxies, (b) cutting COSMOS profiles classed as artefacts or oversized relative to their postage stamps and (c) the same as (b), but also removing galaxies drawn within 100 pixels of a bad COSMOS profile. We find the computed biases are stable to well within 1σ in all apart from the upper redshift bin. Here we lose the bulk of the galaxies flagged by this cut, which is perhaps unsurprising given that these objects tend to be small, faint and thus most susceptible to deblending failures. The change in all scenarios is at the level of the 1σ statistical error at $\Delta m \sim 0.005 - 0.0075$. Though small, this is non-trivial and we incorporate this uncertainty as a systematic contribution to our m prior (see Section 3.5.5).

The use of the Y1 detection catalogue to source the positions of simulated galaxies is intended to capture the galaxy clustering patterns across the survey. It does have some drawbacks, chiefly that it omits undetected or strongly blended galaxies (see Section 3.4.2). A second potential limitation is this: not all detections in the Y1 source catalogue correspond to real galaxies. Spurious detections can be produced by CCD chip edges and by image artefacts such as satellite trails and ghosts. It should be stressed that the real images undergo rigorous processing and vetting prior to shape measurement, and so these features are not included in the final gold catalogues described by Drlica-Wagner et al. (2017). The raw detection catalogues, which are used as inputs to our simulations, however, do not provide sufficient information to distinguish real from false detections during runtime. We tried a simple detection algorithm to flag these features using boxcar averaged source densities along rows of pixels, but this was not found to reliably detect diagonal or curved streaks. The HOOPOE images consequentially include infrequent but visually striking lines of COSMOS galaxies in these locations. To quantify the impact, a second crowdsourcing exercise was implemented analogous to the one described in Appendix B. Participants were asked to inspect approximately half of the simulated tiles, each of which was split into 5×5 square patches. Patches in which the detection positions exhibited structure were flagged for removal. As before we then divide HOOPOE galaxies into DES-like redshift bins and recompute m and c_i , including and then excluding the flagged regions. Using all galaxies (no redshift binning) we find a shift $\Delta m = 3.7 \times 10^{-5}$, which is equivalent to less than 2% of the 1σ statistical uncertainty on m . In four redshift bins we remeasure the bias before and after the cut, and find a difference m by at most $\Delta m = 0.0007$. Since any systematic shift is comfortably subdominant to statistical uncertainty, we do not consider spurious detections further as a source of systematic calibration error.

3.4.7 Sensitivity to Observable Distributions

Finally, we seek to quantify the impact of the differences in observable characteristics seen in Section 3.4.3. The only quantities that show systematic differences are flux (note that we treat bulge and disc fluxes as separate parameters here) and R_{gp}/R_p . The raw distributions of R_{gp}/R_p and flux are shown by the solid lines in Figure 3.8, with the parent DES data shown by the shaded histograms.

We assess the importance of these differences by reweighting the HOOPOE measurements to match the real data. In the case of R_{gp}/R_p we simply divide galaxies into bins of size and assign a uniform weight to each bin, such that the simulated distribution $p(R_{gp}/R_p)$ matches the data. In the second case we carry out the same procedure for galaxy flux. This time, however, an independent set of weights is computed for bulge and disc galaxies, such that they each match the corresponding sub-populations of the data. The reweighted distributions are shown by the dashed lines in Figure 3.8.

As pointed out by Fenech Conti et al. (2016), who carried out a similar test for KiDS, reweighting can be problematic if the quantities in question are covariant with ellipticity. In such cases reweighting to match a 1D projected distribution $p(q)$ may be inadequate to correct (or even worsen) differences in the 2D joint distribution $p(q, e)$. In each case we check both the 2D distributions (not shown here) and the 1D $p(e)$ histograms (shown in the right-hand panels of Figure 3.8). Neither reweighting operation is found to produce such spurious differences.

Finally, galaxies are divided into four Y1-like tomographic bins, as before, the fiducial calibration is applied, and the residual m is calculated in each bin. The results are shown in Figure 3.9. The maximum change under reweighting Δm in both cases is $\mathcal{O}(10^{-3})$. This is not found to have a coherent direction across z bins, and is well within both the statistical error margin (the blue shaded boxes) and the 1σ width of our prior (the dashed horizontal lines).

3.5 Using the Shear Catalogues

3.5.1 Mean Shear

Both catalogues show a non-zero mean ellipticity over the entire Y1 survey, with a value $e_{1,2} = (2.9, 1.8) \times 10^{-4}$ for METACALIBRATION and $e_{1,2} = (-1.5, 2.7) \times 10^{-4}$ for IM3SHAPE. This is marginally too large to be the mean of cosmic shear over the field: in log-normal simulations we find a standard deviation of the mean $e_{1,2} \sim 1 \times 10^{-4}$ over our region. An added constant shear




Figure 3.8: Histograms of flux (upper left), size (lower left) and ellipticity (upper/lower right) in the data and simulations used for IM3SHAPE calibration. In the upper panel we show bulge and disc galaxies separately in red and blue respectively. In all panels the filled distributions are the DES Y1 data and the solid lines show the full HOOPOE simulations. The dashed lines in the upper two panels represent same simulation sample, but reweighted to match the flux distributions of the data, as described in the text. The lower panels show the same, but reweighted for size.

will appear as a constant offset in correlation function measurements. We thus have the choice of either empirically subtracting a constant value from the datavector, or marginalising over an additional nuisance parameter in cosmological parameter estimation.

The origin of the mean shear is unknown, and is likely to be the combination of several effects. Charge self-interaction effects in the DECam CCDs are expected to cause mean shears in the e_1 direction that are of the order of a few times 10^{-4} (see Table 1 of Gruen et al. 2015). Residual shear-PSF correlations might also be expected to contribute a directional shear to the measurement.




Figure 3.9: Change in the residual IM3SHAPE bias resulting from reweighting the simulations prior to calibration, shown in the four tomographic bins used in the DES Y1 shear 2pt analysis. The purple circles show the result when reweighting to compensate for the excess of small galaxies shown in Figure 3.4 and the bottom right-hand panel of Figure 3.4, while the blue diamonds are reweighted for bulge/disc flux. The blue bands mark the 1σ statistical error on m , while the vertical dotted lines are the $\pm 1\sigma$ bounds of the m prior for IM3SHAPE.

3.5.2 Catalogue Flags

Each catalogue uses its own flagging scheme to determine which galaxies can safely be used in science applications. The scheme adopted for IM3SHAPE is similar to that in J16. A small number of “error flags” that identify the most extreme objects, and a larger number of “info flags” are designed to select the highest quality shape measurements. These shape measurement flags are combined into a single `FLAGS_SELECT=0` value in our released catalogues. The requirement of zero-valued flags is imposed when computing the calibration scheme, and should be likewise applied in precision applications that make use of it.

The flag values are described in Table 3.4. The main changes implemented since J16 are a reduction in minimum S/N from 15 in SV to 12 here, and our minimum R_{gp}/R_p from 1.15 to 1.13. These decisions reflect the improved treatment of small faint objects in the calibration simulations. The following cuts are strongly recommended for the IM3SHAPE catalogue:

$$\begin{aligned} S/N &> 12, \\ R_{gp}/R_p &> 1.13. \end{aligned}$$

(3.6)

Value	Info Flag Meaning	Error Flag Meaning
2 ⁰	Area masked out in the gold catalogue	IM3SHAPE failed completely
2 ¹	Region flagged in the gold catalogue	Minimizer failed to converge
2 ²	MODEST classifies as star	$e < 1e - 4$: im3shape fit fail
2 ³	Mask fraction > 0.75	e_1 or e_2 outside $(-1, 1)$
2 ⁴	Minimizer Iterations > 10000	Radius > 20 arcsec
2 ⁵	r-band extractor flag 0x1, (bright neighbours)	$R_{gp}/R_p > 6$ or NaN
2 ⁶	r-band extractor flag 0x2, (blending)	Negative or NaN R_{gp}/R_p
2 ⁷	More than 25% of flux masked	$S/N < 1$ or NaN
2 ⁸	$S/N < 12$	χ^2 per effective pixel > 3
2 ⁹	$S/N > 200$	Normalized residuals < -20 in any pixel
2 ¹⁰	$R_{gp}/R_p < 1.13$	Normalized residuals > 20 in any pixel
2 ¹¹	$R_{gp}/R_p > 3.0$	RA more than 10 arcsec from nominal
2 ¹²	Radius > 5 arcsec	Dec more than 10 arcsec from nominal
2 ¹³	Radius < 0.1 arcsec	Failed to measure the FWHM of psf or galaxy
2 ¹⁴	Centroid more than 1 arcsec from nominal	r-band extractor flag has 0x4 or above
2 ¹⁵	χ^2 per effective pixel < 0.5	
2 ¹⁶	χ^2 per effective pixel > 1.5	
2 ¹⁷	Normed residuals < -0.2 somewhere	
2 ¹⁸	Normed residuals > 0.2 somewhere	
2 ¹⁹	Very large PSF	
2 ²⁰	Negative PSF FWHM	
2 ²¹	One or more error flags is set	

Table 3.4: IM3SHAPE info flags (right) and error flags (left) marking objects with undesirable features. These flags are included in the Y1 data release.

Deciding the exact numerical values of these cuts is not an exact science. The lower bounds are imposed on these quantities due to limitations in the input COSMOS sample we use in the simulations (see Mandelbaum et al. (2015)). It is worth reiterating that additional selection beyond the recommended cuts should be approached with extreme caution. This point is elucidated in the following section, where we describe the correct application of the simulation-derived calibration.

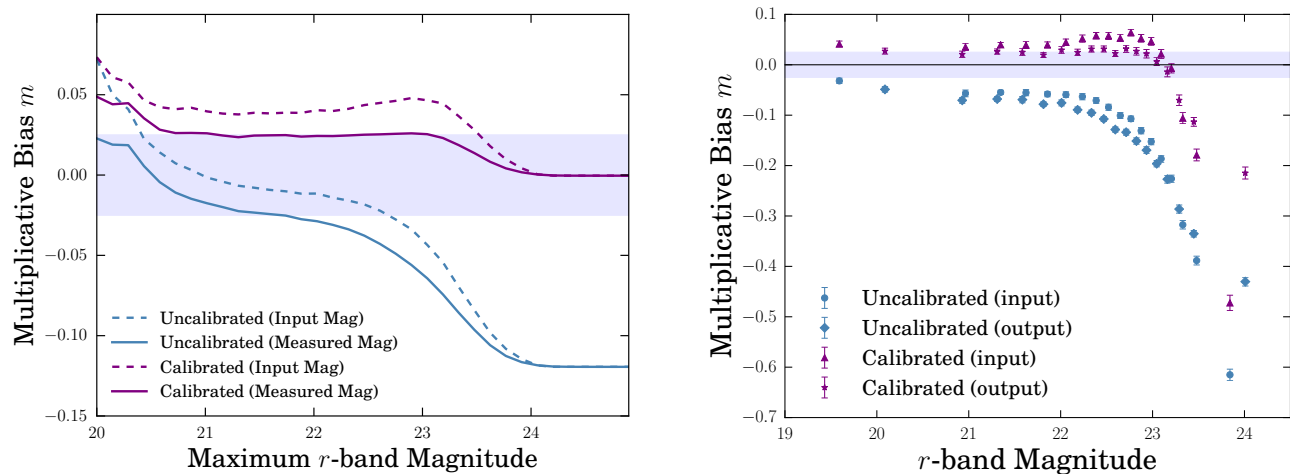


Figure 3.10: Multiplicative bias for IM3SHAPE, measured from the HOOPOE image simulation. The solid lines show the measured bias after imposing a maximum r -band magnitude, using the measured values from the SExtractor run on the simulation. Dashed lines show the same, but defining the cut using the input magnitudes. In the purple curves the shear measurements are corrected using the fiducial scheme described in Section 3.4. Blue curves are uncalibrated. This illustrates the danger of selection biases when cutting on any observable which correlates with ellipticity, as does magnitude. The points in the right-hand panel show the same, but dividing galaxies into narrow bins of magnitude instead of imposing an upper-limit cut.

3.5.3 Applying the IM3SHAPE Calibration

The IM3SHAPE calibration yields m and c values for each object, but because they include corrections for selection biases these values are only valid if the recommended set of IM3SHAPE cuts is imposed. Further selection can induce biases due to noise that is correlated between ellipticity and other quantities. We have verified in Section 3.4.5 that division into the four tomographic bins recommended for the Y1 DES analyses does not induce a significant bias. This cannot be guaranteed for arbitrary subdivisions of the catalogue. In the right hand panel of Figure 3.10 we show a dramatic example of this. Dividing the simulated galaxies into narrow bins by magnitude




Figure 3.11: Multiplicative bias for NGMIX shear measurements on GREAT-DES simulated data as a function of redshift. The red circles show the bias calculated using all galaxies that pass the NGMIX selection criteria. The blue triangles show the bias when also including the recommended SV IM3SHAPE selection. The grey band represents a margin of $\pm 3\%$ in m .

(either observed or input) induces multiplicative biases of 30% or more. A perhaps more practical example is shown in the left-hand panel. An upper-limit magnitude cut (denoted by the horizontal axis) introduces a selection biases of 2 – 6%, even after calibration. Indeed, depending on the cut, the calibration can actually worsen the bias considerably. Cuts on measured quantities are not the only concern here: one should be wary of *any* change in the population of galaxies beyond that which the calibration is tuned to. Figure 3.11 shows the result of a demonstrative exercise carried out using the SV catalogues. In the analysis of J16 we ran two simulation codes on the same simple simulation, referred to as GREAT-DES. The two codes generate their own flags, fail on different galaxies, and thus have slightly different selection functions. As we can see here, imposing the selection function of one catalogue on another (IM3SHAPE on NGMIX in this case) without regenerating the calibration can induce a selection bias of several percent. Direct comparison of two shear catalogues is clearly complicated by the difference in selection at any point prior to cosmology constraints.

The IM3SHAPE calibration is applied to the Y1 catalogue using the prescription set out by J16. The mean shear estimator for an ensemble of galaxies is:

$$\langle \gamma_i \rangle = \frac{\sum_{\alpha} w^{\alpha} (e_i^{\alpha} - c_i^{\alpha})}{\sum_{\alpha} w^{\alpha} (1 + m^{\alpha})} \quad (3.7)$$

where $i = 1, 2$ and α runs over rows in the catalogue, and w^{α} denotes the per-galaxy weights described in Section 3.3.2. For a shear two-point estimator, the additive c correction should first be applied, then the galaxy pairs rotated to the tangential and cross directions e_+ and e_{\times} . The weights and multiplicative corrections should be applied to these rotated values:

$$\xi_{\pm}(\theta) = \frac{\sum_{\alpha} \sum_{\beta} w^{\alpha} w^{\beta} (e_{+}^{\alpha} e_{+}^{\beta} \pm e_{\times}^{\alpha} e_{\times}^{\beta})}{\sum_{\alpha} \sum_{\beta} w^{\alpha} w^{\beta} (1 + m^{\alpha})(1 + m^{\beta})} \quad (3.8)$$

where the sums run over pairs of galaxies (α, β) , whose angular separation falls within a bin of some finite width $\theta \pm \Delta\theta$.

3.5.4 Number Density

Values of the (effective) number density and shape variance for three definitions for the two catalogues are shown in Table 3.5. The raw value is simply the total number of selected objects per unit of unmasked survey area Ω :

$$n_{\text{raw}} \equiv \frac{N_{\text{gal}}}{\Omega}. \quad (3.9)$$

This is not the most useful quantity for gauging the power of a catalogue for lensing cosmology. It is instead common to construct weighted density estimators, which upweight the contribution of galaxies for which we have high-quality shape measurements. The estimator devised by Chang et al. (2013) is the a particularly useful one for assessing cosmological constraining power. It is defined as:

$$n_{\text{eff}}^{\text{C13}} \equiv \frac{1}{\Omega} \sum \frac{\sigma_{\text{sh}}^2}{\sigma_{\text{sh}}^2 + \sigma_{m,i}^2}, \quad (3.10)$$

where the sum runs over galaxies and σ_{sh}^2 is the intrinsic scatter in galaxy shapes. The second component $\sigma_{m,i}^2$ is the measurement uncertainty for galaxy i , which naturally depends on the S/N of the object in question. The denominator here is the total shape variance σ_e^2 , or the sum of an intrinsic shape dispersion and a measurement noise term $\sigma_{\text{sh}}^2 + \sigma_{m,i}^2$. In practice we estimate σ_{sh}^2 using the highest signal-to-noise objects, which have vanishing measurement error. We do not attempt to isolate σ_m for each galaxy, but rather evaluate σ_e as a whole using the IM3SHAPE measurement weights (see Section 3.3.2).

A second common definition is presented in Heymans et al. (2013). This estimator has been used in previous shear analyses, and to facilitate comparison we evaluate it for DES Y1:

$$n_{\text{eff}}^{\text{H13}} \equiv \frac{1}{\Omega} \frac{(\sum w_i)^2}{\sum w_i^2}. \quad (3.11)$$

In the case of applying equal weights for all galaxies (as is the case in our METACALIBRATION catalogue) this reverts to the raw number density.

Catalogue	Definition	$n_{\text{gal}} / \text{arcmin}^{-2}$
IM3SHAPE	Raw	4.50
	C13	3.16
	H13	3.71
METACALIBRATION	Raw	6.40
	C13	5.88
	H13	6.28

Table 3.5: Effective number density evaluated for the two Y1 catalogues. The three quantities shown here are differently weighted measures of the statistical power of the catalogues, as defined in equations 3.9, 3.10 and 3.11.

As is apparent from Table 3.5, the METACALIBRATION catalogue is the larger of the two. This is primarily the result of the fact that the fits used multi-band *riz* data, whereas IM3SHAPE uses only the *r* band. The cuts applied in METACALIBRATION are also less stringent, which contributes to the difference.

3.5.5 Choosing a Prior on Residual Calibration Bias

In the following we outline an exercise designed to characterise the budget of post-calibration error in the IM3SHAPE catalogue. From a cosmologist’s perspective, the practical use of this is to gauge (as objectively as possible) our confidence that the data are systematics free. That feeds into cosmological likelihood calculations via a prior on residual m .

We start with the simplest form of uncertainty: statistical error. The volume of our simulations is large but finite, leading to a zero-centred prior contribution of width of $\sigma_m = 0.002$. We take this as our baseline, to which we will add contributions to reflect the systematics tests below.

In SV, where the systematic uncertainty in the multiplicative bias calibration was determined from the (dis)agreement of our two shape measurement pipelines on the GREAT-DES postage stamp simulations. Thanks to the level of detail in HOOPOE and the complications in directly comparing the two codes, the systematics budget for IM3SHAPE is now set by a set of simulation-based tests

In general, when we identify an entirely untreated systematic the prior absorbs the full amplitude of the change in m it induces (we assume a top-hat prior contribution). If a systematic is treated, but imperfectly so, we add 50% of the amplitude to the prior as a zero-centred Gaussian contribution.

We assume Gaussian priors in the following, although the exact functional form is unlikely to

impact the derived cosmology constraints significantly.

There is one final subtlety we should note here. In practice, many science applications of shear catalogues will involve some form of tomographic binning, and typically one wishes to place a prior on the residual m per bin. If the systematic error contributions are just an unknown scatter, which applies independently to the galaxies in each bin, then simply taking the quadratic sum of the error components derived in the following paragraphs and applying it independently in each bin would be a proper treatment of the uncertainties. If they are correlated between bins, however, the situation is slightly less straightforward; assuming an independent prior per bin could under- or overestimate the uncertainties, depending on how the particular statistic used responds to the shear in each redshift bin. To be conservative in the general case, or if this is not known for certain, we thus recommend that the contributions that are believed to correlate between bins should be widened in the calculation of the tomographic m prior. For a more detailed explanation of this argument see Appendix D of Z17 and Hoyle et al. (2017).

Systematic Uncertainties in the Calibration

A significant part of the uncertainty is due to systematic imperfections in the simulations. The dominant systematic uncertainties arise from:

- **The impact of neighbours.** Comparison of IM3SHAPE runs on identical sets of simulations with and without neighbouring galaxies (Samuroff et al. 2017a; see their Figure 13) show a mean shift in calibration corresponding to $\Delta m = -0.034$; mean shears measured in simulations with neighbours are about 3 percent smaller than for a sample of fully isolated galaxies. While the simulation-based calibration is a valid treatment of this effect, and should capture its dominant influence on shape measurement, one cannot say with absolute confidence that all aspects of neighbour bias are accurately reflected in the simulations. We therefore assume half of the neighbour-induced shift in our calibration as an uncertainty, $\sigma_m = 0.017$, which is conservative given the degree of realism in the simulations.
- **Assignment of cutout sizes.** While stamp size in the real data is based on measurements of a source’s size and ellipticity performed on the coadd using SEXTRACTOR, in the bulk of the simulations the code truncated each simulated galaxy’s image at the bounds of a postage stamp of the original source whose position it was taking. The largest and most highly elliptical galaxies are, then, typically assigned smaller boxes than they would in the data. Assessing this problem from the existing simulations is complicated by the lack of a reliable

measure for what the stamp size *should* be. We were unable to devise a cut based on the true input properties of the simulated galaxies that did not significantly alter the ellipticity and size distributions. Reweighting was found to be unreliable, since the cut leaves very few large elliptical galaxies to upweight, and was not found to be robust to binning in S/N and R_{gp}/R_p . Re-running a subset of the simulation with this problem fixed, we find a maximum change in multiplicative bias of 0.025. We assume a tophat prior of $|m| < 0.025$ per redshift bin, corresponding to a Gaussian $\sigma_m = 0.014$. While this is a non-negligible contribution to our overall error budget, rerunning the full simulation with box sizes assigned according to properties measured in the stack, as is done in the data, would require a large computational overhead. We postpone such an undertaking to future work.

- **Removal of bad objects from the COSMOS galaxy sample.** Within the DES collaboration, we have manually identified galaxies among the COSMOS library that show issues potentially affecting multiplicative bias calibration (see Appendix B). The change in calibration when removing flagged galaxies is at most 0.009 among tophat redshift bins (see Section 3.4.5). Despite these efforts, the choice of which galaxies to remove is somewhat subjective. To account for uncertainty in how the cut changes the morphology distribution of galaxies, we assume half of the observed shift, or 0.005, as a systematic uncertainty.
- **Parameterisation of the calibration.** As discussed in the earlier sections, our calibration is generated as a function of S/N and size, and performed separately for galaxies better fit by bulge- and disc-type Sérsic profiles. Noise bias, however, depends on additional galaxy properties, the distributions of which at given signal-to-noise ratio and size vary as a function of redshift. When we apply the calibration derived from the full galaxy sample, we find that then imposing a redshift binning can induce non-zero multiplicative bias. The maximum amplitude of these residuals is found to be $|\Delta m| \sim 0.008$, and is robust to the choices that enter the calibration scheme (interpolation, binning etc). Since lensing analyses virtually always employ some form of redshift-dependent binning or weighting, we assume an additional systematic uncertainty $\sigma_m = 0.01$ in each redshift bin.

These contributions are summarised in the upper portion of Table 3.6. We constrain them as a set of raw amplitudes, by which m is found to change under the various tests described above. When choosing a prior one could simply assume each contribution corresponds to a uniform m prior of the equivalent width. This is less than desirable for two reasons: (a) it does not facilitate easy combination and manipulation of the prior components and (b) it underplays our confidence in the

simulation, implying *no* confidence in the veracity of the simulation within some bounds. We thus convert these amplitudes Δm into Gaussian contributions with the variance of the equivalent tophat distributions. (see Z17). Combined in quadrature, these effects amount to a Gaussian systematic uncertainty of $\sigma_m = (0.017^2 + 0.014^2 + 0.005^2 + 0.010^2)^{0.5} = 0.0247$.

PSF Size Inaccuracy

In Y1 there is a mean residual between the predicted PSF size and that measured at the positions of stars due to imperfect interpolation (see Z17). While the dominant effect of PSF error will be directional, and thus contribute a mean shear, isotropic size residuals will manifest themselves as an excess m . To assess the impact we run IM3SHAPE on a set of very simple noise-free simulations devised expressly for this purpose. The simulations consist of analytic galaxies under constant shear $\mathbf{g} = (0.01, 0.00)$, convolved with highly elliptical analytic Moffat PSFs. Two versions were created, first with a constant reference PSF, and then with a constant fractional PSF dilation of $\Delta T_{\text{PSF}}/T = 8.3 \times 10^{-4}$. Assuming the same reference PSF used in shape measurement for both simulations, the mean multiplicative bias is seen to vary by $\Delta m = 0.006$ between the two simulations. While these PSF simulations likely capture the dominant part of the effect, realistic galaxy morphology could conceivably change the result marginally. Adopting a conservative approach, we scale this observed change by a factor of 1.5 before incorporating it into our m prior. After conversion to Gaussian width, maintaining variance, the total impact is $\sigma_m = 0.005$. Note that we find no change in additive biases between the two simulations.

Stellar Contamination

Stellar contamination of the IM3SHAPE source sample is believed to have negligible impact due to the stringent cuts to remove point sources. We test this using the COSMOS field, where robust flagging of stars using HST photometry has been carried out. The total impact of misidentified stars on m is estimated to be less than 10^{-3} , and including it has no impact on our prior width, after rounding.

Final Prior

Adding the effects listed above in quadrature, (summarised in Table 3.6), we arrive at a multiplicative bias of $m = 0.000 \pm 0.025$.

This value is the baseline Y1 prior recommended by Z17, and is valid for the full catalogue,

Effect	Centre	Gaussian Width σ_m	Correlated Between Bins
Neighbour Bias	0.0	0.017	✓
Stamp Size	0.0	0.014	×
COSMOS Sample	0.0	0.005	×
Morphology	0.0	0.010	×
Calibration Statistical	0.0	0.002	✓
Stellar Contamination	0.0	0.001	✓
PSF Size Bias	0.0	0.004	✓
Total	0.0	0.025	

Table 3.6: Multiplicative bias budget for IM3SHAPE. The upper part of the table lists the systematic uncertainties, arising from potential sources of shear bias that are not fully treated in the image simulation or the calibration. These systematics include potentially uncalibrated neighbour biases, truncation bias in the rendering of COSMOS profiles, removal of bad objects from the COSMOS sample, and variation in morphology as a function of redshift beyond that captured by S/N and size. The three rows below the divider show other known uncertainties in the calibration, and the final row is the total prior width, given by the quadratic sum of the listed uncertainties. The right-most column indicates whether or not we expect these contributions to be correlated between redshift bins.

as presented. It is finally worth making a cautious note with regards to tomographic binning. A number of the systematics described will not act independently in disjoint populations of galaxies. For reasons set out in mentioned above, if one divides the catalogue into N redshift bins, it is advisable to widen the prior contributions marked in the last column of Table 3.6 by a factor of $\sim \sqrt{N}$ to account for inter-bin correlations in m . We refer the reader to Appendix D of Z17 for details on the calculation of the precise factor.

3.6 Summary and Discussion

In this chapter we have discussed the construction of one of two independent shear catalogues in Year One of the Dark Energy Survey. The IM3SHAPE catalogue covers 1500 square degrees of the Southern sky and contains 21.9 million galaxies. It has passed a battery of tests that demonstrate that, when appropriately used with calibration and error models, it is suitable for weak lensing science. These data have been used in some of the most powerful leading-edge weak lensing studies to date. A raft of science papers due for release with Z17 demonstrate that IM3SHAPE and its sister catalogue METACALIBRATION, provide consistent cosmological constraints. Troxel et al.

(2017) focuses on constraints from cosmic shear alone, Prat et al. (2017) measure and explore the constraining power of galaxy-galaxy lensing, and DES Collaboration et al. (2017) uses both, in addition to galaxy density correlations.

This work made use in the IM3SHAPE catalogues of the most sophisticated simulations currently used in practice for shear bias calibration. They account for a range of systematic effects including blending, and the impact of spatial masking. We have emphasised the importance of matching the simulations to the data with extreme care and as much detail as possible. We have also demonstrated that even a catalogue with zero residual bias when taken as a whole can very easily become inaccurate for shear measurement under additional cuts.

Like all weak lensing catalogues, the DES Y1 results have some inherent uncertainty in the calibration in the form of a multiplicative bias m . Determining appropriate priors on this quantity is a vital part of characterising a lensing catalogue. Considering a number of factors that could contribute systematic bias, which we quantified using the simulations, led us to a zero-centred Gaussian prior of width $\sigma_m \sim 2.5 \times 10^{-2}$ for Y1 IM3SHAPE.

The data we have presented here comprise only 20% of the full Dark Energy Survey. Work to analyse the subsequent two years of data has already begun, and to fully exploit that new dataset will require the methods described in this chapter to be refined in a number of ways.

Chapter 4

The Impact of Image Plane Neighbours on Shear Cosmology

There are many ways in which ensemble shape measurements may be biased relative to the underlying cosmic signal. In the previous chapter we presented a detailed guide to shape measurement, simulating and calibrating the estimates for Year 1 of DES. That process is largely agnostic to the mechanism and potential impact of bias; We are concerned with the practical matter of generating catalogues fit for use in cosmology analyses. In the following we present a more detailed examination of a particular source of bias: the impact of image plane neighbours. There is relatively little existing literature on this subject, and yet it will become increasingly important to understand the impact of blending, as lensing datasets become inexorably deeper and probe higher redshifts. The following study has been submitted for publication in Samuroff et al. (2017a).

4.1 Introduction

A standard and well tested prediction of General Relativity is that a concentration of mass will distort the spacetime around it, and thus produce a curious phenomenon called gravitational lensing. The most obvious manifestation is about massive galaxy clusters, where background galaxies can be elongated into crescent-shaped arcs. So-called strong lensing was first observed in the late 1980s and has been confirmed many times since. A far subtler, but from a cosmologist's perspective far more powerful, consequence of gravitational lensing is that background fluctuations in the density of dark matter will induce coherent distortions to photons' paths. This effect is known as cosmic shear, and it was first detected by four groups at around the same time close to two decades

ago (Bacon et al., 2000; Van Waerbeke et al., 2000; Kaiser, 2000; Wittman et al., 2000).

Cosmic shear has the potential to be the single most powerful probe in the toolbox of modern cosmology. The spatial correlations due to lensing are a direct imprint of the large scale mass distribution of the Universe. Thus it allows one to study the total mass of the Universe and the growth of structure within it (Maoli et al., 2001; Jarvis et al., 2006; Massey et al., 2007a; Kilbinger et al., 2013; Heymans et al., 2013; Dark Energy Survey Collaboration, 2016; Jee et al., 2016; Hildebrandt et al., 2016), or to map out the spatial distribution of dark matter on the sky (eg Kaiser 1994; Van Waerbeke et al. 2013; Chang et al. 2015). As a probe of both structure and geometry, cosmic shear is also attractive as a method for shedding light on the as yet poorly understood component of the Universe known as dark energy (Albrecht et al., 2006; Weinberg et al., 2013). Alternatively, lensing will allow us to place ever more stringent tests of our theories of gravity (Simpson et al., 2015; Harnois-Déraps et al., 2015; Brouwer et al., 2017). Unlike comparable probes it is theoretically very clean, responding directly to the power of dark matter, which is affected by baryonic physics only on very small scales, and avoids recourse to poorly-understood phenomenological rules such as the mass-observable scaling used in cluster cosmology. Indeed galaxy number density enters only at second order as a weighting of the observed shear due to the fact that one can only sample the shear field where there are real galaxies (Schmidt et al., 2009).

Though well modelled theoretically, cosmic shear is technically highly challenging to measure; as with all these probes it is not without its own sources of systematic error. It also cannot be reiterated too many times that the shear component of even the most distant galaxy's shape is subdominant to noise by an order of magnitude. Indeed, the ambitions of the current generation of cosmology surveys will require sub-percent level uncertainties (both systematic and statistical) on what is already a tiny cosmological signal $g \sim 0.01$.

It was realised early on how significant the task of translating photometric galaxy images into unbiased shear measurements would be. In response came a series of blind shear measurement challenges, designed to review, test and compare the best methods available. The first of these, called STEP1 (Heymans et al., 2006) grew out of a discussion at the 225th IAU Symposium in 2004. The exercise was based around a set of simple SKYMAKER simulations (Bertin & Fouqué, 2010), which were designed to mimic ground based observations but with analytic galaxies and PSFs and constant shear. The algorithms at this point represented a first wave of shear measurement codes and included several moments-based algorithms (Kaiser, 1995; Kuijken, 1999; Rhodes et al., 2001), some early forwards modelling methods (Bridle et al., 2002), as well as a technique called shapelets, which models a light profile as a set of 2D basis functions (Bernstein & Jarvis, 2002; Refregier & Bacon, 2003).

The simulations and the codes themselves steadily grew in complexity. STEP2 was followed by series of GREAT challenges (Massey et al., 2007b; Bridle et al., 2008; Kitching et al., 2010; Mandelbaum et al., 2014), which focused on different aspects of shape measurement bias and have been essential in quantifying a number of significant effects. In recent years the drive to find ever more accurate ways to measure shear has intensified, with many novel approaches being suggested. For example Fenech Conti et al. (2016) use a form of self-calibration, which repeats the shape measurement on a test image based on the best-fitting model for each galaxy. A related approach, named metacalibration, involves deriving a shear response directly from the data, using counterfactual copies of the image with additional shear (Huff & Mandelbaum, 2017; Sheldon & Huff, 2017). More advanced moments-based approaches include the BFD method (Bernstein & Armstrong, 2014), which derives a prior on the ensemble ellipticity distribution using deeper fields, and SNAPG (Herbonnet et al., 2017), a similar approach which builds ensemble shear estimates using shear nulling.

The following is intended as a companion study to Zuntz et al. (2017) (Z17 henceforth), where we present two shear catalogues derived from DES Y1 dataset. Containing over 24 million galaxies, these catalogues are the product of two independent maximum likelihood codes. The first, called IM3SHAPE, implements simultaneous fits using multiple models and we calibrate externally using simulations. The second implements a Gaussian model fitting algorithm, NGMIX, supplemented by shear response corrections using metacalibration. Whereas that paper focuses on the catalogues themselves, presenting a raft of calibration tests and a broad overview of the value-added data products, here we use the same resources to explore a narrower topic: the impact of image plane neighbours on shear measurement. Specifically we use the image simulations described in Z17, from which the Y1 IM3SHAPE calibration is derived, to explore the mechanisms for neighbour bias, and then propagate the results to mock shear two-point data to investigate the consequences for weak lensing cosmology.

This chapter is structured as follows. we first briefly review the formalism of lensing, and the observables discussed in this work. In Section 4.3 we present a series of numerical calculations using a toy model to characterise neighbour bias. Section 4.4 describes the simulated DES Y1 datasets, generated using our HOOPOE simulator. we test the earlier predictions under more typical observing conditions in Section 4.5, and extend them into a quantitative set of results using the more extensive Y1 HOOPOE dataset. Section 4.6 then presents a numerical analysis designed to test the cosmological implications of neighbour bias of the nature and magnitude found in our simulations. we conclude in Section 4.7.

4.2 The Shear Measurement Problem

The problem of shape measurement is far more intricate than it might first appear. Any cosmological analysis based on cosmic shear is reliant on a series of technical choices, which can have a non-trivial impact on measurement biases, precision and cosmological sensitivity. Specifically we must choose (a) how to parameterize each galaxy’s shape, (b) which measurement method to use to estimate that shape, (c) what selection criteria are needed to obtain data of sufficiently high quality for cosmology and (d) how biased is the measurement and what correction is needed? These choices should be made on a case-by-case basis, since the optimal solutions are dependent on a number of survey-specific factors. we discuss each briefly in turn below.

4.2.1 Shape Measurement with IM3SHAPE

The shape measurements upon which the following analyses are based make use of the maximum likelihood model fitting code IM3SHAPE¹ (Zuntz et al., 2013). It is a well tested and understood algorithm, which has since been used in a range of lensing studies (Dark Energy Survey Collaboration, 2016; Whittaker et al., 2015; Kacprzak et al., 2016; Clampitt et al., 2016). It was also one of two codes used to produce shear catalogues in the Science Verification (SV) stage and Year 1 of the Dark Energy Survey. We refer the reader to Jarvis et al. (2015) (hereafter J16) and Z17 for the most recent modifications to the code.

we use the definition of the flux signal-to-noise ratio of Z17, J16 and Mandelbaum et al. (2015):

$$S/N \equiv \frac{\left(\sum_{i=1}^{N_{\text{pix}}} f_i^m f_i^{\text{im}} / \sigma_i^2 \right)}{\left(\sum_{i=1}^{N_{\text{pix}}} f_i^m f_i^m / \sigma_i^2 \right)^{\frac{1}{2}}}. \quad (4.1)$$

The indices $i = (1, 2 \dots N_{\text{pix}})$ run over all pixels in a stack of image cutouts at the location of a galaxy detection. The model prediction and observed flux in pixel i are denoted f_i^m and f_i^{im} respectively and σ_i is the RMS noise. This signal-to-noise measure is designed to favour maximal agreement between the model and the image pixel fluxes. Note that if the best-fitting model f^m is identical for two different postage stamps, S/N will favour the image with the greater total flux.

A useful size measure is defined as

¹<https://bitbucket.org/joezuntz/im3shape-git>

$$R_{gp}/R_p \equiv \frac{P[R_g]}{R_{\text{PSF}}}, \quad (4.2)$$

or the Full Width at Half Maximum (FWHM) of the galaxy after PSF convolution, normalised to the PSF FWHM. The operation P maps the size of the pre-convolution galaxy R_g onto that of the PSF convolved one, and typically does not have a simple analytic form. Single-number size estimates are obtained from asymmetric multi-exposure data by circularising (azimuthally averaging) the best-fitting profile for each galaxy and then computing the weighted quadrupole moments of the resulting image.

4.2.2 Shear Measurement Bias

There are many ways bias can enter an ensemble shear estimate based on a population of galaxies. Although the list is not exhaustive, a handful of mechanisms are particularly prevalent, and have been extensively discussed in the literature.

- **Noise Bias:** On addition of pixel noise to an image, the best-fitting parameters of a galaxy model will not respond linearly. This is as an estimator bias as much as a measurement bias, and results in an asymmetric, skewed likelihood surface (Hirata & Seljak, 2003; Refregier et al., 2012; Kacprzak et al., 2012; Miller et al., 2013). Any code which uses the point statistics of the distribution (either mean or maximum likelihood) as a single-number estimates of the ellipticity results in a bias. This is true even in the idealised case where the galaxy we are fitting can be perfectly described by our analytic light profile. The bias is sensitive to the noise levels and also the size and flux of the galaxy, and thus is specific to the survey and galaxy sample in question. For likelihood-based estimates one solution would be to impose a prior on the ellipticity distribution and propagate the full posterior. However, the results can become dependent on the accuracy of that prior, and such codes require cautious testing using simulations (Bernstein & Armstrong, 2014; Simon & Schneider, 2016)
- **Model Bias:** In reality galaxies are not analytic light profiles with clear symmetries. For the purposes of model-fitting, however, we are constrained to use models with a finite set of parameters. A model which does not allow sufficient flexibility to capture the range of morphological features seen in the images will produce biased shape measurements (Lewis, 2009; Voigt & Bridle, 2010; Kacprzak et al., 2013).

- **Selection Bias:** Even if we were to devise an ideal shape measurement algorithm, capable of perfectly reconstructing the histogram of ellipticities in a certain population of galaxies, our attempts to estimate the cosmological shear could still be biased. If any part of the process of converting photometric survey images to galaxy shape catalogue prefers round objects or those aligned in a particular direction, the result would be a net alignment that could be mistaken as having cosmological origin. In practice such preferment commonly results from imperfect correction of PSF asymmetries (eg Kaiser 2000; Bernstein & Jarvis 2002), and the simple fact that many detection algorithms fail less frequently on rounder galaxies, and so the detection sample will tend to underrepresent cases where the shear and intrinsic pre-lensing shape are aligned (Hirata et al. 2004). It is such effects that make post facto quality cuts on quantities such as signal-to-noise or size (both of which naturally correlate with ellipticity) particularly delicate.

4.3 Toy Model Predictions

To develop a picture of how image plane neighbours affect shear estimates with IM3SHAPE, we build a simplified toy model. Using GALSIM² we generate a 48×48 pixel postage stamp containing a single exponential disc profile convolved with a tiny spherically symmetric PSF (though we confirm that our results are insensitive to the exact size of the PSF). One can then apply a small shear along one coordinate axis prior to convolution and use IM3SHAPE to fit the resulting image. In the absence of noise or model bias the turning point in the likelihood of the measured parameters coincides exactly with the input values. The basic setup then has four adjustable parameters: the flux and size of the galaxy plus two ellipticity components, denoted f_c , r_c , g_1^{tr} and g_2^{tr} . At this stage we do not model miscentering error between the true galaxy centroid and the stamp centre. Unless otherwise stated we fix these to the median values measured from the DES Y1 IM3SHAPE catalogue.

It is worth noting that neither this basic model nor the more complex simulations that follow accurately model the fact that the cosmological shear of close neighbours will be correlated at some level (due to the shared line of sight). Nor do we model correlations between the intrinsic shapes of neighbours which are both close together in the image plane and at similar radial distances. These spatial correlations will likely amplify the impact of blending, and are worthy of future investigation. This is, however, likely a second-order effect of neighbours, and we postpone such

²<https://github.com/GalSim-developers/GalSim>

study to a future date.

4.3.1 Single-Galaxy Effects

To explore the interaction in single neighbour-galaxy instances we next draw a second galaxy into the postage stamp, convolved with the same nominal PSF. This adds four more model parameters: neighbour size, flux, radial distance from the stamp centre and azimuthal rotation angle relative to the x coordinate axis. We will call these r_n , f_n , d_{gn} and θ . At this stage the neighbour has zero ellipticity.

We show this setup at four neighbour positions in Figure 4.1. Under zero shear, the system has perfect rotational symmetry, and the measured ellipticity magnitude $\tilde{g}(\theta|g_1^{tr} = 0)$ is independent of θ^3 . As a first exercise, we generate a circular central galaxy with a circular Gaussian neighbour, which is gradually shifted outwards from the stamp centre. Following the usual convention for galaxy-galaxy lensing, tangential shear is defined such that it is negative when the major axis of the measured shape is points radially towards the neighbour. The measured two-component ellipticity shown by the solid and dot-dashed lines in Figure 4.2a.

The decline to zero at small separations is understandable, as there is no reason to expect drawing one circular profile directly on top of another should induce spurious non-zero ellipticity. In the regime of a few pixels, however, strong blending can increase the flux of the best-fitting model.

Next, we repeat the calculation, now applying a moderate cross-component shear to the neighbour ($g_2 = 0.1$). The result is shown by the blue lines in Figure 4.2. Unsurprisingly the measured tangential shear is unaffected by a true shear along an orthogonal axis. In cases where the objects share a large portion of their half-light radii, we are fitting a strongly blended pair with a single profile, and the neighbour/central distinction becomes difficult to define. The best-fitting ellipticity recovered from the blended image is not a pure measurement of either galaxy's shape; rather it is a linear combination of the two. We repeat the zero-offset measurement using a range of neighbour fluxes and find that the best-fitting e_i follows roughly as a flux-weighted sum over the two galaxies $\tilde{g}_i \approx (f_c g_{i,c}^{tr} + f_n g_{i,n}^{tr}) / (f_n + f_c)$.

³Although the zero-shear ellipticity is, of course, sensitive to the other model parameters. For complete rigour we would include the other toy model parameters parameters on the right of the | and write the measured shape as $\tilde{g}(\theta|g^{tr}, d_{gn}, f_n, r_n, f_c, r_c)$. For the sake of clarity we'll omit the other parameters, under the assumption that they have the fiducial values noted above, unless explicitly stated otherwise.




Figure 4.1: Postage stamp snapshots of the basic two-object toy model described in Section 4.3. The overlain ellipse shows the maximum likelihood fit to the image. The panels show three neighbour positions in the range $\theta = [0, \pi/2]$ rad. The best fit ellipticity and half light radius are shown above each image. In all cases the input values are $\mathbf{e} = (0, 0)$, $r = 0.5$.

4.3.2 Ensemble Biases

While useful for understanding what follows, the impact of neighbours on individual galaxy instances is not particularly informative about the impact on cosmic shear measurements. Even significant bias in the per-object shapes could average away over many galaxies with no residual impact on the recovered shear. More important is the *collective* response to neighbours. To explore this we build on the toy model concept. To estimate the ensemble effect, we measure a neighbour-central image repeatedly at 70 positions on a ring of neighbour angles. Again, under zero shear $g^{tr} = 0$ the measured shape is constant in magnitude, and simply oscillates about 0 with peaks of amplitude $|\tilde{g}(\theta|0, d_{gn})|$. This sinusoidal variation is shown by the dotted lines in Figure 4.2b at two values of d_{gn} (7 and 8 pixels). By averaging over a (large) number of neighbours one is effectively marginalising over θ , which results in an unbiased measurement of the shear $\langle \tilde{g}(\theta|g^{tr} = 0, d_{gn}) \rangle_{\theta} = g^{tr} = 0$. A non-zero shear $g^{tr} \neq 0$, however breaks the symmetry of the system. A galaxy sheared along one axis will not respond to a neighbour in the same way irrespective of θ , which can result in a net bias. To show this we fix $g^{tr} = -0.05$ and proceed as before. The solid lines in Figure 4.2b show the periodicity in the measured shear at two d_{gn} . The mean value averaged over θ is shifted incrementally away from the input shear, shown by the horizontal dot-dashed line. Specifically we should note that the peaks below g^{tr} at $\pi/2$ and $3\pi/2$ radians are deeper and narrower than those above it. To understand how this arises, we refer the reader to the cartoon diagram in Figure 4.3. The purple lines are iso-light contours in a strongly sheared Sérsic disc profile ($g_1 = -0.3$). Clearly rotating the neighbour from position A to C carries it from

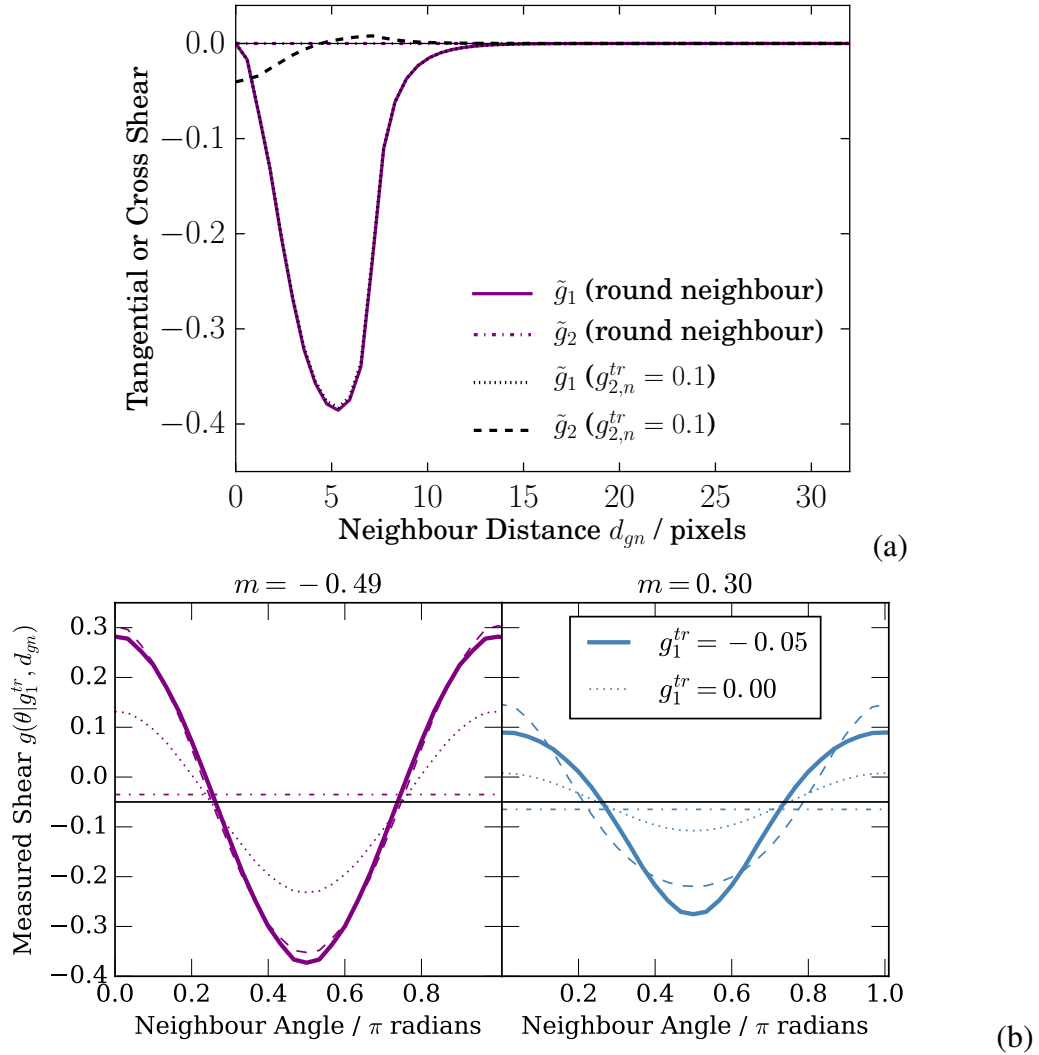


Figure 4.2: **Top:** Tangential shear measured using the numerical toy model described in Section 4.3.1 as a function of radial neighbour distance. The solid purple line shows the shape component aligned with the central-neighbour separation vector and the dot-dashed line is measured along axes rotated through 45° . Note that the latter is smaller than 10^{-6} at all points on this scale. The dashed and dotted black lines show the same ellipticity components when the neighbour is sheared in the e_2 direction by $g_2 = 0.1$. **Bottom:** Best-fit galaxy ellipticity as a function of neighbour position angle at fixed neighbour distance d_{gn} from the toy model described in the text. The two panels (left, right) show the same central-neighbour system ($g^{tr} = -0.05$), but with different d_{gn} (7 and 8 pixels) and biases m (shown above each panel). The solid line in each case is the recovered galaxy shape at each θ , and the integrated mean along this range is shown by the horizontal dot-dashed line. The dotted lines show the zero-shear shape (i.e. the ellipticity that would be measured if the input shear were zero), but shifted downwards such that the mean is at -0.05 . Finally, to illustrate the asymmetry of the system we show the solid line flipped about $y = g_1^{tr}$ and shifted by $\pi/2$ radians as a dashed curve.




Figure 4.3: Cartoon diagram of a neighbour-central system. The purple contours show the lines of constant flux in a Sérsic disc profile with extreme negative ellipticity ($g_1 = -0.3$). The blue crosses labelled A, B and C are points on a ring of equal distance from the centre of the profile. The blue arrows show the local unit vector along a tangent to the ring.

the relatively flat low wings of the central galaxy’s light profile closer to the core. Perturbing an object about C by a small angle results in a much greater change in the local gradient, $\nabla f_c(x, y)$ than doing the same about A. All other parameters fixed, an incremental shift along the blue tangent vector will have a larger impact at $\theta = 0$ than at $\pi/2$, resulting in asymmetry in the width of the positive and negative peaks in Figure 4.3. The depth of the peak can be explained qualitatively by similar arguments. At C a neighbour of given flux is closer to the centre of the light distribution and thus has a greater flux overlap with the central galaxy than at A. Naturally, then, one might expect neighbour A to have less impact than C. Returning to Figure 4.2, one can see that the two effects are in competition. Depending on the exact neighbour configuration, the simultaneous narrowing and deepening the negative peaks can result in a bias in the neighbour-averaged ellipticity towards large or small values.

The level of this effect will clearly correlate with the magnitude of the shear, and so induce a multiplicative bias. To illustrate this point the above exercise is repeated with a range of different input shears. The results for our fiducial setup are shown in in Figure 4.4. Each point on these axes corresponds to a ring of neighbour positions for a given input shear. The equivalent measurements without the neighbour are indistinguishable from the x axis. At small shears, the neighbour induced bias $\tilde{g} - g^{tr}$ is well approximated as a linear in g^{tr} (shown in the inset panel). We leave exploration of the possible nonlinear response at large ellipticities for future investigations. Though the above numerical exercise demonstrates that it is *possible* for significant multiplicative bias to arise as a result of neighbours, it does not make a clear prediction of the magnitude or even the sign. Indeed, our toy model is effectively marginalised over θ , but there is nothing to guarantee that fixing the other neighbour parameters to the median measured values is representative of the real level of neighbour bias in a survey like DES. Motivated by this observation we add a final layer




Figure 4.4: Measured shear minus input shear plotted as a function of input shear. The purple points show the recovered \tilde{g}_1 from averaging over ring of 70 neighbour positions. The light purple lines show the linear relation $\tilde{g} - g^{tr} = m g^{tr}$ at $m = (-0.4, -0.45, -0.5)$. The dotted line shows what would be measured using the same central profile in the absence of the neighbour, and is near indistinguishable from the x axis line on all points within this range of g^{tr} .

of complexity to the model, as follows. A single neighbour-central realisation is created as before, defined by a unique set of model parameters. Now, however, the values of those parameters $\mathbf{p} = (d_{gn}, f_n, r_n, f_c, r_c)$ are drawn randomly from the DES data. As these quantities will, in reality, be correlated we sample from the 5-dimensional joint distribution rather than each 1D histogram individually. We then fit the model at 70 neighbour angles and two input shears $g_{\pm} = \pm 0.05$ (a total of 140 measurements), and estimate the multiplicative bias as a 2-point finite-difference derivative:

$$m + 1 = \frac{\langle \tilde{g}(\theta|g_+) \rangle_{\theta} - \langle \tilde{g}(\theta|g_-) \rangle_{\theta}}{g_+ - g_-}. \quad (4.3)$$

This process is repeated to create 1.33 M unique toy model realisations. Binning by neighbour distance one can then make a rough prediction for the level of neighbour-induced bias and the angular scales over which it should act. The result is shown in Figure 4.5, where full results using all model realisations are indicated by the dashed blue line. The majority of cases yield a negative bias, particularly at low neighbour separation (referring back to Figure 4.3, the broadening of the peak around position A dominates over the increased flux overlap at C). In the real data, of course, we apply a quality based selection and überseg object masking (J16), both of which are neglected here. One can, however, test the impact of selecting on fitted quantities that respond to neighbour

bias. Imposing a flat prior on the centroid offset $\Delta r_0 = (x_0^2 + y_0^2)^{\frac{1}{2}}$ (i.e. discarding randomly generated model realisations where the galaxy centroid is displaced from the stamp centre by more than a fixed number of pixels) changes the shape of this curve significantly, as illustrated by the thick purple line.

One can understand the difference between the results with and without the centroid cut as a form of selection bias, whereby the cut preferentially removes toy model realisations in which the neighbour is bright relative to the central galaxy. At any given d_{gn} we are left with a relative over-representation of galaxies with $f_n/f_c \ll 1$. Faint neighbours, which in reality tend to be compact high redshift objects, have little impact when they sit on the outskirts of the central profile (A in the cartoon picture in Figure 4.3; the regime which produces negative m). The same faint galaxy has a stronger impact if it is rotated to a position closer to the centre of the central's flux profile. Thus one might expect a selection on Δr_0 to make the mean m in a particular bin less negative (or even positive) by preferentially removing brighter galaxies.




Figure 4.5: Multiplicative bias estimated using the Monte Carlo toy model described in the text. For each neighbour realisation, defined by a particular distance, flux and size we compute the average of the measured ellipticity components over 70 rotations on a ring of neighbour angles. To estimate the bias we perform this averaging twice at two non-zero shears, g_+ and g_- , and compute the finite-difference derivative using equation 4.3. The dashed thin blue line shows the result of using all measurements, while the bold purple line has a cut based on the offset between the centroid position of the best-fitting model and the stamp centre.

4.4 A More Realistic Simulation

In this section we provide a brief overview of the simulation pipeline. The process is the same as that described in §5 of Z17, and we refer the reader to that work for more detail. The end point of the pipeline is a cloned set of survey images with many of the observable characteristics of a chosen set of parent images, but for which we know the input noise properties and galaxy population perfectly. The simulated images inherit the pixel masking, PSF variation and noise maps measured from the progenitor data. Each simulated galaxy is then inserted into a subset of overlapping exposures and into the coadd at the position of a real detection in the DES Y1 data. Object detection is rerun on the new coadd images and galaxy cutouts and new segmentation masks are extracted and stored in the MEDS format described by J16. The mock survey footprint is shown in Figure 4.6. In the lower panels we show an example of a simulated coadd (left) and the spatial variation in PSF orientation within the same image (right).

4.4.1 Parent Data

We use reduced images from Year One of the Dark Energy Survey (DES Y1) as input to the simulations discussed in this thesis. The Dark Energy Survey is undertaking a five year programme with the ultimate aim of observing ~ 5000 square degrees of the southern sky to ~ 24 th magnitude in five optical bands, *grizY*, covering $0.40 - 1.06$ microns. The dataset is recorded using a 570 megapixel camera called DECam (Flaugher et al., 2015), which has a pixel size of 0.26 arcseconds. In full it will consist of ~ 10 interwoven sets of 90s exposures in the *g*, *r*, *i* and *z* bands and 45s in *Y*.

The Y1 data were collected over August 2013 and February 2014, and cover a substantially larger footprint than the preliminary Science Verification (SV) stage at 1500 square degrees, albeit to a reduced depth. Details of the reduction and processing are presented in Z17. Our HOOPOE simulations use a selection of the total 3000 0.75×0.75 degree coadded patches known as “tiles”.

4.4.2 Input Galaxy Selection

For populating the mock survey images HOOPOE we use a sample of real galaxy profiles from the HST COSMOS field, imaged at significantly lower noise and higher resolution than DES by the Hubble Space Telescope Advanced Camera for Surveys (HST ACS) (Scoville et al., 2007). It extends significantly deeper than the Y1 detection limit of $M_{r,\text{lim}} = 24.1$, extending to roughly

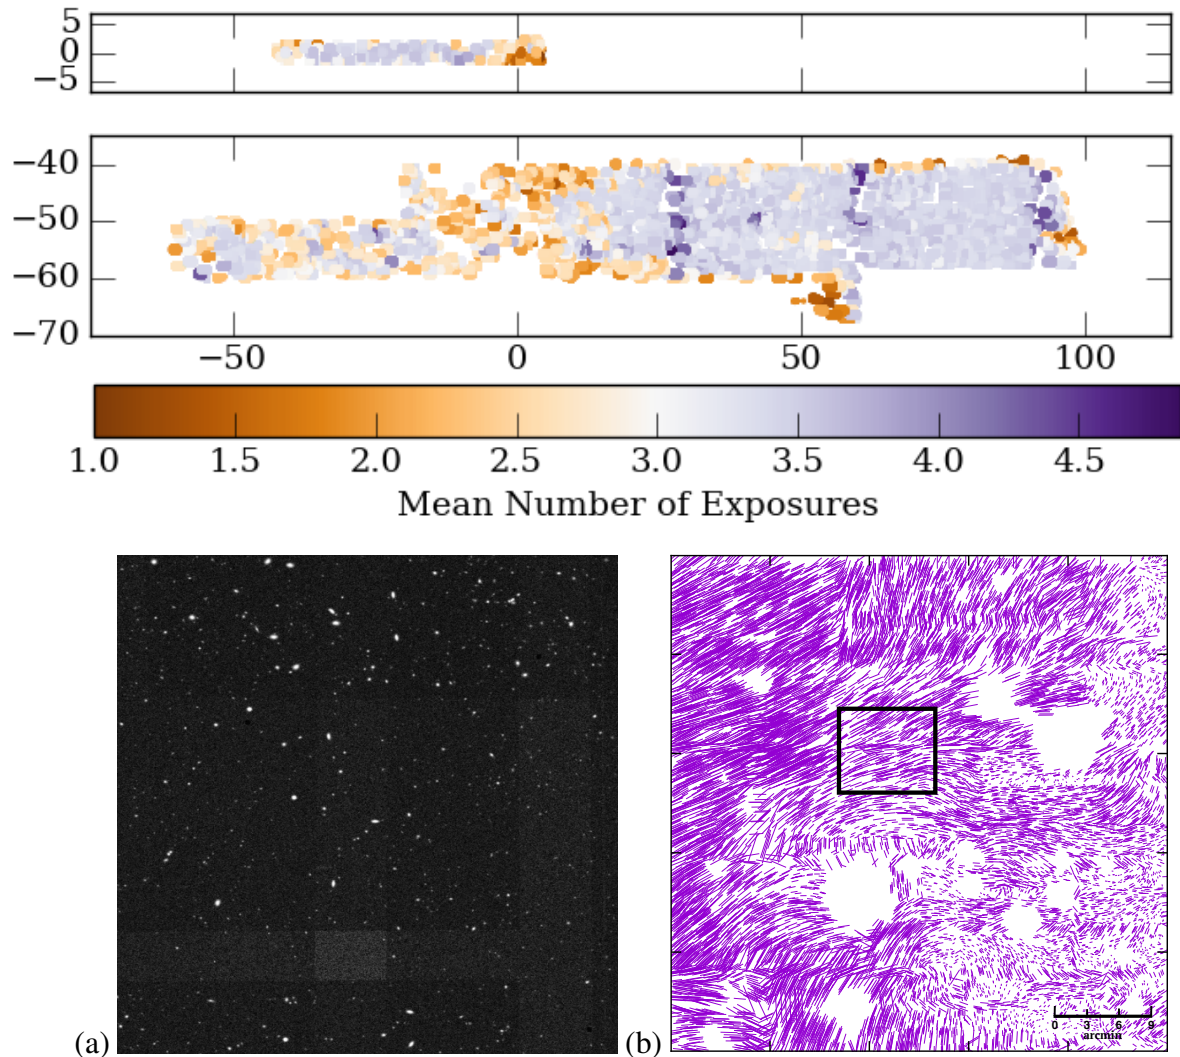


Figure 4.6: **Top:** Cartesian projection of the simulated survey footprint. The axes shown here are the right ascension and declination in units of degrees. Each pixel represents a simulated galaxy coloured according to the mean number of exposures in that tile. The full simulation comprises 1824 0.75×0.75 degree tiles drawn randomly from the DES Y1 area. **Bottom:** Images of a random tile selected from the HOOPOE area. The left panel (a) shows a square subregion of the simulated image plane for coadd tile DES0246-4123 of approximately 9×9 arcminutes. The right hand panel (b) shows a PSF whisker plot covering the full 0.75×0.75 tile. The length (orientation) of the lines represents the magnitude (position angle) of the spin-2 PSF model ellipticity at each simulated galaxy position. Note that we plot only galaxies which pass IM3SHAPE quality cuts. The circular patches in the right-hand panel are example of spatial masking inherited from the DES GOLD catalogue, and correspond to the positions of bright stars in the parent data.

27.9 mag in the SDSS r -band. A main sample for our DES Y1 simulations is defined by imposing a cut at < 24.1 mag.

Since the DES images do not cut off abruptly at 24th magnitude, in reality they contain a tail of fainter galaxies which contribute flux but may not be identifiable above the pixel noise. To approximate the impact of these objects we simulate a population of sub-detection galaxies in addition to the main sample. In brief we use the full histogram of COSMOS magnitudes to estimate the number of faint galaxies within a given tile. The required profiles are selected randomly from the faint end of the COSMOS distribution. Each undetected galaxy is paired with a detection, and inserted at a random location within the overlapping bounds of the same (subset of) single-exposure images. A more detailed description of this process can be found in Z17.

If these galaxies were present in the data they would enter the background flux calculation, and thus the subtraction applied would change due to their presence. Since the simulation pipeline produces images effectively in a post-background subtraction state this is not captured by construction. To test this we rerun the SExtractor background calculation on a handful of tiles drawn with and without the faint galaxies. The impact was found to be well approximated as a uniform shift in the background correction. A flux correction equal to the pixel-averaged flux of the subdetection galaxies over each image plane is, then, applied to postage stamps prior to shape measurement.

In reality the overdensity of sub-threshold galaxies will be coupled to the density of detectable objects, which is clearly not the case in our simulations. To gauge the impact of this we perform the following test. Each tile is divided into a 6×6 grid, and the mean multiplicative bias is calculated in each sub-patch. We bin sub-patches according to the ratio $f_{\text{faint}} \equiv N_{\text{faint}}/N_{\text{det}}$, or the total number of faint galaxies relative to the number of detectable ones. The impact is significant, but not leading order; excluding patches outside the range $0.9 < f_{\text{faint}} < 1.1$ induces a shift of $\Delta m \sim -0.005$.

An independent noise realisation is generated for each exposure using the weight map from the parent data. We reify the noise in each pixel by drawing from a Gaussian of corresponding width. The coaddition process is not rerun, but rather we compute an independent noise field by drawing the flux in each pixel from a zero-centred Gaussian of width determined by the measured variance in that pixel.

4.4.3 Neighbour-Free Resimulations

For the purpose of untangling the impact of image plane neighbours we use the simulated HOOPOE images to create a new spin-off dataset. In a subset of a little over 500 tiles we store the (convolved)




Figure 4.7: An example of an object in the main DES Y1 calibration simulation and the neighbour free resimulation. The upper panels show the coadd pixel data for this object in the original simulated images (left, labelled HOOPOE) and in the neighbour-subtracted resimulation (right, labelled WAXWING). The lower panels are the segmentation masks for the same galaxy. We choose this object as an example as it has multiple neighbours, both masked (upper left and centre left) and unmasked (lower right).

input profile for each object and the noise-only cutout, taken from the same position in the image plane prior to objects being drawn. By adding together these two components one can generate a suite of spin-off MEDS files, which are equivalent to the results of a simpler neighbour-free simulation (e.g. Miller et al. 2013, J16). The pixel noise realisation, COSMOS selection and input shears, however, are identical to the progenitor HOOPOE simulations.

We will call this process “resimulating”, and the basic concept is illustrated in Figure 4.7. The 506-tile set of neighbour free data are named the WAXWING resimulations. Finally the (now empty) segmentation masks corresponding to the subtracted neighbours are also removed. In subsequent IM3SHAPE runs on these data we ignore the SExtractor flags obtained from the main simulations.

4.5 Quantifying Neighbour Bias with HOOPOE

Equipped with qualitative predictions from Section 4.3, we now turn to a more complete simulation. The mock survey was designed to capture as much of the complexity of shape measurements on real photometric data as possible. We refer to Section 4.4 of this chapter for a short overview and to §5 of Z17 for a more detailed discussion of the simulation pipeline and validation tests. The simulated galaxy catalogue used in the following is identical to the one used to calibrate the DES Y1 IM3SHAPE catalogue, including quality cuts and selection masks.

4.5.1 Single-Galaxy Effects

The most straightforward way to assess the impact of neighbours on individual shape measurements in our simulations is to rotate the measured shapes into a frame defined by the central-neighbour separation vector. Whereas in the earlier toy model we had only one neighbour per galaxy, we now have a crowded image plane containing many objects simultaneously. For simplicity, in the earlier case we included no masking. For HOOPOE we wish to mimic the process of shape measurement on real data as closely as possible. We generate new segmentation maps by running SExtractor on the simulated images, and incorporate them into our shape measurements using the überseg algorithm (J16). Each simulated galaxy is allocated a nearest neighbour using a k -d tree matching algorithm constructed on the coadd pixel grid using every galaxy simulated at r -band magnitude $M_r < 24.1$. Note that the neighbours pairing includes all simulated galaxies, not just those that pass quality cuts. The quantities d_{gn} and θ are now redefined slightly as nearest-neighbour distance and angle. We define the tangential shear of a galaxy relative to its nearest neighbour using the standard convention,

$$e_+ = - [e_1 \cos(\theta) + e_2 \sin(\theta)], \quad (4.4)$$

and the cross shear

$$e_\times = - [e_2 \cos(\theta) - e_1 \sin(\theta)]. \quad (4.5)$$

Note that negative values of e_+ imply a net tangential alignment of the measured shapes towards neighbours. By analogy, we define $e_{1,n}$ and $e_{2,n}$, which are the measured ellipticity components, rotated into a reference frame defined by the major axis of the neighbour. Non-zero $e_{i,n}$ would indicate leakage of the neighbour's shape into the measurement, which might conceivably be induced by inadequate deblending of very close neighbours or by extensive non-circular masking. We first divide the main simulated catalogue into bins according to d_{gn} , and measure the tangential shear about nearest neighbours in each bin. The result is shown by the purple curve in Figure 4.8. Note that the statistical uncertainty, if plotted on this scale, is within the width of the line in all bins. The results here show qualitative agreement with the numerical predictions in Figure 4.2. As we found earlier, the exact shape of this curve is sensitive to the properties of both the neighbour and the central galaxy. Despite small differences, the range of variation is comfortably within the scale of the postage stamp for the bulk of galaxies in DES Y1. Repeating the measurement, rotated into the plane of the neighbour shape results in the dotted and dot-dash lines in this figure. As there are not




Figure 4.8: Tangential shear around image plane neighbours in the full HOOPOE simulation. The purple solid line shows the mean component of the measured galaxy shapes radial to the nearest image plane neighbour. The component rotated by 45° , which we have no reason to expect should be non-zero, is shown by the dot-dashed line. The dotted and dashed black lines are the measured ellipticity components when rotated into a coordinate frame defined by the major axis of the neighbour. The inset panel covers the same x range, with a magnified y axis.

necessarily reliable ellipticity measurements for each neighbour, we instead use the sheared input ellipticities. Both components of $e_{i,n}$ are seen to be negligible over all scales.

4.5.2 Neighbour Ensemble Biases

As we suggested in Section 4.3, the pertinent question for lensing cosmology is not whether the single galaxy ellipticities are offset relative to the underlying truth, but whether we are able to accurately recover an input shear by averaging over many galaxies. To explore this we again divide the catalogue into bins according to neighbour distance, but we now fit for multiplicative and additive bias in each bin. We show the measurement as the purple points in Figure 4.9, which can be compared with the earlier numerical model prediction in Figure 4.5. The horizontal band on these axes shows the 1σ mean m measured using all galaxies in the HOOPOE catalogue, and sits at $m \sim -0.12$. We note a steeper decline than in the bold line (without the centroid cut), more akin to the case with the centroid cut ($\Delta r_0 < 1$ arcsec). This is not surprising given that the INFO_FLAG quality selection includes exactly this cut. We do not report a local peak at ~ 11 pixels, which we saw in Figure 4.5. We suggested previously that effect was the result of positive m in galaxies where the nearest neighbour is relatively faint and at middling distance. It is likely that many of these objects manifest themselves as large changes in other quantities to which IM3SHAPE’s INFO_FLAG is sensitive such as ellipticity magnitude and fit likelihood, or are flagged by the SEXTRACTOR deblending cuts.




Figure 4.9: Multiplicative bias as a function of separation from the nearest image plane neighbour. The purple points show the bias calculated in bins of neighbour distance using the main HOOPOE simulated shape catalogue. The coloured bands show the same dataset divided into four equal-number bins according to the r -band magnitude of the neighbour. As shown in the legend, the median values in the four bins are 21.5, 23.0, 23.5 and 24.0. The mean bias and its uncertainty across all distance bins is indicated by the horizontal band.

When divided into broad bins according to the r -band magnitude of the nearest neighbour $M_{r,\text{neigh}}$ (the coloured stripes in Figure 4.9) we find the surviving objects show relatively weak dependence on neighbour brightness, except at the smallest neighbour distances, where bright neighbours have a slightly stronger (negative) impact than faint ones.

We measure the additive bias components in the same bins, but find no systematic variation with d_{gn} above noise.

Finally, we show the analogous measurement in bins of galaxy magnitude in Figure 4.10. The steep inflation of $|m|$ at the faint end of this plot has been seen elsewhere (e.g. Zuntz et al. 2017; Fenech Conti et al. 2016), and is easily understandable as the result of noise bias. We find that splitting by neighbour magnitude does not reveal any obvious trend here.

4.5.3 Untangling the Knot of Neighbour Bias

A central plank of this analysis rests on a comparison of the main HOOPOE simulations with the neighbour-free WAXWING resimulations described in Section 4.4.3. The simplest comparison would be between multiplicative bias values, calculated using all galaxies in each catalogue after cuts. These values are shown by the two upper-most lines (purple) in Figure 4.11. The difference is an indicator of the net impact of neighbours through any mechanism, which we find to be $\Delta m \sim$




Figure 4.10: Multiplicative bias as a function of r -band magnitude. As in Figure 4.9 the four coloured bands represent equal number bins of neighbour magnitude. Purple points show the full catalogue, with no magnitude binning. The mean bias and its uncertainty are shown by the purple horizontal band.

−0.05.

To untangle the various contributions to this shift, we construct a matched catalogue. Galaxies in the overlap between HOOPOE and WAXWING (12 M galaxies over 183 square degrees) are matched by ID; quality cuts are calculated for each set of measurements (see Appendix E from Z17). GOLD catalogue geometric masking and SExtractor deblending flags are included for HOOPOE. Since the latter are irrelevant to WAXWING, we omit them from quality flags on that dataset. For conciseness we will refer to the two measurements as “matched HOOPOE” and “matched WAXWING”, and their cuts as “HOOPOE cuts” and “WAXWING cuts”. Since the images are identical in all respects, but for the presence of neighbours, the statistical noise on the *change* in measured quantities should be smaller than the face-value uncertainties.

The appropriate cuts are first applied to each catalogue, then the results are divided into equal number signal-to-noise bins and fitted for the multiplicative bias in each. The result is shown by the points in the upper left-hand part of Figure 4.12. The equivalent in bins of PSF-normalised size is shown on the right. The difference between the blue and the purple points gives an indication of the total effect of all neighbour-induced effects on m , indicated by the solid purple line in the lower panel.

The generic shift attributed to “neighbour bias” is in reality a collection of distinct effects. By comparing the matched catalogues we identify four main mechanisms: direct contamination, selection bias, S/N bin shifting and neighbour dilution. The following paragraphs seek to explain each of these in turn. For reference each of the components that we describe is shown by one of




Figure 4.11: Graphical illustration of the measured multiplicative bias in the various scenarios considered in this chapter. The upper two lines show the mean m in the main DES Y1 HOOPOE simulations and a spin-off neighbour-free resimulation named WAXWING, as described in Section 4.4.3. The middle section (green) shows results using only galaxies which appear in both the HOOPOE and WAXWING simulations. The matching process alone does not imply any quality-based selection function. The final three lines in red are from a similar matching between a smaller rerun of the simulation with and without sub-detection limit galaxies. See the text for details about each of these cases.

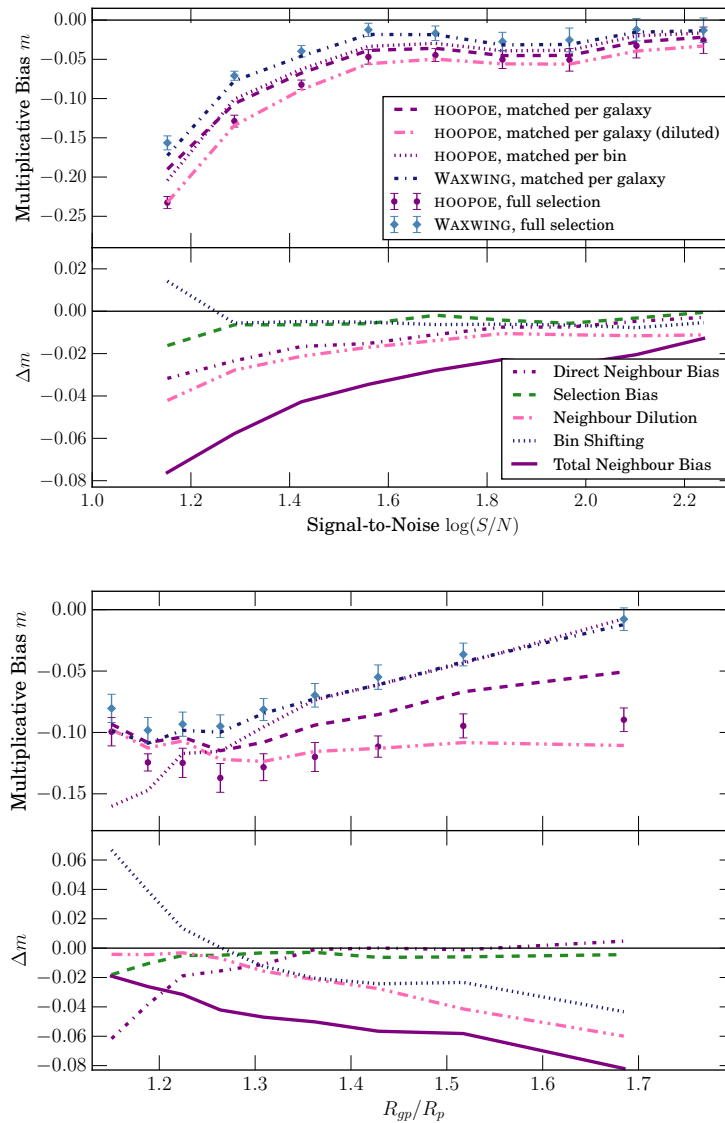


Figure 4.12: **Top half of each panel** Multiplicative bias as a function of signal-to-noise and size. The purple dots show the measured bias using the main simulation and the blue diamonds show the neighbour-free version. The result of applying both the HOPOE and WAXWING selection functions to the neighbour-free catalogue is indicated by blue dash-dot line. The downwards shift relative to the blue diamonds is the result of selection effects only. In purple dotted we show the results using this same galaxy selection but now with neighbours drawn, forcing each galaxy to fall into the same bins as before. The purple dashed line then shows the impact of relaxing the forced binning. Finally, the pink (dot-dot-dashed) line indicates “neighbour dilution” in the same galaxy selection, by adding back any galaxies which shift by more than $\sim 20\%$ in S/N or R_{gp}/R_p between the two catalogues, with random shears. **Bottom half of each panel** The change in bias due to the effects described above. The green (dashed) line shows the impact of selection effects only (the difference between the blue points and the dashed line in the top panel). Purple dash-dotted shows the direct impact of image plane neighbours on the measured shear (purple dotted minus blue dash-dot, top). The impact of shifting between S/N (or R_{gp}/R_p) bins is shown by the blue dotted (purple dashed minus purple dotted, top). The pink dot-dot-dashed line shows the neighbour dilution described above and in the text below. The solid line illustrates the full effect of neighbours, including all effects described.

the lines in Figure 4.12. For a visual summary of the various tests designed to isolate them see Figure 4.11.

Direct Flux Contamination

The most intuitive form of neighbour bias arises from the fact that, even after masking, neighbours contribute some flux to the cutout image of a galaxy. We term the effect of overlapping light profiles direct flux contamination. To gauge its impact we take the common sample of galaxies, which pass cuts in both datasets. The comparison is complicated somewhat by binning in measured S/N or R_{gp}/R_p ; for this test, we divide both sets of galaxies using the WAXWING-derived quantities. The resulting m measured using the HOOPOE galaxies is counterfactual, in the sense that we are binning measurements made in the presence of neighbours by quantities derived from neighbour-subtracted images. This exercise does, however, isolate the impact of the neighbour flux on the measured ellipticity. The result is shown by the purple dotted and purple dot-dashed lines in the upper and lower panels of Figure 4.12. The effect scales significantly with signal-to-noise and size. Faint small galaxies are affected strongly by neighbour light, while larger brighter ones are relatively immune.

Neighbour-Induced Selection Bias

We have noted already at the single-galaxy level that applying the HOOPOE selection mask to WAXWING can induce a link between the RMS ellipticity and distance to the nearest neighbour. To gauge the neighbour-induced selection effect, we take the WAXWING catalogue but now impose, in addition to its own quality cuts, the selection function derived from the with-neighbour HOOPOE dataset. The double masking removes an additional 0.5 M galaxies, which survive cuts in WAXWING but would be cut from the HOOPOE catalogue. The resulting change in m is shown by the dot-dash blue lines in the upper panels of Figure 4.12 (dashed green in the lower). The multiplicative bias arising from this cut is less than one percent in all but the faintest and smallest galaxies, where it can reach up to $m \sim -0.02$.

Bin Shifting

The above two tests encapsulate the impact on the measured ellipticities, and the selection flags from neighbour flux. An additional subtlety arises from the fact that the measured quantities used to bin galaxies (S/N and R_{gp}/R_p) are themselves affected by the presence of neighbours. To test

this we recalculate m using the same galaxy selection as in Section 4.5.3 (i.e. passing both sets of cuts), but now binned by the appropriate *measured* S/N . For clarity, the bin edges are unchanged, defined to contain equal numbers of WAXWING galaxies. The result is shown by the dashed lines in Figure 4.12. The difference compared with the case using fixed binning is purely the result of galaxies moving between bins. This shifting contributes multiplicative bias if one bins galaxies by observed quantities such as S/N , as we do in order to calibrate IM3SHAPE’s shear estimates. The amplitude of this is illustrated by the blue dotted line in the lower panels. Such neighbour-induced shifting is noticeable if we plot out the S/N of objects in HOOPOE against the S/N of the same objects in WAXWING. Objects which are strongly shifted in S/N are more likely to scatter upwards than downwards. A similar skew can be seen in the R_{gp}/R_p plane; when galaxies are scattered in size they tend to be thrown further and more often upwards than downwards. Small galaxies (which we know already are more strongly affected by noise bias) are shifted strongly upwards by the presence of neighbour flux in the HOOPOE images. The result is a net negative m added to the upper R_{gp}/R_p bins, and a simultaneous upwards shift in the lowest bins from which galaxies are lost. In the case of galaxy size we see the effects of bin scatter and direct neighbour bias almost negate each other, although the degree of cancellation is likely dependent on the specifics of the measurement code and the dataset.

Neighbour Dilution

A final point can be gleaned from Figure 4.12: that applying the WAXWING cuts to HOOPOE induces a shift in m . Naively one might expect the HOOPOE selection function, which includes neighbours, to remove the same galaxies as the WAXWING selection, plus some extra strongly blended galaxies. It is true that a sizeable number of galaxies are cut in the presence of neighbours, but would otherwise not be. There is also, however, a smaller population that survive cuts *because* they have image plane neighbours.

We can see this clearly from the fact that the purple points and the dashed purple lines Figure 4.12 are non-identical. We identify three separate (but partially overlapping) galaxy selections in this figure: (a) galaxies passing both sets of cuts, (b) galaxies passing cuts in the absence of neighbours, but cut by the HOOPOE selection and (c) galaxies which pass cuts in the presence of neighbours, but cut by the WAXWING selection. We find that populations (b) and (c) have much smaller mean neighbour separation than the full population (the histograms of d_{gn} show a sharp peak at under 10 pixels). In contrast, both the full catalogue and population (a) objects exhibit a much broader distribution ($\bar{d}_{gn} \sim 24$ pixels).

Based on the toy model predictions in Section 4.3 we set out a working proposal: that population (c), objects cut out only when neighbours are removed, are extreme blends dominated by a superposed neighbour. We will assume these objects are boosted considerably in size, S/N or both, such that what would otherwise be a small faint galaxy is now sufficiently bright to pass quality cuts. In these cases the measured shape of a simulated galaxy might be expected to be only weakly linked with the input ellipticity. To approximate this effect we take population (a) HOOPOE galaxies, subject to both sets of cuts, and add back some of the population (c) galaxies. Specifically, we include any objects shifted in S/N or R_{gp}/R_p by more than 20%. The true shears associated with these galaxies are now randomised to eliminate any correlation with the measured ellipticity. The result is shown as a pink dot-dot-dashed line in Figure 4.12. We can see that this effect, which we call neighbour dilution, to good approximation accounts for the residual difference between the population (a) and (c) samples. Particularly in the upper size bins of the right hand panel the differences are not eliminated entirely. This is thought to be the result of residual (albeit weakened) covariance between the measured shapes of strongly blended objects and the input shears. Clearly the scenario in which a neighbour totally overrides the original shape of a galaxy is extreme, and there will be an indeterminate number of moderate blends which are boosted sufficiently to survive cuts but which retain some correlation with their original unblended shapes. Such cases are, however, extremely difficult to model accurately with the resources available for this investigation.

4.5.4 Isolating the Impact of Subdetection Galaxies

A small number of previous studies have attempted to quantify the impact of galaxies below, or close to, the limiting magnitude of the survey. For example, Hoekstra et al. (2015) and Hoekstra et al. (2016) suggest they can induce a non-trivial multiplicative bias, which is dependent on the exact detection limit. They recommend using a shear calibration sample at least by 1.5 magnitudes deeper than the survey in question (which ours does). Their findings, however, make exclusive use of the moments-based KSB algorithm (see Kaiser 1995); such techniques are known to probe a galaxy's ellipticity at larger radii than other methods, which could in principle make them more sensitive to nearby faint galaxies. It is thus a worthwhile exercise to gauge their impact in our case with IM3SHAPE.

Impact on Multiplicative Bias

We first draw what conclusions we may by again comparing our HOOPOE simulations with the neighbour-free WAXWING resimulations. The logic here is as follows. Since WAXWING postage stamps consist of only a single profile added to Gaussian pixel noise, they are unaffected by neighbours of any sort (faint or otherwise). We have seen that the impact of neighbours is strongly localised, with the excess m converging within a nearest neighbour distance d_{gn} of a dozen pixels or so. Thus, selecting galaxies that are well separated from their nearest *visible* neighbour will isolate the impact of the *undetected* ones.

A further cut is thus imposed on $d_{gn} < 20$ pixels. Relative to the case with quality cuts only, the global multiplicative bias now shifts from $m \sim -0.12$ to $m \sim -0.065$ (the first and second lines in green on Figure 4.11). This measurement is in mild tension with the value measured from WAXWING (again under its own cuts, with the selection on d_{gn}). This difference, which amounts to a negative shift in m of ~ 0.01 is, we suggest, the net effect of the sub-detection galaxies. From these numbers alone one cannot tell if this is a result of selection effects, flux contamination, bin shifting or some combination thereof.

Interestingly we find that imposing both the HOOPOE and WAXWING selection functions in addition to the cut on d_{gn} brings m into agreement to well within the level of statistical precision (compare the final and penultimate lines in green in Figure 4.11). That is, when restricted to a subset of galaxies that pass quality cuts in both simulations the flux contributed by the faint objects has little impact. Their main impact is rather that they allow a population of marginal faint galaxies which would otherwise be flagged and removed by quality cuts to pass into the final HOOPOE catalogue.

To test this idea further we rerun a subset of 100 random tiles from the simulated footprint, without the final step of adding sub-detection galaxies. To minimise the statistical noise in this comparison we enforce the same COSMOS profiles, shears and rotations as well as the per-pixel noise realisation as before. SExtractor source detection is applied and the blending flags are propagated into the postprocessing cuts.

The raw m values calculated from the rerun and the main HOOPOE simulations, matched to the same tiles, are shown by the upper two red lines on Figure 4.11. The downward shift of ~ 0.01 is consistent with the previous result based on the main simulation. This comparison should encapsulate the full effect of the faint objects (since there are no other differences between these datasets).

For each galaxy we next measure the distance to the nearest faint object d_{gf} , the distribution of




Figure 4.13: Histogram of radial distances between galaxies in our measured shape catalogues (the full HOOPOE simulations) and the nearest object below the DES detection limit. The dotted line includes all objects prior to quality cuts, while the solid line shows the impact of applying IM3SHAPE’s INFO_FLAG cuts (see J16). The dashed blue line shows the population of galaxies which survive cuts only in the presence of the faint galaxies.

which is shown under various selections in Figure 4.13. Like in the comparison in Section 4.5.3, there is a population of galaxies that survive cuts *only* in the simulation with the subdetection objects, and these galaxies tend to be ones with extremely close faint neighbours. Interestingly the inverse population surviving only when they are removed do not preferentially have small d_{gf} . This is intuitively understandable: a faint object might boost its neighbour’s apparent size or S/N if it were centred within a few pixels. Otherwise it would act as a source of background noise, which would reduce the quality of the fit.

Finally we find that if we apply both selection functions to the with-faint galaxies, the measured biases become roughly consistent. These findings, combined with the observations in the previous section lead us to an interesting conclusion: the major effect of the faint galaxies in the DES Y1 IM3SHAPE catalogue is to allow a population of small faint galaxies to pass quality cuts, where otherwise they would have been removed. This is analogous to the neighbour dilution effect described above, but is subdominant to the influence of visible neighbours.

Impact on Background Flux Subtraction

As a test of the robustness of this result we recompute our IM3SHAPE fits on the faint-free images, with and without the correction for the shift in the background flux that *would* have been applied if the sub-detection galaxies had been drawn. The mean per-tile correction is $\Delta f \sim 0.05$, against typical noise fluctuations $\sigma_n \sim 6.5$. Matching galaxies and examining the histograms of $\Delta S/N$ and $\Delta R_{gp}/R_p$ reveals weak downwards scatter in both quantities (i.e. the flux subtraction alone makes galaxies appear smaller and fainter). The magnitude of the shift is, however, tiny, peaking at ~ -0.1 and -0.005 respectively. This is logical given the definition in equation 4.1. If the change is small enough such that the best-fitting model is stable, then an incremental reduction in flux will reduce the signal-to-noise of the measurement. Looking at the best-fit shapes, we find a small shift towards high ellipticities, which can likewise be understood as a numerical effect; imposing a flat positive field of zero ellipticity will dilute the measured shear, producing a bias towards round $|e|$. The reverse logic applies with the flux correction, and subtracting a flat value from all pixels will make galaxies appear slightly *more* elliptical. In practice we find a sharp peak at $\Delta e \sim 0.001$.

4.5.5 Suppressing Neighbour Bias

There is no universal definition for the shape-weighted effective number density commonly used as proxy for cosmological constraining power in a shear catalogue. One which is particularly useful in the context of weak lensing, and which has been adopted in DES Y1 is the prescription of Chang et al. (2013), which is designed to account for shape noise and fitting error (see equation 7.3 in Zuntz et al. 2017). A second useful definition is set out by Heymans et al. (2012) in terms of the (see also Z17). We compute a neighbour distance d_{gn} for every object in the simulation which allows us to cut on this quantity. Removing any galaxy with a neighbour detected within a radius of $d_{gn} = 20$ pixels reduces the effective number density of sources using either definition to about 70% of its initial value, from $n_{\text{eff}}^{\text{H13}} = 5.48$ to $n_{\text{eff}}^{\text{H13}} = 3.68 \text{ arcmin}^{-2}$ using Heymans et al. (2012)'s definition. Using the prescription of Chang et al. (2013), the equivalent density drops from $n_{\text{eff}}^{\text{C13}} = 3.18$ prior to cuts and $n_{\text{eff}}^{\text{C13}} = 2.18 \text{ arcmin}^{-2}$ afterwards. This cut is stringent, as we have shown that beyond ~ 12 pixels the multiplicative bias becomes insensitive to further selection on d_{gn} . There are, however, a number of limitations in our analysis, including the fact that d_{gn} is defined using the true input positions, and indeed that we are using only the detected positions in DES to draw our simulated $M_r < 24.1$ galaxies. We thus judge that a level of conservatism is appropriate here. Relaxing the cut to $d_{gn} > 14$ pixels leaves n_{eff} at 84% of its full value.

4.6 Cosmological Implications

As we have shown in the previous sections, if ignored completely image plane neighbours can induce negative calibration biases in IM3SHAPE of a few percent or more. The earlier part of the investigation focused on *when* and *how* neighbour bias can arise, first in the context of single-galaxies and then on ensemble shear estimates. We now turn to a more pressing question from the general cosmologist’s perspective: *how far should we be concerned about these effects in practice?*

In the following paragraphs we present a set of numerical forecasts using a Markov Chain Monte Carlo algorithm to sample trial cosmologies. Our basic methodology here follows previous MC forecasts (eg Joachimi & Bridle 2010; Krause et al. 2015, 2017). We construct mock DES Y1 cosmic shear measurements using a matter power spectrum derived from the Boltzmann code CAMB⁴ with late-time modifications from HALOFIT. The cosmic shear likelihood surface is sampled at trial cosmologies using COSMOSIS⁵. The final data used for the likelihood calculation have the form of real-space ξ_{\pm} correlations and are shown by the purple points in Figure 4.14. For the photometric redshift distributions we use the measured estimates in four tomographic bins, obtained from runs of the BPZ code on the Y1 IM3SHAPE catalogue. Note that for our purposes it is unimportant if the $n(z)$ used to construct the mock data are somewhat inaccurate; we simply wish to mimic the qualitative (non-analytic) shape and mean redshift in each bin. In all MC chains which follow we marginalise over two nuisance parameters (an amplitude and a power-law in redshift) for intrinsic alignments, photo- z bias and shear calibration bias. In total this gives 10 extra free parameters in addition to six for cosmology ($\Omega_m, \Omega_b, A_s, n_s, h, \Omega_{\nu}h^2$), which are also allowed to vary. The analysis choices are identical to those used in the DES Y1 cosmic shear analysis of Troxel et al. (2017), and we refer the reader to that paper for details of the priors and scale cuts. Finally, the following adopts shear-shear covariance matrices derived from the analytic halo model calculations of Krause et al. (2017). We assume a fiducial Λ CDM cosmology $\sigma_8 = 0.81$, $\Omega_m = 0.315$, $\Omega_b = 0.049$, $n_s = 0.97$, $h = 0.688$, $\tau = 0.08$, with non-zero comoving neutrino density $\Omega_{\nu}h^2 = 0.00062$.

4.6.1 Mean Multiplicative Bias

We first seek to quantify the bias that would be present a cosmic shear analysis in a survey like DES, if we were to use a simple postage stamp simulation of the sort presented in J16 and Miller

⁴<http://camb.info>

⁵<https://bitbucket.org/joezuntz/cosmosis>

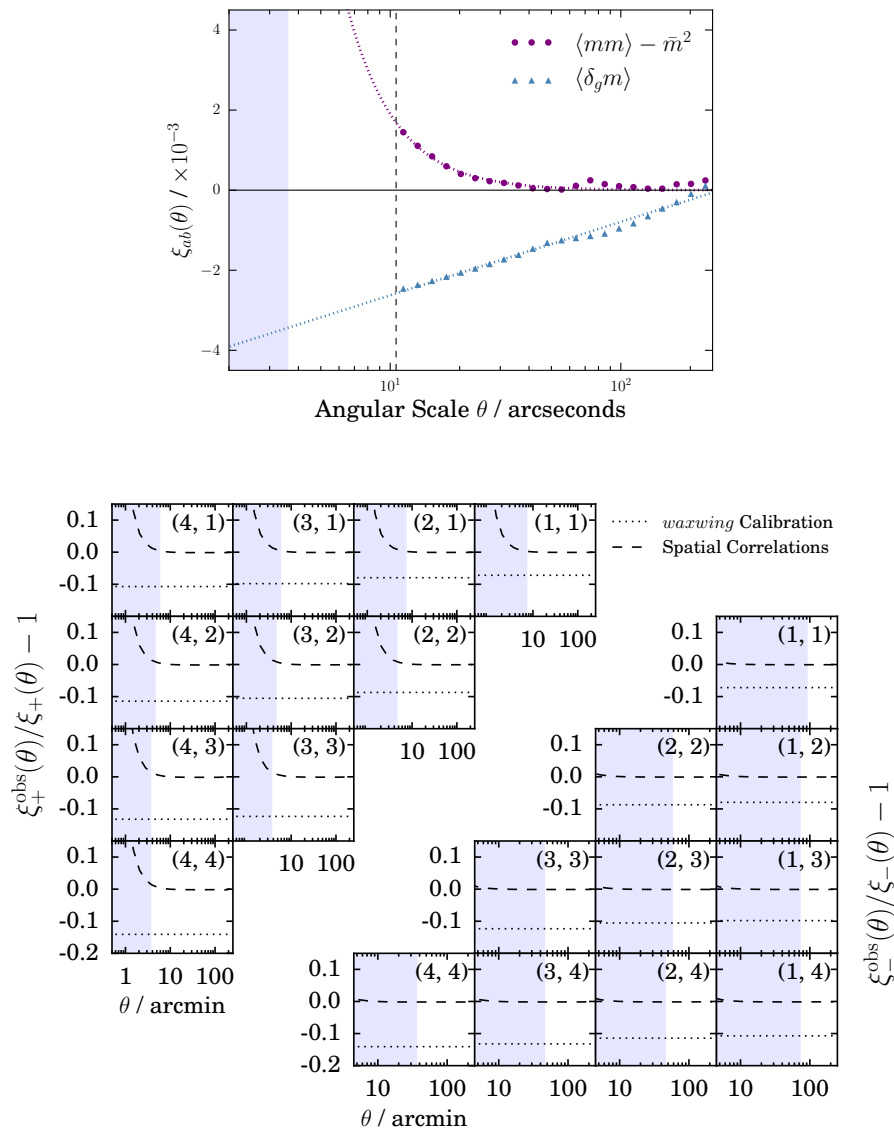


Figure 4.14: **Top:** The observed two point correlation of multiplicative bias, as measured from the main HOOPOE simulation set presented in this chapter. Sub-patches are used to compute m in spatial patches of dimension 0.15×0.15 degrees and the correlation function calculated as described in the text. The dashed vertical line shows the diagonal scale of the sub-patches, below which we do not attempt to directly measure spatial correlations. The shaded blue bands show the minimum and maximum scales used in the DES Y1 cosmic shear analysis of Troxel et al. (2017). **Bottom:** Residuals between the mock two point shear-shear data used in this chapter, before and after different forms of bias have been applied. The upper and lower subpanels show the ξ_+ and ξ_- correlations respectively, calculated using the redshift distributions of Hoyle et al. (2017). The dotted black lines, which are flat on all sim but vary between panels, show the result of calibrating our Y1 shear measurements with a simple postage stamp simulation without image plane neighbours. The dashed lines illustrate the impact of ignoring scale-dependent selection effects, which are not captured by our simulation-based calibration. The shaded blue regions of each panel show the excluded scales for each particular tomographic bin pairing.

et al. (2013). To this end we use the neighbour-free WAXWING dataset to construct an alternative shear calibration. In Z17 we compared three methods for shear calibration using the HOOPOE simulations and found our results to be robust to the differences. We now use the fiducial (grid-based) scheme to derive an alternative set of bias corrections from WAXWING. These are then applied to the same galaxies in the matched HOOPOE simulation, and residual biases are measured in four DES-like tomographic bins. The process is very similar to the diagnostic tests in §5 of Z17, and so we defer to that work for details of the redshift bin assignment of simulated galaxies.

Using the neighbour-free simulation we under-correct the measurement bias by several percent in each bin. The remeasured residual bias after calibration provides an estimate for the level of systematic that would be present were we to calibrate DES Y1 using the simpler WAXWING simulations. In the four tomographic bins used in DES Y1 we find

$(\Delta m^{(1)}, \Delta m^{(2)}, \Delta m^{(3)}, \Delta m^{(4)}) = (-0.037, -0.044, -0.064, -0.073)$, and apply these biases to our mock data. The resulting shift in the shear two-point correlations is shown by the black dotted lines in the lower panel of Figure 4.14. Since the calibration scheme does not explicitly include neighbour distances, but rather orders galaxies into cells of S/N and R_{gp}/R_p , this test does not include any scale dependent neighbour effects. The calibration effectively marginalises out d_{gn} , and the residual biases are an average over the survey. For the moment we will assume this mean shift in m is sufficient, and return to the question of scale dependence in the following section.

Our predicted cosmology constraints with weak lensing alone in DES Y1 are shown in Figure 4.15. In purple we show the results of the fiducial analysis, in which the shear calibration fully captures all neighbour effects and leaves no residual multiplicative bias. The contours are, naturally, centred on the input cosmology, marked in this figure by the black cross. The blue (solid) contours then show the impact of residual neighbour biases per bin at the level described. As one can see, even when marginalising over m^i with an (erroneously) zero-centred Gaussian prior of width $\sigma_m = 0.035$, our cosmology constraints are shifted enough to place the input cosmology outside the 1σ confidence bounds. We reiterate here that this calculation highlights the bias that *would* arise were we to naively apply a calibration of the sort used in DES-SV based on neighbour-free simulations to the Y1 data. Since we are confident that the HOOPOE code captures the effects of image plane neighbours correctly (at least to first order) this is a hypothetical scenario only and not a prediction of actual bias in DES Y1.




Figure 4.15: Expected cosmology constraints from DES Y1 cosmic shear only. The purple (solid) contours show the results of calibrating using a simulation which fully encapsulates all biases in the data, leaving no residual m in the final catalogue. In blue (dash-dotted) we show the result of calibrating with an insufficiently realistic simulation, which leaves a residual bias between -0.03 and -0.08 in each of the redshift bins. For reference the input cosmology is marked with a black cross.

4.6.2 Scale Dependence

It is not trivial that including an mean neighbour-induced component to m over the entire survey will be sufficient to mitigate all forms of neighbour bias. The local mean m on a patch of sky is sensitive to spatial fluctuations in source density, which could induce scale dependent bias on arcminute scales. Clearly, when correlating galaxy pairs on small scales one can expect a larger fraction in which the objects come from a similar image plane environment, and more often than not that environment will be densely populated. Thus the true multiplicative bias should become more negative on small scales.

Two subtly different effects emerge from this thought experiment. First, the multiplicative bias of galaxies will be spatially correlated (i.e. a correlation involving two galaxy populations $\langle m^i m^j \rangle$ is not just the product of the means $\bar{m}^i \bar{m}^j$). Second, in the small θ bins one is selecting galaxies




Figure 4.16: The same as Figure 4.15, but now showing the impact of residual scale dependent selection bias. The two sets of confidence contours represent different assumptions about the small scale extrapolation of the ξ_{mm} correlation, as outlined in the Section 4.6.2. In green (dashed) we show a mildly optimistic case, using the linear fit shown in Figure 4.14. The pink dotted contours show a (strongly pessimistic) power law extrapolation. The dark blue solid line makes identical assumptions to the pink, but incorporates small-scale information, to a minimum separation of $\theta_{\min}^+ = 0.5$ arcminutes in $\xi_+(\theta)$ and $\theta_{\min}^- = 4.2$ arcminutes in $\xi_-(\theta)$. As in Figure 4.15 the input cosmology is shown by a black cross.

with close partners with which to correlate, and thus oversampling the dense parts of the image. To gauge the level of these effects, we divide each simulated coadd tile into a grid of 25 square sub-patches with dimension 0.15×0.15 degrees. we fit for m using the galaxies in each sub-patch and assign the resulting value to these objects. While this only allows a noisy measurement of m , it should capture the spatial variations in number density to the level of a few percent. We next measure the two-point correlation function of multiplicative bias values assigned in this way using TREECORR⁶, excluding galaxy pairs at angular separation smaller than the scale of the sub-patches. We refer to this bias-bias autocorrelation as ξ_{mm} , which we show as a function of angular scale in the upper panel of Figure 4.14. Analogously one could use the subpatches to construct

⁶<http://rmjarvis.github.io/TreeCorr>

correlations between m and galaxy number density ξ_{gm} or density with density ξ_{gg} . The statistical noise on these correlations is significantly lower than that on the individual subpatches by virtue of the large simulation footprint. Note that in Figure 4.14 we subtract \bar{m}^2 , measured from all galaxies in the simulation, from the measured ξ_{mm} . If there were no θ dependence the correlation $\langle m^i m^j \rangle$ should simply average to the square of the global mean in all scale bins. As one can see from the circular points in this figure, scales larger than the diagonal size the sub-patches (shown by the vertical dashed line) exhibit non-negligible excess ξ_{mm} . One obvious question is whether this could be the result of finite binning error, which scatters galaxies in the same sub-patch into different θ bins. To verify this is not the case we repeat the measurement as before, but halve the parameter controlling binning error tolerance (“bin slop”) and obtain the same results.

To extend this measurement down to scales below the sub-patch we must make some assumptions about the functional form of the mm correlation. We fit a power law, $\Delta\xi_{mm}(\theta) \equiv \xi_{mm}(\theta) - \bar{m}^2$:

$$\Delta\xi_{mm}(\theta) = \beta\theta^{-\alpha}, \quad (4.6)$$

which is shown by the dotted purple line in this figure. This provides a qualitatively good fit to the measured points, but as one can see implies a rather dramatic inflation on small scales.

In the limited range over which we have a nonzero measured correlation, however, a linear function of θ (truncated at $\theta = 27$ arcminutes) also provides a reasonable by-eye fit. The small-scale extrapolation in this case is significantly milder. The $\langle m\delta_g \rangle$ and $\langle \delta_g\delta_g \rangle$ measurements are linear with θ to good approximation, and so we use linear fits to extrapolate them below the patch size.

Assuming the bias per tomographic bin can be written as the sum of a redshift dependent contribution (i.e. a scale invariant mean in each bin), and a scale dependent term, one can write the correlation per bin as $m^i m^j = \bar{m}^i \bar{m}^j + \Delta\xi_{mm}(\theta)$. A more complete derivation of this expression can be found in Appendix C. The first part can be extracted from the DES Y1 calibration, and one can fit for $\Delta\xi_{mm}(\theta)$ as described above. A set of modified ξ_{\pm}^{ij} are thus computed. These appear in the lower panels of Figure 4.14 as dashed lines. As can be seen, the scale cuts of Troxel et al. (2017) (excluded scales are shaded in blue) are sufficiently stringent to remove almost all of the visible scale dependence. Though reassuring for the immediate prospects of DES Y1, this will not trivially be true for all future (or indeed ongoing) lensing surveys. It is thus important that the effects we identify here are properly understood at a level beyond the resources of the current analysis. These biased data are then passed through our likelihood pipeline to gauge the cosmological impact,

which is shown in Figure 4.16. In the linear case (dashed green) there is no discernable bias in the σ_8 Ω_m pair; even the much harsher power-law extrapolation (pink dotted) induces only an incremental shift along the degeneracy direction. In both cases the input cosmology still sits comfortably within the 1σ confidence contour.

Finally we test the impact of relaxing the stringency of our scale cuts. The minimum scales used for ξ_+^{ij} and ξ_-^{ij} are shifted downwards to 0.5 and 4.2 arcminutes respectively, irrespective of bin pair, which are the cut-off values used in fiducial cosmic shear analysis of Hildebrandt et al. (2016). This increases the size of our datavectors considerably. Incorporating smaller angular scales will clearly improve the constraining power of the data to an extent. Primarily the effect is to shorten the lensing degeneracy ellipse, cutting out much of the peripheral curvature, but it also reduces the width in the S_8 direction. These scales, however, contain biased information, which induces tension between the small and large angular scales. With the strongest (power law) scale dependence considered, the input cosmology is displaced moderately along the degeneracy curve, though it remains within the 1σ confidence bound.

4.7 Conclusions

The Dark Energy Survey is the current state of the art in cosmological weak lensing. Multi-band imaging down to 24th magnitude across 1500 square degrees of the southern sky has yielded hitherto unparalleled late-time constraints on the basic parameters of the Universe (see Troxel et al. 2017 and DES Collaboration et al. 2017). This an exciting time for lensing cosmology and, of course, requires galaxy shear measurements of exquisite accuracy.

In this chapter we have used one of two DES Y1 shear catalogues, and large-area simulations based upon them, to quantify the impact of image plane neighbours. The present study is one of a small handful on this important subject, and the first to approach it in the context of a realistic simulated image pipeline and a cutting-edge model-based shape measurement algorithm.

In order to properly mitigate the influence of galaxy neighbours, and thus avoid drawing flawed conclusions about cosmology from the data, it is important to first understand the mechanisms by which they enter the shape measurement. Using a simple toy model of the galaxy-neighbour system we have shown that shear bias can arise even when the distribution of neighbours is isotropic (i.e. there is no preferred direction). This is the result of a small difference in the impact of the same neighbour, when it is placed on or away from the axis of the shear. We have furthermore shown that the resulting multiplicative shear bias m can be either positive or negative, depending on the

model parameters. With slight modifications to the toy model, whereby we Monte Carlo sample input parameters from the joint distribution of the equivalent properties measured in DES Y1, we have shown that a mild negative m is dominant when marginalising over a realistic ensemble of neighbours. This was seen to be strongly dependent on the distance of the neighbour, and to be mitigated but not eliminated by basic cuts on the centroid position of the best-fitting model.

Using the DES Y1 HOOPOE simulations, which were also used to derive shear calibration corrections for the Y1 IM3SHAPE catalogue of Jarvis et al. (2015), we have presented a detailed study of the ensemble effects of galaxy neighbours. In this analysis we have identified four mechanisms for neighbour bias, which we call flux contamination, selection effects, bin shifting and neighbour dilution. All can be understood in intuitive terms, resulting from close-by or moderately close neighbours. Our results from the full simulation are consistent with the toy model calculation. Though we have shown strong dependence on distance to the nearest neighbour (and thus on number density) we found only weak sensitivity to neighbour brightness, when averaged across broad bins of magnitude. In addition to this, cuts on the DES Y1 catalogue sufficient to null the impact of the detectable neighbours would result in a degradation of over 20% in source number density. one cannot recommend such measures for a code like IM3SHAPE, in part because the data contains correlations between shear and number density. Unless the link is preserved in the calibration simulations, such selection could conceivably induce additional bias towards low shear⁷.

Our investigation also assessed the impact of the faintest galaxies, which are not reliably detected but nonetheless contribute flux to the survey images. Via two different routes, first using a spin-off neighbour-free resimulation, and also using a subset of images simulated again with the sub-detection galaxies missing, our findings suggested a net contribution to the multiplicative bias budget of $m \sim -0.01$.

Unlike most earlier works on shear measurement, we have propagated these findings to the most meaningful metric for cosmic shear: bias on the inferred cosmological parameters. The study we have presented here uses the DES Y1 cosmology pipeline, as well as non Gaussian shear covariance matrices and real photometric redshift distributions to implement MCMC forecasts. In the first case considered, the data included a (different) multiplicative bias in each redshift bin, designed to approximate the residual m that would arise were we to calibrate DES Y1 with a simple neighbour free simulation. Even marginalising over m with a prior of $\mathcal{N}(0, 0.035^2)$ this scenario was demonstrated to result in a shift in the favoured cosmology towards low clustering amplitude of more than 1σ .

⁷Although the sister catalogue to Y1 IM3SHAPE uses a form of self-calibration, which should allow one to correct for the additional selection bias.

Finally, we have explored a second possible source of measurement bias arising from the link between number density and neighbour bias. This enters two-point measurements as an additional correlation between the multiplicative bias in galaxy pairs at small angular separation. In the final section we have measured such a correlation from the HOOPOE mock images. With the most pessimistic small-scale extrapolation, this was found to result in a shift in the best-fitting cosmology of under 1σ in the negative S_8 direction, which is not remedied by marginalising over m . A less dramatic, though still considerable, increase in the correlation strength on small scales was demonstrated to result in no discernable cosmological bias.

Both of these effects are of primary concern for the next generation of cosmological surveys. By the end of their lifetime KiDS, DES and HSC are set to offer lensing-based cosmological constraints comparable to the CMB. The first, dominant, effect can be remedied relatively easily by calibrating our shear measurements with sufficiently complex image simulations. Indeed, the most recent shear constraints of Hildebrandt et al. (2016), Köhlinger et al. (2017) and Troxel et al. (2017) have done just that. Unfortunately, the correct treatment of scale dependent bias is not as clear. Though further statements about the likely small scale dependence of the mm correlation are beyond the scope of the present study, understanding this intricate topic will be crucial for future surveys if we are to fully exploit the constraining power of the data. The massive simulation efforts of LSST and Euclid, combined with advancement in neighbour mitigation using techniques such as multi-object fitting will be invaluable in this task. With the enhanced understanding these will provide and the exquisite data of the next generation surveys, the coming decade will be an exciting time for cosmology.

Chapter 5

Constraining Redshift Systematics Using Lensing & Galaxy Clustering

In this short chapter we present a set of numerical forecasts designed to explore the complementary constraining power of cosmic shear and galaxy clustering measurements. The following has been accepted for publication in MNRAS Letters as Samuroff et al. (2017b).

5.1 Introduction

Cosmic shear is potentially the most powerful tool available to cosmologists today. As an unbiased probe of the mass distribution, it offers powerful constraints on the mean density of the Universe and the clustering of dark matter. It is also expected to shed new light on the late-time accelerated expansion of the Universe and thus measure the dark energy equation of state and test General Relativity on the largest scales. A three decade programme aiming to extract unprecedented constraints on our cosmological model from cosmic shear is now midway to completion. It began soon after the first detection in 2000 (Bacon et al., 2000; Van Waerbeke et al., 2000; Wittman et al., 2000; Kaiser et al., 2000) using ~ 10000 galaxies and will culminate in catalogues of more than a billion galaxies by the end of the coming decade (Stage IV, Albrecht et al. 2006). Logarithmically, we are halfway there, with ongoing analyses of the preliminary Stage III datasets containing $\sim 10\text{M}$ galaxies (Dark Energy Survey Collaboration 2016; Hildebrandt et al. 2016, see also Heymans et al. 2013; Jee et al. 2016). The increase in the number of galaxies with reliable shape measurements has allowed tighter cosmology constraints, but also requires better control of systematic biases. In this chapter we focus on a potential Achilles' heel of galaxy imaging sur-




Figure 5.1: Components of the fiducial datavector. Shown are angular power spectra of cosmic shear (purple solid), galaxy clustering (red dotted) and the cross-correlations (green dot dashed). Each panel corresponds to a unique redshift bin pairing. In the panels where it is not visible, the $\delta_g\delta_g$ spectrum is below the range shown. All values shown are positive, apart from $C_{\gamma\delta_g}^{1,3}$ (upper right), which becomes negative and is smaller than the lowest point on this scale at $\ell < 900$.

veys for cosmology: the use of photometric redshifts (photo- z) to estimate distances to galaxies. Tomographic cosmic shear analyses bring a number of benefits (Hu, 1999), but place stringent requirements on our knowledge of galaxy redshift distributions. Amara & Réfrégier (2007); Abdalla et al. (2008); Jouvel et al. (2009) present detailed studies of the spectroscopic follow-up needed for Stage IV, while Ma et al. (2006); Huterer et al. (2006); Bernstein (2009) use numerical forecasts to explore cosmological impact of photo- z biases. Many others (e.g. Bordoloi et al. 2012; Cunha et al. 2014) present detailed studies of specific photo- z systematics, albeit with less focus on the ultimate cosmology.

Tightening systematics requirements have sparked interest in spatial cross-correlations between photometric and spectroscopic galaxies within the survey volume as a method for calibrating photo- z (Newman, 2008; Ménard et al., 2013; de Putter et al., 2014; Choi et al., 2015). Given the limited amount of spectroscopic information available, several authors have speculated about calibrating redshift error from the imaging survey itself. Huterer et al. (2006) show that cosmic shear alone affords a limited capacity for self-calibration. Schneider et al. (2006) and Sun et al. (2015) investigated the photo- z calibration information available from Stage IV galaxy clustering,




Figure 5.2: Redshift distributions considered in this chapter. The upper panels show the shear $n(z)$ s, taken from DES SV (Bonnert et al., 2015): SKYNET (solid purple; fiducial), SKYNET with a 0.05 bias (dashed green) and BPZ (dotted blue) without the shift of 0.05 in redshift used in Bonnert et al. (2015). The lower panel displays the galaxy density catalogue, DES SV redMaGiC (Rozo et al., 2016) in bins defined by Clampitt et al. (2016).

and Zhan (2006) explore the constraining power on w_0 using a similar technique with cosmic shear plus clustering constraints. Zhang et al. (2010) point out that shear-density cross-correlations (between shear and galaxy counts, also referred to as tangential shear or galaxy-galaxy lensing) can help to constrain photo- z error when combined with galaxy clustering.

All the studies mentioned above make a crucial assumption, which is unlikely to be realised in practice, that the galaxies used for cosmic shear have a systematics-correctable galaxy clustering signal. In practice regions of the sky with better (worse) seeing conditions are likely to contain a higher (lower) number density of galaxies usable for cosmic shear (e.g. see Appendix C of Choi et al. 2015). A large spurious clustering signal will arise as a result, rendering standard galaxy clustering analyses useless. Thus, in practice, there is usually a different galaxy sample selection for the shear and clustering samples. This is widely accepted in galaxy-galaxy lensing studies and was also the case in the first combined cosmic shear plus large scale structure analysis with

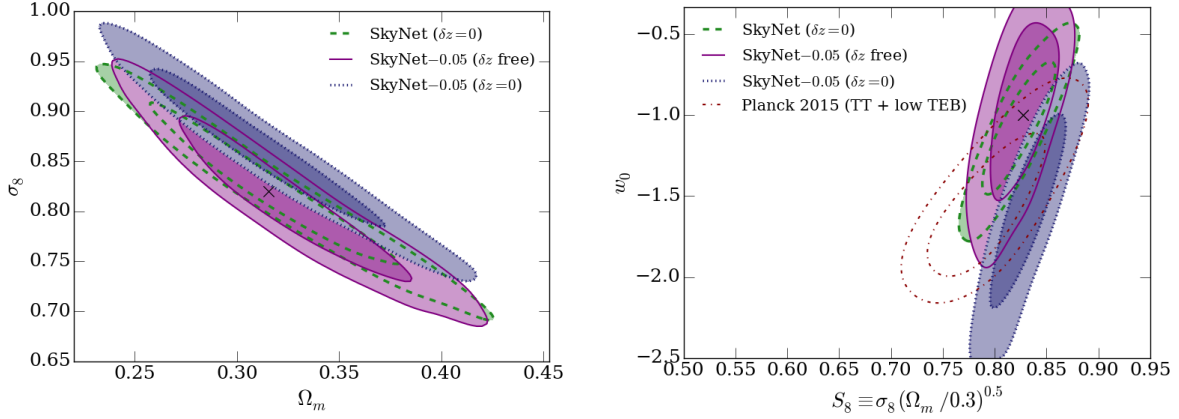


Figure 5.3: **Left:** Forecast constraints on the matter density and clustering amplitude in Λ CDM and **Right:** dark energy equation of state in w CDM for various assumptions about photo- z . For reference Planck 2015 constraints (temperature + low frequency polarization; Planck Collaboration 2015a) are shown in red (dot-dashed). The colours in each panel indicate three photo- z scenarios. In green are the result of using the SKYNET $n(z)$ in the theory calculation and fixing $\delta z = 0$. we show this as an unrepresentative ideal case, where the photometric estimates provide a perfect representation of the true galaxy distribution. Overlain are the same, but using SKYNET biased downwards by 0.05 in redshift under the (erroneous) assumption of no bias (blue dotted) and varying three δz^i nuisance parameters marginalised with a Gaussian prior of width $\Delta\delta z = 0.1$ (purple solid). The input cosmology is shown by the black cross.

real data (Nicola et al., 2016), and was considered for Stage IV in the forecasts of Krause & Eifler (2016). Though one has twice as many redshift distributions to understand as in a shear-only analysis, this offers an opportunity: we can choose a galaxy clustering sample with well controlled photo- z , which in turn helps to calibrate the redshift distribution of the weak lensing sample. In this chapter we explore the potential for simultaneously constraining photo- z error and cosmology using cosmic shear, galaxy clustering and shear-density cross-correlations. Unlike previous studies we consider a scenario in which the redshift distribution of the shear catalogue and galaxy clustering catalogues differ significantly. We assume the clustering sample is highly homogeneous and dominated by luminous red galaxies, which yield high quality photo- z .

This chapter is structured as follows. Section 5.2 outlines the analysis with a description of the simulated data vectors, redshift distributions and the photo- z uncertainties considered. In Section 5.3 we investigate the power of these data to internally constrain photo- z biases. Finally a series of robustness tests are presented to explore the limits of this effect. The following calculations assume a fiducial flat Λ CDM cosmology with $\sigma_8 = 0.82$, $\Omega_m = 0.32$, $h = 0.67$, $w_0 = -1$, $\Omega_b = 0.049$.

5.2 Methodology & Assumptions

We follow a method similar to Joachimi & Bridle (2010) (see also Duncan et al. 2014 for a similar analysis) to implement a forecast of the three weak lensing plus large-scale structure two-point functions: cosmic shear, galaxy clustering and shear-density cross-correlations. We carry out an MCMC forecast by simulating a datavector and covariance at a fiducial cosmology, then fitting a series of trial cosmologies and computing the likelihood of each. The fiducial datavector Figure 5.2 contains three types of correlation, each with 25 logarithmically spaced top hat bins over $10 < \ell < 3000$. We use COSMOSIS (Zuntz et al., 2015) to MCMC sample parameter space and compute matter power spectra using CAMB (Lewis et al., 2000) with nonlinear corrections from Takahashi et al. (2012).

The fiducial analysis assumes a galaxy catalogue typical of a preliminary Stage III survey. We use the galaxy number density of the Dark Energy Survey Science Verification (DES SV; 6.8 arcmin^{-2} Jarvis et al. 2015), with $\sigma_\epsilon = 0.2$ and an area of 1500 sq. degrees. This gives 37M galaxies, which is a little larger than or comparable to CFHTLenS and the KiDS, DES and HSC preliminary analyses. We use the SKYNET $n(z)$ presented in Bonnett et al. (2015), as used in the DES SV shear analysis, and marginalise over multiplicative shear bias with a Gaussian prior ($\Delta m = 0.02$) (see also Jarvis et al., 2015; Fenech Conti et al., 2016; Jee et al., 2016). For conservatism we model intrinsic alignments using the nonlinear alignment model (Bridle & King, 2007) with an additional power law in redshift (e.g. Joachimi et al. 2011; Dark Energy Survey Collaboration 2016), and unlike previous analyses allow the amplitude and power law index to differ for GI and II. This gives four IA parameters.

To model a realistic galaxy clustering catalogue the $n(z)$ of the DES SV redMaGiC luminous red galaxy catalogue (Rozo et al., 2016; Clampitt et al., 2016) are adopted. A linear galaxy bias is applied in each bin and marginalised with a wide flat prior. To avoid the nonlinear bias regime we impose conservative scale cuts to the clustering sample and discard information below a scale determined by rescaling the prescription of Rassat et al. (2008). We tune this rescaling approximately match the smallest scales found by Kwan et al. (2017) to be unaffected nonlinear bias. This leads to a factor of three increase over the Rassat et al. (2008) cuts, giving $\ell_{max} = (91, 203, 432)$ for each bin respectively.

Redshift distribution error is parameterized by δz^i , which describes a uniform linear translation,

$$\tilde{n}^i(z) = n^i(z + \delta z^i). \quad (5.1)$$

For the galaxy clustering sample we marginalise over δz^j with a Gaussian prior of standard deviation $\Delta\delta z^j = 0.01$. The photo- z for the shear catalogue are somewhat lower in quality due to the number and type of objects required for shear measurement and for the fiducial analysis we apply a conservative Gaussian prior ($\Delta\delta z^i = 0.1$). In addition to the nuisance parameters described above we vary 5 cosmological parameters with no external priors. The fiducial analysis then has 21 degrees of freedom,

$$\mathbf{p} = (\sigma_8, \Omega_m, h, \Omega_b, n_s, A_{\text{GI}}, A_{\text{II}}, \eta_{\text{GI}}, \eta_{\text{II}}, m^i, \delta z^i, b_g^j, \delta z^j).$$

5.3 Simultaneous Constraints on Cosmology and Photometric Redshift Bias

Figure 5.3a shows constraints on the matter density Ω_m and clustering amplitude σ_8 for various assumptions about the shear photo- z uncertainties. The green dashed lines are from the fiducial analysis assuming the shear redshift distributions are known precisely ($\Delta\delta z^i = 0$). We investigate the effect of using an incorrect redshift distribution in the simulation, expanding on a similar technique developed in Bonnett et al. (2015). A bias of 0.05 in redshift is applied to $n(z)$ in the simulated data vector, a value inspired by the calibration applied to BPZ in Bonnett et al. (2015) to match simulations, but we assume the redshift distributions are known perfectly. The result (blue dotted, Figure 5.3a) is now incompatible with the true cosmology at greater than 95% confidence. Finally, we allow freedom in the value of the photo- z biases δz^i , marginalising with the fiducial prior ($\Delta\delta z^i = 0.1$). The purple solid contours are shifted back close to the true input cosmology, despite the erroneous redshift distributions. The width of the purple contours is not greatly degraded relative to the case where the distributions are perfectly known (green dashed and blue dotted): we find an increase in error on S_8 of 40%. For Figure 5.3b we carry out the same calculation as in Figure 5.3a, but additionally vary the dark energy equation of state w_0 . Qualitatively similar results are obtained.

In Figure 5.4 we investigate in more detail how the prior width $\Delta\delta z^i$ on the redshift distribution bias δz^i affects the uncertainty on S_8 . The results from cosmic shear alone (blue) can be contrasted with those from the combination of cosmic shear, galaxy clustering and shear-density cross-correlations (purple). We show results using our fiducial systematics assumptions (purple) and using less conservative assumptions (no multiplicative shear or IA uncertainty; green). We see that cosmic shear alone cannot self-calibrate photo- z uncertainties, whereas the combination with galaxy clustering and shear-density correlations weakens the sensitivity of the constraint to prior

width, for both fiducial and optimistic systematics. S_8 is significantly biased (δS_8) at all values of $\Delta\delta z^i$ for cosmic shear (dotted lines) and is biased very little by $\Delta\delta z^i = 0.1$ for the combination of datasets (dot-dashed).

Figure 5.5 gives some insight into how the self-calibration works, using the biased redshift distribution (SKYNET -0.05) as an illustration. The blue (dot-dash) contours show the degeneracy between cosmology and photo- z uncertainties from cosmic shear alone. The contours are closed because we have applied a conservative prior on photo- z uncertainties ($\Delta\delta z^i = 0.1$). In the absence of additional information the cosmology constraints from cosmic shear will be biased because the prior on δz^i is centered on zero whereas the truth is at $\delta z^i = 0.05$. The galaxy clustering and shear-density cross-correlations constrain a different degenerate combination of cosmology and redshift bias (pink dotted). Thus, when these three two-point functions are combined, they produce the purple (solid) contours, which are now centered close to the true cosmology and offer much tighter constraints on cosmology and photo- z uncertainties than either cosmic shear alone or galaxy clustering plus shear-density correlations alone. In physical terms, Figure 5.5 demonstrates how a tight constraint on the clustering redshift distributions constrains the lensing sample. Since we are correlating each of our biased lensing kernels with well known $n(z)$ s in multiple clustering bins, the combination of the $\gamma\delta_g$ and $\delta_g\delta_g$ correlations contain enough information to constrain the angular diameter distance to each lensing bin (and thus δz). The $\gamma\gamma$ correlations are then freed to further constrain cosmology.

We investigate the robustness of these results to perturbing the fiducial assumptions and calculate the degradation

$$D \equiv \sigma S_8(\Delta\delta z = 0.1)/\sigma S_8(\Delta\delta z = 0) - 1 \quad (5.2)$$

of 40% for the fiducial analysis. We characterise stochastic bias in the relation between mass and light with a scale-independent parameter per redshift bin $r_g^j \equiv C_{gG}^{ij}/C_{GG}^{ij}$ (see e.g. Dekel & Lahav 1999; Joachimi & Bridle 2010). We find that the three extra free parameters, $|r_g^j| < 6$, make little difference ($D = 45\%$), but increase $\sigma S_8(0)$ by 6% relative to the fiducial case. We next marginalise over an additional photo- z uncertainty parameter per bin, which stretches the redshift distributions $\tilde{n}^i(z) = n^i(z + S_z^i[z - z_p])$, where z_p is the peak in bin i . The cosmology constraint is weakened at all δz , which reduces the relative degradation slightly ($D = 28\%$). A similar reduction ($D = 24\%$) occurs if ℓ_{max} is increased by a factor of 3. We also rerun our analysis using a wide prior on the shear measurement bias ($\Delta m = 0.05$) and find $D = 16\%$, due to a $\sim 72\%$ degradation in the error on S_8 , independent of $\Delta\delta z^i$. The biggest impact arises when we use a conservative value for the prior on the photo- z uncertainties of the density sample ($\Delta\delta z^j = 0.05$), which gives a factor of

two degradation ($D = 103\%$), with $\sigma S_8(0.1)$ and $\sigma S_8(0)$ increased by 75% and 6% respectively relative to fiducial. In all instances considered, we find no significant residual bias δS_8 when δz^i are marginalised.




Figure 5.4: Uncertainty σS_8 on $S_8 \equiv \sigma_8(\Omega_m/0.31)^{0.5}$ for different prior widths $\Delta\delta z$. Blue circles show cosmic shear alone and purple diamonds indicate shear, galaxy clustering and cross-correlations. The green triangles also show the latter combination, but assuming intrinsic alignments and multiplicative shear calibration are known perfectly. In the first two cases we also show the bias δS_8 induced by marginalising with an erroneously zero centred prior of width $\Delta\delta z$ (dotted and dot-dashed).

To test the self-calibration result in a situation where the true redshift distribution differs from the assumed $n(z)$ by more than a simple bias we take the DES SV redshift distributions from an alternative code (BPZ). Of the SV codes BPZ was the most discrepant with our fiducial choice (SKYNET). To provide a relatively stringent test we choose not to apply a 0.05 calibration used in DES SV. Figure 5.6 shows the result. By construction the green contours in Figure 5.6 and 5.3 are identical and use the fiducial $n(z)$ in the simulation and the fit. The blue contours use the qualitatively different BPZ $n(z)$ in the simulation, but assume perfect photo- z ($\delta z = 0$, $\Delta\delta z^i = 0$) in the fit. Unlike in Figure 5.3, marginalising over photo- z bias ($\Delta\delta z^i = 0.1$) no longer trivially moves the contours (solid purple) onto the input cosmology (cross). This suggests a uniform bias in redshift may not always sufficiently account for differences between the true and estimated




Figure 5.5: Degeneracies between photo- z bias in the uppermost redshift bin (δz^3) and cosmology (S_8) for cosmic shear alone (blue dash-dot), galaxy clustering and shear-density cross-correlations (pink dotted) and the combination (purple solid). The input parameters are shown by the black cross.

distributions, depending on the survey specifications. Finally, we repeat the fiducial BPZ analysis twice, once fixing IA parameters and once additionally varying S_z , but find no qualitative change in the residual bias. In the case of BPZ the truth is within the 68% confidence contour, but a more detailed investigation is necessary to account for the possible range of redshift errors.

5.4 Conclusion

We have investigated the potential for current galaxy imaging surveys to self-calibrate photo- z distribution uncertainties, for the first time considering the case in which the shear sample is different from the galaxy clustering sample and has substantial calibration uncertainties. The focus was on a preliminary Stage III dataset with ~ 40 M galaxies, in which the galaxy clustering sample has well-understood photo- z ($\Delta\delta z^j = 0.01$). The combination of cosmic shear, galaxy clustering and shear-density cross-correlations was found to be significantly more robust to errors and uncertainties in the redshift distribution than cosmic shear alone. We showed that the uncertainty on $S_8 \equiv \sigma_8(\Omega_m/0.31)^{0.5}$ is increased by only 40% on marginalising over three free independent bias




Figure 5.6: As Figure 5.3 but now using a more realistic realisation of the discrepancy between the estimated and true galaxy redshift distributions. Here we use the $n^i(z)$ from an alternative photo- z code (BPZ) in the simulated data vector and the $n^i(z)$ of the fiducial photo- z code (SKYNET) in the fit.

parameters with prior $\Delta\delta z^i = 0.1$, relative to the case $\Delta\delta z^i = 0$. This contrasts with more than a factor of two degradation for cosmic shear alone. We illustrated that this is because cosmic shear constrains a different degenerate combination of cosmology and photo- z calibration parameters to clustering and shear-density cross-correlations. We have shown that the combination of all three two-point functions can correct even a substantial bias (of 0.05) in the $n(z)$ to accurately recover the input cosmology. This result is robust to a basic stochastic bias parameter and strengthened by less conservative scale cuts in the galaxy clustering analysis. We note two crucial differences between our results and established methods: (a) the current work do not rely on spectroscopic data, allowing a self-contained analysis of a single photometric survey and (b) the result is not contingent on having well matched samples in redshift coverage, allowing use of a (realistically) shallower clustering catalogue. The self-calibration result is weakened if redshift bias in the clustering sample is poorly controlled ($\Delta\delta z^j = 0.05$). Using an alternative $n(z)$ estimate (BPZ) We have demonstrated that this result may change if the deviation of the $n(z)$ from the truth is not fully captured by a uniform translation. By inspection of the distributions, we can see that the most prominent qualitative differences arise from secondary peaks or outliers. Unfortunately, such errors vary between photo- z methods and cannot be characterised analytically for a generic implementation. In practice the validity of our findings should be verified for specific realisations of the photo- z error. We note, however, that the uncalibrated BPZ results fail basic photo- z requirements, even for Stage III surveys. They should thus be considered as an extreme case and not a realistic prediction of photo- z performance for Stage IV.

This investigation advances on most previous numerical forecasts in implementing MCMC sampling rather than Fisher analyses and assumes a low-density clustering sample with relatively well known redshifts. We do, however, assume Gaussian covariance matrices, which tend to underestimate the uncertainties for cosmic shear and could thus make our forecast constraining power over-optimistic (although). Investigation of non-Gaussian covariances is beyond the scope of this analysis. We also assume the Limber approximation holds on the scales used and ignore redshift-space distortions. The results suggest that self-calibration may be a practical solution for current cosmological surveys, assuming reliable photo- z can be obtained for the galaxy clustering catalogue if the weak lensing redshift distributions cannot be easily calibrated via a different route.

Chapter 6

Colour Space Divisions and Intrinsic Alignments in the Dark Energy Survey

In this final chapter we now turn to another of the major challenges facing modern high-precision cosmology. The following presents a set of interim results from an ongoing investigation. The work is based around the DES Y1 dataset, and focuses on the impact of intrinsic alignments between galaxies on cosmic shear. We present new constraints on cosmology and IA parameters derived from the DES Y1 dataset. These results, we stress, are preliminary and have not yet been submitted for internal collaboration review. At the time of writing the shear catalogues used in this analysis are blinded. As an additional level of blinding, we remove the scale labels from any cosmologically sensitive axes in the following chapter.

6.1 Theory & Background

This section sets the context to the results presented in this chapter. Since they were first theorised as a contaminant to cosmic shear, there have been a number of attempts to measure the IA signal in real data. We provide a brief overview of previous published observational studies on the subject. The final paragraphs then set out the theoretical models for IA used in this analysis.

6.1.1 Observational Constraints on Intrinsic Alignments

Attempts to constrain IA in measurements of cosmic shear broadly fall into two categories. The first are *direct* constraints. These typically use galaxies at low to intermediate redshift, and often

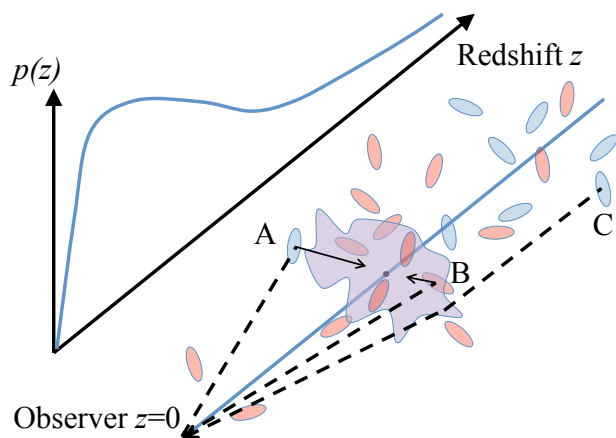


Figure 6.1: Graphical illustration of a toy model for a red/blue intrinsic alignment split. Galaxy C is a background object at high redshift, and is lensed by a foreground clump of dark matter (shown in purple). Objects physically close to the dark matter clump interact with it directly; elliptical red galaxies (such as B) tend to reside within massive clusters, and their interactions with their host halo are dominated by gravity. They tend to align radially towards the centre of mass, inducing an anticorrelation with the lensed background galaxy. Blue galaxies, on the other hand, tend to be supported by rotation rather than pressure, and will thus become aligned with the spin axis of the halo. This tends to create much milder distortions in the tangential direction.

impose colour cuts to isolate well-measured red galaxies and assume some fixed known cosmology. Correlation statistics used in these measurements are explicitly designed to maximise the IA signal (e.g. Faltenbacher et al. 2009).

The second class of IA measurements are implicit, or *simultaneous* constraints. Generally they use statistics designed to be sensitive to cosmic shear such as ξ_{\pm} , and use faint high-redshift galaxies in which the cosmological signal is strongest. These studies include IA parameters alongside various others, including cosmology, in a set of free parameters with which the data are fit. Whether or not they acknowledge as much, all cosmic shear studies in which IA are marginalised rather than suppressed in some other way fall into this category (Heymans et al., 2013; Dark Energy Survey Collaboration, 2016; Jee et al., 2016; Hildebrandt et al., 2016; Köhlinger et al., 2017). The assumptions about IA have differed slightly between these studies, but they have all essentially assumed the same basic model (the nonlinear alignment model), sometimes with additional scaling with redshift or luminosity.

There is some direct evidence for differences in the IA contamination, depending on the nature of the galaxy sample. Broadly there are two paradigms: early-type ellipticals, which tend to be redder and structurally pressure dominated, and late type spirals, which tend to be bluer and rotation dominated. A simplified illustrative diagram of IA in these two populations is shown in Figure 6.1. The former are thought to align through tidal interactions with the background large scale structure of the Universe. If a dark matter halo sits in a local gradient in the gravitational field, it will be sheared along that gradient. Naturally, the visible light-emitting matter will trace

the halo. Spatially neighbouring galaxies will, then, become aligned since they sit in the same background tidal field. If the distortion is small, the induced ellipticity can be assumed to be linear in the gravitational potential.

The first observational constraints on intrinsic alignments in red galaxies were published midway through the first decade of the Twenty First Century. The Sloan Digital Sky Survey was a major resource for such studies, as it afforded relatively high signal-to-noise correlation measurements, and high-quality spectroscopic overlap (see e.g. Mandelbaum et al. 2006; Hirata et al. 2007; Okumura et al. 2009; Joachimi et al. 2011; Li et al. 2013; Singh et al. 2015). The highest maximum redshift reached by the red galaxy catalogues constructed from SDSS data extend to $z \sim 0.6$. They found relatively strong IA signals, but no statistically significant detection of redshift dependence.

The picture for late-time galaxies is very different. These objects form galactic discs which, depending on the orientation, may have some apparent ellipticity. The inclination of the disc is dictated by its angular momentum. The mechanism by which that angular momentum is induced and may become correlated between galaxies, however, is not immediately apparent. One common picture is that galaxy spin is generated by tidal torquing. That is, the background gravitational field exerts a moment on a halo in its early stages of development. At the point where a galactic disc is eventually formed, after collapse and then cooling, the stellar disc takes on the ellipticity of its host halo.

Measurements of large scale correlations of the apparent (projected) ellipticities of blue galaxies are generally relatively weak. Hirata et al. (2007) used a blue sample from SDSS, but could impose only upper limits on the IA signal at $z \sim 0.1$. Beyond SDSS, there have been notable examples of measurements using spectroscopic galaxies from the WiggleZ survey. Mandelbaum et al. (2011) used these data to attempt an IA detection in blue (starburst) galaxies at $z \sim 0.6$, but likewise set upper limits rather than a definite constraint.

6.1.2 Theory Predictions

Shear Two-Point Functions

The cosmological analysis in this chapter uses real-space angular correlation functions in four tomographic bins. The measurements map onto the angular power spectrum of shear via Hankel transforms:

$$\xi_{\pm}^{ij}(\theta) = \frac{1}{2\pi} \int \ell J_{0,4}(\ell\theta) C_{\gamma}^{ij}(\ell) d\ell \quad (6.1)$$

where the indices ij indicate a pair of tomographic bins, and J_0 and J_4 are Bessel functions. The angular spectrum C_γ is related to the dark matter power spectrum under the Limber approximation as,

$$C_\gamma^{ij} = \int_0^{\chi_{\text{hor}}} \frac{1}{S_K^2(\chi)} g^i(\chi) g^j(\chi) P_\delta \left(k = \frac{\ell}{S_K(\chi)}, z \right) d\chi. \quad (6.2)$$

As in the introductory chapters, S_K is the transverse angular diameter distance, χ_{hor} is the comoving horizon distance, and the lensing kernel in each bin is given by

$$g^i(\chi) = \frac{3}{2} \frac{H_0^2 \Omega_m}{c^2} \frac{S_K(\chi)}{a(\chi)} \int_\chi^{\chi_{\text{hor}}} n^i(\chi') \frac{S_K(\chi' - \chi)}{S_K(\chi')} d\chi' \quad (6.3)$$

The redshift distributions $n(z)$ are assumed to be normalised over the depth of the survey, and defined such that $n(z)dz = n(\chi)d\chi$. As in many previous studies, we do not transform the measured correlations into angular spectra. Likelihoods for trial cosmologies are calculated by generating theory angular spectra, which are convolved with the Bessel kernels, resampled at the appropriate angular scales, and then compared with the measurements of ξ_\pm^{ij} .

Although previous studies have explored the benefits of using different estimators, we choose not to pursue this subject here. Though useful for identifying and removing B-modes (e.g. van Uitert et al. 2017), recent results have shown no benefit to cosmological constraining power from alternative statistics such as Complete Orthogonal Sets of E-/B-mode Integrals (COSEBIs; Schneider et al. 2010) or band powers (e.g. Dark Energy Survey Collaboration 2016; Köhlinger et al. 2017).

To extract cosmological parameters from the data, we used the EMCEE Markov Chain Monte Carlo package released with COSMOSIS. The dark matter power spectrum is estimated at each cosmology using CAMB, with late-time corrections generated by HALOFIT. We do not explicitly model baryonic effects or scale dependent IA, but as noted in the next section our choice of scale cuts is relatively conservative and is expected to remove the régime in which such effects are significant.

Modelling Intrinsic Alignments

Even highly accurate galaxy shape correlations are not a pure measurement of the cosmic shear spectrum. Specifically the correlations between the intrinsic (pre-shear) shapes contribute an unknown additive terms of the form

$$C_\epsilon^{ij}(\ell) = C_{\text{GG}}^{ij}(\ell) + C_{\text{II}}^{ij}(\ell) + C_{\text{GI}}^{ij}(\ell). \quad (6.4)$$

The “GI” term represents correlations induced by foreground matter overdensities, which lenses background galaxies and influences the unlensed orientation of galaxies physically close to them (see Figure 6.1). The “II” contribution, which tends to be subdominant and opposite in sign to GI, represents direct local interactions between nearby galaxies. Note that the cosmological component C_{GG} is synonymous with C_γ . We will adopt the former notation for the rest of this chapter for consistency with the IA literature.

Commonly the power spectra of these correlations, defined analogously to $P_\delta(k)$, are written as

$$P_{\text{GI}}(k, z) = F(z)P_\delta(k, z), \quad (6.5)$$

and

$$P_{\text{II}}(k, z) = F^2(z)P_\delta(k, z). \quad (6.6)$$

The most common approach is to assume a single universal intrinsic alignment model governed by a handful of parameters. Owing to its good performance in matching data and simulations, a model known as the nonlinear alignment model (Bridle & King, 2007) has become particularly popular. This is an empirical modification to the linear alignment model of Hirata & Seljak (2004), whereby the linear matter power spectrum is replaced by the nonlinear spectrum. The result is

$$F(z) = -A_{\text{IA}}C_1 \frac{3H_0^2\Omega_m}{8\pi G} D^{-1}(z) \left(\frac{1+z}{1+z_0} \right)^\eta. \quad (6.7)$$

The dimensionless amplitude A_{IA} is an unknown scaling parameter governing the strength of the IA contamination for a particular sample of galaxies, and is generally left as a free parameter to be constrained. As before G is the gravitational constant and $D(z)$ is the linear growth factor. The normalisation constant C_1 is typically fixed at a value obtained from the SuperCOSMOS Sky Survey by Brown et al. (2002) of $C_1 = 5 \times 10^{-14} M_*^{-1} h^{-2} \text{Mpc}^3$. The redshift evolution is governed by a power law index η , which has been measured in low redshift LRG samples (Joachimi et al., 2011). The denominator $1 + z_0$ sets a pivot redshift, for which we assume $z_0 = 0.62$ whenever equation 6.7 is used in this chapter.

In this chapter we introduce a slight generalisation to this formalism. Let us start with the statement that the shear of a galaxy i is the sum of its intrinsic and cosmological parts $\gamma^i = \gamma_{\text{G}}^i + \gamma_{\text{I}}^i$. Correlating galaxy pairs gives the expression,

$$\langle \gamma \gamma^* \rangle = \langle \gamma_G \gamma_G^* \rangle + \langle \gamma_I \gamma_G^* \rangle + \langle \gamma_G \gamma_I^* \rangle + \langle \gamma_I \gamma_I^* \rangle. \quad (6.8)$$

Let us now assume that the galaxies used in the correlation are drawn from two populations R and B , with the former accounting for a fraction f_R of the random draws. We will say, for the sake of argument, that galaxies from both populations exhibit intrinsic shape correlations, but that those correlations may be arbitrarily different in their strength and redshift dependence. In this case one finds,

$$\begin{aligned} P_{\text{II}}(k) &= \langle \gamma_I \gamma_I^* \rangle \\ &= f_R^2 \langle \gamma_{I,R} \gamma_{I,R}^* \rangle + (1 - f_R)^2 \langle \gamma_{I,B} \gamma_{I,B}^* \rangle + 2f_R(1 - f_R) \langle \gamma_{I,R} \gamma_{I,B}^* \rangle \\ &= f_R^2(z) P_{\text{II}}^{RR}(k, z) + (1 - f_R(z))^2 P_{\text{II}}^{BB}(k, z) + 2f_R(z)(1 - f_R(z)) P_{\text{II}}^{RB}(k, z). \end{aligned} \quad (6.9)$$

A similar exercise for GI gives

$$P_{\text{GI}}(k) = \langle \gamma_G \gamma_I^* \rangle = f_R(z) P_{\text{GI}}^R(k, z) + (1 - f_R(z)) P_{\text{GI}}^B(k, z), \quad (6.10)$$

with $\langle \gamma_a \gamma_b \rangle$ denoting averaging galaxy pairs drawn from populations a and b . Note that only one index is specified here, since the shear signal by foreground matter is the same, irrespective of what sort of galaxy is being lensed. We may have a mixed population of galaxies governed by different IA models. In this toy model, however, one can construct effective GI and II spectra relevant for the mixed sample as a linear combination of auto- and cross-spectra for the two sub-samples.

The Limber projections of these spectra in bins ij are simply expressed in the form

$$C_{\text{II}}^{ij}(\ell) = \int \frac{1}{S_K^2(\chi)} n^i(\chi) n^j(\chi) P_{\text{II}} \left(k = \frac{\ell}{S_K(\chi)}, \chi \right) d\chi \quad (6.11)$$

and

$$C_{\text{GI}}^{ij}(\ell) = \int \frac{1}{S_K^2(\chi)} g^i(\chi) n^j(\chi) P_{\text{GI}} \left(k = \frac{\ell}{S_K(\chi)}, \chi \right) d\chi, \quad (6.12)$$

where the composite power spectra used in the intergrals are given by equations 6.9 and 6.10. Note that the form of the blend fraction f_R as a function of redshift must be modelled appropriately for the galaxies in each tomographic bin.

6.2 Data

In this section we define the galaxy samples used in this chapter. We define disjoint samples from the DES Y1 WL catalogue, intended to isolate morphological differences relevant to IA. The following paragraphs discuss the practical details of the split, including how we manage selection effects.

6.2.1 Year One of The Dark Energy Survey

The Dark Energy Survey is a rich and extremely powerful dataset for exploring this subject. We have already considered at some length the survey characteristics and shear calibration process, and refer the reader to Z17, Troxel et al. (2017), and Chapter 3 of this thesis for details. Where the earlier parts of this thesis focused on one of two shear catalogues, IM3SHAPE, this section will make use of the other, METACALIBRATION. This decision was partly due to the larger number density afforded by this catalogue, but primarily because the method of per-galaxy shear correction allows more flexibility for dividing the data. Though it does not give us complete free reign, the response correction can accommodate any cut based on a quantity which has been remeasured on a set of sheared copies of the data (see Sheldon & Huff 2017 and Section 6.2.3 below).

6.2.2 Definitions of Red & Blue Galaxies

There are a number of terms used in the literature to class galaxies, which are broadly analogous but non-identical. This chapter makes use of three, which are summarised in Table 6.1. Though these names are often used somewhat interchangeably in the literature, in the following analysis the terms “early-type”, “red”, “bulge” etc have distinct meanings, as set out below.

A common metric is spectral class. Template-based photo- z codes such as BPZ work by red-shifting a library of spectral templates, and each galaxy is associated with one of these templates (or some interpolated combination). Each redshift estimate, then, comes with an explicit galaxy type estimate, defined by the best-fitting template T_{BPZ} . This quantity has been used in previous studies to divide galaxies expected to have different systematics (Simon et al., 2013; Heymans et al., 2013). We follow those papers and define a boundary at $T_{BPZ} = 1$ to separate “early-type” and “late-type” galaxies.

Another metric is photometric colour, defined by differences in magnitude in different optical bands. One could, for example, define red and blue galaxies according to some value in $r - z$. The 2D histograms of galaxies in colour magnitude space is bimodal, as is shown for the DES




Figure 6.2: The 2D colour-magnitude distribution of galaxies in the DES Y1 METACALIBRATION shear catalogue. The two panels show photometric colour as defined by difference between g and z , and r and z band magnitudes respectively. The colours show the distributions of early- and late-type galaxies. Note that the red contours account for $\sim 20\%$ of galaxies in DES Y1.

Split	Defined By	Division
Early/Late	Spectral Type	$T_{\text{BPZ}} < 1 / T_{\text{BPZ}} > 1$
Red/Blue	$r - z$ Colour	equation 6.13
Bulge/Disc	Sérsic Fit Likelihood	$n = 1 / n = 4$

Table 6.1: Definitions of red and blue galaxies used in this chapter. Though roughly equivalent, these three catalogue splits define different subsets of DES Y1.

METACALIBRATION catalogue in Figure 6.2. The red and blue distributions show early-type and late-type populations, as defined above. Perhaps unsurprisingly, the $g - z$ band difference shows the clearest separation. Since we do not have shape measurements on the g -band images, we cannot define cuts on the METACALIBRATION catalogue using $g - z$. In the following we use a boundary in the $r - z$ plane (right) to define red and blue galaxies. Specifically, we find a line

$$(r - z) = a_{rz} \times r + b_{rz}, \quad (6.13)$$

with $a_{rz} = 0.115$ and $b_{rz} = -1.327$ approximates the division seen in the right-hand panel of Figure 6.2.

The third definition we consider is specific to IM3SHAPE, and cannot be implemented in METACALIBRATION. Since IM3SHAPE performs simultaneous fits to bulge and disc Sérsic profiles, we have a crude morphology measure for each galaxy based on the relative likelihoods of the fits.




Figure 6.3: Distributions of photometric colour, defined as ratios of galaxy fluxes in the g , r and z optical bands f_g , f_r and f_z . The blue and red shaded curves show early- and late-type galaxies alone. The dashed purple line indicates the full sample, without morphological splits. For ease of comparison, the histograms are not re-normalised in these panels.

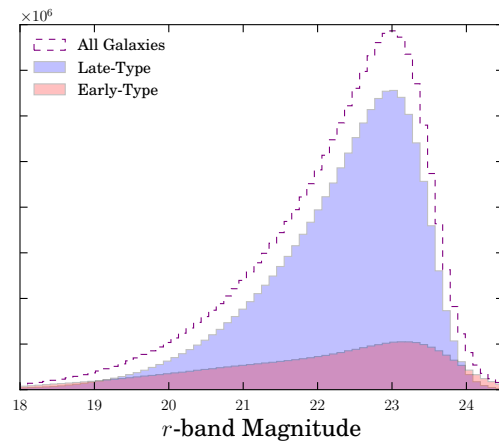


Figure 6.4: MOF r -band magnitude distributions of the galaxy samples defined in this chapter. The dashed purple line is defined by the METACALIBRATION selection flag only, and corresponds to the dark red histogram in Figure 3 of Z17.

Sample	N_{gal}	z^{med}	\bar{r}	$g - z$	$r - z$
All Galaxies (METACALIBRATION)	34.8 M	0.55	22.0	1.64	0.75
Early	6.8 M	0.65	22.1	2.67	1.28
Late	26.8 M	0.53	21.7	1.35	0.62
Red	5.4 M	0.81	22.3	2.93	1.52
Blue	29.4 M	0.51	22.0	1.42	0.62
All Galaxies (IM3SHAPE)	21.9 M	0.69	22.5	1.52	0.69
Bulge	4.0 M	0.69	22.8	1.64	0.76
Disc	17.9 M	0.84	22.4	1.50	0.68

Table 6.2: Observational characteristics of the sub-populations defined in this Table 6.1. The upper and lower sections relates to subsets of the two catalogues, METACALIBRATION and IM3SHAPE. Note that these statistics are weighted appropriately for each shape code, using the mean galaxy response $R = (R_{11} + R_{22})/2$ and the mean multiplicative bias $1 + m$ respectively.

These two populations are calibrated independently, and so one can split on this quantity without inducing selection bias. Clearly the IM3SHAPE bulge/disc categorisation is somewhat degenerate with redshift, and so we do not expect this to match up precisely to more sophisticated measures of galaxy morphology.

For any given definition we can define a quantity called the “red fraction”, $f_R^j(z)$, which varies as a function of redshift for a population of galaxies j . If we had perfect knowledge of the redshift of each galaxy we could simply sort the catalogue into fine bins and count the number of red galaxies in each redshift slice. In practice we instead have redshift distributions, quantifying the uncertainty in our photometric redshift estimates. We thus define the red fraction as:

$$f_R^j(z) = \frac{p(z, R)}{p(z)} = \left(\frac{N_R^j}{N_R^j + N_B^j} \right) \left(\frac{p(z|R)}{p(z|R) + p(z|B)} \right), \quad (6.14)$$

where N_c^j is the number of galaxies in population j observed to be of colour c , and $p(z|c)$ is the conditional probability of a galaxy lying at redshift z , given its colour. That is, the red fraction in this definition is the probability of a random galaxy from the catalogue being red *and* at redshift z , divided by the chance of a random galaxy of any colour being at z .

The estimated redshift distributions of early- and late-type galaxies in four tomographic bins, and the derived red fraction are shown in Figure 6.5.




Figure 6.5: The measured redshift distributions of the galaxy samples described in this Chapter in the four tomographic bins defined by Troxel et al. (2017). The left-hand panels show our fiducial subsets of the DES Y1 dataset, split according to best-fitting template type T_{BPZ} . The solid curves show estimates derived from the photometric redshift code BPZ. The dashed lines are the histograms of redshift estimates for a reweighted COSMOS sample, based on the 30-band photometry available to that survey. The right panel shows the red fraction, as defined by equation 6.14 in each of the tomographic bins shown.

6.2.3 Shape Measurements

For this work we use two different shear catalogues in DES Y1. Conventionally, shape measurements rely on large suites of image simulations to calibrate measurement biases. We provide a detailed discussion of the processes involved in constructing and testing such a calibration in Chapter 3. As we point out in that paper, additional selection can very easily induce multiplicative shear bias. Propagating the same cuts to the simulation and recalibrating is a time-consuming undertaking.

Our fiducial choice for this analysis, however, is the second shape catalogue. These measurements use a technique known as METACALIBRATION, the basis of which is to derive the calibration from the data itself using counterfactual copies of each galaxy image with additional shear applied. The algorithm remeasures the shear for each galaxy and computes a quantity known as the *response*:

$$\mathbf{R}_{\gamma_i, j} = \frac{e_i^+ - e_i^-}{\Delta\gamma_j}, \quad (6.15)$$

where e^+ and e^- are the measured values of the ellipticity obtained from images sheared by $+\gamma$

and $-\gamma$, and $\Delta\gamma = 2\gamma$. The galaxy response must be included whenever a shape-derived statistic is calculated. We refer the reader to Sheldon & Huff (2017); Huff & Mandelbaum (2017) for a full explanation of the algorithm and to Z17 for details of the implementation used in DES Y1, and a recipe for applying response corrections.

It is also possible to correct for selection bias using a similar calculation. To do this we must measure the response of the mean ellipticity to the cut. Imagine for example, we wish to make a cut on galaxy type T_{BPZ} . Since T_{BPZ} is not independent of ellipticity, the raw cut will induce shear selection bias. We must, then, estimate T_{BPZ} three times per galaxy: once in the original images, and in each set of counterfactual sheared images. If we choose to keep or exclude galaxies based on a measured quantity that changes in the sheared images, the cut will preferentially remove either round or elliptical galaxies. A mean response $\langle \mathbf{R}_S \rangle$ contributed by a selection alone is defined as the change in ellipticity

$$\langle \mathbf{R}_S \rangle_{i,j} \approx \frac{\langle e_i \rangle^{S+} - \langle e_i \rangle^{S-}}{\Delta\gamma_j}, \quad (6.16)$$

where $\langle e \rangle^{S\pm}$ denotes the mean ellipticity measured from the unsheared images, after cuts on the sheared quantity $T_{\text{BPZ}}^{S\pm}$. The full response for the mean shear is then given by the sum of the shear and selection parts,

$$\langle \mathbf{R} \rangle = \langle \mathbf{R}_\gamma \rangle + \langle \mathbf{R}_S \rangle. \quad (6.17)$$

This must be recalculated each time galaxies are split in any way, including for tomographic binning.

6.2.4 Photometric Redshifts

We derive estimates for the redshift distribution of our samples using the BPZ code. The per-galaxy PDFs are generated by the DES Redshift Working Group. They have been tested against simulations, against a limited spectroscopic sample, and against alternative redshift algorithms (Hoyle et al., 2017). For each sample used in this chapter the PDFs are re-stacked in tomographic bins.

In addition to shear selection bias, splitting on T_{BPZ} risks rendering the $n(z)$ estimates inaccurate. In DES Y1 we have a Monte Carlo sample and the mean of the full redshift PDF of each galaxy. It is thus relatively easy to recompute the redshift distribution for arbitrary subsample of galaxies; this is simply a matter of re-binning the galaxies by mean redshift and taking the his-

togram of the MC samples in each bin. We show the measured $n(z)$ obtained by the BPZ code for $T_{\text{BPZ}} < 1$ and $T_{\text{BPZ}} > 1$ galaxies in Figure 6.5.

The raw $n(z)$, however, are known to be systematically biased. The main shear selection defined by Z17 has been subjected to a rigorous set of tests designed to constrain this redshift bias. This information is incorporated into cosmic shear analyses via (non-zero centred) priors on redshift nuisance parameters. Unfortunately, one cannot guarantee that these priors will be robust to arbitrary division of the data. If we propose to use any subset of the catalogue for tomographic shear measurements, it is necessary to re-derive appropriate photo- z priors.

To do this we use galaxies from the partially overlapping COSMOS field. The low-noise 30-band photometry from HST provides high-quality point redshift estimates for these galaxies. In the following we will take these as “true” redshifts. In principle we can test for bias for a particular sample by comparing the distribution of the COSMOS redshifts to our ensemble estimates for the same galaxies in the DES images. Selecting the galaxies in the COSMOS overlap, however, can itself induce selection effects, since the COSMOS galaxies are somewhat unrepresentative of DES in magnitude, colour and size. The COSMOS catalogue is thus resampled such that the resulting sample matches the DES Y1 data. The process is described in more detail by Hoyle et al. (2017), and results in a set of 200,000 DES galaxies matched to COSMOS counterparts with similar flux in four bands $griz$ and size.

We divide these galaxies into four tomographic bins according to mean redshift, as estimated from a re-run of BPZ on the artificially noisy COSMOS fluxes. In each bin we compute a weighted mean

$$\langle z \rangle^{(i)} = \frac{\sum_j^{N_{\text{gal}}^{(i)}} w_j z_j^{\text{C}}}{\sum_j^{N_{\text{gal}}^{(i)}} w_j}, \quad (6.18)$$

where z_j^{C} is the COSMOS redshift estimate for galaxy j , $w_j \equiv (R_{1,j} + R_{2,j})/2$ is its mean response, which serves here as a weight on the shape measurement, and the sum runs over all galaxies placed in redshift bin i . The offset between the mean COSMOS redshift and the equivalent weighted mean using the BPZ MC samples provides a constraint the level of systematic bias in the latter. We derive δz in this way for our early, late and full samples, as defined by T_{BPZ} . The result is shown in Table 6.3.

These values set the central values of the redshift priors. In order to decide on an appropriate prior width we must consider a number of sources of uncertainty in this measurements. We subject the reweighted COSMOS dataset described to a series of tests designed to constrain systematic

Selection	$\delta z^{(1)}$	$\delta z^{(2)}$	$\delta z^{(3)}$	$\delta z^{(4)}$
All Galaxies	-0.006 ± 0.018	-0.014 ± 0.018	0.018 ± 0.017	-0.018 ± 0.018
Late-Type	-0.003 ± 0.020	-0.007 ± 0.023	0.0297 ± 0.020	-0.010 ± 0.023
Early-Type	-0.022 ± 0.020	-0.040 ± 0.012	-0.008 ± 0.012	-0.044 ± 0.014

Table 6.3: Priors on the redshift error derived from a matched sample of COSMOS field galaxies described in the text.

uncertainties in a way analogous to the exercise in Section 3.5.5. This includes redshift error contributions for statistical uncertainty, cosmic variance, and the limited matching process using flux and size only. The resulting prior widths in each sample are also shown in Table 6.3.

In the following we adopt fiducial Gaussian priors for each sample centred according to Table 6.3 and with widths given by the above calculation.

6.2.5 Two-Point Correlations

The Y1 shear catalogues are used to construct two-point correlation functions of cosmic shear. Our method and choice of statistics and binning follows Troxel et al. (2017). In brief, we use the public code TREECORR to measure ξ_+ and ξ_- in 20 log-spaced bins in angular scale. Galaxy ellipticities are rotated, weighted and averaged in each bin according to equation 3.8.

To avoid the effects of theoretical uncertainties on small scales, we impose a lower-level angular scale cut in each bin. These bounds are relatively stringent compared with contemporary shear analyses, and are set out in more detail in Troxel et al. (2017). An upper cut of $\theta < 250$ arcminutes is imposed on all bins to remove scales on which additive shear biases become dominant. The correlation is corrected with an overall selection response, as prescribed by Sheldon & Huff (2017) and Z17.

These measurements have passed a raft of null tests, and show no indication of significant B-modes. We measure the two-point correlations separately in the full catalogue, and also in our fiducial early-type and late-type samples. The results are shown in Figure 6.6.

6.2.6 Covariance Matrix

The covariance matrix of the two-point data is estimated using the COSMOLIKE software package (Krause & Eifler, 2016). This code has been tested against log-normal simulations which include the DES survey mask (Krause et al., 2017), and against N-body simulations (MacCrann et al.,

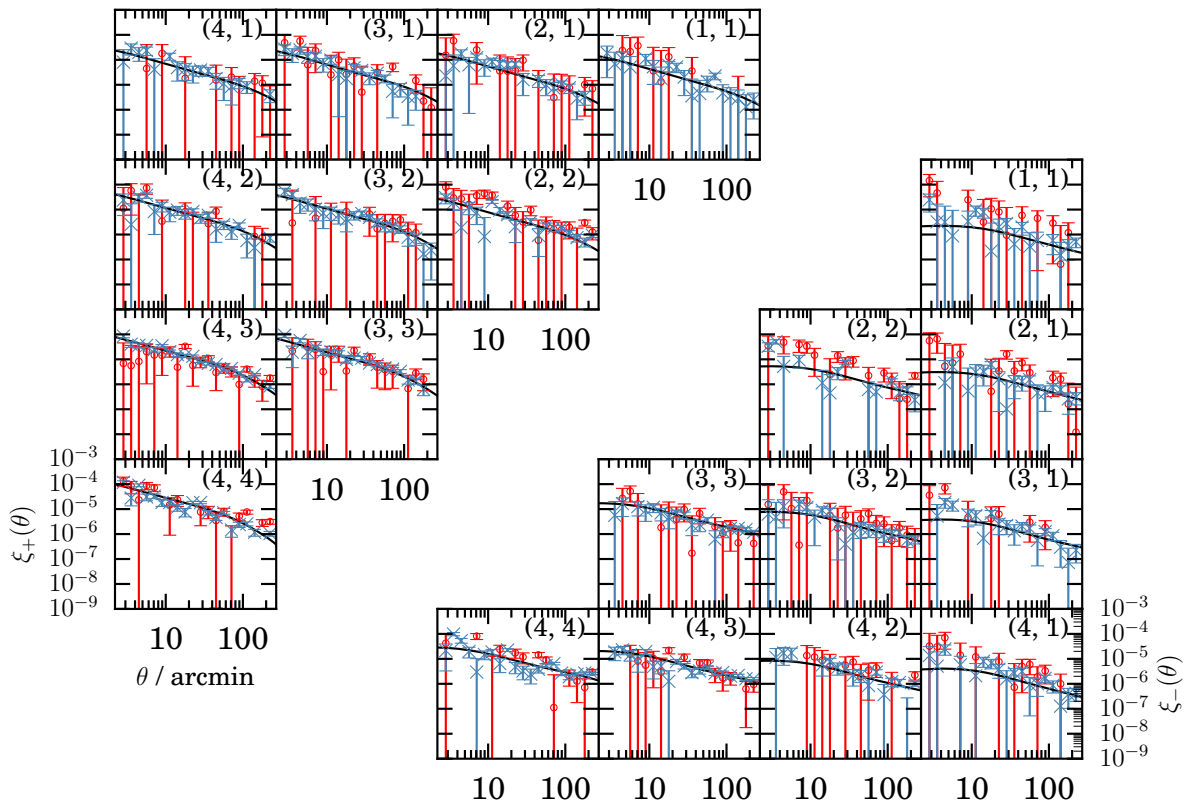


Figure 6.6: Measured two-point shear correlation functions from the DES Y1 shape catalogue in four tomographic bins. The red circles show correlations between early-type galaxies, and blue crosses are the equivalent late-type correlations. The errorbars are obtained from the covariance matrix described in Section 6.2.6. In each panel the solid line shows a theory prediction of the relevant ξ_{\pm} with all parameters (cosmological and nuisance) fixed at their best-fitting values from the full Y1 catalogue.

2017) for a DES Y1-like configuration and is seen to perform favourably. For computational reasons, given the relatively large number of different samples discussed, we choose to use Gaussian covariance matrices in this chapter. The COSMOLIKE code has the capacity for more advanced halo model estimates, which we intend to incorporate into this analysis in future.

Like almost all previous studies of cosmic shear our covariance matrix does not include the impact of intrinsic alignments. In a similar analysis based on CFHTLenS, Heymans et al. (2013) justify this in two ways. First, the galaxies used in cosmological weak lensing measurement are typically not dominated by the low redshift red population where IAs are known to be strong. Indeed, most simultaneous constraints using the full catalogues of contemporary shear surveys have found $A_{IA} \sim 0 - 1$. The impact on the covariance for the full sample is thus expected to be small. Second, extracting redder galaxies for IA measurements leaves $\sim 20\%$ or less of the full catalogue. In such a small sample the covariance is likely dominated by shot noise.

Since DES Y1 is significantly larger and slightly shallower than CFHTLenS, we seek to verify these assumptions. To test this we use a fast analytic code to generate Gaussian covariances for the shear-shear angular power spectrum $C_{\gamma\gamma}$ in DES Y1-like tomographic bins. The IA power spectra are modelled using the NLA model with a range of amplitudes. The results are shown in Figure 6.7.

Unsurprisingly (since the dominant GI term will tend to suppress power in the cosmic shear signal) on most scales ignoring IA in the covariance matrix leads one to overestimate the uncertainties. This is particularly true in the autocorrelation of the lower redshift bins. The negative sign on the largest scales (small ℓ) indicates that we may be underestimating our errorbars here. In all but the most extreme IA scenarios, however, the impact is comfortably under half of the 1σ contour.

6.3 Results

In this section we set out the main results of this investigation. The following paragraphs describe the fiducial analysis, in which we impose simultaneous constraints in IA and cosmology using different populations of galaxies in DES. We then explore the dependence of the IA contamination on redshift and the observed magnitude of the galaxies, a poorly understood topic at any level, and test the robustness of our results to various forms of error.

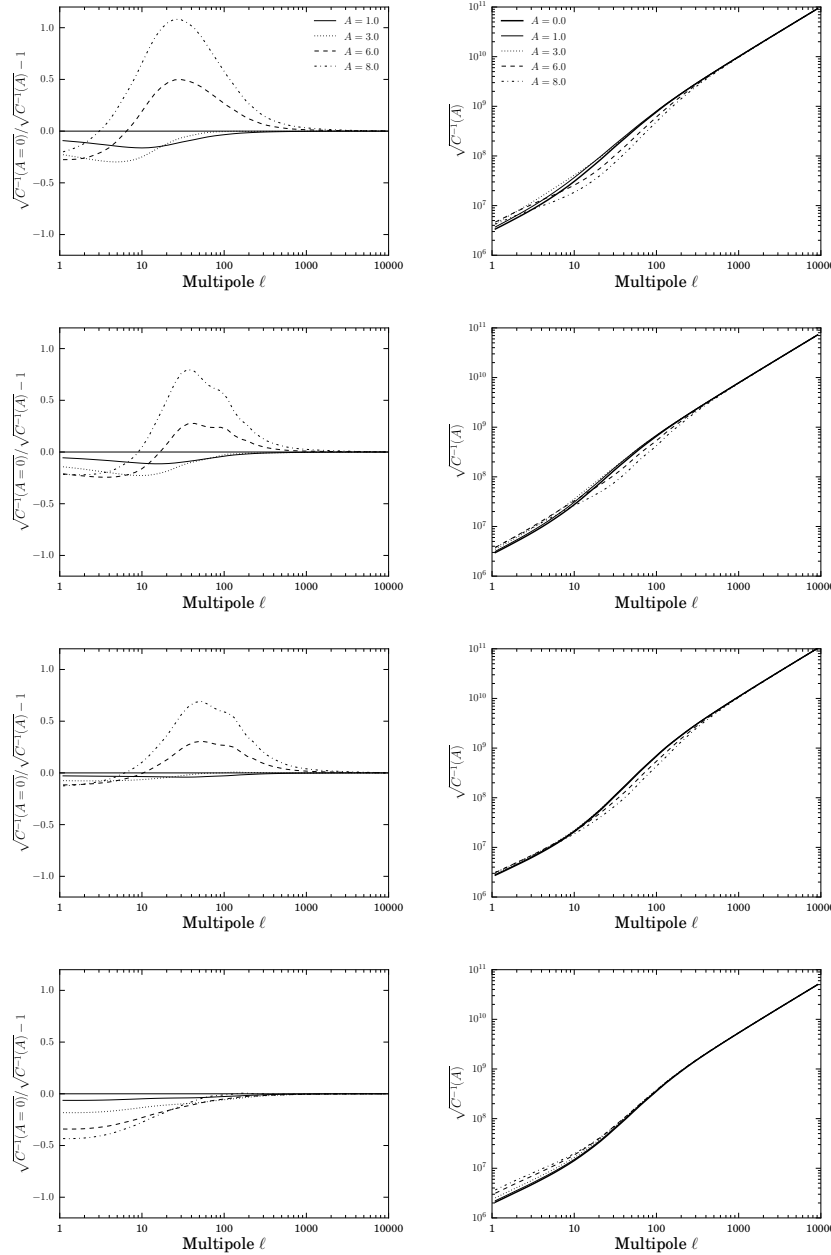


Figure 6.7: A comparison of cosmic shear covariance as a function of angular scale, as calculated with and without intrinsic alignments. In each case a covariance matrix of was computed for the shear-shear angular angular power spectrum, assuming Gaussian density fluctuations. The redshift distributions and number density in each of four tomographic bins and the total survey area were taken directly from the early-type DES Y1 sample presented in the main body of this work. **Left:** Fractional change in the inverse of the diagonal elements of the $\langle C_{\gamma}^{ii} C_{\gamma}^{ii} \rangle$ covariance matrix when intrinsic alignments are omitted from the calculation. The four lines show different amplitudes of the IA power spectrum, as modelled using the NLA model. **Right:** Raw (unnormalised) values of the same inverted covariance elements in scenarios with the IA amplitudes shown.

6.3.1 Simultaneous Constraints

The fiducial analysis in this chapter includes no external information beyond the Y1 shear catalogues. We consider three samples: (a) the full undivided catalogue, (b) late-type galaxies ($T_{\text{BPZ}} > 1$) which account for roughly 80% of the total 35M, and (c) early-type galaxies ($T_{\text{BPZ}} < 1$), which make up the other $\sim 20\%$. Monte Carlo chains are run independently to generate a likelihood of cosmological and nuisance parameters for each of these datasets.

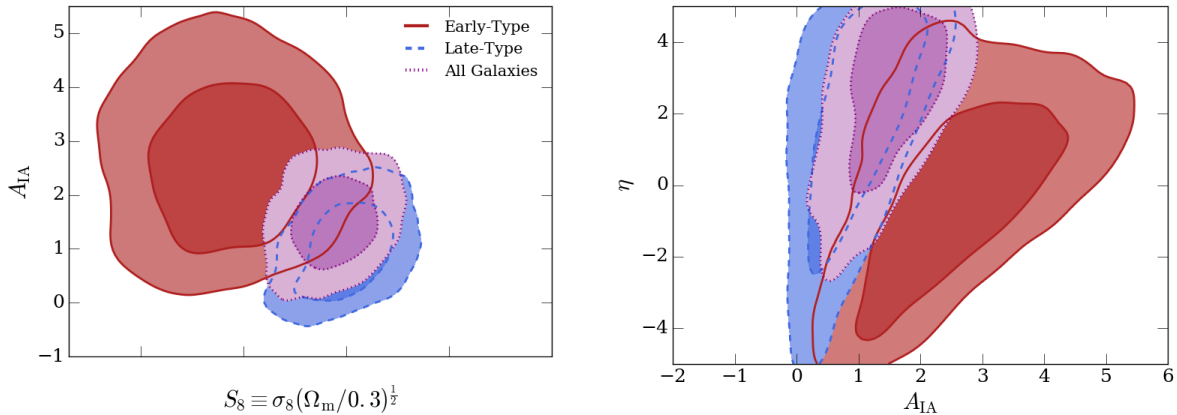


Figure 6.8: Joint constraints on cosmology and intrinsic alignment parameters from sub-populations of the DES Y1 fiducial shear catalogue. Note that the shape catalogues are blinded at the time of writing. The two sets of filled contours are defined by a split according to best-fitting SED ($T_{\text{BPZ}} < 1$ and $T_{\text{BPZ}} > 1$) to isolate early- and late-type galaxies. The purple dotted contours show the equivalent constraints from the full catalogue. Each population is modelled independently using a version of the NLA IA model with a free amplitude A_{IA} and redshift power law index η .

We use a Gaussian prior on residual shear bias of width $\sigma_m = 0.025$ in each tomographic bin, as recommended in Z17. The prior width here is dominated by blending, which we do not expect to depend on galaxy type (at least to first order). We thus do not alter the m prior between subsets of the catalogue. Selection responses are, however, recalculated as described earlier.

The same tomographic binning is used, and single-galaxy redshift PDFs are restacked for each new sample. As the redshift distribution changes significantly under different selections (see Figure 6.5), we rerun the tests using the resampled COSMOS catalogue described in Section 6.2.4 to generate a new set of priors on the photo- z bias parameters.

Intrinsic alignments are modelled using the NLA model (equation 6.7) with a single amplitude and redshift index.

Sample	A_{IA}	η
All Galaxies	$1.48^{+0.48}_{-0.47}$	$2.42^{+1.13}_{-1.62}$
Early-Type	$2.60^{+0.92}_{-1.04}$	$-0.22^{+1.59}_{-1.77}$
Late-Type	$0.97^{+0.52}_{-0.46}$	$1.73^{+1.52}_{-2.29}$
Early-Type, COSMOS $n(z)$	$2.36^{+0.86}_{-0.79}$	$0.03^{+1.53}_{-1.62}$
Late-Type, COSMOS $n(z)$	$0.93^{+0.53}_{-0.42}$	$2.11^{+1.32}_{-2.40}$
Early-Type, Refined	$1.19^{+1.15}_{-1.71}$	$-0.94^{+2.56}_{-1.88}$
All Galaxies, $\eta = 0$	$1.04^{+0.34}_{-0.43}$	0.00 (Fixed)

Table 6.4: Best-fitting Nonlinear Alignment Model parameters from subsets of DES Y1. The parameter value given in each case is the mean value of the marginalised posterior distribution. The errorbars indicate the upper and lower 68% confidence bounds relative to the mean. The upper section present the results from our fiducial split samples, and the lower part show a number of robustness tests described in the text. but entries below the line

In each chain we vary six cosmological parameters, a shear and a photo- z bias parameter per tomographic bin, and two IA parameters. This gives a total of 16 free parameters: $\mathbf{p} = (\Omega_m, A_s, n_s, \Omega_b, h, \Omega_\nu h^2, m^{(i)}, \delta z^{(i)}, A_{IA}, \eta)$.

The results are shown in Figure 6.8, and summarised in Table 6.4. The best fitting cosmological parameters from the analyses on the two disjoint datasets (early- and late-type galaxies) are consistent with each other within the 1σ confidence bounds. Given that these are separate sets of galaxies with uncorrelated shape noise (albeit drawn from the same cosmological volume) one would not expect the mean values to coincide exactly. In the IA parameters, where we *prima facie* would expect to see a difference there is indeed a more prominent disagreement. The early-type sample favours a stronger amplitude, and a much weaker redshift evolution. Though the power of shear alone in coarse tomographic bins to constrain the latter is notably limited, we see a noticeable separation and different degeneracy with A_{IA} in the right-hand panel of Figure 6.8. The likelihood of this parameter is also notably non-Gaussian. We quote the mean of the distribution and the upper and lower 68% confidence bounds for reference in Table 6.4, but note that in this case the one-point statistics are a somewhat reductive way of assessing the constraint.

6.3.2 Redshift and Magnitude Dependence

In general IA are poorly understood. To marginalise their impact on shear we rely on empirical models, without solid physical basis or tight constraints on a dataset representative of the galaxy catalogues used in modern shear cosmology. An example of this is the dependence of the IA spectra on redshift and luminosity. If such dependence is modelled at all it is commonly with




Figure 6.9: Intrinsic alignment amplitude constrained from selected bin combinations of the Y1 catalogue. Each point shows the constraint using only two-point measurements that include a particular foreground bin. The dashed green and dotted purple lines show linear and power law fits to these points.

somewhat arbitrary power laws (see, for example, equation 8 from Hildebrandt et al. 2016, and Jee et al. 2016). It is also, however, true that our ability to explore this subject is limited by the data. Most importantly, we lack the high-quality spectroscopic information that have made the SDSS data, for example, such a rich source of IA-based investigations.

As a simple first exercise we consider the fiducial dataset (all galaxies), and look for internal tensions between redshift bins. To do this we rerun our earlier likelihood calculations, but using only a subset of redshift bin pairings. We do this four times, each time using only ξ_{\pm}^{ij} with a particular foreground bin. Cosmology and all nuisance parameters are fixed at the best-fitting values from the fiducial run on DES Y1, and A_{IA} is allowed to vary between ± 16 . The result is shown in Figure 6.9, as a function of redshift. The right-hand panel shows the 1D likelihood distributions for the four bins. Naturally, the upper bin provides the weakest constraint, since we are using only autocorrelation $i = j = 4$. To test the robustness of the upwards trend seen in this figure, we repeat the process using the best-fitting Planck cosmology rather than that favoured by the DES data, which are non-identical. The result is shown by the shaded curves in the right-hand panel.

A natural question that arises from this exercise is the validity of fixing nuisance parameters. Particularly in the case of photometric redshift error, the DES data constrains these weakly; one finds that the shape of the posteriors on these parameters are prior-dominated. It is also true that the IA amplitude can be strongly degenerate with redshift error. Both of these statements make

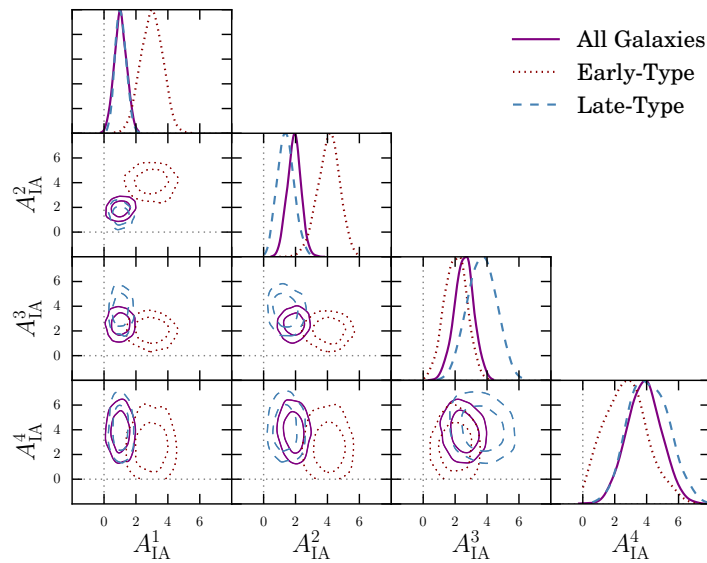


Figure 6.10: Intrinsic alignment amplitude constraints in four redshift bins using the DES Y1 METACALIBRATION catalogue. The purple solid lines show the full catalogue, with early- and late-type split samples shown by the dotted and dashed lines respectively. Note that these analyses were fixed to the best-fitting cosmology from the main DES analysis, but marginalise over photometric redshift error.

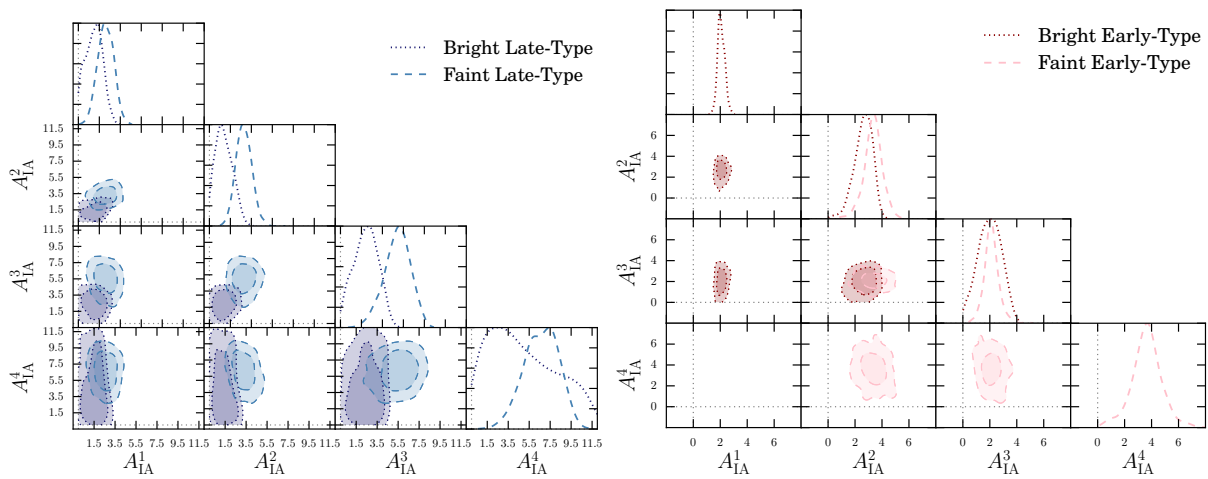


Figure 6.11: The same as Figure 6.10, but additionally split by magnitude. The two panels show constraints from late-type (left, blue) and early-type (right, red/pink) samples.

the fixed δz a major source of doubt in the redshift dependence seen in Figure 6.9. We thus carry out the following exercise. We rerun the fiducial chains, again with cosmology fixed, but allowing the four δz parameters to vary. The IA spectrum is predicted using the NLA model, but in each bin we rescale the IA contributions with an independent free amplitude $A_{\text{IA}}^{(i)}$. The result is shown in Figure 6.10 for the full sample (purple) and the split early and late-type samples (red and blue). Interestingly, though the early-type galaxies dominate the IA contribution at low redshift, the IA amplitude actually decline slightly in the upper two bins. Conversely, we see a monotonic increase in A_{IA} in the late-type sample. Although not conclusive, the tests in Section 6.3.3 indicate that this is not simply due to increased leakage of early-type galaxies into the late-type sample as we move to higher redshift.

Finally, we repeat the above likelihood calculation, but now additionally splitting the early- and late-type samples about the median observed r -band magnitude. Ideally we would use k -corrected rest frame luminosities, since the apparent brightness of a galaxy is sensitive to its intrinsic magnitude, but also its redshift within each broad tomographic bin. At the time of writing, such information is not available within DES. The results are shown in Figure 6.11.

Note that we do not show the results for the brightest early-type galaxies in the upper redshift bin, and the faintest in the lower bin. In these chains the likelihood is strongly bimodal, with divergent positive and negative A_{IA} . This is thought to be because of the poor sampling of galaxies in these bins. The $n(z)$ estimates are visibly very noisy, and the COSMOS-derived priors on redshift error are extremely wide. We do not consider the bimodal results here to be physically meaningful.

6.3.3 Colour Leakage & Redshift Error

Though the results of the tests described above show some differences between early- and late-type samples, there are a number of possible systematics entering the measurement. Firstly, we do not know the redshift distributions of our samples perfectly. If we wish only to constrain cosmological parameters and we can assume minimal IA contamination, then the mean redshift in each tomographic bin is the most salient form of redshift uncertainty we need to marginalise over (Amara & Réfrégier, 2007). This is not necessarily true if we wish to accurately disentangle the IA signal (including an implicit self-calibration analogous to the method described in the previous chapter). The IA spectra are much more sensitive to the shape of the redshift distribution. To test the impact of inaccuracies in the shape of the estimates of the redshift distributions from BPZ we rerun the two split chains (early and late-type galaxies). Instead of using the fiducial photometric

redshifts provided by Hoyle et al. (2017), however, we use the histograms of COSMOS redshifts shown by the dashed lines in the left-hand panel of Figure 6.5. Clearly these are noisy estimates, and will not perfectly represent the distribution of galaxies in DES Y1. They do, however, provide an alternative estimate with equivalent mean redshift in each bin; any difference in our results after switching to the COSMOS $n(z)$ will imply a potential uncertainty in our results. We show the constraints from this exercise as the filled contours in Figure 6.12.




Figure 6.12: Joint constraints on cosmology and IA parameters from DES Y1. The filled (pink and dark blue) contours are the same as in Figure 6.8, but using the reweighted histogram of 30-band COSMOS redshifts described in Section 6.2.4 in place of the photometric estimates from BPZ. The qualitative match to the earlier results demonstrate the sensitivity of our results to error in the shape of the redshift distributions. The open contours show the results using stringent cuts on T_{BPZ} , beyond those used in the fiducial analysis, and are presented as a test of galaxy type leakage, as described in the text.

The amplitude of the intrinsic alignment power spectra is also potentially strongly degenerate with another sort of redshift uncertainty: error in the mean redshift in each tomographic bin. To test this we widen our redshift priors to account for (a) correlations in error between bins and (b) disagreement with alternative constraints on the $n(z)$ from another group within DES, derived from galaxy clustering measurements. This leaves the Gaussian priors on δz with widths between $\Delta\delta z \sim 0.02$ and $\Delta\delta z \sim 0.05$. We rerun our chains with these widened priors. This exercise moderately weakens our constraints, but does not significantly shift the means of the posterior distributions.

Finally we assess the level of error in our measurement due to leakage between our two galaxy type bins. We are dividing galaxies into early-type and late-type samples based on a single measured parameter T_{BPZ} . Due to the limited quality of the DES photometry, however, T_{BPZ} will

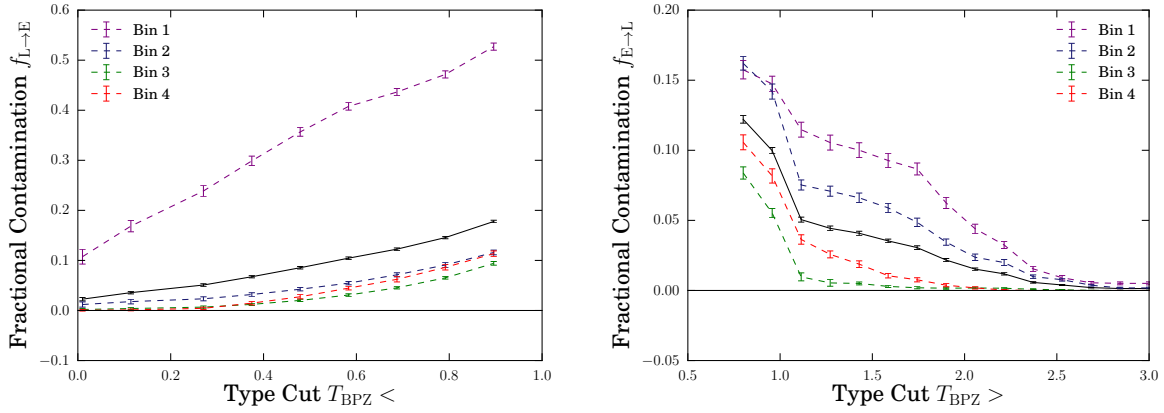


Figure 6.13: Galaxy type leakage in our early-type (left) and late-type (right) samples. We quantify the leakage as the total lensing weight of galaxies included in the sample as defined by the DES T_{BPZ} measurements, but which would have been removed by the equivalent cut on the “true” HST T_{BPZ} values. The four coloured (dashed) lines show the results in the four tomographic bins used in this chapter, and shown in Figure 6.5. The solid black line shows the type leakage for the full sample, without tomographic binning.

be noisy and potentially biased. To quantify the level of misclassification we use a set of photo- z measurements obtained by running BPZ on the reweighted COSMOS galaxies described in Section 6.2.4, using the original HST fluxes as input. Since the noise in the HST fluxes is much smaller than that in DES, we take this as a reference “true” measurement of T_{BPZ} for each galaxy. That is, the value it *should* have been assigned, if DES had access to photometry at the quality of HST. We then define our sample (early or late) using a cut on the measured DES T_{BPZ} , and in each tomographic bin we calculate the number of galaxies which are erroneously included (i.e. the objects that should have been removed by the type cut, according to the HST T_{BPZ}). We show the total lensing weight of these misclassified galaxies as a function of the cut used to define the sample in Figure 6.13. The two panels show early-type and late-type galaxies as a function of the maximum and minimum T_{BPZ} cut.

Although in most cases, the misclassifications account for less than 10% of the lensing weight, in the lowest redshift bin, we find a potentially much larger cross-type contamination. To assess the level at which this enters our results, we define new “refined” samples with more stringent cuts on T_{BPZ} . For the refined early-type sample we shift the upper T_{BPZ} cut to $T_{BPZ} < 0.5$, and for the new late sample we impose $T_{BPZ} > 1.5$. We rerun our joint constraints using these adjusted samples, which are shown by the unfilled (red dashed and blue dot-dashed) contours in Figure 6.12. Note that for each of these new samples we recompute the $n(z)$, redshift priors, two-point functions and




Figure 6.14: Constraints on cosmology and IA parameters from the unsplit Y1 METACALIBRATION shape catalogue. The dark purple contours show the result using our fiducial IA model, which assumes the NLA IA power spectra with a free amplitude A_{IA} and a redshift power law index η . The lighter dashed constraints assume the same, but with η fixed at zero.

covariances, and then run independent Monte Carlo chains with the 16 free parameters used in the fiducial calculation.

6.3.4 Intrinsic Alignment Models

Though previous numerical studies have found significant biases from failing to account for IA in shear-based analyses (Kirk et al., 2012; Krause et al., 2015), the exact impact depends on the nature of the true signal, and how inaccurately one models it. We have found using numerical forecasts, that modelling a mixed population with different IA signals does not necessarily bias one’s cosmological constraints. We construct an analytic Y1-like shear datavector using equations 6.9 and 6.10, with the measured red-fractions. We consider two cases, where the blue galaxies have (a) no alignments, or (b) GI and II power spectra given by equation 6.7 with non-zero amplitude, which is different from that of the red galaxies. In both these cases it is possible to model the total IA contamination using a single model and recover the input cosmology without significant bias. The inferred A_{IA} becomes an effective amplitude, approximately equal to the weighted average of the red and blue IA amplitudes. We cannot, however, guarantee this will be the case if the blue galaxies have IA power spectra that are both non-zero and significantly different from the NLA prediction.

To explore this, we rerun our fiducial cosmology chain on the full dataset (without type splits), but now without marginalising over the redshift power law. As has commonly be assumed in

previous studies, we assume $\eta = 0$.¹ The result of this exercise is shown in Figure 6.14. Switching between these models does not induce significant shift in the favoured cosmology, although we report a marginally lower value of A_{IA} .

6.4 Discussion & Comparison with Previous Results

This work sits alongside a number of simultaneous constraints analyses carried out within the large lensing surveys.

Though useful for understanding the nature of the IA contaminant in cosmic shear measurements, Heymans et al. (2013) are the only previous authors to attempt a direct red/blue split of the type we present here. That work was based on the the CFHTLenS dataset, which is has a slightly deeper magnitude limit than DES, and used a finer binning in redshift (six bins to our four), albeit over only 154 square degrees. They defined early/late-type galaxies in the same way as in this analysis, and about the same boundary². A major result of their analysis was that, when combined with WMAP, they found a mildly negative A_{IA} (although statistically consistent with zero). Their split samples showed a clear separation, with a best-fitting amplitude of ~ 5 for early-type galaxies, and a measurement consistent with no alignments in late-types. They speculated in that work that the negative result could be the result of simultaneously fitting the GI and II spectra with a single amplitude. Subsequent discussion has gradually converged on the idea that it was more likely an artefact of redshift error, an idea reinforced by the findings of Choi et al. (2016). A comprehensive re-analysis of CFHTLenS by Joudaki et al. (2017) found a stronger preference towards negative values ($A_{IA} = -3.6 \pm 1.6$). Later analysis of the KiDS dataset by Hildebrandt et al. (2016), found a value of around unity, marginalising also over a redshift index η . Interestingly, they tested their results using a series of alternative photo- z methods, and found certain cases that could push A_{IA} below zero. Most recently, van Uitert et al. (2017) report positive A_{IA} , again at approximately 1.0 using a similar but independent set of two-point measurements in Fourier space. They found similar results using both cosmic shear alone, and when combining shear, clustering and galaxy-galaxy lensing.

It is clear that these results are somewhat variable. The IA contamination is dependent on the shape measurement technique and on the galaxy sample in question. Moreover, as pointed out by van Uitert et al. (2017), the IA parameters can become less physical “nuisance” parameters when

¹Although we have found already that the DES Y1 data favours $\eta \neq 0$ at $\sim 1\sigma$

²Note that the quoted division in their analysis is $T_{BPZ} = 2.0$, whereas ours is at 1.0. This is purely a difference in indexing notation, and we can confirm that the two are equivalent.

implemented in a high dimensional analysis, as they can very easily absorb other systematics such as photometric redshift error. That said, it is still useful (and, indeed, possible) to understand the IA signal in its pure astrophysical sense. If we do not attempt to do this, it is difficult to answer basic questions about whether our nuisance parameterisation is sufficiently flexible, and whether our assumptions are in fact resulting in a cosmological bias. It is also true that imposing informative priors on the IA contamination could potentially curtail what could become a major limiting source of systematic uncertainty.

Perhaps the most directly comparable previous result is that from the DES Science Verification (DES SV) data. Though narrower (comparable to CFHTLenS) and deeper, the measurements presented in Becker et al. (2016) and analysed in Dark Energy Survey Collaboration (2016) share many features of the Y1 analysis presented in this chapter. The basic shape measurement code is essentially the same (NGMIX, although the calibration method is new), and our current dataset is a superset of the area presented there. Interestingly, their fiducial analysis, which marginalised over η , found a bimodal A_{IA} posterior with peaks at $\sim \pm 1$ to 2. This was seen to disappear when using only the basic NLA model with no redshift dependence. Unlike in our case, where the cosmological result is relatively robust to inclusion of η , they found that allowing this parameter to vary significantly degraded their constraint on S_8 . This was largely a result of the extended bimodal feature in the $S_8 - A_{IA}$ plane, which is now gone (compare their Figure 8 with Figure 6.14 above).

It has been suggested elsewhere that removing red galaxies could be an effective and relatively inexpensive way to null the IA contamination in future surveys. We confirm the result of Krause et al. (2015) that the quality of the cosmic shear result is relatively robust to the cut on galaxy type (since it leaves around 80% of galaxies, and preferentially removes objects at lower redshift, which provide relatively little cosmological information). Unfortunately our result suggests that the IA contamination in these remaining galaxies is neither negligible, nor unchanging with redshift. Indeed, our preliminary results suggest a worsening of the blue IA contamination in the upper redshift bins. Neither do our findings suggest that this is purely the result of type leakage. Though this clearly depends on how we choose to divide our galaxy set, the overall implication is that it would be inadvisable to rely on selecting by galaxy type as a means to null the IA signal in cosmological analyses.

Finally, it is worth mentioning here that the ability to meaningfully constrain intrinsic alignments is strongly dependent on (accurate) redshift tomography. We have found that widening our photo- z priors to the extent that they are essentially uninformative can easily lead to unphysical bimodalities in the IA likelihood surfaces. Our slightly finer binning (4 bins to the 3 in SV) is likely to

have contributed a small improvement in the constraints presented here. As Dark Energy Survey Collaboration (2016) have shown (see also Jee et al. 2013 and Kilbinger et al. 2013) the benefit is dramatic from just a small number of redshift bins relative to the non-tomographic case.

6.5 Conclusion

In this chapter we have presented a new set of joint constraints on cosmology and intrinsic alignments, using the largest late-time lensing dataset to date. We have used a shear catalogue of 35 million galaxies over 1500 square degrees with median redshift 0.59, and corrected for selection effects in both the shape measurements and redshift distributions. We have found that splitting by galaxy type, as defined by a best fitting SED template T_{BPZ} , yields populations with significantly different IA properties (both amplitude and redshift dependence). Using a more flexible model with a free amplitude per redshift bin, we have found a monotonic increase in IA strength with redshift. Though the early-type signal is relatively stable between bins, our late-type (blue) sample exhibits a more significant dependence. It is difficult to interpret this in terms of competing II and GI contributions, and this topic is clearly of interest for future analyses.

We have tested our results for sensitivity to the details of the photometric redshift distribution, and found that they are robust to reasonable changes in the priors on redshift error and the shape of the distributions. Though we have found potentially significant galaxy type leakage between samples in the lowest redshift bin, rerunning our analysis with more stringent cuts on T_{BPZ} does provide evidence for significant bias in our findings. A future step in this analysis will repeat the likelihood MCMC calculation using the full parameter space, but excluding the lowest bin where the contamination is much stronger.

Finally, we have tested the impact of switching between IA models, in a limited case of NLA with and without a redshift power law. Unlike in the analysis of (Dark Energy Survey Collaboration, 2016) on the DES SV lensing dataset, we have found no dramatic degradation in the quality of our shear constraints due to the inclusion of the redshift parameter.

As stated in the introduction to this chapter, the analysis presented here is ongoing. We intend to include a number of extensions in the near future. First, we will test the impact of including non-Gaussian analytic halo model covariance matrices. Limited computational resources have prevented us from doing this for the various sub-samples presented here. Second, we will repeat the calculations in Section 6.3.4 using a more physically motivated IA treatment for blue galaxies, that includes the different mechanisms that are thought to operate (e.g. tidal torquing, see Blazek

et al. 2015), to look for signs of bias. It will also be interesting to implement an IA treatment akin to that in the previous chapter, where we split the GI and II amplitudes and fit them separately. Although we found no shift in cosmology when we removed the redshift dependence from our IA model, the test we presented does not extend beyond the NLA model. It is conceivable that the GI and II spectra for blue galaxies could both have a significant non-zero amplitude and a differ in shape from the NLA prediction. The level of possible cosmological bias due inadequate modelling of such a combination is not currently known. Third, we will explore the alternatives to T_{BPZ} set out at the start of this chapter. As we said, there are a number of alternatives quantities one could use to split the DES catalogue that could be accommodated within our framework for selection effects. It is conceivable that, for example, a colour-magnitude-based cut, or T_{BPZ} in combination with such a cut could provide a more effective means for isolating a sample with small or null intrinsic alignments. This information could clearly have practical implications for future studies. Finally, we will consider combining our results with other measurements derived from the DES data. The combination of galaxy-galaxy lensing and galaxy clustering, for example, could potentially significantly enhance our ability to constrain intrinsic alignments.

As datasets get deeper and other systematic uncertainties are gradually reduced, IAs could very easily become a limiting factor in future cosmological analysis. Whereas baryonic effects can be mitigated by scale cuts, intrinsic alignments operate on all scales where shear is also significant. Improving our understanding of this systematic will thus be crucial if we are to fully exploit the cosmological information of the datasets that will shortly become available to the lensing community.

Chapter 7

Conclusions

This chapter provides a synoptic overview of the ideas presented in this thesis, and seeks to set them within the context of the wider field. We start with a brief overview of the prospects for the immediate future, and the challenges we must address. The following paragraphs consider each of the analyses presented in this thesis. We discuss the scope and the limitations of the various strands of work. The final intention, beyond providing a summary, is to explore the possible avenues for future investigation and identify the most promising extensions in each area.

7.1 Remarks About the Current Work

The aim of this manuscript was to set out a practical example of how a large late-time galaxy lensing survey can be used to place constraints on cosmology. We do not claim our implementation to be perfect, nor that it does not depend on a large number of collaborators with expertise far beyond our own. Rather we sought to present a demonstration of principle, detailing how the task was approached within DES. This thesis aimed to demonstrate that cosmic shear measurement is subject to a range of biases of comparable level, none of which can be safely ignored. But, equally important, we showed that it *is possible* to successfully address these biases to place accurate constraints on cosmology, even at the statistical precision allowed by the largest lensing datasets that exist today.

In Chapter 3 we outlined the process used to construct shape catalogues and, ultimately, the weak lensing correlation functions. The text focused on one of the two science-ready catalogues in Year One of the Dark Energy Survey. Though they take very different approaches to the shear measurement problem, the two catalogues (IM3SHAPE and METACALIBRATION) are derived from

the same set of galaxies, and give consistent results in all analyses that have used them so far (Prat et al., 2017; Troxel et al., 2017; Chang et al., 2017; DES Collaboration et al., 2017). The first, and in some regards the simplest, step is conducting fits to the individual galaxies. For this we used a two-stage fit, which assigns each galaxy a best-fitting Sérsic light profile with a radial index of either $n = 0$ (exponential disc-like) or $n = 4$ (de Vaucouleurs bulge-like). Fits such as this are near-ubiquitous in the field of lensing cosmology. Though inevitably a simplified representation of real galaxy morphologies, it has been shown elsewhere that the bias that results is predictable and almost entirely dependent on the size of the galaxy and the signal-to-noise (S/N) of the measurement (Voigt & Bridle, 2010; Kacprzak et al., 2012). We then introduced the HOOPOE simulations, a set of real-sky mock images tailored to the DES Y1 dataset. We described how HOOPOE was used to constrain the level of bias in the DES Y1 IM3SHAPE catalogue, and how per-galaxy shear corrections were derived from ensemble measurements on the simulation. It is important to note that this analysis is one of the most advanced attempts at shear calibration to date, in the detail included in the simulations, the level of redundancy in the methodology, and the rigour of the process which led to the suggested prior on residual bias. Alongside the (equally rigorous) treatment of the METACALIBRATION catalogue (also detailed in Zuntz et al. 2017) and its nearest contemporary, based on weak lensing measurements from the Kilo Degree Survey (Fenech Conti et al., 2016), it will set a benchmark for treatment of shear estimation bias in future cosmic shear datasets. This analysis discovered a residual multiplicative bias of around $m = 0.01$ after calibration and realistic tomographic binning of the simulated galaxies. Our findings suggest a form of redshift selection bias, due to the evolution in galaxy morphology with redshift. Though not the dominant form of systematic uncertainty in the Y1 IM3SHAPE calibration, this could represent a fundamental challenge for future parametrised simulation-based calibrations, considering the stringent requirements of the upcoming Stage IV surveys. The history of shear estimation, however, has been a story of continuous and conscious improvement and there are reasons to be optimistic about continued feasibility of the simulation-based method.

The focus in Chapter 4 then turned to a more specific form of measurement bias: the impact of image plane neighbours on shear cosmology. This is a question that, until recently, has commonly been ignored in shear-based cosmology analyses. Using a simplified toy model of two neighbouring galaxies we showed how neighbours can leave a residual imprint on ensemble shear estimates, even in the absence of a preferred direction to their spatial distribution. We found that the resulting shear bias could be either positive or negative, depending on the configuration of the neighbours. The investigation was then extended to use the HOOPOE simulations, which accurately represent the level of blending in DES Y1, at least at first order. Comparing direct bias measurements with

a second suite of neighbour-free simulations under different selection functions, we identified four mechanisms by which neighbours influence shear measurements. In combination, and measured in tomographic bins, we found the net impact could lead to artificially low shear estimates at a level of 3% to 9%. Furthermore, we found that the inclusion of galaxies that are too faint to be detected, but which nonetheless appear in real survey images, impacts the derived shear calibration at the level of approximately one percent.

The final part of the analysis in Chapter 4 then sought to propagate the effect of neighbours into cosmological parameter space. A number of earlier studies on the impact of various effects on the constraining power and accuracy of cosmic shear have relied on an assumption of Gaussian parameter likelihoods. Though the Fisher formalism is convenient for its computational speed, it has been argued that full simulated likelihood analyses are necessary to capture the complexities of high-dimensional parameter analyses in modern cosmology (Joachimi et al., 2011; Krause et al., 2015, 2017). We thus followed the latter approach for our analyses. With the cosmological framework set up for analysing the real DES Y1 data we constructed analytic shear datavectors with different forms of bias. The aim was to model the level of contamination that our study of the HOOPOE simulations suggested would be present if we had failed to include neighbours in the DES Y1 IM3SHAPE calibration. With the covariances calculated for the Y1 shear cosmology analysis, we found a net bias of $\sim 1.5\sigma$ towards low values of the best-constrained cosmological parameter combination S_8 .

The position information in the HOOPOE simulations also provides a way to test for additional forms of bias that would be present in the data, but will not be captured in parameterised bias fits. Our analysis identified a potential effect due to spatial correlations in the neighbour-induced biases. We developed an analytic method for propagating measurements of these correlations into a real space two point shear data. The magnitude of these effects, however, was seen to be small: even at significantly less stringent scale cuts than planned for DES Y1, we found no discernable bias in cosmology beyond the current level of statistical uncertainty.

There has been significant attention in recent years towards the complementary measurements possible within a photometric galaxy like DES. Particularly useful are the analogues to pure cosmic shear correlations using galaxy density. Cosmology with galaxy clustering alone is limited by gaps in our understanding of galaxy bias. The combination of lensing, clustering and galaxy-galaxy lensing, however offers significant promise for internally calibrating common observable systematics in both sets of measurements. Put simply, galaxy clustering constrains galaxy bias, which frees galaxy-galaxy lensing to constrain redshift error and intrinsic alignments, which in turn frees the shear-shear correlations to constrain cosmology. In Chapter 5 we explored this con-

cept using analytic data designed to mimic DES Y1. We identified a valuable complementarity in the parameter degeneracies in cosmic shear and the combination of galaxy-galaxy lensing and clustering. Provided sufficiently accurate redshift information in the sample used for clustering measurements, we found that one could effectively self-calibrate redshift inaccuracies in the lensing galaxies. Without significant degradation in the quality of the constraints, this allows one to accurately recover the correct cosmology even in the presence of moderate redshift errors. In a similar analysis using cosmic shear alone, we showed that internal calibration of this sort is not possible. Such techniques are highly valuable, and emerge naturally from a combined analysis that includes these datavectors simultaneously. Clearly complementary to external calibration of systematics like redshift error and measurement bias using simulations or additional data, they are an important reason to include such additional data in future cosmology analyses.

In the final chapter of this thesis, we turned to another major systematic in lensing cosmology. The term Intrinsic Alignments (IAs) describes a set of complex astrophysical processes by which the pre-lensing shapes of galaxies can become spatially aligned. Commonly recognised as a lensing systematic, they have already been shown to induce significant offsets in cosmology if ignored (Kirk et al., 2012). Whereas the earlier chapters made use of simulations and analytic two point data, this analysis used the real DES Y1 data. The chapter detailed the route from shear catalogues to eventually extracting cosmology constraints from the lensing data, which we presented in blinded form. It provided a practical demonstration of how one can correct shear calibration selection and redshift biases in DES Y1, using corrections derived from the METACALIBRATION algorithm and a reweighted sample of galaxies from the HST COSMOS survey. There is physical motivation for the idea that red and blue galaxies have different levels of intrinsic alignment contamination. By fitting a common IA model separately to early-type and late-type sub-samples of the DES Y1 data, we detected a statistically significant difference in IA strength. Our analysis suggested a non-zero IA amplitude in late-type galaxies, though at the time of writing we have not conclusively ruled out cross-sample contamination or photometric redshift error induced by the split as the cause of this signal. Whatever the cause, our findings cast doubt on the viability of galaxy type cuts as a means to cleanly eliminate IAs as a contaminant in lensing cosmology. Finally, we repeated our analysis using a slightly simpler model, which does not allow for redshift evolution in the IA amplitude. There was no resulting shift in the derived cosmology, and unlike in the SV analysis no dramatic change in the strength of the constraint.

7.2 Prospects for the Future

In these final paragraphs we will consider the future and how the findings presented in this thesis might be marshalled to the benefit of the next generation of lensing cosmology analyses.

7.2.1 Shear Estimation

At the time of writing DES has entered its fifth year of data collection. Though a raft of science analyses of the DES Y1 dataset are still awaiting publication, a complete set of Y3 images have been taken, have undergone photometric processing and are available to the collaboration. As of the start of 2017 METACALIBRATION has been successfully run on Y3, and IM3SHAPE fitting is ongoing. There is clearly an imperative to digest the lessons of Y1 quickly. In the Y1 data METACALIBRATION appears to have yielded high-accuracy shear estimates. A deblending technique known as multi-object fitting, which was trialled in Y1, also looks promising, both as an alternative route for neighbour mitigation and as more robust method for extracting single-galaxy fluxes.

A question remains, however, as to how we approach the challenge in the case of IM3SHAPE in Y3 and beyond. These challenges are important for DES, but they are also generic to the conventional simulation-derived calibration techniques implemented by virtually all previous shear cosmology analyses. One significant lesson from the analyses of DES SV and Y1 is that redundancy is important, and having two alternative shear catalogues has been highly useful in verifying the robustness in each stage of analysis. In many ways the lessons are plain. One significant form of calibration uncertainty came from a choice made for computational reasons in constructing the HOOPOE simulations (i.e. the decision to retain the original cutout size for each galaxy). Though too late for the Y1 catalogues, the issue was resolved at the code level in late 2016. Likewise, the exercise outlined in Appendix B of this thesis identified problematic HST profiles, which can be readily omitted from future simulations. These two improvements alone would have reduced our non-tomographic prior to $\sigma_m = 0.02$. Though sufficient for the current surveys, this level of systematic calibration uncertainty will almost certainly become cause for concern in future. Our assumptions about uncertainty in the neighbour treatment were conservative. The development of a multi-object fitting algorithm for IM3SHAPE may prove costly in terms of already non-trivial computing time, but would be scientifically very useful in this regard. It is troubling that the persistent redshift-dependent bias of around 1% we found in Y1, if understood correctly, could be an inherent limitation in the most commonly implemented methodology for shear calibration. The

COSMOS dataset used in our simulations and the low-noise deep supernova fields within DES offer the necessary resources for constraining this effect. For the duration of DES, the 1% overhead will fall within tolerance, and not be the limiting factor in our cosmology analyses. Understanding the effect, however, will be important if we are to achieve the accuracy in simulation-based calibration methods required for the lensing datasets of LSST and Euclid.

7.2.2 The Intrinsic Alignment of Galaxies

Though evidence does not point towards inadequacies in our modelling of intrinsic alignments currently being a limiting factor in cosmic shear analyses this is not guaranteed to hold true indefinitely. Fundamentally, however, we can and should be more ambitious in our treatment of this systematic. In essence an overlapping large-scale astrophysical signal, IA correlations do not present an insurmountable challenge in terms of detailed understanding. On a practical level, devising appropriate parameters and imposing informative priors on them will place IAs on a footing with shear calibration as an understood and controlled systematic.

We have made a small step along the path to that understanding in the final chapter of this thesis. Our preliminary results suggest a statistically significant, if relatively small, IA signal in late-type galaxies. This is not inherently a concern if it can be adequately modelled, although the most commonly used model for IAs (the nonlinear alignment model) was built around observations of a quite different population of late-type red galaxies. The degeneracy with photo- z error is also a question which must be rigorously tested. Before we can claim to fully understand these results, there are essential follow-up tests that should (and will) be carried out. Repeating our analyses with explicit separation of GI and II amplitudes in the IA model will be one way to deconstruct the mechanisms behind our results. For eliminating the impact of possible photo- z error (as we have seen) and raw constraining power, combining with galaxy-galaxy lensing and galaxy clustering correlations would also be both very feasible and scientifically useful.

A more ambitious and slightly longer term project with a similar aim will soon begin now the DES Y1 analyses are being laid to rest. A cross-collaboration project, of which the author is a member, between DES and eBOSS (SDSS-IV) will allow us to undertake a data-driven study aiming for a direct detection of intrinsic alignments in blue galaxies. There are current upper limits on fainter blue populations (Mandelbaum et al., 2011), though these are limited to relatively low redshifts, where cosmic shear is weak. The DES-eBOSS project will extend significantly deeper, utilising emission line galaxies (ELGs) for their high-quality redshifts. Such measurements are unprecedented, and would represent a step forwards in our modelling of IA contamination in shear

cosmology, since this sample is far more typical of the galaxies that dominate in cosmic shear measurements. It would also be a valuable opportunity to explore putative redshift and luminosity dependence in the IA contamination.

7.2.3 Self-Calibration & Multi-Probe Cosmology

Although in some ways simplified, the analysis in Chapter 5 lays the ground for future studies. Indeed, the implicit self-calibration described has already been demonstrated in KiDS (van Uitert et al., 2017), and DES (DES Collaboration et al., 2017), in the presence of all the complexities of real data. The ultimate, if still somewhat distant aim is a fully integrated late-time multiprobe analysis incorporating supernova, baryon acoustic oscillations (BAO), redshift space distortions and optical cluster counts. The benefits in terms of constraining power and internal robustness against systematics would be enormous. Nonetheless, there is a significant challenge ahead, which will start with an analytic forecasting framework of the sort described in this thesis. The methodology of simulated likelihood analyses will be vital in understanding our data, its parameter degeneracies, and which of a host of new systematics are of concern for our eventual inferences about the properties of the Cosmos.

Before ending, we should finally turn our thoughts towards the higher-level future of lensing in the coming years. As we have sought to show in this thesis, weak lensing is more than just another “potentially powerful” probe. It can be, and indeed has been, used to place truly competitive constraints on the properties of the late-time Universe with existing tools. True, it has its subtleties, but one would be hard pushed to name a cosmological probe that does not, given the precision cosmologists have come to expect. Nor have we yet reached the statistical floor imposed by cosmic variance; wider, deeper datasets will bring significant gains in precision. There are inherent systematics, but all are understood relatively well at the theory level. A glance through the history of lensing provides ample reason to be optimistic our ability, as a community, to develop novel ways to mitigate such effects.

The natural question, then, for lensing cosmology might be *is this it?* From one perspective, if years of collaborative effort amount to nothing more than an incremental tightening of some already Gaussian parameter contours about the standard values, perhaps we would be well advised to apply our efforts elsewhere. This is a view that has been voiced by professional cosmologists. No doubt variants of this question have been echoed by funding agencies across the world. There is, however, a more optimistic and perhaps more nuanced case to be made here. The mode of scientific practice evolves, as our understanding of the world changes. If the Twentieth Century

was a time of model building, we have perhaps reached a time of model *breaking*. As a community we have spent decades developing a framework for thinking about the Universe, and building a picture of what we *should* see. In Λ CDM we have a remarkably successful predictive theory and the means to test it, but it is a finite one; this is arguably how science works: a theory is conceived, disseminated and finally tested until we find a regime where it is insufficient. Almost certainly, our picture of the Universe is imperfect. It is the cosmologist's job now to find its edges.

Bibliography

- Abdalla F. B., Amara A., Capak P., Cypriano E. S., Lahav O., Rhodes J., 2008, MNRAS, 387, 969
- Aihara H., et al., 2017, preprint, ([arXiv:1704.05858](https://arxiv.org/abs/1704.05858))
- Albrecht A., Bernstein G., Cahn R., Freedman W. L., Hewitt J., Hu W., Huth J., Kamionkowski M., Kolb E. W., Knox L., Mather J. C., Staggs S., Suntzeff N. B., 2006, ArXiv Astrophysics e-prints,
- Alimi J.-M., Bouillot V., Rasera Y., Reverdy V., Corasaniti P.-S., Balmes I., Requena S., Delaruelle X., Richet J.-N., 2012, preprint, ([arXiv:1206.2838](https://arxiv.org/abs/1206.2838))
- Amara A., Réfrégier A., 2007, MNRAS, 381, 1018
- Angulo R. E., Pontzen A., 2016, MNRAS, 462, L1
- Aubourg É., et al., 2015, PRD, 92, 123516
- Bacon D. J., Refregier A. R., Ellis R. S., 2000, MNRAS, 318, 625
- Baldauf T., Smith R. E., Seljak U., Mandelbaum R., 2010, PRD, 81, 063531
- Bardeen J. M., Bond J. R., Kaiser N., Szalay A. S., 1986, ApJ, 304, 15
- Bartelmann M., 2010, Classical and Quantum Gravity, 27, 233001
- Bartelmann M., Schneider P., 2001, Physics Reports, 340, 291
- Becker M. R., et al., 2015, preprint, ([arXiv:1507.05598](https://arxiv.org/abs/1507.05598))
- Becker M. R., et al., 2016, PRD, 94, 022002
- Benjamin J., et al., 2013, MNRAS, 431, 1547
- Bennett et al. 1996, Astrophys. J. Lett., 464
- Bernstein G. M., 2009, ApJ, 695, 652
- Bernstein G. M., Armstrong R., 2014, MNRAS, 438, 1880

- Bernstein G. M., Jarvis M., 2002, *AJ*, 123, 583
- Bernstein J. P., et al., 2012, *ApJ*, 753, 152
- Bertin E., 2011, in Evans I. N., Accomazzi A., Mink D. J., Rots A. H., eds, *Astronomical Society of the Pacific Conference Series Vol. 442, Astronomical Data Analysis Software and Systems XX*. p. 435
- Bertin E., Fouqué P., 2010, *SkyMaker: Astronomical Image Simulations Made Easy*, *Astrophysics Source Code Library* (ascl:1010.066)
- Betoule M., et al., 2014, *A&A*, 568, A22
- Blake C., Collister A., Bridle S., Lahav O., 2007, *MNRAS*, 374, 1527
- Blazek J., Vlah Z., Seljak U., 2015, *JCAP*, 8, 015
- Bond J. R., Efstathiou G., 1987, *MNRAS*, 226, 655
- Bonnett C., et al., 2015, PRD accepted, ArXiv e-prints,
- Bonvin V., et al., 2017, *MNRAS*, 465, 4914
- Bordoloi R., Lilly S. J., Amara A., Oesch P. A., Bardelli S., Zucca E., Vergani D., Nagao T., Murayama T., Shioya Y., Taniguchi Y., 2012, *MNRAS*, 421, 1671
- Bridle S., King L., 2007, *New Journal of Physics*, 9, 444
- Bridle S. L., Kneib J.-P., Bardeau S., Gull S. F., 2002, in Natarajan P., ed., *The Shapes of Galaxies and their Dark Halos*. pp 38–46, doi:10.1142/9789812778017_0006
- Bridle S., et al., 2008, preprint, (arXiv:0802.1214)
- Bridle S., et al., 2010, *MNRAS*, 405, 2044
- Britton D., Cass A. J., Clarke P. E. L., Coles J., Colling D. J., Doyle A. T., Geddes N. I., Gordon J. C., Jones R. W. L., Kelsey D. P., Lloyd S. L., Middleton R. P., Patrick G. N., Sansum R. A., Pearce S. E., 2009, *Philosophical Transactions of the Royal Society of London Series A*, 367, 2447
- Brouwer M. M., et al., 2017, *MNRAS*, 466, 2547
- Brown M. L., Taylor A. N., Hambly N. C., Dye S., 2002, *MNRAS*, 333, 501
- Brown M., Bacon D., Camera S., Harrison I., Joachimi B., Metcalf R. B., Pourtsidou A., Takahashi K., Zuntz J., Abdalla F. B., Bridle S., Jarvis M., Kitching T., Miller L., Patel P., 2015, *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, p. 23

- Cai Y.-C., Padilla N., Li B., 2015, MNRAS, 451, 1036
- Carlstrom J. E., Holder G. P., Reese E. D., 2002, ARAA, 40, 643
- Carr B., Kühnel F., Sandstad M., 2016, PRD, 94, 083504
- Carroll S. M., Press W. H., Turner E. L., 1992, ARAA, 30, 499
- Castander F. J., 2008, in *Ground-based and Airborne Instrumentation for Astronomy II*. p. 70145C, doi:10.1117/12.789497
- Cavuoti S., Tortora C., Brescia M., Longo G., Radovich M., Napolitano N. R., Amaro V., Vellucci C., La Barbera F., Getman F., Grado A., 2017, MNRAS, 466, 2039
- Chang C., et al., 2012, MNRAS, 427, 2572
- Chang C., Jarvis M., Jain B., Kahn S. M., Kirkby D., Connolly A., Krughoff S., Peng E.-H., Peterson J. R., 2013, MNRAS, 434, 2121
- Chang C., et al., 2015, preprint, (arXiv:1505.01871)
- Chang C., et al., 2017, preprint, (arXiv:1708.01535)
- Chevallier M., Polarski D., 2001, *International Journal of Modern Physics D*, 10, 213
- Choi A., Heymans C., Blake C., Hildebrandt H., Duncan C. A. J., Erben T., Nakajima R., Van Waerbeke L., Viola M., 2015, preprint, (arXiv:1512.03626)
- Choi A., Heymans C., Blake C., Hildebrandt H., Duncan C. A. J., Erben T., Nakajima R., Van Waerbeke L., Viola M., 2016, MNRAS, 463, 3737
- Clampitt J., et al., 2016, preprint, (arXiv:1603.05790)
- Clowe D., Randall S. W., Markevitch M., 2007, *Nuclear Physics B Proceedings Supplements*, 173, 28
- Collett T. E., 2015, ApJ, 811, 20
- Crill B. P., et al., 2008, in *Space Telescopes and Instrumentation 2008: Optical, Infrared, and Millimeter*. p. 70102P (arXiv:0807.1548), doi:10.1117/12.787446
- Cunha C. E., Huterer D., Lin H., Busha M. T., Wechsler R. H., 2014, MNRAS, 444, 129
- DES Collaboration et al., 2017, preprint, (arXiv:1708.01530)
- Dark Energy Survey Collaboration 2016, PRD, 94, 022001
- De Vicente J., Sánchez E., Sevilla-Noarbe I., 2016, MNRAS, 459, 3078

- Dekel A., Lahav O., 1999, *ApJ*, 520, 24
- Dewdney P. E., 2015, Technical Reports SKA-TEL-SKO-DD-001, Revision 1
- Di Valentino E., Melchiorri A., Linder E. V., Silk J., 2017, preprint, ([arXiv:1704.00762](https://arxiv.org/abs/1704.00762))
- Dodelson S., 2003, *Modern Cosmology*. Academic Press, Academic Press, <https://books.google.co.uk/books?id=3oPRxdXJexcC>
- Drlica-Wagner A., et al., 2017, submitted to *Astrophys. J. Suppl. Ser.*
- Duncan C. A. J., Joachimi B., Heavens A. F., Heymans C., Hildebrandt H., 2014, *MNRAS*, 437, 2471
- Dyson F. W., Eddington A. S., Davidson C., 1920, *Philosophical Transactions of the Royal Society of London Series A*, 220, 291
- Efstathiou G., 2014, *MNRAS*, 440, 1138
- Eifler T., Krause E., Dodelson S., Zentner A., Hearin A., Gnedin N., 2014, preprint, ([arXiv:1405.7423](https://arxiv.org/abs/1405.7423))
- Eisenstein D. J., Hu W., 1998, *ApJ*, 496, 605
- Eisenstein D. J., Hu W., 1999, *ApJ*, 511, 5
- Faltenbacher A., Li C., White S. D. M., Jing Y.-P., Shu-DeMao Wang J., 2009, *Research in Astronomy and Astrophysics*, 9, 41
- Fenech Conti I., Herbonnet R., Hoekstra H., Merten J., Miller L., Viola M., 2016, preprint, ([arXiv:1606.05337](https://arxiv.org/abs/1606.05337))
- Feroz F., Hobson M. P., Bridges M., 2009, *MNRAS*, 398, 1601
- Flaugher B., et al., 2015, preprint, ([arXiv:1504.02900](https://arxiv.org/abs/1504.02900))
- Freedman W. L., Carnegie Supernova Project 2005, in Wolff S. C., Lauer T. R., eds, *Astronomical Society of the Pacific Conference Series Vol. 339, Observing Dark Energy*. p. 50 ([arXiv:astro-ph/0411176](https://arxiv.org/abs/astro-ph/0411176))
- Friedman A., 1922, *Zeitschrift für Physik*, 10, 377
- Garcia-Fernandez M., et al., 2016, preprint, ([arXiv:1611.10326](https://arxiv.org/abs/1611.10326))
- Giannantonio T., et al., 2016, *MNRAS*, 456, 3213
- Goobar A., Leibundgut B., 2011, *Annual Review of Nuclear and Particle Science*, 61, 251

- Grego L., Carlstrom J. E., Reese E. D., Holder G. P., Holzapfel W. L., Joy M. K., Mohr J. J., Patel S., 2001, *ApJ*, 552, 2
- Gruen D., Bernstein G. M., Jarvis M., Rowe B., Vikram V., Plazas A. A., Seitz S., 2015, preprint, (arXiv:1501.02802)
- Gruen D., et al., 2016, *MNRAS*, 455, 3367
- Gupta A., Mathur S., Krongold Y., Nicastro F., Galeazzi M., 2012, *ApJL*, 756, L8
- Guy J., et al., 2010, *A&A*, 523, A7
- Hand N., et al., 2015, *PRD*, 91, 062001
- Harnois-Déraps J., Munshi D., Valageas P., van Waerbeke L., Brax P., Coles P., Rizzo L., 2015, preprint, (arXiv:1506.06313)
- Harnois-Déraps J., et al., 2017, preprint, (arXiv:1703.03383)
- Harrison I., Camera S., Zuntz J., Brown M. L., 2016, *MNRAS*, 463, 3674
- Harvey D., Kneib J. P., Jauzac M., 2016, *MNRAS*, 458, 660
- Herbonnet R., Buddendiek A., Kuijken K., 2017, *A&A*, 599, A73
- Heymans C., et al., 2006, *MNRAS*, 368, 1323
- Heymans C., et al., 2012, *MNRAS*, 427, 146
- Heymans C., et al., 2013, *MNRAS*, 432, 2433
- Hildebrandt H., et al., 2016, *MNRAS*, 463, 635
- Hinshaw G., et al., 2013, *ApJS*, 208, 19
- Hirata C., Seljak U., 2003, *MNRAS*, 343, 459
- Hirata C. M., Seljak U., 2004, *PRD*, 70, 063526
- Hirata C. M., Mandelbaum R., Seljak U., Guzik J., Padmanabhan N., Blake C., Brinkmann J., Budávári T., Connolly A., Csabai I., Scranton R., Szalay A. S., 2004, *MNRAS*, 353, 529
- Hirata C. M., Mandelbaum R., Ishak M., Seljak U., Nichol R., Pimbblet K. A., Ross N. P., Wake D., 2007, *MNRAS*, 381, 1197
- Hoekstra H., 2013, preprint, (arXiv:1312.5981)
- Hoekstra H., Herbonnet R., Muzzin A., Babul A., Mahdavi A., Viola M., Cacciato M., 2015, *MNRAS*, 449, 685

- Hoekstra H., Viola M., Herbonnet R., 2016, preprint, (arXiv:1609.03281)
- Hoyle B., 2016, *Astronomy and Computing*, 16, 34
- Hoyle B., et al., 2017, preprint, (arXiv:1708.01532)
- Hu W., 1999, *ApJL*, 522, L21
- Huff E., Mandelbaum R., 2017, preprint, (arXiv:1702.02600)
- Huterer D., Takada M., Bernstein G., Jain B., 2006, *MNRAS*, 366, 101
- Jarvis M., Jain B., Bernstein G., Dolney D., 2006, *ApJ*, 644, 71
- Jarvis M., Schechter P., Jain B., 2008, preprint, (arXiv:0810.0027)
- Jarvis M., et al., 2015, preprint, (arXiv:1507.05603)
- Jee M. J., Tyson J. A., Schneider M. D., Wittman D., Schmidt S., Hilbert S., 2013, *ApJ*, 765, 74
- Jee M. J., Tyson J. A., Hilbert S., Schneider M. D., Schmidt S., Wittman D., 2016, *ApJ*, 824, 77
- Joachimi B., Bridle S. L., 2010, *A&A*, 523, A1
- Joachimi B., Mandelbaum R., Abdalla F. B., Bridle S. L., 2011, *A&A*, 527, A26
- Johnson A., et al., 2017, *MNRAS*, 465, 4118
- Joudaki S., Blake C., Heymans C., Choi A., Harnois-Deraps J., Hildebrandt H., Joachimi B., Johnson A., Mead A., Parkinson D., Viola M., van Waerbeke L., 2017, *MNRAS*, 465, 2033
- Jouvel S., Kneib J.-P., Ilbert O., Bernstein G., Arnouts S., Dahlen T., Ealet A., Milliard B., Aussel H., Capak P., Koekemoer A., Le Brun V., McCracken H., Salvato M., Scoville N., 2009, *A&A*, 504, 359
- Kacprzak T., Zuntz J., Rowe B., Bridle S., Refregier A., Amara A., Voigt L., Hirsch M., 2012, *MNRAS*, 427, 2711
- Kacprzak T., Bridle S., Rowe B., Voigt L., Zuntz J., Hirsch M., MacCrann N., 2013, preprint, (arXiv:1308.4663)
- Kacprzak T., et al., 2016, *MNRAS*, 463, 3653
- Kaiser N., 1994, in Durret F., Mazure A., Tran Thanh Van J., eds, *Clusters of Galaxies*. p. 269 (arXiv:astro-ph/9407004)
- Kaiser N., 1995, *ApJL*, 439, L1
- Kaiser N., 2000, *ApJ*, 537, 555

- Kaiser N., Squires G., 1993, *ApJ*, 404, 441
- Kaiser N., Squires G., Broadhurst T., 1995, *ApJ*, 449, 460
- Kaiser N., Wilson G., Luppino G. A., 2000, *ArXiv Astrophysics e-prints*,
- Kannawadi A., Mandelbaum R., Lackner C., 2015, *MNRAS*, 449, 3597
- Kilbinger M., 2015, *Reports on Progress in Physics*, 78, 086901
- Kilbinger M., Benabed K., Cappe O., Cardoso J.-F., Coupon J., Fort G., McCracken H. J., Prunet S., Robert C. P., Wraith D., 2011, preprint, ([arXiv:1101.0950](https://arxiv.org/abs/1101.0950))
- Kilbinger M., et al., 2013, *MNRAS*, 430, 2200
- Kilbinger M., Heymans C., Asgari M., Joudaki S., Schneider P., Simon P., Van Waerbeke L., Harnois-Déraps J., Hildebrandt H., Köhlinger F., Kuijken K., Viola M., 2017, preprint, ([arXiv:1702.05301](https://arxiv.org/abs/1702.05301))
- Kirk D., Rassat A., Host O., Bridle S., 2012, *MNRAS*, 424, 1647
- Kirk D., Laszlo I., Bridle S., Bean R., 2013, *MNRAS*, 430, 197
- Kirk D., Lahav O., Bridle S., Jouvel S., Abdalla F. B., Frieman J. A., 2015, *MNRAS*, 451, 4424
- Kirk D., et al., 2016, *MNRAS*, 459, 21
- Kirkman D., Tytler D., Suzuki N., O'Meara J. M., Lubin D., 2003, *ApJS*, 149, 1
- Kitching T., et al., 2010, preprint, ([arXiv:1009.0779](https://arxiv.org/abs/1009.0779))
- Kitching T. D., et al., 2012, *MNRAS*, 423, 3163
- Kitching T. D., Alsing J., Heavens A. F., Jimenez R., McEwen J. D., Verde L., 2016a, preprint, ([arXiv:1611.04954](https://arxiv.org/abs/1611.04954))
- Kitching T. D., Verde L., Heavens A. F., Jimenez R., 2016b, *MNRAS*, 459, 971
- Köhlinger F., Viola M., Joachimi B., Hoekstra H., van Uitert E., Hildebrandt H., Choi A., Erben T., Heymans C., Joudaki S., Klaes D., Kuijken K., Merten J., Miller L., Schneider P., Valentijn E. A., 2017, preprint, ([arXiv:1706.02892](https://arxiv.org/abs/1706.02892))
- Krause E., Eifler T., 2016, preprint, ([arXiv:1601.05779](https://arxiv.org/abs/1601.05779))
- Krause E., Eifler T., Blazek J., 2015, preprint, ([arXiv:1506.08730](https://arxiv.org/abs/1506.08730))
- Krause E., et al., 2017, preprint, ([arXiv:1706.09359](https://arxiv.org/abs/1706.09359))
- Kuijken K., 1999, preprint ([astro-ph/9904418](https://arxiv.org/abs/astro-ph/9904418))

- Kwan J., et al., 2017, MNRAS, 464, 4045
- LSST Science Collaboration 2009, preprint, (arXiv:0912.0201)
- Levi M., Bebek C., Beers T., Blum R., Cahn R., Eisenstein D., Flaugher B., Honscheid K., Kron R., Lahav O., McDonald P., Roe N., Schlegel D., representing the DESI collaboration 2013, preprint, (arXiv:1308.0847)
- Lewis A., 2009, MNRAS, 398, 471
- Lewis A., Bridle S., 2002, Phys. Rev., D66, 103511
- Lewis A., Challinor A., Lasenby A., 2000, ApJ, 538, 473
- Li C., Jing Y. P., Faltenbacher A., Wang J., 2013, ApJL, 770, L12
- Liu J., Hill J. C., 2015, PRD, 92, 063517
- Ma Z., Hu W., Huterer D., 2006, ApJ, 636, 21
- MacCrann N., et al., 2017, *in preparation*
- MacKay D. J. C., 2002, Information Theory, Inference & Learning Algorithms. Cambridge University Press, New York, NY, USA
- Mandelbaum R., Hirata C. M., Ishak M., Seljak U., Brinkmann J., 2006, MNRAS, 367, 611
- Mandelbaum R., et al., 2011, MNRAS, 410, 844
- Mandelbaum R., Slosar A., Baldauf T., Seljak U., Hirata C. M., Nakajima R., Reyes R., Smith R. E., 2013, MNRAS, 432, 1544
- Mandelbaum R., et al., 2014, ApJS, 212, 5
- Mandelbaum R., et al., 2015, MNRAS, 450, 2963
- Mandelbaum R., et al., 2017, preprint, (arXiv:1705.06745)
- Maoli R., Van Waerbeke L., Mellier Y., Schneider P., Jain B., Bernardeau F., Erben T., Fort B., 2001, A&A, 368, 766
- Massey R., et al., 2007a, ApJS, 172, 239
- Massey R., et al., 2007b, MNRAS, 376, 13
- Massey R., Kitching T., Richard J., 2010, Reports on Progress in Physics, 73, 086901
- Mauskopf P., et al., 1999, preprint (astro-ph/9911444)

- Melchior P., Sutter P. M., Sheldon E. S., Krause E., Wandelt B. D., 2014, MNRAS, 440, 2922
- Ménard B., Scranton R., Schmidt S., Morrison C., Jeong D., Budavari T., Rahman M., 2013, preprint, (arXiv:1303.4722)
- Milgrom M., 1983, ApJ, 270, 365
- Miller L., et al., 2013, MNRAS, 429, 2858
- Munshi D., Valageas P., van Waerbeke L., Heavens A., 2008, Physics Reports, 462, 67
- Myers A. D., Shanks T., Outram P. J., Frith W. J., Wolfendale A. W., 2004, MNRAS, 347, L67
- Newman J. A., 2008, ApJ, 684, 88
- Nicola A., Refregier A., Amara A., 2016, preprint, (arXiv:1607.01014)
- Okumura T., Jing Y. P., Li C., 2009, ApJ, 694, 214
- Patel P., Harrison I., Makhathini S., Abdalla F. B., Bacon D., Brown M., Heywood I., Jarvis M., Smirnov O., 2015, Advancing Astrophysics with the Square Kilometre Array (AASKA14), p. 30
- Peacock J., 1999, *Cosmological Physics*. Cambridge Astrophysics, Cambridge University Press, <https://books.google.co.uk/books?id=t80-yylU0j0C>
- Penzias A. & Wilson R., 1965, ApJ142, 419
- Percival W. J., Samushia L., Ross A. J., Shapiro C., Raccanelli A., 2011, *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 369, 5058
- Perlmutter S., et al., 1998, preprint (astro-ph/9812133)
- Perlmutter S., et al., 1999, ApJ, 517, 565
- Planck Collaboration 2015a, arXiv:astro-ph/1502.01582,
- Planck Collaboration 2015b, preprint, (arXiv:1502.01597)
- Planck Collaboration 2016, A&A, 594, A15
- Pogosyan D., Starobinsky A., 1995, in Mückel J. P., Gottloeber S., Müller V., eds, *Large Scale Structure in the Universe*. p. 118 (arXiv:astro-ph/9502019)
- Prat J., et al., 2017, preprint, (arXiv:1708.01537)
- Quartin M., Marra V., Amendola L., 2014, PRD, 89, 023009

- Rassat A., Amara A., Amendola L., Castander F. J., Kitching T., Kunz M., Refregier A., Wang Y., Weller J., 2008, preprint, (arXiv:0810.0003)
- Refregier A., 2003, *Monthly Notices of the Royal Astronomical Society*, 338, 35
- Refregier A., Bacon D., 2003, *MNRAS*, 338, 48
- Refregier A., Douspis M., 2008, in *Space Telescopes and Instrumentation 2008: Optical, Infrared, and Millimeter*. p. 701018 (arXiv:0807.4036), doi:10.1117/12.789730
- Refregier A., Kacprzak T., Amara A., Bridle S., Rowe B., 2012, *MNRAS*, 425, 1951
- Refsdal S., 1964, *MNRAS*, 128, 307
- Rhodes J., Refregier A., Groth E. J., 2001, *ApJL*, 552, L85
- Riess A. G., et al., 1998, *AJ*, 116, 1009
- Riess A. G., Macri L., Casertano S., Sosey M., Lampeitl H., Ferguson H. C., Filippenko A. V., Jha S. W., Li W., Chornock R., Sarkar D., 2009, *ApJ*, 699, 539
- Riess A. G., Macri L., Casertano S., Lampeitl H., Ferguson H. C., Filippenko A. V., Jha S. W., Li W., Chornock R., 2011, *ApJ*, 730, 119
- Riess A. G., Macri L. M., Hoffmann S. L., Scolnic D., Casertano S., Filippenko A. V., Tucker B. E., Reid M. J., Jones D. O., Silverman J. M., Chornock R., Challis P., Yuan W., Brown P. J., Foley R. J., 2016, *ApJ*, 826, 56
- Romualdez L. J., et al., 2016, preprint, (arXiv:1608.02502)
- Rowe B. T. P., Jarvis M., Mandelbaum R., Bernstein G. M., Bosch J., Simet M., Meyers J. E., Kacprzak T., Nakajima R., Zuntz J., Miyatake H., Dietrich J. P., Armstrong R., Melchior P., Gill M. S. S., 2015, *Astronomy and Computing*, 10, 121
- Rozo E., et al., 2016, *MNRAS*,
- Rubin V. C., Thonnard N., Ford Jr. W. K., 1978, *ApJL*, 225, L107
- Ryden B., 2003, *Introduction to Cosmology*. Addison-Wesley, <https://books.google.co.uk/books?id=z27vAAAAMAAJ>
- Samuroff S., et al., 2017a, preprint, (arXiv:1708.01534)
- Samuroff S., Troxel M. A., Bridle S. L., Zuntz J., MacCrann N., Krause E., Eifler T., Kirk D., 2017b, *MNRAS*, 465, L20
- Sánchez C., et al., 2017, *MNRAS*, 465, 746

- Schaan E., Krause E., Eifler T., Doré O., Miyatake H., Rhodes J., Spergel D. N., 2016, preprint, (arXiv:1607.01761)
- Schmidt F., Rozo E., Dodelson S., Hui L., Sheldon E., 2009, *ApJ*, 702, 593
- Schmidt F., Leauthaud A., Massey R., Rhodes J., George M. R., Koekemoer A. M., Finoguenov A., Tanaka M., 2012, *ApJL*, 744, L22
- Schneider P., van Waerbeke L., Mellier Y., 2002, *A&A*, 389, 729
- Schneider M., Knox L., Zhan H., Connolly A., 2006, *ApJ*, 651, 14
- Schneider P., Eifler T., Krause E., 2010, *A&A*, 520, A116
- Schrabback T., Erben T., Simon P., Miralles J.-M., Schneider P., Heymans C., Eifler T., Fosbury R. A. E., Freudling W., Hettterscheidt M., Hildebrandt H., Pirzkal N., 2007, *A&A*, 468, 823
- Scovacicchi D., Nichol R. C., Macaulay E., Bacon D., 2017, *MNRAS*, 465, 2862
- Scoville N., et al., 2007, *ApJS*, 172, 1
- Seitz C., Schneider P., 1997, *A&A*, 318, 687
- Sheldon E. S., Huff E. M., 2017, preprint, (arXiv:1702.02601)
- Simon P., Schneider P., 2016, preprint, (arXiv:1609.07937)
- Simon P., et al., 2013, *MNRAS*, 430, 2476
- Simpson F., Harnois-Déraps J., Heymans C., Jimenez R., Verde L., 2015, preprint, (arXiv:1507.04862)
- Singh S., Mandelbaum R., More S., 2015, *MNRAS*, 450, 2195
- Sivia D., Skilling J., 2006, *Data analysis: a Bayesian tutorial*. Oxford science publications, Oxford University Press, <https://books.google.co.uk/books?id=608ZAQAIAAJ>
- Smith R. E., Peacock J. A., Jenkins A., White S. D. M., Frenk C. S., Pearce F. R., Thomas P. A., Efstathiou G., Couchman H. M. P., 2003, *MNRAS*, 341, 1311
- Sollerman J., et al., 2006, in *Beyond Einstein - Physics for the 21st Century*. p. 14.1 (arXiv:astro-ph/0510026)
- Spergel D. N., et al., 2003, *ApJS*, 148, 175
- Spergel D., et al., 2015, preprint, (arXiv:1503.03757)
- Springel V., et al., 2005, *Nature*, 435, 629

- Sullivan M., et al., 2011, *ApJ*, 737, 102
- Sun L., Zhan H., Tao C., 2015, preprint, (arXiv:1512.00600)
- Suyu S. H., et al., 2016, preprint, (arXiv:1607.00017)
- Suzuki N., et al., 2012, *ApJ*, 746, 85
- Takahashi R., Sato M., Nishimichi T., Taruya A., Oguri M., 2012, *ApJ*, 761, 152
- Treu T., 2010, *ARAA*, 48, 87
- Treu T., Marshall P. J., 2016, *A&A Rev*, 24, 11
- Trotta R., 2017, preprint, (arXiv:1701.01467)
- Troxel M. A., et al., 2017, preprint, (arXiv:1708.01538)
- Van Waerbeke L., Mellier Y., Erben T., Cuillandre J. C., Bernardeau F., Maoli R., Bertin E., McCracken H. J., Le Fèvre O., Fort B., Dantel-Fort M., Jain B., Schneider P., 2000, *A&A*, 358, 30
- Van Waerbeke L., Mellier Y., Hoekstra H., 2005, *A&A*, 429, 75
- Van Waerbeke L., et al., 2013, *MNRAS*, 433, 3373
- Verlinde E. P., 2016, preprint, (arXiv:1611.02269)
- Voigt L. M., Bridle S. L., 2010, *MNRAS*, 404, 458
- Wambsganss J., 2005, in Mellier Y., Meylan G., eds, *IAU Symposium Vol. 225, Gravitational Lensing Impact on Cosmology*. pp 321–332, doi:10.1017/S1743921305002127
- Weinberg D. H., Mortonson M. J., Eisenstein D. J., Hirata C., Riess A. G., Rozo E., 2013, *Physics Reports*, 530, 87
- Whittaker L., Brown M. L., Battye R. A., 2015, preprint, (arXiv:1505.08131)
- Wittman D. M., Tyson J. A., Kirkman D., Dell’Antonio I., Bernstein G., 2000, *Nature*, 405, 143
- Wong K. C., Suyu S. H., Auger M. W., Bonvin V., Courbin F., Fassnacht C. D., Halkola A., Rusu C. E., Sluse D., Sonnenfeld A., Treu T., Collett T. E., Hilbert S., Koopmans L. V. E., Marshall P. J., Rumbaugh N., 2017, *MNRAS*, 465, 4895
- Wraith D., Kilbinger M., Benabed K., Cappé O., Cardoso J.-F., Fort G., Prunet S., Robert C. P., 2009, *PRD*, 80, 023507
- Zel’dovich Y. B., 1970, *A&A*, 5, 84

Zhan H., 2006, JCAP, 8, 008

Zhang P., Pen U.-L., Bernstein G., 2010, MNRAS, 405, 359

Zuntz J., Kacprzak T., Voigt L., Hirsch M., Rowe B., Bridle S., 2013, IM3SHAPE: a maximum likelihood galaxy shear measurement code for cosmic gravitational lensing (arXiv:1302.0183), doi:10.1093/mnras/stt1125

Zuntz J., Paterno M., Jennings E., Rudd D., Manzotti A., Dodelson S., Bridle S., Sehrish S., Kowalkowski J., 2015, Astronomy and Computing, 12, 45

Zuntz J., et al., 2017, preprint, (arXiv:1708.01533)

Zwicky F., 1937, ApJ, 86, 217

de Jong J. T. A., Verdoes Kleijn G. A., Kuijken K. H., Valentijn E. A., 2013, Experimental Astronomy, 35, 25

de Putter R., Doré O., Das S., 2014, ApJ, 780, 185

van Uitert E., et al., 2017, preprint, (arXiv:1706.05004)

Appendix A

Simulating Galaxies Below the Survey Detection Limit

In this appendix we seek to justify the flux correction applied in the earlier chapters to the simulated HOOPOE images. The decision to subtract a uniform flux contribution across each tile is designed to account for the impact of galaxies below the detection threshold on the background sky subtraction that occurs during image reduction in the real data. It was devised using simulation based tests, which we will describe here.

To accurately capture the blending and noise properties of the parent images, it is important to include a population below the nominal detection limit of the survey $M_{r,\text{lim}}$. Each time the simulations described in Chapter 3 define an input sample of nominally detectable COSMOS galaxies for a particular tile, we also construct a cache of faint objects below the survey detection limit. The number of subdetection objects to simulate is decided by constructing a smooth function $p(M_r)$ by cubic spline interpolation of the histogram of magnitudes in the relevant DES band in the full COSMOS catalogue. We then integrate $p(M_r)$ above and below the magnitude limit. The number of faint objects is then related to the detected sample size as:

$$N_{\text{faint}} = \frac{f_{\text{faint}}}{(1 - f_{\text{faint}})} \times N_{\text{det}}, \quad (\text{A.1})$$

where $f_{\text{faint}} \equiv \int_{M_{\text{lim}}}^{\infty} p(M) dM$ is the fraction of the weight of the normalised magnitude distribution $p(M_r)$ above the nominal DES detection limit, and M is the aperture magnitude in the relevant band. These objects are then allocated random sky positions (α, δ) within the coadd bounds, and drawn into the single-epoch images at the corresponding pixel coordinates.

The HOOPOE code simulates the image plane without background sky flux. The result is a

Gaussian noise field plus nominally detectable COSMOS profiles in each of the positions a galaxy was detected by SExtractor in the parent coadd. This is assumed to represent the real image plane, after background sky subtraction has been carried out. The additional step of drawing a random scatter of sub-detection galaxies, however, will leave a non-zero background flux in the images. To accurately represent the data we must, then, correct the images somehow to account for the fact that the background subtraction would have included these objects were they there in the real data.

In practice, the sky subtraction algorithm performs the follows calculations. The image is first broken down into square sub-regions of $\sim 100 \times 100$ pixels and those pixels flagged by SExtractor as associated with a detected object are removed. The rest of the pixels are then flattened and the histogram of the pixel vector is calculated. The background algorithm applies an iterative clipping process until that distribution converges to $\pm 3\sigma$ about its median. If the width σ varies by $> 20\%$ during this process the field is considered crowded. In this case the mode is estimated and this value is taken as a single-number background level within this region. In the case that the field is not deemed to be crowded, the arithmetic mean is used. SExtractor then interpolates these sub-regions across the image plane to produce a smooth background flux map.

The question is, then, are the subdetection galaxies in HOOPOE sufficiently faint and low in density to avoid triggering the mode calculation? If so we can describe the impact of the faint galaxies as a flat shift in the mean of the sky flux. Our aim here is to decide whether this is, in fact, the case. If this were true then a simple uniform correction sufficient to mimic the impact of the faint objects on the background subtraction.

We test this by simulating a single tile from the simulation, with and without the faint galaxies. All other properties, including the noise realisations in the coadd and input single-epoch images and the input sample of detectable COSMOS profiles are kept fixed. Visual inspection of the coadd images supports, on a qualitative level at least, the notion that the main impact of faint galaxies is to modulate the background noise levels. We show a random 75×75 pixel window taken from the simulated coadd images in Figure A.1 below.

We next rerun the SExtractor background algorithm to produce a new sky map for each simulation version, using the configuration parameters chosen by the DES Data Management division for use on the real Y1 data. We can see here that the spatial pattern of the sky flux across the image plane is maintained. However, the maps *are* visibly different, being brighter in the case with the faint objects. We show the background maps in each case, and the spatial pattern of the residuals in Figure A.2.

The next question is, can we correct for this difference with a uniform shift in the background

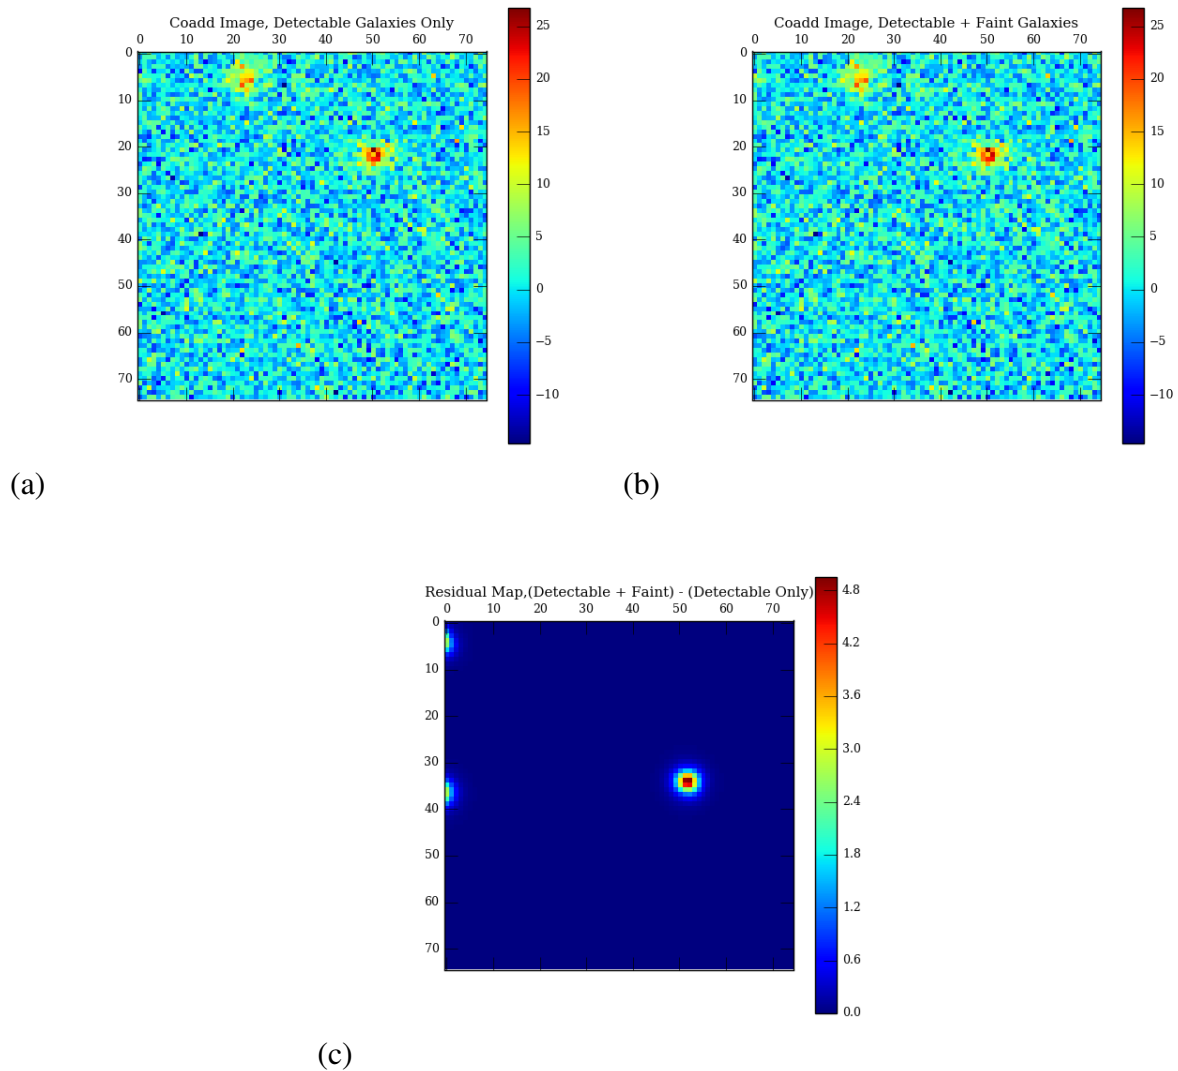


Figure A.1: **Top:** A sub-patch from two simulated instances of a coadd tile from DES, created with (right) and without (left) randomly place subdetection objects. All other properties of the two images are identical. **Bottom:** A map of pixel flux residuals (image with faint minus image without faint) for the same sub-patch.

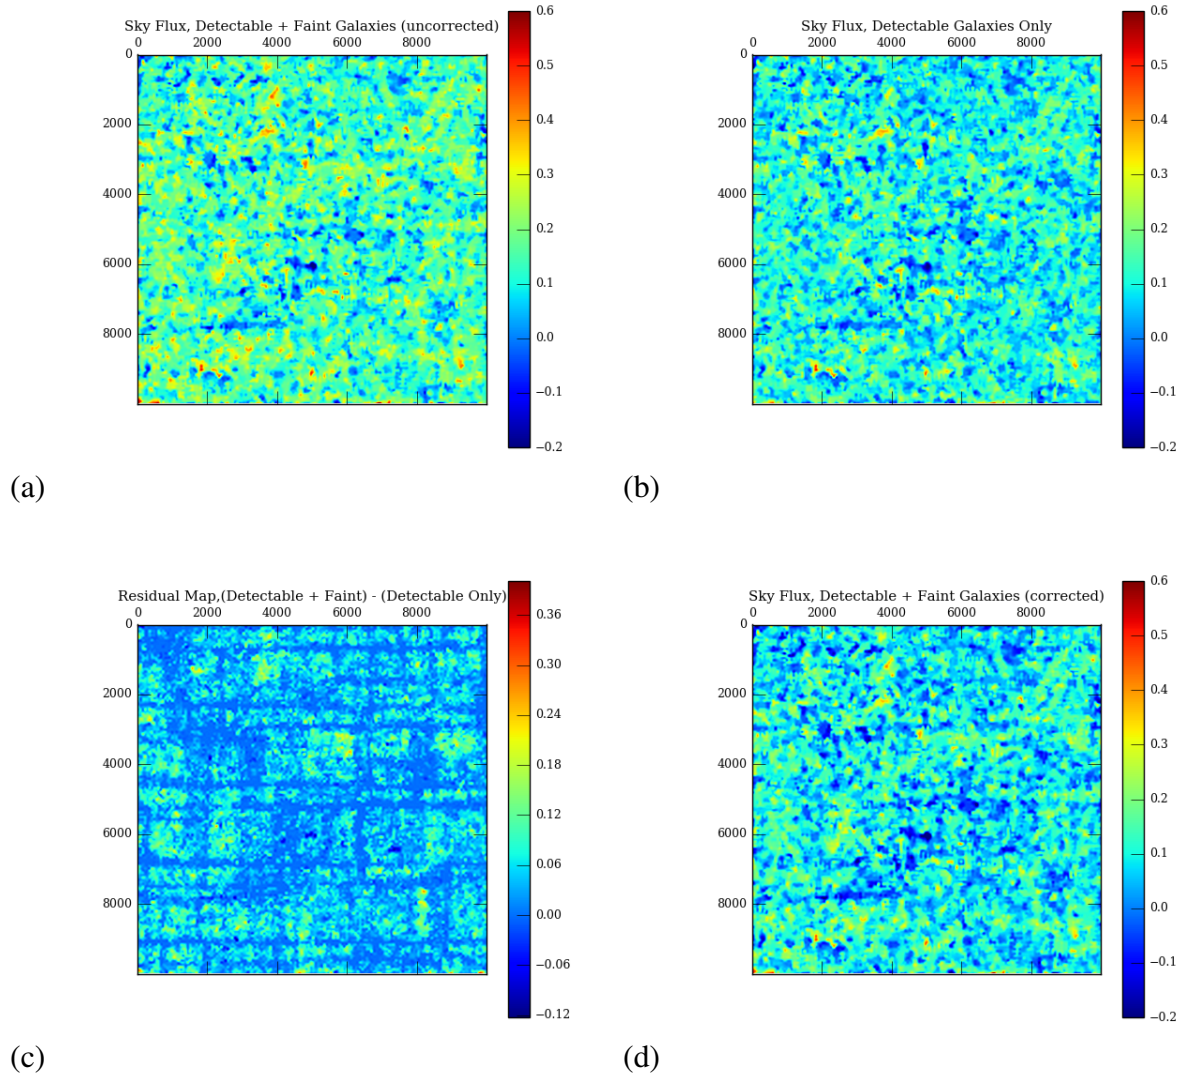


Figure A.2: The impact of sub-detection galaxies on background flux subtraction. The upper two panels show the sky-level background map estimated by running SExtractor on the simulated coadds. The residual between these two maps is shown in the lower left. The lower right-hand panel shows the result of imposing a uniform downwards shift on the background map that includes faint objects as described in equation A.2.

flux? If the main impact of adding galaxies below the detection limit is simply to shift the mean of the calculated sky flux (and thus increase the per-pixel value subtracted by DESDM), then we can apply a uniform correction to the images. To test this we plot out the histograms of the sky flux computed by SExtractor, shown by the solid blue and dot-dash purple lines in Figure A. We next apply a simple additive correction to each pixel in the background map that includes the faint galaxies:

$$\Delta f^{\text{faint}} = \sum_j^{N_{\text{faint}}} f_j^{\text{faint}} / N_{\text{pix}}, \quad (\text{A.2})$$

where the sum $j = (0, 1 \dots N_{\text{faint}})$ runs over the subdetection objects, f_j^{faint} is the total flux contribution of faint galaxy j , and N_{pix} is the total number of pixels in the coadd image. That is, we correct all pixels in the sky map by a flat value equal to the arithmetic mean of the flux from these galaxies. We show the corrected sky map in panel (d) of Figure A.2 and the histogram of pixel fluxes in Figure A (the solid purple line). As we can see, this does not perfectly reproduce the impact of faint galaxies, particularly in the negative tail of the distributions. A more accurate correction may be possible. As a leading-order correction to account for this effect, however, we judge this simple treatment to be sufficient.

Based on these results we apply a correction prior to the flux in pixel i in the simulation prior to shape measurement:

$$\tilde{f}_i = f_i - \Delta f^{\text{faint}}. \quad (\text{A.3})$$

The same additive correction Δf^{faint} is applied to each pixel. Given that $\sum_j^{N_{\text{faint}}} f_j^{\text{faint}}$ is small relative to the total flux of the image, in practice we expect this correction to have relatively little impact. We test this quantitatively at the level of galaxy shape estimates in Chapter 4, and verify it is the case.




Figure A.3: Histogram of sky map pixel fluxes. The solid blue line is the result of running the SEXTRACTOR background algorithm on an image with no subdetection galaxies. The purple dash-dotted line shows the impact of a background of faint galaxies. The solid purple is the same, but shifted by the mean flux contribution of the faint objects over the image plane.

Appendix B

Quality Checking COSMOS Images Using the DES Scientific Community

In this appendix we describe a web-based exercise set up to provide a means to quality control the IM3SHAPE calibration simulations described in Chapters 3 and 4. It was noted post-production that the simulated images contained a small number of artefacts, originating from defects in the input COSMOS profiles. These included deblending failures, and objects with diffuse light profiles truncated at the edges of the postage stamp. Two such objects are shown in Figure B.1. To assess the level to which these objects affect shape measurements on the simulations we initiated a small-scale crowdsourcing project within the scientific community of the Dark Energy Survey. Our specific aim here was to compile a list of COSMOS galaxies in our input catalogues that are qualitatively “bad”, and so should be excluded from our simulations.

The exercise was set up as follows. Each deconvolved COSMOS galaxy was reconvolved with a small nominal PSF and rendered into a postage stamp image at HST pixel resolution with no additional noise. The images were compiled in random order, and via a simple web interface¹, users were assigned batches of ~ 100 images. Galaxies were assigned to the categories show in Table B.1.

To test the impact of the aberrant COSMOS profiles on the IM3SHAPE calibrations we fit for multiplicative and additive bias in the HOOPOE dataset three times with different selection criteria: (a) IM3SHAPE quality cuts only; (b) removing any objects classed as “bad” for any reason; and (c) the same as (b), but additionally cutting any galaxies that fall within a circular aperture of 100 pixels around each flagged COSMOS profile. The results, in four DES Y1-like tomographic bins,

¹https://www.slac.stanford.edu/~dgruen/cosmos_eyeball




Figure B.1: Examples of profiles flagged as “bad” by the COSMOS classification exercise described in the text. The two galaxies shown here were classed as *artefact* (left) and *box too small* (right). For a breakdown of the number in each category see Table B.1.

Category	COSMOS Profiles	Galaxies in HOOPOE	COSMOS Profiles in HOOPOE
Total	87624	17.97 M	27612
Good	76707	16.93 M	25878
Box Too Small	3743	0.16 M	424
Artefact	1024	0.35 M	410
Two Galaxies	542	0.40 M	375
Galaxy Missing	4212	0.08 M	354
Off Centre	915	0.05 M	171
Other	481	0.10 M	127

Table B.1: The number of input galaxies in the Y1 DES image simulations presented in Chapter 3 falling under each category in the profile inspection exercise described. The first three columns show (left to right) the total number of COSMOS galaxies in each category from the full source catalogue from which the simulation draws profiles; the number of simulated galaxies affected; and the corresponding number of COSMOS profiles (note that the second and third columns are not identical since each COSMOS profile is drawn into multiple positions). The final column shows the change in mean multiplicative bias when galaxies in each category are removed.

are shown in Figures B.2.

The straightforward cut (b) induces a shift Δm that is comfortably within the level of statistical error of the fit. The second test suggests the corrupted profile may induce a small neighbour bias on surrounding profiles, which manifests as a modulation in m . It is worth pointing out that some of the categories listed in Table B.1 may be benign. Off-centred galaxies and those with neighbours, for example, should not cause a problem, since we re-run SEXTRACTOR object detection and deblending on the simulations. Our final cut on the simulation rejects instances of COSMOS profiles categorised under “artefact”, “box too small”, or “galaxy missing”. We test that additionally cutting the other categories does not induce a statistically significant change in bias. Based on the results in Figure B.2, we also incorporate a Gaussian component of width $\sigma_m = 0.005$ in the residual m prior for used in DES Y1 lensing analyses based on IM3SHAPE.




Figure B.2: The change in multiplicative and additive bias before and after removal of bad COSMOS profiles using the categorisation exercise described in the text. In the right-hand panel we show the two components c_1 and c_2 respectively as filled and open points. In each case the purple points show the impact of removing the affected COSMOS profiles. The blue points show the results when also cutting out galaxies within a circle of radius 100 pixels around each of them.

Appendix C

Derivation of a Two-Point Modifier for Scale Dependent Bias

In the following we set out a brief derivation of the analytic modifications to account for scale-dependent neighbour effects the shear-shear two-point correlations used in the earlier section. We do not claim that this is a precise calculation of the sort that could be used to derive a robust calibration. Rather it is an order of magnitude estimate to allow us to assess the approximate size of the cosmological bias these effects could induce in the data.

First, with complete generality it is possible to write the i component of the measured shear at angular position θ as

$$\gamma_i^{\text{obs}}(\boldsymbol{\theta}) = [1 + m_i(\boldsymbol{\theta})] \gamma_i(\boldsymbol{\theta}), \quad (\text{C.1})$$

where γ_i is the underlying true shear, which is sensitive to cosmology only. Extending this to the level of a two-point correlation between two populations α and β this implies:

$$\begin{aligned} \xi_i^{\text{obs},\alpha\beta}(\theta) &\equiv \left\langle \gamma_i^{\text{obs},\alpha}(\boldsymbol{\theta}') \gamma_i^{\text{obs},\beta}(\boldsymbol{\theta}' + \boldsymbol{\theta}) \right\rangle_{\theta} \\ &= \left\langle [1 + m_i^{\alpha}(\boldsymbol{\theta}')] [1 + m_i^{\beta}(\boldsymbol{\theta}' + \boldsymbol{\theta})] \tilde{\gamma}_i^{\alpha}(\boldsymbol{\theta}') \tilde{\gamma}_i^{\beta}(\boldsymbol{\theta}' + \boldsymbol{\theta}) \right\rangle_{\theta}. \end{aligned} \quad (\text{C.2})$$

Note that the observed shear used in a particular bin correlation is now weighted by the overdensity of galaxies in the image, in addition to the calibration bias, such that

$$\tilde{\gamma}_i^\alpha(\boldsymbol{\theta}) \equiv [1 + \delta_g^\alpha(\boldsymbol{\theta})] \times \gamma_i^\alpha(\boldsymbol{\theta}). \quad (\text{C.3})$$

Expanding each of the terms one finds:

$$\begin{aligned} \xi_i^{\text{obs},\alpha\beta}(\theta) = & \left\langle \gamma_i^\alpha(\boldsymbol{\theta}') \gamma_i^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \right\rangle_\theta \\ & + \left\langle m_i^\alpha(\boldsymbol{\theta}') \gamma_i^\alpha(\boldsymbol{\theta}') \gamma_i^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \right\rangle_\theta + \left\langle m_i^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \gamma_i^\alpha(\boldsymbol{\theta}') \gamma_i^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \right\rangle_\theta \\ & + \left\langle \delta_g^\alpha(\boldsymbol{\theta}') \gamma_i^\alpha(\boldsymbol{\theta}') \gamma_i^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \right\rangle_\theta + \left\langle \delta_g^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \gamma_i^\alpha(\boldsymbol{\theta}') \gamma_i^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \right\rangle_\theta \\ & + \left\langle m_i^\alpha(\boldsymbol{\theta}') m_i^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \gamma_i^\alpha(\boldsymbol{\theta}') \gamma_i^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \right\rangle_\theta \\ & + \left\langle \delta_g^\alpha(\boldsymbol{\theta}') \delta_g^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \gamma_i^\alpha(\boldsymbol{\theta}') \gamma_i^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \right\rangle_\theta. \end{aligned} \quad (\text{C.4})$$

The terms contributing to the measured two-point shear correlation, then, is sensitive to both spatial correlations between the m in different galaxies and to the correlations with the source density. Note that we have chosen to neglect a higher-order (six-point) term. In reality there will also be a connection between galaxy density and shear, but we will follow the convention and assume the contribution is small enough to be neglected. In simple terms, an excess in the $\langle mm \rangle$ term above the produce of the mean m values independently could arise because galaxy pairs separated on small scales tend to come from *similar* image plane environments. In contrast the density weighted correlations $\langle \delta_g m \rangle$ would be zero, but for a simple observation; selecting a random galaxy with a suitable correlation pair at a distance θ is *not* the same as unconditionally selecting a random galaxy. In the small scale bins we will over-sample the dense regions, where m tends to be larger (see Section 4.6.2).

The angular brackets here indicate averaging over all galaxy pairs separated by θ . If we can assume the bias is independent of the underlying cosmology the above expression simplifies significantly:

$$\begin{aligned} \xi_i^{\text{obs},\alpha\beta}(\theta) &= (1 + \bar{m}_i^\alpha + \bar{m}_i^\beta + \langle m_i^\alpha(\boldsymbol{\theta}')m_i^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta) \\ &\quad + \langle \delta_g^\alpha(\boldsymbol{\theta}')m_i^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta + \langle m_i^\alpha(\boldsymbol{\theta}')\delta_g^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta \\ &\quad + \langle \delta_g^\alpha(\boldsymbol{\theta}')\delta_g^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta \times \xi_i^{\alpha\beta}(\theta|\mathbf{p}), \end{aligned} \quad (\text{C.5})$$

with $\xi_i^{\alpha\beta}$ being the true correlation function of cosmological shears $\langle \gamma_i \gamma_i \rangle$, which is contingent on the underlying cosmological parameters \mathbf{p} . It can be shown that

$$\begin{aligned} \xi_+(\theta) &\equiv \langle \gamma_+(\boldsymbol{\theta}')\gamma_+(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta \pm \langle \gamma_\times(\boldsymbol{\theta}')\gamma_\times(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta \\ &= \langle \gamma_1(\boldsymbol{\theta}')\gamma_1(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta \pm \langle \gamma_2(\boldsymbol{\theta}')\gamma_2(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta \\ &= \xi_1(\theta) + \xi_2(\theta), \end{aligned} \quad (\text{C.6})$$

and so one can use equation C.5 to construct the observed ξ_\pm correlation functions

$$\begin{aligned} \xi_\pm^{\text{obs},\alpha\beta}(\theta) &= \left(1 + \bar{m}^\alpha + \bar{m}^\beta + \langle m^\alpha(\boldsymbol{\theta}')m^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta \right. \\ &\quad + \langle \delta_g^\alpha(\boldsymbol{\theta}')m^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta + \langle m^\alpha(\boldsymbol{\theta}')\delta_g^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta \\ &\quad \left. + \langle \delta_g^\alpha(\boldsymbol{\theta}')\delta_g^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta \right) \xi_\pm^{\alpha\beta}(\theta|\mathbf{p}). \end{aligned} \quad (\text{C.7})$$

The i subscript has been discarded here under the assumption that m_1 and m_2 are approximately equal for a given set of galaxies.

Next, let's say imagine that we have a measured datavector. Our measurements are biased, but we will assume it is possible to devise a correction that recovers the true cosmological signal precisely. Our observed datavector is then just,

$$\xi_\pm^{\text{obs},\alpha\beta}(\theta) = \Upsilon^{\text{tr},\alpha\beta} \xi_\pm^{\alpha\beta}(\theta|\mathbf{p}), \quad (\text{C.8})$$

which follows trivially from equation C.7. Since we do not trivially know $\Upsilon^{\text{tr},\alpha\beta}$ *ab initio* (this is why we need simulations!) we can only construct a best-estimate approximation. By applying a correction factor to the raw measurements we construct a best-estimate datavector:

$$\xi_{\pm}^{\text{BE},\alpha\beta}(\theta) = \frac{1}{\Upsilon_{\text{BE},\alpha\beta}} \xi_{\pm}^{\text{obs},\alpha\beta}(\theta) = \frac{\Upsilon^{\text{tr},\alpha\beta}}{\Upsilon_{\text{BE},\alpha\beta}} \xi_{\pm}^{\alpha\beta}(\theta|\mathbf{p}). \quad (\text{C.9})$$

Of course, if our best correction is perfect then the ratio goes to unity, and we recover the underlying cosmology. Since we apply corrections to the single-galaxy shears we will assume $\Upsilon^{\text{BE},\alpha\beta}$ includes the $\langle \delta_g \delta_g \rangle$ term, but neglects the correlations involving m . We then can write:

$$\Upsilon^{\text{BE},\alpha\beta} = \left(1 + \bar{m}^\alpha + \bar{m}^\beta + \bar{m}^\alpha \bar{m}^\beta + \langle \delta_g^\alpha(\boldsymbol{\theta}') \delta_g^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta \right). \quad (\text{C.10})$$

We can measure the mean bias in each bin that would be obtained from the calibration directly. As we show in Z17, using the full DES-Y1 HOOPOE catalogues, these biases are ~ -0.08 to -0.20 .

Finally, assume that although m clearly varies between redshift bins, the strength of the correlation does not. That is, the bias-bias term is the product of the mean values of m in each redshift bin. plus a scale dependent shift (which doesn't). One then has:

$$\langle m^\alpha(\boldsymbol{\theta}') m^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta = \bar{m}^\alpha \bar{m}^\beta + \Delta \xi_{mm}(\boldsymbol{\theta}). \quad (\text{C.11})$$

The additive part can be measured directly from the simulation using sub-patches, as described earlier. The density-density correlation can be obtained in the same way. This, then, leaves only the $m \times \delta_g$ cross-correlation. This should vanish in the case of zero correlation, but it also seems reasonable to assume that the magnitude should be proportional to the mean bias \bar{m}^α in a particular bin. This allows the scale dependent (non-tomographic) cross correlation measured from HOOPOE to be rescaled appropriately for each bin pair:

$$\langle \delta^\alpha(\boldsymbol{\theta}') m^\beta(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle_\theta = \left(\frac{\bar{m}^\beta}{\bar{m}} \right) \xi_{gm}(\boldsymbol{\theta}), \quad (\text{C.12})$$

where \bar{m} is the global multiplicative bias and $\xi_{gm}(\boldsymbol{\theta}) \equiv \langle m \delta_g \rangle$, each measured using all simulated galaxies.

With the above equations in hand, along with our fiducial calibration and three measured correlations, one can derive a scale dependent modification to shear-shear two-point correlation data using equation C.8.